Luminosity studies and a search for heavy resonances decaying into a pair of Z bosons with the ATLAS detector

by

Joseph William Snarr Carter

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy

Department of Physics
University of Toronto

Luminosity studies and a search for heavy resonances decaying into a pair of Z bosons with the ATLAS detector

Joseph William Snarr Carter
Doctor of Philosophy

Department of Physics
University of Toronto
2022

Abstract

Many theories beyond the Standard Model predict additional, heavy Higgs bosons or other resonances. This thesis presents a search for such new states decaying into a pair of Z bosons. Two final states are considered, corresponding to the $X \to ZZ \to \ell^+\ell^-\ell^+\ell^-$ and $X \to ZZ \to \ell^+\ell^-\nu\bar{\nu}$ decay channels, where X stands for the hypothetical heavy resonance and ℓ stands for either an electron or muon, with the focus of this thesis on the $\ell^+\ell^-\ell^+\ell^-$ channel. The search uses proton–proton collision data at a centre-of-mass energy of $\sqrt{s} = 13\,\text{TeV}$ collected between 2015 and 2018 by the ATLAS detector during Run 2 of the Large Hadron Collider, corresponding to an integrated luminosity of $139\,\text{fb}^{-1}$. Different mass ranges spanning from 200 GeV to 2000 GeV for the hypothetical resonances are considered, depending on the final state and model. No significant excess over Standard Model predictions is observed, therefore the results are interpreted as upper limits on the production cross section of a spin-0 or spin-2 resonance under a number of benchmark scenarios. In the case of a spin-0 resonance, the observed upper limits on the production cross section times branching ratio at 95 % confidence level are between 215 fb and 2.0 fb for the gluon–gluon fusion production mode, and between 87 fb and 1.5 fb for the vector-boson fusion production mode. These results are also interpreted in the context of Type-I and Type-II two-Higgs-doublet models, and are used to constrain the Randall–Sundrum model with an extra dimension giving rise to spin-2 Kaluza–Klein graviton excitations, which are excluded up to a mass of 1830 GeV at 95 % confidence level.

This thesis also presents studies of a novel luminosity-monitoring technique used for evaluating the systematic uncertainties of a track-counting luminosity algorithm used in the nominal calibration of the primary ATLAS luminosity measurement provided by the LUCID detector. This new luminosity algorithm, called LAr energy flow, exploits the observation that the average amount of energy deposited in the liquid-argon calorimeters per bunch crossing is proportional to the instantaneous luminosity. The energy-flow luminosity is evaluated in a set of special pp runs in 2017 and 2018, and is compared against track counting to evaluate their relative nonlinearity, with the two algorithms agreeing to within 0.5 % in fills with sufficiently isolated bunches.

Acknowledgements

The most widely recognized reason to pursue a life of physics is undoubtedly to gain some deeper understanding of nature. There is another often underappreciated reason, which is the admiration for the people in the field and a desire to be a part of that enterprise. I would like to express my deepest gratitude to all of those who came before me who laid the foundation for the work contained in this thesis and to the extraordinary community of scientists and engineers who continue this work today. I have walked among the giants, and I am a better scientist and a better human for it.

First, I am indebted to my advisor Peter Krieger for his guidance and long discussions on my work, the history of particle physics and nearly anything else. His unwavering support and encouragement since I arrived in Toronto was invaluable, and for that I am profoundly thankful. I would also like to extend my thanks to my advisory committee members Bob Orr and Michael Luke for following my research all this time, and to the University of Toronto EHEP faculty, especially Pierre Savard, Pekka Sinervo, William Trischuk and Miriam Diamond for the many insightful discussions over the years.

To my fellow ATLAS graduate students, postdocs and engineers at the University of Toronto, it has been a privilege to work and to learn alongside you. I would especially like to thank Aaron, Adriana, Albert, Ben, Bianca, Chav, Chris, Dylan, Garrin, Haider, Joel, Karola, Kyle, Laurelle, Lukas, Matt, Michael, Olivier, Riccardo, Rob, Sahibjeet, Sana, Tae, Tim, Tina, Vince and Yi Fie for their warmth and openness, their inspiration and commiseration, and above all for their friendship. Special thanks are due to Haider Abidi for first setting me down the path of the heavy ZZ search, and to Lukas Adamek for all our adventures in the Alps. Thanks also to Jesse Velay-Vitow for keeping the wind in my sails with many a swashbuckling sea shanty session and for reminding me that you can and *should* do more than just physics during a physics PhD, and to Pierre Savaria for passing on the unwritten lore of CERN over coffee. *Grazie mille* to my *amici italiani* Francesco, Daniele and Sara for brightening our home in Geneva, and for welcoming me into yours after we had gone our separate ways.

I am extremely grateful to RD Schaffer, Giacomo Artoni, Andrea Gabrielli, Gaetano Barone and Will Leight who as the HZZ convenors during this time tactfully saw our analysis from its inception through to publication, and whose steadfast leadership pushed us to achieve the result we did. I would also like to offer my sincerest thanks to the high mass team, in particular to Lailin Xu, Marc Cano Bret, Heling Zhu, Denys Denysiuk, Xiangyang Ju, Rosy Nikolaidou and Ricardo Wölker for their input and assistance, and for their diligence and rigour in every area of their work.

The LAr community felt immediately like home since my first arrival at CERN, and I would especially like to thank Emmanuel Monnier, Luis Hervas, Martin Aleksa, Steffen Staerz, Clement Camincher, Nikiforos Nikiforou, Maddie McKay, Adriana Milic and Chris Hayes for showing me the ropes, offering a helping hand when I got stuck as a new on-call expert, and for the many wonderful evenings sampling the local cuisine and the world-famous LArBQs.

In the luminosity group, my gratitude to Witold Kozanecki and Richard Hawkings knows no bounds. I sincerely thank them for their guidance and mentorship, for sharing their mastery of the subject with me, and for teaching me the value and joy of a good precision measurement. I also thank Guillaume Unal for first introducing me to LAr energy flow and for his expertise and suggestions throughout the project.

I owe a great deal to my former mentors and colleagues at Carleton and back home in Ottawa.

Thanks to Professors Paul Johns and Thomas Koffas for first introducing me to physics research, to David Di Valentino and Ioannis Nomidis for guiding me along as I took my "first steps" in ATLAS, and to Steven Weber and Steven J. Lee for keeping in touch after my move to Toronto. Thanks also to Ms. Fiset for first showing me that physics was something worth studying.

If I have seen farther, it is because my parents Chris and Vicky and my sister Genny have lifted me up to this point. Thank you for always being there, for supporting me as I trotted around between Ottawa, Toronto and Geneva, and most of all for believing in me.

Above all, to Jill, thank you for all of your love and support these past years—your tenderness, humour and patience has meant the world to me. Completing a PhD under normal circumstances is difficult enough, but it would have been downright *inconceivable* during a pandemic had you not been by my side throughout it all. I dedicate this thesis to you.

我爱你胡小鹿

Contents

List of Tables

List of Figures

Personal Contributions to the ATLAS Experiment

The work I present in this thesis was performed as part of the ATLAS Collaboration, and as such represents but one piece of the highly interconnected cumulative effort of several thousand physicists, engineers and technicians, whose decades of planning, construction and commissioning of the experimental apparatus have allowed me to carry out this work. Furthermore, much of this work represents the combined effort of several subgroups within the larger Collaboration that relied on each others' tools and technical expertise in order to carry out our analyses. While I will focus on my contributions to these analyses throughout this thesis, it is impossible to disentangle them fully from the work of my collaborators while still presenting a sensible result. I list below my personal contributions to this body of work, and more broadly to the ATLAS Experiment in general.

The analysis presented in Chapter 6 on the LAr energy-flow luminosity measurement is entirely my own, although relied on the input from many members of the ATLAS luminosity and LAr calorimeter communities. The search for heavy resonances decaying into a pair of Z bosons presented in Chapter 7 constitutes a larger group effort. While decisions pertaining to the analysis strategy were made as a group, each member of the analysis team was responsible for a specific step (or steps) in the analysis chain. My group was uniquely responsible for the analysis of the $ZZ \to \ell^+\ell^-\ell^+\ell^-$ channel; a separate group performed the analysis of the $ZZ \to \ell^+\ell^-\nu\bar{\nu}$ channel. I include a discussion of the $ZZ \to \ell^+\ell^-\nu\bar{\nu}$ analysis in Section 7.4 as it is necessary to explain the final combined results of the two channels. My main responsibilities in the $ZZ \to \ell^+\ell^-\ell^+\ell^-$ analysis were threefold: development and validation of the signal and background modelling, production of the group's analysis Ntuples, and integrating the outputs of other steps in the analysis chain into our statistical analysis framework. First, I developed the narrow-width spin-0 and graviton signal models, determined by fitting analytical distributions to signal Monte Carlo samples, and developed a novel probability distribution function to describe the non-resonant ZZ background, whose free parameters I determined in fits to background ZZ Monte Carlo samples. I also developed and performed closure tests to validate these modelling methodologies and employed a parameter-sampling technique to evaluate the systematic uncertainties associated with the parameterization of the models. Second, I maintained the group's Ntuple production codebase and was responsible for processing the data samples, in which raw collision events from data collected by ATLAS and from Monte Carlo simulation are reconstructed and calibrated, and the events of interest are identified, selected and processed into a format that can be used for the analysis. Third, I developed general analysis-infrastructure software to facilitate the storage, parsing and integration of the analysis parameters and data for input into our statistical analysis framework. I also served as co-editor of the interal support note documenting the full analysis procedure in detail. The development of the neural-network classifier and large-width signal models, and running the statistical framework in combination with the $ZZ \to \ell^+\ell^-\nu\bar{\nu}$ inputs to extract the final results and interpretations were performed by other members of the analysis team.

In addition to the analyses presented in this thesis, I had a number of other responsibilities as a member of the ATLAS Collaboration. These responsibilities have mainly centred on LAr calorimeter operations and luminosity monitoring, two groups in which I have been an active member since the summer of 2017. I spent roughly a year and a half on site at CERN until the end of 2018 and participated in data-taking operations in a number of roles. First, I took a number of shifts in the ATLAS Control Room and oversaw the calorimeter and forward-detector systems to monitor their

status during data-taking and to report any problems to the relevant system experts. Later, I became one of those experts, first as an on-call LAr software expert and then by the end of 2018 as LAr Run Coordinator. My primary responsibilities as an on-call expert and LAr Run Coordinator were to ensure good detector performance during data-taking, coordinate calibration runs and interventions on the system with ATLAS Run Coordination, and to act as a first point of contact in case of problems with the LAr calorimeters. During this time, I also maintained, developed and documented a number of LAr software applications for use during data-taking operations, and later, for detector commissioning. My main contribution in this respect was as lead developer of a new tool with a graphical user interface designed to take so-called "single-HFEC (Half Front-End Crate)" calibration runs. This tool allows for finer granularity in the choice of calorimeter region participating in the calibration run, and it has been used extensively for commissioning of the LAr calorimeters during the long-term shutdown of the LHC following the end of Run 2 while electronics in the front end crates were being refurbished or replaced. Finally, I also contributed to the detector-simulation effort by updating and running a piece of software to compute the FCal sampling fractions, which serve as important inputs to the the ATLAS Fast Simulation program to model shower development in the ATLAS calorimeters.

Throughout my time at CERN and since returning to Toronto, I was the LAr contact to the ATLAS luminosity group, responsible for processing the LAr "gap-current" luminosity measurements using the high-voltage-line currents from the electromagnetic endcap (EMEC) and forward calorimeters (FCal). This luminosity-monitoring technique, first developed in Run 1, is used to study the long-term stability of the primary ATLAS luminosity algorithm. In 2018, I developed a correction to the FCal luminosity to account for a nonlinearity in the detector's response, which was prompted by the FCal being the uncertainty-limiting detector in the 2017 long-term stability analysis as a consequence of the nonlinearity. This effort was ultimately an important factor in reducing the luminosity stability uncertainty in the 2017 and 2018 pp datasets from approximately 1.3 % to 0.8 %. I include a complete summary of my work on the LAr gap-current luminosity measurements in Appendix A.

I list below a selection of publications to which I have contributed directly as an analyst or as co-author, as well as the conference presentations I have given during my time as a graduate student at the University of Toronto.

Selected publications

ATLAS Collaboration, *Measurement of inclusive and differential cross sections in the $H \to ZZ^* \to 4\ell$ decay channel in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, JHEP **10** (2017) 132, arXiv: 1708.02810 [hep-ex].

ATLAS Collaboration, *Search for heavy ZZ resonances in the $\ell^+\ell^-\ell^+\ell^-$ and $\ell^+\ell^-\nu\bar{\nu}$ final states using proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Eur. Phys. J. C **78** (2018) 293, arXiv: 1712.06386 [hep-ex].

ATLAS Collaboration, *Measurement of the Higgs boson coupling properties in the $H \to ZZ^* \to 4\ell$ decay channel at $\sqrt{s} = 13$ TeV with the ATLAS detector*, JHEP **03** (2018) 095, arXiv: 1712.02304 [hep-ex].

ATLAS Collaboration, *Measurement of the Higgs boson mass in the $H \to ZZ^* \to 4\ell$ and $H \to \gamma\gamma$*

channels with $\sqrt{s} = 13$ *TeV pp collisions using the ATLAS detector*, Phys. Lett. B **784** (2018) 345, arXiv: 1806.00242 [hep-ex].

ATLAS Collaboration, *Measurement of the four-lepton invariant mass spectrum in* 13 *TeV proton–proton collisions with the ATLAS detector*, JHEP **04** (2019) 048, arXiv: 1902.05892 [hep-ex].

ATLAS Collaboration, *Search for heavy resonances decaying into a pair of Z bosons in the* $\ell^+\ell^-\ell'^+\ell'^-$ *and* $\ell^+\ell^-\nu\bar{\nu}$ *final states using* $139\,fb^{-1}$ *of proton–proton collisions at* $\sqrt{s} = 13$ *TeV with the ATLAS detector*, Eur. Phys. J. C **81** (2020) 332, arXiv: 2009.14791 [hep-ex].

G. Aad et al., *The Phase-I Trigger Readout Electronics Upgrade of the ATLAS Liquid Argon Calorimeters*, JINST **17** (2022) P05024, arXiv: 2202.07384 [physics.ins-det].

ATLAS Collaboration, *Luminosity determination in pp collisions at* $\sqrt{s} = 13$ *TeV using the ATLAS detector at the LHC*, (2022), in preparation.

Conference presentations

J. W. S. Carter, *Luminosity determination in pp collisions at* $\sqrt{s} = 13$ *TeV using the ATLAS detector at the LHC*, Poster presentation at the 29th International Symposium on Lepton Photon Interactions at High Energies, Toronto, Canada, 2019, URL: https://indico.cern.ch/event/688643/contributions/3457956/.

J. W. S. Carter, *Luminosity measurement with the ATLAS experiment at the LHC*, Talk at the 40th International Conference on High Energy Physics, Prague, Czech Republic, proceedings: 10.22323/1.390.0726, 2020, URL: https://indico.cern.ch/event/868940/contributions/3813703/.

J. W. S. Carter, *Searches for Higgs-like or diboson resonances with the ATLAS detector*, Talk at the XXVIII International Workshop on Deep-Inelastic Scattering and Related Subjects, 2021, URL: https://indico.bnl.gov/event/9726/contributions/46367/.

J. W. S. Carter, *Search for heavy resonances decaying into a pair of Z bosons with the ATLAS detector*, Talk at the 2021 Canadian Association of Physicists Congress, 2021, URL: https://indico.cern.ch/event/985448/contributions/4295772/.

Chapter 1

Introduction

One of the preeminent goals of fundamental physics is to formulate a complete mathematical description of the elementary constituents of matter and their interactions. Humanity's understanding of the physical laws characterizing these interactions has progressed immensely over the past several hundred years, from Isaac Newton's laws of motion and of universal gravitation, to James Clerk Maxwell's "unification" of electricity and magnetism, to the discovery of quantum mechanics and relativity in the early twentieth century. The pursuit of an understanding of the Universe at its smallest scales has led scientists to a description of nature in which matter is composed of smaller *fundamental* particles whose interactions with one another are mediated by a number of fundamental forces. More recently, physicists have explored these questions through the study of high-energy particle collisions, which has revealed an elegant structure and symmetry of nature at this scale, hidden at the normal "low" energies that pervade everyday human life. The extensive experimental and theoretical work since the mid 1950s has culminated in the *Standard Model of Particle Physics*, a remarkably successful and precise mathematical description of the elementary particles of our Universe and their interactions at high energy [1]. A timeline of the Standard Model particle discoveries is shown in Figure 1.1.

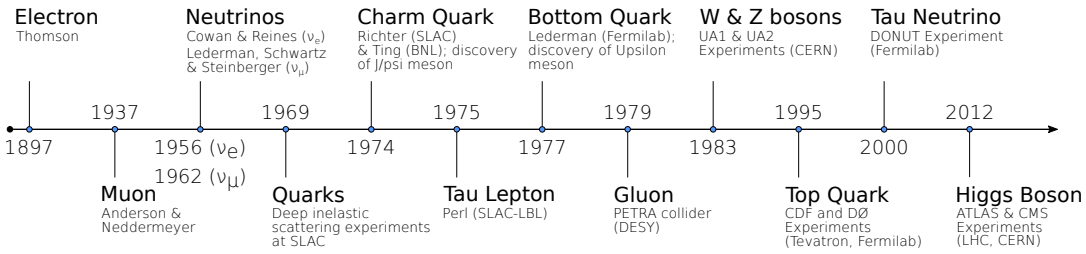

Figure 1.1: Timeline of Standard Model particle discoveries.

The need for experimentation at ever-increasing energies has led to the construction of larger and larger particle accelerators and colliders, with a corresponding increase in size, complexity and cost of the associated detectors required to observe the particle interactions. The Large Hadron Collider (LHC), operated by the European Organization for Nuclear Research (CERN[1]), is the world's largest

[1]The acronym *CERN* is derived from the French *Conseil Européen pour la Recherche Nucléaire*.

and most powerful particle collider, able to reach the highest energies currently accessible and therefore probe the constituents of matter at the smallest scale. The LHC accelerates beams of protons or heavy ions and puts them into collisions in the bellies of four particle detectors located around the LHC ring: ATLAS, CMS, ALICE and LHCb. ATLAS (**A** **T**oroidal **L**HC **A**pparatu**S**) and CMS (**C**ompact **M**uon **S**olenoid) are general-purpose detectors designed for broad physics programs. ALICE (**A** **L**arge **I**on **C**ollider **E**xperiment) and LHCb (**LHC** **b**eauty) are specialized detectors optimized primarily for studies of quark–gluon plasma in heavy-ion collisions and heavy-flavour ("charm" and "beauty") physics, respectively.

No significant deviation from the Standard Model has yet been observed, and it might appear that since the discovery of the Higgs boson in 2012 by the ATLAS and CMS Collaborations that particle physics has uncovered all it can. However, there are compelling theoretical reasons, affirmed by indirect experimental evidence, to suggest that nature is considerably more complicated than the Standard Model purports. The primary focus of this work is to report on a search for physics beyond the Standard Model using data collected by the ATLAS Experiment during LHC Run 2 in the years 2015–2018.

This thesis is organized as follows. Chapter 2 provides historical context and a discussion of the theory of the Standard Model, its shortcomings and proposed extensions. Chapter 3 gives an overview of the LHC and the ATLAS detector and some history of its operation, from its inception to the end of LHC Run 2. Chapter 4 describes the anatomy of proton–proton collisions and the current methods used for event and detector simulations. Chapter 5 provides an overview of the Run 2 ATLAS dataset and the method employed by ATLAS to reconstruct the particles traversing the detector. Chapter 6 describes the nominal ATLAS luminosity determination and studies of a new luminosity measurement based on the total reconstructed energy in the ATLAS liquid-argon calorimeters. Chapter 7 reports on the search for heavy resonances decaying to a pair of Z bosons and comprises the main work of this thesis. Finally, concluding remarks and an outlook for future studies are given in Chapter 8.

Chapter 2

Theory and Historical Context

2.1 Introduction

The history of particle physics is one of experimental discovery and theoretical model building. This chapter will first present an overview of the Standard Model,[1] both in descriptive terms and from a historical perspective, highlighting the theoretical and experimental discoveries that have led us to our current understanding of the fundamental particles in nature and their interactions at high energy. The Standard Model is a *predictive* theory, therefore some time will be given to expound its mathematical structure, emphasizing the elements of the theory relevant to this thesis. The Standard Model is not a complete theory of nature and a number of its noteworthy shortcomings will be discussed at the end of the chapter, followed by a discussion of the selected extensions to the Standard Model that were investigated experimentally for this thesis.

2.2 The Standard Model of Particle Physics

The Standard Model posits that the matter in the Universe is composed of elementary fermions whose interactions are mediated by the exchange of elementary force-carrying gauge bosons. Each particle in the Standard Model is treated as an excited state, or *quantum*, of an underlying quantum *field*.

Four types of interaction fields have been identified in nature, often referred to as the four *Fundamental Forces*, summarized in Table 2.1. On the scales of contemporary particle physics, the gravitational force is insignificant and the Standard Model excludes the gravitational field from consideration. The associated gauge boson of the electromagnetic field mediating the interaction between electrically charged particles is the massless photon γ. The weak interaction is responsible for such things as the radioactive decay of atoms, and the quanta of the weak fields are the charged W^+ and W^- bosons and neutral Z boson. The strong interaction binds protons and neutrons together in the nuclei of atoms, and the quanta of the strong fields are the eight massless gluons g.

The elementary fermions of the Standard Model have spin $\frac{1}{2}$ and are of two types: *leptons*, which interact only through the electromagnetic interaction (if they are charged) and the weak interaction,

[1] I will occasionally abbreviate "Standard Model" to "SM" where appropriate. Similarly, theories that predict physics beyond the Standard Model will be labelled as "BSM".

Table 2.1: The interaction fields and their associated gauge bosons. The electromagnetic, weak and strong fields are described in the Standard Model, while the gravitational field is not.

Interaction Field	Gauge boson(s)	Spin	Relative strength	Range [m]
Strong	8 gluons; g	1	1	10^{-15}
Electromagnetic	photon; γ	1	10^{-3}	∞
Weak	weak bosons; $W^{\pm}, Z$	1	10^{-8}	10^{-18}
Gravitational	"graviton" (postulated); G	2	10^{-37}	∞

and *quarks*, which interact through the electromagnetic, weak and strong interactions. Antiparticle partners for all the fermions also exist, which have the same mass and spin but opposite electric charge, as predicted by the Dirac equation for massive spin-$\frac{1}{2}$ particles; such particles with distinct antiparticle partners are called *Dirac fermions*. The leptons and quarks are further categorized into three *generations*, each of which follows a doublet structure of two particles, as shown in Figure 2.1, with the species of a lepton or quark referred to as its *flavour*.

The three generations of leptons are the electron (e), muon (μ) and tau (τ) particles, and their associated neutrinos: ν_e, ν_μ and ν_τ. The e, μ and τ leptons have charge $-1e$ and the neutrinos are neutral. Charged leptons and neutrinos interact through the weak force, which allows leptons to change flavour, as in the decay of a muon:

$$\mu^- \to \nu_\mu + e^- + \bar{\nu}_e \,.$$

The lepton generations follow a mass hierarchy where $m_e < m_\mu < m_\tau$, though it has not yet been confirmed if the same holds true for the corresponding neutrinos [2, 3]. It is also not yet known if neutrinos are in fact Dirac fermions as the Standard Model describes them, or if they are instead *Majorana fermions*, in which case each neutrino is its own antiparticle. A number of experiments [4–7] are currently underway in search of *neutrinoless double beta decay*—a yet-unobserved signature process that would reveal neutrinos to be Majorana fermions—to address this question.

Similarly, the quarks are arranged into three generations, each containing a doublet of *up-type* and *down-type* quarks. Unlike leptons, which carry an integer electric charge, quarks carry a fractional electric charge. The up-type quarks are the up (u), charm (c) and top (t) quarks and have charge $+\frac{2}{3}e$, and the down-type quarks are the down (d), strange (s) and bottom (b) quarks and have charge $-\frac{1}{3}e$. Like the different lepton flavours, it is possible for an up-type quark to transform into a down-type quark through the weak interaction, and vice versa. Quarks also differ from leptons in that they carry an additional "charge" that participates in the strong interaction, called *colour charge*. Each flavour of quark occurs as a colour triplet of three colour states, typically labelled red, green and blue,[2] and the antiparticle partners of quarks carry anticolour charge (antired, antigreen and antiblue). Analogous to electric charge, the sum of colour and anticolour charge produces a neutral "colourless" state, and the same is also true of the sum of three colours, red+green+blue and anitred+antigreen+antiblue. Colour-neutral bound states of three quarks (qqq) or three antiquarks ($\bar{q}\bar{q}\bar{q}$) are known as *baryons*, and bound states of a single quark and single antiquark ($q\bar{q}$) are known as *mesons*. Baryons and mesons are collectively known as *hadrons*. The gluons themselves also

[2]These colour labels became popular because of the loose analogy to the three primary colours, but should not be confused with the familiar colours of light.

carry colour charge (specifically colour–anticolour charge), so are able to interact with other gluons. Furthermore, colour-charged particles do not exist in isolation at low energies and are said to be *confined* in a colourless bound state [8]. If a quark is liberated from a hadron—during a high-energy collision, for example—the strong interaction is so great that gluons radiating from the quark produce cascades of new quark–antiquark pairs, which then combine to form stable, colour-neutral particles in a process known as *hadronization*. This cascade of strongly interacting particles is called a *jet*.

Investigations of deep-inelastic scattering of leptons from nucleons at the Stanford Linear Accelerator Center (SLAC) established the existence of quarks experimentally by providing evidence of point-like, spin-$\frac{1}{2}$ constituents inside the nucleon [9, 10]. Similarly, the existence of colour charge is supported by a number of experimental observations. For instance, theoretical predictions of observables that depend on the number of quarks participating in an interaction, such as the ratio $\sigma(e^+e^- \to \text{hadrons})/\sigma(e^+e^- \to \mu^+\mu^-)$ and the decay of neutral pions to two photons $\pi^0 \to \gamma\gamma$, which proceeds through a quark loop, disagree with experiment unless each quark is considered as a member of a triplet with three distinct colour states. Finally, experiments conducted at the PETRA collider established the existence of gluons through the observation of three-jet events in which gluons were produced via gluon bremsstrahlung in electron–positron annihilation: $e^+e^- \to q\bar{q}g$ [11–14].

The $W^\pm$ and Z bosons of the weak interaction were considerably more difficult to detect experimentally due to their large mass. Fermi's original theory of beta decay via the weak interaction treated the process as a point-like, four-fermion interaction [15], however shortly thereafter it was noted that the probability of such an interaction occurring diverges at high energies. The solution to this unphysical behaviour was given through the introduction of massive intermediate vector bosons, which were eventually identified as the $W^\pm$ bosons in the unified electroweak theory of Glashow [16], Weinberg [17] and Salam [18]. Their theory predicted an additional, neutral Z boson, which is responsible for *neutral-current* interactions such as the elastic scattering of muon neutrinos from electrons, $\nu_\mu/\bar{\nu}_\mu + e^- \to \nu_\mu/\bar{\nu}_\mu + e^-$, and of neutrinos from nucleons, $\nu/\bar{\nu} + p/n \to \nu/\bar{\nu} + \text{hadrons}$. These neutral-current interactions were first observed experimentally in 1973 at the Gargamelle bubble chamber at CERN [19, 20], however it was not for another decade that the $W^\pm$ and Z bosons were observed directly by the UA1 [21, 22] and UA2 [23, 24] experiments in $p\bar{p}$ collisions at the CERN Super Proton Synchrotron, firmly establishing electroweak theory in the Standard Model.

The final elementary particle of the Standard Model is the scalar (spin-0) Higgs boson. The associated *Higgs field* plays an important role in the Standard Model in which the other elementary particles acquire their mass through a process of spontaneous symmetry breaking. The Higgs boson is of important relevance to this thesis, so will be discussed in greater detail in the sections below.

2.2.1 The mathematical formulation of the Standard Model

The Standard Model is a theory of interacting fields that is consistent with both quantum mechanics and special relatively, and its mathematical formalism is described by *Quantum Field Theory*.

In non-relativistic quantum mechanics, single particles are described as wavefunctions that satisfy the appropriate wave equation. In quantum field theory, particles are described as excitations of a quantum field that satisfies the appropriate quantum mechanical field equations. The dynamics of a quantum field theory can be expressed in terms of the Lagrangian density $\mathcal{L}$ of a number of

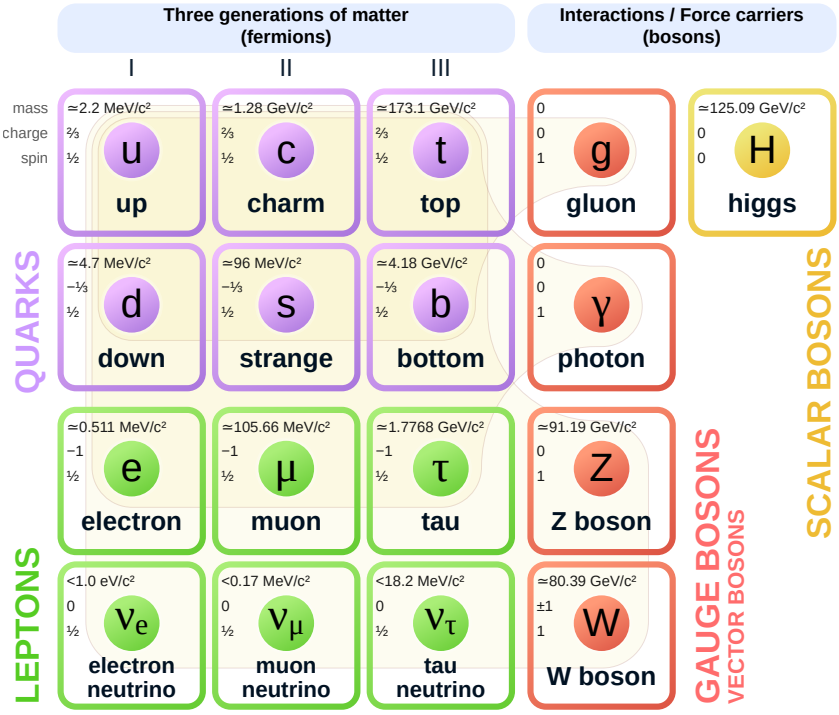

Figure 2.1: The elementary particles of the Standard Model. The shaded areas indicate the interactions between the fermions and the corresponding force-carrying gauge bosons. Adapted from Ref. [25].

continuous fields $\phi_i(x^\mu)$ given as a function of their space-time coordinates $x^\mu = (t, x, y, z)$.[3] The integral over all space gives the Lagrangian L itself:

$$L = \int \mathcal{L}(\phi_i, \partial_\mu \phi_i)\, d^3\mathbf{x}\,, \tag{2.1}$$

where ∂_μ is the four-gradient, defined as

$$\partial_\mu \equiv \frac{\partial}{\partial x^\mu} = \left(\frac{\partial}{\partial t}, \frac{\partial}{\partial x}, \frac{\partial}{\partial y}, \frac{\partial}{\partial z}\right). \tag{2.2}$$

Interactions between fields are introduced to the theory by adding the appropriate interaction terms to the free-field Lagrangian. The interacting fields are said to be *coupled*, and the strength of the coupling is given by a *coupling constant* g. If g is small, the interaction term can be treated as a small perturbation of the free-field theory and the equations of motion can be solved as a power series in g using perturbation theory.

This Lagrangian formalism encodes the rules from which scattering cross sections and particle decay rates are calculated. The *cross section* σ is the measure of the probability of a given interaction occurring during the collision of two (or more) particles and can be thought of as the effective area in the plane transverse to the particle velocities in which these particles will meet to undergo a

[3] The notation adopted here is to use Latin indices for three-vectors, e.g. $\mathbf{x} = x^i = (x^1, x^2, x^3)$, and Greek indices for four-vectors, e.g. $x^\mu = (t, \mathbf{x}) = (x^0, x^1, x^2, x^3)$, where x^0 is the time component of the four-vector. Vector and tensor indices appearing twice in an expression are summed over following the Einstein summation convention. Natural units will be used, in which $\hbar = c = 1$. In this system, [energy] = [momentum] = [mass] = [time]$^{-1}$ = [length]$^{-1}$.

scattering interaction. Accordingly, cross sections have dimensions of area and are generally written in units of *barns*, where $1\,\mathrm{b} = 10^{-28}\,\mathrm{m}^2$. Similarly, the *decay rate* Γ is the measure of the probability of an unstable particle to decay to a final state of two (or more) particles. The decay rate is related to the lifetime of the particle in its rest frame τ according to $\Gamma = 1/\tau$, and has dimensions of energy. From the uncertainty relation $\Delta E \Delta\tau \geq \hbar/2$, the uncertainty on the lifetime of an unstable particle may therefore be re-expressed as a "width" around its central mass value, thus Γ is often referred to as the particle *decay width*. Since an unstable particle can decay to a number of distinct final states, the relative frequency of a particular decay mode j is expressed as a *branching ratio* $BR(j) = \Gamma_j/\Gamma$, where here Γ is the *total* decay rate and by construction, the branching ratios for all decay modes sum to 1. Both the cross section and decay rate are proportional to their respective scattering and decay *matrix elements* $|\mathcal{M}|^2$, where $\mathcal{M}$ is the corresponding scattering or decay *amplitude*. Since the interactions between fields are introduced via perturbation theory, computations of the scattering and decay amplitudes involving the fewest interaction vertices are said to be at *leading order* or *tree level*. Expanding the interaction terms in higher powers of the coupling constant introduces corrections to the leading-order calculation, which are said to be *next-to-leading order* (NLO), *next-to-next-to-leading order* (NNLO), and so on. The interaction terms can be expressed graphically using *Feynman diagrams*, pictorial aids to represent the components of the interaction amplitudes; these diagrams will be used to illustrate scattering and decay processes throughout this thesis. The mathematical procedure to perform such calculations is well documented in many particle physics and quantum field theory textbooks [26–28], and only a discussion of the underlying quantum field theories will be given in the following sections.

Quantum Electrodynamics

Quantum electrodynamics (QED) is the quantum field theory describing the electromagnetic force and the interactions of charged particles mediated by the massless photon γ. The theory was first formulated by Tomonaga [29], Schwinger [30, 31], Feynman [32–34] and Dyson [35, 36] in the late 1940s following the work of Dirac who described the quantization of the electromagnetic field as an ensemble of harmonic oscillators via the introduction of the particle *creation and annihilation operators* [37], and who later derived the equations governing the dynamics of fermions, formulated as a relativistic wave equation [38, 39]. This equation is known as the *Dirac equation*, and for a free fermion it reads

$$(i\gamma^\mu \partial_\mu - m)\psi = 0\,, \tag{2.3}$$

where γ^μ are the Dirac matrices, ψ is the Dirac spinor for the fermion and m is the fermion mass.

Expressed in Lagrangian form, the Dirac equation can be written as

$$\mathcal{L}_{\mathrm{Dirac}} = \overline{\psi}(i\gamma^\mu \partial_\mu - m)\psi\,, \tag{2.4}$$

where ψ is now the Dirac spinor for the fermion *field* and $\overline{\psi} \equiv \psi^\dagger\gamma^0$ is the corresponding adjoint spinor. This Lagrangian is invariant under the transformation

$$\psi(x) \to \psi'(x) = e^{i\alpha}\psi(x)\,, \tag{2.5}$$

where α is a constant phase. These transformations form the unitary group of degree 1, denoted $U(1)$,

corresponding to the symmetry of rotation about a unit circle in the complex plane. Furthermore, since α is a constant, such a transformation is the same at every point in space-time, therefore it is said to be a *global* $U(1)$ transformation.

Interactions between charged fermions and the electromagnetic fields described by Maxwell's equations can be introduced by writing the Lagrangian as follows:

$$\mathcal{L}_{\text{QED}} = \underbrace{-\frac{1}{4}F_{\mu\nu}F^{\mu\nu}}_{\text{Maxwell}} + \underbrace{\overline{\psi}(i\gamma^\mu\partial_\mu - m)\psi}_{\text{Dirac}} + \underbrace{eQ_e\overline{\psi}\gamma^\mu\psi A_\mu}_{\text{interaction}}, \tag{2.6}$$

where A_μ is the electromagnetic four-vector potential, $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the electromagnetic field tensor, and eQ_e is the electric charge of the fermion field, normalized so that $Q_e = -1$ for electrons. This Lagrangian is also invariant under the global $U(1)$ transformation, and by Noether's theorem [40], this symmetry yields a conserved current $j^\mu = \overline{\psi}\gamma^\mu\psi$, which satisfies the continuity condition $\partial_\mu j^\mu = 0$ and thus yields the familiar charge-density continuity equation of classical electromagnetism. Another crucial property of this Lagrangian is that it is invariant under the following *gauge transformations:*

$$\psi(x) \to \psi'(x) = e^{i\alpha(x)}\psi(x), \quad \text{and} \quad A_\mu(x) \to A'_\mu(x) = A_\mu(x) + \frac{1}{eQ_e}\partial_\mu\alpha(x)\,. \tag{2.7}$$

Since the phase $\alpha(x)$ is now a function of space-time, these transformations are said to be *local*. This local gauge transformation prompts the definition of the *gauge covariant derivative*

$$D_\mu \equiv \partial_\mu - ieQ_eA_\mu(x)\,. \tag{2.8}$$

From this definition, the QED Lagrangian can be written more compactly as

$$\mathcal{L}_{\text{QED}} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \overline{\psi}(i\gamma^\mu D_\mu - m)\psi\,. \tag{2.9}$$

Unlike the global $U(1)$ symmetry of the QED Lagrangian, which yields a conserved current, gauge invariance is not physical and does not reveal a symmetry of nature. One must choose a gauge to compute observables, but one is free to choose whichever gauge is most convenient for the problem at hand.

We can also approach the introduction of the electromagnetic interaction into the theory from another perspective. Instead of having the property of local gauge invariance as a consequence of the interaction term in the QED Lagrangian, one might instead begin with the free-field Lagrangian and *impose* local gauge invariance on the theory. From this point of view, we are forced into introducing the gauge field A_μ with the transformation property $A_\mu(x) \to A'_\mu(x) = A_\mu(x) + \frac{1}{eQ_e}\partial_\mu\alpha(x)$ in order to cancel out the additional terms that arise, and the electromagnetic interaction appears as a consequence of the invariance of the Lagrangian under a local gauge transformation with the appropriate choice of covariant derivative. Requiring local gauge invariance in QED has strong theoretical motivations, particularly in its utility to provide a quantum field theory that is *renormalizable* [41], and in which massless spin-1 fields, such as photons, have the required two polarizations. The following sections will show how imposing other local symmetries on a free-particle Lagrangian can naturally give rise to the strong and weak interactions.

Yang–Mills theory

Yang–Mills theory [42] is a generalization of QED in which multiple massless spin-1 particles can interact among themselves. Just as there was a strong motivation to preserve local gauge invariance in the Lagrangian description of QED, the Lagrangians for Yang–Mills theories are similarly constrained but generalized to include symmetry groups whose generators do not commute. Specifically, Yang–Mills theories that are invariant under local $SU(N)$ transformations will be considered, where $SU(N)$ is the special unitary group of $N \times N$ unitary matrices with determinant 1.

First, consider a multiplet of N Dirac spinor fields $\psi_1, \psi_2, \ldots, \psi_N$, each with mass m, whose dynamics follow from the free-field Dirac Lagrangian of Eq. (2.4). The generalized Yang–Mills theory postulates that the free-field Lagrangian is invariant under a local $SU(N)$ gauge symmetry as follows:

$$\psi(x) \to \psi'(x) = U(x)\psi(x)\,, \quad \text{with} \quad U(x) = e^{i\alpha^a(x)t^a}\,, \tag{2.10}$$

where $U(x)$ is an element of the $SU(N)$ group and t^a are the generators of the symmetry group represented by $N \times N$ traceless, Hermitian matrices. The index a runs from 1 to $N^2 - 1$ and the sum over a in this notation is implied. By convention, the matrices t^a are normalized such that

$$[t^a, t^b] = if^{abc}t^c\,, \tag{2.11}$$

where f^{abc} is a set of numbers called *structure constants*. In the *fundamental representation*, the generators for $SU(2)$ are the three Pauli matrices, $t^a = \frac{1}{2}\sigma^a$, and for $SU(3)$ the generators are the eight Gell-Mann matrices, $t^a = \frac{1}{2}\lambda^a$. Field theories generated from non-commuting gauge symmetries are called *non-Abelian*, while the commuting symmetries in which $f^{abc} = 0$, like the $U(1)$ gauge symmetry of QED, are called *Abelian*.

To impose the required $SU(N)$ gauge symmetry on the free-field Lagrangian, we promote the original derivatives to the gauge covariant derivatives defined by

$$D_\mu = \partial_\mu - igA^a_\mu t^a\,, \tag{2.12}$$

which contains one vector field for each independent generator of the local symmetry, where g is the coupling constant. The fermion fields are then required to transform as

$$\psi(x) \to U(x)\psi \simeq [1 + i\alpha^a(x)t^a]\psi\,, \tag{2.13}$$

with the infinitesimal expansion $U(x) = e^{i\alpha^a(x)t^a} \simeq 1 + i\alpha^a(x)t^a$.[4] Similarly, the gauge fields are required to transform as

$$A^a_\mu(x)t^a \to U(x)[A^a_\mu(x)t^a]U^\dagger(x) - \frac{i}{g}[\partial_\mu U(x)]U^\dagger(x)\,, \tag{2.14}$$

or in infinitesimal form:

$$A^a_\mu(x) \to A^a_\mu(x) + \frac{1}{g}\partial_\mu\alpha^a(x) - f^{abc}\alpha^b(x)A^c_\mu(x)\,. \tag{2.15}$$

[4] We use the infinitesimal expansion here since terms involving the derivative $\partial_\mu U(x)$ are not well defined, as the exponent does not necessarily commute with its derivative.

The gauge-invariant kinetic energy term for the fields A^a_μ is constructed from the field strength tensors $F^a_{\mu\nu}$, which are modified from the Abelian case to be

$$F^a_{\mu\nu} = \partial_\mu A^a_\nu - \partial_\nu A^a_\mu + g f^{abc} A^b_\mu A^c_\nu \,, \tag{2.16}$$

and by writing the Lagrangian as

$$\mathcal{L} = -\frac{1}{2}\mathrm{tr}\left[\left(F^a_{\mu\nu} t^a\right)\left(F^{a\mu\nu} t^a\right)\right] = -\frac{1}{4} F^a_{\mu\nu} F^{a\mu\nu} \,. \tag{2.17}$$

Inserting the covariant derivative into the free-field Lagrangian of Eq. (2.4) and combining with Eq. (2.17) gives the Yang–Mills Lagrangian:

$$\mathcal{L}_{\mathrm{YM}} = -\frac{1}{4} F^a_{\mu\nu} F^{a\mu\nu} + \overline{\psi}(i\gamma^\mu D_\mu - m)\psi \,. \tag{2.18}$$

Expanding this Lagrangian reveals that it contains terms cubic and quartic in A^a_μ, implying interactions among the vector fields. This is an important feature of Yang–Mills theory as no such interactions exist for Abelian gauge theories.

Quantum Chromodynamics

Quantum Chromodynamics (QCD) is the quantum field theory that describes the strong interaction between quarks mediated by the massless gluons g. QCD evolved from the *eightfold way* organization of hadrons developed Gell-Mann [43] and Ne'eman [44], and subsequently the quark model of Gell-Mann [45] and Zweig [46]. The three colour charges of the strong interaction naturally leads to an $SU(3)$ Yang–Mills theory, upon which QCD is based. The field strength tensors of the Yang–Mills Lagrangian Eq. (2.18) are replaced with the corresponding gluon field strength tensors $G^a_{\mu\nu}$, where the vector fields A^a_μ now represent the gluon fields. Including the sum over the six quark flavours f with mass m_f, it follows that the QCD Lagrangian is

$$\mathcal{L}_{\mathrm{QCD}} = -\frac{1}{4} G^a_{\mu\nu} G^{a\mu\nu} + \sum_{f=1}^{6} \overline{q}_f (i\gamma^\mu D_\mu - m_f) q_f \,, \tag{2.19}$$

where q_f are the colour triplets of Dirac spinors for each quark flavour and the covariant derivative is $D_\mu = \partial_\mu - i g_s A^a_\mu \frac{\lambda^a}{2}$. The eight generators of $SU(3)$ give rise to the eight massless gluons and the Lagrangian includes terms for the cubic and quartic gluon self interactions arising from the non-commuting nature of the generators discussed in the previous section.

The non-Abelian nature of the QCD gauge fields results in the coupling strength of the strong interaction decreasing at high energy, a behaviour known as asymptotic freedom [47, 48]. The "running" of the QCD coupling constant $\alpha_s = g_s^2/4\pi$ at lowest order takes the form

$$\alpha_s(Q^2) = \frac{12\pi}{(11N_C - 2N_f)\ln\frac{Q^2}{\Lambda}} \,, \tag{2.20}$$

where Q is the four-momentum transfer of the interaction, $N_C = 3$ is the number of colour charges, $N_f = 6$ is the number of quark flavours and Λ is an integration cutoff, typically normalized at the scale of the Z boson mass m_Z for which $\alpha_s(m_Z^2) \simeq 0.12$ [1]. At high energies, the QCD

coupling is sufficiently small such that perturbation theory is reliable, and explains why the quarks observed indirectly in deep-inelastic scattering appear to act as free particles. However, at low energies typically below $O(1\,\text{GeV})$, the coupling constant becomes large ($\alpha_s \sim 1$) and perturbative calculations cannot be made.

The weak interaction and electroweak unification

In 1956, Lee and Yang questioned whether the laws of physics were invariant under parity transformation[5] as they had long been assumed to be [49]. The following year, Wu showed experimentally that parity invariance was violated in the weak interaction through studies of β-decay of cobalt-60, ${}^{60}\text{Co} \rightarrow {}^{60}\text{Ni}^* + e^- + \bar{\nu}_e$, in which the electron decay products were preferentially emitted in the direction opposite the nuclear spin of the cobalt-60 atom [50]. Applying the parity transformation to this scenario would lead to one where the electrons were instead preferentially emitted in the *same* direction as the nuclear spin, which was not observed.

The observation of parity violation led physicists to a description of the weak interaction as a *chiral gauge theory*. In this description, a spinor field ψ can be decomposed in terms of its left- and right-handed chiral eigenstates ψ_L and ψ_R using the chiral projection operators as

$$\psi_L = \frac{(1-\gamma^5)}{2}\psi\,, \quad \text{and} \quad \psi_R = \frac{(1+\gamma^5)}{2}\psi\,, \tag{2.21}$$

where γ^5 is the product of the five Dirac matrices: $\gamma^5 = i\gamma^0\gamma^1\gamma^2\gamma^3\gamma^4$. The weak charged-current interaction is defined by an $SU(2)$ symmetry and couples only to the left-handed states of the Standard Model fermions, while the right-handed states do not participate in the interaction. The charged-current interaction is flavour-changing and couples together fermions that differ by one unit of electric charge. The left-handed fermion states are therefore organized into $SU(2)$ doublets and the right-handed states into singlets uncharged under the weak interaction. There are six such left-handed doublets, one for each lepton generation:

$$\begin{pmatrix}\nu_e \\ e\end{pmatrix}_L, \begin{pmatrix}\nu_\mu \\ \mu\end{pmatrix}_L, \begin{pmatrix}\nu_\tau \\ \tau\end{pmatrix}_L,$$

and one for each generation of quarks:

$$\begin{pmatrix}u \\ d'\end{pmatrix}_L, \begin{pmatrix}c \\ s'\end{pmatrix}_L, \begin{pmatrix}t \\ b'\end{pmatrix}_L.$$

The primes on the down-type quark fields arise because the mass eigenstates of the quarks are not the same as the weak-interaction eigenstates, but are related by the transformation

$$\begin{pmatrix}d' \\ s' \\ b'\end{pmatrix} = \begin{pmatrix}V_{ud} & V_{cd} & V_{td} \\ V_{us} & V_{cs} & V_{ts} \\ V_{ut} & V_{cb} & V_{tb}\end{pmatrix}\begin{pmatrix}d \\ s \\ b\end{pmatrix}. \tag{2.22}$$

This transformation matrix is known as the Cabibbo–Kobayashi–Maskawa or CKM matrix [51]. For antiparticles, the weak interaction couples only to the right-handed chiral states. The "charge"

[5] The parity transformation P inverts a quantum state's three spatial components as $P\psi(x,y,z) = \psi(-x,-y,-z)$.

associated to the weak doublets is called *weak isospin* I_W, which has three components corresponding to the three generators of the $SU(2)$ symmetry group. The left-handed fermion doublets have total weak isopsin $I_W = \frac{1}{2}$, with the upper components (u, c, t and the neutrinos) having $I_W^{(3)} = +\frac{1}{2}$ and the lower components (d, s, b and the charged leptons) having $I_W^{(3)} = -\frac{1}{2}$.

The Glashow–Weinberg–Salam unified electroweak theory postulates that the electromagnetic and weak interactions are different manifestations of the same underlying force. The theory is constructed from an $SU(2)_L \times U(1)_Y$ gauge symmetry, where $SU(2)_L$ indicates the non-Abelian weak-isospin coupling of left-handed states and $U(1)_Y$ refers to the Abelian *weak-hypercharge* coupling to both left- and right-chiral states. By imposing this gauge symmetry, the $U(1)_Y$ requirement introduces a new gauge field B_μ and the $SU(2)_L$ requirement introduces three new gauge fields W_μ^a via the covariant derivative

$$D_\mu = \partial_\mu + ig_W W_\mu^a \frac{\sigma^a}{2} + ig' \frac{Y}{2} B_\mu \,, \tag{2.23}$$

where σ^a are the Pauli matrices, and g_W and g' are the $SU(2)_L$ and $U(1)_Y$ coupling constants, respectively. Defining the field strength tensors

$$W_{\mu\nu}^a = \partial_\mu W_\nu^a - \partial_\nu W_\mu^a - g_W \epsilon_{abc} W_\mu^b W_\nu^c \,, \tag{2.24}$$

where $a = 1, 2, 3$ and the Levi-Civita symbol ϵ_{abc} encodes the structure constants of the $SU(2)$ symmetry, and

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu \,, \tag{2.25}$$

the unified electroweak Lagrangian takes the form

$$\mathcal{L}_{\text{EW}} = -\frac{1}{4} W_{\mu\nu}^a W^{a\mu\nu} - \frac{1}{4} B_{\mu\nu} B^{\mu\nu} + i\overline{\psi}\gamma^\mu D_\mu \psi \,. \tag{2.26}$$

Expanding this Lagrangian reveals the fermion–gauge-boson interactions:

$$\mathcal{L}_{\text{int}} = -\overline{\psi}_L \gamma^\mu \left(g' \frac{Y}{2} B_\mu + g_W W_\mu^a \frac{\sigma^a}{2} \right) \psi_L - \overline{\psi}_R \gamma^\mu \left(g' \frac{Y}{2} B_\mu \right) \psi_R \,, \tag{2.27}$$

which has the required property that the three $SU(2)_L$ gauge bosons couple to left-handed fermions and the $U(1)_Y$ gauge boson couples to both left- and right-handed fermions.

In this description, the first and second $SU(2)$ gauge fields, $W_\mu^{(1)}$ and $W_\mu^{(2)}$, mix to form the physical $W^\pm$ fields, and the third gauge field $W_\mu^{(3)}$ mixes with the B_μ gauge field to produce the physical Z and photon fields of the Standard Model (this mixing will be described in detail in Section 2.2.2). Furthermore, in this unified formulation, weak isospin and hypercharge are related to the electric charge q according to

$$q = I_W^{(3)} + \frac{Y}{2} \,. \tag{2.28}$$

How the Standard Model accounts for the mass of the $W^\pm$ and Z bosons in light of the requirement of gauge invariance is the subject of the following section.

2.2.2 Spontaneous symmetry breaking and the Higgs mechanism

The $W^\pm$ and Z bosons being massive posed a significant problem in the development of the Standard Model, since the inclusion of mass terms for the corresponding gauge fields (and for the fermions)

breaks the gauge invariance of the theory and makes it non-renormalizable. Since the requirement of gauge invariance is understood to produce the interaction itself, and a non-renormalizable theory contains infinitely many free parameters above some cut-off energy scale Λ, any theory that adds the mass terms "by hand" and overlooks the broken gauge invariance cannot provide a functional, internally consistent description of nature. Using the principles of spontaneous symmetry breaking described by Nambu [52], Goldstone [53, 54] and Anderson [55], a mechanism was proposed by Englert and Brout [56], Higgs [57, 58], and Guralnik, Hagen and Kibble [59] in which the $W^{\pm}$ and Z bosons acquire mass while maintaining the required gauge symmetry. Higgs identified that this process of spontaneous symmetry breaking would imply the existence of an additional massive scalar boson, thus for brevity this mechanism is often referred to simply as the *Higgs mechanism*.

Symmetry breaking for a complex scalar field

To illustrate the principles of spontaneous symmetry breaking in the Higgs mechanism, consider the complex scalar field,

$$\phi = \frac{1}{\sqrt{2}}(\phi_1 + i\phi_2)\,, \tag{2.29}$$

whose dynamics are governed by the Lagrangian

$$\mathcal{L} = (\partial_\mu \phi)^*(\partial^\mu \phi) - V(\phi)\,, \quad \text{with} \quad V(\phi) = \mu^2(\phi^*\phi) + \lambda(\phi^*\phi)^2\,. \tag{2.30}$$

This Lagrangian is invariant under the global $U(1)$ transformation $\phi \to \phi' = e^{i\alpha}\phi$ since $\phi'^*\phi' = \phi^*\phi$. For the potential to have a finite minimum, λ must be positive, while the shape of the potential depends on the sign of μ^2, as shown in Figure 2.2. When $\mu^2 > 0$, the minimum of the potential occurs where both fields are zero. If $\mu^2 < 0$, the potential has an infinite set of minima defined by

$$\phi_1^2 + \phi_2^2 = \frac{-\mu^2}{\lambda} = v^2\,, \tag{2.31}$$

as shown by the ring of radius v in Figure 2.2b. The constant v, called the *vacuum expectation value*, is the lowest energy state of the field ϕ, and in this case is degenerate due to the symmetry about the potential minimum. The *physical* vacuum state will correspond to a particular point on this circle, breaking the global $U(1)$ symmetry of the Lagrangian. Without loss of generality, the vacuum state can be chosen to be in the real direction, $(\phi_1, \phi_2) = (v, 0)$, and the complex scalar field ϕ can be expanded as a small perturbation about the vacuum state by writing $\phi_1(x) = \eta(x) + v$ and $\phi_2(x) = \xi(x)$:

$$\phi = \frac{1}{\sqrt{2}}(\eta + v + i\xi)\,. \tag{2.32}$$

The Lagrangian of Eq. (2.30), written in terms of the fields η and ξ, is

$$\mathcal{L} = \frac{1}{2}(\partial_\mu \eta)(\partial^\mu \eta) + \frac{1}{2}(\partial_\mu \xi)(\partial^\mu \xi) - V(\eta, \xi)\,, \tag{2.33}$$

where the potential $V(\eta, \xi)$ is given by

$$V(\eta, \xi) = -\frac{1}{4}\lambda v^4 + \lambda v^2\eta^2 + \lambda v\eta^3 + \frac{1}{4}\lambda\eta^4 + \frac{1}{4}\lambda\xi^4 + \lambda v\eta\xi^2 + \frac{1}{2}\lambda\eta^2\xi^2\,. \tag{2.34}$$

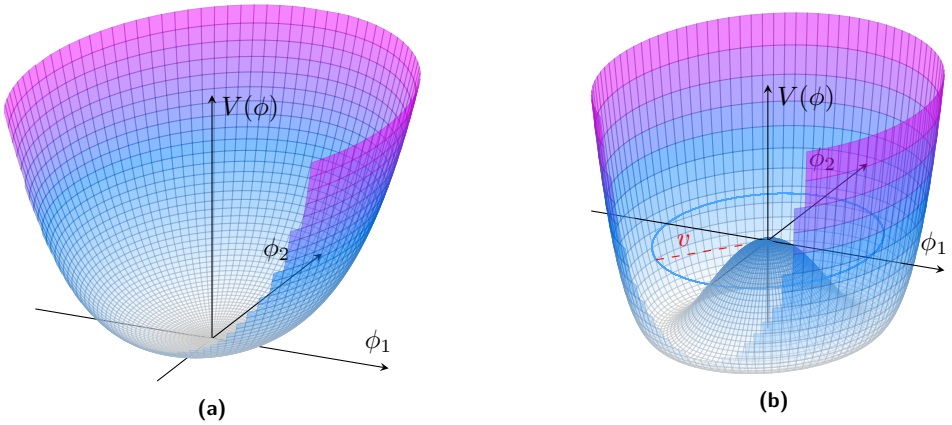

Figure 2.2: The potential $V(\phi) = \mu^2(\phi^*\phi) + \lambda(\phi^*\phi)^4$ for a complex scalar field where (a) $\mu^2 > 0$ and (b) $\mu^2 < 0$.

The term quadratic in the field η can be identified as a mass term with $m_\eta = \sqrt{2\lambda v^2}$, and the terms with either three or four powers of the fields can be identified as interaction terms corresponding to the cubic and quartic couplings of the fields η and ξ. This Lagrangian therefore represents a massive scalar field η and a massless scalar field ξ. The excitations of the massive field η are in the direction where the potential is quadratic. In contrast, the particles described by the massless scalar field ξ correspond to excitations in the direction where the potential does not change. This massless scalar particle is known as a *Goldstone boson*.

The Higgs mechanism

The Higgs mechanism combines the principles of spontaneous symmetry breaking of a complex scalar field and of local gauge invariance to formulate a gauge-invariant theory with massive bosons. We will first consider a $U(1)$ local gauge symmetry to introduce the main principles of the Higgs mechanism, and in the following section we will promote this symmetry to the $U(1)_Y \times SU(2)_L$ symmetry of the unified electroweak interaction.

The Lagrangian for the complex scalar field in Eq. (2.30) can be made gauge invariant under a local $U(1)$ transformation by introducing the covariant derivative

$$\partial_\mu \to D_\mu = \partial_\mu + igB_\mu \,, \tag{2.35}$$

provided the new gauge field B_μ transforms as

$$B_\mu \to B'_\mu = B_\mu - \partial_\mu \alpha(x) \,. \tag{2.36}$$

The combined Lagrangian for the complex scalar field ϕ and the gauge field B_μ is

$$\mathcal{L} = -\frac{1}{4} F_{\mu\nu} F^{\mu\nu} + (D_\mu \phi)^* (D^\mu \phi) - \mu^2 \phi^2 - \lambda \phi^4 \,, \tag{2.37}$$

with $F_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu$. Here, the gauge field B_μ is required to be massless, since the mass term $\frac{1}{2} m_B B_\mu B^\mu$ would break the gauge invariance.

Upon spontaneous symmetry breaking, where the scalar-field potential has $\mu^2 < 0$, the vacuum

state is degenerate and as before, the physical vacuum state is chosen to be $\phi_1 + i\phi_2 = v$. Expanding the complex scalar field ϕ about the vacuum state as

$$\phi(x) = \frac{1}{\sqrt{2}}[v + \eta(x) + i\xi(x)] \tag{2.38}$$

and substituting into the Lagrangian of Eq. (2.37) leads to

$$\begin{aligned}\mathcal{L} = &\underbrace{\frac{1}{2}(\partial_\mu\eta)(\partial^\mu\eta) - \lambda v^2\eta^2}_{\text{massive } \eta} + \underbrace{\frac{1}{2}(\partial_\mu\xi)(\partial^\mu\xi)}_{\text{massless } \xi} - \underbrace{\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \frac{1}{2}g^2v^2B_\mu B^\mu}_{\text{massive gauge field}} \\ &- V_{\text{int}}(\eta, \xi, B) + gvB_\mu(\partial^\mu\xi)\,,\end{aligned} \tag{2.39}$$

where $V_{\text{int}}(\eta, \xi, B)$ contains the three- and four-point interaction terms of the fields η, ξ and B. As before, the breaking of the symmetry of the Lagrangian produces a massive scalar field η and a massless Goldstone boson ξ. In addition, the previously massless gauge field B_μ has acquired a mass term $\frac{1}{2}g^2v^2B_\mu B^\mu$, achieving the intent of giving a mass to the gauge boson associated with the local gauge symmetry.

However, in this form there appears a direct coupling between the Goldstone field ξ and the gauge field B_μ in the term $gvB_\mu(\partial^\mu\xi)$, in which the spin-1 gauge field can transform into a spin-0 scalar field, suggesting these fields are not physical. The solution lies in the gauge-invariant construction of the theory, and the Goldstone field ξ can be eliminated from the Lagrangian altogether by a careful choice of $U(1)$ gauge transformation:

$$B_\mu(x) \to B'_\mu(x) = B_\mu(x) + \frac{1}{gv}\partial_\mu\xi(x)\,. \tag{2.40}$$

This gauge transformation corresponds to taking $\alpha(x) = -\xi(x)/gv$ in Eq. (2.36), called the *unitary gauge*, which in turn corresponds to a gauge transformation of the complex field ϕ such that it becomes real-valued at every point and thus eliminates the Goldstone field ξ:

$$\phi(x) = \frac{1}{\sqrt{2}}[v + \eta(x)]\,. \tag{2.41}$$

Moreover, this choice of gauge removes the degree of freedom associated to the Goldstone boson ξ and replaces it with the degree of freedom associated to the longitudinal polarization of the now massive gauge field B_μ. The field η is the physical Higgs field, and henceforth will be denoted as h.

Writing $\mu^2 = -\lambda v^2$, working in the unitary gauge where $\phi(x) = \frac{1}{\sqrt{2}}[v + h(x)]$, and ignoring constant terms, the Lagrangian of Eq. (2.39) can be written

$$\begin{aligned}\mathcal{L} = &\underbrace{\frac{1}{2}(\partial_\mu h)(\partial^\mu h) - \lambda v^2h^2}_{\text{massive scalar } h} - \underbrace{\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \frac{1}{2}g^2v^2B_\mu B^\mu}_{\text{massive gauge boson } B} \\ &+ \underbrace{g^2vB_\mu B^\mu h + \frac{1}{2}g^2B_\mu B^\mu h^2}_{h,\,B \text{ interactions}} - \underbrace{\lambda vh^3 - \frac{1}{4}\lambda h^4}_{h \text{ self-interactions}}\,.\end{aligned} \tag{2.42}$$

This Lagrangian describes a massive scalar Higgs field h and a massive gauge boson B associated with the $U(1)$ local gauge symmetry. It contains interaction terms between the Higgs boson and the

gauge boson, and Higgs boson self-interaction terms. The mass of the gauge boson,

$$m_B = gv\,,$$

is related to the strength of the gauge coupling and the vacuum expectation value of the Higgs field, and the mass of the Higgs boson is given by

$$m_H = \sqrt{2\lambda}v\,.$$

The Standard Model Higgs boson

Following Glashow's unification of the electromagnetic and weak interactions [16], Weinberg [17] and Salam [18] embedded the Higgs mechanism into the $U(1)_Y \times SU(2)_L$ local gauge symmetry of the electroweak theory of the Standard Model. In the Weiberg–Salam model, three Goldstone bosons are required to provide the longitudinal degrees of freedom of the $W^{\pm}$ and Z bosons. In addition, after symmetry breaking, there will be a massive scalar particle corresponding to the field excitations in the direction picked out by the choice of the physical vacuum. The simplest Higgs model, which has the necessary four degrees of freedom, consists of two complex scalar fields placed in a weak-isospin doublet:

$$\phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix}\,, \tag{2.43}$$

where ϕ^0 denotes a neutral scalar field, and ϕ^+ denotes a charged scalar field, with $(\phi^+)^* = \phi^-$. The corresponding Lagrangian with the Higgs potential is

$$\mathcal{L} = (\partial_\mu \phi)^\dagger (\partial^\mu \phi) - V(\phi)\,, \quad \text{with} \quad V(\phi) = \mu^2 \phi^\dagger \phi + \lambda(\phi^\dagger \phi)^2\,. \tag{2.44}$$

For $\mu^2 < 0$, the potential has an infinite set of degenerate minima satisfying

$$\phi^\dagger \phi = \frac{1}{2}(\phi_1^2 + \phi_2^2 + \phi_3^2 + \phi_4^2) = -\frac{\mu^2}{2\lambda} = \frac{v^2}{2}\,. \tag{2.45}$$

The neutral photon is required to remain massless after symmetry breaking, therefore only the neutral scalar field ϕ^0 acquires a non-zero vacuum expectation value:

$$\langle 0|\phi|0\rangle = \frac{1}{2}\begin{pmatrix} 0 \\ v \end{pmatrix}\,. \tag{2.46}$$

The fields can then be expanded about this minimum by writing the scalar-field doublet as

$$\phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi_1(x) + i\phi_2(x) \\ v + \eta(x) + i\phi_4(x) \end{pmatrix}\,. \tag{2.47}$$

By "gauging-away" the Goldstone fields as before using the unitary gauge, the Higgs doublet can be immediately written as

$$\phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix}\,. \tag{2.48}$$

The mass terms can be identified by writing the Lagrangian of Eq. (2.44) such that it respects the $SU(2)_L \times U(1)_Y$ local gauge symmetry of the electroweak model by promoting the derivatives to the appropriate covariant derivatives from Eq. (2.23), which act on the Higgs doublet ϕ as

$$D_\mu \phi = \frac{1}{2}\left[2\partial_\mu + \left(i g_W \sigma^a W_\mu^a + i g' B_\mu\right)\right]\phi\,, \tag{2.49}$$

where here D_μ is a 2×2 matrix acting on the weak-isospin Higgs doublet with weak hypercharge $Y = 1$. The masses of the gauge bosons arise from the $(D_\mu\phi)^\dagger(D^\mu\phi)$ term in the Lagrangian, given by

$$\begin{aligned}(D_\mu\phi)^\dagger(D^\mu\phi) = &\frac{1}{2}(\partial_\mu h)(\partial^\mu h) + \frac{1}{8}g_W^2(W_\mu^{(1)} + iW_\mu^{(2)})(W^{(1)\mu} - iW^{(2)\mu})(v+h)^2 \\ &+ \frac{1}{8}(g_W W_\mu^{(3)} - g' B_\mu)(g_W W^{(3)\mu} - g' B^\mu)(v+h)^2\,,\end{aligned} \tag{2.50}$$

and the gauge boson mass terms are those which are quadratic in the gauge boson fields:

$$\frac{1}{2}v^2 g_W^2 (W_\mu^{(1)} W^{(1)\mu} + W_\mu^{(2)} W^{(2)\mu}) + \frac{1}{8}v^2(g_W W_\mu^{(3)} - g' B_\mu)(g_W W^{(3)\mu} - g' B^\mu)\,. \tag{2.51}$$

The physical $W^\pm$ fields can be identified in the first term of Eq. (2.51) as a linear combination of the $W_\mu^{(1)}$ and $W_\mu^{(2)}$ gauge fields:

$$W_\mu^\pm = \frac{1}{\sqrt{2}}\left(W_\mu^{(1)} \mp i W_\mu^{(2)}\right)\,, \tag{2.52}$$

with mass terms appearing as $\frac{1}{2}m_W^2 W_\mu^{(1)} W^{(1)\mu}$ and $\frac{1}{2}m_W^2 W_\mu^{(2)} W^{(2)\mu}$. It follows that the mass of the W boson is determined by the coupling constant of the $SU(2)_L$ gauge interaction g_W and the vacuum expectation value of the Higgs field, and is given by

$$m_W = \frac{1}{2} g_W v\,. \tag{2.53}$$

The second term in Eq. (2.51), quadratic in the neutral $W_\mu^{(3)}$ and B_μ fields, can be written in matrix form as

$$\frac{1}{8}v^2 \begin{pmatrix} W_\mu^{(3)} & B_\mu \end{pmatrix} \underbrace{\begin{pmatrix} g_W^2 & -g_W g' \\ -g_W g' & g'^2 \end{pmatrix}}_{\text{“mass matrix” } \mathbf{M}} \begin{pmatrix} W^{(3)\mu} \\ B^\mu \end{pmatrix} \tag{2.54}$$

The off-diagonal elements of the *mass matrix* $\mathbf{M}$ couple together the $W_\mu^{(3)}$ and B_μ fields, allowing them to mix. The physical boson fields A_μ and Z_μ correspond to the basis in which the mass matrix is diagonal, which is obtained from the eigenvalues of $\mathbf{M}$: 0 and $g_W^2 + g'^2$. Hence, in the diagonal basis, Eq. (2.54) becomes

$$\frac{1}{8}v^2 \begin{pmatrix} A_\mu & Z_\mu \end{pmatrix} \begin{pmatrix} 0 & 0 \\ 0 & g_W^2 + g'^2 \end{pmatrix} \begin{pmatrix} A^\mu \\ Z^\mu \end{pmatrix}\,. \tag{2.55}$$

The physical fields are given by the normalized eigenvectors of the mass matrix,

$$A_\mu = \frac{g' W_\mu^{(3)} + g_W B_\mu}{\sqrt{g_W^2 + g'^2}} \quad \text{and} \quad Z_\mu = \frac{g_W W_\mu^{(3)} - g' B_\mu}{\sqrt{g_W^2 + g'^2}}, \tag{2.56}$$

and the masses of the physical gauge bosons can be identified as

$$m_A = 0 \quad \text{and} \quad m_Z = \frac{1}{2} v \sqrt{g_W^2 + g'^2}\,. \tag{2.57}$$

Through this change of basis, we have identified two physical neutral boson fields: a massless gauge boson A, which can be identified as the photon, and a massive gauge boson, which can be identified as the Z boson.

