

SEARCHING FOR DARK MATTER WITH THE ATLAS DETECTOR IN EVENTS WITH
AN ENERGETIC JET AND LARGE MISSING TRANSVERSE MOMENTUM

by

Steven Schramm

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Physics
University of Toronto

CERN-THESIS-2015-038
26/02/2015



© Copyright 2015 by Steven Schramm

Abstract

Searching for Dark Matter with the ATLAS Detector in Events with an Energetic Jet and Large Missing Transverse Momentum

Steven Schramm
Doctor of Philosophy
Graduate Department of Physics
University of Toronto
2015

Hadron colliders, such as the Large Hadron Collider (LHC), principally produce events involving hadronic activity. Such activity is typically modelled by jets, which provide a useful representation of the underlying physics. Given the ubiquity of jets in LHC events, it becomes important to ensure that their properties and performance are well understood. The full approach to jet reconstruction and calibration, as used by the ATLAS Experiment, is detailed with a focus on recent improvements. The systematic uncertainties associated with jets are quantified, with the procedures and resulting reductions in uncertainties thoroughly detailed. Extra attention is placed on the treatment of high energy jets, and particularly the impact of inactive calorimeter regions and calorimeter non-containment (punch-through).

The mono-jet topology is presented as an analysis where high energy jets are particularly relevant. This search makes use of very high missing transverse momentum balanced purely by jets, enabling measurements of the production cross-section of new physics processes producing weakly interacting particles. The full Standard Model background determination is shown, and the data are seen to be consistent with the Standard Model expectations. Limits are set on the visible cross-section for new physics processes.

The results are interpreted as a search for the pair-production of Dark Matter (DM), both through Effective Field Theories (EFTs) and simplified models. Scenarios where the DM is either a scalar or fermionic particle are both considered. The validity of the EFT approach is thoroughly investigated, providing the first complete collider study into all relevant interaction types. Limits are then set on the EFT suppression scale, the WIMP-nucleon scattering cross-section, and the WIMP annihilation cross-section in order to compare with other types of DM experiments. The simplified model is used to conduct a full parameter-space scan of the mono-jet sensitivity to a wide range of conditions.

First projections for the mono-jet analysis at an upgraded LHC are presented, demonstrating the significant gain in both discovery potential and limit sensitivity that accompanies higher collision energies. The analysis is projected to double in sensitivity in the coming year, hinting at the exciting times to come.

Acknowledgements

During my doctoral studies, I had the fortune of meeting and working with a remarkable group of people. This thesis would not be the same without their support. That said, I must begin by thanking my family, and especially my parents, for their boundless patience and encouragement. I may have struggled to explain the intricacies of what I was working on, or confused them by rambling about incomprehensible topics, but they always supported and believed in me nonetheless.

Next, I wish to thank my supervisor, Pierre Savard. He gave me the freedom to pursue my own interests, yet no matter what I was doing, he could always provide valuable insight. His support enabled me to work at CERN for the majority of my degree, which was instrumental in furthering my professional development. When I set myself ambitious goals, he helped to ensure they became a reality.

I was fortunate to have the generous support of multiple benefactors during my studies. Thank you to the National Sciences and Engineering Research Council of Canada (NSERC) for repeatedly supporting my research, and the W. Garfield Weston Foundation for funding a year of my time at CERN. It is due to both of these groups that I was able to conduct the research which went into this thesis.

There are many people at the University of Toronto who helped me along the way. Thanks to Jen for all the emails and for making me look forward to the occasional Toronto visits. Thanks also to Ray and Sylvia for seeing me when I was around. Thank you to my fellow ATLAS students: Chav, Garrin, Joe, and Santiago. A special thanks goes to Santiago for the fascinating discussions and the sushi, on the rare cases where we were in Toronto at the same time. Additionally, thank you to Krystyna for making bureaucracy manageable and for repeatedly helping me no matter which country I was in. The Toronto group at CERN also deserves my thanks: Lily (honorary), Nina, Nicola, Reyhaneh, and Richard.

Thank you to the $\text{jet}/E_T^{\text{miss}}$ performance group as a whole for generally being supportive, enthusiastic, and full of people with interesting ideas. Special thanks goes to Bogdan, Caterina, Cigdem, Dag, Frederik, Michael, and Pierre-Antoine for teaching me everything I know about jets. Similarly, thanks to Cate, Claire, Peter, Rosa, and TJ for teaching me about E_T^{miss} . In particular, thank you to Rosa for answering my innumerable questions, and for laying the foundation of my knowledge. Her patience at all hours of the day and night, her continual support, and her friendship are all greatly appreciated.

The exotics physics group is also full of wonderful people working on a wide range of fascinating subjects. Thank you first and foremost to Philippe for conveying years of insight into mono-jet final states and the data-driven estimation thereof, and for always being available to cross-check anything I was doing. Thanks to my Dark friends, Johanna and Ruth, for sharing your thoughts and insights over the years as we figured out Dark Matter models together. Thanks also to Antonio and Caterina (again) for being fantastic subconvenors and for furthering the needs of the ATLAS Dark Matter community. A special thanks also goes to Thomas for all of his help in understanding the validity of our Dark Matter models, and in developing techniques to properly address the concerns of the global community.

Finally, last but definitely not least, I must reiterate my thanks for Caterina (for a third time). I cannot overstate how much I have learned from her during the past years. Regardless of whether it was related to jet performance, the mono-jet analysis, Dark Matter, birds, or any other topic, she was always there at all hours of the day to help.

Contents

1	Introduction	1
2	Contributions	3
3	Theoretical Background and Motivation for Dark Matter	5
3.1	Introduction to the Standard Model	5
3.1.1	Fundamental particles and forces	5
3.1.2	Electroweak symmetry breaking	8
3.1.3	Limitations of the Standard Model	9
3.2	Dark matter relic density	10
3.3	Astrophysical evidence for dark matter	12
3.3.1	Galactic scales	12
3.3.2	Galactic rotation curves	13
3.3.3	The Bullet Cluster	13
3.3.4	The Cosmic Microwave Background	15
3.4	Weakly interacting massive particles	18
4	The ATLAS Experiment	19
4.1	The CERN Large Hadron Collider	19
4.2	The ATLAS Detector	23
4.3	Inner detector	26
4.4	Calorimetry	28
4.4.1	Particle showers	29
4.4.2	Electromagnetic calorimeters	31
4.4.3	Hadronic calorimeters	33
4.4.4	Forward calorimeters	36
4.5	Muon spectrometers	37
4.6	Triggering system	39
4.7	Simulation	40
4.7.1	Event generation	41
4.7.2	Detector simulation and digitization	41
4.8	Computing and data management	42

5	ATLAS Reconstruction and Performance	44
5.1	Introduction	44
5.2	Electrons and photons	44
5.2.1	Reconstruction	44
5.2.2	Performance	45
5.3	Muons	46
5.3.1	Reconstruction	46
5.3.2	Performance	48
5.4	Tau leptons and b-jets	48
5.4.1	Tau lepton reconstruction and performance	48
5.4.2	b-jet reconstruction and performance	50
5.5	Missing transverse momentum	51
5.5.1	Reconstruction	51
5.5.2	Performance	54
6	Jet Reconstruction and Performance	57
6.1	Introduction	57
6.2	Topo-clustering and reconstruction	59
6.2.1	Topo-cluster formation	59
6.2.2	Jet reconstruction	62
6.2.3	Jet associations	65
6.2.4	Jet cleaning	66
6.3	Jet energy scale calibration	66
6.3.1	Introduction	66
6.3.2	Overview	67
6.4	Jet energy scale uncertainties	68
6.4.1	Introduction	68
6.4.2	Overview	69
6.5	Jet energy resolution	69
6.6	Conclusions	71
7	The Mono-jet Analysis	73
7.1	Introduction	73
7.2	Standard Model backgrounds	75
7.3	Signal and control region overview	76
7.4	Event selection and cleaning	77
7.4.1	Triggering events	78
7.4.2	Event quality	79
7.4.3	Object definitions	81
7.4.4	Overlap removal	84
7.4.5	Jet cleaning	84
7.4.6	Masked Tile calorimeter modules	85
7.4.7	Isolated track rejection	86
7.4.8	Pre-selection	87

7.4.9	Signal and control region specific selections	87
7.5	Dataset information	89
7.5.1	Data samples	89
7.5.2	Simulated samples	90
7.6	Background determination	91
7.6.1	Data-driven $W + \text{jets}$ and $Z + \text{jets}$ estimation technique	91
7.6.2	$Z \rightarrow \ell\ell$ control regions	94
7.6.3	$W \rightarrow \ell\nu$ control regions	98
7.6.4	Coherence of $Z \rightarrow \nu\nu + \text{jets}$ data-driven estimates	101
7.6.5	Multi-jet control region	105
7.6.6	Non-collision background estimation	107
7.6.7	Top control region	112
7.6.8	Estimating other backgrounds	115
7.7	Systematic uncertainties	115
7.7.1	Jet energy scale and resolution	115
7.7.2	Lepton identification and scale	116
7.7.3	$E_{\text{T}}^{\text{miss}}$ SoftTerm scale and resolution	116
7.7.4	Trigger efficiency	117
7.7.5	Data taking conditions	117
7.7.6	Event cleaning	118
7.7.7	Top normalization	118
7.7.8	Diboson normalization	118
7.7.9	Other background normalizations	122
7.7.10	Theoretical systematics on the transfer factor	122
7.7.11	Summary tables for dominant signal region processes	124
7.8	Results	126
7.8.1	Combined $Z \rightarrow \nu\nu + \text{jets}$ data-driven estimates	126
7.8.2	Event yields	128
7.8.3	Model-independent limits	128
7.9	Conclusions	132
8	Mono-jet Dark Matter Interpretation	133
8.1	Introduction	133
8.1.1	Effective field theories	133
8.1.2	Simplified models	137
8.2	Effective field theory results	139
8.2.1	Limits on the EFT suppression scale	139
8.2.2	Model validity	140
8.2.3	Collider limits	142
8.2.4	Comparisons to other experiments	142
8.3	Simplified model results	148
8.3.1	Production regimes	150
8.3.2	Parameter space scan	152
8.4	Conclusions	154

9 Mono-jet Prospects at an Upgraded LHC	155
9.1 Introduction	155
9.2 Simulation	156
9.3 Event selection	158
9.4 Expected sensitivity	159
9.4.1 Event yields	159
9.4.2 Effective field theories	161
9.4.3 Simplified models	164
9.5 Conclusions	165
10 Conclusions	167
Bibliography	170
Glossary	181
Appendices	185
A Introduction to Cosmology	186
A.1 The standard model of cosmology	186
A.2 History of the universe	189
B Review of Direct and Indirect Detection Experiments	191
B.1 Direct detection experiments	191
B.2 Indirect detection experiments	194
C Jets and Masked Tile Calorimeter Regions	197
C.1 Introduction	197
C.2 The impact of masked modules on jet kinematics	198
C.3 The impact of masked modules on other quantities	201
C.4 Cleaning bad jets	202
C.4.1 Cleaning cuts for jets adjacent to masked tile modules	203
C.4.2 Cleaning cuts for jets centred within masked tile modules	204
C.4.3 Cleaning cut limitations	206
C.5 Summary	206
D Jet Energy Scale Calibration	209
D.1 Pileup suppression	209
D.1.1 Jet area	211
D.1.2 Event density	211
D.1.3 Applying the correction	212
D.2 Origin correction	212
D.3 Absolute jet energy scale	214
D.4 Global sequential calibration	215
D.5 <i>In situ</i> jet measurements	217
D.5.1 Di-jet η -intercalibration	217

D.5.2	Z/ γ -jet balance	218
D.5.3	Multi-jet balance	219
D.5.4	<i>In situ</i> combination	221
D.6	Single hadron response measurements	222
E	Jet Energy Scale Uncertainties	224
E.1	Uncertainty components	224
E.1.1	Common to Z/ γ -jet balance	224
E.1.2	Z-jet balance	228
E.1.3	γ -jet balance	229
E.1.4	Multi-jet balance	230
E.1.5	Di-jet η -intercalibration	231
E.1.6	Pileup suppression	232
E.1.7	Flavour differences	232
E.1.8	Other sources	234
E.2	Uncertainty comparisons	236
E.2.1	The <i>Moriond2013</i> release	237
E.2.2	The <i>Final2012</i> release	237
E.2.3	Comparing releases	242
E.3	Reductions via matrix eigendecomposition	242
E.3.1	Eigendecomposition description	249
E.3.2	Global reductions	250
E.3.3	Category reductions	250
E.4	Alternative correlation scenarios	254
E.4.1	The <i>Moriond2013</i> alternative scenario	257
E.4.2	The <i>Final2012</i> alternative scenarios	258
E.4.3	Reductions	267
F	Jet Punch-through	268
F.1	Introduction	268
F.2	Identifying punch-through jets	269
F.3	Correcting for lost energy	272
F.4	Performance of the punch-through correction	274
G	Effective Field Theory Validity	279
G.1	Introduction	279
G.2	Connecting the EFT suppression scale to the mediator mass	280
G.2.1	Simple operators: D5, D8, C3, D9	280
G.2.2	Chirally suppressed operators: D1, C1	281
G.2.3	Gluon operators: D11, C5	282
G.2.4	Using the M_* to M_{med} relations	282
G.3	Rescaling limits	283
G.3.1	The procedure	283
G.3.2	Dependence of validity on the signal region	284

G.3.3	Dependence of validity on the DM mass	286
G.3.4	Dependence of validity on the centre of mass energy	287
G.4	Results	287
G.4.1	Full study at 8 TeV	287
G.4.2	Preliminary study at 8 and 14 TeV	288
H	Mono-jet Dark Matter Kinematic Distributions	300
H.1	Kinematic plots	300
H.2	Acceptance plots	305
I	Mono-jet Dark Matter Interpretation Signal Sample Details	308
I.1	Production and validation	308
I.1.1	Sample production	308
I.1.2	Model validation	309
I.1.3	Fast simulation validation	310
I.1.4	Acceptance and efficiency	312
I.2	Theoretical systematic uncertainties	313
I.2.1	PDF uncertainties	315
I.2.2	Showering uncertainties	317
I.2.3	Factorization and renormalization scales	317
I.2.4	Matching scales and procedure	322
I.3	Experimental systematic uncertainties	328
I.3.1	Detector uncertainties	328
I.3.2	Beam energy uncertainties	328
J	Pileup and the Mono-jet Analysis at an Upgraded LHC	333
J.1	Procedure	334
J.2	Pileup jet multiplicity	334
J.3	Matching method and JVF cut checks	336

Chapter 1

Introduction

One of the largest remaining open questions in physics is the topic of Dark Matter (DM). First postulated in the 1930s [1, 2], many independent astrophysical experiments have observed the effects of DM. Cosmology has even measured its abundance to be approximately five times that of the visible matter which makes up the universe [3], but so far there exists no experimentally verified theory. While numerous experiments have been designed to search for DM, and some have claimed observations consistent with expectations for such a phenomenon [4], the nature of DM remains unknown.

Astrophysical studies have determined that the new state of matter is non-luminous (providing the name Dark Matter), meaning that it has no electromagnetic interaction, and non-baryonic, meaning there is no observed interaction under the strong nuclear force. The only force that it is known to interact with is gravity, which is how it has been observed and studied to date. However, distinguishing DM interactions from background processes within a localized detector and using current technology requires some force stronger than gravity. As such, the baseline for almost every DM search adds one additional constraint, namely that Dark Matter interacts with particles from the Standard Model (SM) through some weakly interacting force. This interaction may be the weak nuclear force, or it may be part of some as-of-yet undiscovered phenomenon. In this general class of theories, DM particles are known as a type of Weakly Interacting Massive Particle (WIMP).

Three general classifications of experimental searches for WIMP DM exist. The first type exploits the presence of DM in the local solar system, and watches for existing DM recoiling off of particles within the sensitive volume of a detector. These experiments are watching for very rare events, and thus must work in very specialized locations to avoid contamination, typically located deep underground. These types of experiments are known as direct detection searches, as they are looking for a direct interaction between DM and a normal particle. Until recently, the primary constraints on WIMP models of DM came primarily from such experiments.

The second type of experiment watches for particles coming from regions outside of the local solar system. The number and energy distribution of such particles is typically well known, and thus a deviation in either is a sign that some unexpected process is occurring. This can be from unexpected astrophysical sources, or it can be from the annihilation of dark matter, which is particularly likely when considering regions where large DM populations are known to exist. These types of experiments are typically located at the South Pole, on the tops of mountains, suspended from balloons in the upper atmosphere, or even better, in space. This is necessary as they need to have the best access possible to

particles coming from outside of the solar system, and thus the Earth's atmosphere typically interferes with their observations. These types of experiments are known as indirect detection searches, as they are looking for particles which came from the annihilation of DM particles rather than the DM itself.

The third area of research watches for the production of DM particles in very high energy collisions, rather than relying on observations of pre-existing DM. If the DM particles can annihilate to produce normal particles, then collisions between normal particles can also produce DM, so long as sufficient energy is available. The production of DM at colliders can be from different types of processes, such as direct pair-production of DM or decay chains involving other new particles which decay to give DM. As such, these experiments typically benefit from increasing the energy to higher and higher levels, although alternate possibilities and interpretations also exist.

Now that the LHC is operational, it is possible to probe entirely new regions of the DM parameter space in a complementary way to previous and ongoing experiments of all three classifications. This document will focus on the search for pair-production of DM in collisions at a new energy frontier. This is very challenging to do, as any produced DM particles will not deposit energy as they pass through the detector, making them impossible to directly observe. Instead, experiments must watch for indirect signatures, such as visible particle(s) recoiling off of the invisible DM particles. In such a case, the conservation of momentum states that the visible particles must be balanced by some non-interacting process, which can be measured and quantified as the missing transverse momentum of the collision.

A particularly interesting signature, given that the LHC collides hadrons, is to watch for a visible jet (the hadronization of a quark or gluon) recoiling off of the invisible particles. This is known as the mono-jet plus missing transverse momentum signature, and is one very promising channel with which Dark Matter production can be probed at the LHC. Such a search has been carried out by the ATLAS Experiment, leading to both complementary and competitive results when compared to the myriad of other experiments currently searching for Dark Matter.

The search for DM is a very competitive field. In order to remain at the sensitivity frontier, as much as possible must be gained from the available data, and thus it is crucial to understand the procedures used to the highest possible precision. For the mono-jet analysis, this means that the performance of methods for reconstructing both jets and missing transverse momentum is of paramount importance, especially at the very high momentum regime where standard techniques are pushed to their limits. This presents interesting challenges, which have led to a better understanding of both the techniques used for reconstruction of these objects and the underlying ATLAS detector.

In the near future, the LHC will be turned back on at a higher energy, which will be accompanied by a large increase in the amount of available data. This increased energy is highly beneficial for collider-based DM searches, and will ensure that the ATLAS Experiment remains at the forefront of the WIMP DM program in the coming years.

This thesis begins by considering the motivation for Dark Matter as a whole, the WIMP interpretation thereof, and existing searches for WIMP Dark Matter, as detailed in Chapter 3. The focus then shifts to the ATLAS Experiment, beginning with a description of the detector in Chapter 4 and general physics object reconstruction and performance in Chapter 5. A specific focus is then placed on one of the most important performance aspects of the mono-jet analysis, jets, in Chapter 6. The core of the mono-jet analysis is discussed in Chapter 7, while the interpretation of the analysis as a search for Dark Matter is contained in Chapter 8. The thesis concludes with an outlook on the possible sensitivity of the analysis under future LHC conditions in Chapter 9 and an overall conclusion in Chapter 10.

Chapter 2

Contributions

Experimental particle physics has evolved into a field which requires very intricate detectors and complicated research questions, where no one person can contribute to every aspect of the experiment. The ATLAS Collaboration is comprised of approximately 3000 members from across the world who have all contributed. This involves building detector components, monitoring the detector operation, translating detector signals into physics information, deriving calibrations, managing the computing resources, and any number of other tasks. In reality, any individual study relies on a significant amount of work done by other people, thus the full ATLAS author list is appended to all ATLAS publications.

This document also cannot exist in isolation, as it involves both direct and indirect contributions from many individuals. Major contributions of the author within this document are listed below.

Chapter 5: ATLAS Reconstruction and Performance

The author contributed to the treatment of missing transverse momentum in ATLAS. This includes updating the reconstruction software and investigating possible cleaning criteria for the rejection of events with large amounts of fake missing transverse momentum from improperly reconstructed muons.

Chapter 6: Jet Reconstruction and Performance

The author was involved in the development and implementation of jet reconstruction software, both during Run-I activities and in preparation for Run-II. The author also contributed to the validation of jet reconstruction and associated property calculations.

The author was one of two principal investigators studying the impact of masked Tile calorimeter regions on the performance of jets, and was responsible for determining the optimal cleaning cut criteria. The author also wrote the associated software package for general use, and was the principal contact for the full ATLAS experiment in matters relating to jet performance around these masked regions. The technical studies which went into the creation of the cleaning selection are provided in Appendix C.

The author made several small contributions to the jet calibration procedure as a whole, with a particular emphasis on corrections for jets which are not contained within the ATLAS calorimetry. The technical details of the full calibration procedure and the specific correction for non-contained jets are provided in Appendices D and F respectively.

The author was the primary contact, developer, and maintainer of all aspects of the systematic uncertainties associated with the jet energy scale. This covers a large number of contributions, including

(but not limited to) the creation of software tools to facilitate ease of access to the uncertainties, eigen-decompositions of the uncertainties to create a more accessible format, and the creation of alternative correlation scenarios to quantify the assumptions which went into the derivation of the uncertainties. All of this is thoroughly documented in Appendix E.

Chapter 7: The Mono-jet Analysis

The author contributed to many different aspects of the mono-jet analysis from the start to the end of the process. The author regularly provided input and feedback on all stages of the analysis, and provided cross-checks to the main result through the independent data-driven estimation of the dominant $Z \rightarrow \nu\nu + \text{jets}$ and $W \rightarrow \ell\nu + \text{jets}$ backgrounds, with a focus on the muon control regions. The author was also responsible for deriving the diboson normalization uncertainties.

In addition to these purely analysis-related contributions, the author also significantly improved performance-related aspects of the analysis. This includes optimization of the jet-related uncertainties, studies on the impact of punch-through jets, cleaning cuts for jets near masked Tile calorimeter modules, and optimization of the missing transverse momentum variation to use for the analysis.

Chapter 8: Mono-jet Dark Matter Interpretation

The author has contributed to essentially every aspect of the mono-jet DM interpretation. The author was solely responsible for the production of all of the DM signal samples, both for this analysis [5], the previous iteration of the analysis [6], and an analysis with a different but related topology [7]. This includes both the production of central values for typical analysis use and the production of samples with the variations required for the derivation of systematic uncertainties associated with the signal models. The production and validation of DM signal samples and associated systematic uncertainty derivations is provided in Appendix I due to its very technical nature.

The author was also the primary signal validation expert, contributed towards the evaluation of systematic uncertainties, and was jointly responsible for the calculation of limits for all of the DM models with an emphasis on the effective field theory samples.

The author led the ATLAS effort to understand the experimental implications of the validity of DM effective field theory models at the LHC, initially investigated for one scenario in Reference [8], fully studied in Reference [5], and now utilized by multiple other ATLAS searches such as References [7, 9]. Despite being a significant example of the author's contributions, the studies are very technical, and thus are documented in Appendix G.

Chapter 9: Mono-jet Prospects at an Upgraded LHC

The author is one of four principal analyzers in Reference [8], and thus was partially involved in most stages of the analysis. The author was solely responsible for the investigation of the impact of increased pileup conditions on jet performance, detailed in Appendix J due to its technical nature, as well as the aforementioned preliminary studies into effective field theory validity criteria. Beyond these two responsibilities, the author also contributed to the DM limit derivations for use in the sensitivity studies.

Chapter 3

Theoretical Background and Motivation for Dark Matter

3.1 Introduction to the Standard Model

The Standard Model (SM) has proven to be remarkably robust and adept at explaining the physics that is accessible at collider energy scales. With the recent discovery of the Higgs boson, which was the last missing piece of the SM, the theory is now complete. Despite being one of the most robust and successful physical models, the SM is not without its limitations, but overall the theory remains a useful starting point for understanding the background to searches for new physics.

3.1.1 Fundamental particles and forces

The SM of particles physics describes the fundamental particles and their interactions, with the exception of gravitational interactions. The model states that all matter is made up of two types of particles, namely fermions and bosons. Fermions, governed by Fermi-Dirac statistics, have half-integer spin and obey the Pauli exclusion principle. Bosons, on the other hand, have integer spin and follow Bose-Einstein statistics. This division also splits the fundamental building blocks of matter (fermions) from the mediators of the forces which control their interactions (bosons).

The SM also dictates that every particle, whether a fermion or a boson, has its own antiparticle. This antiparticle is identical in mass, but has opposite quantum charge. In some cases, a particle is its own antiparticle, but in most cases a separate antiparticle exists. The universe is currently dominated by matter, while antimatter has essentially vanished. This is known as Charge-Parity Violation (CPV), and is a major open topic of investigation. The SM does provide parameters which can provide CPV, but it is not yet known if the parameters are sufficient to represent the level of matter-antimatter asymmetry in the universe.

The different particles of the SM and their properties will now be explained, and are summarized in Figure 3.1.




Figure 3.1: An overview of the Standard Model and the associated properties of each particle [10].

Bosons

The bosons are divided up to account for the different forces, with each boson representing a different aspect of the respective force. The photon is the mediator of the electromagnetic force, and couples to all fermions which have a non-zero electromagnetic charge. It has a spin of 1, is electromagnetically neutral, massless, and is its own antiparticle. This was the first boson to be discovered.

The gluon is the mediator of the strong force, coupling to all fermions which carry a colour charge (quarks), and itself carries colour charge. This means that gluons not only mediate the strong force, but also act in it, and three-gluon and four-gluon vertices are allowed. This is significantly different from the photon, which is neutral under the charge of the relevant force. Gluons are also massless, and come in eight varieties representing the different combinations of colour. The possible ninth gluon representing the singlet state of $(r\bar{r} + b\bar{b} + g\bar{g})/\sqrt{3}$ does not exist, thus removing the long-range interaction potential of the gluon and limiting them to short distance scales [11]. The gluon has a spin of 1.

The W^\pm and Z^0 bosons are the mediators of the weak force, and couple to all fermions. The W^\pm bosons are each others antiparticles, and carry the weak charged current, responsible for generation-changing and flavour-changing processes to be discussed shortly. The W^\pm bosons have an electromagnetic charge of ± 1 , and a mass of 80.385 ± 0.015 GeV [12]. The Z^0 on the other hand is its own antiparticle, and is the mediator of the weak neutral current. It is electromagnetically neutral, couples to the same particles as the photon (with the notable addition of neutrinos), and has a mass of 91.1876 ± 0.0021 GeV [12]. Both the W^\pm bosons and the Z^0 boson have a spin of 1.

The final boson of the SM is the recently discovered Higgs boson (officially known as the boson of

the Brout-Englert-Higgs mechanism). It has a mass currently measured to be 125.7 ± 0.4 GeV [12], is electrically neutral, and is the first fundamental scalar particle (spin 0) to be observed. The Higgs boson can be thought of as the carrier of *boson mass*, as it is the reason that the W^\pm and Z^0 bosons are not massless. This will be discussed further in Section 3.1.2.

Fermions

The fermions are divided into two groups to account for radically different properties. The two groups are the quarks and leptons, all of which have a spin of $1/2$. Both quarks and leptons are further divided into three generations, where the first generation corresponds to what exists in common matter, and the second and third generations are accessed with increasingly higher energies. Going from the first to second to third generation always increases the mass of the fermion, while leaving invariant the other quantum properties such as the electromagnetic charge. In general, the fermion lifetime also decreases with increasing generation, but this is not a strict rule. One notable exception to this trend is the charm (second generation) quark as compared to the bottom (third generation) quark. In this case, the charm lifetime is actually shorter, which has practical implications for the identification of charm and bottom quarks in collisions.

Quarks exist in a further two types, named *up-type* and *down-type* after the name of the first generation quarks. The two types differ in their electromagnetic charge, where *up-type* quarks (anti-quarks) have a charge of $2/3$ ($-2/3$) while *down-type* quarks (anti-quarks) have a charge of $-1/3$ ($1/3$). The masses increase significantly as the generation increases, varying from $2.3_{-0.5}^{+0.7}$ MeV (up) to 1.275 ± 0.025 GeV (charm) to $173.21 \pm 0.51(\text{stat}) \pm 0.71(\text{syst})$ GeV (top) for the *up-type* quarks and $4.8_{-0.3}^{+0.5}$ MeV (down) to 95 ± 5 MeV (strange) to 4.18 ± 0.03 GeV (bottom) for the *down-type* quarks [12].

Leptons also exist in a further two categories, which are referred to as charged leptons and neutrinos. The charged leptons (charged anti-leptons) have a charge of 1 (-1), while neutrinos are electromagnetically neutral. The charged lepton masses are $0.510998928 \pm 0.000000011$ MeV (electron), $105.6583715 \pm 0.0000035$ MeV (muon), and 1776.82 ± 0.16 MeV (tau) [12]. Neutrino masses are more difficult to quantify, as the traditional SM considers them to be massless, but in reality they have been observed to be very small but nonzero. This will be discussed further in Section 3.1.3.

Gauge theories

Within the SM, the interactions are described by gauge theories, meaning that the associated Lagrangian is invariant under a given set of transformations - the symmetries of the problem. The representation matrices of these invariant transformations are the generators of the associated symmetry group. The symmetry groups of the SM are $SU(3)_c \times SU(2)_L \times U(1)_Y$. The $SU(3)_c$ group has $3^2 - 1 = 8$ generators, which are the eight gluons. The $SU(2)_L \times U(1)_Y$ group determines the electroweak sector, with the resulting $(2^2 - 1) + 1 = 4$ generators representing the massive W^\pm and Z^0 bosons as well as the massless photon γ .

However, local gauge symmetries require that the generators are massless. As mentioned, this is not the case for the W^\pm and Z^0 bosons. The answer to this contradiction comes from electroweak symmetry breaking, where the scalar Higgs boson gives mass to the otherwise massless W^\pm and Z^0 bosons. This is discussed in more detail in Section 3.1.2.

Up until now, the representation of the fermions has not been mentioned. To match the gauge theory, the left-handed fermions exist as doublets under $SU(2)_L$, while the right-handed fermions are singlets.

There are thus three fermion doublets and three fermion singlets, representing the three generations. Furthermore, there is no right-handed neutrino, and equivalently no left-handed anti-neutrino. The resulting lepton doublets and singlets are in Equation 3.1, while quark doublets and singlets are in Equation 3.2.

$$\begin{pmatrix} \nu_e \\ e \end{pmatrix}_L \quad \begin{pmatrix} \nu_\mu \\ \mu \end{pmatrix}_L \quad \begin{pmatrix} \nu_\tau \\ \tau \end{pmatrix}_L \quad e_R \quad \mu_R \quad \tau_R \quad (3.1)$$

$$\begin{pmatrix} u \\ d' \end{pmatrix}_L \quad \begin{pmatrix} c \\ s' \end{pmatrix}_L \quad \begin{pmatrix} t \\ b' \end{pmatrix}_L \quad u_R \quad d_R \quad c_R \quad s_R \quad t_R \quad b_R \quad (3.2)$$

It is worth noting that the quark representation has primed *down-type* quarks. This is because they are weak eigenstates rather than mass eigenstates, where the Cabibbo-Kobayashi-Maskawa (CKM) matrix parametrizes the rotation:

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = V_{\text{CKM}} \begin{pmatrix} d \\ s \\ b \end{pmatrix} \quad V_{\text{CKM}} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} = \begin{pmatrix} 0.974 & 0.225 & 0.004 \\ 0.225 & 0.973 & 0.041 \\ 0.009 & 0.040 & 0.999 \end{pmatrix} \quad (3.3)$$

The transition within a given generation (diagonal elements) is nearly 1, showing the dominance of the same generation. However, the emission/absorption of a W^\pm can change the generation (off-diagonal elements). This is a unique property of the weak force, as both strong and electromagnetic forces work only within a single generation.

Furthermore, this matrix demonstrates the existence of flavour-changing processes. The decay of the top quark (up-type) to a bottom quark (down-type) via the emission of a W is a perfect example, for which the relevant matrix element is 0.999. This is another example of a process which is only possible with the weak force.

3.1.2 Electroweak symmetry breaking

Spontaneous electroweak symmetry breaking, also known as the Brout-Englert-Higgs mechanism, is the means by which the massless weak bosons acquire a mass. As the symmetry of the Brout-Englert-Higgs potential is spontaneously broken, Goldstone bosons representing the broken degrees of freedom are generated, and each weak boson acquires one. This new degree of freedom corresponds to the longitudinal polarization mode, and comes with an associated mass, thus the weak bosons obtain mass in a way which does not break local gauge invariance. To do this, start by adding a generic left-handed isospin complex scalar doublet to the Lagrangian:

$$\phi = \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix} \quad \mathcal{L}_\phi = (\partial^\mu \phi)^\dagger (\partial_\mu \phi) - V \quad V = \mu^2 \phi^\dagger \phi + \lambda (\phi^\dagger \phi)^2 \quad (3.4)$$

The shape of the potential V depends on the parameters μ^2 and λ . The parameter λ must be above zero for the theory to remain stable, while μ^2 can be positive or negative. Picking the positive solution is uninteresting, while picking $\mu^2 < 0$ results in the so-called *Mexican Hat potential*. This potential is symmetric, with a local maximum at the origin. The distribution then falls to the minimum potential

before increasing again, thus creating a ring of local minima. The local minimum is identical around the ring of the local potential, thus while the Lagrangian remains invariant under $SU(2)_L \times U(1)_Y$, choosing any particular ground state breaks the symmetry. Parametrizing a small excitation of the complex scalar doublet about the chosen value, and working in the unitary gauge, results in [13]:

$$\phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + \eta(x) \end{pmatrix} \quad (3.5)$$

In this Equation, the *vacuum expectation value* (VEV) of the theory is v , with $v = \sqrt{\mu^2/\lambda}$. The Higgs boson is the real scalar written as η , with a mass $m_\eta = \sqrt{2\lambda}v$. Note that the other three degrees of freedom have disappeared, creating three massless Goldstone bosons. It can be shown that these three degrees of freedom are absorbed by the weak bosons.

Calculating the kinetic term of the Lagrangian from the covariant derivative of the field $\phi(x)$ then gives the following mass terms [13]:

$$m_W = \frac{gv}{2} \quad m_Z = \frac{v}{2}\sqrt{g^2 + g'^2} \quad m_\gamma = 0 \quad (3.6)$$

This is how the W bosons and Z boson obtain mass from the Higgs mechanism, while the photon remains massless. Counting the number of degrees of freedom, there were initially four in the complex scalar doublet, and now there is one for the scalar Higgs boson, and one adding a longitudinal component to each of the W^\pm and Z^0 bosons, for a total of four. The number of degrees of freedom of the system is conserved.

3.1.3 Limitations of the Standard Model

While the SM provides an excellent starting point, there are still a few open questions. First off, as mentioned, the SM states that neutrinos are massless. However, at this point neutrino oscillations have been thoroughly confirmed. This phenomenon means that neutrino flavour is not an eigenstate, rather there are mass eigenstates and the flavour of the neutrino is free to oscillate under a rotation matrix similar to the CKM matrix for quarks, called the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix.

Given that the neutrinos have mass, the next step is to quantify their masses. They have not yet been experimentally quantified, but the relative magnitude of the masses has been measured. In particular, it is known that two of the masses are similar in value ($\Delta m_{21}^2 = 7.53 \pm 0.18 \times 10^{-5} \text{ eV}^2$), while the third is significantly different ($\Delta m_{32,\text{normal}}^2 = 2.44 \pm 0.06 \times 10^{-3} \text{ eV}^2$, $\Delta m_{32,\text{inverted}}^2 = 2.52 \pm 0.07 \times 10^{-3} \text{ eV}^2$) [12]. However, it is not known which orientation the order follows. It is natural to assume that the masses get heavier as the generation number is increased, just like the quarks and the leptons, but this is not yet known. It is possible that neutrinos have an inverted mass hierarchy, where increasing the generation results in decreased masses. More studies are needed to determine which is the case.

The matter-antimatter asymmetry is another open question. The creation of the universe and thermal equilibrium should have resulted in the production of an equal amount of matter and antimatter. However, at the present day, there is a negligible amount of remaining antimatter in the visible universe. It can still be produced in collisions, but it is clear that the matter content of the universe significantly outweighs the antimatter content. The PMNS matrix and the CKM matrix provide CPV terms, which can account for part of the asymmetry, but it is not yet clear if they can account for all of it. This decision will have to wait until the neutrino CPV phase is better measured.

The strong-CP problem relates to the question of why there is no CPV in the strong force sector. While the leptonic sector can involve small CPV contributions from rotation phases, the strong sector naturally allows for CPV, but such a process has not been experimentally observed. A particle called the axion has been theorized to explain this, as the existence of the axion removes the possibility for CPV in the strong sector. So far, there is no clear indication as to whether the axion exists, or if there is another explanation which does not require exact cancellations of large numbers (fine-tuning). If the axion does exist, it is also a DM candidate, and thus the axion could solve two problems at once.

The next open question relates to the Higgs boson. The discovery of the Higgs boson completed the model, and solved the question of electroweak symmetry breaking. However, it is unknown why the Higgs boson mass is at the level of 125 GeV. This means that the electroweak scale is $\mathcal{O}(100 \text{ GeV})$, while the Planck scale is at $\mathcal{O}(10^{19} \text{ GeV})$. The electroweak symmetry breaking mechanism provides no explanation for this difference in scales. In particular, the Higgs propagator correction (such as a virtual top loop) is large, and is expected to increase the Higgs mass to the Planck scale. Given that this has been observed to not be the case, either the universe is considerably fine-tuned, or there is some form of new physics which cancels the divergences and stabilizes the Higgs boson mass.

Even if no new physics appears before the grand unification or even Planck scale, it is still known that the SM cannot be complete, as it does not provide a quantum description of gravity. Integrating a quantum gravity in with the other forces is a necessary step before a model can be considered truly complete.

The most directly relevant limitation of the SM is that it provides no candidate to account for the considerable DM fraction of the universe. The neutrino, as an electromagnetically neutral and weakly interacting particle with a non-zero mass seems like it could be the answer, but the amount of neutrino matter is a tiny fraction of what is required to match the observed abundance [14]. The SM also does not provide any indication into the nature of dark energy. As such, while the SM works well for ordinary matter, that matter represents less than 5% of the universe, and only 1/6 of the matter content. Assuming DM is a particle, the SM is clearly not complete.

3.2 Dark matter relic density

When discussing searches for DM, regardless of how the search is conducted, it is important to understand how much DM is expected to remain at this point in the history of the universe. An overview of the standard cosmological model, also known as the Λ CDM model, is provided in Appendix A. The theory is aptly named, as it includes a non-zero cosmological constant Λ , and considers Dark Matter as a Cold (non-relativistic) particle. What this means is that the *freeze-out* or decoupling of DM during the formation of the universe occurred when DM was non-relativistic. This is only one choice, as it is possible that DM is *hot* (decoupling at relativistic energies) or even non-particulate. These latter two scenarios are strongly disfavoured, but are not fully excluded. It is impossible to determine exactly which interpretation is correct without further evidence, preferably in the form of a direct observation of the non-gravitational properties of DM. Given that the assumption of cold DM has been made, it is important to discuss the *relic density*, or the abundance of DM remaining in the universe today.

The *freeze-out* of DM can be understood by considering the point of thermal equilibrium of a system. If a given particle has a suitable strong interaction rate, then the annihilation of existing particles is balanced by the creation of particles from the ambient energy of the system, and thus the abundance

of the particle is in a state of equilibrium. This occurred for example with hadrons in the hadron epoch, and leptons in the lepton epoch. As the universe expands and the temperature decreases, at some point the ambient energy is insufficient to produce the particle. The strength of the interaction then determines how much the particle continues to interact and annihilate before the expansion of the universe is dominant, and the oppositely charged particles no longer meet and annihilate, where the charge refers to a generic quantum property and not necessarily the electric charge.

The point at which the particles no longer interact is called the *freeze-out*, and is also known as decoupling. The energy at which this occurs is the *freeze-out* temperature T_F . The Boltzmann Equation can then be used to determine the number density of a given particle:

$$\frac{dn}{dt} + 3H_0 n = -\langle\sigma v\rangle (n^2 - n_{eq}^2) \quad (3.7)$$

In Equation 3.7, n is the current number density and n_{eq} is the number density at thermal equilibrium for the particle in question. The thermal average of the total annihilation cross-section times the particle velocity is denoted as $\langle\sigma v\rangle$. This equation implicitly assumes natural units, where $\hbar = c = 1$, thus v is relative to c and is a small number for non-relativistic particles. Making substitutions and solving the equation in the distinct regimes of much before T_F and long after T_F , as per Reference [15], leads to the following equation for the relic density of a generic particle χ :

$$\Omega_\chi h_0^2 \approx \frac{1.07 \times 10^9 \text{ GeV}^{-1}}{M_{\text{Pl}}} \cdot \frac{m_\chi}{T_F \sqrt{g_F^*}} \cdot \frac{1}{a + 3bT_F/m_\chi} \quad (3.8)$$

In the above, h_0 is the rescaled Hubble constant, $H_0 = (100h_0) \text{ kms}^{-1}\text{Mpc}^{-1}$, and a and b are coefficients of the expansion of $\langle\sigma v\rangle$ in powers of v^2 for non-relativistic particles. The final parameter g_F^* is the number of relativistic degrees of freedom at the time of *freeze-out*. Properly estimating this equation requires the calculation of the mass-dependent factors a and b , thus allowing for solving for the value of m_χ/T_F . Given the complexity of this approach, the following estimate is frequently used, which is accurate within an order of magnitude:

$$\Omega_\chi h_0^2 \approx \frac{3 \times 10^{-27} \text{ cm}^3 \text{ s}^{-1}}{\langle\sigma v\rangle} \quad (3.9)$$

This approximation makes clear the dependence of the relic density on the annihilation rate at *freeze-out*. In particular, the larger (smaller) the annihilation rate, the smaller (larger) the residual density of DM remaining in the universe today. This dependence is shown in Figure 3.2. The number density is seen to decrease exponentially as the ratio of the mass to temperature increases, stopping abruptly as the *freeze-out* temperature is reached. Lower values of the annihilation rate cause the *freeze-out* to occur sooner, and thus the remaining number density (relic density) of the particle is larger. This phenomenon where weaker annihilation rates leads to increased relic densities is often referred to as the *survival of the weak*.

It is important to note that these equations were all formed under a series of simplifying assumptions, all of which can be varied to cause significant differences in the end result. Some of the possibilities of such cases are:

- Coannihilations, where there is more than one state in the dark sector and the annihilation occurs between two different states, usually of a similar mass




Figure 3.2: The comoving number density (moving with respect to the expansion of the universe) as a function of the DM mass divided by the temperature. Increasing the thermally averaged annihilation rate times velocity $\langle\sigma v\rangle$ leads to higher *freeze-out* temperature and a lower remaining number density, which is the DM relic abundance. [16]

- The presence of a scalar particle in the early universe

These options are discussed in References [15, 17], but are not considered further here.

3.3 Astrophysical evidence for dark matter

The body of evidence for the existence of DM is continually growing. While the exact nature of DM remains a subject for debate, it is generally accepted that some form of non-interacting relic particle does exist. This Section presents a subset of the many independent observations which hint at the existence of DM.

3.3.1 Galactic scales

There are a large number of observations at the galactic and inter-galactic scale which support the existence of DM. Arguably the first evidence for the existence of DM came from such studies. One example is the Oort Discrepancy, which was first noted by Oort in 1932 and expanded on by Reference [1]. Oort studied the local galactic disk, and specifically the gravitational potential therein. He determined that the amount of visible matter within the disk was not enough to account for their distribution, and thus some unobserved matter must be make up the necessary balance.

A more striking example was observed the following year, in 1933, when Zwicky investigated the Coma cluster in Reference [2]. His measurements of velocity dispersion of galaxies within the cluster required a mass-to-light ratio of 400 solar masses per solar luminosity. This was two orders of magnitude larger than the local solar neighbourhood, and so Zwicky concluded that some other type of non-luminous matter existed, thus setting the stage for the modern notion of DM.

In the time after these initial findings, numerous other similar experiments have been conducted. Oort's observations have been enhanced [1], and Zwicky's technique has been applied to many other clusters [18]. Thanks to velocity dispersion techniques, it now appears that spiral galaxies are associated with dark halos which extend out to 200 kPc from the galactic centre, which is well beyond all of the visible matter. Other techniques have also been developed, and gravitational lensing experiments have made observations which cannot be accounted for by the visible matter content [19].

3.3.2 Galactic rotation curves

Arguably the most compelling evidence for DM comes from the circular velocity of stars and gas as a function of their distance from the galactic centre, also known as galactic rotation curves. The velocity measurements are typically obtained from the 21 cm hydrogen spectral line, and thus it relies on the presence of visible matter. The expected trend is well-known from Newtonian gravity, where the speed should fall off as:

$$v(r) = \sqrt{\frac{G_N M(r)}{r}} \quad M(r) = 4\pi \int \rho(r)r^2 dr \quad (3.10)$$

In the above, $\rho(r)$ is the visible matter density. If all matter is visible, then the speed should decrease as $\sqrt{1/r}$ when moving beyond the centre of the galaxy, where the majority of the visible galactic mass is clustered. This expected behaviour appears to be reproduced in some cases, but frequently observations of galaxies defy this expectation. Two galaxies for which this is very clearly not the case are shown in Figure 3.3. This shows how the circular velocity can be flat as the distance is increased, or in more rare cases it can even increase, well beyond the visible extent of the galaxy. This is compelling evidence for the existence of an invisible halo of DM within the galaxy, thus contributing to the velocity distribution in ways that are not predicted by purely the visible matter content [20]. In some cases, this can be taken even further to galaxies which appear to be entirely dominated by DM. The visible component is minimal, and the rotation curves of the visible component do not agree with what is expected from Newton's Equation [20].

Taking the more common example of the galaxy NGC6503 from Figure 3.3, where the rotation velocity appears to be constant with respect to the distance to the galactic centre, a form for the DM distribution can be inferred. If the expectation is that $v(r) \propto \sqrt{1/r}$ but instead $v(r) \propto 1$ is observed, then Equation 3.10 implies that the DM must have a mass distribution of $M_{\text{DM}}(r) \propto r$, or equivalently $\rho_{\text{DM}}(r) \propto 1/r^2$. This supports the existence of spherical DM halos within galaxies, although the exact structure of the distribution near the centre of the galaxy is less well understood.

3.3.3 The Bullet Cluster

The Bullet Cluster, officially named 1E 0657-558, is one of the most energetic known galaxy clusters in the universe. The cluster is comprised of two galaxies which, when viewed from Earth, are in the process of colliding. In addition to providing the unique opportunity to test many astrophysical predictions and models, the cluster has also proven to be what some call the definitive proof that dark matter is particulate.

The NASA Chandra X-ray Observatory has studied the cluster in detail, providing sufficient precision to determine the mass distribution of the underlying galaxies through weak gravitational lensing. As




Figure 3.3: Rotation curves for the NGC6503 dwarf spiral galaxy (left) and NGC1560 spiral galaxy (right), which are respectively 17 and 10-12 million light years away. The circular velocity is shown as a function of the distance from the galactic centre, with the gas content (dotted lines), matter content (dashed lines), and measured values shown. The difference between data and expectation gives rise to the addition of a curve representing a halo of DM which is fit to match the data (dash-dotted lines). [20]

seen in Figure 3.4, they made four main observations relating to the mass distributions:

1. Due to the large distance scales in question, stellar matter was only moderately affected. For the most part, the stars from each galaxy simply passed through the other galaxy without any inelastic interactions. The only visible effect is a velocity reduction through gravitational forces, with the occasional inelastic meeting of stars.
2. As usual, the gaseous component of the galaxies is much more spread out. The meeting of two gas clouds results in a significant contribution from the electromagnetic force, due to the shorter length scales. As such, significantly more interactions are possible, and the velocities of the two gas clouds is significantly reduced.
3. The centre of total mass of the galaxies, observed through weak gravitational lensing, is offset from the stellar and gaseous matter. This suggests the presence of additional invisible matter.
4. The DM distributions can be inferred from the total mass contours, and they remain mostly spherical in shape. This implies that the DM particles did not interact in any significant fashion beyond gravity, but rather the two DM structures passed through each other.

The observation of the DM content of each galaxy passing through without interactions is seen by Reference [22] as direct evidence that DM exists and is particulate, as modified theories of gravity predict that the DM-like phenomenon would instead interact in a similar manner to the gas. The authors further claim that this is 8σ evidence for particulate dark matter over modified theories, although Reference [23] disagrees with this conclusion following simulations which were capable of producing results similar to the Bullet Cluster observations. In short, while the Bullet Cluster provides compelling evidence for the existence of particulate DM, it cannot yet be taken as definitive proof.




Figure 3.4: The Bullet Cluster (1E 0657-558). Left: the visible matter content is shown in pink, and inferred dark matter content in blue [21]. Right: the visible matter content is shown in colours with higher intensity on the red side of the spectrum, while the weak gravitational lensing contours are superimposed, showing an offset centre of mass with respect to the visual centre for each galaxy [22].

3.3.4 The Cosmic Microwave Background

While all of the observations mentioned so far paint a compelling picture of DM, they do not quantify the amount of DM that exists on a larger scale. In particular, they provide no indication of the relative amounts of luminous matter and DM in the universe. In addition to measuring several other important astrophysical quantities, the Cosmic Microwave Background (CMB) provides an answer to this question.

As mentioned in Appendix A.2, the CMB is produced when the energy scale of the universe drops to the level of 1 eV (11605 K, using the Boltzmann constant $k_B \approx 8.617 \times 10^{-5}$ eV/K), approximately 3.8×10^5 years after the *Big Bang*. At this point, it becomes energetically favourable for the free electrons and protons to form hydrogen atoms, which is known as recombination. This immediately increases the range of photons from very short length scales to very long length scales, and is quickly followed by their decoupling. Due to this decoupling, the photons retain information about the state of the universe at the recombination timescale, and thus carries remnant information about the general properties of matter in the early universe.

At the present date, the CMB has a nearly uniform temperature of approximately 2.7255 K [12] (2.3×10^{-4} eV). This holds when mapping the CMB temperature of the full sky, with local fluctuations on the level of 10^{-4} K and the root mean square variation at approximately 2×10^{-5} K, as seen in Figure 3.5. The remaining structure is known as the anisotropy of the CMB, and matches what would be expected if small thermal variations from quantum fluctuations in a restricted spatial area had been diluted through expansion to the current size of the universe.

A precision measurement of the anisotropies of the universe provides a significant amount of information, particularly when parametrized as the power spectrum in terms of the angular scale (or equivalently the angular moment of spherical harmonics). This parametrization is:

$$\frac{\delta T}{T}(\theta, \phi) = \sum_{\ell=2}^{+\infty} \sum_{m=-\ell}^{m=+\ell} a_{\ell m} Y_{\ell m}(\theta, \phi) \quad (3.11)$$

In Equation 3.11, $Y_{\ell m}$ are the spherical harmonics and $a_{\ell m}$ are the multipole moments. The variance of a given moment is then:




Figure 3.5: The full sky map of the CMB after foreground subtraction as measured by the Planck Collaboration. The colour scheme varies from $-300 \mu\text{K}$ (blue) to $300 \mu\text{K}$ (red). Local temperature anisotropies are seen to be at the level of 10^{-4} K, while the global anisotropy is at the level of 10^{-5} K. [24, 25]

$$\langle |a_{\ell m}| \rangle = \frac{1}{2\ell + 1} \sum_{m=-\ell}^{m=+\ell} |a_{\ell m}|^2 \quad (3.12)$$

So long as the temperature fluctuations are Gaussian, which appears to be a reasonable approximation, this provides a means of accurately representing the information from the CMB as a function of the multipole moment. Fitting a distribution to the resulting curve, typically containing 6 or 7 parameters, allows for the determination of cosmological parameters. Each peak of this distribution can be related to cosmological parameter(s), thus providing a means of constraining cosmology through measurements of the CMB. The most recent such measurements come from the Planck Collaboration, where the resulting power spectrum and a six-parameter fit thereof is shown in Figure 3.6, both in terms of the angular scale and the spherical harmonic number.

Figure 3.6 demonstrates the remarkable precision to which the CMB power spectrum is currently measured, where the error bars represent the experimental uncertainties. The limitation of a single universe becomes particularly pronounced at large angles, or equivalently low multipole moments, where the number of possible relations between different areas in the sky rapidly decreases due to the large required angle of separation.

It is now important to understand the causes and meaning of the underlying anisotropies which give rise to the so-called acoustic peaks in the power spectrum. They are primarily the result of a competition between baryons and photons from the time of the baryon-photon plasma. The pressure of the relativistic photons works to erase temperature anisotropies, while the heavy non-relativistic baryons tend to form dense halos of matter, thus creating sizable local anisotropies. The competition between




Figure 3.6: The CMB power spectrum of spherical harmonics as measured by the Planck Collaboration, comparing temperature fluctuations between different points in the sky over the specified angular scale. The distribution is fit with a model including six parameters. [24, 3]

these two effects creates acoustic waves in the baryon-photon plasma, and is responsible for the observed acoustic oscillations.

The first acoustic peak is sensitive to the curvature of the universe, Ω_K . The fact that it doesn't occur until a multipole moment of approximately 200 signifies that the universe is relatively flat. The odd peaks are related to how much the medium compresses, which depends on the the baryonic content of the medium. Increased baryon numbers results in increased grouping, and thus further compression. The even peaks are instead related to how quickly the regions of compression are smoothed out, and thus are related to the photon content. The ratio of the first to second peak thus provides information on the baryon content of the universe with respect to the photons, where the photon content is known by measuring the CMB. The third peak then provides information on the total matter density of the universe. Using both of these values allows for a means of determining the fraction of baryonic versus non-baryonic matter. The 68% confidence level results from Planck for the resulting abundances, taken from Reference [3], are (in the nomenclature used for Appendix A.1):

$$\Omega_B h_0^2 = 0.02207 \pm 0.00033 \quad \Omega_{DM} h_0^2 = 0.1196 \pm 0.0031 \quad (3.13)$$

Planck further measures the ratio of baryonic matter, non-baryonic matter, and dark energy in Reference [3]. By the Planck measurements, the energy content of the universe is comprised of 68.3% dark energy, 26.7% dark matter, and 4.9% normal matter (baryonic and radiation). This provides compelling evidence for the existence of DM in large abundances throughout the universe, and points toward a larger open question to follow.

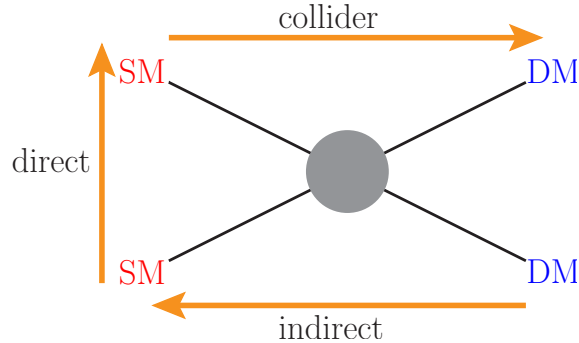


Figure 3.7: The three types of DM searches which become available when the WIMP assumption is made. Annihilation of SM particles to form DM particles is the collider approach, annihilation of DM particles and observing the SM products is the indirect detection approach, and the scattering of a DM particle off of a SM detector is the direct detection approach.

3.4 Weakly interacting massive particles

While many observations indicate the presence of DM, and the Bullet Cluster provides tantalizing hints that DM is particulate, there is as-of-yet no observation of any coupling to DM with a force other than gravity. Furthermore, the strong and electromagnetic forces have already been ruled out as possibilities. This poses a challenge for experiments aiming to search for DM, as the gravitational force is simply too weak for localized experiments. Instead, the assumption is made that DM has a very weak coupling to the SM, thus enabling several different types of searches. When this assumption has been made, DM is referred to as a type of Weakly Interacting Massive Particle (WIMP).

While the WIMP model is chosen without any clear indication of whether or not it is valid, it does have motivation. The first and most convincing motivation is the WIMP miracle, in which a particle with a mass on the weak scale ($\mathcal{O}(\text{GeV})$ to $\mathcal{O}(\text{TeV})$) which has a weak coupling to the SM naturally matches the observed relic density measured in CMB experiments [26]. This has historically been further reinforced by the belief in SuperSymmetry (SUSY), where several particles and in particular the neutralino have been proposed as natural WIMP candidates. SUSY is not the only model where this is the case, as there are numerous other theories for which WIMPs are expected [15].

When making the WIMP assumption, there are two generic types of searches. The first are searches for the associated Ultra-Violet (UV)-complete model, such as SUSY searches at colliders. The second classification pertains to generic DM searches, where it is assumed that the mediator connecting the SM to the dark sector is beyond the scale of the interaction and is thus inaccessible, resulting in a contact interaction. Whether or not this is applicable depends on the experiment in question, and will be discussed in more detail later, but for now it is assumed that such a coupling is at least permitted. This then gives rise to three related Feynman diagrams which can be obtained by rotating Figure 3.7 as appropriate. Note that while the Feynman diagram can be easily rotated as shown, there is no guarantee that the same DM particle is used in each process. For example, coannihilations would require two different particles from the dark sector in indirect detection, while the scattering approach in direct detection could involve only one of the two.

The collider production mode will not be further mentioned in this Section, as it is the focus of Chapters 7 and 8. A review of Direct and Indirect Detection experiments are provided in Appendices B.1 and B.2 respectively.

Chapter 4

The ATLAS Experiment

A substantial part of this Chapter is based off of Reference [27] for the Large Hadron Collider (LHC) and Reference [28] for the A Toroidal LHC ApparatuS (ATLAS) detector.

4.1 The CERN Large Hadron Collider

The Centre for European Nuclear Research (CERN, *Conseil Européen pour la Recherche Nucléaire*), established in 1954, is the world's largest institution for particle physics research. Based in the vicinity of Geneva, Switzerland, on the border between Switzerland and France, the organization is currently comprised of 22 full member states and maintains differing levels of relations with numerous other nations from around the world. The primary focus of CERN is the development, construction, and maintenance of a series of high energy accelerator-based experiments, of which the Large Hadron Collider (LHC) is the current flagship.

The LHC, a circular proton-proton synchrotron collider, first became operational in November 2009. With a circumference of 27 km, and a centre of mass collision energy of 7 or 8 TeV, the LHC is the currently the largest and highest energy particle accelerator in the world. The LHC is currently undergoing upgrades, and is expected to enable collisions to occur at a new energy frontier of 13 TeV in the spring of 2015, with a possible further upgrade to attain the design energy of 14 TeV. In addition to producing proton-proton collisions, the LHC is also used for proton-lead and lead-lead collisions. While interesting in their own right, these other types of collisions will not be considered further.

The LHC supports four main independent experimental collaborations, known as ALICE, ATLAS, CMS, and LHCb, which are spread along the circumference of the LHC as shown in Figure 4.1. The layout of the ring itself is broken up into octants, numbered in clockwise order, with the main Meyrin site at the bottom of the ring occupying octant 1. The eight octants each have a primary purpose, with four nominally representing collision points and associated experiments, and four representing purely LHC-related functionality [27]:

1. The ATLAS experiment, a collision point for proton-proton, proton-lead, and lead-lead collisions
2. The ALICE experiment, a collision point for proton-lead and lead-lead collisions
3. Beam cleaning for the LHC beams




Figure 4.1: The locations of the four main LHC experimental collaborations shown along the circumference of the CERN LHC. The main Meyrin site at Point 1 and the Geneva airport at Point 8 are also shown. [29]

4. Radio Frequency (RF) cavities for the actual acceleration of the LHC beams
5. The CMS experiment, a collision point for proton-proton, proton-lead, and lead-lead collisions
6. Beam dumping, for the LHC beams
7. Beam cleaning for the LHC beams
8. The LHCb experiment, a collision point for proton-proton collisions

The LHC is the final step of a multi-stage accelerator system, and is only one of the many experiments powered by the CERN accelerator complex, as shown in Figure 4.2 and detailed in Reference [30]. The injector chain for the LHC starts with a linear accelerator to bring the protons to an energy of 50 MeV. This linear accelerator is used to fill the 157 m circumference Proton Synchrotron Booster, which accelerates the protons to an energy of 1.4 GeV. These are then injected into the Proton Synchrotron, a 628 m circumference accelerator which brings the protons to an energy of 25 GeV. The next stage is the 7 km circumference Super Proton Synchrotron, which accelerates the protons to 450 GeV. These beams are then injected into the LHC, which accelerates the protons using the RF cavities in Point 4 to 3.5 TeV (2010 and 2011), 4 TeV (2012), 6.5 TeV (planned 2015), and potentially higher. Lead ions follow the same steps, with the exception of starting with a different linear accelerator, and replacing the Proton Synchrotron Booster with the 78 m Low Energy Ion Ring.

Within the LHC, two separate beams are run in opposite directions, which are then set on a collision course using quadrupole magnets at the four interactions points. These beams consist of *bunches* of protons, where the size and frequency of these bunches is directly relevant to the rate of interesting physics produced in collisions. This measure is known as the luminosity \mathcal{L} , which relates the number of events in a given amount of time to the cross-section of a desired process, as per Equation 4.1. The luminosity thus depends on the size of the bunches, the frequency with which bunches interact (frequency of bunch crossings), and the beam cross-sectional area of overlap \mathcal{A} , as defined in Equation 4.2.

$$\frac{dN_{\text{events}}}{dt} = \mathcal{L}\sigma \quad (4.1)$$

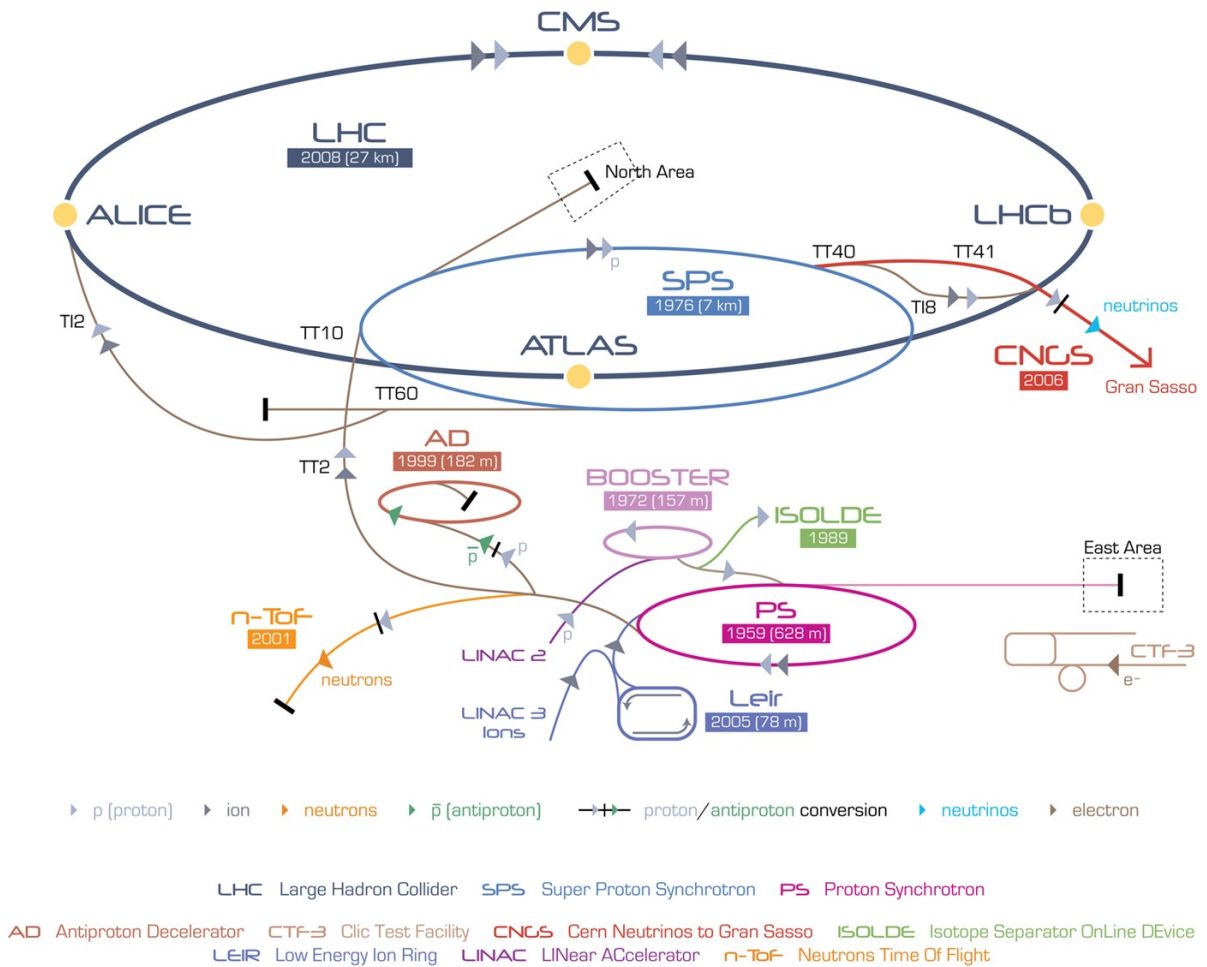


Figure 4.2: The CERN accelerator complex, showing the multi-stage system of injectors which are used to fill the LHC. A subset of the many experiments supported by these accelerators is also shown.[31]




Figure 4.3: The 2011 dataset switched the value for β^* partway through the year. A fit to the two values against the mean number of interactions per bunch crossing demonstrates the impact of squeezing beams on the mean number of collisions per crossing, and thus the luminosity. [32]

$$\mathcal{L} = \frac{f_{\text{crossing}} N_{\text{bunch1}} N_{\text{bunch2}}}{\mathcal{A}} \approx \frac{f_{\text{crossing}} N_{\text{bunch1}} N_{\text{bunch2}}}{4\varepsilon\beta^*} \quad (4.2)$$

The beam cross-sectional area of overlap \mathcal{A} is related to the beam cross-sectional size at the interaction point β^* and the beam emittance ε . The envelope that can be formed around a particle's trajectory as it undergoes betatron oscillation while traversing the accelerator is known as the amplitude function β , and the value at the interaction point is β^* . The beam emittance is independent of the location of the particle in the accelerator, and is instead related to the area of the phase ellipse in which particles can circulate.

Given the need to maximize the sensitivity to new physics signals, it is important to increase the luminosity as much as possible, thus providing as much integrated luminosity $\int \mathcal{L} dt$ as possible for analysis use. This can be done by increasing the number of particles in a bunch, increasing the crossing frequency, or reducing the cross-sectional area of overlap between the beams. Two different values for β^* are shown in Figure 4.3, demonstrating how reducing the amplitude function (squeezing the beam) leads to a higher number of average interactions per crossing. The overlap between the beams is unlikely to further change in a significant fashion, as this is limited by the quadrupole magnets, which are already used at near design values.

The time between bunches was typically 50 ns in Run-I of the LHC at both the ATLAS and Compact Muon Solenoid (CMS) experiments. This will be retained for the first data in Run-II, in order to ensure that the collider is stable with beams at 6.5 TeV, but is planned to be halved to 25 ns bunch spacing by the end of 2015.

Alternatively, the luminosity can be increased by adding more particles to each bunch. Given that the transverse area of each bunch is approximately fixed, this increases the probability that particles from opposing bunches will collide when crossed. This was the preferred technique for increasing the luminosity in the later stages of Run-I, most notably in 2012.

Increasing the luminosity from any of these techniques leads to the creation of pileup, which takes two forms. In-time pileup is the phenomenon by which multiple collisions happen in a single crossing, and thus the detector must be able to separate particles from multiple simultaneous interactions. Reducing the beam size or increasing the number of particles per bunch contributes to increasing the number

Variable		2010	2011	2012	Design value
Per-beam energy	[TeV]	3.5	3.5	4	7
β^* (amp. at int. point)	[m]	3.5,2.0	1.5,1.0	0.6	0.55
Bunch spacing ($1/f_{\text{crossing}}$)	[ns]	150	75,50	50	25
Max number of bunches		368	1380	1380	2808
Max number of protons	[10^{11} /bunch]	1.2	1.45	1.7	1.15
Peak luminosity	[$\text{cm}^{-2}\text{s}^{-1}$]	2.1×10^{32}	3.7×10^{33}	7.7×10^{33}	1×10^{34}
Max $\langle\mu\rangle$ (int/crossing)		4	17	37	19

Table 4.1: Luminosity-related LHC beam conditions in 2010, 2011, and 2012 as well as the design values for the LHC [33, 32]. The design conditions listed correspond to the LHC program itself, not the recently approved HL-LHC program. Max $\langle\mu\rangle$ corresponds to the peak value measured when averaging over a full luminosity block, as shown in Figure 4.4.

of simultaneous collisions, and thus the in-time pileup. The ability of different types of detectors to differentiate between simultaneous collisions varies, as will be discussed later in this Chapter.

Out-of-time pileup relates to the finite time required to for the signals to develop in and be read from the detector. Residual energy from a previous bunch crossing or energy from a future bunch crossing which is deposited before the read-out process is completed can contribute to the electrical signals associated to a given bunch crossing. This effect is primarily controlled by the amount of time between subsequent bunch crossings, and thus the amount of energy overlap. Some types of detectors can be read out at a rate faster than the bunch crossing frequency, rendering this irrelevant, but others take a longer amount of time. The effect on different types of detectors will be discussed later in this Chapter.

Pileup is typically parametrized in terms of the number of reconstructed primary vertices N_{PV} for in-time contributions and the average number of interactions per bunch crossing $\langle\mu\rangle$ for out-of-time contributions. N_{PV} is a per-event quantity as calculated from the number of reconstructed vertices in the tracking detectors, as will be discussed in Section 4.3, and which comes coupled with a vertex identification efficiency. On the other hand, $\langle\mu\rangle$ is averaged over a block of events, and is calculated from the luminosity. As such, the two values provide useful independent estimates of the amount of pileup activity.

As the LHC operations team becomes more familiar with the accelerator, they have successively increased the luminosity, and thus the associated pileup. A summary of the beam conditions as delivered to the ATLAS experiment is provided in Table 4.1. Figure 4.4 shows the peak instantaneous luminosity and peak instantaneous interactions per crossing, demonstrating the strong correlation between the two variables. This is then integrated over time to obtain Figure 4.5.

The ATLAS detector was designed to handle and excel under these high levels of pileup, and has been able to make full use of the data provided by the LHC.

4.2 The ATLAS Detector

The ATLAS detector is a large general-purpose experiment located at LHC Point 1, underneath the CERN main site located in Meyrin, Switzerland. The experiment is sensitive to many different areas of particles physics, from both proton-based or heavy ion-based collisions, including searches for new phenomena, investigations into the properties of the observed Higgs boson, and ultra high precision SM measurements.

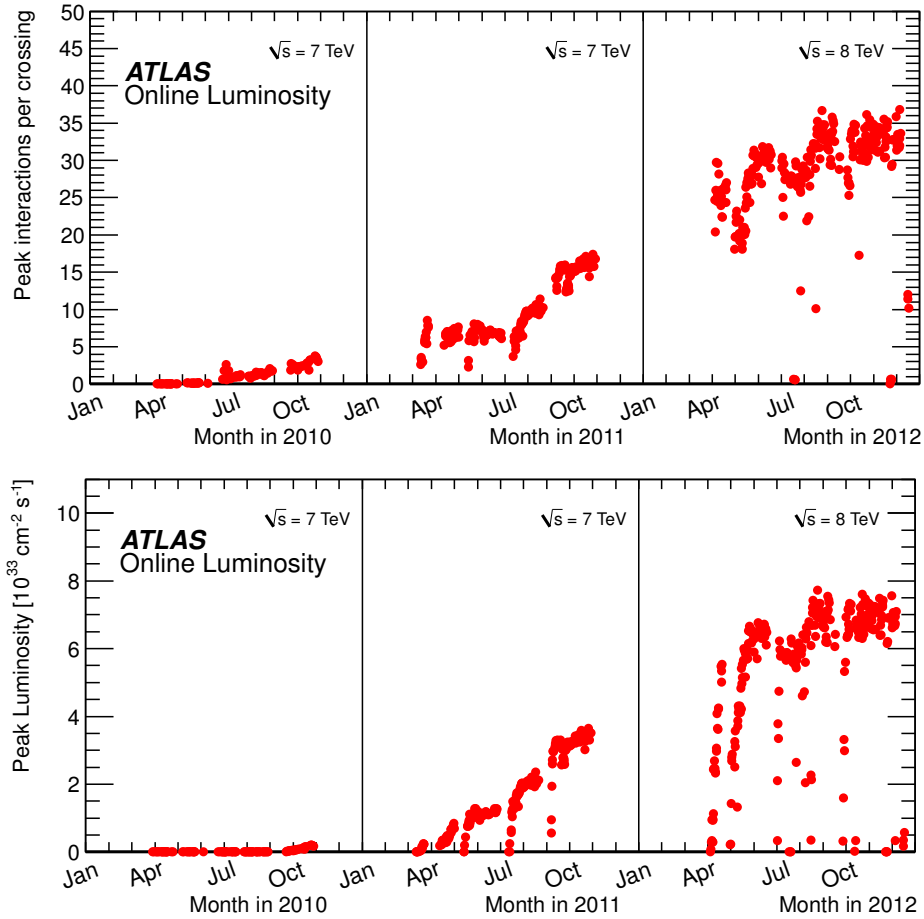


Figure 4.4: The peak of the number of interactions per bunch crossing averaged over the luminosity block $\langle\mu\rangle$ (top) and instantaneous luminosity (bottom) delivered by the LHC as observed by ATLAS. The 2010, 2011, and 2012 data taking periods are shown as a function of time, demonstrating the significant increase in available data as time progresses. [32]

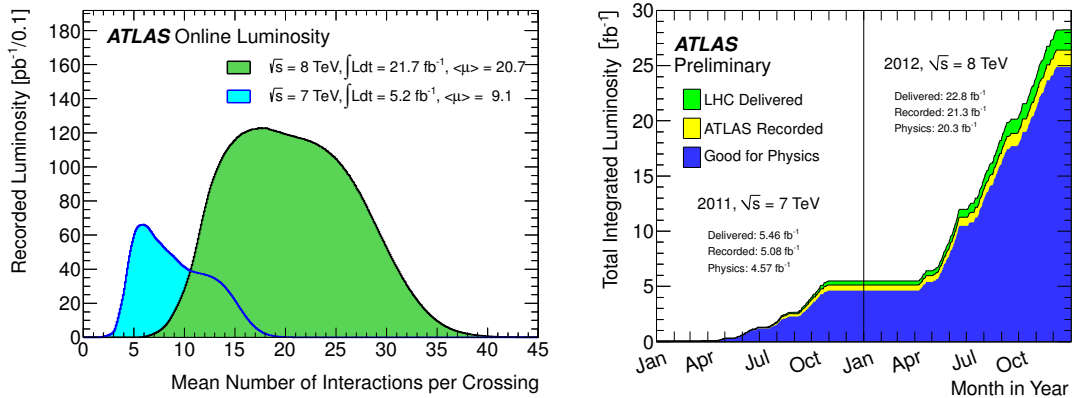


Figure 4.5: The mean number of interactions per bunch crossing as calculated from the instantaneous luminosity (left) and the integrated luminosity (right) delivered by the LHC as observed by ATLAS, for the 2011 and 2012 datasets. The doubling of $\langle\mu\rangle$ and halving of β^* leads to approximately four times more integrated luminosity in 2012. [32]




Figure 4.6: A computer-generated schematic representation of the ATLAS detector, with the various sub-detectors indicated and people shown for scale. [34]

Cylindrically shaped, the ATLAS detector is designed to be as close to hermetic as possible, thus enabling reconstruction and identification of nearly every energetic particle produced in the event. The detector has been designed to handle the expected high luminosity and high pileup conditions, with multiple optimizations designed to counteract the contribution of pileup, as will be discussed where relevant.

ATLAS uses a right-handed coordinate system with its origin at the nominal interaction point (IP) in the centre of the detector and the z -axis along the beam pipe. The x -axis points from the IP to the centre of the LHC ring, and the y -axis points upward. Cylindrical coordinates (r, ϕ) are used in the transverse plane, ϕ being the azimuthal angle around the beam pipe. The pseudorapidity is defined in terms of the polar angle θ as $\eta = -\ln \tan(\theta/2)$.¹

The ATLAS detector is made up of several sub-detectors. The innermost layers, immersed within a magnetic field from a 2 T solenoid, serve as tracking detectors with a coverage of up to $|\eta| < 2.5$. The middle layer of calorimeters are segmented into three main regions, with an ultimate coverage all the way out to $|\eta| < 4.9$. The outermost layer of the detector is the muon spectrometer, bathed in a toroidal magnetic field, which provides coverage for $|\eta| < 2.7$. These three levels of detectors will be discussed in Sections 4.3, 4.4, and 4.5 respectively. A schematic diagram showing the detector and its many components is provided in Figure 4.6.

Many sub-detectors involve two types of components, designed to maximize the precision of measurements over a wide range of angles. In the central $|\eta|$ region, where particles are approximately

¹This paragraph is directly copied from the ATLAS common coordinate system definition, as recommended by the ATLAS publications committee: <https://twiki.cern.ch/twiki/bin/view/AtlasProtected/CcoordinateSystemFootnote>.




Figure 4.7: A computer-generated schematic representation of the ATLAS inner detector, with the various sub-detectors indicated. [35]

perpendicular to direction of the beam, long cylindrical detectors around the beam pipe are used. Once the particle trajectory becomes more aligned with the beam, wheel-shaped endcap detectors are used, where the axle of the wheel is the beam pipe. These two scenarios are associated with very different levels of particle flux, and thus materials with different properties are often required, hence the typical separation of sub-detectors based on $|\eta|$.

This split between barrel and endcap sub-detectors applies to all three types of ATLAS detector classifications, where the exact $|\eta|$ separation between the two differs. Due to this detector scheme, the central $|\eta|$ region is also sometimes called the barrel region, and similarly for the endcap region. Only the calorimeters have significantly extended $|\eta|$ coverage, providing instrumentation in what is typically referred to simply as the forward region. Note that this classification excludes luminosity detectors and beam conditions monitors, where the former is in the very forward region and the latter exists in multiple regions.

4.3 Inner detector

The ATLAS inner detector, comprised of three sub-detectors, is entirely enclosed by a superconducting solenoid magnet providing a 2 T magnetic field. The tracker is designed to provide excellent track and vertex reconstruction with superb momentum resolution within the full range of $|\eta| < 2.5$. The nominal lower track p_T threshold is thus 500 MeV, although in some cases lower thresholds are possible. Electron identification is also provided by the outermost sub-detector over a large kinematic range of 500 MeV to 150 GeV within $|\eta| < 2.0$. A schematic diagram of only the ATLAS inner detector is provided in Figure 4.7. A technical diagram of the inner detector is provided in Figure 4.8.

Pixel Detector

The innermost layer of the inner detector is the Pixel Detector, which contains three layers of high-precision and high-granularity semiconductor modules. These 1744 sensor modules each contain 47232 pixels, which are individually $250 \mu\text{m}$ thick. For technical reasons relating to the read-out system, a

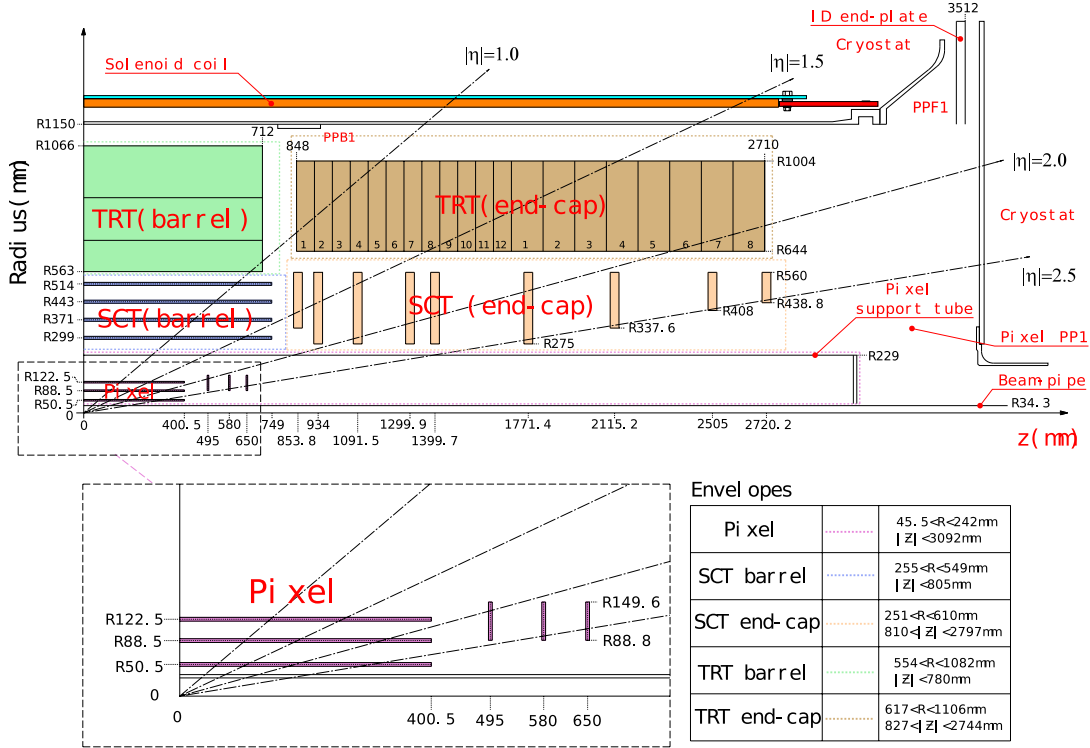


Figure 4.8: A quarter-sectional view of the ATLAS inner detector, showing each of the sub-detectors, their envelopes, and the geometrical space that they cover. [28]

small number of pixels per module are ganged, thus the total number of independent read-out channels is 46080 per module. This segmentation provides excellent position resolution, amounting to a per-layer resolution of $10 \mu\text{m}$ in the transverse ($r - \phi$) plane and $115 \mu\text{m}$ in the axial plane (z , barrel) or radial plane (r , endcaps).

The three layers within the barrel of the Pixel Detector are located at z of 50.5, 88.5, and 122.5 mm from the interaction point. The three layers within the endcap are located at r of 495, 580, and 650 mm from the interaction point. Due to the proximity of the lowest layer of the Pixel barrel to the interaction point, the lowest layer covers out to $|\eta| < 2.5$. At all points there are at least three layers of pixel active, but there is always one from the barrel and zero through two in the endcap.

SemiConductor Tracker

The second layer of the inner detector is the SemiConductor Tracker (SCT), which contains several layers of microstrip silicon wafers. These 15912 sensors are arranged in four layers within the barrel, and nine in the endcap. The endcap layers are arranged such that there are always four levels of sensors for particles to pass through. This segmentation provides high quality position resolution, amounting to a per-layer resolution of $17 \mu\text{m}$ in the transverse ($r - \phi$) plane and $580 \mu\text{m}$ in the axial (z , barrel) or radial (r , endcaps).

The four layers within the barrel of the SCT are located at z of 299, 371, 443, and 514 mm from the interaction point. The nine layers within the endcap are located at r of 853.8, 934, 1091.5, 1299.9, 1399.7, 1771.4, 2115.2, 2505, and 2720.2 mm from the interaction point. Unlike the pixel detector, the

barrel modules are only available up to $|\eta| \approx 1.5$, after which the endcap modules are used independently. Further details, such as the z extent of the endcap modules, is provided in Figure 4.8.

Transition Radiation Tracker

The third and outermost layer of the inner detector is the Transition Radiation Tracker (TRT), which as the name implies is a combined transition radiation detector and tracking detector. The primary building-block of the TRT is a set of 73 (160) layers of polyimide straw tubes in the barrel (endcap). These straws, each 4 mm in diameter, contain a gas mixture of 70% Xe, 27% CO₂, and 3% O₂. The resulting position resolution is significantly weaker than the Pixel Detector or SCT, with the entire detector providing a resolution of 130 μm .

Any charged particle passing through the straws ionizes the gas, thus producing ionization electrons. The outer wall of the straw, which is clad in a thin layer of metal, is held at a high negative voltage thus forcing the electrons toward the central anode wire. The electrons take up to 48 ns to reach the wire, where the arrival of the first ionization electron can be used to determine how close the particle passed to the centre of the straw.

The transition radiation part of the name comes from electrons, which are typically the only particles with a large enough Lorentz boost factor ($\gtrsim 1000$) emit X-rays as they pass through (transition between) the gas mixture and polyimide straws. The inclusion of 3% oxygen was found to reduce the ionization electron mobility, but increase the transiting electron identification power, and thus was retained. As the energy is increased, other particles (most notably charged pions) also begin to emit X-rays, hence the upper electron identification threshold of approximately 150 GeV.

The barrel of the TRT stretches from 563 to 1066 mm in z with respect to the interaction point, while the endcap covers 848 to 2710 mm in r . The barrel is only available within $|\eta| < 1.0$, and the endcap extent ends at $|\eta| \approx 2.0$. This is why the electron identification power of the inner detector is restricted to $|\eta| < 2.0$, rather than 2.5.

4.4 Calorimetry

After particles make it through the tracking detector, they reach the calorimeters. Calorimeters are designed to provide accurate energy measurements of all particles passing through them which undergo either electromagnetic or hadronic interactions. Only neutrinos and muons (which are Minimally Ionizing Particles (MIPs) at the LHC energy scale) pass through the calorimeters mostly unaffected.

As particle(s) from the hard-scatter collision interact with the calorimeter material, they create cascades of particles, also called showers. There are two types of cascades, depending on the nature of the source, and thus they are named electromagnetic and hadronic showers. These two types of showers have radically different properties, and thus require separate techniques for high-precision detection. Electromagnetic showers are primarily initiated by electrons or photons, while hadronic interactions are initiated by hadrons. Note that the hadrons which reach the calorimeter can come from very different sources. For example, hadronic showers can be the result of a single pion from tau decays, or they can be collections of hadrons from the fragmentation of quarks or gluons produced at the interaction point. These are but two examples of many possible processes, and demonstrate the large degree of variation that can be expected from hadronic showers.




Figure 4.9: A computer-generated schematic representation of the ATLAS calorimeters, with the various sub-detectors indicated. [36]

The ATLAS calorimetry, comprised of several sub-detectors, is between the inner superconducting solenoid magnet and the outer toroidal magnets. The calorimeters are designed to provide excellent electron and photon reconstruction with superb energy resolution and containment, while also providing exceptional hadronic shower reconstruction and performance. In order to meet these needs, there are two levels of detectors throughout the full range of $|\eta| < 4.9$, with an inner layer of electromagnetic calorimeters and an outer layer of hadronic calorimeters. Within each layer, there are three or four sub-detectors, covering the barrel and/or extended barrel, endcap, and forward regions. A schematic diagram of the ATLAS calorimetry is provided in Figure 4.9.

The ATLAS calorimeters are all sampling calorimeters, meaning that alternating layers of absorbing and active material are used. As particles traverse the calorimeters, the absorbing material is intended to force an interaction, thus producing either an electromagnetic or hadronic shower. The following active layers are designed to accurately measure the amount of energy in the particle cascade. This alternating layer design allows for creating compact calorimeters with excellent shower containment, but at the cost of energy lost in the absorption layers. The fraction of observed energy, known as the sampling fraction $f_{\text{samp}} = E_{\text{active}} / (E_{\text{active}} + E_{\text{passive}})$, provides a measure of quantifying this effect. If the sampling fraction for a given calorimeter is known, then the observed shower energy can be rescaled to account for the unobserved contribution.

4.4.1 Particle showers

There are several critical differences between electromagnetic and hadronic showers which control the associated detector requirements. Electromagnetic showers can be broken up into two types of contributions, namely electron/positron and photon interactions. At high energy, electrons emit photons via Brehmsstrahlung radiation while photons produce electron-positron pairs, demonstrating how these two sectors are not distinct. At medium energies, photons can also Compton scatter to reduce their energy, while at low energies the photoelectric effect converts them into low energy electrons. Low energy electrons are subsequently absorbed to form ions, thus ending the showering process.

In the end, the vast majority of the energy from electromagnetic showers is absorbed in the calorime-

ters, so long as a sufficiently deep calorimeter is used. However, not all of the material that the shower passes through is an active (recording) medium. Energy is also deposited in the absorbing (passive) layers, thus it is better to have a larger fraction of active material for high precision measurements, so long as containment is not sacrificed. Typical sampling fractions for electromagnetic sampling calorimeters are at the level of 10%, such as the CDF lead-scintillator calorimeter with $f_{\text{samp}} = 12\%$ [37]. The ATLAS Collaboration decided that a higher-quality electromagnetic barrel calorimeter was worth the increased expense, and thus the sampling fraction is 18% [38].

Hadronic showers, on the other hand, are much more complicated. Approximately 1/3 of hadronic interactions produce a π^0 meson, which then almost always decays to two photons, thus adding an electromagnetic component to the hadronic shower. The remaining 2/3 of hadronic interactions produce additional hadronic activity. This can be the production of charged mesons, such as π^\pm , or the fragments of nuclei, such as protons. The fraction of each process depends on the energy and target material, but these two processes account for approximately 20 and 30% of the hadronic interactions respectively. Another approximately 10% goes into neutrons, while the remaining 40% results in invisible processes [39]. This is primarily due to the loss of nuclear binding energy when an atom is fragmented through a strong interaction, although other sources such as nuclear recoil also contribute.

In other words, while the vast majority of energy in electromagnetic showers is captured, a significant fraction of hadronic shower energy is undetectable. Some types of hadronic calorimeters have a means of correcting, or compensating, for this effect. The ATLAS calorimeter, however, is non-compensating, meaning that the response ratio of $e/\pi < 1$. Furthermore, the fraction of energy which is observed for contained hadronic showers is dependent on the energy, meaning that the response ratio of e/π is also non-linear. An explanation of this energy dependence, and the way in which hadronic systems are calibrated to account for these reconstruction differences, is presented in Section 6.3.

The quantity that describes the strength of the interaction between the detector material and the shower also varies between hadronic and electromagnetic showers, as is expected given that different forces are at work. The property which governs the average distance of material which must be traversed to reduce the energy of an electron to $1/e$ of the starting value (and $7/9$ of the distance for the average high-energy photon to produce an electron-positron pair) is called the radiation length (X_0). This is typically specified in units of g/cm^2 , where dividing by the density of the material gives a radiation length in cm. Typical electromagnetic calorimeters work with enough material to provide more than 20 radiation lengths of material in order to fully contain the electromagnetic showers.

Hadronic showers, on the other hand, are governed by the strong force. Atoms are mostly empty space, and thus a given particle can be expected to travel much further through a typical material before undergoing an hadronic interaction as compared to an electromagnetic interaction. The average distance travelled by a hadron before undergoing an inelastic collision with the surrounding material, also known as the nuclear interaction length (λ_I), should thus be much larger than X_0 . This is observed, with the nuclear interaction lengths for typically used materials 5 to 30 times longer than the corresponding radiation length [12]. Similarly to radiation lengths, this is typically quoted in g/cm^2 , thus dividing by the material density gives a nuclear interaction length in cm. Calorimeters typically have closer to 10 nuclear interaction lengths of active material, which contains the majority of hadronic showers.

This difference in the average distance a particle will travel before interacting makes it relatively easy to fully contain all electromagnetic showers up to reasonably high energies, while it is essentially impossible to do the same for hadronic showers. When a shower passes through the entirety of the




Figure 4.10: The number of interaction lengths X_0 of material present in the barrel (left) and endcap (right) electromagnetic calorimeters. The crack region where the barrel accordion meets the endcap accordion is not shown. [28]

calorimeter, this also reduces the calorimeter response, as a potentially large fraction of the energy is lost. This effect, known as punch-through, becomes more frequent as the energy scale of the shower is increased. This phenomenon and preliminary corrections for such scenarios is discussed in Appendix F, where it is seen that punch-through becomes an important effect at high energies.

Note that while electromagnetic and hadronic showers have different properties, hadronic showers still show up in the electromagnetic calorimeters. The division between electromagnetic and hadronic calorimeters refers to the type of shower that they are intended to fully contain, not the type of shower which can be present in the detector.

4.4.2 Electromagnetic calorimeters

The most forward of the electromagnetic calorimeters will be discussed in Section 4.4.4, as it is substantially different from the other electromagnetic calorimeters discussed here.

Despite being broken up into three separate sub-detectors, the ATLAS electromagnetic calorimetry has a predominantly unified form. All of these detectors make use of Liquid Argon (LAr) as the active material, where LAr was chosen due to its uniform nature, stability, and radiation-hardness. The passive material was chosen to be lead, which has a short radiation length of $X_0 = 0.561$ cm [12], thus enhancing the level of electromagnetic shower containment within the available space. Throughout the detectors, the typical number of radiation lengths of calorimeter material is typically larger than 22, as shown in Figure 4.10.

As particles pass through the calorimeters, they ionize the LAr. A high voltage is applied to the plates enclosing the LAr, causing the electrons to drift to the copper electrodes within a time of approximately 450 ns. The uniformity of the LAr results in a clean triangular signal shape, where the majority of the signal arrives within a short period before decreasing to zero linearly. Waiting for the full signal to arrive is not feasible, as 450 ns spans many bunch crossings. Instead, a signal shaper is applied to the raw triangular distributions, creating a bipolar pulse signal with a much shorter readout time. The shaper was designed to match the LHC design bunch spacing of 25 ns, thus the calorimeter performance is expected to improve in 2015 as the bunch spacing is switched from 50 to 25 ns. In order to ensure that the pulse is optimal for the data taking conditions, the constants used in the shaping function




Figure 4.11: A view of a section of the electromagnetic barrel calorimeter, displaying the accordion geometry structure and detector segmentation. The three layers of calorimeter cells and their corresponding sizes and radiation lengths X_0 are shown. [28]

are frequently updated in dedicated calibration runs. These calibration runs determine the optimal amplitude, noise level (primarily pileup activity), and pulse shape [40].

In order to not bias future events where the charge is still being collected, the pulse shape is set to have a large negative tail after the peak, with the integral of the full pulse constructed to be zero. This tail balances the residual charge being collected, thus mitigating the contribution of residual energy from a given interaction affecting future events. Fluctuations in the pileup activity render this an average rather than absolute correction, as cells can have either a positive or negative energy when read out at a given time due to such fluctuations. Further suppression of pileup related effects in the calorimeter are discussed in Section 6.2.1 and Appendix D.1.

The barrel and endcap sub-detectors both make use of a unique design, by which the various sampling and active layers are aligned in an accordion geometry. This geometry reduces the number of blind regions of the calorimeter, and provides full ϕ coverage, which is an improvement over previous calorimeter geometries. The accordion structure in use in the electromagnetic barrel can be seen in Figure 4.11 along with the segmentation of the detector.

Electromagnetic barrel calorimeter

The ElectroMagnetic Barrel (EMB) calorimeter, as shown in Figure 4.11, consists of three layers with differing granularities and depths, covering $|\eta| < 1.475$. The first layer of the EMB is special in that it provides extremely fine segmentation in the η dimension, with strip cells of 0.0031×0.098 in (η, ϕ) -space. This is useful for improving the η measurement precision, but is critical for discriminating photons from $\pi^0 \rightarrow \gamma\gamma$ decays. The second layer captures the large majority of the energy in the calorimeter, and thus a high level of segmentation is needed in both η and ϕ to achieve optimal resolution, resulting in cells

of 0.025×0.0245 in (η, ϕ) -space. The final layer is much more coarse, and is intended to collect only the last parts of the electromagnetic shower, and thus high precision measurements are less critical. As such, the third layer is 0.05×0.0245 in (η, ϕ) -space. However, having a separate third layer rather than just extending the second layer has other uses, such as providing useful information on the nature of the shower, as exploited in Appendix D.4.

Electromagnetic endcap calorimeter

Two coaxial wheels form the ElectroMagnetic EndCap (EMEC) calorimeter, providing coverage for the range of $1.375 < |\eta| < 3.2$. The EMEC has largely the same structure as the EMB, just geometrically rotated into the radial rather than axial direction, and with the internal accordion geometry rearranged for this orientation. The transition region between the EMB and the EMEC, within $1.375 < |\eta| < 1.52$, is traditionally known as the *crack* region. This region contains a significant quantity of material which provides necessary services to the inner detector. These services amount to several radiation lengths of inactive material, and thus are a source of energy loss and reduce performance. Analyses which require high precision electrons or photons typically remove this region from their analysis selection.

Electromagnetic presamplers

To account for energy lost in material before the electromagnetic calorimeters, finely segmented LAr presamplers are placed in front of other material whenever possible, covering $|\eta| < 1.7$. This includes a presampler before the service material mentioned in the *crack* region, thus reducing but not removing the impact.

4.4.3 Hadronic calorimeters

The most forward of the hadronic calorimeters will be discussed in Section 4.4.4, as it is substantially different from the other hadronic calorimeters discussed here.

Unlike the electromagnetic calorimeters, two radically different technologies are used for the three sub-detectors. The barrel and extended barrel calorimeters share one common type of technology, while the technology used for the endcap calorimeter is closer to what was used for its electromagnetic counterpart. The number of interaction lengths of calorimetry of either type varies throughout the ATLAS detector, from approximately 7 to 16, as shown in Figure 4.12. This non-uniformity and lower value of λ_I compared to X_0 leads to many more situations in which hadronic showers punch-through the full calorimeter, as is studied in Appendix F.

Barrel and extended barrel hadronic calorimeters

The barrel and extended barrel segments of the hadronic calorimeter are known as the Tile Barrel and Tile Extended Barrel due to their use of plastic scintillating tiles as the active medium. The absorbing layers for these calorimeters are comprised of steel, which has a moderate nuclear interaction length of $\lambda_I = 16.8$ cm [12]. The scintillating tiles are arranged radially in a fashion perpendicular to the beamline. The barrel covers out to $|\eta| < 0.9$, while the extended barrel covers $0.8 < |\eta| < 1.7$, as shown in Figure 4.13.

As charged particles traverse the scintillating tiles, photons are produced in proportion to the amount of energy deposited. These photons are then collected with wavelength-shifting fibres within each




Figure 4.12: The number of interaction lengths of material in the ATLAS detector as a function of the pseudorapidity $|\eta|$. The unlabelled cyan region is additional material after the calorimeters and before the muon spectrometer, and is only shown up to $|\eta| = 3.0$. The unlabelled gold area represents additional material before the calorimeters. [28]




Figure 4.13: A technical diagram of the Tile barrel (left) and Tile extended barrel (right) hadronic calorimeters, showing the module placement and pseudorapidity coverage. [28]




Figure 4.14: A technical diagram of the Hadronic EndCap calorimeter, showing the flat copper plates and the larger number of plates in the first layer with respect to the second. [28]

calorimeter cell. Rows of cells are grouped together in η to form a module, which is then read out through PhotoMultiplier Tubes (PMTs) to measure the deposited signal. This procedure means that an entire row of cells is lost if any problems occur within the read-out services. The impact of such failures is discussed in Appendix C.

The scintillator and steel layer design was chosen to keep the cost reasonable, a necessary requirement given that the detector is more than 10 m long. The detector consists of three layers in both the barrel and extended barrel, with additional sensors in the transition region between the two sub-detectors to partially recover the energy lost in the crack. The Tile calorimeter has more coarse granularity than the electromagnetic calorimeters, with cells of $0.1 \times \pi/32$ for the first layers in (η, ϕ) -space. This provides suitable precision for hadronic showers, as they are much less compact than electromagnetic showers, and thus higher granularity is less important.

Hadronic endcap calorimeter

The Hadronic End Cap (HEC) calorimeter technology is similar to that of the EMEC, with an active medium of LAr, but with a passive absorption medium of copper rather than lead. The geometrical arrangement of the layers also differs, where a simple design of flat copper plates is used rather than the accordion geometry used in the EMEC, as shown in Figure 4.14.

The HEC is comprised of two wheels, each with two layers, covering the range of $1.5 < |\eta| < 3.2$. The first wheel has a higher granularity with 24 copper plates per module, while the second wheel uses only 16. Individual cells are roughly 0.1×0.1 in (η, ϕ) space for $|\eta| < 2.5$ and 0.2×0.2 for the rest of the detector.




Figure 4.15: A diagram showing the placement of the forward calorimeter with respect to the surrounding endcap calorimeters and other material. [28]

4.4.4 Forward calorimeters

The Forward CALorimeter (FCal) is substantially different from the rest of the ATLAS detector. While the inner detector and muon systems both have barrel and endcap detectors, only the calorimeters have a forward detector, which covers $3.1 < |\eta| < 4.9$. Placed 4.7 m from the interaction point, the FCal is placed in close contact with the endcap calorimeters, thus lessening the impact of the transition between sub-detectors. A diagram of the placement of the FCal and surrounding endcap material is provided in Figure 4.15.

The FCal consists of three separate layers, of which the first is an electromagnetic calorimeter, and the remaining two are hadronic calorimeters. The active medium for all three is LAr, while the passive medium is copper for the first layer and tungsten for the latter two. This very forward region experiences extreme particle flux at very high energy, and thus must be extremely radiation-hard and use materials which can stop very energetic particles. Very dense materials must be used to maximize the number of interaction lengths, thus limiting both the transverse and longitudinal shower development. This partially motivates the use of copper rather than lead for the electromagnetic calorimeter (X_0 of 1.44 vs 0.561 cm), as copper has a shorter nuclear interaction length (λ_I of 15.3 cm vs 17.6 cm). However, the more important property of copper and the ultimate reason it was selected is its conductivity. The extreme particle flux in the forward region quickly heats the material, which quickly leads to problems if left unchecked. Using very conductive materials such as copper reduces the amount of cooling services required to counter this effect.

The second and third layers of the FCal are already sufficiently shielded by the first layer, and thus the cooling is less critical, albeit still important. At this point, the stronger requirement is the need to stop hadronic showers from exiting the calorimeter. For this reason, tungsten is used, as it has a very short nuclear interaction length of 9.94 cm.

This combination of materials results in 27.6 radiation lengths of electromagnetic calorimeter material, thus ensuring that no electrons or photons escape. Similarly, this configuration provides 9.94 nuclear interaction lengths of material in a very small space. This prevents all but the most energetic particles from escaping the FCal. In order to provide even more protection from escaping high-energy particles and shield the muon endcap region, a large uninstrumented brass plug is placed behind the FCal.

The very high particle flux can lead to ion buildup within the LAr. In order to avoid this, very




Figure 4.16: A computer-generated schematic representation of the ATLAS muon spectrometer, with the various sub-detectors indicated. [41]

narrow LAr gaps are used, thus reducing the drift time from the 450 ns in the EMB to 60 ns for the first module, and longer for the second and third (scaling with the LAr gap width). This is achieved by using gaps as small as $269\ \mu\text{m}$, $376\ \mu\text{m}$, and $508\ \mu\text{m}$ for the first, second, and third modules respectively.

4.5 Muon spectrometers

For particles to make it through the full depth of the calorimeters, they must either be non-interacting, a semi-stable MIP, or a shower that was not contained in the available material. The latter process is rare but important at high energies, and is discussed in Appendix F. In the first case, which includes particles such as neutrinos, nothing can reasonably be done to contain them. At typical LHC energies, the only SM example of the second case is the muon. Muons thus provide very clean signatures for both searches and measurements, and a set of detectors designed to identify and measure them has been developed.

The ATLAS muon spectrometer consists of four sub-detectors, two in the barrel and two in the endcap. One sub-detector for each region provides a coarse, quick measurement, while the second provides a high precision particle momentum measurement. The coarse detectors cover the range of $|\eta| < 2.4$, and provide the ability to trigger on events with muon candidates. The much more precise detectors cover the range of $|\eta| < 2.7$, but are not fast enough to use in a trigger system. In this way, the two sets of two sub-detectors fulfil complementary roles. The full muon system is integrated with a series of toroidal magnets, placed outside of the calorimeter systems, which provide the necessary magnetic field to measure particle momenta via their curvature. A schematic diagram of only the ATLAS muon system is provided in Figure 4.16. A technical diagram of the muon system is provided in Figure 4.17.

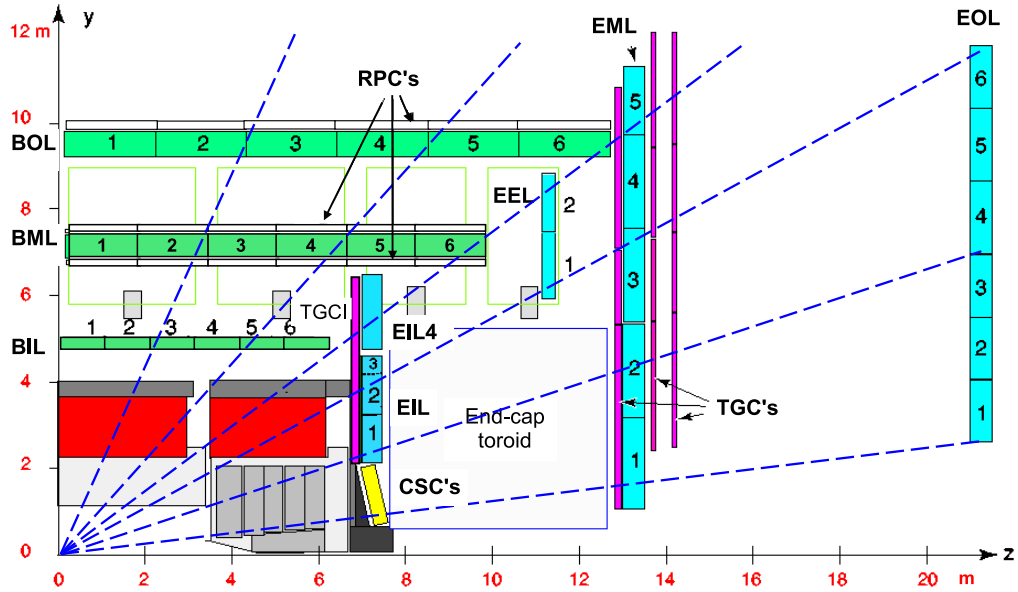


Figure 4.17: A quarter-sectional view of the ATLAS muon spectrometer, showing each of the sub-detectors and the geometrical space that they occupy. Note that the MDT sub-detectors are named after where they are located, rather than the detector type. As such, BOL/BML/BIL is the barrel outer/middle/inner layer of the MDT, and similarly for the endcap EOL/EML/EIL (EEL is extra). The other sub-detectors are named after the technology, namely the RPCs, TGCs, and CSCs. [28]

Resistive Plate Chambers and Thin Gap Chambers

The fast muon systems within the barrel region are made up of Resistive Plate Chambers (RPCs), while the endcap region is made of Thin Gap Chambers (TGCs). Both technologies produce a signal within 15-25 ns, which enables their use as triggering detectors. However, the cost of a very fast response appears in the measurement resolution. The RPCs (TGCs) provide an axial z (radial r) position resolution of 10 mm (2-6 mm) and an azimuthal ϕ resolution of 10 mm (3-7 mm).

The 606 RPCs are split into three layers, and provide a total of 373×10^3 channels. They cover the range of $|\eta| < 1.05$, and are located approximately 7 to 10 m from the interaction point along the radial direction. The 3588 TGCs are split into four layers, and consist of 318×10^3 channels. They cover $1.05 < |\eta| < 2.4$, and are located approximately 7, 13, 13.8, and 14.2 m from the the interaction point along the axial direction.

Monitored Drift Tubes and Cathode Strip Chambers

The slower but much more precise muon detectors are the Monitored Drift Tubes (MDTs). The MDTs are used in both the barrel and endcap regions, covering the full range of $|\eta| < 2.7$. The 1150 modules are split into three barrel layers and four endcap layers, forming a total of 354×10^3 channels. This structure provides an axial or radial (z or r) position resolution of $35 \mu\text{m}$. This is at least 57 times more precise than the trigger chambers, but the maximum MDT charge collection time is 700 ns, thus it is also 28 times slower. This demonstrates the benefit of having separate trigger and precision sub-detectors, which together can provide much more impressive performance than would be possible independently.

Cathode Strip Chambers (CSCs) have a much more specific purpose, and are used only for the innermost module of the innermost layer of the endcap region, where $2.0 < |\eta| < 2.7$. As such, there are only 32 modules, with a total of 30.6×10^4 read-out channels. They provide a similar radial resolution of $40 \mu\text{m}$, but with a maximum collection time of 40 ns, providing a time resolution of approximately 7 ns per plane. This reduced collection time is necessary due to the increased particle flux in the forward region, which would overwhelm MDTs if they were used.

4.6 Triggering system

The LHC bunch crossing rate as delivered to the ATLAS experiment for the 2012 dataset was every 50 ns, providing a collision frequency of 20 MHz. The full ATLAS detector cannot be read out and recorded at this frequency, necessitating the implementation of a triggering system to select only the most interesting events. When specifying the definition of what constitutes an interesting event, care must be taken to both cover all possible topologies required for physics analysis in addition to ensuring the trigger does not bias any particular process.

At a hadron collider, the relevant variable of interest in discriminating interesting inelastic collisions is the transverse momentum or transverse energy of objects in the event. As collisions become harder, the probability of a high transverse momentum object being observed grows, thus allowing for a selection of the most energetic processes which were beyond the reach of previous collider experiments. Typical objects considered are high p_T electrons, muons, taus, photons, jets, and Missing Transverse Momentum (E_T^{miss}). Definitions of all of these objects within ATLAS are provided in Chapter 5.

The probability for a given process to occur, and thus the frequency with which the trigger can be expected to observe such a process, varies dramatically between different final states. For example, the di-jet cross-section dominates all other hard-scatter processes by multiple orders of magnitude at hadron colliders, which in turn is a rare process compared to low-energy scattering events. As such, a naive definition of triggers based purely on the interesting object passing a specific p_T threshold will result in a dataset which is almost completely di-jet events. Defining a separate p_T threshold for each different object helps, but this is just the first step, as it also removes the sensitivity of the experiment to interesting physics which is below the threshold. In order to alleviate this problem, a *prescale* system is used. In this system, after the trigger decides that the event is interesting, a random number is used to decide whether or not to consider the event. This is governed by the prescale factor f_{prescale} , where $f_{\text{prescale}} = 100$ means that 1 in 100 events marked as interesting is recorded, and $f_{\text{prescale}} = 1$ is the special case of an unprescaled trigger.

The majority of ATLAS analyses make use of unprescaled triggers for simplicity, as they are purely interested in the high p_T regime and thus are willing to discard data from lower thresholds where more work is required to perform combinations of prescaled triggers. However, there are several very notable exceptions, such as the di-jet resonance analysis in Reference [42]. This analysis makes use of a combination of many prescaled triggers with f_{prescale} up to 460000 in order to extend its sensitivity to the lower di-jet mass (m_{jj}) regime. While this regime has already been covered by previous collider experiments, the enormous LHC luminosity allows for further improving the sensitivity at the lower energy regime. This is very important, as it is possible that new physics exists in the low energy regime, but that it is a very rare process, and thus it would otherwise be missed by relying purely on unprescaled triggers.

The ATLAS trigger system is comprised of three levels named Level 1 (L1), Level 2 (L2), and Event Filter (EF), and reduces the event rate through both prescaled and unprescaled triggers by a factor of 10^5 from 20 MHz to 200 Hz. In order to maintain this reduction factor, the exact trigger combination, thresholds, and prescale factors must all be updated on a regular basis as the LHC data conditions change, particularly when the centre of mass energy is increased.

If the trigger was infinitely fast, then a reduction of the rate from 20 MHz to 200 Hz would be a simple process. However, the trigger system must first wait for the detectors to complete their read-out of the event, and then must take time to build up a picture of the collision from the raw signals. Only once this is done can it make a decision about whether or not the event is interesting. This process takes much longer than the time between collisions, and thus substantial storage buffers are used. This allows for retention of the complete detector read-out for many independent collisions, and thus is a fundamental component of the trigger system.

The first level of the trigger requires some level of activity in either the calorimeters or muon system, as these are the only parts of the detector which can perform a simple reconstruction and trigger decision within $2.5 \mu\text{s}$ on an event-by-event basis, with the aim of reducing the event rate from 20 MHz to 75 kHz. The L1 system makes use of simple calorimeter towers (sums of energy of geometrical groups of calorimeter cells) for electron/photon and jet reconstruction, while muons make use of the trigger-specific chambers. If the L1 trigger decides an event was interesting, it defines one or more Regions of Interest (RoIs) where the η and ϕ coordinates of the interesting object(s) are located, and which specifies the specific trigger requirement that they passed in order to fire the trigger.

This information is then sent to the High-Level Trigger (HLT), which includes both the L2 and EF steps. The L2 trigger matches inner detector information to the RoIs found by the L1 trigger, and makes a decision about whether the event still passes the appropriate trigger criteria within 40 ms, with the intent of reducing the event rate from 75 kHz to 3.5 kHz. Successful events are then passed to the EF, which must reduce the trigger rate to the final level of 200 Hz. This is done by performing a full event reconstruction very similar to what is used for offline analysis, thus including calibrations, alignment corrections, and advanced algorithms. If an event is marked as interesting by the EF, it is recorded to disk in the RAW data format, which includes the information on all of the digitized energy deposits throughout the full ATLAS detector.

It is worth noting that the above description applies to the majority of triggers intended for use in physics analysis. There are also calibration triggers, minimum bias triggers, and other types of triggers which are typically highly prescaled, and select very different types of events. These types of triggers are not discussed further.

4.7 Simulation

The production of simulated events is a major requirement for ATLAS, given that almost all analyses make heavy use of Monte Carlo (MC) samples. MC can be used for many different purposes, ranging from estimating SM backgrounds to probing different signal hypotheses. MC can also be used to determine calibrations, such as for jets where a substantial amount of energy is not visible in the detector, but the MC retains information on the energy of the original particles. Providing a robust and efficient MC framework is thus a priority, and it is important that this simulation match data as closely as possible.

The ATLAS simulation framework is detailed in Reference [43], and makes use of the GEANT4 toolkit

from Reference [44] for detector modelling. There are multiple steps in the generation processes, splitting the truth level step from the interaction between the truth particles and the detector. This allows for filters to be applied after the truth level generation, thus enabling the selection of interesting events before performing the very computationally expensive digitization step.

4.7.1 Event generation

The details of this step depend on the MC generator(s) used. In some cases, as will be discussed in Appendix I.1, the matrix element calculation and the parton showering are split into separate steps which are each handled in an independent generator. In other cases, both steps are performed within a single generator. The end product regardless of the specific generator choice is a set of truth-level events for a given process of interest, where the truth-level events are independent of the detector.

ATLAS makes use of many different Parton Distribution Function (PDF) sets to model the parton distributions within the colliding protons. This becomes particularly important for very high energy processes, where the LHC is operating in a new regime of parameter space. The Leading Order (LO) CTEQ6L1 [45] and MSTW2008LO [46] are particularly common for signal samples, while SM backgrounds are more likely to make use of the Next to Leading Order (NLO) CT10NLO [47] PDF set.

MC generators in ATLAS do not always provide a good description of data if taken as-is. In order to provide a better MC description, the parton shower model used by MC generator is often tuned, such as is detailed for *Pythia6* [48] and *Pythia8* [49] in Reference [50]. In addition to tuning the specific generator showering model, the underlying event model used in the simulation is also adjusted. These tunes are typically derived with respect to a given PDF set, with CTEQ6L1 and CT10NLO again being the most common choices.

4.7.2 Detector simulation and digitization

This step propagates the stable truth particles produced in the previous step through a complete **GEANT4** model of the ATLAS detector, where the current best understanding of interactions between relativistic particles and material is used. This includes the creation of electromagnetic and hadronic cascades from stable truth particles, and their interactions within the calorimeters. Electromagnetic showers are generally reasonably well modelled, but hadronic showers remain a topic of active investigation. As such, there are currently two showering algorithms widely used, with **QGSP_BERT** as the ATLAS default and **FTFP_BERT** as the alternative often used for systematic variation studies. The first is the quark gluon string model of nuclear fragmentation, described in Reference [51], while the second is the **FRITIOF** model, as per Reference [52]. In both cases, the Bertini cascade model is used, which is detailed in Reference [53].

The full **GEANT4** simulation is used wherever possible for MC production in ATLAS. However, this full simulation is very computationally expensive, especially when computing showers, and thus a fast simulation parametrization known as **AFII** is often used where the highest level of precision is not necessary. This framework parametrizes the average **GEANT4** response to a given type of particle in a given kinematic bin within the various calorimeters, thus providing look-up tables for energy deposits and interaction probabilities rather than performing full calculations for each step. The inner detector and muon spectrometer simulations are unchanged, and continue to use the **GEANT4** model. This reduces the required computational time by an order of magnitude, at the cost of lower calorimeter precision.

Many ATLAS analyses make use of MC created by a combination of both the full `GEANT4` simulation and the fast `AFII` parametrization.

The hard process dealt with so far is independent of pileup, which is a significant difference from real collision events. This is addressed by overlaying a set of minimum bias events created with a specially constructed `Pythia8` LHC tune. These simulated minimum bias events add a pileup contribution, but do not encapsulate all of the possible detector effects, such as cavern background contributions to the muon spectrometer. Studies into the use of overlaid zero bias events from data are ongoing, as shown in Reference [54], but have not yet become the standard. An example of the difference between overlaid MC minimum bias and data zero bias events is shown for muon variables in Figure F.1.

All of the particles, whether from the hard-scatter vertex or from the overlaid interactions, are passed through a complete model of the detector. Energy deposits (hits) are then assigned to each active portion of the detector which is affected. The digitization process described in Reference [54] is then applied to convert the simulated energy deposits into electronic signals which would be observed in the actual detector. This is then reconstructed using the same procedure as data, giving rise to the Raw Digital Object (RDO) format, which primarily differs from the RAW format simply by the presence of additional truth information.

4.8 Computing and data management

The enormous amount of data produced either from LHC collisions or MC simulations requires a high-performance flexible distributed computing system. This requirement is met through the implementation of the Worldwide LHC Computing Grid (WLCG), which is composed of four levels of nodes, called Tiers. All of the events are initially recorded and reconstructed at the CERN Tier 0 (T0) data centre, which contains approximately 20% of the computing power of the full WLCG.

The raw datasets and first reconstruction are then distributed around the world to a small number of Tier 1 (T1) nodes, such as TRIUMF in Canada, which represents 5% of the full grid computing power. These datasets in turn can be requested by Tier 2 (T2) computing centres, which are typically hosted by universities or groups, and which provide additional computing resources as needed. The final level of Tier 3 (T3) services have no defined structure within the WLCG, and can vary from local clusters of computers to a single laptop connecting to the WLCG to download data or make processing requests.

Within ATLAS, there are a few main data formats in use, which facilitate easier management. The RAW and RDO formats are essentially restricted to T0 and T1 sites, with very little additional access available due to the very large computing resources and very complicated processing interface required to reconstruct events from raw events. The next level is the Event Summary Data (ESD), which contains all of the reconstructed data, and thus is a very large data format. These are typically not stored, but in some very specialized cases they are required, and thus the functionality exists for users to access and work with ESDs.

An additional processing step by which the majority of the individual cell information is discarded and other space saving steps are taken forms the Analysis Object Data (AOD). This is the intended format for primary use within ATLAS, where all of the necessary information for reconstruction of interesting objects remains. For example, jets or the E_T^{miss} can be properly reconstructed at this level, as their necessary building blocks remain. In practice, this format is rarely used, although the development of the xAOD for Run-II looks set to change this by addressing many of the limitations of the AOD

format.

Users who do work with ESDs or AODs will typically also apply some additional analysis selection to reduce the size, thus producing Derived Event Summary Datas (DESDs) or Derived Analysis Object Datas (DAODs). However, the most common data format in use in Run-I applies a third level of filtering, earning it the name of the Derived³ Physics Data (D3PD), in addition to flattening the data format. This means that rather than storing the data in objects, every variable is stored as a simple type, such as an integer, vector (list) of integers, or so on. This format typically is created on a per-analysis or per-working group level, and applies specific selection cuts, triggers, and so on to further reduce the size of the file. Many groups found this necessary in order to reduce datasets to a manageable level. Even with this amount of reduction, a typical analysis will work with many TeraBytes (TBs) of D3PDs.

Many private formats where even further levels of derivation or with redesigned information management systems exist, but they are not officially supported. The above list is the officially supported set of ATLAS data formats. These supported data formats are typically stored across T0, T1, and T2 sites, while private data formats are typically restricted to T3 clusters or local use.

Chapter 5

ATLAS Reconstruction and Performance

5.1 Introduction

When conducting an analysis, the first step is to define the topology that will be considered. This selection typically requires the presence or absence of certain objects, where an object is either a physical particle (electrons, muons, taus, photons), a representation of an underlying physical process (jets and b -jets), or the contribution of invisible processes creating an imbalance in the event (E_T^{miss}). As such, detector quantities such as cells, tracks, and clusters must be translated into objects, which can then be properly calibrated and used in analyses.

The mono-jet analysis discussed in Chapter 7 primarily depends on jets and missing transverse momentum, but also makes secondary use of electrons and muons. Jet reconstruction and performance is covered separately in Chapter 6, while the other objects are discussed in the following Sections. The reconstruction of taus, photons, and b -jets, which are not directly considered by the analysis, are only briefly mentioned.

5.2 Electrons and photons

5.2.1 Reconstruction

As mentioned in Section 4.4.1, electromagnetic showers produced by electrons and photons are not distinct processes. Brehmsstrahlung leads to electrons emitting photons, while pair-production leads to photons converting into electron-positron pairs. As such, the reconstruction of electron and photon energy deposits within the calorimeter follows the same procedure. The primary difference in tagging electrons versus photons comes from the tracker, where a photon is neutral and leaves no tracks whereas electrons are charged. Further electron identification power by rejecting the contribution of other particles such as π^\pm mesons comes from the TRT, in which electrons emit X-rays, explained in Section 4.3.

The first step in reconstructing an electron or photon is the construction of clusters of calorimeter energy deposits. Electrons and photons make use of fixed-size clusters built with 3×5 cells, or equivalently 0.075×0.1225 in angular units with respect to (η, ϕ) -space, where the size is fixed by the largest (middle)

electromagnetic calorimeter layer [55]. Despite the size being fixed by the central layer, all three layers are then summed together to get the total transverse energy deposit within this window. The window position is then moved until the transverse energy is a maximum, which is known as the *sliding-window* algorithm [56]. If the resulting window transverse energy sum is above 2.5 GeV, then the region is marked as a seed, where this threshold was chosen to optimize reconstruction efficiency while minimizing fake rates from electronic or pileup noise [55].

Once seed clusters are found, their level of consistency with well-reconstructed inner detector tracks is checked. If such a match is not possible, then the cluster is tagged as an unconverted photon. If a match is possible to a vertex which is not at the interaction point (a secondary vertex), then the cluster is tagged as a converted photon. Finally, if a match is possible to a track from the hard-scatter vertex, then the cluster is tagged as an electron [55]. Of the photons which are identified within ATLAS, approximately 30 to 35% are converted photons [57, 58], thus it is important to distinguish converted photons from electrons if a specific topology is to be selected.

After classification of the cluster, electron or converted photon clusters are rebuilt with a size of 3×7 central layer cells in the barrel, while unconverted photons are rebuilt with 3×5 central layer cells to account for their smaller lateral size [55]. In the endcap region, all three types of objects are rebuilt using a 5×5 central layer cell cluster.

Electromagnetic showers matched to no track or to a secondary vertex provides a very strong criterion, and is the primary means of identifying photons. Other neutral particles which are expected to decay to electrons are rare, thus only minimal additional discrimination is required, such as shower profiles as detailed in Reference [55]. In addition to identifying photons, it is also important to determine if the photon is from the hard process. While such events are rare at the LHC compared to the $\pi^0 \rightarrow \gamma\gamma$ background, they can be identified thanks to the extremely narrow η granularity of the first layer of the electromagnetic calorimeter.

Electrons have more backgrounds, as there are many other charged particles which have similar properties. Charged pions are one example, which are suppressed by the use of additional TRT information as mentioned in Section 4.3. In general, it is always possible for jets to fake an electron signature. The probability is relatively low for a given jet to be both very collimated and also contained within the electromagnetic calorimeter, but the enormous cross-section of the di-jet process means that this is a relevant background. In order to suppress fake electrons, three sets of criteria are defined, named *loose++*, *medium++*, and *tight++*, detailed in Reference [59]. Depending on the fake rate that analyses are willing to tolerate compared to the loss of real electrons, a decision can be made between the three options.

5.2.2 Performance

Once reconstructed, electrons and photons are calibrated through a series of steps, as outlined in Figure 5.1. The calibration proceeds separately for data and MC, with some common corrections, and is split into six steps [55].

1. MultiVariate Analysis (MVA) techniques are used to calibrate the electron and photon energy in MC. The MVA makes use of both tracking and cluster information, such as the longitudinal development of the shower.

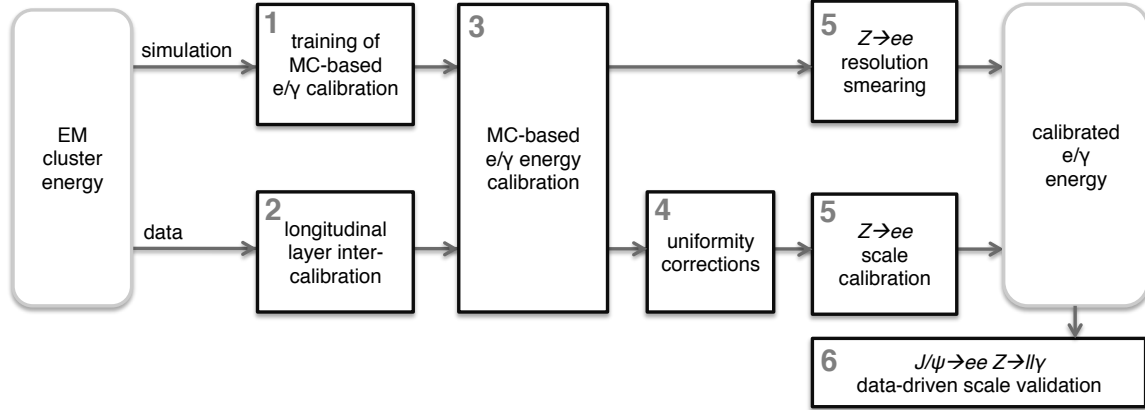


Figure 5.1: A schematic representation of the ATLAS electron and photon calibration chain in both data and MC. [55]

2. The energy scales of the three different layers are equalized (inter-calibrated) in data with respect to simulation.
3. A MC based electron/photon energy scale is applied to both simulated events and data, adjusting the reconstructed scale to the truth scale.
4. Corrections to enforce uniformity across different detector regions are applied to data.
5. *In situ* measurements of electron response in the $Z \rightarrow ee$ topology are used to derive per-electron scale factors, which are then applied to electrons and photons in data. For MC, the same topology is used to derive smearing factors to account for the small observed difference in resolution. These remaining differences are larger for low p_T electrons and photons, particularly near the transition region between the barrel and endcap calorimeters ($|\eta| \approx 1.5$), where it can be as large as 10% for electrons, 8% for converted photons, and 5% for unconverted photons. For more energetic objects or other calorimeter regions, the resolution is typically within the 2-4% range [55].
6. Only for validation, not calibration, the energy scale after the previous step is checked *in situ* with $J/\psi \rightarrow ee$ and $Z \rightarrow \ell\ell\gamma$ events.

As a final step, the cluster energy measurement and track momentum measurement for electrons are combined, which significantly improves the resolution at low energy where the tracker has better resolution, and has less of an impact at higher energies where the calorimeter has better resolution. The resulting performance of the electron calibration purely from clusters as opposed to a combined measurement in reconstructing the J/ψ and Z mass peaks is shown in Figure 5.2.

5.3 Muons

5.3.1 Reconstruction

Muon reconstruction makes extensive use of both the inner detector and muon spectrometers, with contributions from the calorimetry in regions where the spectrometers are not available. The three layers of high precision MDT muon spectrometers each provide six to eight η measurements for a single




Figure 5.2: The reconstruction of the J/ψ (left) and Z (right) mass peaks after electron calibration, showing the improvement gained when both tracking and calorimeter information are combined, particularly in the low energy regime. [55]

muon passing through the detector within $|\eta| < 2.7$, where the coarser muon trigger chambers provide a ϕ measurement. Hits within each layer of the spectrometers are first combined to form local track segments, then the local track segments from each layer are combined to form a global muon spectrometer track [60].

The inner detector provides an independent measurement of the muon close to the interaction point, and with very high precision spatial and momentum resolution due to the high precision available within the inner detector, so long as the muon is within $|\eta| < 2.5$. A typical muon track within the acceptance of the inner detector has 3 pixel hits, 8 SCT hits, and 30 TRT hits (when within $|\eta| < 1.9$) [60]. This allows for additional quality criteria when selecting inner detector tracks to match to muons, where at least 1 pixel hit, 5 SCT hits, and 9 TRT hits (when within $|\eta| < 1.9$) are required. In addition, the muon is only allowed to traverse two pixel or SCT sensors without an associated hit. This criteria provides a high-quality track, which can then be extrapolated to the muon spectrometer and matched to the independently reconstructed muon track.

Using this information, four types of muons are defined, listed in order of decreasing applicability to the typical ATLAS analysis [60]:

1. ComBined (CB): the inner detector and muon spectrometer tracks both exist, and a good match is found between the two, thus providing for a high quality muon candidate. The combination of independent track measurements provides a better resolution than possible independently, and allows for rejection of in-flight decays producing muons. These are the purest variant of muon candidate, and typically the only type considered for analyses requiring precise muon information.
2. Segment Tagged (ST): the inner detector track exists and matches with track segments in the muon chamber(s), but an independent muon track in the muon spectrometer was not reconstructed. This is primarily useful to recover acceptance in rare areas where only one layer of muon chambers is accessible, or where the muon is very low p_T .
3. Stand Alone (SA): the muon spectrometer track exists, but there is no inner detector track with which to match it. This is mostly applicable for the region of $2.5 < |\eta| < 2.7$, where the muon spectrometer exists but the inner detector does not.

4. Calorimeter Tagged (CaloTag): the muon spectrometer track does not exist, but the inner detector track does, and there is an energy deposit in the calorimeter matched to the inner detector track which is consistent with the presence of a MIP. This type of muon has a low purity, and exists primarily for holes in the muon system, most prominently $|\eta| < 0.1$.

It should be noted that there are three separate algorithms for the reconstruction and combination of inner detector and muon spectrometer information. The mono-jet analysis described in Chapter 7 makes use of “first chain” or STACO (STATistical COmbination) muons, and thus this is what has been described. This approach performs a statistical combination of the inner detector and muon spectrometer tracks using the corresponding covariance matrices. Other methods are similar with some variations, such as performing a global track fit to inner detector and muon spectrometer hits rather than using pre-reconstructed tracks from the two detectors.

5.3.2 Performance

Once reconstructed, the muon p_T scale and resolution are calibrated by selecting di-muon events with an invariant mass matching the J/ψ , Υ , and Z mass peaks, thus obtaining a reference object with well-known properties. The resulting muons are used to derive scale factors to ensure agreement between data and MC. The derived correction factor from a given selection to a given detector type (either the muon spectrometer or inner detector) is given in Equation 5.1 [60], where g_m represents normally distributed random variables with a mean of 0 and a width of 1.

$$p_T^{\text{Corr,det}} = \frac{p_T^{\text{MC,det}} + \sum_{n=0}^1 s_n^{\text{det}}(\eta, \phi) \left(p_T^{\text{MC,det}}\right)^n}{1 + \sum_{m=0}^2 \Delta r_m^{\text{det}}(\eta, \phi) \left(p_T^{\text{MC,det}}\right)^{m-1} g_m} \quad s_0^{\text{ID}} = 0 \text{ and } \Delta r_0^{\text{ID}} = 0 \quad (5.1)$$

The extracted momentum resolution smearing factor s_n^{det} and momentum scale correction factor Δr_m^{det} are extracted from the three control regions in bins of η and ϕ . These distributions, both before and after MC is corrected to data, are shown for the J/ψ , Υ , and Z regions in Figure 5.3.

5.4 Tau leptons and b-jets

Neither taus nor b -jets are directly considered by the mono-jet analysis. As such, only a very quick overview is provided.

5.4.1 Tau lepton reconstruction and performance

Tau leptons, with a mass of 1.78 GeV, are the only leptons which are heavy enough to decay into hadrons. This possibility accounts for just under 65% of tau decays, providing a very different behaviour than is observed for the electron and muon. Furthermore, given the very rapid decay rate, the ATLAS detector typically cannot distinguish between electrons/muons from tau decays and the same leptons from other direct production processes. As such, only hadronic tau decays necessitate extra attention. These hadronic decays can be classified into two categories, one-pronged or three-pronged, based on the




Figure 5.3: The reconstruction of the J/ψ (top left), Υ (top right), and Z (bottom) mass peaks after muon calibration. [60]

Decay mode	Decay type	Branching ratio [%]
$\tau^\pm \rightarrow e^\pm \nu_e^\bullet \nu_\tau^\circ$	Leptonic	17.83
$\tau^\pm \rightarrow \mu^\pm \nu_\mu^\bullet \nu_\tau^\circ$	Leptonic	17.41
Sub-total	Leptonic	35.24
$\tau^\pm \rightarrow \pi^\pm \pi^0 \nu_\tau^\circ$	1-prong hadronic	25.52
$\tau^\pm \rightarrow \pi^\pm \nu_\tau^\circ$	1-prong hadronic	10.83
$\tau^\pm \rightarrow \pi^\pm \pi^0 \pi^0 \nu_\tau^\circ$	1-prong hadronic	9.30
$\tau^\pm \rightarrow \pi^\pm \pi^0 \pi^0 \pi^0 \nu_\tau^\circ$	1-prong hadronic	1.05
Sub-total	1-prong hadronic	46.7
$\tau^\pm \rightarrow \pi^\pm \pi^\pm \pi^\mp \nu_\tau^\circ$	3-prong hadronic	9.02
$\tau^\pm \rightarrow \pi^\pm \pi^\pm \pi^\mp \pi^0 \nu_\tau^\circ$	3-prong hadronic	2.70
Sub-total	3-prong hadronic	11.72
Total	All shown	93.66

Table 5.1: The dominant tau decay modes, classified in leptonic, 1-prong hadronic, and 3-prong hadronic categories. The neutrino notation relative to τ^\pm is $\nu^\circ = \bar{\nu}/\nu$ and $\nu^\bullet = \nu/\bar{\nu}$. [12]

number of charged pions produced. Note that more exotic decays such as five-pronged events do exist, but are extremely rare, and thus are not considered further.

The dominant tau decay modes and their respective probabilities are summarized in Table 5.1. This demonstrates that approximately 1 in 5 hadronic tau decays is three-pronged, while the remainder are one-pronged, although this number varies slightly when other processes are added. This single-prong structure is exploited in Section 7.4.7.

Taus are thus reconstructed as jets within ATLAS using the anti- k_t $R = 0.4$ jet definition built from LC topo-clusters, where topo-clusters and jet reconstruction are explained in Section 6.2. MVA taggers allow for the identification of hadronic tau enriched datasets, although the efficiency is relatively low given the challenge of identification. Calibration of taus is applied via a Tau Energy Scale (TES), which is derived in MC. The $Z \rightarrow \tau\tau \rightarrow \mu\nu_\mu\nu_\tau\tau_{\text{had}}$ mass peak is used to both validate this calibration *in situ* and to derive uncertainties.

5.4.2 b -jet reconstruction and performance

The decay of a b -quark proceeds via the emission of a W boson to create a lighter quark (usually a charm). The W then decays either hadronically or leptonically. If the W decays hadronically, the resulting object is said to be hadronic, while if the W decays leptonically it is called semi-leptonic.

While b -jets are not directly considered in the mono-jet analysis, the mono- b topology is strongly related, as described in Reference [7]. Jets initiated by b -quarks are also useful for many other purposes within ATLAS, particularly where coupling to third generation quarks is expected to dominate over the lower generations. Discrimination of b -jets compared to jets initiated by light quarks or gluons is possible due to b -jets having different shower profiles and displaced vertices. This allows for sufficient discrimination to select datasets significantly enriched in b -jets, with an identification efficiency higher than for taus. This can be further refined for semi-leptonic b -jets due to the presence of the lepton.

The first step of b -jet reconstruction is to first find anti- k_t $R = 0.4$ jets built from LC topo-clusters, where topo-clusters and jet reconstruction are explained in Section 6.2. A MVA tagger is then used to select b -jets, with a di-leptonic control region used to calibrate the events [61]. Systematic uncertainties taken from nominal and varied MC samples.

5.5 Missing transverse momentum

Multiple variants of the E_T^{miss} exist, including both track-based and calorimeter-based calculations, in addition to methods which combine the two approaches. This Section will focus on the the use of pure calorimeter-based E_T^{miss} , as used in the mono-jet analysis described in Chapter 7.

5.5.1 Reconstruction

The missing transverse momentum, or E_T^{miss} , is an observable which quantifies the degree to which a given event is balanced in the transverse plane. In essence, the E_T^{miss} is the vectorial sum of the transverse energy from every significant energy deposit in the detector multiplied by -1, thus creating the vector required to balance the event. This allows for a means of selecting or rejecting invisible contributions, where energetic neutrinos or more exotic minimally interacting particles escape the ATLAS detector without leaving any direct signature of their presence.

The E_T^{miss} only exists in the transverse plane, which consists of the two dimensions x and y . As such, it can be built from its components directly as:

$$E_T^{\text{miss}} = \sqrt{(E_x^{\text{miss}})^2 + (E_y^{\text{miss}})^2} \quad (5.2)$$

$$\phi^{\text{miss}} = \arctan(E_y^{\text{miss}}/E_x^{\text{miss}}) \quad (5.3)$$

As the E_T^{miss} is a full event quantity, it involves measurements from all of the detectors, even for the calorimeter-based E_T^{miss} variant. The most striking example is for muons, which are MIPs and thus deposit only a tiny fraction of their energy in the calorimetry. As such, muons reconstructed with muon spectrometer and tracking information are added to the E_T^{miss} . Thus, the calorimeter based E_T^{miss} is defined as below, where i is the x or y component.

$$E_i^{\text{miss}} = E_i^{\text{miss,calo}} + E_i^{\text{miss},\mu} \quad (5.4)$$

There are several prescriptions for the calculation of the calorimeter terms. The default formulation, called `RefFinal` [62], identifies electrons, photons, taus, jets, muons, while the remainder is considered as soft activity in the form of topo-clusters (described in detail in Section 6.2.1). For each identified object, the most recent calibration is applied, thus ensuring that the objects are used at the best known scale possible. Note that a calorimeter muon term is also identified around regions of identified muon trajectories in order to prevent double-counting the muon contribution which was already computed. This results in the following form, where i is the x or y component of the E_T^{miss} .

$$E_i^{\text{miss,calo}} = E_i^{\text{miss},e} + E_i^{\text{miss},\gamma} + E_i^{\text{miss},\tau} + E_i^{\text{miss},\text{jet}} + E_i^{\text{miss,calo},\mu} + E_i^{\text{miss,SoftTerm}} \quad (5.5)$$

If all of the objects were considered as-is, then there would be significant biases in the E_T^{miss} due to double-counting of energy. In particular, note that electrons, photons, and taus are also reconstructed as jets, which is not desired given that they have already been properly identified and calibrated. As such, overlap removal must be performed to ensure that this dangerous double-counting between jets and other terms does not occur. This overlap removal is performed on AODs (ESDs) by assigning clusters (cells) to every calorimeter object, and associating each cluster (cell) with a given object where the association

order is conducted from left to right in Equation 5.5. The overlap removal at ESD level is not described further, despite it being a more rigorous approach, as it is rarely used within ATLAS due to the limited access to ESDs.

The most common example is the overlap of electrons and jets. This scenario will thus be used to explain the procedure. First, note that electrons are formed using fixed size *sliding-window* clusters, as discussed in Section 5.2.1. However, jets are constructed from topo-clusters, as explained in Chapter 6. For now, it is sufficient to know that topo-clusters are a different method for clustering energy within the calorimeter. As such, the first step in dealing with overlapping electrons and jets in the E_T^{miss} is to perform an association between the two different types of calorimeter clusters. In order to facilitate this association, cells near electrons are written to the AOD, which allows for a determination of the degree of energy overlap between a given topo-cluster and sliding window cluster. In all cases, the electron receives the full amount of energy from the sliding window cluster. However, the jet only receives the fraction of energy which is in the topo-cluster but not the sliding window cluster. Furthermore, if this fraction drops below 50% (meaning that more than half of the topo-cluster energy is included in the sliding window cluster), then the topo-cluster is given a weight of 0 within the jet and the residual energy is added to the SoftTerm as will be discussed later.

This procedure is repeated for all such overlaps down to the jet term. The calorimeter muon and SoftTerm are separate. A brief explanation of how each term is specifically defined follows, where a simple selection overview is provided. The general intent of each selection is to reduce the contribution of fakes without being so tight as to throw away sufficient numbers of good objects, as both extremes bias the E_T^{miss} scale.

Electron term

Electrons are required to pass the *medium++* quality criteria, as defined in Reference [59], and to have a $p_T > 10$ GeV. The difference between sliding window clusters and topo-clusters is computed from cell overlaps, as previously discussed.

Photon term

Photons are required to pass the *tight* quality criteria, as defined in Reference [59], and to have a $p_T > 10$ GeV. In the case that the photon was already included as an electron, the electron is retained and the photon discarded. The difference between sliding window clusters and topo-clusters is computed from cell overlaps, as previously discussed.

Hadronic tau term

Hadronic taus are required to pass a *medium* quality criteria, to be within the tracker acceptance, and to have a $p_T > 20$ GeV. In the case of overlap between electrons or photons and taus, the electron or photon is kept, the tau is discarded, and any remaining energy is moved to the SoftTerm. Taus which are retained are calibrated with the Tau Energy Scale.

Jet term

Anti- k_t $R = 0.4$ jets built from LC topo-clusters are used for the E_T^{miss} . Jets are required to be within $|\eta| < 4.5$ and to have a $p_T > 20$ GeV. If more than 50% of the energy of the jet is already accounted

for by other overlapping objects, the jet is moved to the SoftTerm. Otherwise, the jet is assigned scale factor between 0.5 and 1 associated with the fraction of energy in the jet which is not already accounted for in other objects.

Muon term

As previously mentioned, there are two muon terms, one within the calorimeter and the other external to it. So long as identified muons exist, the external term is always non-zero, while the calorimeter term requires more care. Muons are required to be *combined* within $|\eta| < 2.5$, while *stand alone* muons are allowed for $2.5 < |\eta| < 2.7$ to cover the full spectrometer reach as described in Section 5.3. For some regions where the muon spectrometer does not provide full coverage, *segment tagged* muons are also used.

The focus on *combined* muons significantly reduces the contribution of fake or mis-reconstructed objects. In some cases, the contribution from *stand alone* muons can cause significant biases due to mis-reconstruction, even when only the $2.5 < |\eta| < 2.7$ range is used, as was discovered in the mono-jet analysis. Once these muons have been defined, the external muon term is then the negative of the transverse component of their summed four-momentum.

The calorimeter muon term differs between muons which are isolated with respect to jets or not, where the isolation requirement is the presence of a jet within $\Delta R = \sqrt{(\eta_{\text{jet}} - \eta_{\mu})^2 + (\phi_{\text{jet}} - \phi_{\mu})^2} < 0.3$. Isolated muons are known to deposit an average of 2 or 3 GeV of energy in the calorimeters, thus a parametrization is derived and used to account for the average expected isolated muon contribution. The calorimeter contribution of non-isolated muons cannot be resolved from the overlapping calorimeter object(s), thus the muon energy has already been considered as a part of the sliding window or topo-cluster in a previous term. In such cases, the calorimeter muon term is zero.

SoftTerm

The SoftTerm is calculated through the use of an *eflow* algorithm. This algorithm exploits the lower p_T threshold and better low p_T resolution of tracks to both correct the scale of low p_T topo-clusters and account for low p_T particles which didn't make it to the calorimeter. First, a set of good quality tracks is selected down to a p_T of 400 MeV. Additional requirements are added to tracks of above 10 GeV, and any tracks of above 100 GeV are ignored. The tracks are then extrapolated to the middle layer of the electromagnetic calorimeters, selected due to its high-granularity, and the track association to a hard object or topo-cluster is checked through a ΔR matching as defined in Section 6.2.3.

All tracks which are not matched to any topo-cluster or reconstructed object are added to the *eflow* term, thus accounting for missed objects. If the track is matched to a topo-cluster, then the track is added and the topo-cluster is discarded, thus making use of the better tracking resolution at low p_T . If the track is matched to multiple topo-clusters, the track is added and only the highest energy topo-cluster is discarded. Finally, all remaining topo-clusters which have not been discarded by track matching are added to the *eflow* term. This stage includes the remnants of topo-clusters which were part of a jet, but which were removed due to energy overlaps of more than 50%, although these topo-clusters also have an associated energy overlap weight between 0 and 0.5.

Mono-jet definition

The definition provided above is for the standard **RefFinal** algorithm. The mono-jet analysis has made two changes to this algorithm to meet the analysis requirements. The first change is that the tau term has been removed, meaning that hadronic taus are instead collected by the jet term, and thus have the jet energy scale applied rather than the tau energy scale. This is done for coherence, given that the analysis does not identify taus. Treating taus as jets within the analysis but taus as taus within the E_T^{miss} , where the analysis applies cuts on both jets and E_T^{miss} , would be inconsistent and could introduce a bias.

Another change for the mono-jet analysis is the act of ignoring the external muon term. This was done for multiple reasons:

1. The analysis triggers on E_T^{miss} , where the E_T^{miss} trigger is purely calorimeter-based, and thus a E_T^{miss} without muons is closer to the trigger selection. This helps by reducing the amount of different datasets required, as per Section 7.5.
2. The treatment of muons as a source of E_T^{miss} rather than as part of the balance greatly simplifies the data-driven estimation when using $W \rightarrow \mu\nu$ and $Z \rightarrow \mu\mu$ control regions, as will be discussed in Section 7.6.
3. The use of *stand alone* muons in the E_T^{miss} for the region of $2.5 < |\eta| < 2.7$ was found to introduce significant shifts in the tail of the E_T^{miss} distribution from mis-reconstructed muons. Given that the mono-jet analysis focuses on extreme values of the E_T^{miss} , this was observed to introduce significant biases. Cleaning variables were found which successfully vetoed the majority of such events, but in the end it was decided it would be safer to remove the muons altogether.

As a result of these changes, the E_T^{miss} variant of the **RefFinal** algorithm used in the mono-jet analysis is defined as per Equation 5.6, where i is the x or y component. The magnitude and direction of the E_T^{miss} can then be calculated from the components as per Equations 5.2 and 5.3 respectively.

$$E_i^{\text{miss}} = E_i^{\text{miss},e} + E_i^{\text{miss},\gamma} + E_i^{\text{miss},\text{jet}} + E_i^{\text{miss},\text{calo},\mu} + E_i^{\text{miss},\text{SoftTerm}} \quad (5.6)$$

5.5.2 Performance

Due to the E_T^{miss} being the balance between any number of different objects which are present in a given event, the performance is challenging to quantify for all topologies. It is possible that the balance between object of type α and another object of type β (where α may or may not be the same as β) in a certain kinematic region is biased. Rather than attempt to cover every possible topology for all possible kinematic parameter space, the performance is instead studied in a few well-known topologies where systematic effects due to objects external to the E_T^{miss} can mostly be removed. The $Z \rightarrow \mu\mu + \text{jets}$ and $Z \rightarrow ee + \text{jets}$ processes are good examples, where the selection of the Z boson peak provides a very high purity selection of events where the truth E_T^{miss} is zero. While less pure, $W \rightarrow \mu\nu + \text{jets}$ and $W \rightarrow e\nu + \text{jets}$ provide regions where a source of true E_T^{miss} can be examined.

A first check is to consider the $Z \rightarrow \mu\mu + \text{jets}$ topology and to compare data to MC. This is shown for the full E_T^{miss} and the **SoftTerm** in Figure 5.4. Both distributions are observed to agree between data and MC to within 20%, although some clear trends are visible. Reference [63] shows that the decreasing trend in the **SoftTerm** mostly cancels the increasing trend in the jet term.

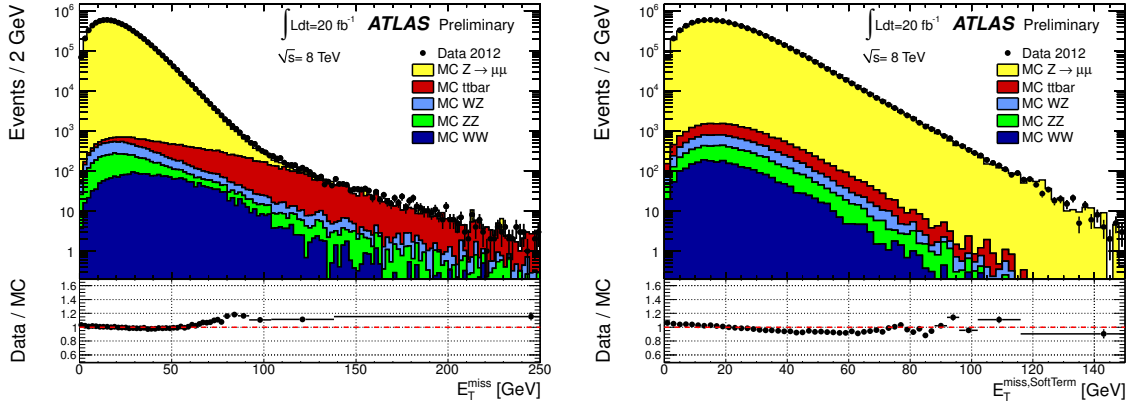


Figure 5.4: The agreement in the E_T^{miss} shapes between data and MC in the $Z \rightarrow \mu\mu + \text{jets}$ topology, shown for the full RefFinal variant of the E_T^{miss} (left) and only the SoftTerm $E_T^{\text{miss,SoftTerm}}$ (right). [63]

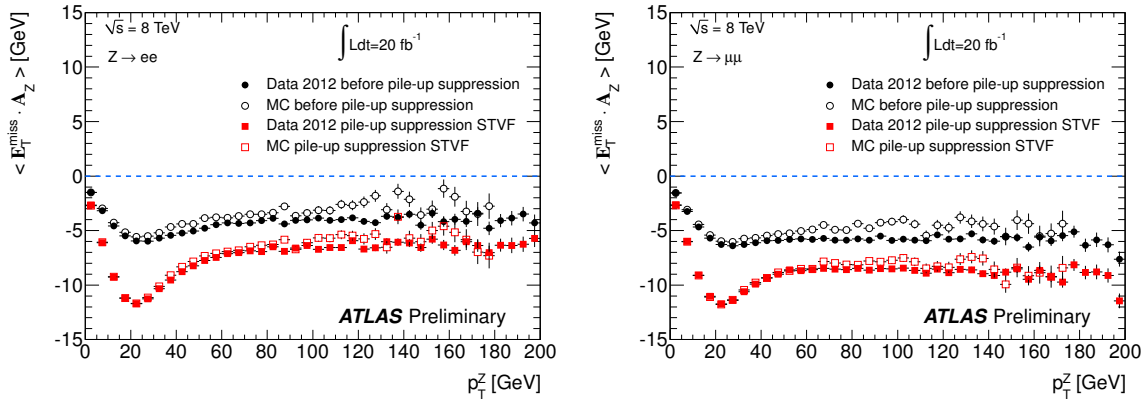


Figure 5.5: The average projection of the E_T^{miss} along the Z boson axis as reconstructed from $Z \rightarrow ee + \text{jets}$ (left) and $Z \rightarrow \mu\mu + \text{jets}$ (right) events. The E_T^{miss} variant used in the mono-jet analysis is closest to the black dots. The negative offset shows a bias in the E_T^{miss} coming from the jet scale. [63]

When considering the balance of the entire event, the poor calibration of any object can lead to undesired imbalances. The jet term in particular is known to bias the scale of the E_T^{miss} preferentially in the direction of jets, as seen when comparing the balance of $Z \rightarrow \mu\mu + \text{jets}$ and $Z \rightarrow ee + \text{jets}$ events in Figure 5.5. These plots show the average value of the projection of the E_T^{miss} on the Z boson axis in events comprised of two electrons or muons (from the Z) and balancing jet(s). If the E_T^{miss} scale was perfect, then the average value would be zero, as the leptonic system would perfectly balance the hadronic system. Instead, an imbalance in the direction of the jet system is observed, indicating the presence of a bias in the E_T^{miss} scale. This is only at the level of 5 GeV for the E_T^{miss} variant considered, thus the bias is negligible when selecting high E_T^{miss} events as in the mono-jet analysis, but it demonstrates the small inherent event imbalance even in the absence of true missing energy.