The mixing of the underlying gauge fields and their relation to the masses of the physical gauge bosons may be written more succinctly by writing the ratio of the couplings of the $U(1)_Y$ and $SU(2)_L$ gauge symmetries as

$$\frac{g'}{g_W} = \tan\theta_W\,, \tag{2.58}$$

where θ_W is the *weak mixing* or *Weinberg angle*. The relationship between the physical fields and underlying fields of Eq. (2.56) can then be written as

$$\begin{aligned} A_\mu &= \cos\theta_W B_\mu + \sin\theta_W W_\mu^{(3)}\,, \\ Z_\mu &= -\sin\theta_W B_\mu + \cos\theta_W W_\mu^{(3)}\,. \end{aligned} \tag{2.59}$$

Furthermore, the mass of the Z boson in Eq. (2.57) can be expressed as

$$m_Z = \frac{1}{2}\frac{g_W}{\cos\theta_W} v\,, \tag{2.60}$$

and the ratio of the W and Z boson masses as

$$\frac{m_W}{m_Z} = \cos\theta_W\,. \tag{2.61}$$

The electromagnetic and weak coupling constants are related by

$$e = g_W \sin\theta_W = g' \cos\theta_W\,. \tag{2.62}$$

Similar analysis of other terms in the Lagrangian after spontaneous symmetry breakings leads to the description of a Higgs boson with a mass of $m_H = \sqrt{2\lambda} v$. While the Higgs potential parameters μ and λ are related to the vacuum expectation value via $v^2 = -\mu^2/\lambda$, the Standard Model offers no independent predictions to compute them directly, and therefore the Higgs boson mass is a free parameter of the Standard Model. In this basis, the electroweak coupling constants g_W and g', and the vacuum expectation value are also free parameters. Using the relation $m_W = \frac{1}{2} g_W v$ and the measured values of g_W and m_W, the vacuum expectation value of the Higgs field is found to be approximately $v \approx 246\,\text{GeV}$, and therefore a measurement of the Higgs boson mass fixes the two Higgs potential parameters.

Finally, the couplings of the Higgs field to the gauge bosons arise in the $(D_\mu\phi)^\dagger(D^\mu\phi)$ term in

the Lagrangian of Eq. (2.50) in the form VVh and $VVhh$, where $V = W^{\pm}, Z$, corresponding to the cubic and quartic couplings between one or two Higgs bosons and the gauge bosons, respectively. The coupling strengths of these interactions are

$$g_{HVV} = 2\frac{m_V^2}{v}\,, \quad \text{and} \quad g_{HHVV} = 2\frac{m_V^2}{v^2}\,. \tag{2.63}$$

Similarly, expanding the interaction term of the Lagrangian reveals the cubic and quartic Higgs self interaction terms with coupling strengths

$$g_{HHH} = 3\frac{m_H^2}{v}\,, \quad \text{and} \quad g_{HHHH} = 3\frac{m_H^2}{v^2}\,. \tag{2.64}$$

The Feynman diagrams corresponding to these interactions are given in Figure 2.3.

Fermion masses

Fermion mass terms are introduced to theory in much the same way as they were for the massive gauge bosons using the Higgs mechanism. Recall that in the framework of the electroweak theory, the left-handed fermions are organized into $SU(2)$ doublets, here labelled L, and the right-handed fermions into $SU(2)$ singlets, here labelled R. Consequently, introducing a mass term such as

$$-m\overline{\psi}\psi = -m(\overline{\psi}_R\psi_L + \overline{\psi}_L\psi_R) \tag{2.65}$$

to the free-field Dirac Lagrangian does not satisfy the required $SU(2)_L \times U(1)_Y$ gauge symmetry. However, combinations with the Higgs doublet of Eq. (2.43) in the form $\overline{L}\phi R$ and its hermitian conjugate $(\overline{L}\phi R)^\dagger = \overline{R}\phi^\dagger L$ *are* invariant under the $SU(2)_L \times U(1)_Y$ gauge transformation. Considering first the charged leptons $\ell = e, \mu, \tau$, the resulting Lagrangians contain interaction terms of the form

$$\mathcal{L}_{\text{Yukawa}} = -g_\ell \left[\begin{pmatrix} \overline{\nu}_\ell & \overline{\ell} \end{pmatrix}_L \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} \ell_R + \overline{\ell}_R \begin{pmatrix} \phi^{+*} & \phi^{0*} \end{pmatrix} \begin{pmatrix} \nu_\ell \\ \ell \end{pmatrix}_L \right], \tag{2.66}$$

where g_ℓ is the *Yukawa coupling constant* of the lepton ℓ to the Higgs field. After spontaneous symmetry breaking and placing the Higgs field in the unitary gauge of Eq. (2.48), this Lagrangian becomes

$$\mathcal{L}_{\text{Yukawa}} = -\frac{g_\ell}{\sqrt{2}} v(\overline{\ell}_L\ell_R + \overline{\ell}_R\ell_L) - \frac{g_\ell}{\sqrt{2}} h(\overline{\ell}_L\ell_R + \overline{\ell}_R\ell_L)\,. \tag{2.67}$$

Thus, the Higgs mechanism introduces mass terms for the charged leptons in a gauge-invariant manner, as well as the interaction terms between the Higgs and lepton fields with a coupling proportional to the lepton mass: $g_\ell = \sqrt{2}m_\ell/v$. The Lagrangian can then be written as

$$\mathcal{L}_{\text{Yukawa}} = -m_\ell\overline{\ell}\ell - \frac{m_\ell}{v}\overline{\ell}\ell h\,. \tag{2.68}$$

Since no right-handed neutrinos exist in the Standard Model, the neutrinos ν_ℓ do not acquire a mass term and are therefore massless after spontaneous symmetry breaking.[6]

[6] This result is in direct contradiction with the experimental observation of neutrino oscillations, which demonstrates that neutrinos do have mass [2, 3]. However, having the Yukawa couplings span many orders of magnitude to generate the small neutrino masses might suggest that neutrinos acquire their mass through some other mechanism. This problem is currently an active area of research.

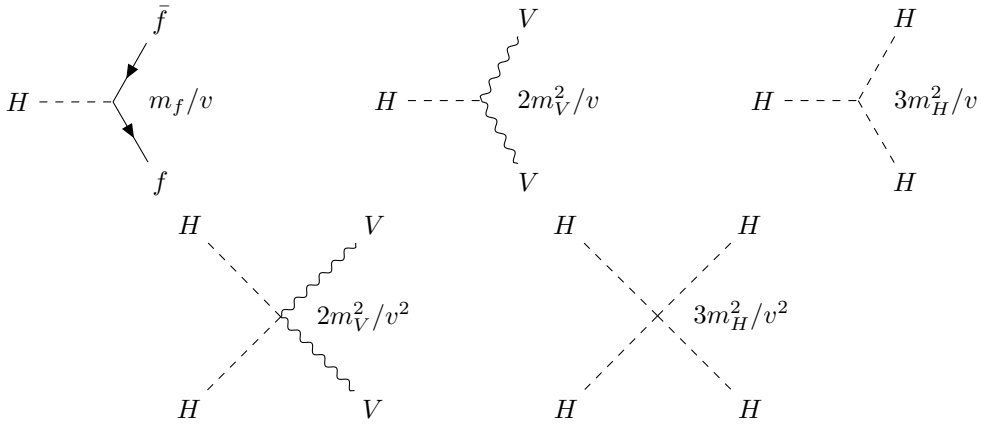

Figure 2.3: Feynman diagrams of the Standard Model Higgs boson couplings to fermions and gauge bosons, where $V = W^{\pm}, Z$, and the Higgs self-couplings. The normalized vertex factors of the Feynman rules are also displayed.

The same procedure can be applied to generate the mass terms for the down-type quarks, however this would incorrectly leave the up-type quarks massless. Replacing the Higgs doublet with its *conjugate doublet* $\phi_c \equiv -i\sigma_2\phi^*$ gives the appropriate mass terms following the above procedure with the invariant quantities $\overline{L}\phi_c R$ and $\overline{R}\phi_c^\dagger L$. In general, the Yukawa couplings of the fermions to the Higgs field are given by

$$g_f = \sqrt{2}\frac{m_f}{v}\,. \tag{2.69}$$

The Feynman diagram corresponding to the Higgs–fermion interaction is included in Figure 2.3.

2.2.3 Higgs boson production and decays

In pp collisions at the LHC, the production of a Standard Model Higgs boson is predominantly, in order of decreasing cross section, through gluon–gluon fusion (ggF), vector-boson fusion (VBF), vector-boson associated production (VH) and top/bottom-quark associated production (ttH/bbH). Figure 2.4 shows the leading-order Feynman diagrams of these production modes. Other Higgs boson production modes also contribute, such as single-top-quark associated production (tH), although with a cross section approximately an order of magnitude less than ttH. Figure 2.5 shows the theoretical predictions of the leading Higgs boson production cross sections in $\sqrt{s} = 13\,\text{TeV}$ pp collision at the LHC as a function of its mass up to $m_H = 3000\,\text{GeV}$. These cross section predictions serve as a basis for BSM Higgs searches (discussed in Section 2.3.1) and assume a narrow Higgs-like resonance with SM-like couplings. State-of-the-art Higgs boson production cross sections include NLO electroweak corrections, while for simplicity these predictions do not [60].

The Higgs boson is a comparatively heavy particle, and consequently tends to decay almost immediately into lighter, more stable particles, with an extremely short lifetime of $\sim 10^{-22}\,\text{s}$ [61]. The decays of the Higgs boson and their branching ratios in the regime where $m_H \approx 125\,\text{GeV}$ reflect the proportionality of the Higgs boson couplings to mass. In particular, final states with heavy particles, or interactions mediated by heavy particles, tend to dominate the Higgs boson total width, while modes with direct couplings to lighter particles are rare. The decay modes can be broadly grouped into three categories: direct Higgs couplings to massive particles (e.g. $H \to b\bar{b}$),

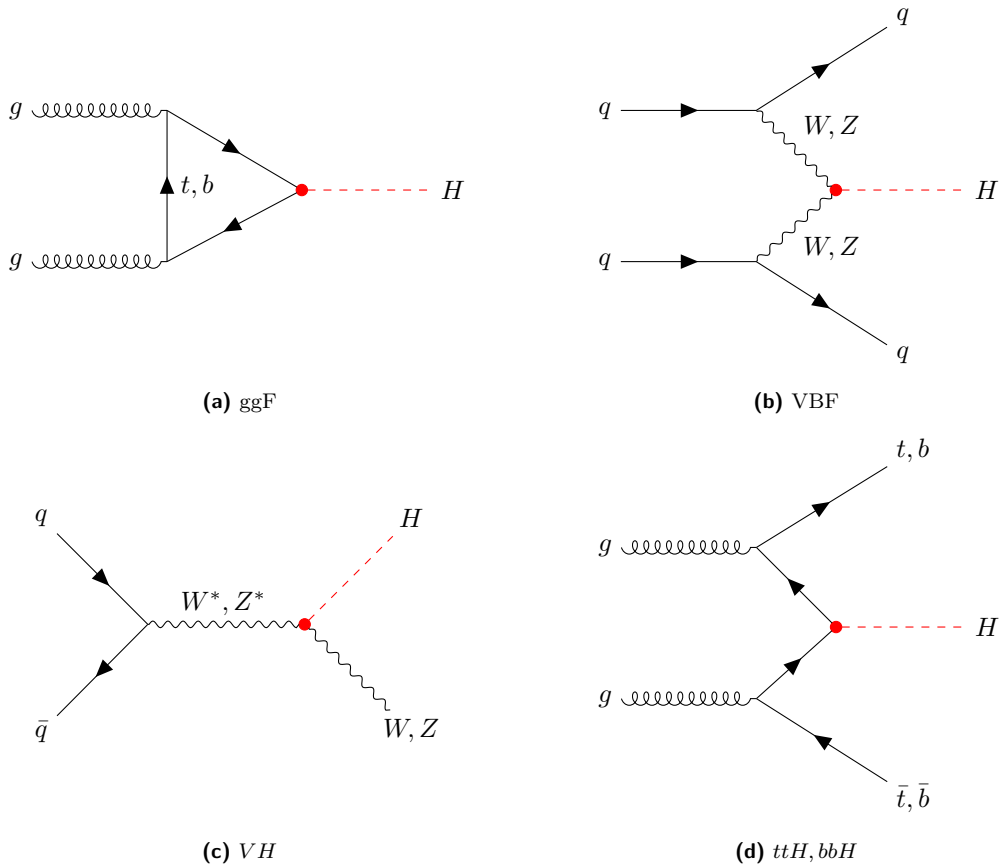

(a) ggF **(b)** VBF **(c)** VH **(d)** ttH, bbH

Figure 2.4: Tree-level Feynman diagrams of the four leading Higgs production modes at the LHC: (a) gluon–gluon fusion (ggF), (b) vector-boson fusion (VBF), (c) associated production with a vector boson (VH) and (d) associated production with a $t\bar{t}$ or $b\bar{b}$ pair ($t\bar{t}H, b\bar{b}H$).

loop-induced decays to massless particles (e.g. $H \to \gamma\gamma$, $H \to gg$), and decays to leptons through pairs of weak bosons (e.g. $H \to ZZ^* \to \ell^+\ell^-\ell^+\ell^-$, $H \to WW^* \to \ell^+\nu\ell^-\bar{\nu}$). Figures 2.6 and 2.7 show the theoretical predictions of the Higgs boson branching ratios and the total decay width as a function of its mass up to $m_H = 1000\,\text{GeV}$.

2.2.4 Discovery of the Standard Model Higgs boson

The mass of the Higgs boson being a free parameter of the Standard Model posed a significant challenge to its detection, since any experimental apparatus would need to be built with a sufficient acceptance to reconstruct its decay products, whose kinematics in turn depend on the Higgs boson mass and could therefore potentially span a large dynamic range. However, early theoretical work and inferences from indirect experimental evidence excluded certain ranges of the Higgs boson mass.[7] The absence of Higgs-boson-related effects in nuclear transitions [64] and neutron-scattering experiments [65] placed an initial lower bound on the Higgs boson mass of $m_H > 18.3\,\text{MeV}$. Later, arguments of unitarity violation in vector-boson scattering at high energies unless a massive boson with the properties of the Higgs is included in the scattering-amplitude calculation placed an upper

[7] For a historical review of Higgs boson searches, see Ref. [63].

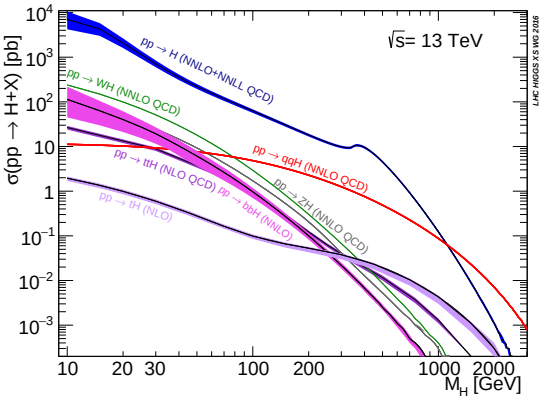

Figure 2.5: Higgs boson production cross sections as a function of its mass in pp collisions at centre-of-mass energy $\sqrt{s} = 13\,\text{TeV}$ [60].

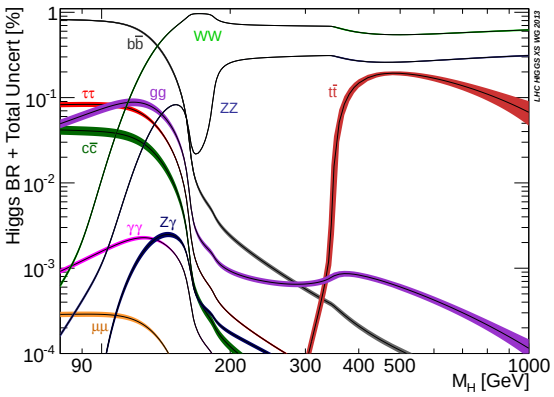

Figure 2.6: Higgs boson decay branching ratios as a function of its mass [62].

bound on the Higgs boson mass of

$$m_H \leq \sqrt{\frac{16\pi}{3}} v \approx 1\,\text{TeV}\,, \tag{2.70}$$

known as the *Lee–Quigg–Thacker bound* [66].

Among the earliest direct searches for the Higgs boson were those conducted at the Large Electron–Positron Collider (LEP) at CERN using e^+e^- collision data, with the most stringent limits set using data from the LEP2 era at centre-of-mass energies between 189 and 209 GeV. The combined LEP results placed a lower bound on the Higgs boson mass of $m_H > 114.4\,\text{GeV}$ at 95 % confidence level (CL), primarily driven by searches for a Higgs boson produced in association with a Z boson: $e^+e^- \to Z^* \to HZ$ [67]. Later searches at the Tevatron $p\bar{p}$ collider in the $H \to b\bar{b}$ and $H \to W^+W^-$ decay channels at a centre-of-mass energy of 1.96 TeV excluded the possibility of a Higgs boson at 95 % CL in the mass range $100 < m_H < 103\,\text{GeV}$ and $147 < m_H < 180\,\text{GeV}$, and furthermore a $\sim 3\sigma$ excess over the predicted background was observed in the range $115 < m_H < 140\,\text{GeV}$ [68]. The Higgs boson mass regions excluded at 95 % CL by the LEP and Tevatron experiments, and in combination with early LHC Run 1 data, are shown in Figure 2.8.

Further theoretical constraints on the Higgs boson mass were provided from a global fit to the electroweak parameters known to high precision, with the strongest constraint arising from the W boson mass and the effective weak mixing angle [69]. The best-fit value for m_H was found to be $m_H = 91^{+30}_{-23}\,\text{GeV}$. Including the results of the direct searches at LEP and the Tevatron, the final best-fit value was $m_H = 120^{+12}_{-5}\,\text{GeV}$.

By July 2012, the ATLAS and CMS Collaborations had recorded a sufficiently large data sample of pp collisions at centre-of-mass energies of 7 and 8 TeV to claim the observation of a new particle consistent with the Standard Model Higgs boson [70, 71] with a mass of approximately 125 GeV. The observed local significance from these searches as a function of the Higgs boson mass is shown in Figure 2.9. Later measurements of the new particle's couplings, spin and CP properties, and production and decay rates using the combined ATLAS and CMS Run 1 datasets [72] provided

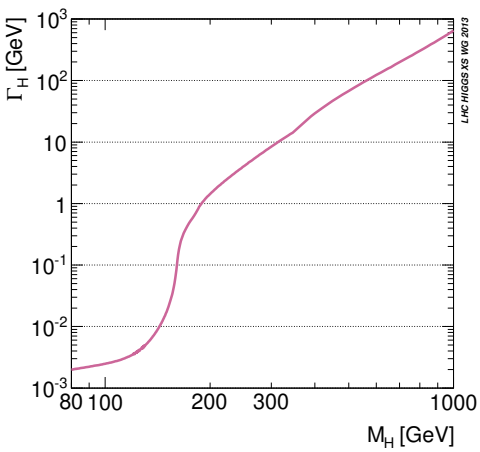

Figure 2.7: Total decay width of the Standard Model Higgs boson as a function of its mass [62].

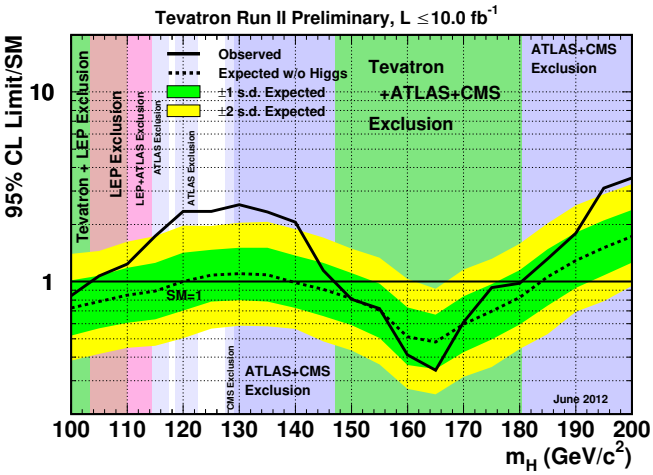

Figure 2.8: Observed and expected upper limits at 95 % CL on the ratios to the SM cross section as a function of mass in early searches for the Higgs boson. The shaded areas indicate the excluded regions from combinations of searches at the Tevatron, LEP and with early LHC Run 1 data [68].

conclusive evidence that the Standard Model Higgs boson had been observed with mass [73]

$$m_H = 125.09 \pm 0.21\,(\text{stat}) \pm 0.11\,(\text{syst})\,\text{GeV}.$$

2.3 Physics beyond the Standard Model

The discovery of the Higgs boson in 2012 concluded the search for the elementary particles contained in the Standard Model. In all, the Standard Model has been and continues to be a remarkably successful theory of nature, both phenomenologically in its description of the fundamental constituents of matter and their interactions, and in the precision of its predictions. For instance, the anomalous magnetic moment of the electron $a_e \equiv (g_s - 2)/2$, which arises as the effect of radiative corrections on the Dirac spin g-factor $g_s = 2$, has been calculated up to $O(\alpha^5)$ [74] and agrees with the experimentally measured value [75] with an accuracy at the level of approximately 1 part per billion.

Despite its success, the Standard Model is an incomplete theory of nature and contains a number of noteworthy limitations. These limitations can be broadly classified into two categories: phenomena and experimental results not explained, and internal theoretical problems.

First, the Standard Model provides no explanation for a number of well-established observed phenomena. The gravitational force, for instance, cannot be explained as a renormalizable gauge theory as the other fundamental interactions are in the Standard Model. The Standard Model also lacks a description of dark matter and dark energy, which have been verified by astrophysical measurements to account for approximately 95 % of the energy density in the Universe [76]. The Universe is also made of matter rather than antimatter, which implies some baryon number non-conserving process occurred in the early Universe, which is again outside of the Standard Model.

Second, the internal theoretical problems of the Standard Model are more subtle, and not "prob-

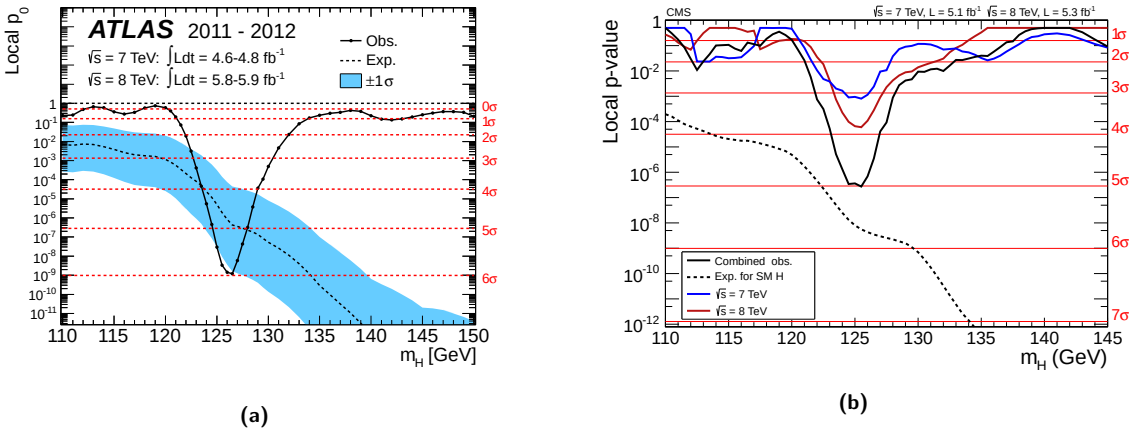

(a) (b)

Figure 2.9: First observation of a new particle consistent with the Standard Model Higgs boson by the (a) ATLAS [70] and (b) CMS [71] Collaborations. Shown is the local significance as a function of the particle mass using data recording during LHC Run 1 at centre-of-mass energies of $\sqrt{s} = 7$ and 8 TeV. Each experiment observed a $\gtrsim 5\sigma$ excess at a mass of approximately 125 GeV.

lems" per se, but rather reflect the large number of free parameters[8] in the Standard Model added in an ad hoc way and the issues of "naturalness" and "fine-tuning". For example, the Standard Model contains distinct Yukawa couplings for each lepton and quark flavour, raising the question of why there is such a spread in the observed mass spectrum; could the fermion masses be related to one another through a more fundamental, underlying theory? Along similar lines, quantum chromodynamics contains a parameter θ that controls the size of the CP-violating effects in the strong interaction. The value of this parameter can be inferred from the size of the neutron electric dipole moment, with experimental results setting an upper bound of $\theta < 10^{-10}$, an "unnaturally" small value. Why this free parameter appears in the Standard Model yet is so close to zero is known as the *strong CP problem*.

A similar issue relating to naturalness and fine-tuning arises in the electroweak sector of the Standard Model, known as the *hierarchy problem*. Broadly speaking, one might ask why the gravitational force is orders of magnitude weaker than the electroweak interaction. To illustrate the nature of this problem more explicitly, consider the radiative corrections to the bare mass[9] of the Higgs boson, which originate from fermion, Higgs and gauge-boson loops, a few examples of which are shown in Figure 2.10. Several of these diagrams are quadratically divergent, with the most significant having a mass counterterm $\delta m_H^2 \propto g_t^2 \Lambda^2$, where g_t is the top-quark Yukawa coupling and Λ is the ultraviolet cutoff of the theory. Thus, if the physical value of the Higgs boson mass m_H is small compared to Λ, $m_H \ll \Lambda$, one has to fine-tune the bare mass to cancel the radiative corrections. If the Standard Model is valid up to very high mass scales, such as that of a Grand Unified Theory $\Lambda_{\mathrm{GUT}} \sim 10^{16}$ GeV or the Planck scale $\Lambda_P \sim 10^{19}$ GeV, these corrections become very large. Consequently, the "amount of fine-tuning" ϵ, conventionally expressed as the accuracy at which the different contributions to the Higgs boson mass must cancel among themselves in order

[8] The Standard Model has nineteen free parameters: nine fermion masses (excluding the neutrino masses), three CKM mixing angles and one CKM CP-violating phase, three gauge couplings, the QCD vacuum angle, the Higgs vacuum expectation value, and the Higgs boson mass.

[9] The bare mass m_0 of an elementary particle is its mass as it appears in the Lagrangian of the theory, and is related to a finite *renormalized mass* m_R via $m_0 = (1 + \delta_m) m_R$, where δ_m is the *mass counterterm* that accounts for the radiative corrections to the bare mass.

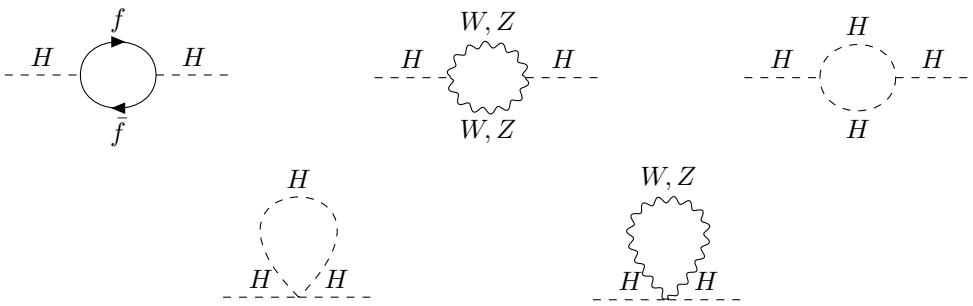

Figure 2.10: Examples of radiative loop corrections to the Higgs boson self-energy. At high energy, these corrections become large and require a significant degree of fine-tuning of the counterterms to keep the mass of the Higgs boson at the electroweak scale of 10^2 GeV.

to reproduce the observed value of $m_H \simeq 125\,\text{GeV}$:

$$\epsilon = \frac{m_0^2 - \delta m_H^2}{\Lambda^2} \sim \left(\frac{100\,\text{GeV}}{4\pi\Lambda}\right)^2 , \tag{2.71}$$

results in $\epsilon \ll 1$ at high energy, another "unnaturally" small number.

The proposed theoretical solutions to the hierarchy problems are plentiful, although a full discussion of each is beyond the scope of this thesis. Supersymmetry, for instance, offers a natural solution to the problem of fine-tuning through the addition of supersymmetric partners, called *sparticles*, which differ from their Standard Model partners by half a unit of spin. For every loop of particles in the radiative corrections to the Higgs self-energy, there is a corresponding loop of sparticles, which provide a correction with the opposite sign, thus eliminating the problem of fine-tuning so long as the sparticle masses appearing in the loops are of approximately the same order as the masses of their Standard Model partners.[10] Other theories extend various sectors of the Standard Model in a variety of ways, predicting a plethora of new particles presumed to be accessible at the energies of the LHC.

With the success of the LHC Higgs physics program, similar analyses of LHC data in an extended mass range lend themselves well to searches for new particles predicted under two categories of theories in particular that attempt to remedy the hierarchy problem and the other issues with the Standard Model. The first class of such *Beyond the Standard Model* (BSM) theories are those with extended Higgs sectors, which predict additional Higgs bosons, and the second are those with warped extra dimensions giving rise to gravity particles. A theoretical overview of each class of theories is given below, along with descriptions of the benchmark scenarios used to interpret the data collected by the ATLAS Experiment in the search for heavy resonances decaying to a pair of Z bosons described in Chapter 7.

[10] In "perfectly unbroken" supersymmetry, the sparticles masses would be the same as their corresponding particles and the fine-tuning would be cancelled exactly. However, if this were the case these sparticles would have already been discovered. Theories of supersymmetry with spontaneous symmetry breaking allow for superpartners to differ in mass while still addressing the issue of fine-tuning.

2.3.1 Extended Higgs sectors

Several extensions to the Standard Model Higgs sector have been proposed, each possessing distinct features and rich phenomenologies.[11] Extended Higgs sectors are by construction generally constrained by measurements of the electroweak ρ parameter, defined in the Standard Model at tree level as

$$\rho_0 \equiv \frac{m_W^2}{m_Z^2 \cos^2 \theta_W} = 1\,, \tag{2.72}$$

Any higher-order modifications to the ρ parameter via the introduction of an extended Higgs sector must therefore be consistent with the experimentally measured value $\rho = 1.000\,38 \pm 0.000\,20$ [1].

A familiar example of a theory containing an extended Higgs sector is the *Minimal Supersymmetric Standard Model* (MSSM), which extends the Standard Model with the fewest number of new sparticle states and new interactions while remaining consistent with current observations. The MSSM predicts additional Higgs bosons arranged in two scalar Higgs doublets and is therefore an example of a broader category of Higgs-sector extensions called *two-Higgs-doublet models*, discussed below.

Two-Higgs-doublet models

Two-Higgs doublet models (2HDM) are constructed by adding a second complex scalar $SU(2)_L$ doublet to the Standard Model field content while requiring that the Lagrangian be invariant under the same symmetries. The first 2HDM was proposed as a means to provide an extra source of CP-violation to explain the observed matter–antimatter asymmetry of the Universe [77], however only CP-conserving 2HDMs will be considered here. In recent years, 2HDMs have appeared in a number of supersymmetric theories and various models of baryogenesis, and more generally they have been used as a benchmark model to aid in identifying possible directions to minimal changes with respect to the Standard Model [78].

For two complex scalar doublets of hypercharge $Y = +1$, here denoted Φ_1 and Φ_2, the most general potential for a 2HDM is

$$\begin{aligned} V_{\text{2HDM}} = {} & m_{11}^2 \Phi_1^\dagger \Phi_1 + m_{22}^2 \Phi_2^\dagger \Phi_2 - m_{12}^2 (\Phi_1^\dagger \Phi_2 - \Phi_2^\dagger \Phi_1) + \frac{\lambda_1}{2} (\Phi_1^\dagger \Phi_1)^2 + \frac{\lambda_2}{2} (\Phi_2^\dagger \Phi_2)^2 \\ & + \lambda_3 \Phi_1^\dagger \Phi_1 \Phi_2^\dagger \Phi_2 + \lambda_4 \Phi_1^\dagger \Phi_2 \Phi_2^\dagger \Phi_1 + \frac{\lambda_5}{2} [(\Phi_1^\dagger \Phi_2)^2 + (\Phi_2^\dagger \Phi_1)^2]\,, \end{aligned} \tag{2.73}$$

By construction, $\rho = 1$ at tree level in 2HDMs since both doublets have the same weak isospin and hypercharge as in the Standard Model. After spontaneous symmetry breaking, each of the two doublets acquires a real vacuum expectation value v_1 and v_2. Expanding the fields about these minima results in eight degrees of freedom, three of which provide the $W^\pm$ and Z bosons their mass and the other five result in as many physical Higgs bosons: two neutral CP-even scalars h and H (with $m_h < m_H$), one CP-odd pseudoscalar A, and two charged scalars $H^\pm$. The two main phenomenological parameters of a 2HDM are the α mixing angle between the neutral scalars, with $|\alpha| \leq \pi/2$, and the β mixing angle between the charged scalars and pseudoscalar, defined as the

[11] For a detailed review of proposed extensions to the Standard Model Higgs sectors, see Section IV of Ref. [60].

ratio of the vacuum expectation values:

$$\tan\beta \equiv \frac{v_2}{v_1}\,. \tag{2.74}$$

The parameters α and β define the rotation angles that diagonalize the corresponding mass matrices, and together with the four scalar masses (m_h, m_H, m_A, $m_{H^\pm}$) and the potential parameter m_{12} that mixes the two Higgs doublets completely specify a 2HDM. The Standard Model Higgs field in this formulation can be identified as the linear combination of h and H:

$$H_{\mathrm{SM}} = h\sin(\alpha-\beta) - H\cos(\alpha-\beta)\,. \tag{2.75}$$

In the limit $\cos(\alpha-\beta)\to 0$, known as the *alignment limit*, the light CP-even Higgs boson h is indistinguishable from the Standard Model Higgs boson with the same mass. Furthermore, in the limit where the masses of the H, A and $H^\pm$ bosons are much greater than the h boson mass, the couplings of h, and its corresponding cross sections and branching ratios, become largely SM-like, and is thus known as the *decoupling limit*.

In general, the 2HDM framework allows flavour-changing neutral currents at tree-level due to the separate Yukawa couplings for each doublet, which leads to mixing between the flavour and mass eigenstates. If tree-level flavour-changing neutral currents exist, they must be highly suppressed to agree with current experimental observations. Most 2HDMs therefore make the simplifying assumption to prohibit flavour-changing neutral currents by imposing an additional discrete symmetry on the 2HDM Lagrangian, or equivalently by specifying to which doublet each right-handed fermion uniquely couples. This discrete symmetry requirement results in four independent combinations in how the two Higgs doublets can couple to leptons and up- and down-type quarks. In *Type-I* models, Φ_2 couples to all quarks and leptons, whereas for *Type-II*, Φ_1 couples to down-type quarks and leptons and Φ_2 couples to up-type quarks. The *lepton-specific* model is similar to Type-I except that the leptons couple to Φ_1, instead of Φ_2, and the *flipped* model is similar to Type-II except that the leptons couple to Φ_2, instead of Φ_1. In all these models, the coupling of the light neutral scalar h to the vector bosons is equal to the Standard Model coupling modified by a factor of $\sin(\beta-\alpha)$, and the coupling of the heavy neutral scalar H to the vector bosons is modified by a factor of $\cos(\beta-\alpha)$. In the $pp\to H\to ZZ$ channel, at no point is there a direct coupling of the Higgs boson to leptons, so only the Type-I and II interpretations are considered in this work. Furthermore, only the neutral scalar bosons will be considered since the charged and pseudoscalar bosons do not decay to pairs of Z bosons.

The Higgs production modes at the LHC for 2HDMs are analogous to those in the Standard Model, with the dominant production modes being gluon–gluon fusion and vector-boson fusion as before. However, their relative contributions to the total cross section depend on the considered region of 2HDM phase space. For instance, the relative contributions from top and bottom quarks in the quark loop of the ggF production mode can change in Type-II 2HDMs, especially for large $\tan\beta$ where the bottom quark Yukawa coupling becomes large. In practice, 2HDMs can be probed experimentally by performing model-independent searches for a SM-like heavy Higgs boson in the ggF and VBF production modes separately and then interpreting the results in the context of Type-I and II 2HDMs under certain benchmark assumptions on the other free parameters of the model [79].

Narrow and large width approximations

While the models discussed above offer valuable benchmarks in searches for additional Higgs bosons, it is also useful to define a model-independent approximation applicable to a broad range of theories with extended Higgs sectors. We will consider both the *Narrow Width Approximation* (NWA), which assumes that the heavy Higgs width is negligibly small with respect to both its mass and the detector resolution, and the *Large Width Approximation* (LWA), which assumes the heavy Higgs width is some fixed fraction its mass. Both these approximations can be used to probe the parameter space of specific extended Higgs-sector models, or to provide model-independent constraints on any number of other BSM Higgs models that predict either a very narrow resonance or a resonance with a specific width.

2.3.2 Warped extra dimensions

Other solutions to the hierarchy problem have also been proposed through the inclusion of additional spatial dimensions. In these models, the gravitational force propagates along all dimensions while the Standard Model interactions are confined to our observable four-dimensional (4D) spacetime surface, resulting in the disparity between the strength of gravity and the other fundamental interactions. One of the first models of this form was proposed by Arkani-Hamed, Dimopoulos and Dvali (ADD) in 1998, in which the extra dimensions are taken to be compact and at the sub-millimetre scale [80–82]. The 4D Plank-scale is then related to the full $4+n$ dimensional Plank-scale by the physical scale R of the compactified dimensions:

$$M_{\mathrm{Pl}}^2 = M_{\mathrm{Pl},4+n}^{2+n} R^n \,. \tag{2.76}$$

Taking $M_{\mathrm{Pl},4+n} \sim 1\,\mathrm{TeV}$ at the electroweak scale yields $R \sim 10^{30/n-17}\,\mathrm{cm}$, thus the hierarchy problem can be resolved if $R \lesssim 1\,\mathrm{mm}$, or equivalently if $n \geq 2$. There is then no longer an energy-scale hierarchy in the overall theory but rather the appearance of one in our 4D universe. At the time when the ADD model was first proposed, the strength of the gravitational force had not been probed experimentally below the millimetre scale, accounting for why such extra dimensions had not yet been observed. However, this model modifies the $1/r^2$ law of gravity below the distance scale R. The accuracy of this law has since been verified experimentally to hold true at the scale of $O(10\,\mu\mathrm{m})$ [83, 84], indicating that up to $n = 2$ additional dimensions are excluded.

An alternative solution of this form is to allow for an additional dimension with a *non-factorizable* geometry, that is, where the metric of the four familiar dimensions and the position in the extra dimension are not independent. In this case, the extra, finite dimension can be described as a line segment between two four-dimensional branes, known as the *Planck* or *gravity brane* (where the gravitational force is concentrated and relatively strong) and the *TeV* or *weak brane* (where the Standard Model applies). Theories of this type are known as *Randall–Sundrum* (RS) models [85, 86]. The total five-dimensional metric for RS models is

$$ds^2 = e^{-2kr_c|\phi|}\eta_{\mu\nu}dx^\mu dx^\nu + r_c^2 d\phi^2 \,, \tag{2.77}$$

where $0 < \phi < \pi$ is the coordinate in the fifth dimension, πr_c defines the size of the new dimension and k is the warp factor. Introducing the Higgs mechanism to this model, where the Higgs doublet

is confined to propagate on the TeV brane and the Higgs boson mass is assumed to be the five-dimensional Planck mass $M_{\mathrm{Pl},5}$, the Higgs four-dimensional vacuum expectation value is given by $v \equiv e^{-\pi k r_c} v_0$, where v_0 is the five-dimensional vacuum expectation value. Taking v_0 on the order of the 5D fundamental mass scale $M_{\mathrm{Pl},5}$, the separation between the Planck mass and electroweak scales is produced by the metric when $kr_c \sim 11$, a more "natural" number, thus offering a solution to the hierarchy problem. The RS equivalent of Eq. (2.76) is then

$$M_{\mathrm{Pl}}^2 = \frac{M_{\mathrm{Pl},5}^3}{k} \left(1 - e^{-2\pi k r_c}\right) . \tag{2.78}$$

In the RS framework, the metric-tensor excitations in the usual four-dimensional space can be identified as spin-2 gravitons, while the excitation in the extra dimension corresponds to a spin-0 field, known as the *radion*, which is massless in the simplest scenario. However, this model lacked a mechanism to stabilize the size of the extra dimension, or equivalently the vacuum expectation value of the radion field. This problem was solved by introducing an additional bulk scalar with its interactions localized on the two ends of the extra dimension, which accounts for both the issue of the stability of the warped dimension and the lack of any observation of additional massless particles [87, 88]. This mechanism dynamically produces a radion vacuum expectation value and causes the radion field to acquire a mass term, which is typically smaller than the graviton mass.

In the original formulation of the RS model, the Standard Model fields are constrained to the TeV brane and only the graviton and radion propagate through the bulk. Since the bulk is finite in the fifth dimension, the wavefunctions can be decomposed into the 4D fields and the Fourier modes in the additional dimension. The visible particle content is then an infinite tower of modes of increasing mass for the graviton and radion, called *Kaluza–Klein* (KK) modes, with the first modes at the electroweak scale.

Furthermore, the Standard Model fields can be allowed to propagate into the fifth dimension bulk, which suppresses the coupling of the graviton and radion to quarks [88]. This is known as the *Bulk RS* model, and accounts for predictions of the original RS model that are inconsistent with constraints from flavour-physics experiments and precision electroweak measurements. Such models also predict KK modes of all Standard Model particles with higher masses. The graviton decay modes in the Bulk RS model are predominantly to pairs of top quarks, Higgs boson, and W/Z bosons [89], making it a natural benchmark model to consider alongside searches for heavy Higgs bosons decaying to a pair of Z bosons. The production cross sections and branching ratios for the graviton are shown in Figure 2.11.

The overall phenomenology of bulk RS models, besides the new KK masses, is determined by two free parameters: the size of the extra dimension r_c and the warp factor k. In the context of graviton searches at particle colliders, the only relevant parameters are the resonance mass and the warp factor, often given in units of the Plank mass as k/M_{Pl} [89]. Chapter 7 includes a discussion on a direct search for spin-2 bulk gravitons decaying to a pair of Z bosons using the ATLAS Run 2 dataset. Since the radion is phenomenologically similar to a Higgs boson, results from the heavy Higgs boson search can be reinterpreted in the context of a radion search.

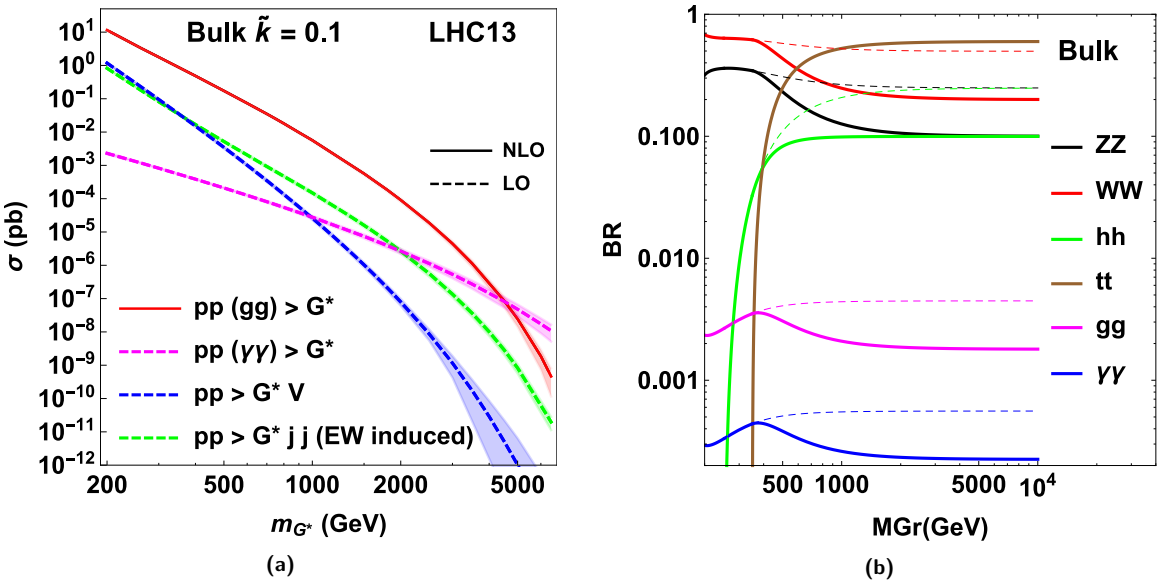

Figure 2.11: Bulk Randall–Sundrum graviton (a) production cross sections in $\sqrt{s} = 13\,\text{TeV}$ pp collisions and (b) branching ratios as a function of the graviton mass. For the branching ratios, the thick curves represent the predictions in a benchmark scenario where right-handed top quarks are localized very near the TeV brane and the left-handed top and other fermions are localized near the Planck brane, while the thin dashed lines represent the predictions where the G^* coupling to top quarks is ignored [89].

2.3.3 Status of heavy Higgs boson searches at the LHC

Thus far, no significant excess over the Standard Model predictions have been observed in direct searches for additional, heavy Higgs bosons at the LHC. Furthermore, precision measurements of the Higgs boson couplings, production cross sections and decay rates offer compelling evidence that the new particle observed by the ATLAS and CMS Collaborations is the Standard Model Higgs boson [72]. Together, these results place strong constraints on theories containing extended Higgs sectors. Figure 2.12, for example, shows exclusion regions using the combination of recent ATLAS results from direct searches for additional Higgs bosons and from measurements of the Standard Model production and decay rates. This combination uses the hMSSM approach as its benchmark, in which the Higgs sector of the MSSM is described in terms of only the parameters entering the tree-level expressions for masses and mixing plus the experimentally known value of $m_h \simeq 125\,\text{GeV}$, therefore the predictions for the properties of the MSSM Higgs bosons do not depend on the details of the unobserved SUSY sector [60].

For the specific case of searches for heavy Higgs-like resonances decaying to pairs of Z bosons, both the ATLAS and CMS Collaborations have reported on direct searches for this signature. The initial results from LHC Run 1 placed limits on heavy Higgs boson production up to a resonance mass of $m_H = 1000\,\text{GeV}$. The ATLAS analysis [91] considered the 4ℓ, $2\ell2\nu$, $2\ell2q$ and $2\nu2q$ final states, while the CMS analysis [92] considered the 4ℓ, $2\ell2\tau$, $2\ell2q$ and $2\ell2\nu$ final states. Later, using early LHC Run 2 data from 2015 and 2016 at centre-of-mass energies of $\sqrt{s} = 13\,\text{TeV}$, ATLAS reported on an extended search using the 4ℓ and $2\ell2\nu$ final states, setting limits up to $m_H = 1200\,\text{GeV}$, examples of which are shown in Figure 2.13 for the narrow-width signal hypothesis for both the ggF and VBF production modes [93]. This analysis also included limits on KK graviton production up to 2 TeV.

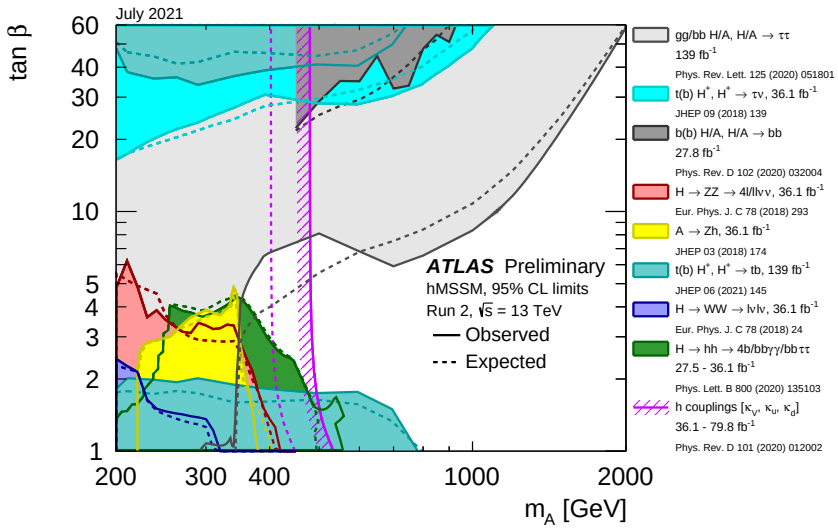

Figure 2.12: Regions of the $(m_A, \tan\beta)$ plane excluded at 95 % CL in the hMSSM via direct searches for heavy Higgs bosons and fits to the measured rates of observed Higgs boson production and decays by ATLAS using early LHC Run 2 data [90].

Similarly, CMS reported on its early Run 2 results using the 4ℓ, $2\ell 2q$ and $2\ell 2\nu$ final states, setting limits up to $m_H = 3000\,\text{GeV}$ [94].

Two small excesses were observed in the early Run 2 ATLAS analysis [93] in the $H \to ZZ \to \ell^+\ell^-\ell^+\ell^-$ channel at 240 and 700 GeV, each having a local significance of 3.6σ, with a global 1.3σ excess after combining with the $\ell^+\ell^-\nu\bar{\nu}$ channel. This result motivated a continuation of the search for an additional heavy Higgs boson in this channel using the complete Run 2 dataset. Moreover, advanced analysis techniques were employed to refine the search strategy, providing further improvements to the experimental sensitivity to these signatures. The objective to either confirm or refute these observed excesses is paramount to this thesis.

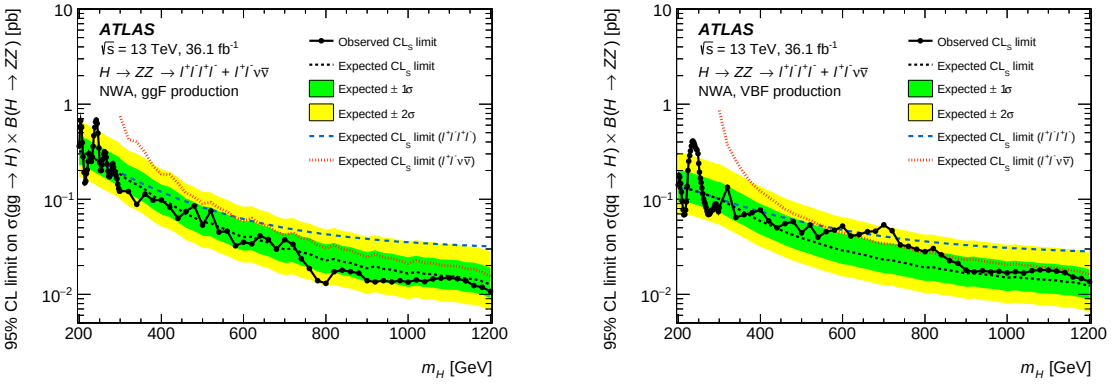

Figure 2.13: Upper limits at 95 % CL on the cross section times branching ratio in the search for a heavy spin-0 resonance as a function of the resonance mass m_H for (a) the ggF production mode and (b) for the VBF production mode using early Run 2 data recorded by ATLAS in 2015 and 2016 [93].

2.4 Summary

Our current understanding of the elementary particles in nature and their interactions via the electromagnetic, weak and strong forces is formulated as the Standard Model of particle physics. The Standard Model is a quantum field theory constructed to have an $SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge symmetry, where $SU(3)_C$ describes the symmetry group of quantum chromodynamics under transformations of colour charge, and where $SU(2)_L \times U(1)_Y$ describes the symmetry group of the unified electroweak theory under transformations of weak isospin for left-handed fermions and under transformations of weak hypercharge Y. The mass of the Standard Model gauge bosons and fermions is introduced to the theory while maintaining the required gauge invariance via the Higgs mechanism. This mechanism predicts an additional, massive scalar boson, the Higgs boson, first observed by the ATLAS and CMS Collaborations at the CERN LHC in 2012 with a mass of approximately 125 GeV.

In its most compact form, the Standard Model Lagrangian reads

$$\begin{aligned}\mathcal{L}_{\text{SM}} = &-\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + i\overline{\psi}\gamma^{\mu}D_{\mu}\psi \\ &+ |D_{\mu}\phi|^2 - \mu^2|\phi|^2 - \lambda|\phi|^2 \\ &+ g_d\overline{\psi}_L\phi\psi_R + g_u\overline{\psi}_L i\sigma_2\phi^*\psi_R + h.c..\end{aligned} \tag{2.79}$$

The first line encodes the kinetic and interaction terms between the fermions and gauge fields, the second describes the Higgs mechanism in relation to the massive gauge bosons, and the final line describes the Yukawa couplings of the down- and up-type fermions to the Higgs field, where *h.c.* denotes the hermitian conjugate of the preceding terms.

Despite the success of the Standard Model, it is not a complete theory of nature. A number of extensions to the Standard Model have been proposed to address several of its limitations, including theories with extended Higgs sectors, and with warped extra dimensions and gravity particles. Direct experimental constraints on a number of benchmark scenarios of these models using the complete ATLAS Run 2 pp dataset in final states with pairs of Z bosons will be presented in Chapter 7.

Chapter 3

The LHC and the ATLAS Detector

3.1 Introduction

Soon after the first proposals for the Large Electron–Positron Collider (LEP) at CERN, John Adams, then technical Director General of CERN, noted in 1977 that the LEP tunnel could be made wide enough to one day accommodate a superconducting proton collider with 3 TeV beam energies [95]. While electron–positron colliders had a history of great success in searches for Standard Model physics and precision measurements, the beam energy loss through synchrotron radiation restricted the size of future e^+e^- colliders, making a proton collider a natural alternative to reach higher energies with a similar geographical footprint.[1] Over the course of the following decades, this idea for a proton collider in the existing LEP tunnel eventually came to fruition as the Large Hadron Collider (LHC).

Around the same time, construction was underway on the Superconducting Super Collider (SSC) in the United States. With a centre-of-mass energy of 40 TeV, the SSC was almost three times more powerful than what could be built at CERN using the existing LEP tunnel, putting the LHC's future into question. Carlo Rubbia, who became the Director General of CERN in 1989, and others at CERN argued that in spite of its disadvantage in energy, the LHC would be competitive with the SSC by having a design luminosity an order of magnitude greater and could be built at a fraction of the cost given the existing accelerator infrastructure at CERN. Moreover, the LHC could be made more versatile by also functioning as a heavy-ion collider at little extra cost as a means to study quark–gluon plasma alongside the LHC *pp* physics program. The SSC was ultimately cancelled in 1993 making the physics case for building the LHC even stronger.

The approval of the LHC was nevertheless a tumultuous time, spanning several years of negotiations regarding the finances required to build and operate the machine. In December 1994, CERN Council approved a two-stage procedure for the construction of the LHC to curtail up-front costs: first, a "partial" machine that left out a third of all magnets would be constructed allowing for operation at a centre-of-mass energy of 9–10 TeV in 2004, followed by an upgrade to 14 TeV operations by 2008 [96]. Over the course of the following three years, CERN secured contributions from a

[1] An accelerating charged particle radiates energy, called *synchrotron radiation*, at a rate given by the *relativistic Larmor formula*. For a particle in a circular orbit with acceleration $a = v^2/R$, the rate of energy loss carried away by this radiation is $P \propto \gamma^4 a^2 = (E/mc^2)^4(v^2/R)^2$. The power loss due to synchrotron radiation for an electron accelerator is therefore $(m_p/m_e)^4 \sim 10^{13}$ times more than a proton accelerator of the same size and particle energy.

number of non-member states, including Japan, India, Russia, the United States and Canada, and in December 1997, CERN Council granted approval for LHC construction in a single stage with commissioning at full energy foreseen for 2005 [97].

Meanwhile, experimental collaborations had emerged and submitted proposals for detectors to be built at the LHC. The first of these proposals were discussed at the 1992 CERN–ECFA meeting "Towards the LHC Experimental Programme" in Evian, which included four general-purpose detectors, three b-physics experiments, three heavy-ion experiments and two neutrino experiments. Of these, four experiments were ultimately approved: the general-purpose ATLAS[2] [98, 99] and CMS [100, 101] experiments in January 1997, the heavy-ion experiment ALICE [102, 103] in February 1997, and the b-physics LHCb experiment [104, 105] in September 1998.

After the final magnet was lowered underground in April 2007, the LHC circulated its first proton beams on September 10, 2008. Just over a week later on September 19, disaster struck when a fault in an electrical bus connection resulted in mechanical damage to the neighbouring magnets and the release of helium from the magnet cold mass into the tunnel. In total, 53 magnets required refurbishing or replacing, the last of which was lowered into the tunnel in April 2009. Later that year on November 20, the LHC was once again circulating proton beams, followed a few days later by the first collisions at the injection energy of 900 GeV. In order to reduce the risk of another incident, the LHC operations team opted to reduce the collision energy to half the design energy of 14 TeV. After a short technical stop, the first 7 TeV collisions were achieved on March 30, 2010, marking the start of LHC Run 1. The LHC continued 7 TeV running until October 2011 before restarting the following year at higher energy with the first 8 TeV collisions taking place on April 5, 2012. A two-year technical stop followed the end of Run 1 in February 2013 to prepare the machine for running at higher energy once more, with the first 13 TeV collisions delivered on June 3, 2015, marking the beginning of LHC Run 2. A collision energy of 13 TeV was chosen to begin operations sooner, as the process of training[3] the LHC magnets to achieve the design energy of 14 TeV would have significantly delayed the start of Run 2 and may have risked another destructive quenching incident.

This chapter describes the experimental apparatus used to investigate the existence of the theorized heavy resonances decaying to pairs of Z bosons discussed in the previous chapter. First, we will describe the LHC, which produces the proton–proton collisions, followed by the ATLAS detector, which is used to detect the products of these collisions. This discussion will focus on the state of the LHC and the ATLAS detector as they were during Run 2 in the years 2015–2018.

3.2 The Large Hadron Collider

The Large Hadron Collider [106–109] is a two-ring superconducting hadron accelerator and collider operated by CERN, installed in the former 26.7 km LEP tunnel on the outskirts of Geneva, Switzerland. The LHC comprises several thousand superconducting magnets, radio-frequency (RF) accelerating cavities, large-scale cryogenics, and other support structures located throughout the

[2]The ATLAS Collaboration came to be through the merger of the EAGLE (**E**xperiment for **A**ccurate **G**amma, **L**epton and **E**nergy Measurements) and ASCOT (**A**pparatus with **S**uper **CO**nducting **T**oroids) Collaborations.

[3]To qualify a superconducting LHC magnet for use, it must be *trained* in a process that involves repeatedly increasing the current through the magnet's coils until it *quenches*, that is, heats beyond its superconducting threshold. After a number of training cycles, the components of the magnet "settle in" and can be operated routinely at their nominal current.

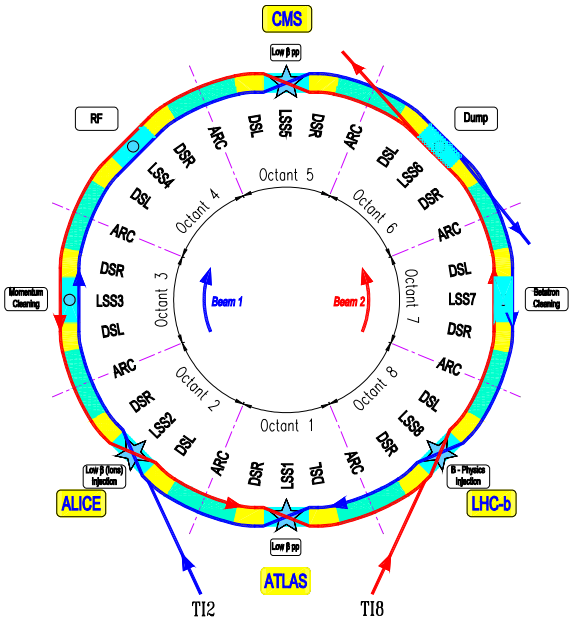

Figure 3.1: Schematic layout of the LHC showing its eight long straight sections and arcs. Beam 1 circulates clockwise and Beam 2 counter-clockwise [106].

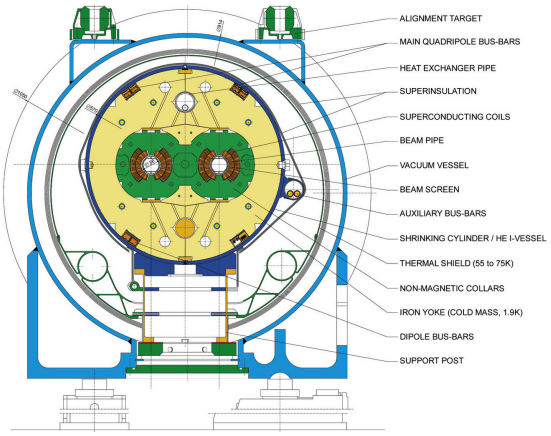

Figure 3.2: Diagram showing the cross section of an LHC dipole magnet with cold mass and vacuum chamber [107].

tunnel. Inside the main LHC ring, two particle beams travel in opposite directions within two separate beam pipes kept in ultrahigh vacuum conditions. These beams cross at four interaction points (IPs) located within the four LHC detectors.

As a circular collider, the LHC requires powerful magnets to maintain the circular trajectories of the high-energy particles it accelerates. Compared to a linear accelerator, the benefit of this circular design is that after injection into the machine, the particle beams pass through each accelerator element many times, allowing the beam energy to be steadily increased. For this reason, colliders such as the LHC are classified as *synchrotrons*, since the magnetic field that bends the particles into their circular orbit is *synchronized* to increase with the beam energy. The circular design also allows for multiple opportunities for the circulating beams to collide, thus such colliders are often also referred to as *storage rings*.

The LHC is in fact not entirely circular, but rather consists of a series of eight alternating *Long Straight Sections* (LSS) and *arc* sections, as shown in Figure 3.1. The regions of the machine spanned between the centres of two adjacent arcs are called *octants* and are numbered 1 through 8 (with ATLAS at Point 1). Each straight section serves as a region for insertion into the experiments or into other LHC utilities, including the two 400.8 MHz RF systems at Point 4 that accelerate the beams, two "beam cleaning" sites that collimate the beams through the removal of particles with large oscillation amplitudes in the longitudinal direction (Point 3; "momentum cleaning") and in the transverse direction (Point 7; "betatron cleaning"), and the beam dump site at Point 6 where the beams are extracted from the machine. The experimental halls are located at Points 1 (ATLAS), 2 (ALICE), 5 (CMS) and 8 (LHCb), with Points 2 and 8 also serving as the insertion regions where beams 1 and 2 are injected into the machine, respectively.

Each LHC arc consists of 46 *arc half-cells* comprising three main twin-aperture superconducting

dipole magnets responsible for bending the beams along their circular trajectories. A cross-sectional view of an LHC dipole magnet is given in Figure 3.2. This twin-aperture design, consisting of two sets of coils and beam channels within the same mechanical structure and cryostat, was employed given the spatial restrictions in the original LEP tunnel that would not have allowed two separate rings of magnets.[4] The magnet coils are made up of NbTi cables that are cooled using superfluid helium to 1.9 K, allowing peak magnetic fields of 8.33 T for beam energies of 7 TeV. The half-cells also contain quadrupole magnets used for focusing the beam, as well as sextupole, octupole, and decapole magnets, responsible for additional focusing and correcting other beam artefacts such as chromatic aberrations. Altogether the periodic arc cells form the *LHC lattice*.

In order to bring the beams into collision at the interaction points, special separation and recombination dipole magnets are placed on either end of an experimental insertion to generate the beam crossing, followed by assemblies of three quadrupole magnets called the *inner triplets*[5] responsible for squeezing the beams to increase the number of collisions in each beam crossing. Horizontal beam dispersions arise as a result of the separation and recombination dipole magnets and from the geometrical structure of the arcs, therefore the transition regions between the LHC arcs and long straight sections are equipped with *dispersion suppressors* (labelled "DSL" and "DSR" in Figure 3.1) to cancel these dispersions using a series of four quadrupole and eight dipole magnets. The insertion regions are also equipped with orbit-correcting dipole magnets used both in regular operation and in special runs to separate the beams at the interaction points (the reasons for doing so will be made clear in the following sections).

The LHC cannot accelerate the particle beams from rest and must instead take advantage of the full CERN accelerator complex (Figure 3.3) to increase the beam energy in stages. The accelerator chain for proton–proton operations is described in the section below.

3.2.1 Proton accelerator chain

The protons circulated in the LHC begin as hydrogen gas molecules that are broken into their constituent atoms and stripped of their electrons by a duoplasmatron located at one end of the LINAC2 linear accelerator.[6] The resulting protons are then accelerated to an energy of 50 MeV by RF cavities in LINAC2 and then injected into the Proton Synchrotron Booster (PSB). The PSB consists of four independent rings that simultaneously accelerate the batches of protons from LINAC2 to energies of 1.4 GeV for injection into the next stage of the accelerator chain. Proton "bunches" begin to take shape in the PSB as these batches of protons synchronize with the RF cavity acceleration frequency.

After the bunching of the protons and acceleration to 1.4 GeV, the proton beam is injected into the Proton Synchrotron (PS), where eight bunches from two PSB injections are accelerated to 26 GeV. As the eight batches of protons are accelerated, they are also merged ("compressed") into four and then split into smaller bunches of protons based on the RF cavity harmonics in the PS in a process known as the *Batch Compression Merging and Splitting* (BCMS) scheme [112]. Under

[4] Separate rings of magnets are unnecessary for oppositely charged particles, as in the case of LEP.

[5] In fact, the inner triplets consist of four quadrupole magnets, not three, but the central two magnets form one functional entity.

[6] LINAC2 was replaced by LINAC4 [111] in 2020 in preparation for LHC Run 3. LINAC4 instead accelerates H^- ions to 160 MeV before stripping the electrons away and injecting into the Proton Synchrotron Booster, allowing for greater beam *brightness* (the ratio of beam current and emittance).

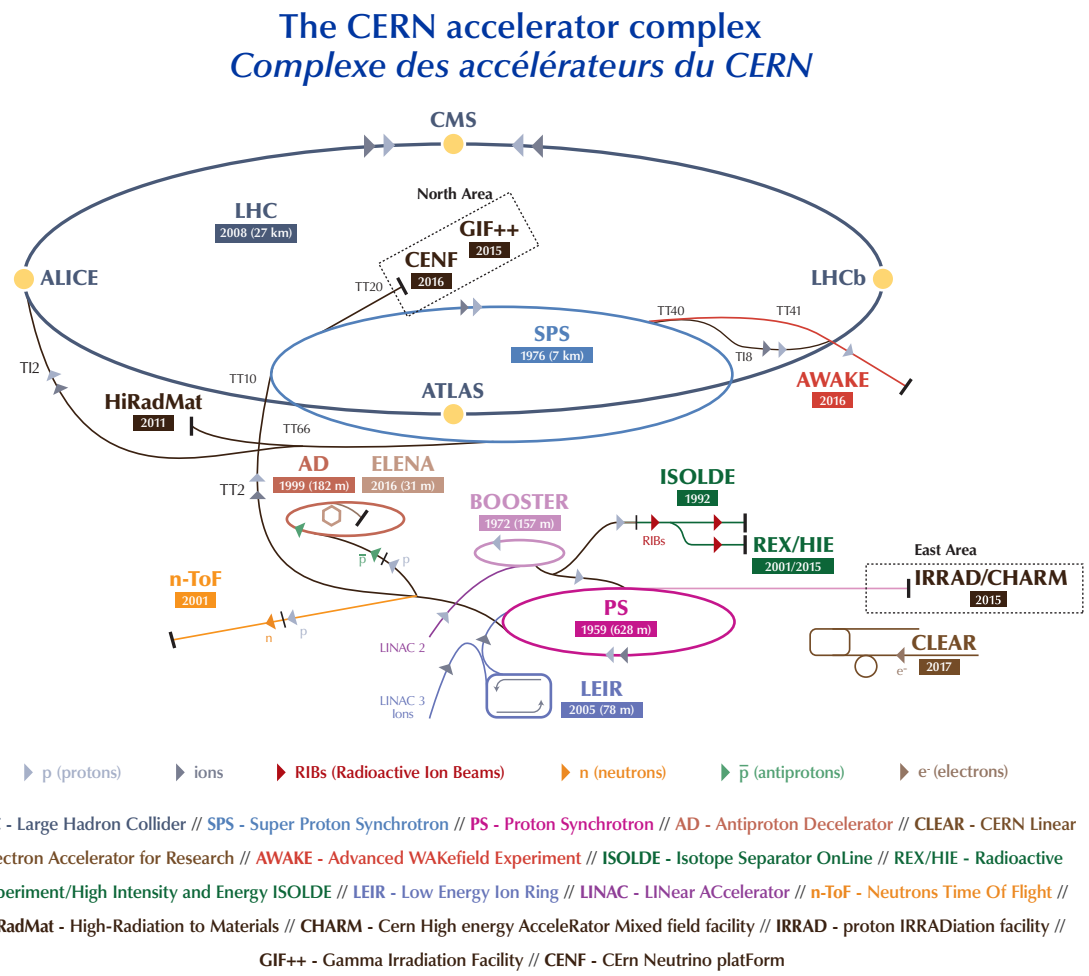

Figure 3.3: The CERN accelerator complex in Run 2 [110].

normal operations, this splitting eventually produces *trains* of 48 bunches with a spacing between the bunches of 25 ns, although fewer bunches may also be produced for special runs or during a ramp-up period after a technical stop.[7] The bunch spacing from the PS defines the final bunch spacing of the beams circulated in the LHC, with the 25 ns spacing resulting in 3564 nominal LHC *bunch slots*. Several of these bunch trains are then injected into the Super Proton Synchrotron (SPS), where the beam is accelerated to 450 GeV. From there, multiple bunch trains are then injected into the LHC via two transfer lines in both the clockwise (Beam 1) and counter-clockwise (Beam 2) directions. Gaps are required in between the trains to account for the rise times of the kicker magnets that inject and dump the beams. All together, this injection scheme allows for a total of 2556 bunches per beam in the LHC, with 2544 of these bunch pairs colliding at the ATLAS and CMS interaction points.[8]

Following their injection into the LHC, the proton beams undergo a ramp period that takes

[7] The BCMS scheme, first deployed in 2016, results in shorter bunch trains than the original design value of 72 bunches, but with the advantage of having a reduced transverse beam emittance.

[8] The original PSB bunching scheme producing 72-bunch trains allowed for a design value of 2808 bunches per beam.

each proton beam from 450 GeV to the target 6.5 TeV. Upon reaching the target energy, the beams must be focused (referred to as the *squeeze* beam mode) and then adjusted (beam mode *adjust*) into their final colliding configuration. Once the beams are colliding in a stable configuration, the LHC declares *stable beams*, and the detectors are able to begin their nominal data collection until the beams are dumped from the machine. The period from the injection of the beams into the machine until the beam dump is referred to as a *fill*.

3.2.2 LHC luminosity

There are two main figures of merit for any particle collider. The first is the energy available in the collisions to produce heavier particles or to probe smaller scales, and is generally quoted in the centre-of-mass system using the Lorentz-invariant Mandelstam variable $\sqrt{s} = |p_1 + p_2|$, where p_1 and p_2 are the four-momenta of the two colliding particles. The second is the *luminosity*,[9] which quantifies the ability of the collider to deliver a statistically significant sample of a given class of events. The luminosity is a property of the colliding beams, measuring the number of particle interactions per unit of time and per unit of area, and is conventionally defined as the ratio

$$\mathcal{L} = \frac{R}{\sigma}\,, \tag{3.1}$$

where R is the rate of events produced per unit of time for a process with cross section σ. Specifically, this equation defines the *instantaneous luminosity*, and is generally expressed in units of $\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$. In terms of LHC beam parameters, the luminosity is given by

$$\mathcal{L} = \frac{N_b^2 n_b f_r \gamma}{4\pi\epsilon_n \beta^*} F\,, \tag{3.2}$$

where N_b is the number of particles per bunch (also called the *bunch population*), n_b the number of bunches per beam, f_r the revolution frequency (11 245.5 Hz for the LHC), γ the relativistic Lorentz factor, ϵ_n the normalized transverse beam emittance, β^* the beta function at the interaction point, and F the geometric luminosity reduction factor due to the crossing angle at the IP, given by

$$F = \left[1 + \left(\frac{\theta_c \sigma_z}{2\sigma^*}\right)^2\right]^{-\frac{1}{2}}\,, \tag{3.3}$$

where θ_c is the full crossing angle at the IP, σ_z the RMS bunch length, and σ^* the transverse RMS beam size at the IP [107]. The beam crossing angle (typically of order 120–160 µrad) avoids so-called *parasitic collision points* during nominal physics running due to the large number of colliding bunch pairs in the shared beam pipe within the experiments. The above expressions assume round beams, with $\sigma_z \ll \beta$, and with equal beam parameters for both beams. In practice, the beam dynamics are considerably more complicated and the individual beam parameters of each beam cannot be measured to high precision, therefore alternative methodologies are required to determine the delivered luminosity. A full discussion of the ATLAS luminosity-determination methodology is given in Chapter 6. Particle colliders such as the LHC run over long periods to produce and record large numbers of events to study rare processes with sufficient statistical precision. The

[9] The word "luminosity" stems from the Latin word *lumen* (light) or *luminosus* (full of light) to describe how "bright" the interaction point is with collision products.

Table 3.1: LHC design and operating parameters for *pp* collisions in Run 2. Design values are taken from Volume 1 of the LHC design report and Run 2 values are as reported in Ref. [113]. The luminosity data in the bottom rows are specific to ATLAS and are taken from Ref. [115]. The values shown are representative of the best accelerator performance during normal physics operation at the beginning of stable beams. The different values of transverse beam emittance in 2016 correspond to the nominal and BCMS injection schemes, respectively. In 2017, the LHC was run with two bunch configurations: standard 25 ns bunch train operation with long trains, and "8b4e", denoting a pattern of eight bunches separated by 25 ns followed by a four bunch-slot gap. Values are given for both configurations.

Parameter	Design	2015	2016	2017	2018
pp centre-of-mass energy ($\sqrt{s}$) [TeV]	14	13	13	13	13
Number of colliding bunch pairs (n_b)	2808	2232	2208	2544/1909	2544
Bunch spacing [ns]	25	25	25	25/8b4e	25
Bunch population (N_b) [10^{11} protons]	1.15	1.1	1.1	1.1/1.2	1.1
Normalized transverse beam emittance (ϵ_n) [µm rad]	3.75	3.5	3.5/2.2	2.2	1.9
Optical β-function at IP β^* [m]	0.55	0.80	0.40	0.30	0.30–0.25
Peak luminosity ($\mathcal{L}_{\text{peak}}$) [$10^{33}\,\text{cm}^{-2}\,\text{s}^{-1}$]	10	5.02	13.8	20.9	21.0
Peak number of inelastic interactions/crossing ($\langle\mu\rangle$)	~ 25	~ 16	~ 41	$\sim 45/60$	~ 55
Luminosity-weighted mean inelastic interactions/crossing	–	13	25	38	36
Total delivered integrated luminosity [fb^{-1}]	–	4.0	38.5	50.2	63.4

figure of merit here to quantify the "size" of the resulting dataset is the *integrated luminosity*, denoted by $\mathcal{L}_{\text{int}} = \int \mathcal{L}\, dt$, which refers to the instantaneous luminosity accumulated over a certain time interval and is typically quoted in units of cm^{-2}, or equivalently in inverse barns, microbarns (μb^{-1}), femtobarns (fb^{-1}), etc.[10]

3.2.3 LHC Run 2 operations

LHC Run 2 saw steadily increasing machine performance from its beginning in 2015 until the last proton beams were circulated in 2018 [113]. These performance improvements were achieved through the implementation of a number of machine optimizations to reduce the transverse beam sizes, thereby increasing the delivered luminosity. For instance, the BCMS scheme discussed above reduced the transverse beam emittance from approximately 3.5 to 2.2 µm rad, and stronger beam focusing in the ATLAS and CMS interaction regions allowed for a reduction in the β^* parameter from approximately 80 cm in 2015 to 30 cm in 2018. Two "anti-levelling" procedures were also introduced to regain some of the lost luminosity as the beam naturally degrades through the loss of protons from collisions and the beam emittance growth over the course of the fill [114]. The first of these was introduced in 2017 to steadily decrease the beam crossing angle, thereby reducing the geometric luminosity reduction factor F of Eq. (3.3). The second procedure was introduced in 2018 to reduce the β^* parameter in two steps from 30 to 25 cm. Together, these procedures provided an integrated luminosity gain of 3–4 % per fill. Table 3.1 shows the typical LHC *pp* operating parameters in each year of Run 2, along with the corresponding design values, together with the total delivered integrated luminosity.

LHC Run 2 was, by and large, remarkably successful, setting world records in both beam energy and delivered luminosity. For instance, the LHC consistently delivered peak luminosities of $2.1 \times 10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$ to both the ATLAS and CMS Experiments towards the end of Run 2, exceeding the design luminosity by over a factor of 2. LHC operations were not without incident, however. In

[10] A useful conversion factor between luminosity units is $10^{33}\,\text{cm}^{-2}\,\text{s}^{-1} = 1\,\text{nb}^{-1}\,\text{s}^{-1} = 0.0036\,\text{fb}^{-1}\,\text{h}^{-1}$.

June 2017, the LHC developed a significant vacuum problem in a magnet interconnection.[11] The problem was eventually identified as originating from small amounts of frozen air that had built up along the beam screen that, in the presence of the high intensity beams and electron clouds, detached from the vacuum chamber and fell towards the beam where they evaporated, creating a plasma that destabilized the beams [116]. This problem lead to frequent, unplanned beam dumps (approximately 70 in 2017) that significantly limited the time available for running with stable beams.

Two approaches were taken to mitigate the issue. First, a new filling scheme of eight bunches separated by 25 ns followed by a four bunch-slot gap, called "8b4e", replaced the standard 25 ns bunch trains to suppress the formation of electron clouds. Second, a solenoid magnet was installed around the effected magnet interconnection to reduce *multipacting*, a beam-induced effect that results in exponentially increasing secondary emissions of electrons that resonate between the vacuum chamber walls with the RF frequency [117]. With these two measures in place, the original performance of the LHC was recovered. However, the 8b4e scheme limited the number of bunches in each fill until the end of 2017 data taking when the air was finally evacuated. While the loss in luminosity due to the reduced number of colliding bunches could be compensated for by increasing the bunch population and lowering the transverse beam sizes, it would have significantly increased the mean number of interactions per bunch crossing, or *pileup*, to the point that the detector performance would be degraded. The solution was to use luminosity levelling via the separation and steady realignment of the beams at the interaction points, which reduces the peak luminosity provided, but maintains high luminosity for a longer time.

In total, the LHC delivered an integrated luminosity of approximately $160\,\text{fb}^{-1}$ in proton–proton collisions with a centre-of-mass energy of 13 TeV to both ATLAS and CMS during Run 2. The LHC also provided collisions to the LHCb and ALICE experiments, but with different luminosities to support their substantially different physics programs. The LHC also dedicated several months of operation throughout Run 2 to heavy-ion running and to a variety of other special runs, including lead–lead collisions, lead–proton collisions, $\sqrt{s} = 5\,\text{TeV}$ proton–proton collisions and low-pileup 13 TeV proton–proton collisions for precision W and Z boson measurements.

3.3 The ATLAS Detector

ATLAS [99] is a multipurpose particle detector installed at LHC Point 1 with a forward–backward symmetric cylindrical geometry measuring approximately 44 m long and 25 m tall. The ATLAS detector is designed to be *hermetic*, that is, it is instrumented over a solid-angle coverage of nearly 4π to surround the collision point as completely as possible. The detector consists of concentric layers of specially tailored subdetectors around the interaction point, each designed to measure specific properties of the plethora of particles emitted from the collisions. The layers are arranged roughly in two groups, the *barrel* and the two *endcaps*, to ensure that the particles traverse the detector in a roughly perpendicular manner. The barrel layers are cylindrical structures surrounding the central portion of the detector aligned coaxially to the beam axis and the endcap layers are disk-like structures placed further along and orthogonal to the beam axis.

Figure 3.4 shows a cut-away view of the full ATLAS detector. Closest to the interaction point

[11] The source of this problem is affectionately referred to by the LHC community as the 16L2 "Gruffalo" after the character from the 1999 children's book by Julia Donaldson of the same name.

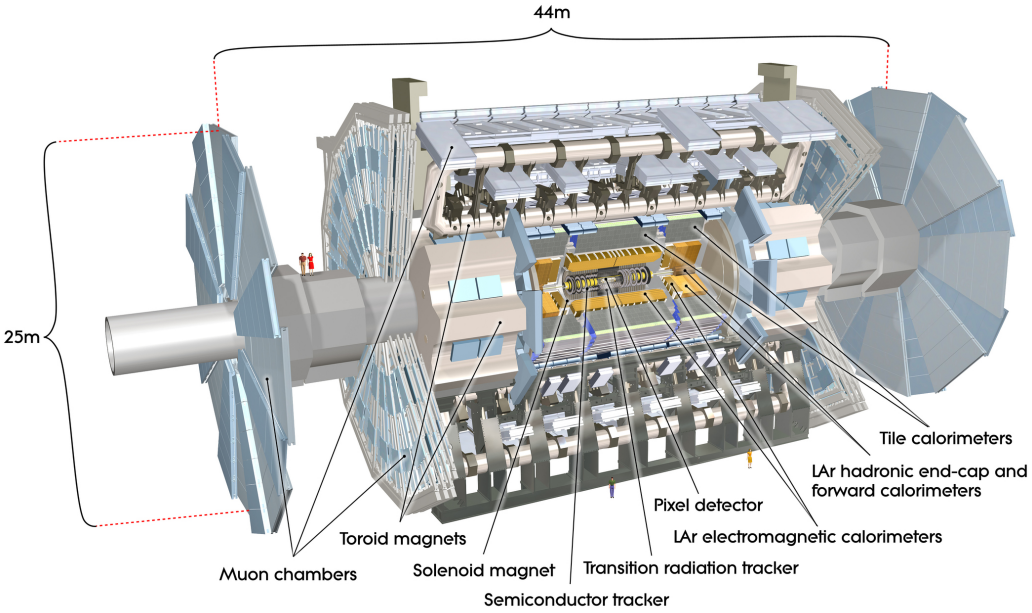

Figure 3.4: Cut-away view of the ATLAS detector [99].

is the inner tracking detector used for recording the trajectories of charged particles with high precision. The Inner Detector is surrounded by a thin superconducting solenoid providing a 2 T magnetic field in the direction of the beam axis. This field causes charged particles leaving the interaction point to bend azimuthally about the beam axis, with the curvature of their tracks providing a means to measure the particles' charge and momenta. Beyond the solenoid are the calorimeter systems used for measuring the particle energies. ATLAS employs sampling-calorimeter technology in which the traversing particles deposit most of their energy in an absorber material that is interleaved with an active medium to produce the signal. First are the finely segmented lead/liquid-argon (LAr) electromagnetic calorimeters, followed by a steel/scintillator-tile hadronic calorimeter in the barrel region and a copper/LAr hadronic calorimeter in the endcap region. The most forward regions in the endcaps are also instrumented with LAr calorimeters, with copper and tungsten as the absorber material. A muon spectrometer system incorporating large superconducting toroidal air-core magnets surrounds the calorimeters. The muon spectrometer includes three layers of wire chambers providing precision muon tracking as well as dedicated fast chambers for triggering purposes. The field of the toroidal magnets is oriented in the azimuthal direction such that the outgoing muons bend in the direction of the beam axis, allowing for measurements of their charge and momenta.

Each detector subsystem, as well as the data-acquisition and triggering systems, are described in detail in the section below. Particular attention will be paid to the liquid-argon calorimeters as they pertain to the LAr "energy-flow" luminosity measurement presented in Chapter 6 (and the LAr "gap-current" luminosity measurements based on the LAr high-voltage system included in Appendix A). We will begin this discussion with a description of the ATLAS coordinate system and its naming conventions.

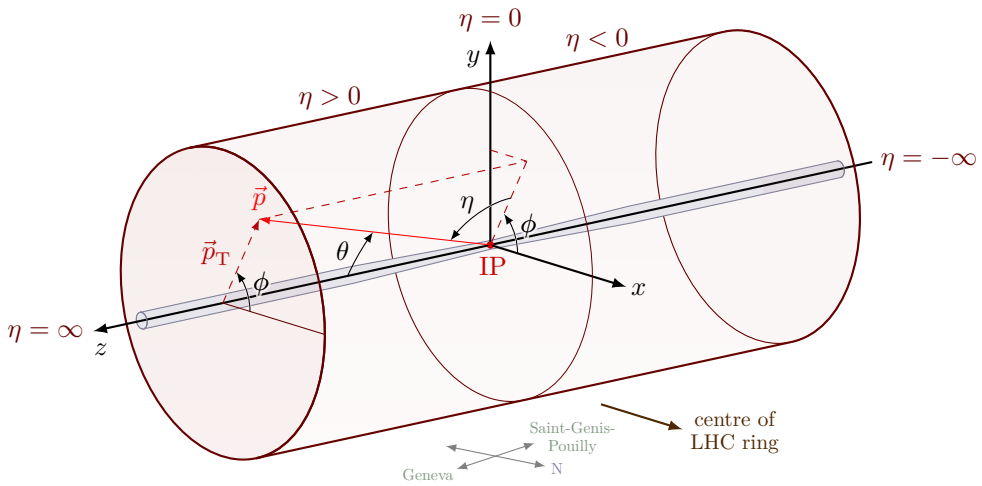

Figure 3.5: ATLAS coordinate system.

3.3.1 Coordinate system and kinematic variables

ATLAS describes collision events within the detector using a right-handed coordinate system with the origin at the nominal interaction point (IP) in the centre of the detector and the z-axis along the beam pipe, as shown in Figure 3.5. The x-axis points from the IP to the centre of the LHC ring and the y-axis points vertically (towards the surface). The $+z$ region of the detector is referred to as the "A-side" and the $-z$ region as the "C-side".[12] It is generally more convenient to describe particle trajectories from the IP using spherical coordinates in terms of the polar angle θ and the azimuthal angle ϕ around the beam axis, defined such that $-\pi < \phi \leq \pi$ with $\phi = 0$ pointing along the x-axis. Cylindrical coordinates are occasionally used for certain detector components in terms of their radial distance R from the beam axis, azimuthal angle ϕ and longitudinal position z.

It is also convenient to define quantities that are invariant under Lorentz boosts along the beam axis in order to describe events whose centre-of-mass frames are boosted by arbitrary velocities with respect to the lab frame. For instance, the *rapidity* of a particle with energy E and component of momentum along the beam axis p_z, defined as

$$y \equiv \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right) , \tag{3.4}$$

is a dimensionless quantity that describes the boost along the z-axis that takes an observer from the lab frame to a frame in which the particle moves only perpendicular to the beam. The difference between the rapidities of two outgoing particles is invariant with respect to Lorentz boosts along the z-axis, making it an especially useful quantity when describing the angular separation between particles. However, precision measurements of a particle's rapidity is challenging as it depends on both the particle's energy and momentum, therefore and alternative quantity is generally used that

[12] This nomenclature originates from labelling the detector's three main segments as A, B, C, where B is the **B**arrel (middle) region and A and C are the two endcaps. The colloquial mnemonic to recall the orientation of the ATLAS detector is that the A-side is towards the Geneva **A**irport and the C-side is towards **C**harly's Pub in Saint-Genis.

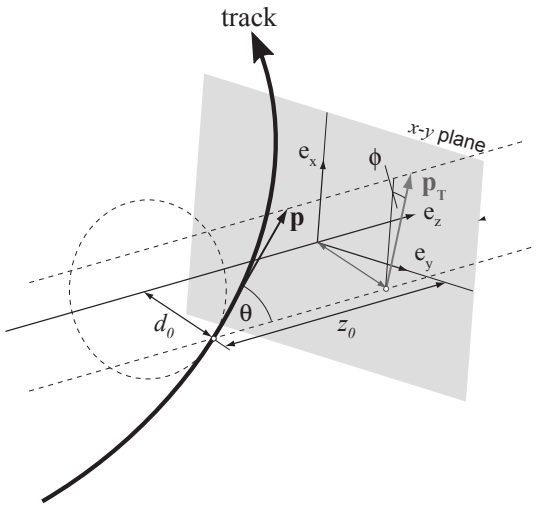

Figure 3.6: ATLAS tracking coordinate system in the "perigee" representation [118]. In this representation, a track's coordinates are given in terms of its polar angle θ, azimuthal angle ϕ, transverse impact parameter d_0 and longitudinal impact parameter z_0, with the origin centred on the primary vertex associated to the track.

depends only on the detector geometry, called *pseudorapidity*, defined as

$$\eta \equiv -\ln\left[\tan\left(\frac{\theta}{2}\right)\right] . \tag{3.5}$$

In the limit that the particle is massless or is at very high energy, $\eta \simeq y$. Other quantities invariant under Lorentz boosts along the beam axis are those defined in the transverse (x–y) plane, including the *transverse momentum* $\vec{p}_{\mathrm{T}}$, *transverse energy* $\vec{E}_{\mathrm{T}}$, and *transverse mass* m_{T}, defined as

$$\begin{aligned} \vec{p}_{\mathrm{T}} &\equiv (p_x, p_y)\,, \\ \vec{E}_{\mathrm{T}} &\equiv E\frac{\vec{p}_{\mathrm{T}}}{|\vec{p}|}\,, \\ m_{\mathrm{T}}^2 &\equiv m^2 + p_x^2 + p_y^2\,. \end{aligned} \tag{3.6}$$

The four-momentum of an outgoing particle in terms of these quantities is

$$(E, p_x, p_y, p_z) = (m_{\mathrm{T}}\cosh y,\ p_{\mathrm{T}}\cos\phi,\ p_{\mathrm{T}}\sin\phi,\ m_{\mathrm{T}}\sinh y)\,. \tag{3.7}$$

Finally, the angular distributions of particles in the detector are normally given in $\eta \times \phi$ space, and the measure of angular distance between particles is given by the quantity

$$\Delta R \equiv \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}\,. \tag{3.8}$$

Tracking coordinate system

Collisions do not necessarily occur at the nominal interaction point, but rather at any number of positions in the overlapping region of the crossing bunches. These interaction points where two particles collide are called *primary vertices*. For reasons discussed in later sections, reconstructed

tracks in the ATLAS Inner Detector may result in particle trajectories associated to but displaced from the primary vertex. The common origin of these displaced tracks is called a *secondary vertex*. The magnitude of this displacement is typically given in terms of a track's transverse and longitudinal *impact parameters* d_0 and z_0, respectively, as illustrated in Figure 3.6. The transverse impact parameter (d_0) is the shortest distance between the primary vertex and a track in the transverse plane. The longitudinal impact parameter (z_0) is the separation along the beam axis between the track and the primary vertex at the point on the track defined by the transverse impact parameter. The longitudinal impact parameter of a track is often reported as $z_0 \sin\theta$, which gives the shortest distance between the track and the primary vertex in the R–z plane. This parameterization of the track coordinates constitutes the *perigee* representation, so-named since it describes the track at the point of closest approach to the primary vertex.

3.3.2 Inner Detector

The ATLAS Inner Detector [119, 120] is responsible for the reconstruction of the trajectories of charged particles and is composed of four concentric subdetectors as shown in Figures 3.7 and 3.8. The Inner Detector is designed to provide hermetic tracking coverage over the range $|\eta| < 2.5$ with excellent momentum resolution and vertex-reconstruction performance.

A particle originating from the interaction point traverses successively the inert 0.8 mm thick beryllium beam pipe, the four layers of the silicon Pixel detector, the four layers of the silicon-microstrip Semiconductor Tracker (SCT), and finally the straw tubes of the Transition Radiation Tracker (TRT). The innermost layer of the Pixel detector was added during the long shutdown period between LHC Run 1 and 2 and is called the Insertable B-Layer (IBL). The entire Inner Detector is immersed in a 2 T field provided by the central superconducting solenoid magnet [121].

The IBL, Pixel and SCT detectors all follow similar detection properties, namely the excitation of p–n junctions in silicon. As charged particles pass through a portion of the silicon, they deposit a small fraction of their energy and ionize the material, creating electron–hole pairs. The detectors contain a reverse bias voltage that causes the electrons to drift through the bulk to the positive voltage connection where the current is measured by external electronics. The current depends on the energy of the incident particle, as a higher momentum particle deposits more energy and produces more electron–hole pairs, therefore a *hit* is only registered if the current is above a given threshold. The TRT operates on a similar principle except its sensitive material is a gas rather than silicon. Moreover, the TRT also makes use of the property of transition radiation for particle identification. A further description of each tracking subdetector is given below.

Insertable B-Layer

The IBL [123, 124] is the first and most-recently installed layer of the Inner Detector. It is designed to provide an additional radiation-hard silicon layer to the pixel detector, allowing for greater robustness against fake-track reconstruction in the high-pileup environment during LHC Run 2 with luminosities greater than $1.0 \times 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, and for improved vertex reconstruction and resolution of the track transverse and longitudinal impact parameters d_0 and z_0. The precision with which these impact parameters can be measured is of particular importance for the identification of jets originating from decays of the relatively long-lived ($\sim 10^{-12}$ s) b hadrons (Section 5.3.6), hence the name "B-Layer".

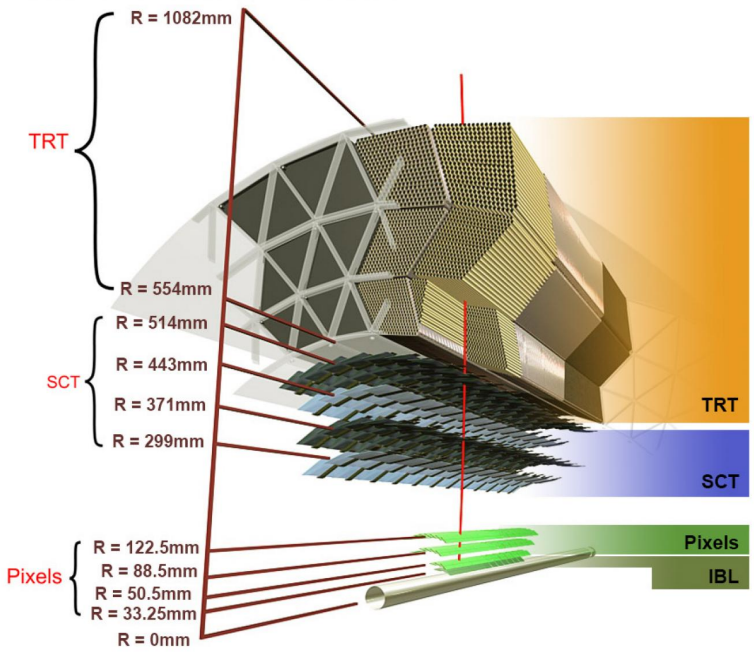

Figure 3.7: Cut-away view of the barrel region of the ATLAS Inner Detector showing the sensor positions [122]. A particle originating from the interaction point (with a trajectory shown by the solid red line) traverses successively the beryllium beam pipe, the Insertable B-Layer (IBL), the three silicon layers of the Pixel detector, the four silicon layers of the Semiconductor Tracker (SCT), and finally the straw tubes of the Transition Radiation Tracker (TRT).

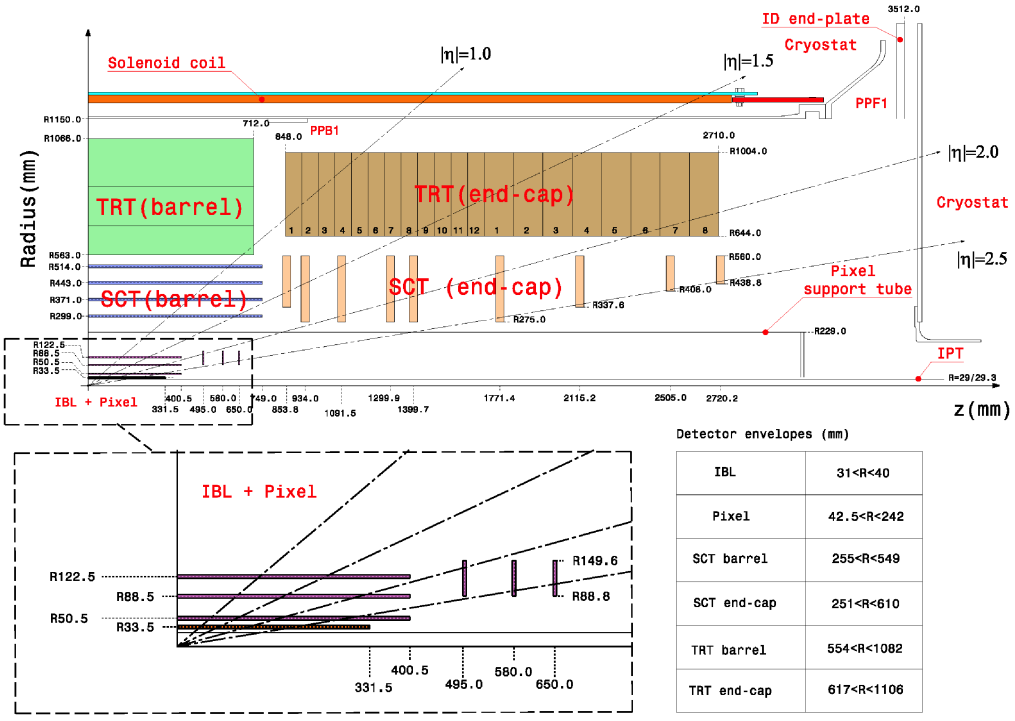

IBL	31<R<40
Pixel	42.5<R<242
SCT barrel	255<R<549
SCT end-cap	251<R<610
TRT barrel	554<R<1082
TRT end-cap	617<R<1106

Figure 3.8: Schematic layout of the ATLAS Inner Detector [124].

The IBL sits just outside of the beam pipe, spanning 330 mm along the beam axis and positioned at a radial distance R between 31 and 40 mm. The IBL contains two types of pixel sensors: planar sensors in its central region and 3D sensors in its high-η region. The size of the nominal IBL pixel sensors is $50 \times 250\,\mu\text{m}^2$ (in $\phi \times z$), providing highly granular and precise measurements of charged particle trajectories.

Pixel

The Pixel detector [125, 126] is comprised of three concentric barrel layers of semiconductor staves at radial distances of 50.5, 88.5 and 122.5 mm extending to $|z| = 400.5$ mm, and two endcap sections consisting of three semiconductor disks held perpendicular to the beam axis at distances of $|z| = 495$, 580 and 650 mm. The Pixel detector provides high-granularity measurements of charged particle positions in the region $|\eta| < 2.5$, with a nominal pixel size of $50 \times 400\,\mu\text{m}^2$ (in $\phi \times z$ in the barrel and in $\phi \times R$ in the endcaps). Together with the IBL, the full Pixel system geometry ensures that a charged-particle track passes through at least four layers of the high-granularity tracking elements within its coverage. In total, the IBL and Pixel systems contain approximately 12 and 80.4 million individual pixels, respectively, accounting for roughly 90 % of all ATLAS readout channels.

Semiconductor Tracker

The SCT detector [127–129] sits just beyond the Pixel layers and employs daisy-chained pairs of silicon-microstrips 80 µm wide and 6 cm long. The SCT consists of four concentric cylindrical barrel layers located at radial distances $R = 299$, 371, 443 and 514 mm and extending to $|z| = 749$ mm, and two endcap sections with nine disk layers each that cover the range $839 < |z| < 2735$ mm. As before, this geometry ensures that a charged-particle track passes through at least four layers within the SCT's coverage. Each barrel SCT module consists of four strip sensors glued back-to-back (two on the top and two on the bottom of the layer), with one set of sensors parallel to the beam axis and the other rotated by 40 mrad to provide measurements of the charged-particle trajectory in both the z- and ϕ-directions. The endcap modules use a similar sensor layout but with the strips aligned radially. The SCT provides coverage over the range $|\eta| < 2.5$ and a spatial resolution of 17 µm in the R–ϕ plane and 580 µm in the z- (barrel) and R-directions (endcaps). The total number of readout channels in the SCT is approximately 6.3 million.

Transition Radiation Tracker

The TRT detector [130–132] consists of carbon-fiber-reinforced Kapton/aluminum drift tubes, also called *straw tubes*, with a diameter of 4 mm. The inner wall of the tube serves as the cathode held at −1530 V with respect to a 31 µm diameter gold-plated tungsten wire running down the centre of each tube, which acts as the anode. The tubes are nominally filled with a 70 % Xe, 27 % CO_2, 3 % O_2 gas mixture, which is ionized by a traversing charged particle and the signal is read out via the anode. After Run 1, several gas leaks were found and some straws were filled with a less expensive mixture replacing the xenon gas with argon [133]. The barrel module consists of 73 layers of 144 cm long straws oriented parallel to the beam axis, organized in three concentric sections extending from $R = 563$ mm to $R = 1066$ mm. The endcap modules each contain 160 layers of 37 cm long straws, aligned radial to the beam axis as wheels, organized in twenty sections along the beam axis covering

the range $848 < |z| < 2710\,\text{mm}$. The TRT provides tracking coverage up to $|\eta| = 2.0$ with a spatial resolution of 130 µm per straw, with a traversing charged particle registering approximately 36 hits along its trajectory.

The TRT tubes are also interleaved with polypropylene or polyethylene fibers in the barrel and polypropylene foils in the endcap. This interleaving of materials gives rise to *transition radiation* (TR) emitted by high-energy charged particles traversing the interface between materials with different dielectric properties. At high energy, transition radiation is primarily emitted as x rays, with the total energy emitted proportional to the relativistic Lorentz factor $\gamma = E/m$ of the traversing charged particle [134]. Low-mass particles such as electrons passing through the TRT will therefore emit more energy as TR photons than other minimum-ionizing charged particles, which are then absorbed in the Xe-based gas mixture and yield much larger signal amplitudes. The distinction between transition-radiation and tracking signals is obtained on a straw-by-straw basis using separate low and high thresholds in the readout electronics. The TRT thus plays an important role in electron identification (Section 5.3.4).

Inner detector performance

The Inner Detector provides full tracking coverage in ϕ, with the IBL+Pixel and SCT detectors providing coverage up to $|\eta| = 2.5$ and the TRT up to $|\eta| = 2.0$. The performance of the Inner Detector is largely governed by the spatial resolution of the sensitive detector elements and the precision of their alignment with respect to the nominal interaction point, the strength of the magnet field and the mean number of interactions per bunch crossing (or pileup). Each of these factors contribute in different ways depending on the specific function of the Inner Detector at hand, such as track reconstruction, vertexing, p_T measurements and particle identification. For instance, collision events contain many primary vertices in high-pileup environments, therefore the spatial resolution of the innermost silicon-based tracking layers is most important for associating particle tracks to a primary vertex and resolving individual vertices. The Inner Detector design results in a typical primary vertex resolution of $\sim 20\,\mu\text{m}$ in the transverse direction and $\sim 40\,\mu\text{m}$ in the longitudinal direction, with the resolution generally degrading for tracks at higher $|\eta|$, in higher-pileup environments and as the number of tracks associated with the primary vertex decreases [135].

Similarly, the track momentum resolution has two main components. The first is proportional to the momentum of the outgoing charged particle and is related to the spatial resolution of the detector and the precision with which the arc of the particle trajectory can be measured. A smaller arc radius and longer track length allow for a greater number of space-point measurements along the particle trajectory, as well as a reduction in fake tracks and tracking inefficiencies due to missing hits and noise. In this case, the strength of the magnetic field provided by the central solenoid and the outermost layers of the Inner Detector are crucial for the momentum resolution, as they determine the ability to reduce and resolve the arc radius of the track, respectively. The second is a constant term that arises from ionization energy loss and multiple scattering within the detector material, which distorts the particle trajectory. Taken together, the design transverse-momentum resolution of the Inner Detector is [99]

$$\frac{\sigma_{p_\text{T}}}{p_\text{T}} = 0.05\,\%\, p_\text{T} \oplus 1\,\%\,, \tag{3.9}$$

where the symbol $\oplus$ denotes that the terms are added in quadrature. The precision of the tracking

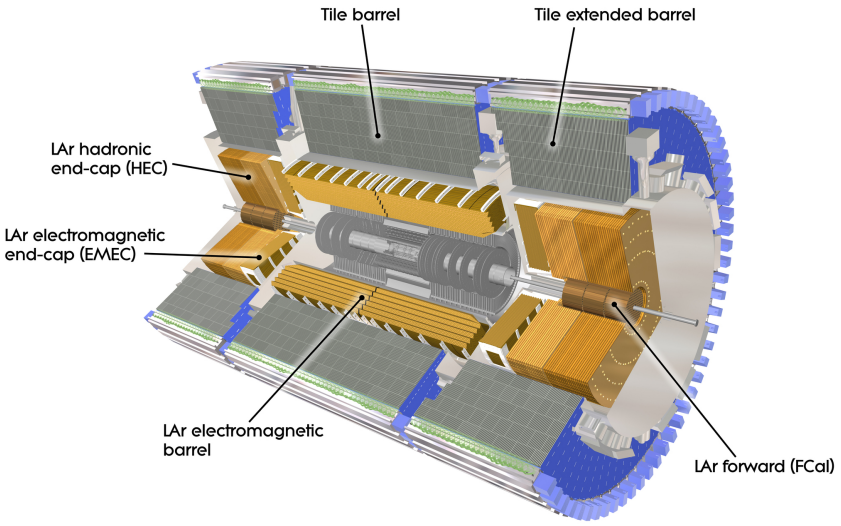

Figure 3.9: Cut-away view of the ATLAS calorimeter systems [99].

system's momentum resolution is therefore greatest for low-p_{T} charged-particle tracks, where the multiple-scattering effects dominate, with the resolution degrading as a function of p_{T} as the tracks straighten and their curvature becomes more difficult to determine.

3.3.3 Calorimeters

The ATLAS calorimeter system covers the pseudorapidity range $|\eta| < 4.9$ and provides particle energy measurements using different techniques suited to the broad requirements of the physics processes of interest and of the radiation environment over this large η range. Sampling calorimeters based on LAr technology are used for the detection of electromagnetic (EM) objects such as electrons and photons over the range $|\eta| < 4.9$, with a finely segmented precision-measurement region implemented in the central region of the detector ($|\eta| < 2.5$), as well as hadronic objects in the range $1.5 < |\eta| < 4.9$. Hadronic calorimetry within $|\eta| < 1.7$ is provided by the Tile calorimeter. A cut-away view of the ATLAS calorimeter systems is shown in Figure 3.9.

Sampling calorimeters consist of alternating layers of an active medium, which generates a signal, and a passive medium, which functions as an absorber. The interaction of an incident particle with the detector (through EM or strong processes) produces a shower of cascading secondary particles with progressively reduced energy as the shower develops and traverses the layers of the calorimeter. Most of the particles' energy is deposited in the absorber material, with only a small fraction of the total energy deposited and detected in the active material, called the *sampling fraction*. The alternating layers of the active medium in the calorimeter sample the shower profile at regular intervals and measure its energy with high granularity. In the case of LAr calorimeters, the layers of absorber material are interleaved with readout electrodes spaced apart by wire-frame structures made from low-density materials. The entire calorimeter is immersed in liquid argon, which fills the gaps between the electrodes and absorbers, and serves as the calorimeter's active medium. Charged particles produced in the shower pass through these LAr gaps and ionize the argon, creating electron–

ion pairs in the LAr medium. The electrons and ions then drift under the influence of an electric field applied across the gaps via a high-voltage system, producing a signal that is collected by the electrodes. The Tile calorimeter operates on a similar principle of sampling the shower profile at regular intervals, except instead it uses plastic scintillators as the active medium, which emit light as the shower traverses them that is collected and read out by a number of photomultiplier tubes.

At high energy, the production of electromagnetic showers is predominantly through *bremsstrahlung*, where an electron radiates a photon, and *pair production*, in which a photon converts to a e^+e^- pair. The main features of an electromagnetic shower are determined by the *radiation length* X_0 and the *Molière radius* R_M of the absorber material. The radiation length represents the average distance that an electron needs to travel in the material to reduce its energy to $1/e \approx 0.368$ of its original energy via bremsstrahlung and is given by [136]

$$X_0\,[\mathrm{g\,cm^{-2}}] \simeq \frac{(716\,\mathrm{g\,cm^{-2}})A}{Z(Z+1)\ln(287/\sqrt{Z})}\,, \tag{3.10}$$

where Z and A are the atomic number and mass of the material, respectively. Similarly, the *Molière radius* R_M is a characteristic constant of a material describing the transverse spread of an EM shower within it, with approximately 90 % of a shower's energy contained laterally in a cylinder of radius R_M, and is approximately

$$R_M \simeq \sqrt{\frac{4\pi}{\alpha}} m_e c^2 \frac{X_0}{E_c} \simeq (21.2\,\mathrm{MeV})\frac{X_0}{E_c\,[\mathrm{MeV}]}\,, \tag{3.11}$$

where $E_c \approx 610\,\mathrm{MeV}/(Z+1.24)$ is the *critical energy* at which the electron ionization losses and bremsstrahlung losses become equal, below which the shower energy is dissipated mainly by ionization and excitation and not in the generation of other particles.

Hadronic showers consisting of strongly interacting particles, on the other hand, deposit some fraction of their energy into the calorimeter material through nuclear processes, such as nuclear excitation, breakup and the production of neutrinos, which are "invisible" to detectors whose active material is sensitive only to ionization effects. Much of the energy deposited through other processes, such as secondary proton and $\pi^{\pm}$ ionization, can however be measured using traditional sampling-calorimeter technology. Moreover, at the centre of a hadronic shower is an electromagnetic component that arises predominantly from the production of neutral pions that decay to photon pairs. Hadronic showers also tend to be more spread out laterally than electromagnetic showers due to the production of fast, secondary hadrons in the direction transverse to the incident particle. Analogous to the radiation length of an EM shower, the longitudinal development of a hadronic shower in a material is characterized by the material's *nuclear interaction length* $\lambda \simeq (35\,\mathrm{g\,cm^{-2}})A^{1/3}$ [136]. The ATLAS hadronic calorimeters are said to be *non-compensating*, meaning that the nature of the absorber material is such that a hadron will on average deposit a smaller fraction of its energy in the active material of the calorimeter than would an electron or photon of equal incident energy. This invisible energy loss is accounted for later in the calibration of hadronic objects (Section 5.3.6).

The ATLAS calorimeters were designed to provide excellent energy resolution and containment of EM and hadronic showers, and to limit punch-through into the muon system. Roughly 20–40 radiation lengths of material in the EM calorimeter is sufficient to contain approximately 95 % of the shower energy deposited by electrons and photons up to $\sim 1\,\mathrm{TeV}$. Similarly, the 8–12 nuclear

Table 3.2: Main parameters of the LAr and Tile calorimeter systems. The calorimeter granularity is given in $\Delta\eta \times \Delta\phi$ space, except for the FCal where it is given in $\Delta x \times \Delta y$. The granularity in the inner and outer edges of each FCal module are approximately four times finer. The nominal high voltage is the potential applied across the LAr gaps and therefore pertains only to the LAr calorimeter subsystems.

Subsystem	Layer	η range	Granularity ($\Delta\eta \times \Delta\phi$)	Nominal high voltage [V]
EMB	Presampler	$\lvert\eta\rvert < 1.52$	0.025×0.1	+2000
	1	$\lvert\eta\rvert < 1.475$	$0.025/8 \times 0.1$, 0.025×0.1	
	2	$\lvert\eta\rvert < 1.475$	0.025×0.025, 0.075×0.025	+2000
	3	$\lvert\eta\rvert < 1.35$	0.050×0.025	
EMEC	Presampler	$1.5 < \lvert\eta\rvert < 1.8$	0.025×0.1	−2000
	1	$1.375 < \lvert\eta\rvert < 3.2$	$0.025/8 \times 0.1$ to 0.1×0.1	
	2	$1.375 < \lvert\eta\rvert < 1.425$	0.025×0.025 to 0.1×0.1	+1000 to +2500
	3	$1.5 < \lvert\eta\rvert < 2.5$	0.050×0.025	
HEC	1–4	$1.5 < \lvert\eta\rvert < 2.5$	0.1×0.1	+1800
		$2.5 < \lvert\eta\rvert < 3.2$	0.2×0.2	
FCal	1	$3.10 < \lvert\eta\rvert < 4.83$	$\Delta x \times \Delta y = 3.0 \times 2.6\,\mathrm{cm}^2$	+250
	2	$3.20 < \lvert\eta\rvert < 4.81$	$\Delta x \times \Delta y = 3.3 \times 4.2\,\mathrm{cm}^2$	+375
	3	$3.29 < \lvert\eta\rvert < 4.75$	$\Delta x \times \Delta y = 5.4 \times 4.7\,\mathrm{cm}^2$	+500
Tile barrel	1–2	$\lvert\eta\rvert < 1.0$	0.1×0.1	N/A
	3		0.2×0.1	
Tile extended barrel	1–2	$0.8 < \lvert\eta\rvert < 1.7$	0.1×0.1	N/A
	3		0.2×0.1	

interaction lengths in the hadronic calorimeters contain the shower energy deposited by strongly interacting particles and jets. The cumulative amount of material in interaction lengths as a function of $|\eta|$ in front of and within the calorimeters is shown in Figure 3.10.

Each of the calorimeter subsystems is described below, with a summary of each component's main parameters given in Table 3.2.

Electromagnetic calorimeters

The electromagnetic calorimeter consists of accordion-shaped kapton-insulated copper electrodes and lead absorber plates clad between sheets of stainless steel for structural support. The accordion geometry provides complete ϕ symmetry without azimuthal cracks. It is divided into the electromagnetic barrel (EMB) calorimeter, covering $|\eta| < 1.475$, and two electromagnetic endcap (EMEC) calorimeters, covering $1.375 < |\eta| < 3.2$, each housed in their own cryostat. As illustrated in Figures 3.11 and 3.12, both the EMB and the EMEC are segmented into two or three layers in depth (front, middle and back, as viewed from the interaction point), with each segment's electrode granularity chosen to optimize the resolution of electromagnetic showers and rejection of neutral hadron decays, while keeping the total number of readout channels to a minimum given the spatial constraints of the cryogenic feedthroughs [137] and to reduce the overall complexity of the data-acquisition system.

Since a particle originating from the interaction point will have already passed through 1–4 X_0 of upstream material in the Inner Detector and central solenoid, EM showers may begin to develop before reaching the calorimeter. The electromagnetic calorimeters are therefore also complemented with LAr presampler detectors placed in front of their inner surface in order to correct for these energy losses.

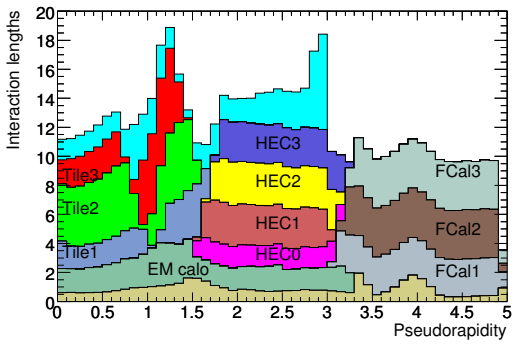

Figure 3.10: Cumulative amount of material, in units of interaction length, as a function of $|\eta|$, in front of and within each of the calorimeter subsystems [99]. Also shown (in cyan) is the total amount of material in front of the first active layer of the muon spectrometer (up to $|\eta| < 3.0$).

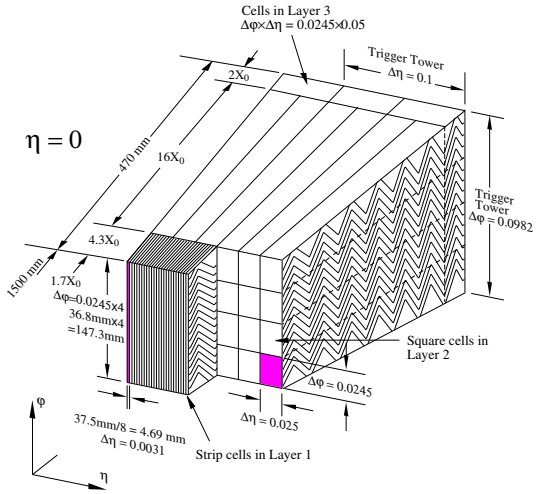

Figure 3.11: Diagram of a LAr EMB module showing the different layers through the ganging of electrodes in ϕ [99]. The granularity in η and ϕ of the cells of each of the three layers and of the trigger towers is also shown.

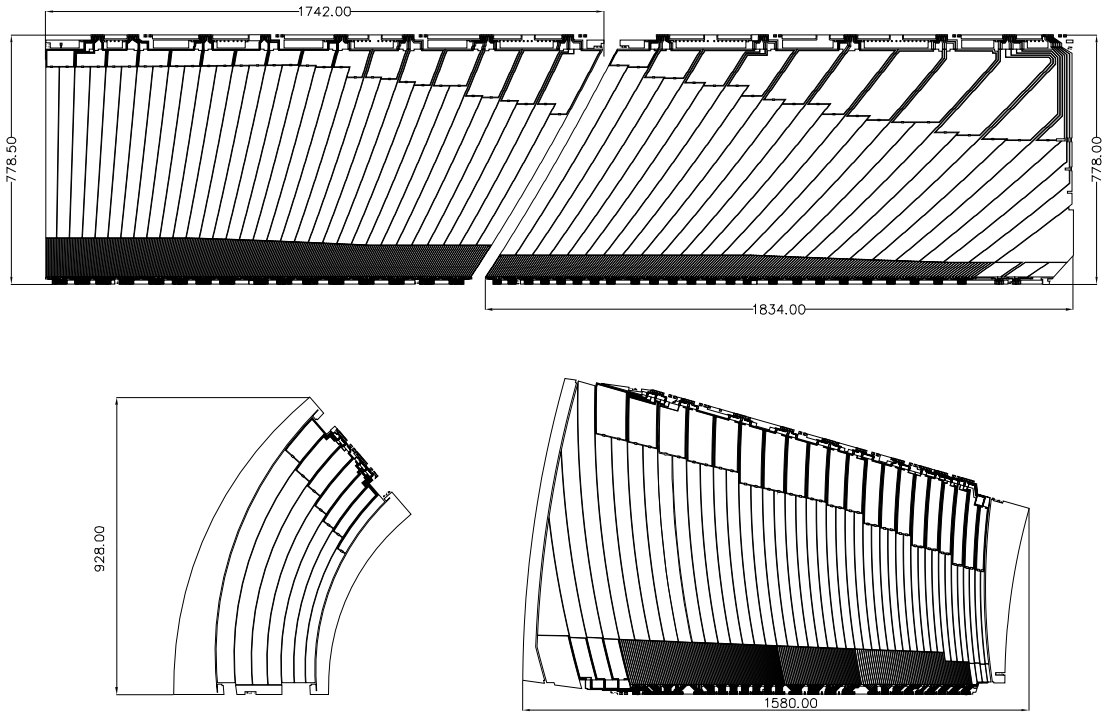

Figure 3.12: Layout of the LAr electromagnetic calorimeter layers for the four different types of electrodes before folding into the accordion shape [99]. The two top electrodes are for the barrel and the two bottom electrodes are for the endcap inner (left) and outer (right) wheels. Dimensions are in mm and the drawings are all at the same scale. The two or three different layers in depth are clearly visible.

Electromagnetic barrel calorimeter

The electromagnetic barrel calorimeter [138] is made of two half-barrels, centred around the z-axis. One half-barrel covers the region with $z > 0$ ($0 < \eta < 1.475$) and the other one the region with $z < 0$ ($-1.475 < \eta < 0$). The length of each half-barrel is 3.2 m, and their inner and outer diameters are 2.8 m and 4 m respectively.

A half-barrel is made of 1024 accordion-shaped absorbers, interleaved with readout electrodes. The electrodes are positioned in the middle of the gap by honeycomb spacers. The size of the drift gap on each side of the electrode is 2.1 mm, which corresponds to a total drift time of about 450 ns for an operating voltage of 2000 V.

Electromagnetic endcap calorimeter

The electromagnetic endcap calorimeters [139] consist of two wheels, one on each side of the electromagnetic barrel. Each wheel is 63 cm thick, with internal and external radii of 330 mm and 2098 mm, respectively. In order to improve the energy measurement in the transition region between the barrel and the endcap calorimeters, a LAr presampler is implemented in front of the endcap calorimeter, covering the range $1.5 < |\eta| < 1.8$.

Each endcap calorimeter consists of two coaxial wheels. The boundary between the inner and the outer wheel, which is 3 mm wide and located at $|\eta| = 2.5$, is mostly filled with low-density material. Each endcap wheel is further divided into eight wedge-shaped modules without introducing any discontinuity along the azimuthal angle owing to the accordion geometry. Each endcap contains 768 absorbers interleaved with readout electrodes in the outer wheel and 256 absorbers in the inner wheel. The electrodes are positioned in the middle of the gaps by honeycomb spacers.

In contrast to the EMB, the drift gap in the EMEC is not constant, but is a function of the radial distance from the beam axis R. In the outer wheel it varies from 2.8 mm at $R = 200$ cm to 0.9 mm at $R = 60$ cm. In the inner wheel it varies from 3.1 mm at $R = 70$ cm to 1.8 mm at $R = 30$ cm. Since the detector signal depends on the drift velocity and liquid-argon gap thickness, a high voltage continuously varying with η would be needed to achieve an η-independent detector response (shown as the open circles in Figure 3.13). In practice, such a variable high voltage could not be implemented and instead is approximated with a variation by steps (shown as the closed triangles in Figure 3.13). This design results in a drift time decreasing with $|\eta|$ in each of the inner and outer wheels, ranging from about 600 ns to 200 ns [140]. Figure 3.14 shows the flat electrode structure of the EMEC and explicitly gives the nominal high voltage values as a function of the η range, defining the seven HV sectors of the outer wheel and the two sectors of the inner wheel.

Hadronic endcap calorimeter

The hadronic endcap (HEC) calorimeter [141] consists of two independent wheels per endcap, located directly behind the endcap electromagnetic calorimeter and sharing the same LAr cryostats. Each wheel is built from 32 identical wedge-shaped modules, and is divided into two segments in depth, for a total of four layers per endcap. The front two wheels closest to the interaction point (HEC1) are built from 25 mm parallel copper plates, while the two rear wheels (HEC2) use 50 mm copper plates (the first plate of each wheel is half-thickness). This structure, along with the layout of the readout electrodes, is shown schematically in Figure 3.15. The copper plates are interleaved with

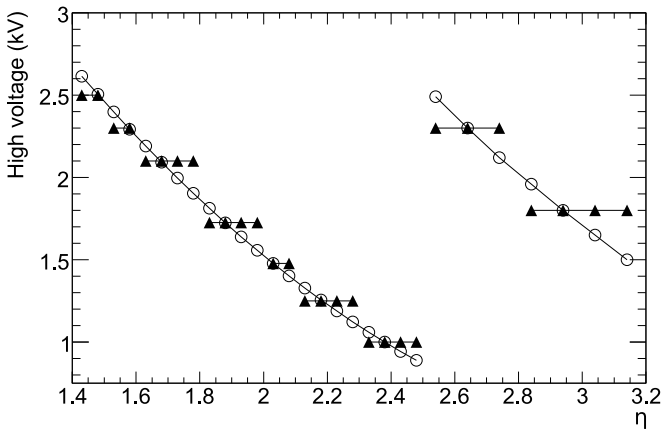

Figure 3.13: Distribution of the EMEC high voltage as a function of $|\eta|$ [99]. A uniform calorimeter response requires a high voltage that varies continuously as a function of $|\eta|$, as shown by the open circles. This has been approximated by a set of discrete values shown as full triangles.

8.5 mm liquid-argon gaps, which are divided by three electrodes into four separate LAr drift zones of 1.8 mm width each, as shown in Figure 3.16. The space between the electrodes is maintained using a honeycomb sheet, which is held in place by seven tie-rods. Each drift zone is individually supplied with a nominal high voltage of 1800 V, resulting in typical drift times of 430 ns.

Forward calorimeter

The forward calorimeters (FCal) [142] are located in the same cryostats as the endcap calorimeters and provide coverage over $3.1 < |\eta| < 4.9$. As the FCal modules are located at high η, at a distance of approximately 4.7 m from the interaction point, they are exposed to high particle fluxes, and required special design considerations to achieve the desired performance under these harsh conditions. The FCal consists of three modules in each endcap, as illustrated in Figure 3.17: the first (FCal1), made of copper, is optimized for electromagnetic measurements, while the other two (FCal2 and FCal3), made of tungsten, measure predominantly the energy of hadronic interactions. Each module consists of a metal matrix with regularly spaced longitudinal channels filled with the electrode structure consisting of concentric rods and tubes parallel to the beam axis. The gap between the inner solid rod and the outer tube is maintained by a helically wound PEEK fiber. The FCal1 electrode arrangement, and the effective Molière radius for this module, is shown in Figure 3.18.

The liquid-argon gaps in the FCal are smaller than the usual 2 mm gap of the other LAr calorimeters to avoid ion-buildup problems and to provide at the same time a faster signal and greater density. In the electromagnetic layer (FCal1), with a LAr gap width of 0.269 mm, the triangular current pulse at the electrode has a full drift time of 60 ns. For FCal2 and FCal3, with LAr gap widths of 0.376 mm and 0.508 mm, respectively, the full drift time scales with the gap width since the field in the gaps is similar for all three modules.

In each FCal electrode, the tube is held at ground and the rod at a positive high voltage such that the electric field in the gap is approximately 1.0 kV/mm. The electrodes act like short transmission lines, open at one end and connected to a readout coaxial cable at the other. The majority of

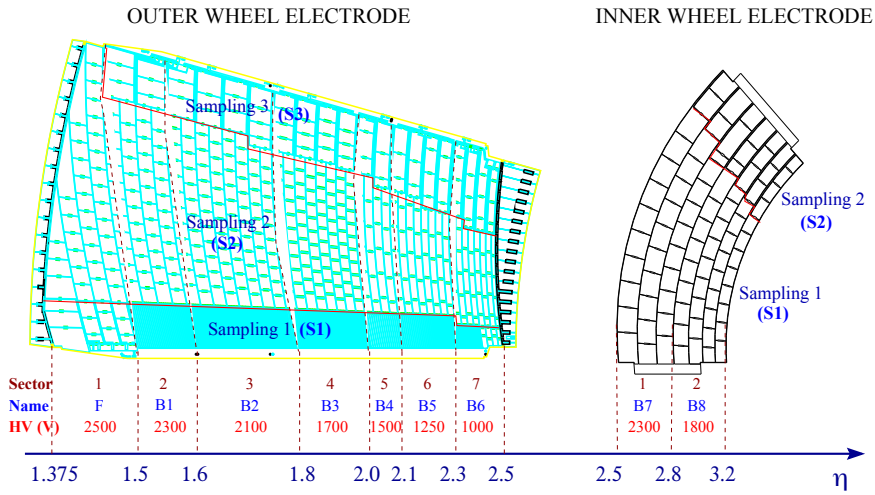

Figure 3.14: Schematic view of the EMEC outer- and inner-wheel flat electrodes [139]. Sampling regions are represented as well as high-voltage sectors. High voltage values are indicated for each sector at liquid argon temperature.

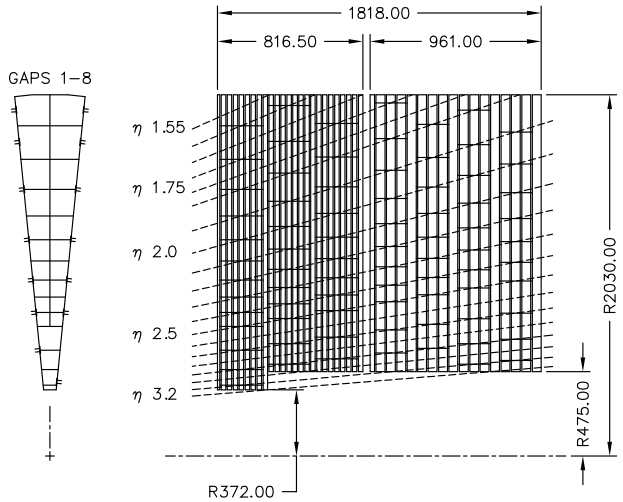

Figure 3.15: Schematic view of the hadronic endcap calorimeter in the R–ϕ (left) and R–z (right) planes [99]. The layout of the readout electrodes is indicated by the dashed lines. All dimensions are in mm.

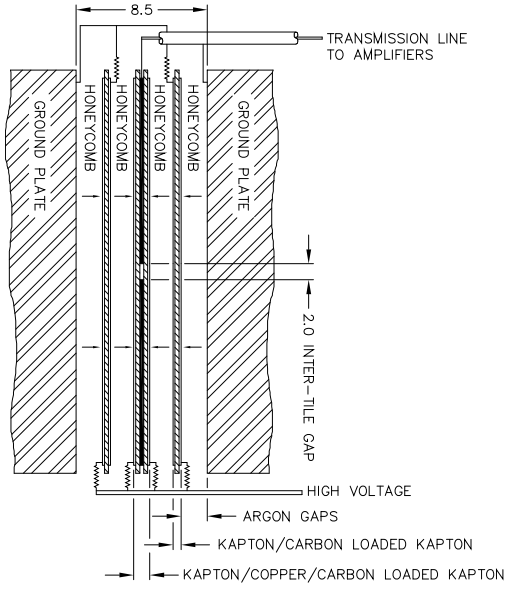

Figure 3.16: Schematic view of the arrangement of the HEC readout structure in the 8.5 mm interplate gap [99]. All dimensions are in mm.