This provides a way to quantify the bias in the E_T^{miss} scale. However, this is done for topologies where there is no real E_T^{miss} . A different approach is to consider the $W \rightarrow \mu\nu + \text{jets}$ and $W \rightarrow e\nu + \text{jets}$ topologies, where the neutrino provides a source of real E_T^{miss} . The linearity of such events, defined as the difference between the truth and reconstructed E_T^{miss} divided by the truth value, is shown in Figure 5.6. This demonstrates how E_T^{miss} cuts of below 50 GeV are poorly modelled in MC.

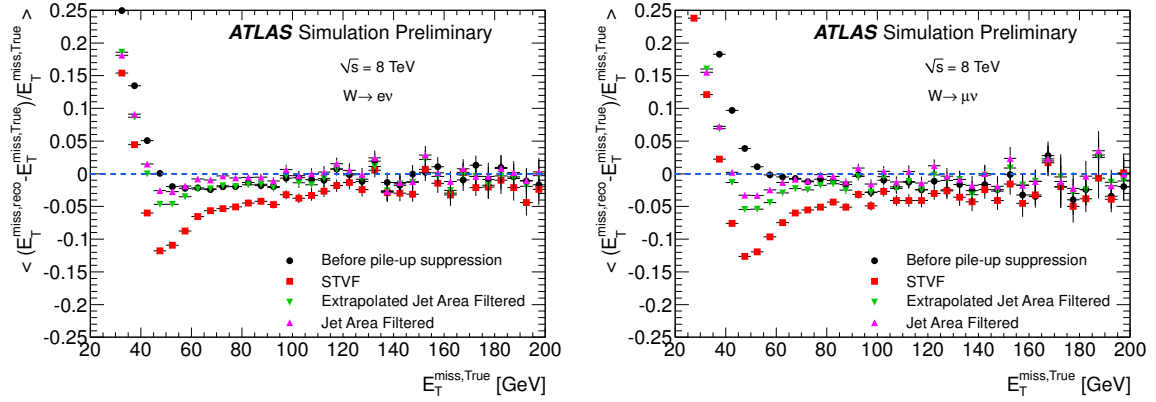


Figure 5.6: The linearity of the E_T^{miss} as reconstructed from $W \rightarrow e\nu + \text{jets}$ (left) and $W \rightarrow \mu\nu + \text{jets}$ (right) events. The E_T^{miss} variant used in the mono-jet analysis is closest to the black dots. Both topologies are observed to reach percent-level linearity between 50 and 100 GeV, depending on the required level of agreement. [63]

The mono-jet analysis makes use of a similar selection for some of the control regions, as will be discussed in Chapter 7, but with a threshold of 150 GeV. This threshold is in the regime where the data and MC are observed to agree within the statistical uncertainty, with the possibility of a small bias of $\mathcal{O}(1\%)$. As such, the mono-jet search is working well within the regime in which the reconstructed E_T^{miss} is a good representation of the true value.

Chapter 6

Jet Reconstruction and Performance

6.1 Introduction

Unlike the other objects discussed in Chapter 5, jets are not a physical particle from the SM, rather they are a tool designed to represent an underlying physical process. While electrons exist as an independent entity, jets are defined only by the method which was used to build them. When a muon is observed, all of the different reconstruction algorithms aim to reproduce the single muon to the best precision with a single set of detector observations. On the other hand, the same set of detector observations can be used to build radically different types of jets, many of which are independently useful. Examples of the independent uses for different types of jet algorithms will be provided.

Given the large degree of flexibility available within the jet framework, they are typically used for many distinct purposes. Arguably the most common use of jets is to represent the hadronic shower formed by a hard-scatter parton from the interaction point, making them a very common object at hadron colliders like the LHC. This stochastic process involves the production of a large number of particles, produced in a collimated cone, with a multiplicity and occupancy well beyond the current detector thresholds for the independent reconstruction of every individual particle. However, this is typically not a limitation, as the common interest is in the kinematics of the original hard-scatter parton, not the observed decay products. As such, measuring the entire ensemble of particles as a single entity is a suitable means of representing the inherently probabilistic underlying process. This is where the jet enters, providing a means of defining algorithms for the reconstruction of collimated streams (or jets) of particles as a single representative object.

The stream of particles produced from the hadronization of a parton is collimated due to the kinematic boost of the system, where the boost is the amount of kinematic energy associated to the object. Particles decay or radiate other particles back to back (for a two body split) in the rest frame. Moving back to the detector frame results in the products primarily aligned along the direction of the boost. As the centre of mass energy is increased, the hard-scatter partons become more energetic, increasing the boost of the system and thus resulting in increased collimation. Note that while the top quark is a type of parton, it is usually excluded for the purposes of jet reconstruction, as it is significantly different than other partons at the LHC energy scale. The top quark decays to a W and a b -quark, where only the latter fits within the most common ATLAS jet definition. Furthermore, due to the very large top quark mass, the W boson is typically not collimated with the b -quark. As such, the top quark cannot typically be

represented as a single jet of particles as it was defined above.

The top quark provides an ideal means of demonstrating an alternate use of jets beyond the typical definition above. The degree of collimation, which is what was used to define a jet in the first place, varies with $2m/p_T$. However, m is the mass of the original object, which does not have to be a parton. This phenomenon has given rise to a growing field of interest in using a different style of jet to represent other objects, most notably top quarks, W bosons, and Z bosons. As the amount of accumulated LHC data has increased, even the newly discovered Higgs boson has joined this list.

These types of objects are typically studied using much *larger* jets, in the sense that the required degree of collimation is much less restrictive. This is required for all of the decay products to fall within a single object. At the same time, this is not possible at low energies. A large boost is still required for the objects to be collimated enough such that the particles are produced in the same general area, thus studies involving *large* jets are typically conducted at high p_T . With this done, advanced substructure techniques can be used to exploit the differences between these *large* jets and the aforementioned typical *small* jets. For example, consider the decays $W \rightarrow qq$ and $Z \rightarrow qq$. Both of these processes produce two partons, which can then be reconstructed as two independent *small* jets, both of which typically have a similar energy. However, by using a *larger* jet and finding two sub-jets within it, a hadronic boson decay can be inferred. Further requiring that both of the sub-jets contain a similar energy provides further confidence that the two jets are from a W or Z decay. These methods are known as jet substructure techniques, and this is just one simple example of the type of power which accompanies the use of *large* jets for tagging entire bosons.

It is worth noting that the collimation relation is once again important for this sample scenario. As the process becomes more and more boosted, the entire top quark or boson can be contained within a typical *small* jet. This is shown in Figure 6.1, which demonstrates how the angular distance $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$ between the decay products of a top quark or W boson shrinks as the transverse momentum is increased. Clearly, there is a large degree of interplay between these two different classes of jets. They both have their own advantages and limitations, and techniques developed for one can be applied to the other. In particular, the procedure for calibration of *small* jets is extremely refined and could be applied to *large* jets, while the substructure techniques from *large* jets could also provide very powerful background rejection techniques for *small* jets.

The remainder of this Chapter will focus on the use of *small* jets unless otherwise mentioned, due to their use in the mono-jet analysis, interpretation, and upgrade projections, as will be discussed in the following Chapters. However, it is important to emphasize that this is only one of the options for representing the underlying physical processes via the use of jets.

In addition, given the versatility of jet reconstruction algorithms, it is important to specify that this Chapter will focus on so-called calorimeter jets. That is, the set of inputs to the jet algorithm will, unless otherwise stated, be a collection of calorimeter energy deposits represented as topo-clusters (discussed in the next Section). Other common inputs are tracks (leading to track jets) and truth particles (leading to truth jets). The actual process of jet reconstruction with other objects is similar, although calibration, uncertainties, and other properties differ by input type. Topo-clusters are the most common input type to jet reconstruction, and is what the mono-jet analysis uses.

The creation of topo-clusters and the reconstruction of jets is explained in Section 6.2. Jet calibration is discussed in Section 6.3, while the associated scale uncertainties and resolution are discussed in Sections 6.4 and 6.5 respectively. A calibration step specifically aimed at very energetic jets which are

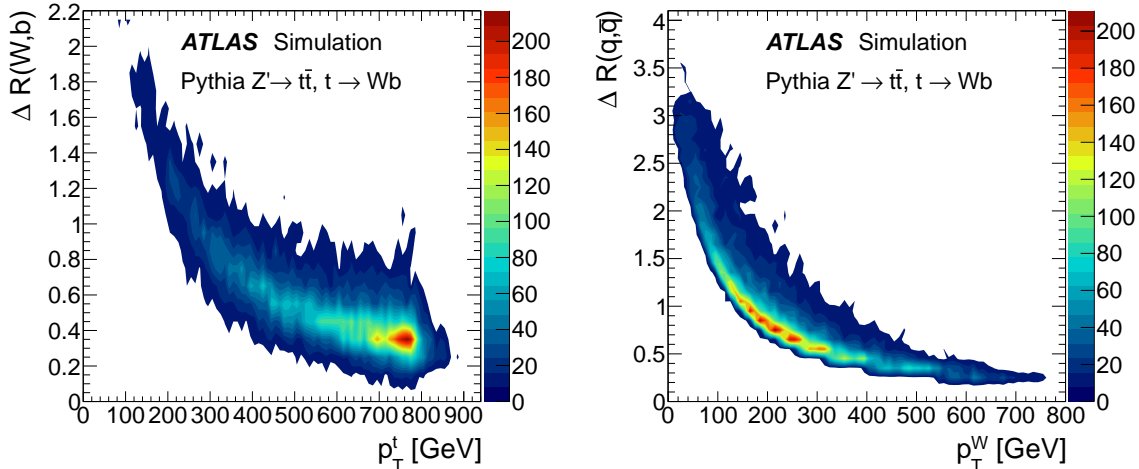


Figure 6.1: The angular distance $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$ between the W and b -quark from top decays (left) and between the two quarks from W hadronic decays (right) as a function of the transverse momentum of the system. As observed, the degree of separation falls with increasing energy, corresponding to more boosted initial states. The dependence of the degree of collimation ΔR on the mass is clear when comparing the top to W at a fixed p_T . For example, at 200 GeV the top collimation is $\Delta R \approx 1.5$, while the W collimation is closer to $\Delta R \approx 0.8$. This approximately matches the mass difference between the top quark and W boson. Combining these two effects demonstrates the dependence of the collimation on the boost and mass of the system, $\Delta R \propto m/p_T$. [64]

not contained within the calorimeter is provided in Appendix F. The impact of jets passing through non-functional regions of the calorimeter is studied in Appendix C. Finally, concluding statements on jet performance are listed in Section 6.6.

6.2 Topo-clustering and reconstruction

6.2.1 Topo-cluster formation

Building topo-clusters

Topological cell clustering, or topo-clustering, is the process by which the individual calorimeter cells are topologically combined in order to find groupings of hard-scatter activity. The resulting three-dimensional groups of cells are known as topo-clusters. Reference [56] contains a full description of the topo-clustering procedure, while the below is a general introduction.

The topo-clustering algorithm is designed to suppress calorimeter contributions from noise-related effects, while retaining activity from real physical processes. In order to do this, it is necessary to have a reasonable estimate of the average amount of noise in a typical calorimeter cell, σ_{noise} . In very early LHC data taking conditions, σ_{noise} was essentially given by the average amount of electronic noise in a given cell of the calorimeter. Now that the pileup level has increased considerably, the associated amount of energy spread throughout the calorimeter from pileup activity dominates the electronic noise, and $\sigma_{\text{noise}} \approx \sigma_{\text{noise}}^{\text{pileup}}$.

The average expected noise allows for the definition of a cell significance, $s_{\text{cell}} = E_{\text{cell}}/\sigma_{\text{noise}}$. This

quantity defines the probability of which the observed cell energy comes from noise, and thus provides a useful discrimination technique to suppress unwanted contributions. This is then used to construct a series of dynamic topo-clusters using the following three-stage classification, often called the ATLAS 4-2-0 topo-cluster threshold scheme.

1. Seed cells: cells with $|c_{\text{cell}}| > 4$
2. Secondary cells: All cells adjacent to a seed cell or already found secondary cell with $|c_{\text{cell}}| > 2$
3. Tertiary cells: All cells adjacent to a seed cell or secondary cell with $|c_{\text{cell}}| > 0$

In order to build the topo-clusters, all seed cells are identified and sorted in decreasing significance. Then, starting with the highest significance cell, all adjacent cells are tested for secondary status. This works in all three spatial dimensions, thus the adjacent cells can be in the same layer of the same calorimeter, different layers of the same calorimeter, or even different calorimeters. If the adjacent cell passes the secondary requirement, but was already in the collection of seed cells, it is removed from the seed collection. This procedure is then iterated, looking for new secondary cells adjacent to already identified secondary cells, and continues until no new secondary cells are added. The topo-cluster is then completed by adding all cells adjacent to the seed or secondary cell(s) that pass the tertiary requirement, which in this case corresponds to all adjacent cells. This step is not iterative. This entire procedure is then repeated for each seed cell until none remain.

Once the raw topo-clusters have been built, the next step is to search for local maxima within each topo-cluster. This is important for events such as collimated Z boson decays, where two energetic quarks can be produced in close proximity. These two quarks can then form a single raw topo-cluster with local maximum energy peaks representing the leading and sub-leading quarks. This is important for both reconstructing unbiased typical *small* jets as well as to ensure that the advanced substructure techniques alluded to in Section 6.1 can be used.

In order to avoid fluctuations, local maxima are required to have at least $E_{\text{cell}} > 500$ MeV, at least four neighbours, and no neighbours with a larger energy signal. The local maximum is also required to occur within the middle layer of the electromagnetic or hadronic calorimeters, which is much deeper than the first and last layers, and which can thus provide a more reliable energy measurement. Once split, boundary cells are shared between at most two local maxima, defined by the two highest adjacent energy clusters. The energy sharing weight is then defined with respect to the energy of the two clusters and the distance of the cell to the centroid of each cluster. This accounts for the overlapping contributions of two independent underlying hard objects sharing the same calorimeter cells.

After topo-clusters are constructed, the energy of all of the constituent cells (with associated weights if appropriate) are summed to get the total energy of the topo-cluster. The angles are then set to match the energy-weighted centroid of the topo-cluster, and the mass is set to zero. This then sets the full four-momentum of the topo-cluster.

The topo-cluster forming procedure is designed to adapt to the stochastic process of electromagnetic and hadronic showers, and thus does not require the energy deposits to be distributed in a uniform shape. Instead, the iteration of the second stage allows for topo-clusters to naturally follow the energy deposits rather than assuming a square shape as is done in *sliding-window* clusters.

There are multiple relevant energy scales for a cluster to exist at. First is the electromagnetic scale, in which the electronic signals are translated into an energy value, and the sampling fraction of the

calorimeter is used to estimate the actual energy of the shower from the observed deposits. This works very well for electromagnetic processes, but hadronic showers have a large degree of variability. Identical hadronic processes have the potential to lead to significantly different measured energies depending on various factors, such as the fraction of π^0 vs π^\pm produced. In order to reduce these dependencies, another energy scale for topo-clusters is defined.

Local calibration of topo-clusters

In order to further refine the performance of topo-clusters, a Local Calibration Weighting (LCW or LC) scheme has been constructed. This correction is based purely on MC, and thus it is important to ensure that any correction applied is well-modelled in order to not introduce a bias between data and MC.

The LC scheme first classifies the cluster as either electromagnetic-like or hadronic-like, based on a likelihood involving the depth of the shower centre and the cluster energy density. This classification determines which corrections are appropriate to apply, which is a critical step given the very different nature of electromagnetic and hadronic showers. The following three sets of correction types are then applied, and are described below.

1. Calorimeter non-compensation (hadronic-like clusters)
2. Signal losses due to energy outside of the topo-cluster (all clusters)
3. Signal losses due to energy deposited in inactive material (all clusters)

The ATLAS calorimeters are non-compensating, thus so-called invisible energy processes such as the energy required for nucleus break-up decrease the amount of energy observed in the calorimeter. For this reason, electromagnetic particles, such as electrons or photons, and hadronic particles, such as pions, will provide different measured calorimeter signals even when the real particles have identical energies. In other words, $E_{\text{truth}}^{\text{EM}} = E_{\text{truth}}^{\text{Had}}$ implies $E_{\text{clus}}^{\text{EM}} > E_{\text{clus}}^{\text{Had}}$. The application of local corrections to the measured cell energies of hadronic-like clusters allows for this to be partially alleviated, with $E_{\text{clus}}^{\text{EM}} \approx E_{\text{clus,LC}}^{\text{Had}}$. This is a very useful change, both for improving agreement between data and MC and for improving the measurement of derived quantities, such as the use of LC clusters in the $E_{\text{T}}^{\text{miss}}$ SoftTerm described in Section 5.5.

While the topo-cluster scheme is designed to keep signal contributions and reject noise, the rejection scheme can be too rigid, and not all of the signal energy is necessarily contained within the topo-cluster. This is particularly true as the amount of pileup increases, and thus the energy tails of the shower profile drop below the required noise threshold. Corrections to adjust the shower profile near the edges of the topo-cluster can partially resolve this problem. While the same procedure is used for both electromagnetic and hadronic clusters, the magnitude of the correction differs to account for the difference in the expected shower profiles.

Whenever a particle passes through material, whether it is active or not, it deposits energy. If the material is not instrumented, then the deposited energy is lost, and the resulting reconstructed cluster energy decreases with respect to the true value. With a full MC simulation accounting for all of the inactive and uninstrumented material, the amount of missing energy as a function of the observed energy for a given region of the detector can be estimated. This can then be used to calibrate the topo-clusters on average for energy deposits within inactive regions.

After these three corrections are applied, the topo-clusters are referred to as having been calibrated from the raw ElectroMagnetic (EM) scale to the Local Calibration weighted (LC) scale. The topo-clusters are then the primary input to calorimeter-based jet reconstruction, and both topo-cluster scales are used depending on the intended uses of the resulting jets. The mono-jet group uses jets built from LC topo-cluster inputs, but both types will be discussed in the following Sections.

6.2.2 Jet reconstruction

Once the jet inputs have been defined, whether they are truth particles, tracks, calorimeter topo-clusters, or something else, an algorithm must be defined which takes the inputs and groups them into collimated collections (often thought of as cones). This mapping of a collection of inputs to a collection of jets which represents the underlying physical process is known as jet reconstruction, or alternatively jet finding. Before selecting any particular algorithm to fulfil this purpose, it is important to understand what is required of the algorithm. Reference [65] provides a very clear outline of what is needed to make a useable jet algorithm, and their list is quoted verbatim below.

1. Simple to implement in an experimental analysis;
2. Simple to implement in a theoretical calculation;
3. Defined at any order of perturbation theory;
4. Yields finite cross section at any order of perturbation theory;
5. Yields a cross section that is relatively insensitive to hadronization.

The first two points are purely practical, as an algorithm is not useful if it is not straightforward to both calculate theoretical results and make experimental measurements. The remaining points deal with the requirement that the jet be both Infra-Red (IR) and collinear safe, meaning that the addition of a soft particle or the collinear splitting of a hard particle should not change the number of reconstructed hard jets in the event. These two requirements require further explanation, which is provided in Section 2.1.4 of Reference [66]. A short summary is provided below.

IR safety requires hard jets to not change if an additional soft emission occurs. This is a necessity given the frequency with which soft gluons are produced in hadronic showers. Figure 6.2 provides a graphical representation of this effect, in which a soft emission causes two jets from a W boson to merge into a single jet, which would not have occurred if the soft emission was not present. More generally, the addition of soft radiation which contributes only negligible energy to the jet can cause the observed hard jet multiplicity to change. This occurs due to the soft emissions modifying the way in which the hard energy deposits are grouped.

A simple example of an algorithm where this would happen would start with the softest partons and then combine objects based purely on the angular distance, with a fixed maximum distance cutoff. The addition of a soft object between two harder objects can then cause them to merge if they are both within the fixed distance with respect to the soft parton, but they are not within the fixed distance with respect to either hard parton.

Collinear safety represents stability against the natural process of a single parton splitting. This can be such processes as a quark radiating a gluon, a gluon fragmenting into two quarks, or other such processes. Collinear safety is a crucial consideration when building jets, as hard-scatter partons




Figure 6.2: The decay of a W boson into two hard quarks, which are independently resolvable for the given jet algorithm parameters (a). After the emission of a soft gluon, an IR safe jet algorithm still observes two separate jets (b), while IR unsafe algorithms merge the two jets into one (c). [66]

each undergo numerous such stochastically determined splittings during the process of a single hadronic shower. Figure 6.3 provides a schematic diagram outlining the way in which this can change the number of reconstructed hard jets in algorithms which are not stable against this phenomenon. In this example, replacing a quark with a virtual gluon correction with a distinct quark and gluon causes the original single jet to be split in two. More generally, the splitting of a hard object into multiple softer objects causes the observed hard jet multiplicity to change. Once again, this occurs due to a change in the way in which the hard energy deposits are grouped.

A simple example of an algorithm where this occurs starts with the hardest partons and then combines objects based purely on the angular distance, with a fixed maximum distance cutoff. The splitting of the hardest object into two softer objects thus changes the starting point (the seed), and the new hardest parton becomes the geometric centre of the jet. Then, the two parts of the previously hardest object are combined into this jet as they are within the distance cutoff, while the residual hard partons are too far away from the seed to be included. This remaining energy is then collected into a second jet.

As demonstrated, jet algorithms which are either IR or collinear unsafe can cause the reconstructed jet multiplicity to vary, which changes the observed physics. This is very dangerous, and must be avoided. It is thus important to find a set of algorithms which are easy to use for both theorists and experimentalists, and which additionally are both IR and collinear safe. There are currently multiple such jet algorithms available, including both cone-jets and k_t -family jets, the majority of which are provided by **FastJet** as per Reference [67].

The ATLAS jet reconstruction software makes extensive use of the **FastJet** software library due to its simple and uniform interface, as well as the large selection of functionality and continued development. ATLAS primarily makes use of k_t -family jet algorithms, with typical uses of anti- k_t for *small* jets, Cambridge/Aachen (C/A) for *large* jets, and the original k_t for sub-jets within *large* jets. The focus of this Section will be on the use of the anti- k_t jet algorithm due to its use in the mono-jet analysis, but the other k_t -family jets will be briefly mentioned as a contrast.

Algorithms from the k_t -family follow a sequential recombination procedure, in which ATLAS uses the full four-momentum during each step of the recombination. This means that the final jet momentum (before calibration) is the sum of the momenta of the input topo-clusters that it contains. The iterative approach results in the construction of pseudo-jets, where a pseudo-jet is a collection of one or more topo-clusters (or other inputs), potentially combined in previous steps. The final pseudo-jets after the

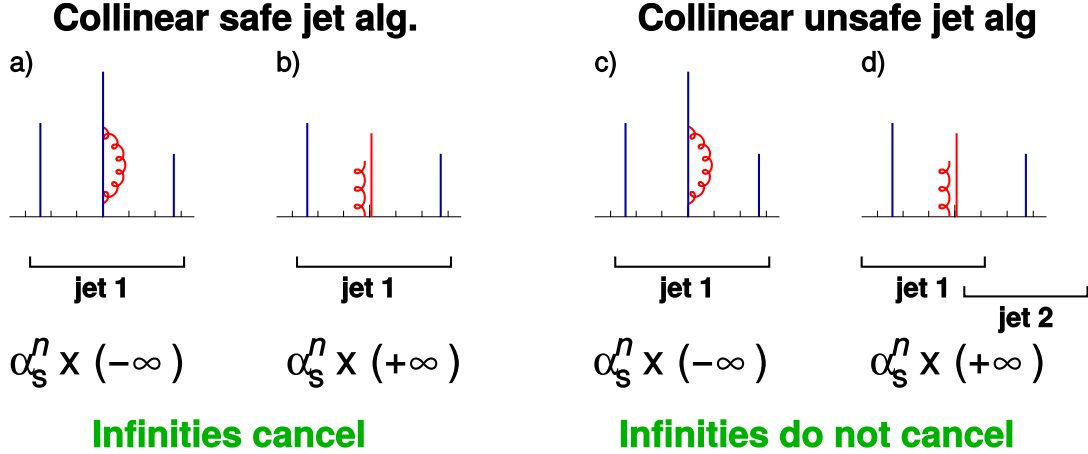


Figure 6.3: A comparison of two processes, demonstrating the difference between collinear safe and unsafe algorithms. Parton p_T is represented by the height of the line, while the horizontal axis represents the parton rapidity. All algorithms are stable against the gluon correction to the quark propagator (a and c), which is by construction given that no particle is emitted. However, the gluon correction should be cancelled in perturbation theory by gluon radiation, where the gluon and parton share the p_T . This feature is provided by collinear safe algorithms (b), but not collinear unsafe algorithms (d), where the perturbative approach breaks down. The presence of an extra parton in the collinear unsafe algorithm leads to the reconstruction of two separate jets rather than one, thus changing the observed hard jet multiplicity. [66]

full event has been processed are then the resulting jets. The algorithms define two distance parameters, one with respect to the distance between the pseudo-jet and each topo-cluster d_{ij} , and one with respect to the distance between the pseudo-jet and the beam d_{iB} . These two distance parameters are defined in Equations 6.1 and 6.2 respectively. Note that the traditional forms of the equation use k_t rather than p_T for the transverse momentum, but p_T is used here to be consistent with the rest of this document.

$$d_{ij} = \min \left(p_{T,i}^{2p}, p_{T,j}^{2p} \right) \frac{\Delta_{ij}^2}{R^2} \quad (6.1)$$

$$d_{iB} = p_{T,i}^{2p} \quad (6.2)$$

The angular distance parameter is defined in terms of the rapidity and azimuthal angle as $\Delta_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$. The variable R is the distance parameter of the jet algorithm, which is a configurable value that controls the resulting size of the reconstructed jets. Typical values in ATLAS are $R = 0.4$ or $R = 0.6$ for *small* jets, and $R = 1.0$ or $R = 1.2$ for *large* jets. The power of the transverse momentum, p , controls the ordering of the inputs. The general properties of a given algorithm for $p \geq 1$ are similar, and the same is true for $p \leq -1$, as the ordering is already set by the sign as detailed in Reference [68], thus only values of 1, 0, and -1 are considered. The traditional k_t algorithm corresponds to a choice of $p = 1$, the Cambridge/Aachen algorithm is $p = 0$, and the anti- k_t algorithm is $p = -1$.

The actual execution of the algorithm begins by identifying the smallest distance between two of the inputs (typically topo-clusters). If $d_{ij} < d_{iB}$, then the objects i and j are combined. Objects i and j are then removed from the collection of inputs, and object ij is added as a new input. If instead $d_{ij} > d_{iB}$, then the object i is designated as a jet, and is removed from the collection of inputs. In either case,

after an action has been taken, the distances d_{ij} and d_{iB} are recalculated for all combinations and the procedure is repeated until no objects remain.

In picking between these three options, it is important to consider the behaviour with respect to soft radiation. While all three are IR safe, the exact treatment is not identical, and thus they behave differently. The k_t jet definition is soft-adaptive, while the anti- k_t definition is soft-resilient. This means that the jet shapes in angular parameter space either depend on or are unaffected by soft radiation respectively. Soft-adaptability means that the jet algorithm is able to better approximate the stochastic evolution of hadronic showers, while soft-resilient jets are expected to miss some of the deposited energy. However, this missed energy can generally be fixed on average through jet calibrations. Soft-resilience, on the other hand, has numerous benefits which more than outweigh the drawbacks when used in typical collider environments. In particular, this shape consistency significantly enhances the ability of experiments to remove contributions from pileup and the underlying event, both of which are very important effects to mitigate at the LHC. As such, anti- k_t jets are the primary reconstruction algorithm used in ATLAS.

It is worth noting that Equation 6.1 controls the order in which objects are combined. Positive values of p favour the softest inputs, negative values of p favour the hardest inputs, and $p = 0$ orders purely by angular separation. This has an important effect on the use of jet substructure techniques, and is the reason that Cambridge/Aachen jets are typically used for *large* jets, while k_t jets are typically used for sub-jets within *large* jets. The angular-only ordering of Cambridge/Aachen jets is ideal to build a large jet composed of several possible hard inputs, as the resulting highly irregular shapes can naturally follow the multiple independent showers contained within it.

The k_t style of merging from soft to hard objects is then very useful for sub-jets, as the jet can be built, and then it can be split into multiple hard sub-jets by reversing the combination chain. This is particularly useful for scenarios like $Z \rightarrow qq$, where each of the quarks is hard at matrix element level, but then will radiate during the parton shower. A recombination with the k_t algorithm will typically result in the construction of two pseudo-jets representing each quark, and the final step is to merge the two pseudo-jets (quarks) into a single jet (Z boson). This can then be reversed by one step to get the two sub-jets. This can be further reversed to see the momentum split within an individual sub-jet, which can hint at further structure if it exists, such as $t \rightarrow Wb \rightarrow qqb$. In other words, the ability to reverse the jet building step by step with the k_t algorithm provides very powerful information about the underlying structure of the jet.

While Cambridge/Aachen and k_t jets definitely have their uses, the calibration and pileup suppression of such objects is much more difficult, and thus they are only used when necessary. Cambridge/Aachen $R = 1.2$ jets are the most common approach taken in studies involving jet substructure, with k_t $R = 0.4$ sub-jets. Instead of using these types, the typical ATLAS analysis relies exclusively on anti- k_t $R = 0.4$ jets (and to a much lesser extent, anti- k_t $R = 0.6$ jets). These options are both fully calibrated through a complicated procedure, as is described in the next Section.

6.2.3 Jet associations

Once a jet is identified, it is a common question to ask whether it is close to or even contains other objects. The most common example is tracks from the inner detector, which can be extrapolated to the calorimeter. It is then useful to know if these extrapolated tracks point to a given jet or not. These associations between jets and other objects are conducted through a procedure known as *matching*.

There are two common types of matching procedures, one which is simplistic but limited, while the other is more technically complicated and robust.

The simplistic approach is known as ΔR matching. In this procedure, the angular distance between the centroid of the jet and the object(s) in question is calculated as $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$. If this value is below an arbitrary cutoff, the jet is considered to be associated to the given object, otherwise they are not associated. The arbitrary cutoff limits the viability of this procedure, particularly in topologies where multiple jets are close to each other, but it performs adequately for many simple uses.

6.2.4 Jet cleaning

Note that jet reconstruction is conducted under the assumption that the inputs are all taken from detectors which are fully operational, and with no extraneous activity which does not come from the collision. This assumption is generally very reasonable, but occasionally breaks down. Calorimeter problems such as noise bursts in the HEC, coherent noise in the electromagnetic calorimeters, and non-collision sources can result in the reconstruction of fake jets. It is important to remove these fake jets such that events are not selected or rejected based on activity which is not actually from the hard-scatter collision. These sources are typically controlled by simple cuts as defined in Reference [69], in a procedure which is known as jet cleaning.

While these fake jet rejection criteria are typically suitable for the average analysis, the mono-jet topology is particularly sensitive to non-collision events. As such, additional criteria were found to be necessary in addition to these cuts, as is demonstrated in Section 7.6.6.

These are the most commonly considered sources of bad jets in ATLAS, but the 2012 dataset had another major calorimeter problem which requires careful attention. The high voltage power supplies used in the Tile calorimeter would frequently fail, creating time periods where part of the calorimeter was inactive. This individually represents only a small fraction of the dataset and thus could be possibly ignored, but occasionally the power supplies would permanently fail. In this case, entire regions of the calorimeter would be inactive for the remainder of the dataset.

The inactive Tile calorimeter regions were not added to the MC simulations as the regions became inoperable, thus the simulated events and actual data had a substantial difference in calorimeter capabilities. Additional cleaning cuts were devised to reject events where jets near these dead regions were substantially biased. The studies leading to the cleaning criteria, as well as the resulting performance, are documented in Appendix C.

6.3 Jet energy scale calibration

6.3.1 Introduction

Hadronic showers contain a large variety of particles, and involve many types of processes. While much of the energy is observable in the detectors, a large fraction is not, as discussed in Section 4.4.1. In particular, the break-up of nuclei is a standard part of the hadronic shower which is invisible to calorimeters. This means that a significant fraction of the energy associated with the hadronic shower is missed, and thus the reconstructed jet momentum is not the same as the true process. Given that the ATLAS calorimeters are non-compensating, and given that the aforementioned topo-cluster local calibration is limited, the only way to fully account for this is to calibrate the jets after reconstruction.

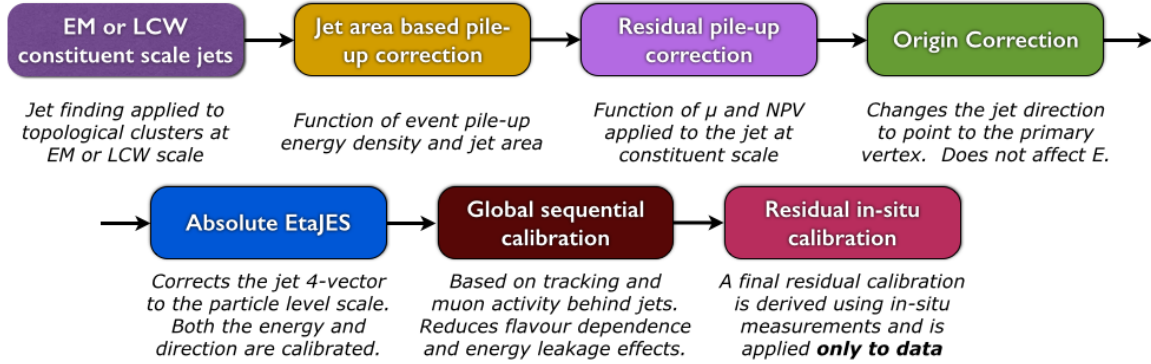


Figure 6.4: An overview of the ATLAS jet calibration procedure, showing the steps which are used to bring jets from the EM or LC scale (the scale of the input topo-clusters) to the final corrected scale. The origin correction and global sequential calibration are new, and were not included in previous versions of the full calibration chain derived for $\sqrt{s} = 8$ TeV data, such as the variant used by the mono-jet search.

Jet calibration is typically referred to as the Jet Energy Scale (JES), even though more than just the energy component of the four-momentum is corrected. This is because the only inputs are the observed energy in the calorimeters, and thus the primary correction is a factor which scales the observed (reconstructed) energy to the true jet energy.

The ATLAS jet calibration scheme is comprised of several parts, designed to account for different effects. The derivation and validation of this many-step process takes considerable time from start to finish, and thus the calibration is constantly evolving as more data or MC becomes available, as new techniques are studied, and as corrections for new effects are derived and added to the calibration chain. The sequence of steps must be followed in order, as the output of one calibration step is the input to the next. As such, there is a considerable time delay between completing the derivation of one calibration and the next. Each completed calibration chain is referred to as a particular release of the jet calibration, and is used for an extended period of time. This Section will thus focus on two of the ATLAS jet calibrations for $\sqrt{s} = 8$ TeV data, the so-called `Moriond2013` JES (released in January 2013) and `Final2012` JES (released in October 2014). The former was used for the mono-jet analysis due to differences in timescales, while the latter represents the most complete understanding of jet performance in the ATLAS detector to date.

A very brief overview is provided below, while the full details of the procedure, the motivation for each step, and demonstrative plots are provided in Appendix D.

6.3.2 Overview

The outline of the full ATLAS jet calibration chain is shown in Figure 6.4. Note that this procedure applies to jets built from either EM or LC scale topo-clusters and reconstructed using the anti- k_t jet algorithm with a distance parameter R of either 0.4 or 0.6. This covers the majority of jet usage within ATLAS.

Starting from the EM or LC scale (the scale of the topo-clusters), the first step is to suppress the contribution of pileup. Pileup can increase both the jet multiplicity and the energy of jets of interest. The multiplicity is increased primarily due to the addition of jets from pileup vertices, while the topo-clusters used to build jets from the hard-scatter vertex can contain contributions from both the hard-scatter

interaction and pileup sources, thus biasing the energy scale.

Once the pileup suppression has been applied, the next step is to correct the angle of the jet. ATLAS jets are typically reconstructed to point toward the nominal interaction point of $(0, 0, 0)$. Fixing this to point back to the actual hard-scatter vertex rather than the nominal point provides a significant improvement in the jet angular resolution. Note that this step was not available for the `Moriond2013` JES, but is included in the `Final2012` JES.

The next stage is the single largest correction in the chain, and is designed to bring jets from the EM or LC scale to the hadronic/particle energy scale, often referred to as the EM+JES or LC+JES. This correction is derived purely from MC, where reconstructed jets are matched to their corresponding truth jet. A series of distributions of $E_{\text{reco}}/E_{\text{truth}}$ are then created in bins of both energy and $|\eta|$, covering the primary dependencies of the correction. For any given jet, the relevant bin can then be found, and the inverse of the mean of this distribution can be used to scale jets to their true energy, on average.

The global sequential calibration is the next stage of the calibration chain, which is new to the `Final2012` JES, and thus is not used by the mono-jet analysis. This step is instrumental in reducing the differences between jets coming from different sources, namely light quark initiated jets as compared to gluon initiated jets. It also provides a means of correcting for jets which punch-through the full ATLAS calorimetry and enter the muon spectrometer, which becomes increasingly important as the jet energy is increased. The ability to calibrate punch-through jets is new, and is detailed separately in Appendix F.

The final stage is the most time consuming step of the chain, and is necessary due to the use of correction factors derived purely from MC information during the previous two steps. By calibrating data jets with correction factors from MC, biases from MC mis-modelling may be introduced. By conducting *in situ* comparisons between data and MC in several well-known topologies, it is possible to derive a residual *in situ* correction to be applied only to data in order to fix such mis-modelling. This step also provides uncertainties accounting for any remaining differences between data and MC across a wide parameter space of jet kinematic values.

6.4 Jet energy scale uncertainties

6.4.1 Introduction

The JES calibration scheme discussed in Section 6.3 and detailed in Appendix D consists of several steps, and thus systematic biases can be introduced in many places during the derivation. The *in situ* analyses measure the uncertainties with respect to $p_{\text{T}}^{\text{jet}}$ and η by quantifying the calorimeter jet response differences observed between data and simulation in well-understood topologies. This accounts for the majority of the possible differences. Additional systematic uncertainties accounting for the MC simulation framework used (such as `GEANT4` or `AFII`), pileup suppression modelling, flavour dependencies, and punch-through are also independently evaluated. All of these sources are detailed in Appendix E.1.

The JES calibration, and thus the associated JES uncertainties, are constantly evolving and improving. A set of uncertainties is only applicable when matched to the jet calibration it was derived for, thus the two aforementioned calibration chains result in two independent sets of jet uncertainties, referred to as uncertainty releases. A comparison of the different uncertainty releases is provided in Appendix E.2, showing the reduction in the JES uncertainties obtained by extending and improving the calibration procedure.

The set of jet uncertainty sources is large, and is far too complicated for the average ATLAS analysis to handle. In order to provide a more manageable set of uncertainties, an eigenvector decomposition is performed, thus significantly reducing the number of JES uncertainty nuisance parameters with minimal information loss. This procedure is explained in Appendix E.3.

When constructing the full set of uncertainty sources, occasionally the level of correlation between different components is not clear, and an assumption is made that the components are either correlated or uncorrelated. In order to understand the impact of this correlation uncertainty, alternative correlation scenarios are provided in Appendix E.4.

6.4.2 Overview

There are many different uncertainty configurations, depending both on the associated calibration used and the level of detail required by the analysis. Only the so-called *category reduction* uncertainty configuration for the Moriond2013 JES with anti- k_t $R = 0.4$ jets built from LC topo-clusters is briefly shown here, as that is what was used for the mono-jet analysis. The full details, comparisons, reductions, and alternative scenarios for the different configurations are left to Appendix E.

This configuration began with a total of 58 components, where 47 were from absolute *in situ* sources. These sources were all assigned a classification based on how they were evaluated, where the options were modelling (MC), detector, mixed modelling and detector, and statistical/method uncertainties. Eigenvector decompositions were then applied to reduce the 47 absolute *in situ* sources to 12, where the procedure is detailed in Appendix E.3. The resulting reduced parameters correspond to four modelling components, three detector components, two mixed components, and three statistical components. The reduction procedure in this scenario retains more than 95% of the correlations between jets across the full p_T^{jet} parameter space, thus the reduced number of parameters do a very good job of representing the full set.

The remaining 11 components cover relative *in situ* corrections (two terms), pileup uncertainties (four terms), flavour systematics (three terms, only two of which are relevant due to not identifying b -jets), a calibration non-closure term, and a high p_T jet uncertainty. The calibration non-closure term is zero when using the full GEANT4 simulation, and non-zero when using the fast AFII parametrization. Summary plots showing the uncertainty overview for a given jet in this configuration are provided in Figure 6.5.

6.5 Jet energy resolution

The JES calibration as described in Section 6.3 provides the method used to correct jets to the hadronic scale by considering the mean of the distribution for $E_{\text{reco}}/E_{\text{truth}}$ in bins of energy and $|\eta|$. While the mean of this distribution is clearly important for the jet energy scale and corrects jets to the correct scale on average, the width quantifies how much of a spread remains, which is also an important measure. Tighter widths correspond to a higher fraction of jets being calibrated to the correct scale, thus the smallest width possible is desired. The width of this Gaussian jet response distribution is known as the Jet Energy Resolution (JER).

A number of factors can give rise to deviations in the jet response for identical input conditions. The stochastic nature of hadronic showers is one source which can never be removed, as is the contribution of out-of-cone radiation due to the fixed jet distance parameter R of 0.4 or 0.6. Both of these are examples

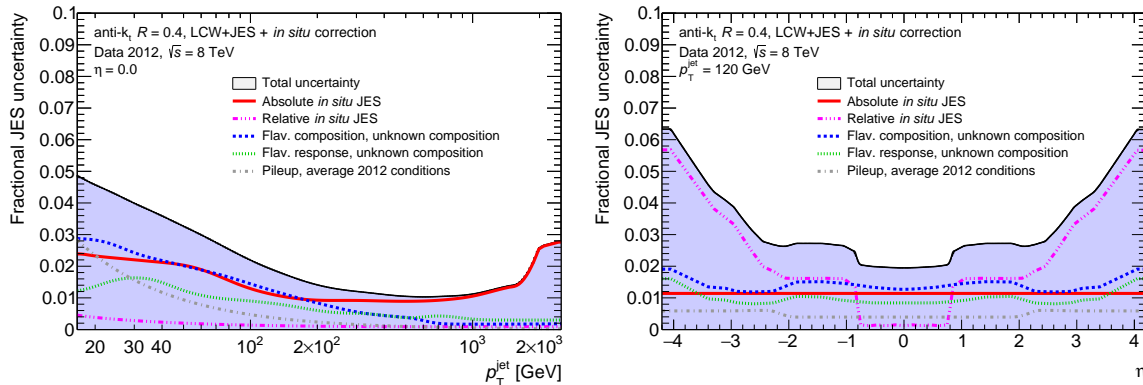


Figure 6.5: Moriond2013 JES uncertainties for anti- k_t $R = 0.4$ jets built from LC topo-clusters, shown vs p_T (left) and η (right), as used in the mono-jet analysis. The mono-jet analysis requires that the leading jet have $p_T > 120$ GeV and $|\eta| < 2.0$, hence the left plot is for $|\eta| = 0$ and the right plot is for the p_T threshold. The total uncertainty is below 2-3% for the majority of the parameter space relevant to the leading jet of the mono-jet analysis. The individual lines represent different components of the JES uncertainty, as detailed in Appendix E.1.

of resolution effects which are expected and cannot be countered, although that does not mean that the contribution to the JER from these is equivalent in data and MC. Beyond that, there are also calorimeter limitations, such as electronic noise or the larger contribution from pileup. While the pileup suppression methods described in Appendix D.1 are designed to minimize these effects, parts of the procedure are derived for the full calorimeter. As such, it is possible that the local pileup activity near the jet in question is either above or below the average for the full calorimeter, thus pileup remains a significant resolution effect.

Given that one of the primary differences between the truth and reconstructed JER is calorimeter resolution, it is important to recall the dependence of the resolution on energy. First there is the noise term, which is large but varies with the inverse of energy. Then, there is the stochastic term, which is moderate and varies with the inverse square root of energy. Finally, there is the underlying constant term, which is independent of energy. As such, the JER is largest at low energy, especially under high pileup conditions. As the energy is increased, the difference between truth and reconstructed resolution shrinks, until it becomes a relatively small effect at high energy. This indicates that, from a practical standpoint, the JER is more relevant for analyses which select or veto on low energy jets rather than those which work in the higher energy regime.

Despite being called the Jet Energy Resolution, it is typically parametrized in terms of the jet p_T , as this is a more generically useful quantity at hadron colliders. The JER with respect to p_T is then measured *in situ* in order to determine what portion of the resolution is irreducible from physical processes, and what portion comes from the detector and limitations of the reconstruction procedure.

For the Moriond2013 JER release, the di-jet balance and bisector methods were used, both of which are detailed in Reference [70]. There are no public plots available at this time, but this procedure matches what was done in 2011, as shown in Figure 6.6.

For the Final2012 JER, which is still under development, a more advanced estimation of the JER is conducted following a full set of *in situ* studies, analogous to what was done for the calibration. The Final2012 JER should provide a very detailed picture of the differences between truth and reconstructed

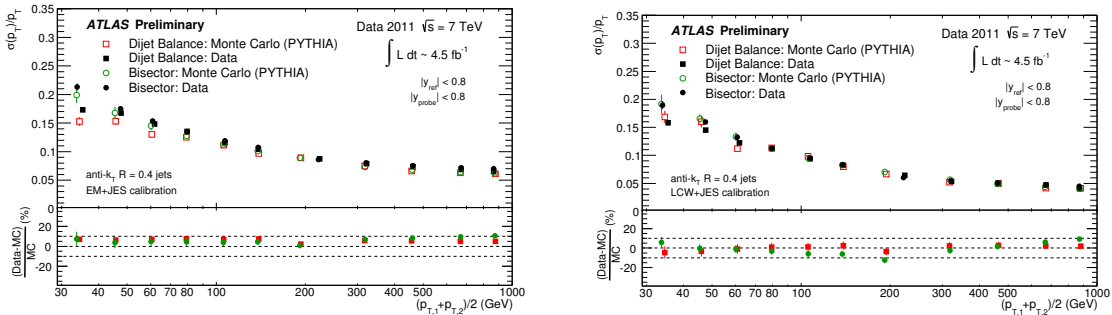


Figure 6.6: The Jet Energy Resolution as derived from the di-jet balance and bisector methods, using $\sqrt{s} = 7$ TeV data and the anti- k_t $R = 0.4$ jet algorithm built from EM topo-clusters (left) or LC topo-clusters (right). Error bars shown only include statistical uncertainties. The two methods for estimating the JER are seen to be in good agreement. The lower sub-plot shows the relative difference between data and MC, which is near the 5-10% level for EM topo-cluster jets and varies between $\pm 10\%$ for LC topo-cluster jets. [71]

jet resolution in ATLAS.

6.6 Conclusions

Jets are a very interesting object, especially at hadron colliders like the LHC where the cross-sections of jet-related processes are enormous. The reconstruction of jets at ATLAS is primarily conducted by first finding topo-clusters, or topologically grouped sets of cells which pass certain noise-reduction criteria. These topo-clusters are then used, either as-is or with an extra local calibration weighting, as the inputs to a jet algorithm. The most common jet algorithms in ATLAS belong to the k_t -family, and the anti- k_t algorithm in particular is widely used.

Hadronic showers involve a substantial portion of energy which is not detected by the ATLAS calorimeters. As a result, the default scale of reconstructed jets is significantly offset from the true values. In order to correct this, the Jet Energy Scale (JES) is calculated through a series of steps involving both MC corrections and subsequent *in situ* studies. This correction fixes the scale of the average jet. The width of the scaling distribution, or the Jet Energy Resolution (JER) has historically been calculated through two different *in situ* methods, but this is currently being expanded to include new approaches.

The uncertainty on the JES has been significantly improved over the previous years, reaching down to a level of precision near 1% for a significant part of the parameter space. In the process, the limiting factors for the JES has switched from being jet systematics, to the uncertainties on the reference objects used for calibrating the jets. This is a tremendous achievement, and demonstrates the success of the dedicated effort conducted over the past two years.

The JES uncertainties have been further simplified through the means of a re-parametrization which reduces the number of components needed to accurately represent the full set. This allows users to more easily handle the systematics, and to determine the impact of the JES uncertainties on their analysis. In addition to providing such a default set of uncertainties, the assumptions that went into building the uncertainties have now been tested. This new approach provides a very good grasp on the degree to which the correlations within the JES uncertainties are understood, which is a necessary step for

furthering the high-precision ATLAS measurements of SM processes which are currently underway.

In the process of the latest derivation of the JES calibration, a new phenomenon was investigated in the form of jets which are not fully contained within the ATLAS calorimetry. This has the potential to become a significant problem for very energetic jets at the TeV scale, with the fraction of such non-contained events reaching 10% for $p_T^{\text{jet}} \approx 2 \text{ TeV}$. A first correction for such jets has been provided, which has been shown to fix the scale of such jets on average. Such corrections and studies will become more important in the near future, as the LHC turns back on at a new energy scale, and the fraction of such jets increases.

The impact of regions of the Tile calorimeter becoming unreliable, due to power supply failures or other sources, has been investigated. It is found that very low p_T jets are essentially unaffected, as the JER smears out any contribution of such an effect. However, for $p_T^{\text{jet}} \gtrsim 50 \text{ GeV}$, and definitely for $p_T^{\text{jet}} > 100 \text{ GeV}$, it is clear that a notable portion of jets which include such regions have significantly under or over-estimated energies. A set of cleaning criteria has been determined which rejects the majority of such bad jets, while retaining most of the good jets, thus providing for a way to save as much of the data as possible.

Chapter 7

The Mono-jet Analysis

7.1 Introduction

This Chapter is based entirely on Reference [5]. All ATLAS material, unless otherwise mentioned, is from this document.

The original intent of the mono-jet analysis, as the name implies, is the search for topologies involving a single jet and no other visible activity. Due to the conservation of transverse momentum at hadron colliders, the single jet is then assumed to be balancing invisible particle(s). More explicitly, the transverse momentum p_T of the jet is balanced by missing transverse momentum E_T^{miss} . This is motivated by several theories for physics beyond the SM which predict the existence of new massive states which only weakly interact with the SM, and which are at least stable enough to escape the detector volume without decaying. These new states would thus be observed as E_T^{miss} balancing the jet's p_T . There are then two sub-sets within this scheme, where the jet either comes from Initial State Radiation (ISR) and the new state is pair-produced, or where the jet is produced in association to or from the decay of a new state. Example Feynman diagrams for several relevant models of Beyond the Standard Model (BSM) processes which give rise to the mono-jet topology are provided in Figure 7.1.

This simple single-jet topology was first investigated by the UA1 Collaboration [72], and was continued at the Tevatron by the CDF Collaboration [73]. The first iteration of the ATLAS mono-jet search using the 2010 dataset also made use of a single-jet selection, as detailed in Reference [74]. However, CMS had already switched to allowing for a second jet, as per Reference [75]. CMS had noticed that insisting on a final state with only a single jet was often too restrictive.

Mono-jet processes may only consist of a single parton in the final state at truth level, as per the Feynman diagrams in Figure 7.1, but these partons frequently split into multiple observed jets. There can also be additional ISR emissions beyond the hard process, which increases the jet multiplicity. Given that these signals tend to be more energetic processes requiring larger momentum transfers than the dominant SM backgrounds, due to reasons such as the production of heavy states or requirement for a mediator much heavier than the W and Z bosons, the E_T^{miss} spectrum of the new phenomena is usually harder. As such, the frequency with which an additional ISR or split jet is above analysis p_T thresholds is higher for the signal than for the SM backgrounds. The sensitivity of the analysis to many generic signals was thus enhanced by allowing for A second jet, and as such both the ATLAS and CMS 2011 publications made use of this modified mono-jet topology, as detailed in References [76, 77] respectively.




Figure 7.1: A subset of the Feynman diagrams for Beyond the Standard Model physical processes which give rise to the mono-jet topology. The production of new states in association with a jet are shown on top, for the ADD Large Extra Dimensions Graviton (top left) and SUSY Gravitino (top right). The pair-production of DM with ISR jet emission is shown for quark EFT operators (middle left), gluon EFT operators (middle right), and a Z' mediator (bottom).

This was continued in 2012, with the ATLAS preliminary publication on the first half of the dataset and the CMS final publication on the full dataset making use of such a selection, as per References [6, 78].

In June 2013, the DM contributors to the ATLAS mono-jet group pushed this even further, and investigated the sensitivity of the analysis with a more inclusive jet multiplicity selection. It was discovered that the DM interpretation benefited from removing the jet multiplicity requirement. This was later noticed independently by the theory community, as first detailed in Reference [79]. As such, the mono-jet group began to investigate the possibility of removing the jet veto, and carried out the full analysis for both the typical 1 or 2 jet case and the inclusive scenario. By July 2014, it had been observed that all of the interpretations considered had a higher expected sensitivity when using the inclusive jet multiplicity selection, and thus the old 1 or 2 jet analysis was discontinued.

The name of the analysis has thus become a misnomer, as the final state has changed to become a mono-jet-like topology. The idea remains the same, the final state is comprised of a single dominant jet balancing the E_T^{miss} , but now there can be any number of additional low-energy jets present in the event. While this Chapter describes the mono-jet-like selection, which was eventually published, the terms mono-jet-like and mono-jet will typically be used interchangeably. The more typical mono-jet selection is a subset of the mono-jet-like selection, and is simpler on several fronts, as will be pointed out where appropriate.

The primary SM backgrounds will be discussed in Section 7.2. An overview of the selection and region definitions is provided in Section 7.3, while full details are provided in Section 7.4. The dataset and simulation details are provided in Section 7.5. The determination of the relevant backgrounds is discussed in Section 7.6 and associated systematics are presented in Section 7.7. The results and conclusions of the analysis are in Sections 7.8 and 7.9.

7.2 Standard Model backgrounds

The primary source of SM backgrounds resulting in mono-jet events is the $Z \rightarrow \nu\nu + \text{jets}$ final state. This accounts for between 50 and 75% of the total background, depending on the E_T^{miss} threshold used. It is an irreducible background, given that neutrinos are weakly interacting particles, and thus the signature is very similar to the expected signals. A precision measurement of the $Z \rightarrow \nu\nu + \text{jets}$ background is of paramount importance, and this is done through the combination of four separate semi data-driven estimations using four orthogonal control regions. These control regions are designed to select $Z \rightarrow \mu\mu + \text{jets}$, $W \rightarrow \mu\nu + \text{jets}$, $Z \rightarrow ee + \text{jets}$, and $W \rightarrow e\nu + \text{jets}$ events. These are all very similar processes to $Z \rightarrow \nu\nu + \text{jets}$ and, with the right treatment, provide an excellent estimate of the number of $Z \rightarrow \nu\nu + \text{jets}$ events to expect. The data-driven estimation procedure is detailed in Section 7.6.1, and the level of agreement between the four estimates is specifically discussed in Section 7.6.4.

The second largest source of SM backgrounds to the signal region is the $W \rightarrow \ell\nu$ family of process. The neutrino provides a source of true E_T^{miss} , which allows the events to pass the mono-jet selection. However, the lepton which is produced with the neutrino provides a means of identifying and rejecting these events, making it a reducible background. If the lepton is an electron or muon, then the event can pass the selection if the lepton is either mis-identified, or outside of the detector acceptance. The scenario involving the τ lepton is more complicated. If it decays leptonically, then the same possibilities exist. However, if the τ decays hadronically, then the resulting jet is not distinguished as separate from the expected signal jet, and thus will naturally be selected. These three backgrounds are all estimated

from control regions, with the $W \rightarrow e\nu + \text{jets}$ and $W \rightarrow \mu\nu + \text{jets}$ contributions to the signal region being estimated from the respective control regions, and the $W \rightarrow \tau\nu + \text{jets}$ estimate also coming from a slightly modified variant of the $W \rightarrow e\nu + \text{jets}$ control region. All of these backgrounds will be described in the relevant part of Section 7.6.

The multi-jet (QCD) background is a priori dominant in comparison to even the $Z \rightarrow \nu\nu + \text{jets}$ process simply due to the enormous cross-section at hadron colliders. Such events pass the mono-jet selection when a highly energetic jet is either missed or significantly mis-reconstructed, giving rise to a large source of fake E_T^{miss} . Specialized cuts are found to very efficiently remove this background, rendering it negligible, but a data-driven estimation is nonetheless conducted in order to ensure that the process is well understood. This procedure is detailed in Section 7.6.5.

The Non-Collision Background (NCB) background also has the potential to be dominant. Events with a single jet and a large amount of E_T^{miss} are expected from NCB sources, and in fact a similar topology is often used to study NCB processes. A dedicated set of cleaning criteria has been developed, as detailed in Section 7.6.6. These cuts are found to reduce the level of NCB contamination in the signal region to negligible levels.

With a 1 or 2 jet requirement mono-jet selection, the remaining backgrounds were known to be small simply due to a combination of their cross-section and the large number of jets which must be missed to pass the selection. This has changed due to the use of the new mono-jet-like criteria. The single- t and $t\bar{t}$ backgrounds are a perfect example. These processes result in a large number of jets, which would be typically rejected by a 1 or 2 jet requirement. Now that this requirement has been removed, these backgrounds are no longer naturally rejected. The contribution of these processes is still found to be small in the signal region, but some of the W and Z control regions see contamination at the level of 10%. As such, a top control region was developed to get a better handle on the associated systematics, as described in Section 7.6.7.

The diboson background of WW , WZ , or ZZ is another example of a final state which becomes a more relevant background due to relaxing the third jet veto. It remains only a small background in the signal region, thanks to the lepton vetoes and high E_T^{miss} requirement, but can become larger in the control regions. In all cases, the diboson contribution naturally passes the mono-jet selection, especially if one of the bosons is a Z which decays to two neutrinos, and the other boson decays to two jets (or two leptons or a lepton and a neutrino for the control regions). Given that sufficient statistics were not available to construct a control region specifically for the diboson contribution, it is important to quantify the associated uncertainties, as is done in Section 7.6.8.

Of the remaining small SM backgrounds, only $Z \rightarrow \ell\ell + \text{jets}$ is further considered. This is a minor contribution to the signal region, as both leptons must either be mis-identified, out of acceptance, or hadronically decaying τ leptons. This is a small fraction of events, and the process is easily rejected. MC without any data-driven estimation is used for this small background, as detailed in Section 7.6.8.

7.3 Signal and control region overview

The mono-jet analysis makes use of a signal region and several control regions, all of which are evaluated at multiple kinematic thresholds. Before delving into the details in Section 7.4, this overview will provide context for what will be discussed.

The regions are named after their intent. The signal region is defined by the selection used to quantify

the possible presence of new physics signals. The different control regions are instead regions where no new physics is expected, and thus they can be used to estimate (control) the contribution of a given process. The $Z \rightarrow \mu\mu$ and $Z \rightarrow ee$ control regions are optimized to select the topology specified in the name (plus jets), and are used for estimating the contribution of the $Z \rightarrow \nu\nu + \text{jets}$ in the signal region, which is possible due to the similar topologies. The $W \rightarrow \mu\nu$ control region selects $W \rightarrow \mu\nu + \text{jets}$, and is used to estimate both the $Z \rightarrow \nu\nu + \text{jets}$ and $W \rightarrow \mu\nu + \text{jets}$ contributions to the signal region. The $W \rightarrow e\nu [Z]$ control region selects $W \rightarrow e\nu + \text{jets}$ events and is used to estimate the $Z \rightarrow \nu\nu + \text{jets}$ process in the signal region, while the $W \rightarrow e\nu [W]$ control region is used to estimate the $W \rightarrow e\nu + \text{jets}$ and $W \rightarrow \tau\nu + \text{jets}$ processes in the signal region. The difference between the $W \rightarrow e\nu [Z]$ and $W \rightarrow e\nu [W]$ regions will be detailed later, but for now it suffices to know the purpose of the two different but similar regions. These represent the most important control regions used for the data-driven estimation of dominant backgrounds. Other control regions will be defined in the Section to which they are relevant.

All regions require general data quality and cleaning requirements. The primary observable varies, with the E_T^{miss} being used for the signal region and the W or Z boson p_T for the respective $W \rightarrow \ell\nu$ and $Z \rightarrow \ell\ell$ control regions. All regions require a central energetic jet of high quality which is not aligned with the primary observable. Either a E_T^{miss} or single electron trigger is used to select data. The signal region requires there to be zero electrons or muons, while the control regions require the presence of two appropriate leptons for the $Z \rightarrow \ell\ell$ regions or one appropriate lepton for the $W \rightarrow \ell\nu$ regions. The control regions also apply boson-related selections in order to improve the purity of the selected process, specifically Z boson invariant mass $m_{\ell\ell}$ cuts or W boson transverse mass m_T^ℓ cuts.

The following primary observable cuts are considered, defining nine sets of signal and control regions: $\{150, 200, 250, 300, 350, 400, 500, 600, 700\}$ GeV. The cuts are applied on either the E_T^{miss} (signal regions) or boson p_T (control regions). The lowest threshold is motivated by the point at which the E_T^{miss} trigger becomes 98% efficient, as detailed in Section 7.4.1. The remaining values are chosen to maximize the sensitivity to many different models with different E_T^{miss} spectrum shapes, while the final threshold was chosen in order to still have a reasonable amount of data statistics available.

Other assorted cuts are applied for extra background rejection or similar reasons, but the above are the primary selection criteria which form the basis of the region definitions. Now the specific details will be discussed for each step of the selection.

7.4 Event selection and cleaning

Arguably the most critical part of any analysis is its event selection. This selection defines the analysis, controlling the dataset to be investigated, and what topologies and processes are selected. The dominant backgrounds can be shaped and changed with the selection, and signal sensitivity can be enhanced with optimized cuts. The mono-jet analysis is a generic search for new weakly interacting phenomena with several interpretations. For this reason, the selection is optimized to generically reject backgrounds and retain sufficient statistics in the control regions at high E_T^{miss} rather than toward any particular signal interpretation. Additional sensitivity to each signal is retained by the use of numerous E_T^{miss} thresholds, where E_T^{miss} is the dominant observable and different signals prefer different regions of the E_T^{miss} spectrum.

Given that the mono-jet analysis involves 5 control regions and one signal region, each of which has 9 different kinematic thresholds and thus sub-regions, it is important to impose a unified approach

whenever possible. This is reflected in a common trigger, object definition, and baseline selection across all regions except for when specific needs dictate otherwise. This is explained in more detail in the following Sections.

In addition to the event selection, it is important to consider the interrelated topic of event cleaning. While most recorded collisions are useable, there are some defective events. Sometimes the event itself is corrupted in some way, or subsections of the ATLAS detector were not fully operational. Alternatively, objects may be improperly identified or reconstructed. In all cases, it is important to ensure that only well-understood and problem-free events are selected.

7.4.1 Triggering events

The trigger is a crucial first step in the selection of events. The trigger controls which collision events are recorded, and thus any analysis must first find a trigger which suits their topology. For the mono-jet analysis, the trigger chosen is based purely on E_T^{miss} . This is a calorimeter based E_T^{miss} , with no knowledge of the muon system, thus the trigger will also fire on events with energetic muons balanced by a jet. This will be later exploited for the muon control regions. The technical details of the implementation of E_T^{miss} triggers at ATLAS is discussed in Reference [80].

The trigger used for the mono-jet analysis is named `EF_xe80_tc1cw`, which was unrescaled for the entire 2012 dataset. This is a topo-cluster (tc) based trigger in which the topo-clusters have had LC weights (lcw) applied. The trigger threshold is 80 GeV of E_T^{miss} (xe80) at the Event Filter (EF) level, where the EF is a high-level trigger which makes the final decision of what is written to disk and what is rejected. The pre-EF steps are also E_T^{miss} triggers, with respective thresholds of 60 GeV and 65 GeV at trigger levels 1 and 2 respectively. The level 2 trigger is also based on LC topo-clusters, while the level 1 trigger uses a more simplistic reconstruction of the E_T^{miss} .

At level 1, the calorimeter is scanned for clusters of energy which are above pre-set thresholds intended to suppress the contribution of noise. This is done by splitting the calorimeter into fixed-size trigger towers, which are primarily 0.1×0.1 in (η, ϕ) space, although this becomes 0.2×0.2 in the forward region due to reduced calorimeter granularity. These trigger towers are then combined into RoIs, which are 0.4×0.4 in (η, ϕ) for this particular trigger. All of the observed RoIs are then summed to get the total transverse momentum of the event, where the E_T^{miss} has the same magnitude but points in the opposite direction [28].

The trigger threshold of 80 GeV does not mean that all analysis events of 80 GeV of E_T^{miss} are retained. The E_T^{miss} calculation in the online EF is not the same as what is used in offline reconstruction, in which objects receive proper calibrations rather than only making use of LC topo-clusters. This difference means that an analysis requirement of $E_T^{\text{miss}} > 150$ GeV is necessary in order for the specified EF trigger to recognize 98% of events.

Trigger efficiency

In order to measure the trigger efficiency, the lack of muon information in the EF trigger definition is exploited. By using the full Muon data stream (datasets where the muon triggers are applied), and by applying a simple $W \rightarrow \mu\nu + \text{jets}$ selection as detailed in Table 7.1, the E_T^{miss} turn-on curve can be measured. This selection is not pure in $W \rightarrow \mu\nu + \text{jets}$ events in data, thus a combination of $W \rightarrow \mu\nu + \text{jets}$, diboson, and $t\bar{t}$ MC samples were used to measure the turn-on in MC in order to match

Jets		Muons	
At least one jet, medium quality		At least one muon, tight quality, tight isolation	
leading jet p_T	> 100 GeV	Trigger	EF_mu24i_tight OR EF_mu36_tight
p_T^{jet}	> 30 GeV	p_T^μ	> 25 GeV
$ \eta^{\text{jet}} $	< 4.5	$ \eta^\mu $	< 2.5

Table 7.1: A simplified selection which enhances the contribution of $W \rightarrow \mu\nu + \text{jets}$ events. This was used for measuring the E_T^{miss} trigger turn-on curve in both data and MC.

	Data	$W \rightarrow \mu\nu + \text{jets}$	Diboson	$t\bar{t}$	MC combination
Sample Size	2.8×10^6	4.92×10^6	8.79×10^4	4.6×10^5	5.47×10^6
Before Trigger	2.41×10^6	1.73×10^6	3.12×10^4	3.12×10^5	2.07×10^6
After Trigger	2.01×10^6	1.66×10^6	3.00×10^4	3.72×10^5	1.96×10^6
Max Efficiency	99.7%	99.9%	99.2%	99.2%	99.6%
$\epsilon_{E_T^{\text{miss}}} = 98\%$	151 GeV	133 GeV	136 GeV	200 GeV	145 GeV
χ^2/N_{dof}	0.34	0.11	0.40	0.35	0.12

Table 7.2: Event yields, trigger efficiencies, and fit qualities for the EF_xe80_tc1cw trigger. Data reaches the 98% efficiency point at 151 GeV, while the MC combination reaches the same point slightly earlier. Due to these results, $E_T^{\text{miss}} > 150$ GeV was deemed to be a reasonable point for both data and MC to be at the 98% efficiency point or higher and thus above the trigger turn-on.

data expectations. The turn-on curve is then fit to the Fermi function, chosen due to the resulting small χ^2/N_{dof} , over the range of $100 \text{ GeV} \leq E_T^{\text{miss}} \leq 400 \text{ GeV}$.

The trigger efficiency turn-on curves for data as compared to $W \rightarrow \mu\nu + \text{jets}$, diboson, $t\bar{t}$, and the combined MC sample set are provided in Figure 7.2. The resulting efficiency values and points at which the efficiency passes 98% are provided in Table 7.2.

Electron control region triggers

For some of the data-driven estimates from the electron control regions, the E_T^{miss} trigger described above was not appropriate. Instead, a combination of two single electron triggers was used. This particular combination is supported by the associated performance group, and thus the trigger efficiencies are provided, meaning there is no need to determine them independently. The two EF triggers in question correspond to 24 GeV *medium++* electrons with very high isolation requirements or 60 GeV *medium++* electrons with no isolation requirements, named EF_e24vhi_medium1 and EF_e60_medium1 respectively. This trigger combination is used for the $Z \rightarrow ee$ control region and the $W \rightarrow e\nu$ [W] control region. The $W \rightarrow e\nu$ [Z] control region still makes use of the E_T^{miss} trigger, as will be motivated in Section 7.6.3.

7.4.2 Event quality

After the trigger is selected, the subset of recorded events accessible to the analysis has been decided. Unfortunately, not all events were recorded under ideal conditions. Detector malfunctions, corrupted event records, and other such problems are present for approximately 10% of events. These bad events must be removed through the application of a set of event quality requirements, which are mandatory for all typical ATLAS analyses. This includes the application of the aforementioned Good Runs List (GRL) to remove detector problems, as well as various event and detector quality flags which signify other problems.

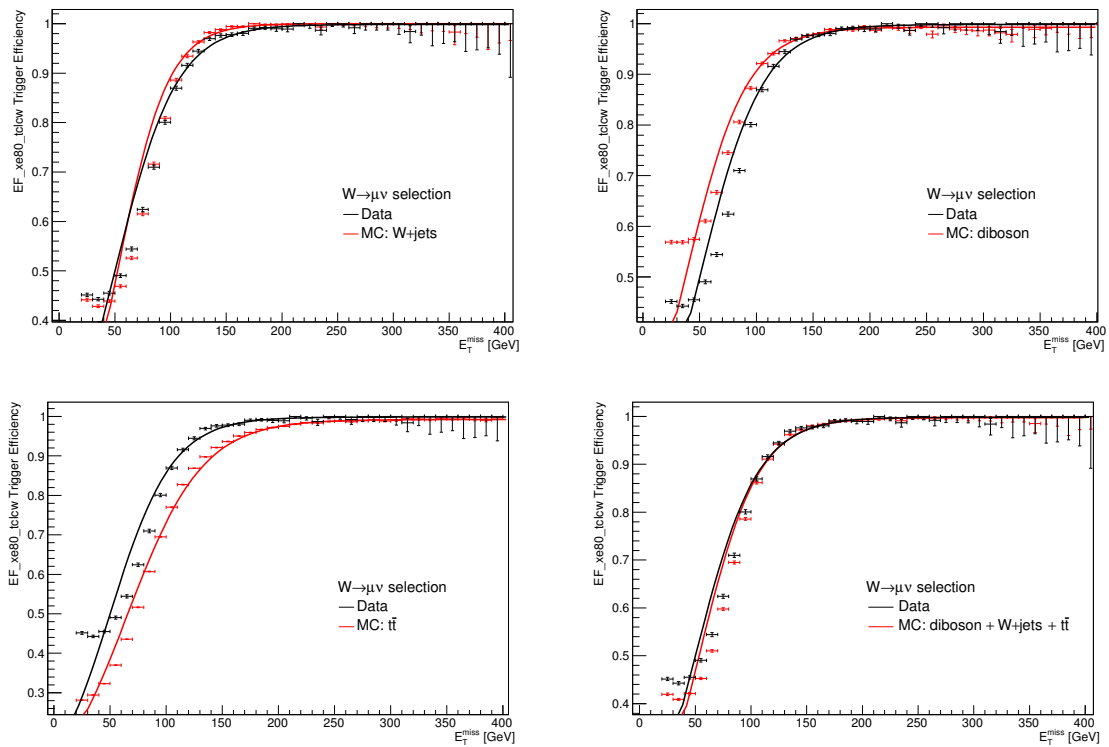


Figure 7.2: Trigger efficiency turn-on curves for the EF_{xe80_tclw} trigger. Data and the combined MC both reach the 98% efficiency point at approximately $E_T^{miss} > 150$ GeV, while sub-dominant MC processes can individually take longer to become fully efficient.

7.4.3 Object definitions

The mono-jet analysis explicitly makes use of electrons, muons, jets, and E_T^{miss} , both for selecting and vetoing events. Multiple definitions are considered for each object, for use in different signal and control regions, as will be described. In general, the signal region should contain only jets and E_T^{miss} (veto on electrons and muons), while the control regions are defined such that they are orthogonal via requiring electrons and muons.

Electron definitions

Three types of electrons are considered in the mono-jet analysis, hereafter referred to as three electron definitions. Each level of selection is applied on top of the previous cuts, thus a collection of electrons passing the most strict definition creates a subset of the least restrictive electron collection. The weakest definition is the veto quality electron, which is used for rejecting events when applicable. This variant of the electron is used to reject events containing electrons in the signal region, or to reject events with too many electrons in the control regions. Given that there should be no additional electrons in these events when selecting a specific topology, a loose definition is suitable for rejecting these potentially dangerous events. As an example, the $Z \rightarrow ee$ control region requires exactly two electrons, as the $Z \rightarrow ee + \text{jets}$ process consists of two electrons. If three veto quality electrons are observed, then the event is rejected to remove the possible contribution from tri-electron processes.

Veto definition electrons require a match between a track in the inner detector and a cluster of energy in the electromagnetic calorimeter. If the cluster falls into an inactive or only partially active region of the calorimeter, then the object is not considered to be an electron. The object is also required to pass the *medium++* quality requirements, as explained in Reference [59]. Beyond that, the only requirements for the veto electron definition is that it have a transverse momentum of $p_T > 7 \text{ GeV}$ in order to be in a well understood kinematic regime, and to have $|\eta| < 2.47$ to be inside the inner tracking detector coverage.

The next level of selection, referred to as a medium electron definition, is used for events where some other factor can be used to suppress dominant backgrounds. This applies to the $Z \rightarrow ee + \text{jets}$ control region, where the requirement that the invariant mass of the two electrons falls under the Z mass peak serves to reject all other backgrounds. This definition provides an identification of good electrons, but without sacrificing statistics to achieve an unnecessary degree of background rejection. This is also used for the $W \rightarrow e\nu + \text{jets}$ control region when estimating the $W \rightarrow e\nu + \text{jets}$ and $W \rightarrow \tau\nu + \text{jets}$ contributions to the signal region, due to the presence of a large truth E_T^{miss} cut for such estimates, which helps to significantly reduce other backgrounds.

The additional selection criteria associated with the medium electron definition raises the p_T cut from 7 to 20 GeV in order to reduce the contribution of fakes and mis-measurements and excludes a crack in the calorimeters between $1.37 < |\eta| < 1.56$. In addition, overlap removal is now performed between electrons and jets, as detailed in Section 7.4.4.

The third and most strict set of criteria, used where high background rejection is required, is the tight definition. This is used for events where electrons are required without any additional variable available to provide significant background discrimination. This is particularly true for the case of the $W \rightarrow e\nu + \text{jets}$ control region when estimating the contribution of $Z \rightarrow \nu\nu + \text{jets}$ to the signal region. In this case, the E_T^{miss} cut is replaced by a cut on the boson p_T . If the electron is energetic, then the

Definition	Veto	Medium	Tight
p_T cut	7 GeV	20 GeV	25 GeV
η cut	$ \eta < 2.47$	$ \eta < 1.37$ or $1.56 < \eta < 2.47$	$ \eta < 1.37$ or $1.56 < \eta < 2.47$
Electron quality	<i>medium++</i>	<i>medium++</i>	<i>tight++</i>
Isolation	None	None	Track and calorimeter
Overlap removal	No	Yes	Yes

Table 7.3: A summary of the three types of electron definitions in use in the mono-jet analysis. Veto electrons are used for rejecting events with more electrons than expected for the topology in question, medium electrons are used for selecting events with electrons that have powerful additional background suppression cuts, and tight electrons are used for selecting events with electrons without additional strong background discrimination. The *medium++* and *tight++* selection criteria are defined in Reference [59].

neutrino from the W boson decay can be low p_T , and these events need to be retained. However, the enormous QCD background enters here. Without the suppression associated with an invariant mass cut or high true E_T^{miss} threshold, the boson p_T cut becomes a cut on the electron p_T , which can then be a cut on a mis-reconstructed jet (such as from the enormous di-jet background). A very strict electron selection criteria allows for a reduction in this fake rate, and is one of multiple techniques that will be explained in Section 7.6.3.

The tight electron definition comes with much more stringent cuts, including the switch from a *medium++* quality to the *tight++* electron quality criterion. The electron p_T threshold is raised to 25 GeV. The most radical difference is with respect to the isolation of the object, which pertains to the amount of energy in the surroundings which is not attributed to the electron. Within the inner detector, the isolation requires that the scalar sum of transverse momenta attributed to tracks within a cone of $\Delta R < 0.3$ around the electron is less than 5% of the measured electron momentum. The calorimeter isolation is similarly defined with EM scale calorimeter clusters around the selected electron cluster, with the difference that the calorimeter clusters are also corrected for energy leakage and pileup calorimeter energy deposits.

A summary of the three levels of electron definitions is provided in Table 7.3.

Muon definitions

Analogous to electrons, there are multiple muon definitions in use by the mono-jet analysis. The veto definition remains, but there is no problem related to jets being reconstructed as fake muons, thus only one alternative definition is required. The veto definition makes use of muons reconstructed with the STACO algorithm, as defined in Reference [60]. Additionally, the muons must be either *combined* (good agreement is observed between momentum measurements in the inner detector and the muon spectrometers) or *segment tagged* (a well measured track in the inner detector points to segments of activity within the muon spectrometer, but without an independent track reconstructed in the muon system). The inner detector track must be of a high quality, passing the following requirements:

- Pixel detector activity: number of pixel hits + number of dead pixel sensors crossed > 0
- SemiConductor Tracker activity: number of SCT hits + number of dead SCT sensors crossed > 4
- Good inner detector region: number of pixel holes + number of SCT holes < 3
- Transition Radiation Tracker activity if applicable (only within $0.1 < |\eta| < 1.9$):

Definition	Veto	Tight
p_T cut	7 GeV	20 GeV
η cut	$ \eta < 2.5$	
Muon quality	<i>segment tagged or combined</i>	

Table 7.4: A summary of the two types of muon definitions in use in the mono-jet analysis. Veto muons are used for rejecting events with more muons than expected for the topology in question, while tight muons are used for selecting events with real expected muons.

- Number of TRT hits + number of TRT outliers > 5
 - * TRT outliers are TRT hits which are only loosely associated with the expected track, rather than providing a good match
- Number of TRT outliers $< 0.9 \times (\text{number of TRT hits} + \text{TRT outliers})$

An isolation requirement on the muons is also required to prevent vetoing on muons from in-flight decays of heavy flavour quarks or similar. The scalar sum of all tracks within $\Delta R < 0.2$ of the muon track must total less than 1.8 GeV of transverse momentum.

Veto muons are required to have a transverse momentum of at least 7 GeV to stay within the well understood regime, while the region of $|\eta| < 2.5$ is defined such that the inner detector is available. The only difference for the tight muon definition is to increase the momentum cut to $p_T > 20$ GeV.

A summary of the two levels of muon definitions is provided in Table 7.4.

Jet definition

The mono-jet analysis applies identical jet selection cuts to all of the signal and control regions, thus only a single type of jet definition is required. The jets considered are built with the anti- k_t algorithm using LC topo-cluster inputs with a distance parameter of $R = 0.4$. The full jet calibration scheme is applied, corresponding to the Moriond2013 release of the JES calibration, as described in Section 6.3. This includes the Jet Areas approach to pileup suppression, the MC JES correction, and the *in situ* combination scale factor. This results in jets at the LC+JES scale, which forms the baseline for all further analysis uses.

Jets are considered to be analysis objects of interest if the LC+JES transverse momentum is at least 30 GeV and the jet is fully contained within the calorimeter for the chosen distance parameter, or in other words $|\eta| < 4.5$. It is possible that jets overlap with other calorimeter objects, such as electrons. In such cases, a choice as to whether the object is an electron or jet must be made, as is discussed in Section 7.4.4. Note that this overlap was already mentioned for the E_T^{miss} reconstruction procedure described in Section 5.5.1, but this must be done again at analysis-level when selecting or rejecting events based on the presence or absence of jets.

Missing transverse momentum and boson transverse momentum definitions

The same definition of the E_T^{miss} is used for all signal and control regions in the mono-jet analysis, although the E_T^{miss} is not always the observable which is used for cuts, as will be explained shortly. The E_T^{miss} algorithm used is based on fully calibrated objects at the optimal scale as described in Section 5.5.1, and is built from electrons, photons, jets, and residual soft topo-clusters. Muons are explicitly excluded to make this a pure calorimeter E_T^{miss} , bringing it closer to what is used in the trigger, and to simplify

the use of the muon control regions. Given that taus are not identified in the mono-jet analysis, it was decided to not explicitly select them when building the E_T^{miss} . As such, the tau contribution to the E_T^{miss} generally is encapsulated in the jet term, thus providing a small precision degradation in events containing taus due to the use of the JES instead of the tau energy scale.

The E_T^{miss} is first built from the selected fully calibrated hard objects (electrons, photons, and jets), before the remaining topo-clusters are combined to form the soft term, as detailed in Section 5.5.1. This provides the consistent definition of the E_T^{miss} used for all regions.

In the $Z \rightarrow \mu\mu$ and $W \rightarrow \mu\nu$ control regions, the E_T^{miss} definition chosen provides a useful equivalence to the boson p_T . For $Z \rightarrow \mu\mu + \text{jets}$ events, the vector sum of the transverse momenta of the two muons is the boson p_T . However, there is also no real source of full-detector E_T^{miss} in these events, thus the calorimeter E_T^{miss} is equivalent to the boson p_T . Similarly, for $W \rightarrow \mu\nu + \text{jets}$ events, the vector sum of the neutrino p_T and muon p_T is the boson p_T . Thus, the vector sum of the full detector E_T^{miss} (the neutrino) and the muon p_T , or equivalently the chosen calorimeter E_T^{miss} , is the boson p_T . In other words, in muon control regions, the E_T^{miss} is equivalent to the relevant boson p_T .

The real observable of interest when estimating $Z \rightarrow \nu\nu + \text{jets}$ is the boson p_T , as the aim is to provide an equivalence between the Z in the signal region and either the W or Z in the control regions. In the signal region, the only possible handle is the E_T^{miss} , as the neutrinos do not interact with the detector. In the muon control regions, as mentioned, the use of a calorimeter E_T^{miss} and the property of muons being MIPs at the LHC energy scale results in an equivalence between boson p_T and E_T^{miss} .

This is not as simple in electron control regions, as the boson p_T is not equivalent to the chosen E_T^{miss} variant. Instead, the transverse momentum of the electron(s) is/are added to the E_T^{miss} , thus turning the sum into the boson p_T . This provides the needed link between the signal and control regions.

In summary, the E_T^{miss} definition is equivalent across all signal and control regions. However, sometimes the related quantity of the boson p_T is the observable, not the E_T^{miss} . The boson p_T and E_T^{miss} are identical for muon control regions, but the electron control regions require the transverse momenta of any electron(s) to be added to the E_T^{miss} to form the boson p_T .

7.4.4 Overlap removal

The jet algorithm groups all topo-clusters without any type of consideration of whether or not any given topo-cluster has already been used in the reconstruction of another type of object, a track from the inner detector, or any other association or identification criteria. For this reason, it is expected that jets and other objects share the same underlying energy deposits. This double-counting of energy is known as overlap, and must be resolved for a consistent treatment of the event kinematic quantities.

The only overlap removal considered in the mono-jet analysis is that of electrons and jets. When a jet and medium or tight definition electron are in the same region of the calorimeter, with $\Delta R(\text{jet}, \text{electron}) < 0.2$, the reconstructed object is determined to be an electron. As such, the corresponding jet is removed in order to resolve the double-counting.

7.4.5 Jet cleaning

As mentioned, jets are general objects built from calorimeter energy deposits with no specific quality criterion built-in. They are not physical objects, but rather representations of the observed activity within the calorimeters. As such, it is important to determine whether the jets considered in the

analysis correspond to a physical process or other unrelated effects.

A standard set of jet cleaning criteria, which is designed to identify undesired calorimeter activity, is applied. This includes rejection of jets from energy spikes in the HEC, coherent noise in the electromagnetic calorimeters, and NCB sources. This selection is applied to all jets as defined previously, except at a p_T threshold of 20 GeV instead of 30 GeV. If any jets fail this cleaning criteria, the event is rejected. This fake jet rejection criteria is referred to as the *looser* selection, and further details are provided in Reference [69].

This standard set of jet cleaning criteria is generally both highly efficient and sufficiently powerful for ATLAS analyses. However, the mono-jet analysis is particularly sensitive to NCB sources due to the reliance on a E_T^{miss} trigger and single-jet topology. Beam backgrounds in particular are a major source of contamination, and mono-jet-like selections are typically used for studying and quantifying these undesired events. For example, beam halo events typically appear as a single very high p_T jet passing through the calorimeter at an angle that does not point back to the interaction point. This single high p_T beam jet is unbalanced, providing a large source of fake E_T^{miss} , and therefore causing the trigger to fire.

In order to reject this otherwise considerable background, a dedicated study into NCB backgrounds was performed, resulting an additional set of cuts to be applied only to the leading jet in the event. If the leading jet failed this requirement, then the event was rejected. The derived cleaning criteria reduces the level of NCB contamination to negligible levels. A summary of these studies and the resulting performance is provided in Section 7.6.6, where the end result is a leading jet cut of $f_{\text{ch}}/f_{\text{max}} > 0.1$.

7.4.6 Masked Tile calorimeter modules

The problems associated with jets containing masked Tile calorimeter modules is detailed in full in Appendix C. In short, when the centroid of a jet is located within a masked module, the jet p_T is under-estimated. This mis-estimation can be up to a factor of 50% at high p_T , and thus it provides a significant source of fake E_T^{miss} aligned with the jet. A cut on $\Delta\phi(E_T^{\text{miss}}, \text{jets})$ provides for an efficient way to remove such fake E_T^{miss} sources, so long as there is no real E_T^{miss} in the event. If there is a source of real E_T^{miss} , then the total E_T^{miss} angle will be balanced between the real E_T^{miss} angle and mis-measured jet angle, thus the effect is less effective. This is particularly important for the mono-jet analysis, given that the dominant $Z \rightarrow \nu\nu + \text{jets}$ background consists of high real E_T^{miss} in addition to high p_T jets which can be mis-measured.

Jets with their centroid located adjacent to the masked module are also mis-estimated by up to 50% at high p_T , but in the opposite direction. This once again creates a source of fake E_T^{miss} , but this time the E_T^{miss} is back-to-back with the jet. As such, the $\Delta\phi(E_T^{\text{miss}}, \text{jets})$ cut doesn't help, as this is the same signature as the expected undetected signal particle(s) recoiling off of a single jet.

Both of these problems are addressed through the use of specialized cleaning criteria defined in Appendix C. Given that the mono-jet analysis makes uses of very high p_T jets, and is particularly sensitive to fake E_T^{miss} , the *tight* (geometric) cleaning is applied to the two leading jets. This ensures that the most energetic jets in the analysis are not significantly mis-modelled due to masked Tile calorimeter modules. The remaining jets, if present, are only required to pass the *medium* cleaning selection. This reduces the amount of integrated luminosity which is rejected by the cleaning cuts, and is reasonable given that jets beyond the leading two are typically relatively soft, where the impact of any bias is smaller than the large low- p_T JER (due to the noise, primarily from pileup) and thus is invisible.

7.4.7 Isolated track rejection

As detailed in Section 7.2, the largest reducible backgrounds in the mono-jet analysis signal region are related to the $W \rightarrow \ell\nu + \text{jets}$ processes where the electron or muon are missed, or the lepton is a tau. The tau background in particular is sizable, accounting for up to 25% of the signal region background, due to the hadronic decay mode naturally evading the lepton veto.

Electrons, muons, and leptonic tau decays (35% of taus) are all notable in that they produce a single isolated charged track in the inner detector. These cases will already be rejected by the signal region lepton veto if the lepton is identified, however this is not always the case. For any number of reasons, the lepton may not be reconstructed properly, but the single isolated track will remain. So long as this track passes a set of quality criteria, it is a useful means of further rejecting leptonic events in the signal region.

The remaining 65% of tau decays are hadronic, primarily to a mix of neutral and/or charged pions, as discussed in Section 5.4.1. Unlike quarks or gluons produced in the hard-scatter interaction, pions are already bound states, and do not undergo hadronization. As such, the track multiplicity for a charged (neutral) pion within the ATLAS inner detector is typically one (zero) as compared to the large multiplicity which can be expected from partons. If the tau decay is single-pronged, meaning that there is only one charged pion (46% of taus), then a single isolated track is expected. Once again, an isolated track veto provides a means of further extending the mono-jet lepton veto.

While isolated tracks indicate the presence of leptons including single-prong hadronic tau decays, they are not the only source. Low p_T jets beyond the dominant leading jet form a substantial background, especially as pileup levels increase. For these jets, the leading hadron in the jet may have enough energy to pass the tracking threshold, while sub-dominant hadrons fall below the p_T threshold. In this case, a single track is also expected. A set of quality criteria must thus be defined such that such signal events with single-track jets are not rejected.

The first set of quality requirements is applied to the candidate lepton track. The track is required to contain at least 10 GeV of p_T , be within the tracking acceptance, and fulfil a set of impact parameter, hit, and fit quality requirements. The second quality requirement is the threshold at which a track is determined to be isolated. The optimal criterion was determined to be a requirement of no other tracks of $p_T > 3 \text{ GeV}$ within $\Delta R < 0.4$ of the candidate track. If a track exists which passes the quality requirements and has no other tracks above this threshold in the surrounding area, and if the selection is not for a control region (in which case the track also cannot be for the selected leptons), then the event is rejected.

A summary of the track veto selection is provided in Table 7.5. This veto is found to reduce the level of $W \rightarrow e\nu + \text{jets}$, $W \rightarrow \mu\nu + \text{jets}$, and $W \rightarrow \tau\nu + \text{jets}$ events in the signal region by 22%, 23%, and 28% respectively. It is found to also reject 5% of $Z \rightarrow \nu\nu + \text{jets}$ events, where no lepton is expected, thus providing an estimate of the inefficiency of the procedure on signal-like events. This is found to improve the signal-like/background event ratio by approximately 5% overall, with the use of $Z \rightarrow \nu\nu + \text{jets}$ as the signal-like process. The systematic uncertainty associated with the procedure is found to be at the sub-percent level, thus a conservative 1% uncertainty is used.

Selection type	Requirements
Kinematics	$p_T > 10 \text{ GeV}, \eta < 2.5$
Impact parameters	$ z_0 < 2 \text{ mm}, d_0 < 1 \text{ mm}$
Track hits	$N_{\text{hits}}^{\text{Pixel}} + N_{\text{hits}}^{\text{SCT}} \geq 5$
Fit quality	$\chi^2/N_{\text{dof}} < 3$
Isolated track	$N_{\text{track}}^{p_T > 3 \text{ GeV}} > 0$
Expected leptons	Ignore one track per expected lepton (control regions)

Table 7.5: Quality requirements for the selection used to determine if an undesired isolated track exists in the event. If a track passing all of the above criteria is observed, the event is rejected. In regions where leptons are selected, such as the $Z \rightarrow \mu\mu$ control region, the number of isolated tracks must be larger than the number of expected leptons in order to veto the event.

7.4.8 Pre-selection

The pre-selection of the mono-jet analysis is constructed such that the full input datasets are reduced to a minimal form, while still retaining sufficient flexibility to be used for any of the signal or control regions. As such, only selection cuts common to all regions considered are included in the pre-selection. The first step is the trigger, in which the event is required to pass either the E_T^{miss} trigger or one of the two single electron triggers. Next, the general event cleaning is applied, ensuring that no corrupted events or bad detector conditions remain.

The analysis objects are then defined. The event must contain at least one central ($|\eta| < 2.0$), highly energetic jet ($p_T > 120 \text{ GeV}$). That jet must not fail the NCB jet cleaning criteria. The leading jet and any other jets in the event are required to pass the jet cleaning requirements. The leading two jets, if applicable, must not be pointing at masked Tile calorimeter regions. Any further jets which are pointing at masked Tile calorimeter regions are required to pass the additional Tile calorimeter cleaning criteria.

The event must contain zero, one, or two veto quality electrons or muons to cover all of the possible control regions. The track veto is then applied, ensuring that the track causing the veto is not associated to one of the electrons or muons which have been identified, in order to suppress backgrounds entering due to missed leptons. Note that it is important that the tracks associated to the identified leptons are not included in the track veto, as these may become good leptons forming the basis of a control region. The track veto is only designed to remove contributions from leptons which were missed by the analysis object definitions.

The E_T^{miss} is then recalculated with the given object inputs, and the boson p_T formed by the vector sum of the E_T^{miss} and electron(s) if present. Either the E_T^{miss} or boson p_T is required to be above the lowest threshold of 150 GeV , thus preserving the inclusion of all events from each of the different possible regions.

This set of criteria constitutes the pre-selection of the mono-jet analysis. This is summarized in Table 7.6.

7.4.9 Signal and control region specific selections

After the pre-selection, the different signal and control regions diverge, although a unified approach is preserved where possible. First, the relevant trigger is applied. This is the E_T^{miss} trigger for the signal region, muon control regions, or the $W \rightarrow e\nu$ control region which is used to estimate $W \rightarrow e\nu$ and $W \rightarrow \tau\nu$ in the signal region. For the $W \rightarrow e\nu$ control region used to estimate $Z \rightarrow \nu\nu$ in the signal

Selection type	SR	Z $\rightarrow \mu\mu$ CR	W $\rightarrow \mu\nu$ CR	Z $\rightarrow ee$ CR	W $\rightarrow e\nu$ [Z] CR	W $\rightarrow e\nu$ [W] CR
Trigger		EF_xe80_tclcw		EF_e24vhi_medium1	EF_e60_medium1	EF_xe80_tclcw
Event cleaning		As recommended by the collaboration				
Leading jet		$p_T > 120 \text{ GeV}$, $ \eta < 2.0$, $f_{ch}/f_{max} > 0.1$				
Jet cleaning		No <i>looser</i> bad jets				
Tile cleaning		Jet 1 and 2 not pointing to masked modules, others pass cleaning cuts				
Leptons	\emptyset	$2\mu, 0e$	$1\mu, 0e$	$0\mu, 2e$	$0\mu, 1e$	$0\mu, 1e$
Track veto		Veto on any track not identified as a lepton				
Observable $> 150 \text{ GeV}$	E_T^{miss}	$p_T^Z \equiv E_T^{\text{miss}}$	$p_T^W \equiv E_T^{\text{miss}}$	$p_T^Z \equiv E_T^{\text{miss}} - \sum p_T^e$	$p_T^W \equiv E_T^{\text{miss}} - p_T^e$	E_T^{miss}

Table 7.6: Pre-selection cuts for the mono-jet analysis. The uses of the different control regions are explained in Section 7.6. Requiring that at least one of the categories (columns) is satisfied for each selection cut (row) ensures that the same input file is useable for all signal and control regions. Additional cuts specific to each region are applied after the pre-selection.

region, or for the $Z \rightarrow ee$ control region, the logical OR of the two single lepton triggers is applied. Next, the E_T^{miss} or boson p_T cut is raised to the desired signal or control region number (1 to 9), corresponding to cuts at 150, 200, 250, 300, 350, 400, 500, 600, or 700 GeV.

The QCD and, to a lesser extent, top backgrounds must be rejected next. For all regions which use a E_T^{miss} cut, this significantly reduces the QCD contamination, but does not entirely remove it. For selections with a boson p_T cut, QCD contamination is an even larger problem. For E_T^{miss} based regions, the QCD background enters primarily due to a high- p_T jet being mis-reconstructed, thus this fake source of E_T^{miss} is aligned with the mis-reconstructed jet. Boson p_T based regions are instead dependent directly on the p_T of the object, so mis-identification of a high- p_T jet as an electron causes significant fake boson p_T , once again leading the observable to be aligned with the mis-identified jet. As such, for all region types, alignment of the primary observable and a jet is strongly correlated to the QCD background. In order to reject such events, a $\Delta\phi(E_T^{\text{miss}}, \text{jets}) > 1.0$ or $\Delta\phi(p_T^{\text{W,Z}}, \text{jets}) > 1.0$ cut is applied as appropriate. This is seen to reduce the QCD contamination to negligible levels, as will be quantified in Section 7.6.5.

In order to further reject top contributions and to enforce a mono-jet-like topology, a cut on the ratio of the leading jet p_T and the observable of interest (E_T^{miss} , p_T^W , or p_T^Z) is applied. For the signal region, this means that the leading jet must carry at least half of the E_T^{miss} , thus ensuring that the topology remains mono-jet-like. This cut was ultimately found to only have a small impact, but was retained nonetheless.

For the signal region, the final cut is to apply is the lepton veto. The event is rejected if any veto definition electrons or muons are present. For W control regions, a single tight definition electron or muon is required, as appropriate. The Z control regions are similar, with the requirement for two medium definition electrons or two tight definition muons as appropriate. For both W and Z control regions, the presence of additional veto quality electrons or muons beyond those required for the selection leads to the event being rejected.

In order to further enhance the purity of the control regions, cuts designed to specifically select Z and W events are applied. For Z control regions, this involves a cut on the invariant mass of the di-electron or di-muon system, $m_{\ell\ell}$. The requirement of $66 \text{ GeV} < m_{\ell\ell} < 116 \text{ GeV}$ selects the Z mass peak, rejecting the largest part of other backgrounds, such as Drell-Yan di-lepton production. The inability to measure the full four-momentum of the neutrino from W decays smears out the W mass peak, so the equivalent procedure is not available. However, a cut on the transverse mass m_T^ℓ of the system at the

Selection type	SR	Z $\rightarrow \mu\mu$ CR	W $\rightarrow \mu\nu$ CR	Z $\rightarrow ee$ CR	W $\rightarrow e\nu$ [Z] CR	W $\rightarrow e\nu$ [W] CR
Trigger	EF_xe80_tclcw			EF_e24vhi_medium1	EF_e60_medium1	EF_xe80_tclcw
Observable	E_T^{miss}	$p_T^Z \equiv E_T^{\text{miss}}$	$p_T^W \equiv E_T^{\text{miss}}$	$p_T^Z \equiv E_T^{\text{miss}} - \sum p_T^e$	$p_T^W \equiv E_T^{\text{miss}} - p_T^e$	E_T^{miss}
Region threshold	Observable $> \{150, 200, 250, 300, 350, 400, 500, 600, 700\}$ GeV, region numbers 1 to 9					
QCD/top rejection	$\Delta\phi(\text{observable}, \text{jets}) > 1.0$ for all jets, (leading jet p_T) / (observable) > 0.5					
Lepton selection	0 μ ,0 e	2 μ ,0 e	1 μ ,0 e	0 μ ,2 e	0 μ ,1 e	0 μ ,1 e
Boson cuts [GeV]	\emptyset	$66 < m_{\mu\mu} < 116$	$40 < m_T^\mu < 100$	$66 < m_{ee} < 116$	$40 < m_T^e < 100$	\emptyset
Additional cuts	\emptyset	\emptyset	\emptyset	\emptyset	$E_T^{\text{miss}} > 25$ GeV	\emptyset

Table 7.7: Selection criteria for the many different signal and control regions used by the mono-jet analysis. The uses of the different control regions are explained in Section 7.6. The pre-selection cuts are defined in Table 7.6.

level of $40 \text{ GeV} < m_T^\ell < 100 \text{ GeV}$ is found to increase the W purity. The transverse mass is defined as in Equation 7.1, where the neutrino observables are determined by vectorially subtracting the observed lepton from the reconstructed W boson p_T .

$$\sqrt{2p_T^\ell p_T^\nu \cdot [1 - \cos(\phi_\ell - \phi_\nu)]} \quad (7.1)$$

This transverse mass cut is applied to the W $\rightarrow \mu\nu$ control region and the W $\rightarrow e\nu$ [Z] control region. It is not applied to the W $\rightarrow e\nu$ [W] control region due to limited statistics. For this case, it was confirmed that the background contamination was still acceptable without the transverse mass cut, as will be shown in Section 7.6.3.

The W $\rightarrow e\nu$ [Z] control region, as mentioned, is very susceptible to QCD contamination. The highly isolated tight electron definition and the $\Delta\phi(p_T^W, \text{jets})$ cuts help with the QCD rejection, but are not always sufficient. One additional cut is applied in this specific control region: $E_T^{\text{miss}} > 25$ GeV. This aids in the rejection QCD backgrounds from jets faking electrons, but retains most real W events thanks to the neutrino from W decays, thus providing a source of real E_T^{miss} .

The region specific cuts for all of the signal regions and control regions are summarized in Table 7.7.

7.5 Dataset information

7.5.1 Data samples

The full ATLAS 2012 dataset consists of proton-proton collision data at a centre of mass energy of 8 TeV, recorded between April and December 2012. The `JetTauEtMiss` stream (datasets where the jet, hadronic τ , and E_T^{miss} triggers are applied) is used for the signal region and muon-based control regions which all make use of the E_T^{miss} trigger, while the `EGamma` stream (datasets where the electron and photon triggers are applied) is used for electron-based control regions which make use of single electron triggers. All of the triggers are unrescaled and provide access to the full dataset.

In total, 22.8 fb^{-1} of data was recorded, but a subset of the recorded data is generally unuseable due to detector problems or other imperfect data taking conditions. When such conditions occur, the associated events are tagged in an ATLAS database, and added to a list which specifies that they should be excluded from typical analysis selections. This is known as a GRL, where the application of the GRL results in a total useable integrated luminosity of 20.3 fb^{-1} . The total dataset is broken down into data-taking periods, with individual integrated luminosities listed in Table 7.8.

period	Dates	$\int \mathcal{L} dt$ [pb $^{-1}$]
A	Apr-04 : Apr-20	892
B	May-01 : Jun-18	5474
C	Jul-01 : Jul-24	1614
D	Jul-24 : Aug-23	3532
E	Aug-23 : Sep-17	2808
G	Sep-26 : Oct-08	1380
H	Oct-13 : Oct-26	1617
I	Oct-26 : Nov-02	1126
J	Nov-02 : Nov-26	2890
L	Nov-30 : Dec-06	961
A-L	Apr-04 : Dec-06	22754

Table 7.8: Data taking periods and the corresponding amounts of integrated luminosity for the 2012 dataset, before the application of data quality requirements.

7.5.2 Simulated samples

The mono-jet analysis makes extensive use of MC samples, both as a ratio of different MC events in the data-driven estimation and as an independent source of estimation for small backgrounds. Both of the ATLAS reconstruction simulation frameworks, **GEANT4** and **AFII**, are used. The **AFII** fast simulation which makes use of a parametrization of the ATLAS detector response is used where possible in order to reduce the computing requirements. The full “correct” simulation of the response of particles passing through a complete **GEANT4** ATLAS detector model are used where high precision is required, or where **AFII** is known to be less reliable.

The dominant Z+jets and W+jets backgrounds are simulated using **Sherpa** [81] with the **CT10NLO** [47] PDF set. **Sherpa** is a LO generator which handles both the Matrix Element (ME) and parton showering steps, and which generates processes in exclusive parton multiplicity bins before merging the results to form a single dataset. For these samples, parton multiplicities of up to five are considered, and the up, down, and strange quarks are treated as massless particles. In order to have sufficient statistics for the full E_T^{miss} spectrum, these MC samples are split at generator level in exclusive bins of truth boson p_T of [70,140], [140,280], [280,500], and 500+ GeV. Less energetic bosons with a p_T below 70 GeV are obtained by applying a truth-level boson p_T cut at analysis level on an inclusive sample. A mix of **GEANT4** and **AFII** samples are used, with **AFII** being used for boson p_T up to either 140 or 280 GeV, while **GEANT4** samples are used for the highest p_T range. The simulated samples for $Z \rightarrow \tau\tau + \text{jets}$ and $W \rightarrow \tau\nu + \text{jets}$ are all based on **GEANT4**, as **AFII** does not accurately replicate the complexity of τ decays. Samples are further sub-divided by jet flavour, with one sub-sample containing only charm quarks, another containing only bottom quarks, and the last one containing a veto on charm and bottom quarks. The resulting sample cross-sections are normalized to Next to Next to Leading Order (NNLO) predictions.

The smaller top-related backgrounds are produced by a combination of **MC@NLO** [82] for $t\bar{t}$ and s -channel single- t , while **AcerMC+Pythia6** [83, 48] was used for t -channel single- t . **MC@NLO**, a NLO generator including both the ME and parton shower steps, was used in conjunction with the **CT10NLO** PDF set. **AcerMC** is a generator for the LO ME generation of complicated processes, which is then matched to **Pythia6** for the parton shower, and where the **CTEQ6L1** [45] PDF set was used for both stages. In both cases, **GEANT4** was used, and the top quark mass was set to 172.5 GeV. The resulting sample cross-sections are normalized to approximate NNLO+NNLL (Next to Next to Leading Logarithm) predictions.

Diboson samples are produced using `Sherpa` with the CT10NLO PDF set. As for the $Z + \text{jets}$ and $W + \text{jets}$ backgrounds, the up, down, and strange quarks are treated as massless particles. However, only parton multiplicities of three or lower are considered. The `GEANT4` simulation framework is used for all diboson processes. The resulting sample cross-sections are normalized to NLO predictions.

All of the above mentioned samples are produced with a `Pythia8` [49] overlay of minimum-bias events in order to simulate the impact of pileup. An event-by-event pileup weight is applied in order to match the simulated pileup distribution to what was observed in data.

7.6 Background determination

7.6.1 Data-driven $W + \text{jets}$ and $Z + \text{jets}$ estimation technique

As mentioned in Section 7.2, the $Z \rightarrow \nu\nu + \text{jets}$ and $W \rightarrow \ell\nu + \text{jets}$ contributions to the signal region are estimated in a semi data-driven fashion from relevant control regions. By applying similar selections which differ only in the treatment of leptons, it is easy to construct orthogonal selections. Taking the ratio of regions then provides an accurate estimate of the contribution of a given process in the signal region, without the need for accurate knowledge of the cross-section of the process, and with the added benefit of many uncertainties cancelling in the ratio. This works very well with minimal assumptions when using $Z \rightarrow \ell\ell + \text{jets}$ control region events to determine the contribution of $Z \rightarrow \nu\nu + \text{jets}$ to the signal region, or when using $W \rightarrow \ell\nu + \text{jets}$ control region events to determine the contribution of the same process in the signal region. The use of $W \rightarrow \ell\nu + \text{jets}$ control region events to estimate the $Z \rightarrow \nu\nu + \text{jets}$ contribution in the signal region is also very useful, but comes with an assumption on the multiplicative factor between $W + \text{jets}$ and $Z + \text{jets}$ processes, and thus an associated uncertainty.

Given that the lepton treatment differs, the lepton identification, acceptance, and efficiencies are important to quantify. These must be known in order to properly estimate the contribution in the lepton veto signal region from a lepton selection control region. This also extends to the $E_{\text{T}}^{\text{miss}}$, where muon regions are almost identical to the signal region given the muons are MIPs and thus are easy to treat as invisible. For electron regions, this is not as simple, as the electrons deposit all of their energy in the calorimeter, unlike neutrinos, thus providing for possible energy overlaps between the electron(s) and jet(s). The means by which the leptons are treated varies by control region, and will be discussed in the relevant parts of this Section.

The data-driven estimation depends on a comprehensive understanding of the control region, signal region, and relation between the two including any differences that arise from different selection criteria. The formula used must account for three main effects:

1. Purity of the control region: the estimation relies on understanding the relation between the control process (such as $Z \rightarrow \mu\mu + \text{jets}$) and the signal process (such as $Z \rightarrow \nu\nu + \text{jets}$). This requires that the control process dominates the event yield in the control region to the maximum extent possible. Any additional events from physics processes unrelated to $W + \text{jets}$ and $Z + \text{jets}$ events, such as diboson, top, or QCD, are directly subtracted from data (using either pure MC information or the QCD data-driven estimate). Contamination from $W + \text{jets}$ and $Z + \text{jets}$ processes beyond the one which the control region was designed for are specially treated later.
2. Understanding of the impact of selection differences: selecting or rejecting leptons results in very different event topologies. The lepton acceptance \mathcal{A}_ℓ and efficiency ε_ℓ differences, as well as scale

factors SF_ℓ which account for differences in lepton identification and reconstruction between data and MC, are thus important to consider. The combination of these contributions is referred to as the transfer factor TF_ℓ .

3. Input stability: if the trigger differs between the signal and control region, then the starting datasets may differ. However, all of the mono-jet signal and control regions work with unrescaled triggers, thus the input datasets are equivalent. Trigger scale factors (ratio of data to MC efficiencies), which are generally small, must be applied in any case to account for data and MC trigger differences. Even though this is generally a small effect, if different triggers are used for the signal and control region, then the difference in the trigger efficiency scale factors must be included in any estimation. This factor will be referred to as $\mathcal{R}_{\text{trig}} = SF_{\text{trig}}^{\text{SR}}/SF_{\text{trig}}^{\text{CR}}$, with $SF_{\text{trig}}^A = \varepsilon_{\text{trig}}^{A,\text{data}}/\varepsilon_{\text{trig}}^{A,\text{MC}}$.

The proper combination of these three requirements gives the relation defined in Equation 7.2 for the data-driven prediction of the event yield $N_A^{\text{SR,DD}}$ of the process A in the signal region as estimated from the process B in the control region.

$$N_A^{\text{SR,DD}} = \left(N_{\text{all}}^{\text{CR,data}} - N_{B,\text{not W/Z}}^{\text{CR,MC}} \right) \cdot TF_\ell \cdot \mathcal{R}_{\text{trig}} \quad (7.2)$$

The full treatment of the transfer factor TF_ℓ depends on the processes A and B . However, it can be generalized as per Equation 7.3, where ‘‘pre-selection’’ implies that only cuts which are equivalent for the signal and control regions have been applied. Here, $\mathcal{A}_\ell^{\text{veto}}$ ($\varepsilon_\ell^{\text{veto}}$) represents the acceptance (efficiency) of the lepton vetoes in the signal region, and $\mathcal{A}_\ell^{\text{req}}$ ($\varepsilon_\ell^{\text{req}}$) represents the acceptance (efficiency) of all of the cuts related to the lepton requirements in the control region. For the control region, this includes quantities derived including lepton information, such as Boson mass cuts. This demonstrates that the differences in the lepton acceptances and efficiencies are covered by taking the ratio of the two processes of interest in their respective region, leaving only the lepton scale factors to be determined independently.

$$TF_\ell = \frac{N_A^{\text{SR,MC,pre-selection}}}{N_B^{\text{CR,MC,pre-selection}}} \cdot \frac{\mathcal{A}_\ell^{\text{veto}} \varepsilon_\ell^{\text{veto}}}{\mathcal{A}_\ell^{\text{req}} \varepsilon_\ell^{\text{req}}} \cdot \frac{SF_\ell^{\text{SR}}}{SF_\ell^{\text{CR}}} = \frac{N_A^{\text{SR,MC,all cuts}}}{N_B^{\text{CR,MC,all cuts}}} \cdot \frac{SF_\ell^{\text{SR}}}{SF_\ell^{\text{CR}}} \quad (7.3)$$

Combining Equations 7.2 and 7.3 results in Equation 7.4, a generic formula for the data-driven estimation of process A from process B (where A may or may not be the same as B).

$$N_A^{\text{SR,DD}} = \left(N_{\text{all}}^{\text{CR,data}} - N_{B,\text{not W/Z}}^{\text{CR,MC}} \right) \cdot \frac{N_A^{\text{SR,MC}}}{N_B^{\text{CR,MC}}} \cdot \frac{SF_\ell^{\text{SR}}}{SF_\ell^{\text{CR}}} \cdot \mathcal{R}_{\text{trig}} \quad (7.4)$$

The E_T^{miss} trigger is used only on the efficiency plateau, where $\varepsilon_{E_T^{\text{miss}}} > 98\%$. For the large majority of parameter space, especially when using higher E_T^{miss} thresholds than 150 GeV, the trigger is at or above the 99% threshold. As such, the scale factor $SF_{E_T^{\text{miss}}}$ is assumed to be 1, and the difference from 1 is taken as a trigger systematic uncertainty. In other words, when using only the E_T^{miss} trigger, $\mathcal{R}_{\text{trig}} = 1$. Given that the E_T^{miss} trigger is calorimeter based, meaning that muons are considered to be a source of missing transverse momentum, this means that any signal region estimates taken from muon control regions have $\mathcal{R}_{\text{trig}} = 1$. This is also true for the $W \rightarrow e\nu + \text{jets}$ and $W \rightarrow \tau\nu + \text{jets}$ estimates from the $W \rightarrow e\nu [W]$ control region, which also makes use of the E_T^{miss} trigger.

Data-driven estimates from electron control regions are more complicated. A combination of two single-electron triggers is used for these regions, where trigger scale factors for the combination are available. As such, despite being the logical OR of two separate triggers, they will frequently be referred

to simply as the single $-e$ trigger. Note that this trigger combination is not necessarily used at the level in which it is fully efficient, and so the same approach of taking $\text{SF}_{\text{single}-e} = 1$ cannot be done. As previously defined, the scale factor is the ratio of the data and MC trigger efficiencies. For the $W \rightarrow e\nu [Z]$ control region, this is simply $\text{SF}_{\text{single}-e}^{W \rightarrow e\nu [Z]} = \frac{\epsilon_{\text{single}-e}^{W \rightarrow e\nu [Z], \text{data}}}{\epsilon_{\text{single}-e}^{W \rightarrow e\nu [Z], \text{MC}}}$. Only trigger efficiencies for MC are provided, but electron identification MC to data scale factors are available, thus $\epsilon_{\text{single}-e}^{W \rightarrow e\nu [Z], \text{data}} = \epsilon_{\text{single}-e}^{W \rightarrow e\nu [Z], \text{MC}} \cdot \text{SF}_e$. The final form of $\mathcal{R}_{\text{trig}}$ for the $W \rightarrow e\nu [Z]$ control region to the signal region is provided in Equation 7.5.

$$\mathcal{R}_{\text{trig}}^{W \rightarrow e\nu [Z]} = \frac{\text{SF}_{\text{E}_T^{\text{miss}}}^{\text{SR}}}{\text{SF}_{\text{single}-e}^{W \rightarrow e\nu [Z]}} = \frac{1}{\epsilon_{\text{single}-e}^{W \rightarrow e\nu [Z], \text{MC}} \cdot \text{SF}_e / \epsilon_{\text{single}-e}^{W \rightarrow e\nu [Z], \text{MC}}} = \frac{1}{\text{SF}_e} \quad (7.5)$$

This is further complicated for the $Z \rightarrow ee$ control region, where the combined single electron trigger is still used, thus the trigger can fire on either (or both) lepton(s) from the Z boson. In this case, it is easier to consider the probability of neither electron firing the trigger, and then subtracting that from 1, which is mathematically equivalent. Thus, we have $\epsilon_{\text{single}-e}^{Z \rightarrow ee, \text{MC}} = 1 - \left(1 - \epsilon_{\text{single}-e, e_1}^{Z \rightarrow ee, \text{MC}}\right) \cdot \left(1 - \epsilon_{\text{single}-e, e_2}^{Z \rightarrow ee, \text{MC}}\right)$. Following the same application of lepton identification scale factors for the data trigger efficiency results in the final $\mathcal{R}_{\text{trig}}$ form for the $Z \rightarrow ee$ control region to the signal region as detailed in Equation 7.6.

$$\begin{aligned} \mathcal{R}_{\text{trig}}^{Z \rightarrow ee} &= \frac{\text{SF}_{\text{E}_T^{\text{miss}}}^{\text{SR}}}{\text{SF}_{\text{single}-e}^{Z \rightarrow ee}} = \frac{1}{\epsilon_{\text{single}-e}^{Z \rightarrow ee, \text{data}} / \epsilon_{\text{single}-e}^{Z \rightarrow ee, \text{MC}}} = \frac{\epsilon_{\text{single}-e}^{Z \rightarrow ee, \text{MC}}}{\epsilon_{\text{single}-e}^{Z \rightarrow ee, \text{data}}} \\ \mathcal{R}_{\text{trig}}^{Z \rightarrow ee} &= \frac{1 - \left(1 - \epsilon_{\text{single}-e, e_1}^{Z \rightarrow ee, \text{MC}}\right) \cdot \left(1 - \epsilon_{\text{single}-e, e_2}^{Z \rightarrow ee, \text{MC}}\right)}{1 - \left(1 - \epsilon_{\text{single}-e, e_1}^{Z \rightarrow ee, \text{MC}} \cdot \text{SF}_{e_1}\right) \cdot \left(1 - \epsilon_{\text{single}-e, e_2}^{Z \rightarrow ee, \text{MC}} \cdot \text{SF}_{e_2}\right)} \end{aligned} \quad (7.6)$$

The lepton data vs MC scale factors SF_ℓ are simple when working with a process where all of the truth leptons have been observed. The requirements of one good quality lepton for the $W \rightarrow \ell\nu + \text{jets}$ regions and two good quality leptons for the $Z \rightarrow \ell\ell + \text{jets}$ regions ensures that this simple case is applicable. The control region scale factors then become $\text{SF}_\ell^{\text{CR}} = \prod_i \text{SF}_\ell^i$ where i is the number of required leptons. When estimating the $Z \rightarrow \nu\nu + \text{jets}$ process in the signal region, which has no truth leptons, this simplifies further to $\text{SF}_{Z \rightarrow \nu\nu + \text{jets}}^{\text{SR}} = 1$.

Scale factors when estimating $W \rightarrow \ell\nu + \text{jets}$ processes in the signal region are more complicated, as there are truth leptons, but they have been missed. A proper treatment requires the separation of the two possible causes, namely the lepton being outside of the detector acceptance (for which there is no scale factor), or the lepton being within acceptance but missed (for which the difference between data and MC must be included). For such $W \rightarrow \ell\nu + \text{jets}$ estimates, the scale factor then becomes $f_{\mathcal{A}_\ell} + (1 - f_{\mathcal{A}_\ell}) \cdot (1 - \text{SF}_{\ell, \text{veto}}^{\text{SR}})$. In this equation, $f_{\mathcal{A}_\ell}$ is the fraction of events where the lepton is outside of the detector acceptance and $\text{SF}_{\ell, \text{veto}}^{\text{SR}}$ is the data vs MC scale factor for leptons passing the signal region lepton definition (and which are thus vetoed).

A summary of this information covering the different data-driven estimates is provided in Table 7.9. All of these estimations depend on the event kinematics, and thus are done bin-by-bin in the relevant observable in order to ensure that the estimate is reliable for all values of the observable in question. The specific observable used is always the E_T^{miss} in the signal region, and is always either the E_T^{miss} or boson p_T in the control region, as will be discussed in the following Sections where relevant.

The data driven estimation formula provided in Equation 7.4 is provided as a means of explaining

Process		Scale Factor SF		$\mathcal{R}_{\text{trig}}$	Observable	
CR	SR	CR	SR		CR	SR
$Z \rightarrow \mu\mu + \text{jets}$	$Z \rightarrow \nu\nu + \text{jets}$	$\text{SF}_{\mu_1} \cdot \text{SF}_{\mu_2}$	1	1	$p_T^Z \equiv E_T^{\text{miss}}$	E_T^{miss}
$W \rightarrow \mu\nu + \text{jets}$	$Z \rightarrow \nu\nu + \text{jets}$	SF_{μ_1}	1	1	$p_T^W \equiv E_T^{\text{miss}}$	E_T^{miss}
$Z \rightarrow ee + \text{jets}$	$Z \rightarrow \nu\nu + \text{jets}$	$\text{SF}_{e_1} \cdot \text{SF}_{e_2}$	1	$\mathcal{R}_{\text{trig}}^{Z \rightarrow ee}$	p_T^Z	E_T^{miss}
$W \rightarrow e\nu + \text{jets}$	$Z \rightarrow \nu\nu + \text{jets}$	SF_{e_1}	1	$\mathcal{R}_{\text{trig}}^{W \rightarrow e\nu [Z]}$	p_T^W	E_T^{miss}
$W \rightarrow e\nu + \text{jets}$	$W \rightarrow \tau\nu + \text{jets}$	SF_{e_1}	$\sum_{\ell=e,\mu} f_{\mathcal{A}_\ell} + [1 - f_{\mathcal{A}_\ell}] \cdot [1 - \text{SF}_{\ell,\text{veto}}^{\text{SR}}]$	1	E_T^{miss}	E_T^{miss}
$W \rightarrow \mu\nu + \text{jets}$	$W \rightarrow \mu\nu + \text{jets}$	SF_{μ_1}	$f_{\mathcal{A}_\mu} + [1 - f_{\mathcal{A}_\mu}] \cdot [1 - \text{SF}_{\mu,\text{veto}}^{\text{SR}}]$	1	$p_T^W \equiv E_T^{\text{miss}}$	E_T^{miss}
$W \rightarrow e\nu + \text{jets}$	$W \rightarrow e\nu + \text{jets}$	SF_{e_1}	$f_{\mathcal{A}_e} + [1 - f_{\mathcal{A}_e}] \cdot [1 - \text{SF}_{e,\text{veto}}^{\text{SR}}]$	1	E_T^{miss}	E_T^{miss}

Table 7.9: A summary of the data-driven estimation procedures used, specifying all of the relevant inputs to Equation 7.4. The Control Region (CR) process to use, Signal Region (SR) process to estimate, lepton scale factors for both processes, trigger scale factor ratio, and the observable used for the bin-by-bin estimation are provided. $\mathcal{R}_{\text{trig}}^{W \rightarrow e\nu [Z]}$ and $\mathcal{R}_{\text{trig}}^{Z \rightarrow ee}$ are defined in Equations 7.5 and 7.6 respectively.

the different contributions to the data-driven estimation. In reality, both $\mathcal{R}_{\text{trig}}$ (when $\mathcal{R}_{\text{trig}} \neq 1$) and SF_ℓ depend on the lepton kinematics, and thus vary event-by-event. The full event-by-event equation which is actually used for the signal region data-driven estimates is provided in Equation 7.7, where W_i is the event weight associated with event i . Note that $\text{SF}_{i,\text{trig}}^{\text{SR}} = 1$ is substituted given that the signal region always uses the E_T^{miss} trigger. Also note that $\text{SF}_{i,\text{trig}}^{\text{CR}}$ is written in the formula to emphasize the event-by-event dependence, but the functional form is $\text{SF}_{i,\text{trig}}^{\text{CR}} = \mathcal{R}_{\text{trig}}^{-1}$ due to the E_T^{miss} trigger scale factor being 1.

$$N_A^{\text{SR,DD}} = \left(N_{\text{all}}^{\text{CR,data}} - N_{B,\text{not W/Z}}^{\text{CR,MC}} \right) \frac{\sum_i^{\text{SR}} W_{i,A}^{\text{SR,MC}} \cdot \text{SF}_{i,\ell}^{\text{SR}}}{\sum_i^{\text{CR}} W_{i,B}^{\text{CR,MC}} \cdot \text{SF}_{i,\ell}^{\text{CR}} \cdot \text{SF}_{i,\text{trig}}^{\text{CR}}} \quad (7.7)$$

7.6.2 $Z \rightarrow \ell\ell$ control regions

$Z \rightarrow \mu\mu$ control region

Of the different control regions considered, the $Z \rightarrow \mu\mu$ control region is closest to the $Z \rightarrow \nu\nu$ process. The two muons are MIPs, meaning they deposit almost no energy in the calorimeters. At the LHC energy scale, the muons are essentially massless. As such, the muons behave similarly to neutrinos, with the added benefit of being visible and measurable in the muon spectrometers. This allows for their identification and quantification, while the similarities to $Z \rightarrow \nu\nu$ allows for a very clean data-driven estimation.

The $Z \rightarrow \mu\mu$ control region is very pure in $Z \rightarrow \mu\mu + \text{jets}$ events. After the selections described in Table 7.7 are applied, the fraction of $Z \rightarrow \mu\mu + \text{jets}$ events in the $Z \rightarrow \mu\mu$ control region varies from 92% to 95%. In particular, the diboson contribution grows as the E_T^{miss} threshold is increased. The E_T^{miss} and relative contribution of each process as a function of the E_T^{miss} can be seen in Figure 7.3.

Kinematic distributions for the control region are provided in Figure 7.4. Generally good agreement between data and MC is observed for all relevant distributions within the statistical and systematic uncertainties. The leading muon is observed to carry the majority of the Z boson p_T , peaking at near the E_T^{miss} threshold. The sub-leading muon is much softer, peaking at the lowest possible p_T required to

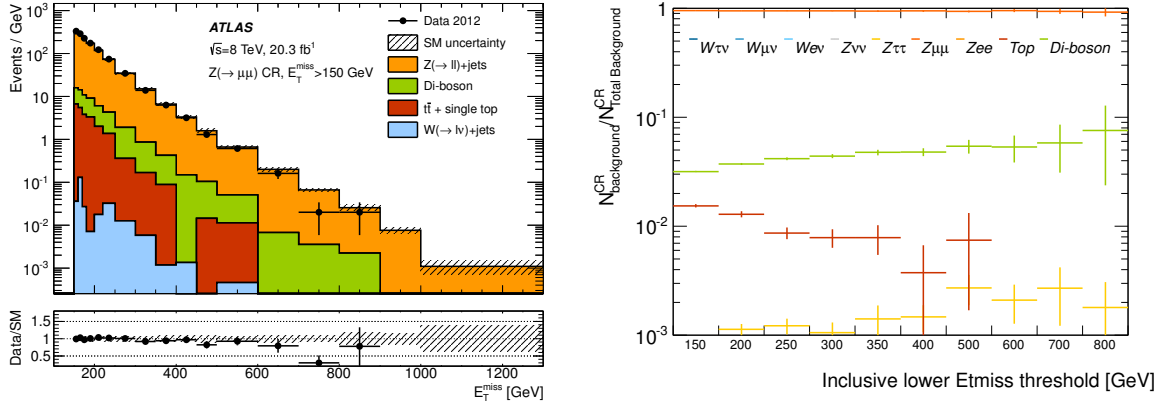


Figure 7.3: The primary observable (E_T^{miss}) and the relative contribution of each process in the $Z \rightarrow \mu\mu$ control region as a function of the primary observable. The region is seen to be very pure in $Z \rightarrow \mu\mu$ events, with a purity between 92% and 95%. Note that E_T^{miss} , E_{miss} , and \cancel{E}_T are different notation conventions for the same quantity.

pass the object definition.

$Z \rightarrow ee$ control region

The $Z \rightarrow ee$ control region is not so simple, as electrons are calorimeter objects, and thus are not as similar to neutrinos. The contribution of electrons to the calorimeter must be removed in order to mimic the $Z \rightarrow \nu\nu$ signature, thus well-identified electrons must be used, and electron isolation is required in order to not remove energy belonging to other calorimeter objects such as jets.

In this case, p_T^Z is not equivalent to the E_T^{miss} , due to the electrons being contained within the electromagnetic calorimeters. Instead, the Z boson p_T is equivalent to the E_T^{miss} (nominally zero, but with some smearing) plus the negative vectorial sum of the p_T of the two well-identified and isolated electrons. For the same reason, the E_T^{miss} trigger is not suitable for this control region. Instead, the logical OR of two single-electron triggers is used. The difference between the signal region E_T^{miss} trigger and control region single-electron triggers must be accounted for in the data-driven estimation.

The $Z \rightarrow ee$ control region is relatively pure in $Z \rightarrow ee$ events, at a similar level to the $Z \rightarrow \mu\mu$ control region thanks to the invariant Z mass cut. After the selection described in Table 7.7 is applied, the fraction of $Z \rightarrow ee + \text{jets}$ events in the $Z \rightarrow ee$ control region varies from 85% to 95%. In particular, the diboson and top backgrounds grow as the p_T^Z threshold is increased, leading to the purity degradation. The p_T^Z and relative contribution of each process as a function of the p_T^Z can be seen in Figure 7.5.

Kinematic distributions for the control region are provided in Figure 7.6. Generally good agreement between data and MC is observed for all relevant distributions within the statistical and systematic uncertainties. The leading electron is observed to carry the majority of the Z boson p_T , peaking at near the p_T^Z threshold. The sub-leading electron is much softer, peaking at the lowest possible p_T required to pass the object definition.

Estimating signal region processes

The transfer factor TF_ℓ as defined in Equation 7.3 is the primary output of the control region. For the $Z \rightarrow \nu\nu + \text{jets}$ estimates from the $Z \rightarrow \mu\mu$ and $Z \rightarrow ee$ control regions, the factor should be

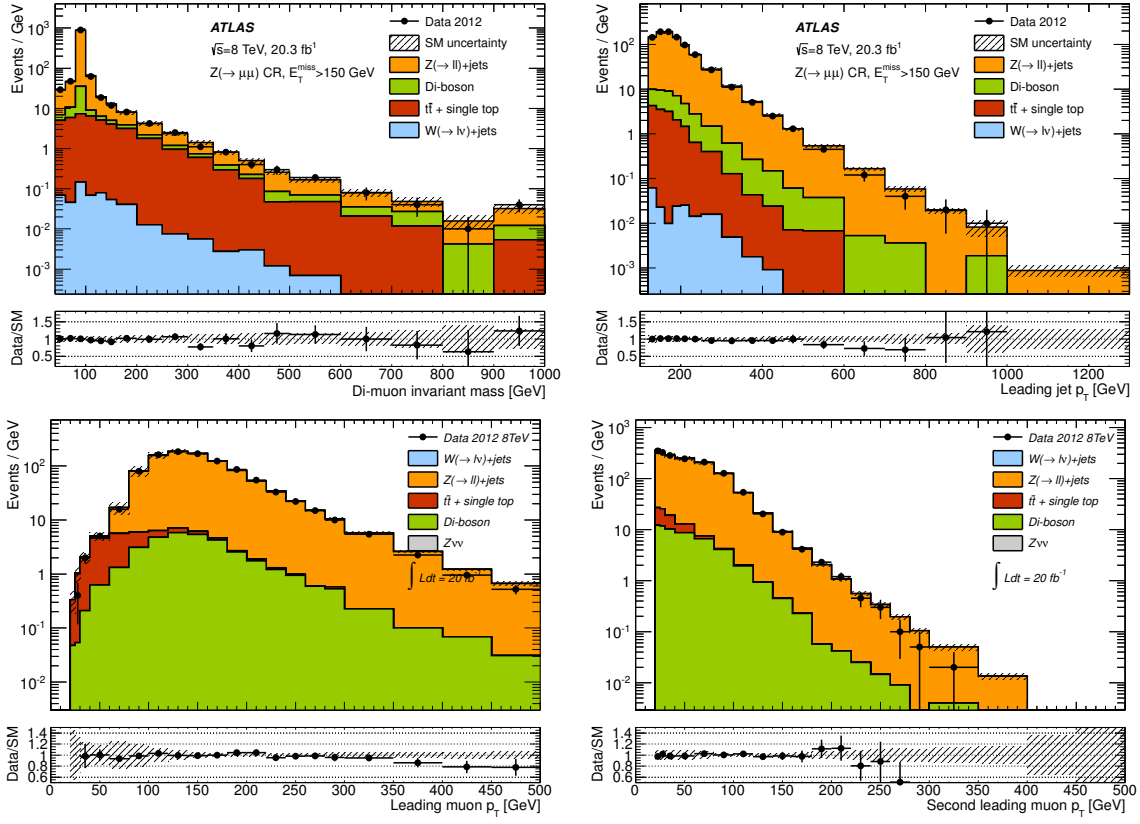


Figure 7.4: Kinematic distributions for the $Z \rightarrow \mu\mu$ control region after applying a 150 GeV E_T^{miss} threshold. The invariant mass before applying the cut $66 < m_{\mu\mu}/\text{GeV} < 116$ (top left), leading jet p_T (top right), leading lepton p_T (bottom left) and sub-leading lepton p_T (bottom right) are shown. Statistical uncertainties on the data are shown as error bars on each data point, while systematic uncertainties are shown as shaded regions.




Figure 7.5: The primary observable ($p_T^Z = E_T^{\text{miss}} - \sum p_T^e$) and the relative contribution of each process in the $Z \rightarrow ee$ control region as a function of the primary observable. The region is seen to be very pure in $Z \rightarrow ee$ events, with a purity between 85% and 95%. Note that E_T^{miss} , E_{Tmiss} , and E_T are different notation conventions for the same quantity.

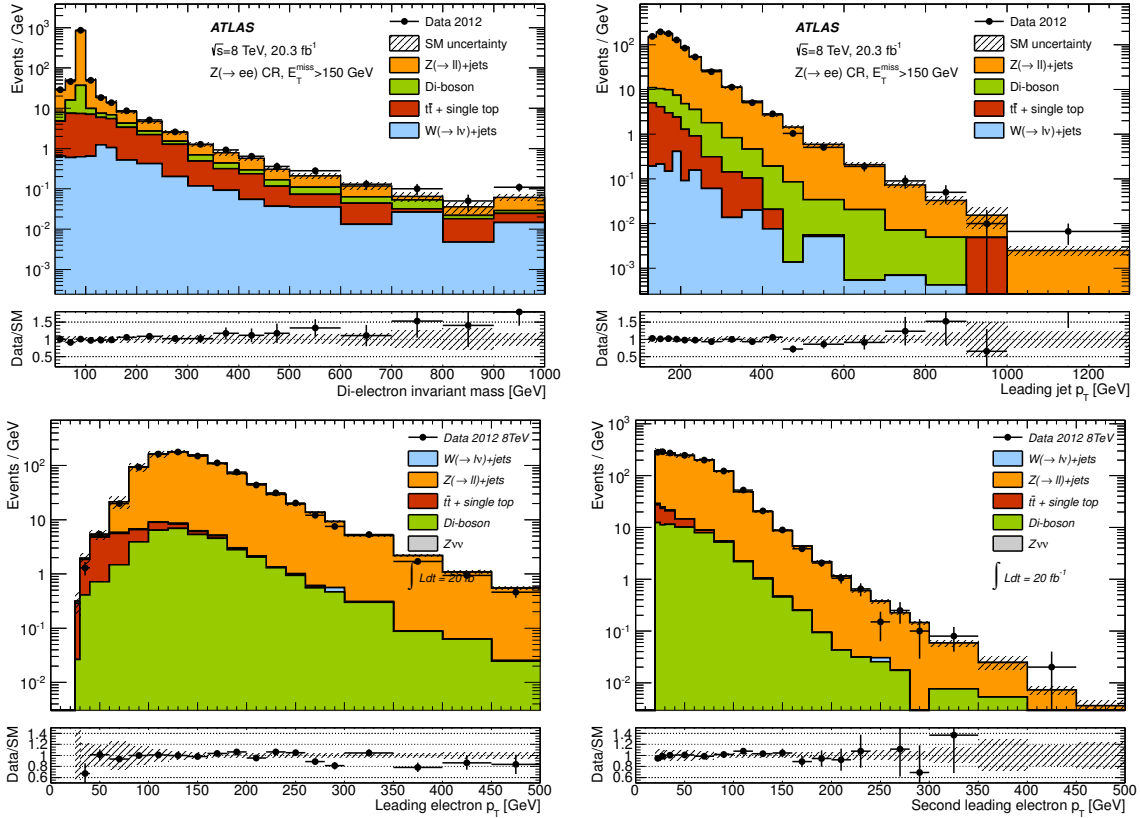


Figure 7.6: Kinematic distributions for the $Z \rightarrow ee$ control region after applying a $150 \text{ GeV } p_T^Z$ threshold. The invariant mass before applying the cut $66 < m_{\mu\mu}/\text{GeV} < 116$ (top left), leading jet p_T (top right), leading lepton p_T (bottom left) and sub-leading lepton p_T (bottom right) are shown. Statistical uncertainties on the data are shown as error bars on each data point, while systematic uncertainties are shown as shaded regions.




Figure 7.7: The transfer factor for the the estimation of the $Z \rightarrow \nu\nu + \text{jets}$ process from the $Z \rightarrow \mu\mu$ (left) and $Z \rightarrow ee$ (right) control regions. Note that $E_{\text{T}}^{\text{miss}}$ and $E_{\text{T}}^{\text{miss}}$ are different naming conventions for the same quantity.

approximately 12. This factor comes from three families of neutrinos vs one family of selected leptons (muons or electrons), a factor of approximately two from the difference in the vector and axial vector coefficients between the Z boson and neutrinos vs leptons, and a factor of approximately two for the detector muon acceptance. This is observed in Figure 7.7.

The $E_{\text{T}}^{\text{miss}}$ is primarily set by the leading lepton with some contribution by the subleading lepton, as seen in Figures 7.4 and 7.6. As such, when the $E_{\text{T}}^{\text{miss}}$ threshold is increased, the lepton p_{T} increases. This leads to a higher acceptance of events, due to the leptons becoming more central.

As the $E_{\text{T}}^{\text{miss}}$ continues to increase in the $Z \rightarrow \mu\mu$ control region, the dependence on the leading muon becomes a limitation. Reconstruction of very high p_{T} muons becomes problematic in the regime of $p_{\text{T}} \gtrsim 500 \text{ GeV}$ due to limitations of the ATLAS inner detector and muon spectrometer systems. In particular, the magnets are not strong enough to bend such energetic muons by a suitable amount between subsequent detector layers, thus reducing the precision of the measurement. As such, the acceptance again decreases, and thus the transfer factor increases.

This is not the case for high $E_{\text{T}}^{\text{miss}}$ values within the $Z \rightarrow ee$ control region. While muon spectrometer resolution degrades at high momentum, calorimeter resolution improves, so long as the object is contained within the calorimeter. As such, the acceptance continues to increase as the $E_{\text{T}}^{\text{miss}}$ is increased, reaching a final factor of approximately 8, suggesting that the acceptance factor is closer to 4/3 rather than 2 for a $E_{\text{T}}^{\text{miss}}$ threshold of 700 GeV.

7.6.3 $W \rightarrow \ell\nu$ control regions

$W \rightarrow \mu\nu$ control region

While the $Z \rightarrow \mu\mu$ control region is the closest to the $Z \rightarrow \nu\nu$ process, it becomes statistically limited for high boson p_{T} thresholds. The $W \rightarrow \mu\nu$ control region increases the available statistics at high kinematic thresholds, but at the cost of an extra required scale factor and associated uncertainty for scaling the W to the Z when estimating the $Z \rightarrow \nu\nu + \text{jets}$ process. The $W \rightarrow \mu\nu$ control region is also crucial for the estimate of the $W \rightarrow \mu\nu + \text{jets}$ contribution to the signal region.

The $W \rightarrow \mu\nu$ control region is relatively pure in $W \rightarrow \mu\nu + \text{jets}$ events, but not to the same extent

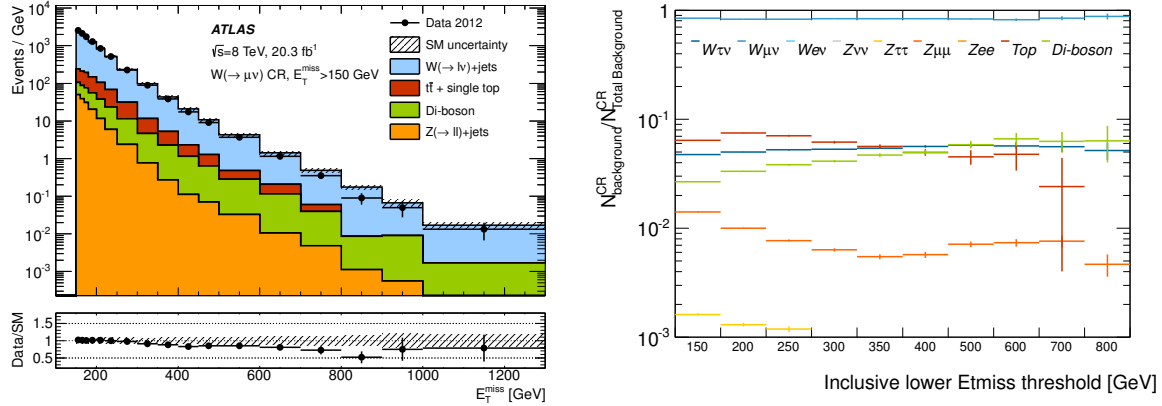


Figure 7.8: The primary observable (E_T^{miss}) and the relative contribution of each process in the $W \rightarrow \mu\nu$ control region as a function of the primary observable. The region is seen to be relatively pure in $W \rightarrow \mu\nu$ events, with a purity between 83% and 90%. Note that E_T^{miss} , E_{miss} , and \cancel{E}_T are different notation conventions for the same quantity.

as the $Z \rightarrow \mu\mu$ control region. This difference is primarily due to the lack of an invariant mass $m_{\mu\mu}$ cut, as the equivalent transverse mass m_T^μ distribution is more smeared out due to the inability to measure the z-component of the E_T^{miss} , thus providing weaker background discrimination. After the selections described in Table 7.7, the fraction of $W \rightarrow \mu\nu + \text{jets}$ events in the $W \rightarrow \mu\nu$ control region varies from 83% to 90%. The diboson contribution is seen to grow as the E_T^{miss} is increased, while the top background decreases. The E_T^{miss} and relative contribution of each process as a function of the E_T^{miss} can be seen in Figure 7.8.

Kinematic distributions for the control region are provided in Figure 7.9. Generally good agreement between data and MC is observed for all relevant distributions within the statistical and systematic uncertainties. The neutrino is observed to carry the majority of the W boson p_T , as seen by the muon p_T distribution being peaked at the lowest possible p_T required to pass the object definition.

$W \rightarrow e\nu$ control region

In addition to increasing the available statistics at high kinematic thresholds, the $W \rightarrow e\nu$ control region is crucial for estimating the $W \rightarrow e\nu + \text{jets}$ and $W \rightarrow \tau\nu + \text{jets}$ contributions to the signal region. For this reason, the $W \rightarrow e\nu$ topology is unique in that there are two sub-regions of interest. When estimating the $Z \rightarrow \nu\nu + \text{jets}$ process, the relevant observable is p_T^W , as it is necessary to mimic the boson p_T . However, when estimating $W \rightarrow e\nu + \text{jets}$ and $W \rightarrow \tau\nu + \text{jets}$, the relevant observable is the E_T^{miss} . This is different than the $W \rightarrow \mu\nu + \text{jets}$ control region due to the electron not being included as a part of the E_T^{miss} , thus requiring separate treatment for the two cases. These two sub-regions thus require fully independent studies, and rely on different triggers. The use of p_T^W as the primary observable necessitates the logical OR of the two single electron triggers, while using E_T^{miss} requires the E_T^{miss} trigger. Both of these two sub-regions are shown in parallel. The region used for estimating $Z \rightarrow \nu\nu + \text{jets}$ is denoted $W \rightarrow e\nu [Z]$, while the region for $W \rightarrow e\nu + \text{jets}$ and $W \rightarrow \tau\nu + \text{jets}$ estimation is called $W \rightarrow e\nu [W]$, as specified in Section 7.3.

The $W \rightarrow e\nu [Z]$ control region is relatively pure in $W \rightarrow e\nu + \text{jets}$ events after the selection described in Table 7.7 is applied, with a fraction of $W \rightarrow e\nu + \text{jets}$ events remaining approximately constant at

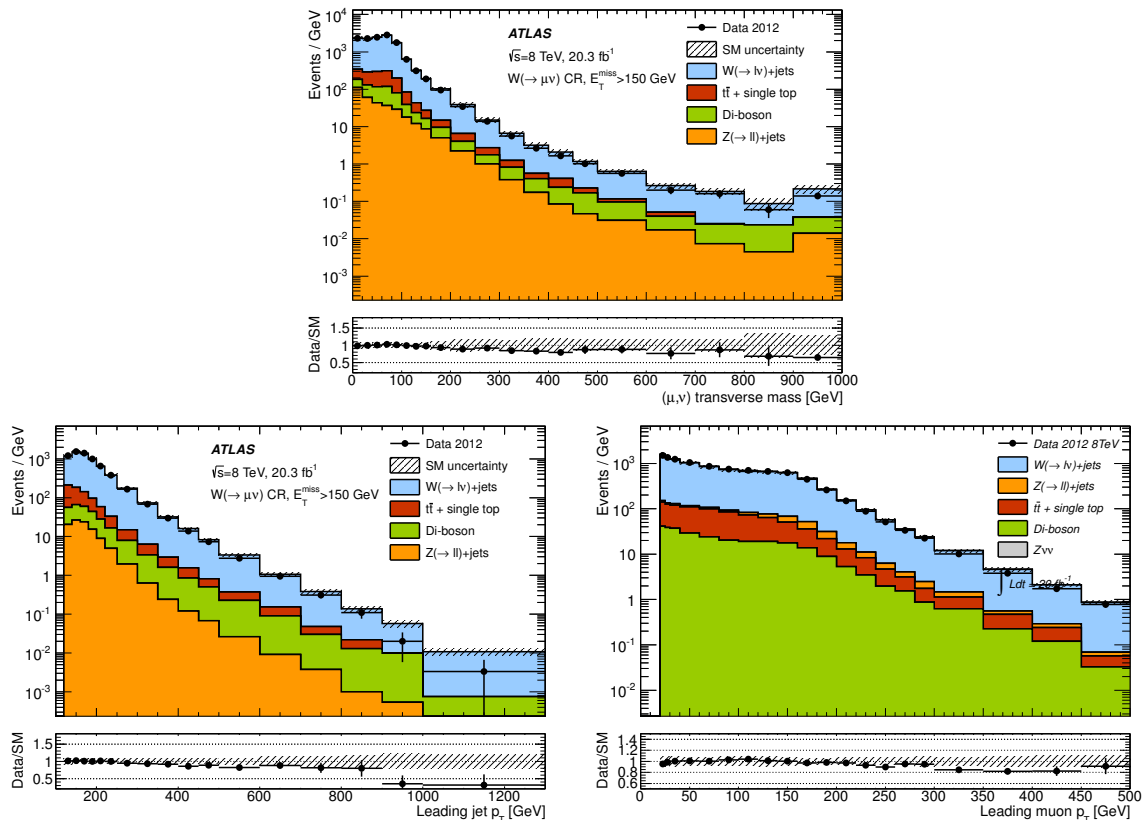


Figure 7.9: Kinematic distributions for the $W \rightarrow \mu\nu$ control region after applying a $150 \text{ GeV } E_T^{\text{miss}}$ threshold. The transverse mass before applying the cut $40 < m_T^{\mu}/\text{GeV} < 100$ (top), leading jet p_T (bottom left), and lepton p_T (bottom right) are shown. Statistical uncertainties on the data are shown as error bars on each data point, while systematic uncertainties are shown as shaded regions.

around 80% to 84%. The $W \rightarrow e\nu[W]$ control region after its respective selection has a much lower purity at the level of 65% to 70%. The observable of interest and the relative contribution of each background can be seen in Figure 7.10.

The difference between the two is primarily due to two effects, the observable used and the lack of a transverse mass cut in the $W \rightarrow e\nu[W]$ control region. The use of $p_T^Z = E_T^{\text{miss}} - p_T^e$ (vectorial) adds a dependence on the electron p_T . While taus can decay to electrons and thus naturally pass the control region selection, W bosons from the tau decay will preferentially boost the neutrino rather than the lepton, thus a cut on p_T^W rejects more tau events than a cut on E_T^{miss} . This can be seen by comparison to the $W \rightarrow \mu\nu$ control region, where the muon was observed to peak at minimal p_T and the neutrino (E_T^{miss}) at high p_T .

Not using a transverse mass cut also results in lower background suppression, although the application of such a cut would only have a limited use for rejecting the extra $W \rightarrow \tau\nu + \text{jets}$ events. Given that the $W \rightarrow e\nu[W]$ control region suffers from statistical limitations, it was decided that this cut was not necessary. This is particularly true given that the control region is used to estimate both $W \rightarrow e\nu + \text{jets}$ and $W \rightarrow \tau\nu + \text{jets}$ in the signal region, and thus the additional $W \rightarrow \tau\nu + \text{jets}$ events in the control region are of a similar process to that which is estimated. The contamination level of other backgrounds remains similar to the other W boson control regions, where the transverse mass cut is applied.

Kinematic distributions for both control regions are provided in Figure 7.11. Generally good agreement between data and MC is observed for all relevant distributions within the statistical and systematic uncertainties. The neutrino is observed to carry the majority of the W boson p_T , as seen by the electron p_T distribution being peaked at the lowest possible p_T required to pass the object definition.

Estimating signal region processes

The transfer factor TF_ℓ as defined in Equation 7.3 is the primary output of the control region. The resulting factors for both $Z \rightarrow \nu\nu + \text{jets}$ and $W \rightarrow \ell\nu + \text{jets}$ estimates are shown in Figure 7.12.

All of the transfer factors are seen to slowly decrease as the E_T^{miss} threshold is increased. This is for the same reason as the $Z \rightarrow \ell\ell$ control regions, where the higher E_T^{miss} threshold corresponds to a higher lepton p_T , and thus increasing the E_T^{miss} leads to more central leptons. This in turn reduces the acceptance loss of the lepton selection. Note that this does not fall off for muons even at very high E_T^{miss} thresholds unlike what is observed for $Z \rightarrow \mu\mu$ because the primary source of the E_T^{miss} is the neutrino, not the muon, as already mentioned with respect to Figure 7.9. As such, the muon does not reach the same energy scale as in $Z \rightarrow \mu\mu$ events with a similar E_T^{miss} value, and thus the limitations of the inner detector and muon spectrometer are less pronounced.

The $Z \rightarrow \nu\nu$ transfer factor is seen to be larger for the $W \rightarrow e\nu[Z]$ control region than for the $W \rightarrow \mu\nu$ control region due to the very tight electron definition which is needed to reduce QCD background. Comparing the $W \rightarrow e\nu + \text{jets}$ and $W \rightarrow \tau\nu + \text{jets}$ estimates from the $W \rightarrow e\nu[W]$ control region demonstrates the significantly higher contribution of the $W \rightarrow \tau\nu + \text{jets}$ process in the signal region due to the hadronic decay mode, even after the track veto has been applied.

7.6.4 Coherence of $Z \rightarrow \nu\nu + \text{jets}$ data-driven estimates

Given that the $Z \rightarrow \nu\nu + \text{jets}$ background in the signal region has four separate estimates from four orthogonal control regions, it is important to ensure that the four methods provide coherent estimates.