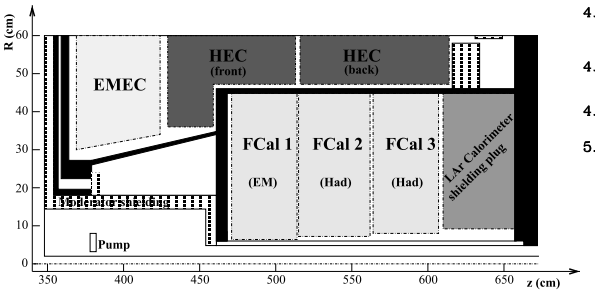

Figure 3.17: Schematic diagram showing the three FCal modules located in the endcap cryostat [99]. The material in front of the FCal and the shielding plug behind it are also shown. The black regions are structural parts of the cryostat. The diagram has a larger vertical scale for clarity.

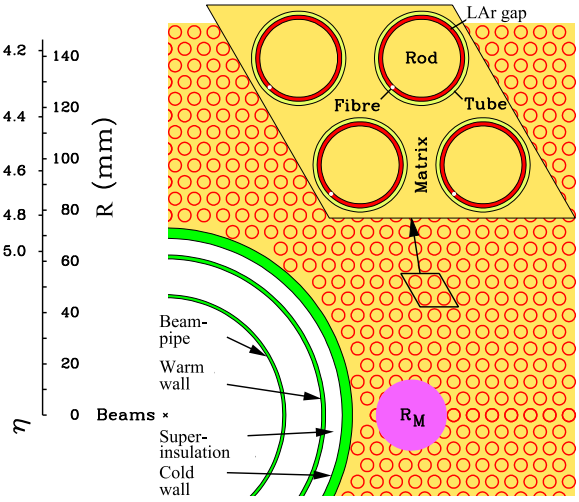

Figure 3.18: Electrode structure of FCal1 with the matrix of copper plates and the copper tubes and rods with the LAr gap for the electrodes [99]. The Molière radius, R_M, is represented by the solid magenta disk.

the electrodes are ganged together via interconnect boards into *tube groups* of four, six and nine electrodes for FCal1, FCal2 and FCal3, respectively. The readout cables carry the current pulse to summing boards mounted behind the HEC calorimeter where radiation levels are much lower than in the region of the FCal. The tube groups at the inner and outer radii of the modules are in general not summed due to geometric constraints and higher counting rates at the inner radius. High voltage is also distributed on the summing boards via a set of protection resistors, which limit the current in case of a spark or short in an electrode, as shown in Figure 3.19 for the specific case of FCal1. Each of the four tube groups shown in Figure 3.19 is fed from a different HV supply. Thus, in case of a high-voltage failure, only one tube group in a readout ensemble is affected.

LAr electronics and signal processing

The LAr readout electronics consist of two main parts: a front-end, mounted on the detector itself, and a back-end, located off-detector in an underground service cavern. The LAr calorimeter signal is a triangular pulse with a fast rise and long tail during the time in which the ionization electrons are drifting towards the calorimeter electrodes. The front-end boards (FEBs), mounted on the outer walls of the detector cryostats in front-end crates [143], receive these raw calorimeter signals and amplify, shape, sample and digitize them, and then finally transmit the processed signals to the back-end electronics. First, a pulse is amplified using three gains—low, medium, and high—and then undergoes bipolar shaping. The resulting bipolar pulse maintains the amplitude of the raw signal but with zero net area such that overlapping pulses from background processes, such as thermal noise and pileup interactions, on average sum to zero. The shaped pulse is sampled every 25 ns, and the measurements—still analog at this point—are stored until the Level-1 trigger decision is made on the event (Section 3.3.5). The triangular "detector pulse" from an EMB cell and the shaped "physics pulse" with 25 ns samplings are illustrated in Figure 3.20. Actual pulse shapes recorded

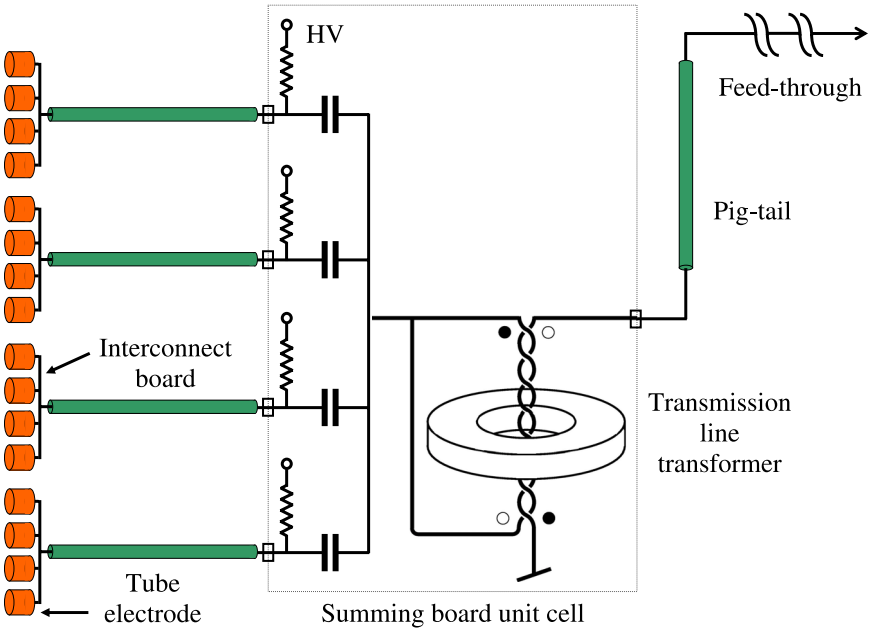

Figure 3.19: Schematic of the FCal1 module cabling from the electrodes to the cryogenic feedthrough [99]. Four groups of four electrodes are shown as short orange tubes on the left. Equal-length coaxial cables take the signals to a summing board where protection resistors introduce HV to the electrodes and blocking capacitors isolate the HV from the remainder of the circuit. The twisted pair transmission line in the transformer is wrapped around a ferrite core. The other modules differ only by the number of electrodes grouped together on the interconnect board (six for FCal2 and nine for FCal3).

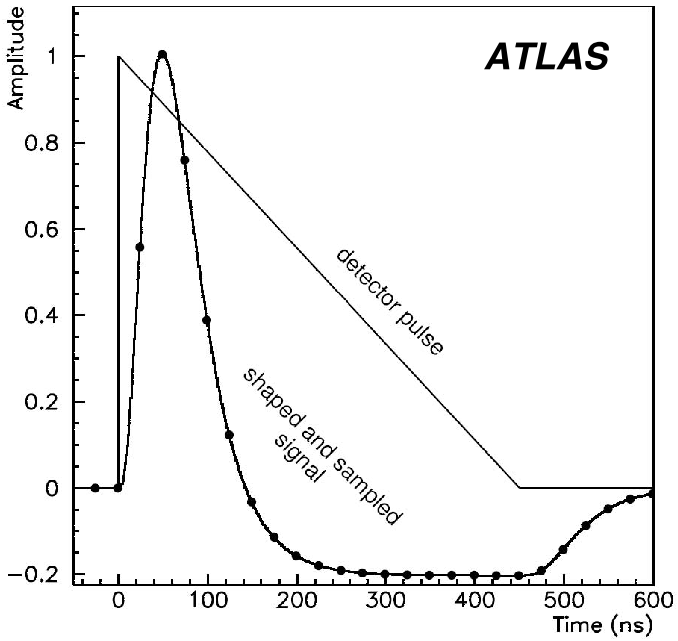

Figure 3.20: Illustration of a typical LAr calorimeter triangular current pulse in the EMB and the signal output after bipolar shaping [99]. The dots indicate ideal sampling positions separated by 25 ns.

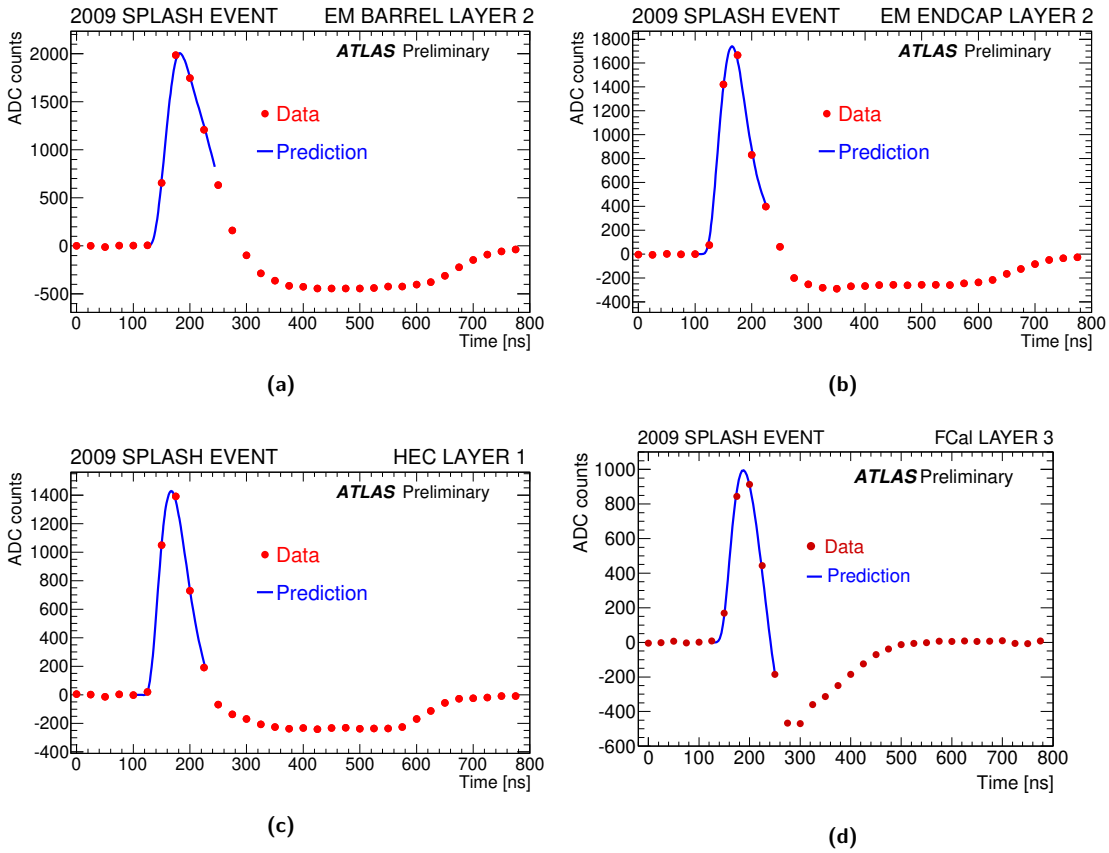

Figure 3.21: Examples of LAr calorimeter ionization pulse shapes for the (a) EMB, (b) EMEC, (c) HEC and (d) FCal. Data were collected during beam-splash events in 2009 using the special 32-sample readout mode. The prediction shown is the one used for energy reconstruction around the peak sample.

during beam-splash events in 2009 are shown in Figure 3.21 for each of the four detector subsystems.

Upon being accepted by the trigger, the optimal gain for the pulse is chosen and four samples—with at least one on the rising edge of the pulse, one near the maximum amplitude and the remaining two mapping out the pulse tail—are passed to an analog-to-digital converter. The FEB transfers the digitized samples to its corresponding back-end readout driver (ROD), which performs the digital signal processing and returns measurements of the signal energy and timing. Optimal filtering coefficients [144], calculated during calibration runs for each channel, are applied to the digitized signal to optimize the energy and timing resolution. The measured energy is finally recorded, to be used later for object reconstruction in normal physics data-taking. Under the right circumstances, the energy measurements may also be used for luminosity monitoring, as described in Chapter 6.

Tile calorimeter

Behind the EMB in the region $|\eta| < 1.7$ is the hadronic Tile calorimeter [145]. It is divided along the beam axis into four partitions: two central barrels covering $|\eta| < 1.0$, approximately 2.6 m in length, and two extended barrels covering $1.0 < |\eta| < 1.7$, each 2.6 m in length, with inner and outer radii of 2.28 m and 4.25 m, respectively. The layout of the Tile calorimeter cells is shown schematically in Figure 3.22. Each barrel part consists of 64 modules, or "wedges", of size $\Delta\phi \approx 0.1$ that are

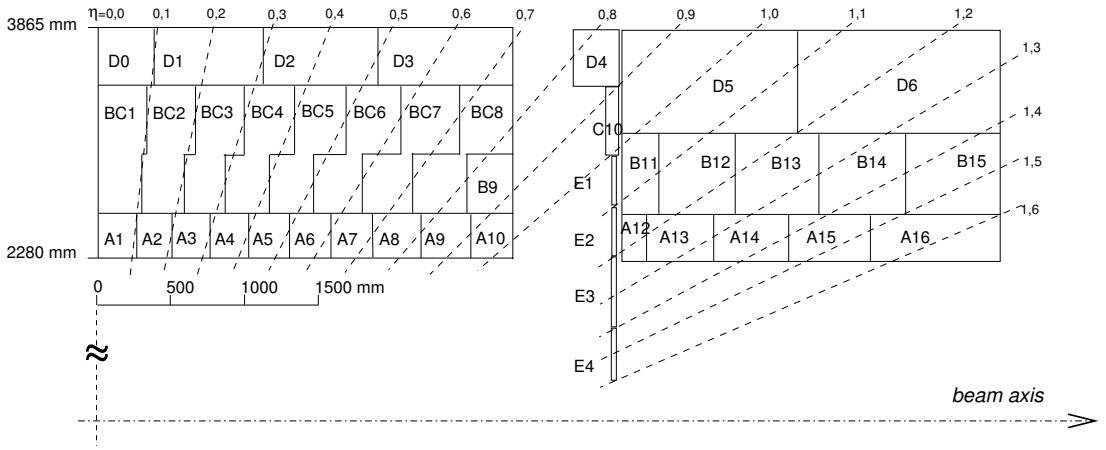

Figure 3.22: Layout of the Tile calorimeter cells, showing their segmentation in depth and in η in the central (left) and extended (right) barrels [99].

made of steel absorber plates and scintillating plastic tiles, as shown in Figure 3.23. In addition, the outsides of individual cryostats are instrumented with scintillators (the Tile E cells) to allow for corrections to the measured particle energy due to upstream energy losses, covering $1.0 < |\eta| < 1.2$ and $1.2 < |\eta| < 1.6$. Steel girders surrounding the Tile calorimeter at radial distance $R \approx 4.25\,\text{m}$ provide mechanical support as well as a return yoke for the central solenoid magnetic field.

When a charged particle passes through the scintillating tiles, ultraviolet light is emitted. This light is collected at two edges of each plastic scintillator, and is then transported via wavelength-shifting fibres to photomultiplier tubes located at the back of each barrel module. Tiles are finally read out in groups to form three longitudinal sampling layers for further processing.

Calorimeter performance

The primary figure of merit for any calorimeter is its energy resolution, conventionally parameterized as:

$$\frac{\sigma_E}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c\,. \tag{3.12}$$

The first term a, called the *sampling* or *stochastic term* is a function of the number of sampling layers in the calorimeter and accounts for random particle fluctuations in the shower development. The second term b, called the *noise term*, arises from electronics and pileup noise in the detector and is generally proportional to the number of channels summed over. Finally, the constant term c accounts for all other energy-independent characteristics of the calorimeter, including spatial inhomogeneities of the detector energy response, dead space, radiation damage and temperature variations of the active material.

The energy resolution of each of the ATLAS calorimeters was measured using test-beam data under a variety of configurations each suited to the requirements of the individual detector subsystems. Electrons beams with energies 1–250 GeV were used for standalone tests of a prototype EMB detector to determine its response and energy resolution [146]. Electron, muon and pion beams with energies 6–200 GeV were used for both standalone and combined tests of EMEC, HEC and FCal prototype detectors to determine the combined endcap calorimeter performance, with dedicated studies in the region $1.6 < |\eta| < 1.8$ and in the transition region around $|\eta| = 3.2$ [147–150]. Finally,

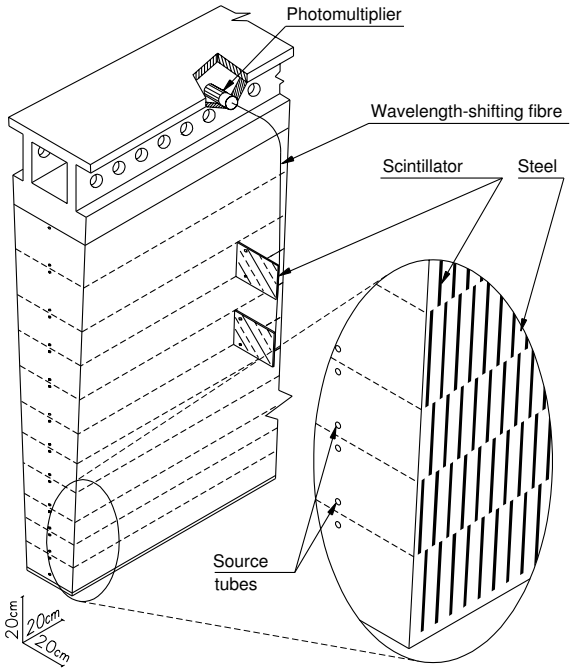

Figure 3.23: Schematic of a Tile calorimeter module showing the integration of the mechanical assembly and the optical readout [99].

electron, muon and pion beams with energies 3–350 GeV were used to determine the response of the hadronic Tile calorimeter [151]. The sampling and constant terms a and c, respectively, were extracted from fits to the noise-subtracted energy resolution as a function of beam energy and are given for each subsystem in Table 3.3 along with their corresponding design value.

Finally, while the accordion geometry of the electromagnetic calorimeters offers near full hermeticity in ϕ, the crack region between the EMB and EMEC, $1.37 < |\eta| < 1.52$, suffers from somewhat degraded performance due to the diminished shower sampling, and therefore this region is generally vetoed in the reconstruction of EM objects such as photons and electrons (Section 5.3.4).

Table 3.3: Design energy resolution of the ATLAS calorimeters and their measured values determined in test-beam studies. The constant term of the HEC (low-η) energy resolution was found to be statistically consistent with zero and is therefore reported as 0 %.

Subsystem	Design resolution	Measured resolution
EMB		$(10.5 \pm 0.4)\,\%/\sqrt{E} \oplus (0.2 \pm 0.1)\,\%$
EMEC (low-η)	$10\,\%/\sqrt{E} \oplus 0.7\,\%$	$(12.1 \pm 0.2)\,\%/\sqrt{E} \oplus (0.4 \pm 0.1)\,\%$
EMEC (high-η)		$(13.5 \pm 0.5)\,\%/\sqrt{E} \oplus (0.7 \pm 0.1)\,\%$
Tile		$(52.9 \pm 0.9)\,\%/\sqrt{E} \oplus (5.7 \pm 0.2)\,\%$
HEC (low-η)	$59\,\%/\sqrt{E} \oplus 3\,\%$	$(84.1 \pm 0.3)\,\%/\sqrt{E} \oplus 0\,\%$
HEC (high-η)		$(88 \pm 5)\,\%/\sqrt{E} \oplus (6.8 \pm 0.4)\,\%$
FCal (EM layer)	—	$(28.5 \pm 1.0)\,\%/\sqrt{E} \oplus (3.5 \pm 0.1)\,\%$
FCal (hadronic layers)	$100\,\%/\sqrt{E} \oplus 10\,\%$	$(94.2 \pm 1.6)\,\%/\sqrt{E} \oplus (7.5 \pm 0.4)\,\%$

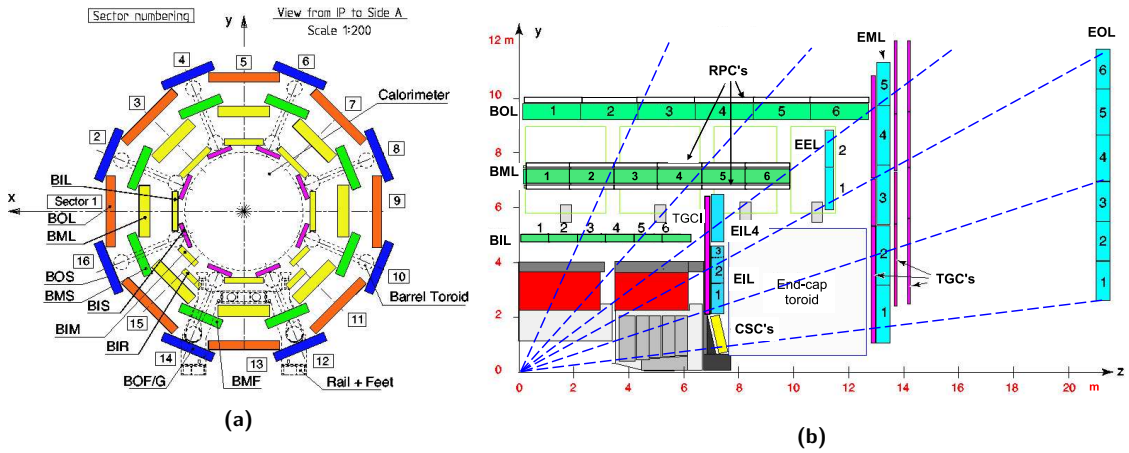

Figure 3.24: Cross-sectional view of the Muon Spectrometer in the (a) transverse (x–y) and (b) longitudinal (z–y) planes [99]. The toroidal magnetic field causes the muons to bend in the longitudinal plane. The naming scheme of the MDTs is as follows: B/E refers to the barrel or endcap positioning, I/M/O/E refers to the inner, middle, outer or "extra" layers positions, and L/S refers to the large or small chambers.

3.3.4 Muon Spectrometer

The Muon Spectrometer (MS) [152] is located just beyond the calorimeters and is responsible for the identification and reconstruction of muons. It comprises several different detector technologies specially designed for two primary uses: chambers dedicated to the precision tracking of each muon candidate and chambers dedicated to fast readout for the ATLAS trigger system (Section 3.3.5). The precision tracking chambers cover the pseudorapidity range $|\eta| < 2.7$, while the trigger chambers are only instrumented in the region $|\eta| < 2.4$. High-energy muons are highly penetrating due to their large mass (relative to electrons) and because they do not participate in the strong interaction, therefore they pass through the calorimeters without depositing much of their energy through ionization. Moreover, their relatively long lifetime ($\sim 2.2\,\mu s$) ensures they traverse the MS chambers before decaying.

The Muon Spectrometer is split into three common regions: a barrel section covering $|\eta| < 1.05$ and two endcaps covering $1.05 < |\eta| < 2.7$. The barrel MS is comprised of three cylindrical layers, concentric to the beam axis, with radii of approximately 5, 7.5 and 10 m from the beam pipe. These barrel layers are composed of eight large chambers and eight small chambers, differing only in lateral length, that follow the azimuthal symmetry of the ATLAS toroidal magnet. These large and small chambers overlap slightly in the ϕ plane to provide full azimuthal coverage. The endcap MS is arranged perpendicular to the beam axis and is arranged in wheel structures located at distances along the beam axis of approximately $|z| = 7.4\,m$ (Small Wheels, the endcap inner layers), 10.8 m (the "extra" layers of the endcap situated around the toroidal magnets), 14 m (first Big Wheel, endcap middle layers), and 21.5 m (last Big Wheel, endcap outer layers). Details of the geometry and positioning of the Muon Spectrometer are shown in Figure 3.24.

The magnetic field for the barrel MS is provided by eight coils of superconducting magnets, referred to as the barrel toroid, which provide a magnetic field of approximately 0.5 T. The eight coils are oriented radially, covering the region $|\eta| < 1.4$, and are spaced approximately equally in the azimuthal direction. Each endcap MS section contains a smaller toroidal magnet consisting of

eight square coils and eight wedge coils that provide a magnetic field of 1.0 T. These endcap toroids provide the magnetic field for the region $1.6 < |\eta| < 2.7$. In the barrel–endcap transition region, $1.4 < |\eta| < 1.6$, the two magnetic fields from the barrel toroid and endcap magnets overlap and provide a combined magnetic field. The toroidal magnetic fields are in the azimuthal direction, causing the outgoing muons to bend in the R–z plane back towards or away from the direction of the beam axis, depending on the sign of the muon charge.

The four detector technologies in the Muon Spectrometer are the Monitored Drift Tubes (MDTs), Resistive Plate Chambers (RPCs), Cathode Strip Chambers (CSCs) and Thin Gap Chambers (TGCs). The detectors are arranged in a complementary fashion, with the MDTs and CSCs providing precision tracking measurements, and the RPCs and TGCs providing fast track readout for the trigger system. The MDTs and CSCs are aligned to precisely measure the primary coordinate of muon tracks in the bending (R–z) plane while the RPCs and TGCs measure the track coordinates in both the bending and non-bending (ϕ) planes. The precision and non-bending coordinate measurements are later matched in order to provide the muon momentum measurement (Section 5.3.5). Descriptions of each of the four detector technologies are given below.

Monitored drift tubes

The MDT chambers are used for precision tracking in the Muon Spectrometer, providing on average a single-tube spatial resolution of 80 µm for a track in the bending plane. The Muon Spectrometer contains three layers of MDT chambers, where each chamber is composed of three to eight layers of drift tubes. The outer two layers of MDTs cover the entire MS detection region $|\eta| < 2.7$, while the innermost layer only covers the region $|\eta| < 2.0$. Each chamber consists of drift tubes with a diameter of 30 mm filled with an Ar/CO_2 gas mixture held at a pressure of 3 atmospheres. After ionization of the gas by a muon, the electrons drift towards a 50 µm tungsten–rhenium anode wire in the centre of the tube held at a voltage of ~ 3 kV. Deformations in the MDT mounting structures occur over time, therefore an internal chamber-alignment system consisting of four optical alignment rays was implemented. These alignment systems continuously monitor potential deformations of the frame in order to allow for corrections to the muon track position.

Resistive plate chambers

The RPCs consist of parallel-plate detectors separated by 2 mm using insulating spacers, with this volume filled with a gas mixture of predominantly tetrafluoroethane ($C_2H_2F_4$). Traversing muons ionize the gas, causing an "avalanche" of charge to drift towards the anode under the ~ 4.9 kV/mm electric field applied across the plates. The RPCs consist of three concentric cylindrical layers of this parallel plate design, with each layer known as a "muon station". The large distance between the innermost RPC station and outermost RPC station allows for identification of muon candidates in the $p_T = 9$–35 GeV range, while low-p_T (4–9 GeV) identification is determined using coincidences between the two layers of the inner chambers.

Cathode strip chambers

The CSC detector replaces the MDTs in the first (innermost) layer in the region $2 < |\eta| < 2.7$, where a high interaction rate during proton collisions would degrade the MDT performance. The

CSC detector is organized into two disks of 16 chambers each, with small and large chambers overlapping in the ϕ coordinate. Each CSC chamber consists of a multi-wire proportional chamber and two segmented cathode strips. One cathode is oriented perpendicular to the wires in order to provide precision-coordinate tracking and the other is oriented parallel in order to capture the transverse component of the muon tracks.

Thin gap chambers

The TGCs are conceptually similar to the multi-wire proportional chambers of the CSCs and consist of 50 µm diameter gold-plated tungsten anode wires separated by a distance of 1.8 mm and graphite planes acting as cathodes, with a cathode–anode separation of 1.4 mm. The TGCs operate with a gas mixture of CO_2 and n-pentane with a 2.9 kV potential allowing for fast measurements of the muon track coordinates in the MS endcaps. In addition to their triggering purposes, the TGCs also provide the azimuthal coordinate of the muon tracks to complement the MDTs. The TGCs consist of seven layers of wire chambers behind the MDT middle layer and two layers behind the MDT inner layer. The TGCs cover the region $1.0 < |\eta| < 2.4$.

Muon Spectrometer performance

The reconstruction efficiency for central muons ($|\eta| < 2.5$) is $\gtrsim 98\,\%$ when the muon systems are used in combination with matched Inner Detector tracks. For these combined central muons, the p_{T} resolution is a few percent for $p_{\mathrm{T}} < 100\,\mathrm{GeV}$, rising to approximately 10 % for $p_{\mathrm{T}} \approx 1\,\mathrm{TeV}$ [99]. In the absence of Inner Detector tracking information, the muon momentum resolution is largely unaffected for muons with p_{T} above 100 GeV, but begins to degrade below this point by up to a factor of approximately two for muons with $p_{\mathrm{T}} \approx 30\,\mathrm{GeV}$. Moreover, in the crack region between the two barrel MS modules ($|\eta| < 0.1$), this performance is significantly degraded and alternative muon-reconstruction methods are required (Section 5.3.5).

3.3.5 Trigger and data-acquisition system

The ATLAS Trigger and Data Acquisition (TDAQ) system is responsible for identifying the collision events containing the signatures of physics processes of interest, reading out the detector and transmitting the event data for processing and further offline analysis. The 40 MHz LHC collision rate is far beyond the bandwidth and storage capacity available to the ATLAS computing infrastructure, therefore one of the main functions of the TDAQ system is to reject the large number of QCD background events and reduce the event-readout rate to 1–2 kHz for permanent storage and analysis.

Since the beginning of Run 2, ATLAS uses a two-level trigger system to select the events of interest. The first of these is the hardware-based Level-1 (L1) trigger, followed by the software-based high-level trigger (HLT). These two systems are described below. A schematic overview of the ATLAS TDAQ chain is shown in Figure 3.25.

Level-1 trigger

The first trigger level is the hardware-based Level-1 trigger, consisting of three main systems: the L1 Muon Trigger (L1Muon), the L1 Calorimeter Trigger (L1Calo), and the L1 Central Trigger

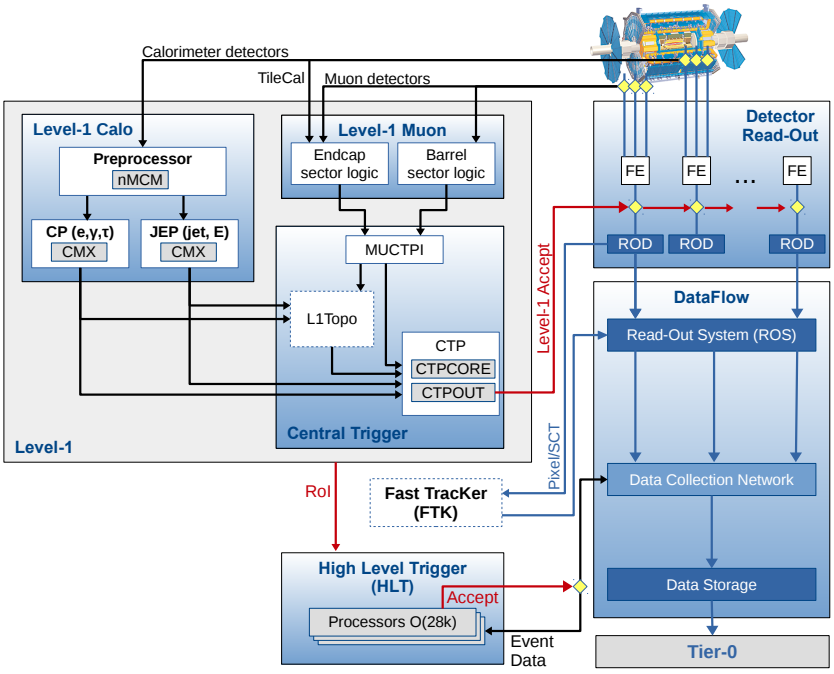

Figure 3.25: The ATLAS Trigger and Data Acquisition (TDAQ) system in Run 2, with emphasis on the components relevant for triggering [153]. L1Topo and FTK were being commissioned during 2015, although the latter was not used during Run 2 operations.

Processor (CTP). The L1Muon trigger processors receive input from the dedicated muon trigger chambers described above in Section 3.3.4, while the L1Calo trigger relies on a reduced-granularity calorimeter readout called calorimeter *trigger towers*. Trigger towers provide energy and timing sums for detector regions of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ for the barrel calorimeters and become less granular (up to $\Delta\eta \times \Delta\phi = 0.4 \times 0.4$) in the forward regions.

During Run 2, a new hardware-based trigger known as the L1 Topological Trigger (L1Topo) was commissioned. L1Topo has several unique abilities, namely that it can calculate geometric and kinematic relationships between different trigger objects and it can receive inputs from both the L1Muon and L1Calo trigger paths. Some of these kinematic quantities used during Run 2 included topological measurements of $\Delta\phi$ and $\Delta\eta$ between objects and invariant-mass measurements of input objects for certain trigger chains [154].

The CTP receives the trigger information from the L1Muon, L1Calo and L1Topo processors and synchronizes the trigger inputs to the LHC collision clock in order to align all signals with the correct bunch crossing [155]. After synchronization, all of the trigger inputs are compared within the CTP to a series of trigger thresholds for each physics process being targeted. If the processed information passes at least one of the L1 thresholds, then the CTP issues a L1 accept (L1A) decision and propagates this decision to all ATLAS subdetectors. Upon receiving the L1A, all of the data from each subdetector for the given event is read out and propagated to the next step in the trigger chain. The choice of thresholds for the L1 trigger are determined by the maximum L1 output rate of approximately 100 kHz, and are typically chosen to select high transverse-momentum signals from the L1Muon chambers and large energy deposits in the L1Calo trigger system.

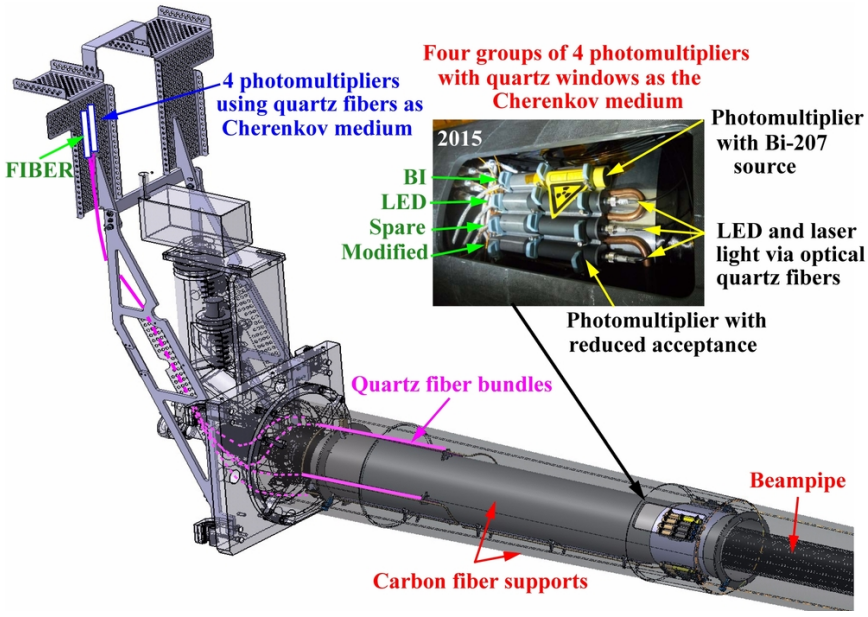

Figure 3.26: The LUCID detector [160].

High-level trigger

Events selected by the L1 trigger are then passed on to the high-level trigger. The HLT is a software-based trigger system that uses the full granularity of the ATLAS detector, including Inner Detector tracks, to make more refined decisions on which events to select for permanent storage. The HLT reduces the L1 event rate from 100 kHz to approximately 1 kHz of output. As a software-based system, the HLT trigger is limited by the processing time of an event. Furthermore, in order to prevent CPU overloads, most trigger decisions for an event must be made in $O(0.1\text{–}1\,\text{s})$, therefore CPU-intensive steps such as track reconstruction are only performed on the regions of interest coming from the L1 trigger. Data from the events selected by the HLT are then sent for permanent storage and processing at the Tier-0 Grid Computing site at the CERN Data Centre.

3.3.6 Forward detectors

The main ATLAS experiment is complemented with a number of specialized detectors installed in the forward regions. These include the Beam Conditions Monitor (BCM) [156], a diamond detector installed around the beam pipe on either side of the interaction point at $z = \pm 1.84\,\text{m}$ used for luminosity monitoring and for triggering beam aborts to protect the ATLAS detector in case of severe beam misalignment, the ALFA (**A**bsolute **L**uminosity **F**or **A**TLAS) detector [157], a set of two scintillating fiber trackers installed at $z = \pm 240\,\text{m}$ that use the Roman-pot technique to measure the elastic proton–proton scattering cross section in dedicated high-β^* LHC fills, the AFP (**A**TLAS **F**orward **P**roton) detector [158] located at $z = \pm 205\,\text{m}$ (near stations) and $\pm 217\,\text{m}$ (far stations), which uses a similar Romon-pot technique to measure diffractive proton–proton collisions, and the Zero Degree Calorimeter (ZDC) [159] located at $z = \pm 140\,\text{m}$, which measures the energies of neutral particles during low-μ proton–proton collisions and heavy-ion collisions.

The final forward detector, and that which is most relevant to the luminosity studies presented

in Chapter 6, is the LUCID (**LU**minosity **C**herenkov **I**ntegrating **D**etector) detector [160], shown in Figure 3.26. LUCID served as the primary luminosity-sensitive detector for ATLAS during Run 2, providing both online monitoring of the instantaneous luminosity and measurements of the total integrated luminosity. The LUCID detector contains 16 photomultiplier tubes (PMTs) in each forward arm of the ATLAS detector, placed at approximately $z = \pm 17\,\text{m}$ from the interaction point at a radial distance $R \approx 10\,\text{cm}$ from the beam line ($|\eta| \approx 5.8$). The PMTs are equipped with quartz windows wherein a traversing high-energy particle emits Cherenkov light that is measured by the PMT. The LUCID detector is read out with dedicated electronics, independent of the main ATLAS TDAQ system, that provide luminosity counts for each of the 3564 nominal LHC bunch slots where a colliding bunch pair could be present. LUCID provides several luminosity "algorithms" to convert the raw signals from the PMTs to a luminosity measurement, combining the information from several PMTs in various ways, which will be discussed in detail in Chapter 6 along with the methodology used to determine its absolute luminosity calibration. The LUCID detector used during Run 2 was upgraded from its Run 1 configuration to improve the linearity of its response and robustness against radiation damage in the higher-luminosity environment of Run 2. The Run 2 configuration used thinner quartz windows as the Cherenkov medium, a reduced window aperture to decrease their acceptance and thus avoid saturation of some luminosity algorithms, and finally some windows were coated with a radioactive ^{207}Bi source to monitor the gain stability of the PMTs.

Chapter 4

Collision Anatomy and Simulation

4.1 Introduction

The flagship ATLAS physics program is largely concerned with the high-luminosity proton–proton data sample. The high collision energy and large integrated luminosity of this data sample offer a robust environment in which to measure the production cross sections of rare Standard Model processes and the properties of the particles participating in theses interactions, and to search for physical phenomena beyond the Standard Model. In order to test the predictions of these theories experimentally to high precision, a full description of a proton–proton collision is required. Such a description offers a framework in which to relate the signatures observed in a particle detector such as ATLAS to the physical processes described by the theoretical model in question, and to disentangle these signatures from background processes that may also take place during the particle collisions. This framework also provides a means to perform *event simulation*, which allows for quantitative comparisons between theoretical predictions and empirical data.

Understanding the final states of high-energy proton–proton collisions at the LHC is, however, an extremely challenging theoretical problem given the composite nature of the proton. Typically hundreds of particles are produced in a collision between a single pair of protons, and in most processes of interest their momenta range over many orders of magnitude. This difficulty is compounded by the fact that the constituents of the proton are strongly interacting and thus involve the intrinsically non-perturbative and unsolved problem of confinement. Given the complexity of the collisions, analytical calculations of production cross sections using quantum field theory prescriptions, from the matrix-element computations to the integration over the vast final-states phase spaces, become so laborious as to be entirely impractical. Advanced modelling and computational techniques are therefore required to address these seemingly intractable problems.

This chapter describes the tools and methodologies currently employed to model the proton–proton collisions at the LHC. This description includes a discussion of the current understanding of the structure of the proton, the anatomy of a high-energy proton–proton collision, and finally how these collisions are simulated for comparisons to observations.

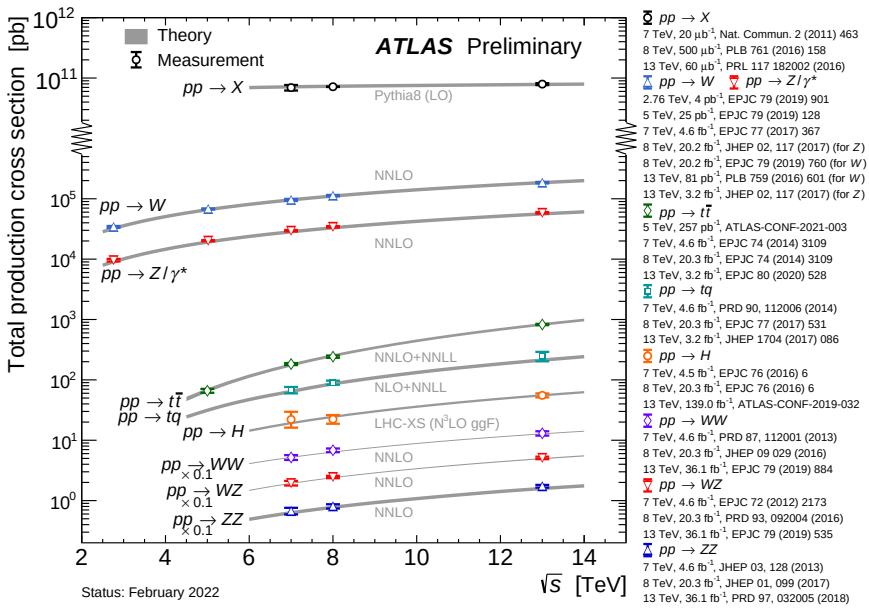

Figure 4.1: Summary of total inelastic proton–proton production cross-section measurements by ATLAS and their associated theoretical predictions presented as a function of centre-of-mass energy from 2.76 to 13 TeV for a few selected processes [163].

4.2 Anatomy of proton–proton collisions

When a pair of high-energy protons collide, their interactions can be broadly classified according to whether or not the initial-state protons remain intact after the collision. Interactions in which both protons stay intact and which contain no other final-state particles are said to be *elastic*. Collision events in which both protons continue their trajectories in the forward direction but with a small deflection relative to the incoming direction are characteristic of elastic scattering, and the rate of such events in governed by the elastic scattering cross section σ_{el}. In contrast, interactions in which one or both protons fragment into a spray of final-state hadrons are said to be *inelastic*, with the rate of these events governed by the inelastic scattering cross section σ_{inel}. The *total* proton–proton cross section, accounting for all possible interactions, is the sum of the elastic and inelastic cross sections: $\sigma_{\text{tot}} = \sigma_{\text{el}} + \sigma_{\text{inel}}$. These cross sections have been measured at the LHC to be approximately $\sigma_{\text{tot}} \simeq 110\,\text{mb}$, $\sigma_{\text{el}} \simeq 31\,\text{mb}$ and $\sigma_{\text{inel}} \simeq 79\,\text{mb}$ at a centre-of-mass energy of 13 TeV [161, 162]. Inelastic collisions thus account for approximately 70 % of all pp collision events at the LHC, and are responsible for the production of other particles, therefore they will be the focus of this discussion. A summary of the production cross sections for various processes and the total inelastic cross section measured by ATLAS as a function of centre-of-mass energy is shown in Figure 4.1.

The scattering properties of proton–proton collisions are governed by the energy at which they collide and by their internal structure. Analogous to the chemical and nuclear properties of an atom being determined by the number of protons and neutrons in its nucleus, the main quantum numbers of a hadron are determined by the composition of its *valence quarks*. In the case of the proton, these are the uud quarks bound together within the proton's volume. Within this volume there also exists a dynamic, interacting system of short-lived $q\bar{q}$ pairs, called *sea quarks*, which arise

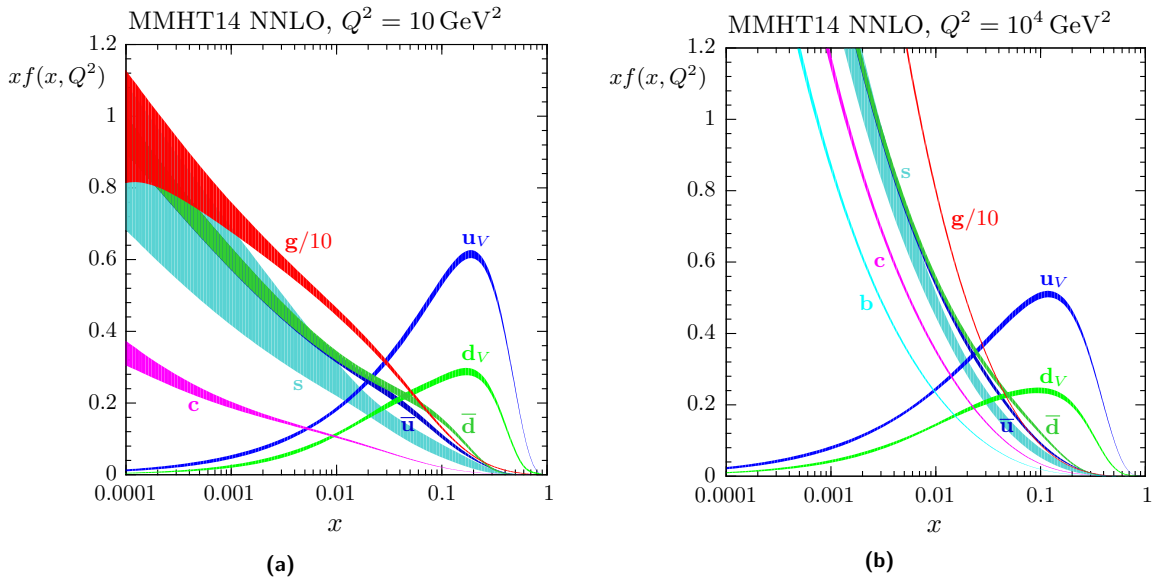

Figure 4.2: Examples of the MMHT2014 NNLO Parton Distribution Functions (PDFs) at (a) $Q^2 = 10\,\mathrm{GeV}^2$ and (b) $Q^2 = 10^4\,\mathrm{GeV}^2$ for different parton species as a function their momentum fraction x [167]. The shaded areas indicate the associated $\pm 1\sigma$ uncertainty bands.

through the continual exchange and splitting of gluons. Together, the valance quarks, sea quarks and gluons are referred to as *partons*. Given this composite nature of the proton, at high energies the collision itself is in fact occurring between the constituent partons of the protons, therefore the interactions in a pp collision can also be between sea quarks and between gluons, in addition to the interactions between valance quarks. The energy of the proton is divided among these many constituents, each carrying an unknown portion x of the proton's initial momentum, called the *momentum fraction*, with $0 \leq x \leq 1$. The distribution of momentum fractions for each constituent parton is expressed as the *parton distribution function* (PDF), which encodes the probability for a parton to carry a given fraction of the proton momentum at a particular value of the collision four-momentum transfer squared, or *collision energy scale*, Q^2. In practice, PDFs are not calculated from first principles. They are highly sensitive to the dynamics of the proton system and are therefore instead determined in global fits to experimental data, in particular from the results of deep inelastic scattering experiments and from jet measurements performed at a variety of centre-of-mass energies. PDFs are normally defined for a particular starting scale Q_0^2 and then "evolved" to other scales via the DGLAP (Dokshitzer–Gribov–Lipatov–Altarelli–Parisi) equations [164–166]. A sample set of PDFs are shown in Figure 4.2 for two different values of Q^2.

If two colliding partons interact at sufficiently high energy, typically with $Q^2 > 100\,\mathrm{GeV}^2$ at the LHC, they undergo what is referred to as a *hard scatter*. Hard-scatter processes are classified according to the final-state or intermediate particles produced in the interaction, including QCD multijet production, top quark processes, $Z/W^{\pm}$ boson production, Higgs boson production, and so on. The large collision energy scale of the hard scatter places the interaction in the perturbative regime of QCD where the strong coupling constant is small, $\alpha_s \sim 0.1$, and where the partons can be treated as free particles. The matrix elements of the hard-scatter process and the associated cross section can therefore be calculated using perturbative QCD. To maintain computational efficiency, these

calculations are typically performed at fixed-order in perturbation theory, with most contemporary event-generation programs implementing these calculations at LO or NLO.

Integrating over the full final-state phase space of the hard-scatter interaction introduces divergences at the high-energy limit (UV) and low-energy limit (IR). The UV divergences occur in virtual loops and are eliminated by introducing a *renormalization scale* μ_R upon which the QCD coupling α_s becomes dependent. The theory is then defined up to the renormalization scale energy. The most common scheme for absorbing this divergence is the modified minimal subtraction scheme ($\overline{MS}$) [168]. The IR limit introduces divergences when a parton is emitted at a small angle (the *collinear limit*) in the initial state before the hard scatter. This divergence can be solved using factorization theory, which states that the cross section can be factorized into a hard-scatter portion and a normalization portion determined by the PDFs. This introduces a factorization scale μ_F below which the physics is absorbed into the PDF of the parton. Above the factorization scale the physics is described by the matrix element of the hard-scatter process, thus the collinear limit is omitted by integrating from μ_F instead of from 0.

Following this prescription, the cross section for the collision of two partons a, b in a hadron collider to some final state particles n, $ab \to n$, can be computed as [169]

$$\sigma = \sum_{a,b} \int_0^1 dx_a dx_b \int d\Phi_n f_a(x_a, \mu_F) f_b(x_b, \mu_F) \frac{1}{2\hat{s}} \left|\mathcal{M}_{ab\to n}\right|^2 (\Phi_n; \mu_F, \mu_R) \,, \tag{4.1}$$

where $f_{a(b)}(x_{a(b)}, \mu_F)$ is the PDF for parton a (b) with momentum fraction x dependent on the factorization scale, and $|\mathcal{M}_{ab\to n}|^2(\Phi_n; \mu_F, \mu_R)$ is the matrix element for the process $ab \to n$ dependent on the final state phase-space momenta Φ_n and the renormalization and factorization scales. Here, $\hat{s} = x_a x_b s$ is the parton centre-of-mass energy squared and s is the hadron centre-of-mass energy squared. The matrix element is computed up to some order in $\alpha_s(\mu_R)$ by summing contributions from all possible $ab \to n$ Feynman diagrams, which defines the precision of the calculation (LO, NLO, NNLO, ...). In LO calculations the dependence on these scales μ_F and μ_R is most significant, while at higher order this dependence is reduced with a careful choice of scales. The uncertainty associated to the choice of μ_F and μ_R is estimated by varying the scales and observing the impact on the cross section, typically varying up and down by a factor of two.

While the hard scatter is the core physics process of interest at a hadron collider, a full description of a proton–proton collision involves a number of other processes that require modelling in order to make comparisons between quantum field theory predictions and observed detector signals. These many different processes are illustrated in Figure 4.3.

First is the *parton shower* in which the strongly interacting initial- and final-state partons "split" through the radiation of gluons, producing a cascade of additional quarks and gluons with progressively decreasing energy. This shower proceeds until the energy of the cascading quarks and gluons reaches the perturbative limit of QCD around 1 GeV. At this energy scale, α_s is of $O(1)$ and non-perturbative effects take over, resulting in *hadronization*. The combination of parton shower splitting and hadronization results in a collimated spray of hadrons observed as *jets* in the detector, with each jet roughly corresponding to a single parton from the hard scatter process. Hard splittings are also referred to as *final state radiation* (FSR), typically from gluons radiating from a final-state quark. When QCD radiation is emitted from the partons before the hard scatter it is similarly referred to as *initial state radiation* (ISR).

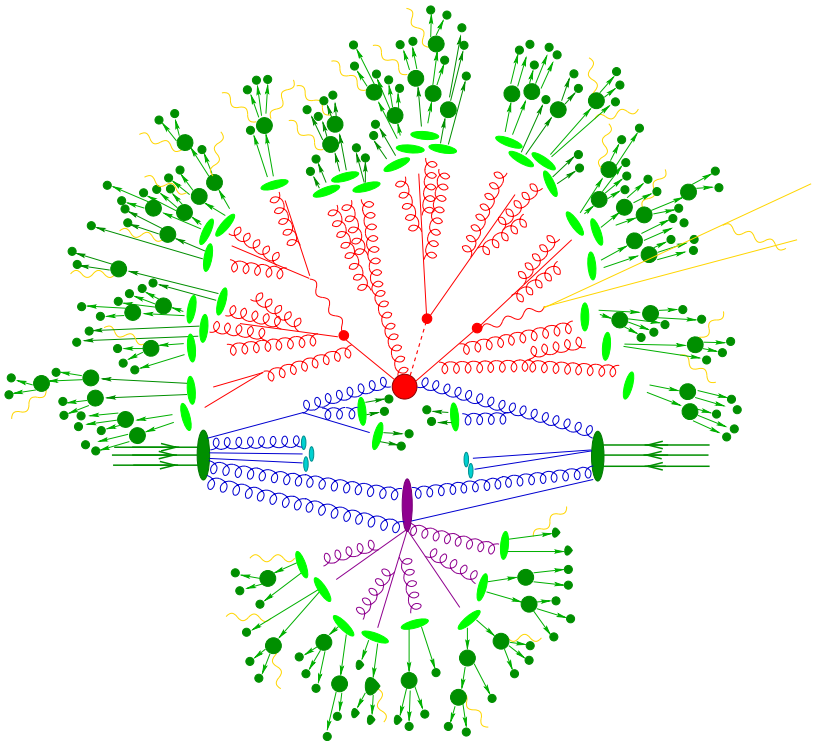

Figure 4.3: Illustration of a typical proton–proton collision at the LHC. The hard-scatter interaction of the event (big red circle) is followed by the decays of the hard-scatter products (small red circles), which in turn release final-state radiation in the form of gluons (red lines). The gluons eventually shower into secondary particles which hadronize (light green blobs) and then decay (dark green blobs). Initial-state radiation is also released as gluons prior to the hard-scatter interaction (blue). The proton remnants may also interact in what is called the *underlying event* and yield hadronizing products (purple). Electromagnetic radiation is released as photons in any stage of the event (yellow lines). Adapted from Ref. [170].

Second is the *underlying event* in which additional "soft" interactions not originating from the primary hard scatter occur, primarily between other partons from remnants of the colliding protons. The products of the underlying event hadronize, creating radiation preferentially in the forward regions, close to the beam line. Finally, the remaining partons in the primary hard scatter can also themselves scatter; these are referred to as *multiple parton interactions* (MPIs). Most MPI are soft and do not give rise to isolated jets like those coming from the hard scatter. Instead, MPIs increase the total amount of scattered energy in the event and cause additional colour exchanges among the remnants, resulting in an overall greater number of particles produced in the hadronization stage [169].

4.3 Simulation

The complexity of a proton–proton collision, from the hard scatter to the parton shower, soft interactions and underlying event, make Monte Carlo (MC) methods the preferred tools for simulating these processes. This section describes the complete event-simulation chain, from the hard-scatter matrix-element calculation to the detector simulation.

4.3.1 Event simulation

The first stage of the event simulation generates events according to the physical theory that models the process in question. Due to the varying levels of precision required by high-energy physics analyses, a broad selection of MC programs have been developed, and can be categorized in one of the following groups: matrix-element generators, parton shower programs, dedicated-purpose programs, and multipurpose programs.

Matrix-element generators are those that provide matrix-element calculations for the hard-scatter process. The calculation of fixed-order scattering amplitudes by these Monte Carlo programs begins from the same underlying equations as any analytic quantum field theory calculation. The relevant Feynman diagrams for a given process are included, and an expression for the differential cross section $d\sigma$ is derived. Instead of evaluating the ensuing integral analytically in order to calculate the total cross section, the differential cross section is used as input for the generation of pseudorandom phase-space points that are sampled or weighted to reproduce the original distribution. Typical matrix-element generators implement either LO or NLO calculation (in QCD or electroweak corrections) in order to maximize computational efficiency, while some specialized generators have also been developed to perform these calculations at higher orders. The PDFs are used as input to the matrix-element calculation to determine statistically which partons interact in a given simulated collision. The predicted cross sections from these generators are generally inconsistent with data or of limited precision, requiring the MC events to be weighted to a higher-order calculation by means of a *k-factor* to accurately describe the data. Most matrix-element calculation programs are interfaced to another program to simulate parton showering and provide a full description of the simulated event.

Parton shower programs are used to simulate the soft gluon radiation coming from partons involved in the hard scatter of the event and also to simulate hadronization. The parton shower algorithm starts with the LO process computed directly via perturbative QCD and generates multiple splitting sequences with assigned probabilities. The LO cross section is thus partitioned into many cross sections for final states with different distributions and multiplicities of final-state quarks and gluons, whose sum is equal to the cross section of the primary process. At the tail end of the shower, hadronization effects dominate, which are simulated using models tuned to data. The most commonly implemented hadronization models are the *Lund string model* [171, 172], in which soft gluons are treated as field lines that are attracted to each other due to the gluon self-interaction that collapse into a narrow tube (or "string") when the separation of the sources becomes much larger than $\sim 1\,\text{fm}$, and the *cluster model* [173], in which hadronization occurs via "preconfinement" of colour-singlet clusters into "proto-hadrons" that decay into the observed final-state hadrons.

Special-purpose generators are used to simulate specific physics processes, such as the decays of b hadrons, τ leptons and Higgs bosons. Multipurpose generators include any generator that can accomplish multiple functions, typically including the matrix-element calculation and the parton showering under one program.

In addition to the theoretical descriptions above, the effect of additional soft interactions from pileup and effects from background radiation in the experimental cavern are also necessary for an accurate description of events. Pileup event modelling often involves the generation of additional proton–proton collisions that are then overlaid on the primary hard-scatter process. For the cavern and machine effects, dedicated triggers with loose selection criteria are used to record *minimum bias events* during regular running to collect a data sample of inelastic collisions with as little bias as

possible. These events are then inserted into the MC simulation to provide an accurate description of these low-energy processes during collisions with hard-scatter interactions.

The final output of an individual event simulation is a list of final-state particles and their four-momenta, including only the particles with a lifetime τ long enough that they can be detected by an experiment (typically such that $c\tau > 1\,\text{mm}$). Observables of these final-state particles from simulation, such as their p_{T}, charge, identity and so on, are referred to as being at *particle-level*. These outputs are then passed to the detector simulation to model how collision events appear in the detector itself, described in the following section. The specific generators used in the search for heavy resonances decaying into pairs of Z bosons are described in Section 7.2.

4.3.2 Detector simulation

Quantum field theory calculations combined with phenomenological models are used to predict the distributions of particles produced in proton–proton collisions at the LHC. To make a comparison with observations, the interactions of these particles in the detector material and the resulting signals produced by the detector must also be understood and simulated. The detector simulation involves propagating the full particle content of the simulated events described in the previous section through a model of the detector volume. As the particles propagate through the detector they interact with it: a particle may leave energy deposits in the detector, while simultaneously having its path or even identity altered by the interactions. The simulated energy deposits then undergo a process called *digitization*, which translates the energy deposits into a set of detector signals, as would be seen in a real data event. These digitized signals can then be treated exactly as real signals in the detector would. The same reconstruction algorithms are applied to the simulated signals as are applied to real data events, producing detector-level objects such as reconstructed electrons, muons, photons and jets, with their corresponding observables said to be at *reconstruction-* or *reco-level*. This allows for a direct comparison of the simulated events with data.

The ATLAS detector simulation [174] is performed using GEANT4 [175]. The GEANT4 software suite allows for definitions of the detector geometry, including the material components, descriptions of the sensitive materials and the readout geometry. The physics models implemented in the software handle the interactions of particles with matter across a wide energy range and include interactions such as particle transport through magnetic fields, ionization, Coulomb scattering, the photoelectric effect, particle annihilation, Compton scattering, photon conversion to e^+e^- pairs, bremsstrahlung radiation, and nuclear interactions. The implementation of the ATLAS detector geometry within GEANT4 is done in detail, including several hundred different materials which form hundreds of thousands of different physical volumes.

The full ATLAS detector simulation also includes a simulation of the trigger and data acquisition systems, and the ability to reproduce actual running conditions, for example by disabling detector elements or updating the detector geometry. Although the detector is nominally stable, many changes to the real detector geometry occur, such as the addition of new subdetectors, or slightly different positions of detectors due to interventions, long-term drift, or changes in temperature. The simulation also includes the overlay of multiple simultaneous collisions, the impact of signals in the detector from previous bunch crossings, and interactions with beam gas. Detector simulation is a very computationally intensive procedure, requiring significant computing resources. The most CPU-intensive simulations of the ATLAS detector are those involving the calorimeters, where complex

showering processes involving cascades of secondary particles occur. To reduce the computational complexity of simulations for analyses that do not require the precision of the full detector simulation, ATLAS has developed fast simulation methods that reduce processing times by approximately an order of magnitude [176]. This performance is achieved by means of a simplified detector geometry and a parametric detector response based on templates derived from a large number of precomputed GEANT4 simulations.

4.4 Summary

In order to test the predictions of the Standard Model and its proposed extensions experimentally to high precision using a data sample of proton–proton collisions, a full description of these collisions is required. The proton is modelled as a "sea" of quarks and gluons, collectively referred to as partons, with its main quantum numbers determined by its valance quarks. The energy of the proton is divided among these partons, with the distribution of each parton's momentum fraction expressed as a parton distribution function (PDF). Generally the collisions of interest at a hadron collider are those containing a hard-scatter interaction, which occurs when two partons collide at sufficiently high energy. The simulation of a proton–proton collision uses Monte Carlo (MC) methods to compute the matrix element of the hard-scatter interaction and the associated scattering cross section, model the ensuing parton showers and hadronization effects, and model other "periphery" interactions among the proton remnants, such as initial- and final-state radiation and the underlying event. In the recorded data sample, the effects of the detector resolution "smear" the reconstructed kinematics of the final-state particles, which is modelled using a full GEANT4 detector simulation or an equivalent, but less complex, "fast" detector simulation. All of these simulation techniques are used extensively in the search for heavy resonances presented in Chapter 7; the specific MC generators and modelling parameters used in this analysis are given in Section 7.2.

Chapter 5

Dataset and Object Reconstruction

5.1 Introduction

Every proton bunch crossing on average produces many interactions, each of which yields many particles emitted from the interaction region. These particles propagate through the detector, interact with the sensitive detector elements, and in certain cases produce many more (secondary) particles. The ATLAS TDAQ system performs initial event reconstruction using the raw information from the detector and searches for physics signatures of interest, as described in Section 3.3.5. More robust object reconstruction and calibration is performed offline for events passing the trigger decisions to identify the collision products, measure their kinematic properties and encapsulate these measurements into *objects* for use in ATLAS physics analyses. The number of recorded collision events accumulates over time as the LHC delivers stable beams to ATLAS, providing greater statistical power with which to measure the properties of the particles produced in the collisions and to search for new phenomena. This chapter first chronicles the ATLAS Run 2 data sample used in this work, and the conditions under which the data were recorded, and then describes the methodologies employed to reconstruct the objects of interest in this data sample and in simulation.

5.2 Dataset

During normal running conditions, ATLAS records data continuously during an LHC fill. Once stable beams are declared, the Pixel, SCT and muon systems begin their high-voltage ramp up, and the Pixel preamplifiers are turned on, after which ATLAS is declared "ready for physics". Each of the datasets taken while ATLAS is continuously recording is referred to as an ATLAS *run*. Each run is further divided into *luminosity blocks* (LBs) during which the instantaneous luminosity, detector and trigger configuration and data-quality conditions are considered constant. One luminosity block typically corresponds to a time interval of 60 s, although their duration is flexible and actions that alter the run configuration or detector conditions may trigger the start of a new luminosity block. In general, the luminosity block serves as a standard unit of time for data-quality monitoring, although the rejection of inferior-quality data is also possible at a much finer granularity (either event-by-event or within O(ms) "time windows") provided that the grounds for rejection are identified in time to be taken into account during the data reconstruction.

The quality of the recorded data is assessed both in real-time *online* monitoring and in the *offline* data reconstruction [177]. The purpose of the online monitoring is to flag, as quickly as possible, any detector failures or reconstruction issues in order to identify any necessary interventions, thereby limiting the amount of unrecoverable data resulting from severe detector coverage losses or data corruption. The offline data-quality assessment ensures that the final dataset used in ATLAS physics analyses only includes the periods of data collection in which the detector is functioning within acceptable parameters.

Data quality is monitored in a number of dedicated *streams*, defined by a set of trigger selections and containing all events that have been stored to disk after satisfying any of those selections. The *physics streams*, for instance, record the events of interest for physics analyses, and the *calibration streams* record events used for alignment and calibration measurements. Most data-quality monitoring is performed using the *express stream*, which contains a representative subset of the data collected in the physics and calibration streams. Other dedicated data-quality-monitoring streams are used alongside the express stream to identify specific detector issues, such as the *SCTNoise* stream, which identifies noisy SCT strips in empty bunch crossings to be masked in the data reconstruction, the *CosmicCalo* stream, which provides an analogous function for noisy calorimeter cells, the *LArCellsEmpty* and *LArNoiseBurst* streams, which allow for detailed monitoring of the LAr calorimeter cells and for identifying and masking events with noise bursts, respectively, and finally a number of "global" Inner Detector streams for precise determination of the position of its sensitive elements and of the beam spot. Problems found in any of these streams that could affect data quality are flagged and logged in a dedicated defect database. These defects are generally classified as being either *tolerable* or *intolerable*. Data with a tolerable defect is flagged and recorded for book-keeping purposes but can still be used for physics analyses. An intolerable defect results in rejection of the effected *interval of validity* (e.g. the luminosity blocks during which the data was affected) and thus cannot be used for analysis. Intolerable defects generally arise as a result of detector hardware issues, such as the noise bursts mentioned above, high-voltage trips, cooling leaks leading to loss of detector coverage, and so on. The defect database is then used to produce a list of luminosity blocks and runs that are declared "good for physics", called a *Good Run List* (GRL). The integrated luminosity of a physics dataset, or the "good for physics" integrated luminosity, is calculated by summing over the luminosity blocks included in the GRL. After these data-quality selections, the full Run 2 pp data sample at centre-of-mass energy of $\sqrt{s} = 13\,\text{TeV}$ corresponds to an integrated luminosity of $139\,\text{fb}^{-1}$, corresponding to a data-taking efficiency of 95.6 % [177]. The cumulative integrated luminosity delivered to ATLAS as a function of time in Run 2, along with luminosity recorded by ATLAS, and the luminosity certified as "good for physics", is shown in Figure 5.1.

5.2.1 Pileup

As the per-bunch instantaneous luminosity increases, so too does the average number of inelastic pp interactions per bunch crossing $\langle\mu\rangle$. Therefore, in addition to the hard-scatter interaction of interest, other soft collisions may also occur simultaneously. These additional collisions are collectively referred to as *pileup*, and present a serious challenge to detector subsystems and physics analyses as their presence obscures the signal from the hard-scatter interaction. The evolution of the pileup conditions over the years of data-taking in Run 2 is shown in Figure 5.2 as luminosity-weighted distributions of $\langle\mu\rangle$ per year and for the total Run 2 dataset.

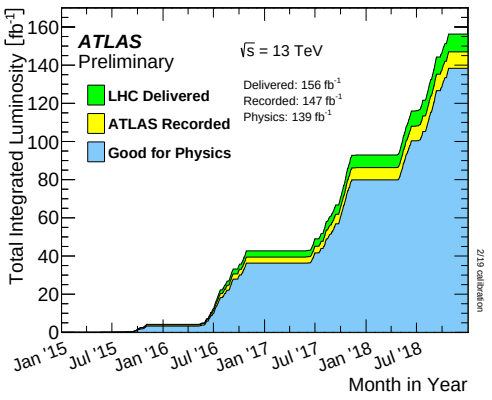

Figure 5.1: Cumulative luminosity versus time delivered to ATLAS (green), recorded by ATLAS (yellow), and certified to be good quality data (blue) during stable beams for *pp* collisions at 13 TeV centre-of-mass energy in the years 2015–2018.

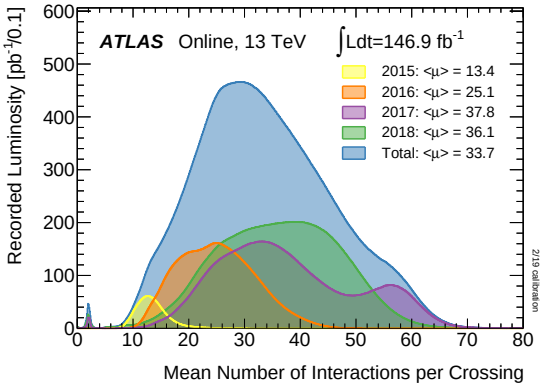

Figure 5.2: Luminosity-weighted distribution of the mean number of interactions per bunch crossing for the Run 2 *pp* collision data at 13 TeV centre-of-mass energy in 2015–2018.

The spatial resolution of the ATLAS readout cells limits the ability to resolve large numbers of particles traversing its volume, therefore the reconstruction performance of nearly every physics object is degraded by pileup to some degree. Specifically, this source of pileup is called *in-time* pileup as it involves multiple simultaneous inelastic *pp* collision occurring in the same bunch crossing. Moreover, many ATLAS subsystems have readout signals much longer than the 25 ns LHC bunch spacing, therefore the signals from adjacent bunch crossings may interfere with one another, further degrading the reconstruction performance; this is known as *out-of-time* pileup and increases with $\langle\mu\rangle$. Other sources of beam-related background exist, including *cavern background*, in which a gas of low-energy particles fills the ATLAS cavern during an LHC fill, contributing mostly to random hits in the muon system,[1] *beam halo backgrounds* originating from protons scraping against an upstream collimator, resulting in showers of predominantly muons nearly parallel to the beam line, and *beam gas backgrounds* resulting from collisions between protons in a bunch and any residual gas inside the beam pipe.

In the high-luminosity conditions at the LHC, the dominant sources of beam-induced background are in-time and out-of-time pileup. In Run 2, each hard-scatter event was on average accompanied by ~ 34 other inelastic *pp* collisions that produce additional tracks in the Inner Detector and showers in the calorimeters, degrading the reconstruction performance of the objects arising from the hard-scatter event of interest. Out-of-time pile-up, on the other hand, affects detectors that have a response time longer than the time interval between consecutive LHC bunch crossings. For example, the LAr calorimeters have a pulse duration of approximately 450 ns and are therefore susceptible to distortions due to energy deposits from collisions in adjacent bunch crossings.

The pileup backgrounds can, however, be modelled and accounted for in physics analyses. Generally the average number of interactions per bunch crossing $\langle\mu\rangle$ is not known *a priori* and the values used for simulation must be corrected during an analysis to reflect the actual running conditions in the data. This correction is performed by reweighting the generated $\langle\mu\rangle$ profile in the simulation to

[1] Cosmic rays also contribute to this background, in which case the rate is constant on average and can be present even without collisions.

the distribution found in the data as a function of data-taking period (and taking into account the good running conditions described above). This *pileup-reweighting* correction is performed on each Monte Carlo sample individually.

5.3 Object reconstruction, calibration and selection

The signals collected from the various ATLAS detector elements are combined in a complementary way to identify and *reconstruct* the particles that generated these detector signals for each recorded collision event. The reconstruction procedure relies on the identification of expected patterns in each of the ATLAS subsystems based on knowledge of particle properties and their interactions with the detector materials, and on simulation studies. Most particle reconstruction begins with identifying low-level objects, including *tracks* left by charged particles in the Inner Detector and Muon Spectrometer and *energy clusters* in the calorimeters. These low-level inputs are then combined in various ways to reconstruct the *physics objects* used in analyses, such as electrons, photons, muons and jets, as shown by the illustration in Figure 5.3. For instance, electron candidates are reconstructed from tracks in the Inner Detector matched to energy clusters recorded in the electromagnetic calorimeter. Photons leave a similar signature in the electromagnetic calorimeter but are not associated to any Inner Detector tracks. Jets and other hadronic particles deposit large amounts of energy in both the electromagnetic and hadronic calorimeters, with charged hadrons also leaving tracks in the Inner Detector. Since muons can potentially traverse all layers of the ATLAS detector, muon candidates can be reconstructed in a variety of ways, combining information from tracks recorded in the Inner Detector and in the Muon Spectrometer. Muons also leave minimal energy deposits in the calorimeter systems, and this information can be combined with Inner Detector tracks in regions with limited Muon Spectrometer coverage. The only long-lived Standard Model particles that cannot be identified directly are neutrinos due to their very weak interactions with the detector material. However, their presence may be inferred in events with large "missing energy" based on energy-balance requirements.

Generally, analyses are interested in the *prompt* physics objects originating directly from the hard-scatter interaction, and require precise reconstruction of their position, energy and momentum. To this end, the object reconstruction methodologies employed by ATLAS are optimized to ensure high reconstruction efficiency and background rejection for these prompt objects over a broad range of final-state phase space, and are carefully calibrated to perform unbiased measurements of the particles' kinematics. Multiple definitions, or *working points*, are generally available to reconstruct any given physics object to cover the wide variety of requirements throughout the ATLAS physics program. These working points apply different selection criteria in order to reject, to varying degrees, objects with poorly resolved kinematics or those prone to misidentification. "Loose" working points generally improve reconstruction efficiency at the cost of reduced signal purity, while "tight" working points preferentially select the "highest quality" objects whose kinematics and identity are known with high confidence, at the cost of reduced reconstruction efficiency.

This section describes the methods employed by ATLAS analyses to reconstruct, calibrate and select the physics objects of interest. Object-reconstruction methods evolve over time as ongoing studies of the detector performance allow for more sophisticated procedures to identify these objects and calibrate the measurements of their kinematics. This discussion will include a general overview

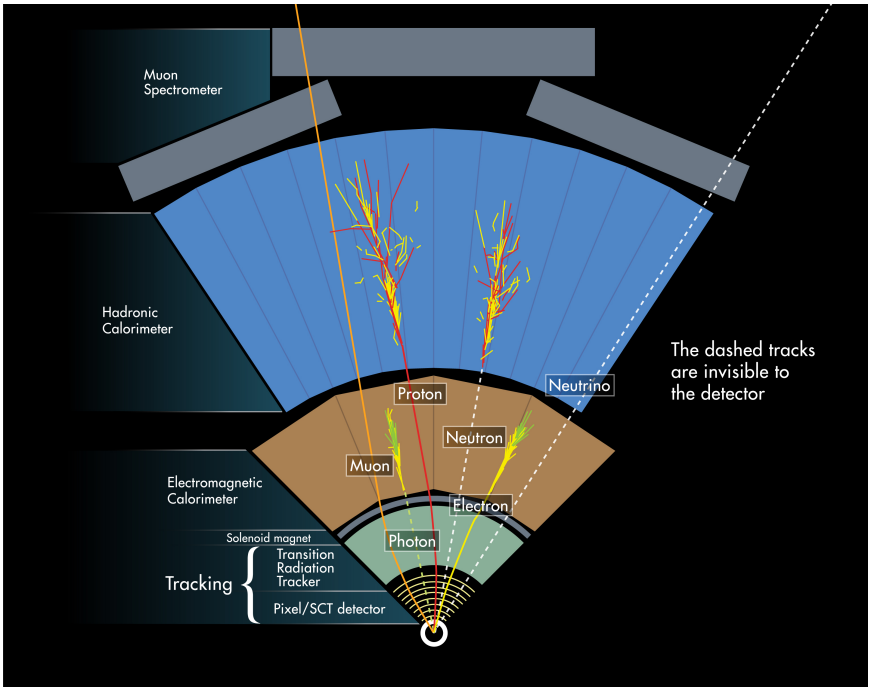

Figure 5.3: Illustration of an ATLAS detector segment in the transverse plane showing the detector elements used for the reconstruction of a variety of particle species emitted from the interaction point [178]. In general, only those particles that are sufficiently long-lived and that interact with the sensitive detector elements can be fully reconstructed.

of these methods in Run 2, with greater attention paid to electron and muon reconstruction, and to the specific methods used in the search for heavy resonances decaying to a pair of Z bosons presented in Chapter 7. The specific physics object selections and working points used in this analysis are given in Section 7.3.1 for the $\ell^+\ell^-\ell^+\ell^-$ channel and in Section 7.4.1 for the $\ell^+\ell^-\nu\bar{\nu}$ channel.

5.3.1 Tracks

The reconstruction of charged-particle tracks primarily follows an "inside-out" algorithm that begins with hits recorded in the Pixel (including the IBL) and SCT detectors, and then extends the reconstructed track segments into the TRT [179].

The measurements from the Pixel and SCT detectors are first clustered by matching hits above threshold where the sensitive detector elements share common edges or corners. A single cluster can be created by one or multiple particles, corresponding to *single-particle* clusters and *merged* clusters arising from highly collimated charged particles, as illustrated in Figure 5.4. The clusters are then used to create three-dimensional *space-point* measurements. Space-points in the Pixel detector are formed from a single cluster, while space-points in the SCT require hits on both sides of a layer to determine the space-point's position along the strip. A combinatorial Kalman filter [180] is then used to build track candidates from track *seeds*, formed from groups of three space-points in depth, by incorporating additional space-points from the remaining layers of the Pixel and SCT detectors compatible with the preliminary trajectory of the seed. Multiple track candidates may result from an initial seed and track candidates may share common clusters, therefore in order to resolve these ambiguities, track candidates are ranked according to a *track score* weighted by their

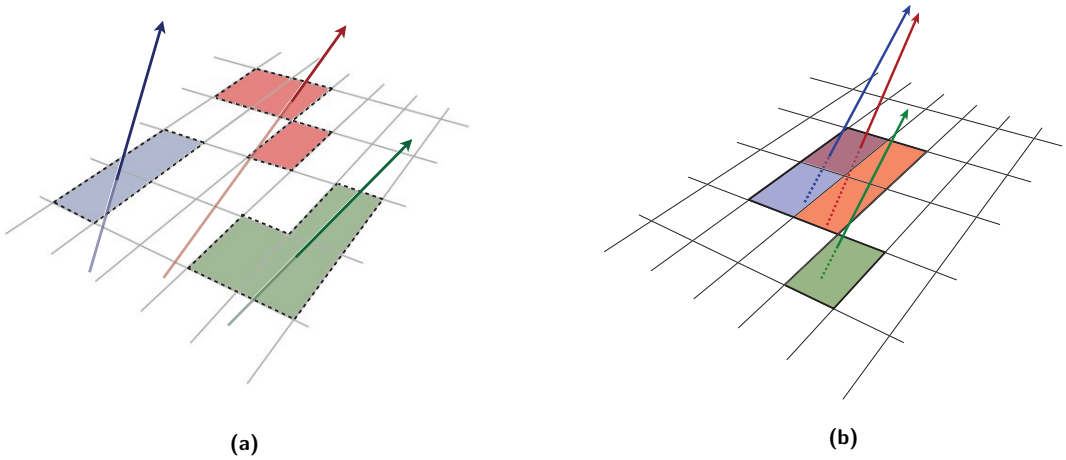

Figure 5.4: Illustration of (a) single-particle pixel clusters on a pixel sensor and (b) a merged pixel cluster as a result of highly collimated charged particles [179]. The particles trajectories are shown as arrows and the different colours represent energy deposits from each of the charged particles traversing the sensor.

cluster multiplicity, χ^2 goodness-of-fit and momentum. The track ambiguity solver proceeds in an iterative manner in which clusters shared between two track candidates are removed from the low-scoring candidate, which is then either scored again and returned to the ordered list of remaining candidates, or rejected if it fails a set of criteria on its momentum, number of hits, number of holes[2], and impact parameters. The selection criteria used in Run 2 are defined in Ref. [179].

The track segments reconstructed in the silicon layers of the Inner Detector selected by the above procedure are then extended into the TRT. First, the trajectory of the silicon-hit segment is extrapolated into the TRT volume and hits compatible with the extrapolation are added to the candidate. Second, a fit is performed on each track candidate, its score is computed using an analogous set of ambiguity-resolving requirements, and the selected candidates are added to the final track collection. Tracks with TRT extensions generally have improved momentum resolution, owing to the extended track length, and are of particular importance for electron identification (Section 5.3.4).

An alternative *outside-in* track-reconstruction algorithm is also used, where hits in the TRT are extended inward and matched to hits in the silicon layers of the Inner Detector. This method recovers some lost track-reconstruction efficiency due to long-lived particles that decay in the silicon tracking layers and from photon conversions, and to recover low-p_{T} electrons with large relative energy losses due to bremsstrahlung interactions [181].

5.3.2 Vertex reconstruction

The nominal interaction point, despite its name, is not a point, but rather a volume in space called the *luminous region* roughly corresponding to the overlapping region of the colliding proton bunches. The position of individual proton–proton interactions within this region (including both the hard-scatter and pileup interactions) can be resolved by grouping together reconstructed charged-particle tracks consistent with originating from the same initial position, called a *vertex*.

[2]Here, a *hole* is defined as an intersection of the reconstructed track trajectory with a sensitive detector element that does not contain a matching cluster.

Vertex reconstruction is performed using an adaptive vertex-finding algorithm [135, 182] and begins by associating tracks to an initial vertex seed position approximately at the centre of the luminous region. Next, compatible tracks are combined with this seed, and an iterative χ^2 fit is performed by assigning a weight to every track reflecting its compatibility with the vertex estimate, where less compatible tracks are down-weighted in the next iteration and the vertex position is recomputed. After the last iteration, the vertex position is determined, and tracks that are incompatible with the vertex by more than 7σ are removed and allowed to be used in the determination of another vertex. The process is repeated using all the removed tracks until no additional vertices can be found, with the requirement that a vertex has at least two tracks associated with it. For events with multiple reconstructed vertices, the *primary vertex* is chosen to be the one with the largest $\sum p_{\mathrm{T}}^2$ of its associated tracks and is used as a proxy for the hard-scatter vertex, with all other vertices assumed to originate from pileup interactions.

5.3.3 Calorimeter clustering algorithms

Clustering algorithms group the signals from neighbouring calorimeter cells into a single entity roughly corresponding to the volume in the calorimeter where energy was deposited from an electromagnetic or hadronic shower. The algorithms are tuned to extract the signal in the calorimeter from a background of electronic noise and other sources of fluctuations such as pileup.

In Run 2, ATLAS employed a dynamic topological cell-clustering algorithm [183] in which adjacent calorimeter cells are grouped into *topological clusters* (or *topo-clusters*) containing significantly more energy than expected from noise alone.[3] The algorithm begins by identifying seed cells with especially high *signal significance* $\varsigma_{\mathrm{cell}}^{\mathrm{EM}}$, defined as the ratio

$$\varsigma_{\mathrm{cell}}^{\mathrm{EM}} = \frac{E_{\mathrm{cell}}^{\mathrm{EM}}}{\sigma_{\mathrm{noise,cell}}^{\mathrm{EM}}}\,, \tag{5.1}$$

where $E_{\mathrm{cell}}^{\mathrm{EM}}$ is the cell energy at the EM scale[4] and $\sigma_{\mathrm{noise,cell}}^{\mathrm{EM}}$ is the expected cell noise, which includes the known electronic noise and an estimate of the pileup noise corresponding to the average instantaneous luminosity delivered in Run 2. The seed threshold is chosen as $|\varsigma_{\mathrm{cell}}^{\mathrm{EM}}| \geq 4$. The algorithm then proceeds iteratively collecting neighbouring cells with significance $|\varsigma_{\mathrm{cell}}^{\mathrm{EM}}| \geq 2$ to form *proto-clusters*, with each cell passing the threshold of $|\varsigma_{\mathrm{cell}}^{\mathrm{EM}}| \geq 2$ becoming a seed cell in the next iteration, collecting each of its neighbours in the growing proto-cluster. If two proto-clusters contain the same cell above the noise threshold, they are merged. The iteration proceeds until all candidate cells with $|\varsigma_{\mathrm{cell}}^{\mathrm{EM}}| \geq 2$ have been exhausted, at which point all cells adjacent to those on the periphery of the proto-cluster are added independent of their energy to form the final cluster. Proto-clusters with two or more local maxima are split into separate clusters. Here, a cell is considered a local maximum when it has $E_{\mathrm{cell}}^{\mathrm{EM}} > 500\,\mathrm{MeV}$, at least four neighbours, and when none of the neighbours has a larger signal. The use of the finely granular first sampling layer of the electromagnetic calorimeters in the topo-cluster splitting, for instance, is particularly important for neutral pion rejection since it improves the photon separation in $\pi^0 \to \gamma\gamma$ decays.

[3] In the case of electron and photon reconstruction, topological cell-clustering replaced a sliding-window algorithm previously exploited by ATLAS in Run 1 and in early Run 2 results for the reconstruction of fixed-size clusters of calorimeter cells [184, 185].

[4] The electromagnetic (EM) scale is the signal scale at which the energy deposited in the calorimeter by electromagnetic showers is correctly reconstructed.

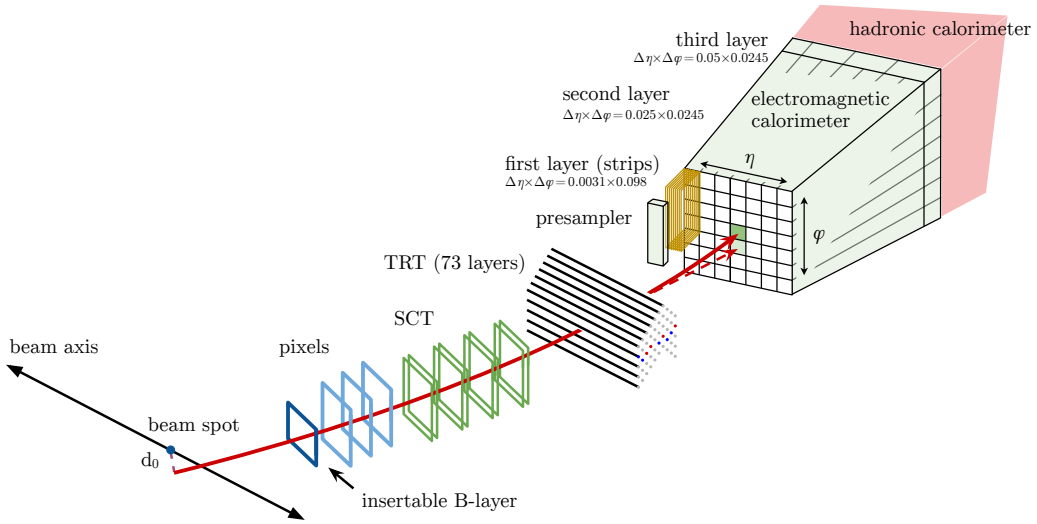

Figure 5.5: Schematic illustration of the path of an electron through the ATLAS detector [184]. The red trajectory shows the hypothetical path of an electron, which first traverses the tracking system and then enters the electromagnetic calorimeter. The dashed red trajectory indicates the path of a photon produced by the interaction of the electron with the material in the tracking system.

Topo-clusters are reconstructed in each recorded event and used as input for photon, electron and jet reconstruction, the latter of which are described in the following sections.

5.3.4 Electrons

Central electron candidates are reconstructed in the precision region of the ATLAS detector ($|\eta| < 2.47$) from clusters of energy deposits in the calorimeter associated with Inner Detector tracks [186]. The path of an electron through this region of the detector is illustrated in Figure 5.5. Forward electrons falling outside the acceptance of the Inner Detector in the range $2.47 < |\eta| < 4.9$, while not considered in the analysis work of this thesis, can also be reconstructed using only calorimeter information [187, 188].

Electron candidates are first identified by the presence of energy deposits clustered into electromagnetic *superclusters*, illustrated in Figure 5.6. Superclusters consist of a main "seed" topo-cluster and any number of "satellite" topo-clusters within a given window around the seed cluster barycentre, which may arise in situations where an electron radiates a bremsstrahlung photon within the volume of the Inner Detector or from topo-cluster splitting. The track–cluster matching is then performed, after the tracks have been fitted with a Gaussian-sum filter (GSF) [189] to account for bremsstrahlung energy losses, by extrapolating candidates tracks into the calorimeter and matching them with the candidate supercluster if they satisfy $|\eta_{\text{track}} - \eta_{\text{clus}}| < 0.05$ and $-0.10 < q(\phi_{\text{track}} - \phi_{\text{clus}}) < 0.05$,[5] where q is the reconstructed charge of the track. If multiple tracks are matched to a cluster, the track with a greater number of hits in the Pixel detector and with a better ΔR match to the cluster is preferred.

The electron's transverse momentum is then computed from the cluster energy and the track direction at the interaction point, with corrections applied to account for energy lost in the material

[5] The requirement on $q(\phi_{\text{track}} - \phi_{\text{clus}})$ is asymmetric as tracks sometimes miss some energy from radiated photons that is measured in the cluster.

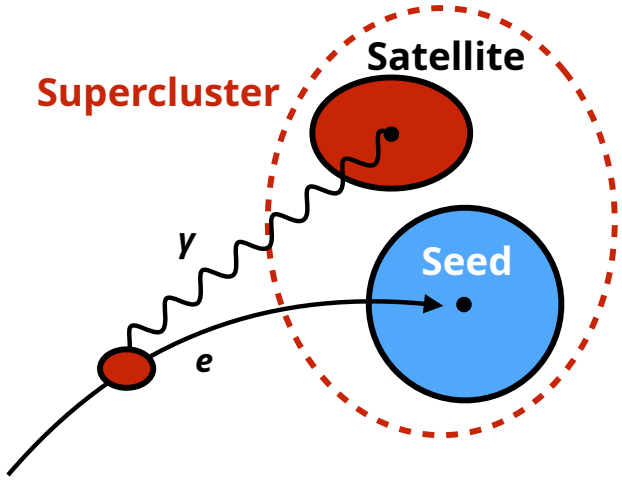

Figure 5.6: Illustration of an electromagnetic supercluster showing a seed electron cluster and a satellite photon cluster. Adapted from Ref. [190].

upstream of the calorimeter, energy deposited in cells neighbouring the supercluster and energy lost beyond the electromagnetic calorimeter. Additional corrections derived from comparisons between data and simulation in the reconstructed invariant mass distribution of $Z \to ee$ events are applied to the electron energy scale in both data and simulation and to the energy resolution in simulation only.

Electron-candidate background rejection relies on the longitudinal and transverse shapes of the electromagnetic showers in the calorimeters, track–cluster matching and properties of tracks in the Inner Detector. All of this information, except for that related to track hits, is combined into a likelihood discriminant for use in electron identification. First, the electron signal and background likelihoods are constructed from the products of the probability density functions (pdfs) for the N (typically 13) discriminating variables:

$$L_{S(B)}(\mathbf{x}) = \prod_{i=1}^{N} P_{S(B),i}(x_i)\,, \tag{5.2}$$

where $\mathbf{x}$ is the vector of input variables, $P_{S,i}(x_i)$ is the value of the signal pdf for variable i at value x_i and $P_{B,i}(x_i)$ is the analogous value of the background pdf. The signal consists of prompt electrons, while the background consists of a combination of jets that mimic the signature of prompt electrons, electrons originating from photon conversions and non-prompt electrons from the decays of heavy-flavour hadrons. Then, for each electron candidate, a discriminant is computed as the inverse sigmoid function

$$d_L' = -\frac{1}{\tau} \ln\left(\frac{L_S + L_B}{L_S} - 1\right)\,, \tag{5.3}$$

where by convention the parameter τ is fixed to 15. The discriminant is used in combination with the number of track hits to define multiple working points by selecting a value of d_L' above which the electron candidates are considered as signal with some predefined target identification efficiency. The "Loose", "Medium" and "Tight" working points, for instance, are defined with target efficiencies of 93 %, 88 % and 80 % respectively. Each of these working points is defined for electrons

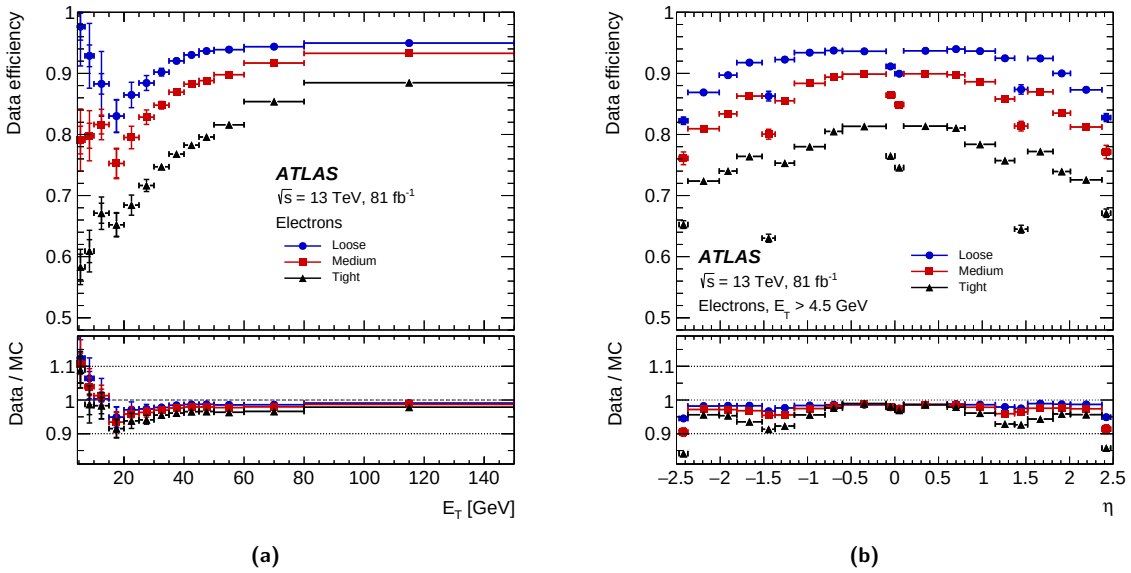

Figure 5.7: Electron identification efficiencies in $Z \to ee$ events in data as a function of (a) E_{T} and (b) pseudorapidity for the Loose, Medium and Tight working points [186]. The reduced efficiencies in the regions $|\eta| \approx 0.1$ and $1.37 < |\eta| < 1.52$ are the result of reduced electromagnetic calorimeter coverage in the crack region between the half barrels and between the barrel and the endcaps, respectively.

with $p_{\mathrm{T}} > 4.5\,\mathrm{GeV}$ and $|\eta| < 2.47$, and require at least two hits in the Pixel detector and seven hits total in the Pixel and SCT detectors combined. The Medium and Tight working points impose the additional requirement that one of these Pixel hits must be in the innermost Pixel layer, which helps to reduce the background from photon conversions. The electron identification efficiency for these three workings points is shown in Figure 5.7 as a function of E_{T} and η.

5.3.5 Muons

Muons are formed from tracks reconstructed in the Inner Detector (*ID tracks*) and Muon Spectrometer (*MS tracks*), and are primarily identified by the presence of the track or track segment in the Muon Spectrometer [191]. If a complete track is present in both the Inner Detector and the Muon Spectrometer, a combined muon track is formed by a global fit using the hit information from both the Inner Detector and Muon Spectrometer detectors (*combined muon*); otherwise the momentum is measured using the Inner Detector, and the MS track segment serves as identification of the muon candidate (*segment-tagged muon*). The segment-tagged muon is limited to the centre of the barrel region ($|\eta| < 0.1$) which has reduced Muon Spectrometer geometrical coverage. Furthermore, in this central region an ID track with $p_{\mathrm{T}} > 15\,\mathrm{GeV}$ is identified as a muon if its calorimetric energy deposition is consistent with a minimum-ionizing particle (*calorimeter-tagged muon*). In the forward region ($2.5 < |\eta| < 2.7$) with limited or no Inner Detector coverage, the MS track formed out of three Muon Spectrometer layers is either used alone (*stand-alone muon*) or combined with silicon-detector hits, if found in the forward region of Inner Detector (also classified as *combined muons*). The ID tracks associated with the muons are required to have at least a minimum number of associated hits in each of the Inner Detector subsystems to ensure good track reconstruction.

Muons tracks are reconstructed in the Muon Spectrometer in a manner similar to charged-particle

tracks in the Inner Detector, although with a few differences given the different detector technologies used by the two systems. First, a search algorithm identifies short straight-line local track segments reconstructed from hits in an individual MDT chamber and nearby trigger chambers, which provide measurements of the track's η and ϕ coordinates, respectively. Track segments in the CSC detectors are built using an analogous but separate combinatorial search with simultaneous measurements of both the η and ϕ coordinates. Segments from multiple stations are then combined intro preliminary track candidates if they are loosely consistent with originating from the luminous region. Finally, an iterative global χ^2 fit of the muon trajectory through the magnetic field is performed, taking into account the effects of possible interactions in the detector material as well as the effects of possible misalignments between the different detector chambers, where on each iteration outlier hits are removed and the track is re-fit. In the case of combined muons, the standalone MS tracks are extrapolated back to the interaction point and matched to ID tracks using a χ^2 fit.

The momentum scale and resolution of muon candidates are calibrated in a process similar to that of electrons using the reconstructed invariant mass distributions of $Z \to \mu\mu$ and $J/\psi \to \mu\mu$ events, with the notable exception that no systematic bias is observed in the invariant mass calibrations. Therefore, no muon momentum correction is applied to the data and the momentum scale and resolution corrections are only applied in simulation.

Identification of prompt muons (and rejection of non-prompt muons produced by pion or heavy-flavour quark decays) is performed by defining selection criteria on a number of variables that have been shown to effectively discriminate between prompt and non-prompt muons. The variables used for combined-muon identification are the *charge-to-momentum* (q/p) *compatibility* (or *significance*), the ρ' parameter and the normalized χ^2 of the combined track fit. The q/p compatibility is defined as

$$q/p \text{ compatibility} = \frac{|q/p_{\mathrm{ID}} - q/p_{\mathrm{MS}}|}{\sqrt{\sigma^2(q/p_{\mathrm{ID}}) + \sigma^2(q/p_{\mathrm{MS}})}}\,, \tag{5.4}$$

where q/p_{ID} and q/p_{MS} are the measurements in the Inner Detector and Muon Spectrometer of the ratio of the muon charge q to its momentum p and $\sigma(q/p_{\mathrm{ID}})$ and $\sigma(q/p_{\mathrm{MS}})$ are the corresponding uncertainties, and the ρ' parameter is defined as

$$\rho' = \frac{p_{\mathrm{T}}^{(\mathrm{ID})} - p_{\mathrm{T}}^{(\mathrm{MS})}}{p_{\mathrm{T}}^{(\mathrm{CB})}}\,, \tag{5.5}$$

where $p_{\mathrm{T}}^{(\mathrm{ID})}$ and $p_{\mathrm{T}}^{(\mathrm{MS})}$ are the muon p_{T} measured in the Inner Detector and in the Muon Spectrometer and $p_{\mathrm{T}}^{(\mathrm{CB})}$ is the value resulting from the combined track fit. No requirements on the q/p compatibility and ρ' variables are considered for muons without an ID or MS track, for which these variables are not defined.

Just as in electron identification, a number of muon-identification working points are defined based on the values of these variables for each muon candidate, as well as additional requirements on the number of hits in the precision layers of the Muon Spectrometer. These include the "Tight", "Medium" and "Loose" working points, whose reconstruction efficiencies are shown in Figure 5.8 as a function of p_{T} and η as measured in $J/\psi \to \mu\mu$ and $Z \to \mu\mu$ events, respectively. The Medium working point, for instance, only includes combined muons and has an efficiency and purity suitable for most physics analyses while keeping the systematic uncertainties in the prompt-muon efficiency and background rejection small. The Loose selection working point, which in addition to the muons

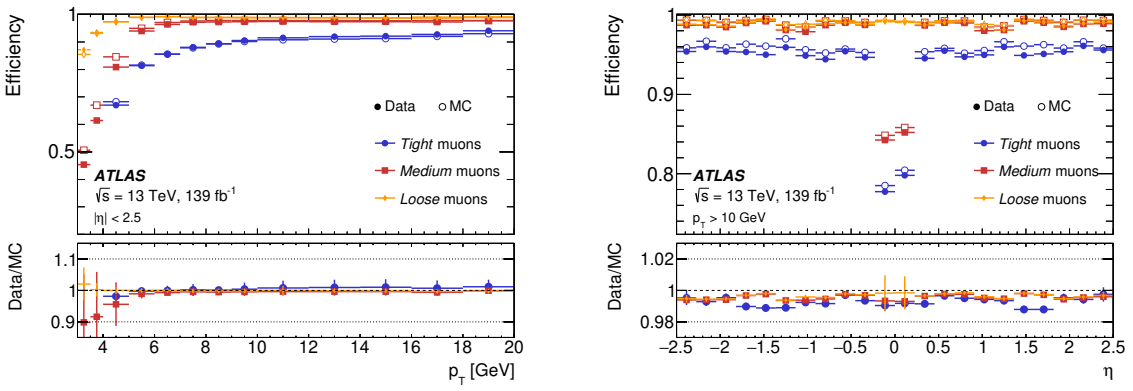

Figure 5.8: Muon reconstruction and identification efficiencies as function of (a) p_{T} in $J/\psi \to \mu\mu$ events and (b) pseudorapidity in $Z \to \mu\mu$ events for the Loose, Medium and Tight working points [191]. The reduced efficiency in the region $|\eta| < 0.1$ is the result of reduced Muon Spectrometer coverage in the crack region between the half barrels.

selected by the Medium working point also accepts segment-tagged and calorimeter-tagged muons in the range $|\eta| < 0.1$, is optimized for the reconstruction of Higgs boson decays in the four-muon final state, which benefits from a higher efficiency at the cost of less purity and larger systematic uncertainties due to the high muon multiplicity and large signal-to-background ratio for these events.

5.3.6 Jets

The reconstruction of jets uses a particle-flow algorithm [192] that combines measurements from both the Inner Detector and the calorimeters. The energy deposited in the calorimeter by all charged particles is removed, and the jet reconstruction is performed on an ensemble of "particle-flow objects" consisting of the remaining calorimeter energy and tracks that are matched to the hard interaction. This improves the accuracy of the charged-hadron measurement, while retaining the calorimeter measurements of neutral-particle energies. These particle-flow objects are then used as input to the anti-k_t jet-clustering algorithm [193] with a radius parameter $R = 0.4$. Compared to only using topological clusters [183], jets reconstructed with the particle-flow algorithm with p_{T} of about 30 GeV have approximately 10 % better transverse momentum resolution. The two different algorithms have similar resolutions for p_{T} above 100 GeV. The jet four-momentum is corrected for the calorimeter's non-compensating response, signal losses due to noise threshold effects, energy lost in non-instrumented regions, and contributions from pile-up [194]. The jets used in this work are required to satisfy $p_{\mathrm{T}} > 30$ GeV and $|\eta| < 4.5$. Jets from pile-up with $|\eta| < 2.5$ are suppressed using a jet-vertex-tagger multivariate discriminant [195, 196].

Jets containing b hadrons, referred to as *b-jets*, are identified by the long lifetime, high mass and decay multiplicity of b hadrons, as well as the hard b-quark fragmentation function. The $\ell^+\ell^-\nu\bar{\nu}$ analysis identifies b-jets of $p_{\mathrm{T}} > 20$ GeV and $|\eta| < 2.5$ using the MV2c10 b-tagging algorithm [197–199] with a working point that achieves an identification efficiency of about 85 % in simulated $t\bar{t}$ events, with a rejection factor for light-flavour jets of about 30 [200].

5.3.7 Missing transverse energy

The missing transverse energy $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$, which accounts for the imbalance of visible momenta in the plane transverse to the beam axis, is computed as the negative vector sum of the transverse momenta of all identified electrons, muons, photons, hadronically decaying τ leptons and jets, as well as a "soft term", accounting for unclassified soft tracks and energy clusters in the calorimeters [201]. The $\ell^+\ell^-\nu\bar{\nu}$ analysis uses a track-based soft term, which is built by combining the information provided by the Inner Detector and the calorimeter, in order to minimize the effect of pile-up, which degrades the $E_{\mathrm{T}}^{\mathrm{miss}}$ resolution. Missing transverse momentum is not considered in the $\ell^+\ell^-\ell^+\ell^-$ analysis as the final state for this analysis can be fully reconstructed.

5.3.8 Overlap removal

As lepton and jet candidates can be reconstructed from the same detector information, a procedure to resolve overlap ambiguities is applied. In the $\ell^+\ell^-\ell^+\ell^-$ case, the overlap ambiguities are resolved as follows. If two electrons have overlapping energy deposits, the electron with the higher p_{T} is retained. If a reconstructed electron and muon share the same ID track, the muon is rejected if it is calorimeter-tagged; otherwise the electron is rejected. Reconstructed jets geometrically overlapping in a cone of size $\Delta R = 0.2$ with electrons or muons are also removed. The overlap removal in the $\ell^+\ell^-\nu\bar{\nu}$ case is similar to that in the $\ell^+\ell^-\ell^+\ell^-$ case, except for an additional criterion that removes any leptons close to the remaining jets with $0.2 < \Delta R < 0.4$. This additional criterion is not imposed in the $\ell^+\ell^-\ell^+\ell^-$ case due to the cleaner environment of this final state and in order to maximize the signal efficiency.

5.4 Summary

The ATLAS Run 2 dataset offers a rich environment in which to measure the properties of the particles produced in the proton–proton collisions at the LHC and to search for new phenomena. The detector signals recorded in these collisions events capture highly detailed and granular "images" of the particles emitted in these interactions, which are then used to reconstructed the physics objects used in ATLAS physics analyses. The physics objects, including electrons, muons, jets, missing transverse energy and so on, are reconstructed using a variety of methods, each suited to the signals that these objects deposit in the detector. The objects are identified by imposing selections on a number of variables that efficiently discriminate between the prompt physics objects and backgrounds. Multiple working points are defined for object identification, each designed to reconstruct these objects with some predefined target reconstruction efficiency.

Collision events accumulate over time as the LHC and ATLAS are in operation, resulting in an integrated luminosity of $139\,\mathrm{fb}^{-1}$ for the Run 2 pp dataset at centre-of-mass energy $\sqrt{s} = 13\,\mathrm{TeV}$. The following chapter describes the ATLAS luminosity-determination methodologies used to arrive at this value, the main sources of systematic uncertainty in the luminosity determination and how they were evaluated in Run 2.

Chapter 6

Luminosity Measurements with LAr Energy Flow

6.1 Introduction

An accurate and precise measurement of the integrated luminosity is a key component of the ATLAS physics program, particularly for cross-section measurements, where the uncertainty in the integrated luminosity is often one of the leading systematic uncertainties. Searches for physical phenomena beyond the Standard Model also rely on accurate information about the integrated luminosity to evaluate background levels and determine sensitivity to the signatures of new phenomena.

The nominal ATLAS luminosity measurement is based on an absolute calibration of the primary luminosity-sensitive detectors in low-luminosity runs with specially tailored LHC conditions using the van der Meer (vdM) method [202, 203]. A "calibration transfer" procedure is then used to transport this calibration to the physics data-taking regime at high luminosity. The vdM calibration is performed once per data-taking year (for each combination of centre-of-mass energy and colliding-particle species), and relative comparisons of the luminosities measured by different detectors are used to set limits on any possible change of the calibration through the year.

The luminosity calibration relies on multiple redundant luminosity detectors and algorithms, which have complementary capabilities and different systematic uncertainties. For Run 2, the primary bunch-by-bunch luminosity measurement was provided by the LUCID Cherenkov detector (Section 3.3.6). This was complemented by bunch-by-bunch measurements from the ATLAS Beam Conditions Monitor (BCM) diamond detectors, and from offline measurements of the multiplicity of charged-particle tracks in randomly selected bunch crossings (track counting). Bunch-integrated measurements were also provided by three of the ATLAS calorimeter subsystems: the photomultiplier currents from the scintillating-tile hadronic calorimeter (TILE), and the LAr gap currents from the electromagnetic endcap (EMEC) and forward (FCal) calorimeters (Appendix A).

The nominal ATLAS luminosity measurement is documented extensively in Refs. [204–206] for Run 1, with the preliminary Run 2 calibration documented in Refs. [115, 207]. This chapter begins with an overview of the ATLAS luminosity-determination methodology and luminosity detectors, with an emphasis on the preliminary Run 2 results and leading systematic uncertainties, followed by a description of a new luminosity measurement using data collected by the liquid-argon calorimeters

in a number of special proton–proton runs at a centre-of-mass energy of $\sqrt{s} = 13\,\text{TeV}$. This new luminosity algorithm, called *LAr energy flow*, was developed at the end of Run 2 and is used to set limits on the relative nonlinearity of track counting in the calibration-transfer procedure.

6.2 ATLAS luminosity

Following the definition of luminosity from Section 3.2.2, the instantaneous luminosity of a hadron collider produced by a single pair of colliding bunches can be expressed in terms of an event rate as

$$\mathcal{L}_b = \frac{\mu f_r}{\sigma_{\text{inel}}}\,, \tag{6.1}$$

where μ is the average number of inelastic interactions per bunch crossing, also called the *pileup parameter*, f_r is the bunch revolution frequency (11 245.5 Hz for the LHC), and σ_{inel} is the inelastic cross section for the given colliding-particle species and centre-of-mass energy, taken by convention to be 80 mb for pp collisions at $\sqrt{s} = 13\,\text{TeV}$.[1] The total instantaneous luminosity is the sum over all colliding-bunch pairs:

$$\mathcal{L} = \sum_{b=1}^{n_b} \mathcal{L}_b = n_b \langle \mathcal{L}_b \rangle = n_b \frac{\langle \mu \rangle f_r}{\sigma_{\text{inel}}}\,, \tag{6.2}$$

where n_b is the number of colliding-bunch pairs, $\langle \mathcal{L}_b \rangle$ is the mean bunch luminosity and $\langle \mu \rangle$ is the bunch-averaged pileup parameter. At the LHC, each colliding-bunch pair is identified numerically by a bunch-crossing identifier (BCID) which labels each of the 3564 possible 25 ns slots in one full revolution of the nominal LHC fill pattern.

ATLAS monitors the delivered luminosity by measuring μ_{vis}, the visible interaction rate per bunch crossing. The bunch luminosity can then be written as

$$\mathcal{L}_b = \frac{\mu_{\text{vis}} f_r}{\sigma_{\text{vis}}}\,, \tag{6.3}$$

where $\mu_{\text{vis}} = \varepsilon\mu$, ε is the efficiency of the detector and algorithm under consideration, and the visible cross section for that same detector and algorithm is defined by $\sigma_{\text{vis}} \equiv \varepsilon\sigma_{\text{inel}}$. Since μ_{vis} is a directly measurable quantity, the absolute luminosity calibration for a particular detector and algorithm amounts to determining its visible cross section σ_{vis}.

6.2.1 Luminosity detectors and algorithms

The ATLAS luminosity detectors and algorithms used in Run 2 are described in detail in Ref. [115]. The LAr energy-flow measurements rely on various LUCID and track-counting algorithms to determine their absolute luminosity calibration, therefore a brief description of each is given below.

The LUCID detector provides luminosity counts for each of the 3564 LHC bunch slots, integrated over luminosity blocks with a typical length of 60 s. Several luminosity algorithms were used to convert the raw signals from the LUCID PMTs to a luminosity measurement, combining the information from several PMTs in various ways. The simplest algorithm uses a single PMT, and counts an *event* if there is a signal in the PMT above a given threshold (a *hit*), corresponding to one

[1] The choice of σ_{inel} value affects the relationship between instantaneous luminosity and μ, but has no effect on the calibrated luminosity values.

or more inelastic pp interactions detected in a given bunch crossing. Assuming that the number of inelastic pp interactions in a bunch crossing follows a Poisson distribution, the probability for such an event P_{evt} in terms of the single-PMT visible interaction rate is given by

$$P_{\text{evt}} = \frac{N_{\text{evt}}}{N_{\text{BC}}} = 1 - e^{-\mu_{\text{vis}}^{(\text{evt})}} , \tag{6.4}$$

where N_{evt} is the (background-subtracted) number of events counted in the luminosity block and N_{BC} is the number of bunch crossings sampled (equal to $f_r \Delta t$ for a single colliding bunch pair, where Δt is the duration of the luminosity block). The μ_{vis} value is then given by

$$\mu_{\text{vis}}^{(\text{evt})} = -\ln\left(1 - \frac{N_{\text{evt}}}{N_{\text{BC}}}\right) , \tag{6.5}$$

from which the instantaneous luminosity can be calculated via Eq. (6.3) once σ_{vis} is known. Several PMTs can be combined in an "EventOR" (or EvtOR) algorithm by counting an event if any of a group of PMTs registers a hit in a given bunch crossing. When $\mu_{\text{vis}} \gg 1$, or equivalently in the limit where $N_{\text{evt}}/N_{\text{BC}} = 1$, event-counting algorithms lose sensitivity as fewer and fewer bunch crossings in a given time interval report zero observed interactions; this is known as *saturation* or *zero starvation*. "HitOR" algorithms provide an alternative method for combining PMTs at high luminosity. For N_{PMT} PMTs, the average probability P_{hit} to have a hit in any given PMT during the N_{BC} bunch crossings of one luminosity block is given from the total number of hits summed over all PMTs N_{hit} by

$$P_{\text{hit}} = \frac{N_{\text{hit}}}{N_{\text{BC}} N_{\text{PMT}}} = 1 - e^{-\mu_{\text{vis}}^{(\text{hit})}} , \tag{6.6}$$

which, for the HitOR algorithm, leads to

$$\mu_{\text{vis}}^{(\text{hit})} = -\ln\left(1 - \frac{N_{\text{hit}}}{N_{\text{BC}} N_{\text{PMT}}}\right) . \tag{6.7}$$

In 2015, the baseline luminosity estimates were derived from LUCID BiEventORA, an algorithm requiring a hit in any of four bismuth-calibrated PMTs on the A-side of the detector.[2] In 2016 and 2017, the BiHitOR algorithm was used, counting the average number of hits per bunch crossing summed over four bismuth-calibrated PMTs on the A-side and four more on the C-side. In 2018, a significant number of PMTs stopped working during the course of the data-taking year, and a single PMT on the C-side (C12) was used for the baseline luminosity estimate, as it showed good stability throughout the year and gave similar results to a more complicated HitOR-type combination of the remaining seven working PMTs.

The track-counting method takes the per-bunch visible interaction rate μ_{vis} from the mean number of reconstructed tracks per bunch crossing averaged over a luminosity block. The measurement is derived from randomly sampled colliding-bunch crossings, where only the data from the SCT and Pixel detectors (including the IBL) were read out, typically at 200 Hz during normal physics data-taking and at much higher rates during vdM scans. These events are saved in a dedicated event stream which is then reconstructed offline. The track selection evolved over the course of Run 2, in order to improve the robustness against time-dependent detector inefficiencies, and to reduce the

[2] Analogous algorithms were available for the C-side and the combination of both sides, but the C-side suffered from significant timing drifts as the PMT high-voltage settings were adjusted.

Table 6.1: Summary of the track-selection working points for the track-counting luminosity measurement by the year of data-taking in which the working point was first used. In all cases, the tracks were required to have transverse momentum $p_{\mathrm{T}} > 0.9\,\mathrm{GeV}$ and to satisfy the "TightPrimary" selection of Ref. [208]. Here, a "pixel hole" refers to a track with a pixel hit missing where one was expected. The $|d_0/\sigma_{d_0}|$ impact-parameter requirements, where d_0 is the transverse impact parameter of the reconstructed track and σ_{d_0} is its associated uncertainty, reject tracks that are inconsistent with originating directly from a pp collision.

Year	2015	2016	2017
$\lvert\eta\rvert$	< 2.5	< 2.5	< 1.0
Allowed pixel holes	$= 0$	$= 0$	≤ 1
Impact parameter	N/A	$\lvert d_0/\sigma_{d_0}\rvert < 7$	$\lvert d_0/\sigma_{d_0}\rvert < 7$

efficiency loss and fake-track contamination at high pileup. These track-selection working points are summarized in Table 6.1. The "baseline" working point is the one developed is 2017 (referred to as the "2017 selection" working point), as it was found to be more reliable in long-term-stability studies [115, 207]. The raw track-counting data from previous years have since been reprocessed using the 2017 selections so that the nominal track-counting luminosity measurements use this working point for all of Run 2.

6.2.2 Absolute luminosity calibration

The absolute luminosity calibration of LUCID, corresponding to the determination of the visible cross section σ_{vis} for each of the LUCID algorithms, is derived using dedicated vdM scan sessions during special LHC fills in each data-taking year. The vdM formalism is described below, followed by a brief description of the data samples and methods used for the calibration procedure.

In terms of colliding-beam parameters, the bunch luminosity is defined (for beams colliding with zero crossing angle) as

$$\mathcal{L}_b = f_r n_1 n_2 \int \hat{\rho}_1(x,y)\hat{\rho}_2(x,y)\,dx\,dy\,, \tag{6.8}$$

where f_r is the bunch revolution frequency, n_1 and n_2 are the bunch populations of beam 1 and 2, and $\hat{\rho}_{1(2)}(x,y)$ is the normalized particle-density distribution in the transverse (x–y) plane of beam 1 (2) at the interaction point. Under the general assumption that the particle densities can be factorized into independent horizontal and vertical components, $\hat{\rho}(x,y) = \rho_x(x)\rho_y(y)$, Eq. (6.8) can be rewritten as

$$\mathcal{L}_b = f_r n_1 n_2 \Omega_x(\rho_{x1},\rho_{x2})\Omega_y(\rho_{y1},\rho_{y2})\,, \tag{6.9}$$

where

$$\Omega_x(\rho_{x1},\rho_{x2}) = \int \rho_{x1}(x)\rho_{x2}(x)\,dx$$

is the beam-overlap integral in the x-direction (with an analogous definition in the y-direction). In the method proposed by van der Meer [202] the overlap integral (for example in the x-direction) can be calculated as

$$\Omega_x(\rho_{x1},\rho_{x2}) = \frac{R_x(0)}{\int R_x(\Delta x)\,d\Delta x}\,, \tag{6.10}$$

where $R_x(\Delta x)$ is the luminosity (or equivalently μ_{vis})—at this stage in arbitrary units—measured during a horizontal scan at the time the two beams are separated by a distance Δx, and $\Delta x = 0$ indicates the case of zero beam separation.

Defining the *convolved beam size* parameter Σ_x as

$$\Sigma_x = \frac{1}{\sqrt{2\pi}} \frac{\int R_x(\Delta x)\, d\Delta x}{R_x(0)}\,, \tag{6.11}$$

and similarly for Σ_y, the luminosity in Eq. (6.9) can be rewritten as

$$\mathcal{L}_b = \frac{f_r n_1 n_2}{2\pi \Sigma_x \Sigma_y}\,, \tag{6.12}$$

which enables the luminosity to be extracted from machine parameters by performing a vdM beam-separation scan. In the case where the luminosity curve $R_x(\Delta x)$ is Gaussian, Σ_x coincides with the standard deviation of that distribution. However, Σ_x and Σ_y, as defined in Eq. (6.11), depend only upon the area under the luminosity curve and make no assumption as to the shape of that curve, therefore Eq. (6.12) holds in general for arbitrary distributions $R_x(\Delta x)$ and $R_y(\Delta y)$. The calibration of a given algorithm (i.e. its σ_{vis} value) can then be determined by combining Eqs. (6.12) and (6.3) to give

$$\sigma_{\text{vis}} = \mu_{\text{vis}}^{\text{max}} \frac{2\pi \Sigma_x \Sigma_y}{n_1 n_2}\,, \tag{6.13}$$

where $\mu_{\text{vis}}^{\text{max}}$ is the visible interaction rate per bunch crossing at the peak of the scan curve.

The above formulation makes a number of simplifying assumptions. The assumption of factorizable particle transverse-density distributions as in Eq. (6.9), for instance, does not necessarily hold true at the level of precision required for the vdM calibration, and required a correction to the value of σ_{vis} on the order of 0.2–0.6 % in Run 2, derived from on- and off-axis beam-separation scans to determine the shape and orientation of the luminous region from the distribution of reconstructed primary collision vertices. Other corrections associated to so-called *beam–beam* effects due to the electromagnetic interaction between the two beams during the scans, which distort the overlap integral of the colliding bunches, are also required. The details of these corrections and others, and their associated systematic uncertainties, are discussed at length in Refs. [115, 206].

The vdM scans were performed in dedicated low-luminosity fills in each data-taking year with special LHC optics settings: a larger β^* parameter of 19.2 m rather than the 0.25–0.8 m used during normal physics running and larger beam emittances of 3–4 µm rad. These parameters were chosen to increase the transverse beam sizes to about 90 µm, thus reducing the uncertainties in evaluating the non-factorization corrections due to the improved vertex resolution in these conditions. Filling schemes with 30–140 isolated bunches were used to avoid the effects of long-range encounters between incoming and outgoing bunches away from the interaction point that occur in normal bunch-train running. This allowed the beam crossing angle to be set to zero, reducing uncertainties relating to orbit-drift and beam–beam effects. Furthermore, reduced bunch populations of $\sim 0.8 \times 10^{11}$ protons/bunch were employed to minimize beam–beam effects. These configurations typically resulted in $\langle\mu\rangle \approx 0.5$ at the peak of the scan curves.

A typical scan curve in the horizontal plane from 2017 using the LUCID BiHitOR algorithm is shown in Figure 6.1. The background-subtracted scan profile was fitted using an analytic function (typically a Gaussian multiplied by a fourth-order polynomial, or "GP4", for the Run 2 *pp* scans), from which the peak interaction rate $\mu_{\text{vis}}^{\text{max}}$ and convolved beam size Σ_x are determined. An analogous scan curve in the vertical plane is used to determine Σ_y. Following Eq. (6.12), a single pair of x–y

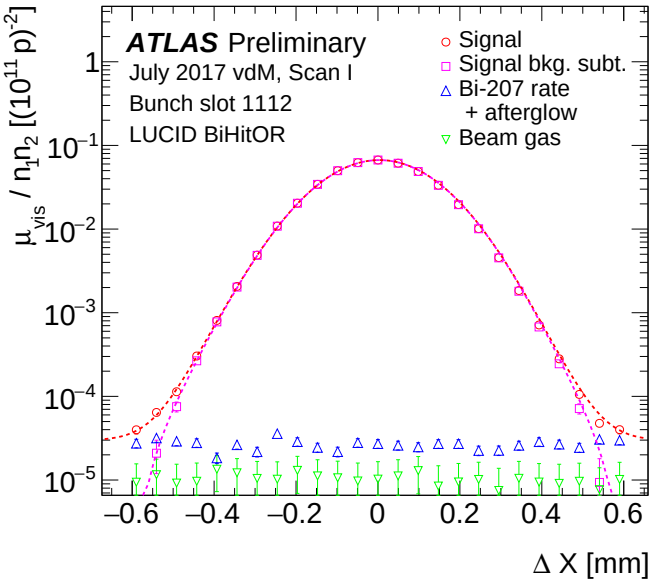

Figure 6.1: Example of an absolute luminosity calibration scan curve [115]. Shown is the visible interaction rate μ_{vis} per unit bunch population product $n_1 n_2$ vs. beam separation Δx in the horizontal plane, as measured by the LUCID BiHitOR algorithm for bunch slot 1112 in scan I of the July 2017 vdM session. The signal and background components of μ_{vis} are shown separately by different markers and the fit functions by the dashed lines.

vdM scans thus suffices to measure σ_{vis} for each luminosity algorithm active during the scan. In practice, the value of σ_{vis} is measured for each colliding bunch pair and in a number of separate vdM scan sessions, with the relative spread of these values used to evaluate the per-year "bunch-by-bunch consistency" and "scan-to-scan reproducibility" uncertainties, respectively.

6.2.3 Calibration transfer to physics data-taking conditions

The procedures discussed above provide the absolute calibration of the LUCID luminosity measurements for low-μ data-taking with a limited number of isolated bunches. However, the LUCID detector suffers from significant nonlinearity, and requires a correction of $O(10\,\%)$ in the physics data-taking regime. This effect can clearly be seen by studying the ratios between instantaneous luminosities measured with LUCID and with other detectors, and how they change as a function of instantaneous luminosity or μ.

The effects of this nonlinearity were corrected by comparing LUCID to track-counting luminosity measurements, which are known from comparisons with EMEC and TILE to have much smaller nonlinearity as a function of μ. A single long high-luminosity physics fill, near in time to the vdM calibration period in each year of data-taking, was used to derive a correction of the form

$$\mu_{\text{corr}} = p_0 \mu_{\text{uncorr}} + p_1 \mu_{\text{uncorr}}^2 \,, \tag{6.14}$$

where μ_{uncorr} is the uncorrected and μ_{corr} the corrected LUCID μ value, and the parameters p_0 and p_1 were obtained from a linear fit to the ratio of μ values measured by track counting and LUCID, $R = \mu_{\text{track}}/\mu_{\text{uncorr}}$, as a function of μ_{uncorr}. The track-counting luminosity was first normalized to

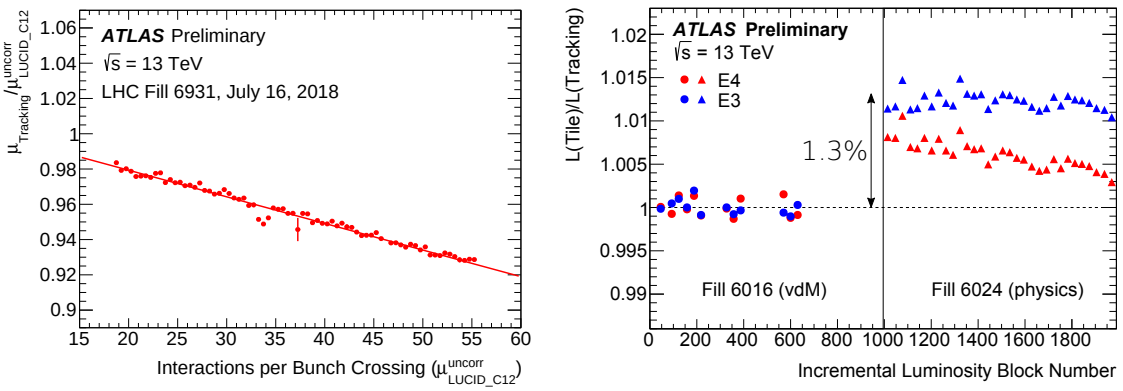

Figure 6.2: The LUCID calibration-transfer procedure [115]. (a) shows the luminosity ratio of track counting to LUCID C12 in LHC physics fill 6931 and the linear fit to the data points. (b) shows the ratios of the instantaneous luminosity measured by TILE E-cell scintillators to that from track counting, in a vdM fill (LHC fill 6016) and a closely following physics fill (LHC fill 6024). The relative difference between these two running conditions is 1.3 %, which is taken as a systematic uncertainty on the calibration transfer.

the absolute luminosity measured by LUCID in parts of the vdM fill with stable, almost-constant luminosity, where the beams were colliding head-on, thus ensuring that $R = 1$ at low luminosity with isolated bunches. The 2018 correction to the C12 algorithm is shown in Figure 6.2a. This calibration-transfer procedure implicitly assumes that track counting suffers from no significant nonlinearity between the low-luminosity vdM and high-luminosity physics regimes. This assumption was verified by comparing the ratio of the track-counting luminosity to that from the TILE E cells, the scintillators installed in the gaps between the barrel and endcap calorimeter assemblies (as shown in Figure 3.22), in the vdM fill and a closely following physics fill. An example of this procedure from 2017 is illustrated in Figure 6.2b. In this case, a relative difference of 1.3 % was observed, which was interpreted as an upper limit on the nonlinearity of the track-counting luminosity measurement and thus assigned as a systematic uncertainty on the correction applied to LUCID at high luminosity.

6.2.4 Long-term stability

The vdM calibration, followed by the calibration transfer to the physics regime, determines the absolute LUCID luminosity scale at one point during the data-taking year. To quantify any possible drifts over the year, the LUCID integrated luminosity estimates were compared for each physics run with those from other subdetectors, after renormalizing the other luminosity estimates to agree with LUCID in a long "reference" run close to the vdM scan. A conservative "stability band" uncertainty is then assigned such that it encloses the bulk of the differences between LUCID and any of the other luminosity measurements. This stability uncertainty is illustrated by the yellow bands in Figure 6.3 for the $\sqrt{s} = 13\,\text{TeV}$ pp datasets in 2017 and 2018. The stability uncertainties in each year of data-taking were 0.7 % (2015+16), 1.3 % (2017)[3] and 0.8 % (2018).

[3]FCal was the uncertainty-limiting luminosity algorithm in 2017. A correction to the FCal luminosity to account for a nonlinearity in its response is described in Appendix A. This correction was applied in 2018 and results in improved stability with respect to the other luminosity detectors and algorithms.

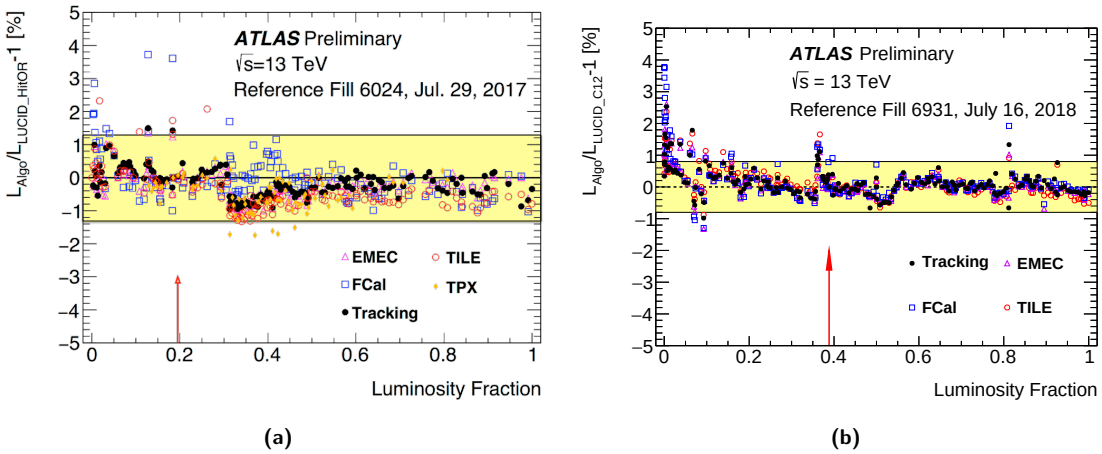

Figure 6.3: Long-term stability: fractional differences in run-integrated luminosity between the LUCID and the track-counting, TILE, EMEC and FCal measurements, plotted as a function of the cumulative delivered luminosity in (a) 2017 and (b) 2018, normalized to the per-year total [115]. The luminosity measurements from the other detectors have been normalized to that of LUCID in the reference run indicated by the arrows. The assigned long-term stability uncertainties are shown by the yellow bands.

6.2.5 Uncertainties, inter-year correlations and results

ATLAS physics analyses typically treat datasets with common centre-of-mass energy and running conditions from different years as a single combined dataset and require a single luminosity uncertainty for this dataset. If the total integrated luminosity $\mathcal{L}_{\text{tot}}$ is given by $\mathcal{L}_{\text{tot}} = \sum_i \mathcal{L}_i$, where $\mathcal{L}_i$ is the integrated luminosity of dataset i, then its absolute uncertainty $\sigma_{\mathcal{L}_{\text{tot}}}$ is given by

$$\sigma^2_{\mathcal{L}_{\text{tot}}} = \mathbf{G}\mathbf{V}_\mathbf{L}\tilde{\mathbf{G}}\,, \tag{6.15}$$

where $\mathbf{V_L}$ is the covariance matrix encoding the absolute luminosity uncertainties for the different years and their correlations and $\mathbf{G}$ is a vector of ones.[4]

Using this formulation, a preliminary luminosity uncertainty of $\delta\mathcal{L}/\mathcal{L} = 1.7\,\%$ was obtained for the Run 2 $\sqrt{s} = 13\,\text{TeV}$ pp dataset [115]. Table 6.2 gives the individual sources of uncertainty in each year of data-taking and their combination using the above formulation.

The calibration transfer resulted in the single largest source of systematic uncertainty in the Run 2 luminosity calibration, with the uncertainty associated to this procedure relying on independent luminosity algorithms to set limits on the relative nonlinearity of track counting. For the preliminary Run 2 calibration, the TILE E cells were employed for this purpose, owing to their large dynamic range in sensitivity, extending from low-luminosity vdM conditions to high-luminosity physics conditions. However, the TILE E cells suffer from material-activation effects and rapid scintillator ageing, degrading their performance as a luminosity-sensitive detector. Two complementary solutions to this issue have been proposed: first, models of the activation and scintillator-ageing effects have been developed to parameterize and correct for this degraded performance, and second, the LAr energy-flow luminosity algorithm, derived from measurements of the energy deposited in the liquid-argon calorimeters, has been developed to provide a second, independent luminosity algorithm to set limits on the relative nonlinearity of track counting. The following section describes the LAr energy-flow

[4] In general, $\mathbf{G} = (d\mathcal{L}_{\text{tot}}/d\mathcal{L}_i,\, i = 1, \ldots, n)$, but as $\mathcal{L}_{\text{tot}} = \sum_i \mathcal{L}_i$, $\mathbf{G}$ reduces to a vector where each element is 1.