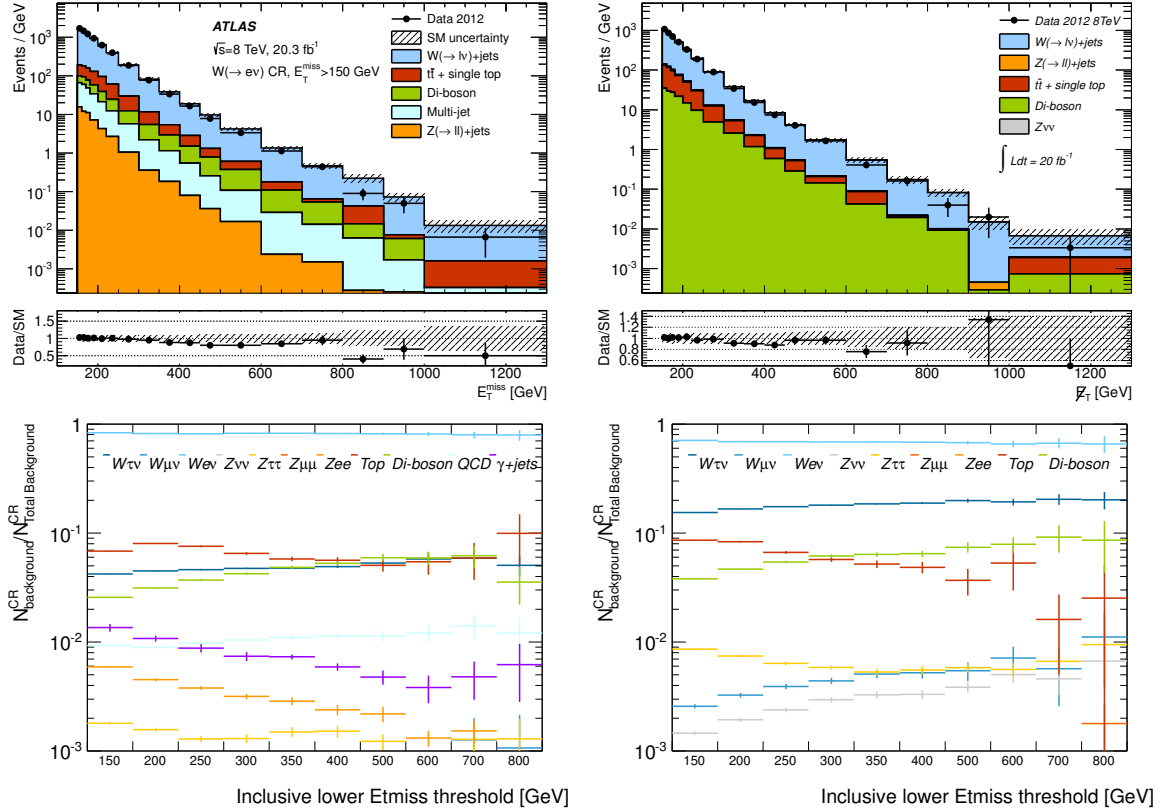


Figure 7.10: The primary observable (top) and the relative contribution of each process in the $W \rightarrow e\nu [Z]$ control region (left) and $W \rightarrow e\nu [W]$ control region (right) as a function of the primary observable (bottom). The primary observable is p_{T}^W for the $W \rightarrow e\nu [Z]$ control region and $E_{\text{T}}^{\text{miss}}$ for the $W \rightarrow e\nu [W]$ control region. The $W \rightarrow e\nu [Z]$ region is seen to be relatively pure in $W \rightarrow \mu\nu$ events, with a purity between 80% and 84%. The $W \rightarrow e\nu [W]$ region is less pure, at a level of 65% to 70%. Note that $E_{\text{T}}^{\text{miss}}$, E_{Tmiss} , and \cancel{E}_{T} are different notation conventions for the same quantity.

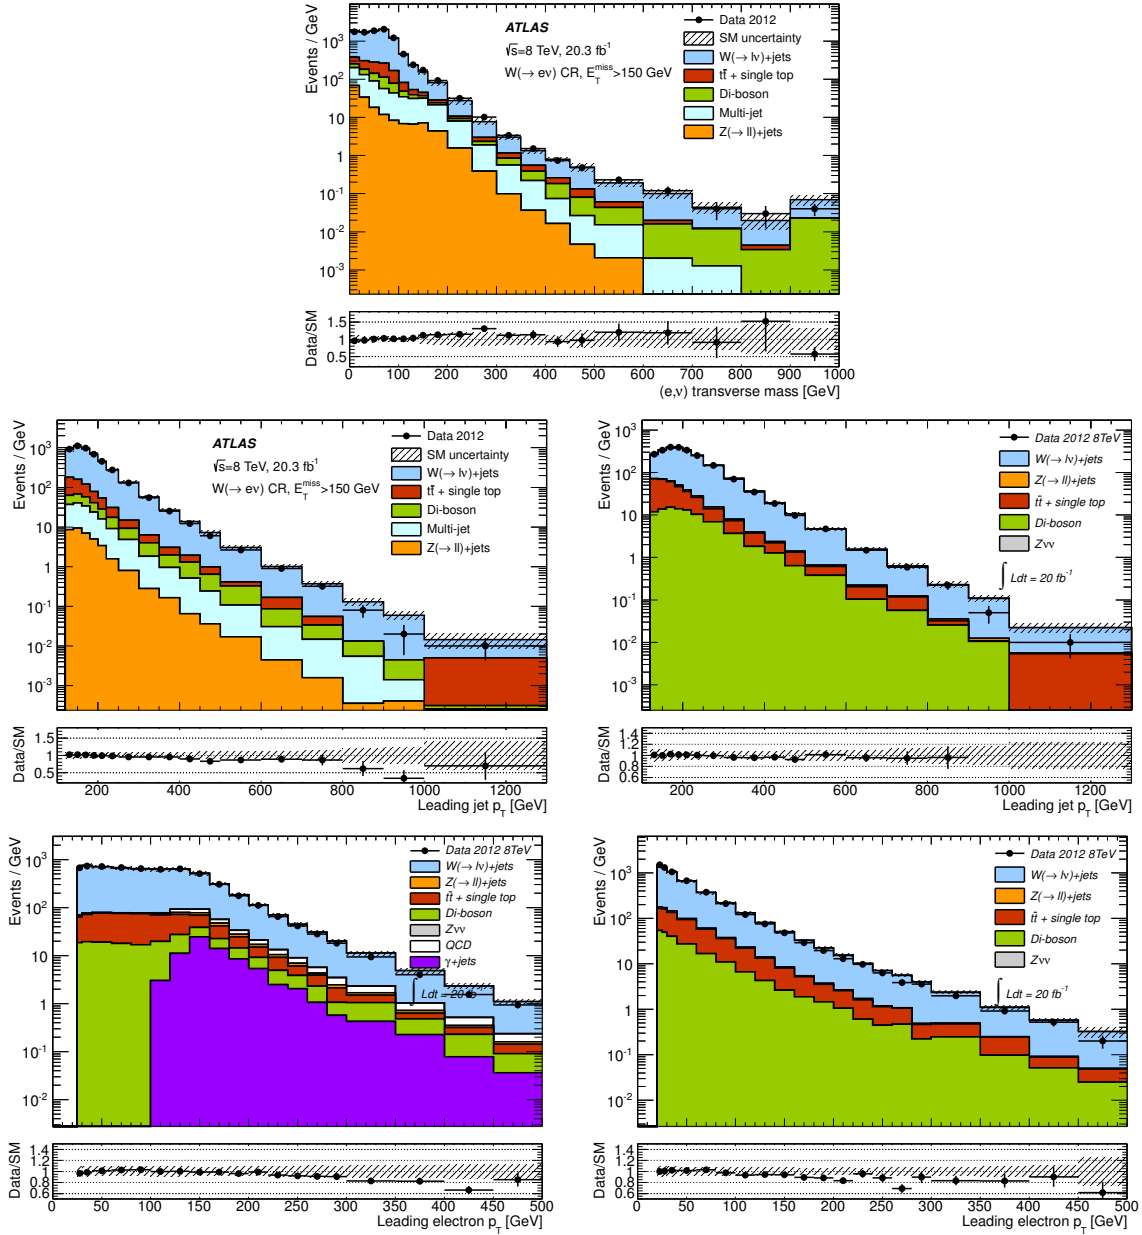


Figure 7.11: Kinematic distributions for the $W \rightarrow e\nu$ control regions after applying a $150 \text{ GeV } p_T^W$ or E_T^{miss} threshold. Both the $W \rightarrow e\nu [Z]$ (left) and $W \rightarrow e\nu [W]$ (right) control regions are shown. The transverse mass before applying the cut $40 < m_T^\mu / \text{GeV} < 100$ (top, only applied for $W \rightarrow e\nu [Z]$), leading jet p_T (middle), and lepton p_T (bottom) are shown. Statistical uncertainties on the data are shown as error bars on each data point, while systematic uncertainties are shown as shaded regions.

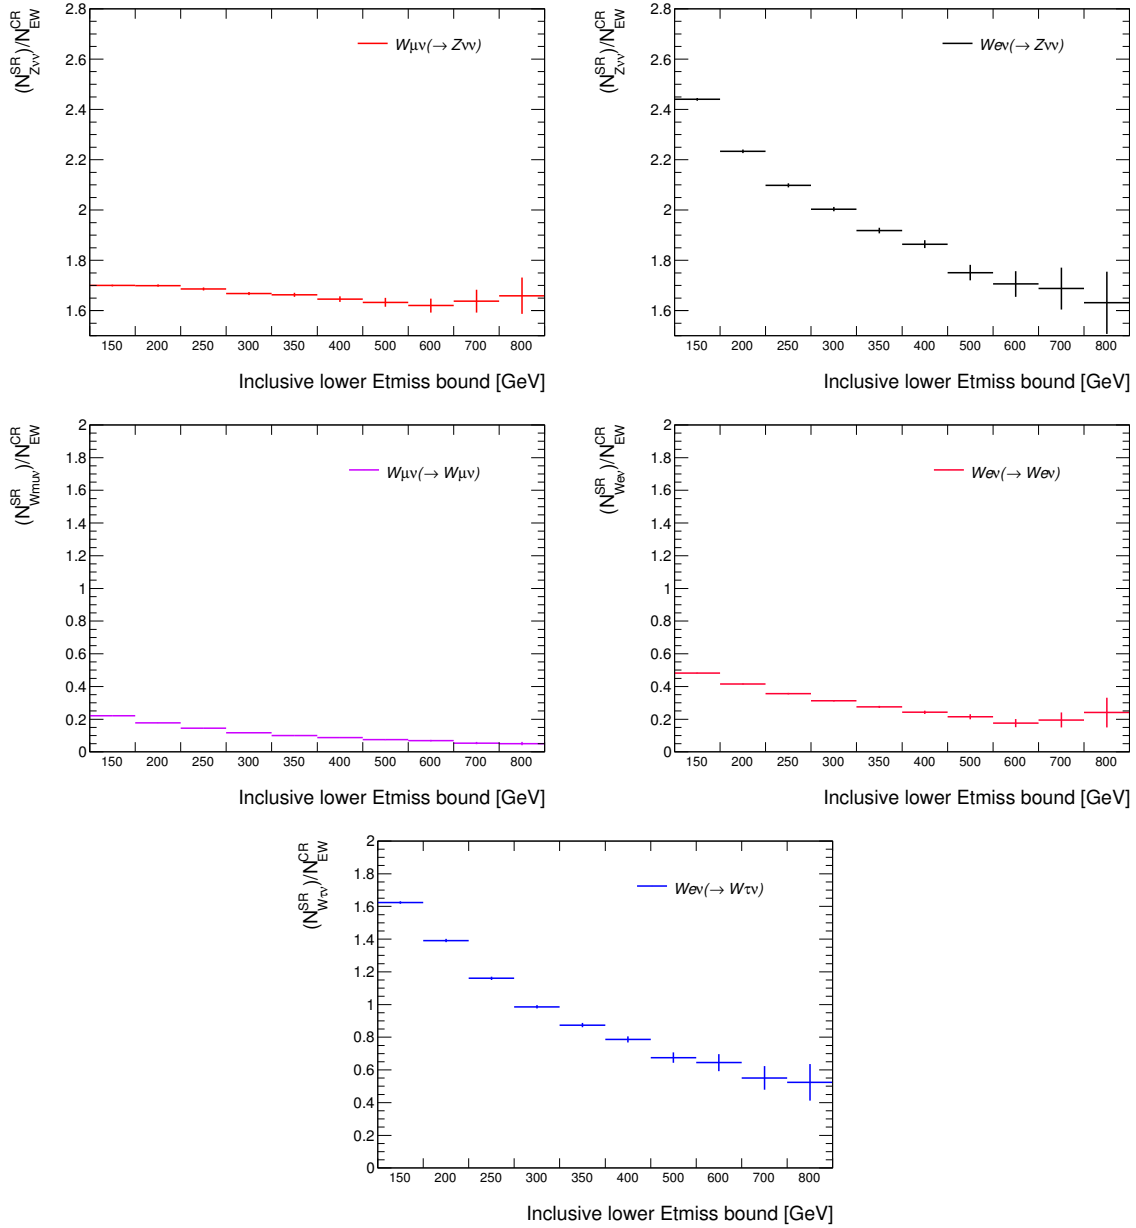


Figure 7.12: The transfer factor for the the estimation of the $Z \rightarrow \nu\nu + \text{jets}$ (top) and $W \rightarrow \ell\nu + \text{jets}$ (middle and bottom) processes from the $W \rightarrow \mu\nu$ (top left and middle left), $W \rightarrow e\nu [Z]$ (top right), and $W \rightarrow e\nu [W]$ (middle right and bottom) control regions. Note that E_{miss} and $E_{\text{T}}^{\text{miss}}$ are different naming conventions for the same quantity.

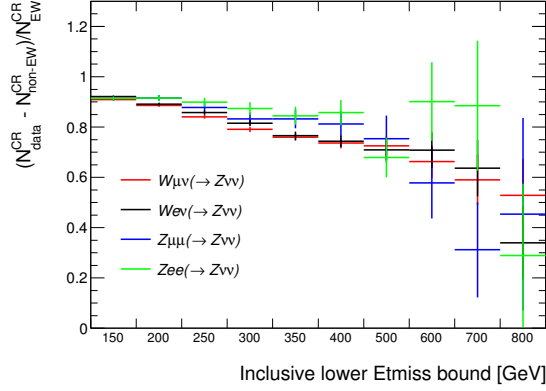


Figure 7.13: A comparison of the normalization factor (Equation 7.8) for estimates of the $Z \rightarrow \nu\nu + \text{jets}$ in the signal region taken from $Z \rightarrow \mu\mu$ (blue), $W \rightarrow \mu\nu$ (red), $Z \rightarrow ee$ (green), and $W \rightarrow e\nu$ [Z] (black) control regions. Error bars are statistical uncertainties. Generally good agreement is observed, although a difference between $W \rightarrow \ell\nu$ and $Z \rightarrow \ell\ell$ processes is possible for higher $E_{T^{\text{miss}}}$ thresholds. As such, an electroweak radiative corrections systematic uncertainty is taken from Reference [84] when estimating $Z \rightarrow \nu\nu + \text{jets}$ from $W \rightarrow \ell\nu$ control regions. Note that $E_{T^{\text{miss}}}$ and $E_{T^{\text{miss}}}$ are different naming conventions for the same quantity.

This is one means of checking for any errors, as well as ensuring that the potential differences between Z and W bosons are negligible. Such cross-checks are particularly important for the high $E_{T^{\text{miss}}}$ regime, where electroweak corrections may cause differences.

In order to make such a comparison in an equivalent way, Equation 7.4 is rearranged to obtain the ratio of the data-driven $Z \rightarrow \nu\nu + \text{jets}$ estimate and the pure MC expected yield. This corresponds to a normalization factor quantifying the degree to which the MC mis-estimates the actual $Z \rightarrow \nu\nu + \text{jets}$ yield in data. If the four different data-driven estimates are to be coherent, they should result in similar normalization factors. The normalization factor for a generic control region B is in Equation 7.8, where the fact that $\text{SF}_{\ell}^{\text{SR}} = 1$ for the $Z \rightarrow \nu\nu + \text{jets}$ process is used.

$$\frac{N_{Z \rightarrow \nu\nu + \text{jets}}^{\text{SR,DD}}}{N_{Z \rightarrow \nu\nu + \text{jets}}^{\text{SR,MC}}} = \frac{\left(N_{\text{all}}^{\text{CR,data}} - N_{B, \text{not } W/Z}^{\text{CR,MC}} \right)}{N_B^{\text{CR,MC}}} \cdot \frac{\mathcal{R}_{\text{trig}}}{\text{SF}_{\ell}^{\text{CR}}} \quad (7.8)$$

A comparison of the four independent $Z \rightarrow \nu\nu + \text{jets}$ estimates is provided in Figure 7.13. Good agreement is observed at low $E_{T^{\text{miss}}}$. As the $E_{T^{\text{miss}}}$ is increased, a slight difference is observed between the $Z \rightarrow \ell\ell$ and $W \rightarrow \ell\nu$ regions. This difference is most significant for a $E_{T^{\text{miss}}}$ threshold of 350 GeV. For higher thresholds, the statistical uncertainty becomes too large to observe any significant differences. This difference, which is expected to be from electroweak radiative corrections, only appears at the level of a few percent and is within the systematic uncertainties. Regardless, an extra electroweak systematic uncertainty from the detailed studies conducted in Reference [84] is included for $W \rightarrow \ell\nu$ control regions when estimating the $Z \rightarrow \nu\nu + \text{jets}$ process, varying from 1% for a 150 GeV threshold up to 6% at 1 TeV.

7.6.5 Multi-jet control region

The multi-jet (QCD) background can enter the mono-jet signal region either by significantly mis-reconstructed jets leading to unbalanced events or highly energetic neutrinos from heavy flavour jet

decays. This background is expected to be small in the signal region due to the high E_T^{miss} cut and the $\Delta\phi(E_T^{\text{miss}}, \text{jets}) > 1.0$ requirement. However, the enormous cross-section for QCD processes and limited MC statistics compared to data necessitates a data-driven study. A technique known as the *jet smearing* method has been borrowed from previous searches for this purpose, as is detailed in the internal ATLAS documentation in References [85, 86]. This method consists of the following five steps:

1. A set of good quality seed events with a low E_T^{miss} is selected. These events are selected using a combination of single-jet triggers (all of which except the last are prescaled) in order to cover the full kinematic range. These single-jet triggers have the following online Event Filter (EF)-level p_T thresholds: 55, 80, 110, 145, 180, 220, 280, 360, and 450 GeV. The cleaning criteria mentioned in Table 7.6 are applied, and a E_T^{miss} significance cut ($E_T^{\text{miss}}/\sqrt{\sum E_T} < 1.0$) is required in order to select well-balanced events. A leading jet p_T cut of 100 GeV is required in order to bring the seed region closer to the region of interest to the analysis. These seed events provide a sample of well-understood data events which can then be smeared by well-controlled functions to provide pseudo-data in extreme regions of parameter space, which a sufficient amount of randomized smearing iterations should provide a reasonable representation of actual data.
2. The jet response $\mathcal{R}_{\text{jet}} = p_T^{\text{jet, reco}}/p_T^{\text{jet, truth}}$ is measured in a QCD control region to estimate the fluctuations that can be expected during the reconstruction of a jet of a given p_T . This is split into two distributions, one for b -jets and one for light quark/gluon jets. This is necessary because the mechanisms by which b -jets and light quark/gluon jets can result in E_T^{miss} is significantly different. In particular, b -jet leptonic decays to muons are a natural source of E_T^{miss} . The measurement of these response distributions is taken from previous studies, as detailed in the previous References.
3. In order to account for differences in data and MC, the derived response function is adjusted by smearing the seed events until the smeared pseudo-data matches the observed data in the QCD control regions. This is done both by widening the response distribution with additional Gaussian smearing as compared with a di-jet QCD control region and by modifying the low-side tail using a tri-jet QCD control region. This is also taken from previous studies, as detailed in the previous References.
4. A mono-jet QCD control region is defined, identical to the signal region, except with an inverted cut on the distance between the jets and E_T^{miss} of $\Delta\phi(E_T^{\text{miss}}, \text{jets}) < 1.0$. This requirement of aligned jets and E_T^{miss} provides a sample of primarily QCD events. The data-adjusted \mathcal{R}_{jet} distribution is then used to smear the seed events numerous times, limited only by the available computing power. 500 smeared events per seed event are produced. The QCD control region is then used to normalize the smeared pseudo-data such that it matches the observed data, where this normalization is performed once per E_T^{miss} threshold.
5. Once the normalization has been derived, the pseudo-data events are used to estimate the contribution of the QCD background to the various mono-jet signal and control regions.

The use of 500 smeared events per seed event provides sufficient statistics for the lower E_T^{miss} thresholds, but breaks down for E_T^{miss} cuts of 350 GeV and higher. As such, the normalization derived in the 300 GeV region is used for all higher thresholds. The 200 GeV threshold QCD control region kinematic distributions post-normalization are shown in Figure 7.14. Differences in the kinematic shapes between

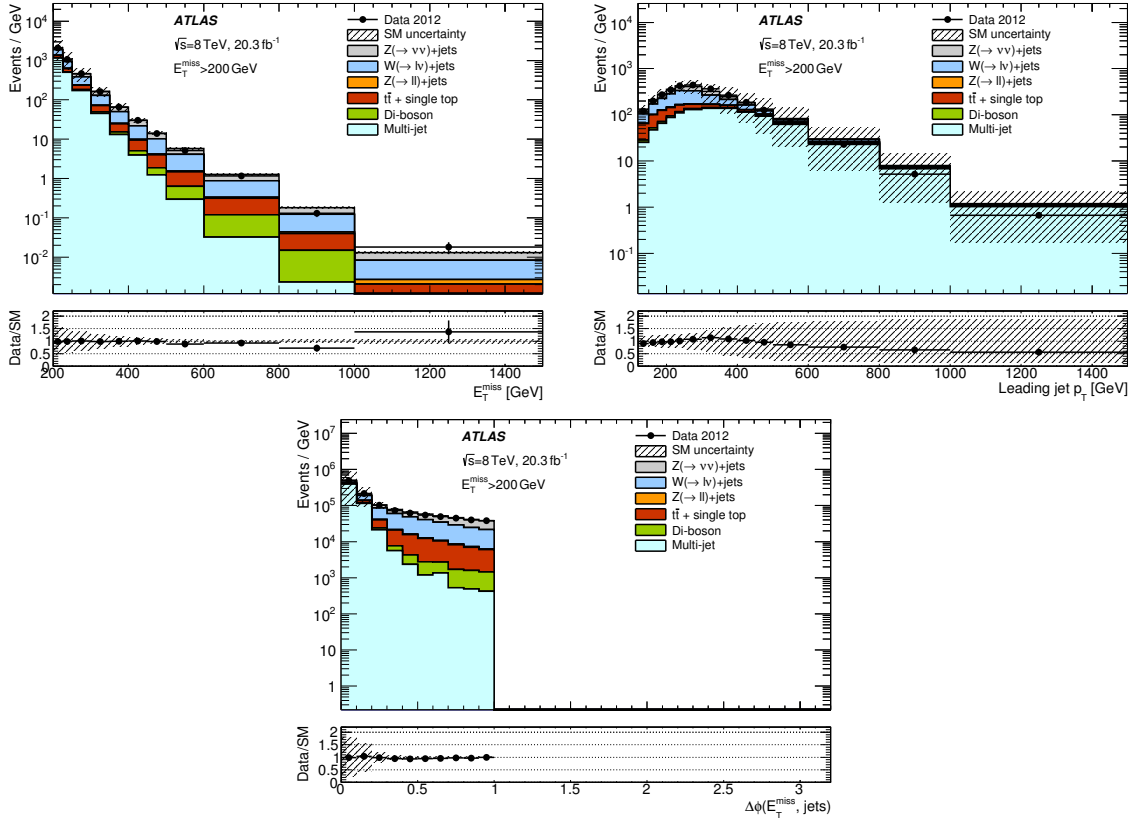


Figure 7.14: Kinematic distributions in the mono-jet QCD control region for a 200 GeV E_T^{miss} threshold. Distributions are shown for the E_T^{miss} (top left), leading jet p_T (top right), and $\Delta\phi(E_T^{\text{miss}}, \text{jets})$ (bottom).

the pseudo-data and the observed data are seen to be reasonable but not perfect. A conservative 100% systematic uncertainty is taken to cover such differences. The QCD contribution is seen to fall off quickly with increasing E_T^{miss} compared to the other backgrounds, while it has a harder leading jet p_T spectrum. The $\Delta\phi(E_T^{\text{miss}}, \text{jets})$ cut at 1.0 is seen to be sub-optimal, and future iterations could benefit from a higher purity QCD control region by moving the cut to a value of 0.5 or smaller.

The derived normalization factors and QCD contamination estimates for the corresponding signal regions are provided in Table 7.10.

7.6.6 Non-collision background estimation

Cleaning

The mono-jet selection, without specialized event cleaning cuts, naturally accepts NCB events. This can be seen by considering four jet variables sensitive to NCB processes:

1. Charged particle fraction, $f_{\text{ch}} = \sum p_T^{\text{track}} / p_T^{\text{reco}}$
 - Only tracks ghost-associated to the jet with $p_T > 500$ MeV are considered.
 - p_T^{reco} is the fully calibrated reconstructed jet p_T .
 - For real hard-scatter jets, tracks can be matched to the correct primary vertex with a high efficiency. The neutral component of real hard-scatter jets still varies, providing distribution

E_T^{miss} threshold	Normalization factor	Expected signal region contamination
150 GeV	0.0035	6453.3 ± 75
200 GeV	0.0033	785.4 ± 26.2
250 GeV	0.0032	176.6 ± 12.4
300 GeV	0.0029	44.3 ± 6.2
350 GeV	0.0029	14.5 ± 3.6
400 GeV	0.0029	5.5 ± 2.2
500 GeV	0.0029	1.1 ± 1.0
600 GeV	0.0029	0.3 ± 0.5
700 GeV	0.0029	0

Table 7.10: Normalization factors and expected QCD contamination in the mono-jet signal region, shown for each E_T^{miss} threshold considered. Uncertainties shown only account for statistics. A systematic uncertainty of 100% is assumed.

spread out between 0 and 1, with a peak near the centre. Values larger than 1 are possible when the jet is significantly mis-reconstructed or passes through a bad region of the calorimeter, resulting in missed energy.

- Jets from peripheral collisions (such as pileup) which have wrongly identified the vertex of interest typically have only a small number of tracks from the selected vertex, thus $f_{\text{ch}} \approx 0$.
- Jets from calorimeter noise bursts typically have no tracks unless a hard-scatter object happens to coincide, resulting in $f_{\text{ch}} \approx 0$.
- Jets from beam-induced backgrounds typically do not come from the interaction point, and thus have no tracks pointing to them, resulting in $f_{\text{ch}} \approx 0$.
- Jets from cosmic muons either have no tracks or a track depending on how close they pass to the interaction point. If the track is not observed, then the cosmic muon appears as $f_{\text{ch}} \approx 0$.
- In general, $f_{\text{ch}} \approx 0$ is a strong indicator of NCB events.

2. Electromagnetic fraction, $f_{\text{em}} = E_{\text{calo}}^{\text{EM}}/E_{\text{calo}}^{\text{EM+Had}}$

- $E_{\text{calo}}^{\text{EM}}$ is the amount of energy deposited in all parts of the EM calorimetry.
- $E_{\text{calo}}^{\text{EM+Had}}$ is the amount of energy deposited in all parts of the EM and hadronic calorimetry.
- Real hadronic jets are expected to deposit their energy in both types of calorimeters. Given that hadronic showers are a stochastic process, it is possible that all of the energy can be contained in the EM calorimeters. Alternatively, very little can be deposited in the EM calorimeters while the vast majority is deposited in the hadronic calorimeters. Both of these options represent uncommon types of jets, but are possible. As such, typical hard-scatter jets have a distribution spread between 0 and 1, with the peak location varying with the p_T of the jet.
- EM coherent noise, HEC spikes, and other sources of calorimeter noise fluctuations generally appear only in a single type of calorimeter. Jets from such noise sources thus typically have either $f_{\text{em}} \approx 0$ or $f_{\text{em}} \approx 1$.
- Beam halo events typically consist of a beam passing through the calorimeter laterally instead of longitudinally, thus depositing energy in one layer of the calorimeter. This leads to either $f_{\text{em}} \approx 0$ or $f_{\text{em}} \approx 1$.

- In general, $f_{em} \approx 0$ and $f_{em} \approx 1$ are both indicators of NCB events.
3. Max energy fraction per calorimeter layer, $f_{max} = \max \left\{ E_{calo}^{layer} / \sum E_{calo}^{layer} \right\}$
- E_{calo}^{layer} is the amount of energy deposited in a single layer of a calorimeter. There are multiple layers per calorimeter.
 - This is essentially a finer version of f_{em} , done per-layer rather than per-calorimeter type, and thus the same arguments apply.
 - In general, $f_{max} \approx 0$ and $f_{max} \approx 1$ are both indicators of NCB events.
4. Timing, t_{jet}
- All reconstructed objects in ATLAS are corrected for the time of flight with respect to the interaction point, where t_{jet} is this corrected time for a given jet.
 - Due to this correction, typical hard-scatter reconstructed jets should have a time $t_{jet} \approx 0$.
 - Beam backgrounds occur either early ($t_{jet} < 0$, jets from before the interaction) or late ($t_{jet} > 0$, jets from after the interaction).
 - Calorimeter noise, cosmic muons, and other such sources can occur at any time, and typically $t_{jet} \neq 0$.
 - In general, $t_{jet} \neq 0$ is a strong indicator of NCB events.

Fake jets from NCB sources are most important when they are the leading jet, providing a large source of fake E_T^{miss} , and thus firing the trigger. As such, the kinematic distributions and respective cuts are only investigated and applied on the leading jet. Distributions of these four variables, before any NCB-specific jet cleaning is applied but after the typical jet cleaning criteria mentioned in Section 7.4.5, are shown in Figure 7.15.

The f_{ch} and f_{max} are two independent variables with high NCB rejection potential. Taking the ratio of the two, f_{ch}/f_{max} , provides a powerful discriminant, as seen in Figure 7.16. This figure shows the ratio in both data and MC, both for the signal region and for the $W \rightarrow \mu\nu$ control region, which is nearly independent of NCB sources. Note that MC does not include NCB sources, thus differences observed between data and MC in regions where NCB is expected is a strong indication of contamination in data. From these distributions, a leading jet cut of $f_{ch}/f_{max} > 0.1$ was determined to be ideal for removing NCB backgrounds with a minimal impact on good jet acceptance. After this cut, good jet rejection is 0.4% as measured in the $W \rightarrow \mu\nu$ control region. This 0.4% is taken as a systematic uncertainty for backgrounds which are taken purely from MC. Backgrounds which are estimated in a data-driven fashion do not require this systematic because the same mis-identification exists in both the numerator and denominator of the transfer factor TF_ℓ .

Additional cuts beyond this ratio, such as $f_{max} < 0.85$, were investigated. These additional cuts were found to further reduce the NCB contamination, but made it harder to estimate the amount of NCB remaining due to being so heavily suppressed. In order to have a means of properly estimating the small amount of remaining NCB, such additional cuts were not applied.

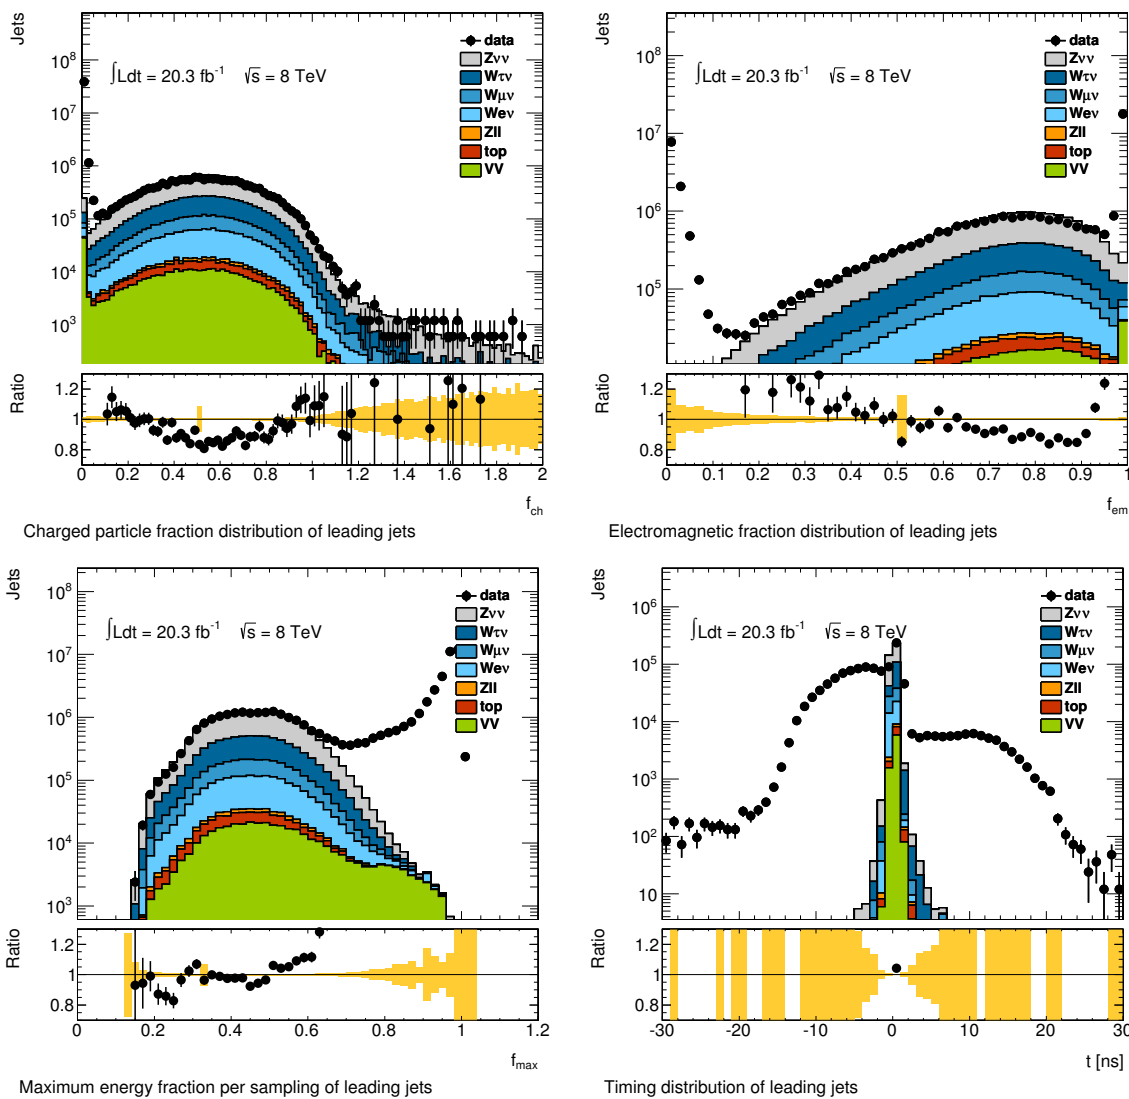


Figure 7.15: The charge fraction f_{ch} (top left), electromagnetic fraction f_{em} (top right), maximum energy fraction per layer f_{max} (bottom left), and timing t_{jet} (bottom right) for the leading jet in the mono-jet signal region with $E_T^{miss} > 150 \text{ GeV}$. Distributions are shown after typical *looser* jet cleaning is applied, but before advanced NCB rejection cuts specifically derived for the mono-jet analysis.




Figure 7.16: 2D plots of the charge fraction f_{ch} and maximum energy fraction per layer f_{max} for the signal region (top) and $W \rightarrow \mu\nu$ control region (bottom) in data (left) and MC (right). NCB is not included in MC, thus differences between data and MC in NCB-sensitive distributions are primarily due to the contribution of such events. The extra events at $f_{\text{max}} = 0$ in MC is primarily due to real hard-scatter jets with a mis-identified primary vertex.

	MC		data				NCB	
	tagged	tagged	tagged, $t < -5$ ns	$t < -5$ ns				
$E_T^{\text{miss}} > 150$ GeV	37	80	26	146	(0.05%)	449 ± 81	(0.12%)	
$E_T^{\text{miss}} > 200$ GeV	16	13	5	18	(0.01%)	47 ± 20	(0.04%)	
$E_T^{\text{miss}} > 250$ GeV	6	2	0	3	(0.01%)	–	–	
$E_T^{\text{miss}} > 300$ GeV	2	2	0	1	(0.01%)	–	–	
$E_T^{\text{miss}} > 350$ GeV	1	2	0	0	(0.00%)	–	–	
$E_T^{\text{miss}} > 400$ GeV	1	0	0	0	(0.00%)	–	–	
$E_T^{\text{miss}} > 500$ GeV	0	0	0	0	(0.00%)	–	–	

Table 7.11: Data-driven Non-Collision Background (NCB) estimates for the mono-jet signal region after the special leading jet cleaning cut of $f_{\text{ch}}/f_{\text{max}} > 0.1$. Numbers in brackets list the percentage of the total signal yield the respective number represents. Tagged MC events represent mis-tagged good jets, given that the MC samples do not include NCB processes.

Remaining contamination estimation

A beam-induced background tagger, detailed in Reference [87], is used to estimate the amount of residual NCB contamination in the signal region after the leading jet cleaning cuts. This tagger’s efficiency is at the level of 10%, with a negligible mis-tag rate. A combination of this tagger and the leading jet timing distribution t_{jet} is used to estimate the amount of NCB contamination.

It is conservatively assumed that all jets which are reconstructed as occurring more than 5 ns before the event are due to NCB. Under this assumption, the NCB contamination is $N_{\text{NCB}}^{\text{DD}} = N_{\text{tag}}^{\text{tag}} \frac{N_{t_{\text{jet}} < -5 \text{ ns}}}{N_{t_{\text{jet}} < -5 \text{ ns}}^{\text{tag}}}$. 100% systematic uncertainties are taken for this procedure, given the absence of a more thorough means of quantification.

One weakness of this procedure is the reliance on the tagger, which is only $\mathcal{O}(10\%)$ efficient. The timing cut of $t_{\text{jet}} < -5$ ns already significantly reduces the available statistics, so a $\mathcal{O}(10\%)$ efficient tagger on top of that results in a small sample available for the data-driven estimation. This is the reason that other cuts beyond the cut on $f_{\text{ch}}/f_{\text{max}}$ were not added, as then the events available for the estimation would be further reduced.

Due to these limitations, $N_{t_{\text{jet}} < -5 \text{ ns}}^{\text{tag}}$ and thus $N_{\text{NCB}}^{\text{DD}}$ is only nonzero for the 150 and 200 GeV E_T^{miss} threshold signal regions, as seen in Table 7.11. However, due to the observed trend and completely negligible contribution where the estimate is available, this is deemed to not be a significant problem. The NCB cleaning cuts in use are sufficient to remove the NCB contamination in the mono-jet signal region. Furthermore, the number of mis-tagged MC events hints that some of the tagged data jets are not actually from NCB processes.

7.6.7 Top control region

The top contamination in the mono-jet signal region remains a small contribution, but the same cannot be said for all of the control regions. The $W \rightarrow e\nu$ [Z] and $W \rightarrow \mu\nu$ control regions both contain up to 7% top contamination for low p_T^W thresholds, and smaller amounts as p_T^W is increased. For this reason, the top processes become very important to the data-driven $Z \rightarrow \nu\nu + \text{jets}$ estimates. In order to ensure that the top normalization is reasonable in this high E_T^{miss} regime, and to estimate the uncertainties on the top normalization, two $t\bar{t}$ control regions are constructed. Single- t is not considered due to a contamination lower than $t\bar{t}$ in the control regions by an order of magnitude.

The two $t\bar{t}$ control regions are identical to the respective $W \rightarrow e\nu$ [Z] and $W \rightarrow \mu\nu$ control regions,

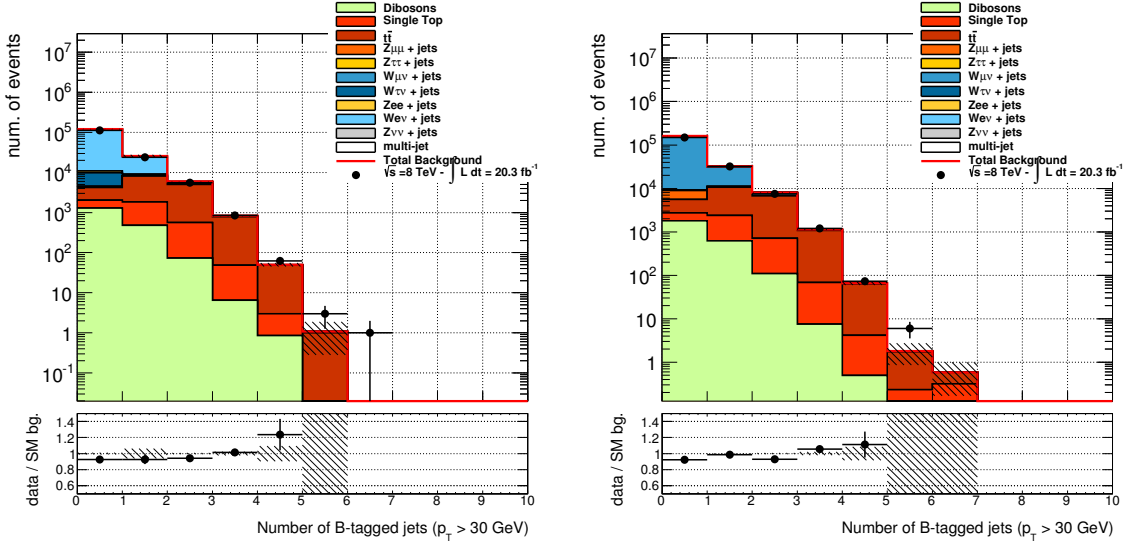


Figure 7.17: The b -jet multiplicity in the $W \rightarrow e\nu[Z]$ (left) and $W \rightarrow \mu\nu$ (right) top control regions, before the $N_{b\text{-jets}} \geq 2$ cut is applied. The $t\bar{t}$ background is observed to be dominant in both cases for this selection.

thus selecting semi-leptonic $t\bar{t}$ events, except that a minimum of two b -jets are required and the cuts related to top suppression are relaxed. In particular, $\Delta\phi(E_T^{\text{miss}}, \text{jets})$ is relaxed from 1.0 to 0.5, and $p_T^{\text{jet1}}/p_T^W < 0.5$ is removed entirely. Note that these regions are not strictly orthogonal to the $W \rightarrow e\nu[Z]$ and $W \rightarrow \mu\nu$ control regions, but they are relatively pure in $t\bar{t}$ and will not be used for any data-driven estimation in the respective $W \rightarrow \ell\nu$ control region, and thus this was deemed acceptable.

The b -tagging criteria applied made the choice of retaining as many real b -jets as possible, rather than focusing on rejecting fake b -jets, due to the limited statistics available for high p_T^W thresholds. The 90% tagging efficiency working point was thus used, and the b -jet was required to be within the tracking acceptance, with $|\eta| < 2.4$. Scale factors for the b -tagging procedure are applied to MC in order to compare with data.

This selection is successful in ensuring that $t\bar{t}$ dominates the event yield. This is shown in Figure 7.17, where the $t\bar{t}$ process is seen to be the largest contribution to the region for $N_{b\text{-jets}} \geq 2$. These regions are then used to scale the MC to the data, thus providing a $t\bar{t}$ normalization factor. The resulting distributions can be seen in Figure 7.18.

In order to assess a systematic uncertainty, the $W + \text{jets}$ normalization is varied by 50% to account for the uncertainty of heavy flavour associated production. The b -tagging weights are also varied and split into the b -jet, c -jet, and light quark/gluon jet components. The resulting normalization factors and associated uncertainties are provided in Table 7.12. The normalization factor is always consistent with one within uncertainties, thus it is not applied. The derived uncertainties are used to bound the required top normalization uncertainty as a function of the E_T^{miss} threshold, resulting in 20% for $150 \leq E_T^{\text{miss}}/\text{GeV} \leq 300$, 50% for $300 < E_T^{\text{miss}}/\text{GeV} < 500$, and 100% for $E_T^{\text{miss}} \geq 500$ GeV.

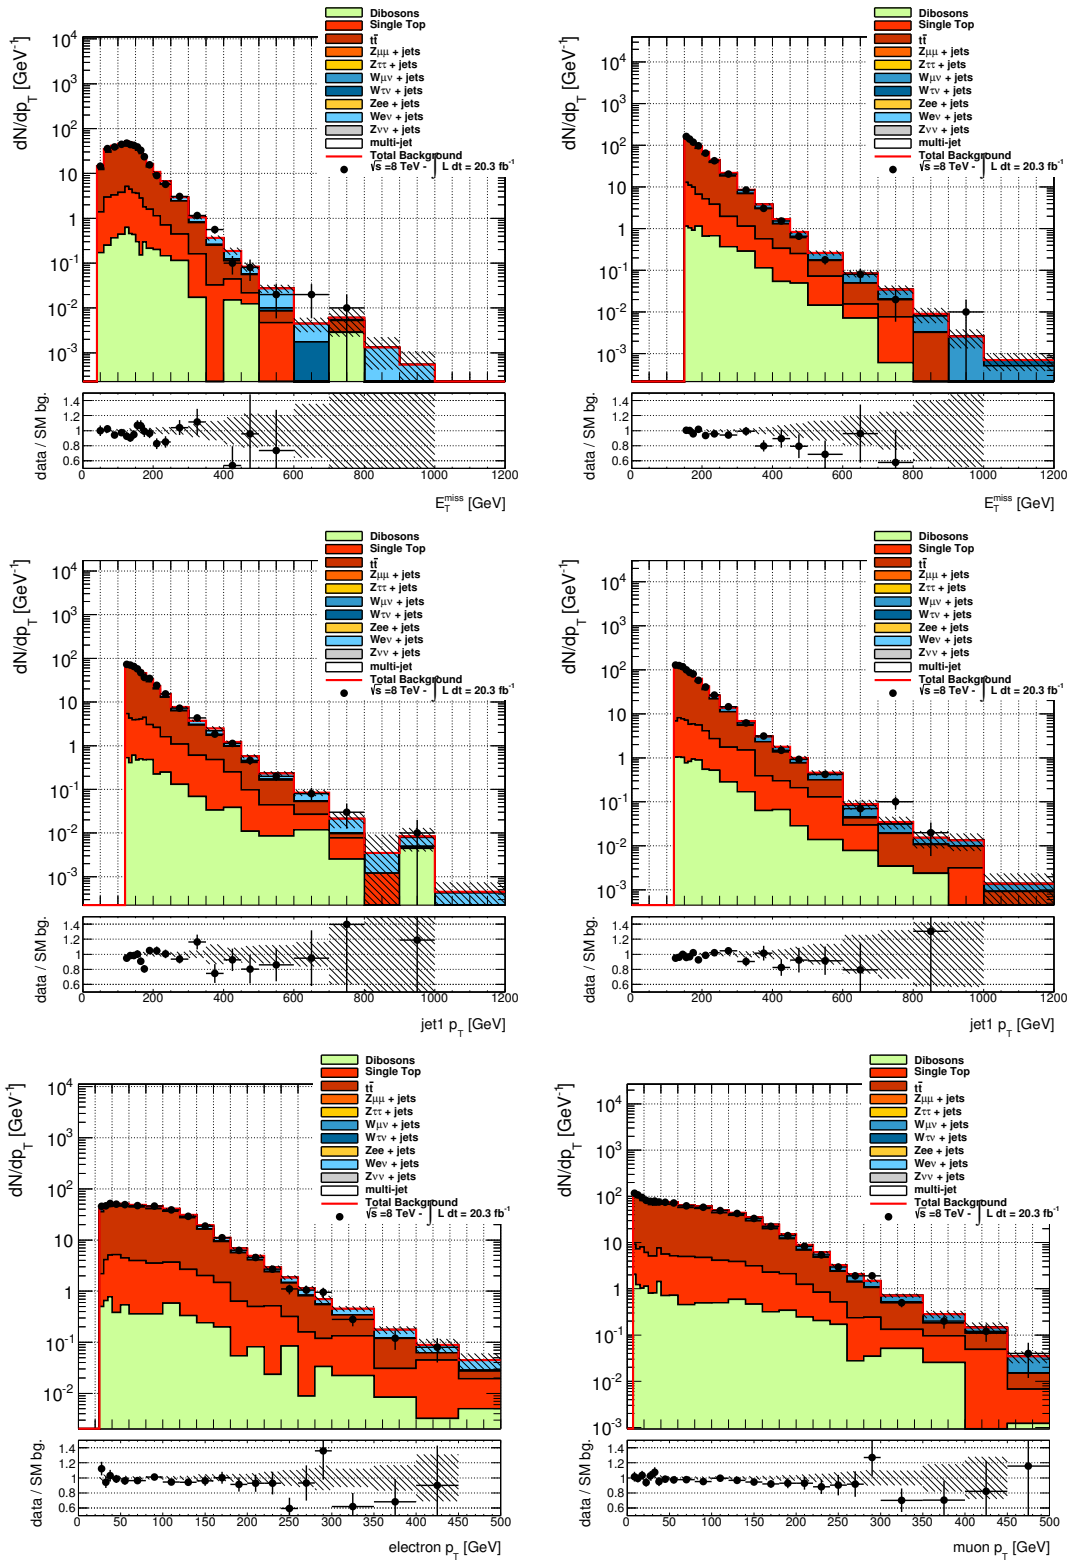


Figure 7.18: The E_T^{miss} (top), leading jet p_T (middle), and lepton p_T (bottom) distributions for the $W \rightarrow e\nu [Z]$ (left) and $W \rightarrow \mu\nu$ (right) top control regions. The $t\bar{t}$ background is observed to always be dominant.

E_T^{miss} threshold [GeV]	Region	Scale Factor	Stat	W+jets	Tagging uncertainties		
					b-jet	c-jet	l-jet
150	$W \rightarrow e\nu [Z]$	0.991 ± 0.129	± 0.019	± 0.074	± 0.086	± 0.026	± 0.029
150	$W \rightarrow \mu\nu$	0.982 ± 0.125	± 0.016	± 0.069	± 0.085	± 0.028	± 0.032
300	$W \rightarrow e\nu [Z]$	0.921 ± 0.184	± 0.067	± 0.126	± 0.089	± 0.049	± 0.050
300	$W \rightarrow \mu\nu$	0.952 ± 0.196	± 0.068	± 0.133	± 0.093	± 0.056	± 0.058
400	$W \rightarrow e\nu [Z]$	1.106 ± 0.372	± 0.204	± 0.225	± 0.122	± 0.123	± 0.125
400	$W \rightarrow \mu\nu$	0.845 ± 0.336	± 0.171	± 0.204	± 0.116	± 0.118	± 0.120
500	$W \rightarrow e\nu [Z]$	0.955 ± 0.776	± 0.507	± 0.415	± 0.234	± 0.239	± 0.245
500	$W \rightarrow \mu\nu$	0.740 ± 0.688	± 0.433	± 0.350	± 0.227	± 0.235	± 0.238

Table 7.12: The top scale factor and associated uncertainties for the $W \rightarrow e\nu [Z]$ and $W \rightarrow \mu\nu$ top control regions, as a function of the E_T^{miss} threshold.

7.6.8 Estimating other backgrounds

The diboson background is a small contribution to the signal region, and is taken directly from MC. However, this contribution is not negligible in the control regions. The $WZ \rightarrow \ell\nu\nu\nu + \text{jets}$ process naturally passes the $W \rightarrow \ell\nu$ control region selection, and grows more important at high E_T^{miss} where the $Z \rightarrow \nu\nu$ becomes more significant. Meanwhile, the $ZZ \rightarrow \ell\nu\nu + \text{jets}$ process naturally passes the $Z \rightarrow \ell\ell$ control regions. The other diboson processes also contribute, but the above two are the simplest explanations for how the diboson background becomes non-negligible in the control regions.

Unfortunately, sufficient statistics were not available to create diboson control regions. This will become more important in Run-II, and should be properly estimated at that time. For this iteration, diboson MC was also used for the control regions, with the systematic uncertainties estimated as described in Section 7.7.8.

The remaining $Z \rightarrow \ell\ell + \text{jets}$ contribution to the signal region is negligible, and is taken from MC. Data-driven estimations from the $Z \rightarrow \ell\ell$ control regions would be possible, but they were not done due to the added time requirement for this truly negligible background.

7.7 Systematic uncertainties

The mono-jet analysis is sensitive to both experimental and theoretical systematic biases, which are addressed through the use of systematic uncertainties. These uncertainties account for a wide range of topics, from experimental effects such as the jet energy scale to theoretical unknowns such as PDF effects. All of these sources will be addressed in the following Sections. Summary tables for each major signal region process are provided in Section 7.7.11.

7.7.1 Jet energy scale and resolution

The Jet Energy Scale (JES) and Jet Energy Resolution (JER) are two important systematic uncertainties for the mono-jet analysis due to the jet requirements. The data-driven estimation procedure and the similarities in jets from $W + \text{jets}$ and $Z + \text{jets}$ events (which dominate the signal and control regions) reduces the impact, as the jet uncertainties largely cancel in the transfer factor. Only small backgrounds taken directly from MC are fully impacted by the JES and JER uncertainties.

The JES uncertainties are discussed in detail in Section 6.4. The JES uncertainty is typically large

at low p_T where pileup dominates, shrinks at medium to high p_T , and then grows for extreme values of p_T where sufficient statistics are not available for the calibration. This analysis does not make use of the Final2012 JES due to the timescale on which it was released. Instead, the previous Moriond2013 JES calibration and uncertainties were used. The category eigendecomposition formulation of the uncertainties as described in Appendix E.3 is used, allowing for the use of a reduced set of eigenvectors to represent the full set of variations, but with information retained as to the general source of the components. This enables the possibility of future combinations with other analyses, while reducing the complexity of the JES uncertainty evaluation within the mono-jet analysis. For simplicity, the JES will typically be listed as a single number where all of the sub-components are added in quadrature, but the actual procedure followed includes the individual variation of each sub-component.

The JER is discussed in Section 6.5. The JER uncertainty is typically large at low p_T where the noise term dominates thanks to the pileup contribution, shrinks at medium p_T where the stochastic term dominates, and then becomes very small at high p_T where the constant term remains. Due to the removal of the jet veto from the mono-jet analysis, the impact of this uncertainty is significantly reduced. This uncertainty is used to smear jets, and the resulting one-sided systematic is then symmetrized to represent a $\pm 1\sigma$ shift.

The application of either the JES or JER to a given jet is also propagated to the E_T^{miss} , which is then re-calculated from the new inputs. This allows for a more coherent treatment of the E_T^{miss} dependence on the jets. The impact of the JES or JER on the E_T^{miss} , and thus the change in selected events, is listed as a part of the respective JES or JER uncertainty.

Thanks to the data-driven background estimation, the impact of the JES and JER uncertainties on the mono-jet analysis enters only at the level of $\mathcal{O}(1\%)$.

7.7.2 Lepton identification and scale

For the signal region, only the veto definition of electrons or muons is relevant. A variation on the electron or muon scale is thus protected by the transfer factor, just like the JES and JER, thanks to the veto on additional leptons in the control regions.

The lepton identification efficiency, on the other hand, is not protected. A change in the lepton identification leads to events migrating between the signal and control regions. As such, the ratio of the signal to control region event yields will change. Lepton identification efficiency uncertainties are applied, resulting in uncertainties up to 1% for muon estimates and between 1 and 2% for electron estimates.

The lepton uncertainties quoted later on correspond to the lepton scales (including the propagation to and re-calculation of the E_T^{miss}) and the identification efficiencies added in quadrature.

7.7.3 E_T^{miss} SoftTerm scale and resolution

Given that the E_T^{miss} is a composite object, there are many sources of systematics. However, a consistent treatment of the E_T^{miss} requires the coherent variation of a given object both for the object itself and as apart of the E_T^{miss} calculation. As such, the individual object terms of which the E_T^{miss} is comprised are dealt with through uncertainties related to the object in question. Only the SoftTerm, which is unique to the E_T^{miss} , is referred to as an independent uncertainty.

The E_T^{miss} SoftTerm is the collection of topo-clusters remaining in the calorimeters which have not

been assigned to any of the other terms. This is primarily composed of but not limited to soft or pileup jets. Unlike the other terms, no object calibration is applied, and the topo-clusters are vectorially summed at the LC scale. This is a very complex term, due to the number of possible contributions, and is the main area of study for future improvements to reduce the dependence of the E_T^{miss} on pileup.

The SoftTerm scale and resolution are thus purely associated to the E_T^{miss} . A proper treatment would provide a means of migration between the SoftTerm and the other hard terms, such as a downward shift in the JES causing a jet to drop below the 20 GeV threshold, and thus becoming part of the SoftTerm. At the time of the mono-jet analysis, this was not possible. As such, the recommended scale and resolution uncertainties were applied. The resulting uncertainties are typically below 1%, thanks to the mono-jet topology selecting events with a significant amount of hard activity, thus decreasing the relevance of the SoftTerm.

7.7.4 Trigger efficiency

As detailed in Section 7.4.1, the first signal region was constructed such that the E_T^{miss} trigger is 98% efficient, and shortly thereafter the trigger reaches 99% efficiency. The data and MC reach the 98% efficiency point at slightly different E_T^{miss} values, with the MC being slightly more efficient. This difference is covered by a 1% trigger systematic uncertainty whenever the E_T^{miss} trigger is used.

When the single-electron trigger combination is used, the efficiencies and uncertainties have been externally derived by electron experts, and their inputs are used. When estimating a process in the signal region using the E_T^{miss} trigger from a process in an electron control region using the single-electron triggers, the two uncertainties are added in quadrature.

7.7.5 Data taking conditions

Luminosity

The luminosity uncertainty for the 2012 dataset is 2.8%, as measured using the same technique as Reference [33]. This uncertainty is irrelevant for the data-driven estimates, due to using the ratio of two regions which both use the full dataset and unrescaled triggers, thus cancelling any luminosity uncertainty. Backgrounds taken purely from MC include this 2.8% normalization uncertainty, but this is typically a small effect compared to other sources of uncertainty applicable to these processes.

Pileup reweighting

The MC used in this analysis is generated with a pileup distribution which is significantly different from that which was observed in data. In order to remedy this, all events are given a pileup weight, thus shifting the distributions in MC to match the data. The uncertainties in this procedure are again mostly cancelled through the use of the data-driven estimation and a ratio of MC. Only the small backgrounds taken purely from MC are significantly effected, but the pileup uncertainties on such backgrounds are small compared to other uncertainty sources.

7.7.6 Event cleaning

Track veto

The track veto described in Section 7.4.7 is observed to have an efficiency difference of 0.4% when comparing data and MC. In order to be conservative, an uncertainty of 0.5% is assumed for the $Z \rightarrow \nu\nu + \text{jets}$ background, while a very conservative uncertainty of 1% is used for all other backgrounds.

Non-Collision Background

The NCB leading jet cleaning cuts described in Section 7.6.6 reject 0.4% of good events, as measured in a control region which is independent of NCB events. This 0.4% fake rate is used as an uncertainty on the cleaning procedure. This does not apply to backgrounds estimated with the data-driven approach due to the ratio of MC, where the same fake rate is expected, thus cancelling the systematic uncertainty. Small backgrounds taken purely from MC are susceptible to this 0.4% uncertainty, but this is a negligible effect compared to the other systematics for such processes.

7.7.7 Top normalization

The top control regions detailed in Section 7.6.7 are used to constrain the normalization uncertainty of the top backgrounds. The quadrature sum of the $W + \text{jets}$ heavy flavour and b -tagging uncertainties results in the following normalization systematics for the top background, as a function of the region E_T^{miss} threshold:

- 20% for $150 \leq E_T^{\text{miss}}/\text{GeV} \leq 300$
- 50% for $350 < E_T^{\text{miss}}/\text{GeV} < 500$
- 100% for $E_T^{\text{miss}} \geq 500 \text{ GeV}$

7.7.8 Diboson normalization

The diboson background is only a small contribution to the mono-jet signal regions, but can be up to $\mathcal{O}(10\%)$ of the control regions at high E_T^{miss} or boson p_T thresholds. It is thus important to have a reasonable estimate for the systematic uncertainties associated with the diboson process. A standard procedure has been devised and is used by many other ATLAS searches for the estimation of diboson uncertainties from splitting the effects into scale variations, PDF uncertainties, and finite parton limitations. The mono-jet analysis has implemented and followed this procedure.

Scale variations

Complete sets of *Sherpa* diboson samples have been produced at truth level with nominal, up, and down variations of the ME to parton shower matching, factorization, and renormalization scales. The maximum difference between the full set of nominal samples and the full set of either up-variation or down-variation samples for a given scale is taken to be the uncertainty associated with the scale in question. The resulting uncertainties from the three different scales are then combined in quadrature to get the total scale uncertainty of the diboson process. The independent scales and quadrature sum for each signal and control region are all shown in Figure 7.19.

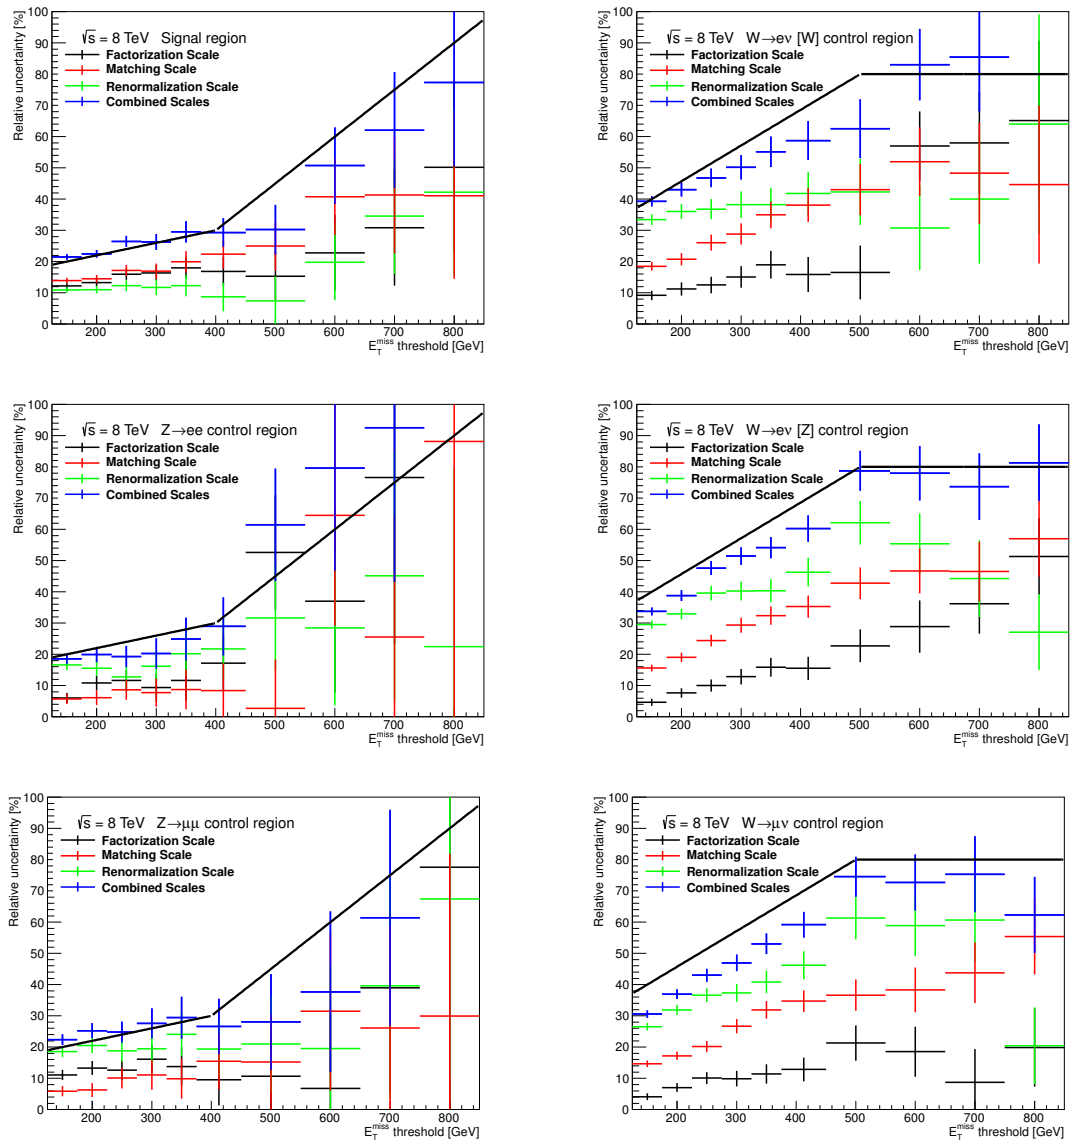


Figure 7.19: Scale uncertainties for the Z-like regions (left) and W-like regions (right). The factorization scale (black), matching scale (red), and renormalization scale (green) are shown, along with the quadrature sum of the three (blue). The uncertainty parametrization used is shown as two thick black lines on each plot.

Region	Total diboson scale uncertainty [%]
$Z \rightarrow \ell\ell$ CRs and SR	$E_T^{\text{miss}} \leq 400 \text{ GeV}$: $\frac{1}{25 \text{ GeV}} E_T^{\text{miss}} + 14$
	$E_T^{\text{miss}} > 400 \text{ GeV}$: $\frac{3}{20 \text{ GeV}} E_T^{\text{miss}} - 30$
$W \rightarrow \ell\nu$ CRs	$E_T^{\text{miss}} \leq 500 \text{ GeV}$: $\frac{100}{875 \text{ GeV}} E_T^{\text{miss}} + \frac{800}{35}$
	$E_T^{\text{miss}} > 500 \text{ GeV}$: 80

Table 7.13: Systematic uncertainty parametrizations due to scale variations for the diboson processes. The parametrization is split into Z-like regions (including the signal region) and W-like regions.

Two clear trends are discernable when comparing the diboson scale uncertainties. The signal region (which is dominated by $Z \rightarrow \nu\nu + \text{jets}$ events) and the two $Z \rightarrow \ell\ell$ control regions have a similar shape with respect to the E_T^{miss} threshold, while the three $W \rightarrow \ell\nu$ control regions have a separate shape. Uncertainty parametrizations are used, with one parametrization for Z-like regions and the other for W-like regions. The resulting parametrizations are listed in Table 7.13.

The clear divide of the uncertainty between two distinct shapes provides an insight into underlying physics behind the scale uncertainties. This can be understood by considering the primary diboson contributions to each region. The diboson processes with the highest cross-section at the LHC are, in order, WW, WZ, and ZZ. The WW process naturally contributes to the three $W \rightarrow \ell\nu$ control regions through the process $WW \rightarrow \ell\nu qq$. The $WZ \rightarrow \ell\nu qq$ process also contributes, but the cross-section is much smaller, so WW dominates. The ZZ process does not naturally pass the selection, and is thus further suppressed.

The invariant mass cut $66 < m_{\ell\ell}/\text{GeV} < 116$ in the $Z \rightarrow \ell\ell$ control regions removes the largest part of contribution of the WW process, as a pair of leptons from two separate W bosons typically does not result in an invariant mass near the Z boson peak. The $WZ \rightarrow qq\ell\ell$ process naturally passes the selection, as does the $ZZ \rightarrow \ell\ell qq$. Due to the difference in cross-section, the WZ process dominates for $Z \rightarrow \ell\ell$ control regions. In the signal region, the WW process is suppressed due to the two leptons which must both be missed in the veto. However, the $WZ \rightarrow qq\nu\nu$ and $ZZ \rightarrow \nu\nu qq$ processes both naturally pass the selection. Once again, the WZ process dominates due to the higher cross-section.

To summarize, the diboson contribution to the $W \rightarrow \ell\nu$ control regions is primarily from the WW process, while the $Z \rightarrow \ell\ell$ control regions and the signal region primarily contain the WZ process. These are very different backgrounds, first of all noting that the first is charge neutral and the second is charged, and come with different diagrams. This is not a solid explanation as to why the WZ process has significantly larger systematic uncertainties than the WW process, but a similarly large set of uncertainties for the WZ process with respect to WW or ZZ has been seen by other groups, who have investigated much further without obtaining a solid explanation.

Note that it is more conservative to assume that this difference is real than to assume that the larger $W \rightarrow \ell\nu$ uncertainties are applicable for all scenarios. The diboson background is the most important at high E_T^{miss} , where the $W \rightarrow \ell\nu$ control regions have the most weight in the $Z \rightarrow \nu\nu + \text{jets}$ estimation. The difference in the uncertainties in the W-like and Z-like regions then results in less cancellation when performing the data-driven background estimates. If the larger $W \rightarrow \ell\nu$ uncertainty is used for the Z-like regions, then the cancellation would be increased, and the final uncertainty in the signal region due to diboson processes would be reduced.

Alternate PDF set	E_T^{miss} cut		
	150 GeV	300 GeV	600 GeV
NNPDF23NLO	+2.4%	+1.4%	-5.5%
MSTW2008NLO	+1.2%	+1.1%	-3.9%

Table 7.14: PDF uncertainties on the diboson samples resulting from the difference between CT10NLO and the specified alternate NLO PDF set central value as a function of the E_T^{miss} threshold in the signal region. The observed systematic uncertainties are negligible in comparison to the diboson scale uncertainties as per Table 7.13.

PDF uncertainties

Given that the diboson samples were produced with CT10NLO, which is internally observed to provide a good estimation of the PDF behaviour at the LHC, the PDF uncertainties are expected to be smaller than the scales. Inter-PDF uncertainties are typically much larger than intra-PDF, and thus the difference between CT10NLO and both NNPDF23NLO and MSTW2008NLO is used to determine the approximate level of the PDF systematic uncertainties.

As expected, the resulting uncertainty is significantly smaller than for the scales, as detailed for the signal region in Table 7.14. These studies have not been conducted for other regions, as it is assumed that they are at a similar level, and thus are negligible compared to the scale uncertainties.

Finite parton multiplicity uncertainties

The procedure used when producing Sherpa MC samples with extra jets, such as the diboson truth level samples in question, is to specify the maximum number of extra partons to consider. The diboson samples used for the systematic uncertainties are generated with up to two extra partons. As such, Sherpa generates the base process such as $ZZ \rightarrow llll$ with 0, 1, or 2 additional partons, and then merges the three independent samples. The mono-jet analysis is mostly insensitive to low- p_T jets due to the removal of the jet multiplicity requirement, while remains dependent on high- p_T jets due to the retention of leading jet p_T cuts and E_T^{miss} cuts. Given that hard partons from matrix element generators are generally considered to be more reliable and to better model data than extra high- p_T jets from the parton shower, it is possible that this limitation introduces a bias. A standard procedure is to estimate this by producing one set of diboson samples with 3 additional partons rather than 2, and using the difference to probe the systematic uncertainty due to the finite number of partons.

This is a reasonable idea, but is not relevant for the mono-jet analysis. The reconstruction level diboson samples used by the mono-jet analysis were produced later than the truth level systematic samples, and with more details. As such, the samples used for the actual background estimation already have up to 3 additional partons. It thus is not sensible to add an uncertainty due to limited parton multiplicity derived from the difference between 2 and 3 parton states to a set of nominal samples with 3 partons.

No set of diboson samples with 4 partons exists, and thus the procedure cannot be simply extended. Regardless, this is expected to be a small contribution, especially in comparison to the scale uncertainties. The mono-jet analysis is dominated by low jet multiplicity events, even without any jet multiplicity veto applied. As such, allowing for 3 extra partons in addition to the diboson process is expected to already cover the large majority of the necessary parameter space. It is thus assumed that any bias coming from the finite parton multiplicity in the generation of the diboson samples is thus negligible in comparison

to the estimated scale uncertainties.

7.7.9 Other background normalizations

$Z \rightarrow \ell\ell + \text{jets}$ normalization

The $Z \rightarrow \ell\ell + \text{jets}$ background is negligible in the signal region and all control regions, except for the control regions where it is the process of interest, and thus is taken directly from MC. The normalization systematic uncertainty assumed for these processes is 40%.

QCD normalization

The QCD data-driven estimation, detailed in Section 7.6.5, makes this background negligible in all but the $W \rightarrow e\nu [Z]$ control region. The systematic uncertainty on the number of QCD events passing the relevant mono-jet region selection is taken to be 100% in order to provide a conservative estimate, while actual estimates from the QCD control region studies all result in uncertainties closer to half of that value.

Non-collision background normalization

The NCB estimation procedure, as described in Section 7.6.6, renders this background negligible. The remaining NCB events passing the signal or control region selections are assumed to have a 100% normalization uncertainty.

7.7.10 Theoretical systematics on the transfer factor

The transfer factor TF_ℓ is the technique used to account for the differences in selection between the control region and signal region when performing data-driven estimations, as detailed in Section 7.6.1. This factor is constructed by taking the ratio between MC events for the desired process in the signal region, and MC events for the control region process. The use of MC means that this ratio is susceptible to various theoretical uncertainties, although most are expected to cancel to first order due to being present in both the numerator and denominator of the ratio. These different possible systematic biases will now be examined in further detail.

Electroweak radiative correction differences between the W and Z Bosons

The electroweak radiative corrections to the $W + \text{jets}$ and $Z + \text{jets}$ production cross-sections can become significant in highly energetic events. This becomes particularly important given that the $W \rightarrow \ell\nu$ control regions have more statistics than the $Z \rightarrow \ell\ell$ control regions at high boson p_T , and thus are the primary contribution to the $Z \rightarrow \nu\nu + \text{jets}$ estimate at high E_T^{miss} . If the W/Z ratio changes significantly with Boson p_T , then this can introduce a bias.

This scenario has been considered by Reference [84] at 14 TeV, although the authors clearly state that the same conclusions apply to 7 and 8 TeV collisions. Rather than applying these theoretical scale factors, the differences are taken as a systematic uncertainty, varying from 1% to 6% when moving from a boson p_T of 150 GeV to 1 TeV.

E_T^{miss} cut [GeV]	150	200	250	300	350	400	500	600	700
ktfactup	0.32	0.3	0.29	0.28	0.33	0.37	0.39	0.43	0.5
ktfactdn	-0.49	-0.45	-0.4	-0.33	-0.38	-0.42	-0.44	-0.56	-0.68
qfactup	0.01	0.01	0.03	0.05	0.05	0.05	0.04	0.04	0.05
qfactdn	-0.21	-0.17	-0.13	-0.08	-0.1	-0.12	-0.14	-0.16	-0.19
Total	± 0.53	± 0.48	± 0.42	± 0.34	± 0.39	± 0.44	± 0.46	± 0.58	± 0.71

Table 7.15: Scale uncertainties on the transfer factor to signal region processes, using reweighting factors derived from **Alpgen** and applied to **Sherpa** MC samples. The overall average scale uncertainty is found to be at the level of approximately 0.5%.

PDF, factorization scale, and renormalization scale

Multiple methods were tested for the investigation of the impact of PDF, factorization scale, and renormalization scale variations on the transfer factor. The simple comparison of **Alpgen** [88] vs **Sherpa** proved undesirable due to limited statistics, and thus was discarded. The more rigorous procedure of creating **Sherpa** samples with scale variations, and with additional comparisons to **MCFM** [89] samples where necessary, proved more beneficial. The PDF uncertainties were determined to be 0.64%, while the scale uncertainties were determined to be between 0.6 and 0.7%, for a total combined uncertainty of just under 1%. However, these numbers were not used in the end, again due to limited statistics in the high E_T^{miss} regions.

The procedure which was ultimately used has a basis in previous ATLAS searches, and is similar to what was done for the diboson systematic uncertainty estimation. The uncertainty is again split into scale, PDF, and finite parton effects. However, unlike the diboson process, $W + \text{jets}$ and $Z + \text{jets}$ backgrounds are much better understood and samples are produced with much higher statistics. These combine to result in much smaller uncertainties, especially after taking the ratio of the signal region and control region process.

It was determined that it was easier to apply scale variations to **Alpgen** rather than to **Sherpa**. As such, **Alpgen** samples were generated with the desired variations, then a procedure was determined allowing for the coherent transfer of reweighting factors from **Alpgen** to **Sherpa**, thus benefiting from significantly higher statistics. The observed scale uncertainties resulting from this procedure are at the level of approximately 0.5%, with all of the numbers provided in Table 7.15.

The PDF systematic uncertainties are determined as for the diboson process, namely by comparing the nominal samples produced with **CT10NLO** to reweighted samples using either **NNPDF23NLO** or **MSTW2008NLO**. As before, this assumes that the inter-PDF uncertainties are larger than the intra-PDF uncertainties, as is generally the case. The resulting systematic uncertainties obtained from inter-PDF comparisons are provided per process in Table 7.16.

The finite parton multiplicity systematic, in which samples with up to 4 versus 5 partons are compared, is the final remaining component. This is not particularly relevant to the mono-jet analysis, given that the vast majority of the events have four or fewer jets in the final state. Given that very few of the 5 parton events pass the mono-jet selection, and thus there is insufficient statistics to derive any realistic uncertainties, this potential systematic source is neglected as irrelevant.

Process		Lower bound on E_T^{miss} [GeV]								
SR	CR	150	200	250	300	350	400	500	600	700
$Z \rightarrow \nu\nu$	$W \rightarrow \mu\nu$	0.222	0.241	0.265	0.344	0.373	0.346	0.603	0.481	0.255
$Z \rightarrow \nu\nu$	$W \rightarrow e\nu$	0.314	0.317	0.324	0.41	0.489	0.598	0.889	1.11	0.624
$Z \rightarrow \nu\nu$	$Z \rightarrow \mu\mu$	0.297	0.143	0.117	0.058	0.022	0.213	0.182	1.062	0.728
$Z \rightarrow \nu\nu$	$Z \rightarrow ee$	0.448	0.44	0.428	0.387	0.256	0.357	0.23	0.468	0.631
$W \rightarrow \tau\nu$	$W \rightarrow e\nu$	0.253	0.35	0.703	0.942	1.191	1.541	2.143	2.191	2.378
$W \rightarrow e\nu$	$W \rightarrow e\nu$	0.495	0.33	0.793	0.935	1.015	1.007	2.156	2.384	4.589
$W \rightarrow \mu\nu$	$W \rightarrow \mu\nu$	2.332	2.033	1.429	0.984	0.294	0.453	0.604	0.449	3.26

Table 7.16: PDF uncertainties in percent on the various signal region processes as a function of the E_T^{miss} threshold. The numbers are the maximum difference between CT10NLO and the two alternate PDF sets of NNPDF23NLO and MSTW2008NLO.

SR: $Z \rightarrow \nu\nu$ CR: $Z \rightarrow \mu\mu$	Lower bound on E_T^{miss} [GeV]								
	150	200	250	300	350	400	500	600	700
Stat (data)	0.75	1.2	1.9	3.0	4.3	6.1	12	24	59
Stat (MC)	0.31	0.44	0.62	0.82	1.2	1.5	2.0	3.7	9.8
JER	0.069	0.19	0.36	0.24	1.4	0.13	1.4	3.5	9.2
JES	0.71	0.74	0.60	0.84	0.40	0.61	1.3	2.6	7.3
E_T^{miss}	0.29	0.56	0.39	0.69	0.49	0.72	1.1	0.48	0.62
Top	0.37	0.33	0.29	0.41	0.58	0.40	1.1	1.5	0.0
Diboson	0.79	1.0	1.3	1.6	1.9	2.3	3.9	5.9	15
Other CR bgd	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Lepton	0.88	0.91	0.95	0.99	1.0	1.1	1.1	1.2	1.4
Trigger	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
NCB	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Theory	0.81	0.77	0.76	0.76	0.76	0.78	0.78	1.3	1.0

Table 7.17: Summary of uncertainties, in percent, on the $Z \rightarrow \nu\nu + \text{jets}$ process in the signal region as estimated from the $Z \rightarrow \mu\mu$ control region. All numbers are shown to two significant digits.

7.7.11 Summary tables for dominant signal region processes

The uncertainties on the dominant $W + \text{jets}$ and $Z \rightarrow \nu\nu + \text{jets}$ processes are provided in summary form in this Section. Most of the names are self-explanatory given the previously mentioned systematics, but some combinations are noted here:

- Other CR background: uncertainties on the QCD and $\gamma + \text{jets}$ processes, relevant for the $W \rightarrow e\nu [Z]$ control region. The $\gamma + \text{jets}$ uncertainty is taken to be 50%.
- Lepton: includes the lepton identification and scale, as well as the track veto.
- NCB: includes both the 100% uncertainty on the NCB estimate, as well as the 0.4% uncertainty on the cleaning procedure where relevant.
- Theory: includes all of the theoretical uncertainties affecting the transfer factor, as explained in Section 7.7.10.

The summary of uncertainties for the $Z \rightarrow \nu\nu + \text{jets}$ process as estimated from the $Z \rightarrow \mu\mu$, $W \rightarrow \mu\nu$, $Z \rightarrow ee$, and $W \rightarrow e\nu [Z]$ control regions are provided in Tables 7.17, 7.18, 7.19, and 7.20 respectively. The same summaries are provided for the $W \rightarrow e\nu + \text{jets}$, $W \rightarrow \mu\nu + \text{jets}$, and $W \rightarrow \tau\nu + \text{jets}$ in Tables 7.21, 7.22, and 7.23 respectively.

SR: $Z \rightarrow \nu\nu$ CR: $W \rightarrow \mu\nu$	Lower bound on E_T^{miss} [GeV]								
	150	200	250	300	350	400	500	600	700
Stat (data)	0.30	0.50	0.82	1.3	1.9	2.7	5.0	9.2	16
Stat (MC)	0.22	0.30	0.41	0.48	0.68	0.94	1.7	3.3	5.3
JER	0.34	0.062	0.29	0.68	0.89	0.63	0.37	0.23	1.0
JES	1.5	1.9	2.0	1.7	2.3	1.4	2.0	5.0	4.9
E_T^{miss}	0.83	0.83	0.53	0.61	0.53	0.77	1.9	2.2	2.9
Top	1.6	2.1	2.4	3.1	4.6	4.8	6.7	7.9	4.4
Diboson	1.5	2.1	2.9	3.6	4.6	5.4	7.1	9.2	9.9
Other CR bgd	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Lepton	0.52	0.56	0.63	0.60	0.57	0.55	1.1	0.85	2.3
Trigger	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
NCB	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Theory	1.3	1.3	1.3	1.3	2.2	2.2	3.2	5.1	6.1

Table 7.18: Summary of uncertainties, in percent, on the $Z \rightarrow \nu\nu + \text{jets}$ process in the signal region as estimated from the $W \rightarrow \mu\nu$ control region. All numbers are shown to two significant digits.

SR: $Z \rightarrow \nu\nu$ CR: $Z \rightarrow ee$	Lower bound on E_T^{miss} [GeV]								
	150	200	250	300	350	400	500	600	700
Stat (data)	0.78	1.2	1.9	2.8	4.1	5.6	12	17	29
Stat (MC)	0.33	0.45	0.64	0.81	1.1	1.5	1.8	2.8	5.6
JER	0.087	0.12	0.067	0.48	0.56	0.80	0.94	0.067	0.34
JES	0.58	0.43	0.25	0.45	0.54	0.22	0.56	1.7	1.2
E_T^{miss}	0.034	0.17	0.36	0.26	0.36	0.36	0.66	0.73	1.5
Top	0.40	0.36	0.31	0.31	0.39	0.14	0.51	1.2	3.1
Diboson	0.94	1.3	1.6	1.9	2.3	2.7	5.2	6.4	9.3
Other CR bgd	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Lepton	1.9	2.0	2.1	2.1	2.1	2.1	2.2	2.1	2.2
Trigger	1.0	0.18	0.15	0.14	0.12	0.099	0.066	0.052	0.046
NCB	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Theory	0.88	0.87	0.87	0.85	0.80	0.84	0.79	0.89	0.98

Table 7.19: Summary of uncertainties, in percent, on the $Z \rightarrow \nu\nu + \text{jets}$ process in the signal region as estimated from the $Z \rightarrow ee$ control region. All numbers are shown to two significant digits.

SR: $Z \rightarrow \nu\nu$ CR: $W \rightarrow e\nu$	Lower bound on E_T^{miss} [GeV]								
	150	200	250	300	350	400	500	600	700
Stat (data)	0.36	0.58	0.92	1.4	2.0	2.9	5.4	9.1	16
Stat (MC)	0.29	0.35	0.47	0.54	0.75	1.1	2.2	3.9	6.8
JER	0.18	0.046	0.17	0.21	0.098	0.61	0.43	0.19	0.75
JES	2.5	2.8	3.0	3.3	3.9	2.8	3.1	4.3	6.1
E_T^{miss}	0.13	0.12	0.32	0.54	0.76	0.85	1.6	2.0	2.3
Top	1.8	2.3	2.6	3.2	4.9	5.5	8.1	7.5	8.2
Diboson	1.4	2.0	2.8	3.7	4.8	5.8	7.5	7.4	8.4
Other CR bgd	1.9	1.8	1.9	2.0	2.2	2.3	2.4	2.5	3.1
Lepton	1.2	1.3	1.3	1.3	1.3	1.3	1.4	1.4	2.1
Trigger	1.2	0.62	0.58	0.54	0.51	0.49	0.46	0.43	0.44
NCB	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Theory	1.3	1.3	1.3	1.3	2.2	2.2	3.2	5.2	6.1

Table 7.20: Summary of uncertainties, in percent, on the $Z \rightarrow \nu\nu + \text{jets}$ process in the signal region as estimated from the $W \rightarrow e\nu [Z]$ control region. All numbers are shown to two significant digits.

SR: $W \rightarrow e\nu$ CR: $W \rightarrow e\nu$	Lower bound on E_T^{miss} [GeV]								
	150	200	250	300	350	400	500	600	700
Stat (data)	0.49	0.84	1.3	2.1	3.0	4.1	7.8	16	25
Stat (MC)	0.53	0.77	1.1	1.4	2.2	3.9	6.5	13	22
JER	0.41	0.019	0.38	0.69	0.76	0.077	1.45	4.4	5.1
JES	1.3	2.3	2.4	2.1	1.8	3.6	3.0	12	20
E_T^{miss}	0.34	0.35	0.14	0.37	0.57	0.87	0.92	1.2	4.2
Top	2.3	2.4	2.3	2.8	4.2	4.1	5.6	10	4.1
Diboson	2.1	3.1	4.1	5.5	6.2	6.7	8.6	11	13
Other CR bgd	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Lepton	6.3	6.1	6.0	6.1	6.2	6.4	6.5	7.8	7.5
Trigger	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
NCB	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Theory	0.90	0.82	1.1	1.2	1.3	1.3	2.3	2.5	4.7

Table 7.21: Summary of uncertainties, in percent, on the $W \rightarrow e\nu + \text{jets}$ process in the signal region as estimated from the $W \rightarrow e\nu [W]$ control region. All numbers are shown to two significant digits.

SR: $W \rightarrow \mu\nu$	Lower bound on E_T^{miss} [GeV]								
CR: $W \rightarrow \mu\nu$	150	200	250	300	350	400	500	600	700
Stat (data)	0.30	0.51	0.83	1.3	1.9	2.7	5.0	9.4	16
Stat (MC)	0.47	0.68	0.99	1.1	1.6	2.3	3.9	7.1	12
JER	0.39	0.28	0.94	0.99	1.3	2.9	1.2	0.69	6.7
JES	0.79	0.53	0.56	0.38	1.2	2.7	3.9	8.8	2.2
E_T^{miss}	0.99	1.1	0.94	0.81	0.59	1.1	0.92	4.0	1.6
Top	1.5	2.0	2.3	2.9	4.6	4.7	6.8	8.3	4.6
Diboson	1.3	2.0	2.7	3.4	4.4	5.3	7.0	9.2	10
Other CR bgd	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Lepton	4.1	5.1	6.2	7.7	9.1	11	13	15	20
Trigger	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
NCB	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Theory	2.5	2.2	1.6	1.2	0.81	0.88	0.97	0.88	3.3

Table 7.22: Summary of uncertainties, in percent, on the $W \rightarrow \mu\nu + \text{jets}$ process in the signal region as estimated from the $W \rightarrow \mu\nu$ control region. All numbers are shown to two significant digits.

SR: $W \rightarrow \tau\nu$	Lower bound on E_T^{miss} [GeV]								
CR: $W \rightarrow e\nu$	150	200	250	300	350	400	500	600	700
Stat (data)	0.49	0.84	1.3	2.1	3.0	4.1	7.7	16	25
Stat (MC)	0.38	0.55	0.75	0.93	1.3	1.9	3.6	7.4	10
JER	0.26	0.025	0.72	0.44	0.26	0.86	0.67	1.2	1.1
JES	1.6	2.1	2.6	2.4	3.4	4.2	4.3	3.2	6.4
E_T^{miss}	0.21	0.26	0.32	0.24	0.35	0.61	0.35	0.69	3.5
Top	2.3	2.4	2.3	2.8	4.2	4.2	5.9	10	3.2
Diboson	2.1	3.1	4.1	5.5	6.2	6.7	8.7	11	13
Other CR bgd	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Lepton	1.7	1.7	1.7	1.8	1.8	1.8	1.9	2.3	3.3
Trigger	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
NCB	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0
Theory	0.80	0.83	1.0	1.2	1.4	1.7	2.3	2.3	2.5

Table 7.23: Summary of uncertainties, in percent, on the $W \rightarrow \tau\nu + \text{jets}$ process in the signal region as estimated from the $W \rightarrow e\nu[W]$ control region. All numbers are shown to two significant digits.

7.8 Results

7.8.1 Combined $Z \rightarrow \nu\nu + \text{jets}$ data-driven estimates

The data-driven background estimation techniques described in Section 7.6 provide four separate estimates of the dominant $Z \rightarrow \nu\nu + \text{jets}$ background. These four estimates are taken from orthogonal control regions, but they are not truly independent measurements. Several of the systematic uncertainties, especially the JES and JER variations, are expected to be highly correlated between all of the different control regions. Lepton scale and identification systematics for regions selecting the same type of lepton should also be correlated, while theoretical systematics relating to the differences between W and Z bosons apply to both of the $W \rightarrow \ell\nu$ control regions identically. Most importantly, all of the estimates share the same signal region, and thus the numerator of the transfer factor and all of the uncertainties therein including statistics should be correlated between the four separate estimates.

Despite these complications, a combination of the four estimates is desired. The $Z \rightarrow \ell\ell$ control regions are typically more pure and come with lower systematic uncertainties, but run out of statistics at high p_T^Z , and thus are primarily useful for the signal regions with E_T^{miss} thresholds below approximately 400 GeV. On the other hand, the $W \rightarrow \ell\nu$ control regions suffer from extra uncertainties such as electroweak radiative correction differences between W and Z processes, but have more statistics, and thus are primarily useful at high p_T^W . A scheme which enables the analysis to take advantage of the different strengths of each control region, thus always providing the best possible estimate of $Z \rightarrow \nu\nu + \text{jets}$ in the signal region for any desired E_T^{miss} threshold, would be very beneficial. In this form, the best estimate is the combination of the four input estimates resulting in the minimum uncertainty.

Any means of combining these four estimates must properly take all of the different correlations into account. The Best Linear Unbiased Estimate (BLUE) method, as outlined in Reference [90], provides a procedure designed to carry out such a combination. The procedure forms a linear combination of the input estimates, where the sum of the weights is one, providing an unbiased result. For input estimates y_i , corresponding weights w_i , and an optimal combined estimate y , the resulting formula for combining four $Z \rightarrow \nu\nu + \text{jets}$ estimates is in Equation 7.9.

$$y = \sum_{i=1}^4 w_i y_i \equiv \vec{w}^T \cdot \vec{y} \quad \left(\text{for } \sum_{i=1}^4 w_i = 1 \right) \quad (7.9)$$

The systematic uncertainty associated with a source a can be propagated to the uncertainty on the final estimate y as the uncertainty $\sigma_{y,a}$. This propagation makes use of the covariance matrix V^a associated with the source a and defined with respect to each of the input estimates. The resulting uncertainty on the final estimate y for a source a is defined in terms of the variance $\sigma_{y,a}^2$ in Equation 7.10.

$$\sigma_{y,a}^2 = \sum_{i=1}^4 \sum_{j=1}^4 w_i w_j V_{ij}^a \quad (7.10)$$

Forming the full covariance matrix V with all of the uncertainty sources then analogously defines the total uncertainty on the final estimate y as per Equation 7.11.

$$\sigma_y^2 = \sum_{i=1}^4 \sum_{j=1}^4 w_i w_j V_{ij} \equiv \vec{w}^T V \vec{w} \quad (7.11)$$

The covariance matrix V is the core aspect of the combination, and it is worth defining the structure. The diagonal elements are simply the variance of a given source, while the off-diagonal elements are the covariance between two independent sources. This is written mathematically in Equation 7.12.

$$V_{ii} = \sigma_i^2 \quad V_{ij}^{i \neq j} = \rho_{ij} \sigma_i \sigma_j \quad (7.12)$$

Once the covariance matrix has been defined, the value of the weights w_i for each estimate can be solved by minimizing the χ^2 distribution in Equation 7.13, resulting in the weight functional form listed in the same Equation. These relations are taken directly from Reference [90]. Note that the weights can become negative if one or more of the estimates have significantly larger uncertainties with respect to the other estimates.

$$\chi^2 = [x\vec{e} - \vec{x}]^T V^{-1} [x\vec{e} - \vec{x}] \quad \vec{w} = \frac{V^{-1} \cdot \vec{e}}{\vec{e}^T \cdot (V^{-1} \cdot \vec{e})} \quad (7.13)$$

The resulting linear combination weights and combined uncertainties after the application of the BLUE method to the four $Z \rightarrow \nu\nu + \text{jets}$ estimates are provided in Figure 7.20. The $W \rightarrow e\nu [Z]$ control region has negative weights for the lower half of the E_T^{miss} parameter space due to the significantly larger uncertainties with respect to the other processes, mostly due to the contribution of QCD and the tight electron definition involving strong QCD rejection cuts at the expense of statistics. The BLUE combination is shown to have a smaller associated uncertainty than any of the individual estimates over the majority of the parameter spaced, as intended by design.

When looking at Figure 7.20, it appears that the independent use of the $Z \rightarrow \mu\mu$ control region

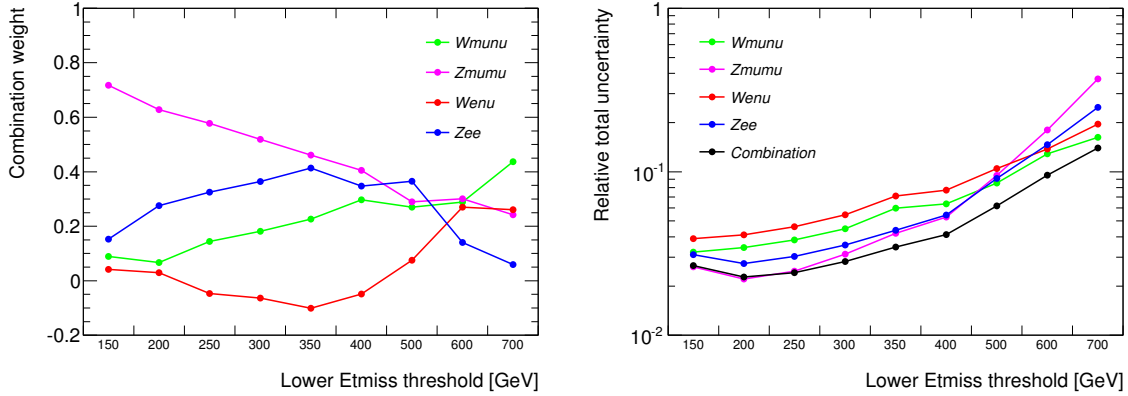


Figure 7.20: The linear combination weights (left) and total uncertainty (right) after the application of the BLUE method to the four $Z \rightarrow \nu\nu + \text{jets}$ data-driven estimates to form one combined result. The legends in the plot refer to the control region used to estimate the $Z \rightarrow \nu\nu + \text{jets}$ contribution in the signal region. The $W \rightarrow e\nu[Z]$ control region suffers from significantly increased uncertainties at low $E_{\text{T}}^{\text{miss}}$ with respect to the other regions, most notably due to QCD contamination, and thus attains a negative weight. In general, the $Z \rightarrow \ell\ell$ control regions are more useful at low $E_{\text{T}}^{\text{miss}}$ due to their lower systematics, while the $W \rightarrow \ell\nu$ control regions are more useful at high $E_{\text{T}}^{\text{miss}}$ due to their increased statistics. Note that $E_{\text{T}}^{\text{miss}}$ and $E_{\text{T}}^{\text{miss}}$ are different naming conventions for the same quantity.

provides a smaller uncertainty on the $Z \rightarrow \nu\nu + \text{jets}$ estimation for the two lowest $E_{\text{T}}^{\text{miss}}$ thresholds. This is misleading, and is not the case. The uncertainty on the $Z \rightarrow \nu\nu + \text{jets}$ process as determined by BLUE is still lower than the uncertainty on the same process as determined by the $Z \rightarrow \mu\mu$ control region. This Figure is instead showing the total uncertainty on the full background, and there is an anti-correlation between some of the systematics in the $Z \rightarrow \nu\nu + \text{jets}$ background and some of the other signal regions background. This anti-correlation works out slightly more favourably for the lowest $E_{\text{T}}^{\text{miss}}$ thresholds when using purely the $Z \rightarrow \mu\mu$ control region to estimate the $Z \rightarrow \nu\nu + \text{jets}$ contribution. This can be seen from the full background breakdown by process provided in the following Section, and specifically detailed in Table 7.24.

7.8.2 Event yields

The final expected and observed event yields in the signal region, including statistical and systematic uncertainties, are detailed in Table 7.24. Various kinematic distributions comparing the expected background to the observed data in the first signal region are provided in Figure 7.21. The $E_{\text{T}}^{\text{miss}}$ distribution trivially changes with the signal region number, as only the $E_{\text{T}}^{\text{miss}}$ threshold is modified, thus the same plots minus the $E_{\text{T}}^{\text{miss}}$ distribution are shown in Figure 7.22 for signal region seven of nine ($E_{\text{T}}^{\text{miss}} > 500$ GeV) and signal region nine of nine ($E_{\text{T}}^{\text{miss}} > 700$ GeV), which is where most of the limits are set.

7.8.3 Model-independent limits

No significant deviations from the SM expected background are observed, thus model-independent limits on the visible cross-section of new physics processes are derived. Limits are provided on $\sigma \times \mathcal{A} \times \varepsilon$, or the combination of the process cross-section, the detector acceptance, and the reconstruction efficiency

Process	Estimation method	Signal region E_{T1}^{miss} threshold [GeV]									
		150	200	250	300	350	400	500	600	700	
$Z \rightarrow \nu\nu + \text{jets}$	$W \rightarrow \mu\nu$ CR	216900 \pm 6862	77831 \pm 3096	28889 \pm 1346	11999 \pm 660	5451 \pm 414	2649 \pm 219	757 \pm 90	228 \pm 39	75 \pm 16	
$Z \rightarrow \nu\nu + \text{jets}$	$W \rightarrow e\nu$ [Z] CR	219023 \pm 9833	77953 \pm 3883	29407 \pm 1675	12399 \pm 831	5484 \pm 489	2673 \pm 260	738 \pm 101	247 \pm 40	84 \pm 19	
$Z \rightarrow \nu\nu + \text{jets}$	$Z \rightarrow \mu\mu$ CR	217770 \pm 4178	80402 \pm 1864	30142 \pm 871	12614 \pm 494	5959 \pm 317	2907 \pm 201	778 \pm 102	195 \pm 50	38 \pm 24	
$Z \rightarrow \nu\nu + \text{jets}$	$Z \rightarrow ee$ CR	217919 \pm 6124	80346 \pm 2348	30904 \pm 1070	13313 \pm 568	6082 \pm 333	3091 \pm 212	699 \pm 91	313 \pm 59	115 \pm 36	
$Z \rightarrow \nu\nu + \text{jets}$	BLUE	217768 \pm 3907	80143 \pm 1674	30244 \pm 762	12771 \pm 406	5943 \pm 244	2906 \pm 147	741 \pm 64	235 \pm 29	71 \pm 13	
$W \rightarrow \tau\nu + \text{jets}$	$W \rightarrow e\nu$ [W] CR	79344 \pm 3303	23830 \pm 1204	7717 \pm 468	2761 \pm 205	1166 \pm 107	542 \pm 57	133 \pm 19	34 \pm 8	11 \pm 3	
$W \rightarrow e\nu + \text{jets}$	$W \rightarrow e\nu$ [W] CR	23547 \pm 1729	7129 \pm 559	2374 \pm 198	875 \pm 83	367 \pm 39	168 \pm 21	43 \pm 7	9 \pm 3	3 \pm 1	
$W \rightarrow \mu\nu + \text{jets}$	$W \rightarrow \mu\nu$ CR	28340 \pm 1558	8211 \pm 532	2517 \pm 193	852 \pm 79	332 \pm 38	141 \pm 20	35 \pm 6	10 \pm 2	2 \pm 1	
$Z \rightarrow \tau\tau + \text{jets}$	MC	783 \pm 323	184 \pm 76	45 \pm 19	14 \pm 6	5 \pm 2	2 \pm 1	0 \pm 0	0 \pm 0	0 \pm 0	
$Z \rightarrow ee + \text{jets}$	MC	0 \pm 0	0 \pm 0	0 \pm 0	0 \pm 0	0 \pm 0	0 \pm 0	0 \pm 0	0 \pm 0	0 \pm 0	
$Z \rightarrow \mu\mu + \text{jets}$	MC	527 \pm 224	97 \pm 42	19 \pm 8	7 \pm 3	4 \pm 2	3 \pm 1	2 \pm 1	1 \pm 1	1 \pm 0	
Top	MC	6901 \pm 1407	2320 \pm 491	681 \pm 163	198 \pm 66	75 \pm 41	32 \pm 20	7 \pm 7	1 \pm 1	0 \pm 0	
Diboson	MC	7983 \pm 1737	3491 \pm 838	1507 \pm 402	685 \pm 204	351 \pm 117	183 \pm 70	65 \pm 35	23 \pm 16	8 \pm 7	
Multi-jet	Data-driven	6453 \pm 6453	785 \pm 785	177 \pm 177	44 \pm 44	15 \pm 15	6 \pm 6	1 \pm 1	0 \pm 0	0 \pm 0	
Non-collision	Data-driven	449 \pm 449	47 \pm 47	0 \pm 0	0 \pm 0	0 \pm 0	0 \pm 0	0 \pm 0	0 \pm 0	0 \pm 0	
Total expected	$Z \rightarrow \nu\nu$ from $W \rightarrow \mu\nu$	371226 \pm 11986	123925 \pm 4254	43926 \pm 1679	17435 \pm 780	7765 \pm 465	3725 \pm 237	1044 \pm 89	307 \pm 39	101 \pm 16	
Total expected	$Z \rightarrow \nu\nu$ from $Z \rightarrow \mu\mu$	372097 \pm 9733	126496 \pm 2787	45178 \pm 1113	18050 \pm 565	8274 \pm 348	3983 \pm 210	1065 \pm 101	273 \pm 49	64 \pm 24	
Total expected	$Z \rightarrow \nu\nu$ from $W \rightarrow e\nu$ [Z]	373349 \pm 14557	124047 \pm 5099	44444 \pm 2049	17835 \pm 973	7798 \pm 554	3749 \pm 289	1025 \pm 107	325 \pm 45	110 \pm 22	
Total expected	$Z \rightarrow \nu\nu$ from $Z \rightarrow ee$	372245 \pm 11599	126440 \pm 3474	45941 \pm 1392	18750 \pm 668	8397 \pm 368	4168 \pm 226	986 \pm 90	391 \pm 57	141 \pm 35	
Total expected	$Z \rightarrow \nu\nu$ from BLUE	372094 \pm 9941	126237 \pm 2864	45281 \pm 1094	18207 \pm 514	8257 \pm 286	3982 \pm 164	1028 \pm 64	313 \pm 30	97 \pm 14	
Data		364378	123228	44715	18020	7988	3813	1028	318	126	

Table 7.24: Expected event yields in each signal region of the mono-jet analysis, broken down into each background, and listing multiple means of estimating the dominant $Z \rightarrow \nu\nu + \text{jets}$ process. The number of observed events in data for each signal region is also listed.

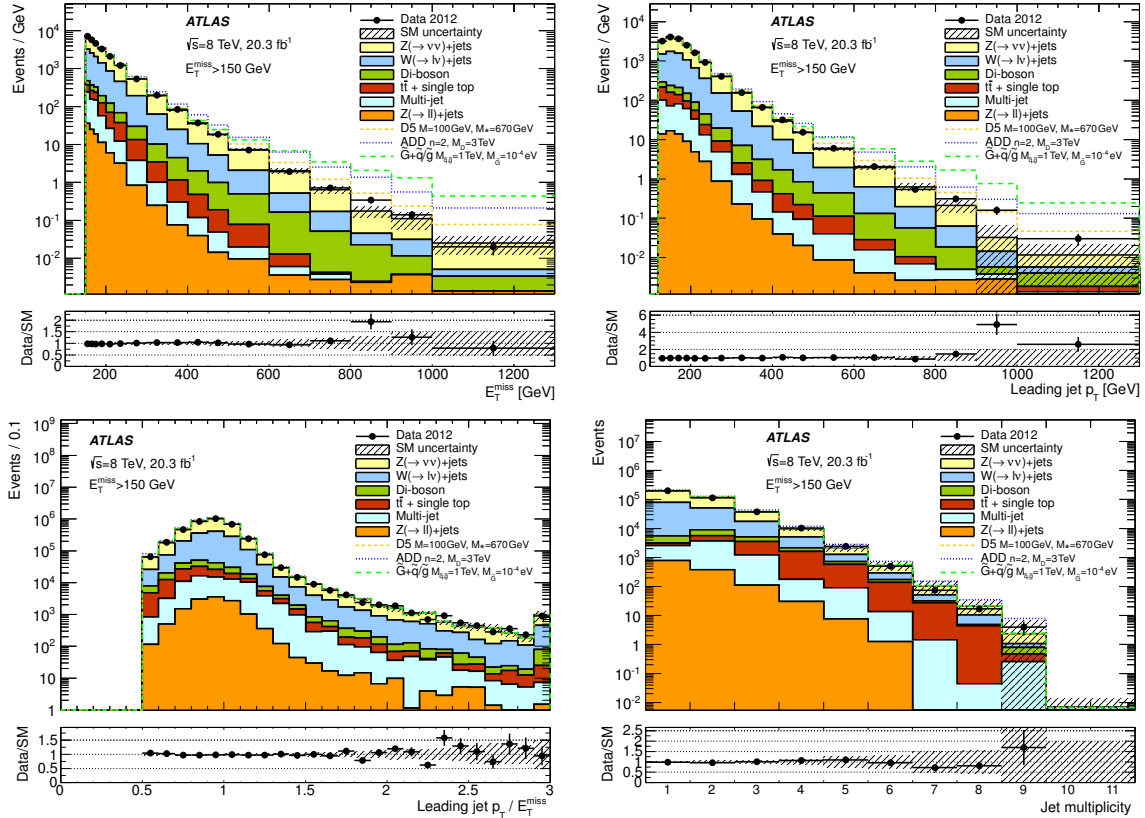


Figure 7.21: A comparison of the expected and observed kinematic distributions in the lowest mono-jet signal region, with a cut at $E_T^{\text{miss}} > 150$ GeV. The E_T^{miss} (top left), leading jet p_T (top right), ratio of leading jet p_T to E_T^{miss} (bottom left) and jet multiplicity (bottom right) are shown. Error bars on the points represent the statistical uncertainty on the data, while hatched regions represent both the statistical and systematic uncertainties on the expected background. Three signal scenarios are shown for demonstration.

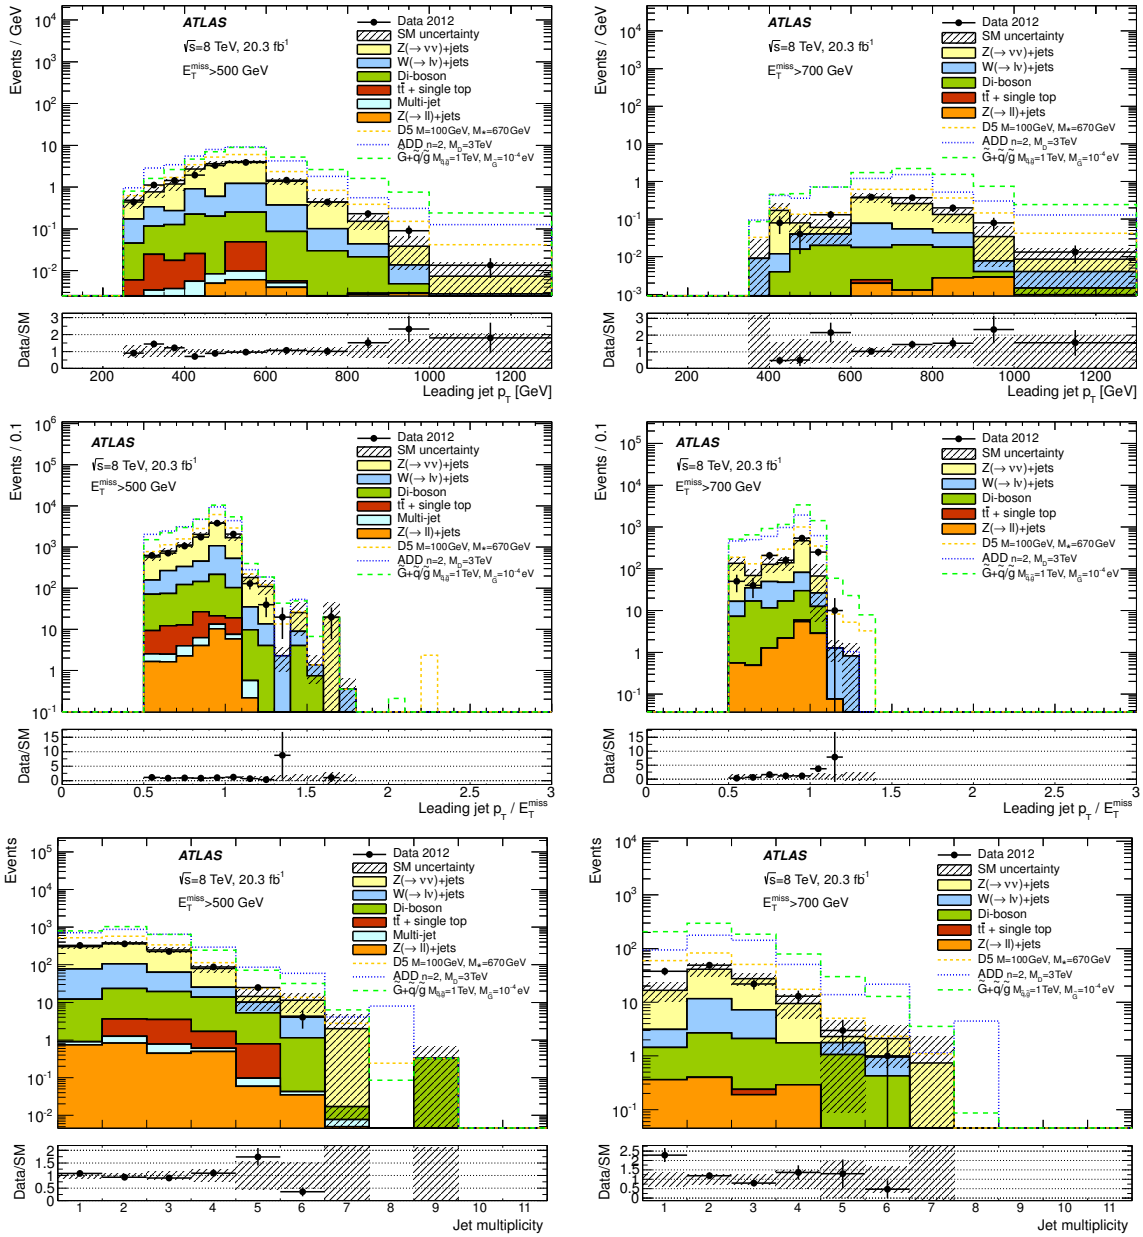


Figure 7.22: A comparison of the expected and observed kinematic distributions in the generally most sensitive mono-jet signal regions, with cuts of $E_T^{\text{miss}} > 500$ GeV (left) and $E_T^{\text{miss}} > 700$ GeV (right). The leading jet p_T (top), ratio of leading jet p_T to E_T^{miss} (middle), and jet multiplicity (bottom) are shown. Error bars on the points represent the statistical uncertainty on the data, while hatched regions represent both the statistical and systematic uncertainties on the expected background. Three signal scenarios are shown for demonstration.

Sigma \times Acceptance \times Efficiency limits		
E_T^{miss} threshold	Limit, 90% CL	Limit, 95% CL
[GeV]	[fb]	[fb]
150	590.3 (778.8)	717.3 (924.4)
200	150.7 (218.7)	184.7 (261.1)
250	67.2 (81.9)	81.4 (97.2)
300	32.1 (37.1)	38.6 (44.0)
350	13.2 (19.4)	16.2 (23.0)
400	7.3 (11.1)	9.0 (13.2)
500	4.4 (4.4)	5.2 (5.2)
600	3.2 (3.0)	3.8 (3.6)
700	2.9 (1.5)	3.4 (1.8)

Table 7.25: Model-independent limits on the visible cross-section $\sigma \times \mathcal{A} \times \varepsilon$ of new physics processes at both 90% and 95% confidence levels, as a function of the E_T^{miss} threshold used for the mono-jet signal regions. Numbers listed are for the observed (expected) limits.

respectively. The CLs [91] procedure is used for setting limits at both the 90% and 95% confidence levels. The typical collider convention is to use the 95% confidence level, but the 90% convention is used by other types of experiments, which will be discussed and compared to in Section 8.2.

The resulting model-independent limits at 90% and 95% confidence levels are provided in Table 7.25.

7.9 Conclusions

The mono-jet analysis is sensitive to the production of weakly interacting physics with several well-motivated BSM interpretations. This iteration of the analysis involves, for the first time, no upper jet multiplicity requirement. The mono-jet-like topology is instead retained through requiring the E_T^{miss} to not be aligned with any jet, and through the criteria that the leading jet p_T is at least half as large as the observed E_T^{miss} magnitude.

The analysis requires a careful estimation of the backgrounds, which are dominated by the $Z \rightarrow \nu\nu + \text{jets}$ and $W \rightarrow \ell\nu + \text{jets}$ final states, all of which are estimated by data-driven means. Smaller backgrounds include NCB, QCD, top, diboson, and $Z \rightarrow \ell\ell + \text{jets}$. NCB and QCD are estimated in a data-driven fashion, and the top contribution is constrained through control regions. Only the diboson and $Z \rightarrow \ell\ell + \text{jets}$ backgrounds are taken purely from MC, both of which are small or negligible contributions to the signal region. These contributions grow for higher kinematic thresholds, and ultimately their lack of data-driven estimation becomes one of the limiting factors of the analysis.

No significant deviation is observed between the expected background distributions and the observed data, and thus limits are set on the model-independent visible cross-section. The impact of these limits on the Dark Matter interpretation is the focus of the next Chapter.

Chapter 8

Mono-jet Dark Matter Interpretation

8.1 Introduction

This chapter is based on Reference [5]. All ATLAS material, unless otherwise mentioned, is from this document.

An introduction to the DM models considered and associated results are provided in this Chapter, while additional plots are in Appendix H. The details of producing, validating, and deriving uncertainties for the signal samples are in Appendix I.

Final states containing one jet and E_T^{miss} can be interpreted as the recoil of an ISR jet against a pair of WIMPs, as per Figure 8.1. The addition of more jets can be attributed to either a single ISR jet splitting, or the emission of additional independent ISR jet(s). There are numerous different WIMP models for such a process, varying from UV complete theories along the lines of SUSY to minimal contact interaction representations. Which model is relevant depends on the theoretical assumptions on the coupling between the SM input particles to the dark sector, which cannot be known before DM is both directly observed and studied. Instead of focusing on any particular complete theoretical description, it is useful to consider a very simple coupling between two SM input partons and two output DM particles via a very massive s-channel mediator. This process can be formulated as a contact interaction in the framework of an EFT, so long as the mediator being integrated out is beyond the energy scale of the LHC. The break-down of this assumption is discussed in detail in Appendix G.

8.1.1 Effective field theories

EFTs are very powerful tools, as they allow for the reduction of a complicated process to the minimal number of degrees of freedom. The variables of interest are reduced to two parameters per mediator type (operator), the mass of the Dark Matter particle (m_χ) and the suppression scale of the interaction (M_*). This allows for a full coverage of a large parameter space with a minimal set of MC. In this case, all possible couplings via spin 0, 1, and 2 mediators between partons and both fermionic and scalar DM are considered. The 14 possible couplings of partons to fermionic DM are listed in Table 8.1, while the 6 scalar DM operators are listed in Table 8.2. All of these operators are provided by Reference [92].




Figure 8.1: The pair-production of DM in events with a single ISR jet. Diagrams shown do not cover all possible means of producing DM, but rather focus on the methods that were investigated. The quark-quark EFT vertex (top left), gluon-gluon EFT vertex (top right), and Z' simplified model mediator (bottom) are considered.

Name	Operator	Name	Operator
D1	$\frac{m_q}{(M_*)^3} \bar{\chi} \chi \bar{q} q$	D2	$\frac{i m_q}{(M_*)^3} \bar{\chi} \gamma^5 \chi \bar{q} q$
D3	$\frac{i m_q}{(M_*)^3} \bar{\chi} \chi \bar{q} \gamma^5 q$	D4	$\frac{m_q}{(M_*)^3} \bar{\chi} \gamma^5 \chi \bar{q} \gamma^5 q$
D5	$\frac{1}{(M_*)^2} \bar{\chi} \gamma^\mu \chi \bar{q} \gamma_\mu q$	D6	$\frac{1}{(M_*)^2} \bar{\chi} \gamma^\mu \gamma^5 \chi \bar{q} \gamma_\mu q$
D7	$\frac{1}{(M_*)^2} \bar{\chi} \gamma^\mu \chi \bar{q} \gamma_\mu \gamma^5 q$	D8	$\frac{1}{(M_*)^2} \bar{\chi} \gamma^\mu \gamma^5 \chi \bar{q} \gamma_\mu \gamma^5 q$
D9	$\frac{1}{(M_*)^2} \bar{\chi} \sigma^{\mu\nu} \chi \bar{q} \sigma_{\mu\nu} q$	D10	$\frac{i}{(M_*)^2} \epsilon_{\mu\nu}^{\alpha\beta} \bar{\chi} \sigma^{\mu\nu} \chi \bar{q} \sigma_{\alpha\beta} q$
D11	$\frac{\alpha_s}{4(M_*)^3} \bar{\chi} \chi G_{\mu\nu} G^{\mu\nu}$	D12	$\frac{i \alpha_s}{4(M_*)^3} \bar{\chi} \gamma^5 \chi G_{\mu\nu} G^{\mu\nu}$
D13	$\frac{i \alpha_s}{4(M_*)^3} \bar{\chi} \chi G_{\mu\nu} \tilde{G}^{\mu\nu}$	D14	$\frac{\alpha_s}{4(M_*)^3} \bar{\chi} \gamma^5 \chi G_{\mu\nu} \tilde{G}^{\mu\nu}$

Table 8.1: EFT operators coupling SM partons to Dirac fermion WIMP DM. Couplings to Majorana fermion WIMP DM can be obtained by a simple scaling of the cross-section. All operators are discussed in more detail in Reference [92].

Name	Operator	Name	Operator
C1	$\frac{m_q}{(M_*)^2} \chi^\dagger \chi \bar{q} q$	C2	$\frac{i m_q}{(M_*)^2} \chi^\dagger \chi \bar{q} \gamma^5 q$
C3	$\frac{1}{(M_*)^2} \chi^\dagger \partial_\mu \chi \bar{q} \gamma^\mu q$	C4	$\frac{i}{(M_*)^2} \chi^\dagger \partial_\mu \chi \bar{q} \gamma^\mu \gamma^5 q$
C5	$\frac{\alpha_s}{4(M_*)^2} \chi^\dagger \chi G_{\mu\nu} G^{\mu\nu}$	C6	$\frac{i \alpha_s}{4(M_*)^2} \chi^\dagger \chi G_{\mu\nu} \tilde{G}^{\mu\nu}$

Table 8.2: EFT operators coupling SM partons to complex scalar WIMP DM. Couplings to real scalar WIMP DM can be obtained by a simple scaling of the cross-section. All operators are discussed in more detail in Reference [92].

Name	Coupling	Parton type	DM type	Operator
D1	Scalar coupling	Quarks	Fermion	$\frac{m_q}{(M_*)^3} \bar{\chi} \chi \bar{q} q$
C1	Scalar coupling	Quarks	Scalar	$\frac{m_q}{(M_*)^2} \chi^\dagger \chi \bar{q} q$
D5, C3	Vector coupling	Quarks	Fermion, Scalar	$\frac{1}{(M_*)^2} \bar{\chi} \gamma^\mu \chi \bar{q} \gamma_\mu q, \frac{1}{(M_*)^2} \chi^\dagger \partial_\mu \chi \bar{q} \gamma^\mu q$
D9	Tensor coupling	Quarks	Fermion	$\frac{1}{(M_*)^2} \bar{\chi} \sigma^{\mu\nu} \chi \bar{q} \sigma_{\mu\nu} q$
D11	Scalar coupling	Gluons	Fermion	$\frac{\alpha_s}{4(M_*)^3} \bar{\chi} \chi G_{\mu\nu} G^{\mu\nu}$
C5	Scalar coupling	Gluons	Scalar	$\frac{\alpha_s}{4(M_*)^2} \chi^\dagger \chi G_{\mu\nu} G^{\mu\nu}$

Table 8.3: EFT operators coupling SM partons to Dirac fermion or complex scalar WIMP DM which represent the full set of possible EFT interactions. All other operators listed in Reference [92] have identical kinematic acceptances, and thus can be obtained with simple cross-section scale factors. Note that both D5 and C3 are listed as while they have the same kinematic shapes, the operators do still differ in where they are applicable due to different DM interpretations. Of these two operators, D5 was used wherever kinematic information was necessary.

While the use of EFTs has reduced the parameter space significantly, there are still 20 types of couplings (operators) to consider. The form of the operator is important, as it controls the structure of the momentum transfer distribution, which in turn controls the kinematics of the DM particles. Different kinematic distributions will show up as different signals in the detector, and thus it is important to cover all such possibilities. However, while there are 20 operators, there are strong similarities between several of them. While the chirality of operators matters to Direct Detection (DD) and Indirect Detection (ID) searches for WIMPs, the LHC is unaffected. This feature does still control the cross-section of the process, but the actual kinematic distributions are negligibly modified. This can be seen by focusing on the two most important variables in the analysis, the E_T^{miss} and leading jet p_T . Figure 8.2 shows how the 14 fermionic operators can be broken down to 4 families, the 6 scalar operators can be broken down to 3 distinct operators, and how one operator between the fermionic and scalar types have the same kinematic shape. Table 8.3 lists the six representative operators, while Figure 8.2 shows their kinematic distributions. These six representative operators provide a means of generically scanning the possible parameter space, while cross-section reweightings provide translations from the representative operators to the other 14 operators. Additional plots of alternative DM masses are shown in Appendix H.1.

The second parameter in the EFT representation is the Dark Matter mass, m_χ , which is a limiting factor for pair-production to be possible at colliders. If the mass becomes too large, then the LHC does not have enough energy to produce a pair of DM particles. The value of m_χ also influences the momentum transfer, E_T^{miss} , and jet p_T distributions, and thus it is important to vary for any scan of the WIMP parameter space. The LHC has sensitivity to a wide range of DM masses, from the $\mathcal{O}(\text{GeV})$ range to $\mathcal{O}(\text{TeV})$. At higher masses than around 1 TeV, more energetic collisions are necessary in order to probe the production of such massive DM particles with sufficient statistics. This was not feasible at the 8 TeV LHC, but should become more accessible during the LHC Run-II. For m_χ at the level of a few GeV, the DM particles become essentially massless at the LHC and thus the kinematics stabilize. The limits set at the LHC actually reach down below $\mathcal{O}(\text{MeV})$ due to the stabilized kinematics, but typically 1 GeV is the lightest range considered.

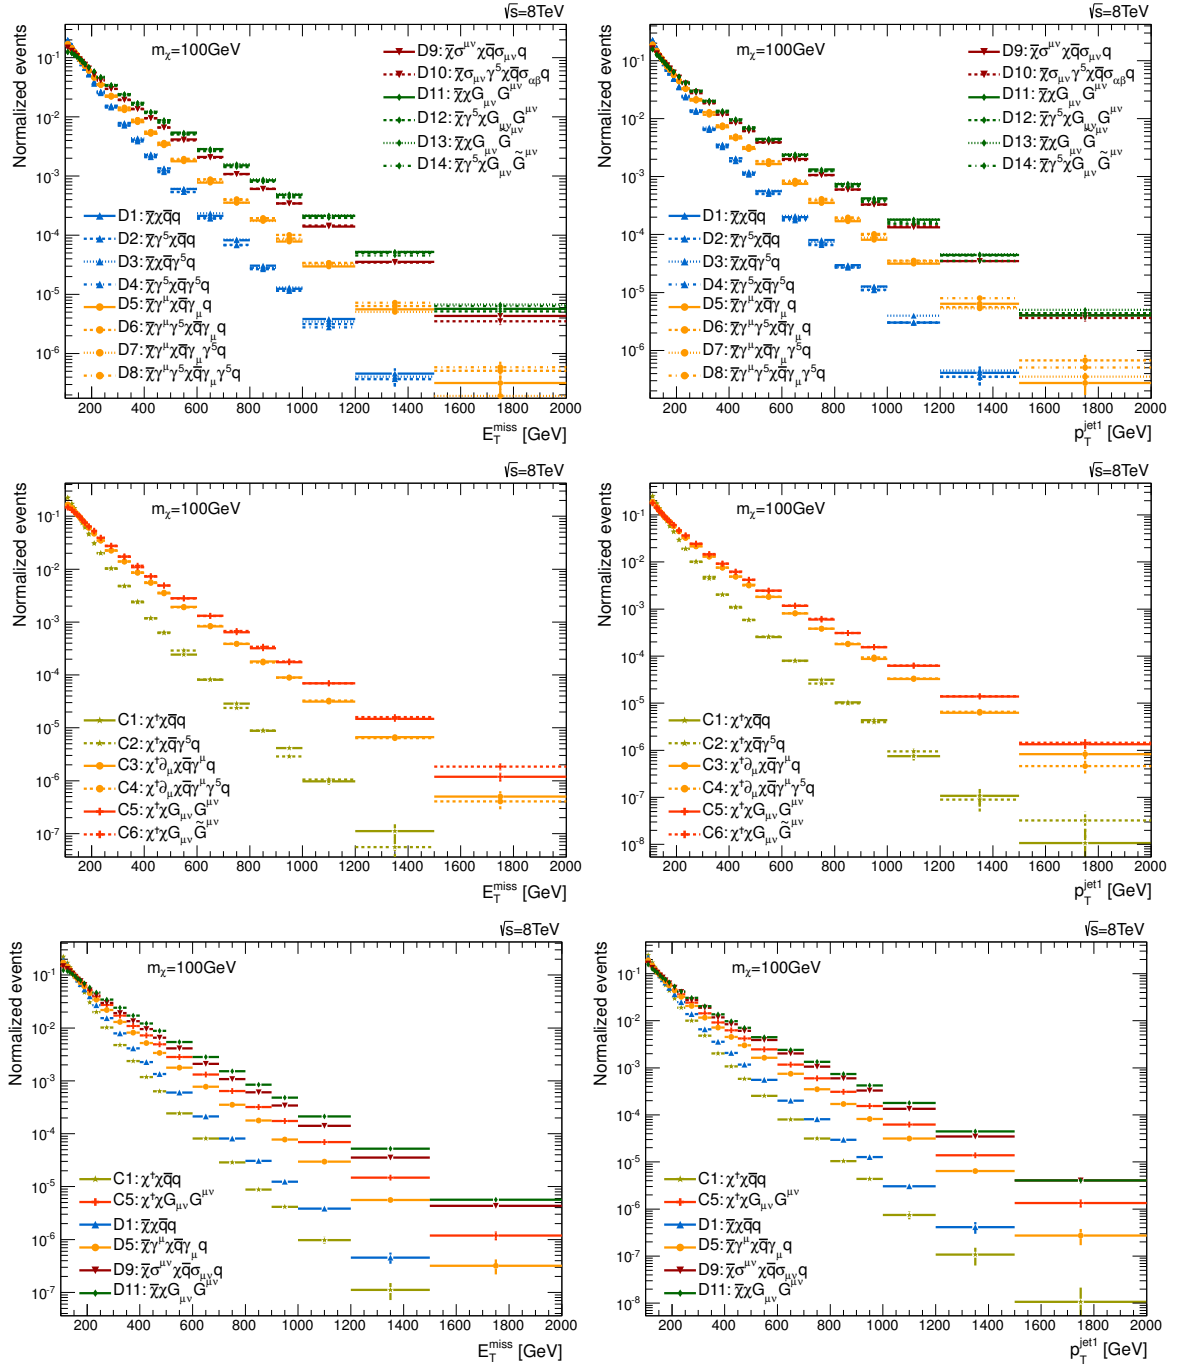


Figure 8.2: E_T^{miss} (left) and leading jet p_T (right) distributions for EFT operators coupling to fermionic DM (top), scalar DM (middle), and the set of representative operators (bottom). A DM mass of 100 GeV is shown in all cases. These kinematic distributions reveal there are four types of fermionic operators and three types of scalar operators, for a total of six types of kinematically distinct operators overall (as there is one overlap between the two types of DM). Operators of the same type are given the same colour in order to be more visible, but different line styles, as all 14 (6) operators are plotted in the top (middle).

The same set of mass points as was used for previous publications was retained for consistency, with a focus on DM masses of 10, 50, 100, 200, 400, 700, 1000, and 1300 GeV. By 10 GeV, it was found that the particles had become essentially massless at LHC energies for all operators, thus lighter mass points were not required. In order to refine this further and not produce unnecessary MC samples where a simple cross-section rescaling would be applicable, studies were done in order to determine where the particle mass became negligible for each operator. The results of these studies can be seen in Figure H.5 of Appendix H.2. It was found that in some cases, the 10 GeV mass point was not required and 50 GeV was essentially massless, while other operators required an additional MC sample at 10 GeV.

The third and final parameter required for the EFT is the suppression scale, M_* . The model used makes the assumption that the EFT is valid, whether this is realistic or not, and thus the value of M_* is a scale factor which controls the cross-section but does not affect the event kinematics. For this reason, limits on WIMP EFTs are typically set on the value of M_* . A full discussion of where this assumption breaks down, and the suppression scale begins to alter the kinematics, is left to Appendix G.

8.1.2 Simplified models

Another approach considered is the simplified model, where the mediator is not integrated out, and thus it can be within the energy scale of the LHC. This is the natural solution to questions relating to the validity of EFTs, but involves significantly more degrees of freedom, and thus poses challenges of its own. While simplified models will play a significant role in the future of DM searches at the LHC [93, 94], initial mono-jet studies considered only two similar variants, corresponding to pure vector and pure axial vector Z' mediators. These two models were both provided by Reference [95]. For mediator masses beyond the scale of the LHC, these two models would correspond to the D5 and D8 EFT operators respectively.

Given the EFT studies previously mentioned and the identical kinematics of the D5 and D8 operators, it was expected that the LHC should be independent of the chirality of the mediator. Studies have shown that this is mostly true, but breaks down for very light mediators with masses of $\mathcal{O}(100 \text{ GeV})$ or below combined with light DM particles, where the available phase space is significantly restricted. This is demonstrated through the kinematic acceptance of the mono-jet analysis in Figure H.6 of Appendix H.2. At this point, a choice was taken to favour the vector simplified model for the initial set of studies, thus reducing the amount of MC samples required.

The vector Z' simplified model has several parameters of interest, which must all be probed in order to conclusively cover all of the available parameter space. The available parameters are:

1. DM mass, m_χ
 - As for the EFT scenario, the DM mass affects the kinematics, and is thus important to vary.
2. Couplings between the mediator and other particles
 - Couplings to the quarks: $g_u, g_d, g_s, g_c, g_b, g_t$ (generically g_q)
 - Couplings to the DM particle, g_χ
 - Couplings between the mediator and other particles (leptons, bosons, etc) was not included as a part of this model.
3. Mediator mass, M_{med}

- The mediator mass is the most significant addition provided by simplified models which is not included in EFTs. Changing the mediator mass will be described in much more detail below.

4. Mediator width, Γ_{med}

- The mediator width is not fully independent of the couplings, but is semi-decoupled until a specific model is chosen. This will be discussed in more detail below.

The impact of the mediator mass, width, and couplings can be seen directly when considering the mediator component of a simple cross-section as in Equation 8.1, where Q_{tr} is the average momentum transfer of the collision ($\sqrt{\hat{s}}$). All three of the production types are discussed in References [95, 96].

$$\sigma \propto \frac{g_{\text{q}}^2 g_{\chi}^2}{(M_{\text{med}}^2 - Q_{\text{tr}}^2)^2 + M_{\text{med}}^2 \Gamma_{\text{med}}^2} \quad (8.1)$$

For small mediator masses, Equation 8.1 reduces to Equation 8.2, showing how the off-shell production regime primarily depends on the momentum transfer and the couplings. However, a very light mediator is always off-shell for such highly energetic collisions. Thus the mediator mass and width do not significantly change the kinematic distributions. The momentum transfer itself is independent of the model used, and is not tunable. All that remains is the coupling and DM mass. The coupling only shapes the cross-section and not the kinematics, thus replacing the suppression scale M_* as a means of parametrizing the limit strength. The DM mass sets the required lower bound on the Q_{tr} distribution in order to pair-produce DM particles, and thus the DM mass is an interesting parameter to vary for low mediator masses.

$$M_{\text{med}} \ll Q_{\text{tr}} \implies \sigma \propto \frac{g_{\text{q}}^2 g_{\chi}^2}{Q_{\text{tr}}^4} \quad (8.2)$$

The contact interaction regime, with heavy mediators, takes Equation 8.1 to Equation 8.3. Here, the dependence on Q_{tr} has disappeared, and the cross-section falls off with increasing mass. However, the mediator is too heavy, and thus is always off-shell for such low energy collisions. Thus the mediator mass and width do not significantly change the kinematic distributions. Once again, the coupling and DM mass remain, and the same logic applies for treating them as the limit strength and a kinematic dependency respectively.

$$M_{\text{med}} \gg Q_{\text{tr}} \implies \sigma \propto \frac{g_{\text{q}}^2 g_{\chi}^2}{M_{\text{med}}^4} \quad (8.3)$$

The resonant production regime is the remaining case, where Equation 8.1 cannot be further simplified. In this case, all of the parameters come into play. The DM mass still shapes the Q_{tr} distribution by setting the minimum possible energy required to pair-produce DM. The mediator mass is now accessible and becomes the most probable state, thus the Q_{tr} distribution becomes peaked around the resonance. The mediator width indirectly can shape the kinematics by controlling how broad the resonance is, although this is a secondary effect compared to the mass, as can be seen in Equation 8.1. The coupling remains a good means of quantifying the strength of limits without changing kinematic quantities.

The relation between the couplings and width is not clearly defined until a model is chosen. Leptophilic and leptophobic Z' models would lead to significantly different widths for the same couplings

relevant to the mono-jet analysis. Furthermore, the quark couplings are often considered to be equivalent to significantly reduce the number of independent parameters, thus $(g_u = g_d = g_s = g_c = g_b = g_t) = g_q$.

This degree of arbitrariness is treated by varying the widths from $\frac{1}{8\pi}$ to $\frac{1}{3}$ in Reference [95]. The first width is chosen such that only a single SM quark couples to the mediator with a strength of $g_q = 1$, where the second is chosen such that the assumption of the mediator being a narrow resonance does not break down. These remain arbitrary choices, and can lead to questionable behaviour. In particular, while it is reasonably safe to allow extra width to be ignored as a part of the leptonic sector, it is not reasonable to use a width which is smaller than necessary to accommodate for the chosen quark-sector couplings. Reference [97] provides width equations for this scenario, where the only particles which couple to the Z' mediator are the quarks and DM. This was not realized at the time of the MC sample production, and thus widths of $\frac{1}{8\pi}$ and $\frac{1}{3}$ are used.

8.2 Effective field theory results

8.2.1 Limits on the EFT suppression scale

As discussed in Section 8.1.1, the three parameters of interest in the EFT interpretation are the operator, DM mass m_χ , and suppression scale M_* . The operator and DM mass affect the kinematic distributions, and thus a scan over these variables is performed. However, the suppression scale simply modifies the cross-section. As such, the model independent limits on the visible cross-section in Section 7.8.3 can be converted to limits on the suppression scale for a given operator- m_χ combination.

All experimental systematic uncertainties, whether cross-section or acceptance based, are included in the limit setting procedure. The acceptance portion of the experimental systematic is included in the nominal CLs [91] limit derivation, while the cross-section component scales the uncertainty band around the nominal limit. This uncertainty band was found to be consistent with the nominal band within the width of the line on plots, and thus was not included, despite being studied and considered. The two experimental systematics are the beam energy systematic and the detector systematics. The detector systematics (quadratic sum of the lepton, JES, JER, and E_T^{miss} soft term systematics) are treated as a fully correlated component between the signal and background. The beam energy uncertainty is only considered for signal samples, as the data-driven approach used for background estimation cancels the contribution.

As explained in Appendix I.2, only the theoretical uncertainties affecting the acceptance are included in the limit setting procedure. The following acceptance-related systematics are thus included:

- Acceptance portion of the PDF uncertainty (value varies)
- Would normally include the showering systematic, but this was found to be negligible
- Acceptance portion of the factorization/renormalization scale uncertainty (3% for all scenarios)
- Matching scale systematics (3% for $E_T^{\text{miss}} < 350$ GeV and 5% for $E_T^{\text{miss}} > 350$ GeV)

Both 90 and 95% Confidence Level (CL) limits on M_* are calculated, in order to compare with other categories of DM searches following different conventions. This is done for all nine mono-jet signal regions, which are explained in Section 7.3. Different EFT operators are more or less sensitive to different E_T^{miss} regimes, and thus this large number of signal regions enables the analysis to cover the optimal

m_χ [GeV]	D1	D5	D8	D9	D11	C1	C5
10	39 (39)	971 (972)	955 (955)	1491 (1822)	372 (418)	9 (9)	198 (237)
50	39 (39)	971 (972)	955 (955)	1491 (1822)	372 (418)	9 (9)	198 (237)
100	39 (39)	974 (975)	930 (930)	1467 (1792)	372 (420)	9 (8)	193 (230)
200	37 (37)	960 (961)	874 (875)	1379 (1681)	359 (404)	7 (7)	181 (216)
400	32 (32)	862 (863)	724 (725)	1162 (1415)	323 (363)	5 (4)	144 (172)
700	24 (24)	678 (679)	521 (522)	813 (978)	260 (292)	2 (2)	97 (115)
1000	17 (17)	486 (486)	344 (344)	548 (655)	199 (224)	1 (1)	61 (73)
1300	12 (12)	328 (328)	218 (218)	359 (422)	145 (163)	0 (0)	37 (45)

Table 8.4: 90%CL lower limits on the suppression scale M_* in GeV for the EFT operators. Numbers are for the observed (expected) limits.

m_χ [GeV]	D1	D5	D8	D9	D11	C1	C5
10	38 (38)	924 (925)	908 (909)	1442 (1733)	362 (404)	9 (9)	190 (224)
50	38 (38)	924 (925)	908 (909)	1442 (1733)	362 (404)	9 (9)	190 (224)
100	37 (37)	927 (928)	885 (885)	1417 (1703)	363 (405)	8 (8)	185 (218)
200	36 (36)	914 (915)	832 (833)	1332 (1598)	350 (390)	7 (7)	174 (205)
400	31 (31)	821 (822)	690 (690)	1121 (1343)	315 (351)	4 (4)	139 (163)
700	23 (23)	647 (647)	497 (497)	788 (927)	253 (282)	2 (2)	93 (109)
1000	17 (17)	462 (462)	327 (328)	533 (620)	194 (216)	1 (1)	59 (69)
1300	11 (11)	311 (311)	207 (207)	336 (398)	141 (158)	0 (0)	36 (42)

Table 8.5: 95%CL lower limits on the suppression scale M_* in GeV for the EFT operators. Numbers are for the observed (expected) limits.

regime for each operator. The specific signal region to use for a given model is decided by comparing the expected limits for each region, and then selecting the region with the strongest expected limit, as shown in Figure 8.3. Based on the expected limits, the following signal regions were found to be optimal and thus were used for the final limits:

- C1: signal region 4 ($E_T^{\text{miss}} > 350$ GeV)
- D1, D5, D8: signal region 7 ($E_T^{\text{miss}} > 500$ GeV)
- D9, D11, C5: signal region 9 ($E_T^{\text{miss}} > 700$ GeV)

These results match the trends shown in the E_T^{miss} and leading jet p_T distributions in Figure 8.2. The most soft distributions belong to C1, which is optimal for the lowest E_T^{miss} cut. The hardest distributions are the D11, D9, and C5 operators, which prefer the highest E_T^{miss} cut. The remaining operators of D1, D5, and D8 (a scaling of D5) are found in between these other two classes of operators, and prefer a medium E_T^{miss} cut. The resulting limits for 90 and 95%CL are shown in Tables 8.4 and 8.5 respectively.

8.2.2 Model validity

By using EFTs, implicit assumptions about the properties of the mediator which connects the SM to the dark sector are made. When working purely with colliders, it can be argued that the models are simply benchmarks, and thus it is not as important if the model is valid for the full parameter space. However, when comparing to other experiments, the validity becomes crucial. Other types of experiments work in different energy regimes, and thus they do not necessarily rely on the same assumptions, thus any

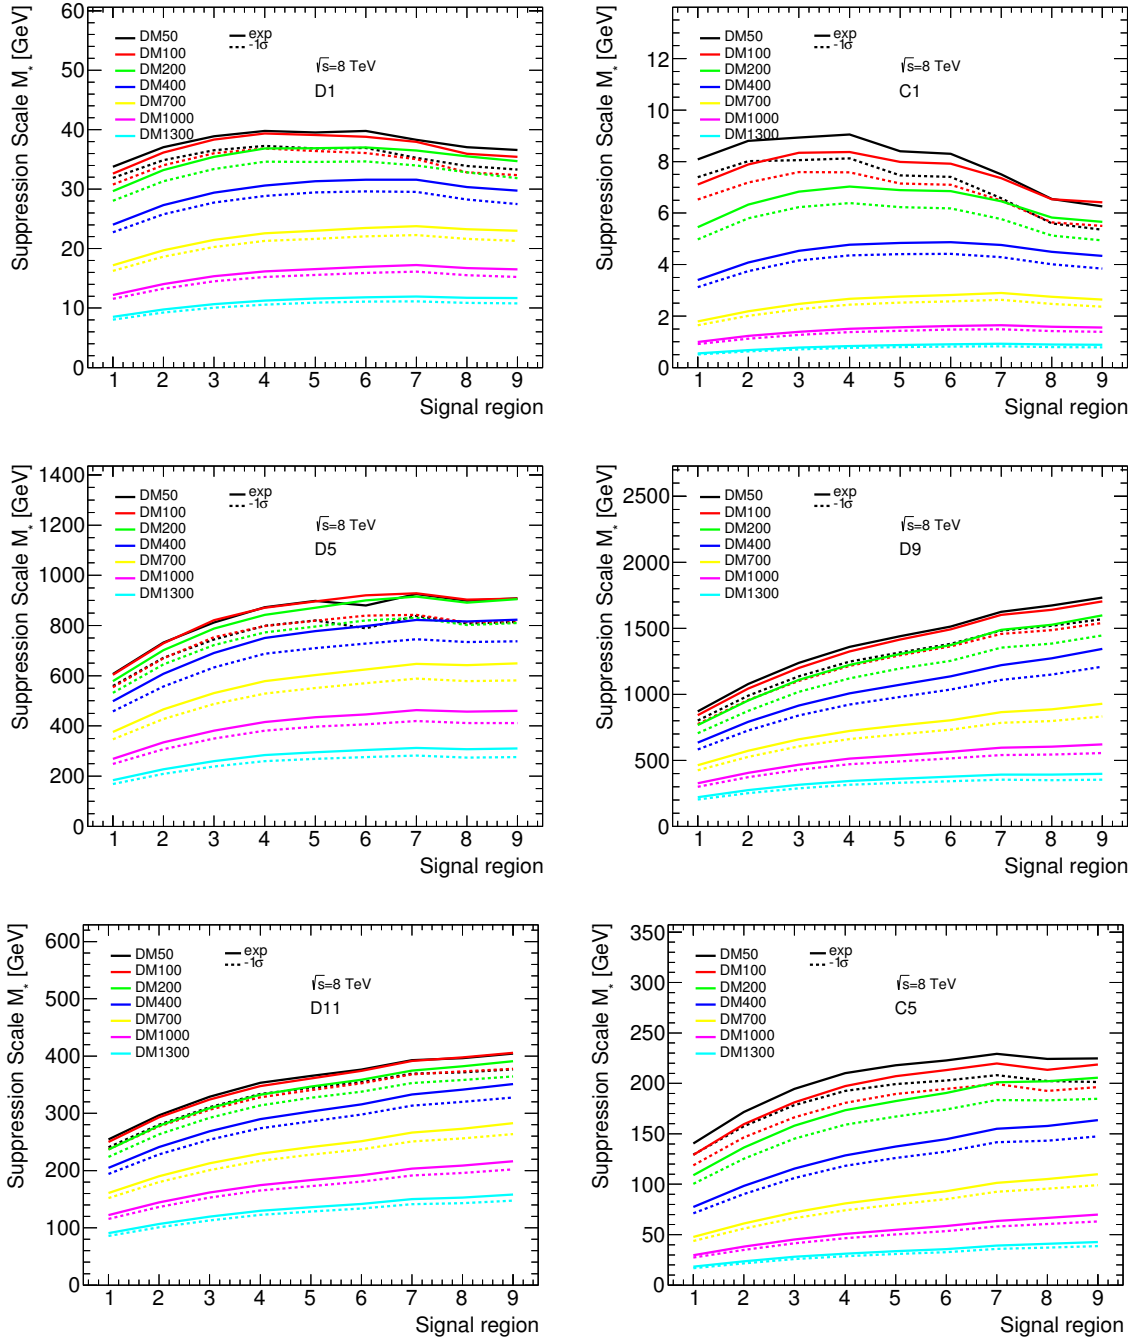


Figure 8.3: Lower limits on the suppression scale M_* for the EFT operators, shown as a function of the mono-jet signal region, for 95%CL expected limits.

comparison ignoring this may be biased. In order to provide a fair comparison, the validity of the DM model at the LHC must be investigated.

The way in which this can be done is very technical, and is fully detailed in Appendix G. In short, a validity criterion is defined, and the fraction of events which fail this criterion is used to truncate (rescale) the resulting limits to remove the contaminated parameter space. The full treatment for each of the operators considered in this particular analysis is presented in Appendix G.4.1 for both 90 and 95% confidence level observed limits. Both of these sets of limits are relevant to future comparisons with other types of experiments, hence the necessity of covering both confidence levels. A quick summary of the truncated limits at the 95%CL is included in Figure 8.4. The values contained in these plots are later used to add truncated limit lines to other figures in the following Sections.

8.2.3 Collider limits

It is important to convey to the external community that we have studied and understood the effects that arise from problems pertaining to the EFT validity. This is critical in order to convince the theory community that we have properly considered the limitations of the models, and to convey to the DD community that our results are comparable to theirs within a given region of parameter space. At the same time, it must be made clear to the theory community that the common assumption of $g_q = g_\chi = 1$ is itself a major assumption, and a proper treatment necessitates a full scan. It is one option which makes calculations trivial, but it is no more justified than any other guess. Emphasis on what areas of coupling phase space are valid, and which are not, leads to a much more robust understanding of the validity associated to collider EFT limits.

In order to do this, the M_* limit plots will focus on the limits we can set ignoring the EFT validity. These are useful benchmark models, and these plots display our sensitivity to such a benchmark. Extra lines are then added in order to show the impact of validity considerations, where the presence of a line means that the operator is valid at that value of m_χ for that particular coupling assumption, while the absence of a line means it is not valid. The plots in Figure 8.5 show this information for both fermionic and scalar DM operators. Note that the relevance of annihilation channels varies by operator, and thus the turn-on of the $t\bar{t}$ process is only relevant for the scales considered for the D8 operator. This is the cause of the bump in the relic density in the associated plot for D8.

8.2.4 Comparisons to other experiments

Given the elusive nature of DM, it is commonly stated that DM will not be considered to be directly measured by multiple independent experiments. Other new phenomena may appear in DM-like signatures, as for example colliders can't confirm an observed invisible particle is stable beyond the volume of the detector and ID experiments could discover new pulsars or other such sources. It is thus particularly important for at least two types of the three search categories (DD, ID, and collider production) to observe consistent results.

For this reason, results from collider experiments are mapped onto the WIMP-nucleon scattering cross-section plane and the WIMP-WIMP annihilation cross-section plane to compare to DD and ID experiments respectively. These translations make use of the EFT's simple form, which can be rotated from collider production ($SM + SM \rightarrow DM + DM$) to WIMP-nucleon scattering ($SM + DM \rightarrow SM + DM$) to WIMP-WIMP annihilation ($DM + DM \rightarrow SM + SM$). This does oversimplify the picture, as there is no

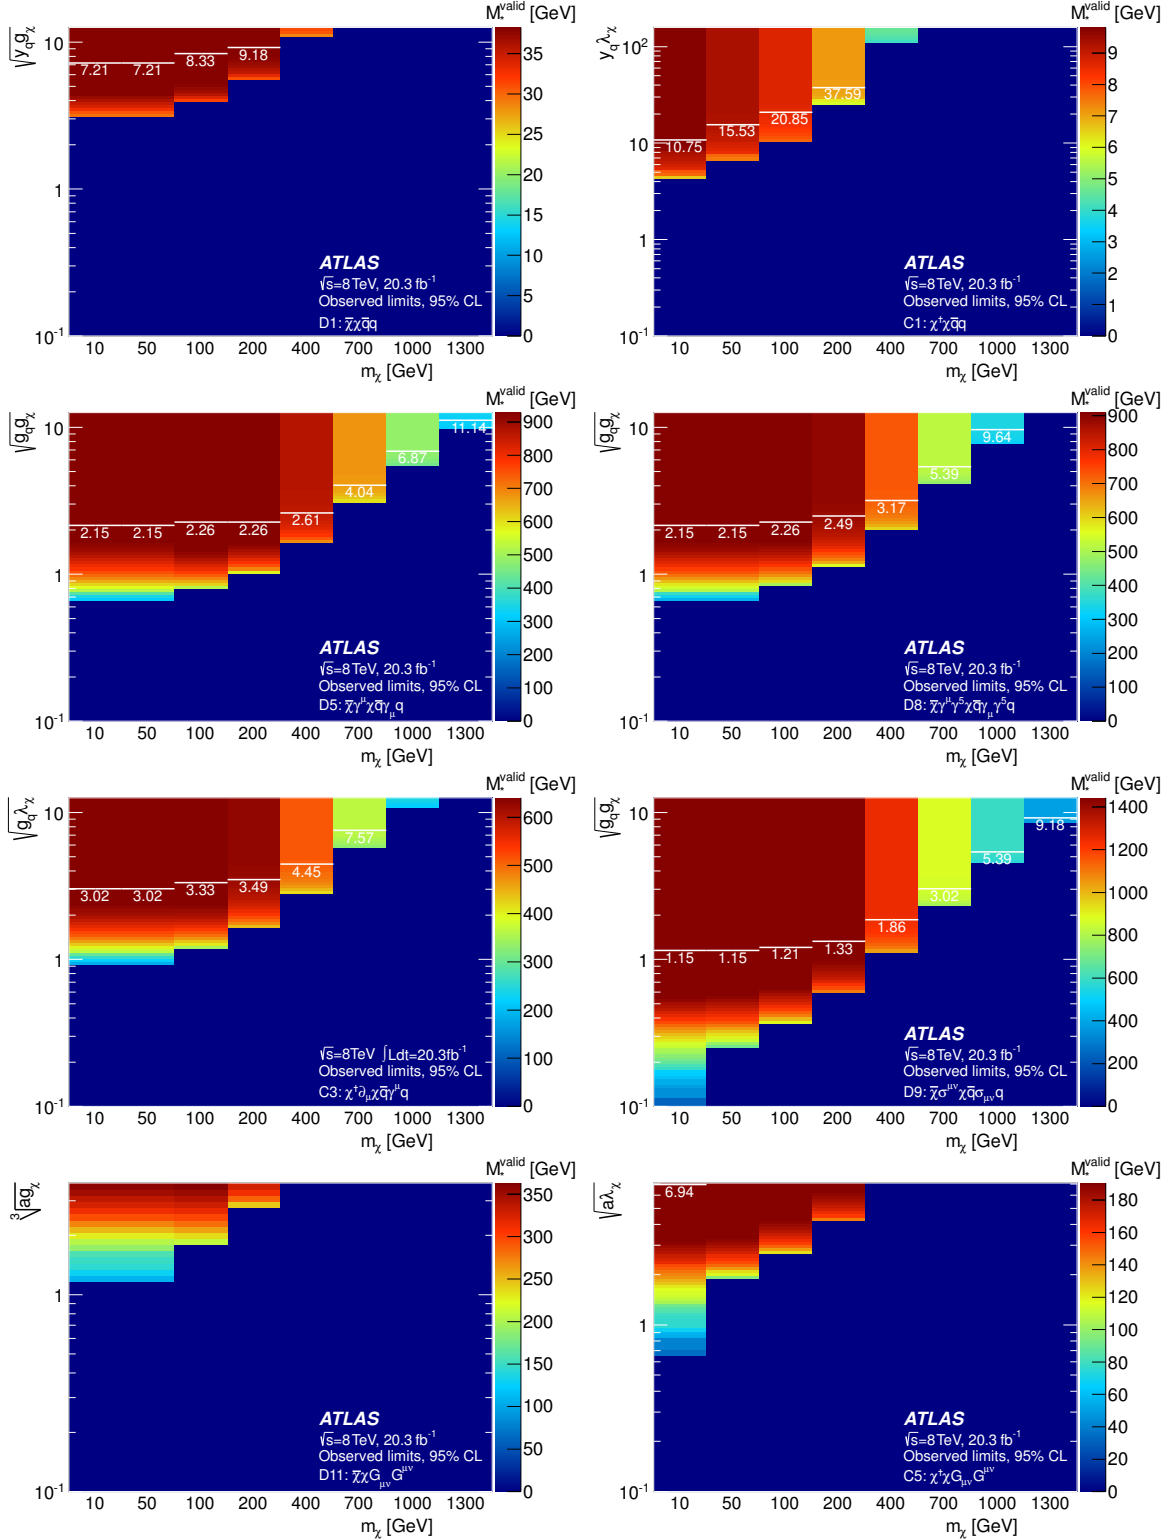


Figure 8.4: The rescaled/truncated limits for the D1 (top left), C1 (top right), D5 (upper middle left), D8 (upper middle right), C3 (lower middle left), D9 (lower middle right), D11 (bottom left), and C5 (bottom right) operators at 95%CL observed limits as a function of the WIMP mass (x-axis) and couplings (y-axis). White lines and associated white numbers correspond to the minimum coupling value for which the truncated limit is within 1% of the nominal limit, or equivalently the scale factor is at least 0.99. Absent lines and numbers (particularly for D11) means that the given operator-mass combination has no valid coupling where this condition is met. A complete breakdown of each operator including the validity fractions, scale factors, and truncated limits at both 90%CL and 95%CL is provided in Appendix G.4.1.