Table 6.2: Summary of the integrated luminosities and uncertainties for the preliminary calibration of each individual year of the Run 2 $\sqrt{s} = 13\,\text{TeV}$ pp data sample and the full combined sample [115]. Contributions marked * are considered fully correlated between years, those marked † are considered partially correlated, and the other uncertainties are considered uncorrelated.

Data sample	2015+16	2017	2018	Comb.
Integrated luminosity (fb^{-1})	36.2	44.3	58.5	139.0
Total uncertainty (fb^{-1})	0.8	1.0	1.2	2.4
Uncertainty contributions (%):				
Absolute vdM calibration†	1.1	1.5	1.2	–
Calibration transfer†	1.6	1.3	1.3	1.3
Afterglow and beam-halo subtraction*	0.1	0.1	0.1	0.1
Long-term stability	0.7	1.3	0.8	0.6
Tracking efficiency time-dependence	0.6	0.0	0.0	0.2
Total uncertainty (%)	2.1	2.4	2.0	1.7

luminosity measurements and its performance in a select number of special runs in the context of the calibration-transfer procedure.

6.3 LAr energy-flow luminosity

The average amount of energy deposited in the LAr calorimeters per bunch crossing is proportional to the instantaneous luminosity. In bunch trains with 25 ns spacing, any signal that results from this effect is washed out by the bipolar shaping (with a drift time of $\sim 450\,\text{ns}$ in the EMB and HEC, 200–600 ns in the EMEC and 60–120 ns in the FCal) and the electronics response time. However, for sufficiently isolated bunches, this "energy-flow" signal from the LAr calorimeters is sensitive enough to perform a luminosity measurement with a precision and linearity comparable to other luminosity detectors and algorithms.

The energy flow is measured as the energy deposited in each calorimeter cell per bunch crossing, E_{cell}, summed over all cells in a LAr detector subsystem, and averaged over all events per bunch crossing in a luminosity block:

$$E_{\text{flow}} \equiv \left\langle \sum_{\text{cells}} E_{\text{cell}} \right\rangle_{\text{LB}} . \tag{6.16}$$

To use this quantity as a measurement of the luminosity, it must be corrected for the electronics pedestals E_{ped}, discussed in detail below in Section 6.3.2, and calibrated to convert the measured energies to units of μ:

$$\mu = \zeta \left(E_{\text{flow}} - E_{\text{ped}} \right) , \tag{6.17}$$

where ζ is the per-bunch calibration constant for the detector subsystem. Typically ζ is computed by anchoring the energy-flow measurement to a reference luminosity algorithm μ_{ref}, such as track counting, in a single, high-μ luminosity block, such as the luminosity block at peak μ within a period of interest:

$$\zeta = \frac{\mu_{\text{ref}}^{(\text{peak})}}{E_{\text{flow}}^{(\text{peak})} - E_{\text{ped}}} . \tag{6.18}$$

It is also convenient to define the pedestal in units of μ using this calibration factor: $q \equiv \zeta E_{\text{ped}}$.

Table 6.3: Summary of the ATLAS runs with recorded LArPEB stream data in 2017 and 2018 at $\sqrt{s} = 13\,\text{TeV}$ used in the LAr energy-flow luminosity measurement. The "Bunch config." column indicates the number of colliding bunches, with the bunch structure in parentheses. The notation "iso." indicates isolated colliding bunches (which are both preceded and followed by a sufficiently large number of empty bunch slots) and "8b4e" signifies a repeating bunch train pattern of eight bunches spaced by 25 ns followed by a four empty bunch-slot gap. $\mathcal{L}_{\text{peak}}$ indicates the peak luminosity in the given run. The "Isolated & recorded BCIDs" column indicates the BCIDs that are both isolated and recorded to the LArPEB stream.

Run number	LHC fill	Date	Bunch config.	$\mathcal{L}_{\text{peak}}$ ($10^{33}\,\text{cm}^{-2}\,\text{s}^{-1}$)	Isolated & recorded BCIDs
336506	6239	Sep 24, 2017	1346 (2 iso. + 8b4e)	9.03	1738, 1944
339197	6336	Oct 26, 2017	140 (iso.)	0.928	all colliding
354124	6847	Jun 26, 2018	140 (iso.)	0.930	1409, 2770
354309	6854	Jun 27, 2018	1214 (2 iso. + 25 ns)	8.26	823, 2608

6.3.1 Data samples

A special data stream was configured for a select number of ATLAS runs in 2017 and 2018 at $\sqrt{s} = 13\,\text{TeV}$ in which only the LAr calorimeter cells were read out. This data stream, called the LAr partial-event-building (LArPEB) stream, recorded raw energy measurements for each LAr calorimeter cell in events selected by random triggers in both colliding and empty BCIDs. Of these runs, four were of particular interest for the LAr energy-flow luminosity measurement, and are summarized in Table 6.3. These are runs 336506, 339197, 354124, and 354309. The online luminosity profiles for these four runs are shown in Figure 6.4, and detailed chronologies are given in Table 6.4. Each of these runs contain a period in which the beams were separated in steps of progressively increasing and decreasing μ, known as μ *scans*. These μ scans provide data over a large range in μ, over a relatively short time frame to ensure similar detector conditions, and are useful for studying the μ-dependence of a luminosity detector or algorithm.

A μ scan progresses in alternating steps of so-called "quiescent"[5] and "non-quiescent" luminosity blocks. A quiescent luminosity block is one in which the beams are held in a steady state and data can be recorded at constant luminosity. A non-quiescent luminosity block is one in which the beams are moving to the next scan step. Since the luminosity is not constant during non-quiescent luminosity blocks, the data recording during these times are typically rejected. Before 2018, the transitions between scan steps were recorded independently of the standard ATLAS data-acquisition system for times in which the LHC scan protocol was activated. The timestamps of these transitions mark the boundaries of the quiescent and non-quiescent scan steps, forming so-called "pseudo luminosity blocks" (PLBs). In 2018, the ATLAS data-acquisition system was updated to synchronize the luminosity block counter with the LHC scan protocol so that regular ATLAS luminosity blocks could be used to demarcate the scan-step boundaries.

6.3.2 Pedestals

Electronics pedestals must be estimated in each of the LAr detector subsystems and subtracted to ensure an accurate luminosity measurement, especially at low μ. These pedestals can drift over time and must therefore be estimated in each run. While the "nominal" LAr pedestals are evaluated regularly in special calibration runs during LHC downtime between most data-taking runs, it is possible to achieve greater precision by manually re-evaluating the pedestals in each run during

[5] **quiescent** /kwaɪˈɛsn̩t/ (*adj.*) marked by inactivity or repose; tranquilly at rest. From the Latin *quiescens:* resting, reposing, sleeping.

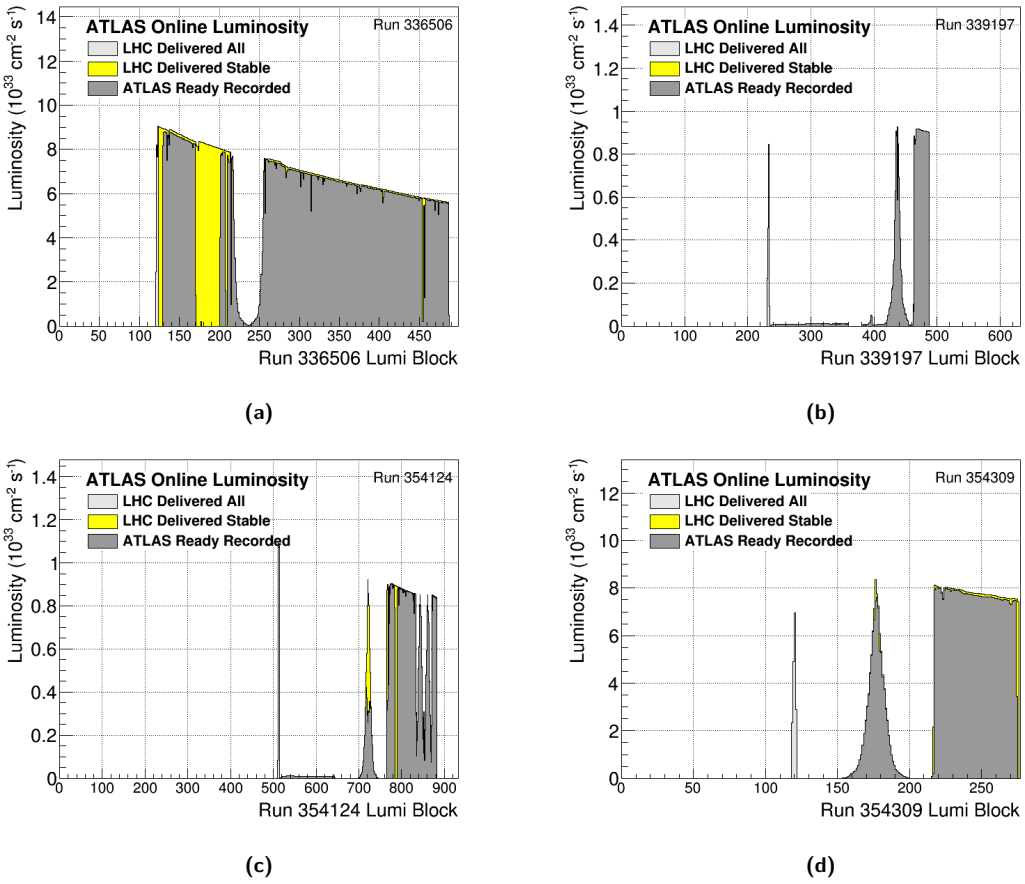

Figure 6.4: Online luminosity profiles for ATLAS runs (a) 336506, (b) 339197, (c) 354124, and (d) 354309.

a period without collisions, effectively anchoring the measurement at zero where the pedestal was evaluated. The pedestal estimate in the context of the energy-flow luminosity measurement can be thought of as a "residual" pedestal with respect to the nominal pedestal, and can therefore have either a positive or negative value. There is also clear evidence of pedestal BCID dependence, as illustrated in Figure 6.5, indicating that it is also necessary to estimate the pedestals separately for each of the colliding BCIDs in a run.

Ideally these pedestals would be evaluated during periods without collisions, however this was not always possible given the running conditions when the LArPEB stream data were recorded. In runs 339197 and 354124, this was not an issue thanks to the long periods of beam separation before (and after, in the case of run 354124) the μ scan. In these cases, the pedestals were measured in the same way as the energy flow (as in Eq. (6.16)), except the cell-summed energies are averaged over multiple luminosity blocks during these periods in which there were no collisions:

$$E_{\text{ped}} = \left\langle \sum_{\text{cells}} E_{\text{cell}} \right\rangle_{\text{no collisions}} . \tag{6.19}$$

However, in run 336506, there was no such period of beam separation, and in run 354309, the beams were only separated in the horizontal plane, resulting in a non-zero luminosity during this

Table 6.4: Detailed chronologies of runs 336506, 339197, 354124 and 354309. The LHC was in stable beams for all times listed below. LB ranges are inclusive and timestamps are given in local time (CEST) starting on the date indicated in parentheses beneath the run number.

Run number	LB range	Timespan	Conditions
336506 (Sep 24, 2017)	124–213	04:56–06:20	Head-on collisions, $\langle\mu\rangle \approx 47$–$42$
	214–257	06:20–07:04	"Down-up" μ scan by horizontal beam separation, $\langle\mu\rangle \approx 42 \to 0.16 \to 40$
	258–486	07:04–10:50	Head-on collisions, $\langle\mu\rangle \approx 40$–$30$
339197 (Oct 26, 2017)	240–359	22:24–00:21	Levelled, $\langle\mu\rangle \approx 0.5$
	360–381	00:21–00:43	Beams separated in x and y, $\langle\mu\rangle \approx 0$
	382–416	00:43–01:18	Lumiscan problems; ignore data
	417–457	01:18–01:57	"Up-down" μ scan by horizontal beam separation, $\langle\mu\rangle \approx 0.5 \to 45 \to 0.5$
	458–462	01:57–02:02	Moving to head-on collisions
	463–486	02:02–02:25	Head-on collisions, $\langle\mu\rangle \approx 46$–$45$
354124 (Jun 26, 2018)	515–642	08:55–10:59	Levelled, $\langle\mu\rangle \approx 0.5$
	648–695	11:04–11:42	Beams separated in x and y, $\langle\mu\rangle \approx 0$
	696–747	11:42–12:22	"Up-down" μ scan by horizontal beam separation, $\langle\mu\rangle \approx 0 \to 21 \to 0$
	748–763	12:22–12:34	Beams separated in x and y, $\langle\mu\rangle \approx 0$
	764–765	12:34–12:36	Moving to head-on collisions
	766–833	12:36–13:36	Head-on collisions, $\langle\mu\rangle \approx 45$–$43$
	834–870	13:36–13:41	Emittance scan in ATLAS
	871–882	13:41–13:53	Head-on collisions, $\langle\mu\rangle \approx 43$
354309 (Jun 27, 2018)	125–150	22:57–23:20	Beams separated in x only, $\langle\mu\rangle \approx 0.02$
	151–200	23:20–23:57	"Up-down" μ scan by horizontal beam separation, $\langle\mu\rangle \approx 0.02 \to 49 \to 0.02$
	201–214	23:57–00:07	Beams separated in x only, $\langle\mu\rangle \approx 0.02$
	215–216	00:07–00:09	Going to head-on collisions
	217–274	00:09–01:04	Head-on collisions, $\langle\mu\rangle \approx 47$–$44$

time from collisions at the peripheries of the bunch pairs. Additionally, the pedestals could not be evaluated during the period without collisions immediately before the LHC declared stable beams or immediately following the beam dump, as the LArPEB stream data are not recorded during these times, nor could they be evaluated using empty bunch-slot data, due to the pedestal BCID dependence. Each run therefore required specially tailored methodologies to evaluate the pedestals, which are described below.

Pedestals for run 339197

In run 339197, the beams were fully separated in both the horizontal and vertical planes for approximately 22 min before the μ scan. The pedestals were measured following Eq. (6.19) during this period of full beam separation per BCID and are shown in Figure 6.5. A "flat" pedestal subtraction was performed in this run, meaning this single pedestal estimate was subtracted uniformly in each luminosity block.

Pedestals for run 354124

In run 354124, the beams were fully separated in both the horizontal and vertical planes for approximately 38 min before the μ scan and for approximately 12 min immediately after the μ scan. The pedestals were measured following Eq. (6.19) separately during each of these periods of full beam separation, corresponding to independent "pre-scan" and "post-scan" pedestal estimates, and are given in Table 6.5.

The pre- and post-scan pedestal estimates were not statistically consistent and differed by several percent. This indicates a drift in the pedestal value over a period of about an hour. In the absence of

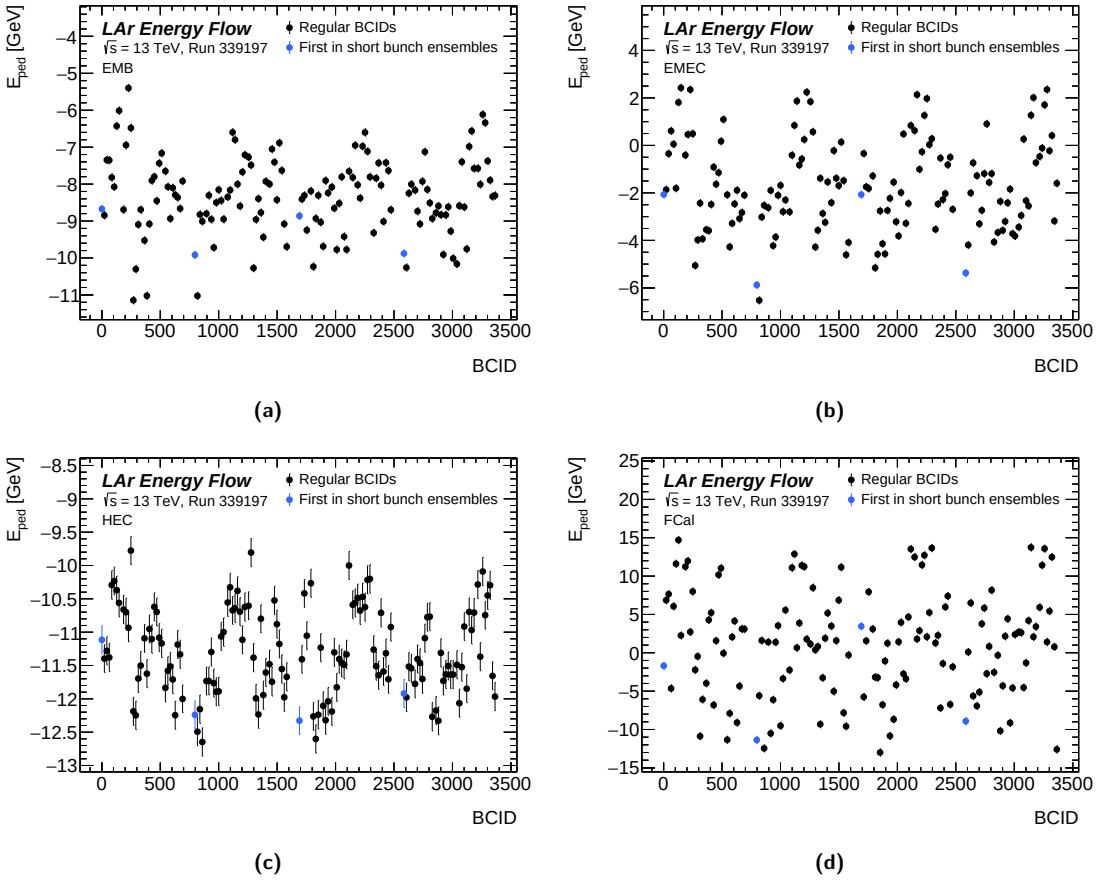

Figure 6.5: LAr energy-flow pedestals for each of the 140 colliding bunches in run 339197 for the (a) EMB, (b) EMEC, (c) HEC and (d) FCal. Regular BCIDs are shown in black, while the first BCIDs in the short bunch ensembles (1, 797, 1691 and 2585) are highlighted in blue. Here, a "short bunch ensemble" refers to the fill pattern used in run 339197 in which bunches were grouped into four "ensembles" of 32, 36, 36 and 36 colliding-bunch pairs, with approximately 21 empty bunch slots between colliding-bunch pairs within the ensemble, and at least 100 empty bunch slots between ensembles. The uncertainties are statistical only.

Table 6.5: Summary of pedestal estimates for run 354124 for each of the LAr detector subsystems and isolated BCIDs. Here, the pre-scan pedestals were evaluated during the period of full beam separation before the μ scan, and the post-scan pedestals during the period of full beam separation after the μ scan. The uncertainties are statistical only.

	Pre-scan E_{ped} (MeV)		Post-scan E_{ped} (MeV)	
Subsystem	BCID 1409	BCID 2770	BCID 1409	BCID 2770
EMB	1719.7 ± 9.1	2535.0 ± 9.5	1870 ± 17	2853 ± 17
EMEC	3253 ± 14	-2079 ± 14	2779 ± 28	-2381 ± 28
HEC	-4090 ± 18	-3593 ± 18	-5116 ± 37	-4594 ± 37
FCal	1810 ± 43	8266 ± 43	-1587 ± 87	5106 ± 87

a specific model for this apparent pedestal time dependence, it was decided to evaluate the pedestal at each step in the μ scan by linearly interpolating between the pre- and post-scan pedestals. The validity of this assumption that the pedestal varies linearly with time is verified in Appendix B. A flat pedestal subtraction was used for the periods before and after the μ scan using the pre- and post-scan pedestals, respectively. Taking the start time of the μ scan as $t = 0$ and the end time as $t = T$, and defining the normalized scan time $w_t \equiv t/T$, the pedestal q at time w_t is:

$$q(w_t) = \begin{cases} q_1, & w_t < 0 \\ (1 - w_t)q_1 + w_t q_2, & 0 \leq w_t \leq 1 \\ q_2, & w_t > 1 \end{cases}, \tag{6.20}$$

where q_1 and q_2 are the measured pre- and post-scan pedestals per BCID in units of μ, respectively.

Pedestals for run 336506

There was no period of full beam separation in run 336506, therefore an alternate method was needed. Here, a linear-extrapolation procedure was employed by fitting the energy-flow measurements as a function of μ measured by LUCID BiHitOR, in the form

$$E_{\text{flow}} = p_0 + p_1 \mu_{\text{LUCID}} ,$$

using the luminosity blocks at the lowest point of the μ scan. Since LUCID requires a μ correction for μ values greater than ~ 5,[6] only luminosity blocks satisfying $\mu_{\text{LUCID}} < 5$ were selected. Examples of these fits are shown in Figure 6.6 for BCID 83. The pedestal value was then evaluated by extrapolating to $\mu = 0$, i.e. by taking the p_0 value from the fit. The corresponding uncertainty on this pedestal was taken to be the uncertainty on the fitted p_0 value, σ_{p_0}.

Pedestals for run 354309

In run 354309, the beams were separated for approximately 23 min before the μ scan and for approximately 10 min immediately after the μ scan. However, as stated above, the beams were only separated in the horizontal plane, resulting in typical μ values of ~ 0.002 during these times for the BCIDs in question. A modified pedestal-subtraction and calibration method was therefore needed for this run. While the standard pedestal-subtraction method effectively anchors the energy-flow measurement at zero when there are no collisions, Eqs. (6.17) and (6.18) can be modified to instead anchor the energy-flow measurement to this non-zero luminosity μ_{ped} during these periods of partial beam separation:

$$\mu = \frac{\mu_{\text{ref}}^{(\text{peak})} - \mu_{\text{ped}}}{E_{\text{flow}}^{(\text{peak})} - E_{\text{ped}}} \left(E_{\text{flow}} - E_{\text{ped}}\right) + \mu_{\text{ped}} . \tag{6.21}$$

This procedure is illustrated in Figure 6.7. In this case, μ_{ped} was measured using LUCID C12, averaged over the same luminosity blocks in which E_{ped} was evaluated, and $\mu_{\text{ref}}^{(\text{peak})}$ was measured in the luminosity block at the peak of the μ scan using track counting.

[6] Any μ correction for LUCID below $\mu = 5$ is expected to be negligible.

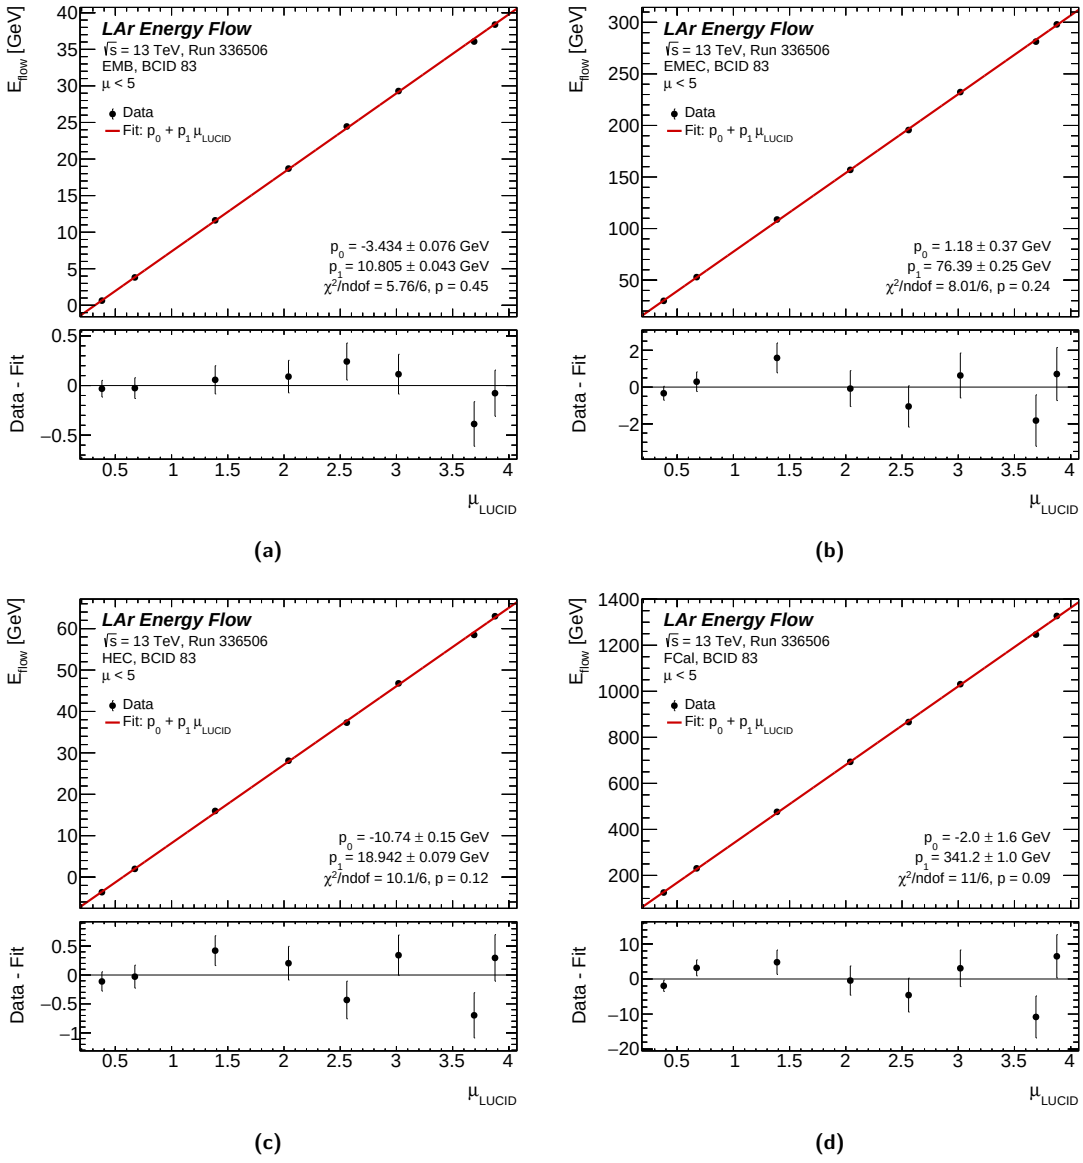

Figure 6.6: LAr energy-flow pedestal-extrapolation procedure in run 336506. Shown are the LAr energy-flow measurements in the run 336506 μ scan as a function of μ measured by LUCID BiHitOR for the (a) EMB, (b) EMEC, (c) HEC, and (d) FCal. Only luminosity blocks satisfying $\mu_{\mathrm{LUCID}} < 5$ were selected, and only data for BCID 83 is shown here. The red lines show linear fits to the data and the lower panels show the fit residuals. The uncertainties are statistical only, and are too small to be visible in the upper panel.

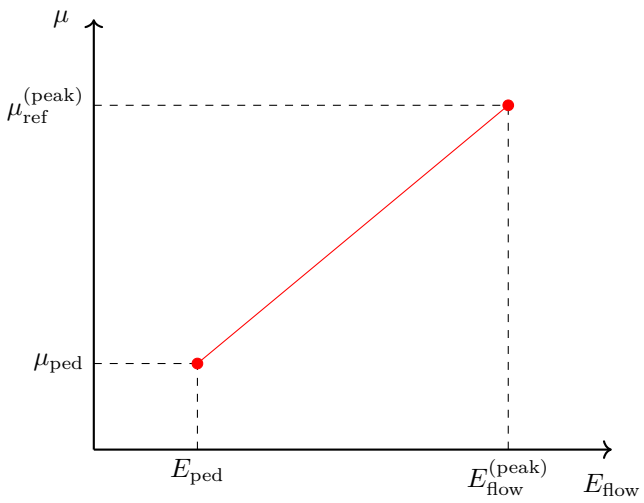

Figure 6.7: Illustration of the pedestal-subtraction and calibration procedure for the LAr energy-flow luminosity measurement in run 354309.

Table 6.6: Summary of pedestal estimates for run 354309 for each of the LAr detector subsystems and isolated BCIDs. Here, the pre-scan pedestals were evaluated during the period of partial beam separation before the μ scan, and the post-scan pedestals during the period of partial beam separation after the μ scan. The values of μ_{ped} measured by LUCID C12 in these pre- and post-scan periods are also given in the final row. The uncertainties in all cases are statistical only.

	Pre-scan E_{ped} (MeV)		Post-scan E_{ped} (MeV)	
Subsystem	BCID 823	BCID 2608	BCID 823	BCID 2608
EMB	-1858 ± 27	805 ± 26	-1832 ± 37	610 ± 36
EMEC	-605 ± 43	-886 ± 44	-978 ± 59	-1486 ± 60
HEC	-3837 ± 54	-3274 ± 54	-6060 ± 76	-5665 ± 76
FCal	-8130 ± 140	-880 ± 14	-4900 ± 190	1950 ± 190
	Pre-scan μ_{ped}		Post-scan μ_{ped}	
LUCID	$0.002\,591 \pm 0.000\,057$	$0.003\,754 \pm 0.000\,073$	$0.001\,927 \pm 0.000\,056$	$0.002\,891 \pm 0.000\,093$

As in run 354124, the pre- and post-scan pedestal estimates, given in Table 6.6, were not statistically consistent and differed by several percent, indicating a drift in the pedestal value. The formalism developed for run 354124, and outlined in Eq. (6.20), to linearly interpolate the pedestal between the pre- and post-scan periods was also employed here. The values of μ_{ped} also differ between the pre- and post-scan periods, however, since μ_{ped} is assumed to be the result of a small number of peripheral collisions, and therefore a function of the colliding-beam parameters and not an electronics pedestal, no interpolation is done. Instead, the first half of the μ scan, up to the luminosity block at peak μ, uses the pre-scan value of μ_{ped}, and the second half uses the post-scan value. Since μ_{ped} is small (order 0.005 % of $\mu_{\text{ref}}^{(\text{peak})} \approx 60$), the discontinuity this step introduces is negligible.

6.3.3 Results

The results from the LAr energy-flow luminosity measurements and their comparisons to track counting are shown in Figures 6.8 (run 339197), 6.9 (run 354124), 6.10 (run 336506) and 6.11 (run

354309). These plots show the luminosities measured by LAr energy flow and by track counting as a function of luminosity block number (and/or pseudo luminosity block number in runs 339197 and 336506), as well as their ratios. For run 339197, the data are averaged over all 140 colliding-bunch pairs, while for the other runs, data for only a single isolated colliding-bunch pair are shown. In all cases, the track-counting data use the 2017 selection working point (Table 6.1).

These results demonstrate that the LAr energy-flow and track-counting luminosity algorithms generally agree, with a few notable exceptions. First, the HEC generally gives poor agreement with track counting and with the other three LAr subsystems under nearly all running conditions: levelled low μ, head-on collisions at high μ, and during the μ scans. Second, the EMEC appears to drift by a few percent with respect to track counting during the levelled low-μ period of run 339197. Third, there is a $\sim 1\,\%$ step in run 339197 between the luminosity at the peak of the μ scan and the period of head-on collisions after the scan for both the EMB and, to a lesser extent, the EMEC, despite them being at roughly the same μ. Finally, the luminosity-block-to-luminosity-block scatter during the levelled low-μ periods of run 339197 and 354124 is several times larger than the statistical uncertainty on both the energy-flow and track-counting measurements. These disagreements between the EMB, EMEC, FCal and track counting suggest additional systematic uncertainties that have not been accounted for in the LAr energy-flow measurement, such as short-term and long-term drifts in the pedestals. This possibility is explored in Section 6.3.4 below.

6.3.4 Interpretations

To quantify the relative nonlinearity between LAr energy flow and track counting, fits of their ratio as a function of μ were performed using the μ-scan data in each of the four runs. Since the objective of the LAr energy-flow luminosity measurement is to validate the assumption that track counting is linear over a large range in μ, from vdM scan running conditions to physics data-taking conditions, only luminosity blocks with $\mu > 0.5$ were considered.

The "baseline" model is a linear fit in the form:

$$\mu_{\mathrm{LAr}} = p_0 \mu_{\mathrm{Tracks}} + p_1 \mu_{\mathrm{Tracks}}^2 \,,$$

or equivalently:

$$\frac{\mu_{\mathrm{LAr}}}{\mu_{\mathrm{Tracks}}} = p_0 + p_1 \mu_{\mathrm{Tracks}} \,,$$

where the p_0 parameter accounts for the relative scale of the two luminosity algorithms and the p_1 parameter accounts for the relative nonlinearity between them at first order. To assess the effect of pedestal variations on the p_0 and p_1 parameters, nuisance parameters were introduced to the fit model according to the pedestal-subtraction method for that run.

The most general model with pedestal nuisance parameters is for the pedestal-subtraction method used in run 354124, with both pre- and post-scan pedestal measurements and a linearly interpolated pedestal during the μ scan portion of the run, as shown in Eq. (6.20) for $0 \leq w_t \leq 1$. In this case, nuisance parameters ν_1 and ν_2 were introduced, which quantify the variation of the true pedestals before and after the scan from their measured values. The pedestal, in units of μ, becomes

$$q(w_t) = (1 - w_t)(q_1 + \nu_1) + w_t(q_2 + \nu_2) \,.$$

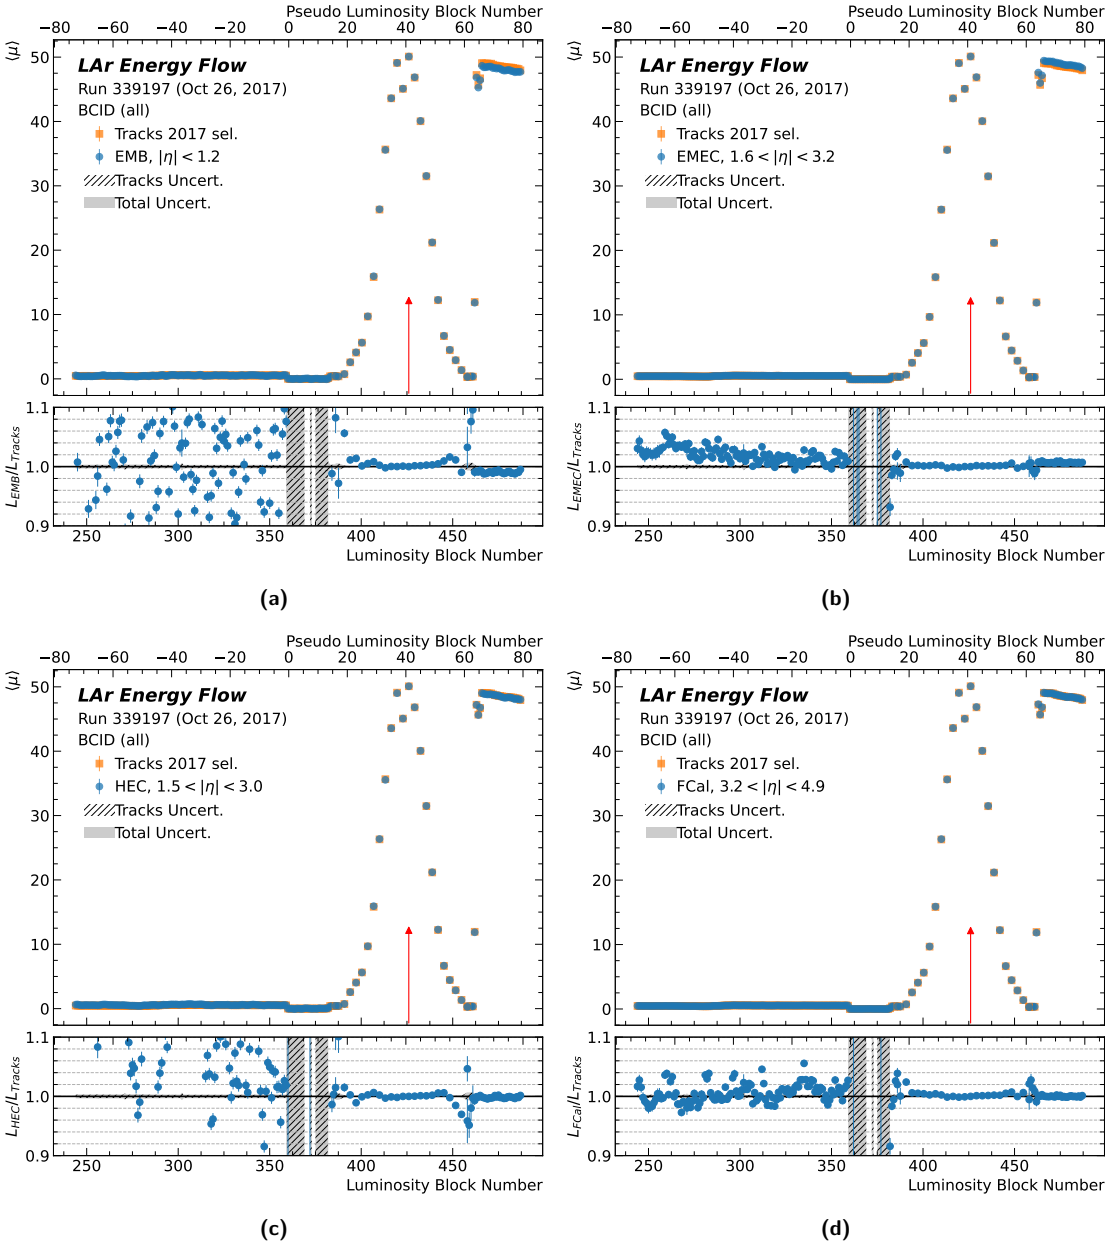

Figure 6.8: Luminosities measured by LAr energy flow and by track counting, and their ratio, in run 339197 for the (a) EMB, (b) EMEC, (c) HEC, and (d) FCal, plotted as a function of luminosity block number. Pseudo luminosity blocks were used for the μ-scan portion of the run, with the pseudo luminosity block number shown on the top axes. The data are averaged over all 140 colliding-bunch pairs. The red arrow indicates the pseudo luminosity block at the peak of the scan in which the energy-flow luminosity was anchored to track counting. The error bars show the statistical uncertainties. The hatched-band in the ratio panel shows the track-counting uncertainty and the shaded band shows the "total uncertainty": the track-counting uncertainty added in quadrature with the LAr pedestal uncertainty.

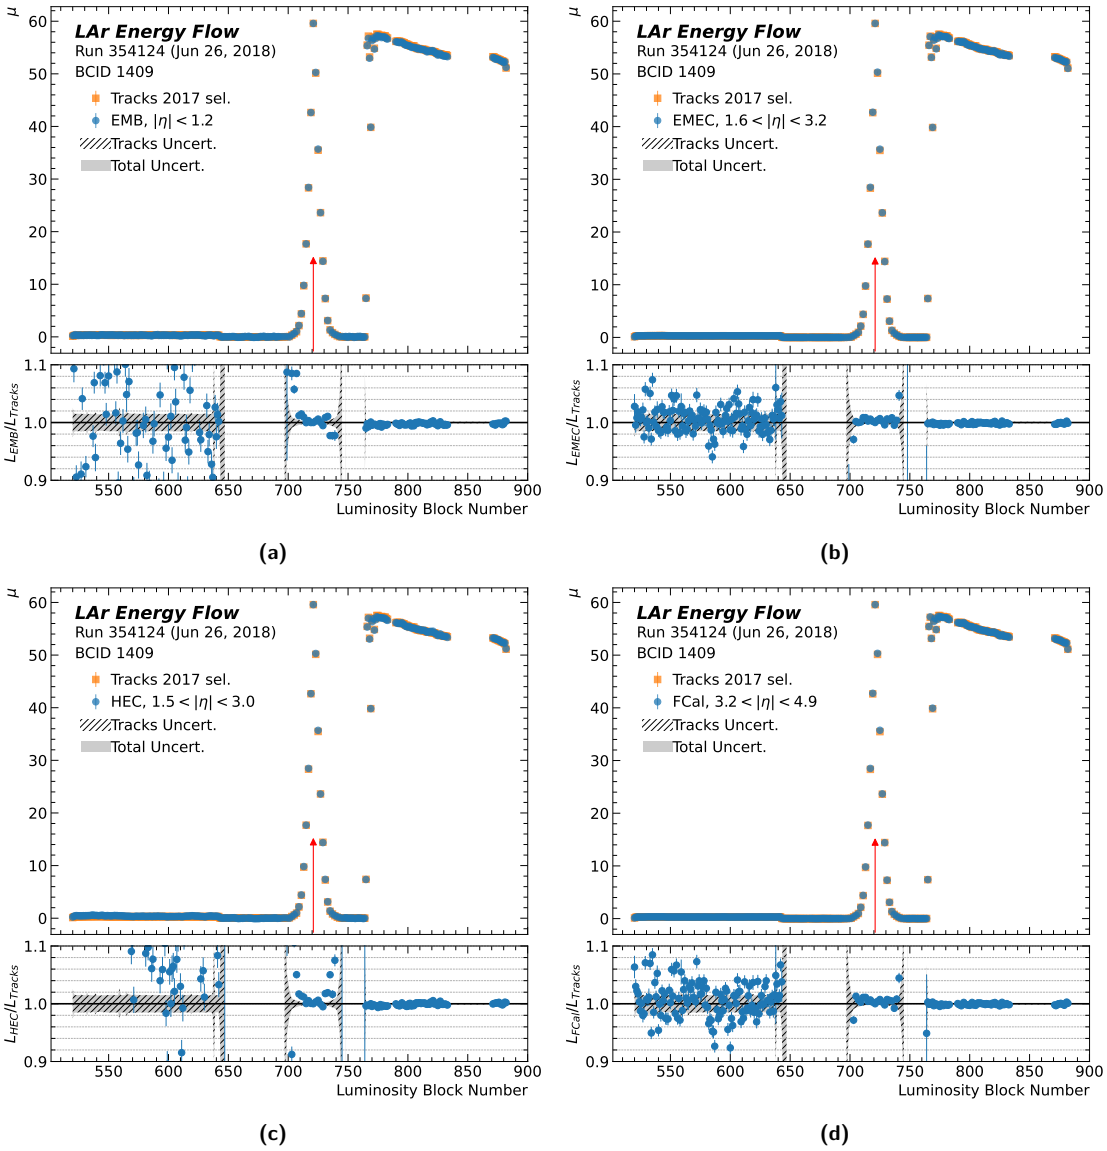

Figure 6.9: Luminosities measured by LAr energy flow and by track counting in BCID 1409, and their ratio, in run 354124 for the (a) EMB, (b) EMEC, (c) HEC, and (d) FCal, plotted as a function of luminosity block number. The red arrow indicates the luminosity block at the peak of the scan in which the energy-flow luminosity was anchored to track counting. The error bars show the statistical uncertainties. The hatched-band in the ratio panel shows the track-counting uncertainty and the shaded band shows the "total uncertainty": the track-counting uncertainty added in quadrature with the LAr pedestal uncertainty.

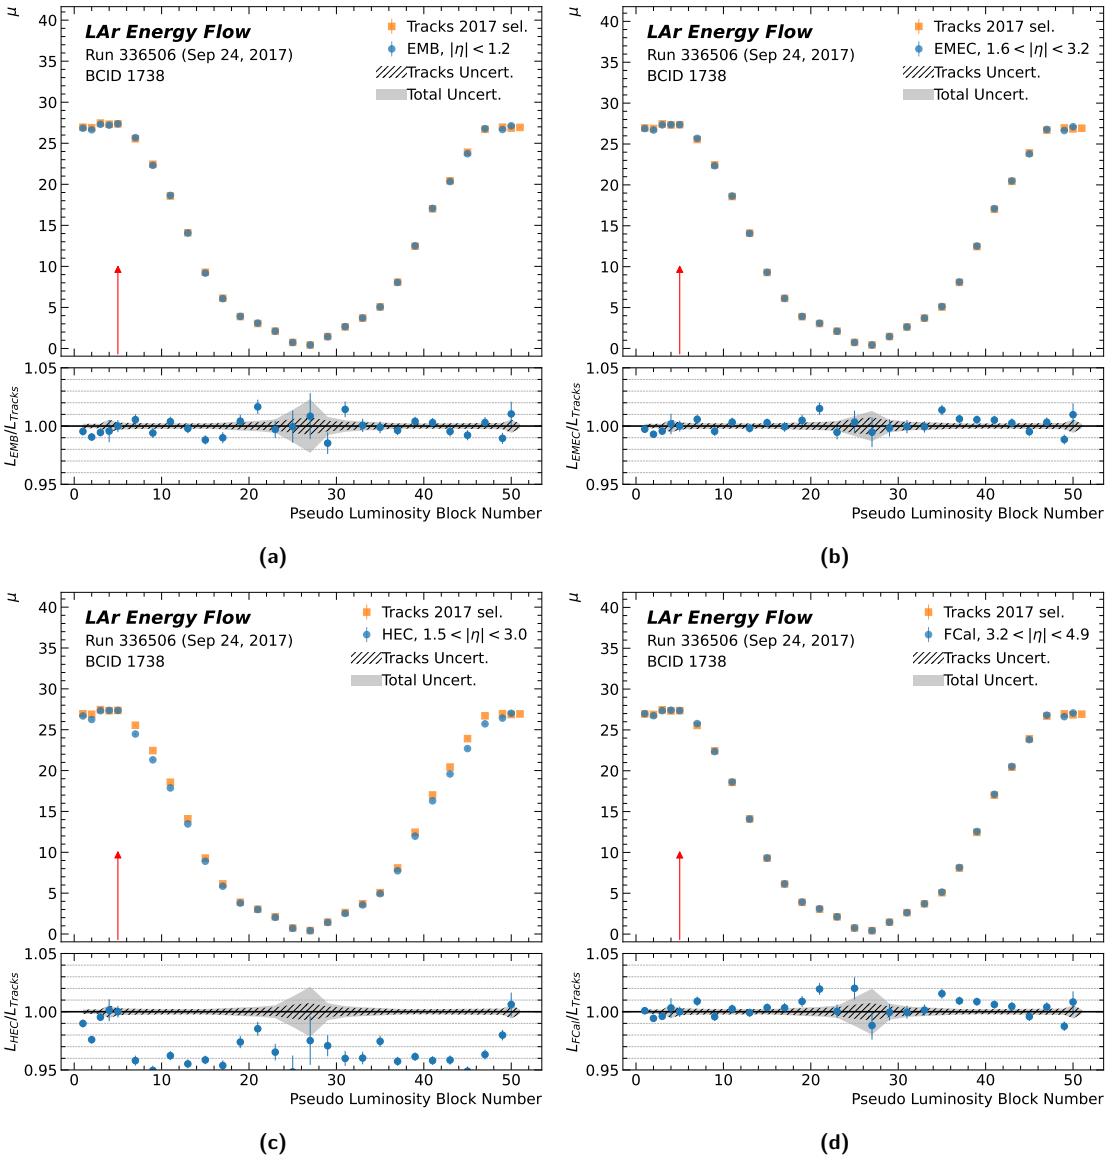

Figure 6.10: Luminosities measured by LAr energy flow and by track counting in BCID 1738, and their ratio, in run 336506 for the (a) EMB, (b) EMEC, (c) HEC, and (d) FCal, plotted as a function of pseudo luminosity block number. Only the μ-scan portion of the run is shown here. The red arrow indicates the pseudo luminosity block in which the energy-flow luminosity was anchored to track counting. The error bars show the statistical uncertainties. The hatched-band in the ratio panel shows the track-counting uncertainty and the shaded band shows the "total uncertainty": the track-counting uncertainty added in quadrature with the LAr pedestal uncertainty.

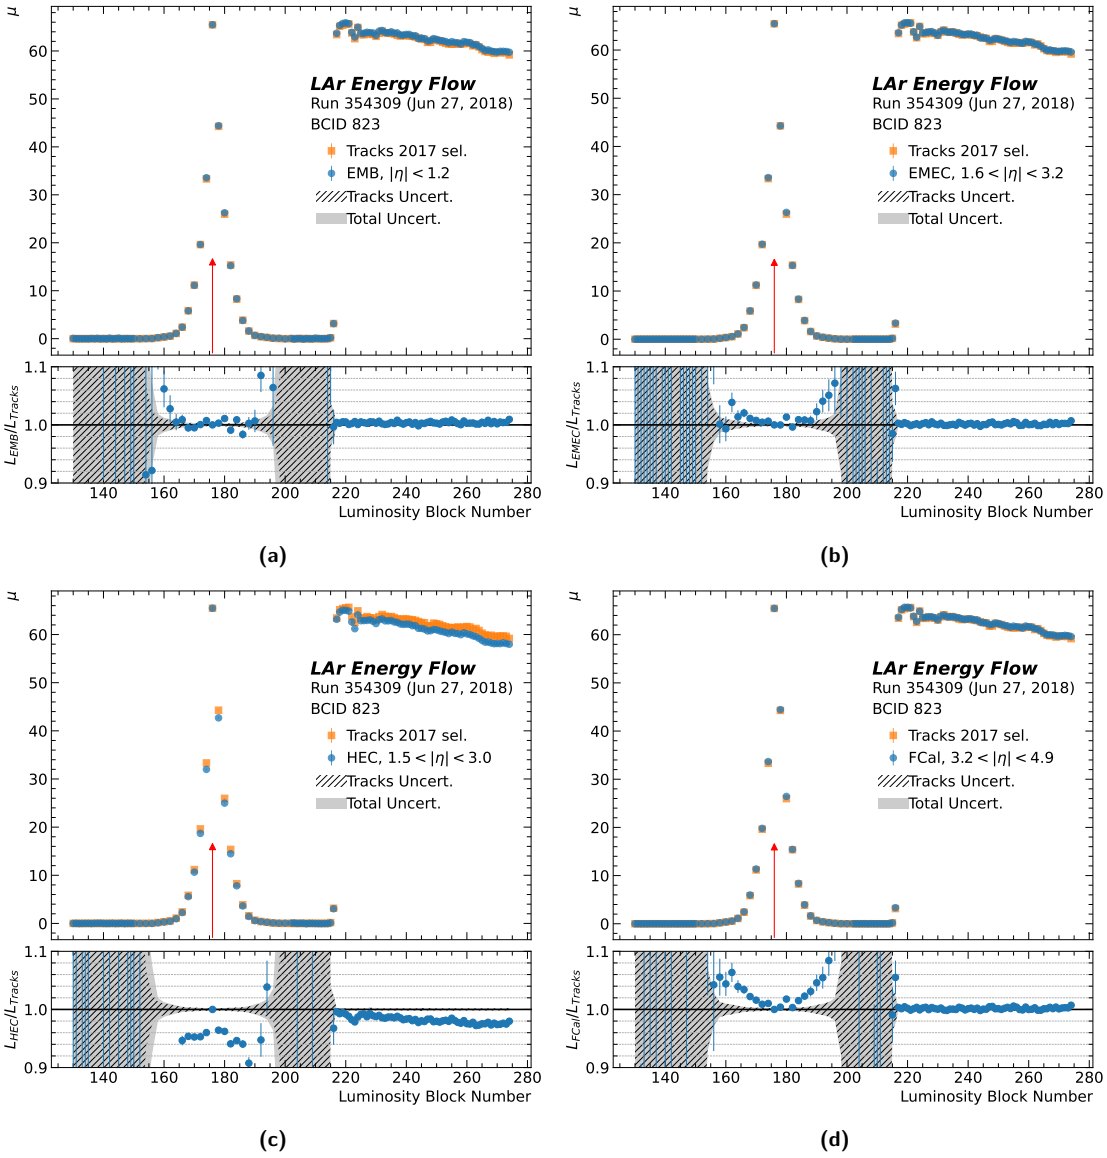

Figure 6.11: Luminosities measured by LAr energy flow and by track counting in BCID 823, and their ratio, in run 354309 for the (a) EMB, (b) EMEC, (c) HEC, and (d) FCal, plotted as a function of luminosity block number. The red arrow indicates the luminosity block at the peak of the scan in which the energy-flow luminosity was anchored to track counting. The error bars show the statistical uncertainties. The hatched-band in the ratio panel shows the track-counting uncertainty and the shaded band shows the "total uncertainty": the track-counting uncertainty added in quadrature with the LAr pedestal uncertainty.

If we consider LAr energy flow data that has already been fully calibrated with the pedestal subtracted, as in Eq. (6.17), the model can be reformulated in terms of a "residual" pedestal:

$$q_{\text{resid}}(w_t) = (1 - w_t)\nu_1 + w_t\nu_2 \,.$$

Given this form of the pedestal, a fit model can be constructed as before with the assumption that μ_{LAr} follows the form

$$\mu_{\text{LAr}} = p_0\mu_{\text{Tracks}} + p_1\mu_{\text{Tracks}}^2 + q_{\text{resid}}(w_t) \,.$$

The model function to fit is then the ratio of μ_{LAr} to μ_{Tracks}:

$$\begin{aligned} f(\mu_{\text{Tracks}}, w_t; p_0, p_1, \nu_1, \nu_2) &= \frac{p_0\mu_{\text{Tracks}} + p_1\mu_{\text{Tracks}}^2 + q_{\text{resid}}(w_t)}{\mu_{\text{Tracks}}} \\ &= p_0 + p_1\mu_{\text{Tracks}} + \frac{q_{\text{resid}}(w_t)}{\mu_{\text{Tracks}}} \,. \end{aligned} \tag{6.22}$$

If the pre- and post-scan pedestals are measured to be $q_1 \pm \sigma_1$ and $q_2 \pm \sigma_2$, this is equivalent to measuring $\nu_1 = 0 \pm \sigma_1$ and $\nu_2 = 0 \pm \sigma_2$, so the χ^2 function to minimize is

$$\chi^2(y; p_0, p_1, \nu_1, \nu_2) = \frac{\nu_1^2}{\sigma_1^2} + \frac{\nu_2^2}{\sigma_2^2} + \sum_{i=1}^{N} \frac{[y_i - f(\mu_{\text{Tracks}}^i, w_t^i; p_0, p_1, \nu_1, \nu_2)]^2}{\sigma_i^2} \,, \tag{6.23}$$

where $y_i \pm \sigma_i$ are the ratios of the LAr energy-flow and track-counting measurements, and their uncertainties, for each luminosity block $i = 1, \ldots, N$. The ν_1^2/σ_1^2 and ν_2^2/σ_2^2 terms are referred to as the pre- and post-scan *penalty terms*, respectively. The pedestal nuisance parameter formalism can be modified to accommodate the pedestal-subtraction method of each run, which are described below along with the results of the fits.

Run 339197

While a constant pedestal subtraction was used in run 339197, using the pedestals evaluated in the pre-scan period of the run in which the beams were fully separated, the formalism described above may still be used, but with the post-scan ν_2 penalty term removed since no measurement of the post-scan pedestal q_2 could be made. The ν_2 nuisance parameter is still included in the model and is constrained by the data during the scan.

Examples of the fit results are shown in Figure 6.12 for BCID 1 and 22, for each of the EMB, EMEC and FCal (results for the HEC are not shown since it has already been demonstrated that it gives unreliable measurements). The p_1 parameters extracted from these fits for each BCID are shown in Figure 6.13, as well as the weighted mean. The p_1 parameters are also summarized in Table 6.7, along with the equivalent values extracted from data recorded in the remaining runs.

Run 354124

Since both pre- and post-scan pedestal measurements were taken in run 354124, the formalism described by Eqs. (6.22) and (6.23) may be directly applied. The fit results are shown in Figure 6.14 for both isolated colliding-bunch pairs (BCIDs 1409 and 2770), for each of the EMB, EMEC and FCal.

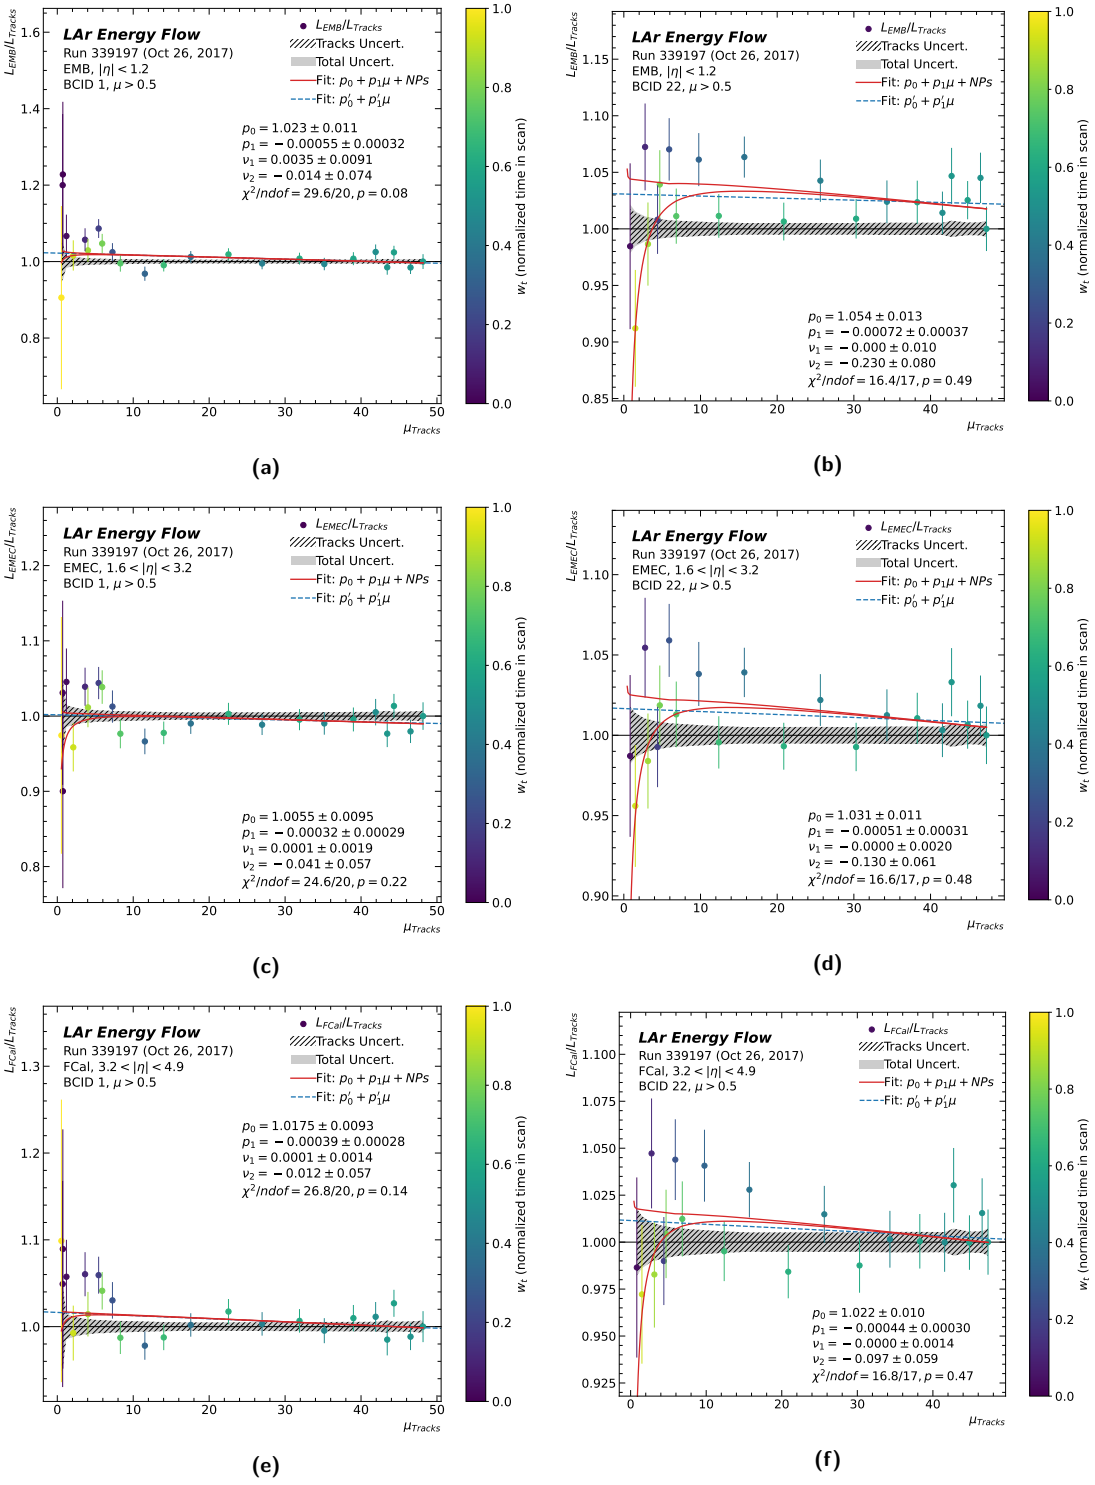

Figure 6.12: Ratios of the LAr energy-flow luminosity to track counting as a function of μ during the μ-scan portion of run 339197 for the (a and b) EMB, (c and d) EMEC, and (e and f) FCal. The left column shows the results for BCID 1, and the right column shows the results for BCID 22. The result of the fit, which is a function of μ_{Tracks} (shown on the horizontal axis) and the normalized scan time w_t (shown according to the colour scale), is shown by the red line, which traces out the three-dimensional surface at each scan step, and projected onto the μ_{Tracks}-$\mu_{\text{LAr}}/\mu_{\text{Tracks}}$ plane. A linear fit without pedestal nuisance parameters is shown for reference by the blue dashed line.

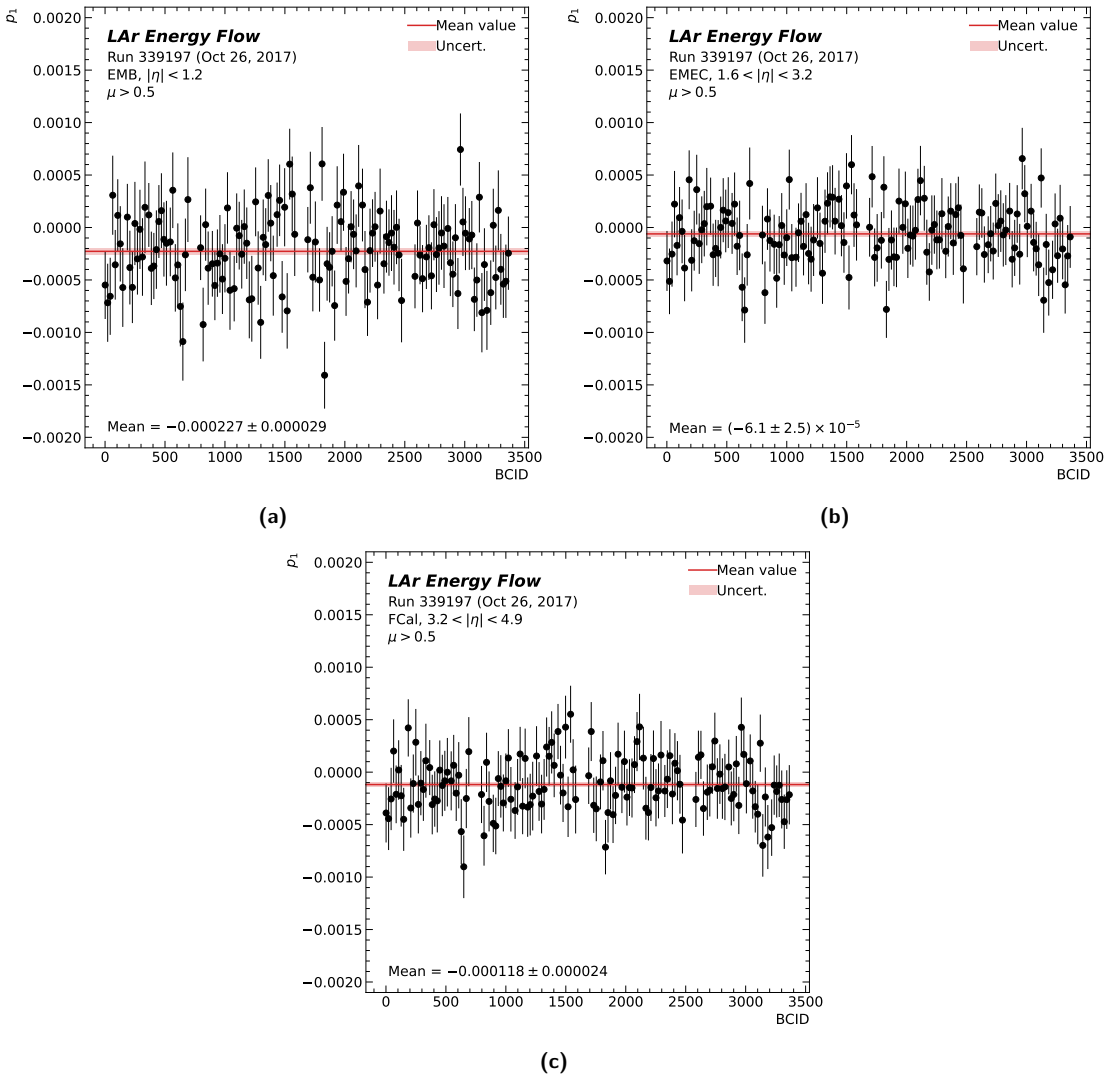

Figure 6.13: The p_1 parameters extracted from the fits to the LAr energy-flow and tracking-counting luminosity measurements for each of the 140 isolated colliding-bunch pairs in run 339197. The weighted mean over all colliding-bunch pairs is shown by the solid red line, with the corresponding uncertainty shown by the red shaded area.

Run 336506

Since the run 336506 μ scan is a "down-up" scan,[7] pedestal drifts of several percent over the course of $\sim 1\,\mathrm{h}$ are negligible given that the tails of the μ scan are at $\mu \approx 25$. Therefore, a simplified, time-independent fit model was used for this run. In this case, the fit function becomes

$$f(\mu_{\mathrm{Tracks}}; p_0, p_1, \nu_1) = p_0 + p_1 \mu_{\mathrm{Tracks}} + \frac{\nu_1}{\mu_{\mathrm{Tracks}}} ,$$

[7] A "down-up" μ scan refers to one in which the beams begin in head-on collisions, are incrementally separated in each step of the "downward arm" of the scan and then brought back into head-on collisions in the "upward arm" of the scan.

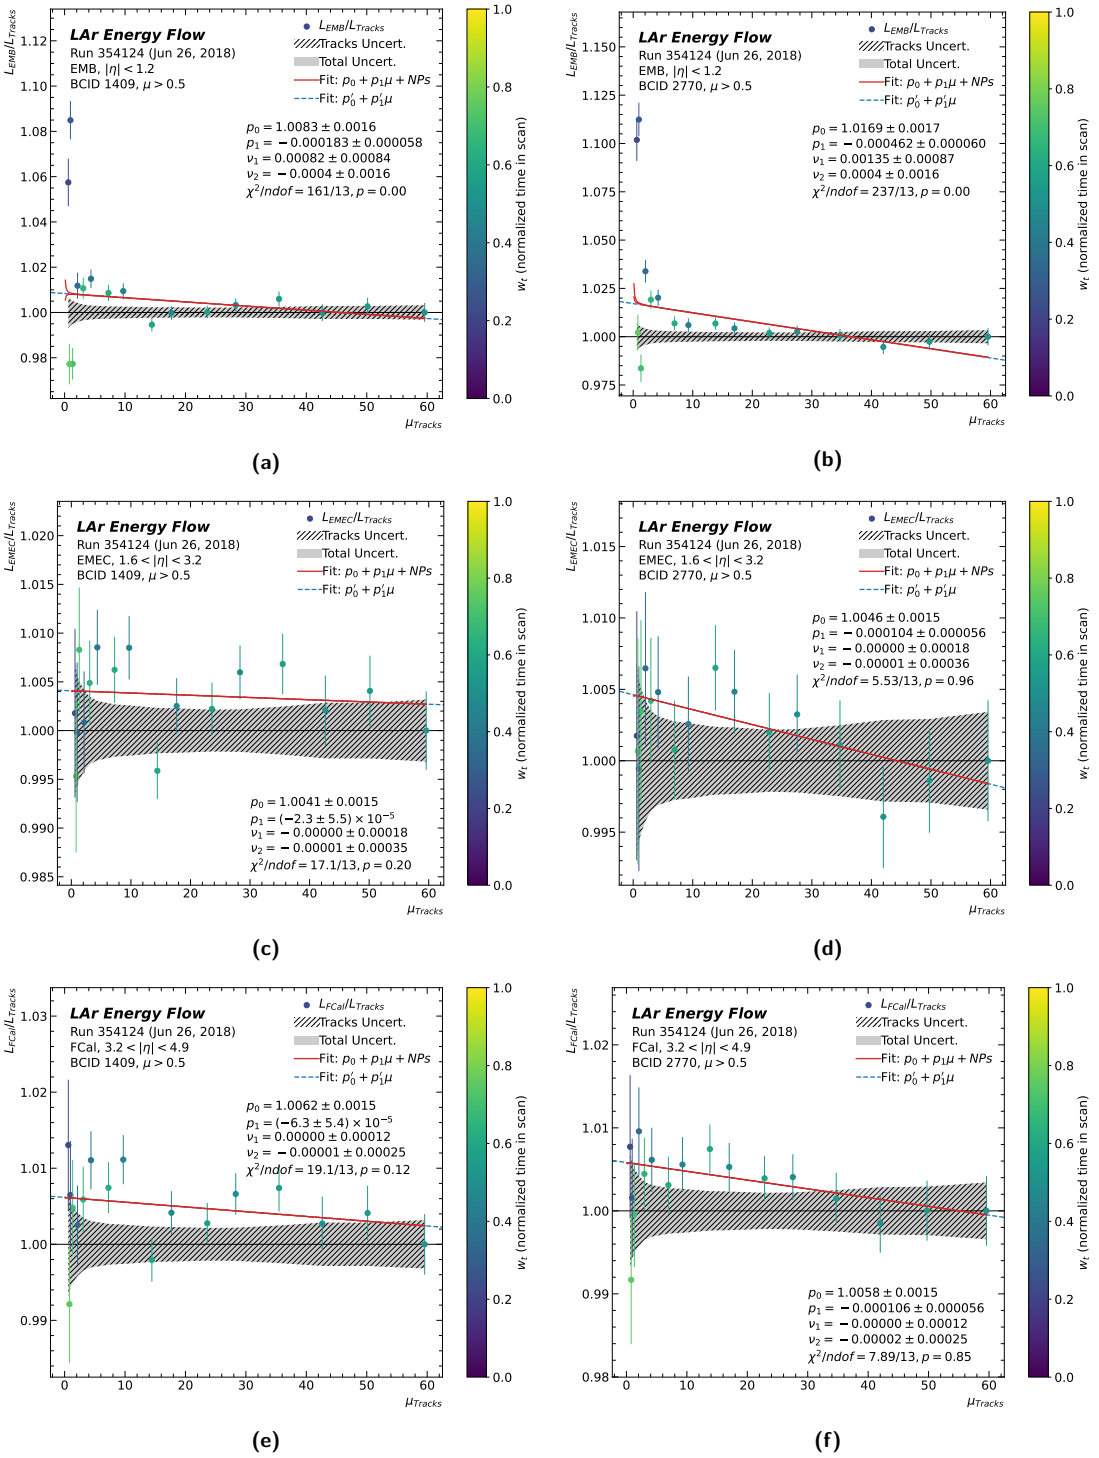

Figure 6.14: Ratios of the LAr energy-flow luminosity to track counting as a function of μ during the μ-scan portion of run 354124 for the (a and b) EMB, (c and d) EMEC, and (e and f) FCal. The left column shows the results for BCID 1409, and the right column shows the results for BCID 2770. The result of the fit, which is a function of μ_{Tracks} (shown on the horizontal axis) and the normalized scan time w_t (shown according to the colour scale), is shown by the red line, which traces out the three-dimensional surface at each scan step, and projected onto the μ_{Tracks}-$\mu_{\text{LAr}}/\mu_{\text{Tracks}}$ plane. A linear fit without pedestal nuisance parameters is shown for reference by the blue dashed line.

Table 6.7: Summary of the p_1 parameters extracted from the fits to the LAr energy-flow and tracking-counting luminosity measurements in each of the four LArPEB-stream runs and for each detector subsystem. Results are given for each individual BCID as well as their weighted mean. For run 339197, in which all 140 BCIDs were recorded to the LArPEB stream, only the weighted mean is given.

Run number	BCID	p_1, EMB [$\times 10^{-3}$]	p_1, EMEC [$\times 10^{-3}$]	p_1, HEC [$\times 10^{-3}$]	p_1, FCal [$\times 10^{-3}$]
336506	83	-0.20 ± 0.14	-0.51 ± 0.13	0.20 ± 0.14	-0.74 ± 0.12
	1738	-0.07 ± 0.16	-0.27 ± 0.13	0.18 ± 0.15	-0.38 ± 0.14
	Mean	-0.15 ± 0.11	-0.397 ± 0.092	0.19 ± 0.10	-0.586 ± 0.094
339197	Mean	-0.227 ± 0.029	-0.061 ± 0.025	-0.137 ± 0.027	-0.118 ± 0.024
354124	1409	-0.183 ± 0.058	-0.023 ± 0.055	-0.354 ± 0.058	-0.063 ± 0.054
	2770	-0.462 ± 0.060	-0.104 ± 0.056	-0.393 ± 0.059	-0.106 ± 0.056
	Mean	-0.318 ± 0.042	-0.063 ± 0.039	-0.373 ± 0.041	-0.084 ± 0.039
354309	823	0.040 ± 0.073	-0.175 ± 0.067	0.963 ± 0.070	-0.371 ± 0.065
	2608	0.001 ± 0.075	-0.203 ± 0.069	1.132 ± 0.071	-0.386 ± 0.067
	Mean	0.021 ± 0.053	-0.189 ± 0.048	1.046 ± 0.050	-0.378 ± 0.047

and the χ^2 minimization function is

$$\chi^2(y; p_0, p_1, \nu_1) = \frac{\nu_1^2}{\sigma_1^2} + \sum_{i=1}^{N} \frac{[y_i - f(\mu^i_{\text{Tracks}}; p_0, p_1, \nu_1)]^2}{\sigma_i^2} \,.$$

The fit results are shown in Figure 6.15 for both isolated colliding-bunch pairs (BCIDs 83 and 1738), for each of the EMB, EMEC and FCal.

Run 354309

As with run 354124, both pre- and post-scan pedestals were measured in run 354309, therefore the formalism described by Eqs. (6.22) and (6.23) may also be applied for this run. The fit results are shown in Figure 6.16 for both isolated colliding-bunch pairs (BCIDs 823 and 2608), for each of the EMB, EMEC and FCal.

Comparison between runs

The fitted p_1 parameters from each of the four runs and for each of the EMB, EMEC and FCal are shown in Figure 6.17. Results from the HEC are also shown for completeness. With $p_1 = 0$ indicating no relative nonlinearity between LAr energy flow and track counting, the EMEC and the FCal show the best agreement with track counting in the two 140-bunch fills (runs 339197 and 354124). This figure also shows the same data scaled by a factor of $\mu_{\text{max}} = 50$, in percent. This quantity gives the relative nonlinearity from typical vdM conditions at $\mu = 0.5$ to typical physics data-taking conditions at $\mu = 50$. The plot shows the agreement between the EMEC, FCal and track counting are on the order of 0.5 % in the 140-bunch runs, and 1–3 % in the $\sim$ 1200-bunch fills (runs 336506 and 354309).

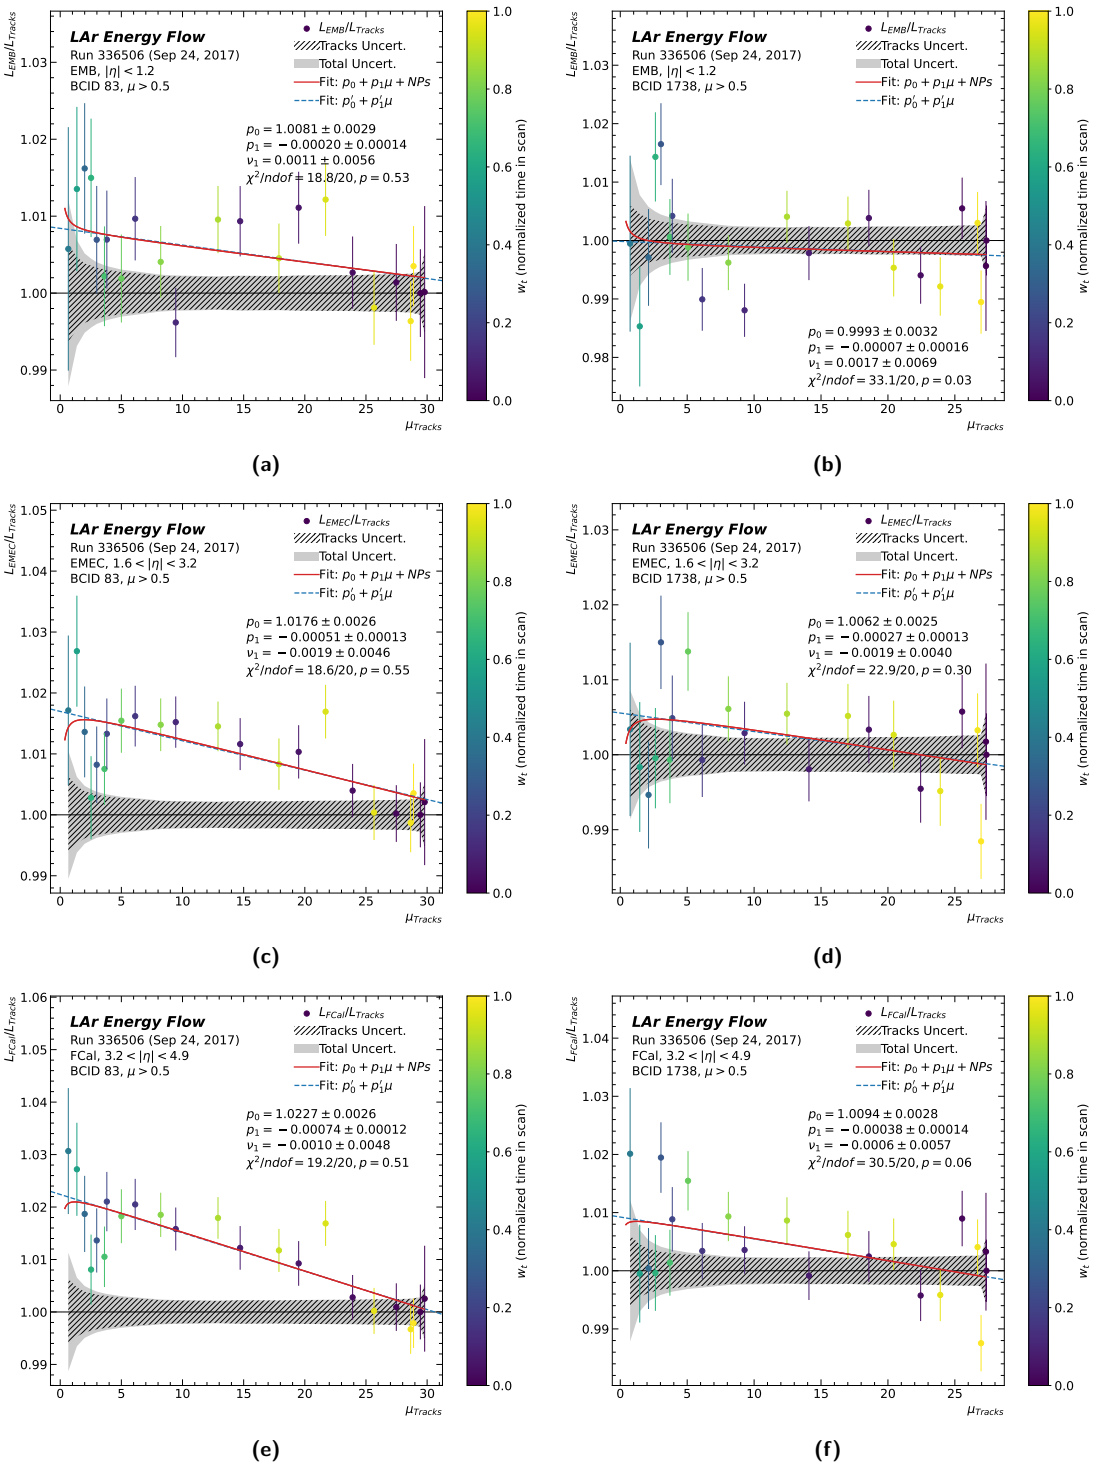

Figure 6.15: Ratios of the LAr energy-flow luminosity to track counting as a function of μ during the μ-scan portion of run 336506 for the (a and b) EMB, (c and d) EMEC, and (e and f) FCal. The left column shows the results for BCID 83, and the right column shows the results for BCID 1738. The result of the fit, which is a function of μ_{Tracks} only, is shown by the red line. A linear fit without pedestal nuisance parameters is shown for reference by the blue dashed line.

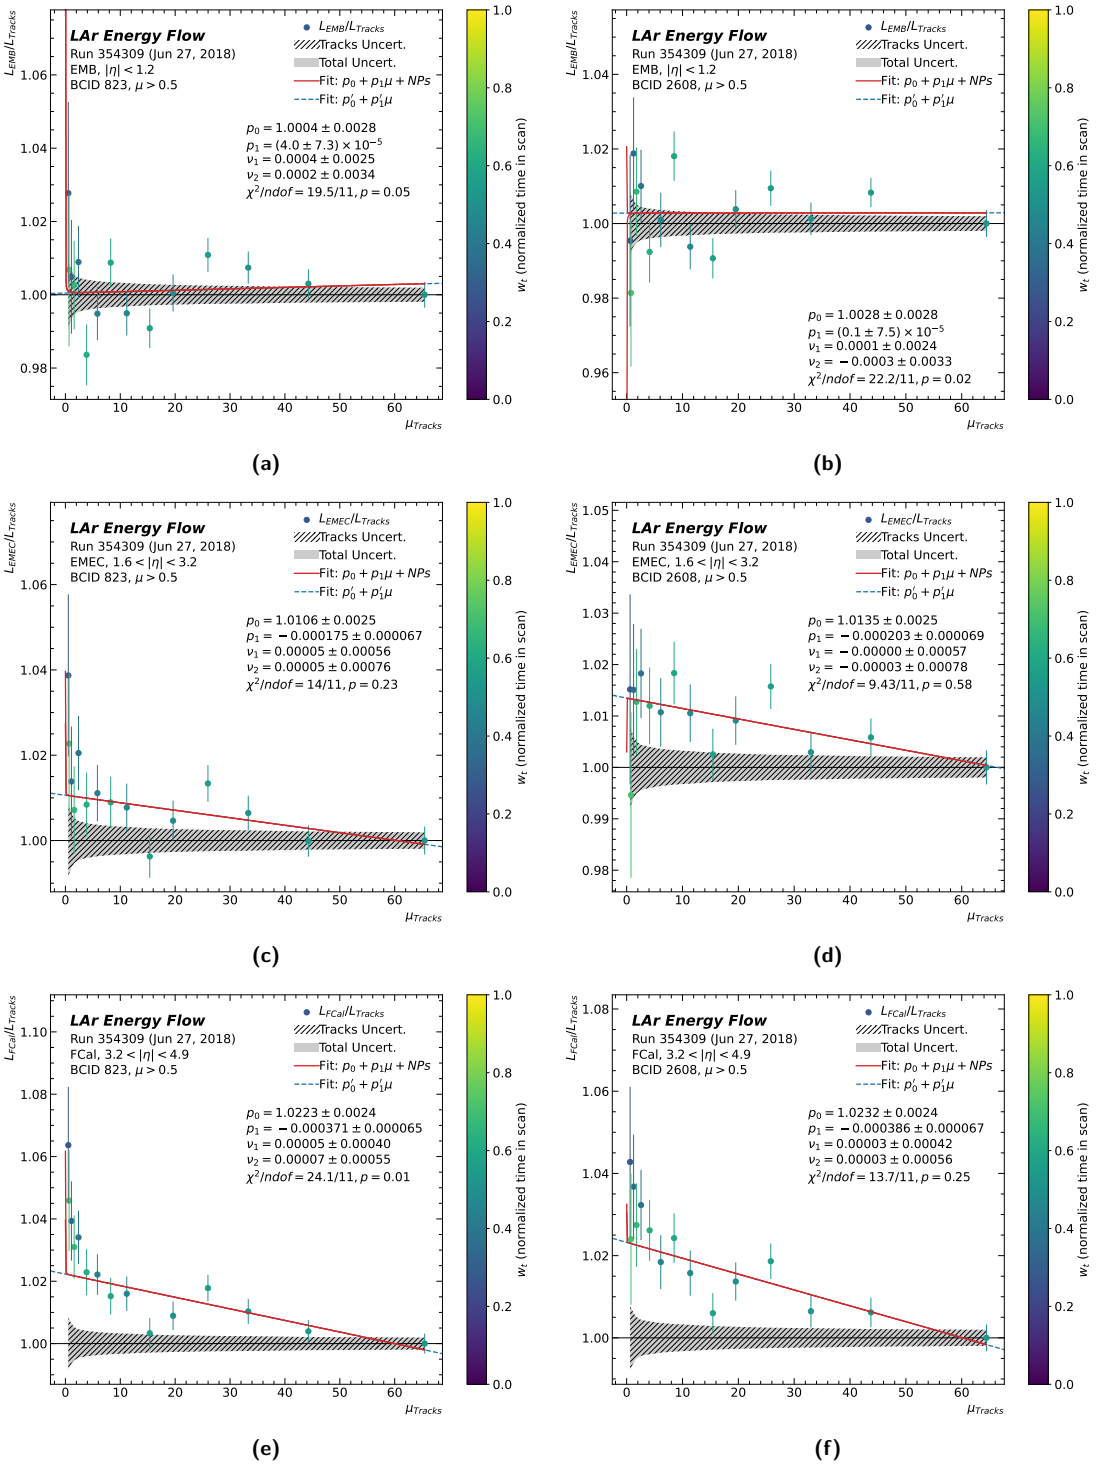

Figure 6.16: Ratios of the LAr energy-flow luminosity to track counting as a function of μ during the μ-scan portion of run 354309 for the (a and b) EMB, (c and d) EMEC, and (e and f) FCal. The left column shows the results for BCID 823, and the right column shows the results for BCID 2608. The result of the fit, which is a function of μ_{Tracks} (shown on the horizontal axis) and the normalized scan time w_t (shown according to the colour scale), is shown by the red line, which traces out the three-dimensional surface at each scan step, and projected onto the μ_{Tracks}-$\mu_{\text{LAr}}/\mu_{\text{Tracks}}$ plane. A linear fit without pedestal nuisance parameters is shown for reference by the blue dashed line.

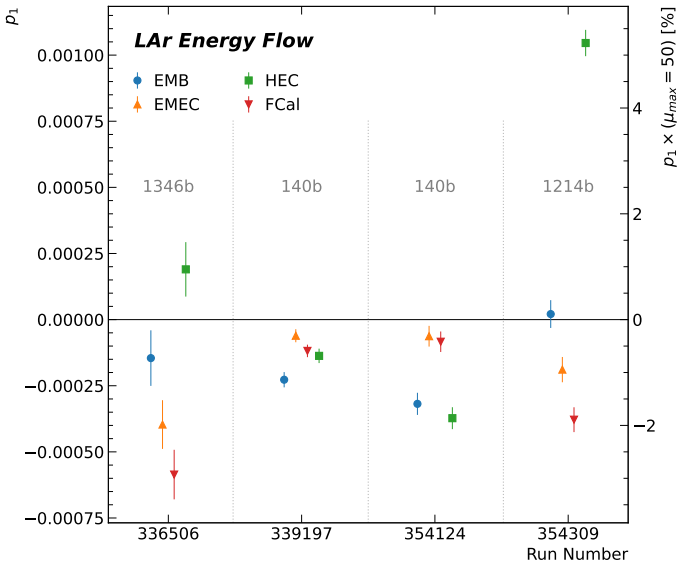

Figure 6.17: Summary of the bunch-averaged p_1 parameters extracted from the fits to the LAr energy-flow and tracking-counting luminosity measurements in each of the four LArPEB-stream runs. The left-hand vertical axis shows the raw p_1 values, while the right-hand vertical shows the p_1 values scaled by a factor of $\mu_{\text{max}} = 50$, in percent. The number of colliding-bunch pairs in each run is also indicated.

6.4 Summary and outlook

A new luminosity algorithm, called *LAr energy flow*, has been developed for complementary measurements of the ATLAS luminosity in a number of special data-taking runs in 2017 and 2018 during LHC Run 2. For LHC fills with sufficiently isolated bunches ($\gtrsim 450\,\text{ns}$ bunch spacing), the total energy deposited in a group of LAr calorimeter cells and averaged over a luminosity block is proportional to the instantaneous luminosity (after pedestal subtraction). A dedicated data stream reading out only the LAr calorimeter data in randomly sampled bunch crossings was used in these runs to derive a LAr energy flow luminosity measurement. Comparisons of the LAr energy-flow and track-counting luminosity measurements in special μ-scan runs with only isolated bunches, parameterized as the ratio $\mu_{\text{LAr}}/\mu_{\text{Tracks}}$ as a function of μ, result in small negative slopes corresponding to a negative deviation in the track-counting luminosity algorithm of up to 0.5 % at $\langle\mu\rangle \approx 50$.

These LAr energy-flow luminosity measurements can be exploited to constrain potential nonlinearity in the track-counting luminosity measurement used in the LUCID calibration-transfer procedure using a new method to set the calibration-transfer uncertainty called the *uncertainty ladder approach*. The ladder approach considers the transition from the vdM to physics regimes in three steps, (i) low-μ to high-μ isolated bunches, (ii) high-μ isolated bunches to high-μ bunch trains, then (iii) an increasing number of bunches in trains until the nominal LHC fill pattern is reached. The 0.5 % deviation between the LAr energy-flow and track-counting measurements at $\langle\mu\rangle \approx 50$ can be taken as the uncertainty in step (i) of the uncertainty ladder. Ongoing studies of material-activation effects and rapid scintillator ageing in the TILE E cells can be used to improve the comparisons between these measurements and track counting to set the uncertainties in remaining steps. Taken together, the uncertainties evaluated in the ladder approach will offer a substantial improvement to the preliminary calibration-transfer uncertainty of 1.3 % as described in Section 6.2.3.

Chapter 7

Search for Heavy Resonances Decaying into a Pair of Z Bosons

7.1 Introduction and analysis strategy

The discovery of a scalar particle by the ATLAS and CMS collaborations [70, 71] in 2012, with measured properties [209–213] consistent with those of the Standard Model Higgs boson, was a major milestone in the understanding of electroweak symmetry breaking. One important remaining question is whether the discovered particle is part of an extended scalar sector as postulated by various extensions to the Standard Model such as the two-Higgs-doublet model (2HDM) and others described in Section 2.3. These extensions predict additional Higgs bosons, motivating searches in an extended mass range.

This chapter reports on a search for heavy resonances decaying into a pair of Z bosons. To perform this search, events consistent with such decays must first be selected from the dataset. As the heavy resonances are assumed to be short-lived, both they and the pair of Z bosons decay before interacting with any part of the detector. Two final states, or *decay channels*, are considered: $X \to ZZ \to \ell^+\ell^-\ell^+\ell^-$ and $X \to ZZ \to \ell^+\ell^-\nu\bar{\nu}$, where X stands for the heavy resonance in question, with the focus of this thesis on the $\ell^+\ell^-\ell^+\ell^-$ channel. Here, ℓ denotes either an electron or muon, as these final-state leptons can be fully reconstructed with high precision and with excellent energy resolution, while ν denotes all three neutrino flavours. Tau leptons have comparatively worse energy resolution and larger systematic uncertainties, which limits their utility to improve the experimental sensitivity [214], and are not considered in this analysis.

The presence of four leptons in the final state serves as a powerful discriminant to select events consistent with a new heavy resonance decay given that very few physics processes in proton–proton collisions lead to events with a large lepton multiplicity. However, this clean signal generally comes at the cost of the low $ZZ \to \ell^+\ell^-\ell^+\ell^-$ branching ratio and the requirement that all four leptons be within the acceptance of the detector. Combining the results of the $\ell^+\ell^-\ell^+\ell^-$ analysis with those of the $\ell^+\ell^-\nu\bar{\nu}$ analysis partially compensates for the low rate of $ZZ \to \ell^+\ell^-\ell^+\ell^-$ events, due to the larger $ZZ \to \ell^+\ell^-\nu\bar{\nu}$ branching ratio, while conversely the poor mass resolution of the $\ell^+\ell^-\nu\bar{\nu}$ final state resulting from the presence of two neutrinos is compensated by the improved mass resolution of the fully reconstructable $\ell^+\ell^-\ell^+\ell^-$ final state, particularly at low resonance mass.

The additional Higgs boson (spin-0 resonance), denoted by H throughout this chapter, with the Standard Model Higgs boson denoted by h, is assumed to be produced mainly via gluon–gluon fusion (ggF) and vector-boson fusion (VBF) processes with the ratio of the two production mechanisms unknown in the absence of a specific model. The results are interpreted separately for the ggF and VBF production modes, with events being classified into ggF- and VBF-enriched categories in both final states, as discussed in Sections 7.3 and 7.4. The searches cover a wide mass range from 200 GeV up to 2000 GeV and look for an excess of events with respect to background predictions in the distribution of the four-lepton invariant mass $m_{4\ell}$ for the $\ell^+\ell^-\ell^+\ell^-$ final state, and the transverse mass m_{T} for the $\ell^+\ell^-\nu\bar{\nu}$ final state, as the escaping neutrinos do not allow the full reconstruction of the final state. The transverse mass is defined as:

$$m_{\mathrm{T}} \equiv \sqrt{\left[\sqrt{m_Z^2 + \left(p_{\mathrm{T}}^{\ell\ell}\right)^2} + \sqrt{m_Z^2 + \left(E_{\mathrm{T}}^{\mathrm{miss}}\right)^2}\right]^2 - \left|\vec{p}_{\mathrm{T}}^{\,\ell\ell} + \vec{E}_{\mathrm{T}}^{\mathrm{miss}}\right|^2}, \tag{7.1}$$

where m_Z is the mass of the Z boson, $\vec{p}_{\mathrm{T}}^{\,\ell\ell}$ and $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ are the transverse momentum of the lepton pair and the missing transverse momentum with magnitudes of $p_{\mathrm{T}}^{\ell\ell}$ and $E_{\mathrm{T}}^{\mathrm{miss}}$, respectively. The chosen mass range is based on the sensitivity of the analysis as determined by the selection criteria and the size of the data and simulation samples.

In the absence of such an excess, limits on the production rate of different signal hypotheses are obtained from a simultaneous likelihood fit in the two final states. The hypothesis of a heavy Higgs boson in the narrow-width approximation (NWA) is studied. The upper limits on the production rate of a heavy Higgs boson are also translated into exclusion contours in the context of selected benchmark scenarios of the two-Higgs-doublet models, described in Section 2.3.1. As several theoretical models favour non-negligible natural widths, large-width assumption (LWA) models, assuming widths of 1 %, 5 %, 10 % and 15 % of the resonance mass, are examined only for ggF production, which dominates in many scenarios over the next-largest contribution (VBF) in the search range. Results are also interpreted assuming the bulk Randall–Sundrum (RS) model with a warped extra dimension giving rise to a spin-2 Kaluza–Klein (KK) excitation of the graviton G_{KK}, described in Section 2.3.2. Finally, in order to limit potential implicit bias through the tuning of analysis parameters that favour a particular result, the signal region was blinded in the data sample until all analysis procedures were established.

This analysis was recently published in Ref. [215] and builds upon the previous publication based on the early Run 2 36.1 fb^{-1} proton–proton data sample recorded by ATLAS in the years 2015 and 2016 [93]. The main improvements relative to the previous search are the following: i) the full 139 fb^{-1} ATLAS Run 2 data sample is used; ii) both $\ell^+\ell^-\ell^+\ell^-$ and $\ell^+\ell^-\nu\bar{\nu}$ analyses profit from improved lepton reconstruction and isolation selection to mitigate the effects of pileup; iii) the reconstruction of jets uses a particle-flow algorithm that combines measurements from the tracker and the calorimeter for improved jet calibration; iv) the normalization of the non-resonant ZZ background is derived from data rather than being estimated from Standard Model predictions; v) event classification targeting different production processes is optimized using machine learning algorithms in the case of $\ell^+\ell^-\ell^+\ell^-$ final state; vi) the m_{T} distribution is used to search for signals in the VBF-enriched category in the case of the $\ell^+\ell^-\nu\bar{\nu}$ final state, in addition to the use of m_{T} in the ggF-enriched category; and vii) the search range is extended from 1200 GeV to 2000 GeV in signal mass.

The chapter is organized as follows. First, the data and simulated samples are described in Section 7.2. The analysis strategies for the $\ell^+\ell^-\ell^+\ell^-$ and $\ell^+\ell^-\nu\bar{\nu}$ final states are described in Sections 7.3 and 7.4, respectively. Section 7.5 describes the theoretical and experimental systematic uncertainties, Section 7.6 the final results, and Section 7.7 the interpretation of these results under the various models and benchmark scenarios.

7.2 Data and simulation

The proton–proton (pp) collision data used in these searches were collected by the ATLAS detector at a centre-of-mass energy of 13 TeV with a 25 ns bunch-spacing configuration from 2015 to 2018. The data are subjected to quality requirements, as described in Section 5.2. After these quality requirements, the total accumulated data sample corresponds to an integrated luminosity of 139 fb^{-1}, as discussed in previous chapters.

Simulated events are used to determine the signal acceptance and some background contributions. The events produced by each Monte Carlo (MC) event generator were processed through the ATLAS detector simulation within the GEANT4 framework, described in Section 4.3.2. Additional inelastic pp interactions were overlaid on the simulated signal and background events. The MC event generator used for pile-up is PYTHIA 8.186 [216] with the A2 set of tuned parameters [217] and the MSTW2008LO [218] parton distribution function (PDF) set. The simulated events are weighted to reproduce the observed distribution of the mean number of interactions per bunch crossing in data (*pile-up reweighting*).

Heavy spin-0 resonance production was simulated using the POWHEG BOX v2 [219] MC event generator. The gluon–gluon fusion and vector-boson fusion production modes were simulated separately, with matrix elements calculated to next-to-leading-order (NLO) accuracy in quantum chromodynamics (QCD). POWHEG BOX was interfaced to PYTHIA 8.212 [220] for parton showering and hadronization with the AZNLO set of tuned parameters [221], and for decaying the Higgs boson into the $H \to ZZ \to \ell^+\ell^-\ell^+\ell^-$ or $H \to ZZ \to \ell^+\ell^-\nu\bar{\nu}$ final states. The event generator was interfaced to the EVTGEN v1.2.0 program [222] for the simulation of bottom and charm hadron decays. The leading-order (LO) CT10 PDF set [223] was used for the hard-scattering process. Events from ggF and VBF production were generated in the resonance mass range of 300 GeV to 2000 GeV in the NWA, using a step size of 100 GeV up to 1000 GeV and 200 GeV above. For the $\ell^+\ell^-\ell^+\ell^-$ final state, due to the sensitivity of the analysis at lower masses, events were also generated for $m_H = 200$ GeV. In addition, events from ggF heavy Higgs production with a width of 15 % of the Higgs boson mass m_H were generated at NLO accuracy in QCD with MADGRAPH5_AMC@NLO v2.3.2 [224], which was interfaced to PYTHIA 8.210 for parton showering and hadronization with the A14 set of tuned parameters (A14 tune) [225], and for decaying the Higgs boson into the two leptonic final states. The properties of bottom and charm hadron decays were simulated by EVTGEN v1.2.0. Events were generated in the resonance mass range of 400 GeV to 2000 GeV using a step size of 100 GeV up to 1000 GeV, and a step size of 200 GeV above this point. Similarly, events with a width of 5% or 10% of $m_H = 900$ GeV were generated for validating the analytic parametrization of the $m_{4\ell}$ distribution used in the $\ell^+\ell^-\ell^+\ell^-$ final state as described in Section 7.3.3. For the $\ell^+\ell^-\nu\bar{\nu}$ final state, a reweighting procedure as described in Section 7.4.3 is used on fully simulated events to obtain the reconstructed m_{T} distribution at any value of mass and width tested.

Spin-2 Kaluza–Klein gravitons from the bulk Randall–Sundrum model [88, 226] were generated with MadGraph5_aMC@NLO at LO accuracy in QCD with the NNPDF2.3 LO PDF set with $\alpha_s = 0.130$ [227], which is then interfaced to Pythia 8.210 for parton showering and hadronization with the A14 tune and for decaying the heavy ZZ resonance into the two leptonic final states. The properties of bottom and charm hadron decays were simulated by EvtGen v1.2.0. The dimensionless coupling $k/\overline{M}_{\mathrm{Pl}}$, where $\overline{M}_{\mathrm{Pl}} = M_{\mathrm{Pl}}/\sqrt{8\pi}$ is the reduced Planck scale and k is the curvature scale of the extra dimension, is set to 1. The width of the resonance is correlated with the coupling $k/\overline{M}_{\mathrm{Pl}}$ and in this configuration it is around $\sim 6\%$ of its mass. Mass points between 600 GeV and 2000 GeV with 200 GeV spacing were generated for both final states.

The $q\bar{q} \to ZZ$ background was simulated by the Sherpa v2.2.2 [228] generator, in which the NNPDF3.0 NNLO PDF set [227] was used for the hard-scattering process, achieving NLO accuracy in the matrix-element calculation for 0- and 1-jet final states and LO accuracy for 2- and 3-jet final states with the Comix [229] and OpenLoops [230–232] matrix-element generators. The merging with the Sherpa parton shower [233] was performed using the MEPS@NLO prescription [234]. NLO electroweak (EW) corrections were applied as a function of $m_{4\ell}$ for the $\ell^+\ell^-\ell^+\ell^-$ final state [235, 236], and as a function of the transverse momentum of the Z boson that decays into two neutrinos for the $\ell^+\ell^-\nu\bar{\nu}$ final state [230, 237–240]. The EW production of a ZZ pair and two additional jets via vector-boson scattering up to $O(\alpha_{\mathrm{EW}}^6)$ was generated using Sherpa v2.2.2 for both the $\ell^+\ell^-\ell^+\ell^-$ and $\ell^+\ell^-\nu\bar{\nu}$ final states, where the process $ZZZ \to 4\ell qq$ is also taken into account. In addition, the WZ diboson events from both QCD and EW production, with the subsequent leptonic decays of both the W and Z bosons, were simulated by Sherpa with a similar set-up. The WZ events with Z boson decaying leptonically and W boson decaying hadronically were modelled with Sherpa v2.2.1.

The $gg \to ZZ$ process was modelled by Sherpa v2.2.2 at LO accuracy in QCD for both final states, including the off-shell SM h boson contribution and the interference between the h and ZZ processes. The higher-order correction factor accounting for up to NLO accuracy in QCD for the $gg \to ZZ$ continuum production was calculated for massless quark loops [241–243] in the heavy-top-quark approximation [244], including the $gg \to h^* \to ZZ$ process [245]. Based on these studies, a constant factor of 1.7 is used, and a relative uncertainty of 60 % is assigned to the normalization in both searches.

For the $\ell^+\ell^-\nu\bar{\nu}$ final state, the contribution from WW production was removed in the Sherpa simulation of the $q\bar{q} \to ZZ$ and $gg \to ZZ$ processes by requiring the charged leptons and the neutrinos to have different lepton flavours. The $q\bar{q} \to WW$ and $gg \to WW$ processes were then modelled with Powheg Box v2 and Sherpa v2.2.2, respectively. The interference between WW and ZZ production is expected to be negligible [238] and is therefore not considered.

Events containing a single Z boson with associated jets were simulated using the Sherpa v2.2.1 event generator. Matrix elements were calculated for up to two partons at NLO and four partons at LO using the Comix and OpenLoops matrix-element generators and merged with the Sherpa parton shower using the MEPS@NLO prescription. The NNPDF3.0 NNLO PDF set was used in conjunction with dedicated parton-shower tuning developed by the Sherpa authors. The Z+jets events are normalized using the NNLO cross sections [246].

The triboson backgrounds ZZZ, WZZ, and WWZ with fully leptonic decays and at least four prompt charged leptons were modelled using Sherpa v2.2.2 with LO accuracy of the QCD calculations and the CT10 PDF set. The simulation of $t\bar{t}V$ production ($V = W$ or Z) with both

top quarks decaying semileptonically and the vector boson decaying inclusively was performed with MadGraph5_aMC@NLO interfaced to Pythia 8.210 for parton showering and hadronization with the A14 tune and to EvtGen v1.2.0 for the simulation of bottom and charm hadron decays. The total cross section is normalized to the prediction of Ref. [247], which includes the two dominant terms at both LO and NLO in a mixed perturbative expansion in the QCD and EW couplings. The $t\bar{t}$ background, as well as single-top and Wt production, were modelled using Powheg Box v2 interfaced to Pythia 8.230 with the A14 tune and to EvtGen v1.6.0 for the simulation of bottom and charm hadron decays.

In order to study the interference effects for the LWA case, samples containing the $gg \to ZZ$ continuum background (B) as well as its interference (I) with a hypothetical heavy Higgs signal (S) were used and are referred to as *SBI* samples hereafter. In the $\ell^+\ell^-\ell^+\ell^-$ final state the MCFM NLO event generator [248], interfaced to Pythia 8.212, was used to produce *SBI* samples where the width of the heavy scalar is set to 15 % of its mass, for masses of 200, 300, 400, 500, 600, 800, 1000, 1200, and 1400 GeV. Background-only samples were also generated with the MCFM event generator, and are used to extract the signal-plus-interference term (*SI*) by subtracting them from the aforementioned *SBI* samples. For the $\ell^+\ell^-\nu\bar{\nu}$ final state, the *SBI* samples were generated with the gg2VV event generator [249, 250]. The samples include signal events with a scalar mass of 400, 700, 900, 1200, and 1500 GeV.

7.3 Analysis of the $\ell^+\ell^-\ell^+\ell^-$ final state

7.3.1 Event selection and categorization

This section first describes the four-lepton event selection, from the online triggers used, to the object-definition working points and finally to the event-level kinematic selections optimized to extract signal events while suppressing contamination from background processes. After this selection, events are further split into several categories, in order to probe different signal production modes, such as ggF and VBF production. To enhance the search sensitivity to the NWA signals, multivariate classifiers are optimized for the event categorization. In order to also obtain results that are more model-independent (since the training of the multivariate classifiers is usually based on a specific signal model), a cut-based event categorization that enhances the sensitivity in the VBF production mode is also considered.

In the search for LWA signals, due to the complexity of modelling the interference between heavy Higgs boson and SM Higgs boson processes, only the ggF-enriched categories of the cut-based analysis are used. The same strategy is adopted in the search for a Kaluza–Klein graviton excitation.

Common event selection

Four-lepton events are selected and initially classified according to the lepton flavours: 4μ, $4e$ and $2e2\mu$, called "channels" hereafter. They are selected using a combination of single-lepton, dilepton and trilepton triggers with different transverse momentum thresholds. The single-lepton triggers with the lowest p_{T} thresholds had tighter isolation requirements than the high p_{T} threshold single-lepton triggers and the multilepton triggers. Due to an increasing peak luminosity, these p_{T} thresholds increased during the data-taking periods [251, 252]. For single-muon triggers, the

p_{T} threshold increased from 20 GeV to 26 GeV, while for single-electron triggers, the p_{T} threshold increased from 24 GeV to 26 GeV. The overall trigger efficiency for signal events passing the final selection requirements is about 98 %.

In each channel, four-lepton candidates are formed by selecting a lepton-quadruplet made out of two same-flavour, opposite-sign lepton pairs. Both electrons and muons are identified using the corresponding "Loose" working points as described in Section 5.3. Each electron (muon) must satisfy $p_{\mathrm{T}} > 7$ (5) GeV and be measured in the pseudorapidity range of $|\eta| < 2.47$ (2.7). The highest-p_{T} lepton in the quadruplet must satisfy $p_{\mathrm{T}} > 20$ GeV, and the second (third) lepton in p_{T} order must satisfy $p_{\mathrm{T}} > 15$ (10) GeV. In the case of muons, at most one calorimeter-tagged, segment-tagged or stand-alone ($2.5 < |\eta| < 2.7$) muon is allowed per quadruplet.

If there is ambiguity in assigning leptons to a pair, only one quadruplet per channel is selected by keeping the quadruplet with the invariant mass of the lepton pairs closest (leading pair) and second closest (subleading pair) to the Z boson mass, with invariant masses referred to as m_{12} and m_{34} respectively. In the selected quadruplet, m_{12} must satisfy $50\,\mathrm{GeV} < m_{12} < 106\,\mathrm{GeV}$ and m_{34} must satisfy $50\,\mathrm{GeV} < m_{34} < 115\,\mathrm{GeV}$.

Selected quadruplets are required to have their leptons separated from each other by $\Delta R > 0.1$. For 4μ and $4e$ quadruplets, if an opposite-charge same-flavour lepton pair is found with $m_{\ell\ell}$ below 5 GeV, the quadruplet is removed to suppress the contamination from J/ψ mesons. If multiple quadruplets from different channels are selected at this point, only the quadruplet from the channel with the highest signal acceptance is retained, in the order: 4μ, $2e2\mu$, $4e$.

The Z+jets and $t\bar{t}$ background contributions are reduced by imposing impact-parameter requirements as well as track- and calorimeter-based isolation requirements on the leptons. The transverse impact-parameter significance, defined as the impact parameter calculated relative to the measured beam-line position in the transverse plane divided by its uncertainty, $|d_0|/\sigma_{d_0}$, for all muons (electrons), is required to be lower than 3 (5). The track-isolation discriminant is calculated from the tracks with $p_{\mathrm{T}} > 500$ MeV that lie within a cone of $\Delta R = 0.3$ around the muon or electron and that either originate from the primary vertex or have a longitudinal impact parameter z_0 satisfying $|z_0 \sin(\theta)| < 3$ mm if not associated with any vertex. Above a lepton p_{T} of 33 GeV, the cone size falls linearly with p_{T} to a minimum cone size of 0.2 at 50 GeV. Similarly, the calorimeter isolation is calculated from the positive-energy topological clusters that are not associated with a lepton track in a cone of $\Delta R = 0.2$ around the muon or electron. The sum of the track isolation and 40 % of the calorimeter isolation is required to be less than 16 % of the lepton p_{T}, as defined in the "FixedCutPFlowLoose" isolation working point. The calorimeter isolation is corrected for electron shower leakage, pileup, and underlying-event contributions. Both isolations are corrected for track and topological cluster contributions from the remaining three leptons. The pileup dependence of this isolation selection is reduced compared with that of the previous search by optimizing the criteria used for exclusion of tracks associated with a vertex other than the primary vertex and by the removal of topological clusters associated with tracks.

An additional requirement based on a vertex-reconstruction algorithm, which fits the four-lepton candidates with the constraint that they originate from a common vertex, is applied in order to further reduce the Z+jets and $t\bar{t}$ background contributions. A cut of $\chi^2/\mathrm{ndof} < 6$ for 4μ and < 9 for the other channels is applied, with an efficiency larger than 99 % for signal in all channels.

The QED process of radiative photon production in Z boson decays is well modelled by sim-

ulation. Some of the final-state-radiation (FSR) photons can be identified in the calorimeter and incorporated into the $\ell^+\ell^-\ell^+\ell^-$ analysis. The strategy to include FSR photons into the reconstruction of Z bosons is the same as in Run 1 [91]. It consists of a search for collinear (for muons) and non-collinear FSR photons (for muons and electrons) with only one FSR photon allowed per event. After the FSR correction, the four-momenta of both dilepton pairs are recomputed by means of a Z-mass-constrained kinematic fit [253]. The fit uses a Breit–Wigner Z boson lineshape and a single Gaussian function per lepton to model the momentum response function with the Gaussian width set to the expected resolution for each lepton. The Z-mass constraint is applied to both Z candidates.

Events that pass the common event selection (as described above) which are not yet split according to lepton flavours, form a category which is called "inclusive" hereafter. A comprehensive summary of all the cuts and requirements used in the event selection is given in Table 7.1.

Event categorization: multivariate analysis

In order to improve the sensitivity in the search for an NWA Higgs boson signal produced either in the VBF or the ggF production mode, two multivariate classifiers, namely a "VBF classifier" and a "ggF classifier", are used. These classifiers are built with deep neural networks (DNN) and use an architecture similar to that in Ref. [254], combining a multilayer perceptron (MLP) and one or two recurrent neural networks (rNN) implemented with the `Keras` [255] and `TensorFlow` [256] software packages. For both classifiers, the outputs of the MLP and rNN(s) are concatenated and fed into an additional MLP that produces an event score, denoted by $\mathrm{NN_{VBF}}$ and $\mathrm{NN_{ggF}}$, with $0 \leq \mathrm{NN_{VBF}}, \mathrm{NN_{ggF}} \leq 1$.

The VBF classifier uses two rNNs and an MLP. The two rNNs have as inputs the p_{T}-ordered transverse momenta and the pseudorapidities of the two leading jets and the transverse momenta and the pseudorapidities of the four leptons in the event. The MLP uses as inputs the invariant mass of the four-lepton system, the invariant mass and the transverse momentum of the two-leading-jets system, the difference in pseudorapidity between the $\ell^+\ell^-\ell^+\ell^-$ system and the leading jet, and the minimum angular separation between either of the $\ell^+\ell^-$ pairs and a jet.

The ggF classifier uses one rNN and an MLP. The rNN has as inputs the p_{T}-ordered transverse momenta and the pseudorapidities of the four leptons in the event. The MLP uses as inputs the following variables: (1) the four-lepton invariant mass; (2) the transverse momentum and the pseudorapidity of the four-lepton system; (3) the production angle of the leading Z defined in the four-lepton rest frame, $\cos\theta^*$; (4) the angle between the negative final-state lepton and the direction of flight of leading (subleading) Z in the Z rest frame, $\cos\theta_1$ ($\cos\theta_2$); (5) the angle between the decay planes of the Z bosons defined by four final-state leptons and expressed in the four-lepton rest frame, Φ; and (6) the transverse momentum and the pseudorapidity of the leading jet.

The two classifiers are trained separately using the above-listed discriminating variables on all simulated NWA signal events from their corresponding production mode, and the SM ZZ background events. The VBF classifier is trained on events with at least two jets while the ggF classifier is trained on events with fewer than two jets. In order to represent the relative importance of the signal and background events, weights that scale the events to the same luminosity according to their production cross sections are used in the training. Furthermore, in order to achieve good discriminating power of the classifiers over a large range of signal mass hypotheses, the signal events are reweighted such

Table 7.1: Summary of the physics object definitions and event selection requirements for the $\ell^+\ell^-\ell^+\ell^-$ analysis.

Physics Objects	
Electrons	
Loose Likelihood quality electrons with hit in innermost layer, $E_{\mathrm{T}} > 7\,\mathrm{GeV}$ and $\lvert\eta\rvert < 2.47$	
Interaction point constraint: $\lvert z_0 \sin\theta\rvert < 0.5\,\mathrm{mm}$ (if ID track is available)	
Muons	
Loose identification with $p_{\mathrm{T}} > 5\,\mathrm{GeV}$ and $\lvert\eta\rvert < 2.7$	
Calo-tagged muons with $p_{\mathrm{T}} > 15\,\mathrm{GeV}$ and $\lvert\eta\rvert < 0.1$, segment-tagged muons with $\lvert\eta\rvert < 0.1$	
Stand-alone and silicon-associated forward restricted to the $2.5 < \lvert\eta\rvert < 2.7$ region	
Combined, stand-alone (with ID hits if available) and segment-tagged muons with $p_{\mathrm{T}} > 5\,\mathrm{GeV}$	
Interaction point constraint: $\lvert d_0\rvert < 1\,\mathrm{mm}$ and $\lvert z_0 \sin\theta\rvert < 0.5\,\mathrm{mm}$ (if ID track is available)	
Jets	
anti-k_T jets with *bad-loose* identification, $p_{\mathrm{T}} > 30\,\mathrm{GeV}$ and $\lvert\eta\rvert < 4.5$	
Overlap removal	
Jets within $\Delta R < 0.2$ of an electron or $\Delta R < 0.1$ of a muon are removed	
Vertex	
At least one collision vertex with at least two associated track	
Primary vertex	
Vertex with the largest p_{T}^2 sum	
Event Selection	
Quadruplet Selection	- Require at least one quadruplet of leptons consisting of two pairs of same-flavour opposite-charge leptons fulfilling the following requirements: - p_{T} thresholds for three leading leptons in the quadruplet: 20, 15 and 10 GeV - Maximum one calo-tagged or stand-alone muon or silicon-associated forward per quadruplet - Leading di-lepton mass requirement: $50 < m_{12} < 106\,\mathrm{GeV}$ - Sub-leading di-lepton mass requirement: $50 < m_{34} < 115\,\mathrm{GeV}$ - $\Delta R(\ell, \ell') > 0.10$ for all leptons in the quadruplet - Remove quadruplet if alternative same-flavour opposite-charge di-lepton gives $m_{\ell\ell} < 5\,\mathrm{GeV}$ - Keep all quadruplets passing the above selection
Isolation	- Contribution from the other leptons of the quadruplet is subtracted - FixedCutPFlowLoose WP for all leptons
Impact Parameter Significance	- Apply impact parameter significance cut to all leptons of the quadruplet - For electrons: $d_0/\sigma_{d_0} < 5$ - For muons: $d_0/\sigma_{d_0} < 3$
Best Quadruplet	- If more than one quadruplet has been selected, choose the quadruplet with highest Higgs decay ME according to channel: 4μ, $2e2\mu$, $2\mu2e$ and $4e$
Vertex Selection	- Require a common vertex for the leptons: - $\chi^2/\mathrm{ndof} < 5$ for 4μ and < 9 for others decay channels

that their overall four-lepton invariant mass spectrum matches that of the SM background events. As a result of this reweighting method the classifiers do not produce a bias towards a specific mass point. Extensive checks are performed to ensure such treatment does not create a local excess of background events that would fake a signal. Figure 7.1 shows the ggF classifier and VBF classifier output for the data, the SM background and an example signal with $m_H = 600\,\mathrm{GeV}$. The signal distributions in this figure are normalized to an arbitrary value for visualization purposes, with the signal cross sections set to one hundred times the observed limit (see Section 7.7.1) for the ggF classifier and thirty times the observed limit for the VBF classifier.

After the common event selection, events with at least two jets ($N_{\mathrm{jets}} \geq 2$) and a VBF classifier score value greater than 0.8 form the VBF-MVA-enriched category. Events failing to enter the VBF-

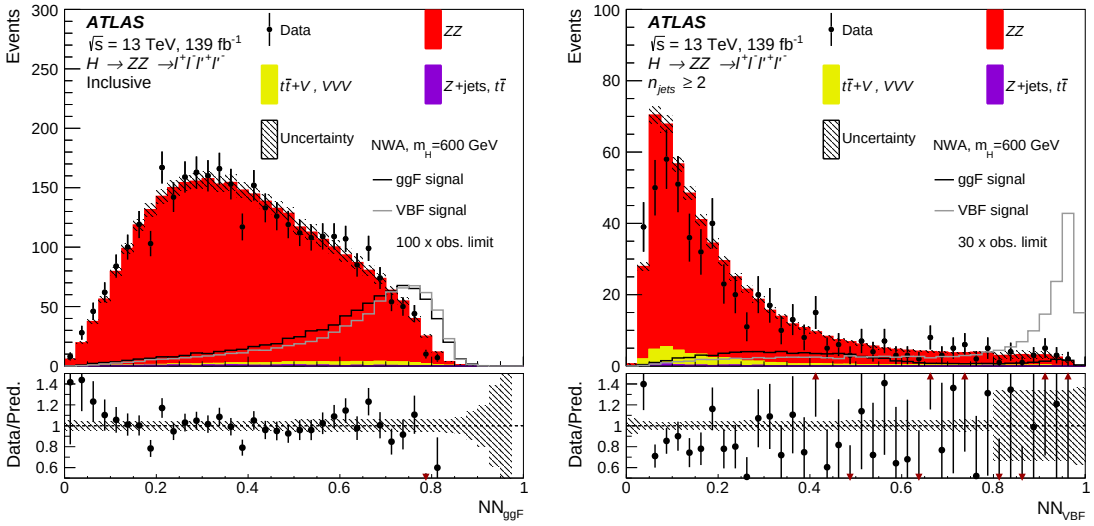

Figure 7.1: The output of (a) the ggF classifier ($\mathrm{NN_{ggF}}$) and (b) the VBF classifier ($\mathrm{NN_{VBF}}$) for the events passing the common event selections for the data, the SM background and NWA signal events with a resonance mass of 600 GeV. For (b) the VBF classifier output, an additional requirement of at least two jets in the event, is applied. The signal cross section is set to one hundred times the observed limit for the ggF classifier and thirty times the observed limit for the VBF classifier. The ZZ background is scaled by the normalization factors shown in Table 7.3. The lower panels show the ratio of data to prediction. Only statistical and experimental systematic uncertainties are included.

MVA-enriched category are classified into the ggF-MVA-high category if the ggF classifier score value is greater than 0.5; these events are further split into three distinct categories according to the lepton composition of the $\ell^+\ell^-\ell^+\ell^-$ system. Finally, events failing both classifiers form the ggF-low category. The optimal classification cuts were chosen on the basis of good overall performance across the entire mass range, while retaining a high signal selection efficiency and background rejection. Figure 7.2 shows the significance improvement with respect to the corresponding cut-based categories as a function of VBF and ggF signal mass m_H in the VBF and ggF categories for a number of different cuts on the classifier scores. The measure of significance used to quantify the improvement is given by

$$Z = \sqrt{2\left(n\ln\left[\frac{n(b+\sigma^2)}{b^2+n\sigma^2}\right] - \frac{b^2}{\sigma^2}\ln\left[1+\frac{\sigma^2(n-b)}{b(b+\sigma^2)}\right]\right)}, \tag{7.2}$$

where n is the number of observed events for $b \pm \sigma$ expected events [257].

Overall, five mutually exclusive categories are formed: VBF-MVA-enriched, ggF-MVA-high-4μ, ggF-MVA-high-2e2μ, ggF-MVA-high-4e and ggF-MVA-low, as shown in Figure 7.3. This categorization is used in the search for a heavy scalar with the NWA and in the search in the context of a CP-conserving 2HDM.

Finally, the signal acceptance, defined as the ratio of the number of reconstructed events after all selection requirements to the total number of simulated events, is found to be between 30 % (15 %) and 46 % (22 %) in the ggF (VBF)-enriched category for the ggF (VBF) production mode depending on the signal mass hypothesis. The signal acceptance is computed for each generated mass m_H and then fitted with a polynomial in m_H to interpolate between the generated mass points.

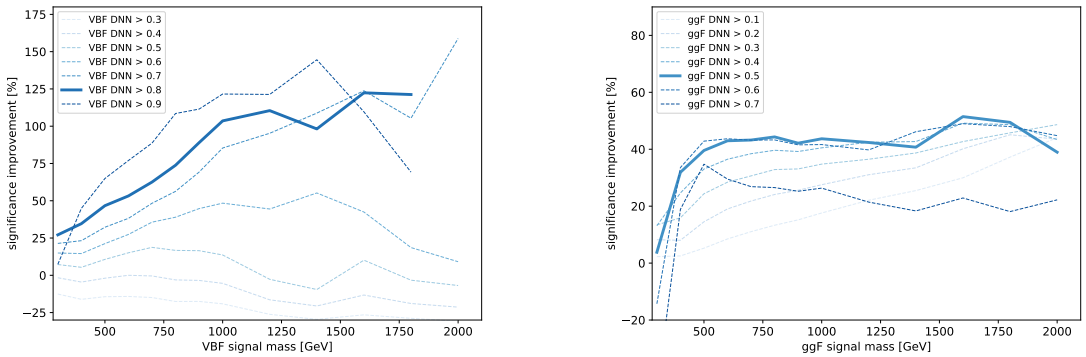

(a) **(b)**

Figure 7.2: Significance improvement of the DNN-based classification with respect to the cut-based definition of (a) the VBF-enriched category and (b) the ggF-enriched category for signal samples between $m_H = 300$ and 2000 GeV for various cuts on the DNN event score. The optimal cuts of 0.8 (VBF classifier) and 0.5 (ggF classifer) are represented by the solid lines, while other alternative cuts are plotted with dashed lines. Entries at 2000 GeV for cuts of 0.8 and 0.9 in the case of the VBF classifier are missing due to a lack of background events passing this selection.

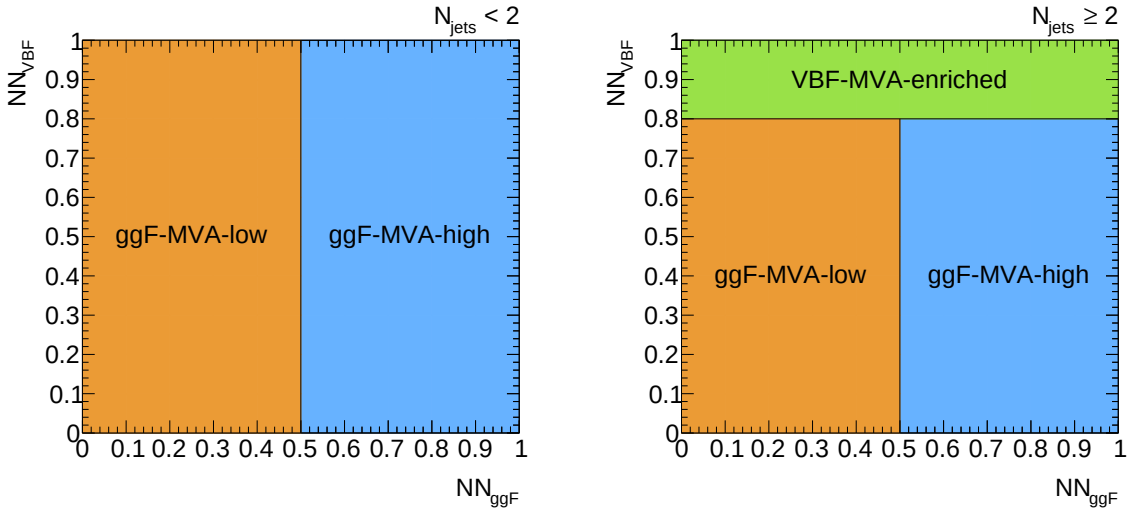

(a) **(b)**

Figure 7.3: Diagram showing the DNN-based VBF and ggF event categories given the VBF classifier score $\mathrm{NN_{VBF}}$, ggF classifier score $\mathrm{NN_{ggF}}$ and jet multiplicity N_{jets} for events with (a) $N_{\mathrm{jets}} < 2$ and (b) $N_{\mathrm{jets}} \geq 2$.

Event categorization: cut-based analysis

As in the previous publication [93], a cut-based analysis (CBA) is also performed to probe the sensitivity in the VBF production mode. If an event has two or more jets with $p_{\mathrm{T}} > 30\,\mathrm{GeV}$, with the two leading jets being well separated in η, $\Delta\eta_{\mathrm{jj}} > 3.3$, and having an invariant mass $m_{\mathrm{jj}} > 400\,\mathrm{GeV}$, this event is classified into the VBF-enriched category; otherwise the event is classified into one of the ggF-enriched categories further split according to the lepton flavour of the $\ell^+\ell^-\ell^+\ell^-$ system. Four distinct categories are formed, namely VBF-CBA-enriched, ggF-CBA-4μ, ggF-CBA-2e2μ and ggF-CBA-4e. The ggF-enriched categories are used in the search for a heavy large-width scalar and the search for a Kaluza–Klein graviton excitation. The cut-based analysis is also carried out in parallel with multivariate-based analysis for the search for a heavy scalar with the NWA; the corresponding results are described in Appendix C.

7.3.2 Background estimation

The main background source in the $\ell^+\ell^-\ell^+\ell^-$ final state is non-resonant SM ZZ production, accounting for 97 % of the total background events in the inclusive category. It arises from quark–antiquark annihilation $q\bar{q} \to ZZ$ (86 %), gluon-initiated production $gg \to ZZ$ (10 %), and a small contribution from EW vector-boson scattering (1 %). The last of these is more important in the VBF-enriched category using the DNN-based categorization, where it accounts for 20 % of the total background events. While in the previous publication [93] the SM ZZ background was exclusively estimated from simulation for both the shape and the normalization, in this analysis its normalization is derived from the data in the likelihood fit used in the statistical treatment of the data as explained in Section 7.6. The shapes of the $q\bar{q} \to ZZ$ and $gg \to ZZ$ invariant mass distributions are parameterized with analytic functions as described in Section 7.3.3. Additional background comes from the Z+jets and $t\bar{t}$ processes. These contribute to the total background yields at the percent level and decrease more rapidly than the non-resonant ZZ contribution as a function of $m_{4\ell}$. These backgrounds are estimated using data where possible, following slightly different approaches for final states with a dimuon ($\ell\ell + \mu\mu$) or a dielectron ($\ell\ell + ee$) subleading pair [258, 259].

The $\ell\ell + \mu\mu$ non-ZZ background comprises mostly $t\bar{t}$ and Z+jets events, where in the latter case the muons arise mostly from heavy-flavour semileptonic decays and to a lesser extent from π/K in-flight decays. The normalizations of the Z+jets and $t\bar{t}$ backgrounds are determined by fitting the invariant mass of the leading lepton pair in dedicated data control regions. The control regions are formed by relaxing the χ^2 requirement on the four-lepton vertex fit, and by inverting and relaxing isolation and/or impact-parameter requirements on the subleading muon pair. An additional control region ($e\mu\mu\mu$) is used to improve the $t\bar{t}$ background estimate. Transfer factors, defined as the number of events in the signal region divided by the number of events in the control region, are obtained separately for $t\bar{t}$ and Z+jets using simulated events and are used to extrapolate the yields from the control regions to the signal regions.

The main non-prompt background for the $\ell\ell + ee$ process arises from three sources: light-flavour jets misidentified as electrons; photon conversions; and semileptonic decays of heavy-flavour hadrons. The $\ell\ell + ee$ control-region selection requires the electrons in the subleading lepton pair to have the same charge, and relaxes the identification and isolation requirements on the electron candidate, denoted X, with the lower transverse momentum. The heavy-flavour background is found to be neg-

ligible, whereas the light-flavour and photon-conversion background is obtained with the sPlot [260] method, based on a fit to the number of hits in the innermost ID layer in the data control region. Transfer factors for the light-flavour jets and converted photons, obtained from simulated samples, are corrected using a $Z + X$ control region and then used to extrapolate the extracted yields to the signal region. Both the yield extraction and the extrapolation are performed in bins of the transverse momentum of the electron candidate and the jet multiplicity.

The WZ production process is included in the data-driven estimates for the $\ell\ell + ee$ final states, while it is added from simulation for the $\ell\ell + \mu\mu$ final states even though its contribution to the total background is at the per-mille level. The contributions from $t\bar{t}V$ (where V stands for either a W or a Z boson) and triboson processes are minor and taken from simulated samples.

7.3.3 Signal and background modelling

The reconstructed four-lepton invariant mass $m_{4\ell}$ distribution is used as the discriminating variable for the $\ell^+\ell^-\ell^+\ell^-$ final state. It is extracted from simulation for signal events and for most background components ($t\bar{t}V$, VVV, $\ell\ell + \mu\mu$ and heavy-flavour hadron component of $\ell\ell + ee$), except for the light-flavour jets and photon conversions in the case of $\ell\ell + ee$ background, which are taken from the control region as described in Section 7.3.2. To obtain statistical interpretations for each mass hypothesis, the $m_{4\ell}$ distribution for signal is parameterized as a function of the mass hypothesis m_H.

Narrow-width scalar signal model

In the case of a narrow resonance, the width in $m_{4\ell}$ is determined by the detector resolution, which is modelled by the sum of a Crystal Ball ($\mathcal{C}$) function, consisting of a Gaussian core and a power-law low-end tail [261, 262], and a Gaussian ($\mathcal{G}$) function:

$$P_s(m_{4\ell}) = f_{\mathcal{C}} \times \mathcal{C}(m_{4\ell}; \mu, \sigma_{\mathcal{C}}, \alpha_{\mathcal{C}}, n_{\mathcal{C}}) + (1 - f_{\mathcal{C}}) \times \mathcal{G}(m_{4\ell}; \mu, \sigma_{\mathcal{G}}). \tag{7.3}$$

The Crystal Ball and Gaussian functions share the same peak value of $m_{4\ell}$ (μ), but have different resolution parameters, $\sigma_{\mathcal{C}}$ and $\sigma_{\mathcal{G}}$. The $\alpha_{\mathcal{C}}$ and $n_{\mathcal{C}}$ parameters control the shape and position of the non-Gaussian tail, and the parameter $f_{\mathcal{C}}$ ensures the relative normalization of the two probability density functions. To improve the stability of the parameterization in the full mass range considered, the parameter $n_{\mathcal{C}}$ is set to a fixed value. Examples of the simulated $m_{4\ell}$ distributions for the 4μ, $4e$ and $2e2\mu$ final states and the fitted signal model for resonance mass hypotheses of $m_H = 300$, 1000 and 1800 GeV are shown in Figure 7.4 The bias in the extraction of signal yields introduced by using the analytic function is below 2 % and treated as a systematic uncertainty of the signal parameterization. The function parameters are determined separately for each final state using the simulated events for each generated mass m_H, and then fitted with a polynomial in m_H to interpolate between the generated mass points; examples of these polynomial fits for the $2e2\mu$ case are given in Figure 7.5. The order of the polynomial is determined by first fitting a third-order polynomial and decreasing its order until the magnitude of the highest-order term is greater than its associated error (to avoid over-fitting). The use of this parameterization for the function parameters introduces a bias in the signal yield and m_H extraction of about 1 %. The extra bias is included in the systematic uncertainties of the signal acceptance.