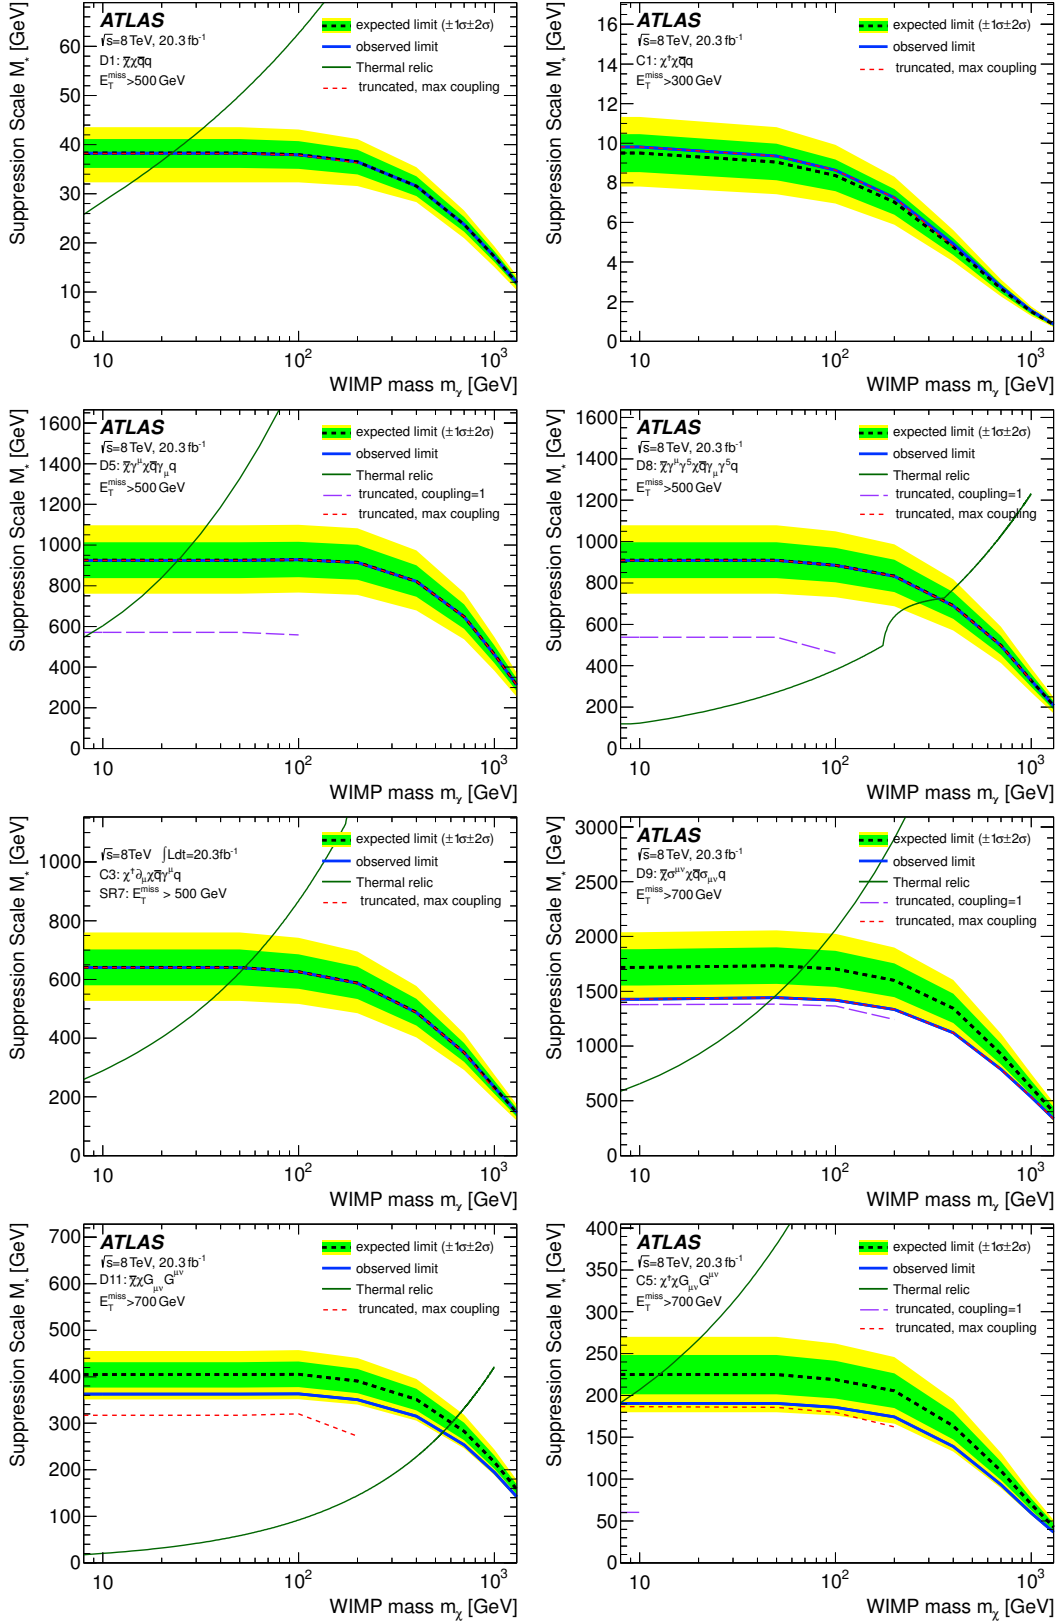


Figure 8.5: 95% CL lower limits on the suppression scale for each EFT operator. Expected (observed) limits are shown as a dashed black (solid blue) line. Experimental and theoretical acceptance 1σ (2σ) uncertainties on the expected limit are shown in green (yellow). The remaining uncertainty on the observed limit is within the line width, and thus is not shown. The green line is the value of (M_*, m_χ) resulting in the correct relic density as measured by WMAP [98], assuming a single DM type with only one SM coupling. Purple long-dashed (red short-dashed) lines show truncated limits for a coupling of 1 (max perturbative). Absent truncation lines mean the operator is not valid for the given coupling.

m_χ [GeV]	D1	D5	D8	D9	D11	C1	C5
10	2.15×10^{-39}	9.28×10^{-40}	3.38×10^{-41}	5.69×10^{-42}	1.07×10^{-44}	2.88×10^{-36}	4.61×10^{-41}
50	2.48×10^{-39}	1.07×10^{-39}	3.90×10^{-41}	6.56×10^{-42}	1.23×10^{-44}	1.33×10^{-37}	2.12×10^{-42}
100	2.52×10^{-39}	1.08×10^{-39}	4.41×10^{-41}	7.13×10^{-42}	1.25×10^{-44}	3.38×10^{-38}	5.99×10^{-43}
200	3.49×10^{-39}	1.15×10^{-39}	5.71×10^{-41}	9.21×10^{-42}	1.57×10^{-44}	2.33×10^{-38}	1.96×10^{-43}
400	8.39×10^{-39}	1.78×10^{-39}	1.22×10^{-40}	1.84×10^{-41}	2.97×10^{-44}	2.25×10^{-38}	1.23×10^{-43}
700	4.72×10^{-38}	4.67×10^{-39}	4.55×10^{-40}	7.68×10^{-41}	1.09×10^{-43}	2.88×10^{-37}	1.95×10^{-43}
1000	3.74×10^{-37}	1.77×10^{-38}	2.40×10^{-39}	3.72×10^{-40}	5.44×10^{-43}	2.26×10^{-36}	6.11×10^{-43}
1300	3.03×10^{-36}	8.52×10^{-38}	1.49×10^{-38}	2.02×10^{-39}	3.64×10^{-42}	0	2.67×10^{-42}

Table 8.6: Upper 90% confidence level observed limits on the WIMP-nucleon scattering cross-section in units of cm^2 . Conversions to the scattering cross-section from M_* follow Reference [92].

guarantee that the same SM and/or DM particle(s) is/are available for each of these processes. However, it provides a useful comparison, where the additional subtleties are left to the theory community when performing combinations of independent measurements.

These different types of experiments have different conventions when setting exclusion limits. DD experiments tend to use 90% confidence level exclusion bounds, while colliders and ID experiment work with 95% confidence levels. In order to compare with all cases, both 90 and 95%CL bounds on the suppression scale M_* are calculated as previously listed in Tables 8.4 and 8.5 respectively. These are then converted to 90%CL bounds on WIMP-nucleon scattering cross-sections and 95%CL bounds on WIMP-WIMP annihilation cross-sections as per Reference [92]. The resulting WIMP-nucleon scattering cross-sections are listed in Table 8.6.

WIMP-nucleon scattering cross-section comparisons

The DD DM community has been very vocal in their condemnation of the use of EFTs when comparing to their results. They have rightfully pointed out that the validity is not guaranteed, and that it must be properly considered when comparing to their results, which do not have the same limitations. Now that a thorough study has been conducted, it is important to demonstrate the degree to which limits can be compared, and the regions in which the limits are applicable. In particular, it is important to realize that $\sigma_{\text{valid}} \propto R_{\text{M}_{\text{med}}}^{\text{tot}}$, meaning that a 20% validity fraction means that the limit is multiplied by a factor of 0.2. This is a large change, but does not change the general picture when considering the typical plots spanning many orders of magnitude.

In order to convey all of this information, and to provide fair comparisons, only operators which are valid for some point in the relevant parameter space will be shown. Furthermore, points will only be plotted where they are valid, thus emphasizing the dependence on the DM mass. Two lines will be shown per operator, one for the “natural scale” of a coupling of 1, and the other for the maximum possible coupling. The interpretation of exactly where a model is trusted and where it is not is then left to the reader, as the choice of coupling remains arbitrary until the interaction strength is observed and quantified.

The following comparisons to DD experiments are considered:

- Spin-Independent:
 - Collider operators: C1, C5, D1, D5, D11

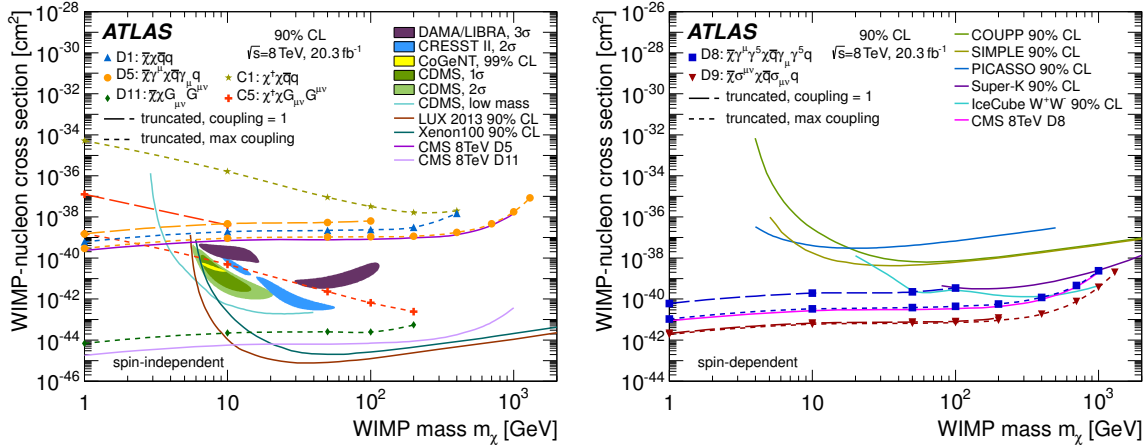


Figure 8.6: ATLAS mono-jet 90% confidence level limits on the WIMP-nucleon scattering cross-section in comparison to the current relevant Spin-Independent (left) and Spin-Dependent (right) Direct Detection experiments. ATLAS lines have the EFT validity taken into account, and truncated limits for couplings of 1 (the natural scale, long dashes) and the maximum before perturbativity breaks down (4π for quark operators, $\sqrt[3]{16\pi}$ for gluon operators, short dashes) are shown. The C1 and D1 operators are not shown on the left plot because they are invalid for all points in the displayed parameter space. CMS mono-jet limits are shown when available, although these results are not truncated, and thus a direct comparison is not possible.

- DD experiments: CDMS-II [99], CRESST-II [100], CoGeNT [101], CoGeNT low mass [102], DAMA/LIBRA [4, 103], LUX [104], Xenon100 [105]
- Spin-Dependent:
 - Collider operators: D8, D9
 - DD experiments: COUPP [106], IceCube [107], PICASSO[108], SIMPLE [109], Super-K [110]

The following comparisons to other collider experiments in the WIMP-nucleon scattering cross-section plane are considered:

- CMS mono-jet [78]
- ATLAS mono- $W/Z(qq)$ [111]
- ATLAS mono- $W(\ell\nu)$ [112]
- ATLAS mono- $Z(\ell\ell)$ [113]
- ATLAS mono-photon [9]
- ATLAS Heavy Flavour + DM [7]

Comparisons of these Spin-Independent (SI) and Spin-Dependent (SD) DD experiments to the ATLAS mono-jet truncated results are shown in Figure 8.6.

DD constraints generally fall significantly at low DM masses. This is due to the nuclear recoil energy of DD experiments falling to levels below a few keV, and thus pushing the limits for detection of current technologies. On the other hand, colliders are the most sensitive to low DM masses because the DM

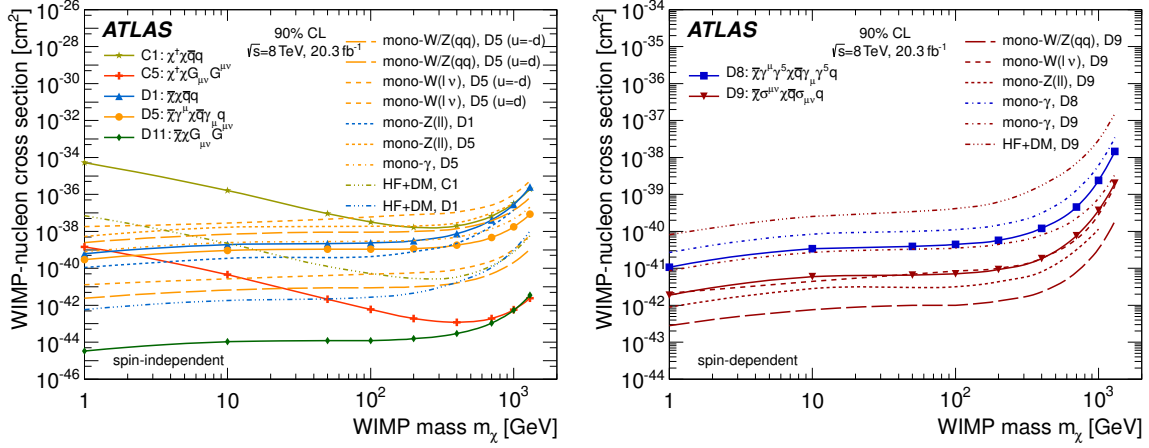


Figure 8.7: ATLAS mono-jet 90% confidence level limits on the WIMP-nucleon scattering cross-section in comparison to other ATLAS mono-X Spin-Independent (left) and Spin-Dependent (right) Direct Detection searches for DM. In order to provide a fair comparison across the many topologies, the EFT validity is not considered in this figure.

becomes effectively massless. This interplay leads to colliders providing the strongest limits in this region of parameter space for both SI and SD experiments.

SI DD experiments are in general less sensitive than SD experiments due to the coupling to the nucleon spin rather than nucleon mass. This suppression is important at the nucleon level, but the equivalent notion of chirality at colliders has a much smaller impact. In particular, primary kinematic distributions considered are independent of the chirality as shown in Section 8.1, and the change only results in minor cross-section factors. As a result, collider limits are generally stronger than SI DD limits over the full parameter space considered.

Comparisons of the numerous independent ATLAS mono-X searches is a useful exercise, and clearly demonstrates the strengths of each topology, as shown in Figure 8.7. This comparison does not consider the validity of the EFT, as most of the analyses did not include such criteria. In general, the quark mass dependence allows the Heavy Flavour + DM analysis to set the strongest limits on the D1 and C1 operators, and the possibility of constructive interference gives mono- $W/Z(qq)$ the strongest limits on D5 (mono-jet dominates if this is not true). Mono-jet sets the strongest limits for D8, and is the only analysis capable of probing the gluon operators, D11 and C5. D9 is more complicated as the mono- $W/Z(qq)$, mono- $W(\ell\nu)$, mono- $Z(\ell\ell)$ and mono-jet analyses all have similar sensitivity. However, some topologies observe more SM background events than expected, while others observe fewer, thus the current iteration shows the mono- $W/Z(qq)$ topology to be dominant for the D9 tensor operator.

Figure 8.8 compares ATLAS to CMS mono-jet results, without validity considerations, thus treating both sets of results equitably. ATLAS and CMS results are observed to generally be comparable, with CMS observed limits always slightly stronger. This is due to the small excess of events observed by ATLAS compared to the small deficit observed by CMS. When comparing the expected limits, ATLAS limits are generally stronger due to the lack of a veto on events with more than two jets.




Figure 8.8: ATLAS mono-jet 90% confidence level limits on the WIMP-nucleon scattering cross-section in comparison to CMS mono-jet results [78]. In order to provide a fair comparison between collider experiments, the EFT validity is not considered in this figure.

WIMP-WIMP annihilation cross-section comparisons

ID results are provided in the form of limits on the thermally averaged WIMP-WIMP annihilation cross-section. Comparisons of the ATLAS mono-jet D5 and D8 pure vector and pure axial vector operators are made to Fermi-LAT[114] and HESS[115] results. The cross-section which leads to the correct relic abundance, along with the comparisons, is shown in Figure 8.9. Collider limits are particularly powerful for low DM masses, and reach well below the necessary relic density cross-section for m_χ below approximately 20 GeV (10 GeV) and 100 GeV (45 GeV) for D5 and D8 respectively for assumed couplings of 4π (1). It is important to keep in mind that these comparisons make the assumption that there is only one type of interaction (one possible EFT operator) which connects the SM to the dark sector. Possible combinations of interactions, which would lead to different results, are not considered.

8.3 Simplified model results

Explicit inclusion of a mediator coupling the SM to the dark sector allows for much more detailed studies into the structure of DM production at the LHC at the cost of many additional degrees of freedom. As discussed in Section 8.1.2, the parameters of interest for the Z' simplified model used are the mediator mass M_{med} , mediator width Γ_{med} , DM mass m_χ , and couplings. The primary contributors to the shapes of kinematic distributions are the mediator and DM mass, while a focus will be applied to a single width. The model independent limits on the visible cross-section in Section 7.8.3 can thus be converted into limits on the geometric mean of the mediator's couplings to the SM and dark sector, $\sqrt{g_q g_\chi}$, for a given $M_{\text{med}}-m_\chi$ combination.

While this approach does naturally address the question of EFT validity, it does not encompass the full set of models which reduce to a given EFT, and it comes with its own assumptions. In particular, both the coupling of the SM quark to the mediator g_q and the coupling of the mediator to the dark sector g_χ must exist and be perturbative. This provides a natural validity constraint of $0 < g_q, g_\chi < 4\pi$, or $0 < \sqrt{g_q g_\chi} < 4\pi$. Without knowledge of the full UV-complete theory which encapsulates this interaction, this coupling range cannot be further restricted.




Figure 8.9: ATLAS mono-jet 95% confidence level limits on the thermally averaged WIMP-WIMP annihilation cross-section in comparison to Fermi-LAT and HESS results. ATLAS lines have the EFT validity taken into account, and truncated limits for couplings of 1 (the natural scale, long dashes) and the maximum before perturbativity breaks down (4π , short dashes) are shown.

This theoretical limitation has a profound impact on collider searches for DM. The lower bound of zero implies that such searches can never rule out a particular point as contributing the amount of DM in the universe. However, at some point the contribution from a given type of WIMP to the total DM content can be determined to be negligible in comparison to the observed relic density in the universe. The relic density should not be used as a perfect solution, as any comparison to a particular model requires the addition of assumptions.

The relic density provides constraint on what coupling $\sqrt{g_q g_\chi}$ would be required for a single dark matter particle of mass m_χ and a single mediator of mass M_{med} . If either of these constraints is relaxed, such that the dark sector is comprised of multiple particles or there are multiple annihilation channels, then the derived coupling bound to fulfil the relic density can be moved in either direction. In particular:

- If there is only one type of DM particle and only one annihilation channel, then the relic density should be precisely matched, and the assumptions made are valid.
- If the dark sector is comprised of multiple particles and only a subset couple to the standard model, then the relic density is a lower limit on the coupling. Larger coupling values would imply additional annihilation, leaving space for additional dark particles. Smaller coupling values would result in too little annihilation, implying too much dark matter, and thus are not permitted by relic density constraints.
- If the dark sector is comprised of a single particle which can couple to the standard model through multiple interactions, then the relic density is an upper limit on the coupling. Smaller coupling values from multiple different interactions could be combined to give the expected relic density, while larger couplings would lead to a deficit of dark matter compared to the observed relic density.
- If the dark sector is comprised of multiple particles which can couple to the standard model through multiple interactions, then it is a mix of the above two effects, and no strong statement can be made about upper vs lower limits. In general, the relic density remains a good indication of the




Figure 8.10: Limits on the coupling between quarks and a Z' -like particle in the di-jet topology. [42]

approximate area of greatest interest, but regions below and above the line cannot be immediately excluded.

The use of simplified models also provides access to the resonance structure, which is very beneficial at colliders. In addition to increasing the cross section, it also allows for new channels to contribute. If the mediator can be produced by standard model particle(s), then it can also decay back to them. Di-jet resonance searches are one type of topology which can take advantage of this new feature. Given the enormous di-jet cross-section, the immediate next question is whether this will completely overshadow any mono-jet result. However, the aforementioned couplings once again come into play. If the coupling between the standard model and the mediator is large while the coupling of the dark sector to the mediator is weak, then the di-jet search will provide stronger bounds. On the other hand, a weak standard model coupling to the mediator and a strong dark sector coupling would be very hard to discover as a di-jet resonance. In particular, one can consider the scenario of $g_q = \frac{1}{4\pi}$ and $g_\chi = 4\pi$. Such a scenario would be nearly invisible to di-jet searches, and easily evades the current best di-jet limits, while providing a reasonably strong signature in mono-jet topologies. This is demonstrated in Figure 8.10, where the current strongest di-jet limits are at the level of $g_q \gtrsim 0.5$ or weaker for all mediator masses considered.

This is one of many examples, and demonstrates the complementary nature of simplified models in collider searches. It is important to note that this is still a simplified model and not a full theory. It may be that the full theory has additional complications which either reduce or increase the impact of particular signatures, and thus it is still important and interesting to conduct searches in similar phase space regions for multiple possible topologies. Thus interpreting the mono-jet results in the framework of a Z' simplified model are complementary to other search strategies, and covers regions of parameter space which are otherwise untouched.

A full list of limits set on the Z' simplified model is provided in Table 8.7. Graphical representations and analysis of these results will be provided in the following sections.

8.3.1 Production regimes

One of the uses of simplified models with respect to EFTs is the ability to study how the mediator mass effects the DM production and resulting limits, as done in References [95]Khachatryan:2014rra.

M_{med}	Γ_{med}	m_χ [GeV]	$\sqrt{g_q g_\chi}$	M_* [GeV]
10 GeV	$M_{\text{med}}/3$	10	0.52 (0.52)	19 (19)
50 GeV	$M_{\text{med}}/3$	10	0.36 (0.36)	140 (140)
		50	0.5 (0.5)	99 (99)
		400	1.39 (1.39)	35 (35)
	$M_{\text{med}}/8\pi$	50	0.55 (0.55)	91 (91)
		400	1.38 (1.38)	36 (36)
100 GeV	$M_{\text{med}}/3$	10	0.36 (0.36)	278 (278)
		50	0.42 (0.42)	240 (240)
		400	1.39 (1.39)	71 (71)
	$M_{\text{med}}/8\pi$	50	0.46 (0.46)	216 (216)
		400	1.36 (1.36)	73 (73)
300 GeV	$M_{\text{med}}/3$	10	0.42 (0.42)	709 (709)
		50	0.43 (0.43)	700 (701)
		200	0.68 (0.68)	442 (442)
		400	1.32 (1.32)	227 (227)
	$M_{\text{med}}/8\pi$	50	0.24 (0.24)	1255 (1256)
		400	1.32 (1.32)	226 (226)
600 GeV	$M_{\text{med}}/3$	10	0.55 (0.55)	1090 (1091)
		50	0.58 (0.58)	1038 (1038)
		200	0.57 (0.57)	1047 (1047)
		400	1.13 (1.13)	531 (531)
	$M_{\text{med}}/8\pi$	50	0.31 (0.31)	1930 (1931)
		400	1.1 (1.1)	547 (547)
1 TeV	$M_{\text{med}}/3$	10	0.76 (0.76)	1309 (1310)
		50	0.78 (0.78)	1285 (1286)
		200	0.77 (0.77)	1294 (1295)
		400	0.88 (0.88)	1134 (1135)
		1000	4.2 (4.2)	238 (238)
	$M_{\text{med}}/8\pi$	50	0.44 (0.44)	2275 (2277)
		400	0.47 (0.47)	2141 (2142)
3 TeV	$M_{\text{med}}/3$	10	2.68 (2.68)	1120 (1120)
		50	2.71 (2.71)	1106 (1106)
		200	2.69 (2.69)	1115 (1116)
		400	3.1 (3.1)	967 (967)
		1000	3.99 (3.98)	752 (753)
		1300	5.3 (5.3)	565 (565)
	$M_{\text{med}}/8\pi$	50	2.06 (2.06)	1458 (1459)
		400	2.14 (2.14)	1403 (1404)
6 TeV	$M_{\text{med}}/3$	50	5.93 (5.93)	1011 (1011)
		400	6.63 (6.63)	904 (905)
	$M_{\text{med}}/8\pi$	50	5.87 (5.86)	1022 (1023)
		400	6.57 (6.57)	912 (913)
10 TeV	$M_{\text{med}}/3$	50	10.22 (10.21)	978 (979)
		400	11.3 (11.29)	885 (885)
	$M_{\text{med}}/8\pi$	50	9.97 (9.96)	1003 (1003)
		400	11.31 (11.3)	884 (884)
30 TeV	$M_{\text{med}}/3$	50	30.34 (30.33)	988 (989)
		400	34.11 (34.1)	879 (879)
	$M_{\text{med}}/8\pi$	50	30.27 (30.26)	991 (991)
		400	33.56 (33.55)	893 (894)

Table 8.7: 95%CL limits on the coupling $\sqrt{g_q g_\chi}$ (upper limits) and interaction scale $M_* = M_{\text{med}}/\sqrt{g_q g_\chi}$ (lower limits) for the Z' simplified model.




Figure 8.11: 95%CL limits on the Z' simplified model, shown as a function of the mediator mass, for both the interaction scale $M_* = M_{\text{med}}/\sqrt{g_q g_\chi}$ (left) and coupling strength $\sqrt{g_q g_\chi}$ (right). In both cases, the mediator width Γ_{med} is shown to scale the degree of resonant enhancement, but to be a negligible effect for very light or heavy mediators. The resonant production regime (intermediate M_{med}) comes with the strongest limits, followed by the contact interaction regime (high M_{med}), and then the off-shell production regime (low M_{med}). Above the 4π contour (left) or 4π horizontal line (right), the theory breaks down, as the limit on the couplings is set in the non-perturbative regime. The Z' interpretation of the mono-jet search has no sensitivity to these particular points in parameter space at this time.

This allows for direct studies of the impact of off-shell production, resonant production, and contact interaction production on the sensitivity of the analysis to a particular DM model. Figure 8.11 clearly shows these three regions. In general, it reveals that the strongest limits on the interaction scale come from the resonant production regime, medium strength limits come from the contact interaction regime, and the weakest limits come from the off-shell production regime. Note that the interaction scale only holds a physical meaning for the contact interaction regime, although this comparison is useful to show where the EFT limits are relevant, conservative, or optimistic. It also shows how the mediator width serves to enhance or reduce the resonant contribution in the resonance regime, but becomes a negligible contribution in the off-shell or contact interaction regimes.

An alternative view of the same information, given the relation between the scale and the couplings $M_* = M_{\text{med}}/\sqrt{g_q g_\chi}$, is to consider the limits on the couplings. Figure 8.11 also covers this representation, and makes it much more clear where the theory breaks down due to the couplings passing the perturbativity limit of 4π .

8.3.2 Parameter space scan

In addition to understanding the implication of the mediator mass on the limit strength, it is important to understand the interplay between the DM and mediator masses. A full scan over the m_χ - M_{med} parameter space for limits on the coupling was conducted for a fixed width of $\Gamma_{\text{med}} = M_{\text{med}}/3$. This was inspired by Reference [113], which followed a similar procedure for the t -channel exchange of a scalar mediator. However, s -channel exchange of a Z' is very different. In particular, the t -channel method for a mediator η through $q \rightarrow \eta + \chi$ can be turned around to give DM decay of $\chi \rightarrow \eta + q$. This restricts the DM η mediator mass to be above the DM mass so that DM is stable. On the other hand, for the Z' s -channel exchange, such a decay is not possible. In this scenario, the mediator production follows




Figure 8.12: 95%CL coupling limits for the Z' simplified model, as a function of the DM and mediator masses. The black line corresponds to the point where the relic density is naturally satisfied by the m_χ - M_{med} combination, so long as there is only one type of DM particle and only one type of coupling between that particle and the SM.

$Z' \rightarrow \chi\bar{\chi}$. Turning this around for a decay gives $\chi \rightarrow Z' + \bar{\chi}$, which is stable for all mediator masses so long as the DM particle and its antiparticle have the same mass. Thus the Z' interpretation requires a full grid of m_χ - M_{med} rather than just half of the grid.

Figure 8.12 shows that the strongest limits on $\sqrt{g_q g_\chi}$ come from low DM masses and low mediator masses. As either of these is increased, the required momentum transfer is increased. If this continues, the required momentum transfer ends up at the edge or beyond the LHC energy scale, at which point the limits drop dramatically. A line was added to show the points for which the m_χ - M_{med} combination results in the observed relic abundance, given assumptions. As previously mentioned, this should not be taken as an absolute requirement, but rather a guide as to which areas are of the most interest.

The coupling limits provided on the Z' simplified are the maximum values that have not been excluded by the mono-jet analysis. Any value of the coupling larger than the value shown has been ruled out for this particular model. If the assumption is made that the dark sector is comprised of multiple DM particles, but only one interaction with the SM, then the relic density constraint is still satisfied if the coupling between the dark sector and the mediator grows. However, this region has already been ruled out by the collider limits. As such, the region above the relic density line (low m_χ and high M_{med}) has been ruled out under the assumption that there is only one coupling between the SM and the dark sector, regardless of the number of independent states the dark sector is made of.

Weakening the other assumption is less precise. A single DM particle with multiple interactions to the SM results in weaker allowed couplings while still satisfying the relic density, but lower couplings are already beyond the reach of the mono-jet analysis. At this time, nothing more can be said about this scenario. Relaxing both assumptions, where the dark sector can contain multiple particles with multiple interactions, removes the ability to rule out any region of parameter space based on a combination of the relic density and the mono-jet sensitivity.

8.4 Conclusions

The full set of EFT operators provided by Reference [92] has been studied and broken down into 6 representative operators which cover the kinematic distributions of all 20 EFT operators. The DM masses considered cover the range of $\mathcal{O}(\text{GeV})$ to $\mathcal{O}(\text{TeV})$. Systematic uncertainties relating to both the theoretical and experimental uncertainties are derived, and limits are set on the EFT suppression scale M_* . For the first time, the EFT validity is thoroughly investigated, with a representative UV completion chosen for every operator to show the dependence of the validity on the relevant coupling terms. The fraction of valid events and truncated limits in which the degree of validity is incorporated are provided. These truncated limits are then translated into WIMP-nucleon scattering cross-sections and WIMP-WIMP annihilation cross-sections for comparison to relevant DD and ID experiments. ATLAS mono-jet results are typically stronger than SI DD experiments at the low DM mass range ($m_\chi \lesssim 10 \text{ GeV}$), stronger than SD DD experiments over the full mass range, and stronger than ID constraints over the low to medium DM mass range ($m_\chi \lesssim 100 \text{ GeV}$).

The pure vector Z' simplified model from Reference [95] is also considered, with a full systematic uncertainty treatment as done for the EFT samples. Limits are set on the interaction scale M_* for comparison to the EFT, showing the off-shell production, resonant production, and contact interaction regimes. This clearly delineates the range over which the D5 operator is either correct or conservative, which only breaks down for low mediator masses $M_{\text{med}} \lesssim 1 \text{ TeV}$. Limits are also set on the geometrical average of the mediator couplings $\sqrt{g_q g_\chi}$, which provides an alternative picture of the same information in a form more useful for excluding particular models with given coupling assumptions. A full coupling limit scan over the DM mass and mediator mass parameter space is then performed, providing a detailed summary of which regions remain for a given model. The relic density is also shown on this plot as a guide of what has been excluded by the analysis, although different assumptions on the structure of the dark sector can change the overlap region from excluded to still untouched.

Chapter 9

Mono-jet Prospects at an Upgraded LHC

9.1 Introduction

This chapter is based on Reference [8].

The mono-jet topology is arguably the most generically sensitive mono-X search for DM at the LHC, and thus a natural question is how changing the centre of mass energy \sqrt{s} will affect the analysis. In particular, with the LHC due to begin collisions at a new energy of 13 TeV in mid 2015, a pertinent question is how much data is needed to surpass the current sensitivity. If the analysis is significantly more sensitive with only a small amount of data, then it is an excellent candidate for the ATLAS high-priority searches which will form the vanguard of early Run-II studies. In order to answer these questions, a simple mono-jet sensitivity study was conducted.

The nature of such preliminary studies requires a series of assumptions based on reasonable projections, most notably the LHC collision energy. At the time of the studies, it was believed that the centre of mass energy would be 14 TeV, which is now known to be beyond the initial capacity of the dipole magnets. While 14 TeV may be achieved in future years, the first dataset will be taken at 13 TeV. Other assumptions, such as pileup conditions and expected luminosity have also been updated, although these will remain unknown until data is actually taken. As such, these studies are intended to provide a useful means of understanding the general performance of the mono-jet analysis under conditions similar to Run-II.

When these studies were conducted, the 8 TeV mono-jet result was still in flux, and the analysis selection had not been frozen. As such, a simplistic 8 TeV study was included to form a baseline and to demonstrate what changes could be expected when going from \sqrt{s} of 8 to 14 TeV. Everything was done with simulated collisions, the description of which is in Section 9.2. The event selection and analysis strategy for both the simple 8 and 14 TeV analyses will be described in Section 9.3. Sensitivity projections, including both representative EFT limits and validity as well as simplified model limits, are provided in Section 9.4. A conclusion of the study and the relevance of the analysis in early Run-II is provided in Section 9.5.

9.2 Simulation

Both 8 and 14 TeV MC samples were used for the relevant sensitivity studies. While an effort was made to follow the same simulation procedure for both sets, this was not always possible. The 8 TeV simulated MC samples used are identical to those considered in the mono-jet paper, as previously described in Section 7.5.2. The details of the 8 TeV simulation will thus only be discussed when relevant to contrast with the procedure for 14 TeV simulation.

One of the largest differences between the two sets of samples is the type of simulation that was used. The 8 TeV samples were a mix of the full **GEANT4** detector simulation and the **AFII** parametrized detector response. The event selection criteria was all based on reconstructed quantities, following the relevant recommendations. The 14 TeV samples, on the other hand, were all produced with the full **GEANT4** detector simulation. However, a bug was found in one of the calorimeter constants used in this sample production. Given that the mono-jet analysis depends heavily on the calorimeters, both for the jets and E_T^{miss} , this was not acceptable. As such, while the samples were run through the **GEANT4** simulation, in the end only truth-level information was used, and the reconstructed information was ignored. The truth smearing parametrizations derived for studies aiming at the 2013 European Community for Future Accelerators (ECFA) High Luminosity LHC Experiments October Workshop, detailed in References [116, 117], were applied to the jets and E_T^{miss} . All event selection criteria were then applied to these truth-smearred objects.

The other main differences are intentional, and deal with the different data taking conditions present in 8 and 14 TeV collisions. In addition to changing the per-beam energy from 4 to 6.5 TeV in all simulations, the pileup conditions were also varied. 8 TeV MC samples contain overlaid **Pythia8** minimum-bias events. These events are generated for a series of different $\langle\mu\rangle$ values, in order to account for the differing number of average interactions of collisions in the 8 TeV dataset. Furthermore, the 8 TeV MC samples are then reweighted as a function of $\langle\mu\rangle$ in order to match the distributions observed in data, which is peaked at approximately $\langle\mu\rangle = 20$.

This very detailed approach to pileup modelling is not yet possible for 14 TeV MC, as the pileup conditions distributions are not yet known. The samples were instead generated for a fixed value of $\langle\mu\rangle$, but this was rendered irrelevant by the aforementioned calorimeter constant bug. By using only truth-level primary vertex information, pileup is explicitly removed. The truth smearing parametrizations were thus derived as a function of $\langle\mu\rangle$, thereby allowing one to consider many different pileup conditions with a single set of MC samples. Two pileup scenarios were considered for 14 TeV collisions, corresponding to expected LHC Run-II or HL-LHC conditions of $\langle\mu\rangle = 60$ and $\langle\mu\rangle = 140$ respectively.

In addition to varying the pileup conditions, the amount of luminosity was also varied, providing an estimate of how much the mono-jet analysis will benefit from the first few months of data (5 fb^{-1}), the first year of data (25 fb^{-1}), the full LHC dataset (300 fb^{-1}), and the full HL-LHC dataset (3000 fb^{-1}). The amount of integrated luminosity expected in the first few months and first year of data have since been revised to smaller values, but these were the best estimates at the time. The full set of pileup conditions and luminosity scenarios considered, along with the approximate timeframe in which the data is expected to be taken by, is provided in Table 9.1.

In order to conduct sensitivity studies, the SM background must be understood. As such, simulated samples for all relevant SM processes must be generated. For the mono-jet analysis, this consists primarily of $Z + \text{jets}$ and $W + \text{jets}$ events, with smaller contributions from single- t , $t\bar{t}$, and diboson processes. A priori, multi-jet (QCD) and NCBean also be large, but appropriate cleaning criteria were found to reduce

Dataset	Timeline	\sqrt{s} [TeV]	$\langle\mu\rangle$	$\int\mathcal{L}dt$ [fb^{-1}]
Run-I comparison	2012	8	20	20
First months of Run-II	2015 (mid)	14	60	5
First year of Run-II	2015 (end)	14	60	25
Full LHC program	2021	14	60	300
Full HL-LHC program	2030	14	140	3000

Table 9.1: The expected timeline, nominal centre of mass energy \sqrt{s} , average number of interactions per bunch crossing $\langle\mu\rangle$, and integrated luminosity $\int\mathcal{L}dt$ for each data taking scenario considered.

these both to negligible at 8 TeV. It was assumed that they will remain negligible given appropriate updated cuts at 14 TeV, and thus these processes were not considered further.

For 8 TeV, all of the necessary SM backgrounds had dedicated MC samples. This was not possible for 14 TeV at the time of these studies due to limited computing power. As such, a combination of 14 TeV samples and 8 TeV scale factors were used as necessary. The method for emulating each individual background will be discussed in the following paragraphs in more detail.

The single largest background is the $Z \rightarrow \nu\nu + \text{jets}$ process, which comprises at least 50% of the signal region SM event yield for low $E_{\text{T}}^{\text{miss}}$ cuts, and significantly more with high $E_{\text{T}}^{\text{miss}}$ requirements. Unfortunately, no such samples existed at 14 TeV. The $Z \rightarrow \nu\nu + \text{jets}$ process was instead taken from $W \rightarrow \mu\nu + \text{jets}$, with the muon considered as a part of the $E_{\text{T}}^{\text{miss}}$, and scaled by the ratio of the relevant cross-sections. This is known to provide a very good estimate, as seen during data-driven estimation of the $Z \rightarrow \nu\nu$ contribution to the signal region, both in the 8 TeV paper and previous mono-jet publications.

The second largest background is the $W \rightarrow \tau\nu + \text{jets}$ process, where the τ either decays hadronically or is missed. 14 TeV MC samples exist for this process, and thus they were directly used. The same goes for the $W \rightarrow \mu\nu + \text{jets}$ process, but samples did not exist for $W \rightarrow e\nu + \text{jets}$. In this case, the $W \rightarrow \mu\nu + \text{jets}$ samples were again used. This estimation is not as clean as for $Z \rightarrow \nu\nu$, as the electron interacts with the calorimetry. A scale factor from the ratio of 8 TeV $W \rightarrow \mu\nu$ and $W \rightarrow e\nu$ samples is used to scale the 14 TeV $W \rightarrow \mu\nu$ MC, providing an estimate of $W \rightarrow e\nu + \text{jets}$.

The much less significant $Z \rightarrow \ell\ell + \text{jets}$ process is strongly suppressed by the requirement that both of the leptons are missed, or that both τ leptons decay hadronically. It only contributes to the signal region at the sub-percent level at 8 TeV. It is assumed that lepton identification and performance will remain at similar levels in the future, and that the background will remain negligible, thus $Z \rightarrow \ell\ell + \text{jets}$ was not further considered.

Top-related backgrounds, consisting of single- t and $t\bar{t}$, are only a small contribution in this study due to the use of a the third jet veto. It is thus much less important than for the analysis described in Chapter 7. Nonetheless, it is not a negligible contribution, and must be considered. Only part of the necessary samples were available at the time, covering approximately 60% of the $t\bar{t}$ contribution. This was scaled with a ratio of the full set of single- t and $t\bar{t}$ 8 TeV samples compared to just the 8 TeV $t\bar{t}$ samples available at 14 TeV, providing a rough estimate of the contribution of top-related processes. This was deemed sufficient given that the total contribution in the signal region is expected to be at the percent level.

The final relevant background comes from diboson processes. As with top contributions, the diboson signal region contamination is reduced with respect to Chapter 7 due to the use of a third jet veto. The signal region event yield at 8 TeV contains $\mathcal{O}(1\%)$ diboson events, classifying it as a small background. Unfortunately no diboson samples existed at 14 TeV at the time, and thus it was decided to entirely

Background	8 TeV estimation	14 TeV estimation
$Z \rightarrow \nu\nu + \text{jets}$	Sherpa MC	Emulated from $W \rightarrow \mu\nu + \text{jets}$
$W \rightarrow \tau\nu + \text{jets}$	Sherpa MC	Sherpa MC
$W \rightarrow \mu\nu + \text{jets}$	Sherpa MC	Sherpa MC
$W \rightarrow e\nu + \text{jets}$	Sherpa MC	Emulated from $W \rightarrow \mu\nu + \text{jets}$
$Z \rightarrow \ell\ell + \text{jets}$	Sherpa MC	Neglected as a small background, $\mathcal{O}(< 1\%)$
single- t	MC@NLO MC	Subset of $t\bar{t}$ samples scaled with 8 TeV ratio
$t\bar{t}$	AcerMC+Pythia6 MC	
Diboson	Sherpa MC	Neglected as a small background, $\mathcal{O}(1\%)$

Table 9.2: Relevant Standard Model backgrounds and means of estimating them for both the 8 and 14 TeV analyses.

neglect the background. 8 TeV cross-checks showed that the kinematic shapes of the diboson samples was similar to that of the $Z + \text{jets}$ and $W + \text{jets}$ samples, so excluding diboson is primarily a scaling of the normalization, not a change in key variable shapes. Given that the diboson background is expected to transition from 8 to 14 TeV in a manner similar to the dominant $Z + \text{jets}$ and $W + \text{jets}$ backgrounds [118], thus remaining an $\mathcal{O}(1\%)$ contribution, it was deemed acceptable to neglect the background entirely. The impact of this assumption will be quantified in Section 9.4.2.

A summary of the backgrounds and means of estimation at 8 and 14 TeV is presented in Table 9.2.

Production of DM MC signal samples for the 8 TeV study has already been detailed in Appendix I.1. 14 TeV MC samples are produced following the same procedure, and with higher generator-level cuts to be relevant at the higher energy scale. While 8 TeV samples were generated with two matching scales of 80 and 300 GeV, the 14 TeV samples were generated with three scales of 200, 400, and 600 GeV. This allows for sufficient statistics at much higher kinematic thresholds.

One of the main objectives of the studies was to provide a first look at a proper EFT validity treatment, and the possibility of simplified models as a means of improving robustness of the DM interpretation in future ATLAS searches. It appears that both the ATLAS and CMS experiments will be heading in this direction, with the support of the theory community, as outlined in References [93, 94]. In order to do this, both EFT and simplified model signal points were produced. At the time of production, 14 TeV samples required a substantial overhead, and thus only a small subset of the points considered in Chapter 8 were included. In particular, the D5 EFT operator and the corresponding pure vector Z' simplified model were considered, for DM masses of 50 and 400 GeV. This mass choice was selected to cover the low and medium mass range, which are the most relevant regions of parameter space for collider searches.

9.3 Event selection

For these preliminary studies, a simple analysis along the lines of previous ATLAS mono-jet publications such as References [76, 6] was followed, rather than the more complicated selection in Chapter 7. In particular, the traditional mono-jet requirement of at most two jets was used. The selection also differs significantly between the 8 and 14 TeV selections to account for the different simulation types and the different conditions. Most notably, the 8 TeV analysis applies minimal cleaning cuts to reject bad events, necessary due to the use of reconstructed data with overlaid pileup. The 14 TeV analysis applies no such requirements due to the use of truth-smearred observables.

One of the core requirements of the mono-jet analysis is the presence of zero identified leptons in the signal region, satisfied by the application of a veto on all electrons and muons. This veto is dependent on the lepton identification efficiency, as the lepton must be properly categorized in order to be vetoed on. Identification efficiencies are available for 8 TeV, but the detector performance at 14 TeV has not yet been measured. The recommendation of the relevant experts was to assume that lepton identification efficiencies at 14 TeV identical to 8 TeV was followed. This was done by determining the veto efficiency at 8 TeV for our selection, and applying the same efficiency factor to the 14 TeV event yields.

The jet p_T and E_T^{miss} thresholds of 120 and 150 GeV respectively for the 8 TeV analysis have been set by the point at which the E_T^{miss} trigger used becomes 99% efficient. At this time, it is not known what threshold will be necessary under the significantly different Run-II conditions. Reasonable estimates put the new threshold for a pure E_T^{miss} trigger at 250 to 300 GeV. In order to ensure that the trigger is not a limitation, a more conservative threshold of 400 GeV of E_T^{miss} and 300 GeV for the leading jet p_T is applied. Additional signal regions with the same jet p_T but with higher E_T^{miss} cuts of 600 and 800 GeV are also defined.

One of the primary means by which the QCD background is controlled at 8 TeV is the use of a $\Delta\phi(E_T^{\text{miss}}, \text{jets})$ cut. This suppresses events where the E_T^{miss} source is the mis-measurement of very energetic jets, and has a minimal impact on the signal selection efficiency as the signal jet(s) are expected to be back-to-back with the E_T^{miss} . This cut is retained at 14 TeV.

The last key aspect of the simplified mono-jet selection is the jet definition, which controls when a collection of calorimeter energy is deemed a jet, and when it is ignored. This affects both the 3rd jet veto and to a lesser extent the $\Delta\phi(E_T^{\text{miss}}, \text{jets})$ cut. In general, jet definitions involve both a p_T and $|\eta|$ cut. For the mono-jet analysis, this corresponds to 30 GeV and 4.5 for 8 TeV, and 50 GeV and 3.6 at 14 TeV. The eta cut difference is simply due to the range in which the truth smearing parametrization is valid, and should make little difference. The change in the p_T cut is more complicated.

One of the main factors that the p_T threshold is designed to counter is pileup contamination, in which a jet from a different interaction within the same bunch crossing (in-time pileup) or from a different bunch crossing (out-of-time pileup) is incorrectly included as a part of the relevant event. This is unavoidable, and at low p_T the rate of pileup jets dwarfs the number of jets from the hard-scatter vertex in question. As the collision energy is increased, the pileup jet frequency for a given p_T value is also increased. The jet definition p_T threshold change from 30 to 50 GeV was done in order to ensure that the pileup contamination remains a minimal effect at 14 TeV. The studies which went into this choice are detailed in Appendix J.

The simplified selection applied to both 8 and 14 TeV samples are summarized in Table 9.3.

9.4 Expected sensitivity

9.4.1 Event yields

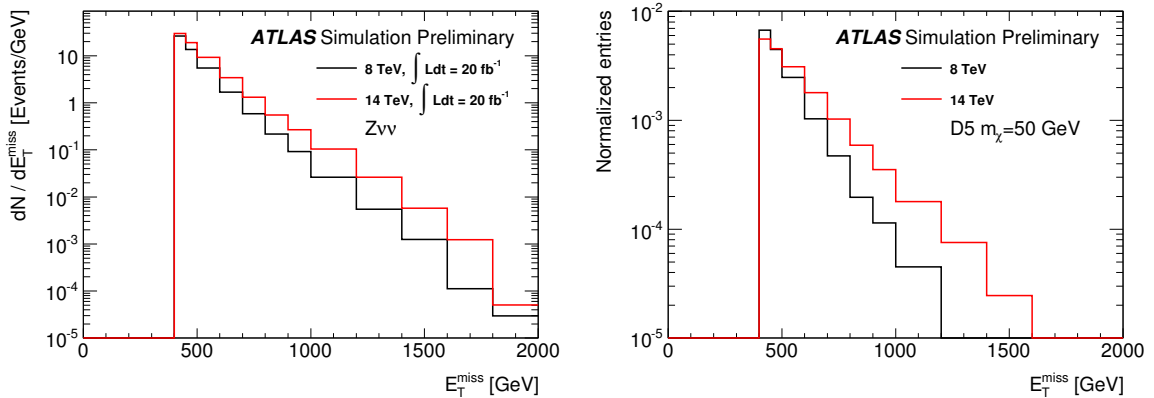
The raw event yields obtained after the analysis event selection is applied to both 8 and 14 TeV samples, normalized to the same integrated luminosity of $\int \mathcal{L} dt = 20 \text{ fb}^{-1}$, are shown in Table 9.4. Figure 9.1 provides a comparison of the 8 TeV and 14 TeV E_T^{miss} distributions for the $Z \rightarrow \nu\nu + \text{jets}$ process and one DM signal process, while Figure 9.2 covers the full set of SM backgrounds considered.

The increase in event yield for $Z \rightarrow \nu\nu + \text{jets}$ events is only moderate, as the Z mass is already

Cut purpose	$\sqrt{s} = 8 \text{ TeV}$	$\sqrt{s} = 14 \text{ TeV}$
Data quality	trigger primary vertex jet cleaning	primary vertex
Lepton veto	muons: $p_T > 7 \text{ GeV}, \eta < 2.5$ electrons: $p_T > 7 \text{ GeV}, \eta < 2.47$	emulated using the lepton selection efficiency from 8 TeV data
Jet and E_T^{miss} cuts	jet definition: $p_T > 30 \text{ GeV}, \eta < 4.5$ leading jet: $p_T > 300 \text{ GeV}, \eta < 2.0$ $E_T^{\text{miss}} > 400, 600, 800 \text{ GeV}$ $\Delta\phi(E_T^{\text{miss}}, \text{jets}) > 0.5$ $N_{\text{jets}} \leq 2$	$p_T > 50 \text{ GeV}, \eta < 3.6$ $p_T > 300 \text{ GeV}, \eta < 2.0$ $E_T^{\text{miss}} > 400, 600, 800 \text{ GeV}$ $\Delta\phi(E_T^{\text{miss}}, \text{jets}) > 0.5$ $N_{\text{jets}} \leq 2$

Table 9.3: A summary listing the event selection used for the 8 and 14 TeV sensitivity studies.

Signal region	Sample type	$\sqrt{s} = 8 \text{ TeV}$	$\sqrt{s} = 14 \text{ TeV}$	
		$\int \mathcal{L} dt = 20 \text{ fb}^{-1}$	Phase-I $\langle \mu \rangle = 60$	Phase-II $\langle \mu \rangle = 140$
$E_T^{\text{miss}} > 400 \text{ GeV}$	$Z \rightarrow \nu\nu + \text{jets}$	2800	3600	3900
	D5 EFT, $m_\chi = 50 \text{ GeV}$	200	3300	3300
	D5 EFT, $m_\chi = 400 \text{ GeV}$	120	2500	2600
$E_T^{\text{miss}} > 600 \text{ GeV}$	$Z \rightarrow \nu\nu + \text{jets}$	260	510	580
	D5 EFT, $m_\chi = 50 \text{ GeV}$	39	1100	1100
	D5 EFT, $m_\chi = 400 \text{ GeV}$	26	910	960
$E_T^{\text{miss}} > 800 \text{ GeV}$	$Z \rightarrow \nu\nu + \text{jets}$	37	100	110
	D5 EFT, $m_\chi = 50 \text{ GeV}$	8.5	390	400
	D5 EFT, $m_\chi = 400 \text{ GeV}$	6.6	340	350

Table 9.4: The number of events passing the selection outlined in Section 9.3, for the dominant $Z \rightarrow \nu\nu + \text{jets}$ background and the two D5 EFT signal samples considered, assuming an EFT suppression scale M_* of 1 TeV. All numbers are normalized to the same integrated luminosity for ease of comparison. Different pileup conditions at 14 TeV corresponding to different LHC upgrade phases are considered [8].Figure 9.1: E_T^{miss} distribution for the dominant $Z \rightarrow \nu\nu + \text{jets}$ background and the D5 (left) EFT operator for the pair-production of DM (right). 8 and 14 TeV E_T^{miss} distributions are overlaid, normalized to the same amount of integrated luminosity (left) or the same area (right). 14 TeV lines assume $\langle \mu \rangle = 140$ [8].

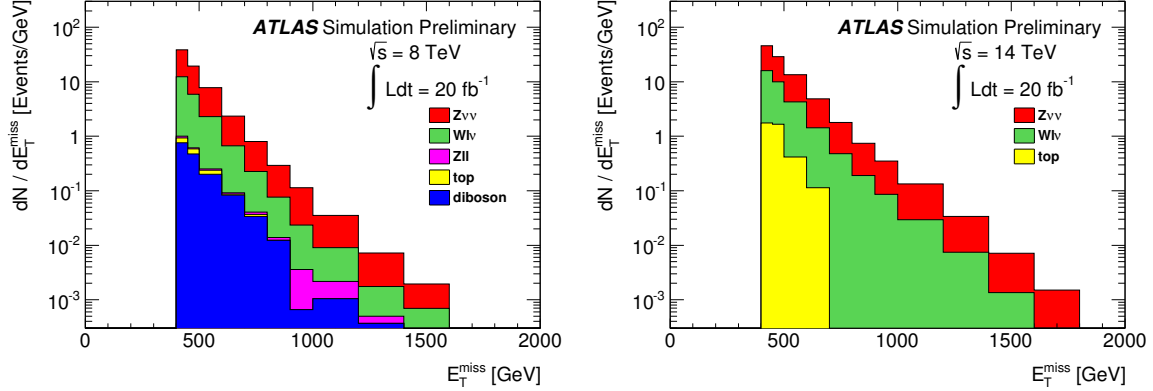


Figure 9.2: E_T^{miss} distribution for all of the Standard Model backgrounds considered, at 8 TeV (left) and 14 TeV with $\langle\mu\rangle = 140$ (right). The two plots are normalized to the same amount of integrated luminosity, demonstrating the increase in the average E_T^{miss} of the background processes as the centre of mass energy is increased. The $Z \rightarrow \ell\ell + \text{jets}$ and diboson processes are not considered at 14 TeV, as explained in Section 9.3, hence their exclusion from the plot [8].

significantly lower than the LHC energy scale, and thus the Z is produced significantly off-shell in order to satisfy the required E_T^{miss} threshold. The gain in event yields primarily comes from the higher \sqrt{s} enabling higher Z boosts, hence the increased gain at higher E_T^{miss} cuts. This provides an increase by a factor of $\mathcal{O}(1)$ in the event yields.

The DM sample event yields grow much more substantially (when moving from 8 TeV to 14 TeV as these models are contact interactions, where it is assumed that the scale of the mediator has not yet been reached. As such, the cross-section greatly depends on the collision energy. There is then an additional gain in event yields when moving to higher E_T^{miss} thresholds due to the higher \sqrt{s} allowing for more boosted DM particles. The cross-section dependence on \sqrt{s} provides an increase in the event yield for DM samples by a factor of $\mathcal{O}(10)$, while the E_T^{miss} increase provides an additional factor of $\mathcal{O}(1)$.

For the two different pileup scenarios considered at $\sqrt{s} = 14$ TeV, a small difference is observed despite normalization to the same integrated luminosity. This is understood as pileup raising the energy scale of the event, and further smearing out the resolution of both the jet p_T and E_T^{miss} . Given that these are both exponentially falling distributions, and increase in the resolution results in a larger number of events with lower kinematic values being shifted to higher kinematic ranges as compared to the opposite direction.

9.4.2 Effective field theories

An introduction to the WIMP DM EFT is provided in Section 8.1. For this sensitivity study, only the D5 operator, which connects a pair of quarks to a pair of DM particles through a vector mediator, is considered. Only DM masses of 50 and 400 GeV are investigated, corresponding to the most interesting mass region for collider searches. Limits and discovery potential results are set with respect to the suppression scale of the EFT, M_* . For the D5 operator, the simple UV completion of a pure vector mediator is easily re-integrated through $M_{\text{med}} = \sqrt{g_q g_\chi} M_*$, as explained in Appendix G.2. All results work under the assumption that the EFT is valid at the LHC energy scale, which is shown in Appendix G.4.2 to be a reasonable assumption so long as the geometric mean of the couplings $\sqrt{g_q g_\chi}$ satisfies the relation

$$\pi < \sqrt{g_q g_\chi} < 4\pi.$$

Four scenarios will be considered for the 14 TeV results, covering different expected data taking conditions, with one 8 TeV result provided for comparison. The different conditions studied follow those outlined in Table 9.1, with systematics assumed to be a simple flat uncertainty at 5% for all cases. This is expected to be reasonable to first order given that previous mono-jet studies typically have smaller systematic uncertainties for the low to medium E_T^{miss} thresholds, while at higher signal regions the statistical uncertainty dominates [76, 6].

An additional possibility of a flat 1% systematic is considered for the 300 fb^{-1} and 3000 fb^{-1} scenarios. This is an optimistic assumption, but it may be possible on the 10 to 15 year timescale given a concerted effort to improve the knowledge of the detector response, SM backgrounds, and theoretical unknowns. The last type of uncertainty pertaining to theoretical uncertainties is particularly important to improve over this timeframe, given that they constitute the dominant systematic uncertainty source in the most recent ATLAS mono-jet analysis, as described in Chapter 7.

Luminosity uncertainties on the background are not considered, as a standard mono-jet analysis employs data-driven techniques in which any uncertainty on the luminosity is cancelled. Systematic uncertainties on the signal MC is neglected for this simple study. The CLs [91] frequentist approach is used for both setting limits and probing the discovery potential. The impact of the missing $\mathcal{O}(1\%)$ diboson contribution is probed by scaling the background event yield up by a conservative 4%, resulting in differences to the limits at the level of $\mathcal{O}(0.1\%)$. It was thus concluded that the lack of diboson samples does not pose a serious problem for these studies.

Projected limits

The upgraded LHC will increase the sensitivity to new physics both by increasing the collision energy and by increasing the amount of available data. In order to determine which benefit will provide more substantial gains, the similar dataset sizes of 20 fb^{-1} at 8 TeV and 25 fb^{-1} at 14 TeV were compared, shown in Figure 9.3. A factor of 2 improvement in the limit strength is observed, indicating that the increase in \sqrt{s} provides an enormous boost to the sensitivity of the mono-jet DM interpretation.

The increase in luminosity is less beneficial for the mono-jet analysis, as shown in Figure 9.4. The assumed 5% systematic quickly becomes the limiting factor, resulting in no observed difference between 300 fb^{-1} and 3000 fb^{-1} of data. The much more optimistic assumption of a 1% systematic allows the analysis to continue to benefit from the increased luminosity. Comparing 5 fb^{-1} with a 5% systematic to 3000 fb^{-1} with a 1% systematic results in a factor of 2 increase in sensitivity. The change in centre of mass energy from 8 to 14 TeV with a single year of data is thus an approximately equivalent gain to increasing the luminosity by a factor of 600 and reducing the systematic uncertainty from 5% to 1%. This is a very important statement, as it confirms that the first period of data taking immediately after the LHC turns back on is crucial, providing the best hope for a discovery for years to come.

It must be noted that the final E_T^{miss} threshold of 800 GeV contributes to the reduced benefits from increasing the luminosity. 800 GeV is only a bit above what was available at 8 TeV with 20 fb^{-1} of data. 800 GeV is thus a reasonable threshold for 14 TeV with only 25 fb^{-1} of data, but the increase of luminosity would provide sufficient statistics to increase the E_T^{miss} threshold to much higher values. This would provide additional sensitivity with added luminosity. At the same time, the 1% systematic assumption is probably an over-optimistic goal. As a rough approximation, it is likely that the increased E_T^{miss} thresholds providing increased sensitivity and the less optimistic systematic uncertainties will




Figure 9.3: 95% confidence level lower limits on the EFT suppression scale M_* for the D5 operator, for DM masses of 50 GeV (solid lines) and 400 GeV (dashed lines), comparing 8 TeV (orange lines) to 14 TeV (blue lines). Three signal regions are considered, identical except for the E_T^{miss} threshold, which varies from 400 to 600 to 800 GeV. The results assume EFT to be fully valid, which is reasonable for $\pi < \sqrt{g_q g_\chi} < 4\pi$. A factor of 2 improvement in the limits is observed when moving from 8 to 14 TeV. [8]

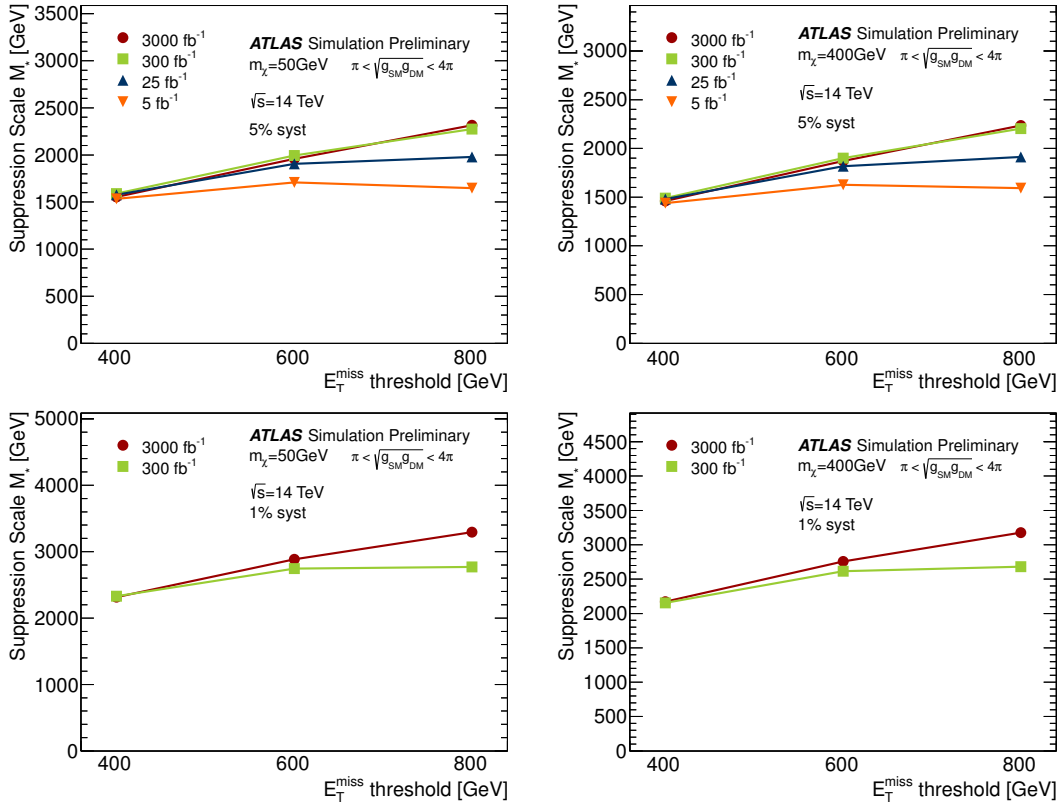


Figure 9.4: 95% confidence level lower limits on the EFT suppression scale M_* for the D5 operator at $\sqrt{s} = 14$ TeV, for DM masses of 50 GeV (left) and 400 GeV (right), assuming flat systematics of 5% (top) and 1% (bottom). Four different luminosity scenarios are considered for the 5% systematic, while only the latter two are studied for the 1% systematic. Three signal regions are considered, identical except for the E_T^{miss} threshold, which varies from 400 to 600 to 800 GeV. The results assume the EFT to be fully valid, which is reasonable for $\pi < \sqrt{g_q g_\chi} < 4\pi$. A factor of 2 improvement in the limits is observed when moving from 5 fb^{-1} and a 5% systematic to 3000 fb^{-1} and a 1% systematic. [8]




Figure 9.5: Discovery (5σ) and observation (3σ) potential for the D5 EFT operator at $\sqrt{s} = 14$ TeV, for DM masses of 50 GeV. DM masses of 400 GeV were found to provide similar results. Luminosity scenarios considered correspond to the first year of data taking (top), the full LHC run (bottom left), and the full HL-LHC run (bottom right). The results assume the EFT to be fully valid, which is reasonable over the full M_* range considered in each plot for couplings $\pi < \sqrt{g_{\text{SM}} g_{\text{DM}}} < 4\pi$. An approximately 75% improvement is seen when moving from 25 fb^{-1} and a 5% systematic to 3000 fb^{-1} and a 1% systematic. [8]

counteract each other, providing a net result similar to what has been observed. Further studies are necessary to confirm this.

Discovery potential

While limits are one measure of estimating the future sensitivity of the mono-jet analysis, the aim of any search for exotica is to discover new physics, not exclude it. As such, it is very interesting to study the discovery potential of the mono-jet analysis in Run-II and beyond. Figure 9.5 presents a scan over the signal significance as a function of the suppression scale M_* for 50 GeV DM particles. 400 GeV DM particles are found to result in similar significances, and thus are not shown here. The discovery potential is seen to be at the level of 1.5 TeV with the first year of 14 TeV data with 5% systematics, and only 1.8 TeV after the full HL-LHC dataset with the same level of uncertainty. Reducing the systematic uncertainty to the level of 1% allows the HL-LHC discovery potential to reach approximately 2.6 TeV. This clearly demonstrates the significant impact of systematic uncertainties on the discovery potential, and emphasizes the importance of the first year of $\sqrt{s} = 14$ TeV data in regards to the possible discovery of DM at the LHC.

9.4.3 Simplified models

Given the limitations of EFTs at the LHC with respect to their validity, a natural solution is the use of simplified models in which the mediator is not integrated out, thus naturally solving the problem. A more thorough study of this possibility has already been discussed in Section 8.3, but these sensitivity projections precede the studies done for the 8 TeV mono-jet paper. In these studies, the same pure vector Z' simplified model is used, but the study is restricted to DM masses of 50 and 400 GeV. These masses were chosen to match the EFT, thus allowing for a comparison between the EFT and Z' simplified model at high mediator masses. The added benefit of this investigation into simplified models is the comparison

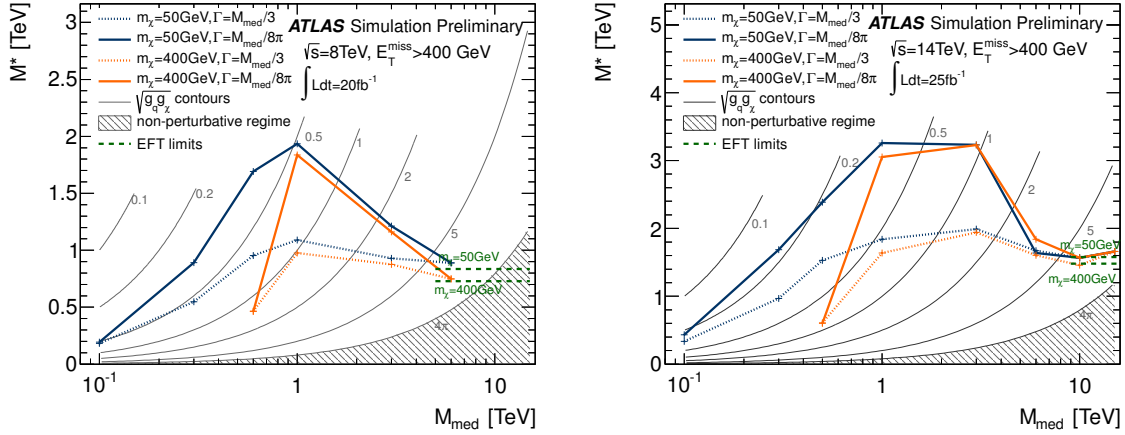


Figure 9.6: 95% confidence level lower limits on the pure vector Z' simplified model at 8 TeV (left) and 14 TeV (right) on the interaction/suppression scale M_* as a function of the mediator mass M_{med} . DM masses of 50 GeV (blue) and 400 GeV (orange) are considered, for mediator widths Γ_{med} of $M_{\text{med}}/3$ (dashed lines) and $M_{\text{med}}/8\pi$ (solid lines). The shaded region where the coupling $\sqrt{g_q g_\chi} > 4\pi$ is where the model becomes non-perturbative. All results are from the lowest signal region with a 400 GeV E_T^{miss} threshold due to limited statistics in the 8 TeV samples at higher thresholds. The three regimes of off-shell production (low M_{med}), resonant production (medium M_{med}), and contact interactions (high M_{med}) are clearly visible, with the corresponding EFT limits (green dashed lines) meeting with the simplified model at high enough values of M_{med} . The increase in the centre of mass energy results in the resonant regime being shifted up by due to the increase in the average momentum transfer. Limits are seen to improve by approximately a factor of 2 when comparing 8 to 14 TeV across the full range of M_{med} , consistent with what was seen for the EFT in Figure 9.3. [8]

between 8 and 14 TeV, which provides extra insight into the future of simplified models at the LHC.

Limits on the scale of the interaction at 8 and 14 TeV are provided in Figure 9.6. The three regimes of DM production are clearly visible, demonstrating the off-shell production, resonant production, and contact interaction mechanisms. As expected, the simplified model and EFT meet at high values of M_{med} , where the simplified model reduces to an EFT. The point at which the strongest limits are set, the peak of the resonant regime, shifts up in M_{med} when moving from 8 to 14 TeV due to the increase in the average momentum transfer of collisions, thus enabling higher mediator masses to be produced on-shell with an increased frequency. The strength of the limits is seen to grow by approximately a factor of two over the full range of mediator masses, in agreement with what was seen for the EFT in Figure 9.3.

9.5 Conclusions

The mono-jet topology is perfectly situated to benefit from the upcoming increase of the LHC centre of mass energy. With data from the first few months of 14 TeV collisions, the sensitivity of the mono-jet analysis to the D5 EFT operator for the pair-production of DM, with masses of 50 and 400 GeV, improves by a factor of two with respect to the full 8 TeV dataset. Additional luminosity provides further benefits, albeit at a reduced rate, and the analysis quickly becomes limited by the systematic uncertainties rather than the amount of available data. If the systematic uncertainties are reduced from the level of 5% to 1%, it is possible to gain an additional factor of two in sensitivity when comparing the first few months

of 14 TeV data to the end of the full HL-LHC program in the 2030s. A similar gain in sensitivity with added luminosity is observed for the discovery potential of the mono-jet analysis. The ultimate projected limit strength on M_* after the full HL-LHC program is at the level of 3.3 TeV, while the 5σ discovery threshold is approximately at 2.6 TeV of M_* .

The use of 1% systematic uncertainties is likely an overly optimistic goal, even by the end of the HL-LHC program. At the same time, the highest E_T^{miss} threshold considered of 800 GeV will not be optimal once sufficient statistics are available to push further into the tails of the distribution. Given that higher E_T^{miss} thresholds typically provide stronger limits due to the SM backgrounds falling off more quickly than the signal, these two effects are expected to approximately balance, and the sensitivity improvement of a factor of two with added luminosity and reduction of systematic uncertainties is expected to be a reasonable estimate.

The validity of the EFT at the LHC energy scale has been investigated by imposing the minimal requirement of $Q_{\text{tr}} < M_{\text{med}} = \sqrt{g_q g_\chi} M_*$. Without knowledge of the full UV theory, the only bounds that can be placed on the couplings is $0 < g_q, g_\chi < 4\pi$, and thus a scan over the possible coupling values is performed. The transition region in which the D5 operator goes from invalid to valid is found to be approximately $1 < \sqrt{g_q g_\chi} < \pi$, depending on the DM mass and centre of mass energy. Furthermore, all points considered are found to be $\mathcal{O}(100\%)$ valid for the coupling range $\pi < \sqrt{g_q g_\chi} < 4\pi$, providing a validity criteria for the limit sensitivity and discovery potential studies. There is therefore a sizable region of coupling parameter space over which the model is valid, but there is also a significant region over which it is not, and no further conclusion on the validity can be made without the observation of new physics.

The EFT validity is seen to depend strongly on the DM mass, where increasing the mass results in reduced validity. The validity is also seen to improve with the increase in the centre of mass energy, due to the significant increase in the limit strength, associated with the SM backgrounds falling off faster than the signal.

A pure vector mediator simplified model is also investigated, where the improvement in sensitivity due to the change from \sqrt{s} of 8 to 14 TeV averages at approximately a factor of two, consistent with the EFT observations. The mediator mass which is the most sensitive also increase, due to the higher average momentum transfer associated with the increase in the centre of mass energy. A comparison of the EFT and simplified model shows that the EFT is either correct or conservative for all but the lower mediator mass ranges of $M_{\text{med}} \lesssim 1$ TeV, where the off-shell production of DM results in reduced limits.

While the EFT validity improves with the increased in \sqrt{s} , questions remain as to what coupling values represent reality (if any), and thus the EFT is disfavoured. The many benefits of the simplified model approach, including the ability to benefit from the resonance structure where available, provide a compelling reason to more fully investigate such models in the future. This is particularly important in the coming year, given the expected overall factor of two improvement in sensitivity observed in all models considered.

The announcement that the LHC will be running at 13 TeV rather than 14 TeV in 2015 and possibly later will reduce the quoted sensitivities. However, the general trends remain unchanged. The centre of mass energy provides the single largest boost to the mono-jet DM interpretation sensitivity, while additional luminosity and reduced systematics provide slower and more marginal benefits.

Chapter 10

Conclusions

The Standard Model (SM) provides an excellent formulation of physics involving the known fermions and their interactions via the strong, weak, and electromagnetic forces. With the recent discovery of the Higgs boson at the LHC by both the ATLAS and CMS collaborations, the SM is now complete. However, there remain several open questions relating to observed phenomena for which the SM provides no explanations, such as the unknown nature of Dark Matter (DM).

While numerous independent astrophysical observations of DM have confirmed the existence of the phenomenon, it is not yet known whether DM is affected by any force other than gravity. The electromagnetic and strong forces have already been ruled out, but it remains possible for there to be a weak coupling between DM and the SM, resulting in Weakly Interacting Massive Particles (WIMPs). This particular assumption is motivated by the WIMP miracle, in which such particles at the GeV to TeV scale naturally produce the correct amount of DM during the formation of the universe to account for what is observed in astrophysical measurements. Furthermore, by making the assumption that such a coupling exists, several different types of searches for DM become possible.

One of the growing areas of enquiry in this broad scientific program is the use of colliders, such as the LHC, to probe the DM parameter space. This approach relies on the pair-production of DM particles as a by-product of SM collisions, where the DM particles escape the detector unobserved. A generic mono-X approach has been developed, where the presence of a single unbalanced high energy visible object, most often through initial state radiation, is used to tag such events. The imbalance of the single visible object is used to infer the presence of invisible particles, thus balancing the system and enforcing the conservation of momentum.

A particularly sensitive final state is the mono-jet topology, which is the mono-X search with the largest SM cross-section at hadron colliders such as the LHC. As the name implies, historically this analysis made use of a single jet balanced by a large source of missing transverse momentum. This has since evolved to include a second sub-dominant jet, thus enhancing the sensitivity to DM signals which typically are more energetic than SM processes and thus include more energetic fragmentation resulting in a higher jet multiplicity.

For the first time, this has now been extended to an inclusive jet selection, where a high E_T^{miss} is balanced by a recoil system of jets. The inclusion of additional radiation once again exploits the more energetic nature of DM production processes with the standard model to further improve the analysis sensitivity. In order to retain the original motivation for the search, the leading jet in the event is still

required to be dominant, and none of the jets are allowed to be in the same direction as the missing transverse momentum.

The performance of the mono-jet analysis is tightly tied to the level at which jets are understood, especially now that the number of allowed jets has grown. Not only does jet performance affect selection criteria applied on the jets themselves, but it also dominates the missing transverse momentum calculation. The missing transverse momentum is constructed from the visible objects, which in the mono-jet topology is primarily high energy jets. As such, improving the understanding of jets within the ATLAS detector is important for extending the sensitivity of the mono-jet analysis.

The full jet calibration chain, including new features designed to reduce flavour dependencies and improve the jet resolution, is described. Jets which are not contained in the calorimeters are also investigated, demonstrating the significant degradation in the scale and resolution of such jets. This is particularly important for the mono-jet analysis, where such events both have the wrong scale and can be the source of large amounts of fake missing energy. The first version of the correction for non-contained jets is shown to significantly reduce the impact, which will become more important in the near future as the LHC energy scale is increased and such energetic jets become more common.

Jets which pass through inactive areas of the central hadronic calorimeter are also studied. It is shown that such jets have the potential to introduce significant biases of up to 50% of the p_T at high energies. Such events are dangerous, both due to the incorrect scale and the creation of large amounts of fake missing transverse momentum, and thus a set of cleaning criteria is developed and applied in the mono-jet analysis.

The resulting systematic uncertainties on the jet energy scale are presented in complete detail, demonstrating the recent substantial improvements in the understanding of jets within the ATLAS detector. Advanced procedures to reduce the large number of uncertainty components with many insignificant sources to a more manageable reduced set, without any substantial losses in correlation information, is presented. These procedures have since become the default in ATLAS, and are used in the mono-jet analysis. A further procedure to estimate the uncertainties on the correlations themselves has been developed. While this advanced correlation information is beyond the sensitivity of the mono-jet analysis, it is very useful for further improving the knowledge of jet performance within ATLAS, particularly for SM measurements.

Once the jet performance is under control, the SM backgrounds to the mono-jet search must be considered. The dominant SM irreducible background to the mono-jet topology is the $Z \rightarrow \nu\nu + \text{jets}$ final state, which is estimated from four orthogonal control regions, corresponding to $Z \rightarrow \mu\mu$, $Z \rightarrow ee$, $W \rightarrow \mu\nu$, and $W \rightarrow e\nu$ selections. The sub-dominant $W \rightarrow \ell\nu + \text{jets}$ backgrounds are also estimated in a data-driven fashion from dedicated control regions, thus reducing the dependence on MC. This, combined with the full 2012 LHC dataset, provides sufficient statistics and sensitivity to reach much further into the tails of the missing transverse momentum distribution than previously done. No deviation from the SM expectations is observed, and thus limits are set on the visible cross-section for new physics.

This limit on the visible cross-section is translated into limits on the pair-production of weakly interacting DM processes. Two types of models are used, corresponding to a set of SM-DM contact interactions in which the mediator is assumed to be well beyond the energy scale of the LHC, and one specific simplified model where a pure vector Z' -like mediator is retained, thus allowing for the dependence on the mediator mass to be established. The sensitivity to both scalar and fermionic DM through contact interactions is considered.

The past years have seen repeated comments from the theory community on the use of effective field theories for DM production at the LHC energy scale. The assumption that such contact interactions are a valid picture may break down at such high energies, and thus a validity constraint has been derived and applied. This is the first complete collider treatment of model validity for all relevant contact interaction operators, and it demonstrates the impact of including validity constraints. Any matching between a given contact interaction and a representative theory which retains information on the mediator requires assumptions. The impact of these assumptions has been reduced where possible by providing a full scan over the available parameter space, thus allowing for the assumption to be varied.

A first study into the sensitivity of the mono-jet topology to DM at an upgraded LHC, similar to what is expected in the coming years, has been conducted. This study highlights the importance of improving the systematic uncertainties, as both the discovery potential and limit sensitivity quickly obtain a sufficient amount of statistics and thus become systematics-dominated. This study also clearly demonstrates the significant dependence of the analysis sensitivity on the centre of mass energy. Both the discovery potential and limit sensitivity are expected to grow by approximately a factor of two with the first year of data, suggesting that the coming year could prove to be very exciting for mono-jet searches at the LHC.

Bibliography

- [1] J. N. Bahcall, C. Flynn, and A. Gould, *Local dark matter from a carefully selected sample*, The Astrophysical Journal 389 (1992) 234–250.
- [2] F. Zwicky, *Die rotverschiebung von extragalaktischen nebeln*, Helvetica Physica Acta 6 (1933) 110–127.
- [3] Planck Collaboration, *Planck 2013 results. XV. CMB power spectra and likelihood*, Astronomy and Astrophysics 571 (2014) A15, [arXiv:1303.5075 \[astro-ph.CO\]](#).
- [4] DAMA Collaboration, *First results from DAMA/LIBRA and the combined results with DAMA/NaI*, The European Physical Journal C 56 (2008) 333–355, [arXiv:0804.2741 \[astro-ph\]](#).
- [5] ATLAS Collaboration Collaboration, ATLAS Collaboration, *Search for new phenomena in final states with an energetic jet and large missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, [arXiv:1502.01518 \[hep-ex\]](#).
- [6] ATLAS Collaboration, *Search for New Phenomena in Monojet plus Missing Transverse Momentum Final States using 10fb^{-1} of pp Collisions at $\sqrt{s}=8$ TeV with the ATLAS detector at the LHC*, Tech. Rep. ATLAS-CONF-2012-147, CERN, Geneva, Nov, 2012.
- [7] ATLAS Collaboration, *Search for dark matter in events with heavy quarks and missing transverse momentum in pp collisions with the ATLAS detector*, European Physics Journal C 75 (2015) no. 2, 92, [arXiv:1410.4031 \[hep-ex\]](#).
- [8] ATLAS Collaboration, *Sensitivity to WIMP Dark Matter in the Final States Containing Jets and Missing Transverse Momentum with the ATLAS Detector at 14 TeV LHC*, Tech. Rep. ATL-PHYS-PUB-2014-007, CERN, Geneva, Jun, 2014.
- [9] ATLAS Collaboration, *Search for new phenomena in events with a photon and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, [arXiv:1411.1559 \[hep-ex\]](#).
- [10] A. Purcell, *Go on a particle quest at the first CERN webfest. Le premier webfest du CERN se lance la conquete des particules*, BUL-NA-2012-269. 35/2012, Aug, 2012.
- [11] D. Griffiths, *Introduction to elementary particles*. John Wiley & Sons, 2008.
- [12] Particle Data Group Collaboration, K. Olive et al., *Review of Particle Physics*, Chinese Physics C 38 (2014) 090001.

- [13] M. E. Peskin and D. V. Schroeder, *An introduction to quantum field theory*. Westview, 1995.
- [14] Planck Collaboration, *Planck 2013 results. XVI. Cosmological parameters*, Astronomy and Astrophysics 571 (2014) A16, [arXiv:1303.5076 \[astro-ph.CO\]](#).
- [15] G. Bertone, D. Hooper, and J. Silk, *Particle dark matter: Evidence, candidates and constraints*, Physics Reports 405 (2005) 279–390, [arXiv:hep-ph/0404175 \[hep-ph\]](#).
- [16] D. Hooper, *TASI 2008 Lectures on Dark Matter*, Tech. Rep. FERMILAB-CONF-09-025-A, 2009. [arXiv:0901.4090 \[hep-ph\]](#).
- [17] K. Griest and D. Seckel, *Three exceptions in the calculation of relic abundances*, Physical Review D 43 (1991) 3191–3203. <http://link.aps.org/doi/10.1103/PhysRevD.43.3191>.
- [18] M. Azzaro, F. Prada, and C. Gutierrez, *Motion properties of satellites around external spiral galaxies*, ASP Conference Series 327 (2004) 268, [arXiv:astro-ph/0310487 \[astro-ph\]](#).
- [19] H. Hoekstra, H. Yee, and M. Gladders, *Current status of weak gravitational lensing*, New Astronomy Reviews 46 (2002) 767–781, [arXiv:astro-ph/0205205 \[astro-ph\]](#).
- [20] K. Begeman, A. Broeils, and R. Sanders, *Extended rotation curves of spiral galaxies: Dark haloes and modified dynamics*, Monthly Notices of the Royal Astronomical Society 249 (1991) no. 3, 523–537.
- [21] NASA Chandra X-ray Observatory, *NASA Finds Direct Proof of Dark Matter*, http://chandra.harvard.edu/press/06_releases/press_082106.html.
- [22] D. Clowe et al., *A direct empirical proof of the existence of dark matter*, The Astrophysical Journal 648 (2006) L109–L113, [arXiv:astro-ph/0608407 \[astro-ph\]](#).
- [23] G. W. Angus, B. Famaey, and H. Zhao, *Can MOND take a bullet? Analytical comparisons of three versions of MOND beyond spherical symmetry*, Monthly Notices of the Royal Astronomical Society 371 (2006) 138, [arXiv:astro-ph/0606216 \[astro-ph\]](#).
- [24] Planck Collaboration, *Planck picture gallery*, <http://www.cosmos.esa.int/web/planck/picture-gallery>.
- [25] Planck Collaboration, *Planck 2013 results. I. Overview of products and scientific results*, Astronomy and Astrophysics 571 (2014) A1, [arXiv:1303.5062 \[astro-ph.CO\]](#).
- [26] J. L. Feng, *Dark Matter Candidates from Particle Physics and Methods of Detection*, Annual Review of Astronomy and Astrophysics 48 (2010) 495–545, [arXiv:1003.0904 \[astro-ph.CO\]](#).
- [27] L. Evans and P. Bryant, *LHC Machine*, Journal of Instrumentation 3 (2008) no. 08, S08001. <http://stacks.iop.org/1748-0221/3/i=08/a=S08001>.
- [28] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, Journal of Instrumentation 3 (2008) S08003.
- [29] AC Team, *The four main LHC experiments*, CERN-AC-9906026, Jun, 1999.
- [30] M. Benedikt et al., *LHC Design Report*. CERN, Geneva, 2004.

- [31] C. Lefevre, *LHC: the guide*, CERN-Brochure-2008-001-Eng, Jan, 2008.
- [32] ATLAS Collaboration, *Luminosity Public Results*,
<https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResults>.
- [33] ATLAS Collaboration, *Improved luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC*, The European Physical Journal C 73 (2013) 2518, [arXiv:1302.4393](https://arxiv.org/abs/1302.4393) [hep-ex].
- [34] J. Pequeno, *Computer generated image of the whole ATLAS detector*, Mar, 2008.
- [35] J. Pequeno, *Computer generated image of the ATLAS inner detector*, Mar, 2008.
- [36] J. Pequeno, *Computer Generated image of the ATLAS calorimeter*, Mar, 2008.
- [37] A. V. Kotwal and C. Hays, *Electromagnetic shower properties in a lead-scintillator sampling calorimeter*, Nuclear Instruments and Methods A 729 (2013) 25–35.
- [38] M. Aharrouche et al., *Energy linearity and resolution of the ATLAS electromagnetic barrel calorimeter in an electron test-beam*, Nuclear Instruments and Methods A 568 (2006) no. 2, 601–623.
- [39] C. Grupen and B. Shwartz, *Particle detectors*. Cambridge University Press, 2008.
- [40] N. Ilic, *Performance of the ATLAS Liquid Argon Calorimeter After Three Years of LHC Operation and Plans for a Future Upgrade*, Tech. Rep. ATL-LARG-PROC-2013-016, CERN, Geneva, Dec, 2013.
- [41] J. Pequeno, *Computer generated image of the ATLAS Muons subsystem*, Mar, 2008.
- [42] ATLAS Collaboration, *Search for new phenomena in the dijet mass distribution using pp collision data at $\sqrt{s} = 8$ TeV with the ATLAS detector*, [arXiv:1407.1376](https://arxiv.org/abs/1407.1376) [hep-ex].
- [43] ATLAS Collaboration, *The ATLAS Simulation Infrastructure*, The European Physical Journal C 70 (2010) no. 3, 823–874. <http://dx.doi.org/10.1140/epjc/s10052-010-1429-9>.
- [44] GEANT4 Collaboration, *GEANT4: A simulation toolkit*, Nuclear Instruments and Methods A 506 (2003) 250–303.
- [45] P. M. Nadolsky et al., *Implications of CTEQ global analysis for collider observables*, Physical Review D 78 (2008) 013004, [arXiv:0802.0007](https://arxiv.org/abs/0802.0007) [hep-ph].
- [46] A. Martin, W. Stirling, R. Thorne, and G. Watt, *Parton distributions for the LHC*, European Physical Journal C 63 (2009) 189–285, [arXiv:0901.0002](https://arxiv.org/abs/0901.0002) [hep-ph].
- [47] H.-L. Lai et al., *New parton distributions for collider physics*, Physical Review D 82 (2010) 074024, [arXiv:1007.2241](https://arxiv.org/abs/1007.2241) [hep-ph].
- [48] T. Sjostrand, S. Mrenna, and P. Z. Skands, *PYTHIA 6.4 Physics and Manual*, Journal of High Energy Physics 0605 (2006) 026, [arXiv:hep-ph/0603175](https://arxiv.org/abs/hep-ph/0603175) [hep-ph].
- [49] T. Sjostrand, S. Mrenna, and P. Skands, *PYTHIA 8.1*, Computer Physics Communications 178 (2008) 852.

- [50] ATLAS Collaboration, *ATLAS tunes of PYTHIA 6 and Pythia 8 for MC11*, Tech. Rep. ATL-PHYS-PUB-2011-009, CERN, Geneva, Jul, 2011.
- [51] G. Folger and J. Wellisch, *String parton models in GEANT4*, eConf C0303241 (2003) MOMT007, [arXiv:nuc1-th/0306007](https://arxiv.org/abs/nuc1-th/0306007) [nuc1-th].
- [52] H. Pi, *An event generator for interactions between hadrons and nuclei {FRITIOF} version 7.0*, Computer Physics Communications 71 (1992) no. 12, 173 – 192.
<http://www.sciencedirect.com/science/article/pii/001046559290082A>.
- [53] H. W. Bertini, *Intranuclear-Cascade Calculation of the Secondary Nucleon Spectra from Nucleon-Nucleus Interactions in the Energy Range 340 to 2900 MeV and Comparisons with Experiment*, Physical Review 188 (1969) 1711–1730.
<http://link.aps.org/doi/10.1103/PhysRev.188.1711>.
- [54] J. D. Chapman et al., *The ATLAS detector digitization project for 2009 data taking*, Journal of Physics: Conference Series 219 (2010) no. 3, 032031.
<http://stacks.iop.org/1742-6596/219/i=3/a=032031>.
- [55] ATLAS Collaboration, *Electron and photon energy calibration with the ATLAS detector using LHC Run 1 data*, The European Physical Journal C 74 (2014) no. 10, 3071, [arXiv:1407.5063](https://arxiv.org/abs/1407.5063) [hep-ex].
- [56] W. Lamp et al., *Calorimeter Clustering Algorithms: Description and Performance*, Tech. Rep. ATL-LARG-PUB-2008-002, CERN, Geneva, Apr, 2008.
- [57] ATLAS Collaboration, *Measurements of the photon identification efficiency with the ATLAS detector using 4.9 fb⁻¹ of pp collision data collected in 2011*, Tech. Rep. ATLAS-CONF-2012-123, CERN, Geneva, Aug, 2012.
- [58] ATLAS Collaboration, *Measurement of the inclusive isolated prompt photon cross section in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, Physical Review D 83 (2011) 052005, [arXiv:1012.4389](https://arxiv.org/abs/1012.4389) [hep-ex].
- [59] ATLAS Collaboration, *Electron performance measurements with the ATLAS detector using the 2010 LHC proton-proton collision data*, The European Physical Journal 72 (2012) 1909, [arXiv:1110.3174](https://arxiv.org/abs/1110.3174) [hep-ex].
- [60] ATLAS Collaboration, *Measurement of the muon reconstruction performance of the ATLAS detector using 2011 and 2012 LHC proton-proton collision data*, European Physics Journal C 74 (2014) no. 11, 3130, [arXiv:1407.3935](https://arxiv.org/abs/1407.3935) [hep-ex].
- [61] ATLAS Collaboration, *Calibration of b-tagging using dileptonic top pair events in a combinatorial likelihood approach with the ATLAS experiment*, Tech. Rep. ATLAS-CONF-2014-004, CERN, Geneva, Feb, 2014.
- [62] ATLAS Collaboration, *Performance of Missing Transverse Momentum Reconstruction in Proton-Proton Collisions at 7 TeV with ATLAS*, The European Physical Journal C 72 (2012) 1844, [arXiv:1108.5602](https://arxiv.org/abs/1108.5602) [hep-ex].

- [63] ATLAS Collaboration, *Performance of Missing Transverse Momentum Reconstruction in ATLAS studied in Proton-Proton Collisions recorded in 2012 at 8 TeV*, Tech. Rep. ATLAS-CONF-2013-082, CERN, Geneva, Aug, 2013.
- [64] ATLAS Collaboration, *Performance of jet substructure techniques for large- R jets in proton-proton collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector*, Journal of High Energy Physics 1309 (2013) 076, [arXiv:1306.4945](https://arxiv.org/abs/1306.4945) [hep-ex].
- [65] J. E. Huth et al., *Toward a standardization of jet definitions*, Tech. Rep. FERMILAB-CONF-90-249-E, Dec, 1990.
- [66] G. P. Salam, *Towards Jetography*, The European Physical Journal C 67 (2010) 637–686, [arXiv:0906.1833](https://arxiv.org/abs/0906.1833) [hep-ph].
- [67] M. Cacciari, G. Salam, and G. Soyez, *FastJet user manual*, The European Physical Journal C 72 (2012) no. 3, . <http://dx.doi.org/10.1140/epjc/s10052-012-1896-2>.
- [68] M. Cacciari, G. P. Salam, and G. Soyez, *The anti- k_t jet clustering algorithm*, Journal of High Energy Physics 2008 (2008) no. 04, 063.
- [69] ATLAS Collaboration, *Selection of jets produced in proton-proton collisions with the ATLAS detector using 2011 data*, Tech. Rep. ATLAS-CONF-2012-020, CERN, Geneva, Mar, 2012.
- [70] ATLAS Collaboration, *Jet energy resolution in proton-proton collisions at $\sqrt{s} = 7$ TeV recorded in 2010 with the ATLAS detector*, The European Physical Journal C 73 (2013) 2306, [arXiv:1210.6210](https://arxiv.org/abs/1210.6210) [hep-ex].
- [71] ATLAS Collaboration, *Jet/ $E_{T\text{Miss}}$ Approved 2013 JER 2011*, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/JetEtmissApproved2013Jer2011>.
- [72] UA1 Collaboration, *Experimental observation of events with large missing transverse energy accompanied by a jet or a photon (S) in $p\bar{p}$ collisions at $\sqrt{s} = 540$ GeV*, Physics Letters B 139 (1984) no. 12, 115 – 125. <http://www.sciencedirect.com/science/article/pii/0370269384900467>.
- [73] CDF Collaboration, *Search for large extra dimensions in final states containing one photon or jet and large missing transverse energy produced in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ -TeV*, Physical Review Letters 101 (2008) 181602, [arXiv:0807.3132](https://arxiv.org/abs/0807.3132) [hep-ex].
- [74] ATLAS Collaboration, *Search for new phenomena with the monojet and missing transverse momentum signature using the ATLAS detector in $\sqrt{s} = 7$ TeV proton-proton collisions*, Phys.Lett. B705 (2011) 294–312, [arXiv:1106.5327](https://arxiv.org/abs/1106.5327) [hep-ex].
- [75] CMS Collaboration, *Search for New Physics with a Mono-Jet and Missing Transverse Energy in pp Collisions at $\sqrt{s} = 7$ TeV*, Physical Review Letters 107 (2011) 201804, [arXiv:1106.4775](https://arxiv.org/abs/1106.4775) [hep-ex].
- [76] ATLAS Collaboration, *Search for dark matter candidates and large extra dimensions in events with a jet and missing transverse momentum with the ATLAS detector*, Journal of High Energy Physics 1304 (2013) 075, [arXiv:1210.4491](https://arxiv.org/abs/1210.4491) [hep-ex].

- [77] CMS Collaboration, *Search for dark matter and large extra dimensions in monojet events in pp collisions at $\sqrt{s} = 7$ TeV*, Journal of High Energy Physics 1209 (2012) 094, arXiv:1206.5663 [hep-ex].
- [78] CMS Collaboration, *Search for dark matter, extra dimensions, and unparticles in monojet events in proton-proton collisions at $\sqrt{s} = 8$ TeV*, arXiv:1408.3583 [hep-ex].
- [79] U. Haisch, F. Kahlhoefer, and E. Re, *QCD effects in mono-jet searches for dark matter*, Journal of High Energy Physics 1312 (2013) 007, arXiv:1310.4491 [hep-ph].
- [80] D. Casadei et al., *The implementation of the ATLAS missing Et triggers for the initial LHC operation*, Tech. Rep. ATL-DAQ-PUB-2011-00, CERN, Geneva, 2011.
- [81] T. Gleisberg et al., *Event generation with Sherpa 1.1*, Journal of High Energy Physics 02 (2009) 007, 0811.4622 [hep-ph].
- [82] S. Frixione et al., *The MC@NLO 4.0 Event Generator*, arXiv:1010.0819 [hep-ph].
- [83] B. P. Kersevan and E. Richter-Was, *The Monte Carlo event generator AcerMC version 1.0 with interfaces to PYTHIA 6.2 and HERWIG 6.3*, Computer Physics Communications 149 (2003) 142, arXiv:hep-ph/0201302 [hep-ph].
- [84] J. H. Kuhn, A. Kulesza, S. Pozzorini, and M. Schulze, *Electroweak corrections to hadronic production of W bosons at large transverse momenta*, Nuclear Physics B 797 (2008) 27–77, arXiv:0708.0476 [hep-ph].
- [85] S. Owen, *Data-driven estimation of the QCD multijet background to SUSY searches with jets and missing transverse momentum at ATLAS using jet smearing*, Tech. Rep. ATL-PHYS-INT-2012-008, CERN, Geneva, Feb, 2012.
- [86] S. Amoroso et al., *Search for squarks and gluinos with the ATLAS detector in final states with jets and missing transverse momentum and 20.3 fb^{-1} of $\sqrt{s} = 8 \text{ TeV}$ proton-proton collision data: supporting documentation*, Tech. Rep. ATL-COM-PHYS-2013-1224, CERN, Geneva, Aug, 2013.
- [87] ATLAS Collaboration, *Characterisation and mitigation of beam-induced backgrounds observed in the ATLAS detector during the 2011 proton-proton run*, Journal of Instrumentation 8 (2013) P07004, arXiv:1303.0223 [hep-ex].
- [88] M. L. Mangano, M. Moretti, F. Piccinini, R. Pittau, and A. Polosa, *ALPGEN, a generator for hard multiparton processes in hadronic collisions*, Journal of High Energy Physics 07 (2003) 001, arXiv:hep-ph/0206293.
- [89] J. M. Campbell and R. K. Ellis, *Next-to-leading order corrections to W^+ 2 jet and Z^+ 2 jet production at hadron colliders*, Physical Review D 65 (2002) 113007, arXiv:hep-ph/0202176 [hep-ph].
- [90] L. Lyons, D. Gibaut, and P. Clifford, *How to Combine Correlated Estimates of a Single Physical Quantity*, Nuclear Instruments and Methods A 270 (1988) 110.
- [91] A. L. Read, *Presentation of search results: the CL s technique*, Journal of Physics G: Nuclear and Particle Physics 28 (2002) no. 10, 2693. <http://stacks.iop.org/0954-3899/28/i=10/a=313>.

- [92] J. Goodman et al., *Constraints on dark matter from colliders*, Physical Review D 82 (2010) 116010. <http://link.aps.org/doi/10.1103/PhysRevD.82.116010>.
- [93] J. Abdallah et al., *Simplified Models for Dark Matter and Missing Energy Searches at the LHC*, arXiv:1409.2893 [hep-ph].
- [94] S. Malik et al., *Interplay and Characterization of Dark Matter Searches at Colliders and in Direct Detection Experiments*, arXiv:1409.4075 [hep-ex].
- [95] P. J. Fox, R. Harnik, J. Kopp, and Y. Tsai, *Missing Energy Signatures of Dark Matter at the LHC*, Physical Review D 85 (2012) 056011, arXiv:1109.4398 [hep-ph].
- [96] O. Buchmueller, M. J. Dolan, and C. McCabe, *Beyond Effective Field Theory for Dark Matter Searches at the LHC*, Journal of High Energy Physics 1401 (2014) 025, arXiv:1308.6799 [hep-ph].
- [97] O. Buchmueller, M. J. Dolan, S. A. Malik, and C. McCabe, *Characterising dark matter searches at colliders and direct detection experiments: Vector mediators*, arXiv:1407.8257 [hep-ph].
- [98] WMAP Collaboration, *Seven-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation*, The Astrophysical Journal Supplemental 192 (2011) 18, arXiv:1001.4538 [astro-ph.CO].
- [99] CDMS Collaboration, *Silicon Detector Dark Matter Results from the Final Exposure of CDMS II*, Physical Review Letters 111 (2013) 251301, arXiv:1304.4279 [hep-ex].
- [100] CRESST-II Collaboration, *Results from 730 kg days of the CRESST-II Dark Matter Search*, The European Physical Journal C 72 (2012) 1971, arXiv:1109.0702 [astro-ph.CO].
- [101] CoGeNT Collaboration, *Maximum Likelihood Signal Extraction Method Applied to 3.4 years of CoGeNT Data*, arXiv:1401.6234 [astro-ph.CO].
- [102] CoGeNT Collaboration, *CoGeNT: A Search for Low-Mass Dark Matter using p-type Point Contact Germanium Detectors*, Physical Review D 88 (2013) no. 1, 012002, arXiv:1208.5737 [astro-ph.CO].
- [103] C. Savage, G. Gelmini, P. Gondolo, and K. Freese, *Compatibility of DAMA/LIBRA dark matter detection with other searches*, Journal of Cosmology and Astroparticle Physics 0904 (2009) 010, arXiv:0808.3607 [astro-ph].
- [104] LUX Collaboration, *First results from the LUX dark matter experiment at the Sanford Underground Research Facility*, arXiv:1310.8214 [astro-ph.CO].
- [105] XENON100 Collaboration, *Limits on spin-dependent WIMP-nucleon cross sections from 225 live days of XENON100 data*, Physical Review Letters 111 (2013) no. 2, 021301, arXiv:1301.6620 [astro-ph.CO].
- [106] COUPP Collaboration, *First dark matter search results from a 4-kg CF₃I bubble chamber operated in a deep underground site*, Physical Review D 86 (2012) no. 5, 052001, arXiv:1204.3094 [astro-ph.CO].

- [107] IceCube Collaboration, *Search for Dark Matter Annihilations in the Sun with the 79-String IceCube Detector*, Physical Review Letters 110 (2013) 131302.
<http://link.aps.org/doi/10.1103/PhysRevLett.110.131302>.
- [108] PICASSO Collaboration, *Constraints on Low-Mass WIMP Interactions on ^{19}F from PICASSO*, Physics Letters B 711 (2012) 153–161, [arXiv:1202.1240](https://arxiv.org/abs/1202.1240) [hep-ex].
- [109] SIMPLE Collaboration, *Final Analysis and Results of the Phase II SIMPLE Dark Matter Search*, Physical Review Letters 108 (2012) no. 20, 201302, [arXiv:1106.3014](https://arxiv.org/abs/1106.3014).
- [110] Super-Kamiokande Collaboration, *An Indirect Search for Weakly Interacting Massive Particles in the Sun Using 3109.6 Days of Upward-going Muons in Super-Kamiokande*, The Astrophysical Journal 742 (2011) 78, [arXiv:1108.3384](https://arxiv.org/abs/1108.3384) [astro-ph.HE].
- [111] ATLAS Collaboration, *Search for dark matter in events with a hadronically decaying W or Z boson and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, Physical Review Letters 112 (2014) no. 4, 041802, [arXiv:1309.4017](https://arxiv.org/abs/1309.4017) [hep-ex].
- [112] ATLAS Collaboration, *Search for new particles in events with one lepton and missing transverse momentum in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, Journal of High Energy Physics 1409 (2014) 037, [arXiv:1407.7494](https://arxiv.org/abs/1407.7494) [hep-ex].
- [113] ATLAS Collaboration, *Search for dark matter in events with a Z boson and missing transverse momentum in pp collisions at $\sqrt{s}=8$ TeV with the ATLAS detector*, Physical Review D 90 (2014) 012004, [arXiv:1404.0051](https://arxiv.org/abs/1404.0051) [hep-ex].
- [114] Fermi-LAT Collaboration Collaboration, *Fermi-LAT Collaboration, Dark matter constraints from observations of 25 Milky Way satellite galaxies with the Fermi Large Area Telescope*, Physical Review D 89 (2014) 042001. <http://link.aps.org/doi/10.1103/PhysRevD.89.042001>.
- [115] HESS Collaboration, *Search for a Dark Matter Annihilation Signal from the Galactic Center Halo with H.E.S.S.*, Physical Review Letters 106 (2011) no. 16, 161301, [arXiv:1103.3266](https://arxiv.org/abs/1103.3266) [astro-ph.HE].
- [116] ATLAS Collaboration, *Performance assumptions for an upgraded ATLAS detector at a High-Luminosity LHC*, Tech. Rep. ATL-PHYS-PUB-2013-004, CERN, Geneva, Mar, 2013.
- [117] ATLAS Collaboration, *Performance assumptions based on full simulation for an upgraded ATLAS detector at a High-Luminosity LHC*, Tech. Rep. ATL-PHYS-PUB-2013-009, CERN, Geneva, Sep, 2013.
- [118] J. M. Campbell, R. K. Ellis, and C. Williams, *Vector boson pair production at the LHC*, Journal of High Energy Physics 1107 (2011) 018, [arXiv:1105.0020](https://arxiv.org/abs/1105.0020) [hep-ph].
- [119] L. Bergstrom, *Nonbaryonic dark matter: Observational evidence and detection methods*, Reports on Progress in Physics 63 (2000) 793, [arXiv:hep-ph/0002126](https://arxiv.org/abs/hep-ph/0002126) [hep-ph].
- [120] L. Bergström and A. Goobar, *Cosmology and particle astrophysics*. Springer, 1999.
- [121] F. Arneodo, *Dark Matter Searches*, ArXiv e-prints (2013) , [arXiv:1301.0441](https://arxiv.org/abs/1301.0441) [astro-ph.IM].

- [122] G. Chalons, *Gamma-ray lines constraints in the NMSSM*, arXiv:1204.4591 [hep-ph].
- [123] J. Kopp, *Constraints on dark matter annihilation from AMS-02 results*, Physical Review D 88 (2013) 076013. <http://link.aps.org/doi/10.1103/PhysRevD.88.076013>.
- [124] A. D. Simone, A. Riotto, and W. Xue, *Interpretation of AMS-02 results: correlations among dark matter signals*, Journal of Cosmology and Astroparticle Physics 2013 (2013) no. 05, 003. <http://stacks.iop.org/1475-7516/2013/i=05/a=003>.
- [125] AMS Collaboration, *First Result from the Alpha Magnetic Spectrometer on the International Space Station: Precision Measurement of the Positron Fraction in Primary Cosmic Rays of 0.5-350 GeV*, Physical Review Letters 110 (2013) 141102. <http://link.aps.org/doi/10.1103/PhysRevLett.110.141102>.
- [126] C. Weniger, *A Tentative Gamma-Ray Line from Dark Matter Annihilation at the Fermi Large Area Telescope*, Journal of Cosmology and Astroparticle Physics 1208 (2012) 007, arXiv:1204.2797 [hep-ph].
- [127] Fermi-LAT Collaboration, *Search for gamma-ray spectral lines with the Fermi Large Area Telescope and dark matter implications*, Physical Review D 88 (2013) 082002. <http://link.aps.org/doi/10.1103/PhysRevD.88.082002>.
- [128] ATLAS Collaboration, *Approved Tile Calorimeter Plots - Detector Status*, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ApprovedPlotsTile#Detector_Status.
- [129] K. Perez, *Masked tile calorimeter corrections*, Private communication, internal talk on 2010-11-18.
- [130] ATLAS Collaboration, *Luminosity Determination in pp Collisions at $\sqrt{s} = 7$ TeV Using the ATLAS Detector at the LHC*, The European Physical Journal C 71 (2011) 1630, arXiv:1101.2185 [hep-ex].
- [131] ATLAS Collaboration, *Pile-up subtraction and suppression for jets in ATLAS*, Tech. Rep. ATLAS-CONF-2013-083, CERN, Geneva, Aug, 2013.
- [132] ATLAS Collaboration, *Tagging and suppression of pileup jets with the ATLAS detector*, Tech. Rep. ATLAS-CONF-2014-018, CERN, Geneva, May, 2014.
- [133] T. Gabriel et al., *Energy dependence of hadronic activity*, Nuclear Instruments and Methods A 338 (1994) no. 2, 336–347.
- [134] D. E. Groom, *Energy flow in a hadronic cascade: Application to hadron calorimetry*, Nuclear Instruments and Methods A 572 (2007) no. 2, 633 – 653.
- [135] ATLAS Collaboration, *Jet energy measurement and its systematic uncertainty in proton-proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, arXiv:1406.0076 [hep-ex].
- [136] ATLAS Collaboration, *Properties of Jets and Inputs to Jet Reconstruction and Calibration with the ATLAS Detector Using Proton-Proton Collisions at $\sqrt{s} = 7$ TeV*, Tech. Rep. ATLAS-CONF-2010-053, CERN, Geneva, Jul, 2010.

- [137] S. Batista et al., *Global Sequential Calibration with the ATLAS Detector in Proton-Proton Collisions at $\sqrt{s} = 8$ TeV with ATLAS 2012 data*, Tech. Rep. ATL-COM-PHYS-2014-753, CERN, Geneva, Jun, 2014.
- [138] C. W. Fabjan and T. Ludlam, *Calorimetry in high-energy physics*, Annual Review of Nuclear and Particle Science 32 (1982) no. 1, 335–389.
- [139] M. Bahr et al., *Herwig++ Physics and Manual*, European Physics Journal C 58 (2008) 639–707, arXiv:0803.0883 [hep-ph].
- [140] S. Frixione, P. Nason, and G. Ridolf, *A Positive-Weight Next-to-Leading-Order Monte Carlo for Heavy Flavour Hadroproduction*, Journal of High Energy Physics 0709 (2007) 126, arXiv:hep-ph/0707.3088.
- [141] ATLAS Collaboration, *Jet/EtMiss Approved 2013 JES Uncertainty*, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/JetEtmissApproved2013JESUncertainty>.
- [142] ATLAS Collaboration, *Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV*, The European Physical Journal C 73 (2013) 2304, arXiv:1112.6426 [hep-ex].
- [143] ATLAS Collaboration, *Jet energy scale uncertainty correlations between ATLAS and CMS*, Tech. Rep. ATL-PHYS-PUB-2014-020, CERN, Geneva, Oct, 2014.
- [144] S. Gupta, *Punch-through derivation plots*, Private communications, 2013-2014.
- [145] G. Busoni et al., *On the Validity of the Effective Field Theory for Dark Matter Searches at the LHC, Part II: Complete Analysis for the s-channel*, Journal of Cosmology and Astroparticle Physics 1406 (2014) 060, arXiv:1402.1275 [hep-ph].
- [146] G. Busoni et al., *On the Validity of the Effective Field Theory for Dark Matter Searches at the LHC Part III: Analysis for the t-channel*, arXiv:1405.3101 [hep-ph].
- [147] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer, and T. Stelzer, *MadGraph 5 : Going Beyond*, Journal of High Energy Physics 1106 (2011) 128, arXiv:1106.0522 [hep-ph].
- [148] M. Botje et al., *The PDF4LHC Working Group Interim Recommendations*, arXiv:1101.0538 [hep-ph].
- [149] R. D. Ball et al., *Parton distributions with LHC data*, Nuclear Physics B 867 (2013) 244–289, arXiv:1207.1303 [hep-ph].
- [150] R. D. Ball et al., *Impact of Heavy Quark Masses on Parton Distributions and LHC Phenomenology*, Nuclear Physics B 849 (2011) 296–363, arXiv:1101.1300 [hep-ph].
- [151] M. Whalley, D. Bourilkov, and R. Group, *The Les Houches accord PDFs (LHAPDF) and LHAGLUE*, arXiv:hep-ph/0508110 [hep-ph].
- [152] J. M. Campbell, J. Huston, and W. Stirling, *Hard Interactions of Quarks and Gluons: A Primer for LHC Physics*, Reports on Progress in Physics 70 (2007) 89, arXiv:hep-ph/0611148 [hep-ph].

- [153] P. Z. Skands, *Tuning Monte Carlo Generators: The Perugia Tunes*, Physical Review D 82 (2010) 074018, [arXiv:1005.3457 \[hep-ph\]](#).
- [154] J. Wenninger, *Energy Calibration of the LHC Beams at 4 TeV*, Tech. Rep. CERN-ATS-2013-040, CERN, Geneva, May, 2013.

Glossary

AFII Atlas Fast simulation version II, a parametrized detector response framework.

AOD Analysis Object Data.

ATLAS A Toroidal LHC ApparatuS.

BCH Bad CHannel.

BLUE Best Linear Unbiased Estimate.

BSM Beyond the Standard Model.

CL Confidence Level.

CLs Confidence Levels statistical method for setting frequentist limits.

CMB Cosmic Microwave Background.

CMS Compact Muon Solenoid.

CPV Charge-Parity Violation.

CSC Cathode Strip Chamber.

CT10NLO A recent next to leading order parton distribution function provided by the CT (CTEQ) collaboration.

CTEQ6L1 An old leading order parton distribution function provided by the CTEQ collaboration.

CTEQ6M An old leading order parton distribution function provided by the CTEQ collaboration.

D3PD Derived³ Physics Data.

DAOD Derived Analysis Object Data.

DD Direct Detection.

DESD Derived Event Summary Data.

DM Dark Matter.

ECFA European Community for Future Accelerators.

EF Event Filter.

EFT Effective Field Theory.

EM ElectroMagnetic, usually the scale at which the raw electronic signals are used to determine the jet energy.

EMB ElectroMagnetic Barrel.

EMEC ElectroMagnetic EndCap.

ESD Event Summary Data.

E_T^{miss} Missing Transverse Momentum.

FastJet A software library for building and manipulating jets.

FCal Forward CALorimeter.

Final2012 The final version of the jet calibration, uncertainties, and resolution for the 2012 dataset, released in October 2014 for the JES, with the JER still ongoing.

GEANT4 GEometry ANd TRacking version 4, for the simulation of particles passing through matter.

GRL Good Runs List.

GSC Global Sequential Calibration.

hadron A bound state comprised of either two or three valence quarks held together by the strong nuclear force.

HEC Hadronic End Cap.

Higgs boson The boson associated with the Brout-Englert-Higgs mechanism and the electroweak spontaneous symmetry breaking.

HL-LHC High Luminosity Large Hadron Collider.

HLT High-Level Trigger.

ID Indirect Detection.

IR Infra-Red.

ISR Initial State Radiation.

JER Jet Energy Resolution.

JES Jet Energy Scale.

JVF Jet Vertex Fraction.

JVT Jet Vertex Tagger.

LAr Liquid Argon.

LC Locally Calibrated scale, the scale at which the topo-clusters have been calibrated with a local calibration weight.

lepton Muon, electron, tau, or the the anti-particles thereof. Often taus and anti-taus will not be included when discussing leptons.

LHAPDF A package and convention for the storage and access of parton distribution functions.

LHC Large Hadron Collider.

LO Leading Order.

MC Monte Carlo.

MDT Monitored Drift Tube.

ME Matrix Element.

MIP Minimally Ionizing Particle.

Moriond2013 The jet calibration, uncertainties, and resolution released for 2012 data in January 2013, in time for the Moriond 2013 conference.

MSTW2008LO An old leading order parton distribution function provided by the MSTW collaboration.

MSTW2008NLO An old next to leading order parton distribution function provided by the MSTW collaboration.

MVA MultiVariate Analysis.

NCB Non-Collision Background.

NLO Next to Leading Order.

NNLL Next to Next to Leading Logarithm.

NNLO Next to Next to Leading Order.

NNPDF21LO An old leading order parton distribution function provided by the NNPDF collaboration.

NNPDF23NLO A recent next to leading order parton distribution function provided by the NNPDF collaboration.

parton Quark or gluon.

PDF Parton Distribution Function.

PDF4LHC A working group dedicated to proper parton distribution function treatment at the LHC.

pileup Objects or energy deposits from one collision which are treated as a part of a different collision, whether in the same bunch crossing or a different one, often in the context of calorimeter jets.

PMT PhotoMultiplier Tube.

QCD Quantum ChromoDynamics, the theory of the strong nuclear interactions.

RDO Raw Digital Object.

RF Radio Frequency.

RoI Region of Interest.

RPC Resistive Plate Chamber.

SCT SemiConductor Tracker.

SD Spin-Dependent.

SI Spin-Independent.

SM Standard Model.

SUSY SUperSYmmetry.

TB TeraByte.

TGC Thin Gap Chamber.

TRT Transition Radiation Tracker.

TTC Timing, Trigger, and Control.

UV Ultra-Violet.

VBF Vector Boson Fusion.

VEV Vacuum Expectation Value.

WIMP Weakly Interacting Massive Particle.

Appendices

Appendix A

Introduction to Cosmology

A.1 The standard model of cosmology

The *Big Bang*, in which the universe began to evolve from a highly compressed state, is a generally accepted theory of the formation of the universe. This event, which occurred 10^{10} years ago, predicts and matches observations for the CMB, thermal history, and large-scale structure of the universe [15]. Also known as the standard cosmological model or the Λ CDM model, this model is based on three fundamental concepts, the first of which is the relation between the matter and energy content of the universe with the geometry via the Einstein Equations [15]:

$$R_{\mu\nu} - \frac{1}{2}g_{\mu\nu}R = \frac{-8\pi G_N}{c^4}T_{\mu\nu} + \Lambda g_{\mu\nu} \quad (\text{A.1})$$

The left hand side of Equation A.1 defines the curvature of the universe, while the right hand side defines the energy and momentum. In this Equation, $R_{\mu\nu}$ is the Ricci curvature tensor, a contracted form of the Riemann curvature tensor. The Ricci tensor represents the degree of difference in the geometry between a given Riemannian metric and the metric of flat Euclidean space. Contracting this further produces the Ricci Scalar R , which represents the difference in volume between the given Riemannian metric and the Euclidean metric. The metric tensor of the manifold is then $g_{\mu\nu}$, which generalizes the definition of lengths and angles between tangent vectors. Newton's constant is the same as always, with $G_N = 6.673 \times 10^{-11} \text{ N(m/kg)}^2$, as is the speed of light c . The stress-energy-momentum tensor $T_{\mu\nu}$ defines the density and flux of energy and momentum, while Λ is the cosmological constant.

The cosmological constant is special, and is worth considering further. This constant specifies that, even in the absence of any energy or momentum associated with the matter content of the universe (setting $T_{\mu\nu} = 0$), some form of energy remains so long as $\Lambda \neq 0$. This is the so-called *vacuum energy*, which generates a gravitational field and creates curvature even in the absence of any matter. This was originally abandoned by Einstein when it was discovered that the universe was expanding, but has proven useful in explaining more recently observed phenomena, such as Type Ia supernovae and specific estimates of the CMB.

Once the Einstein Equation has been defined, the symmetries of the problem must be defined in order to solve it, forming the second fundamental piece of the puzzle. One typically assumes isotropy and homogeneity of the universe, thus greatly simplifying the equation. These assumptions are confirmed on the large-scale by several experiments, most notably the CMB for isotropy and and galactic surveys

for homogeneity. In fact, it appears that spheres with diameters on the scale of 100 Mpc or larger placed anywhere in the universe contain roughly the same matter content.

The application of these two symmetries to the Einstein Equation results in the following form for the line element of the metric [15]:

$$ds^2 = -c^2 dt^2 + a^2(t) \left(\frac{dr^2}{1 - kr^2} + r^2 d\Omega^2 \right) \quad (\text{A.2})$$

The spatial curvature k can be 0 or ± 1 , where the simplest case of 0 reduces to the simple flat Euclidean space, and $a(t)$ is the *scale factor*. Solving this equation leads to the Friedmann Equation:

$$\left(\frac{\dot{a}}{a} \right)^2 + \frac{k}{a^2} = \frac{8\pi G_N}{3} \rho_{\text{tot}} \quad (\text{A.3})$$

In Equation A.3, the ratio \dot{a}/a is the Hubble Parameter $H(t)$, and the Hubble constant is currently accepted to be $H_0 = 67.3 \pm 1.2 \text{ kms}^{-1}$ as measured by the Planck Collaboration in Reference [14]. The total average energy density of the universe is ρ_{tot} . It immediately follows that the universe is flat, or equivalently $k = 0$, for the critical density $\rho_{\text{tot}} = \rho_c$:

$$\rho_c = \frac{3H^2}{8\pi G_N} \quad (\text{A.4})$$

A standard cosmological convention is to quote the abundance of a given substance i in terms of the ratio with respect to the critical density, Ω_i , and the total density of the universe Ω as the sum over substances:

$$\Omega_i \equiv \frac{\rho_i}{\rho_c} \quad \Omega \equiv \sum_i \Omega_i \quad (\text{A.5})$$

The Friedmann Equation then takes the following form:

$$\Omega - 1 = \frac{k}{H^2(t) \cdot a^2(t)} \quad (\text{A.6})$$

The form of the Friedmann Equation in Equation A.6 emphasizes the dependence of the sign of the spatial curvature on the universal density. $\Omega = 1$ implies a flat universe, while $\Omega < 1$ is called *open* and $\Omega > 1$ is called *closed*. This Equation implies that Ω is time dependent, where the deviation from 1 can be thought of as the contribution of the curvature Ω_K to the energy density of the universe:

$$1 = \Omega + \frac{-k}{H^2(t) \cdot a^2(t)} = \Omega + \Omega_K \quad (\text{A.7})$$

Considering now the content at the start of the universe t_0 and splitting the total energy density into the matter (Baryonic and Dark), Radiation (photons), and cosmological sources results in:

$$\Omega(t_0) = \Omega_M(t_0) + \Omega_R(t_0) + \Omega_\Lambda(t_0) = \Omega_B(t_0) + \Omega_{\text{DM}}(t_0) + \Omega_R(t_0) + \Omega_\Lambda(t_0) \quad (\text{A.8})$$

Combining Equations A.7 and A.8 at t_0 gives:

$$1 = \Omega_M(t_0) + \Omega_R(t_0) + \Omega_\Lambda(t_0) + \Omega_K(t_0) \quad (\text{A.9})$$

Even if $\Omega_k(t_0) = 0$ and $\Omega(t_0) = \Omega_M(t_0) + \Omega_R(t_0) + \Omega_\Lambda(t_0) = 1$ as predicted by inflation, the different

types of matter content evolve with time differently [119].

Moving back to the Hubble Parameter $H(t) = \dot{a}/a$, the fact that the universe has been and continues to expand since the *Big Bang* means that the universal scale factor $a(t)$ as also been increasing with it. This was first confirmed by the observation of cosmological redshift from distant galaxies by the Hubble telescope, but further evidence has since followed. The formulation of this shift in terms of the emitted wavelength λ_e and the observed wavelength λ_o as related through the redshift parameter z as:

$$1 + z = \frac{\lambda_o}{\lambda_e} \quad (\text{A.10})$$

Thus a redshift of 0 occurs when the observed and emitted light have identical wavelengths, as expected, while values of z that are positive (negative) correspond to a redshift (blueshift). The redshift parameter can be related to the scale factor, as per Reference [120], through the times of observation t_o and emission t_e :

$$1 + z = \frac{a(t_o)}{a(t_e)} \quad (\text{A.11})$$

Now, note that the matter density Ω_M will be diluted as the universe expands and the same amount of matter is spread out over a volume increasing at the rate of $(1 + z)^3$. However, the cosmological term remains constant, per Equation A.1. As such, taking the ratio gives:

$$\frac{\Omega_\Lambda(t)}{\Omega_M(t)} = \frac{\Omega_\Lambda(t_o)}{(1 + z)^3 \Omega_M(t_o)} \quad (\text{A.12})$$

This demonstrates that even though the cosmological term may be small at the formation of the universe (large z), it dominates at longer timescales, where the matter content becomes dilute. Similar formulations for the radiation density result in a $(1 + z)^4$ dependence, as the wavelength of the photon is also red-shifted in addition to the volume dilution, thus confirming that radiation density becomes irrelevant faster than matter density. Furthermore, Equation A.8 shows the dependence of the curvature density on a^2 , and thus it follows from Equation A.11 that the ratio at a generic time to the start of the universe varies as $(1 + z)^2$.

In principle, there can be other sources contributing to the energy density. A generic extra source X with density Ω_X depends on the redshift through $(1 + z)^{3(1+\alpha_X)}$, where α_X relates the density ρ_X to the pressure p_X through $\rho_X = \alpha_X p_X$, as per Reference [119]. The cosmological contribution can be seen as a special case of this equation, with $\alpha_\Lambda = -1$. With this knowledge, a formula for the expansion rate of the universe can be derived:

$$\frac{H^2(z)}{H_0^2} = \Omega_X(1 + z)^{3(1+\alpha_X)} + \Omega_K(1 + z)^2 + \Omega_M(1 + z)^3 + \Omega_R(1 + z)^4 \quad (\text{A.13})$$

This is the third and final fundamental piece of the model. By studying the abundance of each of the different forms of energy at the present day, this equation allows for a determination of what happened in the past, approaching the *Big Bang* itself.

To summarize, the three fundamental pieces of the *Big Bang* cosmological model are [15]:

1. The Einstein Equation, providing a relation between the geometry and energy-momentum content of the universe, as per Equation A.1.

2. The symmetries of the problem, where homogeneity and isotropy of the universe at large distance scales is assumed.
3. The equation of state specifying the evolution of the universe, with dependencies on each of the individual types of energy densities, as per Equation A.13.

A.2 History of the universe

By extrapolating backward with the equation of state, a series of boundaries or otherwise critical points can be defined. A very short summary is provided below.

- $t = 0 \rightarrow 10^{-43}$ s: The Planck epoch, with an energy scale at the Planck mass of $M_{\text{Pl}} = 10^{19}$ GeV and an essentially infinite temperature. In this period, quantum gravitational interactions either dominated or were a part of a grand unified theory. At the end of this epoch, the gravitational force loses relevance, and the strong, weak, and electromagnetic forces were likely unified into a single theory.
- $t = 10^{-43} \rightarrow 10^{-36}$ s: The grand unification epoch, with an energy scale above 10^{15} GeV and a temperature above 10^{27} K. Mass, charge, flavour, and colour are irrelevant at this scale. At the end of this epoch, the strong force separates and the Standard Model of $SU(3)_c \times SU(2)_L \times U(1)_Y$ appears.
- $t = 10^{-36} \rightarrow 10^{-32}$ s: The inflationary epoch, likely marking a phase transition following the break-down of the grand unified theory into the strong and electroweak forces. The volume of the universe is believed to have grown by at least a factor of 10^{78} during this period. While the density of particles existing before the inflationary epoch began were enormously diluted by this growth, the break-down of the enormous potential energy of the inflationary field repopulated the expanded universe with a dense mixture of quarks, antiquarks, and gluons.
- $t = 10^{-32} \rightarrow 10^{-12}$ s: The electroweak epoch, with an energy scale above 10^2 GeV. The strong and electroweak forces are dominant. At the end of this epoch, the Standard Model undergoes electroweak symmetry breaking, leaving the strong and electromagnetic forces $SU(3)_c \times U(1)_q$.
- $t = 10^{-12} \rightarrow 10^{-6}$ s: The quark epoch, with an energy scale above 0.3 GeV. Quarks exist in isolation, as they are too energetic to be bound into mesons or baryons. At the end of this epoch, the quarks undergo confinement, and are bound into hadrons.
- $t = 10^{-6} \rightarrow 1$ s: The hadron epoch, with an energy scale above 1 MeV. Quarks have formed hadrons, and no longer exist in isolation. The hadrons and anti-hadrons exist in thermal equilibrium, which is broken at the end of the epoch.
- $t = 1 \rightarrow 10$ s: The lepton epoch, with an energy scale above 100 keV. The majority of hadrons and anti-hadrons have annihilated, while leptons and anti-leptons still exist in thermal equilibrium.
- $t = 10 \text{ s} \rightarrow 380 \times 10^3 \text{ y}$: The photon epoch, with an energy scale above 1 eV. The majority of leptons and anti-leptons have annihilated, while the universe is dominated by a plasma of free electrons, protons, and photons. The universe is effectively opaque, as the photons cannot

travel far before undergoing Thomson scattering off of the electrons. At the end of this epoch, it becomes energetically favourable to form atoms, thus recombination occurs. This sudden drop in free electrons causes photons to have an approximately infinite range, and is shortly followed by photon decoupling, the origin of the CMB.

- Current day, $t \approx 13.8 \times 10^9$ y: The energy of the universe has dropped to 2.7K, which is approximately 10^{-4} eV.

In addition to the major epochs listed above, there are also two other effects of importance which have the potential to span epochs:

- Variable time: The era of cold Dark Matter freeze-out, which can span many orders of magnitude depending on the nature of DM. Under the WIMP assumption, discussed in more detail in Section 3.2, this is at the scale of 1 GeV to 1 TeV.
- $t = 0.1 \rightarrow 10^3$ s: The era of Big Bang Nucleosynthesis, with an energy scale a bit below 100 keV up to 10 MeV. This regime is the period in which light elements other than hydrogen ^1H were produced. This includes deuterium, helium-3, helium-4, and lithium.

Appendix B

Review of Direct and Indirect Detection Experiments

B.1 Direct detection experiments

DD experiments rely on the scattering of existing DM particles off of target SM particles while traversing the sensitive volume of a detector. This is expected to be a very rare process lost in an enormous background created by numerous other processes. For this reason, DD experiments are typically built deep underground, thus providing substantial shielding from cosmic muons and other surface sources. This is usually only the start of the background rejection procedure, as other sources such as radioactive decays are expected to occur much more frequently than a genuine DM event. As such, all zero-background DD experiments make use of additional detector shielding, specially constructed materials which are made of minimally radioactive isotopes, and often employ additional veto detectors around the sensitive volume to further reject radiation activity or other backgrounds.

All DD experiments make the assumption that WIMPs are non-relativistic and are trapped within galaxies, thus accounting for the observed galactic rotation curves. The velocity of the DM particles is further assumed to be similar to the velocity of stars within the Milky Way, which is at the level of a few hundred kilometres per second in the local system. At these velocities, DM is clearly non-relativistic, and the the typical DM-SM scattering interaction is a nuclear recoil event. Given that the WIMP assumption is for DM particles with masses on the GeV to TeV scale, the energy transfer in such an event is at the level of $\mathcal{O}(1 \text{ keV})$ to $\mathcal{O}(100 \text{ keV})$. At these energies, the aforementioned assumption of the mediator coupling the dark sector to the SM is clearly applicable.

The probability of a DM-SM nuclear recoil event occurring is directly linked the the number of nucleons with which the DM can interact and the local flux of DM. The number of nucleons in the detector is given by the mass of the detector divided by the atomic mass of the particle being used as the sensitive volume, while the DM flux requires further assumptions. In particular, it depends on:

- The local density of DM, ρ_{DM} , typically taken to be $0.3 \text{ GeV}/\text{cm}^3$
- The DM circular velocity, v_c , typically taken to be 220 km/s
- The galactic escape velocity, v_{esc} , typically taken to be 544 km/s

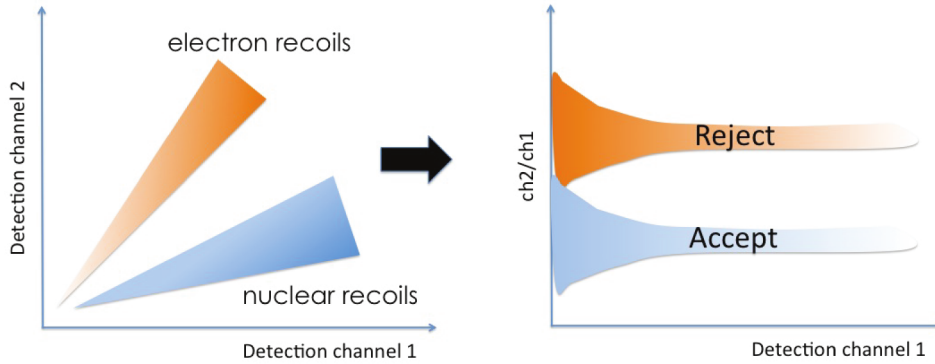


Figure B.1: A schematic representation of the added power that comes with having two independent measurement channels in nuclear recoil detectors. While both electron and nuclear recoils are seen to have similar shapes in a single dimension, the response in the 2D plane is clearly demarcated (left), thus allowing for a simple rejection criterion to be defined (right). [121]

- The DM mass, m_χ , typically used as a free parameter and varied between $\mathcal{O}(\text{GeV})$ and $\mathcal{O}(\text{TeV})$
- The velocity-averaged interaction cross-section between the DM particle χ and the SM nucleon N , $\langle\sigma_{\chi N}\rangle$, an unknown and the variable constrained by experiment measurements

Given these assumptions, the frequency of nuclear scattering events from DM particles passing through a SM detector depends on two parameters, the mass of the DM and the WIMP-nucleon interaction cross-section. For this reason, results are typically set on the cross-section parametrized in terms of m_χ .

Results in DD experiments are typically strongest when the WIMP mass is close to the target mass. If the DM particle is too light, then it only has a negligible impact on the target nucleus, and the resulting nuclear recoil energy is below detector thresholds. If the DM particle is heavier then the nuclear recoil is more apparent, but the WIMP density flux scales with m_χ^{-1} , thus such events are also suppressed.

There are in principle three methods to identify such low energy nuclear recoil events, using ionization, scintillation, or phonons. Many experiments design detectors which are able to make use of two of these measurements simultaneously. This proves very beneficial, as while background events and signal events may behave similarly in any given channel, usually the correlation between two channels works to significantly enhance the separation between signal and background. As such, having two independent measurements of every event significantly enhances the separation power. This notion is summarized schematically in Figure B.1.

While the idea of nuclear recoil occurring between WIMPs and nucleons is clear, there are multiple ways in which this can occur. The first method is known as SI interactions, where the cross-section depends only on the number of nucleons, and thus the particle mass. This motivates the use of massive target particles, such as the common choices of germanium and xenon. The alternate possibility is the SD interaction, in which the coupling between the WIMP and target particle depends on the spin. For SD interactions, there is no motivation to use massive particles, rather elements which are highly polarized (such as fluorine) are used. These two types of interactions can be thought of in terms of

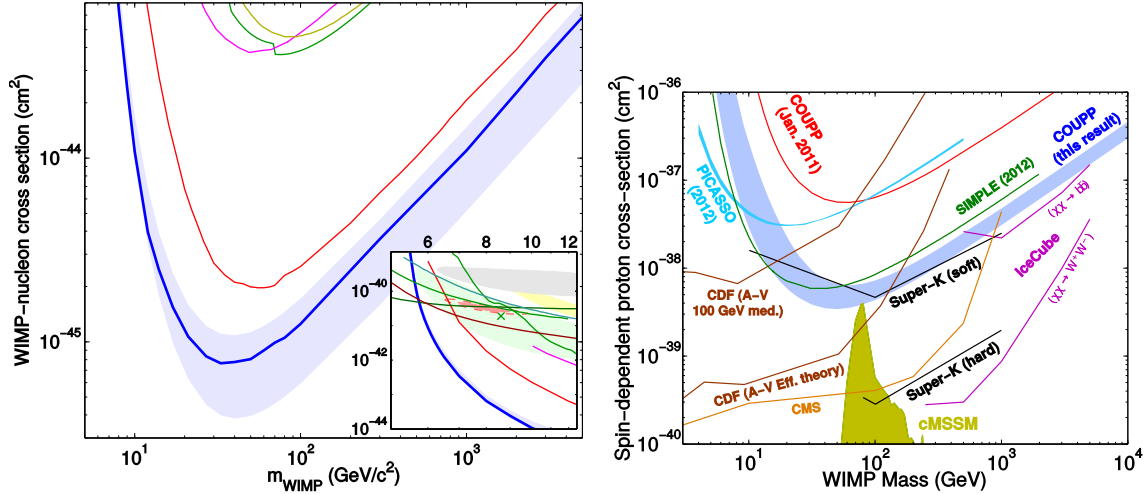


Figure B.2: The Spin-Independent (left, [104]) and Spin-Dependent (right, [106]) WIMP-nucleon cross-section observations and exclusion limits for various DD collider experiments. Left: LUX is the blue result which is the most sensitive at the time of plot creation, while Xenon100 is the red result, and the grey region is the DAMA/LIBRA observation. Further details on the many experiments are provided in Section 8.2.4. Right: the various experiments are labelled, and limits from colliders are shown alongside the DD results.

the difference between vector (SI) and axial-vector (SD) couplings, although other possible high-energy couplings also exist.

The DD search community is very active, with multiple competing experiments active at any given time. Furthermore, there are many observation claims by some experiments which are in contradiction with exclusions from other experiments. Given the different target materials which are typically used, it is possible that DM couples preferentially to some types of materials rather than others, but the viability of such scenarios is increasingly stretched as experimental sensitivity continues to improve. Sample plots for both the SI and SD interaction possibilities are shown in Figure B.2.

A selection of relevant zero-background DD experiments are:

- Spin-Independent:
 - CDMS-II [99]: Ionization and phonon channels, germanium and silicon detectors
 - CRESST-II [100]: Scintillation and phonon channels, CaWO₄ crystal detectors
 - CoGeNT [101, 102]: Ionization channel, germanium detectors
 - LUX [104]: Photon and ionization channels, liquid xenon detector
 - Xenon100 [105]: Photon and ionization channels, liquid xenon detector
- Spin-Dependent:
 - COUPP [106]: Acoustic pulse channel, CF₃I bubble chamber
 - IceCube [107]: Neutrino detector
 - PICASSO[108]: Acoustic pulse channel, C₄F₁₀ bubble chamber
 - SIMPLE [109]: Acoustic pulse channel, C₂ClF₅ bubble chamber




Figure B.3: Annual modulation of the number of observed events at DAMA/NaI and DAMA/LIBRA peaking in late May, as expected from the movement of the Earth into and out of the *WIMP wind*, with a significance of 8.2σ . [4]

– Super-K [110]: Neutrino detector

One experiment is worth specifically mentioning, due to its different construction and long-standing discovery claim. The DAMA/NaI and DAMA/LIBRA detectors are not zero-background experiments. Instead, they rely on a property expected of the DM distribution in the galaxy, namely that it should continue to point in one direction, while the earth moves through this field of DM particles. As such, when the Earth is moving into this *WIMP wind*, more interactions occur. The opposite is then true when the Earth is moving in the same direction of the *WIMP wind*. This results in the prediction of annual modulation, where the number of DM interactions is an oscillation on top of an otherwise flat background of events from cosmic muons, radioactive decays, and so on.

Both the DAMA/NaI and DAMA/LIBRA detectors have observed this modulation with a maximum near the end of May as expected and a significance of 8.2σ [4]. The observed modulation is shown in Figure B.3. This result has been excluded by numerous zero-background experiments to very high confidence levels, but the authors maintain that the difference is due to the type of experiment. Ideally an identical experiment would be conducted, preferably on the other side of the world to avoid local effects. This is currently under development at the South Pole as part of the DM-ICE effort, which is ideal given the location of DAMA/NaI and DAMA/LIBRA in Italy.

B.2 Indirect detection experiments

ID experiments rely on the annihilation of existing DM and watch for the resulting decay products. This process is limited by the annihilation rate, which is proportional to the dark matter density. It is thus natural to first consider regions of strong gravitation where DM is expected to gather in large quantities, such as the centre of galaxies, within stars, or even within the centre of the Earth. Dwarf spheroidal galaxies, where the dark matter content is observed to be significantly larger than the visible matter content, provide another excellent location for searches. The centre of galaxy clusters can also be considered, and is a prime candidate for DM annihilation, but these regions also suffer from significant modelling assumptions to the extent that they are typically avoided.

The decay products can take different forms, but the most common searches are for anti-particle




Figure B.4: The positron fraction as measured by the AMS, Fermi-LAT, and PAMELA experiments. [125]

flux excesses and mono-energetic lines in the photon spectrum. The former suffers from uncertainties on the propagation of the anti-particles through space, where modelling of their effects under the variable magnetic fields is not optimal, and has many astrophysical backgrounds such as pulsars. The latter is a particularly tantalizing signature, as there is no known astrophysical source which can predict mono-energetic photon lines at the WIMP scale, making this a *smoking gun* signature for DM [122].

The AMS and Fermi-LAT Collaborations have both recently presented results on anti-particle flux excesses. The AMS results for the positron fraction are the most precise yet and cover the largest energy range, as seen in Figure B.4. The observed excess of events could be compatible with DM, but there are several limitations of this interpretation. In particular, References [123, 124] discuss the difficulty in fitting DM distributions to the data, given the magnitude of the excess and the required annihilation cross-section. Furthermore, given that positron flux is enhanced but the anti-proton flux is not, any DM interpretation fitting the data must also be leptophilic. It is instead likely that the observed excess is due to nearby astrophysical sources, such as pulsars, but no conclusions can be drawn at this time.

While the AMS results may be the most precise to date, the Fermi-LAT Collaboration has results for many different annihilation channels taken from a survey of 15 dwarf spheroidal galaxies. The results are in Figure B.5, where no significant excess is observed and thus bounds are placed on the six different final states.

The Fermi-LAT Collaboration has also been the subject of discussion over a possible mono-energetic photon line, first mentioned in Reference [126]. The significance of the excess grew as the dataset size increased, and the open data policy allowed numerous individual groups to explore the phenomenon, but has since reduced once the Collaboration released updated calibrations, systematics, and its own analysis on the subject [127].

The AMS and Fermi-LAT Collaborations are both space-based experiments, but there is another class of important ID experiments. Large telescope arrays, such as HESS, can be used to constrain the DM annihilation rate through measurements of the photon spectrum. The results from HESS are shown in Figure B.6. The different properties of these types of experiments show that HESS is more sensitive than Fermi-LAT for high DM masses, while Fermi-LAT is superior for lower masses.




Figure B.5: 95% confidence level constraints on the WIMP-WIMP annihilation cross-section from the analysis of 15 dwarf spheroidal galaxies in several channels by the Fermi-LAT Collaboration. [114]




Figure B.6: 95% confidence level constraints on the WIMP-WIMP annihilation cross-section from photons from the galactic centre halo by the HESS Collaboration. [115]

Appendix C

Jets and Masked Tile Calorimeter Regions

This Appendix is based on work done in conjunction with Caterina Doglioni.

C.1 Introduction

During data taking, it is not uncommon for calorimeter module power supplies to temporarily fail. When this happens, the charge deposited by particles passing through the calorimeter is not properly collected until the power supply is once again fully operational. Due to the frequency with which these temporary failures occur, and the small fraction of the detector which is affected, these failures are not enough to mark an event as bad within the ATLAS GRL framework.

Occasionally, there are more serious problems which result in a full calorimeter module being lost for a longer portion of time, typically until the end of the data taking period due to the need for manual intervention. This is very rare for modules of the LAr calorimeters to fail in such a fashion, thus when it does occur, the MC is typically updated to include the dead region. This ensures that data and MC represent the same detector, and thus the problem is resolved to first order. As such, failures of the LAr calorimetry will not be considered further in this Section.

Failures of the Tile calorimeter high voltage power supplies were more common, with five modules rendered unusable on a permanent or long-duration basis during the 2012 data taking period. Only the one module which was dead for the entire 2012 dataset was added to MC, while the non-inclusion of the other failures introduced a significant source of possible bias. The list of modules which were individually masked for a significant period of time are listed in Table C.1.

In terms of the Tile calorimeter coverage, the permanently masked modules represent a small fraction of the total calorimeter, as can be seen from Figure C.1. Additionally, the transient trips may occur frequently, but again they represent a small region of the calorimeter when they are masked. It is thus tempting to assume that it is a negligible effect. However, the actual impact is much larger than the size of the modules. Jets have a finite size, where the most common use is $R = 0.4$, but which can also be much larger depending on the analysis requirements. As such, despite the masked modules representing a small fraction of the actual calorimeter, the impact on the fraction of jets which will contain some part of these regions is much larger. It is thus important to study what happens in situations where jets fall

Tile module information			Module failure		Masked 2012 pp data	
Name	η range	ϕ range	Type	In MC?	$\int \mathcal{L} dt$	Fraction
LBA05	[0, 1]	$[4\pi/32, 5\pi/32]$	Permanent	Yes	20.28 fb^{-1}	100%
LBC16	[-1, 0]	$[15\pi/32, 16\pi/32]$	Permanent	No	16.72 fb^{-1}	82.4%
LBA18	[0, 1]	$[17\pi/32, 18\pi/32]$	Permanent	No	7.25 fb^{-1}	35.7%
LBA35	[0, 1]	$[34\pi/32, 35\pi/32]$	Permanent	No	5.95 fb^{-1}	29.3%
EBC01	[-1.6, -0.8]	$[0, \pi/32]$	Long-duration	No	0.82 fb^{-1}	4.0%
LBC18	[-1, 0]	$[17\pi/32, 18\pi/32]$	Transient	No	0.11 fb^{-1}	0.05%
LBA38	[0, 1]	$[37\pi/32, 38\pi/32]$	Transient	No	0.10 fb^{-1}	0.05%
LBA26	[0, 1]	$[25\pi/32, 26\pi/32]$	Permanent	No	0 fb^{-1}	0%
LBA63	[0, 1]	$[62\pi/32, 63\pi/32]$	Permanent	No	0 fb^{-1}	0%

Table C.1: The list of tile modules which were masked for at least 0.05% of the pp collision dataset, the amount of pp integrated luminosity $\int \mathcal{L} dt$ that it was masked for, and the fraction of the dataset that this represents. Two modules which were permanently masked after the 2012 pp dataset but before the heavy ion collision dataset are listed at the bottom for completeness.




Figure C.1: Masked Tile calorimeter cells after the end of the 2012 dataset (left, February 2013) and after replacing all of the high voltage power supplies in preparation for Run-II (right, September 2014). The masked regions in the left plot are explained in Table C.1. The new masked module in the right plot is due to a TTC card problem rather than a high voltage power supply failure. It is not known if it will be operational before Run-II collisions begin. [128]

near masked modules.

C.2 The impact of masked modules on jet kinematics

Masked regions of the calorimeter are inherently unreliable, thus the energy value for all cells that are masked is typically set to zero. This ensures that the cell cannot become a seed for topo-cluster formation. This is particularly important for hot cells, where the cell is read to always have the maximum energy, and thus would always seed the creation of a very high energy topo-cluster. For cells within a module which has been masked due to a bad power supply or similar, the behaviour is much more erratic, and thus it is still safer to set the energy to zero.

While this is one solution to deal with masked cells, it is not the last step in the ATLAS reconstruction scheme. Jets which contain masked cells where the energy is set to zero will also be biased, as this ignores the energy which would be actually deposited in this region. As such, an average cell correction has

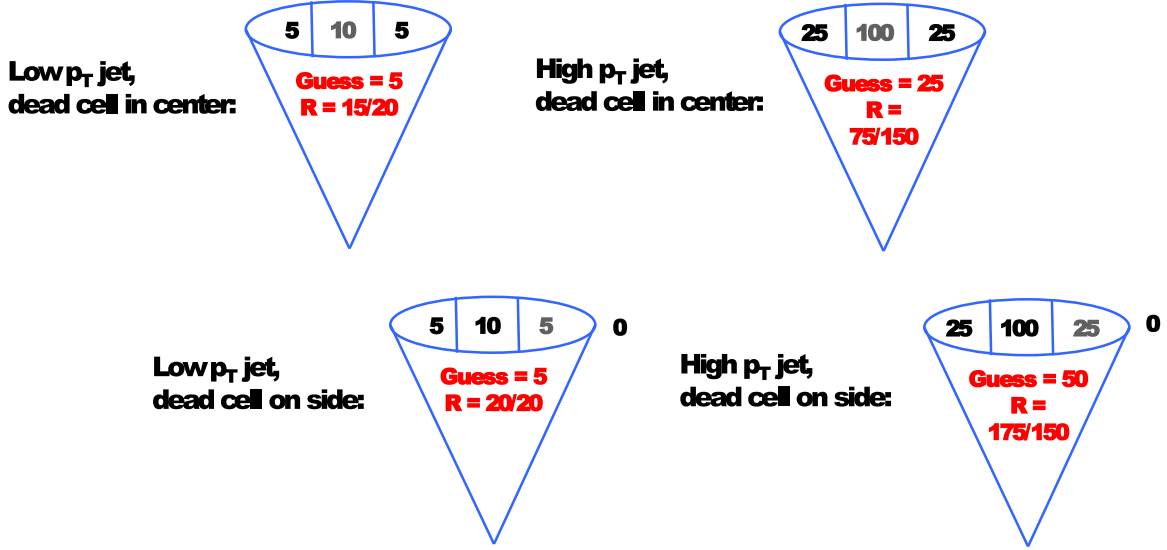


Figure C.2: A schematic representation of the impact of the BCH correction in which the energy of a masked cell is estimated from the average of the neighbours. The impact is split by low p_T (left) vs high p_T (right), and by masked cells in the core of a jet (top) vs masked cells on the edge of a jet (bottom). Values in black are measured values in good cells, while regions in grey are the true value which a masked cell contains. Values in red then provide the value which the BCH correction uses, and the ratio of the estimated energy to the true energy. All numbers shown have arbitrary units, but can be thought of as using GeV. [129]

been used, where the average energy of the neighbouring cells is used. In other words, if the cell has four neighbours with energies in arbitrary units of 10, 15, 15, and 20, then the cell would be filled with an energy value of 15 before topo-clusters are built. This cell energy estimate is known as the Bad CHannel (BCH) correction, and the scale factor when applied to the full jet is known as `BCH_CORR_CELL`, which is simply the sum of energy from masked cells divided by the total energy in the jet (at EM scale).

The BCH correction will never create a new topo-cluster, as the average ensures that at least one of the neighbours contains at least as much energy if not more. As such, if the masked cell can be classified as a seed for topo-cluster formation, then the neighbour is also a seed. In addition to this, built-in software restrictions explicitly forbid the use of masked cells as topo-cluster seeds.

While this average neighbour energy scheme is reasonable as a first estimate of the energy in the cell, and it works well for low p_T jets with only a single masked cell, it actually exacerbates the problem at high p_T or with large masked regions. In order to explain this, Figure C.2 provides a schematic representation of a simple jet where one cell is masked. This demonstrates how masked cells within the core of a jet lead to under-estimations, where the degree of mis-estimation grows as the energy is increased due to the jet becoming more collimated. While this is an undesired effect, it will not destroy the average jet response, as jets are being shifted from high energies (where there are fewer jets) to low energies (where there are many more jets).

The same Figure demonstrates how masked cells on the edge of a jet lead to over-estimations, which again get worse as the energy is increased. This is much more dangerous, as a jet population with large statistics (lower p_T) is being migrated to another population with much less statistics (higher p_T), thus providing a way to shift the mean of the entire jet response distribution. This problem did not exist before the BCH correction, which is why the addition of this correction rather than just setting the cell

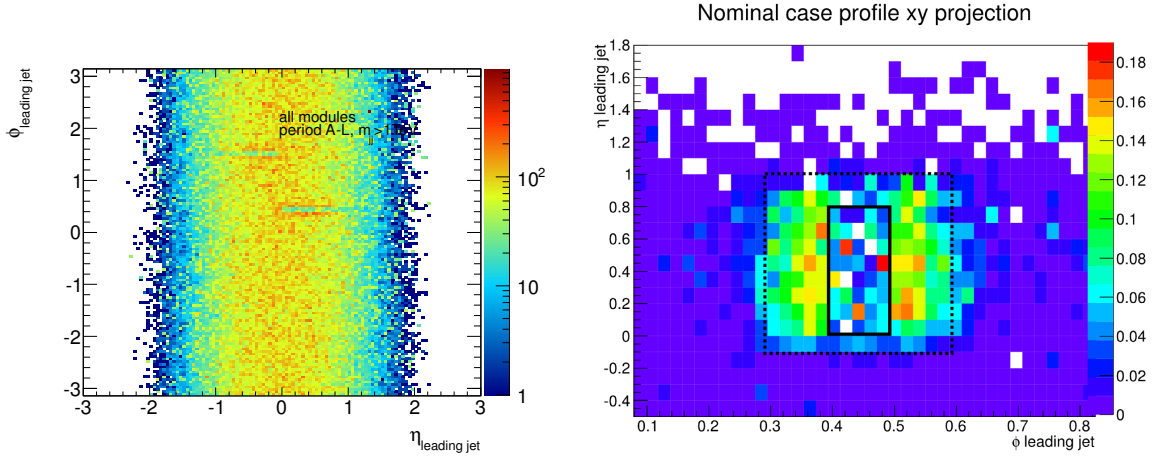


Figure C.3: Left: the Tile calorimeter jet occupancy in data for a di-jet selection. This high p_T scenario is a perfect example of where BCH effects are expected to be significant. The two masked modules which were disabled for the longest period of time (LBA05 and LBC16, see Table C.1 for details) are clearly visible in the data plot, with the expected excess of jets adjacent to the masked modules and deficit of jets within the masked modules. Right: the average BCH correction value in MC around the one masked module added to simulation, for the di-jet topology. The geometrical extent of the module is drawn in solid black lines, while a border one cell wide in all directions is drawn in dashed lines. This extended region of the masked module and a surrounding region of cells is necessary to cover the majority of the region which receives a BCH correction.

energy to zero has the potential to make the situation worse.

These schematic examples were provided to explain the effect of the BCH correction on jets with a single masked cell. Entire masked modules serve to amplify the effect, as a module is a series of sequential cells. Given that a single Tile calorimeter cell is approximately 0.1×0.1 in (η, ϕ) , and typical jets use distance parameters of $R = 0.4$ or larger, it is likely that a jet containing or adjacent to one masked cell will also contain or be adjacent to other masked cell(s) from the same module. Only boundary cases can avoid this geometrical restriction.

The jet p_T is not the only variable affected by this correction. The centroid of the jet is calculated as an energy-weighted sum of the topo-clusters, and the topo-cluster centroid is calculated as the energy-weighted sum of the cells. If masked cells are under-corrected, then the centroid can be shifted outside of the cell. In the case of a jet and a masked module, jets with a truth centroid in the masked module can be shifted to be adjacent to the masked module. This shift serves to reduce the fraction of jets pointing directly at masked modules, and enhance the population of jets pointing adjacent to them. Furthermore, jets which are adjacent to masked modules receive an over-correction. This increase also leads to more jets pointing to regions adjacent to masked modules for a given p_T cut.

In other words, the BCH correction is expected to lead to a deficit of jets within masked modules, and an excess of jets adjacent to masked modules. This is confirmed by looking at the jet occupancy distribution within the Tile calorimeter, as shown in Figure C.3. A region encapsulating the masked module plus a one cell boundary is shown to be required to account for cells which have an irregular occupancy.

Another way to quantify the magnitude of the effect is to consider the jet p_T and m_{jj} response. The JES has already been applied, thus the average response should be at approximately 1. Even if a bias




Figure C.4: The jet response in the di-jet topology as a function of the BCH correction factor BCH_CORR_CELL , in MC, for the p_T (left) and dijet mass m_{jj} (right).

remained in the scale and it was offset from 1, then plotting against the value of the BCH correction should produce a flat line. Instead, a strong trend is observed, as shown in Figure C.4.

This trend provides further proof for the previously explained behaviour where the correction has the potential to both under and over-estimate the true p_T . Two clear trends are visible, representing jets with their centroid in a masked module and jets adjacent to masked modules. These two populations have very different kinematic behaviour, and thus must be dealt with separately.

C.3 The impact of masked modules on other quantities

By manipulating jet quantities at a fundamental level, the BCH correction also biases other distributions. The most obvious example is the E_T^{miss} distribution, where a jet pointing to a masked module loses a significant portion of its energy. This creates an imbalance in the system, and generates a source of fake E_T^{miss} . It is not uncommon to see 20% under-corrections, thus reasonably energetic jets can provide large fake E_T^{miss} sources. Alternatively, the jet can be adjacent to a masked module, and thus get a large over-correction, once again providing a large fake E_T^{miss} source. In the first scenario, the E_T^{miss} will be aligned with the jet, while in the second scenario the E_T^{miss} will be back-to-back with the jet.

This is a particularly important point for analyses which rely purely on high p_T jet(s) and E_T^{miss} , such as the mono-jet analysis which will be detailed in Chapter 7. With only jet and E_T^{miss} related cuts, this provides a way for the enormous QCD background to pass a simple analysis selection. Typical cuts to ensure that the E_T^{miss} and jets are not aligned will reject the first scenario, but not the second. This will be discussed further in Section 7.4.6.

The impact of the BCH threshold on E_T^{miss} or other event-level balance variables is not surprising, but the dependence on other event-level variables is often overlooked. In particular, many analyses which do not apply jet kinematic requirements, and only work with the jet multiplicity N_{jet} . This may be for selecting or rejecting events with a given number of jets, or for binning distributions as a function of the number of jets. Either way, it is a common misconception that these variables are independent of the BCH problem because topo-clusters cannot be seeded by masked cells and because they are not directly using the kinematics.

While groups may not explicitly apply selection cuts on their jet objects, it is easy to forget that an implicit dependence does exist. More specifically, the object definition used to define what is a jet and

what is not always makes use of a p_T cut. This is mandated when using any of the common ATLAS jet techniques, as cleaning, calibrations, uncertainties, and resolution can only be evaluated down to a certain level before the measurements are dominated by pileup or until they are below reconstruction thresholds. These object definition cuts can be sensitive to the BCH correction, and thus analyses making use of the jet multiplicity must also take this into account.

At an even more fundamental level, it is easy to forget that this is not actually a jet correction. The correction is applied before the topo-clusters are built, thus the topo-clusters (and even the cells if that level of detail is available) are affected. This means that analyses which never touch a single jet or E_T^{miss} related variable, but which do use calorimeter-based isolated leptons, are sensitive to BCH effects. This should be a much less significant bias given that isolation thresholds are usually very low, and thus the mere presence of a hadronic topo-cluster is usually enough to qualify the lepton as non-isolated, but the principle remains.

In summary, it is important to remember that the BCH correction has the potential to bias all topo-clusters which include cells in the Tile calorimeter, no matter how they are used. That can be jet-related kinematics, E_T^{miss} or other event balance quantities, event-level jet variables like the jet multiplicity, or even calorimeter-based isolation of leptons. There are very few analyses in ATLAS where no such information is used. This emphasizes the importance of deriving a method for accounting for the possibility of bias accompanying the inclusion of the BCH correction in the ATLAS reconstruction framework.

C.4 Cleaning bad jets

Given the number of different ways that the BCH correction can introduce a bias, it was clear that some means of correcting for the correction must be derived. If the MC included these effects, then the first order fix would be to check if the data and MC modelling was consistent, in which case the problem could be reduced and more possibilities for corrections would be available. However, this is not the case, so a scheme for removing these bad jets must be determined. There are two primary schools of thought on this topic:

1. All jets pointing to masked modules in the Tile calorimeter are bad. Events where selected jets point to these regions should be vetoed.
2. Throwing away all jets pointing to bad regions is an over-reaction. A set of cleaning criteria which remove the worst events and retain the good events should be used.

Note that in both cases, the jets which should be considered are the ones passing the analysis object definition of a jet. These criteria are not meant to be applied on pileup, but rather on analysis-selected jets.

Both of these methods have their advantages and limitations. The first method entails a geometric cut around the masked modules. This is simple to implement, and provides an absolute method for removing the potential bias. The difference between data and MC becomes easy to emulate, as the fraction of events for which a given module was disabled in data was known. The use of a random number generator to determine when a given module is “masked” in MC, and thus to veto the event if a jet is in or adjacent to the module, provides a means of reproducing the same effect in both data and MC. The application of this geometric cleaning cut to the di-jet topology on the key kinematic




Figure C.5: The jet response in the di-jet topology as a function of the BCH correction factor `BCH_CORR_CELL`, in MC, for the p_T (left) and dijet mass m_{jj} (right) after the application of geometric cleaning. A comparison to the pre-cleaning plots in Figure C.4 shows significant improvement.

variables is shown in Figure C.5, which shows a significant improvement over the pre-cleaned variables in Figure C.4.

However, the geometric cleaning procedure can be very restrictive, and the available integrated luminosity after such a cut can suffer significantly. Topologies with a small jet multiplicity may lose only a fraction of their events, but those with a high jet multiplicity will observe a much higher data loss. The loss of 10-15% of data in the di-jet resonance analysis was deemed acceptable given their sensitivity to BCH effects, while the loss of $\mathcal{O}(30\%)$ of data in the fully hadronic $t\bar{t}$ cross-section measurement was not acceptable. In conclusion, the geometric cut can be very powerful for low jet multiplicity analyses, but it is not suitable for everyone.

In devising a set of cleaning criteria, several variables were studied, including both the individual use and combinations of the following:

1. Angular distance between the jet and E_T^{miss} , $\Delta\phi(E_T^{\text{miss}}, \text{jets})$: works well for rejecting jets where the jet centroid falls in the core of a masked module, but can be very topology dependent
2. BCH correction factor, `BCH_CORR_CELL`: works very well, part of the final cut
3. Charge fraction, $f_{\text{ch}} = \sum p_T^{\text{track}} / p_T^{\text{reco}}$: works reasonably well, but not as powerful as f_{em}
4. Electromagnetic barrel energy (raw energy), $E_{\text{calo}}^{\text{EM}} = E_{\text{calo}}^{\text{EM1}} + E_{\text{calo}}^{\text{EM2}} + E_{\text{calo}}^{\text{EM3}}$: not as effective as f_{em}
5. Electromagnetic fraction, $f_{\text{em}} = E_{\text{calo}}^{\text{EM}} / E_{\text{calo}}^{\text{EM+Had}}$: works very well, part of the final cut
6. Maximum energy fraction per calorimeter layer, $f_{\text{max}} = \max \left\{ E_{\text{calo}}^{\text{layer}} / \sum E_{\text{calo}}^{\text{layer}} \right\}$: good, but not as powerful as f_{em}
7. Tile barrel energy (raw energy), $E_{\text{calo}}^{\text{Tile}} = E_{\text{calo}}^{\text{Tile0}} + E_{\text{calo}}^{\text{Tile1}} + E_{\text{calo}}^{\text{Tile2}}$: not as effective as f_{em}

C.4.1 Cleaning cuts for jets adjacent to masked tile modules

The most powerful cut for jets with their centroid adjacent to a masked module was determined to be a combination of the jet electromagnetic fraction and the BCH correction factor. This combination was found to provide significant rejection of jets where the jet response was significantly biased, while

retaining the large majority of jets for which the response was approximately 1, exactly as desired. The specific cut used to reject bad events is listed in Equation C.1. The application of this selection criteria in MC, comparing the one masked module to a separate good module, is shown in Figure C.6.

$$\left(\text{BCH_CORR_CELL} > \frac{2 - f_{\text{em}}}{10} \right) \text{ OR } \left(\text{BCH_CORR_CELL} < 1 - \frac{5}{2} f_{\text{em}} \right) \quad (\text{C.1})$$

These plots are a selection of the many p_{T} bins considered, including both the lowest ($20 \text{ GeV} < p_{\text{T}} < 30 \text{ GeV}$) and the highest ($1500 \text{ GeV} < p_{\text{T}} < 3000 \text{ GeV}$), as well as two intermediary bins where the impact of the BCH correction becomes apparent. The difference between the known masked module and known good module makes it clear that the distributions are significantly different, thus confirming that the behaviour we are devising cuts for is due to the fact that the module is masked. The cut is also shown to remove only a small fraction of the jets at low p_{T} , where it is clear that the JER is larger than any resolution effect coming from the BCH correction. As the p_{T} is increased, a region with a significantly over-corrected response begins to form. The structure starts to appear past 50 GeV, and becomes much more defined for $p_{\text{T}} > 100 \text{ GeV}$, where the correction is observed to remove a jet population which is on average 10% or more over-corrected.

By the highest p_{T} bin considered, the effect of jet collimation has made the difference much more pronounced, with a large population of jets which are being over-estimated by a factor of 50%. This an enormous effect, as the number of jets with p_{T} of 1500 GeV is significantly higher than the number of jets with a p_{T} of $1500 \times 1.5 = 2250 \text{ GeV}$. These migrations, if not accounted for, can cause bumps to occur at high jet p_{T} . The derived cut is still removing the worst parts, although there are many jets with mis-measurements on the level of $\pm 10\%$ which remain, and it is hoped that these fluctuations cancel out on average.

The highest p_{T} bin also reveals the presence of a second distribution with a lower BCH_CORR_CELL value, which gradually emerged as the p_{T} was increased. This distribution is from jets where the truth jet was in the middle of the masked module, but the BCH correction shifted the reconstructed centroid into the adjacent module, making it appear in this plot. In addition to tending to have a mis-estimated response, these jets also worsen the angular resolution, and the cut removes the worst part of them.

C.4.2 Cleaning cuts for jets centred within masked tile modules

All of the above discussion was for jets with a reconstructed centroid adjacent to a masked module. Note that the masked module itself is 10 cells in η by 1 cell in ϕ , for a total of 10 cells. On the other hand, the edge region is 12 cells in η by 3 cells in ϕ , minus the 10 cells which are the core of the masked module, leaving 26 cells. As such, purely by geometrical reasons, the fraction of jets adjacent to a masked module is larger than the fraction of jets within a masked module. Add to this the increased occupancy due to increasing the energy for adjacent jets vs reducing the energy for core jets, and the shift of the jet centroid due to energy being weighted towards the edge of a module, and the number of jets adjacent to a masked module is significantly larger than the number of jets in the module itself.

This is good news, as it turns out that no combination of variables considered was suitable for isolating good jets in the core of a masked module. This is demonstrated in Figure C.7, where the effect of BCH is at first irrelevant due to the dominant JER, but then by 50 GeV the average jet is approaching the point of a 10% under-correction. By 100 GeV it is clear that the large majority of the parameter space is at least 10% under-corrected, and by the high p_{T} ranges, the vast majority of jets

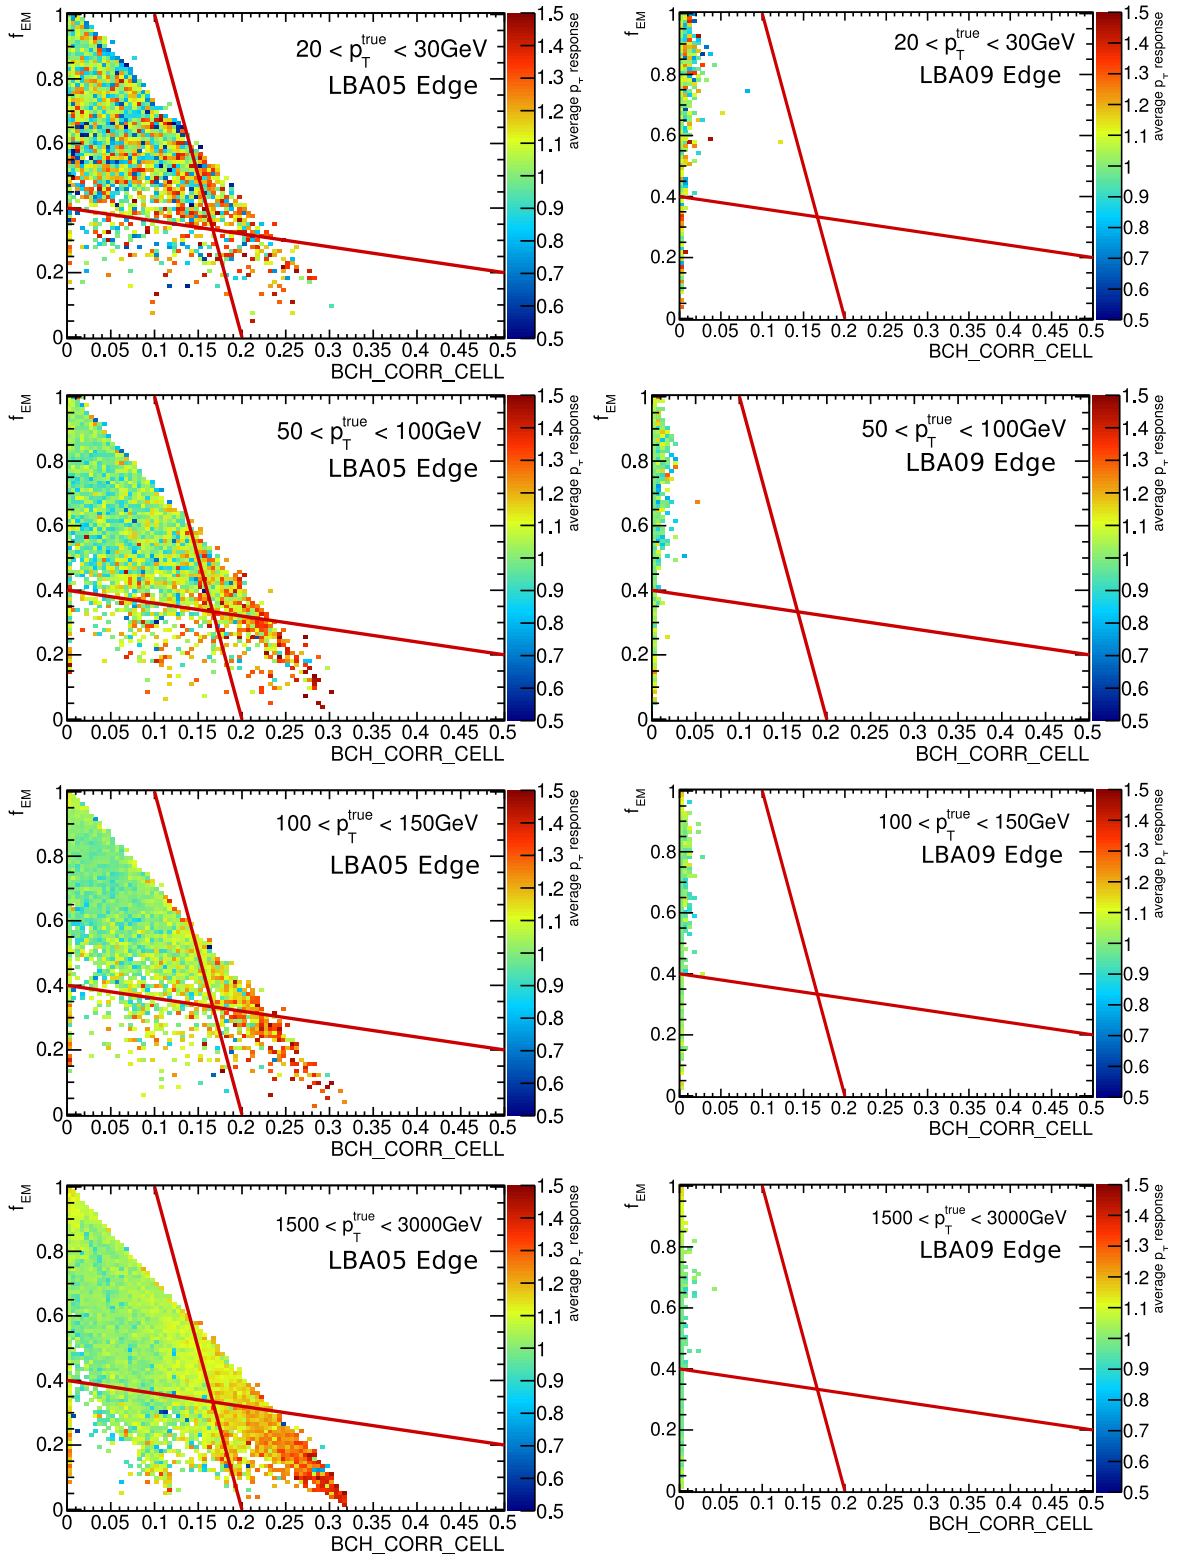


Figure C.6: A sample of MC distributions for various $p_T^{\text{jet, true}}$ of the jet electromagnetic fraction vs the BCH correction factor, for jets adjacent to a module which is masked (left) or fully operational (right). Red lines show the cut values which were selected. Only the top left region of each plot is retained.

are under-corrected by 20 to 50%.

The use of a $\Delta\phi(E_T^{\text{miss}}, \text{jets})$ cut showed some promise, but introduced too much of a topology dependence, as the cut acts very different in events where there is a real source of E_T^{miss} vs where all E_T^{miss} is from fakes. For this reason, it was decided that a geometric cut was the only safe solution.

C.4.3 Cleaning cut limitations

If the dead Tile module distribution in MC simulation mirrored the data, then this set of cleaning criteria would be both simple to use and very powerful. However, this is not the case. This provides a problem, as the cleaning criteria cannot be blindly applied on MC. Even though it is simple to emulate the probability of a given module being masked in MC based off of the fraction of luminosity the same module was masked for in data, this does not change the values of f_{em} and `BCH_CORR_CELL` which the cut is based on. Thus, while this emulation will specify that the module should be considered as masked for a given event, the variables which the user has access to were simulated under the assumption that it was not.

In order to work around this, the one dead module in MC was used to derive an extra layer of emulation. Rather than try to provide what should be a reasonable f_{em} and `BCH_CORR_CELL` distribution for a given jet p_T , η , and ϕ , it was decided that a simpler and safer method was to calculate the fraction of jets which fail the cleaning criteria as a function of p_T . Then, when a jet of a given p_T is marked as falling in an emulated MC masked module, a further random number is used to determine whether or not the event should be rejected.

In the calculation of this probability was found to have a dependence on the flavour of the jet, which is not surprising. Gluons and light quarks have different radiation patterns, and thus the value of f_{em} is expected to be different. In order to quantify the associated uncertainty on this cleaning procedure, a systematic was determined. The nominal value is the probability if all partons are used, while the systematics bracket this with one corresponding to only light quarks and the other only gluons. The resulting probabilities for the rejection of a jet of a given p_T , which falls into a region which the emulation has decided is masked, is provided in Figure C.8.

The cleaning cut probabilities were derived in the di-jet topology, but it is believed to be generally applicable to jets of all types. Users are allowed to determine this probability in their own topology if desired, as the tool will accept their input file.

C.5 Summary

Masked modules in the Tile calorimeter have the potential to negatively affect any analysis which makes use of topo-cluster-derived quantities, whether that is jet kinematics, jet multiplicity, E_T^{miss} , or even calorimeter-based lepton isolation. Any analysis which makes use of any of these types of variables is recommended to apply one of two BCH cleaning methods to determine if their key variables of interest are noticeably affected.

Two cleaning strategies have been proposed, which have different intended uses:

1. Geometrical cleaning (“tight”): the event is rejected if any selected jet is either within or adjacent to a masked Tile module

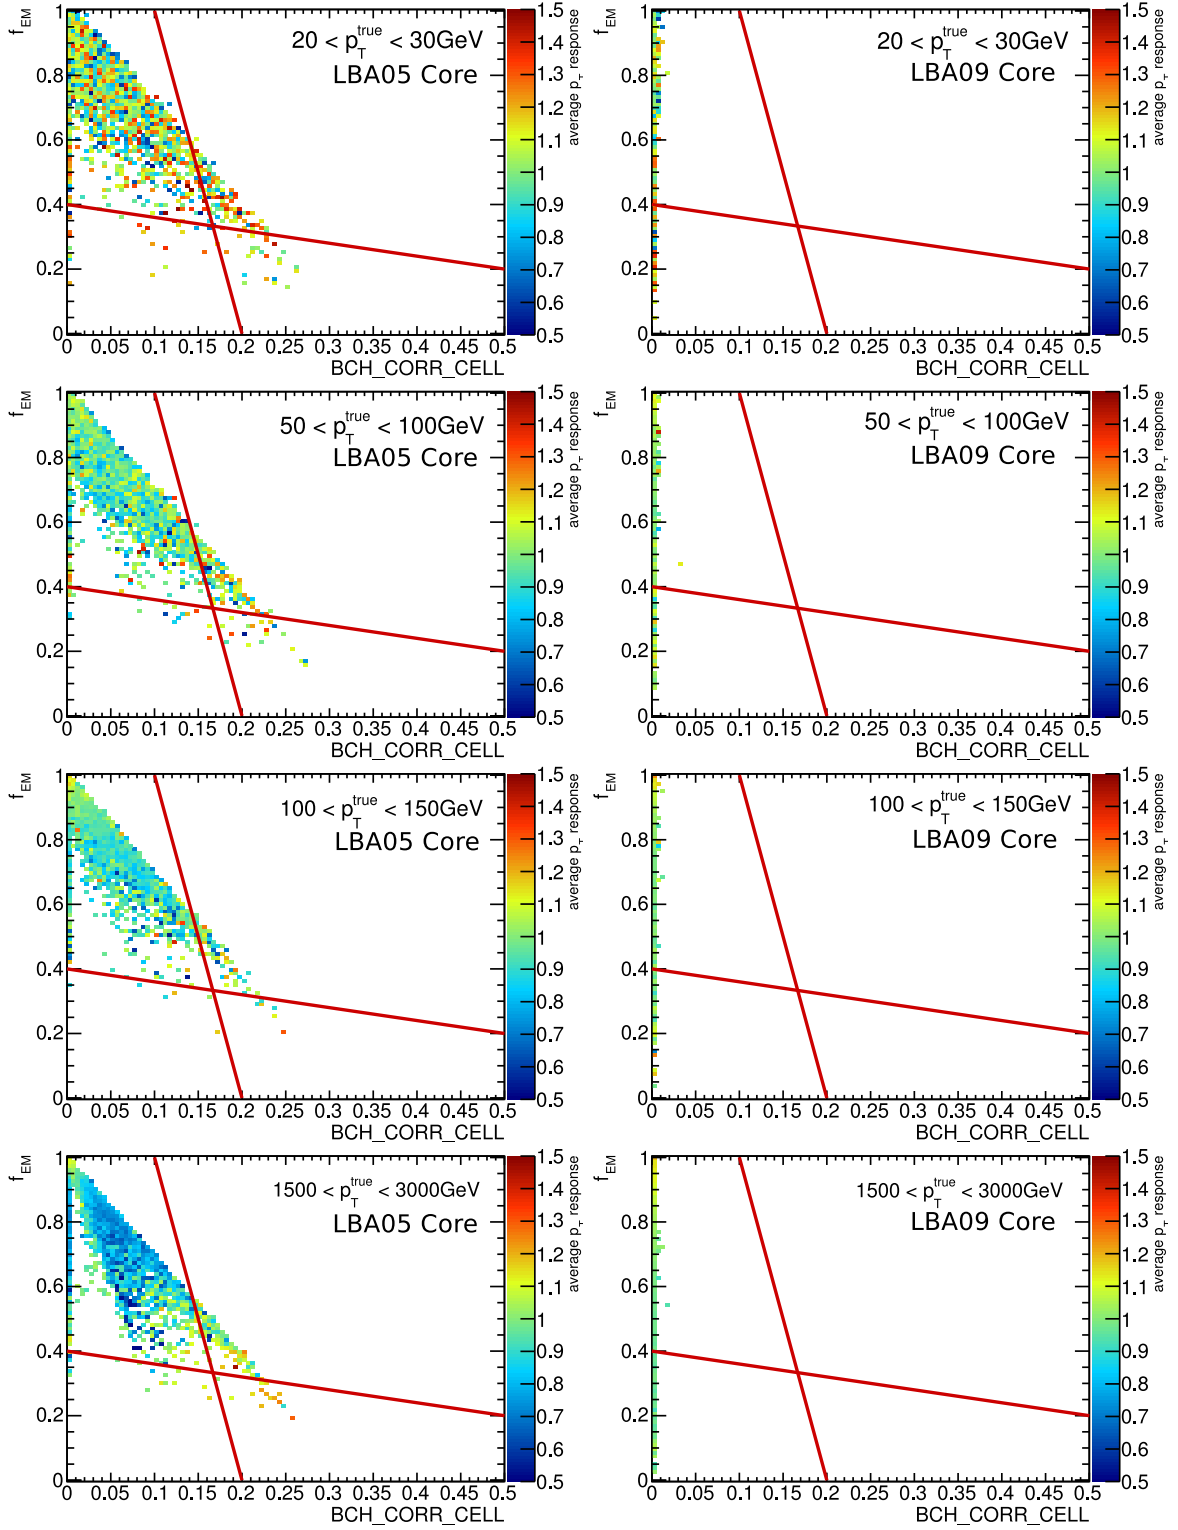


Figure C.7: A sample of MC distributions for various $p_T^{\text{jet, true}}$ of the jet electromagnetic fraction vs the BCH correction factor, for jets where the core is within a module which is masked (left) or fully operational (right).




Figure C.8: The derived probability used to determine whether or not to veto an event in MC, based on a jet of a given p_T , which is in region marked by the emulation as being masked. This is needed in order to match the application of the actual cut on f_{em} and BCH_CORR_CELL in data to the MC where the variables are not simulated in the same way. The probability depends on the type of parton which initiated the jet, due to the different fragmentation patterns causing expected changes in f_{em} . The all flavours probability is used for the nominal cleaning value, while the light quark and gluon probabilities are used for systematic variations.

2. Cut-based cleaning (“medium”): the event is rejected if any selected jet is either within a masked Tile module, or the jet is adjacent to such a module and it fails a set of cleaning criteria

Both strategies require the use of emulation to approximate the probability of a given module in being marked as masked in MC, given that the MC only includes a single masked module while there were several in data. This also takes into account transient trips, which occurred frequently, but for only a very small amount of data.

The geometric strategy requires no further assumptions, and efficiently removes the bias caused by masked modules. While it is a simple procedure, it is only recommended for analyses with low jet multiplicities, or which rely on very high p_T jets. A pure geometric cut is too strong of a requirement for many-jet topologies, where a significant part of the 2012 dataset would be rejected with this procedure.

For users with large jet multiplicities or who only work with low p_T jets, the cut-based cleaning procedure is recommended. This method only rejects the worst jets, and retains the majority of the dataset. However, it comes at the price of extra emulation, given that the variables used for the cleaning cut are not simulated in MC. This extra emulation has a known dependence on the flavour of the jet, and thus a cleaning systematic uncertainty is provided to account for this dependence.

Appendix D

Jet Energy Scale Calibration

An overview of the Jet Energy Scale (JES) calibration procedure is provided in Section 6.3. This Appendix provides a much more detailed view of the steps going into the calibration chain.

D.1 Pileup suppression

The first step of the jet calibration chain is the suppression of contributions from pileup jets. These can be either from in-time or out-of-time pileup effects, where in-time corresponds to energy coming from the wrong hard-scatter vertex (not the primary vertex) and out-of-time corresponds to energy deposits in the calorimeter from a different bunch crossing. In-time effects are thus typically parametrized with respect to the number of primary vertices in the bunch crossing, N_{PV} , while out-of-time effects are parametrized with respect to the average number of interactions per bunch crossing, $\langle\mu\rangle$.

The average number of inelastic interactions per bunch crossing is defined relative to the instantaneous luminosity \mathcal{L} , as defined in Equation D.1 from Reference [130]. In this formula, σ_{incl} is the total inelastic pp cross-section and $N_{\text{bunch}} \cdot f_{\text{LHC}}$ is the average frequency of bunch crossings at the LHC.

$$\langle\mu\rangle = \frac{\mathcal{L} \cdot \sigma_{\text{incl}}}{N_{\text{bunch}} \cdot f_{\text{LHC}}} \quad (\text{D.1})$$

Rather than being defined with respect to the luminosity, the number of primary vertices is a reconstructed quantity, which counts the number of observed interaction points within a given bunch crossing. A primary vertex is typically defined as a location along the beamline which has at least two tracks pointing to it. With an ideal detector, all of the inelastic collisions would be reconstructed, and the instantaneous μ would be equivalent to N_{PV} . However, this is not the case, given that the average over the full 2012 $\sqrt{s} = 8 \text{ TeV}$ dataset is $\langle\mu\rangle = 20.7$ and $N_{\text{PV}} = 12.5$. This suggests that the primary vertex reconstruction efficiency in 2012 was approximately 60%.

For the event to have been recorded, a trigger must have been fired, which generally means that at least one of the primary vertices is a hard-scatter event producing interesting physics. Given the cross-section, it is a good approximation to consider these other inelastic collisions to all be di-jet processes, where the resulting jets are referred to as pileup jets. Due to the exponentially falling jet transverse momentum distribution, these additional di-jet hard-scatter vertices typically produce soft jets. With $N_{\text{PV}} = 12.5$, and one vertex being dedicated to an interesting process, approximately $11.5 \times 2 = 23$

pileup jets can be expected.

Given that the vertices contributing to pileup jets are separate from the hard-scatter process, the pileup contribution to a given event is mostly independent of the topology being studied. As such, studies are typically conducted on very clean signatures, such that the pileup contribution can be clearly identified. The $Z \rightarrow \mu\mu$ selection in particular is often used, as was done in Reference [131], given the use of objects external to the calorimeter for identifying the events of interest.

Applying a jet object definition of $p_T > 20$ GeV to the 2012 dataset already substantially reduces the multiplicity of pileup jets, reducing the average multiplicity to about 5.5 jets/event as measured in Reference [131] with a $Z \rightarrow \mu\mu$ selection. Note that many of these 5.5 jets/event will come from the $Z + \text{jets}$ process rather than from pileup, where 4.5 of the 5.5 expected jets are not from pileup. As such, the expected average of 23 pileup jets/event has been reduced to an average of 1 jet/event just by applying a p_T threshold, before any pileup corrections are considered (95.7% rejection). Reference [131] has also observed that the fraction of pileup jets above 20 GeV which are also above a $p_T > 50$ GeV cut is below 1% in 2012 data, for an overall pileup jet rejection efficiency of 99.96% at 50 GeV. While this means that pileup jets remain at energies above this scale, it is a reasonable upper bound on the typical energy at $\sqrt{s} = 8$ TeV, and thus the pileup jet regime is often considered to be below 50 GeV.

While applying a jet p_T cut of 50 GeV would remove most of the pileup jets, this is far too restrictive for most ATLAS analyses. In the example above, such a cut would also remove a substantial part of the expected hard-scatter jets from the $Z + \text{jets}$ process in addition to the pileup jets. This becomes even more important for other processes such as Higgs boson measurements in the Vector Boson Fusion (VBF) topology, where the use of forward jets results in low transverse momenta but very high energy jets. For this reason, it is important to reduce the impact of pileup jets independent of the application of additional jet p_T cuts.

In addition, while a jet p_T threshold can be used to reduce the pileup jet multiplicity, this does not remove the pileup energy in the calorimeter. In particular, when a pileup jet overlaps with a jet from the primary vertex, the energy of the hard jet will be mis-measured. This means that even jets which are from the hard-scatter vertex of interest are biased, and a correction for the pileup contribution in the calorimeter is necessary.

There are also reasons from the jet performance side to reduce the contamination of pileup jets without p_T cuts. The MC JES calibration as described in Appendix D.3 works by correcting the average jet of a given p_T . If the pileup jet multiplicity is not reduced before this step, then the average correction applied will be derived with respect to correcting pileup jets to truth hard-scatter jets rather than correcting reconstructed hard-scatter jets to truth hard-scatter jets. This is a scenario which must be avoided if the calibration is to be accurate and useable.

Given the importance of pileup suppression, multiple methods have been devised. Techniques such as the JetVertexFraction(JVF) and JetVertexTagger(JVT), detailed in References [131] and [132] respectively, make use of the very precise directional information of the ATLAS tracking detectors in order to match jets to their primary vertex, and thus determine which came from the vertex of interest and which are from pileup. These methods are very powerful, but are limited to the acceptance of the tracking detector, and thus do not remove the large amount of pileup in the forward region. Furthermore, some analyses cannot make use of such cuts due to working with atypical topologies.

The jet calibration procedure instead makes use of a Jet Areas correction technique followed by a residual offset correction, as described in Reference [131]. This approach is particularly powerful for

anti- k_t jets due to their soft-resilient nature, thus resulting in hard jets with shapes (and areas) that are unaffected by pileup. This particular method is dynamic, with the benefit of being a jet-by-jet, event-by-event correction, thus having sufficient flexibility to handle the differing data taking conditions that naturally occur.

D.1.1 Jet area

The jet-by-jet part of the pileup correction relies on the calculation of the area of each jet. This is handled by the `FastJet` software library, in which a dense grid of *ghost* particles is added to the event before jet reconstruction begins. These ghost particles do not actually exist, but provide a useful tool. By setting them to have infinitesimal energies, the reconstructed jets are unaffected, thanks to the use of IR safe jet algorithms. The resulting jets have a list of constituents which they contain. Typically this is a set of topo-clusters for calorimeter jets, but now ghost particles are also included in the list. By counting the number of ghost particles in a given jet, and recalling they were added to the event in a grid, the number of ghosts provides a means of measuring the area of a given jet

Each of the ghosts may have an infinitesimal energy, but they are still particles with a full four-momentum. As such, summing all of the ghosts within a given jet results in a quantity known as the jet area four-momentum, defined in Equation D.2. In this Equation, $v_g \langle g_t \rangle$ is the transverse momentum density of the ghosts. Correcting the full jet four-momentum with the full jet area four-momentum has proven to require special care that was not possible at the time of the studies, thus only the jet transverse area A_T^{jet} was used to correct the jet transverse momentum.

$$A_{\text{jet}} = \frac{1}{v_g \langle g_t \rangle} \sum_{g_i \in \text{jet}} g_i \quad (\text{D.2})$$

D.1.2 Event density

The event-by-event part of the procedure is the calculation of the median calorimeter energy density for each event. Once again, the calorimeter is filled with a dense grid of ghost particles of infinitesimal energy. All of the ghost particles and positive energy topo-clusters within $|\eta| < 2$ in the event are then reconstructed with the k_t jet algorithm and a distance parameter of $R = 0.4$. The requirement of $|\eta| < 2$ is due to the higher calorimeter granularity within this region (the hadronic calorimetry switches from 0.1×0.1 in $\eta - \phi$ to 0.2×0.2 at $|\eta| \approx 2.5$, so to fully contain a jet of $R = 0.4$, the centroid must be approximately $|\eta| < 2$). From these jets, the resulting Jet Areas are calculated, and the collection of jets is sorted by the ratio of the jet transverse momentum to the jet transverse area. This provides an estimate of the average amount of energy in each calorimeter cell within a given jet, and this value should be much larger for real hard-scatter jets than for calorimeter noise. The median of this distribution is then taken to be the average event density. The use of the median rather than the mean allows for a selection of what should be an average calorimeter energy level without being biased by hard activity in the event, which should be in the high energy tail of the distribution. This average calorimeter energy should be due to pileup given that the hard energy should only fill a fraction of the calorimetry within $|\eta| < 2$, leaving the remaining topo-clusters accounting for noise, and pileup is the dominant source of calorimeter noise. In summary, the definition of the average event density ρ is:

$$\rho = \text{median} \left\{ \frac{p_{T,i}^{\text{jet}}}{A_{T,i}^{\text{jet}}} \right\} \quad (\text{D.3})$$

The calculation of ρ is always performed with the k_t jet algorithm, independent of what jets the resulting ρ is going to be applied to. This is because the k_t jet algorithm is soft-adaptive, and thus is expected to provide better modelling of soft jets, which is the primary use of the approach. However, while the k_t jet algorithm is always used, the topo-cluster inputs do have to be at the same scale to obtain a consistent calorimeter energy measurement. As such, two event densities are calculated, ρ_{EM} and ρ_{LC} .

D.1.3 Applying the correction

Now that the quantities have been derived, the jet-by-jet, event-by-event Jet Areas pileup suppression correction can be applied. The general intent is that the transverse area A_T times the average energy density ρ provides an estimate of the amount of pileup activity within a given jet. This can then be subtracted from the jet, providing a better estimate of the pileup-independent energy scale. The correction is thus defined in Equation D.4.

$$p_T^{\text{jet,corrected}} = p_T^{\text{jet}} - \rho \cdot A_T^{\text{jet}} \quad (\text{D.4})$$

This provides a very strong means of suppressing the in-time pileup contribution to the JES, as demonstrated for anti- k_t $R = 0.4$ jets in Figure D.1. However, this does not help with out-of-time pileup, and can actually slightly enhance the pileup sensitivity. This is resolved by applying an average residual offset correction after the Jet Areas technique, resulting in Equation D.5.

$$p_T^{\text{jet,corrected}} = p_T^{\text{jet}} - \rho \cdot A_T^{\text{jet}} + \alpha \cdot (N_{\text{PV}} - 1) + \beta \cdot \mu \quad (\text{D.5})$$

The results of applying this approach to the $Z \rightarrow \mu\mu + \text{jets}$ topology shows a significant reduction in the dependence of the number of jets on the level of pileup activity in the event, as seen in Figure D.2. This clearly demonstrates the power of the Jet Areas plus residual offset correction used by ATLAS. This correction was used for both the `Moriond2013` and `Final2012` releases of the ATLAS JES.

D.2 Origin correction

The jet origin correction was not applied in the `Moriond2013` release, but was included as a part of the `Final2012` JES. The default jet reconstruction procedure in 2012 created jets with respect to the nominal origin at $(0,0,0)$. This correction was designed to adjust the jet angular axis to point back to the actual hard-scatter interaction point, as calculated from the energy-weighted centroids of the constituent topo-clusters.

This correction is important to apply before the primary MC JES due to the dependence of the MC JES on η . While small changes in the pseudorapidity of the jet have little geometric difference, it can result in the jet passing through different regions of the calorimeter, which can result in very different required scale factors. In particular, jets passing in or close to dead regions, cracks, or transition regions between calorimeters are particularly sensitive to the angle. Note that these regions are also inherently

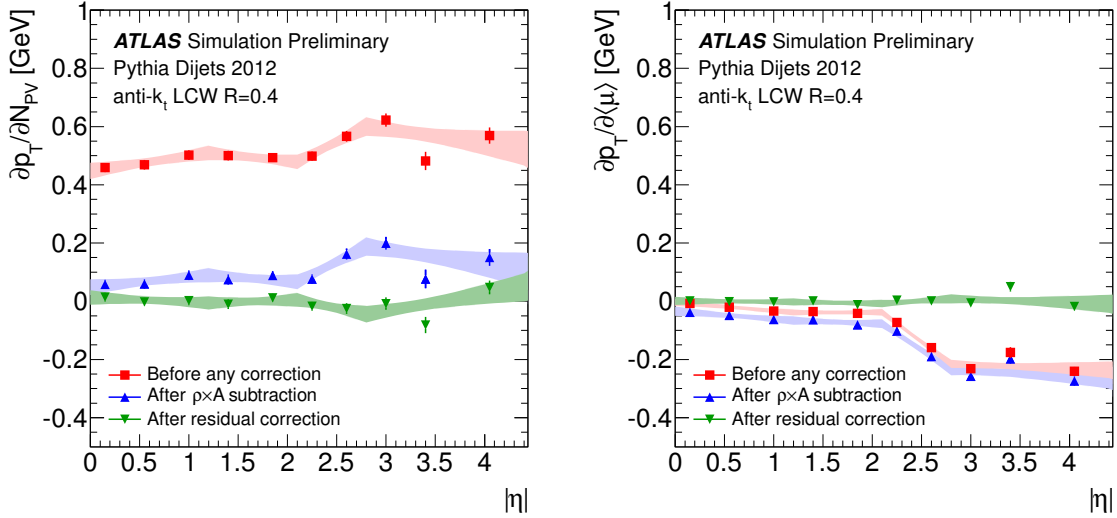


Figure D.1: The dependence of the p_T of anti- k_t $R = 0.4$ jets built from LC topo-clusters on in-time pileup (left) and out-of-time pileup (right). The dependence is shown before any pileup correction is applied (red), after the Jet Areas correction (blue), and after the Jet Areas plus residual offset correction (green). The Jet Areas correction provides near-closure with respect to in-time pileup, and only a minimal further residual correction is required. The Jet Areas correction does not help with out-of-time pileup, while the residual correction provides good closure. [131]

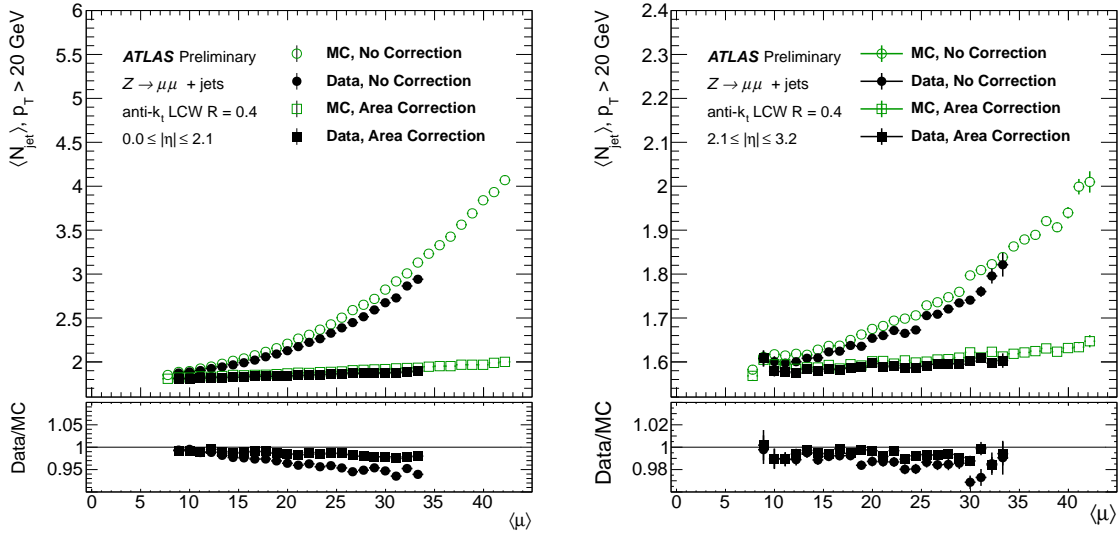


Figure D.2: The reconstructed jet multiplicity for jets of $p_T > 20$ GeV after a $Z \rightarrow \mu\mu$ selection as a function of the out-of-time pileup level $\langle\mu\rangle$, in the higher calorimeter granularity region (left) and lower calorimeter granularity region (right). The correction significantly decreases the dependence of the jet multiplicity on the level of pileup. The correction also brings the behaviour of data and MC with respect to $\langle\mu\rangle$ into closer agreement. [131]

biased, as any loss of energy can shift the jet centroid, and thus the reconstructed direction. The origin correction thus has the highest impact when applied to these regions.

In addition to changing the scale factor applied in the next step of the calibration, the origin correction serves to significantly reduce the jet angular resolution.

D.3 Absolute jet energy scale

As mentioned in Section 6.2, hadronic showers contain a sizeable semi-stable non-ionizing component. While some of this component is detected through the use of strong interactions to break the neutral particles into charged particles, or through the capture of soft neutrons, a large amount of energy is still unobserved by the ATLAS calorimetry. This lost energy causes an imbalance between the reconstructed jet and the true underlying physical energy, which must be accounted for. The dominant portion of this correction comes purely from MC, with possible data to MC differences addressed in the following Sections.

The MC JES is the single largest correction applied to the scale of the jet, and is applied in both the `Moriond2013` and `Final2012` JES releases. The correction is performed in bins of the jet energy and pseudorapidity, for reasons that will be detailed below. This correction is derived independently for jets built from EM or LC topo-clusters, as the LC corrections are a lower-level means of accounting for what the MC JES is designed to correct. Once the MC JES is applied, jets are referred to as being at either the EM+JES or LC+JES scale.

In order to derive a scale factor, truth and reconstructed MC jets from di-jet events are matched. This allows for the determination of the correct scale of the jet, and the ratio of the two forms the jet energy response, $\mathcal{R}_{\text{jet}} = E_{\text{jet}}^{\text{reco}}/E_{\text{jet}}^{\text{true}}$. The inverse of the response is then the scale factor which takes a jet of a given energy and pseudorapidity from the reconstructed scale to the truth scale. Thus, by creating a distribution for each energy and pseudorapidity, a correction can be derived which takes both MC or data jets from the reconstructed level to the truth level on an average basis.

As the jet energy is increased, the fraction of energy observed in the calorimeters grows, and thus the required scale factor is smaller. This phenomenon is explained in References [133, 134], where the fraction of energy in the hadronic sector of a hadronic shower is determined to be $F_h = (E/E_0)^{m-1}$, where $m \approx 0.83$ and E_0 is a constant which depends on the type of hadron. The remainder of the energy is in the electromagnetic sector, primarily from the production of π^0 mesons which decay almost exclusively to pairs of photons, thus initiating electromagnetic cascades. Given that electromagnetic energy is the means by which objects are detected in calorimeters, increasing the fraction of energy in the electromagnetic sector leads to a stronger calorimeter signal, and thus smaller scale factors are required to get back to the truth scale.

The dependence of the correction factors on the jet pseudorapidity varies as a function of the ATLAS detector structure, which is not uniform with respect to η . Some regions of the detector have more inactive or uninstrumented material than others, which introduces a means by which some of the hadronic shower can be either absorbed before detection or escape the detector. This is particularly relevant in the transition regions between calorimeters, such as between the barrel and the endcap at $|\eta| \approx 1.5$ or between the endcap and the forward calorimeter at $|\eta| \approx 3.2$. In these regions where more energy is lost, the jet response is lower, and thus the required scale factor is larger.

The dependence of the jet response on jet energy and pseudorapidity are both visible in Figure D.3.

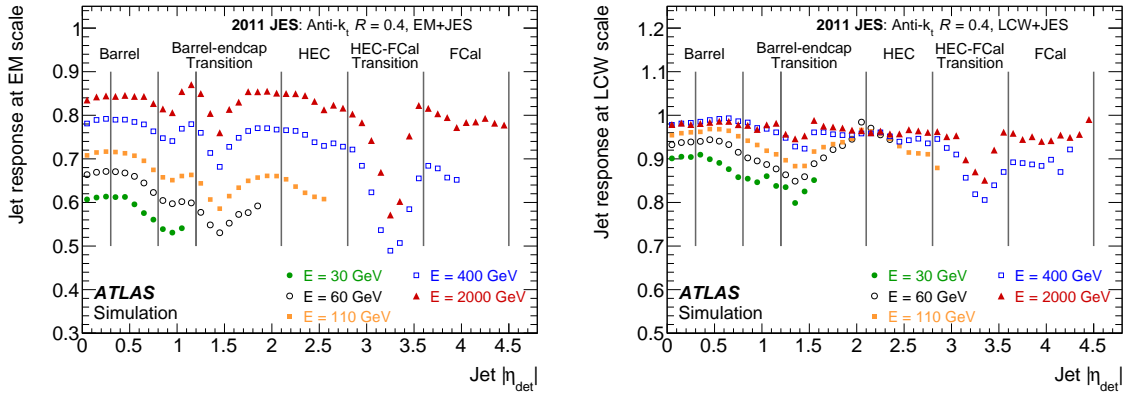


Figure D.3: The MC jet response \mathcal{R}_{jet} , which is the inverse of the MC JES, for anti- k_t $R = 0.4$ jets built from EM topo-clusters (left) or LC topo-clusters (right). These are taken from the 2011 $\sqrt{s} = 7$ TeV dataset, although the same procedure has been applied at $\sqrt{s} = 8$ TeV, with similar observed features. The LC weighting of topo-clusters is designed to account for some of the same effects as the MC JES, hence the response is closer to 1. Increasing the jet energy leads to a larger fraction of energy in the electromagnetic sector, leading to higher jet responses and thus lower scale factors. Varying jet pseudorapidity shows the impact of inactive material and cracks, most notable when jets fall in transition regions between calorimeters. [135]

These plots are from the 2011 $\sqrt{s} = 7$ TeV dataset, although the same procedure applies, and the detector inactive material configuration has not changed significantly.

D.4 Global sequential calibration

The Global Sequential Calibration (GSC) was first investigated with 2010 data with preliminary results published in Reference [136], but has only recently been included in the jet calibration sequence. In fact, the GSC is not a part of the `Moriond2013` JES calibration, although it is now included in the `Final2012` release. Jets which have the GSC applied are referred to as being at either the EM+JES+GSC or LC+JES+GSC scale. The technique is extensible and will likely contain more corrections in the future, but for now, it has three main purposes:

1. Improve the jet energy resolution
2. Reduce the observed jet response differences between light quark and gluon jets
3. Correct for jets which are not fully contained within the calorimeters

The GSC is designed to leave the core of the jet kinematic distributions unaffected, while modifying the tails, thus affecting the jet resolution but not the scale. In order to do this, the GSC focuses on key global jet properties in conjunction with the jet pseudorapidity and either the transverse momentum or energy, depending on which is more relevant. The five properties considered, the general purpose of each correction, the ordering which they are applied, and the $|\eta|$ region in which they are available is provided in Table D.1.

The correction factor for the property χ is derived by determining the average jet response $\mathcal{R}_{\text{jet}}^\chi = \langle p_{\text{T}}^{\text{jet, reco}} / p_{\text{T}}^{\text{jet, truth}} \rangle$ or $\mathcal{R}_{\text{jet}}^\chi = \langle E^{\text{jet, reco}} / E^{\text{jet, truth}} \rangle$ for a given bin of $p_{\text{T}}^{\text{jet}}$ or E^{jet} , $|\eta|$, and χ . The correction

Topo-cluster scale	Shower resolution		Flavour differences		Punch-through
	Correction 1 (f_{Tile0})	Correction 2 (f_{LAR3})	Correction 3 (N_{trk})	Correction 4 ($\text{width}_{\text{trk}}$)	Correction 5 (N_{segments})
EM	$ \eta \in [0, 1.7]$	$ \eta \in [0, 3.5]$	$ \eta \in [0, 2.5]$	$ \eta \in [0, 2.5]$	$ \eta \in [0, 2.7]$
LC			$ \eta \in [0, 2.5]$	$ \eta \in [0, 2.5]$	$ \eta \in [0, 2.7]$

Table D.1: Properties used for the Global Sequential Calibration, shown in the order in which they are applied. The first two corrections improve the resolution by exploiting shower profiles, and are only applied to jets built from EM scale topo-clusters, as the LC procedure has already applied similar techniques. The second set of two corrections reduce the differences in response between jets from gluons vs light quarks, and is applied to both EM and LC jets. The final correction accounts for jets which are only partially contained within the calorimeters, and which thus punch-through into the muon systems. This final correction is only available for use with full-simulation GEANT4 MC due to limitations in the current AFII parametrization. [137]

factor is then the inverse of the response, and is applied as a multiplicative factor to the MC JES. The properties χ considered are:

- f_{Tile0} : the fraction of the total jet energy (before calibration) deposited in the first layer of the Tile calorimeter, $|\eta| < 1.7$. The higher the value, the larger the contribution of the hadronic shower component. This results in a smaller fraction of the actual jet energy being reconstructed, meaning a lower jet response, and thus a larger needed correction factor. This correction is not applied to jets built from LC topo-clusters as they already have a similar correction built in. This correction is parametrized in terms of $p_{\text{T}}^{\text{jet}}$.
- f_{LAR3} : the fraction of the total jet energy (before calibration) deposited in the final layer of the LAR barrel or endcap calorimeters, $|\eta| < 3.5$ (EMEC ends at 3.2, so 3.5 covers jets which are partially within the EMEC). The transition between the barrel and endcap calorimeters contains cracks and inactive material, resulting in missed energy. In such cases, this variable will be larger, and thus a larger correction is needed. This correction is parametrized in terms of $p_{\text{T}}^{\text{jet}}$.
- N_{trk} : the number of tracks ghost-associated to the jet, which requires the tracker to be present, $|\eta| < 2.5$. This variable is sensitive to differences between light quark and gluon initiated jets, and can be used to correct for differences. This correction is parametrized in terms of $p_{\text{T}}^{\text{jet}}$.
- $\text{width}_{\text{trk}}$: the track width, which requires the tracker to be present, $|\eta| < 2.5$. This variable is also sensitive to the differences between light quark and gluon initiated jets. The track width is defined as the p_{T} -weighted average distance between the calorimeter jet axis and the ghost-associated tracks, $\text{width}_{\text{trk}} = \sum_i^{\text{tracks}} p_{\text{T}}^i \cdot \Delta R(i, \text{jet}) / \sum_i^{\text{tracks}} p_{\text{T}}^i$. This correction is parametrized in terms of $p_{\text{T}}^{\text{jet}}$.
- $N_{\text{seg}}^{\text{MS}}$: the number of segments in the muon system behind the jet using a simple ΔR matching, which requires the muon system to be present, $|\eta| < 2.7$. This variable is highly sensitive to jet punch-through, where a jet penetrates the full depth of the calorimeters and escapes into the muon system. This variable can be used to apply a factor to correct for the escaped energy, thus significantly improving the tails of very high energy jets. This is the one correction parametrized in terms of energy rather than $p_{\text{T}}^{\text{jet}}$, as the stochastic punch-through probability is dependent on only the amount of material to pass through and the jet energy. To demonstrate this, Reference [138] provides the equation for 95% longitudinal containment of a jet:

$L(95\%) = \lambda_I (0.6 \ln(E_{\text{jet}}) [\text{GeV}] + 4E_{\text{jet}}^{0.15} [\text{GeV}] - 0.2)$, where λ_I is the interaction length of the material the jet is traversing.

Unfortunately, no public plots are available to demonstrate the impact of this calibration sequence. Regardless, the GSC has been observed to change the central value of the JES by only a small amount, while significantly improving the resolution, thus serving the first purpose it was designed for. The difference in response between light quark and gluon jets is also observed to be reduced by a few percent, demonstrating the impact on the second intent of the GSC. The performance of the punch-through correction and more details on its derivation is provided in Appendix F.

D.5 *In situ* jet measurements

Both the MC JES and GSC calibrations are based on MC. This is guaranteed by construction to calibrate MC reconstructed jets to the truth scale on average, but the same correction is applied to data, which may behave differently. In order to confirm that this is not the case, and to quantify the difference between data and MC, a set of *in situ* studies are conducted in well-understood topologies. Each one of these studies is a full analysis, thus only a minimal summary of the overarching goal for each technique will be provided.

All of these studies were carried out for both the `Moriond2013` and `Final2012` JES releases, although to different levels of detail. The `Moriond2013` release occasionally made an approximation and inflated the corresponding uncertainty due to time constraints. In addition, the `Moriond2013` JES was derived with only approximately half of the 2012 dataset, while the `Final2012` JES includes the full dataset. Other more specific differences will be briefly mentioned where relevant, but in general the same set of techniques have been applied in both cases.

Systematic uncertainties resulting from all of these techniques will be shown and compared in Appendix E.

D.5.1 Di-jet η -intercalibration

Di-jet events have the highest statistics of any process at the LHC due to the enormous QCD cross-section. In addition, the two jets should balance, namely $p_T^{\text{jet1}} = p_T^{\text{jet2}}$ in a pure di-jet event with a perfect detector. By defining a central reference region of $|\eta| < 0.8$, the more forward calorimeter regions can be probed by requiring the central and forward jets to be balanced. Any deviations from this provide a means to quantify the dependence of the jet response on η in a data-driven way. The central reference jets are not modified, while the forward jets receive a correction to enforce the p_T balance. This provides a relative scale correction for the forward region with respect to the central region, guaranteeing that the correction derived for the central region in the following steps will be applicable to the full calorimeter.

Multiple methods have been defined to derive the needed corrections, which will not be discussed further here. The extent to which the forward region is probed is limited by statistics, and thus the jet p_T . The lowest p_T bins calibrate jets out to $|\eta| \approx 4.0$, after which the calibration from $|\eta| = 4.0$ is used. This range decrease significantly as the p_T is increased, with the highest p_T bins reaching out to $|\eta| \approx 1.5$

Uncertainties on the possible bias of applying a MC derived JES to the data are split between statistical and modelling components. The statistical components cover the full p_T and $|\eta|$ range, and

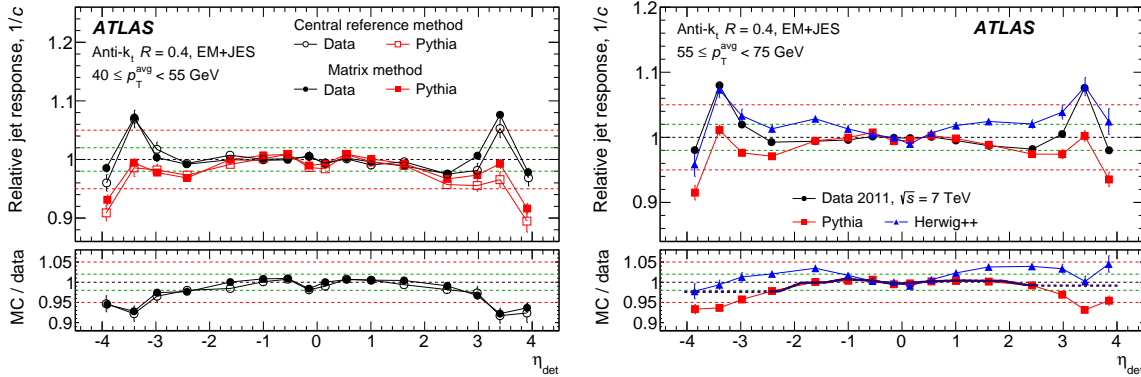


Figure D.4: The η -intercalibration jet response (top sub-plot) and double ratio (bottom sub-plot) using $\sqrt{s} = 7$ TeV data and the anti- k_t $R = 0.4$ jet algorithm built from EM topo-clusters. Error bars show only statistical uncertainties. A comparison of different methods for estimating the scale of the forward regions (left) and a comparison of the different generators used to estimate the modelling dependence (right) are provided. The different estimation methods provide slightly different results on the response, but very similar results on the double ratio, and thus provide a useful cross-check. The different MC generators result in significant differences to both the response and double ratio, resulting in the dominant η -intercalibration systematic uncertainty. [135]

are typically small. The modelling component is defined to be zero within the central region as this region is not touched, and is left to the following *in situ* methods. However, the modelling uncertainty can be very large in the forward regions, becoming the largest JES uncertainty. The dominant source of this uncertainty is due to the difference in radiation modelling between different MC generators.

For the Moriond2013 JES release, Pythia8 [49] and Herwig++ [139] were used to derive the modelling uncertainty. It was later discovered that the primary limitation in this technique was coming from radiation of a third jet, which is not available at ME level in either of the generators. As such, for the Final2012 JES, the Powheg Box [140] and Sherpa [81] generators were used instead, which support higher jet multiplicity generation at ME level. This change resulted in significantly improved overall modelling, although this source remains the dominant systematic uncertainty in the forward region.

No public plots currently exist for the $\sqrt{s} = 8$ TeV η -intercalibration procedure, but a similar approach was followed for $\sqrt{s} = 7$ TeV, and the results are provided in Figure D.4. The multiple methods for the estimation of the dependence of the jet response on η provide nearly identical results, while switching the generator used demonstrates the large modelling dependence, giving rise to the dominant systematic uncertainty.

D.5.2 Z/γ -jet balance

The $Z \rightarrow \ell\ell + \text{jets}$ and $\gamma + \text{jets}$ processes have a much smaller cross-section with respect to di-jets, but they have other benefits in that they provide a reference object with a well-known scale. Events which consist of a Z boson or a photon and a single isolated jet form a perfect topology to calibrate the JES to the reference object in a data-driven way. The $Z \rightarrow ee$ topology depends only on the electron energy scale, the $Z \rightarrow \mu\mu$ topology on the muon p_T scale, and the photon topology on the γ energy scale, all of which are independent of the JES. These three orthogonal estimates are all independently useful.

For each of these methods, the balance of the system is quantified as $p_T^{\text{jet}}/p_T^{Z,\gamma}$. This ratio should be

1 with a perfect reconstruction, although in practice it can be slightly different due to reconstruction effects. Instead, the double ratio of this quantity in data vs MC is used to calibrate the JES and quantify the residual differences. This double ratio is defined as $\mathcal{R}_{\text{jet/ref}}^{\text{data/MC}} = \langle p_{\text{T}}^{\text{jet}}/p_{\text{T}}^{\text{ref}} \rangle_{\text{data}} / \langle p_{\text{T}}^{\text{jet}}/p_{\text{T}}^{\text{ref}} \rangle_{\text{MC}}$, where $p_{\text{T}}^{\text{ref}}$ is the p_{T} of the reference object (Z or γ).

Given the much smaller statistics, these techniques are used to calibrate only central jets of $|\eta| < 0.8$. This correction is then applied to jets of all pseudorapidities, where the validity of this approach is guaranteed by the relative calibration applied in the previous di-jet η -intercalibration step.

The $Z \rightarrow ee$ and $Z \rightarrow \mu\mu$ *in situ* studies are primarily useful at low p_{T} , where the processes have considerable statistics available. The use of both the electron and muon channel allows for this range of usefulness to be slightly extended by doubling the available statistics. On the other hand, the photon topology is the most useful in the medium to high p_{T} regime. A combination of these techniques allows for the calibration of jets in the range of $20 \text{ GeV} < p_{\text{T}} < 800 \text{ GeV}$, where the $Z \rightarrow \ell\ell$ reference is used for the range of 20 to 200 GeV and the photon reference is used from 30 to 800 GeV. In all cases, the residual calibration derived from these methods is at the level of at most a few percent, as the MC samples used for the MC JES models the calorimeter response in data to a high degree of accuracy.

For the Moriond2013 release, only the direct balance of $Z \rightarrow \ell\ell/\gamma$ and a single jet was used. The Final2012 JES repeated the same studies in more detail, in addition to adding an alternate missing projection fraction selection. The missing projection fraction uses the entire hadronic recoil of the event rather than just the dominant balancing jet, and is more robust against both pileup and out-of-cone radiation [135].

Systematic uncertainties for these topologies are broken down much more finely than is done for di-jet η -intercalibration, as there are more possible sources which contribute to similar degrees. A full breakdown of these uncertainties is provided in Appendix E.

No public plots exist at this time for the $\sqrt{s} = 8 \text{ TeV}$ Z/ γ -jet balance calibration, but a similar procedure was followed for $\sqrt{s} = 7 \text{ TeV}$, and the results are provided for Z-jet and γ -jet in Figures D.5 and D.6 respectively. At very low p_{T} , the agreement between data and MC is reduced by the presence of considerable pileup. Excluding the lowest p_{T} range, both methods observe good agreement between data and MC.

D.5.3 Multi-jet balance

While $Z \rightarrow \ell\ell + \text{jets}$ and $\gamma + \text{jets}$ topologies provide an excellent method for the calibration of jets up to transverse momenta at the level of 800 GeV, they become limited at that scale due to a lack of statistics. In the absence of another process with a well-known reference object at a higher scale, a multi-jet balance is performed. This approach consists of a single very high p_{T} jet balanced by a recoil system of lower p_{T} jets. The scale of very high p_{T} jet is then corrected by the recoil system, as the lower p_{T} jets already have well-known scales thanks to the previous $Z \rightarrow \ell\ell/\gamma$ -jet balance studies. Once this is done, the very high p_{T} jet scale becomes well known, and thus the procedure can be repeated iteratively with even higher p_{T} jets balanced by a recoil system involving jets which were calibrated in the previous step. This process continues until the available statistics for such a topology becomes insufficient.

Analogous to what was done for $Z \rightarrow \ell\ell/\gamma$ -jet, the balance of the system is quantified as $p_{\text{T}}^{\text{jet}}/p_{\text{T}}^{\text{recoil}}$, where the ratio is not exactly 1. The same double ratio of data to MC as previously defined, with $p_{\text{T}}^{\text{ref}} = p_{\text{T}}^{\text{recoil}}$, is then used to further calibrate the JES and quantify the residual differences. The multi-jet balance procedure is actually begun for leading jets with a p_{T} of approximately 300 GeV in order to

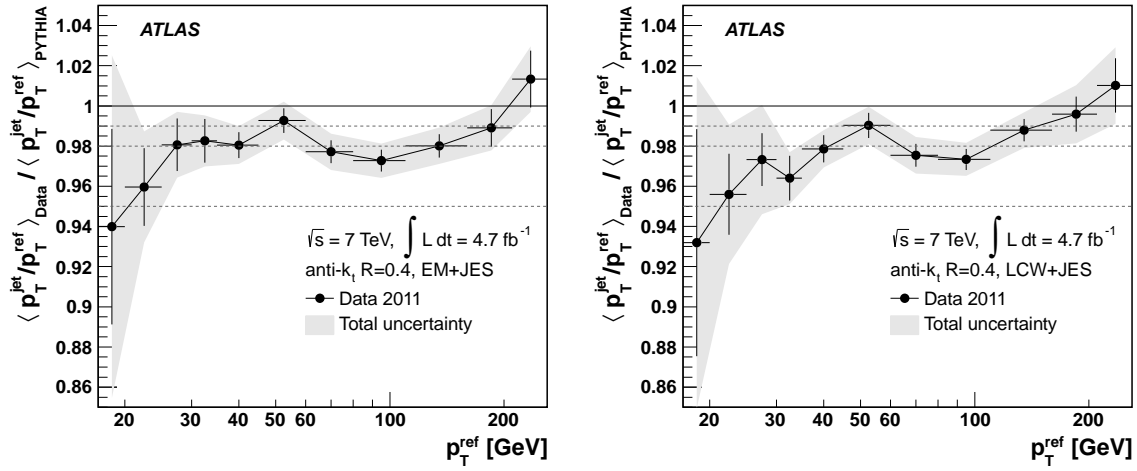


Figure D.5: The Z-jet balance double ratio using $\sqrt{s} = 7$ TeV data and the anti- k_t $R = 0.4$ jet algorithm built from EM topo-clusters (left) or LCW+JES (right). Data and MC are observed to differ by up to 7% in the lowest p_T bin, where pileup effects become important. Error bars show statistical uncertainties, while the shaded region is the associated systematic uncertainty. Dashed lines show residual correction shifts of 1%, 2%, and 5%. [135]




Figure D.6: The γ -jet balance jet response (top sub-plot) and double ratio (bottom sub-plot) using $\sqrt{s} = 7$ TeV data and the anti- k_t $R = 0.4$ jet algorithm built from LCW+JES, following the direct balance (left) and missing projection fraction (right) methods. Data and MC are seen to agree to within a couple percent across the full p_T parameter space. Error bars shown only include statistical uncertainties. The missing projection fraction is seen to perform better, and thus was used as the baseline method for the 2011 dataset, while direct balance was used as a cross-check. [135]

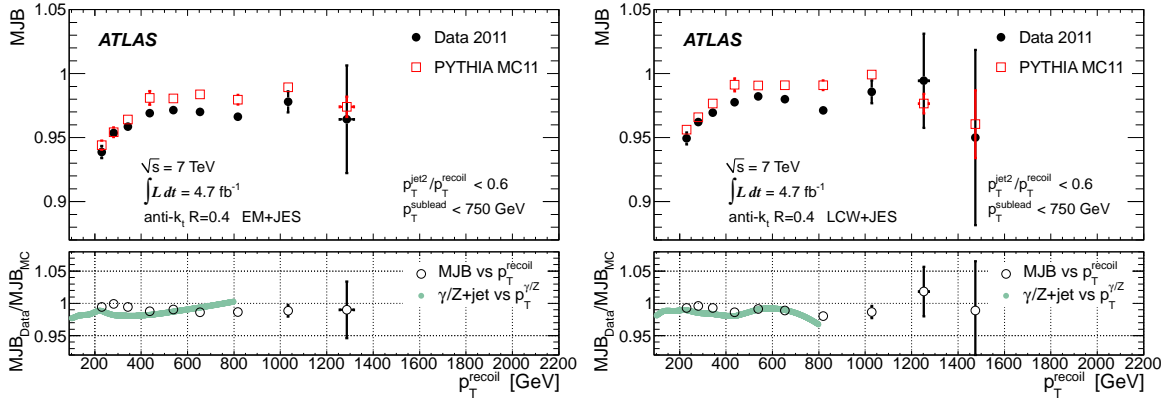


Figure D.7: The multi-jet balance jet response (top sub-plot) and double ratio (bottom sub-plot) using $\sqrt{s} = 7$ TeV data and the anti- k_t $R = 0.4$ jet algorithm built from EM topo-clusters (left) or LC topo-clusters (right). Data and MC are observed to agree to within approximately 1% precision. Error bars shown only include statistical uncertainties. The added solid curve on the bottom sub-plots represent the double ratio coming from Z/ γ -jet balance techniques. [135]

compare with the γ -jet balance studies, thus providing a cross-check to the multi-jet balance approach. Good agreement between the data to MC double ratio is observed over the full p_T space for which both methods are available.

As before, this procedure is carried out for only the central region of $|\eta| < 0.8$. After the full iterative procedure is applied, the jet calibration is extended from a maximum p_T of 800 GeV up to 1.5 TeV. The residual calibration derived over the full parameter space is at the percent level, once again demonstrating the good MC modelling of data and thus the applicability of the MC JES for calibrating data to the truth MC scale.

Given that the multi-jet balance procedure makes use of jets previously calibrated with the $Z \rightarrow \ell\ell/\gamma$ -jet approach, the systematics are more complicated. In addition to systematic uncertainties which arise purely from the multi-jet balance method, the systematic uncertainties from the previous steps are also propagated as necessary. The full description of the resulting uncertainties is provided in Appendix E

No public plots yet exist for the $\sqrt{s} = 8$ TeV multi-jet balance calibration, but a similar procedure was followed for $\sqrt{s} = 7$ TeV, and the results are provided in Figure D.7. Very good agreement is observed between data and MC, leading to a double ratio consistently within 1% of 1.

D.5.4 *In situ* combination

Once the independent results for the residual JES calibration in the form of $\mathcal{R}_{\text{jet/ref}}^{\text{data/MC}}$ are available for the $Z \rightarrow \ell\ell$ -jet, γ -jet, and multi-jet balance techniques, a combination is performed. This combination is conducted in terms of p_T^{jet} , where multiple methods are relevant for the majority of the parameter space. Good agreement is observed between the individual methods, and additionally the required residual JES correction is consistent with 1 within systematic uncertainties, demonstrating the validity of applying the MC JES to both MC and data reconstructed jets. The different *in situ* methods are then weighted to minimize the resulting uncertainties, exploiting each technique where it is the most sensitive, thus providing a more precise estimate of the residual JES correction than is possible from any single technique.




Figure D.8: The jet response double ratio $\mathcal{R}_{\text{jet/ref}}^{\text{data/MC}}$ as a function of $p_{\text{T}}^{\text{jet}}$ for the Moriond2013 JES calibration, showing the *in situ* combination of the Z-jet (blue open squares), γ -jet (purple filled triangles), and multi-jet balance techniques (green upside down open triangles). The resulting combined uncertainty is shown for the statistical (inner yellow band) and statistical plus systematic effects (outer green band). The residual *in situ* JES correction to apply to data is the inverse of $\mathcal{R}_{\text{jet/ref}}^{\text{data/MC}}$. [141]

The p_{T} ranges over which each *in situ* approach is used varies between the Moriond2013 and Final2012 JES releases, but the procedure remains unchanged. No plots for the Final2012 calibration are available at this time, although a public plot of the Moriond2013 is provided in Figure D.8.

D.6 Single hadron response measurements

While there is an insufficient quantity of multi-jet events above 1.5 TeV to provide a residual JES calibration factor for data, there are still collision events which produce jets above this scale, mostly in the di-jet topology. As such, some estimate of the scale must be determined. For this range, the MC JES is taken as-is, but a method for estimating the uncertainties through single hadron response measurements is performed. This method works by using truth information about jets in MC to determine which particles the jet was built from. Then, isolated track and 2004 test beam studies provide the calorimeter response to single particles of $0.5 \text{ GeV} < p_{\text{T}} < 20 \text{ GeV}$ and $20 \text{ GeV} \leq p_{\text{T}} \leq 350 \text{ GeV}$ respectively. By combining the truth particles in the jet with the measured single particle response, an *in situ* measurement of the JES of data can be derived. More details on this method as used for $\sqrt{s} = 7 \text{ TeV}$ collisions are provided in References [142, 135].

This approach has its limitations, given the need for conservative extrapolation measurements when the single pion energy is in excess of 400 GeV. The fraction of jets for which this is the case continues to grow as the centre of mass energy increases, and thus the range not covered by the multi-jet balance studies becomes higher and higher in energy. Regardless, this approach provides a means of estimating the uncertainty of very high p_{T} jets independent of the need for collision data. The resulting uncertainties derived from this approach for the range of $p_{\text{T}}^{\text{jet}} > 1.5 \text{ TeV}$ will be shown in Appendix E.

At this time, no public plots exist for the full treatment of the single hadron response measurements exists for the $\sqrt{s} = 8 \text{ TeV}$ dataset. The result of the same technique applied to the $\sqrt{s} = 7 \text{ TeV}$ dataset as compared to the *in situ* p_{T} balance measurements is provided in Figure D.9. The single hadron response uncertainty is seen to provide an envelope around the *in situ* results, confirming that the method works, and validating its use as a conservative estimate for the regime in which the *in situ* combination is not




Figure D.9: A comparison of the calorimeter response correction measured in the $\sqrt{s} = 7$ TeV dataset for anti- k_t $R = 0.6$ jets built from EM topo-clusters. The small grey band with a solid central line is the correction and associated uncertainty derived from the *in situ* combination of the Z-jet, γ -jet, and multi-jet balance studies. The larger yellow band and points are the estimate provided by the single hadron response measurement. The single hadron response measurement is observed to provide an envelope around the *in situ* combination, thus validating the procedure and justifying its use for uncertainties above the point at which the *in situ* combination is available. [135]

available.

Appendix E

Jet Energy Scale Uncertainties

An overview of the Jet Energy Scale (JES) uncertainties applicable to the mono-jet analysis is provided in Section 6.4. This Appendix provides a much more detailed view of the steps going into the derivation of the uncertainties, eigendecomposition procedures, and alternative correlation scenarios.

E.1 Uncertainty components

Complete lists of all of the uncertainty components for the `Moriond2013` JES and `Final2012` JES are provided in Tables E.1 and E.2 respectively. Note that each uncertainty component is assigned to a category representing a higher-level classification of the type of source responsible for the systematic. In particular, these categories are:

- Detector: effects related to the detector and its composition or response
- Modelling: uncertainties related to physics modelling and MC generator biases
- Mixed: uncertainties which are an intrinsic mix of detector and modelling effects
- Statistics and method: statistical uncertainties or uncertainties related to the method used

When none of these classifications is appropriate, the component is instead labelled as “special”. Beyond these categories, the full set of uncertainty components can be broken down into a groups based on their cause or method of estimation, as will be summarized below.

E.1.1 Common to Z/ γ -jet balance

The Z-jet and γ -jet *in situ* studies of the JES make use of electrons and photons respectively. The energy scale for these two types of physical objects is strongly correlated, and comes from the energy scale of the LAr electromagnetic calorimeters. As such, the electron and photon energy scales for a given source are considered to be one component with two fully correlated sub-components.

For both the `Moriond2013` and `Final2012` JES, this provides three uncertainty components, which quantify the dependence of the e/γ energy scale on the detector material composition, the presampler signal, and the $Z \rightarrow ee$ scale. The `Final2012` JES adds one more pair of correlated sub-components for the e/γ energy smearing dependence. For the `Moriond2013` release, all three energy scale components

Name	Description	Category
Common <i>in situ</i> sources		
e/γ scale, material	Electron/photon energy scale, detector material variations	Detector
e/γ scale, presampler	Electron/photon energy scale, presampler readings	Detector
e/γ scale, $Z \rightarrow ee$	Electron/photon energy scale, $Z \rightarrow ee$ scale uncertainty	Detector
<i>In situ</i> Z-jet balance		
Extrapolation	Extrapolation in $\Delta\phi$ between the jet and Z boson	Modelling
MC generator	MC generator difference between Alpgen/Herwig++ and Pythia8	Modelling
Out-of-cone radiation	Contribution of particles outside of the jet cone	Modelling
Pileup rejection	Pileup jet suppression using Jet Vertex Fraction cut	Mixed
Radiation suppression	Radiation suppression via a second jet veto	Modelling
Statistical components	Statistical uncertainty for each of the 10 bins	Stat/meth
<i>In situ</i> γ-jet balance		
Extrapolation	Extrapolation in $\Delta\phi$ between the jet and photon	Modelling
MC generator	MC generator difference between Herwig++ and Pythia8	Modelling
Out-of-cone radiation	Contribution of particles outside of the jet cone	Modelling
Photon purity	Background uncertainty in the photon purity estimate	Detector
Radiation suppression	Radiation suppression via a second jet veto	Modelling
Statistical components	Statistical uncertainty for each of the 9 bins	Stat/meth
<i>In situ</i> multi-jet balance		
α selection	Angle between leading jet and recoil system	Modelling
β selection	Angle between leading jet and closest sub-leading jet	Modelling
Fragmentation	Jet fragmentation modelling	Mixed
p_T asymmetry selection	p_T asymmetry selection between and sub-leading jet	Modelling
p_T threshold	Jet p_T threshold for selection	Mixed
Statistical components	Statistical uncertainty for each of the 10 bins	Stat/meth
<i>In situ</i> η-intercalibration		
Modelling term	Extra parton radiation generator difference, Pythia8 vs Herwig++	Modelling
Statistics term	Statistical uncertainty, all bins combined into one parameter	Stat/meth
Pileup suppression		
$\langle\mu\rangle$ average offset	<i>In situ</i> validation of MC based average $\langle\mu\rangle$ offset correction	Special
N_{PV} average offset	<i>In situ</i> validation of MC based average N_{PV} offset correction	Special
p_T bias	Dependence of the p_T on $\langle\mu\rangle$ and N_{PV}	Special
ρ topology	Dependence of the average energy density ρ on event topology	Special
Flavour differences		
b -jet scale	Only for b -jets, modelling of b -quark showers	Modelling
Flavour composition	Not for b -jets, uncertainty on fraction of light quarks vs gluons	Modelling
Flavour response	Not for b -jets, difference between light quark and gluon response	Modelling
Other sources		
Calibration non-closure	Only for AFII , MC JES calibration closure with respect to GEANT4	Special
High- p_T	Single particle extrapolation for jets above the <i>in situ</i> p_T range	Detector

Table E.1: Summary of the 58 JES uncertainty components for the **Moriond2013** release. Each uncertainty is also categorized based on its source, as either detector, MC modelling, mixed detector and modelling (mixed), or statistics and method (stat/meth). A small number of uncertainty sources do not fit into any of these classifications, and are thus labelled as special.

Name	Description	Category
Common <i>in situ</i> sources		
e/γ scale, material	Electron/photon energy scale, detector material variations	Detector
e/γ scale, presampler	Electron/photon energy scale, presampler readings	Detector
— e/γ scale, $Z \rightarrow ee$	Electron/photon energy scale, $Z \rightarrow ee$ scale uncertainty	Mixed
== e/γ smearing	Electron/photon energy smearing	Mixed
<i>In situ</i> Z-jet balance		
Extrapolation	Extrapolation in $\Delta\phi$ between the jet and Z boson	Modelling
— MC generator	MC generator difference between Powheg Box/Pythia8 and Sherpa	Modelling
== Muon scale	Muon transverse momentum scale	Detector
== Muon smearing, ID	Muon momentum smearing in the inner detector	Modelling
== Muon smearing, MS	Muon momentum smearing in the muon system	Modelling
Out-of-cone radiation	Contribution of particles outside of the jet cone	Modelling
Pileup rejection	Pileup jet suppression using Jet Vertex Fraction cut	Mixed
Radiation suppression	Radiation suppression via a second jet veto	Modelling
— Statistical components	Statistical uncertainty for each of the 11 bins	Stat/meth
<i>In situ</i> γ-jet balance		
Extrapolation	Extrapolation in $\Delta\phi$ between the jet and photon	Modelling
MC generator	MC generator difference between Herwig++ and Pythia8	Modelling
Out-of-cone radiation	Contribution of particles outside of the jet cone	Modelling
Photon purity	Background uncertainty in the photon purity estimate	Detector
Radiation suppression	Radiation suppression via a second jet veto	Modelling
— Statistical components	Statistical uncertainty for each of the 13 bins	Stat/meth
<i>In situ</i> multi-jet balance		
α selection	Angle between leading jet and recoil system	Modelling
β selection	Angle between leading jet and closest sub-leading jet	Modelling
Fragmentation	Jet fragmentation modelling	Mixed
p_T asymmetry selection	p_T asymmetry selection between and sub-leading jet	Modelling
p_T threshold	Jet p_T threshold for selection	Mixed
Statistical components	Statistical uncertainty for each of the 10 bins	Stat/meth
<i>In situ</i> η-intercalibration		
— Modelling term	Extra parton radiation generator difference, Powheg Box vs Sherpa	Modelling
Statistics term	Statistical uncertainty, all bins combined into one parameter	Stat/meth
Pileup suppression		
$\langle\mu\rangle$ average offset	<i>In situ</i> validation of MC based average $\langle\mu\rangle$ offset correction	Special
N_{PV} average offset	<i>In situ</i> validation of MC based average N_{PV} offset correction	Special
p_T bias	Dependence of the p_T on $\langle\mu\rangle$ and N_{PV}	Special
ρ topology	Dependence of the average energy density ρ on event topology	Special
Flavour differences		
b -jet scale	Only for b -jets, modelling of b -quark showers	Modelling
Flavour composition	Not for b -jets, uncertainty on fraction of light quarks vs gluons	Modelling
Flavour response	Not for b -jets, difference between light quark and gluon response	Modelling
Other sources		
Calibration non-closure	Only for AFII , MC JES calibration closure with respect to GEANT4	Special
High- p_T	Single particle extrapolation for jets above the <i>in situ</i> p_T range	Detector
== Punch-through	Only for GEANT4 MC, <i>in situ</i> validation of punch-through correction	Detector

Table E.2: Summary of the 68 JES uncertainty components for the **Final2012** release. Each uncertainty is also categorized based on its source, as either detector, MC modelling, mixed detector and modelling (mixed), or statistics and method (stat/meth). A small number of uncertainty sources do not fit into any of these classifications, and are thus labelled as special. Uncertainties which are significantly changed with respect to the **Moriond2013** release, either in procedure or categorization, have a name starting with “—”. Uncertainties which are new to the **Final2012** release have a name starting with “==”.

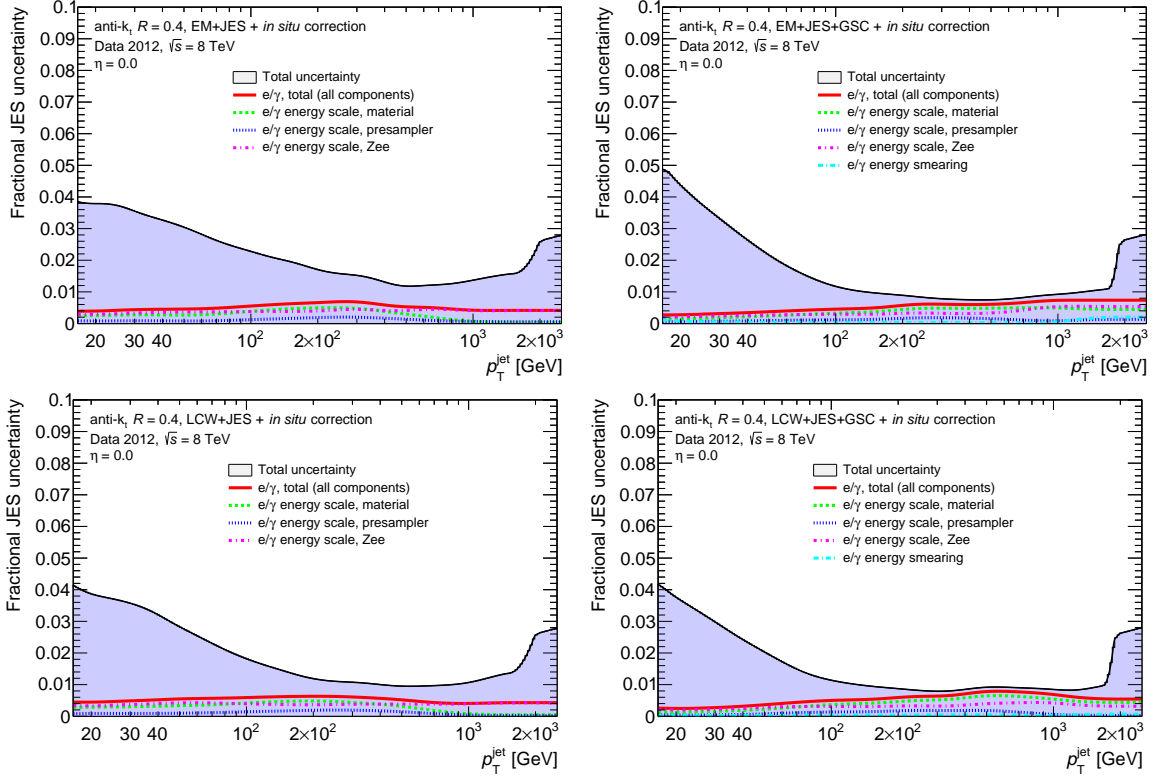


Figure E.1: JES uncertainties common to both the Z-jet and γ -jet *in situ* studies, for the Moriond2013 release (left) and Final2012 release (right), for anti- k_t $R = 0.4$ jets built from EM topo-clusters (top) and LCW topo-clusters (bottom), shown for the di-jet flavour composition topology. The total JES uncertainty including all sources is provided as a blue/grey filled background to demonstrate the relevance of the displayed components. The total uncertainty for all of the components of this type is shown as a solid red line.

are considered to be detector effects, while for the Final2012 release the $Z \rightarrow ee$ scale is changed to be a mixed detector and modelling effect. The e/γ energy smearing component is also referred to as a mixed detector and modelling effect. The Final2012 JES also studies variations of these categorization assumptions, as will be discussed in Appendix E.4.

Given that these components are from the Z-jet and γ -jet *in situ* measurements, which are restricted to the central region, the uncertainties are parametrized purely in terms of the jet p_T . The resulting breakdown of the uncertainty components in this group is provided in Figure E.1. This demonstrates that there is no single dominant e/γ uncertainty component, and they all (except smearing for the Final2012 release) contribute approximately equally. Furthermore, it is seen that the e/γ uncertainties are mostly relevant at the level of $p_T \approx 300$ GeV for the Moriond2013 JES, and other places are sub-dominant. However, for the Final2012 JES, the reduction in other uncertainties has led to a much stronger dependence on the e/γ components over a range of $100 \text{ GeV} \lesssim p_T \lesssim 1.5 \text{ TeV}$. This points out the interesting fact that the JES has become dominated by uncertainties on the photon energy scale over a wide p_T range.

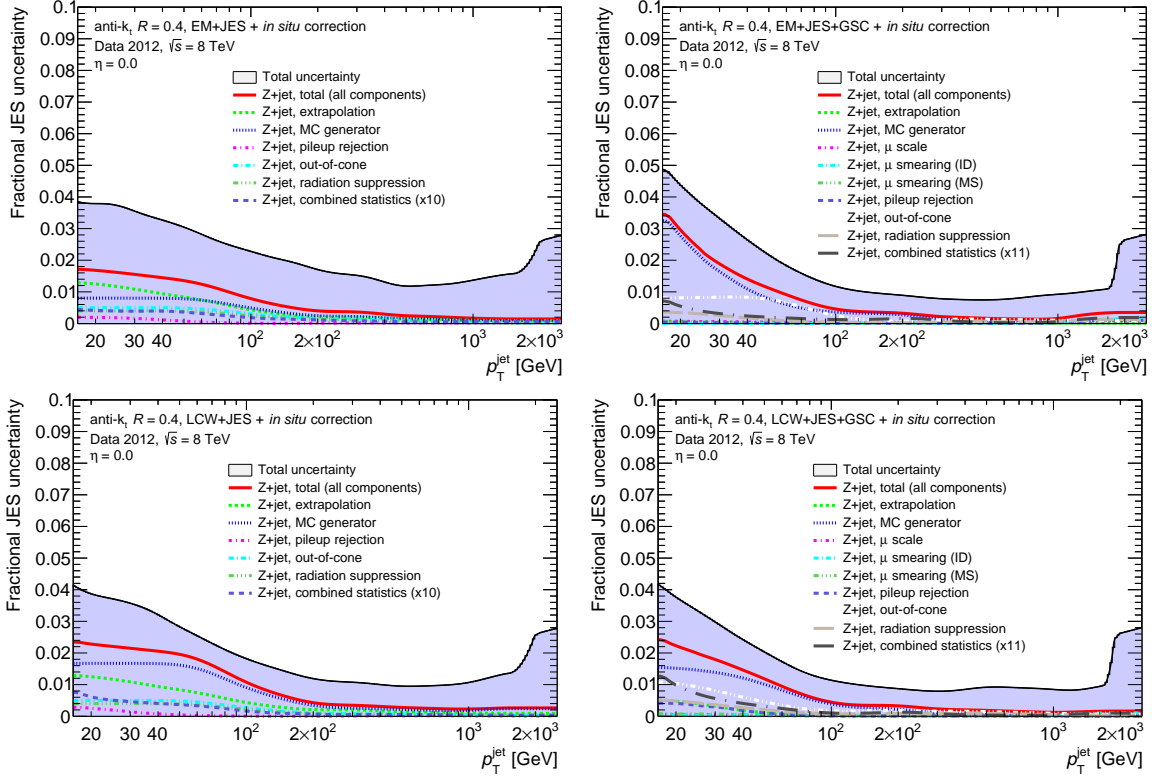


Figure E.2: JES uncertainties associated with the Z-jet balance *in situ* studies, for the `Moriond2013` release (left) and `Final2012` release (right), for anti- k_t $R = 0.4$ jets built from EM topo-clusters (top) and LC topo-clusters (bottom), shown for the di-jet flavour composition topology. The total JES uncertainty including all sources is provided as a blue/gray filled background to demonstrate the relevance of the displayed components. The total uncertainty for all of the components of this type is shown as a solid red line.

E.1.2 Z-jet balance

The Z-jet balance *in situ* measurement of the JES with respect to a reference Z boson has a small number of possible systematic biases. These are the extrapolation of the angle $\Delta\phi$ between the jet and the reference object, the dependence on the MC generator used, the means of pileup suppression, out-of-cone contributions, and the radiation suppression cuts. The `Final2012` JES adds three additional sources by the addition of the $Z \rightarrow \mu\mu$ topology, thus adding a muon scale and two muon smearing parameters. There are also one statistical uncertainty per bin, thus adding an additional 10 or 11 uncertainties for the `Moriond2013` and `Final2012` releases respectively.

Figure E.2 shows the breakdown of the Z-jet balance *in situ* uncertainties, demonstrating the primary importance of the method at low p_T . In general, the MC generator systematic is always large, but this is particularly true for EM topo-cluster based jets in the `Final2012` release. For the `Moriond2013` release, the $\Delta\phi$ extrapolation is also very important, while the out-of-cone contribution is more relevant for the `Final2012` JES.

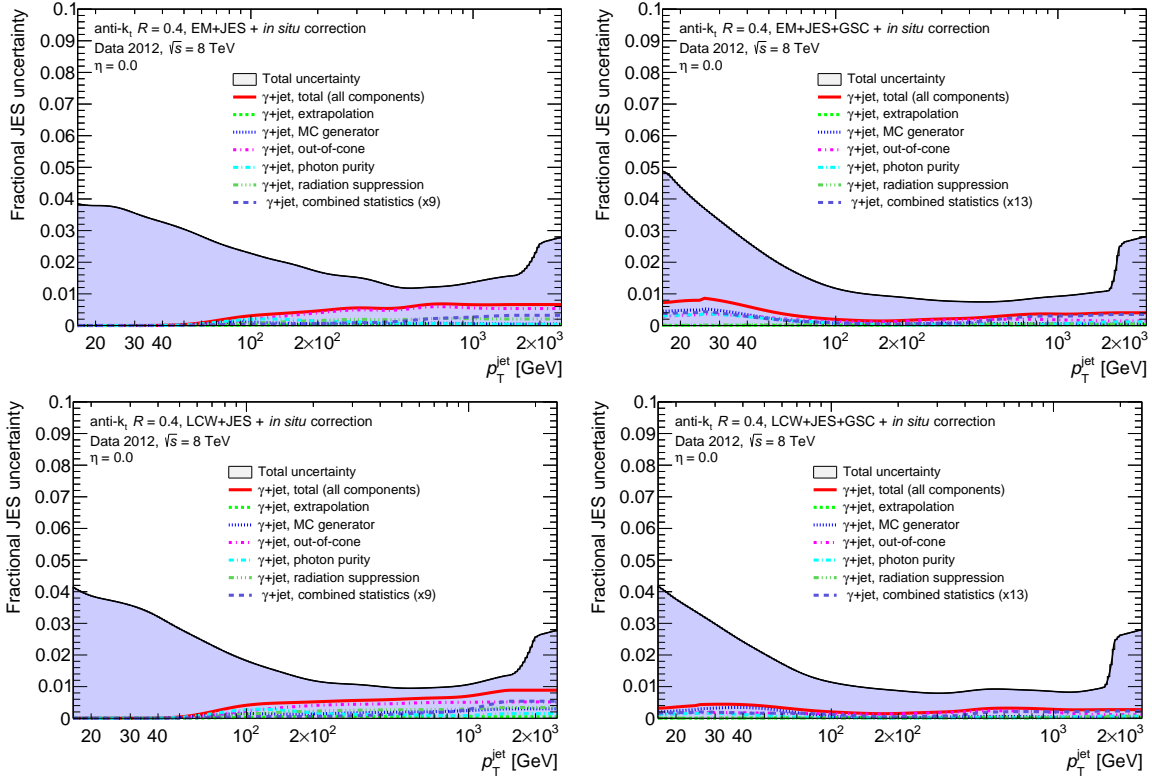


Figure E.3: JES uncertainties associated with the γ -jet balance *in situ* studies, for the `Moriond2013` release (left) and `Final2012` release (right), for anti- k_t $R = 0.4$ jets built from EM topo-clusters (top) and LC topo-clusters (bottom), shown for the di-jet flavour composition topology. The total JES uncertainty including all sources is provided as a blue/gray filled background to demonstrate the relevance of the displayed components. The total uncertainty for all of the components of this type is shown as a solid red line.

E.1.3 γ -jet balance

The γ -jet balance *in situ* measurement of the JES with respect to a reference photon also has a small number of possible systematic biases of a similar nature to the Z-jet balance method. The only difference between Z-jet and γ -jet for the `Moriond2013` release is the removal of the pileup rejection systematic and the addition of a source related to the photon purity. The γ -jet systematic sources are unchanged between the `Moriond2013` and `Final2012` JES, excluding an increase in the number of bins, and thus extra statistical uncertainty components.

Figure E.3 shows the breakdown of the γ -jet balance *in situ* uncertainties, demonstrating how they are relevant over a wide parameter range for the `Moriond2013` release, but have been substantially reduced to the point of being nearly irrelevant for the `Final2012` JES. In both cases, the γ -jet method is the dominant procedure and thus the dominant systematic source over a wide parameter space, the difference is just what specific component is dominant. For the `Moriond2013`, the answer is that both the jet sources (in this Section) and the photon energy scale (previously discussed) are relevant. However, for the `Final2012`, the jet-related systematics were improved to the point that they became negligible, while the photon energy scale uncertainty has become the primary limitation.

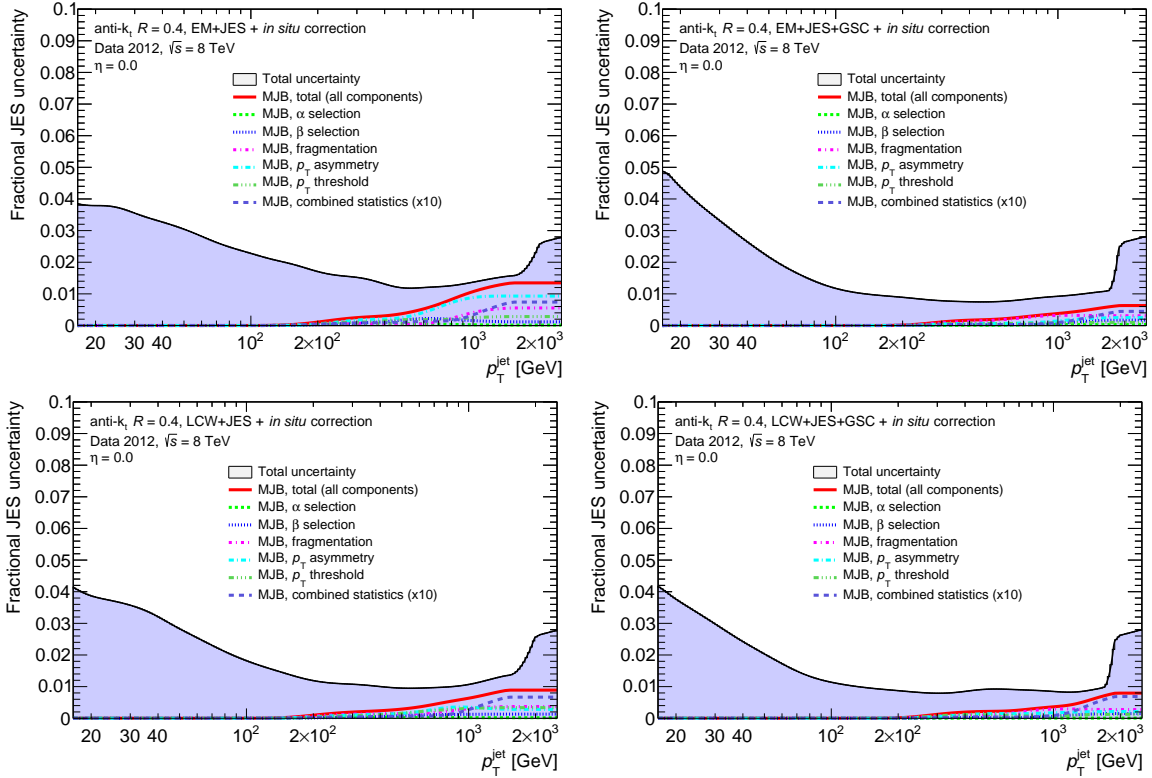


Figure E.4: JES uncertainties associated with the multi-jet balance *in situ* studies, for the `Moriond2013` release (left) and `Final2012` release (right), for anti- k_t $R = 0.4$ jets built from EM topo-clusters (top) and LC topo-clusters (bottom), shown for the di-jet flavour composition topology. The total JES uncertainty including all sources is provided as a blue/gray filled background to demonstrate the relevance of the displayed components. The total uncertainty for all of the components of this type is shown as a solid red line.

E.1.4 Multi-jet balance

The multi-jet balance *in situ* measurement of the JES with respect to a reference recoil system of lower p_T jets has a larger set of systematic uncertainties. However, most of them come from the propagation of the Z-jet and γ -jet balance systematics, and thus have already been discussed in previous Sections. The remaining systematics purely due to the multi-jet balance technique are related to the α and β angles on which cuts are applied, the MC fragmentation modelling, and the p_T asymmetry and threshold cuts. The 10 statistical uncertainty components are also more relevant for this set of studies.

Figure E.4 shows the breakdown of the multi-jet balance *in situ* uncertainties, demonstrating how they become significant only for very high p_T jets. In almost all cases, the statistical uncertainty dominates over the systematic uncertainty, underlining the limitation of the method, but also highlighting its ability to perform in a region where other methods cannot. Only the `Moriond2013` JES for jets built from EM topo-clusters is limited by systematic uncertainties, in the form of the p_T asymmetry requirement. This has been reduced for the `Final2012` JES.

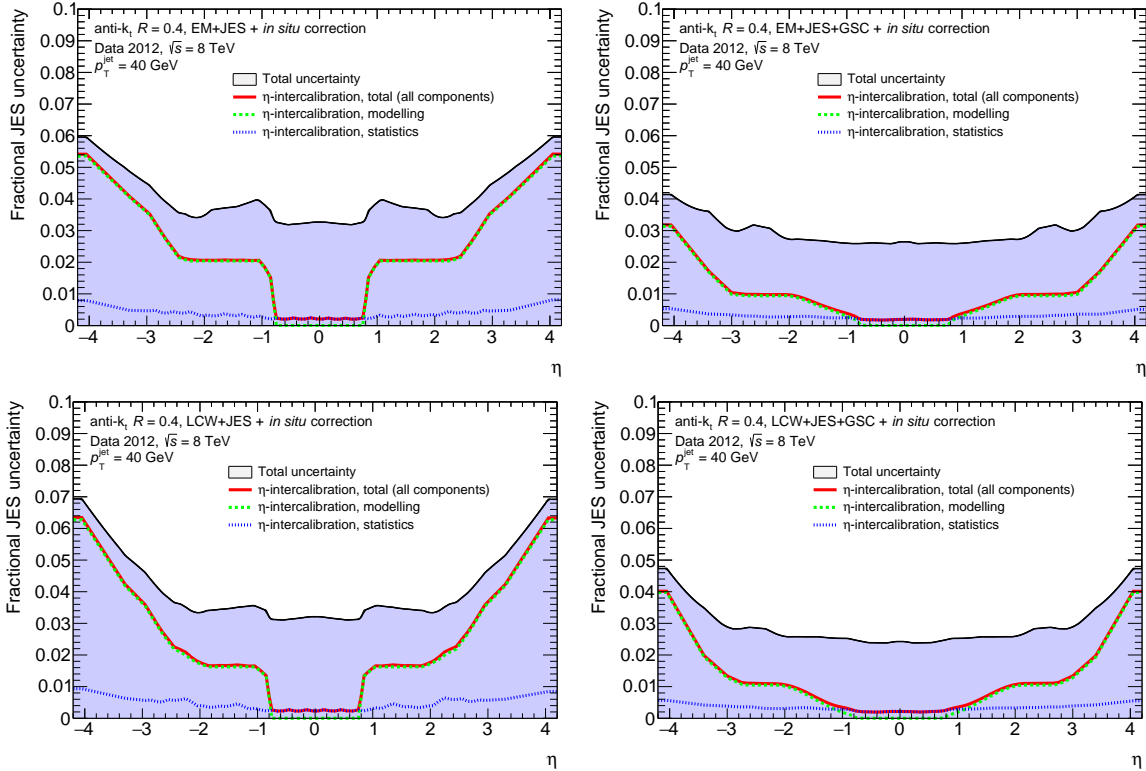


Figure E.5: JES uncertainties associated with the η -intercalibration *in situ* studies, for the Moriond2013 release (left) and Final2012 release (right), for anti- k_t $R = 0.4$ jets built from EM topo-clusters (top) and LCW topo-clusters (bottom), shown for the di-jet flavour composition topology. The total JES uncertainty including all sources is provided as a blue/gray filled background to demonstrate the relevance of the displayed components. The total uncertainty for all of the components of this type is shown as a solid red line.

E.1.5 Di-jet η -intercalibration

The di-jet η -intercalibration *in situ* measurement of the η -dependence of the JES is combined into only two uncertainty components representing the total statistical uncertainty and the total modelling uncertainty. The statistical uncertainty is non-zero everywhere, but is only relevant for $|\eta| < 0.8$, where the modelling term is zero by construction. Once outside of the reference region, the modelling term quickly becomes the dominant part of the JES uncertainty in the forward region.

As mentioned in Appendix D.5, the dominant modelling uncertainty comes from radiation of a third jet, which was not performed at ME level for the Moriond2013 release. By switching from Pythia8 vs Herwig++ to Powheg Box vs Sherpa, where third jet radiation is done at ME level, the resulting modelling uncertainty decreases substantially. This is demonstrated in Figure E.5. This change reduces the modelling uncertainty by an average of a factor of 2 over the entire jet pseudorapidity parameter space. This change means that for low to medium p_T jets, the η -intercalibration uncertainty goes from being the dominant or at least a major uncertainty in the region defined by $|\eta| > 0.8$ to a dominant uncertainty only for $|\eta| \gtrsim 3.5$. This is an enormous improvement, and future iterations of the JES will greatly benefit from this development.

E.1.6 Pileup suppression

The pileup suppression technique described in Appendix D.1 has not been changed between the `Moriond2013` and `Final2012` JES calibrations, and the uncertainties have not been re-evaluated. Four systematic sources were identified, corresponding to the dependence of the event energy density ρ on the topology, the average $\langle\mu\rangle$ offset correction, the average N_{PV} offset correction, and the *in situ* dependence difference of the jet p_T on both $\langle\mu\rangle$ and N_{PV} .

Figure E.6 shows the breakdown of the pileup suppression systematic uncertainties, although it must be noted that all of these uncertainties except for the ρ topology have a dependence on $\langle\mu\rangle$, N_{PV} , or both. As such, a value had to be chosen to make the plot. The values used correspond to the average value of $\langle\mu\rangle$ for the 2012 dataset plus $\sigma_{\langle\mu\rangle}$ and the average value of N_{PV} for the 2012 dataset plus $\sigma_{N_{PV}}$. The below conclusions are drawn from this set of assumed pileup conditions. Note that the pileup components will still be relevant primarily at low p_T regardless of the conditions chosen, as pileup is primarily a source of low p_T jets, as explained in Appendix D.1.

The dependence vs p_T reveals that the dominant uncertainty is due to the ρ topology uncertainty, closely followed by the N_{PV} offset. The $\langle\mu\rangle$ offset and p_T bias are both less relevant. Note that the p_T term is a function of both $\langle\mu\rangle$ and N_{PV} , so the fact that both $\langle\mu\rangle$ and N_{PV} were simultaneously shifted up means that this is smaller by construction. If $\langle\mu\rangle$ and N_{PV} fluctuate in opposite directions, then this term is expected to grow.

The dependence vs the jet pseudorapidity is also shown, demonstrating how the $\langle\mu\rangle$ offset uncertainty becomes relevant in the forward regions. The N_{PV} offset and ρ topology uncertainties remain mostly flat, and thus become sub-dominant for $|\eta| \gtrsim 2.0$. Once again, the p_T bias can vary significantly depending on the relative values of $\langle\mu\rangle$ and N_{PV} .

E.1.7 Flavour differences

The *in situ* JES calibration techniques calibrate jets which are primarily initiated by light quarks. It is well known that light quark jets have different radiation patterns than both b -jets and gluons, and thus a jet response difference is to be expected. This difference is dealt with on two fronts. If the user identifies the jet as a b -jet, then they are given a single scale systematic uncertainty related to the difference in response. This component is taken as a two-point systematic difference between the `Pythia8` and `Herwig++` MC generators.

If the user does not identify the jet as a b -jet, either because they do not apply b -tagging or because the jet is a light quark or gluon, then there are two systematic uncertainty components. The first is related to the difference between the quark and gluon response, and is calculated as per Equation E.1. The second relates to the fraction of light quarks vs gluons after the analysis selection has been applied, and is calculated following Equation E.2. Currently, the user is not required to consider b -tagging, as the uncertainty for light quarks vs gluons is larger, and thus conservative. However, if the difference between light quarks and gluons ever becomes smaller than the b -jet scale uncertainty, then this would need to be revisited.

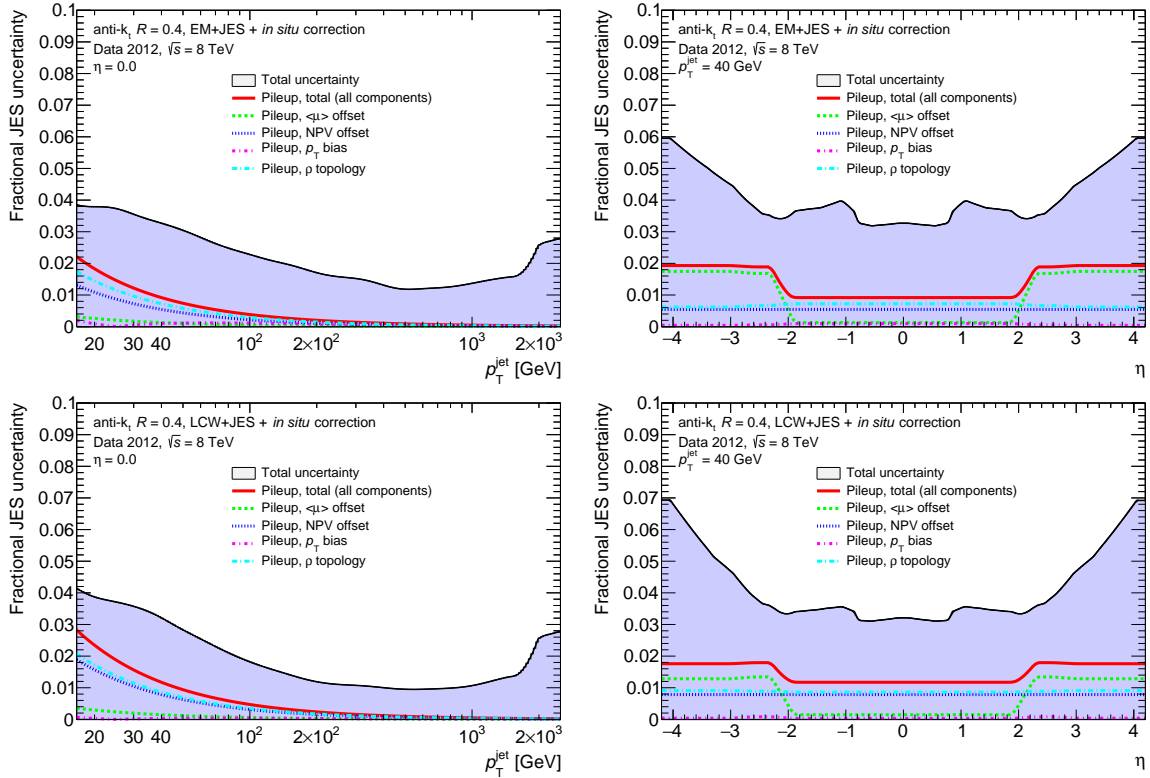


Figure E.6: JES uncertainties associated with the pileup suppression technique, for the Moriond2013 release, for anti- k_t $R = 0.4$ jets built from EM topo-clusters (top) and LC topo-clusters (bottom), plotted vs p_T (left) and η (right), shown for the di-jet flavour composition topology. The total JES uncertainty including all sources is provided as a blue/gray filled background to demonstrate the relevance of the displayed components. The total uncertainty for all of the components of this type is shown as a solid red line. The Final2012 plots are not shown due to using identical values for the pileup components, thus the plots only differ in the total JES uncertainty plotted in the background.

$$\sigma_{\text{flavour}}^{\text{response}} = f_g \cdot (\mathcal{R}_g^{\text{Pythia8}} - \mathcal{R}_g^{\text{Herwig++}}) \quad (\text{E.1})$$

$$\sigma_{\text{flavour}}^{\text{composition}} = \sigma_{f_g} \cdot \frac{|\mathcal{R}_q^{\text{Pythia8}} - \mathcal{R}_g^{\text{Pythia8}}|}{f_g \cdot \mathcal{R}_g^{\text{Pythia8}} + (1 - f_g) \cdot \mathcal{R}_q^{\text{Pythia8}}} \quad (\text{E.2})$$

In these equations, \mathcal{R}_g is the gluon response in the specified MC generator and \mathcal{R}_q is the light quark response in the specified MC generator. Furthermore, f_g is the user-specified value of the fraction of gluons after their selection, which is allowed to depend on p_T , η , and the jet multiplicity. The uncertainty on this quantity is σ_{f_g} . If the user does not specify the value of f_g and σ_{f_g} , the default unknown composition is assumed to be 50% gluons ($f_g = 0.5$), with a 50% uncertainty ($\sigma_{f_g} = 0.5$), thus covering all possible topologies.

Given that the flavour uncertainties constitute one of the larger parts of the JES, some users do determine the gluon fraction and error associated with their topology. As mentioned, the user is able to parametrize this in terms of p_T and η , which is the most common approach. The typical performance plots only parametrize the flavour composition in terms of p_T and η , specifically derived for the di-jet topology.

Some analysis groups, typically performing top measurements, take this one step further. They study and specify the flavour composition as a function of the number of jets in the event, in addition to the p_T and η . This is also supported, and provides a further means of more thoroughly understanding the JES for a given topology.

The flavour uncertainties are known to be smaller for jets built from LC topo-clusters rather than EM topo-clusters due to the shower profile weighting applied as a part of the LC procedure. This is very clear for the `Moriond2013` release, as shown in Figure E.7. With the addition of the GSC to the `Final2012`, LC remains superior in matters of flavour differences, but by a smaller margin. This is due to the addition of track-based corrections in the GSC, which are sensitive to flavour differences. This is also apparent in the same Figure, where the flavour uncertainty for both topo-cluster inputs is clearly reduced within the tracking acceptance region of $|\eta| < 2.5$.

E.1.8 Other sources

There are two or three other sources of JES systematic uncertainties, depending on the release. The first is the MC non-closure. While a separate MC JES is calculated for the full `GEANT4` simulation and the fast `AFII` parametrization, there are occasionally significant differences in the response. The `GEANT4` truth scale is taken to be correct, as it is the complete simulation, and the differences observed in `AFII` are taken as a systematic uncertainty. This uncertainty is typically only relevant for the low p_T region, where it can be sizable. This effect will be demonstrated in Appendix E.2.

The second component applicable to both the `Moriond2013` and `Final2012` release is the single particle high p_T extrapolation, explained in Appendix D.6. This is the dominant uncertainty for jets of $p_T > 1.5$ TeV, where the *in situ* methods are not applicable due to lack of statistics, but is zero everywhere else.

The final uncertainty is only available for the `Final2012` JES calibration, and relates to the punch-through correction within the GSC. This correction is only applied to full simulation `GEANT4` MC, as the `AFII` parametrization currently does not support the necessary calorimeter to muon system energy

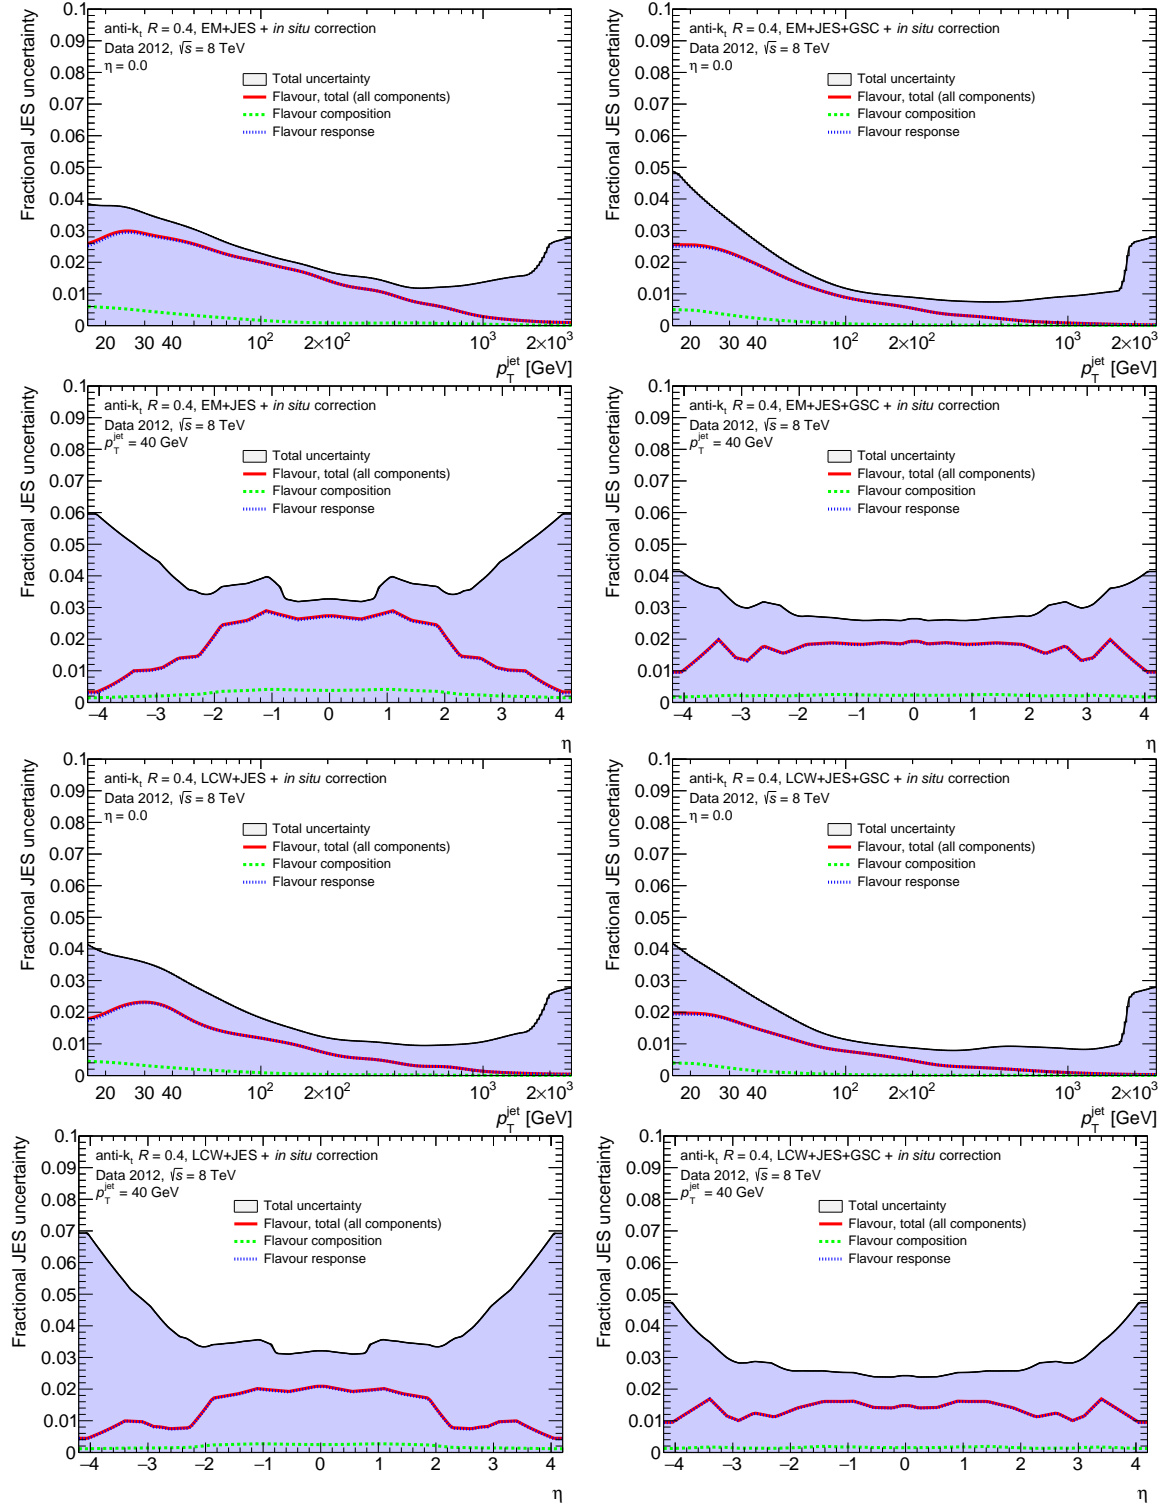


Figure E.7: JES uncertainties associated with jet flavour differences between light quarks and gluons, for the Moriond2013 release (left) and Final2012 release (right), for anti- k_t $R = 0.4$ jets built from EM topoclusters (top half) and LC topoclusters (bottom half). Plots are shown for the di-jet flavour composition topology, which reduces the flavour composition and increases the flavour response uncertainties. The total JES uncertainty including all sources is provided as a blue/grey filled background to demonstrate the relevance of the displayed components. The total uncertainty for all of the flavour components is shown as a solid red line. The impact of the GSC for $|\eta| < 2.5$ is clearly visible.

sharing which is necessary for the correction to be relevant. This correction has an insignificant impact on the average jet, as only a small fraction of jets penetrate the full depth of the ATLAS calorimetry, but it can be a sizable uncertainty for such jets. Much more information on this correction and the associated uncertainty is provided in Appendix F.

E.2 Uncertainty comparisons

The uncertainty components were split into groups in the previous Section in order to explain the sources and to clearly demonstrate the relevance of each individual term. However, the full set of 58 or 68 components is far too much to show at once when trying to compare one release to another, or when trying to convey the primary structure of the JES uncertainty to other groups. To this end, summary plots are produced, which implement the following structure:

- Absolute *in situ* JES: the quadrature sum of all of the p_T *in situ* sources. This includes the Z-jet balance, γ -jet balance, common Z/ γ -jet balance, multi-jet balance, and single hadron response uncertainties. This represents the majority of the uncertainty components by number, but not necessarily by magnitude.
- Relative *in situ* JES: the quadrature sum of the η *in situ* sources. This is currently just the two terms that come from the η -intercalibration.
- Flav. composition: the flavour composition uncertainty. This component is not combined with others as it is independently significant. The topology is typically also specified, where “inclusive jets” refers to the di-jet topology, and “unknown composition” refers to the default format where a composition of 50% gluons with a 50% uncertainty is assumed.
- Flav. response: the flavour response uncertainty. This component is not combined with others as it is independently significant. The topology is typically also specified, where “inclusive jets” refers to the di-jet topology, and “unknown composition” refers to the default format where a composition of 50% gluons with a 50% uncertainty is assumed.
- Pileup: the quadrature sum of the four pileup uncertainties, under the $\langle\mu\rangle$ and N_{PV} assumptions previously outlined in Section 6.4.
- MC non-closure (optional): this term is only shown for plots where the AFII parametrization was used rather than the full GEANT4 simulation. This term is zero by construction for GEANT4 MC.
- Punch-through (optional): this term is only shown for plots of the Final2012 JES release, and only for the full GEANT4 simulation. This correction was not available for the Moriond2013 release, and is not applicable for the AFII parametrization.

The remainder of this Section will cover comparisons of various different aspects of the JES uncertainties which have recently been used by ATLAS. Specifically, four cases will be shown per algorithm, per JES release, corresponding to:

1. The di-jet topology using full GEANT4 simulation MC
2. The di-jet topology using fast AFII parametrization MC

3. The unknown composition using fast AFII parametrization MC
4. The unknown composition using full GEANT4 simulation MC

After a comparison of the JES uncertainties within each of the individual releases, the two will be compared against each other.

E.2.1 The Moriond2013 release

The anti- k_t $R = 0.4$ algorithm using EM topo-cluster inputs is shown in Figure E.8. The use of AFII rather than GEANT4 is seen to be a relatively small effect for this scenario, while the total JES can be reduced by about 1% in the central region by specifying the flavour composition. In the forward region ($|\eta| \gtrsim 2.5$), the jet composition becomes less important due to the dominance of the relative *in situ* JES.

The anti- k_t $R = 0.4$ algorithm using LC topo-cluster inputs is shown in Figure E.9. The same statements can be made, although the gain from specifying the flavour composition is reduced due to a lower overall dependence on the flavour when using LC topo-clusters.

The anti- k_t $R = 0.6$ algorithm using EM topo-cluster inputs is shown in Figure E.10. Here, the use of AFII is significantly limited for jets of $p_T < 25$ GeV. This is a problem given that many analyses veto on jets at a similar scale, thus the use of AFII for jets of this type when applying a jet veto is not recommended. The dependence on the flavour is smaller than for the $R = 0.4$ case due to the increase of other uncertainties, but benefits can still be observed when the topology is specified.

The anti- k_t $R = 0.6$ algorithm using LC topo-cluster inputs is shown in Figure E.11. In this case, the AFII parametrization was not supported at all. The pileup uncertainty is observed to be dominant at low p_T , and there is little benefit observed from specifying the flavour composition.

In general, the uncertainties on central anti- k_t $R = 0$ jets where the user has specified a flavour composition has been reduced to the level of approximately 2% or below for the range of $150 \text{ GeV} < p_T < 1 \text{ TeV}$. At low p_T , the uncertainty in the forward calorimeters is at the level of 6 or 7%. The uncertainties for anti- k_t $R = 0.6$ jets are generally a bit larger than those for $R = 0.4$.

E.2.2 The Final2012 release

The anti- k_t $R = 0.4$ algorithm using EM and LC topo-cluster inputs is shown in Figures E.12 and E.13 respectively. In both cases, the use of AFII rather than GEANT4 is a minor effect. Specifying the flavour composition still helps, but in a much reduced fashion in the central region due to the use of track variables to correct for such effects as a part of the GSC. The main benefit of specifying the flavour composition is now in the forward region ($|\eta| \gtrsim 2.5$), there the uncertainties are much more relevant than before due to the significantly reduced relative *in situ* JES uncertainty.

The anti- k_t $R = 0.6$ algorithm using EM and LC topo-cluster inputs is shown in Figures E.14 and E.15 respectively. In both cases, the use of AFII rather than GEANT4 is a large effect for $p_T \lesssim 40$ GeV. Higher p_T jets are minimally affected by the difference in simulation types. As before, this means that anti- k_t $R = 0.6$ jets should not be used in conjunction with the AFII parametrization if the analysis intends to veto on low p_T jets.

In general, the uncertainties for central anti- k_t $R = 0.4$ jets where the user has specified a flavour composition has been reduced to the level of 1% for the range of $100 \text{ GeV} < p_T < 1 \text{ TeV}$. This is a significant achievement, and demonstrates the impressive precision of the ATLAS calorimeters and

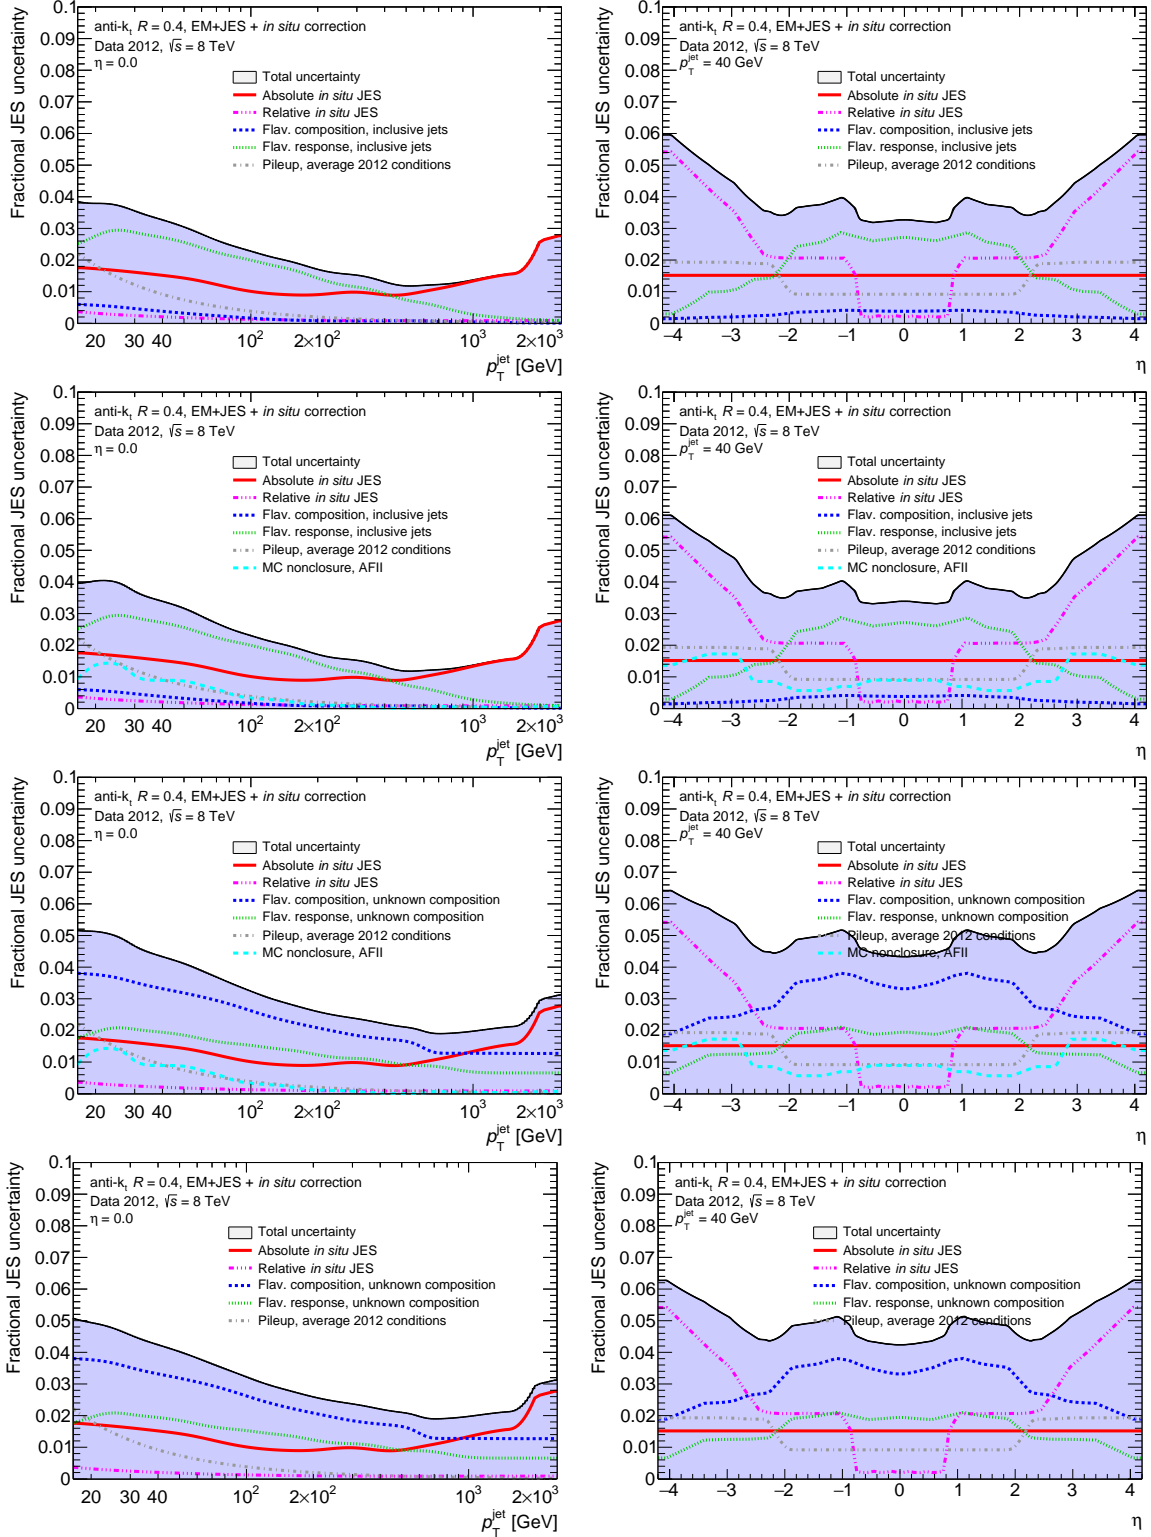


Figure E.8: Moriond2013 JES uncertainties for anti- k_t $R = 0.4$ jets built from EM topo-clusters, shown vs p_T (left) and η (right), for the full simulation di-jet topology (top), fast simulation dijet topology (upper middle), fast simulation unknown topology (lower middle), and full simulation unknown topology (bottom).

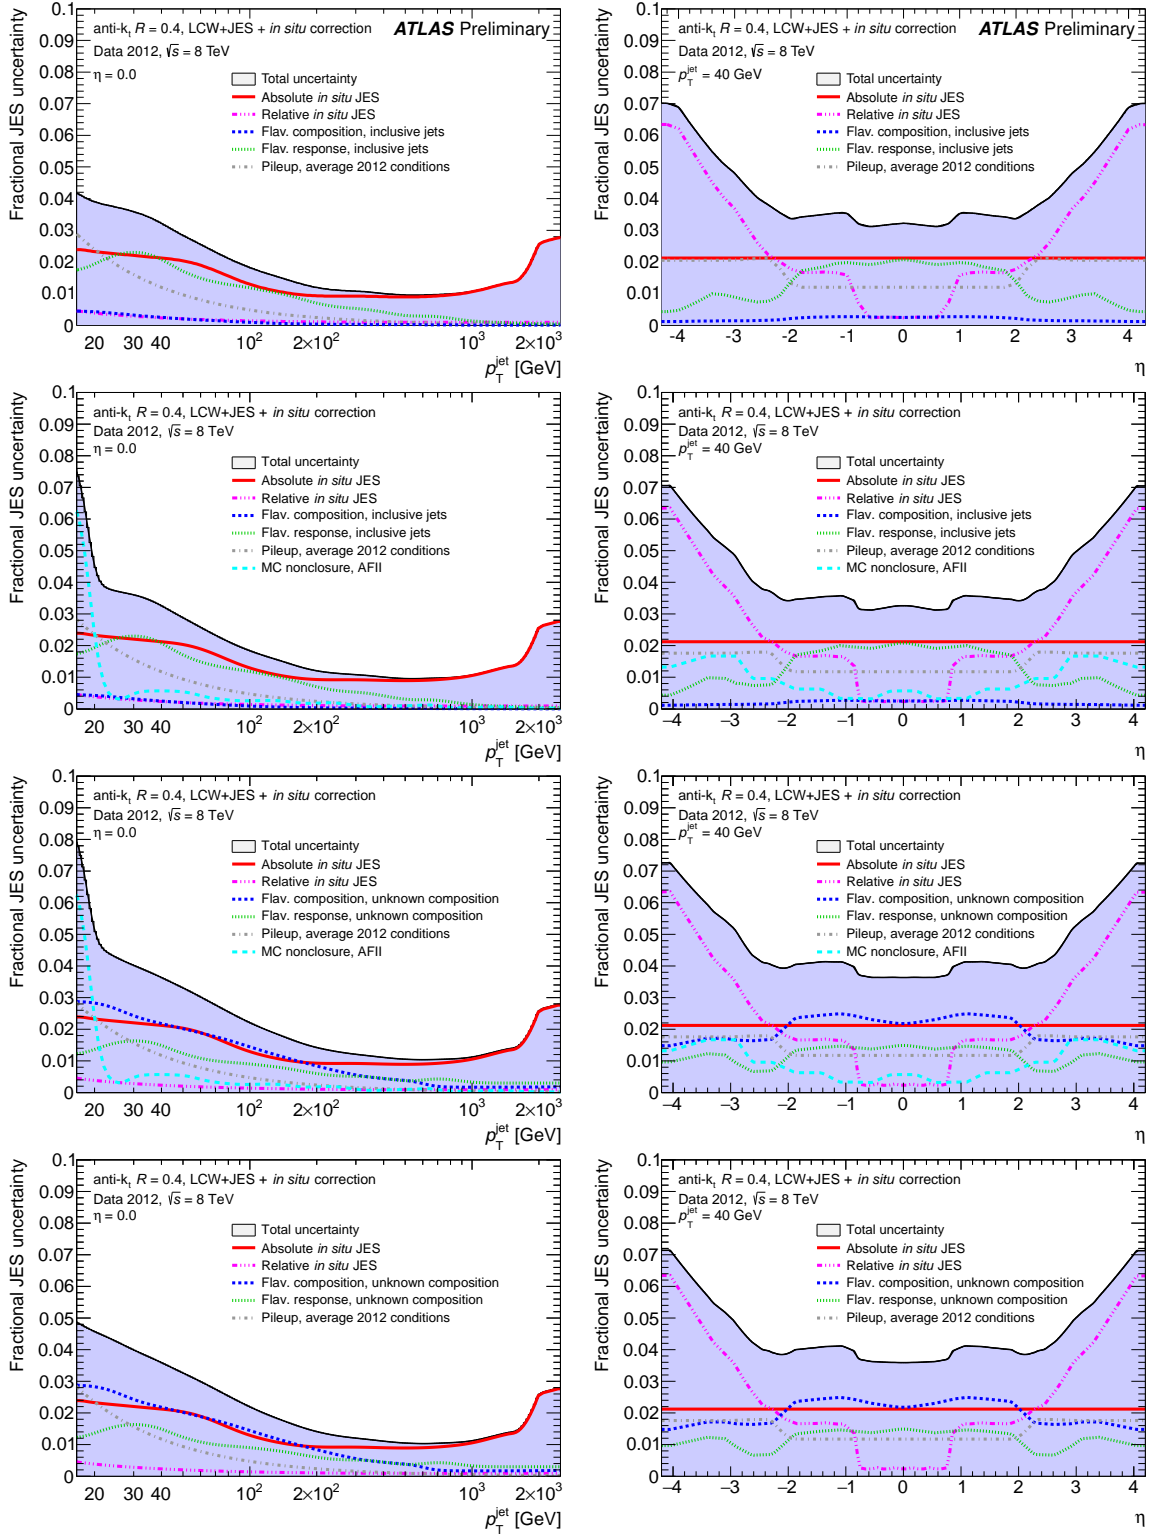


Figure E.9: Moriond2013 JES uncertainties for anti- k_t $R = 0.4$ jets built from LC topo-clusters, shown vs p_T (left) and η (right), for the full simulation di-jet topology (top), fast simulation dijet topology (upper middle), fast simulation unknown topology (lower middle), and full simulation unknown topology (bottom). The change in pileup between the top row and others is due to an update of $\sigma_{\langle\mu\rangle}$ and σ_{NPV} . [141]

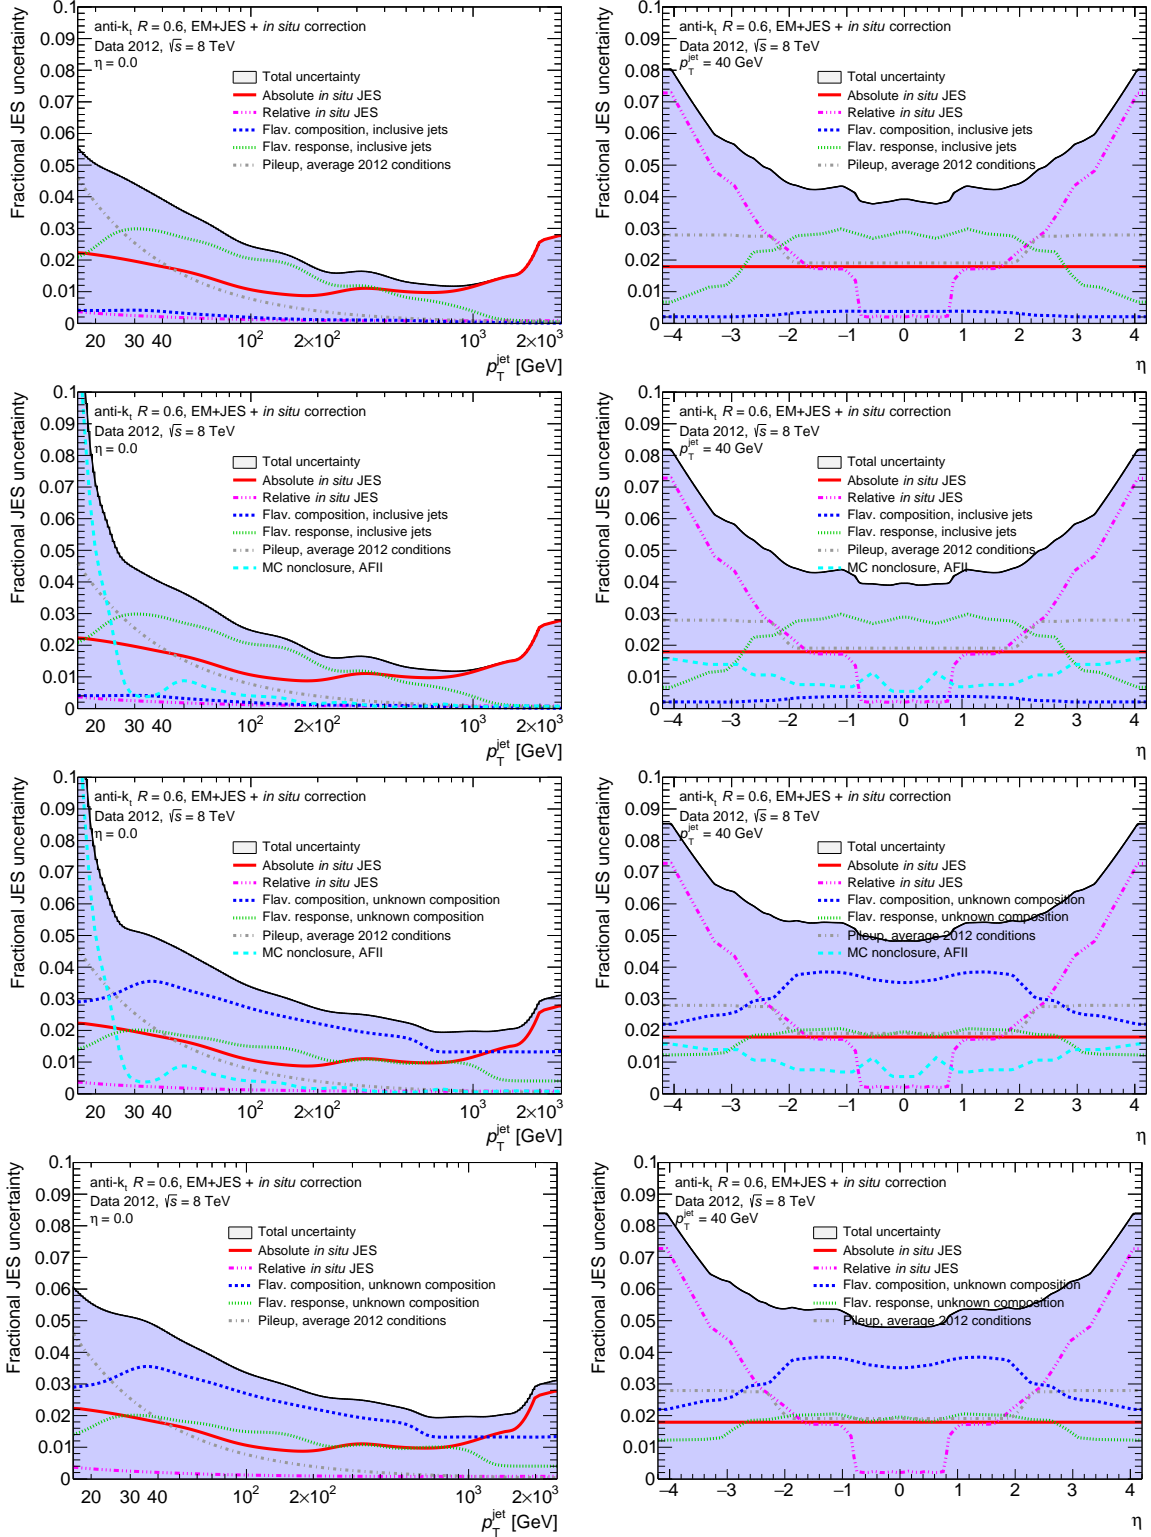


Figure E.10: Moriond2013 JES uncertainties for anti- k_t $R = 0.6$ jets built from EM topo-clusters, shown vs p_T (left) and η (right), for the full simulation di-jet topology (top), fast simulation dijet topology (upper middle), fast simulation unknown topology (lower middle), and full simulation unknown topology (bottom).

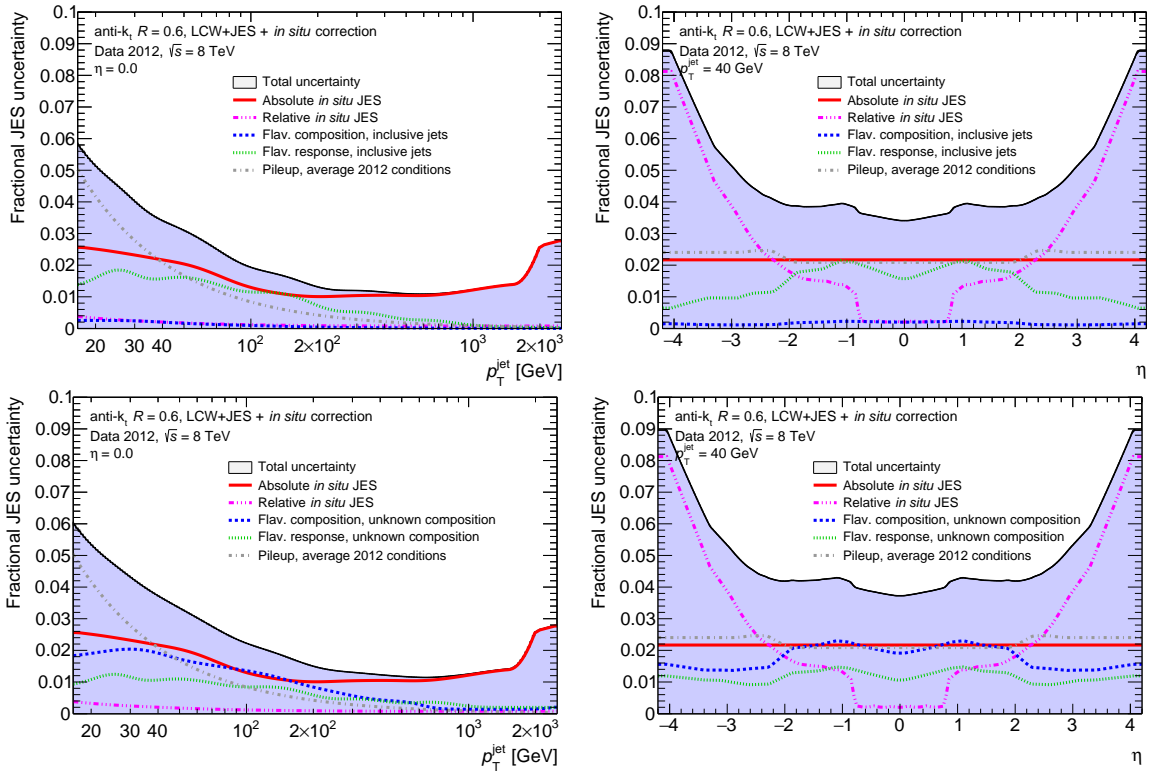


Figure E.11: Moriond2013 JES uncertainties for anti- k_t $R = 0.6$ jets built from LC topo-clusters, shown vs p_T (left) and η (right), for the full simulation di-jet topology (top) and full simulation unknown topology (bottom). Fast simulation for this jet type was not supported in the Moriond2013 JES release.

the *in situ* techniques used. Furthermore, even at low p_T , the uncertainty out to the most forward calorimeter regions is at or below the level of 4%. The uncertainties for anti- k_t $R = 0.6$ jets are only slightly larger than those for $R = 0.4$.

E.2.3 Comparing releases

Given the enormous amount of effort that went into the derivation of the `Final2012` JES calibration and uncertainties, it is important to consider the outcome of this work. In particular, comparing to previous uncertainty releases demonstrates where gains have been made, and where future work could prove to be beneficial. In addition to the `Moriond2013` and `Final2012` releases, the `HCP2012` preliminary release from November 2012 will also be considered as a baseline. This particular JES was from before the *in situ* studies had been run on 2012 data, and thus it was intended for quick results.

A comparison of the different JES releases for anti- k_t $R = 0.4$ jets is provided in Figure E.16. The improvement obtained from the first *in situ* studies as a part of the `Moriond2013` provide a significant improvement with respect to `HCP2012`. For the di-jet topology, a more moderate gain is observed when moving from the `Moriond2013` to `Final2012` release. However, a direct comparison of how far the line moves is not a good comparison. Note that the sum of components is conducted in quadrature, thus the closer to zero the total uncertainty becomes, the harder it is to reduce further by the same factor. In particular, uncertainties which used to be negligible become relevant, and thus any assumptions taken must be re-evaluated every time the uncertainty is reduced by a visible margin.

The same Figure also clearly demonstrates the impact of the GSC on the unknown topology for jets built from EM topo-clusters, where the uncertainty is reduced more when moving from the `Moriond2013` to `Final2012` release than when moving from the `HCP2012` to `Moriond2013` release. Another positive message is that the uncertainties decrease across the entire parameter space (except for very low p_T) with each subsequent JES uncertainties release. This confirms that the assumptions made at each step have always been conservative, and the extra work has always proven to be beneficial.

The same conclusions can be drawn for anti- k_t $R = 0.6$ jets, shown in Figure E.17, although the level of improvement from one release to the next is typically smaller. In particular, the `HCP2012` uncertainties appear to have been a bit optimistic for the information available at the time, although subsequent releases did still improve on the previous results.

E.3 Reductions via matrix eigendecomposition

The 58 (68) independent uncertainty components provided as a part of the `Moriond2013` (`Final2012`) release provide a very detailed but complicated set of systematic uncertainties. This level of precision is usually well beyond what a given analysis wishes to deal with, and can lead to failed fits or other computational limitations. This is particularly true for analysis groups which are conducting searches rather than measurements, where the theoretical or statistical uncertainties are typically larger than the JES systematics, and where a quick and simple approach is required. In order to facilitate such use cases, the set of uncertainty components are reduced via eigendecomposition.

Of the 58 (68) uncertainty components, 47 (56) come from sources related to the absolute *in situ* JES determination. These components are typically small, dominated by a small subset of the full list, and depend purely on p_T . As such, they form a perfect set which can be significantly reduced to only the dominant eigenvectors and a residual term. This has the benefit of preserving the correlation

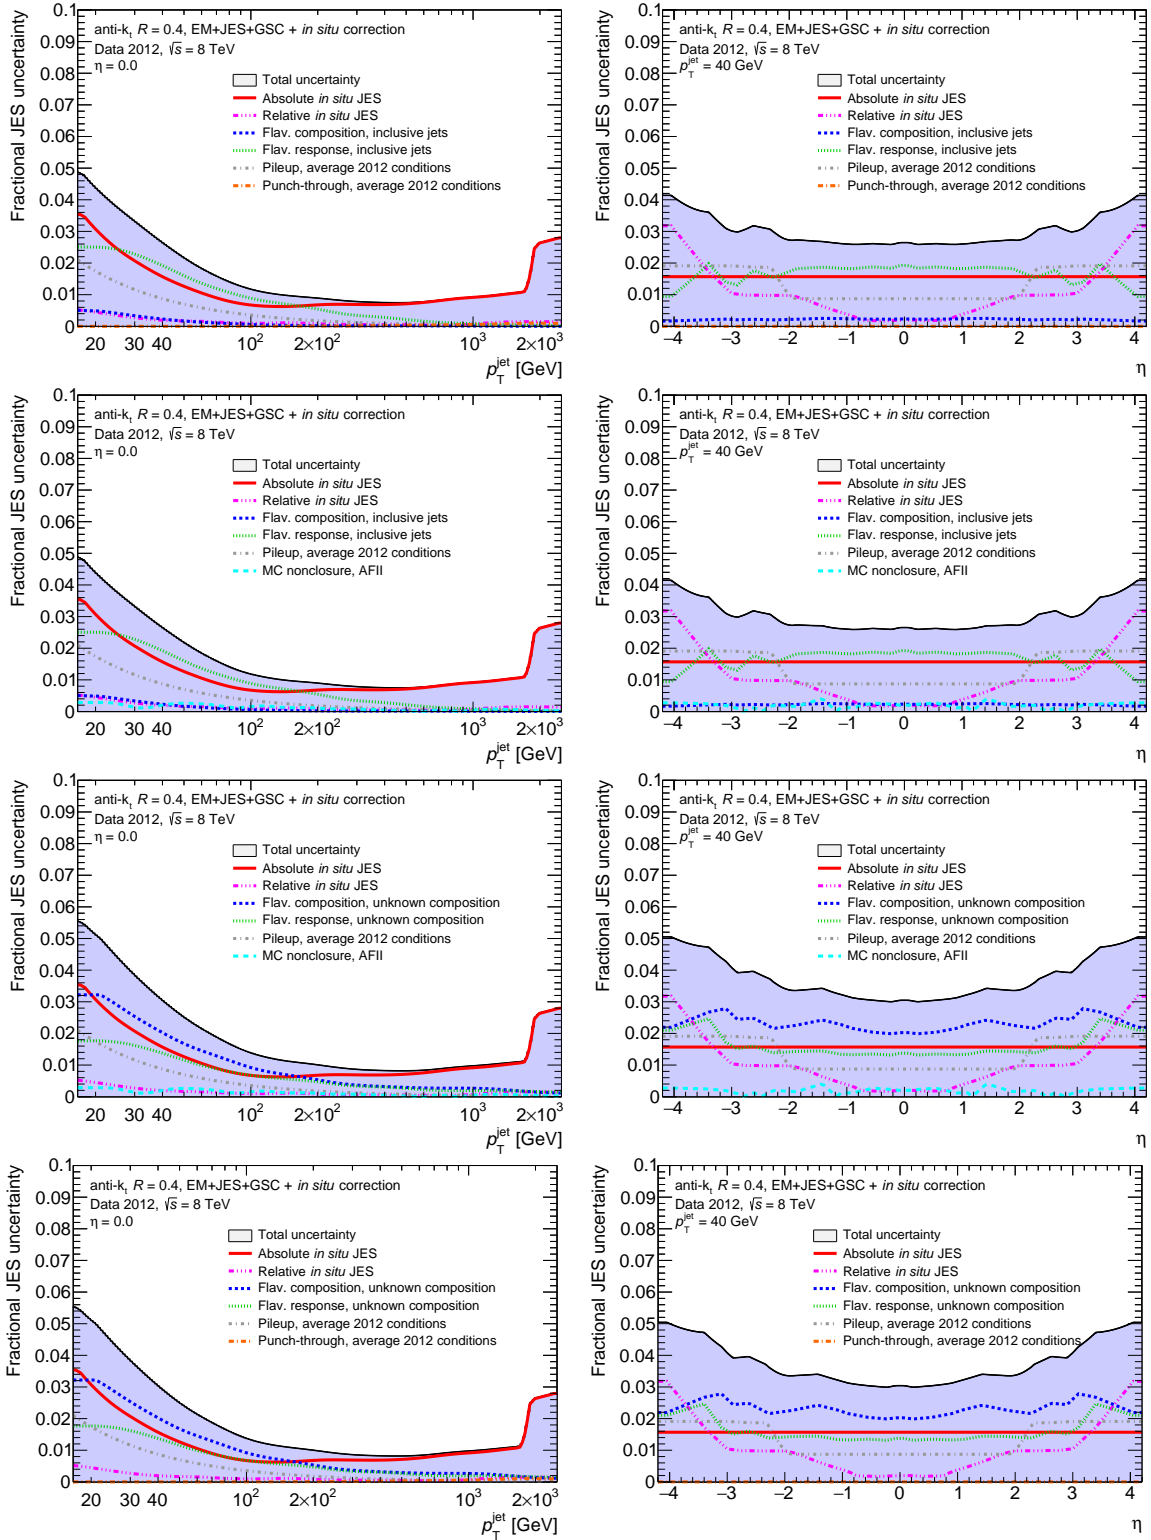


Figure E.12: Final2012 JES uncertainties for anti- k_t $R = 0.4$ jets built from EM topo-clusters, shown vs p_T (left) and η (right), for the full simulation di-jet topology (top), fast simulation dijet topology (upper middle), fast simulation unknown topology (lower middle), and full simulation unknown topology (bottom).

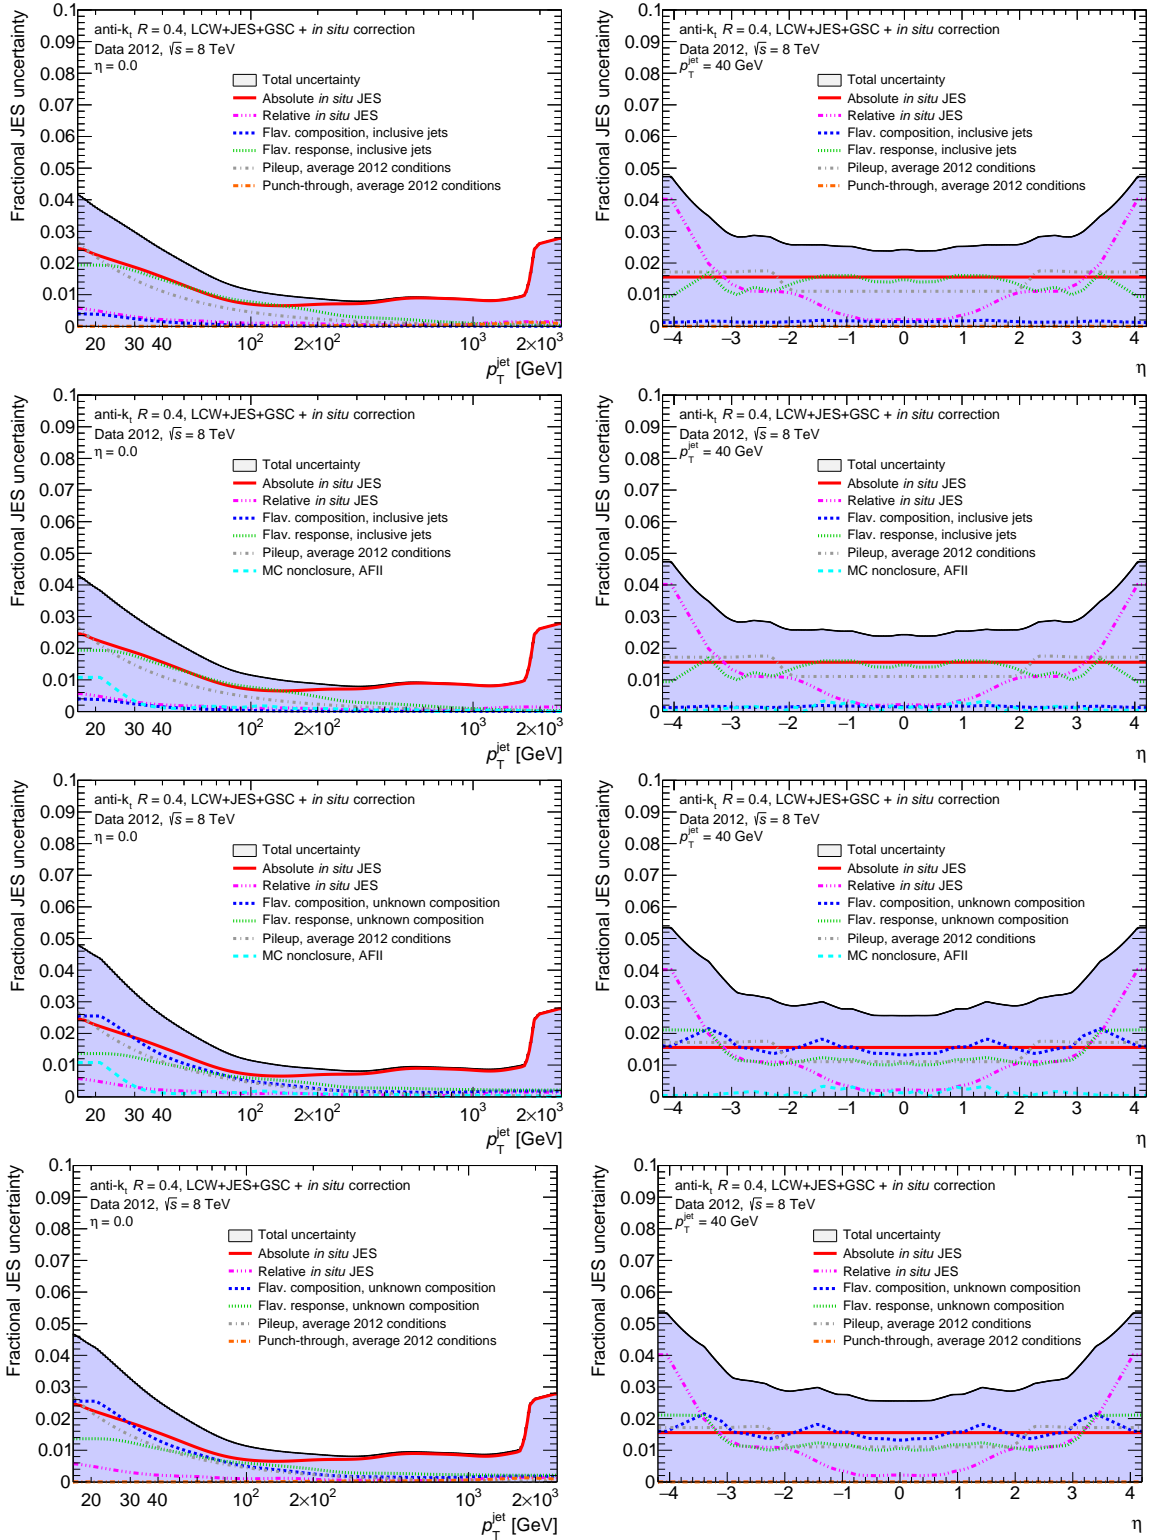


Figure E.13: Final2012 JES uncertainties for anti- k_t $R = 0.4$ jets built from LC topo-clusters, shown vs p_T (left) and η (right), for the full simulation di-jet topology (top), fast simulation dijet topology (upper middle), fast simulation unknown topology (lower middle), and full simulation unknown topology (bottom).

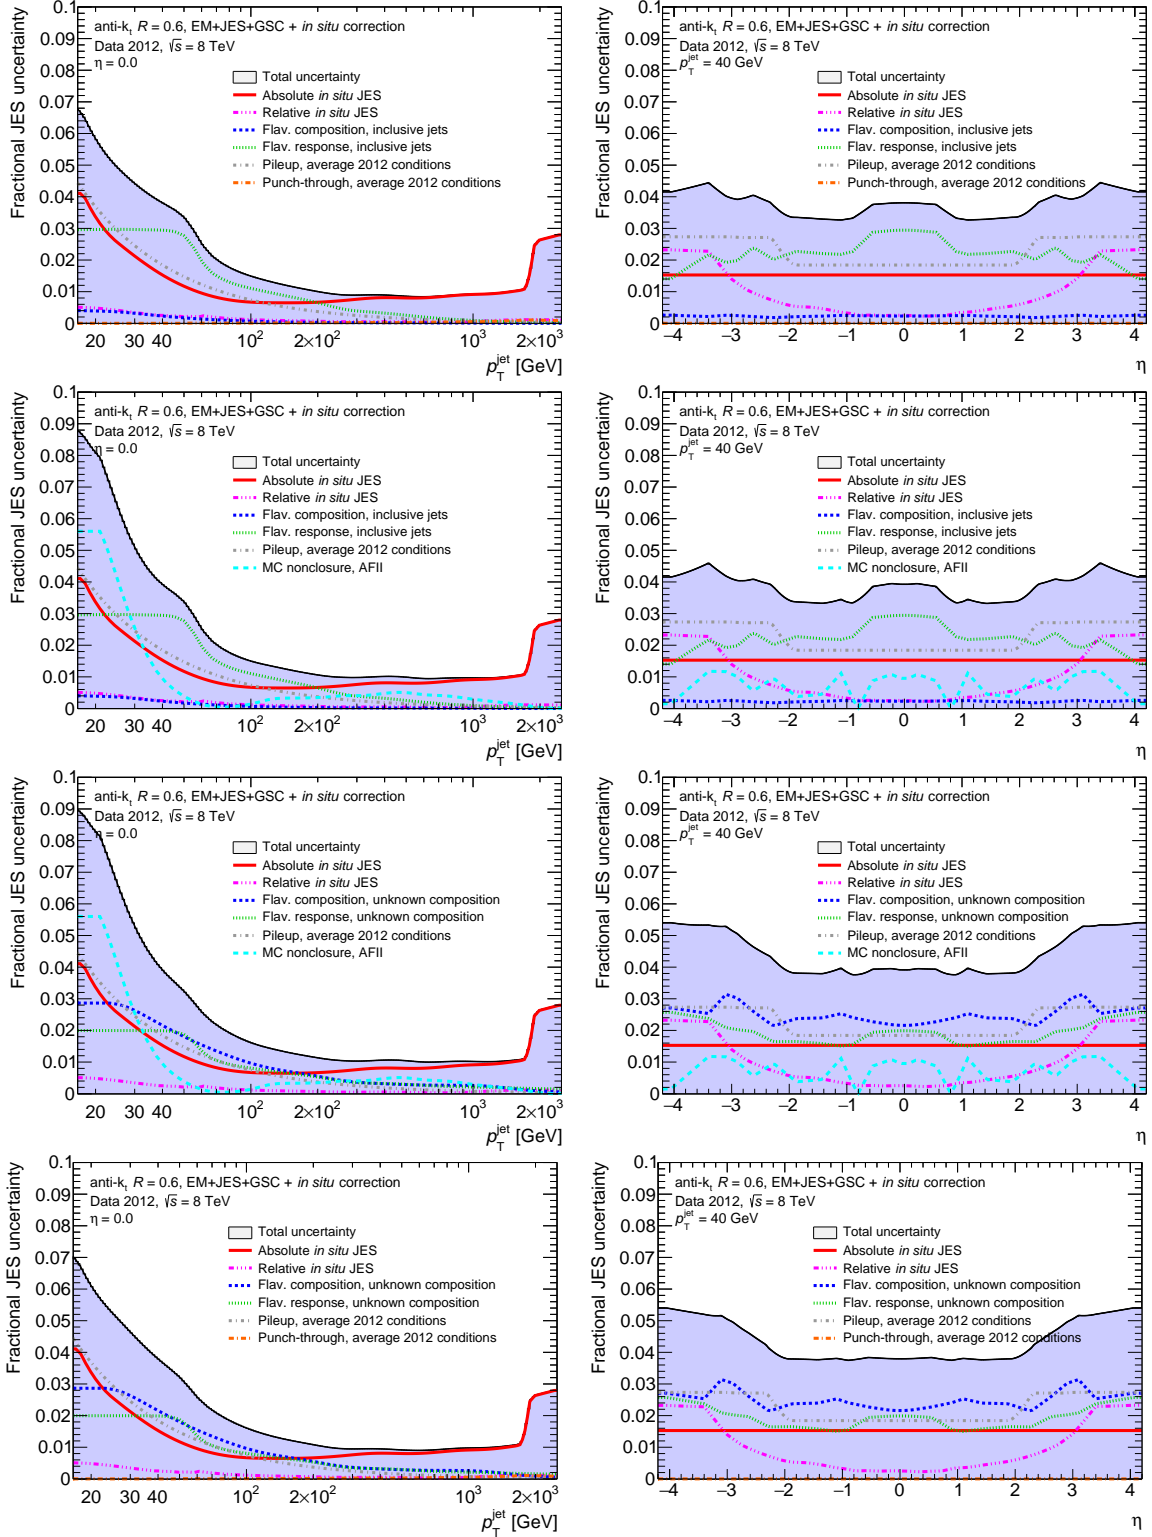


Figure E.14: Final2012 JES uncertainties for anti- k_t $R = 0.6$ jets built from EM topo-clusters, shown vs p_T (left) and η (right), for the full simulation di-jet topology (top), fast simulation dijet topology (upper middle), fast simulation unknown topology (lower middle), and full simulation unknown topology (bottom).

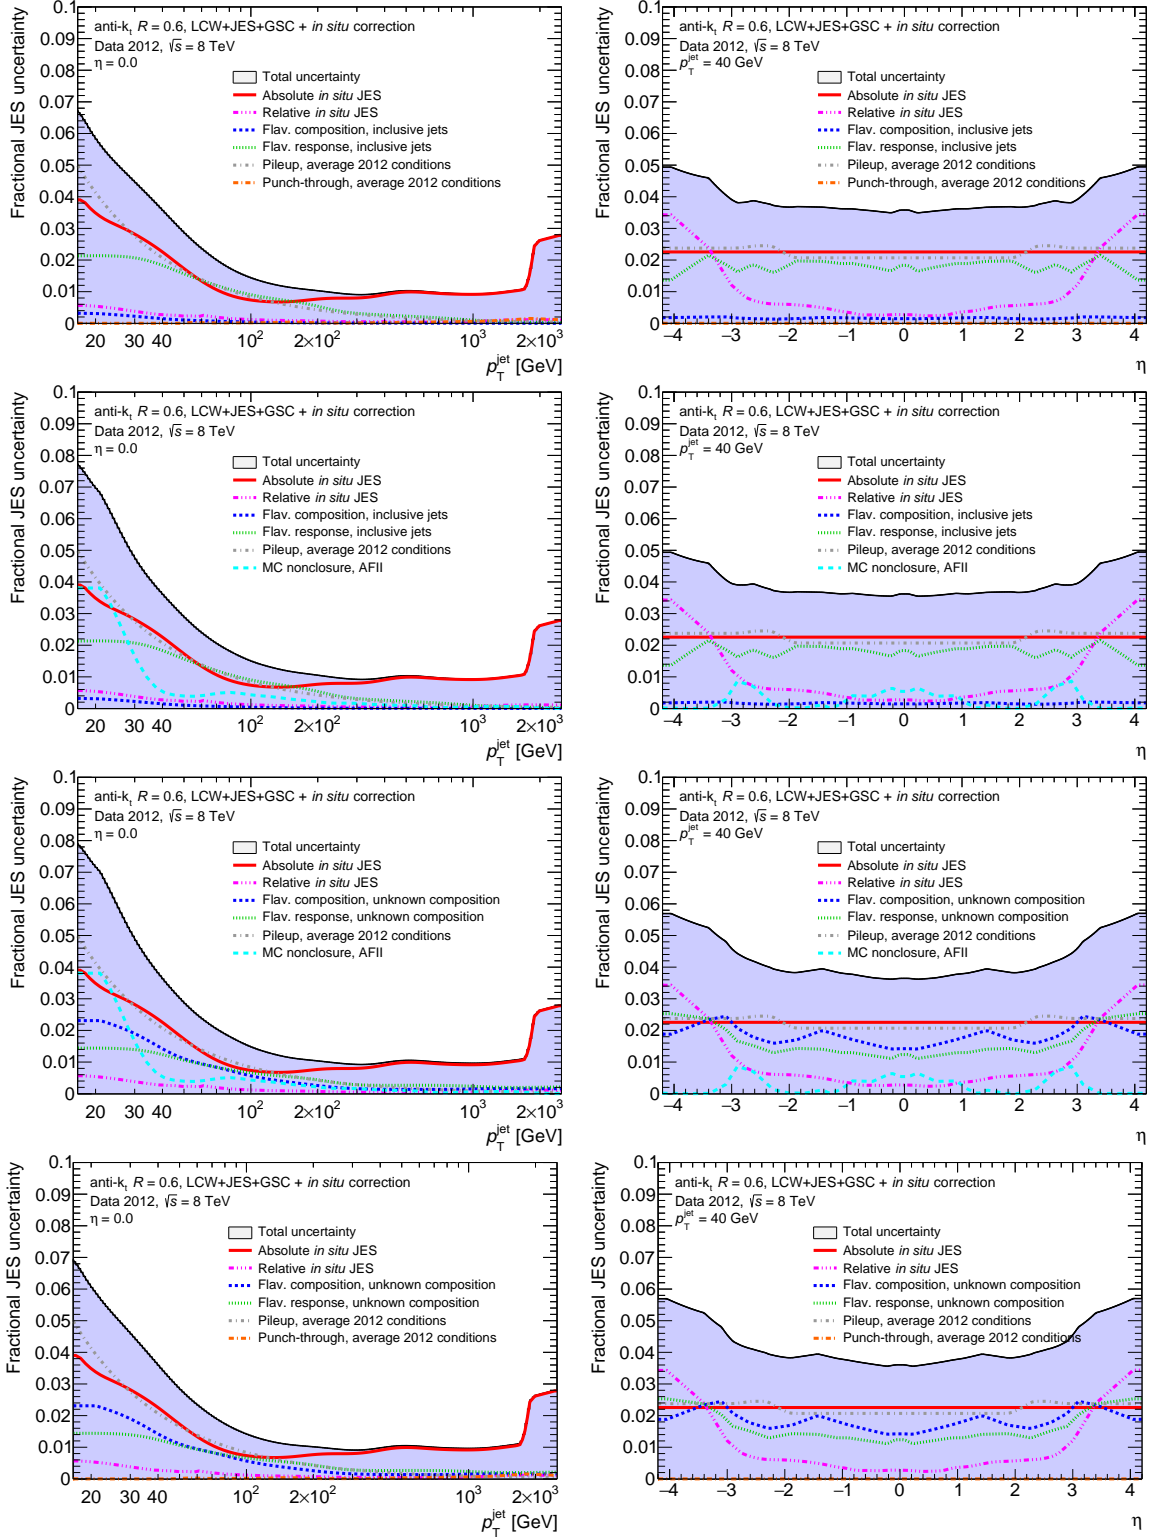


Figure E.15: Final2012 JES uncertainties for anti- k_t $R = 0.6$ jets built from LC topo-clusters, shown vs p_T (left) and η (right), for the full simulation di-jet topology (top), fast simulation dijet topology (upper middle), fast simulation unknown topology (lower middle), and full simulation unknown topology (bottom).

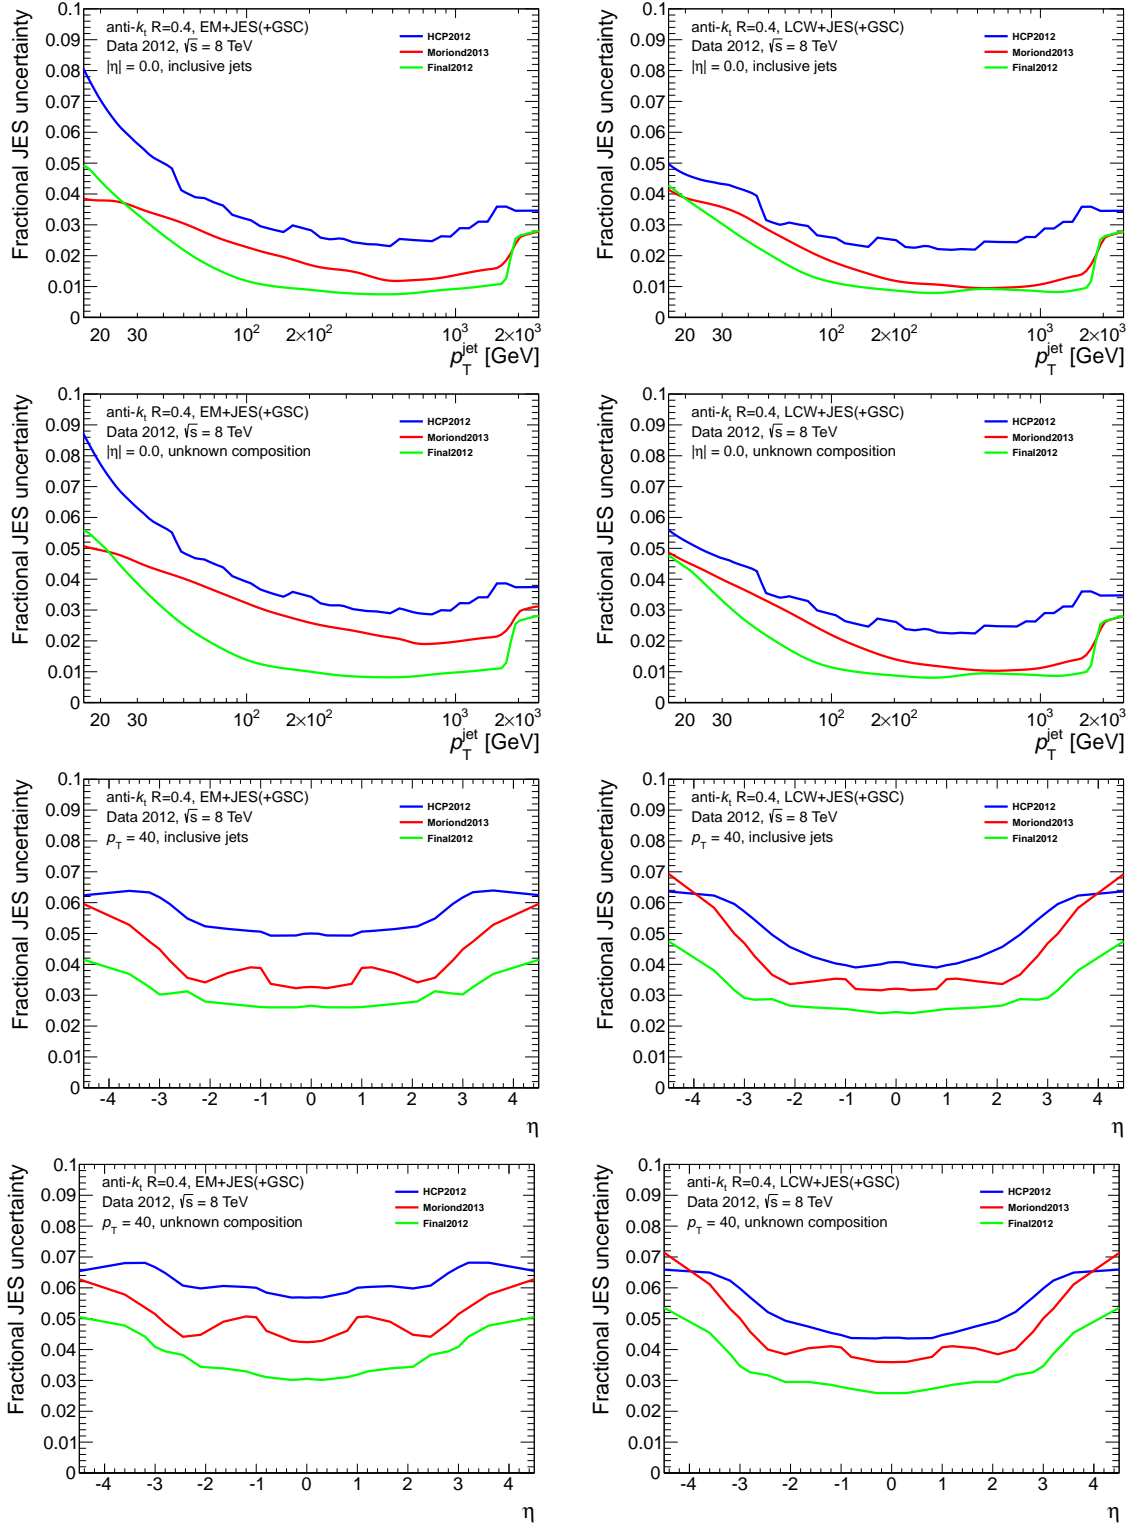


Figure E.16: A comparison of different JES uncertainty releases for anti- k_t $R = 0.4$ jets built from EM topo-clusters (left) and LCW topo-clusters (right) for the di-jet topology (first and third rows) and the unknown topology (second and fourth rows). The first simple release (HCP2012, blue), first proper *in situ* release (Moriond2013, red), and final *in situ* release (Final2012, green) are compared.

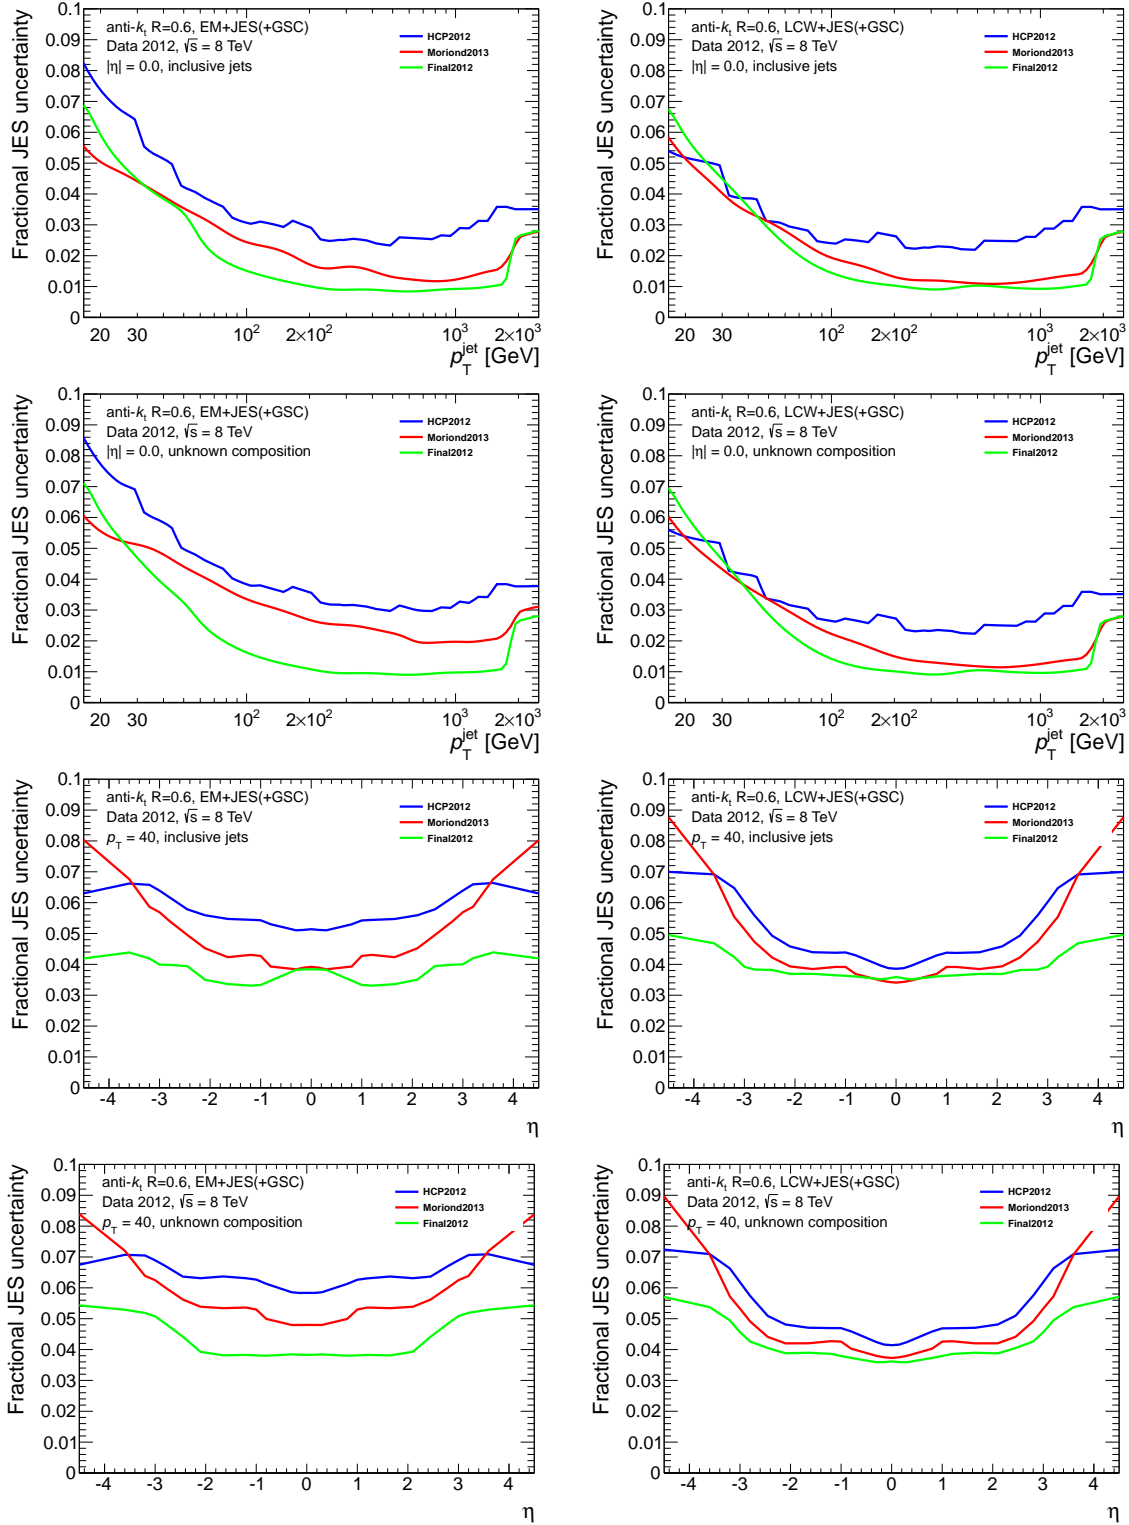


Figure E.17: A comparison of different JES uncertainty releases for anti- k_t $R=0.6$ jets built from EM topo-clusters (left) and LCW topo-clusters (right) for the di-jet topology (first and third rows) and the unknown topology (second and fourth rows). The first simple release (HCP2012, blue), first proper *in situ* release (Moriond2013, red), and final *in situ* release (Final2012, green) are compared.

information over the full parameter space to a configurable level of precision, while leaving the total uncertainty untouched by construction, and significantly reducing the number of required independent variations.

E.3.1 Eigendecomposition description

The eigendecomposition procedure used is detailed in Reference [135]. The intent of this method is to reduce the number of components with minimal loss of precision on the correlations between jets. Note that this is separate from the correlations between uncertainty sources, as they are all assumed to be uncorrelated. The only correlated sources are combined into a single component, which is then uncorrelated from the rest, as is done for the common e/γ energy scale and smearing uncertainties.

All of the absolute *in situ* parameters depend purely on the jet p_T , thus the first step is to construct the covariance matrix from the full set of components as a function of p_T^{jet1} and p_T^{jet2} . This is done in Equation E.3, where it is assumed that a given uncertainty source is fully correlated for all values of p_T , thus the Pearson product-moment correlation coefficient r_p is 1.

$$\begin{aligned} \text{cov}_{\text{tot}}\left(p_T^{\text{jet1}}, p_T^{\text{jet2}}\right) &= \sum_i^{\text{sources}} \text{cov}_i\left(p_T^{\text{jet1}}, p_T^{\text{jet2}}\right) \\ &= \sum_i^{\text{sources}} r_p\left(p_T^{\text{jet1}}, p_T^{\text{jet2}}\right) \cdot \sigma_i\left(p_T^{\text{jet1}}\right) \cdot \sigma_i\left(p_T^{\text{jet2}}\right) \\ \text{cov}_{\text{tot}}\left(p_T^{\text{jet1}}, p_T^{\text{jet2}}\right) &= \sum_i^{\text{sources}} \sigma_i\left(p_T^{\text{jet1}}\right) \cdot \sigma_i\left(p_T^{\text{jet2}}\right) \end{aligned} \quad (\text{E.3})$$

The total covariance matrix derived in Equation E.3 is then diagonalized as in Equation E.4, where D is a diagonal matrix of the eigenvalues (σ_k^2) and columns of S form the eigenvectors (V_k).

$$\text{cov}_{\text{tot}}\left(p_T^{\text{jet1}}, p_T^{\text{jet2}}\right) = S^T D S \quad (\text{E.4})$$

Multiplying each eigenvector by the corresponding eigenvalue creates a new basis of uncertainty components, specified in Equation E.5, corresponding to a rotation from the set of derived uncertainty components to the set of p_T^{jet} bins used in evaluating the covariance matrix. At this point, the rotated basis of uncertainty components has lost its underlying source, and thus is no longer holds a physical meaning. Instead, it is now a nuisance parameter, and can be composed from a linear combination of numerous different related or unrelated sources.

$$\text{cov}_{\text{tot}}(i, j) = \sum_k^{\text{N}_{\text{bins}}^{\text{jet}}} \sigma_k^2 \cdot V_k(i) \cdot V_k(j) \quad (\text{E.5})$$

So far, the number of eigenvectors forming the basis of uncertainty components has gone from the number of sources to the number of p_T^{jet} bins used, which is likely a larger number. This can be reduced by forming an approximation of the covariance matrix, in which the dominant N_{eff} eigenvectors (as determined by the magnitude of the eigenvalues) are retained. The remaining $\text{N}_{\text{bins}}^{\text{jet}} - N_{eff}$ eigenvectors are combined to form a residual component with an associated residual covariance matrix $\text{cov}_{\text{res}}(i, j)$. By combining rather than discarding the residual components, the total uncertainty is preserved. The

form of the approximate covariance matrix is detailed in Equation E.6.

$$\text{cov}_{\text{tot}}(i, j) \approx \sum_k^{N_{\text{eff}}} \sigma_k^2 \cdot V_k(i) \cdot V_k(j) + \text{cov}_{\text{res}}(i, j) \quad (\text{E.6})$$

In this equation, the value N_{eff} can be chosen, thus providing for a means to reduce the number of components to a minimum of 1 (everything in the residual term), at the cost of a loss of correlation information. By manipulating the value of N_{eff} , a balance between correlation precision and a minimal number of components can be found. In order to estimate this loss, the first step is to define the correlation matrix, as per Equation E.7. Then, the difference between the correlation matrix formed by the full set of input uncertainty sources and the correlation matrix formed by the reduced set of nuisance parameters quantifies the loss of correlation information due to the reduction procedure. Examples of the application of this process and the resulting correlation losses will be applied in the following Sections.

$$\begin{aligned} \text{corr}_{\text{tot}}(p_{\text{T}}^{\text{jet1}}, p_{\text{T}}^{\text{jet2}}) &= \frac{\text{cov}_{\text{tot}}(p_{\text{T}}^{\text{jet1}}, p_{\text{T}}^{\text{jet2}})}{\sqrt{\text{cov}_{\text{tot}}(p_{\text{T}}^{\text{jet1}}, p_{\text{T}}^{\text{jet1}}) \cdot \text{cov}_{\text{tot}}(p_{\text{T}}^{\text{jet2}}, p_{\text{T}}^{\text{jet2}})}} \\ \text{corr}_{\text{tot}}(p_{\text{T}}^{\text{jet1}}, p_{\text{T}}^{\text{jet2}}) &= \frac{\sum_i^{\text{sources}} \sigma_i(p_{\text{T}}^{\text{jet1}}) \cdot \sigma_i(p_{\text{T}}^{\text{jet2}})}{\sqrt{\left[\sum_i^{\text{sources}} \sigma_i^2(p_{\text{T}}^{\text{jet1}})\right] \left[\sum_i^{\text{sources}} \sigma_i^2(p_{\text{T}}^{\text{jet2}})\right]}} \end{aligned} \quad (\text{E.7})$$

E.3.2 Global reductions

Reductions from the full set of 47 (56) initial *in situ* components down to $N_{\text{eff}} = 5$ for the `Moriond2013` and `Final2012` releases respectively are referred to as global reductions. This name reflects the fact that all of the components are allowed to be combined into each eigenvector. This will be contrasted with category reductions in the next Section. The resulting global reduced parameters are provided in Figure E.18. This Figure demonstrates how the residual component is very small for this choice of $N_{\text{eff}} = 5$, with a maximum uncertainty value of 0.3%, and demonstrates how 47 or 56 inputs can be reduced down to 5 nuisance parameters with a negligible residual.

While the residual term may have a minimal uncertainty magnitude, this does not guarantee that the loss of correlation information is also negligible. In order to study this, the difference of correlation matrices is built, following Equation E.7. The pre-reduction correlation matrices are provided in Figure E.19 in order to demonstrate the starting point, clearly demonstrating the regimes where each of the *in situ* methods is active (and thus a strong correlation is observed). The resulting correlation losses for the `Moriond2013` and `Final2012` releases are shown in Figure E.20. A summary of the maximum loss of correlation information for each jet algorithm and for both releases is provided in Table E.3.

E.3.3 Category reductions

Category reductions also work with the full set of 47 (56) *in situ* inputs, but do not allow them all to mix. The reduction procedure can create an arbitrary linear combination of any of the uncertainty source inputs that it is given, but often some minimal degree of knowledge is required, and thus these generic nuisance parameters are not useful. For example, when combining different analyses, it is important to

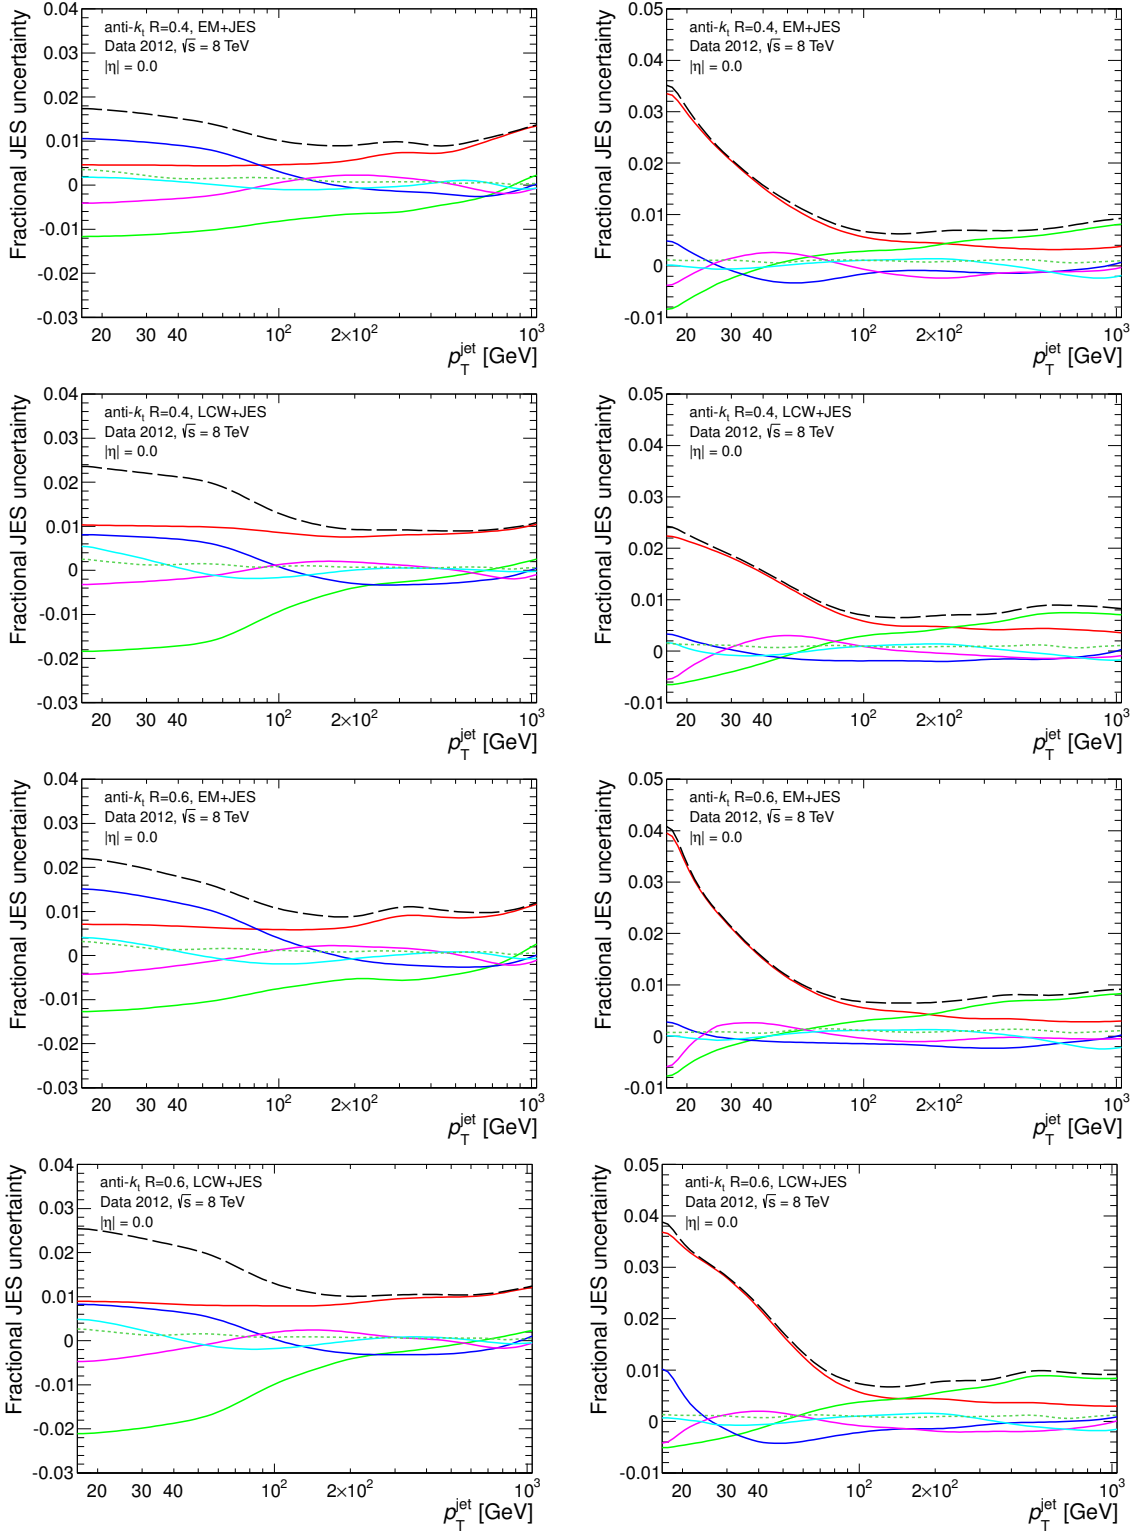


Figure E.18: Reduced approximate representations for $N_{\text{eff}} = 5$ of the Moriond2013 (left) and Final2012 (right) releases, from an input set of 47 or 56 *in situ* components respectively. The solid lines are dominant eigenvectors, the coloured short-dashed line is the residual term, and the black long-dashed line is the total uncertainty of the input components. Similar plots for 2011 data are available in [135].

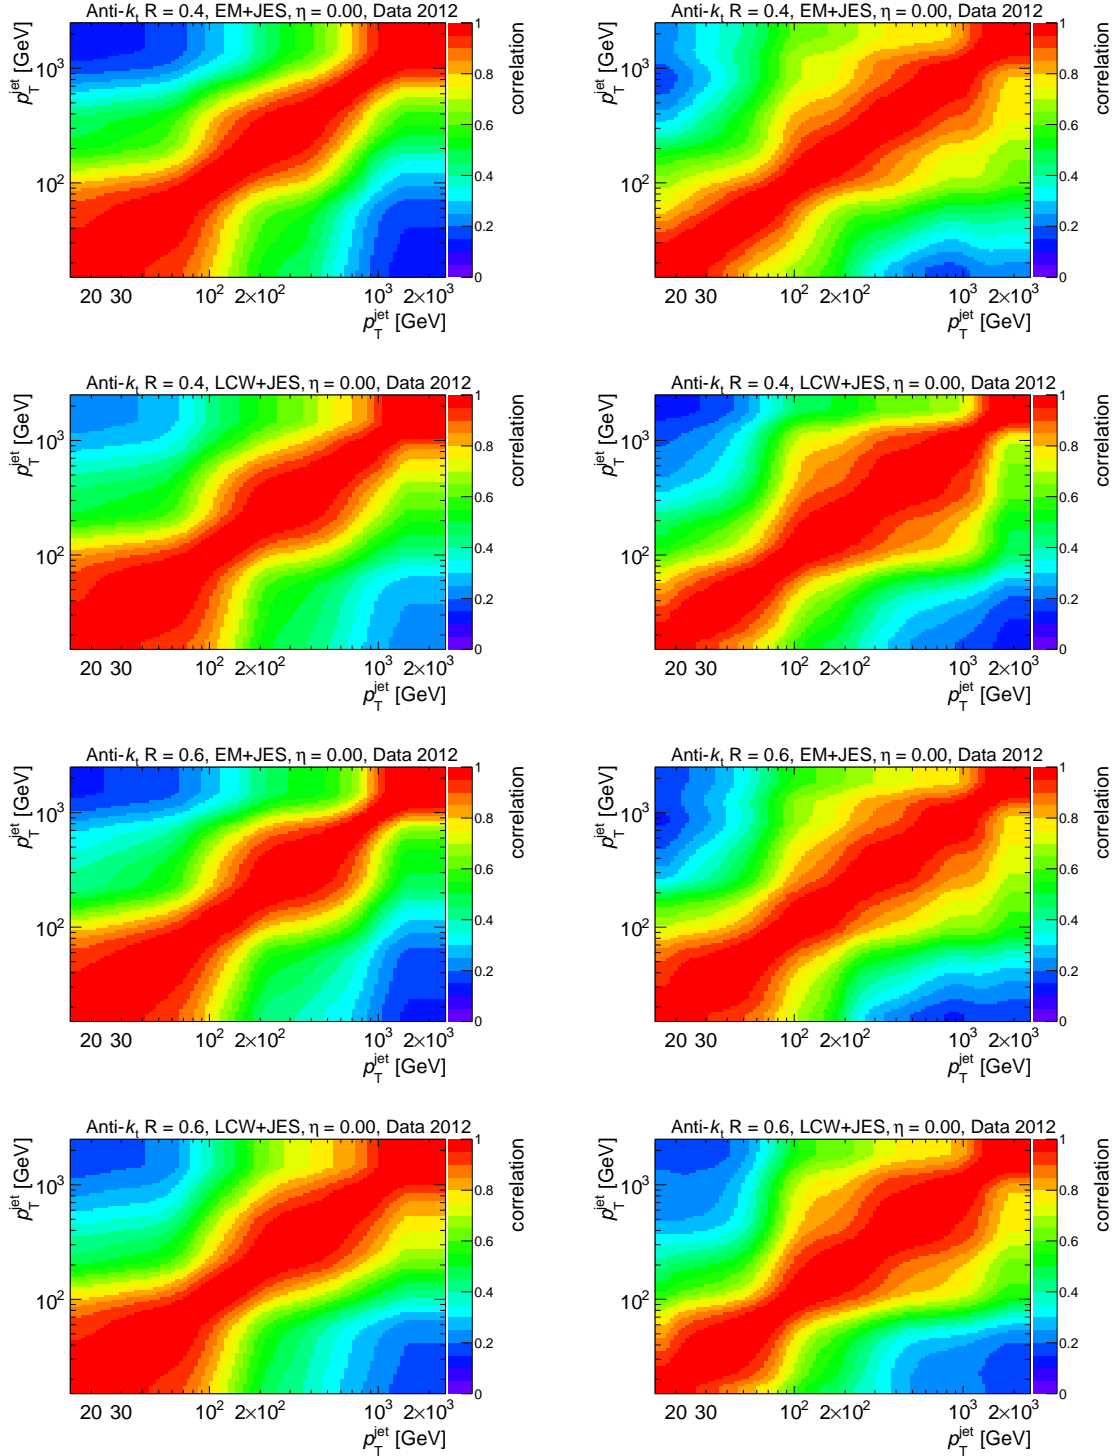


Figure E.19: Nominal correlation matrices before reduction for the Moriond2013 (left) and Final2012 (right) releases, only showing the input set of 47 or 56 *in situ* components respectively. The addition of other unrelated components such as flavour uncertainties slightly reduces the level of correlation, but the general trends remain.

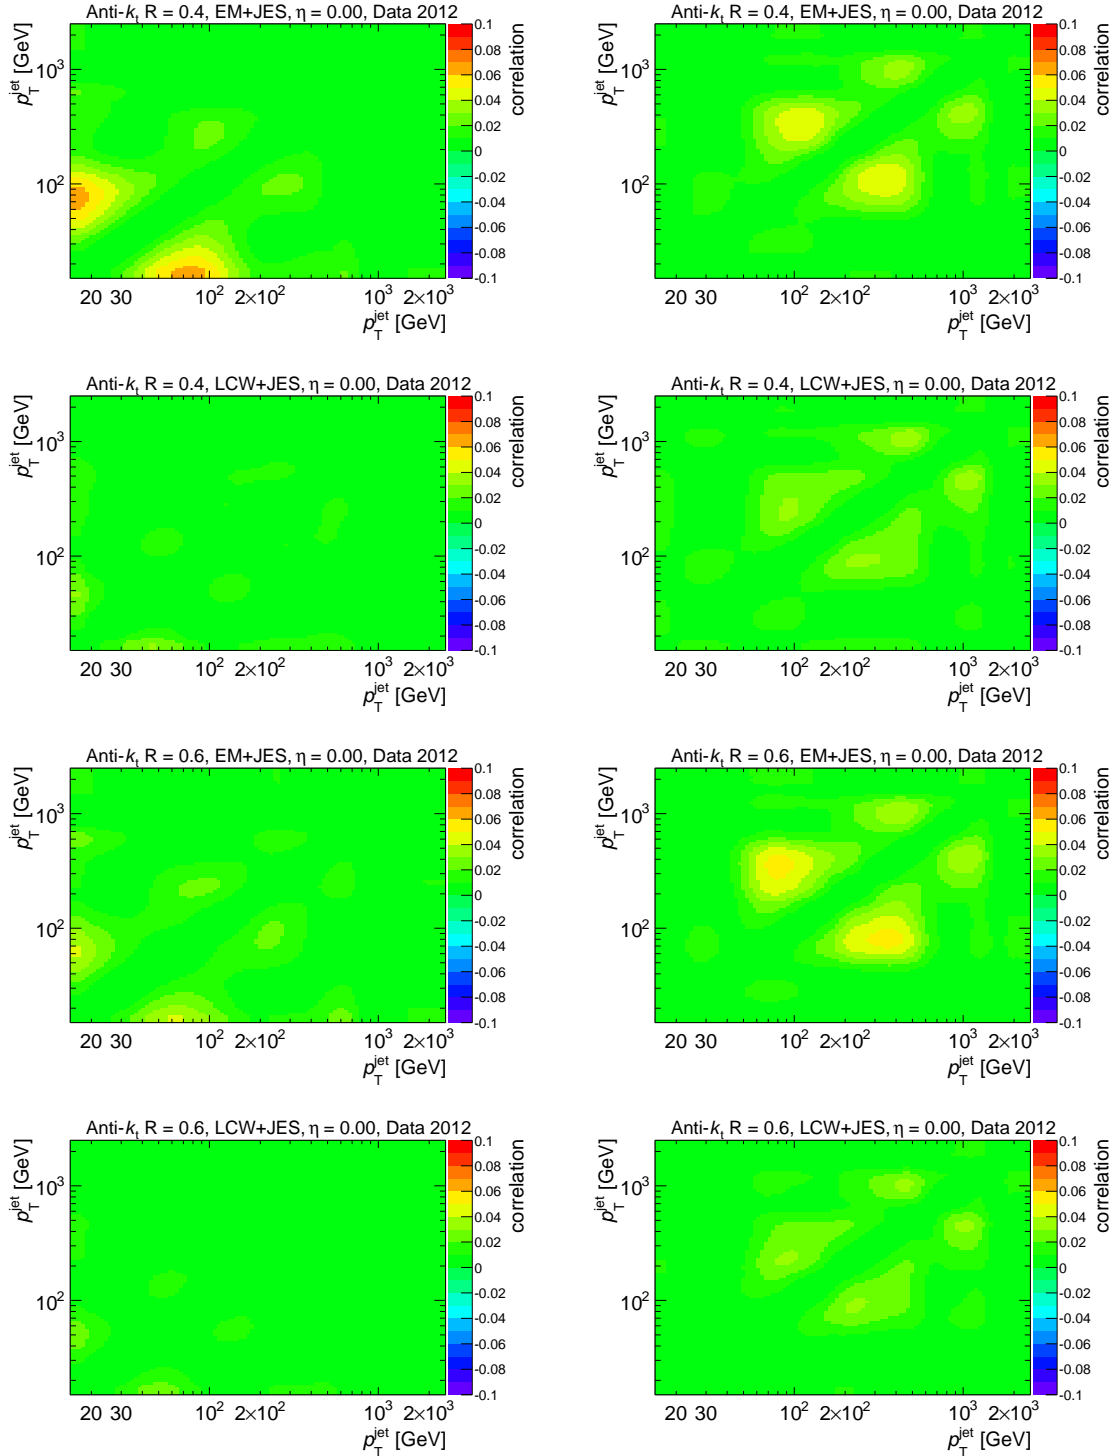


Figure E.20: Global reduction correlation losses (reduced-original) for $N_{\text{eff}} = 5$ for the **Moriond2013** (left) and **Final2012** (right) releases, from an input set of 47 or 56 *in situ* components respectively. Smaller uncertainty magnitudes for the **Final2012** JES make it more challenging to remove components without correlation losses due to sources which were once negligible becoming relevant. Maximum differences are provided in Table E.3.

Jet definition		Max correlation loss (Moriond2013)			Max correlation loss (Final2012)		
R	cluster scale	Diff. [%]	p_T^{jet1} [GeV]	p_T^{jet2} [GeV]	Diff. [%]	p_T^{jet1} [GeV]	p_T^{jet2} [GeV]
0.4	EM	6.5	15	79	4.8	113	349
0.4	LC	3.0	15	45	3.7	474	1075
0.6	EM	4.1	17	61	5.4	79	349
0.6	LC	2.5	16	50	3.2	451	1022

Table E.3: The maximum loss in correlations due to the application of the global eigendecomposition procedure to both the Moriond2013 and Final2012 JES uncertainty releases. The point in p_T^{jet} -space where this maximum loss occurs is also specified.

know what uncertainties are statistical, and so on. Even more importantly, when combining results between ATLAS and CMS such as for the top mass measurement, it is necessary to know what uncertainties are related to the detector, as these systematics should not be correlated across experiments.

In order to satisfy this requirement, category reductions were created. Rather than performing the eigendecomposition simultaneously over the full set of *in situ* inputs, the decomposition is instead performed per category, with the categories specified in Section 6.4 and Tables E.1 and E.2. The optimal effective values determined for the Moriond2013 (Final2012) release are $N_{\text{eff}}^{\text{detector}}$ of 2 (2), $N_{\text{eff}}^{\text{modelling}}$ of 3 (3), $N_{\text{eff}}^{\text{mixed}}$ of 1 (3), and $N_{\text{eff}}^{\text{stat}}$ of 2 (3), resulting in 8 (11) effective parameters and 4 (4) residual parameters, for a total of 12 (15) nuisance parameters. This is an increase with respect to the 6 parameters obtained after the global reduction for both the Moriond2013 and Final2012 releases, but the retained general uncertainty source makes it preferable in many cases. This procedure went on to become the recommendation for combining the JES uncertainties between the ATLAS and CMS collaborations, as recently published in Reference [143]. This procedure has already been used in 2012 by many groups, becoming the default ATLAS recommendation, and as used for the Higgs boson measurement combinations.

The resulting category reduced parameters are provided for the anti- k_t $R = 0.4$ jet algorithm built from LC clusters in Figure E.21. This Figure demonstrates how the residual component is always very small for these choices of $N_{\text{eff}}^{\text{category}}$, with a maximum uncertainty value of below 0.2%, and demonstrates how 47 or 56 inputs can be reduced to 12 or 15 while retaining basic source information with a negligible residual.

While the residual terms may all have a minimal uncertainty magnitude, this does not guarantee that the loss of correlation information is also negligible. In order to study this, the difference of correlation matrices is built, following Equation E.7. Due to the significantly reduced set of inputs for the reduction within a given category, it is more difficult to obtain a high degree of correlation precision without retaining a significant number of components. This is apparent in Figure E.22, where much larger regions of p_T^{jet} -space have suffered a correlation loss at the level of 5% when compared to the global reduction, despite retaining at least twice as many components. A summary of the maximum loss of correlation information for each jet algorithm and for both releases is provided in Table E.4.

E.4 Alternative correlation scenarios

When constructing the full set of uncertainty sources, occasionally the level of correlation between two different components is not clear, and an assumption is made that the components are either correlated or uncorrelated. In order to understand the impact of this correlation uncertainty, alternative correlation

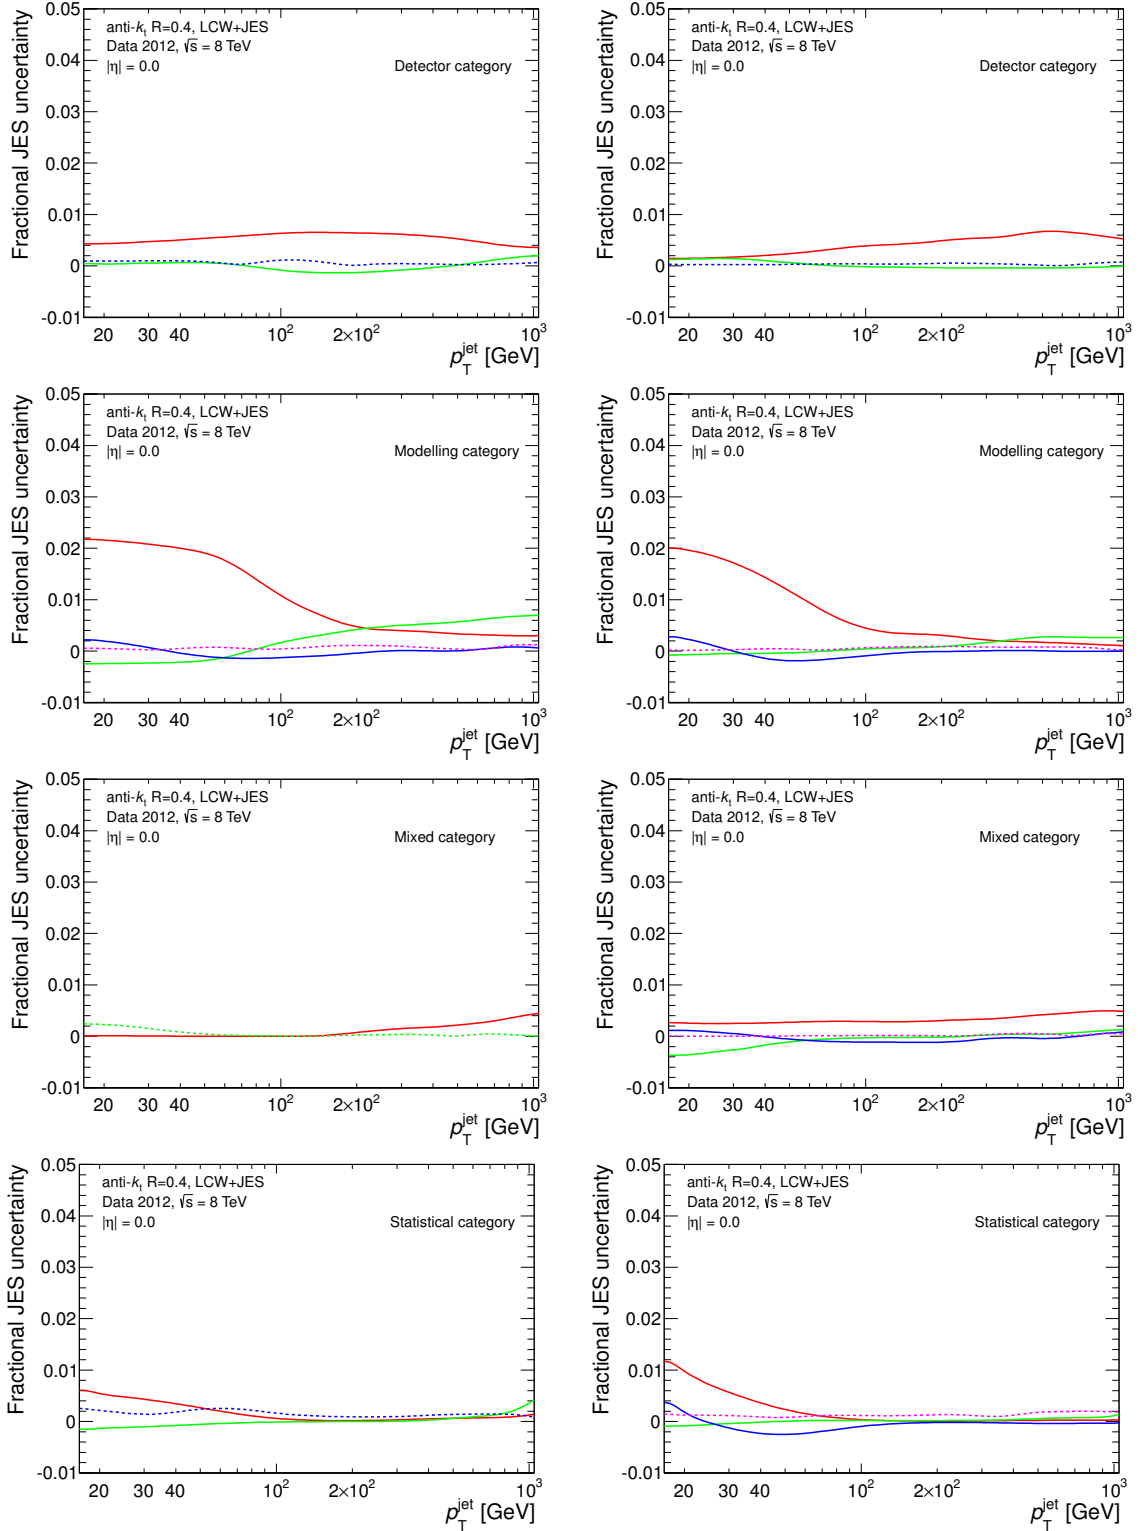


Figure E.21: Approximate representations for category reductions of the Moriond2013 (left) and Final2012 (right) releases, from an input set of 47 or 56 *in situ* components respectively. The effective parameters are shown per category. The solid lines are dominant eigenvectors and the coloured short-dashed line is the residual term. Similar plots for 2011 data are available in Reference [135].

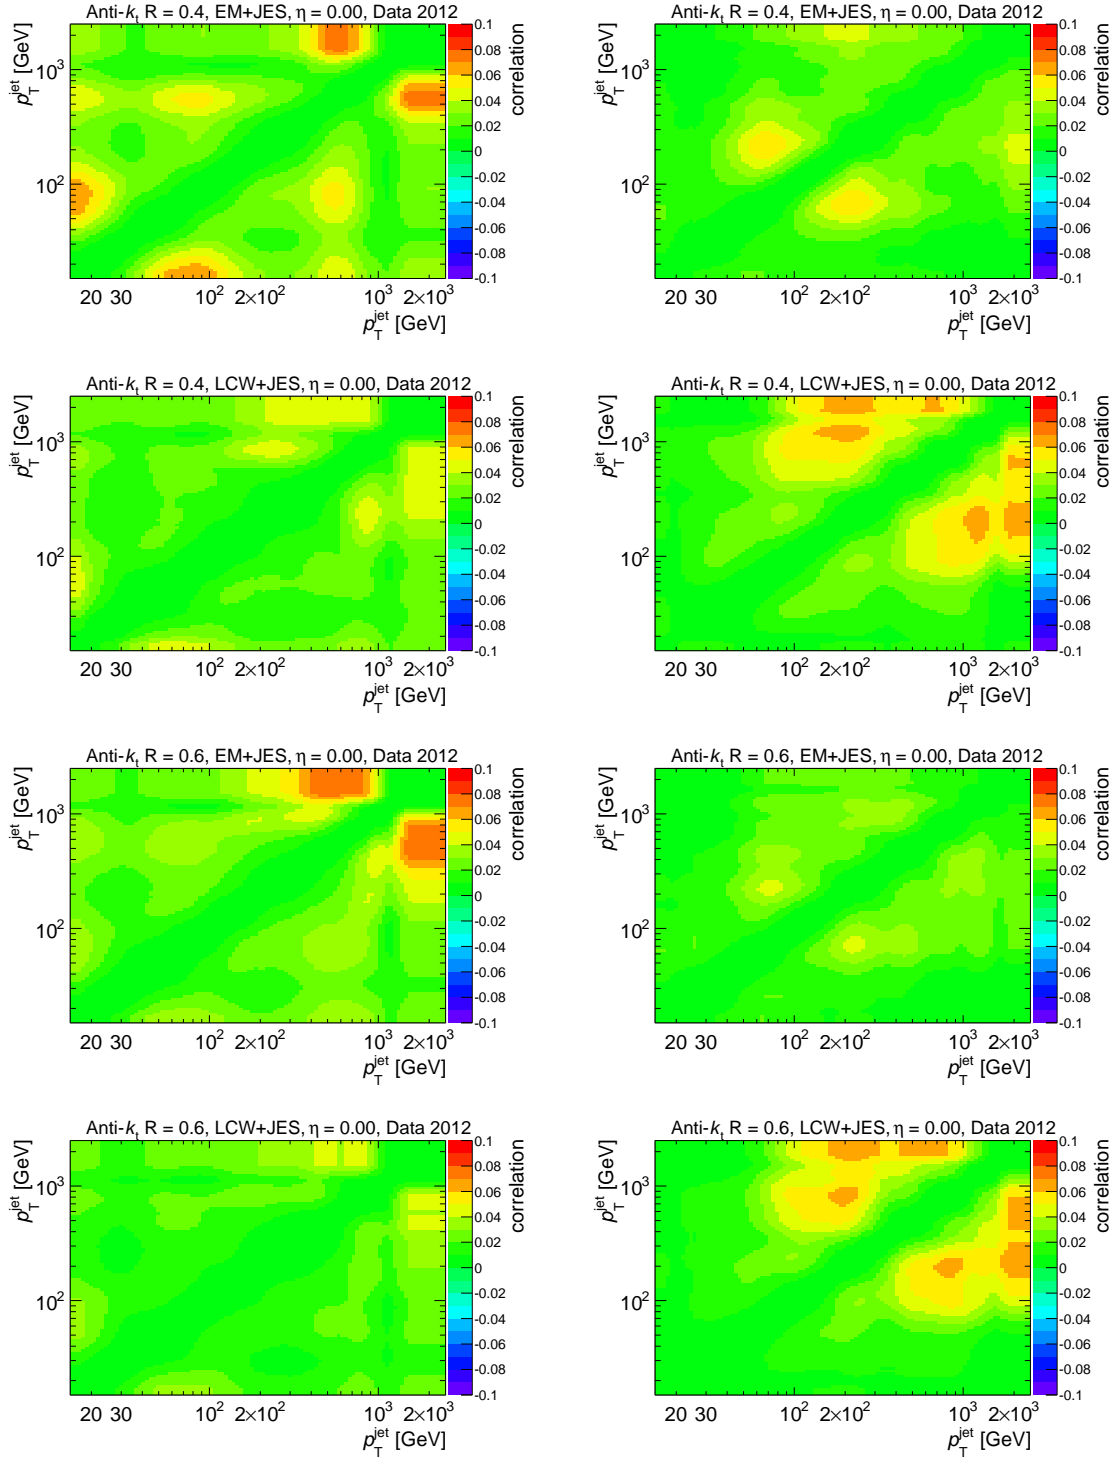


Figure E.22: Category reduction correlation losses (reduced-original) for the Moriond2013 (left) and Final2012 (right) releases, from an input set of 47 or 56 *in situ* components respectively. The category reduction was performed for the Moriond2013 (Final2012) with $N_{\text{eff}}^{\text{detector}}$ of 2 (2), $N_{\text{eff}}^{\text{modelling}}$ of 3 (3), $N_{\text{eff}}^{\text{mixed}}$ of 1 (3), and $N_{\text{eff}}^{\text{stat}}$ of 2 (3). Maximum differences are provided in Table E.3.

Jet definition		Max correlation loss (Moriond2013)			Max correlation loss (Final2012)		
R	cluster scale	Diff. [%]	p_T^{jet1} [GeV]	p_T^{jet2} [GeV]	Diff. [%]	p_T^{jet1} [GeV]	p_T^{jet2} [GeV]
0.4	EM	7.5	582	2091	5.9	68	220
0.4	LC	4.6	232	876	6.8	209	1254
0.6	EM	7.8	525	2316	4.3	71	220
0.6	LC	4.3	752	2200	6.9	220	1888

Table E.4: The maximum loss in correlations due to the application of the category eigendecomposition procedure to both the Moriond2013 and Final2012 JES uncertainty releases. The point in p_T^{jet} -space where this maximum loss occurs is also specified. The category reduction was performed for the Moriond2013 (Final2012) with $N_{\text{eff}}^{\text{detector}}$ of 2 (2), $N_{\text{eff}}^{\text{modelling}}$ of 3 (3), $N_{\text{eff}}^{\text{mixed}}$ of 1 (3), and $N_{\text{eff}}^{\text{stat}}$ of 2 (3).

scenarios are provided.

While every effort is taken to avoid the need for approximations, in reality the derivation of the JES and associated uncertainties occasionally requires the use of assumptions. The largest assumption made is that all of the components listed in Tables E.1 and E.2 are strictly independent, and that there is no degree of correlation between them. In reality, several of the components are expected to be partially correlated, but it is difficult to quantify the degree of partial correlation. In order to understand the impact of these unknown correlations and other assumptions have on the quoted JES uncertainties, alternative correlation scenarios are provided. Users which are performing measurements sensitive to the JES are recommended to fully propagate the uncertainties from both the nominal and alternative scenarios, and to use whichever result is more conservative. The average ATLAS user does not carry out such tests at this time, although it may become more common in the future.

E.4.1 The Moriond2013 alternative scenario

Some of the sources which were assumed to be uncorrelated, but are expected to be partially correlated, are combined into a single fully correlated component. This brackets the expected value of the uncertainty due to these sources (quadrature vs linear addition) and the strength of the correlation (multiple components vs a single component). This is done for the following types of sources:

- MC generator related uncertainties within the Z-jet, γ -jet, and multi-jet balance *in situ* methods
- Sources related to radiation suppression within the Z-jet, γ -jet, and multi-jet balance *in situ* methods
- Out-of-cone biases from the Z-jet and γ -jet balance *in situ* methods

This approach is analogous to what was done for the 2011 dataset, as described in Reference [135], with two differences. In the 2011 estimation, the γ -jet balance approach had a pileup systematic, which was not relevant for the Moriond2013 release. As such, the Z-jet and γ -jet correlated component relating to pileup rejection was removed. The second difference is related to the LAr energy scale. In 2011, the LAr energy scale was only available as a single fully correlated component. It was expected that this was not realistic, and so in 2011 this was split into two components following an ad-hoc split, thus accounting for the partial correlation of the component. In 2012, the LAr energy scale was split into many components, and thus this was no longer required. As mentioned in Section 6.4, there are now sources which are considered to be correlated between Z-jet and γ -jet balance approaches, but there are

Nominal correlations		Stronger correlations	
Name	Category	Name	Category
Z-jet balance: MC generator	Modelling	Correlated MC generators	Special
γ -jet balance: MC generator	Modelling		
Multi-jet balance: fragmentation	Mixed		
Z-jet balance: extrapolation	Modelling	Correlated suppression and selection	Modelling
Z-jet balance: radiation suppression	Modelling		
γ -jet balance: radiation suppression	Modelling		
Multi-jet balance: α selection	Modelling		
Multi-jet balance: β selection	Modelling		
Z-jet balance: out-of-cone radiation	Modelling	Correlated out-of-cone radiation	Modelling
γ -jet balance: out-of-cone radiation	Modelling		

Table E.5: Changes in the component breakdown between the nominal and alternative (stronger) correlation scenario for the `Moriond2013` JES uncertainty. Only components which changed have been listed. For the complete list of components and their individual descriptions, see Table E.1.

Jet definition		Max correlation change (stronger - nominal)		
R	cluster scale	Difference [%]	p_T^{jet1} [GeV]	p_T^{jet2} [GeV]
0.4	EM	24.8	15	1022
0.4	LC	40.7	15	1132
0.6	EM	37.1	15	1191
0.6	LC	53.2	15	1022

Table E.6: The maximum change in correlations for the `Moriond2013` JES stronger correlations scenario, with respect to the nominal.

multiple independent LAr energy scale sources. This provides the desired separation, and thus the split from 2011 was no longer applicable.

A summary of the changes between the nominal and alternative (stronger) correlation scenario for the `Moriond2013` JES uncertainty is provided in Table E.5.

In general, a large impact is observed when switching between the nominal and stronger correlation scenarios. This is primarily due to an overall shift in the JES uncertainty combined with the creation of a fully-correlated link between the dominant high p_T components and relevant but not dominant low p_T components, as seen in Figure E.23. This leads to the creation of decoherence between the high and low p_T^{jet} regime when compared to the nominal configuration, seen as a large change in correlations between such jets in Figure E.24, and quantified at the 25 to 50% level in Table E.6. Note that the largest difference in correlations is always for the lowest p_T bin considered of 15 GeV versus a very high p_T bin of over 1 TeV, thus confirming this observation.

E.4.2 The Final2012 alternative scenarios

The `Final2012` alternative scenarios expand on the previous variations of assumptions with both a “stronger” and “weaker” correlation scenario. The names are a bit misleading, as they do not necessarily bracket the nominal scenario. Instead, they are stronger and weaker with respect to each other. The variations considered can be broken up into three types, depending on what they affect:

1. Varied in both of the alternative correlation scenarios
2. Varied only in the stronger correlation scenario

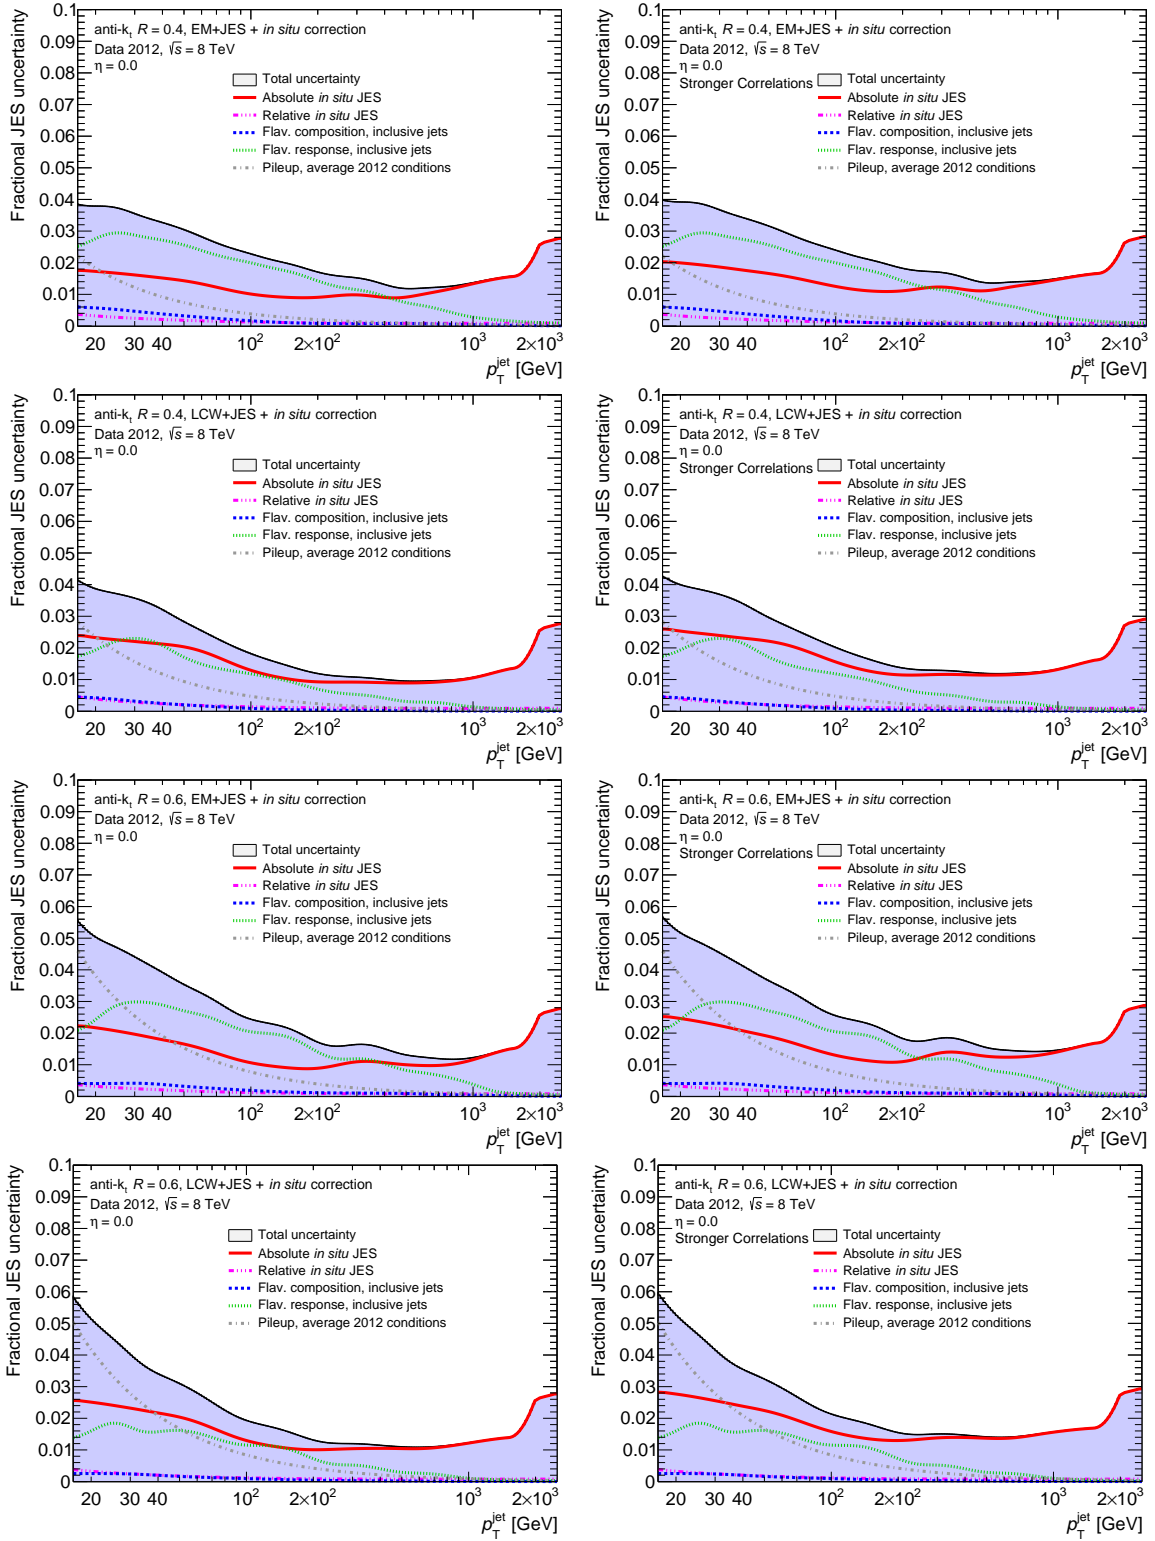


Figure E.23: Summary plots of the total JES uncertainty in the di-jet topology for the Moriond2013 release, for the nominal correlations (left) and stronger correlations (right). The absolute *in situ* JES is the only component modified.

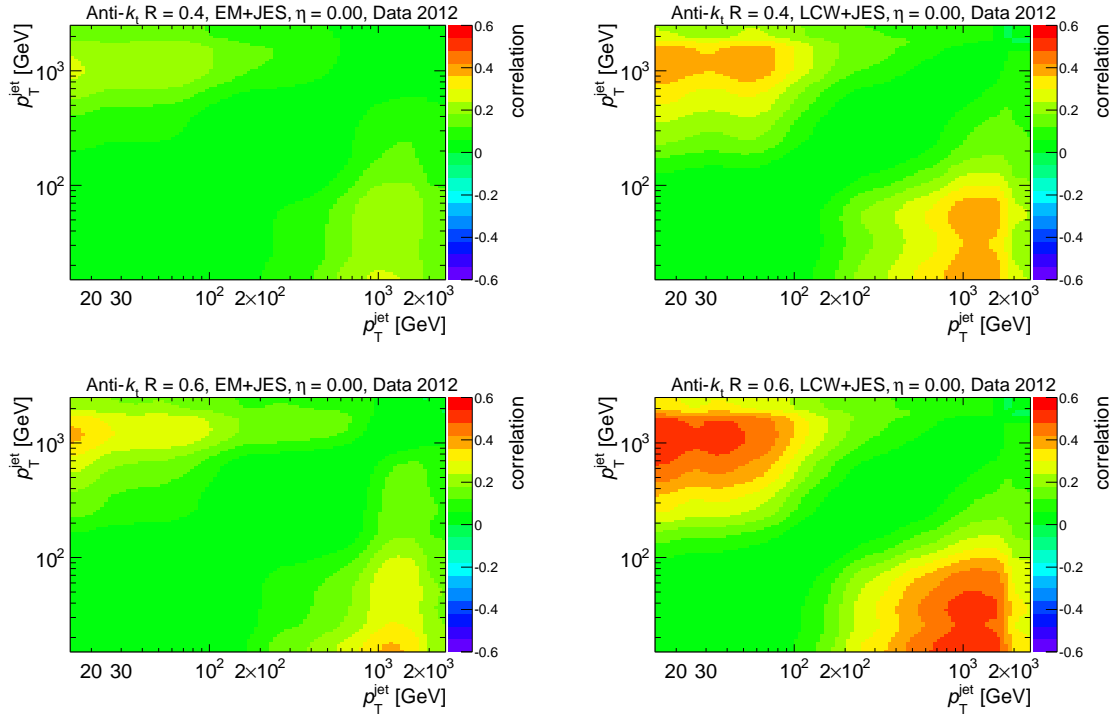


Figure E.24: Correlation changes for the alternative correlation scenario (stronger-nominal) in the di-jet topology for the Moriond2013 JES release. The maximum correlation difference values are in Table E.6.

3. Varied only in the weaker correlation scenario

The first category represents the global variations, of which two are considered. There are two different rebinning choices which were used for reducing the contribution of statistical uncertainties to the *in situ* measurements. They lead to very similar results at medium to high p_T , but moderately different results at low p_T . The nominal uncertainty configuration uses the more conservative (loose) rebinning, while both of the alternative scenarios use the more aggressive (tight) rebinning. The second change is the transition range from the multi-jet balance *in situ* studies to the single-particle response measurement. The nominal scenario uses a quick transition, while the alternative scenarios change more slowly.

The second category of changes, which are only present in the stronger correlation scenario, is mostly analogous to what was done for the Moriond2013 release. In addition to these combinations from correlated sub-components, the categorization is also changed. All components which were listed as “mixed” are instead changed to “modelling”. Furthermore, the baseline value of the light quark and gluon responses for the flavour uncertainties (\mathcal{R}_q and \mathcal{R}_g) are changed from Pythia8 to Herwig++.

For the changes which are only present in the weaker correlation scenario, the opposite categorization change is done, where all “modelling” components are changed to “mixed”. The more radical change relates to a division of the dominant two-point systematics in the analysis. Two-point systematics are the construction of an uncertainty from the difference between two generators. These uncertainties are very dangerous, as the use of two generators significantly impacts the value of the uncertainty. Using a different combination of two generators will almost certainly result in a different value. In addition, it

Weaker correlations		Nominal correlations		Stronger correlations	
Name	Category	Name	Category	Name	Category
—	—	Z-jet: MC generator	Modelling	MC generators	Mixed
—	—	γ -jet: MC generator	Modelling		
—	—	Multi-jet: fragmentation	Mixed		
—	—	Z-jet: extrapolation	Modelling	Suppression/selection	Mixed
—	—	Z-jet: radiation suppression	Modelling		
—	—	γ -jet: extrapolation	Modelling		
—	—	γ -jet: radiation suppression	Modelling		
—	—	Multi-jet: α selection	Modelling		
—	—	Multi-jet: β selection	Modelling		
—	—	Z-jet: out-of-cone radiation	Modelling	Out-of-cone radiation	Mixed
—	—	γ -jet: out-of-cone radiation	Modelling		
η modelling split #1	Modelling	η -intercalibration modelling	Modelling	—	—
η modelling split #2	Modelling				
b -jet split #1	Modelling	b -jet scale	Modelling	—	—
b -jet split #2	Modelling				
Flav comp split #1	Modelling	Flavour composition	Modelling	Switch to $\mathcal{R}_{q,g}^{\text{Herwig++}}$	Mixed
Flav comp split #2	Modelling				
Flav resp split #1	Modelling	Flavour response	Modelling	Switch to $\mathcal{R}_{q,g}^{\text{Herwig++}}$	Mixed
Flav resp split #2	Modelling				
Use tight rebinning (aggressive)		<i>In situ</i> components use loose rebinning		Use tight rebinning (aggressive)	
Slow single-particle transition		Quick multi-jet to single-particle switch		Slow single-particle transition	
“Mixed” become “modelling”		Both “mixed” and “modelling” terms		“Modelling” become “mixed”	

Table E.7: Changes in the component breakdown between the nominal and both alternative correlation scenario for the Moriond2013 JES uncertainty. Only components which changed have been listed. For the complete list of components and their individual descriptions, see Table E.2. Components which were not changed for a given alternative scenario are denoted as such with “—”.

is not clear what the level of correlation between the generators is. The magnitude of the uncertainty cannot be easily addressed, but the correlations can be probed to first order. Uncertainties relating to the η -intercalibration modelling, b -jet scale, flavour composition, and flavour response are all split, as they all are examples of two-point systematics (except for the flavour composition, which is a one-point systematic). The first half of the split is calculated by multiplying the uncertainty by a value ranging from 0, at the minimum p_T^{jet} value considered, to 1, at the maximum value considered. The second half of the split is then calculated by subtracting the first half from the original component in quadrature, thus always maintaining the same total uncertainty.

A summary of all of these changes is provided in Table E.7. It is clear that there are significantly more variations considered for the Final2012 JES uncertainty release than in the past, thus providing a much more rigorous treatment of the correlation uncertainties. These variations also visibly change the total JES uncertainty from one release to the next, as demonstrated in Figures E.25 and E.26 for anti- k_t $R = 0.4$ and $R = 0.6$ respectively.

There are multiple simultaneous effects occurring when comparing these plots, thus it is important to look at the impact piece by piece. First off, the changed transition rate to the single particle extrapolation is visible in all plots when comparing the nominal to either of the alternatives at $p_T^{\text{jet}} \approx 1.5$ TeV. At low p_T , comparing the absolute *in situ* JES of the nominal to weaker scenario demonstrates the significant impact of the changed rebinning scheme from the conservative to more aggressive approach. The total uncertainty is reduced by about 1% for jets built from EM topo-clusters, while jets built from LC topo-clusters are less affected.

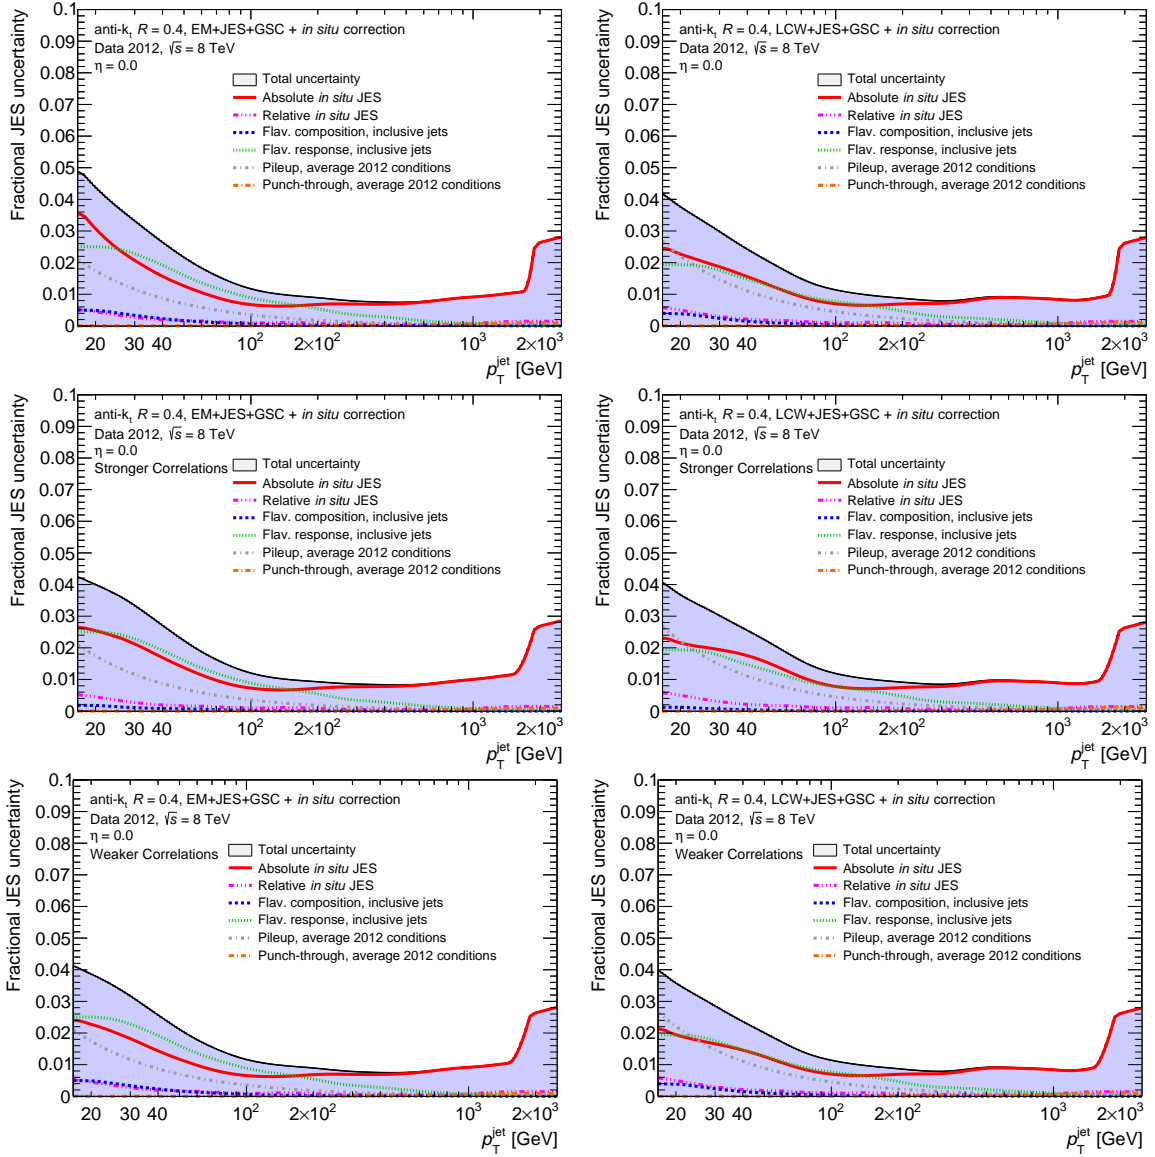


Figure E.25: Summary plots of the total JES uncertainty in the di-jet topology for the Final2012 release, showing anti- k_t $R = 0.4$ jets built from EM topo-clusters (left) and LCW topo-clusters (right). Scenarios shown cover the nominal correlations (top), stronger correlations (middle), and weaker correlations (bottom).

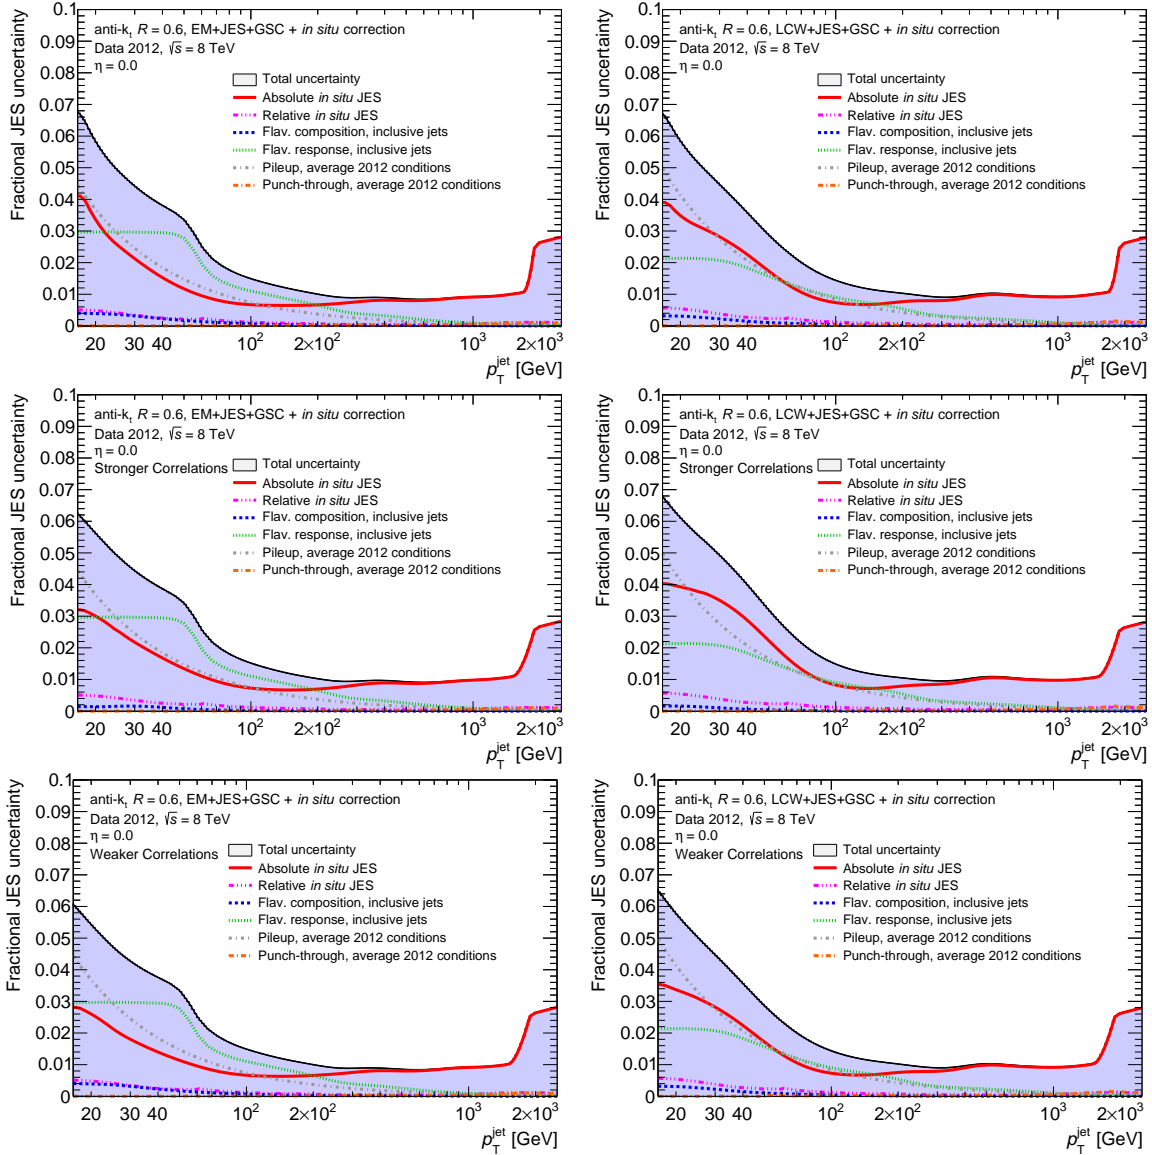


Figure E.26: Summary plots of the total JES uncertainty in the di-jet topology for the Final2012 release, showing anti- k_t $R = 0.6$ jets built from EM topo-clusters (left) and LCW topo-clusters (right). Scenarios shown cover the nominal correlations (top), stronger correlations (middle), and weaker correlations (bottom).

Next, comparing the absolute *in situ* JES between the weaker and stronger scenario quantifies the impact of the correlated treatment of several parameters. This is seen to generally be a smaller effect than the rebinning, which is in contrast to the `Moriond2013` where the same change produced a large effect. The reason for this difference is that the absolute *in situ* JES has shifted from being dominated by jet-related systematics to other effects, such as the photon energy scale, as has been previously discussed. As such, treating sub-dominant components as correlated has a smaller impact than when they were the dominant components.

Comparing the stronger to the nominal or weaker scenarios allows one to understand the impact of the switch from `Pythia8` to `Herwig++` for the baseline flavour response, $\mathcal{R}_{q,g}$. The flavour composition uncertainty is very small for the di-jet topology, but there is a clear difference when comparing this component between the scenarios. This has a much larger impact on the total uncertainty when the unknown composition is used.

Splitting the two-point systematics is not visible in this set of plots, as the split was done in a way such that the total uncertainty was unchanged. Similarly, the change of component categories between mixed and modelling is not visible here, and will prove to be more relevant when combining different results.

The next step is to study the change in actual correlations. The direct differences between the nominal and either stronger or weaker scenarios, as presented in Figure E.27, are dominated by the change of rebinning procedures at low p_T and the transition speed from the multi-jet balance to single-particle approach at high p_T . This is expected, as these changes significantly modify the values of the uncertainties in those regimes. Additional structures linking low to high p_T jets when comparing the stronger to nominal scenarios, due to the correlated combination of components, are present. These have a reduced impact with respect to the `Moriond2013` JES, as expected. A further comparison of the nominal to weaker scenarios also indicates the impact of splitting 2-point systematics, which broadly weakens the correlations across the full p_T^{jet} parameter space.

The change of rebinning strategy and single-particle transition can be thought of as modifications to the actual uncertainty level itself, rather than purely correlation variations. As such, given that the stronger and weaker scenarios both include these differences, a comparison between them provides a better idea of the impact purely due to changes intended to modify the correlations. This is provided in Figure E.28, where it is important to note that the colour scale has been shrunk to see more details now that the single-particle transition has been removed. The largest differences are then seen to once again be due to the combination of correlated components, as for the `Moriond2013` release.

A table summarizing the largest correlation differences observed between the different alternative scenario combinations is provided in Table E.8. The resulting correlation differences in the low to high p_T region, due to the combination of correlated components, is seen to be consistent at the level of approximately 30%. This is a smaller deviation than seen for the `Moriond2013` release, as expected given that the significance of the contributions has reduced. The only exception is the anti- k_t $R = 0.4$ jet algorithm with EM topo-cluster inputs, where the difference in correlations has increased by a small amount. In the `Moriond2013` JES, the flavour uncertainty was dominant at low p_T for this jet type, but the GSC has reduced this uncertainty considerably. As such, the relevance of the absolute *in situ* JES jet-related systematics at low p_T has grown for this scenario.

In summary, the `Final2012` JES uncertainty has two alternative correlation scenarios, which probe the dependence of both the total uncertainty and the jet-to-jet correlations on the assumptions made

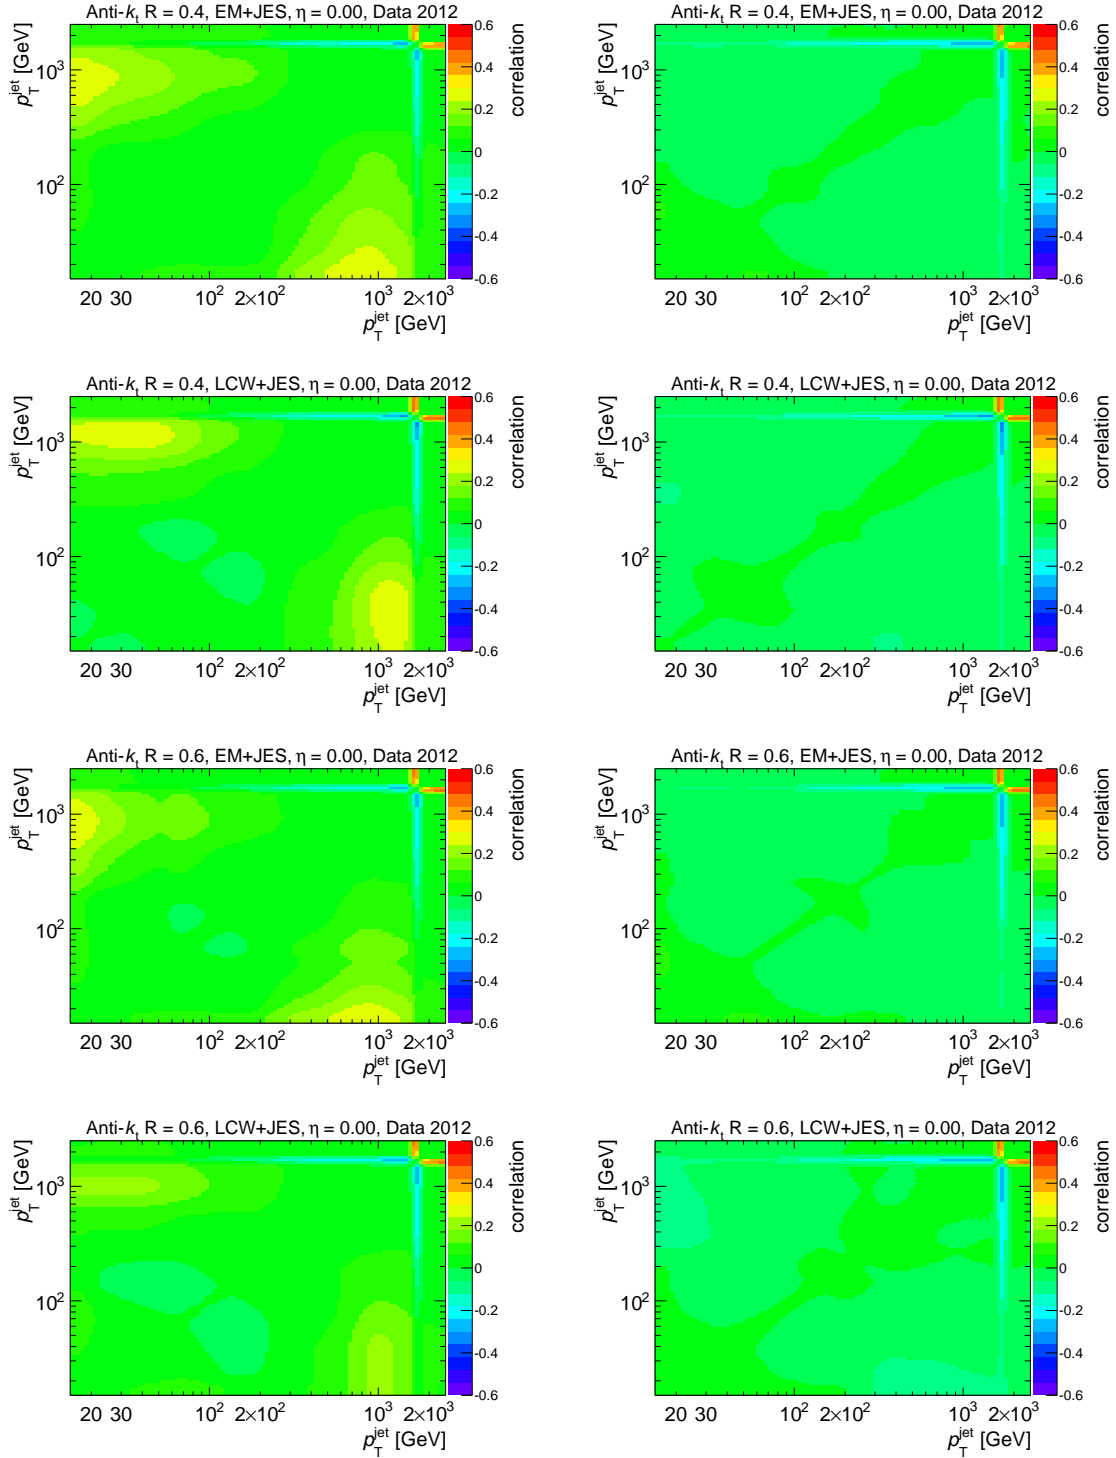


Figure E.27: Correlation changes for the alternative correlation scenarios in the di-jet topology for the Final2012 JES release. Stronger minus nominal correlations (left) and weaker minus nominal correlations (right) are shown. Direct differences between the stronger and weaker scenarios are provided in Figure E.28. The maximum correlation difference values are in Table E.8.

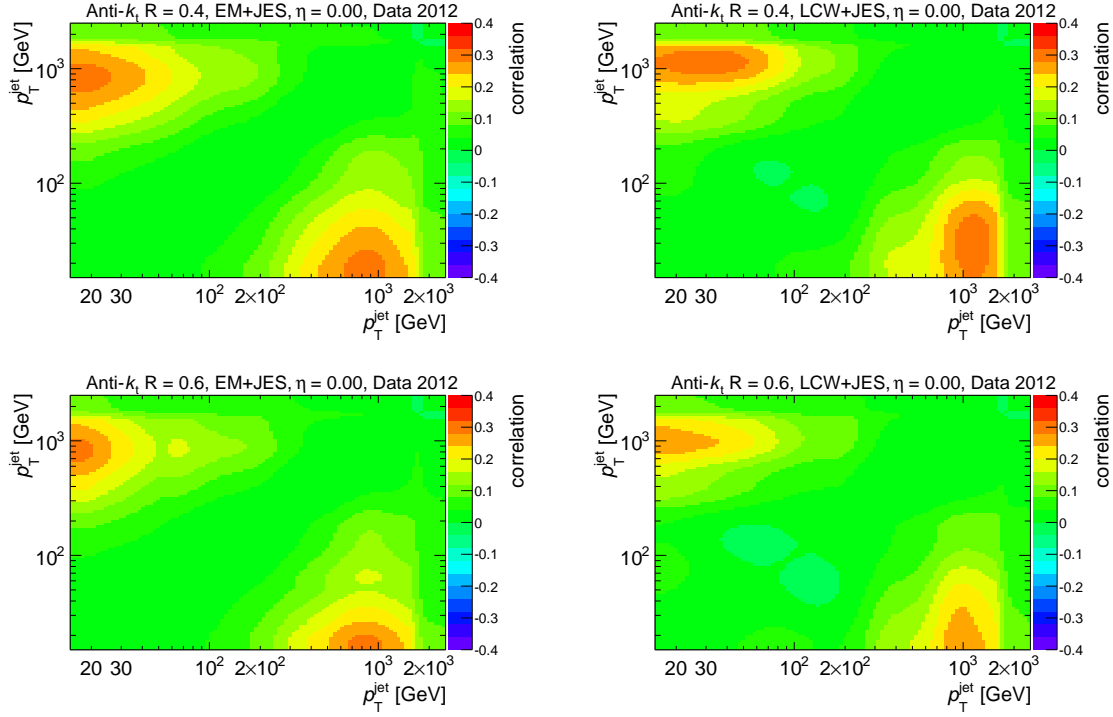


Figure E.28: Correlation changes for the stronger minus weaker alternative correlation scenarios in the di-jet topology for the Final2012 JES release. The maximum difference values are listed in Table E.8.

Jet definition		Corr, stronger - nominal			Corr, weaker - nominal			Corr, stronger - weaker		
		Max diff	p_T^{jet1}	p_T^{jet2}	Max diff	p_T^{jet1}	p_T^{jet2}	Max diff	p_T^{jet1}	p_T^{jet2}
R	cluster scale	[%]	[GeV]	[GeV]	[%]	[GeV]	[GeV]	[%]	[GeV]	[GeV]
0.4	EM	41.7	1619	2438	42.1	1619	2438	30.0	17	876
0.4	LC	46.6	1619	2438	47.1	1619	2438	30.3	33	1132
0.6	EM	42.6	1619	2438	43.0	1619	2438	29.4	15	833
0.6	LC	42.3	1619	2438	42.6	1619	2438	27.3	17	971

Table E.8: The maximum correlations differences between different combinations of the Final2012 JES alternative scenarios. The difference between the nominal and either stronger or weaker scenario is dominated by the changed transition between the multi-jet balance and single-particle measurements, and does not provide an accurate representation of the actual correlation changes across the majority of the parameter space. The difference between the stronger and weaker scenarios provides a better measure of these changes.

R	Jet definition cluster scale	Correlations (Moriond2013)				Correlations (Final2012)			
		N_{eff}	Max loss [%]	$p_{\text{T}}^{\text{jet1}}$ [GeV]	$p_{\text{T}}^{\text{jet2}}$ [GeV]	N_{eff}	Max loss [%]	$p_{\text{T}}^{\text{jet1}}$ [GeV]	$p_{\text{T}}^{\text{jet2}}$ [GeV]
Stronger correlations									
<i>global reduction</i>									
0.4	EM	5	3.5	16	65	5	4.0	125	387
0.4	LC	5	2.7	17	47	5	2.6	451	1075
0.6	EM	5	2.8	15	65	5	5.4	132	387
0.6	LC	5	2.1	15	50	5	2.5	428	1022
<i>category reduction</i>									
0.4	EM	2,2,1,2	6.4	179	2091	2,0,4,3	3.8	79	220
0.4	LC	2,2,1,2	3.6	15	61	2,0,4,3	4.6	209	922
0.6	EM	2,2,1,2	4.8	791	2316	2,0,4,3	3.1	75	220
0.6	LC	2,2,1,2	3.1	146	2090	2,0,4,3	4.7	209	922
Weaker correlations									
<i>global reduction</i>									
0.4	EM	—	—	—	—	5	5.2	119	367
0.4	LC	—	—	—	—	5	3.5	474	1075
0.6	EM	—	—	—	—	5	5.5	97	367
0.6	LC	—	—	—	—	5	3.0	451	1022
<i>category reduction</i>									
0.4	EM	—	—	—	—	2,4,0,4	3.6	209	791
0.4	LC	—	—	—	—	2,4,0,4	4.2	199	922
0.6	EM	—	—	—	—	2,4,0,4	3.5	349	876
0.6	LC	—	—	—	—	2,4,0,4	4.2	209	1022

Table E.9: Global and category reductions for alternative correlation scenarios associated with the Moriond2013 (left) and Final2012 (right) releases. The global reductions are always performed with $N_{\text{eff}} = 5$. The category reductions vary with each scenario in order to retain the maximum amount of correlation information with the minimum number of remaining components. Values are listed in the table as $(N_{\text{eff}}^{\text{detector}}, N_{\text{eff}}^{\text{modelling}}, N_{\text{eff}}^{\text{mixed}}, N_{\text{eff}}^{\text{statistical}})$.

and procedures followed in their derivation. This is expected to provide valuable information to high-precision combinations of results, such as future top mass measurements.

E.4.3 Reductions

After the alternative scenarios are created, both global and category reductions are derived following Appendix E.3. These reduced configurations are checked to ensure that the correlation losses from the reduction process are small compared to the correlation changes intentionally added through the alternative variations. The reduced plots look similar to what has already been provided, thus they are not included here. A summary listing the maximum correlation losses due to the reduction procedure is provided in Table E.9.

Appendix F

Jet Punch-through

This Appendix is based on work done in conjunction with Shaun Gupta and Cigdem Issever; almost all plots were produced by Shaun. This is not intended to completely analyze the correction which has been derived, but rather to provide a general overview and motivation for the choices made. Significantly more details will be published in Shaun’s thesis in the near future.

F.1 Introduction

Jet punch-through is the phenomenon by which a jet is not fully absorbed by the calorimeters, and thus “punches through” the outer edge of the detector. Given that the physical hadronic showers which jets are designed to represent are inherently stochastic, it is not realistic to create calorimeters which are completely impervious to punch-through. To first order, the probability of punch-through depends on the amount of material which the jet must pass through and the energy of the jet. The dependence on these variables equating to 5% longitudinal punch-through, or equivalently 95% longitudinal containment, is defined in Equation F.1 as taken from Reference [138].

$$L(95\%) = \lambda_I (0.6 \ln(E_{\text{jet}}) [\text{GeV}] + 4E_{\text{jet}}^{0.15} [\text{GeV}] - 0.2) \quad (\text{F.1})$$

This cannot reasonably be taken to the limit of complete containment due to the competing needs of shower containment, measurement precision, cost, and size. The size is typically dictated by a combination of the available funding and whatever will go outside of the calorimeters, such as the muon systems in ATLAS. As such, usually the calorimetry has to fit within a certain set of size constraints. For a given size calorimeter, it then becomes a trade-off between hadronic containment and the precision of the energy measurement. Increasing the stopping power by adding more or thicker layers of absorbing material in turn decreases the amount of active material, thus decreasing the fraction of energy which is observed and the resulting precision of the measurement.

Rather than sacrifice the precision of lower energy jet measurements, it is typically accepted that some fraction of punch-through will occur, and it is hoped that reconstruction or calibration algorithms can be designed to reduce the impact of such events. The need for such algorithms at the LHC is set to increase shortly due to the upcoming move from $\sqrt{s} = 8 \text{ TeV}$ to $\sqrt{s} = 13 \text{ TeV}$ collisions. This will significantly increase the fraction of jets with energies at the TeV level, and thus increase the fraction of jets which experience punch-through.

In order to correct punch-through jets, they must first be identified, which is described in Appendix F.2. After a jet has been tagged as experiencing punch-through, a correction must be derived and the performance of that correction must be studied. The derivation and performance of the punch-through correction are detailed in Appendices F.3 and F.4 respectively.

F.2 Identifying punch-through jets

In ATLAS, the focus has so far been placed on identifying and calibrating punch-through jets within $|\eta| < 2.7$. This requirement ensures that the muon system is located behind the calorimeters, and thus provides a means of identifying a sample of punch-through jets. Preliminary studies into punch-through in the more forward regions based on shower profiles have been conducted, but have been left for future investigation, and will not be discussed further here.

The ATLAS muon system is designed to identify and measure charged particles which pass through it. While the reconstruction algorithms are typically optimized to reconstruct muons, the raw hits and track segments remain independent. This is very useful for identifying punch-through, as when a jet escapes from the calorimeter, the resulting highly collimated cone of both charged and neutral particles will pass through the muon system and generate both hits and segments. While the observation of such activity in the muon system behind a jet is not a guarantee of punch-through, it is a strong indication, and serves as the starting point for these studies.

Given that muon activity behind a jet is going to be used to study punch-through, it is important to define what “behind” means. These studies were done independent of the raw data, and thus a complete ghost matching at the per-hit or per-segment level was not possible. Instead, a muon hit/segment container is assigned to each jet of $p_T > 50$ GeV which has muon activity within $\Delta R < 0.7$ of it. The number of hits or segments in this container is then used as the number of hits or segments behind the jet. In the future, a more elegant approach will be possible, as the rewritten software framework for Run-II should provide a ghost matching between jets and muon segments by default.

The next step to identifying punch-through events is to ensure that the identification can be done in both data and MC, as otherwise applying a correction to one but not the other will introduce a bias rather than fixing the problem. To this end, both the number of hits in the muon system behind the jet, $N_{\text{hit}}^{\text{MS}}$, and the number of segments in the muon system behind the jet, $N_{\text{seg}}^{\text{MS}}$, were investigated. Figure F.1 shows that the lack of non-collision sources such as cavern backgrounds is a major limitation for the use of hits, while segments are more stable. In both cases, hits and segments become closer to data when data zero bias events are overlaid with the MC, but without such overlays the hits distribution is unusable.

Punch-through events produce a much higher multiplicity of either hits or segments than other backgrounds, given the cone of charged particles that is expected. As such, the most important part of the distributions is the higher multiplicity region. In this case, the MC segment distribution is closer to data with or without zero bias overlay. Given that the majority of MC in ATLAS is produced without zero bias overlay, it was decided that muon track segment multiplicity was the better variable to use for these studies.

In order to further confirm that punch-through is correlated to the number of segments behind the jet, $N_{\text{seg}}^{\text{MS}}$ is plotted against multiple other variables where the expected punch-through behaviour is known. In particular, the pseudorapidity distribution of material in the ATLAS detector is known.

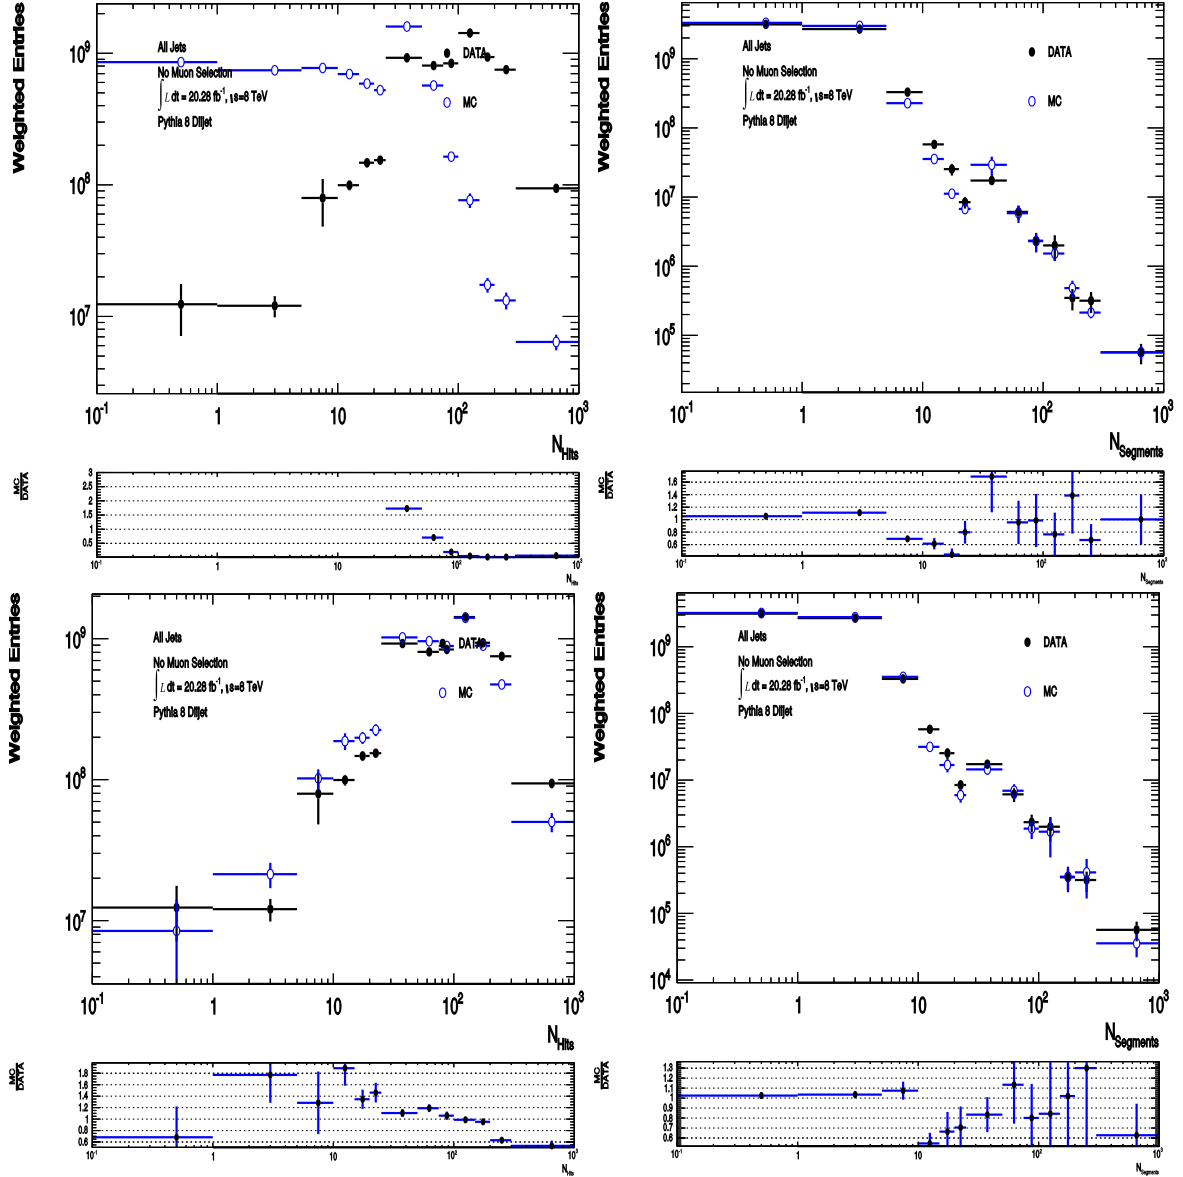


Figure F.1: The number of hits (left) and track segments (right) in the muon system behind a jet from a di-jet-like event, comparing data to MC. Both pure MC simulation (top) and MC with overlaid data zero bias events (bottom) are shown. The zero bias overlay provides a source of cavern backgrounds and other non-collision sources, showing that the disagreement in hits is primarily from such backgrounds, while segments are more stable. [144]




Figure F.2: A comparison of the average number of segments behind the jet in MC as a function of the pseudorapidity $|\eta|$ and the number of interaction lengths of material in the ATLAS detector as a function of the pseudorapidity. The cyan region is additional material after the calorimeters and before the muon spectrometer, and is only shown up to $|\eta| = 3.0$. $N_{\text{seg}}^{\text{MS}}$ is both expected and observed to increase with higher punch-through fractions, as should be the case for regions with less interaction lengths of material before the muon system, such as the transition region between the Tile and HEC. [144, 28]

The fraction of jet punch-through should be larger in regions of lower interaction lengths, and thus the average number of segments should increase. This is exactly what is observed in Figure F.2, where the Tile to HEC transition region at $|\eta| \approx 1.6$ is very clear. The transition from the Tile barrel to the Tile extended barrel is also visible.

When investigating different MC types, it was discovered that the current AFII parametrization does not support energy sharing between the calorimeter and muon spectrometer. As such, while a jet may be marked as punching through by the parametrization and thus have segments in the spectrometer, the energy of the jet recorded in the calorimeter is not reduced to account for the escaping energy. For this reason, the current AFII parametrization is not suitable for these studies, and the correction derived is only applicable for data and MC produced using the full GEANT4 simulation.

One other identification matter of note is the difference between genuine punch-through jets and jets which contain real muons. This can be from processes such as W/Z radiation or, more importantly, heavy flavour decays. This is a real concern, as punch-through jets are not expected to have the same response as b -jets, but they both give rise to muon activity behind the jet. It is not desired to derive a b -jet correction and apply it to punch-through jets or vice versa. Figure F.3 provides an answer, as the fraction of b -jets peaks as a function of $N_{\text{seg}}^{\text{MS}}$ at approximately 10 segments, and then drops and stabilizes by 20 segments. This is taken to mean that the contribution of semi-leptonic b -jets occurs for lower segment multiplicities. For $N_{\text{seg}}^{\text{MS}} \geq 20$, no difference between b -jets and other types of jets signifies that they are all obtaining that number of segments from the same process, namely punch-through. For this reason, a cut on the muon segment multiplicity is applied.

In summary, the selection criteria used to mark a jet as undergoing punch-through is:

- $|\eta^{\text{jet}}| < 2.7$: Detector requirement to have the muon spectrometer behind the calorimeter
- $p_{\text{T}}^{\text{jet}} > 50 \text{ GeV}$: Reconstruction requirement for matching jets to segment containers
- $N_{\text{seg}}^{\text{MS}} \geq 20$: Requirement to suppress heavy flavour decays and enhance punch-through purity




Figure F.3: The fraction of events with a given number of segments $N_{\text{seg}}^{\text{MS}}$ which are tagged as b -jets. The distribution peaking at $N_{\text{seg}}^{\text{MS}}$ of 10 and levelling off by 20 is taken to mean that heavy flavour decays have a lower segment multiplicity than punch-through, as is expected from real muons. [144]

- GEANT4 MC: AFII parametrization does not currently support energy sharing between calorimetry and muon systems

F.3 Correcting for lost energy

Now that a set of criteria for the definition of a punch-through jet has been defined, it can be used to study the phenomenon in more detail. Punch-through jets are expected to have a reconstructed energy which is below the truth value, corresponding to a jet energy response below 1. Ideally, $N_{\text{seg}}^{\text{MS}}$ would serve not only as a variable that identifies candidate jets, but also quantifies the degree of punch-through related to the amount of energy lost. This is indeed the case for $N_{\text{seg}}^{\text{MS}}$, as demonstrated by the dependence of the energy response on the number of segments in Figure F.4.

This correlation implies that the inverse of the jet energy response can be used as a correction factor to fix the scale of punch-through jets on average by fitting a Gaussian to each bin and using the mean to scale the energy. Doing so in MC works quite well, as demonstrated in Figure F.5. This demonstrates that a correction of the inverse of the jet energy response works to remove the JES dependence on $N_{\text{seg}}^{\text{MS}}$ for punch-through jets in all but the last bin. The last bin is still seen to be significantly improved by the application of the correct, but suffers due to the use of an inclusive final bin. The inclusive final bin means that some jets have considerably more segments than others but use the same scale factor, thus pulling down the response.

The next step is to confirm that method also works in data. In this case, the true energy is not known, so an *in situ* tag and probe method is used. This method applies a standard di-jet selection where both jets are within the muon spectrometer acceptance, and then requires that either the leading or sub-leading jet passes the punch-through selection criteria, while the other dominant jet has $N_{\text{seg}}^{\text{MS}} = 0$. This does not strictly rule out the possibility that both jets exhibit punch-through, but the requirement of 0 segments is expected to be relatively pure in good reference jets. The derivation of the jet response with respect to $N_{\text{seg}}^{\text{MS}}$ in data and the comparison to MC is provided in Figure F.6, and reasonable




Figure F.4: The dependence of the jet energy response $\mathcal{R}_{jet} = E_{jet}^{reco} / E_{jet}^{true}$ on the number of segments behind the jet N_{seg}^{MS} . The correlation implies that jets with more segments behind them have also lost more energy, thus the inverse of the response can be used as a first-order correction. [144]




Figure F.5: The jet energy response in MC for the central region of $|\eta| < 1.3$ and for jets with $800 \text{ GeV} < E_{jet}^{true} < 1000 \text{ GeV}$ before and after the correction. The scale factor (lower sub-plot) is seen to vary between approximately 0.98 and 1.1. Good agreement between truth and the corrected MC is seen for all cases except the last bin. [144]

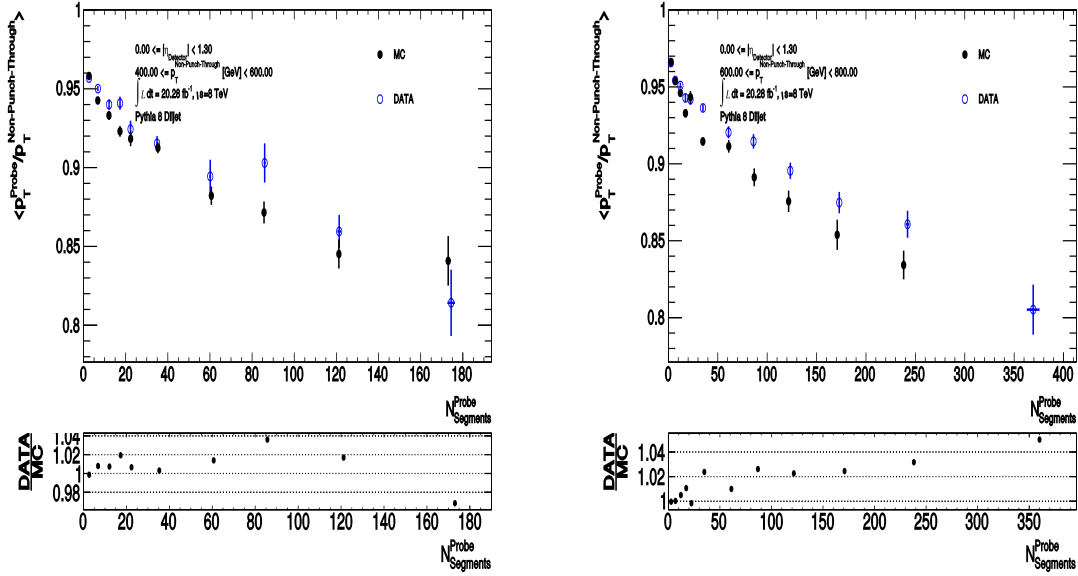


Figure F.6: The *in situ* jet p_T response $\mathcal{R}_{jet} = p_T^{\text{probe}}/p_T^{\text{ref}}$ compared to MC as a function of the number of segments behind the probe jet $N_{\text{seg}}^{\text{MS}}$ for two reference jet p_T bins, $400 \text{ GeV} < p_T^{\text{ref}} < 600 \text{ GeV}$ (left) and $600 \text{ GeV} < p_T^{\text{ref}} < 800 \text{ GeV}$ (right). Perfect agreement is not expected, given that the $N_{\text{seg}}^{\text{MS}}$ distributions do not agree between data and MC, but the observed responses are reasonably similar. [144]

agreement is observed, given that $N_{\text{seg}}^{\text{MS}}$ itself does not perfectly agree between data and MC.

The subsequent application of the MC-derived correction factors to the *in situ* data analysis is provided in Figure F.7. The correction is observed to significantly reduce the dependence of the jet response on $N_{\text{seg}}^{\text{MS}}$, and to bring the probe jet to approximately the same scale as the reference jet, with some residual systematic uncertainty.

F.4 Performance of the punch-through correction

As has already been shown, the punch-through correction derived in MC can be applied to data in order to mostly remove the dependence on $N_{\text{seg}}^{\text{MS}}$ and to approximately fix the scale of punch-through jets. However, the fix is not perfect, and comes with an associated systematic uncertainty. This uncertainty is parametrized in terms of the jet p_T , $|\eta|$, $N_{\text{seg}}^{\text{MS}}$, where the value of $|\eta|$ only serves to select one of three coarse bins: $[0, 1.3)$, $[1.3, 1.9)$, and $[1.9, 2.7)$. This systematic comes from the *in situ* response difference from 1, which is then smeared in $(p_T, N_{\text{seg}}^{\text{MS}})$ -space in order to reduce the large fluctuations from lack of punch-through statistics. The resulting uncertainties for the lower two $|\eta|$ bins, which are less statistics-limited than the most forward bin, are provided in Figure F.8.

The systematic uncertainties on jets which have received the punch-through correction are observed to be at the level of 7% or lower for the typical highly energetic punch-through jet, with an average uncertainty closer to 4%. For lower p_T jets to punch-through the calorimeter, they must have undergone an unusual shower development, where the MC models are not as robust. For such low p_T jets, the uncertainty can reach up to 15%, although it is typically closer to 10%.

While these uncertainties may seem large in comparison to the total JES uncertainty described in

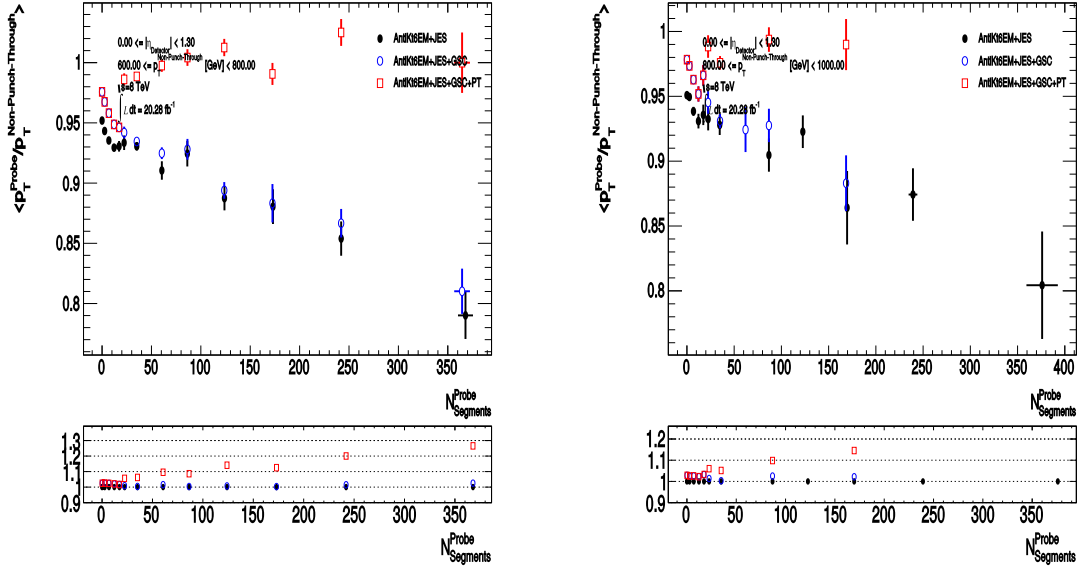


Figure F.7: The *in situ* jet p_T response $\mathcal{R}_{jet} = p_T^{\text{probe}}/p_T^{\text{ref}}$ before and after the punch-through correction as a function of the number of segments behind the probe jet $N_{\text{seg}}^{\text{MS}}$ for two reference jet p_T bins, $600 \text{ GeV} < p_T^{\text{ref}} < 800 \text{ GeV}$ (left) and $800 \text{ GeV} < p_T^{\text{ref}} < 1000 \text{ GeV}$ (right). Shown is the response after the basic JES calibration (black), after the first four steps of the GSC is added (blue), and after the punch-through correction has been applied (red). The application of the MC-derived correction to data is seen to perform well, and to result in a probe jet p_T at approximately the same scale as the reference. [144]

Appendix E, the fraction of jets which undergo punch-through must be considered, as these uncertainties do not apply to the typical jet. The probability for a jet to punch-through as measured *in situ* in data, for the definition of punch-through used in this analysis, is shown in Figure F.9. The fraction of punch-through jets is seen to reach the level of 5% at $p_T \approx 1.3 \text{ TeV}$ and 10% at $p_T \approx 2 \text{ TeV}$. This confirms that punch-through will become significantly more important in the future as the centre of mass energy is increased.

In addition to fixing the scale, it is also important to consider the resolution impact of the correction. Jet punch-through is not yet an average scenario within the current LHC energy reach. As such, punch-through contributes to the broadening of the Gaussian distribution, thus increasing the jet resolution at high energies. This is an important effect, as Section 6.5 mentioned that the highly energetic regime is where the JER matters most, given that only the calorimeter constant term is relevant. This is not strictly true, as that was under the assumption that the jet is contained within the calorimeter. If this is not the case, then the tails of the distribution grow, and punch-through becomes a noticeable widening of the Gaussian fit to the JES.

Given that punch-through events do not yet constitute a significant fraction of the total, the JER of the average highly energetic jet is only improved by a small amount, on the order of 1%. A better means of studying the impact on the resolution is to instead select the punch-through population, and determine how the resolution of that sample changes as the correction is applied. This is done in Figure F.10, where a noticeable improvement is seen with respect to both the jet energy and the associated number of segments. This confirms that the correction is tightening the width of the reconstructed energy distribution for punch-through jets.

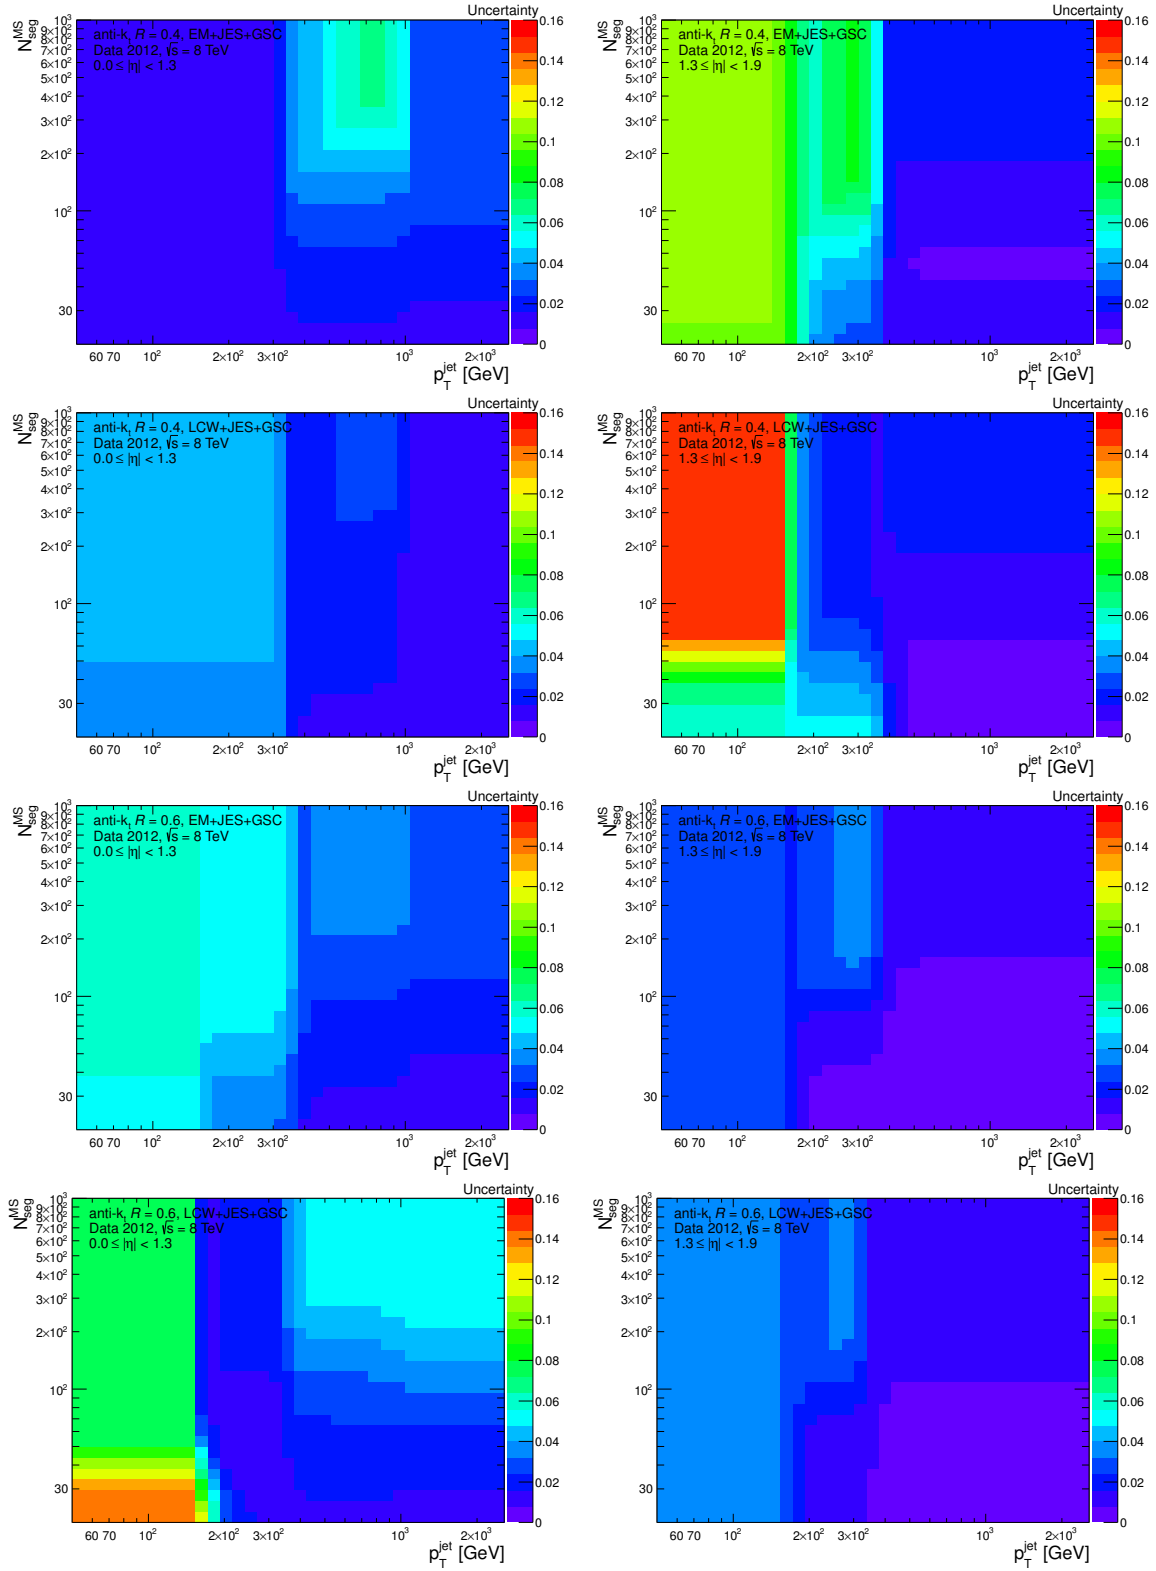


Figure F.8: Systematic uncertainties as a function of p_T^{jet} and $N_{\text{seg}}^{\text{MS}}$ for punch-through-tagged jets within $0 \leq |\eta| < 1.3$ (left) and $1.3 \leq |\eta| < 1.9$ (right), for each standard jet collection (rows), derived from *in situ* studies. Typical high p_T punch-through jets have 7% uncertainties or lower.




Figure F.9: The probability for a jet to punch-through the full ATLAS calorimetry and appear in the muon systems, where punch-through is defined as $N_{\text{seg}}^{\text{MS}} \geq 20$. The fraction is defined within the restricted kinematic parameter space of $p_{\text{T}}^{\text{jet}} > 50$ GeV and $|\eta| < 2.7$.

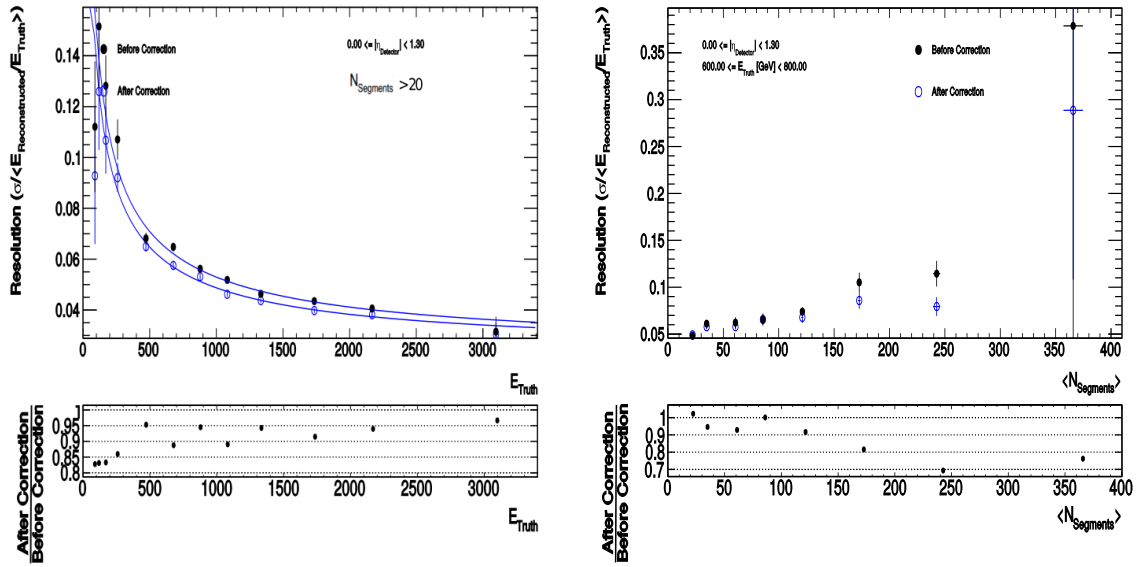


Figure F.10: The improvement in the JER of the punch-through jet population, as a function of the truth jet energy (left) and $N_{\text{seg}}^{\text{MS}}$ (right). Distributions from before (black closed circles) and after the correction was applied (blue open circles) are shown. The resolution noticeably improves in both cases. [144]

The punch-through correction has fulfilled its purpose. It has provided a scale correction, quantified the remaining uncertainties, and significantly reduced the resolution of punch-through jets. It has provided a baseline from which further studies toward Run-II can begin. This is an important step, as the punch-through correction will become much more important in Run-II due to the increased centre of mass energy creating significantly more highly energetic jets.

Appendix G

Effective Field Theory Validity

G.1 Introduction

Using an EFT to represent the production of DM makes the implicit assumption that any new physics coupling the SM inputs to the DM outputs is beyond the energy scale of the interaction. LHC collisions involve momentum transfers at the TeV level, and thus it is not clear if this assumption is valid or not. There are reasonable arguments for a scenario where DM is the only new particle accessible at the LHC, and there are equivalently reasonable scenarios for which this is not the case. Until the underlying physical process is observed and measured, no absolute statement can be made as to whether or not this is a valid assumption.

When comparing collider results on EFT models, the same assumption is made in all cases, and thus the validity is not necessarily a problem for comparing sensitivities. Even if there is new physics within the reach of the LHC, these models can still be treated as benchmarks. However, this is not true when comparing to other types of DM searches. In particular, DD experiments search for nuclear recoil events, where the momentum transfer is typically at the level of a few keV. For the WIMP interpretation of DM, this is reasonably expected to be well below the scale of the new physics coupling to the dark sector, and thus the use of non-relativistic EFTs in the DD community is not questioned.

ID searches are generally considered to be less sensitive to validity constraints than collider searches, as the mediator is typically assumed to be heavier than the DM particles. This allows the annihilation of two DM particles going to two SM particles to not be able to produce the mediator. However, if the mediator and DM are at the same scale, or if the mediator is lighter, then the use of EFTs for ID searches is not valid.

If we consider a generic simple mediator of mass M_{med} which couples to standard model quark(s) with coupling g_q and the dark sector with coupling g_χ , we can represent the suppression scale of an associated simply coupled EFT using $M_* = \frac{M_{\text{med}}}{\sqrt{g_q g_\chi}}$. In order for the couplings to be perturbative, we gain the additional constraint of $g_q, g_\chi < 4\pi$. Additionally, $M_{\text{med}} > 2m_\chi$ is required for the EFT to be valid, as otherwise there is enough energy flow into the final state particles to resonantly produce the mediator. Combining these constraints gives us a region of validity requiring $M_* > \frac{2m_\chi}{\sqrt{(4\pi)^2}}$, or $M_* > \frac{1}{2\pi}m_\chi$. (Note that in previous ATLAS publications, there was no requirement of $M_{\text{med}} > 2m_\chi$ to prevent resonant production, but only $M_{\text{med}} > m_\chi$ such that the DM is lighter. Thus the factor of 2 is removed, and we have $M_* > \frac{1}{4\pi}m_\chi$.)

While this is one requirement for validity, it is not the most rigorous interpretation. In reality, the momentum transfer through the mediator must be below the scale of resonant production, and thus a more realistic minimum constraint is $M_{\text{med}} > Q_{\text{tr}}$ as per Reference [145]. In the case where the two DM particles are produced at rest, which is the most optimistic situation, this is equivalent to the above statement of $M_{\text{med}} > 2m_\chi$. However, in more realistic scenarios where the DM particles have some remaining kinetic energy, a tighter bound is required.

Additionally, note that this is not a generically true validity condition. An assumption that the effective vertex represents a simple interaction was made, which is valid for the D5, D8, and D9 operators where the contact interaction is only integrating out a mediator. This is not true for C1 and D1, where there is chiral suppression, or D11 and C5, where there are actually two different effective processes being combined into one vertex. As such, this assumption does not hold for every operator, and will not be considered further. Instead, a more rigorous interpretation which takes the structure of each operator into consideration has been conducted.

The best solution is to impose a constraint of $M_{\text{med}} \gg Q_{\text{tr}}$, but it is impossible to quantify “much greater than” as a factor. Instead, the minimum constraint of $M_{\text{med}} > Q_{\text{tr}}$ is considered. This is a complicated constraint for an EFT, given that the mediator has been integrated out, and there remains no simple handle on M_{med} . A relation between M_* and M_{med} must first be found for each operator in question.

G.2 Connecting the EFT suppression scale to the mediator mass

By defining a relation between M_* and M_{med} , it is typically impossible to avoid additional assumptions on what forms the UV theory can take. This is because some form of assumption must be made about the behaviour of a given type of UV-complete theories in the contact interaction limit, which can vary by the specific form of the UV theory. The following Section will endeavour to minimize these additional UV assumptions where possible, while providing a set of reasonable relations between M_* and M_{med} for each of the operators considered.

G.2.1 Simple operators: D5, D8, C3, D9

The coupling between quarks and dark matter through vector, axial vector, and tensor interactions is fairly simple. In the standard s-channel UV completion, the two quarks annihilate to form a mediator, which then decays to two WIMPs. In this picture, the following steps are taken when moving from UV to EFT:

$$\sigma(pp \rightarrow \chi\chi + \text{jet}) \propto \frac{g_q^2 g_\chi^2}{(Q_{\text{tr}}^2 - M_{\text{med}}^2)^2 + \Gamma_{\text{med}}^2 M_{\text{med}}^2} \xrightarrow{M_{\text{med}} \gg Q_{\text{tr}}, \Gamma_{\text{med}}} \frac{g_q^2 g_\chi^2}{M_{\text{med}}^4} \xrightarrow{EFT} \frac{1}{M_*^4} \quad (\text{G.1})$$

In this case, the conversion from M_* to M_{med} is simple, with $M_{\text{med}} = \sqrt{g_q g_\chi} M_*$. The couplings can take any value from 0 to 4π , corresponding to no coupling between the SM and Dark matter to the maximum coupling before the theory becomes non-perturbative. Without knowledge of the exact nature of dark matter, it is impossible to further restrict this range.

In order to follow established naming conventions, couplings to fermionic DM will be labelled as g_χ , while λ_χ is used for couplings to scalar DM. According to this naming convention, $M_{\text{med}} = \sqrt{g_q g_\chi} M_*$ applies to the D5, D8, and D9 operators, while $M_{\text{med}} = \sqrt{g_q \lambda_\chi} M_*$ applies to C3.

G.2.2 Chirally suppressed operators: D1, C1

The coupling between quarks and dark matter through scalar interactions is more complicated, and can take one of two forms. It is commonly considered to be a minimally flavour violating vertex with a Yukawa coupling $y_q = g_q m_q / \nu_y$, where ν_y is the Yukawa Vacuum Expectation Value (VEV) (often assumed to be the Higgs VEV). Note that this equation for the break-down of the Yukawa coupling into other terms is valid, but the perturbativity bound is still $y_q < 4\pi$. This is because the form of the Lagrangian is $y_q \phi q \bar{q}$, thus the perturbativity requirement is on the composite y_q rather than the independent scale g_q . For this reason, the following Equations will be left in the form of y_q rather than performing the substitution.

Due to the dependence of the cross-section on the quark mass, the most massive kinematically accessible quark dominates the cross-section. Requiring sea quarks and the related PDF suppression reduces the effect, but for example the inclusion of bottom quarks results in a sample containing $\mathcal{O}(90\%)$ bottom quark interactions. The mono-jet WIMP samples were generated only with light quarks, where charm was the most massive option. Thus, for the mono-jet samples, the charm mass is the scale of interest. Note that this does not mean the charm is what would be the most commonly produced in actual data if such a signal existed - the bottom would be a much larger contribution, as seen in the Heavy Flavour + DM analysis [7], so long as a coupling to the bottom exists. However, generating samples including bottom quarks also increases the cross-section, compensating for this effect. Thus, the most important consideration is to be consistent - when setting limits on samples where charm is the most massive quark, the charm mass scale must be used.

Note the inclusion of a VEV ν_λ in UV completions of the complex scalar operators for the tri-scalar vertex (scalar DM, scalar DM, scalar mediator). To simplify equations, the scale of the VEV is linked to the mediator mass through a scale factor, $\nu_\lambda = \zeta_\lambda M_{\text{med}}$.

$$\begin{array}{cc}
 \text{D1} & \text{C1} \\
 \frac{m_q}{M_*^3} = \frac{y_q g_\chi}{M_{\text{med}}^2} & \frac{m_q}{M_*^2} = \frac{y_q \lambda_\chi \nu_\lambda}{M_{\text{med}}^2} \\
 & \text{Note: } \nu_\lambda = \zeta_\lambda M_{\text{med}} \\
 M_{\text{med}} = \sqrt{y_q g_\chi} \cdot \sqrt{M_*^3 / m_q} & M_{\text{med}} = y_q \lambda_\chi \zeta_\lambda \cdot M_*^2 / m_q \tag{G.2}
 \end{array}$$

As before, $0 < y_q, g_\chi, \lambda_\chi < 4\pi$. The scale factor ζ_λ is an unknown constant, but some value must be used for the creation of a criterion which can be applied to measure the validity. As such, the common assumption of the VEVs being at approximately the scale of the mediator is used, or equivalently $\zeta_\lambda \approx 1$. This assumption can be recovered by changing the range of couplings over which the validity is tested.

G.2.3 Gluon operators: D11, C5

The coupling between gluons and dark matter through scalar interactions is significantly different than for other operators. These are operators which involve two levels of effective interactions combined into a dimension-3 vertex. Two of the dimensions represent the mediator as usual, while the third is the coupling of two gluons to the mediator through a loop. Here, there are numerous possibilities. Some form of assumption must be made about the two relative scales. If the scale of the loop is below the scale of the mediator, then the full completion of the mediator will be analogous to the gluon-gluon fusion to top quark loop production mechanism of the Higgs boson, so long as the top quark is the relevant quark for the loop. This scenario is not considered further at this time.

Alternatively, if the scale of the loop is above the scale of the mediator, then the mediator can be resolved leaving a one-dimensional EFT connecting the two gluons to the mediator [93]. The matching condition here is much more complicated and involves the assumption that the remaining effective scale (Λ_s) is beyond the scale of the LHC, which is generally reasonable given flavour violation constraints.

As for the C1 and C3 operators, the naming convention of λ_χ for scalar couplings is used. The addition of a VEV is also necessary, and the relation to the mediator mass of $\nu_\lambda = \zeta_\lambda M_{\text{med}}$ is used.

D11	C5
$\frac{\alpha_s}{4M_*^3} = \frac{\alpha_s g_\chi}{M_{\text{med}}^2 \Lambda_s}$	$\frac{\alpha_s}{4M_*^2} = \frac{\alpha_s \lambda_\chi \nu_\lambda}{M_{\text{med}}^2 \Lambda_s}$
$\Lambda_s > M_{\text{med}}$ (resolved mediator)	$\Lambda_s > M_{\text{med}}$ (resolved mediator)
Let $\Lambda_s = b \cdot M_{\text{med}}$, $b > 1$	Let $\Lambda_s = b \cdot M_{\text{med}}$, $b > 1$
$\frac{1}{4M_*^3} = \frac{g_\chi}{bM_{\text{med}}^3}$	$\frac{1}{4M_*^2} = \frac{\lambda_\chi \zeta_\lambda}{bM_{\text{med}}^2}$
$M_{\text{med}} = \sqrt[3]{\frac{4g_\chi}{b}} M_*$	$M_{\text{med}} = \sqrt{\frac{4\lambda_\chi \zeta_\lambda}{b}} M_*$
Let $a = 4b^{-1}$, $0 < a < 4$	Let $a = 4b^{-1}$, $0 < a < 4$
$M_{\text{med}} = \sqrt[3]{a g_\chi} M_*$	$M_{\text{med}} = \sqrt{a \lambda_\chi \zeta_\lambda} M_*$

(G.3)

Multiple assumptions were made in the creation of this relation, but it is one possible mapping to the D11 and C5 operators from a more complete theory where the mediator is accessible but the gluon vertex is not. As before, $0 < g_\chi, \lambda_\chi < 4\pi$. The non-dimensional scale factor a was shown above to have a range of $0 < a < 4$.

The scale factor ζ_λ is an unknown constant, but some value must be used for the creation of a criterion which can be applied to measure the validity. As such, the common assumption of the VEV being at approximately the scale of the mediator is used, or equivalently $\zeta_\lambda \approx 1$. This assumption can be recovered by changing the range of couplings over which the validity is tested.

G.2.4 Using the M_* to M_{med} relations

Now that a relation between M_* and M_{med} has been determined for each operator of interest, it is important to note that each relation depends on the value of an unknown coupling term. Without knowledge of the underlying UV theory, it is impossible to know what value of the couplings is justified,

Operator(s)	Relation between M_{med} and M_*	Coupling term range
D1	$M_{\text{med}} = \sqrt{y_q g_\chi} \sqrt{M_*^3/m_q}$	$\sqrt{y_q g_\chi} \in (0, 4\pi)$
C1	$M_{\text{med}} = y_q \lambda_\chi \zeta_\lambda M_*^2/m_q$	$y_q \lambda_\chi \zeta_\lambda \in (0, [4\pi]^2 \cdot \zeta_\lambda)$
D5, D8, D9	$M_{\text{med}} = \sqrt{g_q g_\chi} M_*$	$\sqrt{g_q g_\chi} \in (0, 4\pi)$
C3	$M_{\text{med}} = \sqrt{g_q \lambda_\chi} M_*$	$\sqrt{g_q \lambda_\chi} \in (0, 4\pi)$
D11	$M_{\text{med}} = \sqrt[3]{a g_\chi} M_*$	$\sqrt[3]{a g_\chi} \in (0, \sqrt[3]{16\pi})$
C5	$M_{\text{med}} = \sqrt{a \lambda_\chi \zeta_\lambda} M_*$	$\sqrt{a \lambda_\chi \zeta_\lambda} \in (0, 4\sqrt{\pi \zeta_\lambda})$

Table G.1: Relations between the mediator mass M_{med} and the suppression scale M_* for typical UV theories and their corresponding EFT operators. The scale factor ζ_λ is defined as $\nu_\lambda = \zeta_\lambda M_{\text{med}}$, and serves to quantify the difference between the VEV scale and the mediator mass scale. The common natural assumption of $\zeta_\lambda \approx 1$ is used in all instances, but this assumption can be changed by varying the maximum coupling considered to match any other assumption for the VEV scale factor.

if any. As such, it is important to scan over the possible coupling parameter space. Two benchmark points can be chosen. The first is the common theory assumption of a coupling of 1 ($g_q = g_\chi = 1$ for the simple operator case), which is the so-called natural scale. The second is the maximum coupling value before the theory becomes non-perturbative, the value of which varies by operator. A comparison of these two values and general coupling parameter space scans will be carried out in the following Sections. A summary of the operators, M_* to M_{med} relations, and coupling ranges is provided in Table G.1.

Note that while these relations were all derived for s-channel processes, the same validity arguments can be applied to the t-channel. This can be done by using a sum of s-channel operators to represent a single t-channel operator and then studying the validity, as per Reference [146]. Also note that these truncated limits are always more conservative than those from the simplified model as long as M_{med} is greater than or equal to the value we use for EFT truncation. This can be seen from comparing the D5 operator in Figures 8.4 and 8.5 with the corresponding Z' simplified model in Figure 8.11.

G.3 Rescaling limits

G.3.1 The procedure

Now that a relation between the mediator mass and the suppression scale has been defined, it is possible to understand where the momentum transfer becomes large enough to produce an on-shell mediator (within assumptions). The fraction of valid events passing the validity criteria of $Q_{\text{tr}} < M_{\text{med}}$ can be used to translate our nominal limits M_*^{exp} to rescaled (truncated) limits M_*^{valid} along the lines of Reference [145]. However, note that this condition depends on M_{med} , which in turn depends on M_* . As the limit is rescaled, this condition changes, and more phase space must be cut out. It is thus necessary to scan through M_* until a convergence point is reached, as per the following algorithm:

1. Start with the nominal expected/observed limit on M_* assuming 100% validity, named M_*^{exp} .
2. For each step i , obtain the relative fraction of valid events $R_{M_{\text{med}}}^i$ satisfying $Q_{\text{tr}} < M_{\text{med}}^{\text{in}}$ with respect to all events passing the previous iteration. $M_{\text{med}}^{\text{in}}$ is the mediator mass limit obtained in the previous step, which depends on M_*^{in} . For the first step, $M_*^{\text{in}} = M_*^{\text{exp}}$.
3. Rescale M_* following Reference [145]: $M_*^{\text{out}} = [R_{M_{\text{med}}}^i]^{1/2(d-4)} M_*^{\text{in}}$, noting that D1 and D11 are dimension $d = 7$ operators, while D5, D9, C1, and C5 are dimension $d = 6$.

4. Go to step 2, using the current M_*^{out} as the new M_*^{in} , repeating until the fraction of valid events at a given step $R_{M_{\text{med}}}^i$ reaches 0 or 1.
5. Calculate the total validity fraction $R_{M_{\text{med}}}^{\text{tot}} = \prod_i R_{M_{\text{med}}}^i$ and the final rescaled limit on the suppression scale $M_*^{\text{valid}} = [R_{M_{\text{med}}}^{\text{tot}}]^{1/2(d-4)} M_*^{\text{exp}}$.

This procedure provides very accurate results, but requires substantial statistics. The iteration procedure throws away events in every iteration, particularly if the given scenario is only minimally valid. Typical signal MC sample sizes are not sufficient for this procedure when $R_{M_{\text{med}}}^{\text{tot}}$ is below approximately 15 to 20%.

An alternative procedure which does not suffer from this limitation but which is less precise was also studied:

1. Construct the Q_{tr} distribution for a given MC sample
2. Scan over M_* and compare to the Q_{tr} distribution, determining $R_{M_{\text{med}}}^{\text{tot}}$ for each value of M_*
3. Rescale the nominal expected/observed limit on the cross-section of each signal sample, $\sigma_{\text{exp}}(M_*)$, so that it corresponds to only the valid events: $\sigma_{\text{valid}}(M_*) = \sigma_{\text{exp}}(M_*) \cdot R_{M_{\text{med}}}^{\text{tot}}(M_*)$.
4. Determine the point where σ_{valid} and the experimental limit on the visible cross-section σ_{vis} meet, and take the corresponding M_* value as the truncated limit on the suppression scale M_*^{valid} .

It was checked that applying either procedure leads to the same results, within the precision of the methods. Given sufficient statistics, the iterative procedure is more precise, while the cross-section procedure is more generically stable. The cross-section procedure was used for all plots and tables in the end due to low MC statistics, where the iterative procedure was used as a cross-check for all points.

G.3.2 Dependence of validity on the signal region

The fraction of valid events is larger, or equivalently a particular scenario is more valid, if the original limit is stronger. This is easy to see from any of the relations in Table G.1, given that the validity constraint is $Q_{\text{tr}} < M_{\text{med}}$. As such, there is a dual dependence of the validity on the signal region:

1. Higher signal regions imply higher $E_{\text{T}}^{\text{miss}}$, which is weakly correlated to higher Q_{tr} , and thus worse validity. The correlation is only weak because the $E_{\text{T}}^{\text{miss}}$ smears out the momentum of the WIMPs.
2. Higher signal regions typically lead to stronger limits, up to a point. These stronger limits provide larger values of M_*^{exp} , providing increased validity.

These two competing effects have both been observed, but the increased value of M_*^{exp} from higher signal regions dominates. This can be seen in Figure G.1 as an example of the D5 operator, where the correlation of Q_{tr} and $E_{\text{T}}^{\text{miss}}$ only has a minor impact, and only makes any difference for low coupling values.

Given that the signal region to be used for the observed limits is selected based on the optimal expected limit, it is necessary to question whether this should be based on M_*^{exp} or M_*^{valid} . Figure G.1 shows that the difference is minimal. Selecting based on the best M_*^{valid} will only differ for border cases, where the weak dependence of Q_{tr} on the $E_{\text{T}}^{\text{miss}}$ has enough of an impact to impact the signal region




Figure G.1: The fraction of valid events $R_{M_{\text{med}}}^{\text{tot}}$ (top), scale factor $M_*^{\text{valid}}/M_*^{\text{exp}}$ (bottom left), and truncated limit M_*^{valid} (bottom right) for the D5 operator with $m_\chi = 50$ GeV as a function of the signal region (x-axis) and couplings (y-axis). The value of M_*^{exp} comes from the 95%CL expected limit, which is used to select the signal region for the observed limits. The validity is strongly dominated by the best limit on M_*^{exp} , as the validity improves with higher E_T^{miss} (to a point). The region with the best validity uses a 400 GeV E_T^{miss} threshold (visible in the valid fraction plot), while the region with the best value of M_*^{exp} uses a 500 GeV cut on E_T^{miss} . This difference shows the subdominant impact of the correlation between Q_{tr} and E_T^{miss} , which matters when the fraction of valid events is around 10%. While the E_T^{miss} matters for low valid fractions, for higher valid fractions at the level of approximately 25% (which comes with higher couplings), selecting the best SR by M_*^{exp} or by M_*^{valid} makes no difference.

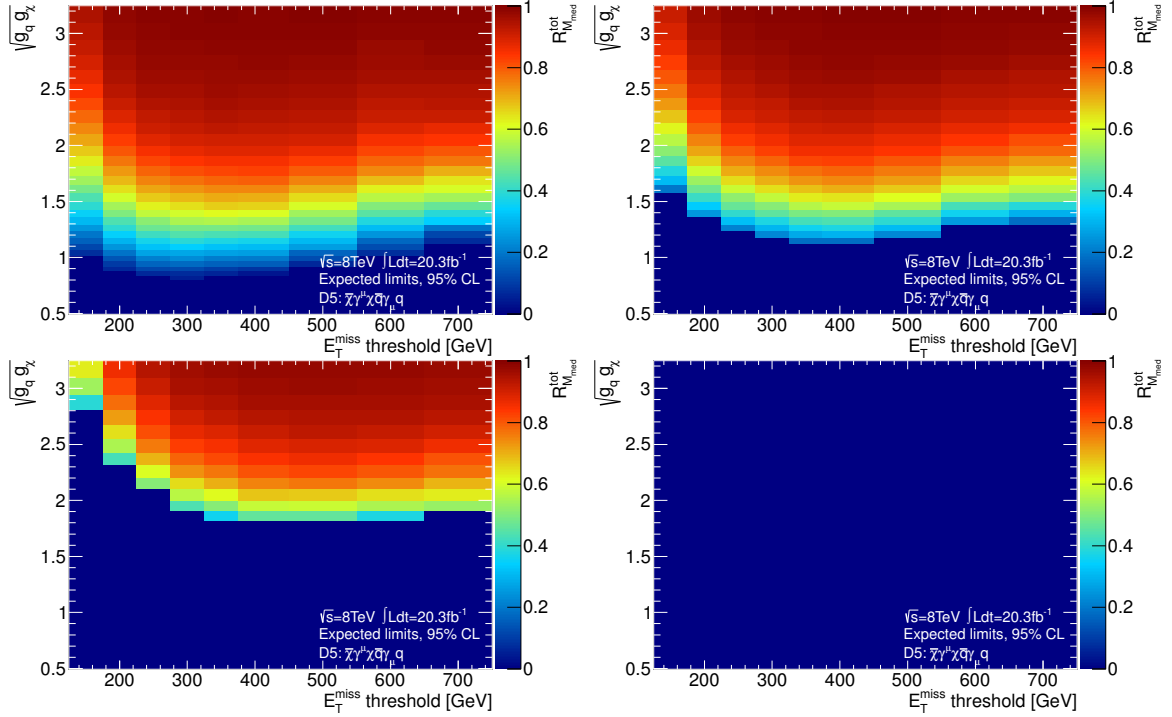


Figure G.2: The fraction of valid events for the D5 operator with m_χ of 100 GeV (top left), 200 GeV (top right), 400 GeV (bottom left), and 700 GeV (bottom right) as a function of the E_T^{miss} threshold and couplings. The fraction of valid events decreases with increasing m_χ , while increasing with the couplings. A limit of $M_*^{\text{exp}} \sim 1$ TeV prevents the pair-production of 700 GeV WIMPs for the coupling range shown.

choice. For even slightly higher coupling values, this becomes irrelevant, as both selections converge on the same region.

In order to have a simple definition of the optimal signal region which can be applied to all interpretations, to be comparable to past publications, and in light of the very small difference between the two methods, it was decided to continue to select signal regions based on the best expected limit before any validity conditions are considered. This has the added benefit of not tying the limits to any given validity procedure, which also carry assumptions about the UV completions.

G.3.3 Dependence of validity on the DM mass

The momentum transfer Q_{tr} and the DM mass m_χ are not independent quantities. In order to pair-produce two WIMPs at rest, $Q_{\text{tr}} = 2m_\chi$. In general, it is expected that the WIMPs are moving, thus $Q_{\text{tr}} \geq 2m_\chi$. This relation is observed, and by increasing the value of m_χ one can observe significantly reduced validity. This is clearly visible in Figure G.2, where the validity degrades with every increase in m_χ . This also shows how the DM mass shapes the Q_{tr} distribution. The distribution begins with a Gaussian-like shape for low m_χ values, as the DM mass is negligible. As the mass is increased, the lower tail disappears. This is visible by the quicker transition from 0% valid to 50% valid when increasing the coupling value. The low tail has been replaced by a sharpened transition (step function for high enough masses), and thus the validity can jump from 0% to 50% with only a small change.




Figure G.3: The fraction of events satisfying the validity condition of $Q_{\text{tr}} < M_{\text{med}} = \sqrt{g_q g_\chi} M_*$ as a function of the coupling $\sqrt{g_q g_\chi}$. DM particles with a mass of 50 GeV are considered, at both 8 TeV (left) and 14 TeV (right). Three signal regions with different E_T^{miss} cuts are shown, with 400 GeV in black, 600 GeV in red, and 800 GeV in blue. [8]

G.3.4 Dependence of validity on the centre of mass energy

The preliminary studies conducted as a part of Reference [8] were only conducted in a simple fashion and only for the D5 operator, but they had one added variable which was not reproducible in Reference [5]. These studies compared different values of the centre of mass energy, thus providing extra insight into the future of EFTs at the LHC. There are multiple competing effects present in the change of energy scale. By increasing the centre of mass energy, the average momentum transfer of the interaction grows, and thus $Q_{\text{tr}} < M_{\text{med}} = \sqrt{g_q g_\chi} M_*$ is more restrictive for a fixed value of M_* . At the same time, increasing the centre of mass energy significantly increases the strength of the analysis, as the SM backgrounds grow slower than the signal at higher energies as detailed earlier in Section 9.4.2. For a fixed value of Q_{tr} , this means that $Q_{\text{tr}} < \sqrt{g_q g_\chi} M_*$ becomes less restrictive.

Figure G.3 demonstrates that the increase in the mono-jet sensitivity to DM dominates, with the fraction of valid events going from $\mathcal{O}(1\%)$ to $\mathcal{O}(40\%)$ for a DM mass of 50 GeV and a coupling of $\sqrt{g_q g_\chi} = 1$ for this simple analysis. An interesting conclusion that can be drawn from these studies is that increasing the centre of mass energy in collider searches for DM with EFTs results in an increased validity of the model, using the stated definition of model validity.

G.4 Results

G.4.1 Full study at 8 TeV

The plots seen in Figure G.1 end up with one redundant dimension for observed limits, as the signal region is fixed by the best expected limit on M_*^{exp} . However, they leave out the aforementioned dependence on m_χ . This can be rectified by combining a full series of such plots, one per DM mass, resulting in plots which show the dependence of the validity on m_χ and the assumed coupling. Such plots are very useful to convey the validity of a given operator in a compact and concise format. The resulting plots for the full breakdown per operator, covering D1, D5, D8, D9, D11, C1, C3, and C5, for both 90% and 95% confidence level observed limits, are all available in Figures G.4, G.5, G.6, G.7, G.8, G.9, G.10, and G.11 respectively. These Figures show the fraction of valid events, truncation scale factor, and resulting

truncated limits for each scenario.

Note that some operators are more valid than others simply because it is possible to set stronger limits on them, and thus it is more likely that $Q_{\text{tr}} < M_{\text{med}}$. This poses an interesting problem, where our stronger limits remain strong, and our weaker limits are even further diminished. This also reinforces the problem that the analysis sets limits at a given M_* value, but that does not necessarily rule out weaker scales. As M_* decreases, the EFT becomes invalid, and thus our limits no longer apply. So EFT results including validity considerations set limit on a band of parameter space, not the full region.

G.4.2 Preliminary study at 8 and 14 TeV

EFT validity studies were conducted for the 8 TeV 20 fb^{-1} and the 14 TeV 25 fb^{-1} scenarios, both under the assumption of a flat 5% systematic uncertainty. As previously derived, the EFT to UV matching condition for the D5 operator implies $M_{\text{med}} = \sqrt{g_q g_\chi} M_*$, where $0 < g_q, g_\chi < 4\pi$, and thus $0 < \sqrt{g_q g_\chi} < 4\pi$. A scan over the couplings is conducted on the transition region of $1 < \sqrt{g_q g_\chi} < \pi$. Figure G.12 shows the resulting fraction of valid events, which is seen to go from $\mathcal{O}(0\%)$ to $\mathcal{O}(100\%)$ valid at 8 TeV for couplings of 1 and π respectively when $m_\chi = 50 \text{ GeV}$.

There are now two dimensions to explore, the first of which is the DM mass. By increasing the DM mass, the momentum transfer spectrum becomes harder, and at some point prevents the process entirely due to the relation $Q_{\text{tr}} \geq 2m_\chi$. The same figure shows how the point at which the model becomes partially valid is for a coupling approaching 2 rather than 1, although both are $\mathcal{O}(100\%)$ valid by $\sqrt{g_q g_\chi} = \pi$. This is consistent with what has been seen in the more detailed studies in Appendix G.4.1.

The second dimension is unique to this study, and provides insight into the impact of the centre of mass energy, as already discussed in Appendix G.3.4. This dimension demonstrates that the improved sensitivity dominates over the increased momentum transfer, and thus the EFT becomes more valid for the definition of validity considered.

Another conclusion for this particular set of studies is that the fraction of valid events is at the level of $\mathcal{O}(100\%)$ for all scenarios considered so long as $\pi < \sqrt{g_q g_\chi} < 4\pi$. This result provided the coupling range added to the legend and captions of all limit projection and discovery potential plots from Section 9.4.2.

The fraction of valid events can be turned into truncated limits, as previously described. For the dimension-6 D5 operator, the conversion is $M_*^{\text{valid}} = [R_{M_{\text{med}}}^{\text{tot}}]^{1/4} M_*^{\text{exp}}$. The resulting truncated limits (and the corresponding limits excluding validity considerations) for both 8 and 14 TeV and for both m_χ of 50 and 400 GeV are in Figure G.13.

Figure G.13 also makes a comparison of the signal regions considered much more direct. The fraction of valid events is not the final observable which the selection of an optimal signal region should be based on. Instead, the optimal signal region selection is based on the best expected limits, which could possibly be extended to the signal region which provides the strongest truncated limits. Given that there is a correlation between the $E_{\text{T}}^{\text{miss}}$ threshold and the Q_{tr} , it could be that increasing the $E_{\text{T}}^{\text{miss}}$ threshold will in the end reduce the truncated limits. However, at the same time, the strongest nominal limit generally grows with the $E_{\text{T}}^{\text{miss}}$ threshold.

Figure G.13 clearly shows the dominance of the improved limit from the higher $E_{\text{T}}^{\text{miss}}$ cut. The region with the strongest nominal limit also has the strongest truncated limit, with a single point exception ($\sqrt{g_q g_\chi} = 1$ for $m_\chi = 50 \text{ GeV}$ at 8 TeV). This is despite the sometimes smaller fraction of valid events which comes with the higher $E_{\text{T}}^{\text{miss}}$ cut, including a range of $\sqrt{g_q g_\chi} < 1.4$ rather than just the single point, as seen in Figure G.12. As such, it is a reasonable choice to continue to select the optimal signal

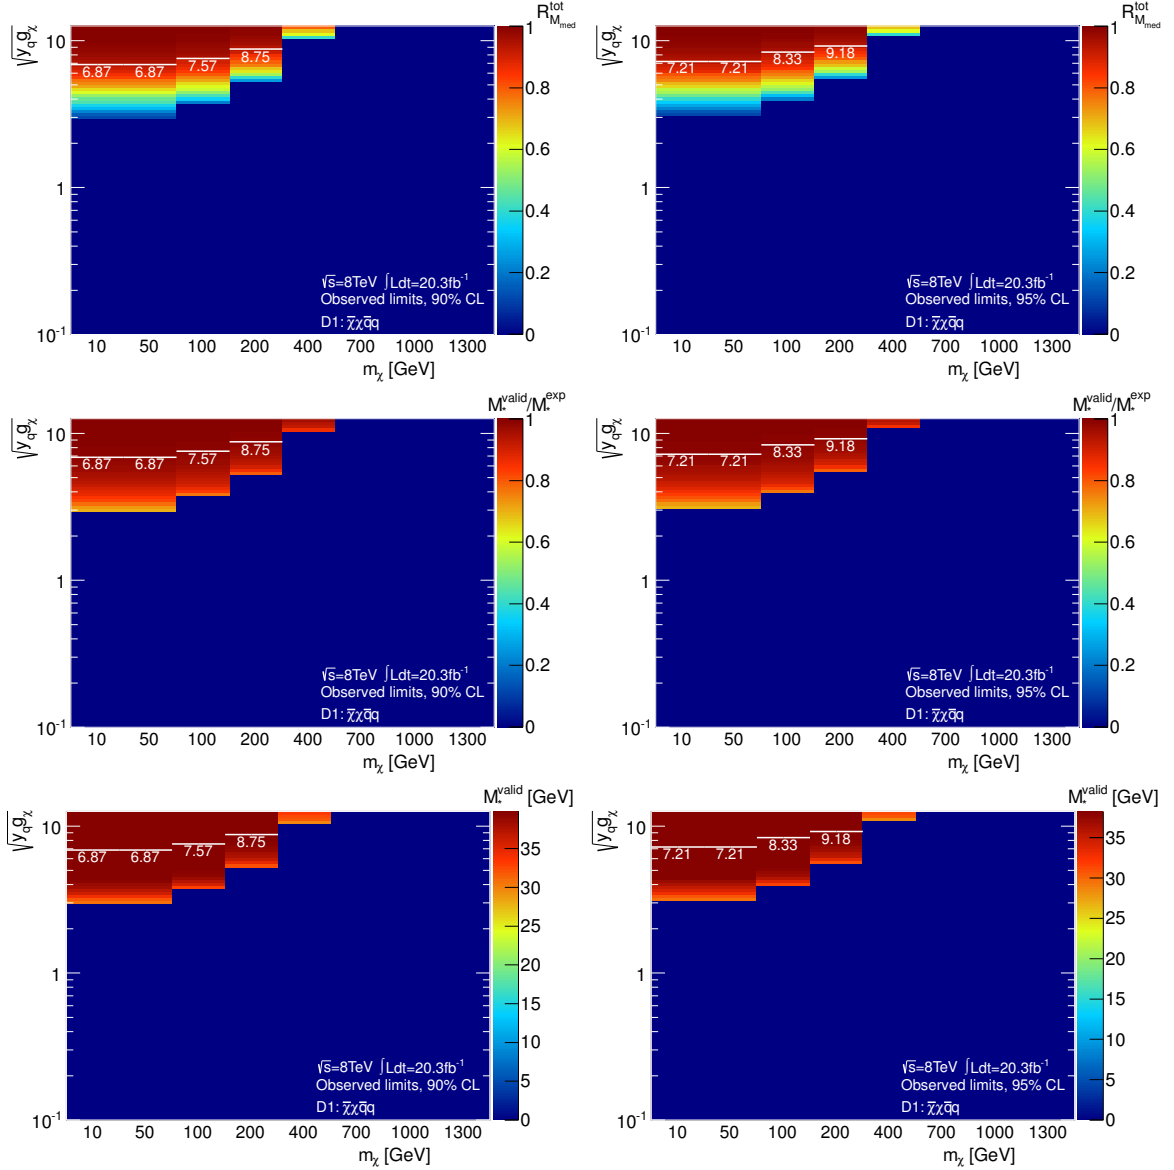


Figure G.4: The fraction of valid events (top), scale factor (middle), and rescaled limits (bottom) for the D1 operator at 90%CL (left) and 95%CL (right) observed limits as a function of the WIMP mass (x-axis) and couplings (y-axis). White lines and associated white numbers correspond to the minimum coupling value for which the truncated limit is within 1% of the nominal limit, or equivalently the scale factor is at least 0.99.

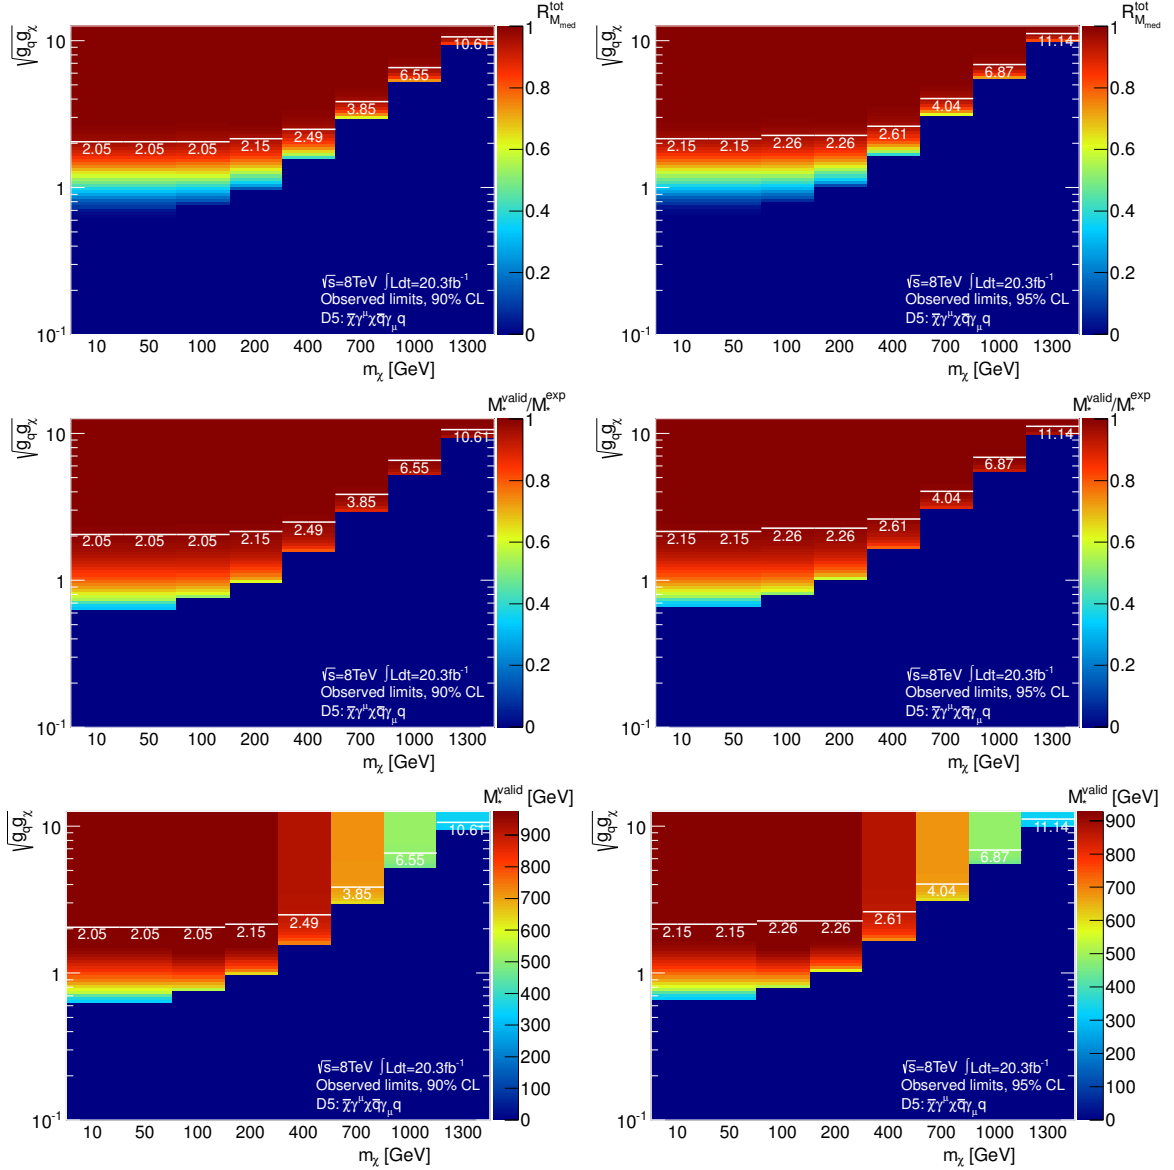


Figure G.5: The fraction of valid events (top), scale factor (middle), and rescaled limits (bottom) for the D5 operator at 90%CL (left) and 95%CL (right) observed limits as a function of the WIMP mass (x-axis) and couplings (y-axis). White lines and associated white numbers correspond to the minimum coupling value for which the truncated limit is within 1% of the nominal limit, or equivalently the scale factor is at least 0.99.

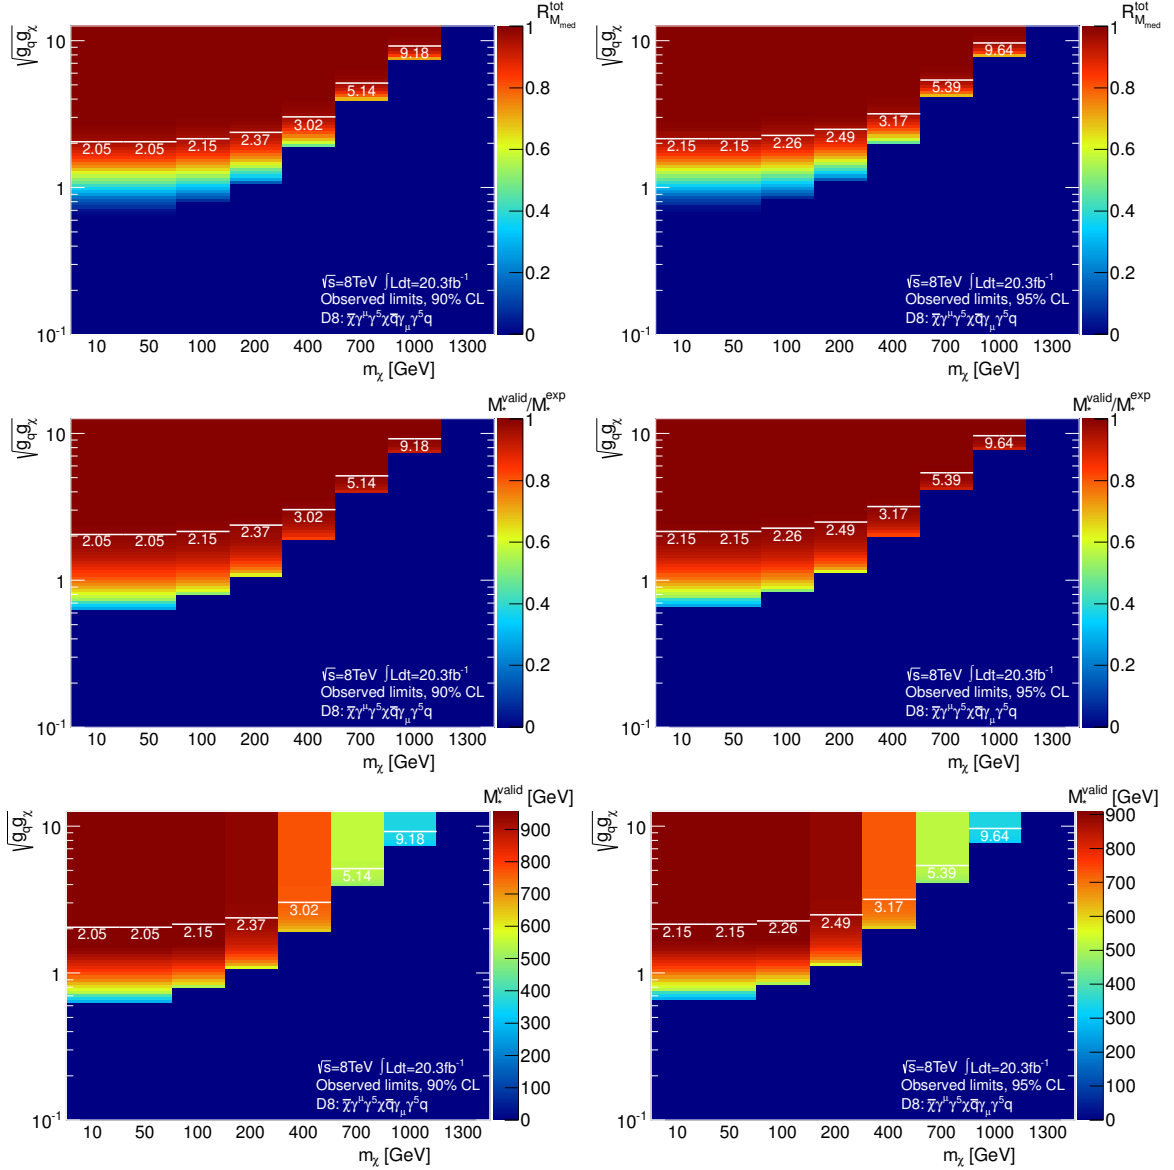


Figure G.6: The fraction of valid events (top), scale factor (middle), and rescaled limits (bottom) for the D8 operator at 90%CL (left) and 95%CL (right) observed limits as a function of the WIMP mass (x-axis) and couplings (y-axis). White lines and associated white numbers correspond to the minimum coupling value for which the truncated limit is within 1% of the nominal limit, or equivalently the scale factor is at least 0.99.

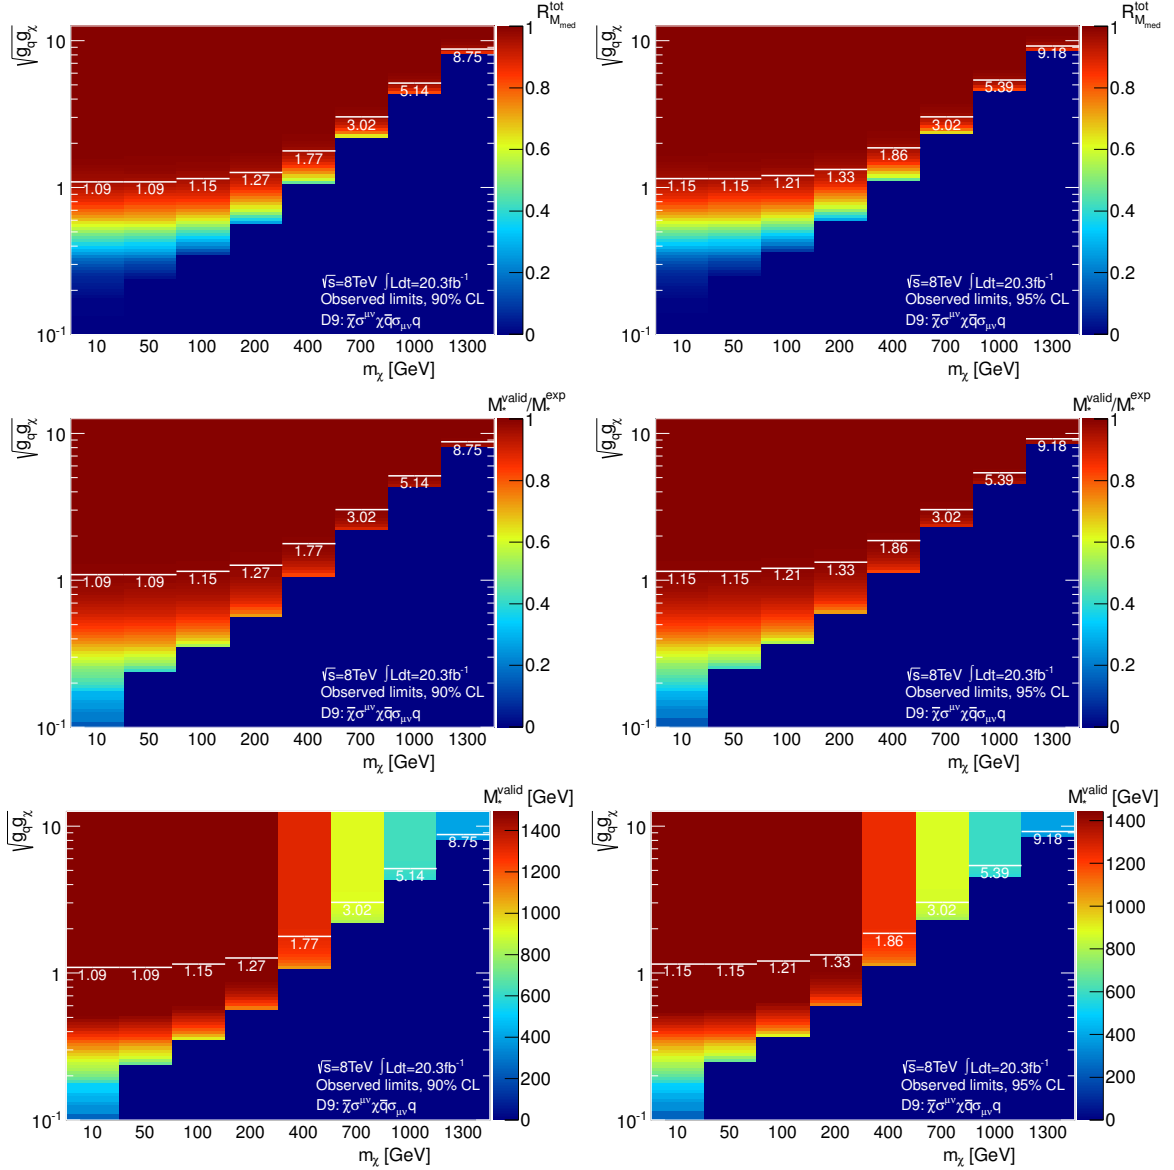


Figure G.7: The fraction of valid events (top), scale factor (middle), and rescaled limits (bottom) for the D9 operator at 90%CL (left) and 95%CL (right) observed limits as a function of the WIMP mass (x -axis) and couplings (y -axis). White lines and associated white numbers correspond to the minimum coupling value for which the truncated limit is within 1% of the nominal limit, or equivalently the scale factor is at least 0.99.

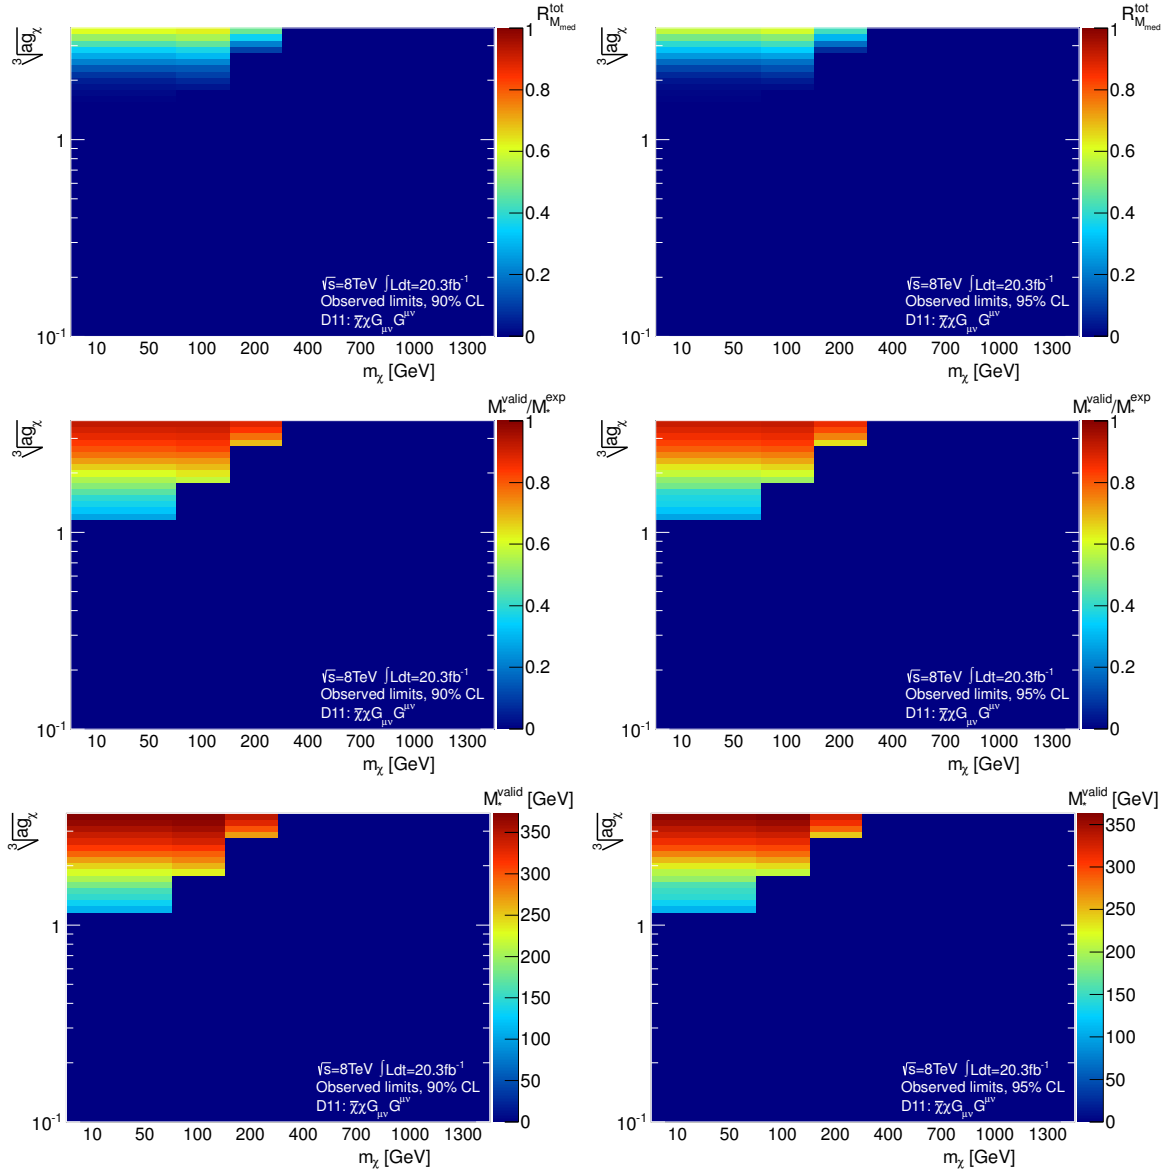


Figure G.8: The fraction of valid events (top), scale factor (middle), and rescaled limits (bottom) for the D11 operator at 90%CL (left) and 95%CL (right) observed limits as a function of the WIMP mass (x-axis) and couplings (y-axis). There is no point where the truncated limit is within 1% of the nominal limit.

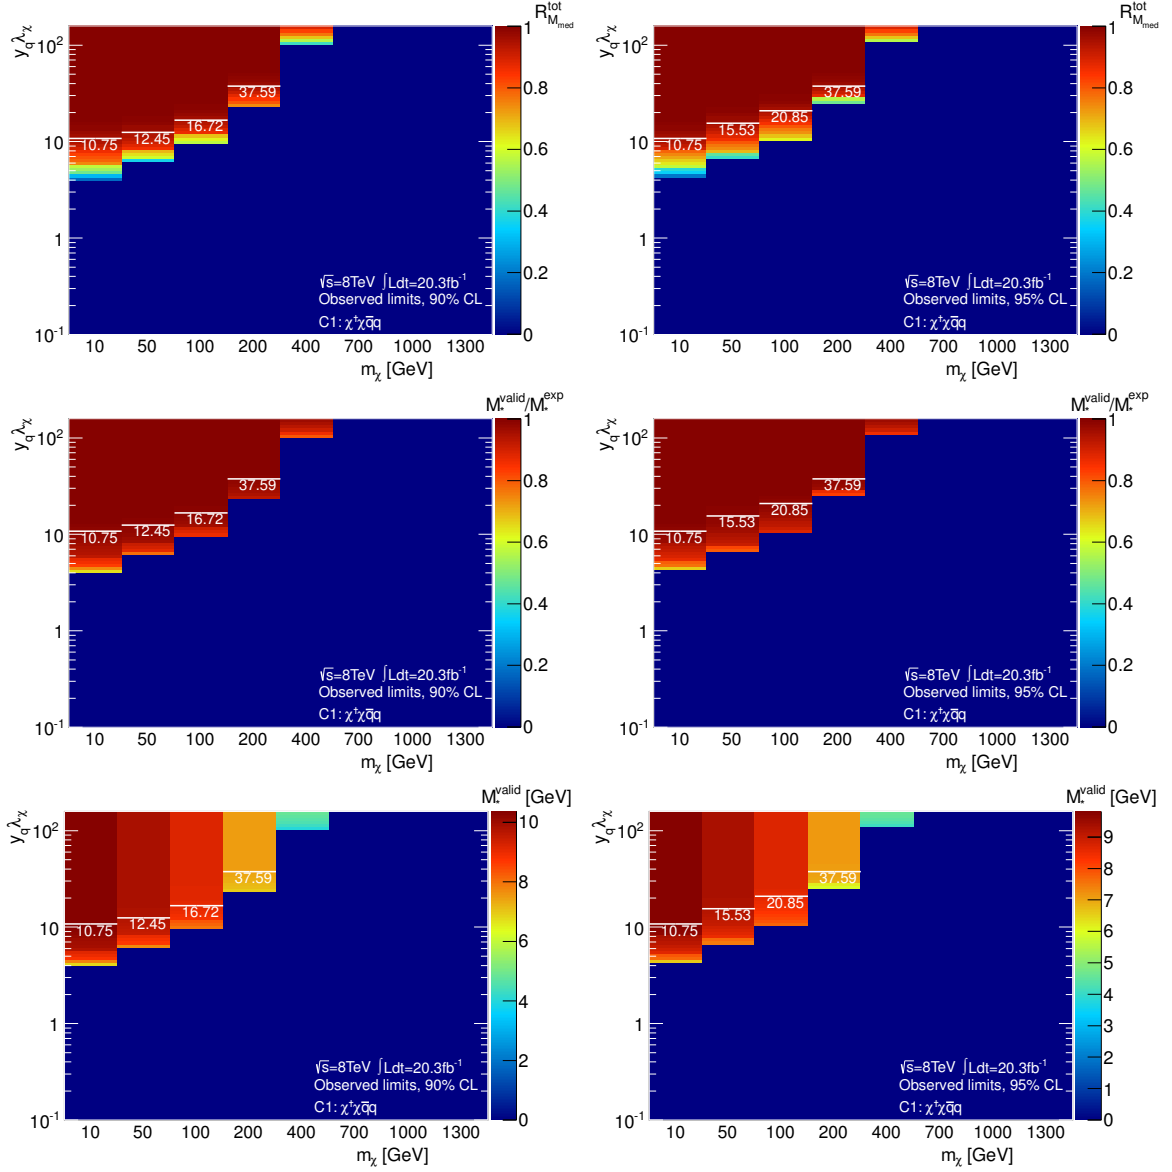


Figure G.9: The fraction of valid events (top), scale factor (middle), and rescaled limits (bottom) for the C1 operator at 90%CL (left) and 95%CL (right) observed limits as a function of the WIMP mass (x-axis) and couplings (y-axis). White lines and associated white numbers correspond to the minimum coupling value for which the truncated limit is within 1% of the nominal limit, or equivalently the scale factor is at least 0.99.

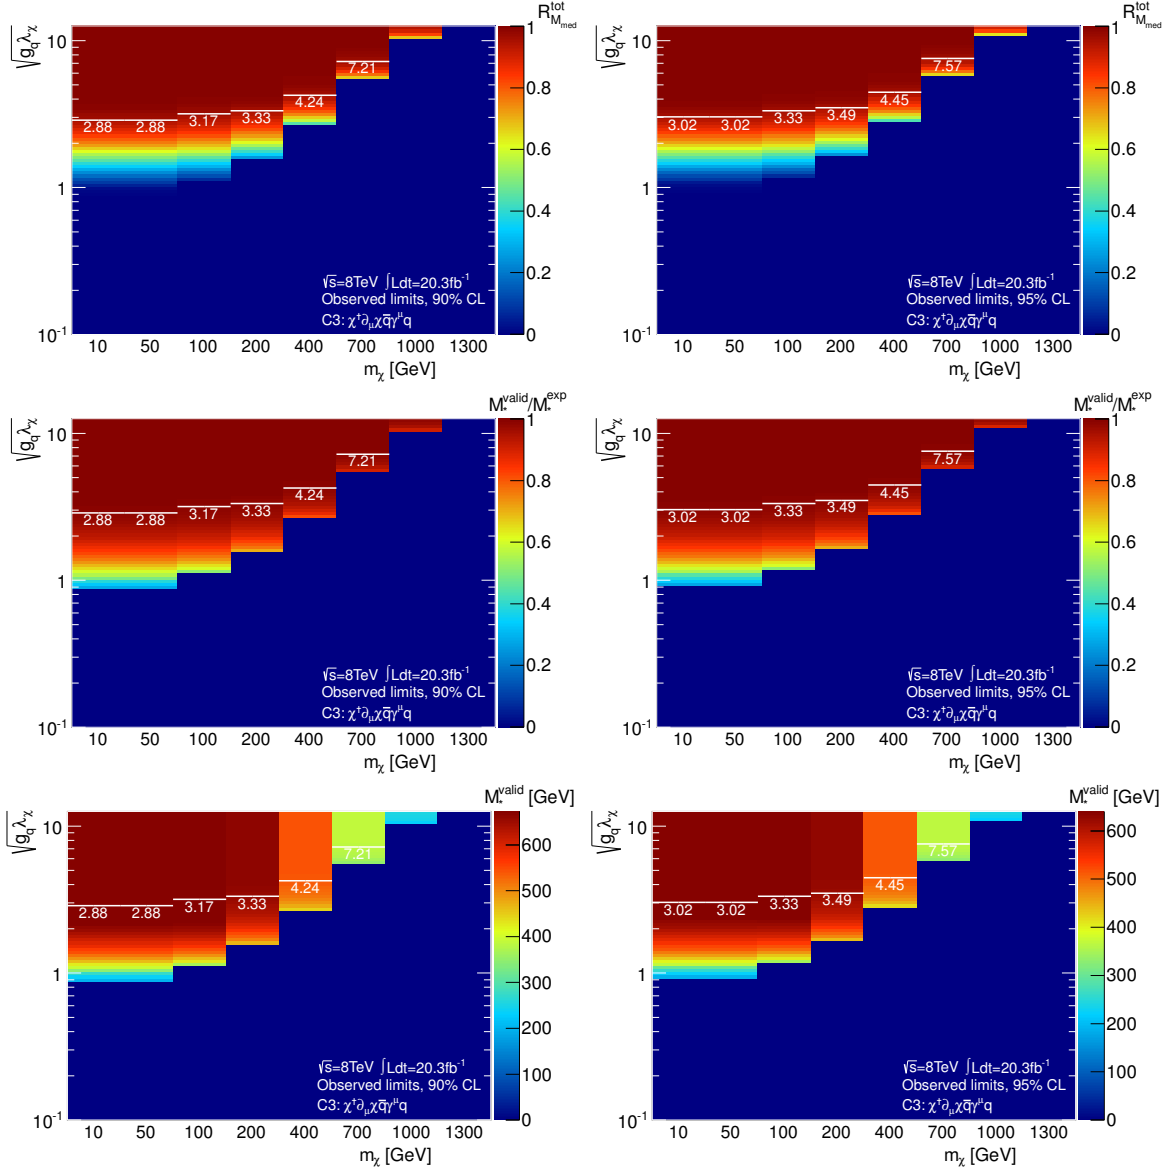


Figure G.10: The fraction of valid events (top), scale factor (middle), and rescaled limits (bottom) for the C3 operator at 90%CL (left) and 95%CL (right) observed limits as a function of the WIMP mass (x-axis) and couplings (y-axis). White lines and associated white numbers correspond to the minimum coupling value for which the truncated limit is within 1% of the nominal limit, or equivalently the scale factor is at least 0.99.

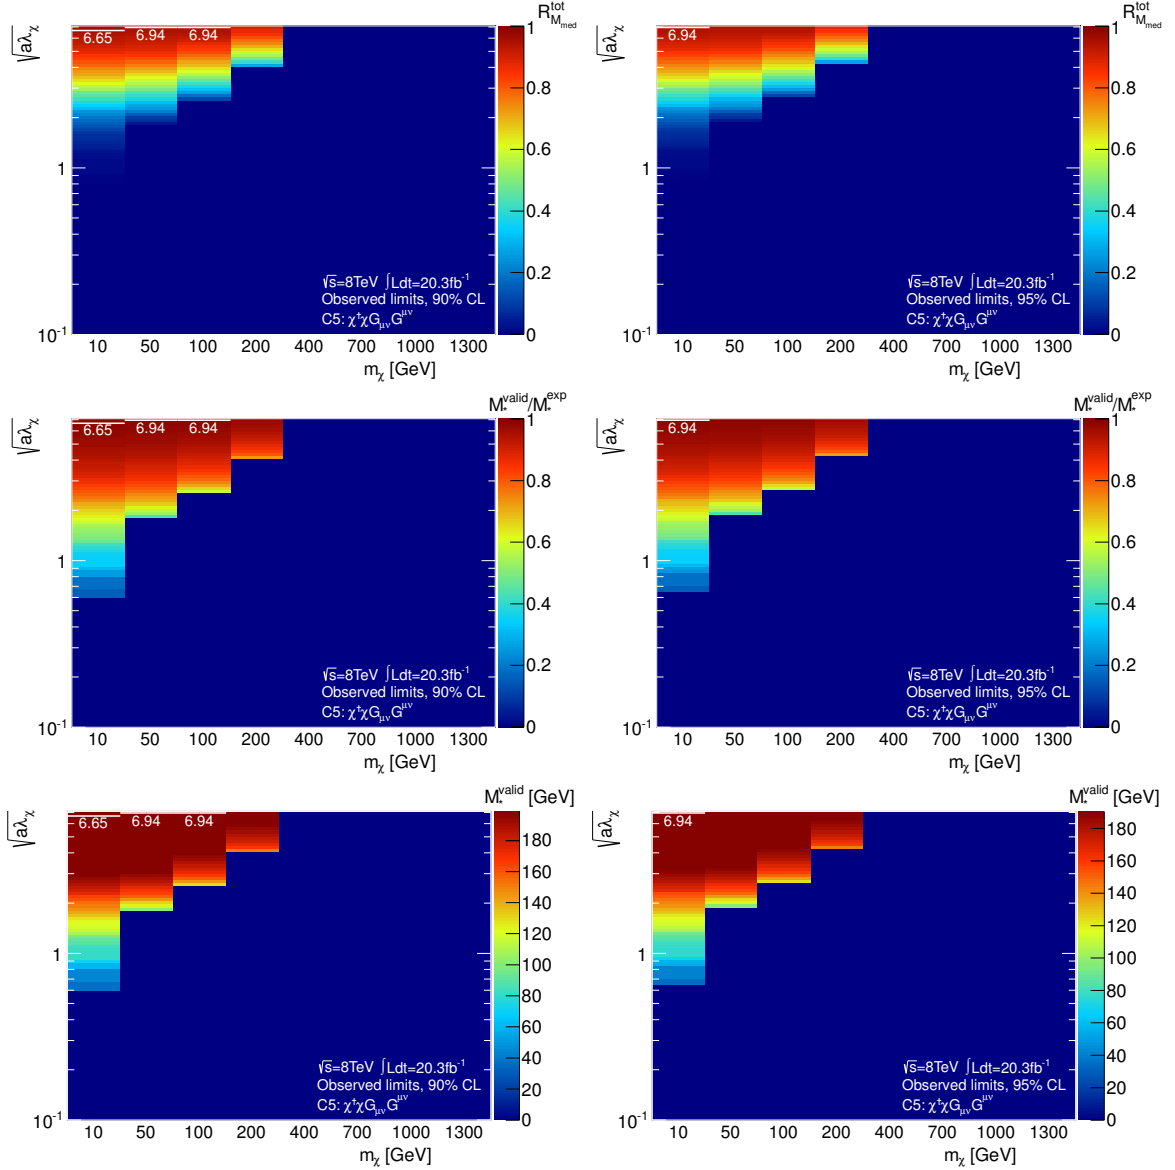


Figure G.11: The fraction of valid events (top), scale factor (middle), and rescaled limits (bottom) for the C5 operator at 90%CL (left) and 95%CL (right) observed limits as a function of the WIMP mass (x-axis) and couplings (y-axis). White lines and associated white numbers correspond to the minimum coupling value for which the truncated limit is within 1% of the nominal limit, or equivalently the scale factor is at least 0.99.

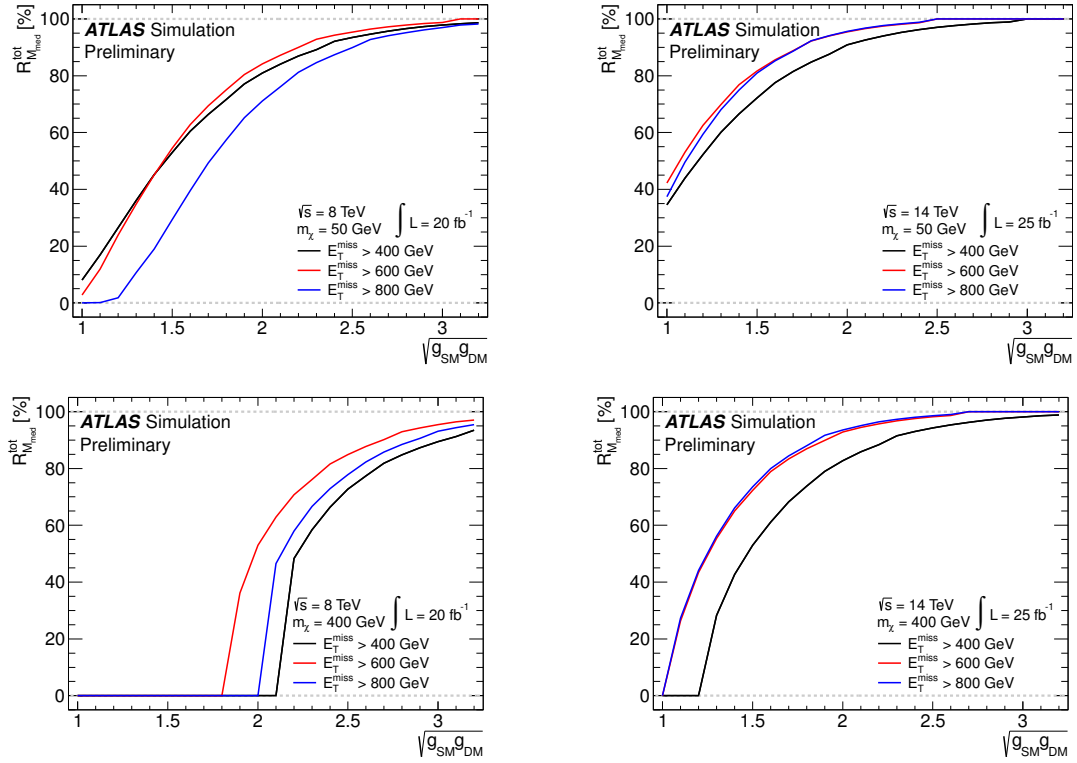


Figure G.12: The fraction of events satisfying the validity condition of $Q_{\text{tr}} < M_{\text{med}} = \sqrt{g_q g_\chi} M_\star$ as a function of the coupling $\sqrt{g_q g_\chi}$. DM particles with a mass of 50 GeV (top) and 400 GeV (bottom) are considered, at both 8 TeV (left) and 14 TeV (right). The three signal regions with different $E_{\text{T}}^{\text{miss}}$ cuts are shown, with 400 GeV in black, 600 GeV in red, and 800 GeV in blue. [8]

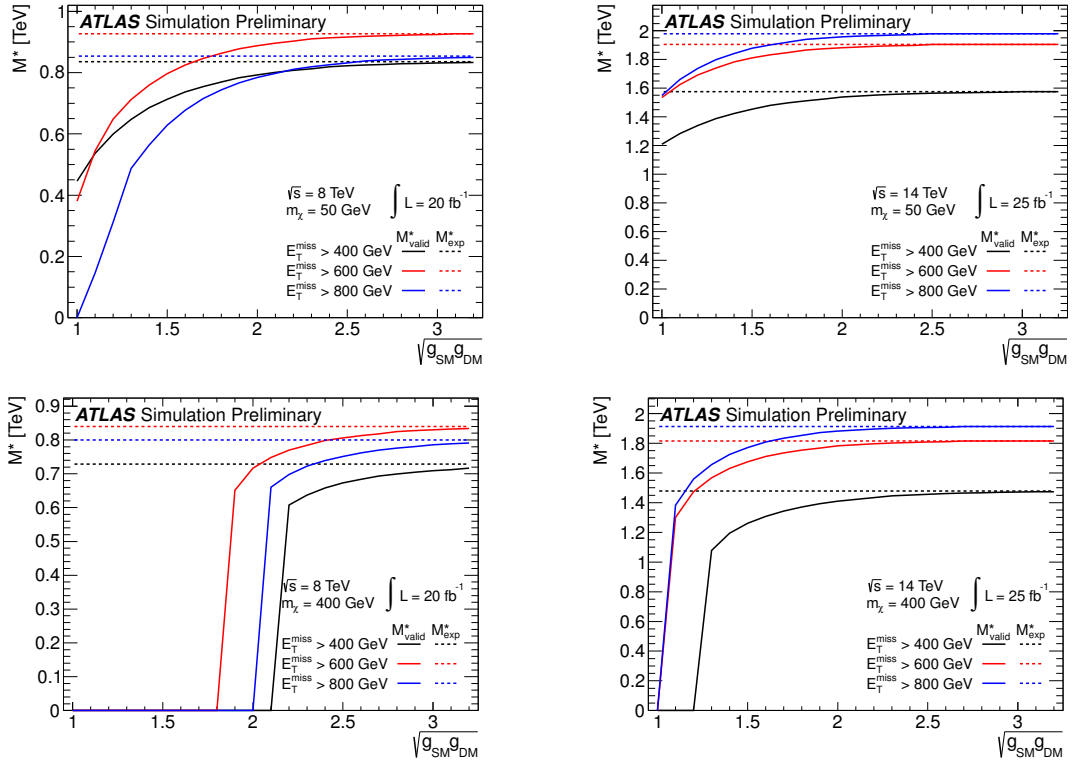


Figure G.13: Truncated limits M_*^{valid} (subject to the condition of $Q_{\text{tr}} < M_{\text{med}} = \sqrt{g_q g_\chi} M_*$) and nominal limits (excluding validity considerations) M_*^{exp} as a function of the coupling $\sqrt{g_q g_\chi}$. DM particles with a mass of 50 GeV (top) and 400 GeV (bottom) are considered, at both 8 TeV (left) and 14 TeV (right). The three signal regions with different E_T^{miss} cuts are shown, with 400 GeV in black, 600 GeV in red, and 800 GeV in blue. [8]

\sqrt{s} [TeV]	m_χ [GeV]	$E_T^{\text{miss cut}}$ [GeV]	M_*^{exp} [TeV]	$R_{M_{\text{med}}}^{\text{tot}}$ and M_*^{valid} [% , TeV]					
				$\sqrt{g_q g_\chi} = 1$		$\sqrt{g_q g_\chi} = 2$		$\sqrt{g_q g_\chi} = \pi$	
8	50	400	0.84	8.1	0.45	81	0.79	98	0.83
8	50	600	0.93	2.8	0.38	84	0.89	100	0.93
8	50	800	0.85	0	0	71	0.78	98	0.85
8	400	400	0.73	0	0	0	0	91	0.71
8	400	600	0.84	0	0	53	0.72	96	0.83
8	400	800	0.80	0	0	0	0	94	0.79
14	50	400	1.6	35	1.2	91	1.5	100	1.6
14	50	600	1.9	42	1.5	95	1.9	100	1.9
14	50	800	2.0	37	1.5	95	2.0	100	2.0
14	400	400	1.5	0	0	83	1.4	99	1.5
14	400	600	1.8	0	0	93	1.8	100	1.8
14	400	800	1.9	0	0	94	1.9	100	1.9

Table G.2: Validity results of the projected 8 and 14 TeV nominal limits M_*^{exp} for the D5 EFT operator with both 50 and 400 GeV DM masses. The fraction of valid events $R_{M_{\text{med}}}^{\text{tot}}$ and truncated limit M_*^{valid} are listed for the coupling choices of 1, 2, and π . All three signal regions and associated E_T^{miss} thresholds are provided for each scenario considered. [8]

region based on the strongest expected nominal limit. The impact of the E_T^{miss} threshold only becomes significant when the model is only marginally valid or when the difference between the strongest signal region and the next strongest is minimal, in which case the final truncated limits are also approximately equivalent but possibly with the order reversed.

Table G.2 lists the fraction of valid events, nominal limits, and truncated limits for a few specific points.

Appendix H

Mono-jet Dark Matter Kinematic Distributions

H.1 Kinematic plots

These plots supplement Figure 8.2 from Section 8.1, demonstrating the dependence of the E_T^{miss} and leading jet p_T distributions for a few different DM masses. The fermionic DM operators are shown in Figure H.1, while scalar DM operators are in Figure H.2, resulting in four and three distinct kinematic shapes respectively. The overlap between the D5 and C3 operators is shown in Figure H.3, bringing the total number of kinematically distinct operators to six, as summarized in Figure H.4.

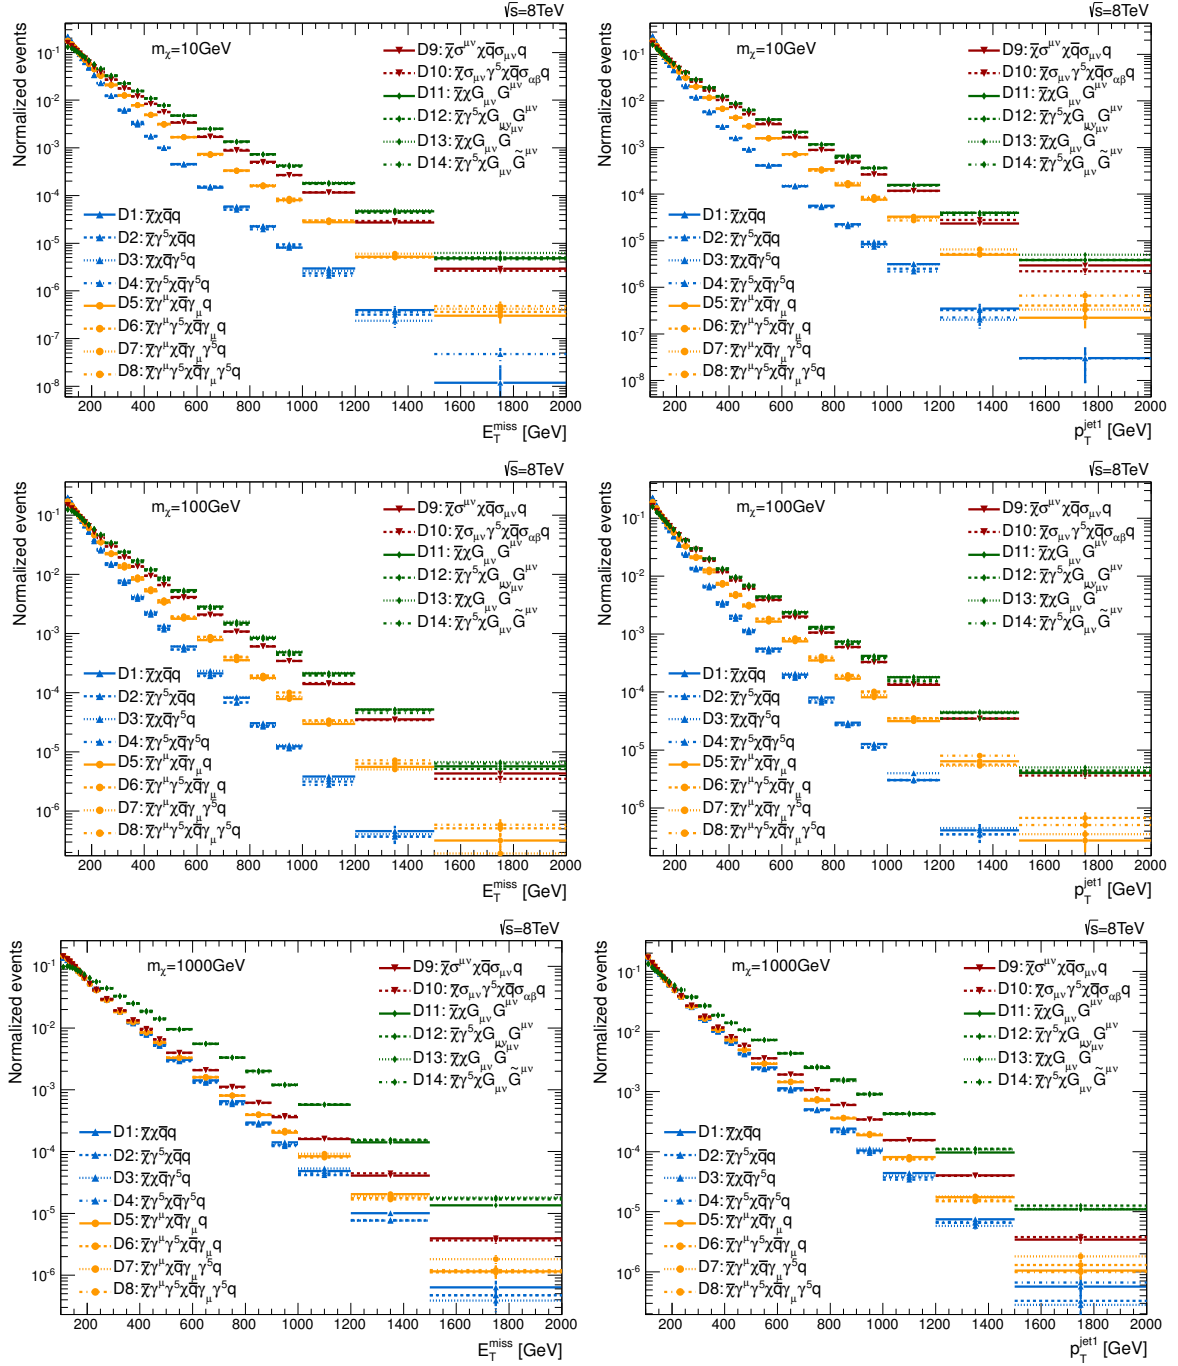


Figure H.1: E_T^{miss} (left) and leading jet p_T (right) distributions for EFT operators coupling to fermionic DM. Dark matter masses m_χ of 10 (top), 100 (middle), and 1000 GeV (bottom) are shown. These kinematic distributions reveal there are four types or classifications of fermion operators, which have the same acceptance as each other. These four types and their representative operators are scalar couplings to quarks (D1), vector couplings to quarks (D5), tensor couplings to quarks (D9), and scalar couplings to gluons (D11). Operators of the same type are given the same colour in order to be more visible, but different line styles, as all 14 operators are plotted.

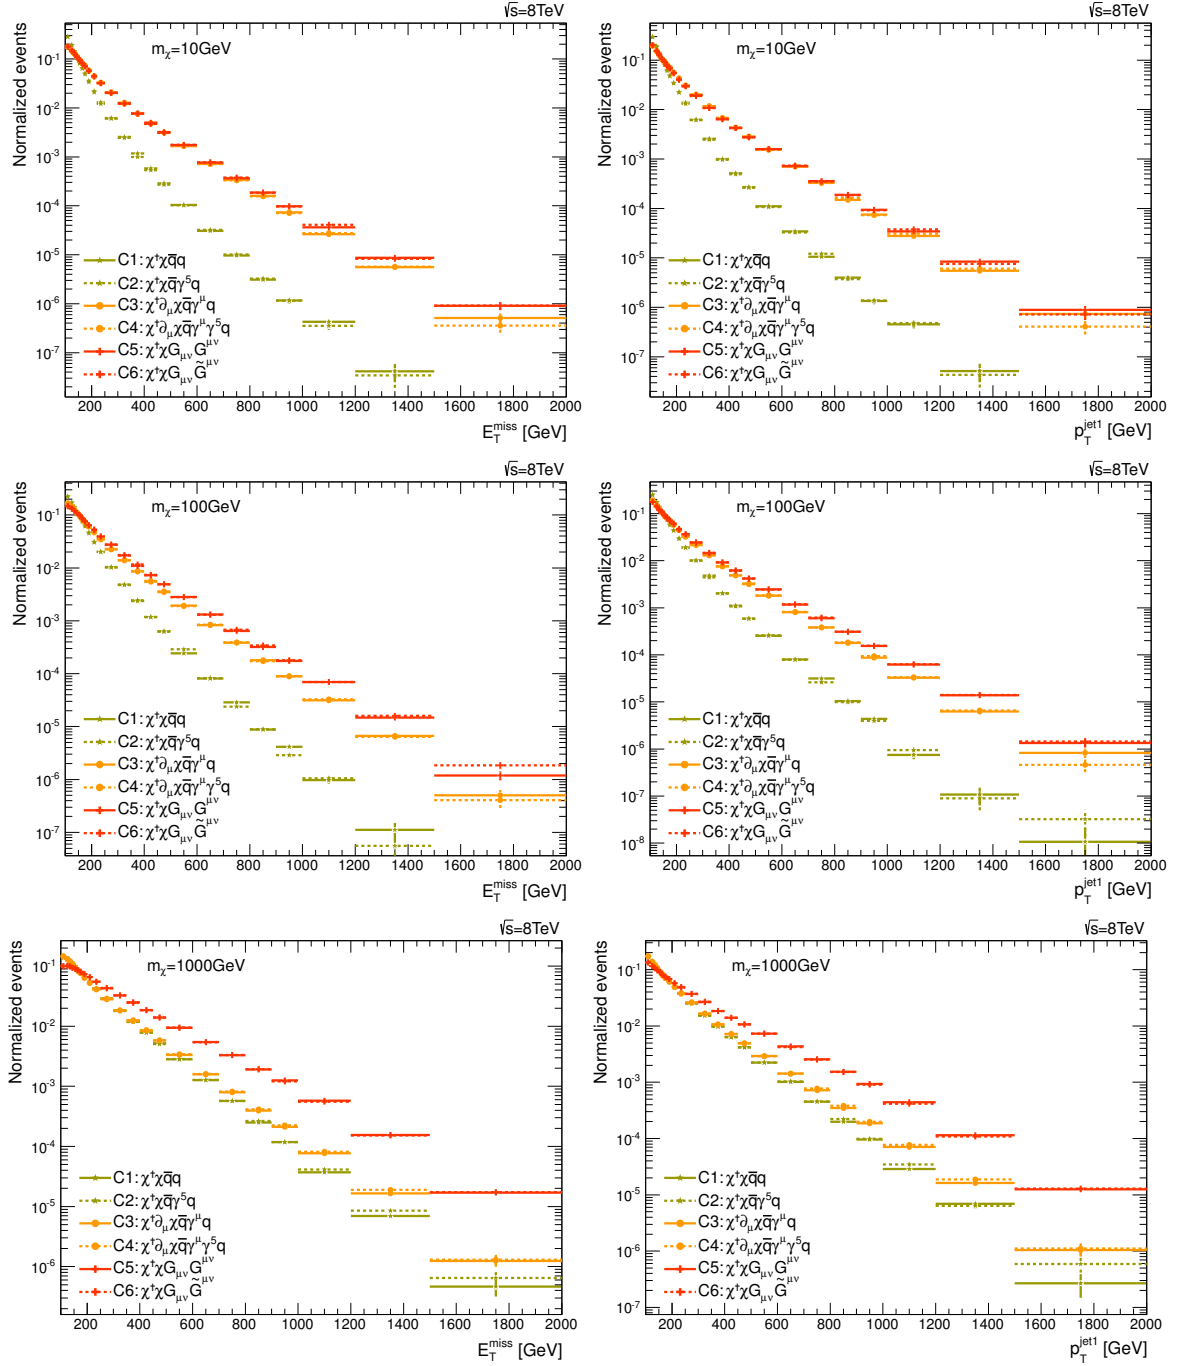


Figure H.2: E_T^{miss} (left) and leading jet p_T (right) distributions for EFT operators coupling to scalar DM. Dark matter masses m_χ of 10 (top), 100 (middle), and 1000 GeV (bottom) are shown. These kinematic distributions reveal there are three types or classifications of scalar operators, which have the same acceptance as each other. These three types and their representative operators are scalar couplings to quarks (C1), vector couplings to quarks (C3), and scalar couplings to gluons (C5). Operators of the same type are given the same colour in order to be more visible, but different line styles, as all 6 operators are plotted.

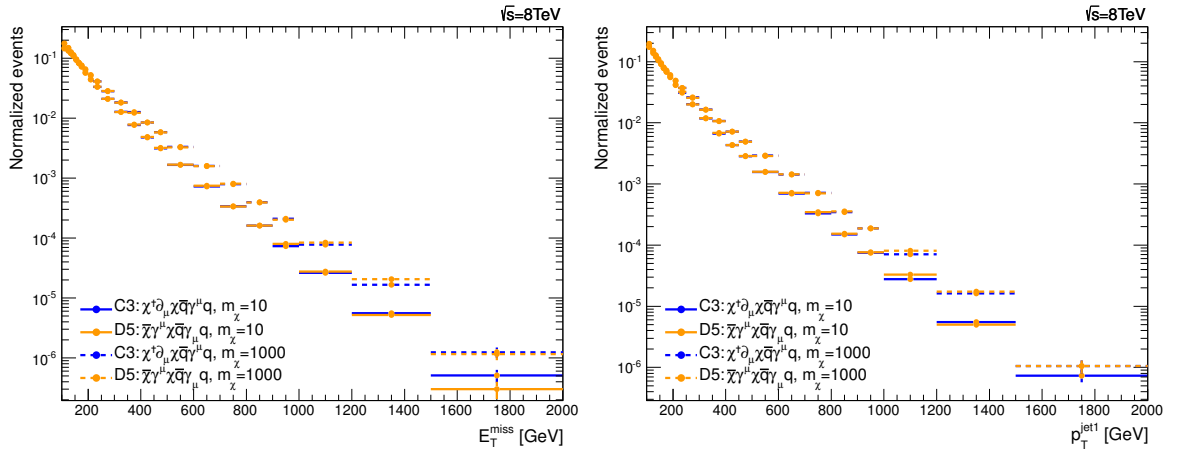


Figure H.3: E_T^{miss} (left) and leading jet p_T (right) comparisons for the D5 and C3 EFT operators. Dark matter masses m_χ of 10 and 1000 GeV are overlaid to show that the two operators agree for a wide range of masses. These kinematic distributions reveal that the D5 and C3 operators have the same acceptance, and thus a simple cross-section scaling transforms one to the other.

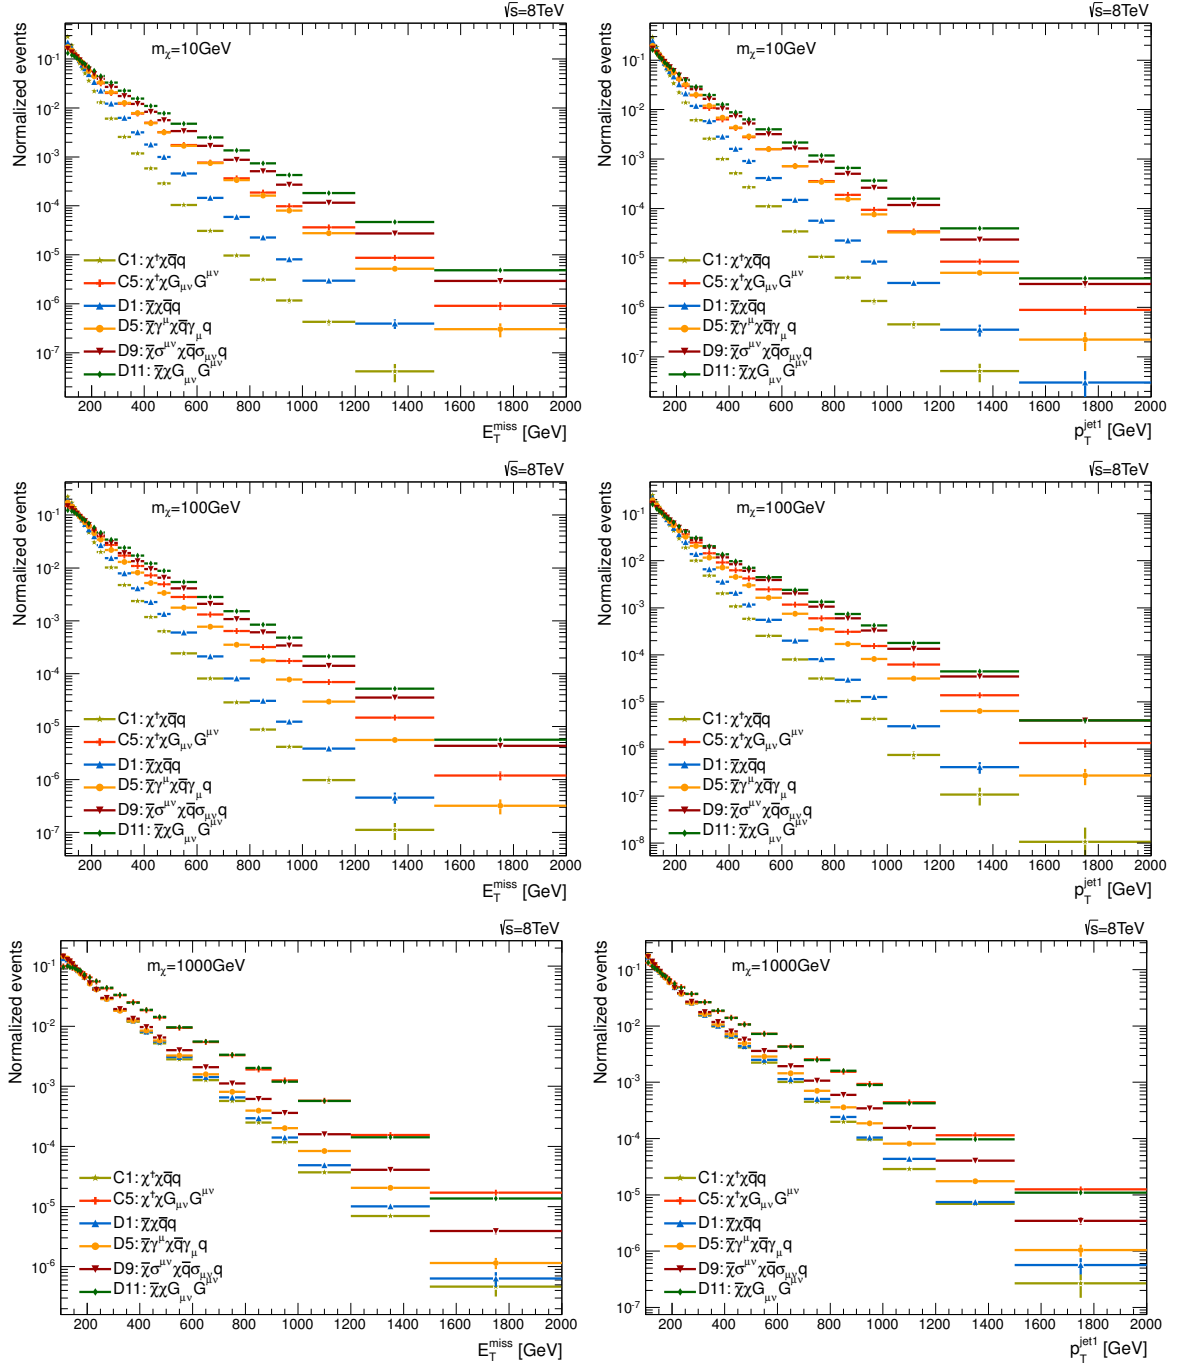


Figure H.4: E_T^{miss} (left) and leading jet p_T (right) distributions for the 6 representative EFT operators coupling to scalar or fermionic DM. Dark matter masses m_χ of 10 (top), 100 (middle), and 1000 GeV (bottom) are shown.

H.2 Acceptance plots

Once the kinematically distinct operators have been identified, it is important to determine where the DM mass becomes negligible. This varies by operator, as can be seen in Figure H.5. This Figure demonstrates that some operators become independent of the DM mass for values below 50 GeV, while others are independent by 10 GeV. In no case is the 1 GeV point needed, due to the energy scale of the LHC rendering 1 GeV particles effectively massless.

Similarly for the simplified models, it is important to determine the level of overlap between the pure vector Z' and pure axial vector Z' models. These are observed to agree for the large majority of the parameter space considered, as shown in Figure H.6. As such, only the pure vector Z' model was further investigated in order to reduce the required computing resources.

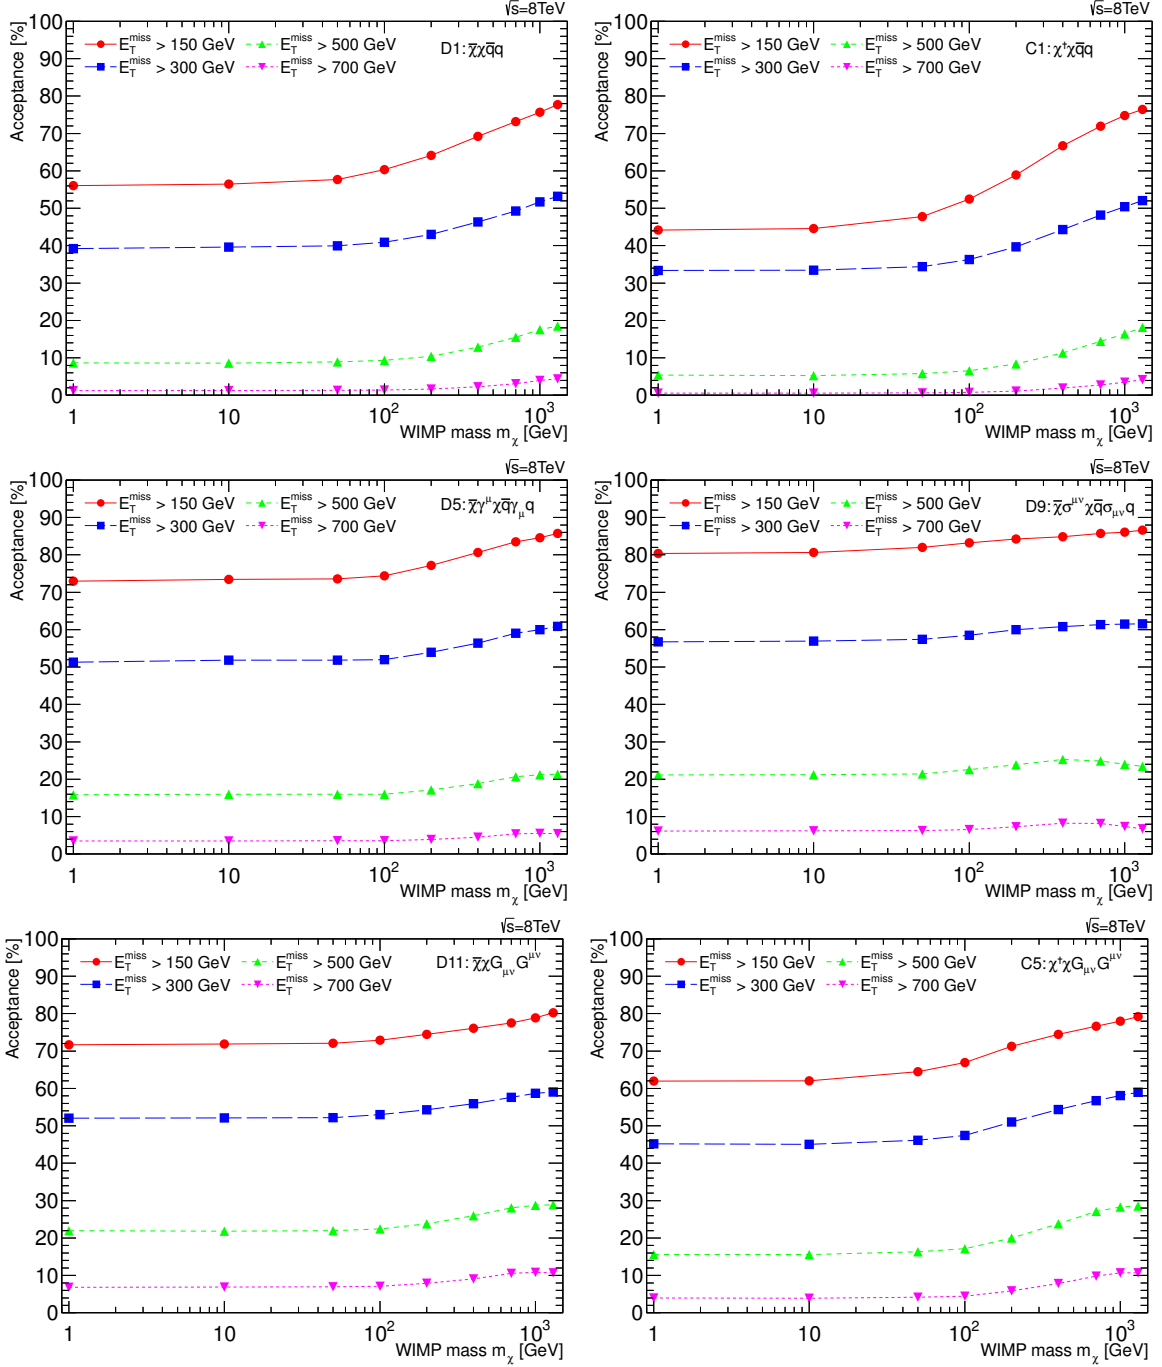


Figure H.5: The acceptance at truth level in percentages for differing DM masses, passing the basic mono-jet analysis selection plus a requirement of $E_T^{\text{miss}} > 150$ GeV (red), $E_T^{\text{miss}} > 300$ GeV (blue), $E_T^{\text{miss}} > 500$ GeV (green), or $E_T^{\text{miss}} > 700$ GeV (purple). The E_T^{miss} cuts were chosen to cover the different signal regions used by the mono-jet analysis. The representative operators shown are D1 (top left), C1 (top right), D5 (middle left), D9 (middle right), D11 (bottom left), and C5 (bottom right). From this, it is clear that m_χ of 1 and 10 GeV are indistinguishable for all operators. It was decided that m_χ of 10 and 50 GeV are also indistinguishable for the D1, D5, and D11 operators, keeping in mind that E_T^{miss} of 150 GeV will not be used for setting optimal limits. For some operators with harder E_T^{miss} spectra, such as D9, the acceptance turns over at high DM masses. This is due to a break-down of the matching scale used, as discussed in Appendix I.2.4.

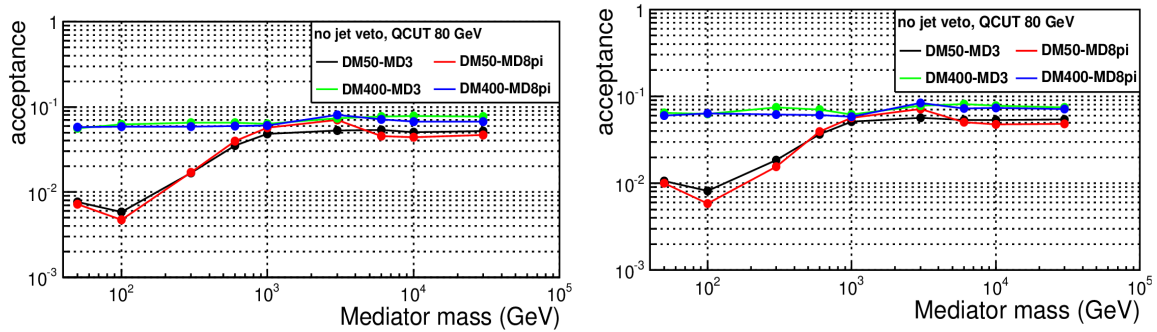


Figure H.6: The acceptance (at truth level) as a function of the mediator mass for mediator widths of the mass divided by 3 (black and green) or 8π (red and blue), for a DM particle with m_χ of 50 GeV (black and red) or 400 GeV (blue and green), for the vector (left) and axial vector (right) Z' simplified models. Small changes, possibly statistically insignificant, are observed for mediator masses of below 1 TeV for m_χ of 400 GeV. Larger and more significant differences are seen for mediator masses of below 1 TeV for m_χ of 50 GeV. The two models appear to have indistinguishable kinematic acceptance at the LHC excluding the region of mediator masses of $\mathcal{O}(100 \text{ GeV})$ combined with smaller values of m_χ .

Appendix I

Mono-jet Dark Matter Interpretation Signal Sample Details

I.1 Production and validation

I.1.1 Sample production

Sample production occurs at LO and is split into two steps. The ME calculation is done using `MadGraph5` [147], in which the final state is required to contain two DM particles and at least one parton with an optional second parton. Showering and hadronization is then conducted with `Pythia6` [48]. At the time of production, `Pythia8` [49] did not support generic stable non-interacting particles, and thus was not suitable for DM samples. Three types of ATLAS detector simulations are considered, full simulation with `GEANT4`, fast simulation with `AFII`, and truth simulation. In general, truth simulation is used to determine when samples are kinematically different and worth producing (such as seen in Section 8.1) or for generator-level systematic uncertainty studies, `AFII` is used for the official samples used to derive limits and other public results, and `GEANT4` is used to cross-check the `AFII` samples. The summary of `AFII` EFT samples produced are detailed in Table I.1, while the corresponding simplified model samples are listed in Table I.2.

When splitting the generation of a physics process into ME and parton shower components, there needs to be an agreement between the two generators in how to treat the boundary. In this case, `MadGraph5` should be generating the energetic partons (the hard process), while `Pythia6` should handle

Operator	DM masses [GeV]
D1	50, 100, 200, 400, 700, 1000, 1300
D5	50, 100, 200, 400, 700, 1000, 1300
D9	10, 50, 100, 200, 400, 700, 1000, 1300
D11	50, 100, 200, 400, 700, 1000, 1300
C1	10, 50, 100, 200, 400, 700, 1000, 1300
C5	10, 50, 100, 200, 400, 700, 1000, 1300

Table I.1: `AFII` EFT samples used for limits and public results. Operators where DM masses of 10 and 50 GeV were found to have equivalent kinematic acceptances under a mono-jet selection exclude 10 GeV points.

DM mass [GeV]	Mediator masses [GeV]	Mediator width factors ($\Gamma_{\text{med}}/M_{\text{med}}$)
10	10, 50, 100, 300, 600, 1000, 2000, 3000, 5000	1/3
50	50, 100, 300, 600, 1000, 3000, 6000, 10000, 30000	1/3, 1/8 π
	10, 2000, 5000	1/3
100	10, 50, 100, 300, 600, 1000, 2000, 3000, 5000	1/3
200	10, 50, 100, 300, 600, 1000, 2000, 3000, 5000	1/3
400	50, 100, 300, 600, 1000, 3000, 6000, 10000, 30000	1/3, 1/8 π
	10, 2000, 5000	1/3
700	10, 50, 100, 300, 600, 1000, 2000, 3000, 5000	1/3
1000	10, 50, 100, 300, 600, 1000, 2000, 3000, 5000	1/3
1300	10, 50, 100, 300, 600, 1000, 2000, 3000, 5000	1/3

Table I.2: AFII simplified model samples used for limits and public results. DM masses of 50 and 400 GeV were chosen as two baselines for a full scan of M_{med} vs couplings, while the remaining DM mass values are necessary for a scan of M_{med} vs m_χ .

the fragmentation and showering of those partons (the soft process). Recall that the samples are produced with either one or two partons at ME level. However, one-parton `MadGraph5` events can be showered into two jets by `Pythia6`, which is an indistinguishable final state from two-parton `MadGraph5` events and which leads to double-counting. By introducing a cutoff scale, known as the matching scale, this can be resolved. The matching scale is the threshold above which `MadGraph5` is allowed to produce jets (partons) and below which `Pythia6` is allowed to produce jets.

All of the samples listed in Tables I.1 and I.2 have two sub-samples, one with a matching scale of 80 GeV and the other with a matching scale of 300 GeV. This is primarily done to increase the statistics in events with more energetic jets and E_T^{miss} , as a matching scale cut of 300 GeV naturally requires that there be at least one parton of $p_T > 300$ GeV in the event. There second purpose of having two matching scales relates to limitations of the matching procedure, as will be discussed in Appendix I.2.4.

I.1.2 Model validation

The pure vector and pure axial vector Z' simplified models from Reference [95] provide a useful means of cross-checking a subset of the EFT operators provided by Reference [92]. For very massive mediators, where the simplified model enters the contact interaction regime, it should directly correspond to the EFT operator. This provides a useful means of validating models against each other. This cross-check was performed for both cases, leading to good agreement. This is shown for the pure vector model and the D5 EFT operator in Figure I.1, and for the pure axial vector model and D8 EFT operator in Figure I.2.

Reference [95] additionally provides a contact interaction model for gluons coupling to fermionic DM, analogous to the D11 EFT operator from Reference [92]. The definition of the two models differs by a factor of 4, where the former absorbs the factor into the suppression scale and the latter leaves it explicitly in the functional form ($4M_*^3 = \Lambda^3$). If this factor of 4 is accounted for, the two can be directly compared, and identical results are expected. This test has also been conducted, and leads to very good agreement within statistics, as seen in Figure I.3.

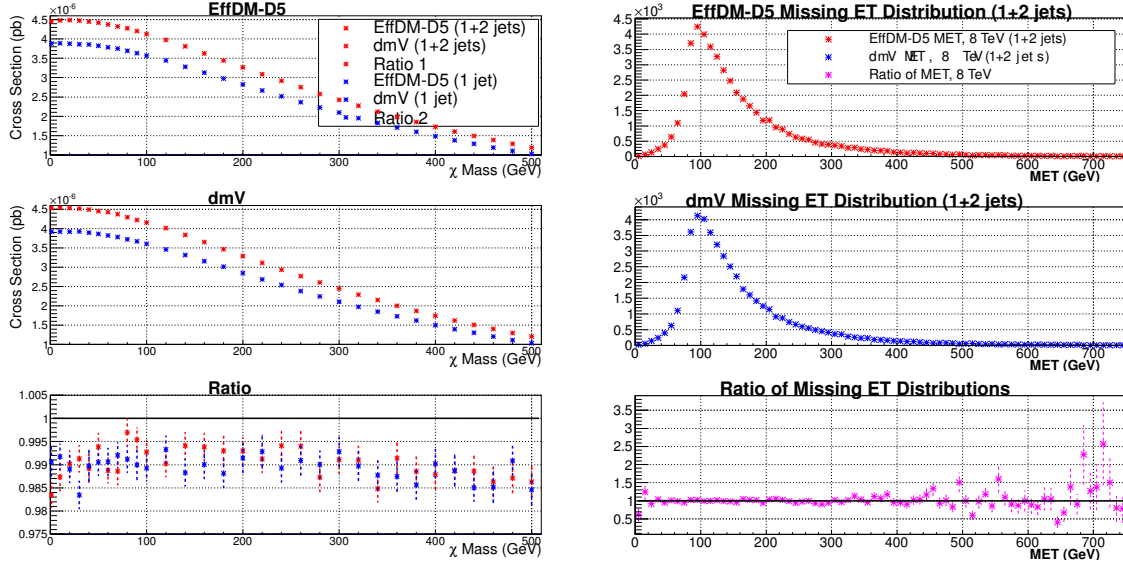


Figure I.1: A comparison of the pure vector Z' simplified model coupling quarks to fermionic DM [95] (middle frame) and the corresponding EFT operator [92] (top frame), for a mediator mass M_{med} of 20 TeV. The ratio of the two, EFT/ Z' , is shown in the bottom frame. The cross-sections (left) and $E_{\text{T}}^{\text{miss}}$ distributions (right) are compared. Cross-sections are compared at both single-parton (1 jet) and single-parton plus an optional second parton (1+2 jets) level. Good agreement is observed in all cases, with the cross-section offset by 1%. This difference is in the expected direction, with the model including the mediator resulting in a slightly higher cross-section due to residual contributions. The $E_{\text{T}}^{\text{miss}}$ distributions are seen to be consistent within statistical uncertainties.

I.1.3 Fast simulation validation

Fast simulation, or AFII, was used for the generation of all DM samples intended to be shown in public results. AFII has been generally shown to be in good agreement with the GEANT4 full simulation, but it is still important to ensure that the fast parametrization provides a good modelling of the detector response after the mono-jet selection. The mono-jet selection depends primarily on jet and $E_{\text{T}}^{\text{miss}}$ related variables, but reconstructed $E_{\text{T}}^{\text{miss}}$ is calculated from the sum of visible particles. As such, jet-related quantities are crucial to properly model, and are the main point of comparison that needs to be made between AFII and GEANT4. The difference between fully simulated and fast simulation jets has already been studied in detail by the relevant experts. Differences in jet performance between the two simulation types are already provided in the form of an additional uncertainty on the JES, as discussed in Appendix E. As such, so long as any difference observed between AFII and GEANT4 DM samples is within this uncertainty and the associated statistical uncertainty, there is no need to add an additional systematic for the simulation type.

In order to study this effect, one or two samples per EFT operator were produced with both AFII and GEANT4. Table I.3 lists the specific samples produced as well as the results of a comparison of the mono-jet acceptance for each pair of samples. A comparison of different operators is not straightforward, particularly given that many of the comparisons are equal to zero within their respective statistical uncertainty. Thus there is no need for an extra systematic. The one operator to clearly stand out is C1.

Studies have shown that the difference between GEANT4 and AFII with respect to jets is largest at low p_{T} and falls off as the jet p_{T} increases, resulting in smaller uncertainties. As such, operators

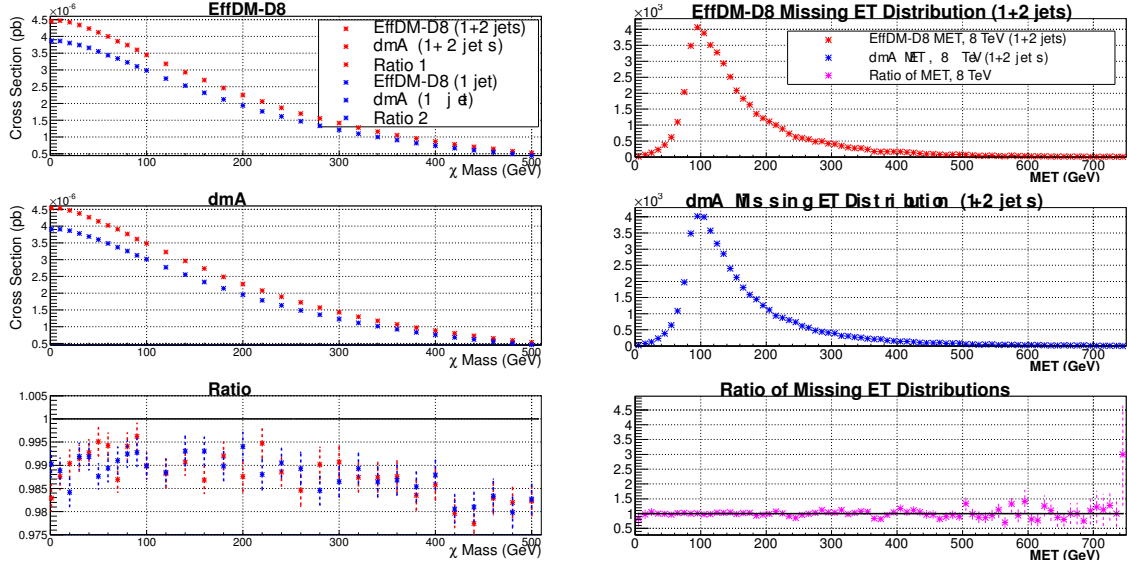


Figure I.2: A comparison of the pure axial vector Z' simplified model coupling quarks to fermionic DM [95] (middle frame) and the corresponding EFT operator [92] (top frame), for a mediator mass M_{med} of 20 TeV. The ratio of the two, EFT/ Z' , is shown in the bottom frame. The cross-sections (left) and E_T^{miss} distributions (right) are compared. Cross-sections are compared at both single-parton (1 jet) and single-parton plus an optional second parton (1+2 jets) level. Good agreement is observed in all cases, with the cross-section offset by 1%. This difference is in the expected direction, with the model including the mediator resulting in a slightly higher cross-section due to residual contributions. The E_T^{miss} distributions are seen to be consistent within statistical uncertainties.

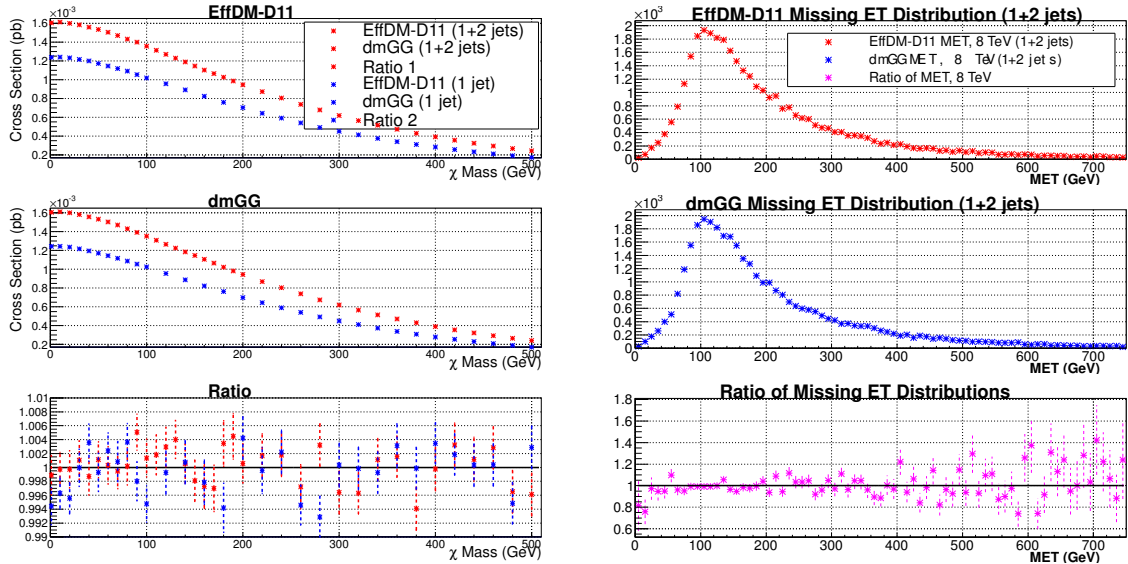


Figure I.3: A comparison of two different EFT models of gluons coupling to fermionic DM, from Reference [95] (middle frame) and Reference [92] (top frame). The ratio of the two is shown in the bottom frame. The cross-sections (left) and E_T^{miss} distributions (right) are compared. Cross-sections are compared at both single-parton (1 jet) and single-parton plus an optional second parton (1+2 jets) level. Agreement within statistical uncertainties is observed in all cases, for both the cross-section and E_T^{miss} distributions.

Operator	DM mass [GeV]	N _{events}		Relative difference [%]
		GEANT4	AFII	
D1	1000	6588	6510	1.19
D5	1000	7072	6927	2.09
D9	50	6933	6829	1.52
	1000	7150	7022	1.82
D11	1000	8533	8477	0.66
C1	50	2492	2374	4.97
C5	50	4908	4801	2.22

Table I.3: EFT samples generated with both AFII and GEANT4 for cross-checking the fast simulation parametrization after the mono-jet selection. Only the 80 GeV matching scale sub-sample was generated for GEANT4, thus only the 80 GeV matching scale sub-sample was used in comparisons. Numbers of events are the raw unweighted and non-normalized numbers, but the AFII and GEANT4 samples start with the same initial number of events before the mono-jet selection. Statistical uncertainties plus the MC-specific JES uncertainty cover most of the observed differences, with the exception of C1.

with softer jet p_T distributions should be more affected. Similarly, lower DM masses result in lower momentum transfer requirements and thus lower jet p_T requirements. The type of jet could also play a role, especially when comparing topologies with gluon or light quark jets (for which the uncertainty was derived) vs heavy flavour jets. None of these effects is clearly visible when comparing the other operators, but the differences are all within statistics, so it is difficult to quantify whether this motivation is valid. However, C1 is the extreme case. It has the softest jet kinematic distributions, the DM mass used is very light, and the operator prefers heavy flavour quark jets. These three effects combined make it clear that the validity of using AFII for the C1 operator is the most questionable, as is observed. The statistics and JES combined do not fully cover the observed difference, but it is unclear if this is a large statistical fluctuation or a real effect. Given that other systematic uncertainties for C1 dwarf any value that could be assigned here, as discussed in Appendix I.2, no extra MC-specific systematic was assigned for the C1 operator.

I.1.4 Acceptance and efficiency

Once samples are produced and validated, the next step is to determine the acceptance and efficiency of the samples for use in limit setting and to allow theorists to re-interpret the results in the form of their own models. The acceptance is the fraction of truth-level events which pass the analysis selection. The efficiency is then the level to which this is replicated when moving from truth information to reconstructed quantities. Thus the observed number of events in data must be compared to the number of events times the acceptance times the efficiency. More explicitly stated:

- Acceptance: $\mathcal{A} = N_{\text{events}}^{\text{truth}}/N_{\text{events}}^{\text{all}}$
- Efficiency: $\varepsilon = N_{\text{events}}^{\text{reco}}/N_{\text{events}}^{\text{truth}}$
- Observed: $N_{\text{events}}^{\text{reco}} = N_{\text{events}}^{\text{all}} \times \mathcal{A} \times \varepsilon$
- $N_{\text{events}}^{\text{all}}$: the weighted number of all events in the sample, before any selection (except for overlap removal if necessary, in order to not double-count events)
- $N_{\text{events}}^{\text{truth}}$: the weighted number of events passing the truth-level mono-jet selection

Operator	m_χ [GeV]	E_T^{miss} threshold [GeV]								
		150	200	250	300	350	400	500	600	700
D1	50	29.63×81.1	15.84×80.4	8.51×78.1	4.64×80.7	2.85×76.6	1.69×80.0	0.64×73.2	0.24×83.6	0.10×89.1
	100	34.40×80.5	19.49×79.2	10.70×79.6	6.17×81.6	3.80×77.4	2.11×78.1	0.85×72.0	0.33×74.7	0.14×73.5
	200	44.03×80.6	26.78×78.3	15.25×79.2	9.40×77.4	5.57×77.5	3.38×78.5	1.32×78.4	0.59×78.4	0.25×76.7
	400	66.03×79.9	42.91×77.4	26.05×76.6	16.43×75.2	10.56×76.7	6.90×74.7	2.88×75.2	1.19×73.8	0.52×72.8
	700	96.21×81.7	65.71×80.3	43.15×78.9	28.57×77.3	19.06×75.8	12.47×77.6	5.73×76.2	2.63×76.6	1.21×73.6
	1000	133.19×81.2	92.27×79.7	61.26×79.8	40.91×78.7	27.01×79.1	18.64×78.5	8.78×78.0	4.05×75.0	1.81×74.8
	1300	161.69×80.8	113.61×78.4	74.83×77.0	49.95×77.4	33.59×79.7	23.15×79.5	10.57×77.2	5.01×75.9	2.39×75.6
D5	50	28.12×81.8	17.41×82.1	10.74×78.3	6.73×79.7	4.50×76.2	2.88×77.7	1.25×76.5	0.55×79.7	0.27×78.5
	100	29.55×81.4	18.90×79.6	11.86×79.7	7.39×77.7	4.68×78.0	3.04×79.6	1.32×78.1	0.61×77.3	0.28×79.8
	200	31.45×81.6	20.43×80.4	12.92×79.1	8.14×78.1	5.40×77.3	3.63×77.8	1.59×77.9	0.73×77.8	0.35×80.2
	400	33.69×81.2	22.73×79.1	14.78×79.2	10.00×78.2	6.59×78.0	4.45×76.3	2.06×77.1	0.99×79.0	0.48×79.6
	700	37.04×81.0	26.01×79.2	17.24×79.1	11.59×79.3	7.80×79.3	5.41×78.4	2.62×76.9	1.31×77.4	0.65×79.2
	1000	38.41×81.7	27.31×79.9	18.12×79.0	12.41×76.7	8.36×78.9	5.74×77.1	2.79×76.5	1.35×78.6	0.67×77.8
	1300	38.33×81.6	26.76×80.2	18.19×78.8	12.33×79.4	8.47×78.7	5.77×78.8	2.74×77.3	1.34×77.5	0.68×76.2
D9	10	34.25×81.9	23.16×80.9	15.98×79.4	10.81×79.9	7.56×79.5	5.31×79.1	2.69×78.7	1.42×79.9	0.77×77.3
	50	35.03×82.6	24.43×82.2	16.61×81.1	11.46×80.0	8.01×80.4	5.63×78.7	2.88×79.1	1.57×78.3	0.88×75.8
	100	36.63×82.0	25.74×80.6	17.58×80.9	12.51×79.5	8.98×79.3	6.44×78.9	3.32×78.7	1.80×76.6	0.99×76.5
	200	37.32×82.2	26.34×81.4	18.94×80.8	13.42×79.8	9.74×79.0	7.03×78.2	3.71×78.3	1.95×79.8	1.14×76.9
	400	38.51×82.3	27.62×81.3	19.79×79.9	13.78×80.2	10.06×79.3	7.28×80.4	3.90×80.1	2.25×78.4	1.31×76.8
	700	38.05×82.6	26.66×81.0	18.59×79.8	13.02×79.1	9.24×79.7	6.77×78.3	3.77×78.9	2.11×77.3	1.19×77.5
	1000	39.02×81.5	27.96×79.7	19.47×79.1	13.47×80.1	9.61×78.1	6.83×78.8	3.60×78.0	1.96×78.8	1.05×77.9
1300	38.99×81.2	27.57×79.5	18.97×80.1	13.03×78.3	9.02×79.0	6.46×78.7	3.29×78.1	1.67×80.5	0.89×77.7	
D11	50	33.99×82.5	25.53×80.5	18.86×79.6	13.88×79.8	9.94×77.2	7.22×77.1	3.83×76.6	2.09×75.3	1.16×74.3
	100	36.24×81.1	27.61×80.1	19.88×80.3	14.73×78.9	10.59×79.8	7.87×79.1	4.24×79.3	2.39×78.7	1.33×77.6
	200	38.44×81.0	30.10×79.9	22.32×78.7	16.44×79.4	12.36×78.4	8.98×78.3	4.89×76.8	2.82×76.4	1.54×76.9
	400	41.40×82.2	33.15×80.2	25.33×78.8	19.15×78.3	14.35×77.6	10.82×78.3	6.30×77.3	3.66×77.5	2.14×76.7
	700	44.67×81.7	36.54×80.2	27.91×79.8	21.22×80.1	16.43×79.1	12.66×78.7	7.51×78.0	4.40×78.5	2.65×78.2
	1000	45.66×81.7	37.42×79.7	29.23×78.7	22.44×78.7	17.30×77.2	13.29×78.2	7.91×77.4	4.71×77.3	2.80×79.1
	1300	44.67×83.6	36.80×79.7	28.88×77.4	22.27×78.8	17.18×78.5	13.01×77.6	7.71×76.5	4.40×76.2	2.65×75.7
C1	10	12.46×86.0	5.20×79.1	2.28×80.3	1.10×78.1	0.54×77.4	0.30×67.8	0.08×77.0	0.03×76.0	∅
	50	16.77×81.6	7.18×82.1	3.44×74.8	1.70×78.8	0.84×72.3	0.44×72.5	0.13×76.0	0.05×70.6	∅
	100	21.72×84.0	10.68×79.8	5.22×78.5	2.53×80.3	1.32×79.2	0.70×81.1	0.23×83.0	0.09×79.0	0.03×91.0
	200	32.31×82.4	18.13×79.8	10.05×76.9	5.36×77.8	3.05×75.5	1.75×75.9	0.60×75.9	0.23×73.5	0.09×74.7
	400	52.87×81.3	32.94×80.7	19.62×80.2	11.94×78.5	7.34×77.6	4.48×79.2	1.77×78.2	0.77×74.6	0.31×75.5
	700	81.57×81.1	55.21×78.5	35.79×78.1	22.84×78.5	15.26×78.1	10.14×76.3	4.57×77.7	1.95×78.2	0.80×75.6
	1000	118.34×82.0	81.98×79.7	53.39×78.4	35.50×77.6	24.18×77.6	16.31×77.6	7.38×77.0	3.33×75.6	1.56×75.3
1300	156.86×81.0	110.16×79.3	72.65×77.5	48.21×77.2	32.53×77.6	22.60×76.8	10.36×76.7	4.83×76.5	2.36×73.5	
C5	10	22.41×81.0	14.34×80.4	9.07×80.1	5.96×78.7	3.88×77.0	2.62×75.6	1.18×76.5	0.58×75.5	0.29×75.1
	50	25.49×81.1	17.00×79.9	11.24×77.9	7.28×78.9	4.86×77.9	3.23×78.1	1.54×75.9	0.72×76.3	0.36×75.0
	100	28.72×80.8	20.17×79.1	13.37×78.0	9.03×77.8	6.30×77.8	4.32×76.5	2.05×76.8	0.96×73.6	0.49×75.2
	200	33.82×81.3	25.02×79.4	17.64×79.5	12.48×77.7	8.78×78.0	6.25×77.5	3.13×79.7	1.59×80.9	0.83×80.6
	400	38.82×81.9	30.37×80.3	22.58×79.8	16.84×79.9	12.40×79.9	9.30×77.6	5.05×78.0	2.81×76.6	1.58×76.1
	700	43.00×82.1	34.84×79.7	26.83×79.6	20.62×77.9	15.80×78.7	12.12×78.6	7.13×78.7	4.25×77.7	2.50×77.2
	1000	44.34×82.1	36.78×79.3	28.34×79.8	21.93×79.1	16.99×78.5	13.17×78.5	7.80×78.0	4.66×79.1	2.81×77.4
1300	47.31×82.2	38.61×81.1	30.17×80.8	22.57×80.2	17.68×78.4	13.54×78.5	7.98×78.5	4.84×78.2	2.84×78.6	

Table I.4: The acceptance \times efficiency, in %, for the EFT operator samples. The value \emptyset denotes that the acceptance and efficiency were not evaluated due to insufficient statistics.

- $N_{\text{events}}^{\text{reco}}$: the weighted number of events passing the reco-level mono-jet selection

The resulting acceptances and efficiencies are listed in Tables I.4 and I.5 for the EFT operators and simplified model respectively.

I.2 Theoretical systematic uncertainties

Working with LO samples for new physics signals, such as DM, comes with large associated theoretical uncertainties. These uncertainties must be studied and quantified, and care must be taken to not overstate the robustness of the approach in comparison to other DM searches which utilize more complete models. As per the standard agreed within the ATLAS SUSY and Exotics groups, the theoretical uncertainties are split into two parts. The part which changes the kinematic acceptance of the sample is included in limit setting, while the part which only scales the cross-section is not. This decision was made because LO signal samples have a questionable cross-section, and it was decided that including

M _{med}	Γ _{med}	m _χ [GeV]	E _T ^{miss} threshold [GeV]								
			150	200	250	300	350	400	500	600	700
10 GeV	M _{med} /3	10	18.60×85.4	7.27×83.9	2.97×80.4	1.22×77.5	0.57×77.6	0.28×76.6	∅	∅	∅
50 GeV	M _{med} /3	10	21.72×84.6	8.28×84.5	3.16×80.9	1.28×85.1	0.60×79.8	0.33×78.8	0.09×74.2	0.03×82.2	0.01×78.2
		50	13.39×84.1	6.25×82.9	2.84×80.1	1.35×79.5	0.71×73.1	0.39×76.6	0.11×77.4	0.04×81.9	0.02×77.3
	M _{med} /8π	400	32.19×82.2	20.96×80.8	13.27×81.3	8.63×80.2	5.78×78.4	3.84×77.1	1.70×76.7	0.78×77.3	0.35×78.5
		50	13.42×85.1	6.09×80.9	2.69×80.1	1.34×82.4	0.64×88.2	0.38×79.0	0.12×72.7	0.04×77.9	0.02×74.0
100 GeV	M _{med} /3	400	32.38×83.0	21.40×80.8	13.64×80.9	8.83×79.8	5.73×80.1	3.80×78.8	1.64×78.8	0.76×77.8	0.36×77.8
		10	33.67×82.5	13.45×82.4	5.62×79.3	2.51×78.5	1.19×78.0	0.65×74.5	0.19×77.5	0.06×79.8	0.02×80.9
	M _{med} /8π	50	11.47×84.7	4.75×82.8	1.97×83.6	0.92×81.2	0.48×78.0	0.25×76.3	0.08×80.4	0.03×79.3	0.01×76.0
		400	32.81×80.8	21.59×79.4	13.74×77.4	8.54×79.5	5.66×78.6	3.73×79.1	1.64×79.3	0.76×82.2	0.37×79.9
300 GeV	M _{med} /3	50	10.12×85.4	4.01×85.0	1.54×80.9	0.67×78.5	0.38×75.2	0.18×74.7	∅	∅	∅
		400	31.87×82.3	20.94×81.6	13.30×81.0	8.63×79.6	5.80×78.9	3.85×78.5	1.72×78.2	0.78×78.0	0.36×80.4
	M _{med} /8π	10	50.01×82.6	25.52×82.1	13.23×80.7	6.63×80.2	3.74×76.4	2.06×78.0	0.72×79.4	0.28×74.9	0.11×76.1
		50	20.81×83.2	10.70×82.8	5.47×80.5	3.00×78.5	1.69×78.5	1.00×73.8	0.32×73.0	0.12×70.9	0.04×77.7
600 GeV	M _{med} /3	200	57.92×82.4	34.67×82.3	19.84×80.7	11.63×80.1	6.79×80.7	4.26×78.2	1.68×78.9	0.66×79.2	0.28×80.5
		400	34.09×83.6	22.35×82.8	14.45×82.0	9.36×79.6	6.11×78.8	3.93×78.3	1.69×76.2	0.78×75.8	0.36×78.2
	M _{med} /8π	50	21.01×83.2	10.86×82.2	5.60×79.2	2.83×80.9	1.59×79.0	0.87×75.3	0.30×75.3	0.10×79.8	0.04×71.4
		400	32.29×81.8	21.31×80.8	13.63×79.8	8.94×78.8	5.69×79.6	3.67×80.2	1.62×79.2	0.77×78.7	0.35×81.2
1 TeV	M _{med} /3	10	57.77×82.0	33.86×81.4	19.34×81.3	11.46×79.0	6.97×80.5	4.28×80.4	1.73×79.0	0.72×80.2	0.32×75.7
		50	27.14×82.0	15.87×80.5	9.44×79.5	5.66×78.0	3.27×78.5	2.01×79.5	0.77×79.4	0.31×77.1	0.14×75.3
	M _{med} /8π	200	60.31×82.2	36.29×81.7	21.35×80.3	12.57×80.3	7.77×80.0	4.96×78.9	1.97×78.4	0.80×76.7	0.35×76.7
		400	33.32×83.5	21.71×81.4	13.78×80.9	8.83×78.7	5.67×79.0	3.72×77.2	1.61×77.3	0.70×78.2	0.33×76.2
3 TeV	M _{med} /3	50	28.62×84.0	17.14×80.6	9.93×81.4	6.04×79.7	3.76×77.8	2.29×77.3	0.87×78.7	0.36×75.3	0.14×76.3
		400	32.07×82.1	20.85×79.9	13.17×80.0	8.39×77.9	5.49×77.5	3.58×78.0	1.55×79.6	0.72×78.2	0.34×79.8
	M _{med} /8π	10	62.88×82.9	39.22×82.2	24.04×81.2	15.17×80.6	9.67×79.9	6.33×80.2	2.73×80.3	1.23×80.1	0.56×82.7
		50	31.27×83.8	19.90×83.3	12.37×80.7	7.60×81.0	4.85×80.8	3.11×80.9	1.32×79.1	0.61×78.8	0.28×79.3
6 TeV	M _{med} /3	200	63.76×82.5	40.91×81.6	25.60×81.0	16.05×79.9	10.15×80.1	6.78×79.4	3.01×78.1	1.33×77.0	0.61×77.3
		400	34.81×83.0	22.68×81.5	14.16×81.2	9.23×79.2	5.91×79.4	3.84×79.2	1.70×79.4	0.78×79.3	0.38×76.8
	M _{med} /8π	1000	71.90×83.1	49.72×81.6	33.52×80.9	22.81×80.0	15.92×78.8	10.95×78.4	5.17×79.2	2.56×78.9	1.26×77.3
		50	33.47×82.6	21.32×79.6	13.23×79.5	8.34×77.6	5.40×77.8	3.54×78.4	1.57×79.5	0.69×76.4	0.32×79.4
10 TeV	M _{med} /3	400	33.70×83.1	21.46×81.4	13.42×81.7	8.61×79.3	5.59×80.1	3.80×78.3	1.60×80.3	0.74×80.3	0.36×79.4
		10	63.56×82.5	40.68×81.0	25.30×80.7	16.35×79.6	10.65×79.8	7.08×78.7	3.18×77.3	1.45×78.7	0.69×78.5
	M _{med} /8π	50	32.01×82.4	20.07×80.0	12.36×79.7	7.91×80.0	5.15×79.1	3.52×77.6	1.54×78.9	0.72×80.4	0.35×78.3
		200	66.04×82.6	42.91×81.8	27.70×79.7	17.66×79.7	11.76×78.9	7.94×78.8	3.68×79.4	1.71×80.9	0.85×80.3
30 TeV	M _{med} /3	400	36.93×84.8	24.73×82.9	15.74×82.6	10.63×79.6	7.20×79.7	4.93×78.3	2.31×79.2	1.16×80.5	0.56×81.6
		1000	72.71×82.8	50.47×81.4	33.68×80.7	22.60×80.1	15.48×79.4	10.66×78.8	5.10×79.3	2.54×78.3	1.23×76.9
	M _{med} /8π	1300	71.54×83.2	44.00×81.6	24.55×77.2	11.58×72.5	3.84×71.4	1.50×70.0	∅	∅	∅
		50	35.73×83.2	23.47×82.3	15.18×81.5	10.19×79.0	6.94×79.7	4.66×79.8	2.21×78.2	1.05×77.2	0.51×78.2
60 TeV	M _{med} /3	400	40.08×83.4	27.23×81.7	17.80×81.0	12.01×81.7	8.18×80.5	5.55×80.8	2.65×78.8	1.32×78.3	0.65×75.6
		50	29.98×83.4	19.03×79.8	11.70×80.1	7.42×78.7	4.73×80.1	3.11×78.1	1.36×77.6	0.60×77.0	0.29×79.1
	M _{med} /8π	400	36.30×83.8	24.46×82.3	16.04×80.7	10.47×80.4	6.89×77.7	4.53×78.9	2.10×81.3	1.00×79.1	0.49×79.2
		50	29.50×82.1	18.55×80.7	11.27×79.5	7.13×78.4	4.46×80.0	2.91×79.4	1.32×77.2	0.61×76.7	0.30×76.1
100 TeV	M _{med} /3	400	34.53×82.0	23.37×81.1	15.21×79.4	10.11×79.7	6.80×79.3	4.50×79.7	2.07×80.7	1.01×79.6	0.49×76.6
		50	30.06×83.9	18.98×81.2	11.38×81.6	7.25×79.1	4.57×80.6	2.94×80.1	1.25×78.2	0.56×77.0	0.27×79.7
	M _{med} /8π	400	36.73×83.2	25.08×80.1	16.35×79.9	10.74×79.4	7.02×78.2	4.74×77.7	2.16×78.6	0.99×78.3	0.46×77.0
		50	28.32×82.6	17.81×81.1	11.01×80.5	7.04×80.1	4.50×79.0	2.86×80.0	1.27×79.7	0.57×78.7	0.28×76.5
300 TeV	M _{med} /3	400	34.52×82.5	23.22×82.3	15.39×80.6	10.06×80.6	6.86×79.0	4.66×78.0	2.07×78.7	0.99×77.1	0.49×77.7
		50	30.18×83.0	18.89×81.2	11.59×78.5	7.29×79.5	4.78×75.6	2.99×78.2	1.27×77.9	0.57×76.9	0.27×78.0
	M _{med} /8π	400	36.74×82.5	24.52×79.1	15.99×79.6	10.57×80.2	7.10×79.5	4.87×77.8	2.11×78.9	1.00×78.4	0.48×77.1
		50	28.29×83.3	17.61×81.0	10.79×80.0	6.85×80.6	4.44×76.5	2.80×77.9	1.21×76.9	0.54×78.5	0.25×77.6
M _{med} /8π	400	34.01×81.6	22.95×80.9	14.98×80.1	10.02×79.2	6.69×78.3	4.51×77.6	2.11×77.4	0.97×76.7	0.49×74.6	

Table I.5: The acceptance \times efficiency, in %, for the simplified model samples. The value \emptyset denotes that the acceptance and efficiency were not evaluated due to insufficient statistics.

this in the limit setting procedure was not a good representation of the true sensitivity to comparable signals.

I.2.1 PDF uncertainties

The use of LO ME samples necessitates the use of LO PDF sets for a consistent treatment of QCD corrections and to not pick up a part of the k -factor associated with moving from LO to NLO calculations. The samples were all generated with `CTEQ6L1` [45], the default PDF set for `MadGraph5` and the set which corresponds to the `Pythia6` tune used. However, `CTEQ6L1` is a very old PDF set, and is known to have problems for events with high momentum transfers in $\sqrt{s} = 8$ TeV collisions. This is exactly the regime which the mono-jet selection emphasizes, and thus it is important to carefully study the PDF uncertainties associated with the DM signal samples.

The procedure established by the `PDF4LHC` working group, as detailed in Reference [148], outlines the proper procedure for estimating PDF uncertainties in such a situation. The report requires the construction of an uncertainty envelope from three NLO PDF error sets, namely `CT10NLO` [47], `MSTW2008NLO` [46], and `NNPDF23NLO` [149]. However, as mentioned, the DM samples are LO and thus use LO PDFs. A modified `PDF4LHC` procedure was followed using `CTEQ6L1`, `MSTW2008LO`, and `NNPDF21LO` [150], where the `NNPDF` variant is an older version because the new release does not contain any LO PDF sets. All PDF sets used come from `LHAPDF`, see Reference [151] for more details.

The `CTEQ6L1` PDF does not include an error set, and thus only the central value could be used. The `MSTW2008LO` and `NNPDF21LO` PDFs do include error sets, but it is important to note that these two PDF sets are conceptually different and care must be taken when constructing an uncertainty band. `MSTW2008LO` contains one central value and 40 error sets, where an uncertainty band is calculated via the asymmetric Hessian procedure, as described in Reference [152]. Equations I.1 and I.2 list the method for calculating the asymmetric Hessian uncertainty, where X_0 is the central value and X_i is the value of the i^{th} error set.

$$\Delta X^+ = \sqrt{\sum_i (X_i - X_0)^2} \quad \forall X_i > X_0, i \neq 0 \quad (\text{I.1})$$

$$\Delta X^- = \sqrt{\sum_i (X_i - X_0)^2} \quad \forall X_i < X_0, i \neq 0 \quad (\text{I.2})$$

`NNPDF21LO` instead contains 99 independent sets and a central value, where the central value is given by the mean of the distribution. The symmetric uncertainty then takes the form of Equation I.3.

$$\Delta X = \sqrt{\frac{1}{99} \sum_i (X_i - X_0)^2} \quad i \neq 0 \quad (\text{I.3})$$

The uncertainty envelope encompassing the different PDF families is formed as a symmetric error of the maximum deviations from the respective central values, per Equation I.4.

$$\Delta X_{\text{env}} = \frac{1}{2} \left[\max \left(X_0^{\text{NNPDF}} + \Delta X, X_0^{\text{MSTW}} + \Delta X^+ \right) - \max \left(X_0^{\text{NNPDF}} - \Delta X, X_0^{\text{MSTW}} - \Delta X^- \right) \right] \quad (\text{I.4})$$

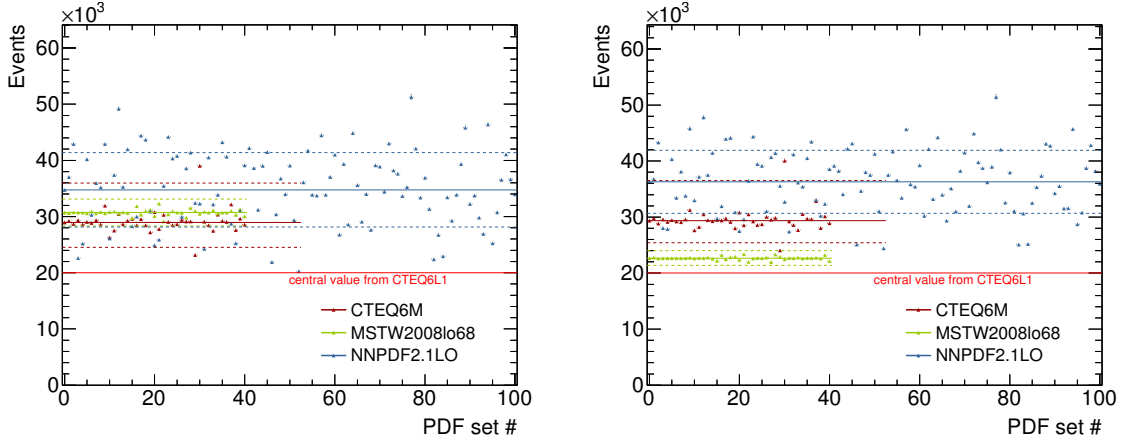


Figure I.4: The number of events after PDF reweighting is applied. The D1 operator with m_χ of 100 GeV (left) and C5 operator with m_χ of 400 GeV (right) are shown. Solid lines are the central values of a given PDF family, while points are the results of error set variations. MSTW2008LO is shown in green, NNPDF21LO in blue, and CTEQ6L1 in red. Note that CTEQ6M is shown as red points for illustration and demonstration that CTEQ6L1 is significantly different, but is not used in the actual PDF uncertainty envelope calculation. The tendency of CTEQ6L1 to be outside of the error bands of all of the other PDF sets led to the decision to not use it, and to instead use the MSTW2008LO central value.

Full generation of a sample with every PDF set is too computationally expensive. Instead, a reweighting procedure is applied on an event-by-event basis, parametrized in the momentum transfer squared (Q_{tr}^2), the momentum fractions (x_1, x_2), and the parton identity. This provides a reasonably accurate means of emulating a sample with a desired PDF set independent of the PDF that was used during the actual generation.

In following this procedure, it was noted that the CTEQ6L1 central value is routinely offset from the NNPDF21LO and MSTW2008LO PDF sets. Furthermore, the addition of CTEQ6M shows that even a PDF set from the same collaboration on a slightly older timescale disagrees with CTEQ6L1. This is shown for the D1 and C5 EFT operators in Figure I.4. Given the aforementioned problems observed by other groups at high momentum transfer when using CTEQ6L1, and the observation that it differs significantly from all of the other sets, it was decided to abandon this baseline. While all DM samples were generated with CTEQ6L1, every time they are used both for systematics and for other results, they are first rescaled to the central value of MSTW2008LO. This decision was taken in conjunction with ATLAS PDF experts who are part of the PDF4LHC working group.

The uncertainties derived using Equation I.4 and used around the MSTW2008LO central value are calculated for both the EFT and Z' simplified model, as listed in Tables I.6 and I.7 respectively. The uncertainties show clear trends. The higher the required momentum fraction, the further the PDF sets are extrapolated, and the larger the associated uncertainties. This is visible both as the DM mass is increased as well as by going to higher signal regions, which in turn require higher $E_{\text{T}}^{\text{miss}}$ and thus harder jets. Additionally, the type of parton involved makes a significant difference. Valence quarks are the best understood, and thus D5 and D9 have the smallest uncertainties. Gluons are next, leaving D11 and C5 with moderate uncertainties. Operators which require access to sea quarks for heavy flavour interactions, notably D1 and C1, are accompanied by very large PDF uncertainties. A comparison of the Z' simplified model at high mediator masses with the corresponding EFT operator shows reasonable

agreement, as expected.

I.2.2 Showering uncertainties

Given that the increased jet multiplicity for DM signals with respect to the dominant SM backgrounds was one of the main motivations for moving to an inclusive mono-jet selection, it is clear that the impact of different showering strengths must be studied. The value of the strong coupling constant, α_s , is important for parton showers and impacts the jet multiplicity in an event. However, α_s and PDFs are inter-related, and it is difficult to study the two independently. This is additionally complicated by our method of generating samples being split between `MadGraph5` and `Pythia6`. In particular, the value of α_s in `MadGraph5` is fixed once you specify a PDF. Some means of factoring the two effects must thus be found.

During the parton shower, the PDF has only minor effects, while the value of α_s becomes very important in controlling the jet multiplicity. As such, for the shower, we can consider α_s to be dominant and assume that the PDF effects are negligible. For ME generation, the PDF controls the parton type(s) produced, while α_s once again influences the jet multiplicity. This is additionally complicated by the value of α_s in `CTEQ6L1` being too large, leading to an artificially increased jet multiplicity after `MadGraph5` generation. It is not generically safe to assume that PDF and α_s effects are decoupled in the ME calculation.

In our particular case, recall that samples are generated at ME level with at least one and optionally a second parton. As such, the fraction of two parton events provides extra information. In particular, the average number of two parton events is approximately 10% when using `CTEQ6L1`. A small change in α_s to a more realistic value should thus lead to suppressed jet production at a level of 1% or less. This cannot actually be tested due to the mentioned limitations of the `MadGraph5` interface, but it seems to be a reasonable approximation, and allows for one to consider the α_s effect to be primarily in the parton shower. Note that the impact of α_s on the cross-section is not as important given the decision to only include theoretical acceptance systematics in the limit setting procedure.

Under this assumption, the parton shower can be varied from $\alpha_s(p_T/2)$ to $\alpha_s(2p_T)$ using `Pythia6` tunes numbered 370, 371, and 372 of the `Perugia2012` family described in Reference [153]. These have the advantage of being tuned with respect to the `CTEQ6L1` PDF set, allowing for holding the PDF constant for both the ME and parton shower while varying the strength of the shower.

For previous iterations of the mono-jet analysis, this would have been a more important systematic. However, the new selection no longer has an explicit cut on the jet multiplicity. This does not remove all of the possible contributions of showering uncertainties, as secondary effects can still enter, but it was sufficient. Studying the showering uncertainties arising from these three tunes results in differences which are consistent with zero within statistical uncertainties, as described in Tables I.8 and I.9 for EFTs and simplified models respectively.

I.2.3 Factorization and renormalization scales

The most obvious limitation of LO generation is the factorization and renormalization scale uncertainty, which is known to reduce substantially when using NLO generators [79]. In `MadGraph5`, the two scales are linked, and can only be coherently varied. Truth samples are produced following the standard procedure of varying the scale up and down by an arbitrary factor of 2. The resulting uncertainties are

Operator	mass [GeV]	$\Delta\sigma$	SR1	SR2	SR3	SR4	SR5	SR6	SR7	SR8	SR9
D1	50	18.98	5.57	9.378	13.36	16.63	19.8	22.68	28.4	31.96	36.09
	100	21.61	5.143	8.525	12.03	15.21	18.15	20.73	25.23	29.68	32.63
	200	28.9	4.2	7.013	10.04	12.17	14.1	15.76	20.07	23.95	26.46
	400	43.11	2.972	4.713	6.582	8.101	9.777	11.15	14	16.83	19.76
	700	60.88	1.855	2.869	4.197	5.62	6.899	8.223	11.81	15.44	20.78
	1000	72.03	2.555	4.193	5.791	7.748	9.178	11.4	14.59	18.97	24.08
	1300	73.43	2.977	4.799	6.894	8.949	10.61	12.91	16.93	21.07	25.56
D5	50	4.919	2.214	3.603	5.142	6.771	8.397	10.01	13.58	17.02	21.45
	100	5.202	2.325	3.814	5.501	7.127	8.757	10.13	13.61	16.9	20.73
	200	6.349	2.58	3.938	5.516	6.853	8.259	9.946	13.25	16.96	19.95
	400	8.921	2.077	3.407	5.155	6.564	8.363	9.831	13.17	16.39	19.99
	700	13.23	2.502	3.898	5.569	6.508	7.83	9.32	11.93	15.59	18.13
	1000	21.47	2.475	3.77	5.541	6.751	8.529	9.619	13.67	16.72	21.32
	1300	35.52	2.5	4.194	6.81	8.793	10.54	12.62	17.34	21.07	28.22
D9	10	4.775	2.12	3.376	4.713	5.874	7.215	8.4	11.21	13.97	16.49
	50	5.122	2.001	3.235	4.589	5.803	7.015	8.246	11.16	13.83	16.33
	100	5.562	2.164	3.479	4.87	6.008	7.395	8.74	11.71	15.01	17.92
	200	6.319	2.429	3.847	5.156	6.624	8.165	9.57	12.24	14.84	16.81
	400	8.825	2.745	4.15	5.837	7.54	9.258	10.91	14.2	16.46	18.99
	700	13.62	2.658	4.173	5.741	7.538	8.924	10.69	13.88	16.55	19.69
	1000	21.83	2.433	3.571	5.466	6.971	8.324	9.345	12.81	16.47	21.05
1300	35.36	2.758	3.906	5.442	7.986	10.06	11.95	15.87	21.7	29.13	
D11	50	37.77	0.9986	1.914	2.818	3.418	4.385	5.004	6.243	7.7	8.773
	100	39.86	0.6845	1.192	1.843	2.496	3.176	3.733	5.227	6.21	7.313
	200	44.79	0.7062	0.8821	1.22	1.521	1.787	2.107	2.952	3.887	4.515
	400	54.04	1.389	1.634	1.897	2.142	2.297	2.294	2.571	3.118	4.061
	700	64.45	2.159	2.59	3.149	3.524	3.934	3.963	4.331	5.058	5.891
	1000	68.44	2.628	2.941	3.551	3.951	4.104	4.346	4.522	5.543	6.404
	1300	69.34	2.778	3.297	3.835	4.417	4.448	4.544	5.187	6.133	6.562
C1	10	8.947	6.466	10.67	16.18	19.07	23.89	27.46	34.49	42.42	48.31
	50	10.77	5.686	9.906	13.86	17.82	21.74	24.78	31.24	40.12	44.28
	100	13.42	4.955	8.597	12.55	16.1	19.12	21.93	29.22	37.24	41.79
	200	19.56	4.059	6.885	9.907	12.5	15.22	18.09	23.72	28.05	32.29
	400	34.25	3.187	5.058	7.425	9.787	11.47	13.5	16.92	20.05	22.18
	700	53.7	2.063	3.195	4.552	5.872	7.229	8.355	11.07	13.43	16.64
	1000	69.58	2.154	3.354	5.016	6.843	8.351	9.908	13.49	17.6	22.71
1300	74.93	3.214	4.923	6.668	8.655	10.55	12.5	16.58	20.64	25.23	
C5	10	19.9	3.108	5.063	6.959	8.54	10.18	11.68	13.78	16.41	17.35
	50	22.52	2.371	4.013	5.451	6.976	8.002	9.405	11.79	13.83	16.21
	100	26.27	1.377	2.516	3.668	5.053	5.963	7.398	9.481	11.42	13.44
	200	33.22	0.7493	1.361	2.144	2.733	3.517	4.4	6.257	7.891	9.182
	400	45.76	0.8561	0.9606	1.14	1.401	1.68	1.954	2.593	3.374	4.087
	700	59.72	1.721	2.041	2.505	2.776	3.205	3.241	3.438	4.113	4.911
	1000	67.68	2.471	2.935	3.602	3.968	4.018	4.173	4.727	5.507	6.589
1300	69.64	2.716	3.22	3.884	4.268	4.652	4.758	5.344	6.551	7.529	

Table I.6: PDF uncertainties on the cross-section σ and the acceptance for each signal region (SR) for the EFT samples. Uncertainties increase with the DM mass, as this increases the required momentum transfer and thus the momentum fraction. Uncertainties also increase with higher signal regions due to higher E_T^{miss} cuts, which indirectly require harder jets to balance the E_T^{miss} , thus also increasing the required momentum fraction. Valence quark operators are the best understood (D5 and D9), followed by gluon operators (D11 and C5), while sea quark operators (D1 and C1) have the largest PDF uncertainties.

M_{med}	Γ_{med}	m_χ [GeV]	$\Delta\sigma$	SR1	SR2	SR3	SR4	SR5	SR6	SR7	SR8	SR9
10 GeV	$M_{\text{med}}/3$	10	3.92	0.54	1.63	3.11	3.27	4.9	6.22	7.88	12.32	16.28
50 GeV	$M_{\text{med}}/3$	10	3.92	0.46	1.4	2.82	3.8	5.76	7.1	9.55	12.89	14.87
		50	2.73	2.12	3.41	4.88	6.42	8.38	9.21	11.86	14.51	16.21
	$M_{\text{med}}/8\pi$	400	6.64	1.94	3.34	5.13	6.8	7.91	9.62	13.08	17.63	21.09
		50	2.64	2.41	3.91	5.27	6.47	7.58	7.78	11.64	16.0	20.74
100 GeV	$M_{\text{med}}/3$	400	6.63	2.2	3.58	5.52	7.13	8.91	10.72	12.97	16.16	19.78
		10	4.14	0.37	1.51	2.38	3.8	5.74	7.0	8.5	12.08	14.02
	$M_{\text{med}}/8\pi$	50	2.5	2.33	3.41	4.57	5.56	6.98	9.36	12.42	15.74	20.67
		400	6.56	2.01	3.29	4.67	6.16	7.68	9.27	12.18	16.0	19.76
300 GeV	$M_{\text{med}}/3$	50	2.55	2.23	2.99	3.38	5.07	6.16	9.12	11.97	14.27	17.36
		400	6.55	2.2	3.47	4.97	6.41	8.13	9.52	13.0	16.83	19.59
	$M_{\text{med}}/8\pi$	10	4.5	0.35	1.35	2.66	3.89	5.08	6.49	9.58	12.64	15.87
		50	3.16	1.67	2.83	4.32	6.16	7.32	8.22	12.4	14.56	16.54
600 GeV	$M_{\text{med}}/3$	200	4.74	0.35	1.01	2.29	3.65	5.25	6.74	9.86	12.77	15.68
		400	6.54	2.11	3.46	4.74	6.08	7.6	9.34	12.17	15.57	19.13
	$M_{\text{med}}/8\pi$	50	3.04	1.92	3.33	4.53	6.01	6.43	7.05	9.46	12.9	16.12
		400	6.57	2.03	3.27	5.0	6.37	7.67	9.77	13.21	16.87	19.73
1 TeV	$M_{\text{med}}/3$	10	4.68	0.44	1.22	2.69	3.91	5.73	7.09	10.5	13.59	16.73
		50	3.18	1.73	3.15	4.3	5.51	7.05	8.19	11.32	14.53	17.62
	$M_{\text{med}}/8\pi$	200	4.74	0.18	1.32	2.69	4.4	5.92	7.36	10.22	13.33	16.43
		400	6.19	1.89	3.14	4.57	6.37	7.99	9.69	14.07	19.71	25.54
3 TeV	$M_{\text{med}}/3$	50	3.43	2.21	3.73	5.17	6.67	8.09	9.64	11.9	14.75	17.85
		400	6.13	1.91	3.5	5.04	6.28	7.67	9.69	13.04	15.99	18.3
	$M_{\text{med}}/8\pi$	10	5.6	0.31	1.32	2.82	4.42	5.46	6.9	10.01	13.3	15.88
		50	5.13	1.69	3.05	4.51	6.02	7.0	9.07	12.06	15.47	18.96
6 TeV	$M_{\text{med}}/3$	200	6.13	0.29	1.24	2.66	3.96	5.45	6.78	10.07	13.46	16.56
		400	6.67	1.81	3.22	4.72	6.31	7.65	9.15	12.51	15.74	18.83
	$M_{\text{med}}/8\pi$	1000	19.65	0.32	1.16	2.4	4.04	5.16	6.19	9.43	13.59	18.07
		50	6.11	1.86	3.29	4.74	6.31	8.01	9.44	13.16	16.68	17.98
10 TeV	$M_{\text{med}}/3$	400	6.31	1.99	3.6	5.2	6.2	8.16	9.69	13.31	17.25	20.71
		10	8.81	0.56	1.95	3.43	4.92	6.91	8.46	12.23	17.38	21.5
	$M_{\text{med}}/8\pi$	50	6.83	2.58	4.53	6.88	9.23	10.14	11.62	14.9	18.55	22.98
		200	9.81	0.43	1.74	3.44	5.52	7.11	8.76	12.04	15.9	20.74
30 TeV	$M_{\text{med}}/3$	400	11.3	2.21	3.48	5.66	6.76	8.58	10.31	12.33	16.1	19.18
		1000	25.28	0.32	1.72	3.38	5.31	6.88	8.74	10.64	15.73	19.82
	$M_{\text{med}}/8\pi$	1300	34.21	0.64	1.45	3.16	4.8	6.62	8.41	12.84	17.56	20.9
		50	16.57	3.67	5.47	6.84	8.45	10.74	12.95	15.57	21.28	25.9
10 TeV	$M_{\text{med}}/3$	400	21.95	2.29	3.46	5.23	6.55	8.71	10.86	12.8	16.33	20.33
		50	5.17	2.33	4.22	5.89	7.64	9.41	11.12	15.05	17.65	19.78
	$M_{\text{med}}/8\pi$	400	9.39	2.13	3.26	5.06	6.61	7.68	10.26	13.59	15.78	19.62
		50	5.38	2.8	4.74	5.68	7.17	8.95	10.54	14.09	18.65	21.44
30 TeV	$M_{\text{med}}/3$	400	9.38	2.48	3.79	5.33	6.67	8.26	10.12	14.15	18.35	22.29
		50	5.07	2.2	3.79	5.21	6.32	8.33	9.98	13.56	16.57	20.19
	$M_{\text{med}}/8\pi$	400	9.17	1.81	3.3	4.75	6.21	8.37	9.76	12.16	16.43	18.52
		50	4.93	2.61	4.06	5.62	6.97	8.7	10.59	13.87	18.17	22.69
30 TeV	$M_{\text{med}}/3$	400	9.16	2.31	3.77	4.85	6.5	8.73	9.73	13.05	16.06	20.68
		50	4.98	2.28	3.67	5.25	6.66	7.8	9.53	13.16	16.49	20.77
	$M_{\text{med}}/8\pi$	400	9.22	2.12	3.36	4.82	6.62	7.34	9.75	13.13	18.47	22.24
		50	5.0	2.41	3.71	5.55	7.12	8.86	10.03	13.86	18.64	23.28
30 TeV	$M_{\text{med}}/8\pi$	400	8.97	2.21	3.6	5.13	6.99	8.95	11.51	13.84	18.25	20.72

Table I.7: PDF uncertainties on the cross-section σ and the acceptance for each signal region (SR) for the Z' simplified model samples. Due to the resonance structure, the picture is more complicated than the EFT. However, in general, uncertainties increase with heavier DM particles and higher E_T^{miss} cuts (from higher signal regions), both of which increase the required momentum fraction. Comparing very heavy mediators to the uncertainties observed for the D5 EFT provides comparable results.

Operator	m_χ [GeV]	$\Delta\sigma$ [%]	SR1 [%]	SR2 [%]	SR3 [%]	SR4 [%]	SR5 [%]	SR6 [%]	SR7 [%]	SR8 [%]	SR9 [%]
C1	10	0.21 (0.72)	1.55 (2.28)	1.49 (3.51)	1.48 (5.13)	2.36 (6.47)	8.91 (6.83)	8.59 (7.37)	1.58 (8.97)	0.97 (5.73)	2.46 (10.07)
	50	0.23 (0.72)	1.23 (1.95)	0.92 (2.92)	3.00 (4.26)	5.46 (5.46)	3.28 (5.69)	7.85 (6.52)	2.02 (8.60)	4.91 (5.68)	8.56 (9.00)
	100	0.18 (0.72)	0.69 (1.75)	2.04 (2.46)	1.80 (3.38)	2.04 (4.17)	2.31 (4.35)	2.22 (4.41)	5.27 (5.93)	6.82 (9.64)	1.69 (8.69)
	200	1.04 (0.73)	0.45 (1.50)	0.22 (2.00)	2.21 (2.63)	0.43 (3.09)	2.25 (3.17)	0.86 (3.37)	1.00 (4.52)	2.86 (5.77)	3.66 (6.77)
	400	0.07 (0.72)	1.20 (1.33)	1.42 (1.65)	1.98 (2.03)	1.08 (2.28)	2.03 (2.26)	1.97 (2.35)	2.29 (2.99)	2.43 (4.23)	3.72 (5.50)
	700	0.54 (0.73)	1.26 (1.23)	1.49 (1.46)	0.20 (1.72)	2.00 (1.91)	0.59 (1.87)	0.88 (1.92)	0.81 (2.45)	2.23 (3.24)	1.76 (4.65)
	1000	0.50 (0.75)	0.83 (1.24)	0.61 (1.44)	1.08 (1.65)	2.34 (1.78)	1.94 (1.76)	0.67 (1.85)	0.71 (2.27)	3.05 (3.02)	3.49 (4.13)
1300	0.31 (0.77)	0.30 (1.26)	0.21 (1.45)	0.53 (1.67)	0.93 (1.79)	1.38 (1.72)	2.46 (1.78)	0.97 (2.22)	2.09 (3.05)	3.10 (4.10)	
C5	10	0.09 (0.70)	0.86 (1.41)	0.69 (1.69)	0.49 (1.99)	1.36 (2.18)	1.20 (2.11)	0.65 (2.18)	1.89 (2.51)	1.60 (2.87)	1.85 (3.65)
	50	0.59 (0.70)	0.41 (1.32)	1.12 (1.55)	2.10 (1.79)	1.58 (1.95)	2.16 (1.93)	2.48 (1.97)	2.50 (2.20)	2.47 (2.60)	3.06 (3.50)
	100	0.28 (0.69)	0.25 (1.25)	0.39 (1.44)	0.63 (1.64)	0.21 (1.77)	0.31 (1.75)	1.57 (1.78)	0.46 (2.07)	0.26 (2.65)	2.30 (3.44)
	200	0.34 (0.71)	0.37 (1.20)	0.28 (1.34)	0.29 (1.49)	1.01 (1.57)	0.43 (1.58)	0.85 (1.61)	0.41 (1.85)	1.00 (2.29)	0.50 (3.02)
	400	0.33 (0.74)	1.12 (1.18)	0.22 (1.27)	0.72 (1.37)	0.64 (1.42)	1.00 (1.41)	1.02 (1.43)	0.96 (1.64)	0.95 (2.00)	0.78 (2.59)
	700	0.47 (0.76)	0.56 (1.16)	0.39 (1.23)	0.88 (1.31)	1.41 (1.36)	1.12 (1.34)	0.41 (1.37)	0.16 (1.53)	1.27 (1.83)	2.25 (2.31)
	1000	0.31 (0.78)	0.54 (1.15)	1.04 (1.23)	1.31 (1.30)	0.86 (1.34)	0.36 (1.32)	0.13 (1.34)	0.78 (1.48)	0.12 (1.75)	0.46 (2.21)
1300	0.58 (0.79)	0.45 (1.16)	0.27 (1.23)	1.33 (1.31)	0.84 (1.36)	0.73 (1.36)	1.51 (1.41)	0.29 (1.62)	0.98 (1.94)	0.83 (2.60)	
D1	10	0.16 (0.73)	0.34 (1.66)	1.48 (2.27)	1.82 (2.98)	0.84 (3.56)	0.80 (3.74)	1.16 (4.21)	4.35 (5.44)	5.52 (4.81)	5.58 (7.64)
	50	0.31 (0.73)	0.28 (1.60)	1.16 (2.15)	2.59 (2.84)	1.34 (3.31)	1.67 (3.33)	2.33 (3.45)	3.89 (3.35)	3.48 (4.26)	5.23 (6.96)
	100	0.24 (0.73)	0.88 (1.52)	1.37 (1.98)	1.29 (2.55)	0.65 (3.05)	2.82 (3.07)	1.69 (3.31)	0.94 (3.87)	7.51 (5.05)	1.40 (6.59)
	200	0.44 (0.73)	0.59 (1.41)	1.00 (1.77)	0.66 (2.24)	1.56 (2.55)	1.69 (2.61)	2.96 (2.70)	0.71 (3.49)	0.78 (5.60)	1.52 (6.01)
	400	0.09 (0.73)	0.67 (1.29)	0.79 (1.56)	1.80 (1.85)	2.26 (2.08)	1.85 (2.09)	1.76 (2.18)	1.21 (2.93)	1.45 (3.65)	2.34 (5.37)
	700	0.19 (0.74)	0.92 (1.24)	1.47 (1.46)	1.50 (1.69)	1.70 (1.86)	1.50 (1.83)	1.31 (1.93)	3.23 (2.30)	0.52 (2.97)	0.50 (4.39)
	1000	0.09 (0.76)	1.02 (1.25)	1.12 (1.44)	0.17 (1.67)	0.62 (1.79)	0.10 (1.75)	1.00 (1.83)	1.48 (2.17)	2.80 (2.97)	1.33 (4.07)
1300	0.18 (0.79)	0.31 (1.29)	0.33 (1.49)	0.93 (1.73)	1.38 (1.91)	3.35 (1.91)	2.61 (2.09)	1.70 (2.43)	0.38 (3.02)	1.78 (4.51)	
D5	10	0.32 (0.70)	0.76 (1.26)	0.50 (1.54)	0.98 (1.81)	0.15 (1.98)	0.64 (1.87)	1.16 (1.87)	2.62 (2.11)	0.42 (2.70)	3.12 (3.76)
	50	0.53 (0.70)	0.91 (1.24)	0.90 (1.51)	1.29 (1.78)	1.77 (1.94)	0.48 (1.87)	0.50 (1.91)	3.44 (2.28)	1.67 (2.81)	0.95 (3.82)
	100	0.16 (0.70)	0.44 (1.23)	0.14 (1.49)	1.21 (1.78)	0.71 (1.93)	0.48 (1.84)	0.36 (1.92)	1.50 (2.47)	2.48 (3.01)	1.65 (3.80)
	200	0.30 (0.70)	0.54 (1.19)	0.82 (1.42)	0.58 (1.65)	1.22 (1.78)	0.47 (1.72)	0.82 (1.78)	1.25 (2.00)	2.56 (2.57)	0.75 (3.61)
	400	0.33 (0.69)	0.44 (1.13)	0.54 (1.33)	0.98 (1.52)	1.80 (1.63)	1.83 (1.56)	1.01 (1.61)	1.65 (2.00)	0.80 (2.52)	2.26 (3.39)
	700	0.26 (0.69)	0.66 (1.11)	0.66 (1.28)	0.30 (1.45)	0.22 (1.54)	0.48 (1.46)	0.67 (1.50)	1.44 (1.82)	1.19 (2.29)	2.97 (3.13)
	1000	0.51 (0.69)	0.35 (1.09)	0.32 (1.25)	0.01 (1.42)	0.69 (1.51)	0.36 (1.44)	1.98 (1.50)	0.88 (1.81)	0.23 (2.26)	1.27 (3.01)
1300	0.22 (0.69)	0.07 (1.09)	0.33 (1.26)	0.69 (1.42)	0.53 (1.51)	0.85 (1.46)	1.34 (1.53)	1.14 (1.85)	1.97 (2.31)	2.65 (3.05)	
D9	10	0.19 (0.69)	0.32 (1.11)	0.31 (1.28)	0.24 (1.43)	0.38 (1.50)	0.41 (1.42)	0.16 (1.46)	0.49 (1.73)	0.78 (2.14)	1.78 (2.83)
	50	0.21 (0.69)	0.31 (1.08)	0.28 (1.24)	0.11 (1.39)	0.58 (1.46)	0.19 (1.40)	0.18 (1.45)	0.75 (1.74)	1.27 (2.18)	1.36 (2.85)
	100	0.09 (0.68)	0.83 (1.06)	0.42 (1.21)	0.08 (1.34)	0.14 (1.41)	0.24 (1.35)	0.18 (1.40)	0.88 (1.69)	0.92 (2.12)	0.63 (2.83)
	200	0.40 (0.68)	0.60 (1.05)	0.41 (1.19)	0.39 (1.31)	0.83 (1.36)	0.29 (1.30)	0.85 (1.34)	1.54 (1.61)	0.34 (2.03)	3.18 (2.67)
	400	0.26 (0.68)	0.80 (1.06)	0.87 (1.19)	0.45 (1.31)	0.86 (1.36)	0.67 (1.31)	0.14 (1.35)	0.44 (1.58)	1.23 (1.95)	3.15 (2.55)
	700	0.38 (0.69)	0.37 (1.06)	0.42 (1.21)	0.28 (1.33)	0.87 (1.38)	0.29 (1.31)	0.73 (1.34)	0.21 (1.59)	0.51 (1.93)	1.10 (2.47)
	1000	0.50 (0.69)	0.89 (1.06)	0.70 (1.21)	0.52 (1.35)	0.90 (1.40)	1.71 (1.34)	0.81 (1.40)	0.53 (1.63)	0.19 (2.01)	0.99 (2.62)
1300	0.27 (0.69)	0.69 (1.07)	0.37 (1.22)	0.36 (1.36)	1.02 (1.43)	0.71 (1.36)	0.40 (1.41)	1.19 (1.67)	0.57 (2.08)	0.43 (2.74)	
D11	10	0.30 (0.71)	0.63 (1.19)	0.36 (1.33)	0.27 (1.46)	1.42 (1.52)	0.66 (1.49)	0.87 (1.52)	0.92 (1.75)	0.96 (2.11)	0.77 (2.80)
	50	0.67 (0.72)	1.13 (1.20)	0.97 (1.33)	0.99 (1.46)	1.57 (1.52)	1.35 (1.52)	1.61 (1.56)	1.11 (1.75)	3.00 (2.14)	1.94 (2.76)
	100	0.28 (0.73)	0.63 (1.18)	0.70 (1.29)	0.49 (1.42)	1.13 (1.48)	0.78 (1.46)	1.10 (1.48)	1.17 (1.69)	2.39 (2.05)	0.84 (2.67)
	200	0.06 (0.73)	0.44 (1.17)	1.42 (1.28)	1.18 (1.39)	0.61 (1.44)	1.41 (1.43)	1.20 (1.44)	1.01 (1.63)	1.25 (1.98)	0.78 (2.58)
	400	0.47 (0.75)	0.41 (1.16)	1.29 (1.24)	1.09 (1.33)	0.91 (1.37)	1.29 (1.35)	1.35 (1.37)	1.17 (1.54)	1.67 (1.86)	0.83 (2.39)
	700	1.65 (0.78)	0.76 (1.16)	1.06 (1.23)	1.65 (1.30)	1.13 (1.34)	0.76 (1.32)	1.03 (1.33)	1.80 (1.48)	1.74 (1.78)	1.93 (2.22)
	1000	0.24 (0.78)	1.20 (1.15)	1.32 (1.21)	1.44 (1.28)	0.83 (1.33)	1.13 (1.32)	1.41 (1.34)	0.35 (1.49)	1.24 (1.78)	1.66 (2.24)
1300	0.45 (0.80)	0.51 (1.16)	1.11 (1.22)	0.82 (1.31)	0.64 (1.36)	0.38 (1.35)	0.28 (1.39)	0.89 (1.55)	2.18 (1.82)	1.37 (2.35)	

Table I.8: Showering systematics in % on the cross-section (σ) and acceptance by signal region (SR) for the EFT operators. Numbers in parentheses are the corresponding statistical uncertainties. The cross-section uncertainty and the acceptance uncertainty for the signal region used for limit setting are bolded. As this is a LO signal model, the cross-section uncertainty is not used for limit setting. The acceptance uncertainty is seen to be consistent with 0 within statistics, and thus no extra showering uncertainty is used.

M_{med}	Γ_{med}	m_{χ} [GeV]	$\Delta\sigma$ [%]	SR1 [%]	SR2 [%]	SR3 [%]	SR4 [%]	SR5 [%]	SR6 [%]	SR7 [%]	SR8 [%]	SR9 [%]
10 GeV	$M_{\text{med}}/3$	10	0.06 (0.45)	1.10 (1.88)	1.83 (2.97)	0.56 (4.48)	1.41 (5.95)	1.75 (6.60)	1.31 (7.03)	2.07 (10.03)	7.98 (9.95)	6.09 (7.48)
		50	0.08 (0.45)	1.23 (1.15)	0.41 (1.69)	0.84 (2.35)	0.77 (2.93)	1.20 (2.95)	1.05 (3.06)	0.94 (3.45)	1.30 (3.47)	1.37 (4.54)
		100	0.07 (0.45)	0.80 (0.94)	0.56 (1.27)	0.67 (1.64)	0.62 (1.95)	0.58 (1.93)	1.43 (1.97)	1.87 (2.39)	4.14 (3.06)	1.43 (3.85)
		200	0.18 (0.44)	0.94 (0.81)	0.60 (1.01)	0.22 (1.23)	0.12 (1.39)	0.46 (1.33)	0.61 (1.37)	2.62 (1.63)	0.80 (2.08)	1.31 (3.01)
		400	0.11 (0.44)	0.29 (0.74)	0.18 (0.88)	0.82 (1.02)	1.34 (1.11)	0.91 (1.05)	0.96 (1.09)	1.72 (1.33)	2.15 (1.71)	0.97 (2.31)
		700	0.05 (0.44)	0.51 (0.70)	0.47 (0.81)	0.81 (0.92)	0.36 (0.98)	0.20 (0.94)	0.22 (0.97)	1.03 (1.19)	0.48 (1.50)	0.23 (2.00)
		1300	0.08 (0.44)	0.33 (0.69)	0.27 (0.79)	0.22 (0.90)	0.38 (0.95)	0.39 (0.91)	0.25 (0.95)	0.37 (1.14)	1.69 (1.40)	1.24 (1.91)
50 GeV	$M_{\text{med}}/3$	10	0.04 (0.45)	1.81 (1.78)	1.44 (2.86)	1.12 (4.36)	6.47 (5.71)	1.29 (6.22)	0.92 (6.20)	2.92 (4.98)	0.80 (4.26)	2.82 (7.06)
		50	0.08 (0.45)	0.94 (1.16)	1.04 (1.69)	1.26 (2.37)	2.04 (2.96)	2.27 (3.17)	2.88 (3.55)	1.04 (4.29)	4.73 (4.54)	2.29 (6.53)
		100	0.09 (0.45)	0.55 (0.94)	0.74 (1.26)	0.22 (1.66)	0.96 (1.97)	0.65 (1.92)	1.08 (2.05)	1.25 (2.35)	0.46 (2.62)	0.34 (3.78)
		200	0.07 (0.44)	0.24 (0.82)	0.26 (1.01)	0.54 (1.23)	0.27 (1.38)	0.79 (1.35)	1.18 (1.41)	0.81 (1.66)	0.30 (2.09)	0.08 (3.02)
		400	0.10 (0.44)	0.30 (0.74)	0.24 (0.88)	0.36 (1.02)	0.50 (1.10)	0.86 (1.06)	0.13 (1.10)	0.27 (1.33)	0.73 (1.69)	0.45 (2.30)
		700	0.05 (0.44)	0.47 (0.70)	0.56 (0.82)	0.27 (0.93)	0.62 (0.99)	0.59 (0.94)	0.81 (0.97)	0.39 (1.18)	0.33 (1.50)	1.22 (2.02)
		1300	0.08 (0.44)	0.21 (0.69)	0.13 (0.79)	0.34 (0.90)	0.05 (0.95)	0.17 (0.91)	0.39 (0.95)	0.20 (1.15)	1.31 (1.43)	1.66 (1.88)
100 GeV	$M_{\text{med}}/8\pi$	50	0.13 (0.45)	1.41 (1.17)	1.25 (1.72)	0.90 (2.41)	1.42 (3.06)	1.93 (3.12)	0.32 (3.10)	1.59 (3.77)	3.59 (3.94)	0.24 (4.60)
		10	0.23 (0.45)	0.46 (1.38)	0.75 (2.17)	2.50 (3.16)	3.89 (4.08)	4.49 (4.19)	0.25 (4.19)	1.47 (3.85)	1.98 (3.53)	1.24 (5.77)
		50	0.03 (0.45)	0.41 (1.26)	1.22 (1.93)	1.71 (2.74)	1.76 (3.49)	1.89 (3.53)	2.81 (3.48)	2.11 (3.66)	2.99 (6.16)	8.15 (5.08)
		200	0.12 (0.44)	0.45 (0.81)	0.25 (1.01)	0.57 (1.24)	1.19 (1.40)	0.71 (1.34)	0.35 (1.39)	0.14 (1.66)	1.62 (2.06)	3.77 (3.09)
		400	0.25 (0.44)	0.22 (0.74)	0.33 (0.88)	1.13 (1.02)	0.81 (1.10)	1.03 (1.06)	0.52 (1.09)	0.23 (1.33)	0.50 (1.71)	1.02 (2.33)
		700	0.14 (0.44)	0.24 (0.71)	0.21 (0.82)	0.59 (0.93)	0.14 (0.99)	0.26 (0.94)	0.37 (0.98)	0.50 (1.20)	0.29 (1.53)	0.05 (2.00)
		1300	0.21 (0.44)	0.52 (0.69)	0.77 (0.79)	0.64 (0.90)	0.54 (0.95)	0.56 (0.91)	0.13 (0.95)	0.74 (1.13)	1.64 (1.42)	0.96 (1.92)
300 GeV	$M_{\text{med}}/3$	10	0.29 (0.45)	0.48 (0.95)	0.54 (1.28)	1.13 (1.68)	1.65 (2.01)	1.40 (2.01)	0.93 (2.03)	1.49 (2.47)	2.03 (3.47)	2.38 (4.10)
		50	0.17 (0.45)	0.60 (0.95)	0.20 (1.28)	1.18 (1.69)	0.96 (2.02)	0.85 (2.01)	0.96 (2.10)	2.67 (2.48)	1.90 (3.02)	1.66 (3.98)
		200	0.01 (0.44)	0.12 (0.82)	0.48 (1.03)	0.60 (1.26)	0.39 (1.44)	1.26 (1.45)	1.63 (1.56)	2.88 (1.86)	1.37 (2.25)	1.25 (3.07)
		400	0.09 (0.44)	0.38 (0.73)	0.40 (0.87)	0.44 (1.02)	0.83 (1.12)	0.79 (1.10)	0.70 (1.18)	1.72 (1.42)	0.50 (1.72)	0.81 (2.27)
		700	0.22 (0.44)	0.18 (0.70)	0.34 (0.81)	0.39 (0.93)	0.18 (0.99)	0.26 (0.95)	0.20 (0.98)	1.91 (1.19)	1.24 (1.49)	0.69 (2.01)
		1000	0.10 (0.44)	0.11 (0.69)	0.29 (0.79)	0.33 (0.89)	0.37 (0.95)	0.45 (0.91)	0.23 (0.95)	0.38 (1.16)	1.20 (1.43)	0.35 (1.91)
		1300	0.13 (0.44)	0.20 (0.69)	0.28 (0.79)	0.53 (0.90)	0.54 (0.96)	0.35 (0.92)	0.32 (0.95)	0.19 (1.12)	0.52 (1.43)	1.11 (1.95)
600 GeV	$M_{\text{med}}/8\pi$	50	0.13 (0.45)	0.42 (0.94)	0.10 (1.28)	1.10 (1.69)	1.33 (2.06)	0.64 (2.04)	1.72 (2.13)	3.30 (2.70)	1.79 (3.63)	1.07 (4.23)
		10	0.16 (0.44)	0.13 (0.82)	0.70 (1.03)	1.87 (1.27)	1.00 (1.42)	1.38 (1.35)	0.32 (1.39)	1.81 (1.64)	0.18 (2.10)	1.60 (2.96)
		50	0.09 (0.44)	0.12 (0.82)	0.64 (1.02)	0.90 (1.25)	0.04 (1.40)	0.55 (1.36)	1.09 (1.38)	0.71 (1.66)	0.49 (2.08)	0.91 (2.97)
		200	0.05 (0.44)	0.21 (0.79)	0.56 (0.99)	0.18 (1.19)	0.31 (1.33)	0.55 (1.28)	0.67 (1.31)	1.16 (1.64)	3.43 (1.98)	3.89 (2.86)
		400	0.16 (0.44)	0.09 (0.73)	0.18 (0.88)	0.37 (1.02)	0.51 (1.13)	0.47 (1.12)	1.10 (1.17)	0.40 (1.42)	0.58 (1.77)	2.58 (2.41)
		700	0.15 (0.36)	0.08 (0.57)	0.04 (0.66)	0.53 (0.76)	0.43 (0.83)	1.46 (0.85)	1.39 (0.92)	1.63 (1.17)	0.81 (1.47)	1.16 (1.97)
		1300	0.46 (0.44)	0.18 (0.68)	0.49 (0.79)	0.19 (0.90)	0.51 (0.96)	0.87 (0.91)	0.28 (0.94)	0.58 (1.13)	0.55 (1.44)	0.88 (1.94)
1 TeV	$M_{\text{med}}/3$	10	0.11 (0.44)	0.15 (0.68)	0.18 (0.79)	0.18 (0.89)	0.69 (0.95)	0.30 (0.90)	1.06 (0.92)	0.39 (1.12)	0.49 (1.42)	1.05 (1.91)
		50	0.11 (0.44)	0.01 (0.80)	0.48 (1.00)	0.54 (1.21)	1.06 (1.37)	0.68 (1.32)	1.12 (1.38)	1.47 (1.68)	1.93 (2.09)	1.09 (2.96)
		50	0.09 (0.44)	0.36 (0.74)	0.20 (0.88)	0.22 (1.03)	0.66 (1.12)	0.45 (1.08)	0.87 (1.12)	1.06 (1.38)	0.84 (1.72)	0.88 (2.36)
		10	0.09 (0.44)	0.30 (0.76)	0.46 (0.92)	0.64 (1.09)	1.25 (1.19)	1.21 (1.13)	0.33 (1.16)	0.38 (1.43)	0.43 (1.78)	1.39 (2.45)
		50	0.05 (0.44)	0.42 (0.76)	0.27 (0.92)	0.39 (1.09)	0.04 (1.19)	0.54 (1.14)	0.79 (1.16)	0.93 (1.40)	2.19 (1.73)	2.38 (2.43)
		200	0.24 (0.44)	0.25 (0.74)	0.23 (0.89)	0.56 (1.05)	0.54 (1.15)	1.19 (1.09)	1.03 (1.12)	0.69 (1.36)	0.69 (1.73)	1.77 (2.43)
		400	0.17 (0.44)	0.44 (0.72)	0.07 (0.87)	0.34 (1.01)	0.58 (1.09)	0.49 (1.05)	0.29 (1.09)	0.36 (1.31)	0.43 (1.65)	1.34 (2.43)
3 TeV	$M_{\text{med}}/8\pi$	50	0.13 (0.45)	0.19 (0.68)	0.35 (0.78)	0.39 (0.89)	0.36 (0.95)	0.36 (0.95)	0.22 (0.98)	0.94 (1.18)	0.89 (1.44)	0.42 (1.90)
		50	0.38 (0.44)	1.13 (0.73)	1.06 (0.88)	1.20 (1.02)	1.08 (1.12)	2.35 (1.07)	1.86 (1.11)	1.39 (1.33)	0.17 (1.69)	0.24 (2.35)
		400	0.11 (0.44)	0.28 (0.73)	0.30 (0.87)	0.48 (1.02)	0.67 (1.11)	0.67 (1.06)	1.23 (1.09)	0.32 (1.36)	0.37 (1.68)	1.86 (2.33)
		10	0.05 (0.44)	0.23 (0.76)	0.29 (0.92)	1.15 (1.09)	0.28 (1.18)	0.24 (1.11)	0.27 (1.13)	0.53 (1.38)	0.93 (1.65)	0.82 (2.22)
		50	0.08 (0.44)	0.28 (0.76)	0.36 (0.92)	0.34 (1.07)	0.81 (1.16)	0.58 (1.11)	0.39 (1.13)	0.61 (1.30)	0.61 (1.61)	0.57 (2.23)
		200	0.23 (0.44)	0.26 (0.72)	0.43 (0.86)	0.63 (1.01)	0.65 (1.08)	0.86 (1.03)	0.79 (1.06)	0.86 (1.29)	1.24 (1.59)	0.96 (2.16)
		400	0.10 (0.44)	0.28 (0.70)	0.67 (0.82)	0.34 (0.94)	0.26 (1.01)	0.11 (0.96)	1.40 (0.99)	1.34 (1.21)	0.73 (1.50)	2.17 (2.02)
6 TeV	$M_{\text{med}}/3$	1000	0.11 (0.44)	0.36 (0.68)	0.31 (0.79)	0.63 (0.89)	0.44 (0.95)	0.10 (0.90)	0.29 (0.93)	0.79 (1.14)	0.71 (1.43)	0.57 (1.91)
		1300	0.07 (0.44)	0.29 (0.68)	0.20 (0.79)	0.26 (0.89)	0.46 (0.94)	0.94 (0.91)	0.21 (0.95)	0.39 (1.15)	0.58 (1.48)	2.14 (2.04)
		50	0.14 (0.44)	0.26 (0.72)	0.43 (0.85)	0.91 (0.97)	0.60 (1.04)	0.43 (0.98)	0.25 (1.01)	0.59 (1.20)	2.23 (1.49)	0.98 (2.04)
		400	0.11 (0.44)	0.10 (0.68)	0.56 (0.80)	0.43 (0.91)	0.65 (0.97)	0.29 (0.92)	0.24 (0.94)	0.38 (1.14)	0.08 (1.46)	0.61 (1.93)
		50	0.08 (0.44)	0.47 (0.78)	0.15 (0.94)	0.44 (1.12)	0.68 (1.24)	1.10 (1.22)	0.43 (1.30)	0.65 (1.55)	0.50 (1.87)	2.07 (2.42)
		400	0.11 (0.44)	0.09 (0.70)	0.81 (0.82)	0.37 (0.95)	0.98 (1.03)	1.12 (1.02)	0.43 (1.09)	0.49 (1.33)	0.29 (1.60)	0.48 (2.11)
		50	0.46 (0.44)	0.47 (0.78)	0.19 (0.95)	0.17 (1.13)	0.34 (1.23)	0.34 (1.17)	0.88 (1.21)	1.00 (1.43)	0.74 (1.78)	1.85 (2.42)
10 TeV	$M_{\text{med}}/8\pi$	400	0.22 (0.44)	0.12 (0.71)	0.41 (0.84)	0.23 (0.96)	1.50 (1.03)	0.84 (0.98)	0.38 (1.02)	0.70 (1.24)	0.79 (1.58)	1.12 (2.11)
		50	0.08 (0.44)	0.32 (0.77)	0.38 (0.95)	0.32 (1.13)	0.19 (1.25)	0.45 (1.24)	0.35 (1.32)	0.50 (1.56)	0.85 (1.88)	1.10 (2.39)
		400	0.23 (0.44)	0.24 (0.71)	0.59 (0.83)	0.38 (0.96)	1.07 (1.05)	0.78 (1.03)	0.33 (1.09)	0.58 (1.33)	1.32 (1.58)	1.74 (2.11)
		50	0.08 (0.44)	0.21 (0.79)	0.71 (0.96)	0.72 (1.14)	1.19 (1.24)	0.76 (1.18)	0.55 (1.19)	0.66 (1.44)	0.42 (1.74)	0.65 (2.40)
		400	0.10 (0.44)	0.16 (0.72)	0.55 (0.84)	0.57 (0.97)	1.19 (1.03)	0.44 (0.99)	0.36 (1.02)	0.62 (1.23)	0.78 (1.56)	1.29 (2.10)
		50	0.05 (0.44)	0.14 (0.78)	0.12 (0.95)	0.19 (1.13)	0.37 (1.25)	0.46 (1.24)	1.51 (1.31)	0.24 (1.51)	0.97 (1.81)	1.22 (2.39)
		400	0.13 (0.44)	0.18 (0.71)	0.04 (0.83)	0.85 (0.96)	0.21 (1.04)	0.46 (1.04)	0.52 (1.10)	0.52 (1.33)	0.04 (1.59)	1.19 (2.11)
30 TeV	$M_{\text{med}}/3$	50	0.11 (0.44)	0.63 (0.79)	0.44 (0.96)	0.87 (1.14)	0.49 (1.24)	0.95 (1.17)	0.80 (1.19)	1.48 (1.39)	1.38 (1.74)	0.86 (2.42)
		400	0.18 (0.44)	0.22 (0.72)	0.26 (0.84)	0.98 (0.96)	0.36 (1.04)	0.56 (0.99)	0.28 (1.03)	0.48 (1.24)	0.79 (1.55)	1.61 (2.08)

Table I.9: Showering systematics in % on the cross-section (σ) and acceptance by signal region (SR) for the simplified model. Numbers in parentheses are the corresponding statistical uncertainties. The cross-section uncertainty and the acceptance uncertainty for the signal region used for limit setting are bolded. As this is a LO signal model, the cross-section uncertainty is not used for limit setting. The acceptance uncertainty is seen to be consistent with 0 within statistics, and thus no extra showering uncertainty is used.

Operator	m_χ [GeV]	$\Delta\sigma$ [%]	SR1 [%]	SR2 [%]	SR3 [%]	SR4 [%]	SR5 [%]	SR6 [%]	SR7 [%]	SR8 [%]	SR9 [%]
C1	10	3.29 (0.74)	5.19 (2.25)	5.69 (3.39)	6.14 (5.17)	2.88 (6.47)	6.95 (7.21)	5.58 (7.73)	7.78 (9.52)	9.94 (6.03)	12.03 (9.58)
	50	1.73 (0.73)	3.53 (2.02)	4.21 (3.00)	6.05 (4.27)	5.23 (5.47)	6.56 (5.54)	5.27 (6.18)	9.09 (8.69)	12.10 (9.48)	18.33 (9.30)
	100	1.29 (0.73)	2.23 (1.78)	4.42 (2.50)	3.22 (3.41)	5.57 (4.33)	2.78 (4.43)	5.67 (4.60)	4.88 (6.23)	10.27 (7.55)	13.55 (8.71)
	200	2.72 (0.73)	1.74 (1.54)	2.47 (2.06)	3.66 (2.69)	3.17 (3.20)	4.36 (3.16)	4.06 (3.26)	3.86 (4.44)	4.59 (6.50)	7.13 (6.91)
	400	5.82 (0.73)	1.26 (1.34)	1.14 (1.66)	2.50 (2.04)	2.70 (2.31)	3.18 (2.29)	2.70 (2.43)	4.57 (3.06)	8.15 (3.73)	11.92 (5.85)
	700	10.00 (0.73)	0.90 (1.27)	0.72 (1.50)	0.43 (1.74)	3.51 (1.94)	4.58 (1.93)	3.75 (2.01)	5.61 (2.49)	4.78 (3.27)	7.00 (4.80)
	1000	12.62 (0.75)	0.88 (1.25)	1.43 (1.43)	0.88 (1.67)	0.80 (1.80)	1.56 (1.79)	0.30 (1.85)	0.98 (2.34)	0.35 (2.98)	1.59 (4.16)
1300	15.82 (0.79)	1.09 (1.29)	0.77 (1.48)	1.61 (1.68)	0.48 (1.79)	1.06 (1.75)	1.02 (1.81)	3.06 (2.19)	4.37 (2.88)	2.40 (4.09)	
C5	10	35.63 (0.77)	1.17 (1.42)	0.80 (1.71)	0.26 (2.01)	1.22 (2.18)	3.11 (2.14)	3.03 (2.17)	0.96 (2.49)	4.45 (2.99)	1.52 (3.90)
	50	36.02 (0.76)	1.51 (1.34)	1.61 (1.58)	1.70 (1.83)	0.80 (2.01)	2.48 (2.00)	3.67 (2.05)	0.61 (2.26)	2.91 (2.69)	1.20 (3.57)
	100	35.12 (0.75)	1.00 (1.26)	1.13 (1.45)	1.41 (1.64)	1.06 (1.78)	1.00 (1.76)	1.38 (1.81)	0.52 (2.09)	0.55 (2.65)	2.20 (3.37)
	200	35.13 (0.77)	0.72 (1.21)	1.38 (1.35)	2.11 (1.50)	1.63 (1.57)	1.28 (1.56)	0.95 (1.60)	0.94 (1.83)	2.88 (2.27)	3.51 (2.97)
	400	36.30 (0.80)	0.24 (1.17)	0.17 (1.27)	2.08 (1.36)	0.64 (1.42)	0.40 (1.42)	1.19 (1.44)	1.45 (1.65)	1.77 (2.02)	2.56 (2.62)
	700	37.45 (0.83)	0.66 (1.17)	0.47 (1.25)	1.40 (1.33)	0.68 (1.36)	0.65 (1.35)	1.51 (1.37)	1.10 (1.54)	1.66 (1.83)	2.35 (2.35)
	1000	38.55 (0.86)	0.53 (1.19)	0.44 (1.27)	0.78 (1.36)	0.94 (1.42)	0.69 (1.32)	0.96 (1.34)	0.31 (1.49)	0.87 (1.78)	3.26 (2.20)
1300	40.42 (0.88)	1.66 (1.18)	0.85 (1.25)	0.89 (1.32)	0.59 (1.36)	0.42 (1.34)	2.26 (1.36)	3.68 (1.54)	6.36 (1.91)	7.97 (2.55)	
D1	10	1.52 (0.73)	2.43 (1.72)	2.78 (2.31)	3.28 (2.97)	4.70 (3.59)	8.44 (3.58)	8.94 (3.78)	8.14 (4.78)	14.80 (4.66)	17.48 (7.19)
	50	2.71 (0.74)	2.08 (1.62)	3.92 (2.19)	4.99 (2.87)	6.45 (3.41)	6.94 (3.52)	3.89 (3.79)	5.05 (3.86)	6.57 (6.31)	11.71 (7.24)
	100	3.20 (0.73)	1.20 (1.55)	2.58 (2.03)	1.81 (2.58)	2.37 (3.08)	3.91 (3.08)	5.52 (3.25)	6.81 (4.54)	6.65 (4.21)	8.19 (6.50)
	200	5.06 (0.73)	1.56 (1.43)	3.34 (1.82)	4.09 (2.27)	3.20 (2.60)	3.91 (2.64)	3.24 (2.74)	2.74 (3.24)	2.57 (4.10)	5.33 (6.08)
	400	7.69 (0.74)	0.53 (1.32)	1.38 (1.58)	2.70 (1.89)	2.24 (2.10)	2.25 (2.11)	2.66 (2.23)	2.07 (2.73)	2.34 (3.47)	5.48 (5.21)
	700	11.19 (0.74)	0.89 (1.26)	2.25 (1.49)	3.10 (1.74)	1.57 (1.88)	2.00 (1.89)	1.70 (1.95)	1.31 (2.39)	1.31 (3.11)	4.18 (4.37)
	1000	14.05 (0.77)	0.67 (1.27)	0.98 (1.47)	1.91 (1.66)	1.96 (1.80)	2.35 (1.77)	1.58 (1.82)	4.01 (2.23)	1.94 (2.28)	2.60 (4.26)
1300	16.43 (1.06)	1.61 (1.30)	1.99 (1.50)	1.44 (1.69)	1.15 (1.82)	1.50 (1.77)	1.52 (1.83)	2.29 (2.24)	2.15 (2.92)	1.47 (4.07)	
D5	10	6.16 (0.70)	1.47 (1.29)	2.67 (1.58)	2.25 (1.85)	1.24 (2.00)	0.37 (1.90)	2.04 (1.96)	3.64 (2.27)	5.01 (2.84)	6.54 (3.97)
	50	6.21 (0.70)	0.96 (1.27)	1.44 (1.55)	1.88 (1.84)	2.45 (2.00)	2.70 (1.88)	2.62 (1.90)	2.46 (2.29)	3.91 (2.83)	5.02 (3.79)
	100	6.45 (0.70)	1.06 (1.25)	1.66 (1.52)	1.82 (1.79)	0.72 (1.93)	0.80 (1.83)	2.89 (1.90)	4.37 (2.34)	5.18 (2.86)	6.74 (3.84)
	200	7.29 (0.70)	0.94 (1.21)	1.37 (1.43)	0.96 (1.68)	1.95 (1.84)	0.84 (1.77)	1.63 (1.83)	2.39 (2.27)	1.73 (3.17)	3.16 (3.72)
	400	8.94 (0.70)	0.35 (1.16)	0.36 (1.35)	1.12 (1.54)	0.56 (1.66)	1.66 (1.59)	1.93 (1.66)	2.70 (2.01)	3.25 (2.50)	1.56 (3.39)
	700	11.30 (0.70)	0.35 (1.12)	0.55 (1.29)	0.67 (1.45)	1.29 (1.53)	0.39 (1.47)	1.49 (1.54)	1.85 (1.83)	2.68 (2.32)	2.12 (3.14)
	1000	13.99 (0.70)	0.51 (1.11)	0.65 (1.27)	0.70 (1.42)	0.54 (1.51)	1.21 (1.45)	1.15 (1.50)	0.62 (1.82)	1.94 (2.28)	0.54 (3.13)
1300	15.51 (0.71)	0.96 (1.12)	1.11 (1.30)	1.97 (1.48)	2.06 (1.59)	1.00 (1.58)	3.70 (1.78)	2.20 (1.81)	2.87 (2.30)	2.37 (3.09)	
D9	10	5.64 (0.69)	1.43 (1.12)	1.98 (1.30)	2.01 (1.46)	3.12 (1.53)	3.18 (1.44)	3.74 (1.48)	3.58 (1.75)	4.64 (2.21)	5.76 (2.90)
	50	6.07 (0.69)	1.32 (1.10)	1.97 (1.27)	2.79 (1.42)	2.82 (1.48)	3.86 (1.42)	3.53 (1.44)	2.88 (1.73)	4.67 (2.18)	6.94 (2.85)
	100	6.64 (0.69)	0.79 (1.08)	1.52 (1.23)	2.30 (1.37)	2.78 (1.43)	2.82 (1.37)	2.84 (1.41)	3.94 (1.71)	4.37 (2.18)	5.72 (2.78)
	200	7.65 (0.69)	0.95 (1.05)	0.94 (1.19)	0.92 (1.31)	0.93 (1.36)	1.85 (1.30)	2.13 (1.34)	2.50 (1.58)	3.98 (2.03)	6.32 (2.73)
	400	9.27 (0.69)	0.29 (1.07)	0.53 (1.20)	0.87 (1.31)	2.00 (1.35)	1.81 (1.29)	1.55 (1.32)	1.74 (1.56)	2.33 (1.92)	2.43 (2.49)
	700	11.54 (0.70)	0.76 (1.07)	1.02 (1.21)	0.91 (1.35)	1.18 (1.40)	0.40 (1.32)	0.94 (1.35)	1.91 (1.58)	1.11 (1.93)	2.10 (2.52)
	1000	13.07 (0.70)	0.88 (1.08)	0.20 (1.22)	0.43 (1.36)	1.26 (1.42)	0.52 (1.35)	1.40 (1.40)	1.52 (1.64)	3.64 (2.02)	4.71 (2.63)
1300	15.17 (0.71)	0.99 (1.08)	1.27 (1.23)	0.87 (1.37)	0.64 (1.46)	0.74 (1.40)	2.03 (1.45)	1.12 (1.74)	1.57 (2.19)	3.80 (2.99)	
D11	10	35.19 (0.78)	0.29 (1.20)	0.24 (1.33)	0.69 (1.45)	0.74 (1.53)	0.67 (1.51)	0.76 (1.54)	0.75 (1.73)	3.03 (2.10)	3.33 (2.73)
	50	35.58 (0.77)	1.62 (1.21)	1.48 (1.33)	0.93 (1.45)	1.12 (1.51)	1.53 (1.49)	0.91 (1.52)	1.86 (1.74)	4.29 (2.10)	5.56 (2.67)
	100	35.83 (0.78)	2.26 (1.17)	1.08 (1.29)	1.01 (1.41)	1.00 (1.48)	0.36 (1.47)	0.69 (1.52)	0.29 (1.71)	2.04 (2.10)	1.64 (2.71)
	200	36.25 (0.79)	0.28 (1.18)	0.43 (1.28)	0.58 (1.38)	0.95 (1.43)	1.01 (1.41)	1.08 (1.44)	0.78 (1.65)	0.11 (2.02)	2.49 (2.56)
	400	36.79 (0.82)	0.95 (1.17)	1.08 (1.26)	0.64 (1.34)	0.76 (1.39)	1.95 (1.38)	1.71 (1.39)	1.06 (1.57)	3.14 (1.91)	1.81 (2.42)
	700	38.44 (0.85)	0.81 (1.17)	0.99 (1.24)	1.15 (1.33)	1.05 (1.36)	0.40 (1.34)	0.39 (1.35)	0.32 (1.51)	0.43 (1.78)	2.14 (2.28)
	1000	39.92 (0.86)	0.09 (1.17)	0.07 (1.24)	0.47 (1.30)	0.32 (1.34)	1.18 (1.34)	1.48 (1.36)	2.53 (1.55)	2.53 (1.82)	2.39 (2.29)
1300	41.14 (0.90)	0.68 (1.18)	1.57 (1.24)	0.57 (1.32)	0.55 (1.37)	0.94 (1.38)	0.28 (1.41)	1.48 (1.55)	2.53 (1.88)	3.84 (2.37)	

Table I.10: Factorization and renormalization scale systematic uncertainties in % on the cross-section (σ) and acceptance by signal region (SR) for the EFT operators. Numbers in parentheses are the corresponding statistical uncertainties. The cross-section uncertainty and the acceptance uncertainty for the signal region used for limit setting are bolded. The cross-section uncertainty is not used for limit setting. A conservative acceptance uncertainty of 3% is used for all EFT samples.

listed in Tables I.10 and I.11 for EFT and Z' simplified models respectively. Scale uncertainties for the simplified model are found to be larger in some parts of parameter space, which is not surprising now that a mediator is included. EFT uncertainties are found to be larger for gluon-initiated states (D11 and C5) than for quark-initiated states (D1, D5, D9, and C1). A conservative scale uncertainty on the acceptance of 3% (6%) is used for all EFT (simplified model) samples.

I.2.4 Matching scales and procedure

As mentioned in Appendix I.1.1, matching is a necessary procedure to ensure that double-counting does not occur when using MadGraph5 for the ME calculation and Pythia6 for the parton shower. The DM samples used in the mono-jet analysis are each comprised of two sub-samples with matching scales of 80 and 300 GeV, chosen in order to provide sufficient statistics at higher energies. This is important as mono-jet limits are typically set in signal regions with higher $E_{\text{T}}^{\text{miss}}$ cuts.

This detail becomes particularly important when carefully investigating the side-effects of matching scales. It is unnatural to impose a low matching scale on a high p_{T} jet, as the matching scale suppresses the radiation of additional jets above that scale. For example, consider two jets, one of 200 GeV and one of 800 GeV. The 200 GeV jet has the ability to radiate a jet of 80 GeV and become a jet of 120 GeV, a process which will be forbidden by the matching scale, but this is a rare process. On the other hand,

M_{med}	Γ_{med}	m_{χ} [GeV]	$\Delta\sigma$ [%]	SR1 [%]	SR2 [%]	SR3 [%]	SR4 [%]	SR5 [%]	SR6 [%]	SR7 [%]	SR8 [%]	SR9 [%]
10 GeV	$M_{\text{med}}/3$	10	4.15 (0.46)	3.20 (1.89)	7.12 (3.03)	7.55 (4.45)	10.94 (5.98)	10.81 (6.37)	11.31 (6.78)	7.07 (8.82)	7.40 (4.63)	9.29 (7.53)
		50	3.72 (0.46)	3.15 (1.18)	3.79 (1.71)	4.57 (2.37)	5.21 (2.97)	5.23 (3.09)	6.32 (3.10)	7.61 (3.51)	5.49 (3.66)	9.34 (4.58)
		100	3.74 (0.45)	2.66 (0.97)	3.10 (1.30)	2.51 (1.69)	4.02 (2.00)	3.87 (1.94)	4.41 (1.99)	4.89 (2.35)	6.49 (3.16)	9.11 (3.95)
		200	5.25 (0.45)	0.88 (0.82)	1.60 (1.02)	2.15 (1.25)	2.35 (1.40)	3.20 (1.36)	4.37 (1.39)	4.35 (1.68)	4.50 (2.13)	4.02 (3.00)
		400	7.69 (0.44)	0.82 (0.75)	0.52 (0.89)	0.98 (1.03)	1.27 (1.11)	1.34 (1.07)	1.97 (1.09)	2.71 (1.31)	3.26 (1.67)	3.39 (2.34)
		700	10.56 (0.44)	0.78 (0.71)	0.80 (0.82)	0.66 (0.94)	1.12 (1.00)	1.08 (0.95)	0.21 (0.98)	0.23 (1.19)	0.92 (1.50)	0.70 (2.02)
		1000	12.66 (0.44)	0.33 (0.70)	0.62 (0.81)	1.36 (0.91)	1.04 (0.96)	0.34 (0.92)	1.19 (0.95)	0.47 (1.14)	0.81 (1.46)	1.28 (1.92)
1300	15.08 (0.45)	0.19 (0.70)	0.15 (0.80)	0.21 (0.90)	0.45 (0.96)	0.71 (0.93)	1.21 (0.98)	1.83 (1.25)	1.15 (1.79)	2.18 (2.83)		
50 GeV	$M_{\text{med}}/3$	10	3.71 (0.46)	4.00 (1.81)	2.77 (2.89)	5.35 (4.46)	2.56 (5.85)	5.40 (6.21)	6.57 (6.86)	7.21 (8.94)	13.50 (8.57)	10.14 (7.44)
		50	3.73 (0.46)	2.54 (1.19)	4.19 (1.75)	5.41 (2.45)	6.78 (3.10)	4.35 (3.24)	6.20 (3.61)	6.10 (4.50)	6.31 (4.38)	11.74 (8.03)
		100	3.72 (0.45)	1.52 (0.96)	2.01 (1.28)	2.91 (1.68)	4.63 (2.01)	4.95 (2.03)	5.17 (2.15)	4.50 (2.60)	2.70 (3.03)	5.06 (3.91)
		200	5.32 (0.45)	1.38 (0.83)	0.62 (1.03)	1.59 (1.24)	1.95 (1.39)	2.52 (1.37)	2.29 (1.42)	2.82 (1.70)	4.06 (2.07)	6.40 (2.99)
		400	7.47 (0.44)	0.22 (0.75)	0.73 (0.89)	0.83 (1.03)	1.64 (1.11)	1.90 (1.06)	1.94 (1.10)	2.17 (1.32)	2.95 (1.69)	2.85 (2.35)
		700	10.45 (0.44)	0.43 (0.71)	0.38 (0.82)	1.02 (0.94)	2.09 (1.00)	2.11 (0.96)	1.72 (0.99)	0.60 (1.18)	1.09 (1.51)	1.77 (2.05)
		1000	12.65 (0.44)	0.25 (0.70)	0.31 (0.80)	0.14 (0.90)	1.16 (0.95)	0.61 (0.91)	1.41 (0.94)	2.10 (1.14)	2.59 (1.44)	2.61 (1.91)
1300	14.79 (0.45)	0.23 (0.70)	0.12 (0.80)	0.29 (0.90)	0.83 (0.96)	1.19 (0.92)	1.06 (0.96)	0.73 (1.15)	1.49 (1.45)	1.02 (1.96)		
100 GeV	$M_{\text{med}}/8\pi$	50	3.54 (0.46)	3.71 (1.21)	4.50 (1.78)	4.37 (2.46)	5.13 (3.06)	5.54 (3.12)	5.51 (3.20)	6.81 (3.26)	4.83 (2.89)	5.26 (4.64)
		10	4.21 (0.46)	3.63 (1.40)	5.16 (2.18)	7.21 (3.24)	5.72 (4.24)	5.75 (4.33)	4.29 (4.22)	15.87 (5.14)	8.05 (4.97)	6.97 (5.93)
		50	3.94 (0.46)	4.37 (1.29)	5.59 (1.96)	8.09 (2.83)	8.57 (3.57)	10.90 (3.75)	10.83 (3.74)	7.16 (3.94)	12.49 (5.63)	10.25 (5.55)
		200	5.39 (0.45)	0.68 (0.82)	1.14 (1.02)	1.93 (1.25)	2.46 (1.42)	2.64 (1.36)	3.21 (1.39)	4.70 (1.73)	5.65 (2.11)	6.10 (3.00)
		400	7.71 (0.44)	0.60 (0.75)	0.86 (0.88)	1.76 (1.02)	1.26 (1.11)	0.57 (1.07)	1.47 (1.11)	1.94 (1.33)	2.18 (1.67)	3.25 (2.31)
		700	10.40 (0.44)	1.21 (0.72)	1.65 (0.83)	1.91 (0.94)	2.20 (1.00)	1.32 (0.96)	0.87 (0.99)	0.74 (1.20)	2.37 (1.49)	1.71 (2.05)
		1000	12.86 (0.44)	0.07 (0.70)	0.28 (0.80)	0.21 (0.90)	0.47 (0.96)	0.25 (0.92)	0.07 (0.95)	0.62 (1.14)	0.86 (1.47)	1.45 (1.91)
1300	14.88 (0.45)	0.24 (0.71)	0.18 (0.80)	0.53 (0.90)	1.06 (0.96)	0.65 (0.91)	0.12 (0.95)	0.82 (1.50)	0.66 (1.44)	1.04 (1.96)		
300 GeV	$M_{\text{med}}/8\pi$	50	3.69 (0.46)	4.50 (1.40)	5.00 (2.20)	9.49 (3.34)	9.99 (4.36)	3.69 (4.44)	6.60 (4.62)	8.50 (5.50)	9.55 (3.78)	6.91 (6.07)
		10	3.62 (0.45)	1.65 (0.97)	2.72 (1.30)	3.69 (1.72)	3.82 (2.04)	4.97 (2.04)	5.45 (2.13)	6.02 (2.45)	7.14 (3.21)	8.62 (3.93)
		50	3.74 (0.45)	1.54 (0.95)	2.26 (1.29)	3.17 (1.70)	3.86 (2.04)	4.33 (2.05)	3.14 (2.10)	5.67 (2.54)	4.96 (2.85)	6.35 (4.15)
		200	4.85 (0.45)	1.55 (0.83)	2.40 (1.04)	2.92 (1.27)	3.90 (1.46)	4.52 (1.49)	4.84 (1.57)	5.43 (1.99)	6.65 (2.39)	9.02 (3.20)
		400	7.46 (0.44)	0.54 (0.74)	0.94 (0.88)	0.79 (1.03)	2.02 (1.14)	0.93 (1.11)	1.17 (1.18)	2.30 (1.41)	3.79 (1.70)	4.51 (2.32)
		700	10.46 (0.44)	0.64 (0.71)	1.14 (0.82)	1.09 (0.94)	0.41 (1.00)	0.36 (0.95)	1.26 (0.97)	0.49 (1.17)	0.80 (1.49)	1.29 (2.03)
		1000	12.88 (0.44)	0.38 (0.70)	0.30 (0.80)	0.15 (0.90)	0.69 (0.96)	0.72 (0.91)	0.27 (0.95)	1.48 (1.16)	1.55 (1.44)	1.37 (1.93)
1300	14.88 (0.45)	0.46 (0.70)	0.34 (0.80)	0.57 (0.90)	0.42 (0.95)	0.54 (0.91)	0.49 (0.95)	0.70 (1.15)	1.85 (1.50)	0.98 (1.94)		
600 GeV	$M_{\text{med}}/8\pi$	50	3.50 (0.45)	1.83 (0.96)	2.95 (1.31)	6.40 (1.77)	6.51 (2.12)	4.52 (2.12)	4.35 (2.22)	6.91 (2.70)	6.88 (3.35)	8.70 (4.41)
		10	5.05 (0.45)	0.36 (0.83)	0.39 (1.04)	1.11 (1.27)	3.40 (1.45)	2.06 (1.39)	3.06 (1.43)	4.60 (1.77)	7.07 (2.20)	8.43 (3.10)
		50	5.11 (0.45)	0.98 (0.83)	1.39 (1.04)	1.73 (1.26)	1.92 (1.41)	1.95 (1.37)	2.91 (1.41)	5.37 (1.66)	6.80 (2.20)	8.35 (3.05)
		200	5.35 (0.44)	0.15 (0.80)	0.88 (0.99)	1.89 (1.20)	1.47 (1.35)	3.32 (1.30)	2.93 (1.36)	4.09 (1.67)	5.59 (2.00)	5.41 (2.88)
		400	7.63 (0.44)	0.33 (0.74)	0.95 (0.88)	0.93 (1.04)	0.45 (1.14)	0.42 (1.13)	0.57 (1.18)	1.14 (1.45)	2.91 (1.76)	2.96 (2.36)
		700	10.18 (0.41)	0.41 (0.64)	0.64 (0.74)	0.98 (0.85)	1.16 (0.93)	0.98 (0.93)	1.18 (0.99)	1.20 (1.21)	1.39 (1.48)	1.76 (1.98)
		1000	12.83 (0.45)	0.29 (0.69)	0.27 (0.80)	0.64 (0.90)	1.60 (0.96)	1.11 (0.92)	0.96 (0.94)	0.60 (1.13)	1.04 (1.44)	0.65 (1.93)
1300	14.98 (0.48)	0.42 (0.78)	0.71 (0.79)	0.67 (0.90)	1.23 (0.95)	0.90 (0.91)	0.33 (0.93)	0.70 (1.15)	1.07 (1.44)	0.27 (1.99)		
1 TeV	$M_{\text{med}}/8\pi$	50	5.34 (0.45)	0.41 (0.81)	1.28 (1.01)	1.38 (1.22)	2.61 (1.36)	2.19 (1.33)	3.22 (1.37)	4.77 (1.69)	5.51 (2.16)	5.34 (3.11)
		400	7.43 (0.44)	0.16 (0.75)	0.34 (0.89)	0.43 (1.04)	0.89 (1.13)	1.69 (1.08)	1.88 (1.12)	1.42 (1.34)	3.46 (1.66)	4.71 (2.37)
		10	6.58 (0.44)	0.78 (0.77)	0.85 (0.94)	1.28 (1.11)	1.98 (1.21)	1.72 (1.16)	2.03 (1.18)	2.52 (1.44)	2.81 (1.79)	3.20 (2.49)
		50	6.63 (0.44)	0.92 (0.77)	1.19 (0.93)	1.04 (1.10)	1.61 (1.21)	1.91 (1.15)	2.61 (1.17)	2.35 (1.40)	3.88 (1.78)	5.32 (2.48)
		200	6.83 (0.44)	0.46 (0.75)	0.36 (0.91)	0.76 (1.06)	1.40 (1.17)	1.79 (1.13)	2.01 (1.16)	3.18 (1.39)	3.79 (1.75)	5.23 (2.47)
		400	7.76 (0.44)	0.18 (0.73)	0.75 (0.87)	1.12 (1.02)	1.20 (1.10)	1.12 (1.06)	0.88 (1.09)	1.51 (1.32)	2.90 (1.69)	3.81 (2.33)
		1000	12.64 (0.45)	0.61 (0.69)	0.93 (0.79)	0.88 (0.90)	0.79 (0.97)	0.53 (0.95)	1.03 (1.00)	0.99 (1.20)	1.50 (1.47)	0.69 (1.93)
1300	14.95 (0.45)	0.43 (0.68)	0.43 (0.79)	0.43 (0.90)	0.60 (0.96)	0.67 (0.94)	0.69 (1.00)	0.46 (1.22)	0.36 (1.52)	0.59 (1.94)		
3 TeV	$M_{\text{med}}/8\pi$	50	7.43 (0.44)	0.48 (0.74)	0.50 (0.89)	0.72 (1.04)	1.38 (1.14)	1.77 (1.09)	2.41 (1.13)	2.62 (1.37)	3.64 (1.72)	3.51 (2.42)
		400	7.28 (0.43)	4.33 (0.70)	7.35 (0.81)	13.14 (0.90)	19.47 (0.91)	25.14 (0.82)	27.69 (0.83)	29.22 (1.03)	30.05 (1.31)	30.87 (1.80)
		10	7.62 (0.44)	1.34 (0.77)	1.55 (0.93)	2.19 (1.09)	2.30 (1.19)	3.14 (1.12)	3.14 (1.13)	4.72 (1.35)	5.04 (1.71)	4.69 (2.28)
		50	7.57 (0.44)	1.37 (0.77)	1.76 (0.93)	1.48 (1.08)	1.16 (1.17)	1.73 (1.13)	1.54 (1.13)	3.50 (1.36)	4.29 (1.69)	5.81 (2.30)
		200	8.46 (0.44)	0.56 (0.73)	0.55 (0.87)	1.14 (1.01)	0.85 (1.09)	1.10 (1.04)	2.05 (1.07)	2.81 (1.28)	2.93 (1.60)	4.24 (2.17)
		400	10.35 (0.44)	0.55 (0.71)	0.80 (0.83)	1.12 (0.95)	0.67 (1.01)	1.02 (0.96)	0.93 (0.99)	0.99 (2.00)	2.27 (1.52)	1.57 (2.07)
		1000	13.82 (0.45)	0.44 (0.68)	0.49 (0.79)	0.12 (0.90)	0.78 (0.96)	1.10 (0.92)	0.79 (0.94)	1.06 (1.13)	2.62 (1.44)	0.92 (1.93)
1300	14.93 (0.45)	0.42 (0.68)	0.31 (0.79)	0.83 (0.90)	0.36 (0.96)	0.83 (0.91)	0.61 (0.95)	0.89 (1.17)	1.27 (1.48)	2.52 (2.01)		
6 TeV	$M_{\text{med}}/8\pi$	50	10.73 (0.45)	0.97 (0.73)	1.59 (0.86)	1.67 (0.98)	1.92 (1.05)	1.68 (0.99)	1.92 (1.01)	3.32 (1.21)	4.31 (1.52)	3.98 (2.06)
		400	13.05 (0.45)	0.63 (0.69)	0.37 (0.81)	0.56 (0.92)	1.02 (0.98)	0.60 (0.93)	0.93 (0.95)	1.04 (1.16)	2.39 (1.47)	2.79 (1.97)
		50	6.66 (0.45)	1.68 (0.79)	1.88 (0.95)	2.20 (1.13)	2.13 (1.25)	2.14 (1.24)	3.48 (1.30)	3.85 (1.54)	4.03 (1.82)	4.08 (2.44)
		400	9.08 (0.44)	0.54 (0.71)	0.80 (0.83)	0.37 (0.97)	0.52 (1.05)	0.79 (1.04)	0.36 (1.09)	0.74 (1.31)	2.44 (1.61)	4.45 (2.15)
		50	6.89 (0.44)	1.39 (0.80)	1.72 (0.96)	1.80 (1.14)	1.75 (1.24)	1.93 (1.18)	2.34 (1.20)	3.29 (1.44)	3.90 (1.78)	3.83 (2.42)
		400	9.33 (0.44)	0.96 (0.73)	0.98 (0.85)	0.94 (0.98)	1.05 (1.04)	0.61 (1.00)	1.53 (1.03)	2.84 (1.24)	3.51 (1.60)	4.37 (2.13)
		50	6.57 (0.45)	0.63 (0.78)	1.22 (0.96)	1.13 (1.14)	2.33 (1.26)	2.27 (1.25)	2.48 (1.32)	3.10 (1.58)	4.43 (1.82)	5.97 (2.36)
10 TeV	$M_{\text{med}}/8\pi$	400	9.35 (0.44)	0.12 (0.71)	0.17 (0.84)	0.93 (0.97)	1.65 (1.06)	2.27 (1.05)	2.89 (1.11)	2.81 (1.35)	2.08 (1.66)	2.71 (2.14)
		50	6.60 (0.44)	0.52 (0.80)	1.46 (0.97)	2.23 (1.14)	2.75 (1.25)	2.91 (1.18)	2.54 (1.21)	2.88 (1.44)	3.93 (1.77)	5.45 (2.45)
		400	9.18 (0.44)	0.04 (0.73)	0.22 (0.85)	0.43 (0.97)	0.89 (1.04)	0.94 (1.00)	0.84 (1.03)	1.99 (1.24)	2.33 (1.57)	2.60 (2.10)
		50	6.61 (0.45)	0.99 (0.78)	1.06 (0.95)	2.13 (1.13)	3.12 (1.25)	2.74 (1.25)	3.57 (1.33)	4.64 (1.58)	4.45 (1.89)	5.03 (2.50)
		400	9.39 (0.44)	0.35 (0.72)	0.16 (0.84)	0.28 (0.97)	0.56 (1.05)	0.25 (1.04)	1.05 (1.10)	1.85 (1.32)	1.49 (1.62)	1.00 (2.11)
		50	6.59 (0.44)	0.86 (0.80)	1.59 (0.97)	2.65 (1.15)	3.50 (1.25)	4.02 (1.19)	3.97 (1.22)	2.91 (1.44)	1.95 (1.81)	3.24 (2.47)
		400	9.14 (0.44)	0.33 (0.73)	0.08 (0.85)	0.31 (0.97)	0.77 (1.04)	1.26 (1.00)	1.17 (1.03)	1.09 (1.25)	2.82 (1.54)	2.85 (2.09)

Table I.11: Factorization and renormalization scale systematic uncertainties in % on the cross-section (σ) and acceptance by signal region (SR) for the simplified model. Numbers in parentheses are the corresponding statistical uncertainties. The cross-section uncertainty and the acceptance uncertainty for the signal region used for limit setting are bolded. The cross-section uncertainty is not used for limit setting. A conservative acceptance uncertainty of 6% is used for all simplified model samples. This covers all but one point, where there is likely a sample problem.

the 800 GeV jet can easily radiate an 80 GeV jet to become a 720 GeV jet. As such, as the jet p_T is increased with respect to the matching scale, the events which remain after matching are artificially tuned to be dominantly one very high p_T jet and maybe other very low p_T jets. This process is shaping the stochastic nature of the parton shower.

Typical matching inefficiencies are on the order of a few percent. In some specific combinations which have particularly hard jet spectra, the inefficiency can grow to approximately 40 or 50%. This is a significant bias which needs to be addressed when working with very heavy DM masses, operators with harder kinematic spectra, or operators which dominantly involve gluons and thus typically result in higher levels of QCD radiation. It is particularly dangerous because it leaves lower energy regions mostly unchanged, while modifying the high energy distributions, and thus the events used for limit setting.

In addition to providing higher statistics at high energies, the 300 GeV matching scale samples help to alleviate this problem. While the use of two matching scales does not completely resolve the problem, it does reduce the matching inefficiency to the level of a few percent. This clear bias must be covered by a systematic. In order to quantify this effect, additional truth-level samples with a matching scale of 190 GeV were generated. The E_T^{miss} , a variable strongly correlated to the jet p_T and the variable used for defining mono-jet signal regions, is then used to study the change in acceptance. Matching scales of 80 and 190 GeV are compared at the level of $E_T^{\text{miss}} > 250$ GeV. This is above the turn-on, and thus any difference beyond statistics is a bias of the matching procedure. The same is done for matching scales of 190 and 300 GeV for $E_T^{\text{miss}} > 350$ GeV. The resulting matching scale uncertainties are listed in Tables I.12 and I.13 for EFTs and simplified models respectively. In both cases, the resulting uncertainty is often by not always consistent with zero within statistics. Conservative uncertainties of 3% for events with a E_T^{miss} of below 350 GeV and 5% for events with a E_T^{miss} of above 350 GeV are used.

One additional check of the matching procedure used was conducted. The typical matching procedure is to generate one sample per ME parton multiplicity and match them to `Pythia6` in an exclusive procedure. Then, the highest ME parton multiplicity sample is matched to `Pythia6` with an inclusive procedure. Instead, our samples are at least one parton with an optional second parton, and then the samples are matched to `Pythia6` with the inclusive procedure. It was also worried that only two ME jets were considered, where the mono-jet selection no longer has a jet multiplicity requirement, and thus the impact of further ME partons may be important. In order to check this, the nominal samples were compared against new samples following the correct matching procedure and split into 1-parton, 2-parton, and 3-parton ME samples. This was done for the D5 EFT operators as a standard benchmark scenario with a DM mass of 100 GeV.

The differences between the nominal samples and this more complete set are listed in Table I.14. The maximum difference observed in the acceptance is 3.6% at very high E_T^{miss} , where statistics are limited, although the kinematic distributions show no distinguishable features. Limits are instead set in lower E_T^{miss} regions, where the difference is below the level of 2%, and is thus consistent with zero within statistics. The leading jet p_T and E_T^{miss} distributions, shown in Figure I.5, confirm the lack of statistically significant deviations. The combination of these two effects is thus a small effect in comparison to other theoretical uncertainties, including the 5% matching scale systematic. The difference is considered to be negligible, and no additional systematic uncertainty is considered for this potential source.

Operator	m_χ [GeV]	$E_T^{\text{miss}} > 250$ GeV [%]	$E_T^{\text{miss}} > 350$ GeV [%]
D1	50	0.06 (1.7)	0.27 (2.07)
	100	1.37 (1.6)	1.71 (1.98)
	200	0.65 (1.42)	3.06 (1.67)
	400	1.36 (1.18)	1.93 (1.43)
	700	0.77 (1.08)	3.06 (1.26)
	1000	0.36 (1.07)	2.88 (1.21)
	1300	1.13 (1.08)	3.08 (1.26)
D5	50	2.28 (1.1)	2.77 (1.26)
	100	0.53 (1.11)	1.6 (1.25)
	200	2.19 (1.03)	4.44 (1.15)
	400	0.18 (0.99)	0.7 (1.14)
	700	1.35 (0.92)	3.8 (1.02)
	1000	1.43 (0.9)	3.33 (1.0)
	1300	0.31 (0.91)	2.17 (1.02)
D9	10	1.32 (0.94)	1.52 (1.02)
	50	0.1 (0.91)	1.23 (1.01)
	100	1.12 (0.88)	1.78 (0.98)
	200	1.15 (0.85)	1.94 (0.94)
	400	2.28 (0.84)	4.76 (0.92)
	700	0.13 (0.87)	3.39 (0.94)
	1000	0.63 (0.88)	3.83 (0.95)
	1300	3.1 (0.86)	3.21 (0.97)
D11	50	0.75 (0.97)	2.7 (1.05)
	100	0.34 (0.93)	1.76 (1.02)
	200	2.26 (0.89)	4.64 (0.96)
	400	1.18 (0.88)	3.69 (0.94)
	700	2.18 (0.86)	4.33 (0.91)
	1000	1.82 (0.88)	3.97 (0.91)
	1300	1.44 (0.87)	3.15 (0.93)
C1	10	0.76 (3.06)	4.8 (3.86)
	50	1.75 (2.46)	2.97 (3.23)
	100	0.85 (2.04)	3.07 (2.78)
	200	1.83 (1.57)	3.98 (1.99)
	400	0.74 (1.27)	0.36 (1.57)
	700	1.56 (1.09)	4.58 (1.27)
	1000	0.79 (1.08)	2.35 (1.23)
	1300	0.66 (1.07)	3.89 (1.2)
C5	10	0.4 (1.27)	4.68 (1.36)
	50	0.18 (1.17)	3.57 (1.3)
	100	1.36 (1.05)	2.92 (1.19)
	200	0.44 (0.98)	2.46 (1.08)
	400	1.47 (0.9)	2.77 (0.97)
	700	2.38 (0.86)	3.31 (0.92)
	1000	0.81 (0.86)	3.6 (0.91)
	1300	2.36 (0.86)	3.51 (0.93)

Table I.12: Matching scale acceptance uncertainties in percent for the EFT operators. Numbers in parentheses are the corresponding statistical uncertainties.

M_{med}	Γ_{med}	m_χ [GeV]	$E_T^{\text{miss}} > 250$ GeV [%]	$E_T^{\text{miss}} > 350$ GeV [%]
10 GeV	$M_{\text{med}}/3$	50	1.03	4.26
	$M_{\text{med}}/8\pi$	50	0.91	2.8
50 GeV	$M_{\text{med}}/3$	50	0.89	1.27
	$M_{\text{med}}/8\pi$	50	2.18	4.76
100 GeV	$M_{\text{med}}/3$	50	0.34	3.08
	$M_{\text{med}}/8\pi$	50	1.02	2.05
300 GeV	$M_{\text{med}}/3$	50	2.35	3.49
		400	0.94	4.6
	$M_{\text{med}}/8\pi$	50	1.42	2.97
		400	1.1	3.01
600 GeV	$M_{\text{med}}/3$	50	1.08	1.4
		400	0.8	2.99
	$M_{\text{med}}/8\pi$	50	1.9	3.37
		400	2.26	4.16
1 TeV	$M_{\text{med}}/3$	50	0.99	2.1
		400	0.96	3.34
	$M_{\text{med}}/8\pi$	50	0.57	5.39
		400	0.53	2.6
3 TeV	$M_{\text{med}}/3$	50	2.14	5.43
		400	3.1	5.47
	$M_{\text{med}}/8\pi$	50	0.29	3.01
		400	0.92	2.53
6 TeV	$M_{\text{med}}/3$	50	2.71	5.21
		400	3.06	4.7
	$M_{\text{med}}/8\pi$	50	1.58	3.52
		400	0.9	4.04
10 TeV	$M_{\text{med}}/3$	50	3.09	4.59
		400	3.22	6.71
	$M_{\text{med}}/8\pi$	50	0.89	1.44
		400	1.49	3.3

Table I.13: Matching scale acceptance uncertainties in percent for the Z' simplified model.

	SR1	SR2	SR3	SR4	SR5	SR6	SR7	SR8	SR9
Difference [%]	0.84	0.51	0.93	0.38	0.39	1.84	1.88	2.09	3.56

Table I.14: The difference in acceptance between the full matching procedure with an extra ME parton and the nominal procedure used for general sample generation. The comparison is for the benchmark D5 EFT operator with a DM mass of 100 GeV. The difference is largest at the highest signal regions, where statistics are limited. The region where limits are set, SR7, appears to be equal to zero within statistics.

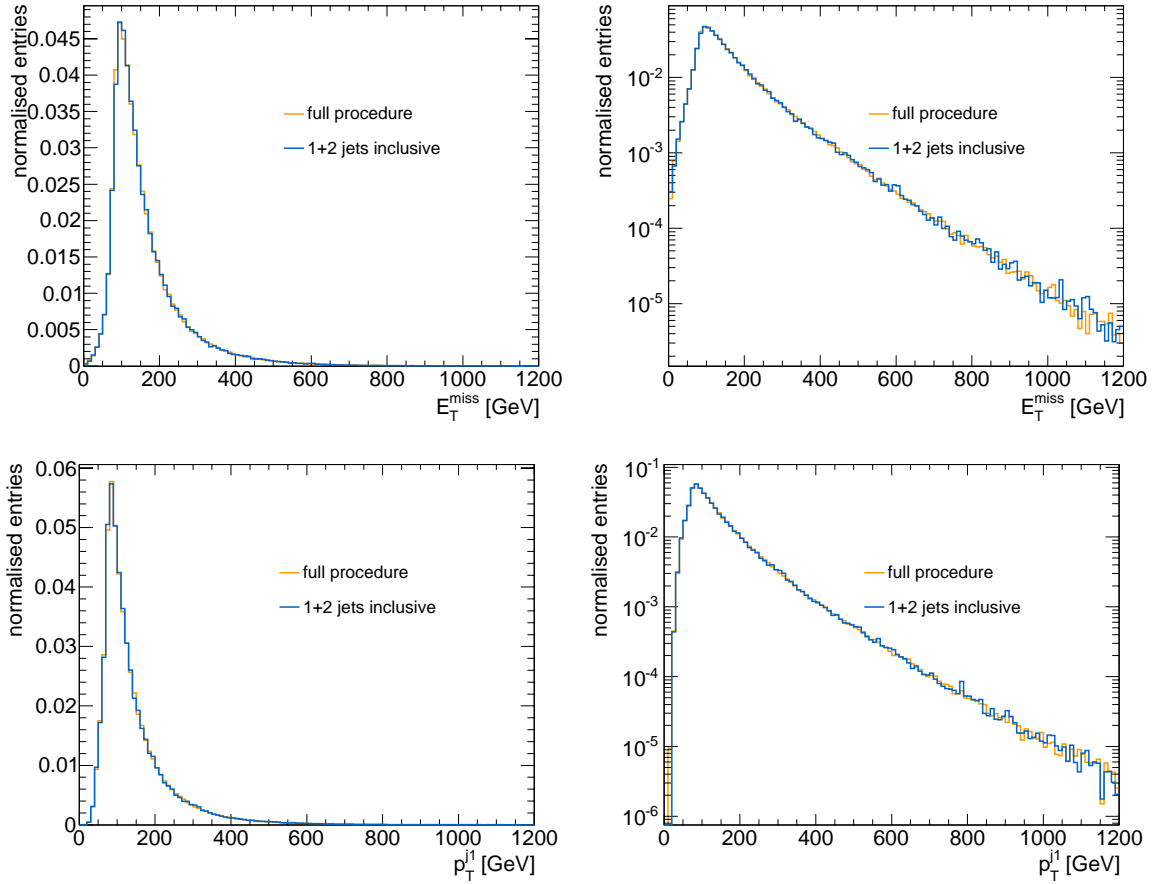


Figure I.5: Comparison of the E_T^{miss} (top) and leading jet p_T (bottom) distributions for the complete matching scale procedure with an extra ME parton (yellow) compared to the nominal procedure used for official sample production (blue). The comparison is for the benchmark D5 EFT operator with a DM mass of 100 GeV. No statistically significant differences are observed.

I.3 Experimental systematic uncertainties

Experimental systematic uncertainties on DM signals are not protected by the data-driven estimation used in the mono-jet analysis. As such, the same uncertainty source such as the JES can have a much larger impact than observed for the background estimation. Regardless, the experimental systematics generally remain sub-dominant due to the frequently large theoretical uncertainties as discussed in Appendix I.2.

I.3.1 Detector uncertainties

There are four main types of uncertainties related to the ATLAS detector which impact signal samples. The first is the lepton scale and resolution, which enters through the mono-jet lepton vetoes. The quadrature sum of the electron and muon energy and scale uncertainties will be referred to simply as the lepton-related uncertainty.

The JES and JER are much more important. The JER is particularly large at low p_T , which impacts the number of jets observed above the p_T threshold of 30 GeV, and thus most notably the $\Delta\phi(E_T^{\text{miss}}, \text{jets})$ cut. At higher p_T it becomes less important, while the JES remains relevant. In particular, the JES impacts both the leading jet p_T cut and the E_T^{miss} cut, as the E_T^{miss} is dominated by the high p_T jet(s) when lepton vetoes are applied.

The remaining large detector uncertainty source is the E_T^{miss} soft term scale and resolution. This is particularly important at low E_T^{miss} , where the soft term makes up a larger portion of the full E_T^{miss} . In very high E_T^{miss} scenarios, the high p_T jets form the majority of the balance, and so the impact of the soft term is reduced.

The quadrature sums of all of these uncertainties are listed in Tables I.15 and I.16 for EFT operators and simplified models respectively. The combined detector uncertainty is frequently at the level of 2 to 6%, making it larger than for the background estimation, but still generally small compared to the theoretical systematics.

I.3.2 Beam energy uncertainties

The LHC data taken in 2012 is generally considered to have a centre of mass energy \sqrt{s} of 8 TeV. However, this is not strictly true. Reference [154] measured the actual per-beam energy, and determined it to be $3988 \pm 5(\text{stat}) \pm 26(\text{syst})$ GeV. For topologies involving low momentum transfer processes, this is a negligible difference. However, for high momentum transfer selections such as mono-jet, this can become very important. As described in Appendix I.2.1, the PDFs used are already required to be in a region of very high momentum transfer and very high momentum fraction. The PDFs are therefore in a region where all results are extrapolated from previous experiments at lower energies, and where the actual distributions are not well known. As such, small changes in the nominal beam energy can lead to notable variations in the required momentum fraction.

It is important to consider this source of uncertainty in addition to the PDF uncertainty already described. While they are both PDF effects, the PDF systematic uncertainty can be thought of as a lack of knowledge of the PDF “in use” by nature, while the beam energy systematic is the uncertainty on the scale of the interaction at which to evaluate the PDF. In order to do this, additional truth samples were produced with beam energies of 3962, 3988, and 4014 GeV to account for the quoted systematic uncertainty on the beam energy.

Operator	m_χ [GeV]	$E_{\text{T}}^{\text{miss}}$ threshold [GeV]								
		150	200	250	300	350	400	500	600	700
D1	50	+3.6	+4.2	+2.6	+3.6	+4.4	+3.4	+5.4	+3.6	+6.8
		-3.7	-4.7	-4.7	-6.4	-3.7	-5.0	-4.5	-6.3	-11.1
	100	+3.3	+4.9	+3.8	+5.0	+4.1	+7.2	+4.6	+4.9	+5.6
		-3.6	-4.5	-3.0	-5.9	-7.9	-3.7	-4.1	-15.6	-8.4
	200	+2.9	+2.4	+3.9	+3.4	+5.3	+3.5	+5.9	+5.4	+9.1
		-3.2	-4.1	-2.3	-2.9	-5.2	-4.5	-5.5	-7.0	-13.0
	400	+2.7	+2.9	+2.0	+2.5	+3.7	+4.9	+5.7	+5.9	+4.4
		-2.9	-3.1	-3.1	-3.2	-2.1	-3.2	-4.5	-3.6	-3.7
	700	+2.3	+2.7	+3.5	+2.6	+5.3	+3.5	+3.8	+6.6	+7.5
		-2.6	-2.1	-1.9	-3.3	-2.9	-4.9	-3.7	-4.4	-8.6
	1000	+2.9	+3.7	+2.4	+4.4	+2.8	+2.2	+3.2	+5.8	+7.7
		-2.3	-3.0	-2.8	-3.5	-5.2	-3.7	-5.3	-3.2	-6.9
	1300	+1.5	+2.2	+2.0	+2.6	+2.1	+3.1	+4.3	+4.2	+6.9
		-2.6	-3.4	-2.5	-2.6	-4.1	-4.7	-5.0	-3.8	-6.7
D5	50	+2.7	+2.7	+3.5	+3.4	+4.4	+4.3	+3.5	+5.1	+3.5
		-3.1	-3.8	-4.3	-4.7	-3.9	-4.2	-7.1	-6.8	-4.1
	100	+2.6	+3.7	+4.1	+4.4	+4.0	+3.5	+3.1	+6.7	+4.5
		-3.8	-1.8	-2.7	-2.0	-3.4	-3.6	-3.9	-3.6	-5.9
	200	+2.4	+2.7	+3.7	+2.0	+4.7	+2.5	+4.6	+3.5	+4.3
		-3.7	-2.2	-2.9	-4.2	-3.1	-4.2	-4.3	-5.8	-7.1
	400	+2.5	+2.5	+2.6	+3.8	+3.4	+2.6	+4.9	+4.8	+4.7
		-2.2	-4.0	-2.6	-3.3	-3.3	-3.1	-3.5	-4.6	-6.0
	700	+2.0	+2.8	+2.9	+4.0	+3.2	+2.4	+4.0	+2.2	
		-2.8	-2.3	-2.9	-4.4	-3.3	-3.5	-3.0	-4.1	-4.0
	1000	+2.1	+1.9	+2.9	+2.6	+2.6	+3.9	+3.4	+4.2	+3.5
		-2.0	-3.1	-2.8	-2.0	-3.1	-4.4	-4.5	-5.2	-5.2
	1300	+2.7	+2.8	+1.7	+3.3	+3.9	+3.5	+5.1	+4.9	+4.8
		-1.9	-2.1	-2.6	-2.6	-3.7	-3.6	-3.3	-3.5	-3.4
D9	10	+1.5	+2.0	+2.1	+3.0	+1.9	+3.7	+3.8	+3.6	+4.4
		-2.5	-2.1	-2.6	-2.7	-3.2	-1.9	-3.7	-3.6	-5.2
	50	+2.5	+2.0	+2.0	+2.1	+3.2	+2.9	+3.1	+2.7	+2.4
		-2.2	-2.6	-2.9	-2.6	-3.7	-2.4	-2.5	-3.7	-3.9
	100	+1.7	+2.6	+2.6	+2.2	+2.3	+2.3	+4.4	+2.8	+3.4
		-1.5	-2.2	-2.1	-2.4	-2.2	-3.0	-3.5	-3.5	-4.1
	200	+1.6	+1.9	+2.0	+2.1	+1.9	+2.7	+2.7	+5.0	+3.7
		-1.8	-2.0	-1.5	-2.2	-2.4	-2.6	-3.8	-4.2	-3.0
	400	+2.1	+2.1	+2.4	+3.3	+3.0	+2.7	+3.0	+3.6	+4.0
		-1.6	-1.3	-2.4	-2.7	-2.5	-2.6	-2.3	-2.6	-4.8
	700	+2.4	+1.6	+1.9	+2.2	+1.8	+2.9	+3.8	+3.7	+3.2
		-1.8	-1.8	-2.0	-2.4	-2.7	-1.7	-3.2	-4.1	-4.8
	1000	+2.0	+2.1	+2.0	+1.7	+2.5	+1.9	+2.8	+4.9	+3.9
		-1.9	-1.9	-2.2	-2.9	-3.1	-3.6	-3.2	-5.0	-4.1
	1300	+1.5	+1.4	+2.2	+2.4	+2.3	+2.9	+5.7	+4.4	+4.3
		-2.3	-2.8	-2.5	-2.9	-2.8	-3.6	-3.0	-5.2	-2.6
D11	50	+2.1	+2.3	+2.9	+2.1	+2.8	+2.3	+3.9	+3.7	+4.1
		-2.2	-1.8	-1.9	-2.9	-2.2	-4.8	-2.9	-2.7	-4.3
	100	+1.4	+1.3	+2.4	+1.8	+2.1	+2.3	+3.2	+3.9	+5.5
		-1.6	-2.1	-2.4	-2.2	-3.6	-3.7	-3.3	-3.2	-6.0
	200	+1.1	+0.9	+1.8	+2.6	+3.2	+2.6	+3.0	+3.3	+5.1
		-1.9	-1.9	-1.7	-2.0	-2.1	-3.9	-3.2	-5.9	-5.3
	400	+1.2	+1.6	+2.0	+2.6	+2.0	+2.2	+2.6	+2.8	+3.8
		-1.8	-1.5	-1.2	-2.5	-2.1	-2.4	-3.6	-2.7	-4.2
	700	+1.3	+1.2	+1.4	+1.5	+2.0	+3.1	+3.5	+3.5	+3.1
		-1.4	-1.6	-2.2	-2.4	-2.3	-2.7	-3.7	-3.6	-4.4
	1000	+1.0	+1.5	+1.8	+1.8	+2.4	+2.2	+2.8	+2.5	+4.0
		-1.4	-1.2	-1.7	-2.1	-2.1	-2.0	-3.5	-3.6	-3.0
	1300	+1.4	+2.0	+1.2	+0.9	+3.8	+2.3	+1.6	+3.6	+3.5
		-1.4	-1.2	-2.0	-2.5	-3.4	-2.2	-3.4	-2.8	-3.3
C1	10	+6.5	+6.7	+7.9	+8.4	+4.6	+6.0	+5.8	+5.3	\emptyset
		-6.0	-6.6	-9.8	-7.1	-13.1	-4.6	-6.8	-7.7	\emptyset
	50	+5.7	+5.9	+11.7	+10.5	+14.1	+4.2	+5.9	+6.9	\emptyset
		-5.9	-5.7	-2.3	-4.0	-4.4	-5.0	-8.7	-6.7	\emptyset
	100	+4.3	+6.8	+5.5	+4.1	+8.9	+5.1	+7.4	+9.8	+9.3
		-5.4	-4.7	-5.4	-2.5	-9.8	-4.3	-3.7	-3.6	-7.0
	200	+4.6	+4.2	+5.5	+7.0	+8.6	+5.4	+4.9	+11.8	+2.3
		-4.6	-3.0	-5.1	-2.7	-4.8	-3.7	-5.4	-6.5	-12.0
	400	+2.4	+3.8	+2.9	+3.6	+4.0	+5.0	+5.4	+7.6	+4.7
		-3.0	-3.0	-2.7	-3.7	-3.2	-4.6	-7.3	-9.3	-4.3
	700	+2.7	+3.2	+4.1	+2.4	+2.6	+3.8	+3.5	+3.5	+4.0
		-2.2	-3.2	-1.8	-3.6	-3.2	-2.6	-5.2	-8.6	-4.0
	1000	+1.9	+2.9	+2.4	+3.2	+3.9	+2.6	+3.3	+3.2	+4.0
		-2.6	-1.9	-2.6	-1.8	-2.0	-2.9	-3.9	-4.3	-7.8
	1300	+1.8	+2.1	+3.0	+2.4	+2.2	+2.8	+4.7	+5.5	+7.2
		-1.8	-2.9	-2.0	-2.6	-1.5	-3.9	-3.7	-3.9	-5.4
C5	50	+1.3	+2.4	+1.9	+3.7	+3.0	+5.4	+4.2	+5.4	+4.8
		-3.1	-3.0	-2.8	-3.4	-2.4	-3.1	-3.2	-4.6	-6.5
	100	+2.7	+2.6	+2.8	+3.0	+3.2	+3.4	+6.0	+6.5	+3.7
		-2.0	-2.0	-3.2	-4.2	-5.4	-4.9	-3.2	-3.3	-5.3
	200	+1.2	+2.1	+2.0	+1.4	+2.1	+3.9	+3.4	+3.5	+5.2
		-2.1	-2.7	-3.3	-3.2	-4.0	-4.3	-4.4	-4.6	-5.0
	400	+1.3	+1.7	+2.0	+2.2	+2.5	+2.5	+3.3	+3.9	+3.2
		-1.5	-1.5	-1.8	-2.7	-2.4	-2.6	-2.8	-4.0	-3.4
	700	+1.4	+1.5	+2.1	+1.8	+2.9	+3.1	+3.0	+3.7	+4.4
		-1.5	-1.3	-1.5	-1.2	-1.5	-1.7	-2.1	-3.1	-3.0
	1000	+1.5	+1.7	+1.1	+1.8	+2.8	+3.2	+3.1	+2.3	+3.2
		-1.8	-1.3	-1.9	-2.7	-2.8	-2.6	-2.3	-3.5	-3.8
	1300	+1.3	+1.7	+1.4	+1.2	+2.2	+1.9	+2.8	+3.0	+3.2
		-1.5	-1.4	-1.5	-2.4	-1.8	-2.7	-2.5	-3.5	-4.9

Table I.15: The quadrature sum of all detector-related uncertainties for the EFT models. The value \emptyset denotes that the uncertainty was not evaluated due to insufficient statistics.

M_{med}	Γ_{med}	m_χ [GeV]	E_T^{emiss} threshold [GeV]									
			150	200	250	300	350	400	500	600	700	
10 GeV	$M_{\text{med}}/3$	10	4.9	7.8	6.8	11.4	10.1	8.7	\emptyset	\emptyset	\emptyset	
		10	5.4	7.0	9.3	7.3	8.2	6.7	29.7	8.0	6.8	
50 GeV	$M_{\text{med}}/3$	50	5.0	6.0	9.1	6.3	7.7	7.6	5.7	8.6	8.4	
		400	3.2	3.8	4.0	4.8	4.2	4.2	6.2	5.4	7.5	
	$M_{\text{med}}/8\pi$	50	5.9	8.4	6.7	7.6	6.0	5.9	18.4	11.0	7.9	
		400	3.0	3.5	2.6	3.9	4.4	4.2	6.2	5.2	6.7	
100 GeV	$M_{\text{med}}/3$	10	5.6	7.6	9.3	8.5	11.2	11.8	7.2	9.3	9.8	
		50	6.4	8.4	11.1	9.1	9.9	8.6	7.6	7.6	12.5	
	$M_{\text{med}}/8\pi$	400	3.0	3.3	3.4	3.0	3.5	4.1	4.6	5.1	6.3	
		50	6.8	5.7	8.0	8.9	8.7	4.7	\emptyset	\emptyset	\emptyset	
300 GeV	$M_{\text{med}}/3$	400	3.2	3.3	3.5	3.4	5.2	5.4	6.7	6.3	5.3	
		10	3.4	4.2	5.6	6.5	5.9	9.6	5.4	11.4	11.7	
		50	5.5	5.3	5.2	8.2	5.7	6.7	5.9	7.1	6.0	
	$M_{\text{med}}/8\pi$	200	3.0	3.7	4.0	4.9	4.9	5.2	4.6	7.3	7.4	
		400	3.3	3.3	3.6	4.1	5.1	4.8	5.8	5.9	7.7	
		50	4.9	5.3	7.7	6.8	5.6	5.3	6.5	7.4	10.6	
600 GeV	$M_{\text{med}}/3$	400	2.9	3.4	3.6	4.4	4.5	3.8	4.4	8.6	7.3	
		10	3.0	4.1	3.1	4.2	4.8	5.5	6.5	6.2	9.2	
		50	89.8	82.2	70.0	49.9	23.7	13.6	5.4	6.5	9.4	
	$M_{\text{med}}/8\pi$	200	3.0	3.3	4.3	5.4	4.5	5.2	6.0	5.1	6.6	
		400	2.7	3.2	4.9	4.7	4.8	4.1	6.6	5.3	10.2	
		50	3.9	4.4	3.8	6.3	5.2	6.2	5.6	5.0	8.5	
1 TeV	$M_{\text{med}}/3$	400	3.2	3.0	3.1	4.4	3.8	5.3	5.0	5.5	7.6	
		10	3.1	3.0	3.8	3.9	3.7	4.4	4.2	5.3	5.3	
		50	3.1	3.3	4.1	4.7	4.5	4.4	5.7	4.6	6.5	
	$M_{\text{med}}/8\pi$	200	2.6	3.0	4.0	3.5	3.5	4.6	4.8	4.9	7.0	
		400	2.9	3.1	3.5	4.8	4.8	5.0	5.7	5.7	6.6	
		1000	2.0	2.5	3.5	3.3	2.8	3.5	4.1	4.7	5.7	
3 TeV	$M_{\text{med}}/3$	50	2.9	3.5	3.9	4.7	4.2	3.8	5.5	5.5	5.7	
		400	2.8	3.2	3.6	3.9	2.9	4.8	4.4	5.8	5.7	
		10	2.3	2.9	3.6	4.0	4.2	3.7	4.5	3.4	7.4	
	$M_{\text{med}}/8\pi$	50	3.4	3.3	3.0	3.8	3.3	4.4	4.4	4.5	8.3	
		200	2.4	3.0	3.4	3.3	3.6	3.3	5.4	4.3	5.7	
		400	2.7	3.1	4.3	3.1	2.7	4.3	5.1	4.9	8.9	
6 TeV	$M_{\text{med}}/3$	1000	2.3	2.6	3.1	3.2	3.8	4.2	4.2	4.0	6.4	
		1300	2.5	3.9	5.0	8.9	12.6	11.1	\emptyset	\emptyset	\emptyset	
		50	3.1	3.0	3.9	3.6	4.2	4.3	4.7	5.1	5.8	
	$M_{\text{med}}/8\pi$	400	2.4	2.6	3.0	3.4	3.4	3.8	3.7	5.2	6.4	
		50	3.6	3.0	3.9	3.9	4.6	4.7	4.7	6.7	4.8	
		400	2.4	3.7	3.2	4.6	3.6	4.0	4.0	4.9	5.3	
10 TeV	$M_{\text{med}}/3$	50	3.6	3.8	4.3	5.0	4.9	4.6	5.5	6.8	7.0	
		400	3.3	3.3	2.5	3.1	3.8	4.3	3.9	5.2	6.6	
	$M_{\text{med}}/8\pi$	50	3.5	4.2	3.3	3.8	4.9	5.4	6.8	5.8	4.8	
		400	2.8	3.0	3.6	3.4	3.7	4.6	5.5	6.4	5.2	
30 TeV	$M_{\text{med}}/3$	50	2.9	3.0	2.9	3.2	5.1	3.5	5.7	6.0	7.5	
		400	2.7	3.1	3.8	3.5	4.4	4.6	4.1	4.4	6.8	
	$M_{\text{med}}/8\pi$	50	3.6	3.8	3.8	3.9	4.9	4.8	5.1	4.5	5.1	
		400	2.8	3.3	3.3	2.9	4.1	4.7	4.4	5.0	5.4	
30 TeV	$M_{\text{med}}/8\pi$	50	3.3	3.7	3.3	3.8	4.2	4.3	6.3	4.9	4.8	
		400	3.0	3.3	3.4	4.8	4.6	4.8	6.5	9.7	7.9	

Table I.16: The quadrature sum of all detector-related uncertainties for the Z' simplified model. The value \emptyset denotes that the uncertainty was not evaluated due to insufficient statistics.

Operator	m_χ [GeV]	$\Delta\sigma$ [%]	SR1 [%]	SR2 [%]	SR3 [%]	SR4 [%]	SR5 [%]	SR6 [%]	SR7 [%]	SR8 [%]	SR9 [%]
C1	10	1.69 (0.72)	0.91 (2.29)	0.19 (3.56)	2.53 (5.16)	1.78 (6.74)	5.27 (7.19)	0.66 (7.61)	11.66 (9.19)	6.42 (5.57)	8.24 (9.32)
	50	1.69 (0.72)	1.30 (2.01)	2.24 (2.99)	3.39 (4.33)	4.92 (5.58)	4.44 (5.73)	4.56 (6.60)	11.80 (10.02)	4.21 (7.81)	4.85 (9.71)
	100	2.22 (0.72)	1.32 (1.77)	1.86 (2.51)	2.30 (3.44)	1.97 (4.26)	3.51 (4.47)	2.33 (4.60)	5.00 (5.60)	4.50 (4.96)	7.07 (8.63)
	200	2.41 (0.72)	1.71 (1.53)	2.43 (2.02)	3.25 (2.61)	3.59 (3.10)	0.81 (3.25)	2.79 (3.41)	0.56 (4.14)	4.81 (6.31)	3.77 (7.42)
	400	3.04 (0.72)	0.13 (1.35)	1.15 (1.67)	0.97 (2.04)	0.37 (2.33)	2.56 (2.31)	3.48 (2.40)	5.01 (2.99)	3.05 (3.78)	3.11 (5.75)
	700	4.42 (0.73)	0.86 (1.27)	0.49 (1.50)	1.58 (1.76)	1.21 (1.92)	1.25 (1.94)	3.53 (2.05)	2.71 (2.48)	3.00 (3.25)	0.90 (4.72)
	1000	5.32 (0.75)	0.52 (1.24)	0.82 (1.44)	0.42 (1.66)	0.49 (1.79)	1.09 (1.78)	1.14 (1.85)	2.34 (2.27)	1.70 (3.05)	3.11 (4.36)
1300	7.31 (0.77)	0.43 (1.27)	0.53 (1.47)	0.23 (1.68)	2.03 (1.78)	1.78 (1.77)	1.92 (1.82)	2.67 (2.30)	1.01 (3.03)	1.52 (4.30)	
C5	10	1.93 (0.70)	1.03 (1.43)	1.22 (1.71)	1.42 (2.00)	1.70 (2.18)	1.32 (2.16)	1.21 (2.17)	1.34 (2.36)	4.16 (2.89)	8.22 (3.95)
	50	1.99 (0.69)	0.54 (1.34)	1.29 (1.57)	3.13 (1.80)	4.18 (1.96)	2.92 (1.96)	2.22 (2.03)	2.81 (2.36)	3.49 (2.81)	2.02 (3.55)
	100	1.87 (0.69)	0.58 (1.26)	0.63 (1.45)	1.35 (1.67)	1.47 (1.78)	1.20 (1.78)	2.02 (1.81)	1.45 (2.07)	2.16 (2.57)	2.21 (3.45)
	200	2.23 (0.71)	0.85 (1.21)	0.80 (1.35)	0.68 (1.50)	1.39 (1.58)	1.29 (1.57)	0.68 (1.59)	1.27 (1.82)	4.28 (2.33)	1.90 (3.05)
	400	3.41 (0.73)	1.18 (1.18)	0.86 (1.28)	0.85 (1.38)	0.98 (1.43)	0.30 (1.41)	0.10 (1.44)	0.46 (1.65)	1.33 (2.02)	2.58 (2.62)
	700	4.01 (0.76)	0.75 (1.17)	1.38 (1.26)	1.45 (1.33)	1.23 (1.37)	2.86 (1.37)	1.71 (1.38)	0.52 (1.53)	1.12 (1.81)	2.15 (2.30)
	1000	5.20 (0.78)	0.55 (1.17)	0.87 (1.25)	0.74 (1.32)	0.61 (1.35)	1.34 (1.34)	1.85 (1.36)	1.04 (1.50)	0.66 (1.81)	1.62 (3.11)
1300	9.98 (2.88)	2.77 (1.15)	2.39 (1.22)	3.45 (1.28)	2.84 (1.32)	1.86 (1.31)	1.28 (1.33)	1.35 (1.48)	1.55 (1.75)	1.33 (2.23)	
D1	10	2.42 (0.73)	1.52 (1.68)	4.68 (2.22)	4.55 (2.90)	7.08 (3.40)	4.54 (3.54)	1.24 (3.60)	2.32 (4.12)	8.43 (6.10)	5.14 (7.09)
	50	2.41 (0.73)	2.43 (1.66)	3.49 (2.22)	1.18 (2.85)	1.67 (3.36)	1.72 (3.45)	1.20 (3.63)	2.54 (5.16)	1.53 (5.33)	1.88 (6.80)
	100	2.64 (0.73)	1.36 (1.55)	1.34 (2.06)	2.23 (2.65)	2.61 (3.10)	2.63 (3.11)	1.15 (3.27)	2.38 (4.19)	2.92 (5.01)	5.39 (6.97)
	200	3.22 (0.73)	0.30 (1.43)	0.50 (1.82)	3.11 (2.22)	1.67 (2.61)	0.89 (2.65)	1.01 (2.83)	2.89 (3.59)	7.49 (4.52)	7.81 (5.94)
	400	4.12 (0.73)	0.76 (1.32)	0.83 (1.60)	2.59 (1.94)	2.34 (2.17)	2.05 (2.15)	2.23 (2.22)	2.21 (2.75)	2.41 (3.50)	5.52 (5.44)
	700	4.54 (0.74)	0.65 (1.27)	1.46 (1.50)	3.79 (1.76)	2.06 (1.90)	1.01 (1.86)	2.57 (1.93)	2.16 (2.41)	1.67 (3.15)	1.83 (4.48)
	1000	6.26 (0.86)	0.21 (1.76)	0.73 (2.26)	1.30 (3.08)	1.30 (4.27)	4.89 (5.80)	7.52 (8.20)	1.22 (2.26)	0.92 (2.90)	0.89 (4.13)
1300	7.80 (0.80)	0.46 (1.30)	0.78 (1.50)	0.65 (1.72)	0.74 (1.86)	0.78 (1.85)	1.68 (1.98)	2.96 (2.84)	6.12 (4.29)	9.73 (7.55)	
D5	10	2.12 (0.70)	0.80 (1.28)	1.31 (1.57)	1.46 (1.86)	2.06 (2.03)	0.73 (1.92)	1.43 (1.97)	1.31 (2.50)	2.65 (2.82)	2.14 (4.02)
	50	2.13 (0.70)	0.12 (1.28)	0.74 (1.55)	0.57 (1.84)	1.16 (1.99)	1.45 (1.89)	0.79 (1.95)	1.40 (2.29)	1.78 (2.73)	2.47 (3.91)
	100	2.23 (0.70)	0.38 (1.25)	0.72 (1.51)	0.58 (1.79)	1.49 (1.94)	2.97 (1.87)	1.25 (1.92)	0.95 (2.29)	0.47 (2.80)	3.49 (3.76)
	200	2.67 (0.70)	1.06 (1.20)	1.08 (1.42)	0.90 (1.66)	1.25 (1.80)	1.12 (1.73)	1.36 (1.76)	0.95 (2.10)	2.89 (2.75)	5.19 (3.77)
	400	2.82 (0.69)	0.29 (1.16)	0.46 (1.35)	0.10 (1.55)	0.81 (1.64)	0.17 (1.57)	2.17 (1.63)	1.73 (1.90)	3.42 (2.41)	4.07 (3.31)
	700	3.83 (0.69)	0.02 (1.12)	0.40 (1.29)	0.59 (1.46)	1.45 (1.54)	1.29 (1.47)	1.28 (1.53)	2.35 (1.82)	2.98 (2.29)	4.56 (3.09)
	1000	5.23 (0.69)	0.53 (1.11)	0.60 (1.27)	0.75 (1.43)	1.27 (1.51)	0.73 (1.46)	1.58 (1.50)	2.64 (1.84)	0.99 (2.99)	1.78 (3.07)
1300	6.44 (0.70)	1.38 (1.11)	0.93 (1.27)	2.05 (1.44)	1.62 (1.52)	1.39 (1.46)	1.22 (1.50)	1.30 (1.83)	2.03 (2.32)	2.35 (3.12)	
D9	10	1.89 (0.69)	0.52 (1.11)	0.15 (1.29)	1.00 (1.44)	0.56 (1.51)	0.29 (1.43)	0.17 (1.46)	1.47 (1.73)	2.26 (2.15)	2.10 (2.88)
	50	1.98 (0.69)	1.22 (1.09)	1.42 (1.25)	0.92 (1.40)	0.99 (1.47)	0.56 (1.40)	0.61 (1.43)	2.03 (1.72)	1.56 (2.20)	1.61 (2.93)
	100	2.32 (0.68)	0.41 (1.08)	0.70 (1.23)	0.43 (1.36)	0.42 (1.42)	0.51 (1.37)	0.01 (1.41)	2.07 (1.67)	1.35 (2.11)	2.30 (2.79)
	200	2.44 (0.68)	0.30 (1.07)	0.54 (1.20)	1.27 (1.33)	1.19 (1.38)	1.35 (1.32)	0.89 (1.36)	2.10 (1.64)	2.60 (2.06)	3.27 (2.72)
	400	2.72 (0.68)	0.56 (1.07)	0.16 (1.20)	0.91 (1.32)	1.16 (1.36)	1.41 (1.29)	0.81 (1.34)	1.46 (1.57)	1.20 (1.96)	2.06 (2.52)
	700	3.75 (0.69)	2.27 (1.09)	2.36 (1.23)	1.38 (1.35)	0.88 (1.39)	1.84 (1.34)	0.92 (1.36)	1.95 (1.62)	0.96 (1.94)	1.43 (2.50)
	1000	5.02 (0.69)	1.09 (1.09)	0.71 (1.23)	0.43 (1.37)	0.85 (1.43)	0.92 (1.35)	1.32 (1.39)	1.44 (1.62)	2.88 (2.02)	3.50 (2.63)
1300	6.71 (0.70)	1.45 (1.11)	1.99 (1.28)	2.06 (1.43)	1.86 (1.46)	1.54 (1.40)	1.32 (1.44)	1.51 (1.73)	2.36 (2.19)	5.39 (2.91)	
D11	10	2.71 (0.72)	0.36 (1.21)	0.31 (1.34)	1.79 (1.46)	1.29 (1.53)	2.01 (1.51)	2.66 (1.54)	2.71 (1.79)	2.34 (2.18)	1.81 (2.83)
	50	2.88 (0.72)	0.19 (1.20)	0.72 (1.33)	0.45 (1.46)	0.19 (1.52)	0.73 (1.51)	0.68 (1.54)	1.41 (1.74)	2.75 (2.13)	2.92 (2.79)
	100	2.85 (0.72)	0.46 (1.19)	0.89 (1.31)	0.84 (1.43)	1.28 (1.49)	1.23 (1.48)	1.32 (1.52)	0.49 (1.72)	1.21 (2.12)	2.59 (2.77)
	200	3.31 (0.74)	1.75 (1.23)	1.38 (1.36)	1.25 (1.53)	1.60 (1.68)	1.41 (1.83)	2.15 (2.09)	1.53 (3.07)	2.92 (4.85)	5.85 (8.03)
	400	3.92 (0.75)	1.79 (1.18)	1.87 (1.27)	2.24 (1.35)	2.08 (1.40)	1.81 (1.38)	1.39 (1.41)	1.68 (1.57)	1.48 (1.88)	2.08 (2.40)
	700	5.17 (0.77)	0.72 (1.17)	0.85 (1.24)	1.03 (1.31)	0.35 (1.35)	0.64 (1.34)	0.84 (1.36)	0.16 (1.51)	0.61 (1.80)	0.58 (2.26)
	1000	6.40 (1.92)	1.10 (1.21)	1.50 (1.28)	0.98 (1.37)	1.84 (1.38)	1.29 (1.37)	1.03 (1.41)	0.26 (1.65)	0.39 (1.76)	2.36 (2.22)
1300	7.70 (0.87)	0.46 (1.38)	1.11 (1.53)	0.68 (1.32)	1.20 (1.37)	1.11 (1.36)	1.01 (1.37)	0.89 (1.54)	0.94 (1.89)	2.85 (2.44)	

Table I.17: Beam energy systematic uncertainties in % on the cross-section (σ) and acceptance by signal region (SR) for the EFT operators. Numbers in parentheses are the corresponding statistical uncertainties. The cross-section uncertainty and the acceptance uncertainty for the signal region used for limit setting are in bold. As the required momentum transfer is increased, such as by increasing the DM mass, the required momentum fraction is increased. Working with more extreme values of the momentum fraction further pushes the PDFs, resulting in larger cross-section uncertainties, as observed. The acceptance uncertainty is generally consistent with zero within statistics, but there are exceptions, so a conservative value of 3% was used in all cases.

The results of these studies are listed in Tables I.17 and I.18 for the EFTs and simplified models respectively. The cross-section uncertainties are seen to increase as the momentum transfer is increased, such as by increasing the DM mass. This matches the expectation by which the requirement of a higher momentum fraction will result in a larger variation from changing beam conditions. The acceptance variations are less pronounced, and are often but not always consistent with zero within statistical uncertainties. A conservative beam energy acceptance systematic uncertainty of 3% is used in all cases in order to cover the outliers.

M_{med}	Γ_{med}	m_χ [GeV]	$\Delta\sigma$ [%]	SR1 [%]	SR2 [%]	SR3 [%]	SR4 [%]	SR5 [%]	SR6 [%]	SR7 [%]	SR8 [%]	SR9 [%]		
10 GeV	$M_{\text{med}}/3$	10	1.22 (0.45)	0.80 (1.87)	1.51 (2.96)	3.08 (4.37)	6.05 (5.77)	5.46 (6.62)	3.83 (7.42)	9.00 (8.79)	11.15 (10.01)	6.69 (7.38)		
		50	1.18 (0.45)	1.80 (1.18)	1.64 (1.72)	1.60 (2.39)	2.57 (2.98)	0.86 (2.99)	0.61 (3.12)	4.30 (3.87)	4.91 (4.02)	3.19 (6.59)		
		100	1.32 (0.45)	0.96 (0.96)	2.17 (1.27)	2.46 (1.65)	0.77 (1.98)	0.62 (1.98)	0.62 (2.07)	1.19 (2.70)	1.96 (2.80)	1.21 (4.13)		
		200	1.73 (0.44)	0.81 (0.83)	1.13 (1.03)	0.94 (1.25)	1.58 (1.40)	0.50 (1.37)	2.50 (1.40)	1.82 (1.70)	3.10 (2.10)	3.02 (2.98)		
		400	2.17 (0.44)	0.60 (0.75)	0.78 (0.89)	1.01 (1.03)	1.16 (1.11)	1.66 (1.07)	1.61 (1.11)	1.59 (1.34)	1.38 (1.70)	3.32 (2.29)		
		700	3.07 (0.44)	0.87 (0.72)	0.56 (0.83)	0.78 (0.94)	0.31 (1.00)	0.33 (0.96)	0.31 (0.99)	0.32 (1.20)	1.29 (1.50)	2.25 (2.02)		
		1000	4.56 (0.44)	0.39 (0.70)	0.28 (0.80)	0.42 (0.90)	0.80 (0.96)	0.25 (0.92)	0.61 (0.96)	0.81 (1.14)	0.87 (1.43)	1.02 (1.92)		
		1300	5.81 (0.44)	0.88 (0.70)	0.79 (0.80)	1.28 (0.90)	1.02 (0.95)	1.72 (0.91)	1.28 (0.94)	1.71 (1.13)	1.99 (1.42)	2.80 (2.00)		
50 GeV	$M_{\text{med}}/3$	10	1.02 (0.45)	1.22 (1.80)	1.42 (2.90)	0.62 (4.41)	5.69 (5.70)	4.71 (6.25)	3.76 (6.80)	5.92 (5.43)	8.36 (10.82)	6.18 (7.24)		
		50	1.26 (0.45)	0.61 (1.18)	1.41 (1.74)	3.53 (2.44)	1.84 (3.06)	5.32 (3.33)	4.96 (3.77)	5.99 (4.93)	3.35 (5.57)	3.22 (4.73)		
		100	1.40 (0.45)	0.16 (0.96)	0.23 (1.29)	0.21 (1.69)	0.70 (2.00)	1.67 (1.98)	1.72 (2.00)	1.57 (2.26)	2.14 (2.69)	2.41 (3.93)		
		200	1.52 (0.44)	0.50 (0.83)	1.23 (1.04)	1.07 (1.26)	1.32 (1.42)	1.37 (1.35)	0.66 (1.39)	0.64 (1.67)	0.56 (2.12)	2.97 (3.05)		
		400	2.44 (0.44)	0.75 (0.76)	0.61 (0.89)	0.65 (1.03)	0.72 (1.11)	0.69 (1.07)	1.51 (1.10)	1.84 (1.35)	2.04 (1.71)	2.67 (2.38)		
		700	3.27 (0.44)	0.44 (0.72)	0.56 (0.83)	0.75 (0.94)	0.79 (1.00)	0.92 (0.95)	1.44 (0.99)	1.73 (1.20)	1.81 (1.58)	3.65 (2.28)		
		1000	4.60 (0.44)	0.55 (0.70)	0.52 (0.80)	0.24 (0.91)	0.45 (0.96)	0.34 (0.92)	0.40 (0.95)	0.89 (1.14)	1.33 (1.45)	1.57 (1.93)		
		1300	5.90 (0.44)	0.51 (0.70)	1.04 (0.80)	1.24 (0.90)	1.94 (0.96)	1.85 (0.91)	2.23 (0.95)	2.69 (1.14)	3.08 (1.44)	3.54 (1.94)		
	$M_{\text{med}}/8\pi$	50	1.25 (0.45)	0.29 (1.20)	0.95 (1.74)	2.00 (2.42)	1.68 (3.07)	2.96 (3.09)	3.14 (3.24)	3.89 (3.76)	3.51 (4.07)	7.65 (4.63)		
100 GeV	$M_{\text{med}}/3$	10	1.29 (0.45)	1.17 (1.38)	0.68 (2.15)	0.67 (3.19)	1.90 (4.18)	3.53 (4.27)	4.13 (4.40)	4.39 (5.11)	5.01 (5.76)	0.59 (5.92)		
		50	1.38 (0.45)	1.84 (1.29)	0.80 (1.93)	2.10 (2.80)	2.62 (3.61)	1.42 (3.64)	0.08 (3.77)	5.07 (4.38)	1.18 (3.24)	4.00 (5.53)		
		200	1.67 (0.44)	0.63 (0.83)	0.74 (1.03)	0.57 (1.26)	0.29 (1.41)	0.24 (1.37)	1.10 (1.41)	1.69 (1.75)	1.74 (2.17)	1.00 (3.02)		
		400	2.25 (0.44)	0.06 (0.75)	0.50 (0.89)	0.26 (1.03)	0.69 (1.11)	0.31 (1.07)	0.57 (1.10)	0.75 (1.31)	2.32 (1.67)	2.02 (2.33)		
		700	3.10 (0.44)	0.42 (0.71)	0.28 (0.82)	0.80 (0.93)	0.59 (0.99)	0.46 (0.95)	1.19 (0.98)	1.81 (1.18)	1.91 (1.48)	2.02 (2.01)		
		1000	4.49 (0.44)	0.28 (0.71)	0.81 (0.81)	0.64 (0.92)	1.40 (0.97)	1.09 (0.93)	1.52 (0.96)	1.68 (1.17)	1.85 (1.45)	1.74 (1.96)		
		1300	5.96 (0.44)	0.19 (0.70)	0.31 (0.80)	0.35 (0.90)	0.68 (0.96)	1.23 (0.91)	1.68 (0.95)	1.23 (1.15)	2.04 (1.43)	5.97 (1.89)		
			$M_{\text{med}}/8\pi$	50	1.17 (0.45)	0.57 (1.38)	1.26 (2.17)	2.07 (3.26)	2.29 (4.27)	2.99 (4.50)	8.93 (4.74)	0.08 (4.84)	3.10 (3.65)	3.58 (6.07)
300 GeV	$M_{\text{med}}/3$	10	1.51 (0.45)	0.60 (0.97)	0.81 (1.31)	1.12 (1.71)	1.33 (2.06)	1.35 (2.06)	2.40 (2.13)	2.80 (2.46)	0.96 (2.80)	4.74 (4.21)		
		50	1.50 (0.45)	1.31 (0.97)	1.69 (1.31)	0.98 (1.70)	2.00 (2.03)	0.52 (2.02)	1.33 (2.10)	2.60 (2.41)	1.55 (2.75)	1.39 (4.02)		
		200	1.71 (0.44)	0.55 (0.83)	0.36 (1.04)	0.75 (1.27)	1.24 (1.46)	1.35 (1.47)	0.88 (1.56)	1.69 (1.98)	2.91 (2.35)	3.36 (3.23)		
		400	2.40 (0.44)	0.37 (0.74)	0.68 (0.88)	1.09 (1.02)	1.09 (1.12)	1.22 (1.12)	0.64 (1.19)	1.51 (1.44)	2.37 (1.76)	4.05 (2.33)		
		700	3.27 (0.44)	0.26 (0.70)	0.27 (0.82)	0.24 (0.93)	0.55 (0.99)	0.74 (0.95)	0.68 (0.98)	1.53 (1.19)	2.23 (1.49)	1.78 (2.03)		
		1000	4.68 (0.44)	0.31 (0.70)	0.24 (0.80)	1.32 (0.90)	1.87 (0.95)	0.93 (0.92)	0.30 (0.96)	1.11 (1.14)	2.06 (1.44)	2.72 (1.96)		
		1300	5.92 (0.44)	0.62 (0.70)	0.59 (0.80)	1.36 (0.91)	0.90 (0.96)	1.00 (0.91)	0.86 (0.94)	1.69 (1.14)	1.85 (1.44)	3.28 (1.93)		
			$M_{\text{med}}/8\pi$	50	1.51 (0.45)	1.43 (0.96)	1.38 (1.30)	0.81 (1.71)	1.55 (2.05)	0.66 (2.09)	2.03 (2.17)	1.09 (2.41)	0.44 (3.19)	2.84 (4.19)
600 GeV	$M_{\text{med}}/3$	10	1.70 (0.44)	0.57 (0.83)	1.44 (1.03)	2.31 (1.25)	2.36 (1.41)	1.09 (1.38)	1.09 (1.39)	1.97 (1.70)	3.13 (2.16)	4.96 (3.08)		
		50	1.68 (0.44)	0.13 (0.83)	1.09 (1.03)	1.33 (1.26)	1.40 (1.42)	0.61 (1.38)	0.90 (1.40)	3.40 (1.66)	2.89 (2.23)	2.30 (3.04)		
		200	1.58 (0.44)	0.90 (0.80)	1.76 (1.00)	1.39 (1.20)	0.99 (1.34)	1.25 (1.31)	0.63 (1.33)	2.40 (1.61)	1.79 (2.11)	2.14 (2.93)		
		400	2.08 (0.44)	0.13 (0.74)	0.66 (0.88)	0.55 (1.03)	0.49 (1.14)	1.46 (1.12)	2.13 (1.18)	1.06 (1.46)	1.61 (1.77)	3.60 (2.29)		
		700	3.20 (0.49)	0.48 (0.77)	0.42 (0.89)	0.32 (1.02)	0.42 (1.10)	0.13 (1.07)	0.06 (1.12)	1.09 (1.31)	1.53 (1.57)	1.88 (1.98)		
		1000	4.42 (0.45)	0.38 (0.69)	0.76 (0.80)	0.64 (0.91)	0.36 (0.96)	0.96 (0.91)	1.45 (0.94)	1.75 (1.15)	1.88 (1.42)	1.67 (1.92)		
		1300	5.84 (0.44)	0.54 (0.68)	0.97 (0.79)	1.08 (0.90)	0.99 (0.95)	0.86 (0.91)	0.77 (0.94)	1.21 (1.13)	1.66 (1.43)	3.48 (1.96)		
			$M_{\text{med}}/8\pi$	50	1.55 (0.44)	0.05 (0.81)	0.25 (1.01)	0.86 (1.23)	1.56 (1.38)	1.26 (1.33)	1.52 (1.34)	0.89 (1.62)	1.57 (2.08)	1.51 (2.93)
		400	2.18 (0.44)	0.54 (0.75)	0.63 (0.89)	0.93 (1.04)	0.48 (1.14)	0.49 (1.08)	0.15 (1.10)	1.27 (1.33)	2.04 (1.73)	3.25 (2.34)		
1 TeV	$M_{\text{med}}/3$	10	1.90 (0.44)	0.32 (0.77)	0.27 (0.93)	0.51 (1.10)	0.90 (1.20)	0.66 (1.16)	0.74 (1.18)	1.90 (1.42)	0.49 (1.84)	3.52 (2.55)		
		50	1.93 (0.44)	0.18 (0.77)	0.19 (0.93)	0.72 (1.11)	0.41 (1.21)	1.77 (1.16)	1.26 (1.19)	1.27 (1.42)	1.06 (1.79)	1.57 (2.47)		
		200	1.95 (0.44)	0.66 (0.75)	0.70 (0.91)	0.66 (1.06)	0.77 (1.16)	1.00 (1.12)	1.77 (1.15)	0.91 (1.37)	1.77 (1.78)	2.41 (2.48)		
		400	2.27 (0.44)	0.19 (0.74)	0.25 (0.88)	0.06 (1.02)	0.06 (1.11)	0.42 (1.06)	1.95 (1.08)	1.75 (1.31)	1.19 (1.69)	1.07 (2.33)		
		1000	4.74 (0.44)	0.88 (0.69)	1.35 (0.80)	0.28 (0.90)	0.56 (0.97)	0.67 (0.96)	0.46 (1.01)	0.97 (1.21)	0.83 (1.49)	2.24 (1.95)		
		1300	5.84 (0.45)	0.67 (0.70)	0.79 (0.82)	0.55 (0.90)	0.34 (0.97)	0.33 (0.95)	0.47 (1.00)	0.88 (1.20)	2.38 (1.45)	2.04 (1.90)		
			$M_{\text{med}}/8\pi$	50	1.82 (0.44)	0.64 (0.74)	0.86 (0.89)	0.78 (1.05)	1.29 (1.14)	1.20 (1.10)	1.56 (1.13)	2.20 (1.36)	2.32 (2.08)	2.33 (2.41)
				400	2.12 (0.49)	0.48 (0.91)	1.24 (1.19)	1.95 (1.64)	0.19 (2.39)	1.86 (4.12)	6.89 (7.14)	30.32 (19.66)	47.65 (67.44)	0.00 (0.00)
3 TeV	$M_{\text{med}}/3$	10	2.52 (0.44)	0.81 (0.77)	0.57 (0.93)	0.51 (1.09)	0.55 (1.18)	1.32 (1.12)	0.82 (1.13)	0.97 (1.36)	2.48 (1.64)	2.18 (2.26)		
		50	2.64 (0.44)	0.29 (0.77)	0.49 (0.92)	0.68 (1.08)	1.22 (1.17)	1.02 (1.11)	1.54 (1.12)	2.12 (1.31)	2.24 (1.62)	2.78 (2.23)		
		200	2.86 (0.47)	0.79 (0.74)	0.77 (0.88)	0.88 (1.02)	1.09 (1.10)	0.90 (1.06)	0.86 (1.10)	1.79 (1.36)	2.96 (1.64)	4.25 (2.20)		
		400	3.33 (0.44)	1.95 (0.71)	1.53 (0.84)	0.91 (0.96)	1.05 (1.02)	1.07 (0.97)	1.43 (1.00)	1.36 (1.22)	2.19 (1.53)	3.84 (2.04)		
		1000	5.24 (0.44)	0.82 (0.68)	0.86 (0.79)	0.26 (0.90)	0.69 (0.95)	1.06 (0.91)	1.17 (0.94)	2.07 (1.12)	2.17 (1.41)	2.43 (1.91)		
		1300	5.98 (0.54)	0.45 (0.83)	0.83 (0.96)	0.92 (1.09)	1.35 (1.14)	1.36 (1.06)	1.58 (1.08)	2.14 (1.21)	3.19 (1.47)	4.29 (1.93)		
			$M_{\text{med}}/8\pi$	50	4.22 (0.44)	0.26 (0.73)	0.75 (0.88)	0.90 (1.02)	1.35 (1.12)	0.54 (1.16)	1.17 (1.34)	1.99 (1.27)	2.21 (1.58)	3.31 (2.08)
				400	5.19 (0.44)	0.45 (0.69)	1.29 (0.81)	1.46 (0.92)	1.38 (0.98)	0.91 (0.93)	1.38 (0.96)	1.68 (1.17)	1.52 (1.47)	2.08 (2.00)
6 TeV	$M_{\text{med}}/3$	50	2.18 (0.44)	0.32 (0.78)	0.58 (0.95)	0.66 (1.13)	1.35 (1.25)	1.67 (1.24)	1.85 (1.30)	2.01 (1.56)	2.94 (1.83)	2.67 (2.39)		
		400	2.90 (0.44)	0.55 (0.71)	0.85 (0.84)	0.81 (0.96)	1.69 (1.04)	0.98 (1.03)	1.82 (1.08)	2.76 (1.29)	2.03 (1.59)	2.58 (2.09)		
	$M_{\text{med}}/8\pi$	50	1.88 (0.44)	0.71 (0.80)	0.68 (0.96)	0.70 (1.14)	0.69 (1.24)	1.22 (1.18)	1.12 (1.21)	0.98 (1.43)	2.61 (1.78)	3.26 (2.49)		
		400	3.07 (0.44)	0.27 (0.73)	0.23 (0.85)	0.81 (0.97)	0.40 (1.04)	1.32 (1.00)	2.19 (1.03)	1.87 (1.23)	3.27 (1.56)	4.36 (2.13)		
10 TeV	$M_{\text{med}}/3$	50	2.04 (0.44)	0.77 (0.79)	1.52 (0.97)	2.32 (1.16)	2.35 (1.28)	2.08 (1.27)	2.27 (1.33)	1.60 (1.58)	2.03 (1.90)	2.67 (2.46)		
		400	2.82 (0.44)	0.89 (0.71)	0.67 (0.84)	0.21 (0.97)	0.35 (1.06)	0.66 (1.04)	0.79 (1.09)	1.32 (1.33)	2.15 (1.62)	2.06 (2.16)		
	$M_{\text{med}}/8\pi$	50	3.33 (0.44)	1.29 (0.81)	2.64 (0.98)	3.62 (1.16)	3.48 (1.27)	3.00 (1.23)	2.88 (1.26)	2.04 (1.49)	1.52 (1.81)	1.90 (2.41)		
		400	3.96 (0.44)	1.08 (0.73)	2.06 (0.86)	3.06 (0.99)	2.92 (1.06)	2.33 (1.02)	2.50 (1.05)	1.43 (1.26)	2.25 (1.57)	3.19 (2.15)		
30 TeV</														

Appendix J

Pileup and the Mono-jet Analysis at an Upgraded LHC

This Section motivates the choice of jet p_T threshold used in the 14 TeV analysis, as listed in Section 9.3.

To first order, the mono-jet analysis minimally sensitive to pileup effects due to the leading jet p_T and E_T^{miss} requirements being at levels well beyond the pileup regime. However, the jet multiplicity cut is very sensitive to both pileup and the p_T threshold used to define when an object is a hard-scatter jet or not. Given that the upgrade study was done using truth smearing parametrizations of jets and E_T^{miss} , it is imperative that the jet threshold used does not lead to a significant difference between the truth parametrization and reconstructed quantities. In particular, the scenario where the truth parametrization contains no pileup and the reconstructed quantities contain significant pileup contamination must be avoided.

A first approximation of a reasonable threshold can be taken from Reference [116]. To obtain a pileup contamination level of $\mathcal{O}(1\%)$, a jet threshold of 40 to 80 GeV is required, depending on the $|\eta|$ range considered. However, these values are for a $\langle\mu\rangle = 140$ operating point, and are an extrapolation from $\langle\mu\rangle = 40$ samples. Preliminary studies were thus conducted in order to determine the dependence at intermediate levels of pileup activity after the mono-jet selection.

Note that the 14 TeV samples available at the time included the aforementioned cluster constant bug, where the LC noise thresholds were overwritten by an old set of constants. For these pileup studies, the reconstructed quantities had to be used, as the truth information did not include pileup. It is assumed that the bug is not a critical problem for these pileup studies, given that the calibrated jet response after the mono-jet selection was found to be off by only a few percent on average, and the intent of these studies is only to provide a rough estimate of a reasonable jet p_T threshold necessary to suppress pileup.

An additional limitation of these pileup studies is the restriction of the truth smearing parametrization to the region where $|\eta| < 3.6$. The pileup contamination is largest in the forward $|\eta|$ region, as seen in Reference [116]. This is also the region where the cluster bug is expected to have the largest impact. This limitation should help to avoid the majority of the problems associated with the cluster bug, but also limits the robustness of these studies, which may need to be revisited with fixed samples before Run-II begins.

J.1 Procedure

Jet threshold studies were carried out on the 14 TeV EFT signal samples and the corresponding 8 TeV samples. The 14 TeV samples are full GEANT4 simulation with the cluster bug, while the 8 TeV samples are fast AFII simulation with no known bugs. Event selection and reconstructed to truth jet matching followed the below procedure, all on reconstructed level quantities unless otherwise stated:

1. Apply the part of the mono-jet selection which does not depend on the jet threshold
 - 8 TeV: $E_T^{\text{miss}} > 150 \text{ GeV}$, $p_T^{\text{jet1}} > 120 \text{ GeV}$, $|\eta| < 2.0$
 - 14 TeV: $E_T^{\text{miss}} > 300 \text{ GeV}$, $p_T^{\text{jet1}} > 300 \text{ GeV}$, $|\eta| < 2.0$
2. Match anti- k_t $R = 0.4$ truth jets to anti- k_t $R = 0.4$ LC jets built from reconstructed topo-clusters
 - Require a minimum of $\Delta R < 0.2$ for matching to succeed, and then take the closest match in ΔR
3. Set the jet threshold level and define jet objects
4. Apply the remaining mono-jet selection cuts
 - $N_{\text{jets}} = \{1, 2\}$, $\Delta\phi(E_T^{\text{miss}}, \text{jets}) > 0.5$
5. Define final jet collections for use in the following plots

J.2 Pileup jet multiplicity

The first quantity to check is the average number of pileup jets in an event, defined as the number of reconstructed jets above the p_T threshold which are not matched to any truth jets. Note that this is a suitable criteria due to the truth collection being comprised solely of jets from the hard-scatter event. Once this is determined, a jet threshold which provides similar pileup contamination at 8 and 14 TeV can be found.

First, it must be confirmed that the E_T^{miss} and leading jet p_T cuts do not shape the pileup distributions, as the 8 and 14 TeV analyses have different thresholds. These cuts are not expected to shape the distributions, as pileup is from other collisions and thus should not be dependent on the topology selected, but the comparison was done as a cross-check. Figure J.1 shows that the impact is indeed negligible.

The next step is to compare 8 TeV samples to 14 TeV samples. Recall that the 14 TeV samples were generated at $\langle\mu\rangle = 80$, while the values quoted in Reference [116] are for $\langle\mu\rangle = 140$. The pileup levels should thus be lower. This is what is observed as shown in Figure J.2, except for in the most forward regions, which is most likely due to the impact of the cluster bug. Note that the approximate pileup contamination with the 30 GeV jet threshold at 8 TeV is at the level of $\sim 4\%$. This is the level we wish to match or surpass for the 14 TeV projections. With no additional techniques, it looks like a threshold cut at $\sim 60 \text{ GeV}$ is required to obtain the same performance at 14 TeV and $\langle\mu\rangle = 80$.

Techniques to improve pileup jet rejection within the acceptance of the tracker are under development, and thus it is reasonable to assume that this will be improved in time for Run-II. For example, the JVT technique has become public since this study was published, as detailed in Reference [132]. As a first

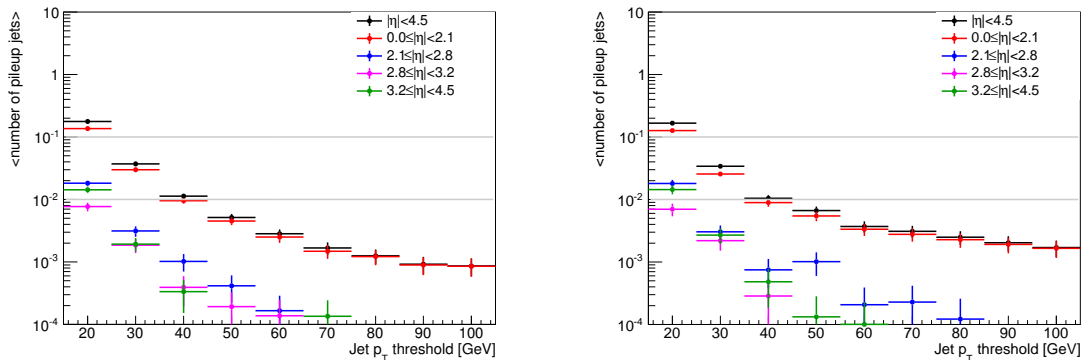


Figure J.1: Comparison of the average number of pileup jets in the 8 TeV EFT signal samples with $E_T^{\text{miss}} > 150 \text{ GeV}$, $p_T^{\text{jet}1} > 120 \text{ GeV}$ (left) and $E_T^{\text{miss}} > 300 \text{ GeV}$, $p_T^{\text{jet}1} > 300 \text{ GeV}$ (right) cuts applied. This demonstrates the independence of pileup with respect to these cuts.




Figure J.2: Comparison of 8 TeV (left) and 14 TeV (right) pileup contamination levels. Values are generally better than what is quoted for $\langle \mu \rangle = 140$ in Reference [116], as expected for $\langle \mu \rangle = 80$ samples. The most forward regions appear to be significantly affected by the cluster bug, but other regions appear to be at about the level expected.

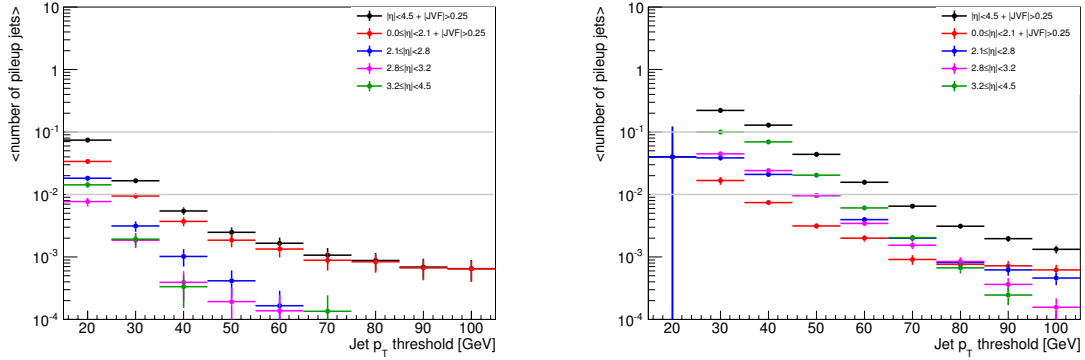


Figure J.3: Comparison of 8 TeV (left) and 14 TeV (right) pileup contamination levels, after applying a $|JVF| > 0.25$ cut to the central region.

approximation, the commonly used JVF technique was used to reject pileup jets. Cut values of 0, 0.25, and 0.5 were tried, where 0.25 was found to be the most reasonable balance between the desired pileup rejection and undesired hard-scatter signal jet rejection. The results of such a $|JVF| > 0.25$ cut can be seen in Figure J.3. Discarding the most forward region to consider only the region allowed by our truth smearing and to reduce the LC cluster bug gives a pileup contamination level of approximately 3% at 40 GeV, which is better than what was used for 8 TeV studies. Given that this is for $\langle\mu\rangle = 80$ samples, a threshold at 50 GeV is a reasonable estimate for a suitably low pileup contamination up to $\langle\mu\rangle = 140$.

In conclusion, it appears that the 8 TeV level of pileup contamination that exists with a 30 GeV jet threshold can be reproduced at 14 TeV up to $\langle\mu\rangle \approx 140$ by using a jet threshold of 50 GeV and additional pileup rejection techniques (such as $|JVF| > 0.25$) in the central region.

J.3 Matching method and JVF cut checks

Given that the studies depend on linking truth to reconstructed jets to determine when a jet is pileup and when it is from the hard-scatter, it is important to check that the matching procedure itself is performing reasonably well. In order to do this, the number of truth jets above the jet threshold which were not matched to a reconstructed jet is checked. Figure J.4 shows that this is at the sub-percent level, confirming that the procedure is working reasonably. For the cut values chosen, the failure rate at 8 TeV is approximately 1%, or about 0.5% at 14 TeV.

The JVF cut value to use is also important to study. Higher cut values provide an increased level of pileup rejection, but at the cost of rejecting more hard-scatter jets. It appears that the increased number of vertices in $\langle\mu\rangle = 80$ samples significantly changes the mis-tag rate of the JVF cut. Figure J.5 shows a comparison between the 14 TeV $\mu = 80$ samples with different JVF cuts applied. It demonstrates that a value of 0.5 kills a large number of hard-scatter jets, while a cut at 0.25 is much more reasonable and provides only a slightly reduced level of pileup rejection. It is hoped that some improved form of track-based pileup suppression will be available for Run-II thus rendering this discussion irrelevant, but for now it's a necessary step to ensure that the threshold used is not too optimistic.

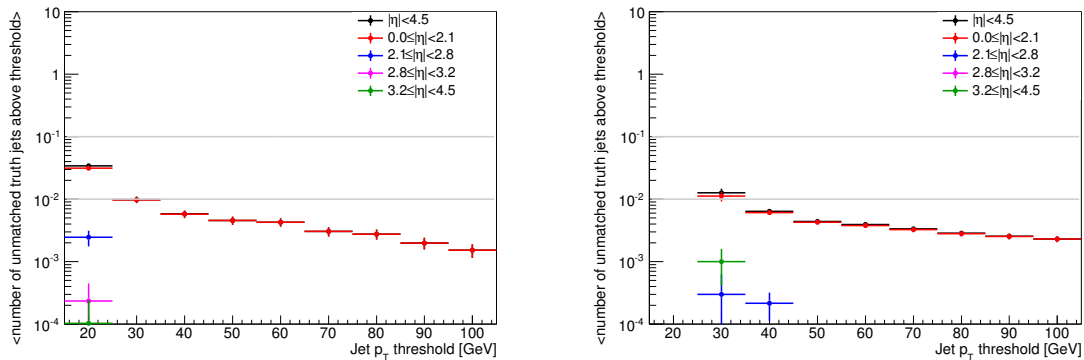


Figure J.4: Comparison of 8 TeV (left) and 14 TeV (right) levels of unmatched hard-scatter truth jets above the jet threshold. Both are seen to be at negligible levels, confirming that the matching procedure works reasonably well.

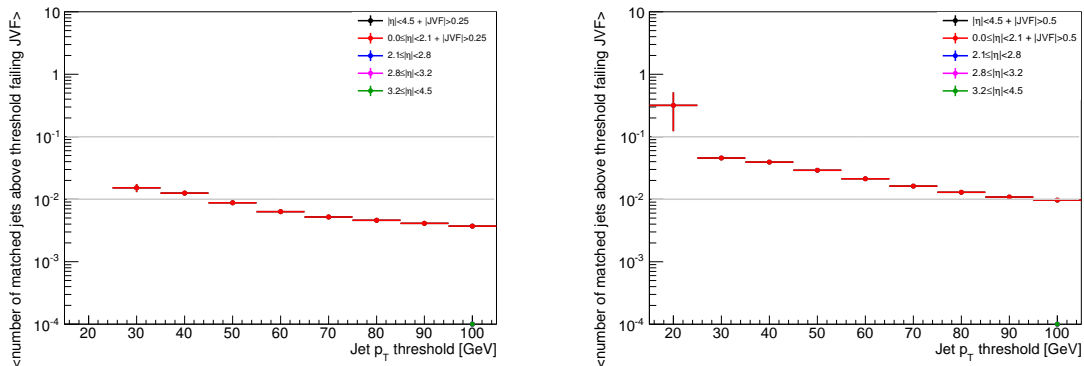


Figure J.5: Comparison of average number of hard-scatter jets rejected in 14 TeV samples with $|JVF| > 0.25$ (left) and $|JVF| > 0.5$ (right) cuts applied. At the chosen jet threshold of 50 GeV, the lower JVF cut rejects $< 1\%$ of hard scatter jets, while the higher JVF cut rejects approximately 3%. This is the reason that a cut at 0.25 was deemed to be more reasonable for these studies, although ultimately it is hoped that JVF will be replaced before Run-II with a more robust track-based pileup suppression scheme, such as JVT [132].