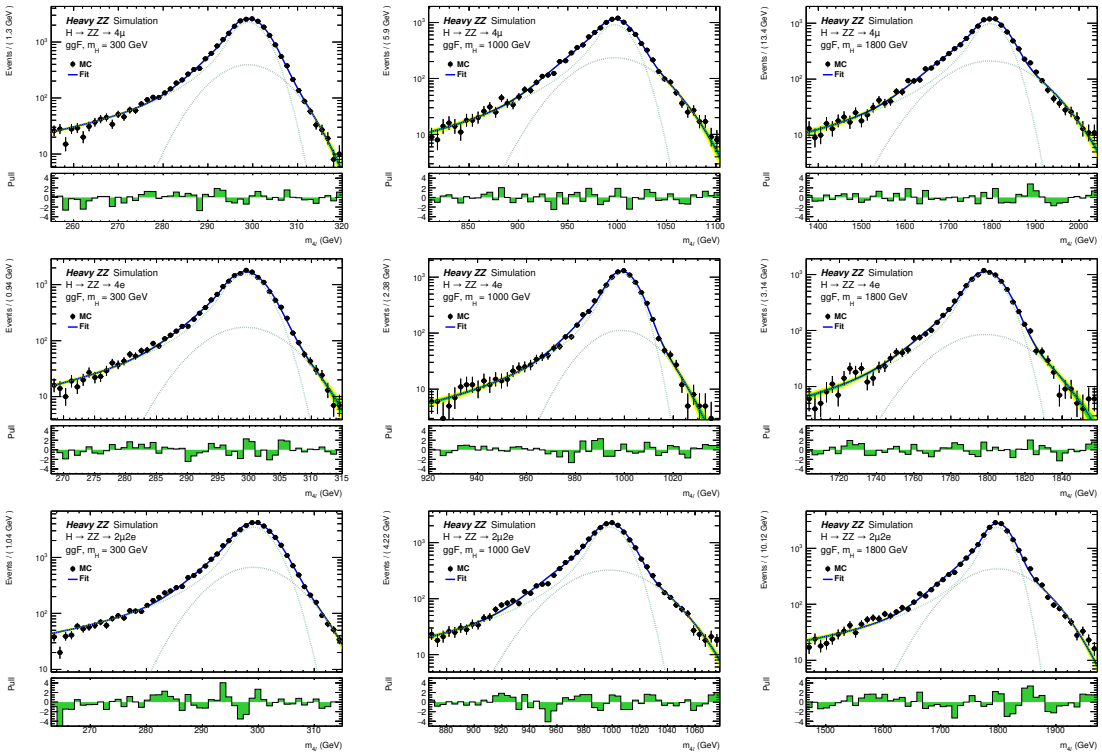

Figure 7.4: Distributions of the simulated four-lepton invariant mass $m_{4\ell}$ for resonance mass hypotheses of $m_H = 300$, 1000 and 1800 GeV in the NWA. The top, middle and bottom rows show the 4μ, $4e$ and $2e2\mu$ final states, respectively. The markers show the simulated $m_{4\ell}$ distribution for these specific values of m_H and the solid blue line shows the signal model fitted to this distribution. The dotted green lines show the Crystal Ball and Gaussian components of the fitted signal model and the green and yellow bands represent the $\pm1\sigma$ and $\pm2\sigma$ uncertainties from the fit. The lower panel in each plot shows the pull in each bin.

Large-width scalar signal model

In the case of a large-width resonance, the reconstructed $m_{4\ell}$ distribution is modelled as the convolution of the parton-level lineshape and the detector resolution function. The parton-level lineshape is derived from a theoretical calculation, described below, and multiplied by the signal acceptance obtained from simulation. The detector resolution is obtained from the narrow-width signal model described above, since the effect of the parton-level lineshape for a narrow resonance on the reconstructed $m_{4\ell}$ distribution is negligible and is therefore dominated by detector effects. Although the detector resolution varies as a function of the resonance mass, these effects are assumed to be constant under the peak of the parton-level distribution. For simplicity, the gluon–gluon fusion production mode is assumed to be dominant, and therefore the other production modes are neglected.

The parton cross section for a heavy Higgs boson produced via the collision of incident partons i, j and decaying into a generic final state F can be written as in Eq. (21) of Ref. [263]:

$$\sigma_{ij\to H\to F}(\hat{s}) = \frac{1}{2\hat{s}} \int d\Omega \left[\sum_{s,c} |A_{ij\to H}(\hat{s},\Omega)|^2\right] \frac{1}{|\hat{s} - s_H|^2} \left[\sum_{s,c} |A_{H\to F}(\hat{s},\Omega)|^2\right], \tag{7.4}$$

where $\hat{s}$ is the square of the centre-of-mass parton collision energy, $A_{ij\to H}(\hat{s},\Omega)$ and $A_{H\to F}(\hat{s},\Omega)$ are the corresponding Higgs boson production and decay amplitudes, respectively, averaged (summed)

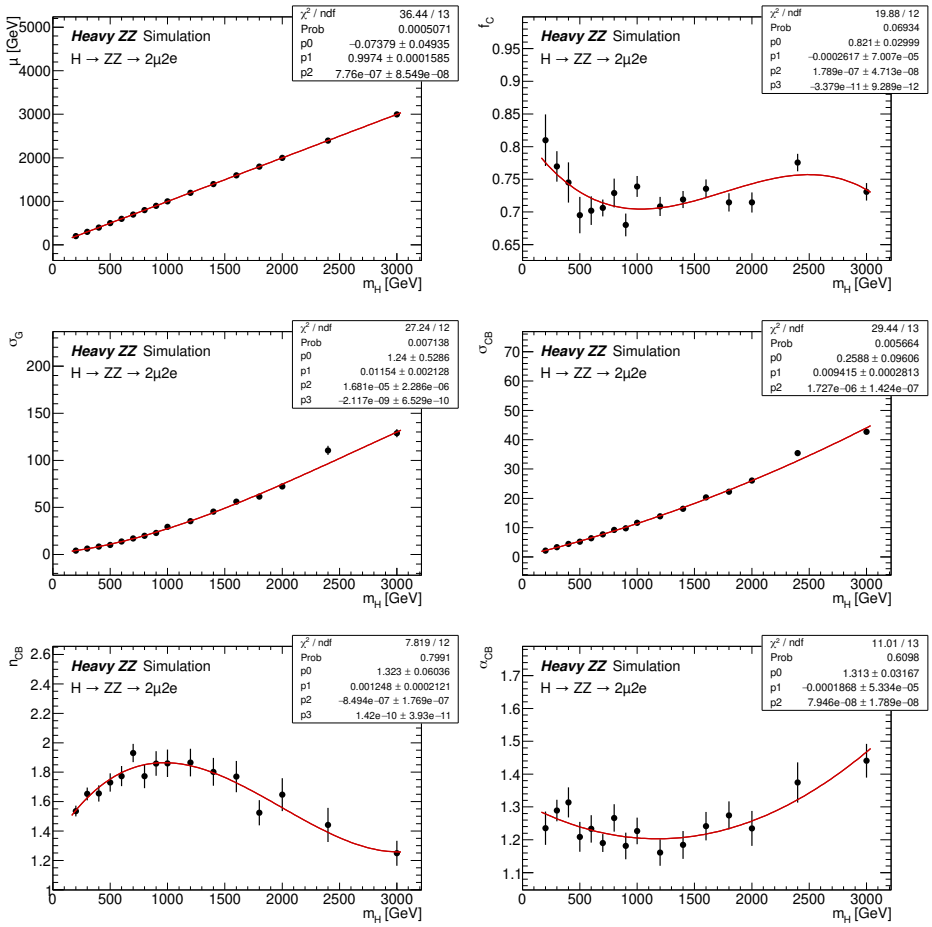

Figure 7.5: Interpolation of the parameters μ, f_C, σ_G, σ_C, n_C and α_C of the signal model in the $2e2\mu$ channel as a function of the resonance mass m_H in the NWA.

over spin and colours in the initial (final) state, $\frac{1}{\hat{s}-s_H}$ is a Higgs propagator, and Ω is the phase space of the process over which the amplitudes are integrated.

Using the definition of the partial width of a Higgs boson decaying to a final state F,

$$\Gamma_{H\to F}(\hat{s}) = \frac{1}{2\sqrt{\hat{s}}} \int d\Omega' \sum_{s,c} |A_{H\to F}(\hat{s},\Omega')|^2 \,, \tag{7.5}$$

both the production and decay components in Eq. (7.4) can be substituted with the corresponding partial decay widths $\Gamma_{H\to ij}$ and $\Gamma_{H\to F}$, since the amplitude is invariant under time inversion, modulo a normalization constant N. Therefore, the parton cross section for a heavy Higgs boson produced via gluon–gluon fusion decaying into a pair of Z bosons can be written as

$$\sigma_{gg\to H\to ZZ}(\hat{s}) = 2N \frac{1}{|\hat{s}-s_H|^2} \Gamma_{H\to gg}(\hat{s}) \Gamma_{H\to ZZ}(\hat{s}) \,. \tag{7.6}$$

According to the signal definition, the hypothetical heavy Higgs boson has the same partial widths as the Standard Model Higgs boson, such that $\Gamma_{H\to F} = \Gamma_{h\to F}$. Therefore, one can use the analytical expressions for $\Gamma_{H\to gg}(\hat{s})$ and $\Gamma_{H\to ZZ}(\hat{s})$ computed in Refs. [263, 264], with the

appropriate kinematic substitution $m_H^2 \to \hat{s}$:

$$\Gamma_{H\to gg}(\hat{s}) = \frac{G_F \alpha_s^2}{36\sqrt{2}\pi^3}\hat{s}^{\frac{3}{2}}\left|\frac{3}{4}\sum_q A_q(\tau_q)\right|^2 , \tag{7.7}$$

where

$$\begin{aligned} A_q(\tau) &= 2\frac{\tau + (\tau-1)f(\tau)}{\tau^2} , \\ \tau_q &= \frac{\hat{s}}{4m_q^2} , \\ f(\tau) &= \begin{cases} \arcsin^2(\sqrt{\tau}) , & \tau \le 1 \\ -\frac{1}{4}\left|\log\left(\frac{1+\sqrt{1-\tau^{-1}}}{1-\sqrt{1-\tau^{-1}}}\right) - i\pi\right|^2 , & \tau > 1 \end{cases} , \end{aligned}$$

and

$$\Gamma_{H\to ZZ}(\hat{s}) = \frac{G_F}{16\sqrt{2}\pi}\hat{s}^{\frac{3}{2}}\sqrt{1-4x}(1-4x+12x^2) , \quad \text{with} \quad x = \frac{m_Z^2}{\hat{s}} . \tag{7.8}$$

The top quark contribution is dominant in the sum over quark loops in Eq. (7.7), therefore the contributions from all other quarks are neglected. Finally, the scalar propagator can be written according to the Bar-schema as described in Ref. [263]:

$$\frac{1}{\hat{s}-s_H} = \frac{1 + i\frac{\overline{\Gamma}_H}{\overline{m}_H}}{\hat{s} - \overline{m}_H^2 + i\frac{\overline{\Gamma}_H}{\overline{m}_H}\hat{s}} , \tag{7.9}$$

where

$$\begin{aligned} \overline{m}_H &= \sqrt{\Gamma_H^2 + m_H^2} , \\ \overline{\Gamma}_H &= \overline{m}_H \frac{\Gamma_H}{m_H} , \end{aligned}$$

and where m_H and Γ_H are the assumed mass and total width of the heavy Higgs boson, respectively. Combing Eqs. (7.6)–(7.9) provides a fully analytical expression for the $gg \to H \to ZZ$ parton cross section as a function of the centre-of-mass parton collision energy, mass and total width of the heavy Higgs boson. The normalization constant N is arbitrary since the signal normalization is determined in the fit to data.

For pp collisions, the final parton-level $m_{4\ell}$ lineshape must account for the parton luminosity, which gives the probability of parton interactions in hadron–hadron collisions, as described in Ref. [265]. Accounting for the gluon–gluon parton luminosity $\mathcal{L}_{gg}$ in pp collisions at a centre-of-mass energy of $\sqrt{s} = 13\,\text{TeV}$, and making the appropriate kinematic substitution $s \to m_{4\ell}^2$, the parton-level cross section becomes

$$\sigma_{pp\to H\to ZZ}(m_{4\ell}) = 2N m_{4\ell}\mathcal{L}_{gg}\frac{1}{|m_{4\ell} - s_H|^2}\Gamma_{H\to gg}(m_{4\ell}^2)\Gamma_{H\to ZZ}(m_{4\ell}^2) . \tag{7.10}$$

Comparisons of the analytical lineshape to the simulated parton-level $m_{4\ell}$ distribution produced using the GG2VV MC generator for various resonance masses and widths are shown in Figure 7.6.

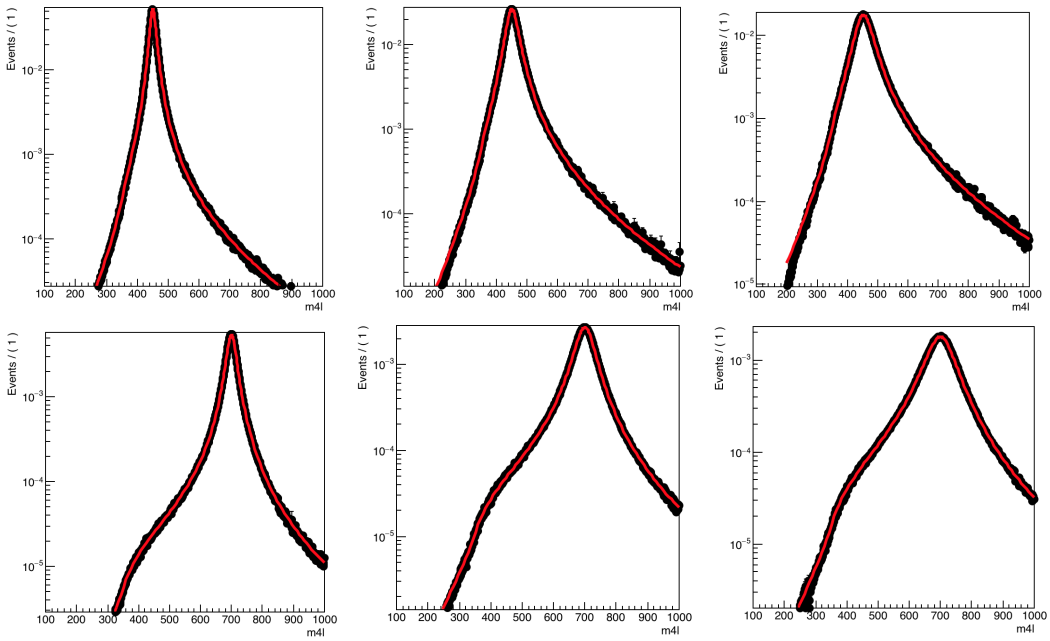

Figure 7.6: Comparisons of the analytical lineshape to the simulated parton-level $m_{4\ell}$ distribution in the GG2VV MC samples for $m_H = 400\,\text{GeV}$ (top row) and 700 GeV (bottom row), and widths equal to 5 % (left column), 10 % (middle column) and 15 % (right column) of the resonance mass.

These comparisons show generally excellent agreement between the analytical lineshape and simulation and serve to verify the validity of the analytical model.

Graviton signal model

The reconstructed RS graviton $m_{4\ell}$ lineshape is determined in a similar manner as for the large-width resonance by convolving the parton-level lineshape with the detector-resolution function. First, the parton-level lineshape is modelled as the product of a relativistic Breit-Wigner (RBW) term, a parton luminosity term $\mathcal{L}$, and a term corresponding to the squared matrix element of the production process as given in Ref. [266]. For example, the $gg \to G^* \to ZZ \to 4\ell$ parton-level lineshape takes the form,

$$m_{4\ell}^{\text{parton}} \sim \mathcal{L}_{gg} \cdot s^2 \cdot \frac{s(1+s)(1+2s+2s^2)}{(s^2 - m_G^2)^2 + m_G^2\Gamma^2}\,, \tag{7.11}$$

where m_G is the graviton mass and Γ the resonance width. Second, the detector resolution is obtained from the narrow-width signal model described above, convolved with the parton-level lineshape, and multiplied by the signal acceptance obtained from the simulated events.

The parton-level $m_{4\ell}$ distributions are presented in Figure 7.7 for graviton masses of $m_G = 600$, 1200 and 1800 GeV. The parton-level graviton signal model is fit to the simulated distribution with the mass m_G and width Γ parameters floating. The final parton-level signal model is then parameterized by fitting a linear function to each of the m_G and Γ parameters as a function of m_G as shown in Figure 7.8.

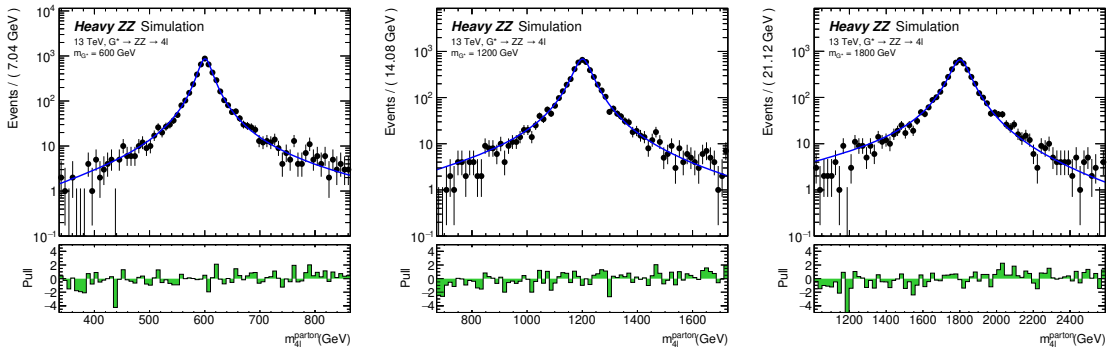

Figure 7.7: Distributions of the simulated parton-level four-lepton invariant mass $m_{4\ell}$ for graviton resonance mass hypotheses of $m_G = 600$, 1200 and 1800 GeV. The markers show the simulated $m_{4\ell}$ distribution for these specific values of m_G and the solid blue line shows the signal model fitted to this distribution. The lower panel in each plot shows the pull in each bin.

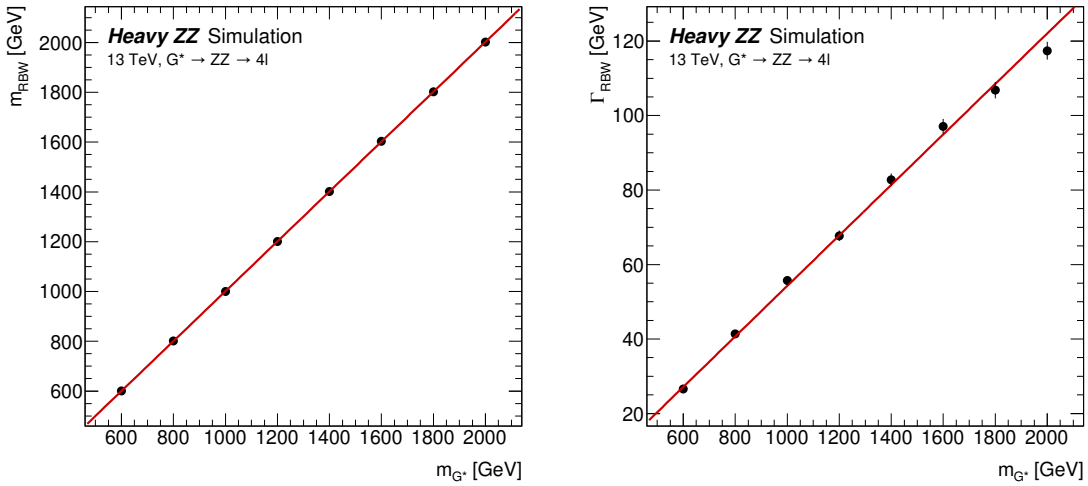

Figure 7.8: Interpolation of the fitted parameters of the graviton signal model, m_{RBW} and Γ_{RBW}, as a function of the graviton resonance mass m_G.

ZZ continuum background model

For the ZZ continuum background, the $m_{4\ell}$ distribution is parameterized by an empirical function for both the quark- and gluon-initiated processes in order to reduce the statistical uncertainties stemming from the limited number of simulated events. The empirical function is described by the following:

$$f_{qqZZ/ggZZ}(m_{4\ell}) = C_0 \times H(m_0 - m_{4\ell}) \times f_1(m_{4\ell}) + H(m_{4\ell} - m_0) \times f_2(m_{4\ell}), \tag{7.12}$$

where,

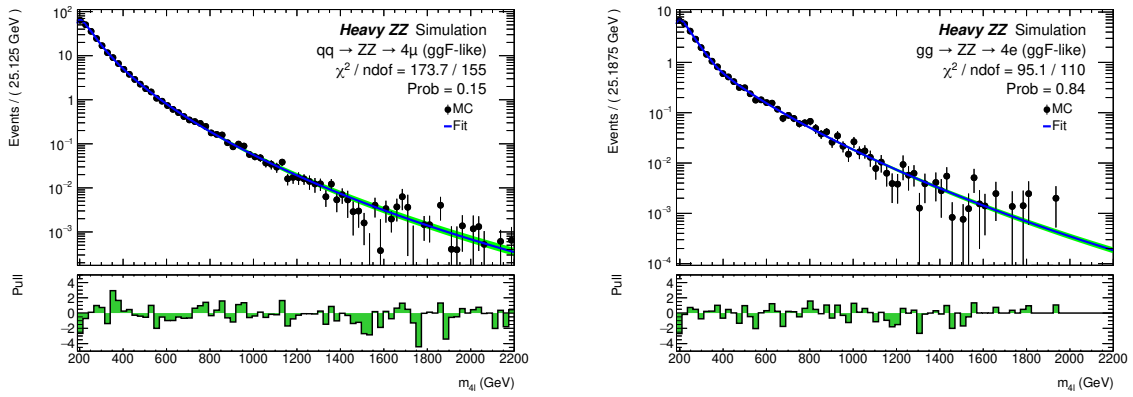

(a) **(b)**

Figure 7.9: Examples of the ZZ continuum background modelling. The markers show the simulated $m_{4\ell}$ distribution for (a) the $q\bar{q} \to ZZ$ process for the 4μ final state and (b) the $gg \to ZZ$ process for the $4e$ final state, and the solid blue line shows the background model fitted to these distributions. The green band represents the fit uncertainty computed using a sampling procedure that captures 68 % of the variations in the fitted parameters.

$$f_1(m_{4\ell}) = \left(\frac{m_{4\ell} - a_4}{a_3}\right)^{a_1-1}\left(1 + \frac{m_{4\ell} - a_4}{a_3}\right)^{-a_1-a_2},$$
$$f_2(m_{4\ell}) = \exp\left[b_0\left(\frac{m_{4\ell} - b_4}{b_3}\right)^{b_1-1}\left(1 + \frac{m_{4\ell} - b_4}{b_3}\right)^{-b_1-b_2}\right],$$
$$C_0 = \frac{f_2(m_0)}{f_1(m_0)}.$$

The function's first part, f_1, covers the low-mass part of the spectrum until the ZZ threshold around $2m_Z$, and the second part, f_2, describes the high-mass tail. The transition between low- and high-mass parts is modelled with the Heaviside step function $H(x)$ around $m_0 = 260\,\text{GeV}$ for $q\bar{q} \to ZZ$ and around 350 GeV for $gg \to ZZ$. The continuity of the function around m_0 is ensured by the normalization factor C_0 that is applied to the low-mass part. Finally, a_i and b_i are shape parameters which are obtained by fitting the $m_{4\ell}$ distribution in simulation for each category. Examples of the simulated $m_{4\ell}$ distributions and the model functions fitted to them are given in Figure 7.9. A large number of $m_{4\ell}$ distributions are calculated from the analytic function with variations of the a_i and b_i values sampled from a multivariate Gaussian distribution that is constructed from their covariance matrix. The uncertainty in the $m_{4\ell}$ distribution is determined by calculating a central interval that captures 68 % of the variations, and is treated as a nuisance parameter in the likelihood fit, namely a ZZ parameterization uncertainty. The ZZ parameterization uncertainty is one of the leading systematic uncertainties for a low-mass signal, as shown in Table 7.2.

Interference modelling

The gluon-initiated production of a heavy scalar H, the SM Higgs h and the $gg \to ZZ$ continuum background all share the same initial and final state, and thus lead to interference terms in the total amplitude. The corresponding Feynman diagrams for these processes are shown in Figure 7.10. Theoretical calculations described in Ref. [267] have shown that the effect of interference could

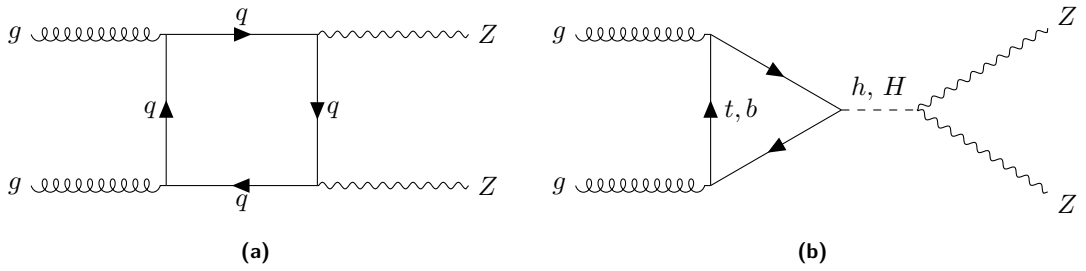

Figure 7.10: Tree-level Feynman diagrams of (a) gluon-initiated ZZ production ($gg \to ZZ$) and (b) gluon–gluon fusion production of a Standard Model or heavy Higgs boson decaying to a pair of Z bosons.

modify the integrated cross section by up to $O(10\,\%)$, and this effect is enhanced as the width of the heavy scalar increases. Therefore, a search for a heavy scalar Higgs boson in the LWA case must properly account for two interference effects: the interference between the heavy scalar and the SM Higgs boson (denoted by H–h) and between the heavy scalar and the $gg \to ZZ$ continuum (denoted by H–B). However, because the width of the KK excitation resonance is relatively small, the interference effect is assumed to be negligible in the graviton interpretation for both final states.

If the H and h bosons have similar properties, they have the same production and decay amplitudes and therefore the only difference between the signal and interference terms in the production cross section comes from the propagator. Hence, the acceptance and resolution of the signal and interference terms are expected to be the same. The H–h interference is obtained by reweighting the particle-level lineshape of generated signal events using the following formula:

$$w(m_{4\ell}) = \frac{2 \cdot Re\left[\frac{1}{s-s_H} \cdot \frac{1}{(s-s_h)^*}\right]}{\frac{1}{|s-s_H|^2}}, \tag{7.13}$$

where $1/\left(s - s_{H(h)}\right)$ is the propagator for a scalar (H or h). The particle-level lineshape is then convolved with the detector resolution function, and the signal and interference acceptances are assumed to be the same.

In order to extract the H–B interference contribution, signal-only and background-only samples are subtracted from the generated *SBI* samples. The extracted particle-level $m_{4\ell}$ distribution for the H–B interference term is then convolved with the detector resolution.

7.4 Analysis of the $\ell^+\ell^-\nu\bar{\nu}$ final state

7.4.1 Event selection and categorization

The $\ell^+\ell^-\nu\bar{\nu}$ final state consists of a pair of high-p_{T} isolated leptons (electrons or muons) and large $E_{\mathrm{T}}^{\mathrm{miss}}$, and is subject to larger background contamination than the $\ell^+\ell^-\ell^+\ell^-$ channel. Candidate events are recorded with a combination of multiple single-lepton triggers, which gives a high efficiency of about 98 % for typical signal processes in the signal region defined in the following.

Candidate events are preselected by requiring exactly two electrons or muons with opposite charges and $p_{\mathrm{T}} > 20\,\mathrm{GeV}$, where the electrons (muons) must have $|\eta| < 2.47$ (2.5). Both electrons and muons are identified using the corresponding "Medium" working points as described in Sec-

tion 5.3. The leading lepton is further required to have $p_{\mathrm{T}} > 30\,\mathrm{GeV}$, well above the threshold of the single-lepton triggers. The selected electrons or muons must have a longitudinal impact parameter satisfying $|z_0 \sin(\theta)| < 0.5\,\mathrm{mm}$. The lepton candidates are required to satisfy the same isolation criteria and the same requirement on the transverse impact-parameter significance as used in the $\ell^+\ell^-\ell^+\ell^-$ channel (see Section 7.3.1), which leads to an efficiency above 98 % for typical prompt leptons with $p_{\mathrm{T}} > 30\,\mathrm{GeV}$. To suppress the WZ background, events containing any additional lepton satisfying the "Loose" identification requirement with $p_{\mathrm{T}} > 7\,\mathrm{GeV}$, in addition to the other requirements, are rejected. Requiring the dilepton invariant mass ($m_{\ell\ell}$) to be in the range between 76 and 106 GeV largely reduces the contamination from the non-resonant-$\ell\ell$ background, originating from $t\bar{t}$, Wt, WW, and $Z \to \tau\tau$ production. The data sample after the preselection is dominated by the Z+jets and non-resonant-$\ell\ell$ processes. To suppress these backgrounds, a further selection based on $E_{\mathrm{T}}^{\mathrm{miss}}$ and event topology is applied.

Candidate events are required to have $E_{\mathrm{T}}^{\mathrm{miss}} > 120\,\mathrm{GeV}$, which suppresses the Z+jets contamination by several orders of magnitude. The number of residual Z+jets events, which have large fake $E_{\mathrm{T}}^{\mathrm{miss}}$, is further reduced by requiring $\mathrm{S}(E_{\mathrm{T}}^{\mathrm{miss}}) > 10$, where $\mathrm{S}(E_{\mathrm{T}}^{\mathrm{miss}})$ is the statistical significance of the $E_{\mathrm{T}}^{\mathrm{miss}}$ value against the null hypothesis of zero-$E_{\mathrm{T}}^{\mathrm{miss}}$ [268]. Additional selection criteria based on angular variables are imposed to further reject the Z+jets and non-resonant-$\ell\ell$ background events. The selection on angular variables is motivated by the desired detector signature, where the $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ is back-to-back with the transverse momentum of the dilepton system. The azimuthal angle difference between the dilepton system and $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$, $\Delta\phi(\vec{p}_{\mathrm{T}}^{\,\ell\ell}, \vec{E}_{\mathrm{T}}^{\mathrm{miss}})$, must be larger than 2.5 radians, and the selected leptons must be close to each other, with the distance $\Delta R_{\ell\ell} = \sqrt{(\Delta\phi_{\ell\ell})^2 + (\Delta\eta_{\ell\ell})^2} < 1.8$. Furthermore, the azimuthal angle difference between any of the selected jets with $p_{\mathrm{T}} > 100\,\mathrm{GeV}$ and $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ must be larger than 0.4 radians. As a consequence of all the requirements, the Z+jets process only constitutes a small fraction of the total background (about 4 %) after the full selection. Finally, events containing one or more b-jets are vetoed to further suppress the $t\bar{t}$ and Wt backgrounds.

The signal region for the VBF production mode (VBF-enriched signal region) is defined for candidate events containing at least two selected jets with $p_{\mathrm{T}} > 30\,\mathrm{GeV}$, where the two leading jets must have $m_{\mathrm{jj}} > 550\,\mathrm{GeV}$ and $\Delta\eta_{\mathrm{jj}} > 4.4$. The remaining events, failing the requirements for the VBF-enriched signal region, are categorized for the ggF-enriched signal region. The signal acceptance in the ggF-enriched signal region for signal events containing a heavy spin-0 resonance from ggF production is about 30 % at $m_H = 400\,\mathrm{GeV}$ and up to 50 % at $m_H = 1.4\,\mathrm{TeV}$. For VBF signal events the signal acceptance in the VBF-enriched signal region is generally lower, ranging from 3 % at $m_H = 400\,\mathrm{GeV}$ to 20 % at $m_H = 1.6\,\mathrm{TeV}$.

7.4.2 Background estimation

In the ggF-enriched signal region, the major backgrounds originate from the ZZ and WZ processes, which account for 60 % and 30 % of the total background contribution, respectively. The non-resonant-$\ell\ell$ background yields a relative contribution of about 5 % to the total background, while the largely suppressed Z+jets background only constitutes a small fraction (4 %). Finally, the remaining contributions from other processes (VVV and $t\bar{t}V$), amount in total to less than 1 % of the total background. A similar composition of background processes is found in the VBF-enriched signal region, where the total background yield is expected to be smaller than 1 % of that in the ggF-enriched signal region, due to the event selection for the VBF phase space. The various background

estimates and their uncertainties are described below.

The main background contribution from ZZ production is estimated using a semi-data-driven method. Similarly to the $\ell^+\ell^-\ell^+\ell^-$ analysis, the predicted ZZ yield is scaled by a floating normalization factor, which is determined in the statistical fit to the signal-region data (see Section 7.6.1). The introduction of the data-driven normalization factor helps constrain the total uncertainty in the ZZ yield, while the theoretical and experimental uncertainties in the transverse mass distribution are evaluated from simulation.

To estimate the background from WZ production in the ggF-enriched signal region, a control region enriched in WZ events, with a purity of over 90 %, is defined using the preselection criteria, except that a third lepton with $p_{\mathrm{T}} > 20\,\mathrm{GeV}$ is required. Several further selections such as $\mathrm{S}(E_{\mathrm{T}}^{\mathrm{miss}}) > 3$, a b-jets veto, and $m_{\mathrm{T}}^{W} > 60\,\mathrm{GeV}$, where m_{T}^{W} is constructed from the third lepton's transverse momentum and the $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ vector,[1] are applied to suppress non-WZ contributions. A normalization factor is calculated in the control region as the number of observed events in data, after subtracting the non-WZ contributions estimated from simulation, divided by the predicted WZ yield. The factor is found to be 1.05 with a total uncertainty of 5 %, which is consistent with a recent WZ measurement [269] performed within a broader fiducial phase space. The statistical uncertainty of the data in the control region leads to a 0.8 % uncertainty in the WZ estimate in the signal region. The main systematic uncertainty is evaluated for the ratio of the WZ predictions in the signal and control regions, and covers the experimental uncertainties and the theoretical ones related to the PDFs and the QCD scales. The uncertainty related to the subtraction of the non-WZ contribution in the control region is estimated by applying cross-section uncertainties for all the relevant processes and is found to be negligible. An additional uncertainty is assigned to the WZ prediction in the signal region, to account for the efficiency mismodelling of vetoing a third lepton in $WZ \to \ell\ell\ell\nu$ events. The total uncertainty in the WZ estimate for the ggF-enriched signal region is about 5 %. A similar method is adopted to estimate the WZ contribution in the VBF-enriched signal region, except that the control region additionally selects two jets with $p_{\mathrm{T}} > 30\,\mathrm{GeV}$. The normalization factor is found to be 0.83 with an uncertainty of 0.27, which is compatible with the results presented in Ref. [270]. The total uncertainty in the WZ estimate for the VBF-enriched signal region is about 30 %. The kinematic distributions are taken from simulation.

To estimate the non-resonant-$\ell\ell$ background, a control region dominated by the non-resonant-$\ell\ell$ processes (with a purity of about 95 %) is defined with all the event selection criteria except that the final state is required to contain an opposite-sign $e\mu$ pair. The non-resonant-$\ell\ell$ contribution in the ee ($\mu\mu$) channel is calculated as one half of the observed data yield after subtracting the contribution from the other background processes in the control region, and then corrected for the difference in the lepton reconstruction and identification efficiencies between selecting an $e\mu$ pair and an ee ($\mu\mu$) pair. The lepton efficiency correction is derived as the square root of the ratio of the numbers of $\mu\mu$ and ee events in data after the preselection. The choice of deriving the correction after preselection minimizes the resulting statistical uncertainty. The total uncertainty in the non-resonant-$\ell\ell$ estimate in the ggF-enriched signal region is about 9 %, including the statistical uncertainty of the data in the control region and the method bias estimated from simulation. The estimation of the non-resonant-$\ell\ell$ background in the VBF-enriched signal region relies on a similar methodology, except that the control region is defined with a jet selection that is looser than in the

[1] $m_{\mathrm{T}}^{W} = \sqrt{2p_{\mathrm{T}}^{\ell}E_{\mathrm{T}}^{\mathrm{miss}}[1-\cos\Delta\phi(\vec{p}_{\mathrm{T}}^{\,\ell}, \vec{E}_{\mathrm{T}}^{\mathrm{miss}})]}$

signal region. The non-resonant-$\ell\ell$ estimate obtained with the looser selection is then scaled by a simulation-based transfer factor to derive the final estimate in the VBF-enriched signal region. The transfer factor is subject to experimental and theoretical uncertainties, and the relative uncertainty in the final estimate in the VBF-enriched signal region is 70 %. The kinematic distributions for the non-resonant-$\ell\ell$ background in the signal region are predicted with simulation, and the assigned systematic uncertainty covers the experimental uncertainty in the simulated shape as well as the difference between data and simulation in the control region.

The Z+jets background contribution is estimated from simulation and scaled by a normalization factor derived in a control region enriched in Z+jets events. The control region is defined with all event selection criteria except that $\mathrm{S}(E_{\mathrm{T}}^{\mathrm{miss}})$ must be less than 9 and no requirements on the azimuthal angle difference between jets with $p_{\mathrm{T}} > 100\,\mathrm{GeV}$ and $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ are made. The normalization factor is found to be close to one. Apart from the statistical uncertainty in the control sample, the experimental and theoretical uncertainties are evaluated for the ratio of the number of simulated events in the signal region to that in the control region. The total uncertainty in the Z+jets estimate is about 40 %. The kinematic distributions for the Z+jets background are modelled with simulation. Finally, backgrounds from the VVV and $t\bar{t}V$ processes, which contribute less than 1 % of the total background, are estimated from simulation.

7.4.3 Signal and background modelling

The modelling of the transverse mass m_{T} distribution for signal and background is based on templates derived from fully simulated events and afterwards used to fit the data. In the case of a narrow resonance, simulated events generated for fixed mass hypotheses as described in Section 7.2 are used as the inputs in the moment-morphing technique [271] to obtain the m_{T} distribution for any other mass hypothesis.

The extraction of the interference terms for the LWA case is performed in the same way as in the $\ell^+\ell^-\ell^+\ell^-$ final state, as described in Section 7.3.3. In the case of the $\ell^+\ell^-\nu\bar{\nu}$ final state a correction factor, extracted as a function of m_{ZZ}, is used to reweight the interference distributions obtained at particle level to account for reconstruction effects. The final expected LWA m_{T} distribution is obtained from the combination of the interference distributions with simulated m_{T} distributions, which are interpolated between the simulated mass points with a weighting technique using the Higgs propagator, a method similar to that used for the interference.

7.5 Systematic uncertainties

The systematic uncertainties can be categorized into experimental and theoretical uncertainties. The first category includes the uncertainties resulting from the integrated luminosity, the trigger efficiencies, the momentum scale and resolution of tracks, the reconstruction and identification of leptons and jets, and their energy scale and resolution calibrations. Systematic uncertainties associated with data-driven methods are also in this category, but described in their corresponding sections: Section 7.3.2 for the $\ell^+\ell^-\ell^+\ell^-$ final state and Section 7.4.2 for the $\ell^+\ell^-\nu\bar{\nu}$ final state. The second category includes the uncertainties in the theoretical descriptions of the signal and background simulations.

These systematic uncertainties evaluated separately for signal and background in each category affect signal acceptances and background yields as well as the probability density distributions of the discriminating variables. They are provided as the inputs for the statistical interpretations described in Section 7.7, in which the impact of these uncertainties on the expected signal yields are also presented.

7.5.1 Experimental uncertainties

The uncertainty in the combined 2015–2018 integrated luminosity is 1.7 %, as described in Section 6.2.

The lepton identification and reconstruction efficiency and energy/momentum scale and resolution are derived from data using $J/\psi \to \ell\ell$ and $Z \to \ell\ell$ events. The uncertainties in the reconstruction performance are computed following the method described in Section 5.3.5 and Ref. [272] for muons, and in Section 5.3.4 and Ref. [184] for electrons. In general, their impact on the signal and background yields is less than 1 % in the $\ell^+\ell^-\nu\bar{\nu}$ final state, and up to 1.5 % in the $\ell^+\ell^-\ell^+\ell^-$ final state. In addition, the lepton isolation uncertainty is estimated to be less than 1 % in both final states.

The uncertainties in the jet energy scale and resolution have several sources, including uncertainties in the absolute and relative *in situ* calibration, the correction for pileup, the flavour composition and response [194]. Each source is treated as an independent component. They vary from 4.5 % for jets with transverse momentum $p_{\mathrm{T}} = 20\,\mathrm{GeV}$, decreasing to 1 % for jets with $p_{\mathrm{T}} = 100$–$1500\,\mathrm{GeV}$ and increasing again to 3 % for jets with higher p_{T}. They are the dominant uncertainties in the VBF-enriched categories for ggF signal production and SM ZZ production in both final states.

Uncertainties in the lepton and jet energy scales are propagated to the uncertainty in the $E_{\mathrm{T}}^{\mathrm{miss}}$ [201]. Additionally, the uncertainties from the momentum scale and resolution of the tracks that are not associated with any identified lepton or jet contribute 8 % and 3 %, respectively, to the uncertainty in the $E_{\mathrm{T}}^{\mathrm{miss}}$ value.

The efficiency of the lepton triggers in events with reconstructed leptons is nearly 100 %, and hence the related uncertainties are negligible. The uncertainties associated with the pileup reweighting are also taken into account; their impact on the signal and background yields is about 1 % for both final states.

These experimental uncertainties are common to the two final states; therefore, they are fully correlated between the two final states.

7.5.2 Theoretical uncertainties

For the simulation-based estimates, the theoretical uncertainties stemming from parton distribution functions (PDFs), missing higher-order QCD corrections, and parton showering are considered.

The PDF uncertainty is evaluated by taking the envelope of variations among alternative PDF choices and the estimate from its internal PDF error sets, following the PDF4LHC recommendation [273]. The missing higher-order QCD corrections are estimated by halving or doubling the factorization and renormalization scales independently, among which the largest effect is taken as the systematic uncertainty.[2] The parton-showering uncertainty is assessed by varying the Pythia con-

[2] In general, a seven-point pairwise variation of the factorization and renormalization scales is used, with $0.5 \leq \mu_R/\mu_F \leq 2$, i.e. $(\mu_R, \mu_F) \times \{(0.5, 0.5), (1, 0.5), (0.5, 1), (1, 1), (2, 1), (1, 2), (2, 2)\}$.

figurations, such as the parameter values of the AZNLO tune, the multi-parton models and the final-state radiation models.

For different signal hypotheses, the impact of these theoretical uncertainties on the signal acceptance and the spectrum of the discriminating variables is evaluated. In total, the theoretical uncertainty in the signal acceptance varies from less than 1 % in the low mass region to 12 % in the high mass region of the $\ell^+\ell^-\nu\bar{\nu}$ final state, and from less than 1 % in the low mass region to up to 20 % in the high mass region of the $\ell^+\ell^-\ell^+\ell^-$ final state.

For the continuum ZZ background, a common floating normalization factor is introduced to scale the number of events for the $q\bar{q} \to ZZ$ and $gg \to ZZ$ processes, while the relative yields of the two processes are estimated from the simulations. Therefore, in addition to the spectrum of the discriminating variables in the ZZ background, the theoretical uncertainties are also propagated to the simulation-based estimation of the relative yields. Moreover, the uncertainty associated with the NLO EW corrections, calculated in Refs. [235, 236, 238], are also taken into account, affecting the discriminating variables by less than 1 % in the low mass region and up to 10 % in the high mass region for both final states.

Because the $\ell^+\ell^-\ell^+\ell^-$ and $\ell^+\ell^-\nu\bar{\nu}$ searches are sensitive to different mass ranges, and therefore also to different energy scales, these theoretical uncertainties are assumed to be completely uncorrelated between the two analyses. A fully correlated scenario is also examined and the differences between the two scenarios in terms of the expected limits on various signal hypotheses are negligible.

7.6 Results

The statistical procedure used to extract the results is described in Section 7.6.1 and the results are presented in Section 7.6.2.

7.6.1 Statistical procedure and impact of systematic uncertainties

The statistical treatment of the data interpretation follows the procedure for the Higgs-boson search combination in 7 TeV data [274, 275]. The test statistic used for limit setting is the profile likelihood ratio $\Lambda(\alpha, \boldsymbol{\theta})$, which depends on one or more parameters of interest α, additional normalization factors and extra nuisance parameters $\boldsymbol{\theta}$. The parameter of interest is the cross section times branching ratio of the heavy resonance decaying into the two final states. The normalization factors, which were not used in the previous publication [93], are introduced separately for each final state to scale the expected number of the SM ZZ background events in each category and are determined by a likelihood fit to the data. This allows the systematic uncertainty to be reduced by removing both the theoretical and luminosity uncertainties contributing to the normalization uncertainty. In the $\ell^+\ell^-\ell^+\ell^-$ final state, three floating normalization factors are introduced for the VBF-enriched, ggF-MVA-high and ggF-MVA-low categories. They are referred to as $\mu_{ZZ}^{\text{VBF-MVA}}$, $\mu_{ZZ}^{\text{ggF-MVA-high}}$ and $\mu_{ZZ}^{\text{ggF-MVA-low}}$, respectively. The use of three ZZ normalization factors for the $\ell^+\ell^-\ell^+\ell^-$ final state is motivated by the different phase spaces defined for the respective signal regions. Only one floating normalization factor μ_{ZZ} is introduced in the $\ell^+\ell^-\nu\bar{\nu}$ final state, due to the limited size of the data sample and the worse signal-to-background ratio in the respective VBF-enriched signal region.

The nuisance parameters represent the estimates of the systematic uncertainties and each of them is constrained by a Gaussian distribution. For each category of each final state, a discriminating

variable is used to further separate signal from background. The number of signal events is extracted from a simultaneous fit to the discriminating variable, $m_{4\ell}$ in the $\ell^+\ell^-\ell^+\ell^-$ analysis and m_{T} in the $\ell^+\ell^-\nu\bar{\nu}$ analysis, in the event categories described in Sections 7.3 and 7.4.

The impact of a systematic uncertainty on the result depends on the production mode and the mass hypothesis. For the ggF production mode, at lower masses the ZZ parameterization for the $\ell^+\ell^-\ell^+\ell^-$ final state and the systematic uncertainty of the Z+jets background for the $\ell^+\ell^-\nu\bar{\nu}$ final state dominate, and at higher masses the uncertainties in the NLO EW correction and parton showering become important, as also seen in VBF production. For the VBF production mode, the dominant uncertainties come from the theoretical modelling of the discriminating variables of the ZZ events in the VBF category. At lower masses, jet-energy-scale uncertainties are also important. Table 7.2 shows the impact of the leading systematic uncertainties on the predicted signal event yield when the cross section times branching ratio is set to the expected upper limit (shown in Figure 7.13), for ggF and VBF production modes. The statistical uncertainty of the data sample dominates in both of the present searches, and the systematic uncertainties impact the searches to a much lesser extent.

Table 7.2: Impact of the leading systematic uncertainties, the data statistical uncertainties and the total uncertainties on the predicted signal event yield with the cross section times branching ratio being set to the expected upper limit, expressed as a percentage of the signal yield for the ggF (left) and VBF (right) production modes at $m_H = 300$, 600, 1000 and 1500 GeV.

ggF production		VBF production	
Systematic source	Impact [%]	Systematic source	Impact [%]
$m_H = 300\,\mathrm{GeV}$			
ZZ parameterization ($\ell^+\ell^-\ell^+\ell^-$)	4.5	Jet flavor composition	3.0
Z+jets modelling ($\ell^+\ell^-\nu\bar{\nu}$)	2.3	$q\bar{q} \to ZZ$ QCD scale (VBF-enriched category, $\ell^+\ell^-\ell^+\ell^-$)	2.8
Parton showering of ggF ($\ell^+\ell^-\ell^+\ell^-$)	2.2	ZZ parameterization ($\ell^+\ell^-\ell^+\ell^-$)	2.3
$e\mu$ statistical uncertainty $\ell^+\ell^-\nu\bar{\nu}$	2.0	Jet energy scale(*in-su* calibration)	1.8
Data stat. uncertainty	53	Data stat. uncertainty	58
Total uncertainty	55	Total uncertainty	60
$m_H = 600\,\mathrm{GeV}$			
Electroweak corrections for $q\bar{q} \to ZZ$ ($\ell^+\ell^-\nu\bar{\nu}$)	4.9	QCD scale of $q\bar{q} \to ZZ$ ($\ell^+\ell^-\nu\bar{\nu}$)	7.6
QCD scale of $q\bar{q} \to ZZ$ ($\ell^+\ell^-\nu\bar{\nu}$)	2.5	Jet energy resolution	5.4
Z+jets modelling ($\ell^+\ell^-\nu\bar{\nu}$)	2.5	Parton showering ($\ell^+\ell^-\nu\bar{\nu}$)	3.3
PDF of $q\bar{q} \to ZZ$ ($\ell^+\ell^-\ell^+\ell^-$)	2.2	Electroweak corrections for $q\bar{q} \to ZZ$ ($\ell^+\ell^-\nu\bar{\nu}$)	3.0
Data stat. uncertainty	54	Data stat. uncertainty	61
Total uncertainty	57	Total uncertainty	63
$m_H = 1000\,\mathrm{GeV}$			
Electroweak corrections for $q\bar{q} \to ZZ$ ($\ell^+\ell^-\nu\bar{\nu}$)	9.3	Parton showering ($\ell^+\ell^-\nu\bar{\nu}$)	6.8
Parton showering ($\ell^+\ell^-\nu\bar{\nu}$)	5.2	Electroweak corrections for $q\bar{q} \to ZZ$ ($\ell^+\ell^-\nu\bar{\nu}$)	4.7
QCD scale of $q\bar{q} \to ZZ$ ($\ell^+\ell^-\nu\bar{\nu}$)	4.8	QCD scale of $q\bar{q} \to ZZ$ ($\ell^+\ell^-\nu\bar{\nu}$)	2.4
Z+jets modelling ($\ell^+\ell^-\nu\bar{\nu}$)	2.4	Jet flavor composition	2.4
Data stat. uncertainty	57	Data stat. uncertainty	58
Total uncertainty	59	Total uncertainty	59
$m_H = 1500\,\mathrm{GeV}$			
Parton showering ($\ell^+\ell^-\nu\bar{\nu}$)	9.6	Parton showering ($\ell^+\ell^-\nu\bar{\nu}$)	9.0
Electroweak corrections for $q\bar{q} \to ZZ$ ($\ell^+\ell^-\nu\bar{\nu}$)	6.8	Electroweak corrections for $q\bar{q} \to ZZ$ ($\ell^+\ell^-\nu\bar{\nu}$)	4.6
PDF of $q\bar{q} \to ZZ$ ($\ell^+\ell^-\nu\bar{\nu}$)	5.4	PDF of $q\bar{q} \to ZZ$ ($\ell^+\ell^-\nu\bar{\nu}$)	3.4
QCD scale of $q\bar{q} \to ZZ$ ($\ell^+\ell^-\nu\bar{\nu}$)	4.6	QCD scale of $q\bar{q} \to ZZ$ ($\ell^+\ell^-\nu\bar{\nu}$)	2.8
Data stat. uncertainty	57	Data stat. uncertainty	55
Total uncertainty	59	Total uncertainty	57

7.6.2 General results

The total number of observed events is 3275 in the $\ell^+\ell^-\ell^+\ell^-$ final state ($m_{4\ell} > 200\,\text{GeV}$) and 2794 in the $\ell^+\ell^-\nu\bar{\nu}$ final state. The expected background yields are obtained from a simultaneous likelihood fit of the two final states under the background-only hypothesis. The fitted normalization factors for the SM ZZ background are summarized in Table 7.3.

Table 7.3: The ZZ normalization factors together with their total (statistical+systematic) uncertainties in each category of the two final states, which scale the number of ZZ events estimated from the simulations, obtained from a simultaneous likelihood fit of the two final states under the background-only hypothesis. For the $\ell^+\ell^-\ell^+\ell^-$ final state, the MVA-based categorization is used.

Final state	Normalization factor	Fitted value
$\ell^+\ell^-\ell^+\ell^-$	$\mu_{ZZ}^{\text{VBF-MVA}}$	0.9 ± 0.3
	$\mu_{ZZ}^{\text{ggF-MVA-high}}$	1.07 ± 0.05
	$\mu_{ZZ}^{\text{ggF-MVA-low}}$	1.12 ± 0.03
$\ell^+\ell^-\nu\bar{\nu}$	μ_{ZZ}	1.04 ± 0.06

The number of observed candidate events with mass above 200 GeV together with the expected background yields for each of the five categories of the $\ell^+\ell^-\ell^+\ell^-$ analysis as described in Section 7.3.1 are presented in Table 7.4. The $m_{4\ell}$ spectrum in each category is shown in Figure 7.11 for illustration, since the backgrounds are determined from a combined unbinned likelihood fit to the data under the background-only hypothesis. Table 7.5 contains the number of observed events along with the obtained background yields for the $\ell^+\ell^-\nu\bar{\nu}$ analysis and Figure 7.12 shows the m_{T} distribution for the electron and muon channels in the ggF-enriched and VBF-enriched categories.

The two previous excesses around 240 GeV and 700 GeV that were observed in the publication [93] using 2015 and 2016 data are not confirmed using the full Run 2 dataset as explained below. The maximum deviation of the data from the background-only hypothesis is evaluated in the context of a NWA signal from the ggF production or from the VBF production separately. For the ggF production, the maximum deviation is for a signal mass hypothesis around 240 GeV, with a local significance of 2.0 standard deviations and a global significance of 0.5 standard deviation. For the VBF production, the maximum deviation is for a signal mass hypothesis around 620 GeV, with a local significance of 2.4 standard deviations and a global significance of 0.9 standard deviation.

7.7 Interpretations

Since no significant excess with respect to the background predictions is found, results obtained from the combination of the $\ell^+\ell^-\ell^+\ell^-$ and $\ell^+\ell^-\nu\bar{\nu}$ final states are interpreted in terms of exclusion limits for different signal hypotheses as presented below.

7.7.1 Spin-0 resonances

Spin-0 resonances with NWA

Upper limits on the cross section times branching ratio ($\sigma \times B(H \to ZZ)$) for a heavy resonance are obtained from the combination of the two final states, as a function of m_H with the CL_s proce-

Table 7.4: Expected and observed numbers of events in the $\ell^+\ell^-\ell^+\ell^-$ final state for $m_{4\ell} > 200\,\text{GeV}$, together with their uncertainties, for the VBF-MVA-enriched, ggF-MVA-high and ggF-MVA-low categories. The expected numbers of events, as well as their uncertainties, are obtained from a combined likelihood fit to the data under the background-only hypothesis. The uncertainties of the ZZ normalization factors, presented in Table 7.3, are also taken into account.

Process	VBF-enriched	ggF-MVA-high			ggF-MVA-low
		4μ channel	$2e2\mu$ channel	$4e$ channel	
$q\bar{q} \to ZZ$	11 ± 4	232 ± 10	389 ± 17	154 ± 7	2008 ± 47
$gg \to ZZ$	3 ± 2	37 ± 6	64 ± 10	26 ± 4	247 ± 19
ZZ (EW)	4.1 ± 0.4	4.5 ± 0.2	7.5 ± 0.4	3 ± 2	14.3 ± 0.7
Z+jets, $t\bar{t}$	0.08 ± 0.02	0.6 ± 0.1	1.7 ± 0.4	0.8 ± 0.1	8.8 ± 2.1
$t\bar{t}V$, VVV	0.96 ± 0.01	9.8 ± 0.2	17.5 ± 0.4	7.7 ± 0.2	21.9 ± 0.5
Total background	19 ± 5	284 ± 12	480 ± 20	192 ± 8	2300 ± 51
Observed	19	271	493	191	2301

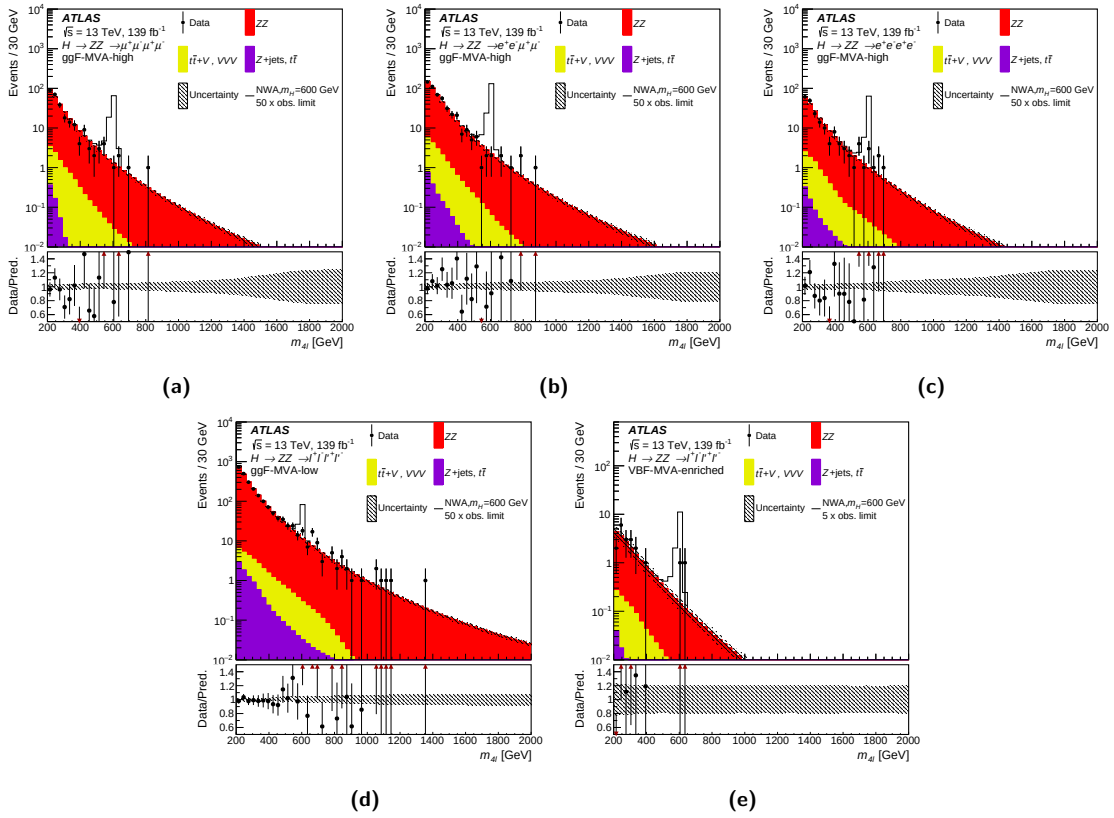

Figure 7.11: Distributions of the four-lepton invariant mass $m_{4\ell}$ in the $\ell^+\ell^-\ell^+\ell^-$ search for the ggF-MVA-high categories ($\mu^+\mu^-\mu^+\mu^-$ (a), $e^+e^-\mu^+\mu^-$ (b), and $e^+e^-e^+e^-$ (c) final states), for the ggF-MVA-low category (d), and for the VBF-MVA-enriched category (e). The backgrounds are determined from a combined likelihood fit to the data under the background-only hypothesis. The simulated $m_H = 600\,\text{GeV}$ signal is normalized to a cross section corresponding to 50 (5) times the observed limit given in Section 7.7.1 for the ggF (VBF) production mode. The error bars on the data points indicate the statistical uncertainty, while the systematic uncertainty in the prediction is shown by the hatched band. The lower panels show the ratio of data to prediction. The red arrows indicate data points that are outside the displayed range.

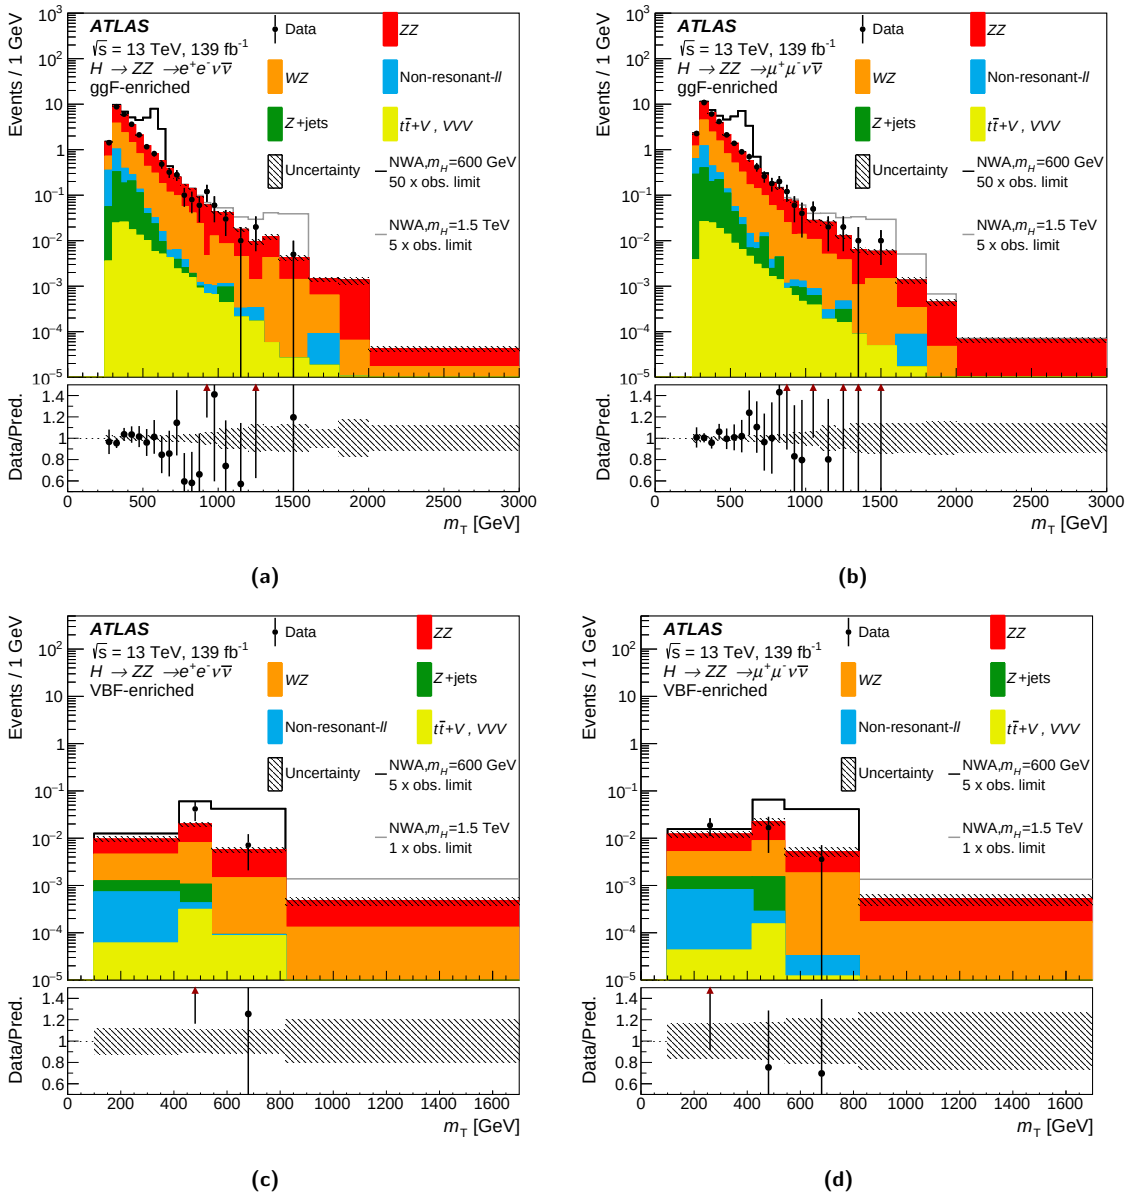

Figure 7.12: The m_{T} distribution in the $\ell^+\ell^-\nu\bar{\nu}$ search for (a),(b) the ggF categories and (c),(d) the VBF categories. Events beyond the upper limit of the histogram are included in the last bin of the distribution. The backgrounds are determined from a combined likelihood fit to data under the background-only hypothesis. The simulated $m_H = 600\,\mathrm{GeV}$ (1.5 TeV) signals are normalized to a cross section corresponding to 50 (5) times the observed limit given in Section 7.7.1 for the ggF production mode and to 5 (1) times the observed limit for the VBF production mode. The error bars on the data points indicate the statistical uncertainty and markers are drawn at the bin centre. The systematic uncertainty in the prediction is shown by the hatched band. The lower panels show the ratio of data to prediction. The red arrows indicate data points that are outside the displayed range.

Table 7.5: Expected and observed numbers of events together with their uncertainties in the $\ell^+\ell^-\nu\bar{\nu}$ final state, for the ggF- and VBF-enriched categories. The expected numbers of events, as well as their uncertainties, are obtained from a likelihood fit to the data under the background-only hypothesis. The uncertainties of the ZZ normalization factors, presented in Table 7.3, are also taken into account.

Process	ggF-enriched		VBF-enriched	
	e^+e^- channel	$\mu^+\mu^-$ channel	e^+e^- channel	$\mu^+\mu^-$ channel
$q\bar{q} \to ZZ$	695 $\pm$ 39	795 $\pm$ 44	2.8 $\pm$ 0.2	3.3 $\pm$ 0.2
$gg \to ZZ$	87 $\pm$ 28	97 $\pm$ 31	1.0 $\pm$ 0.4	1.0 $\pm$ 0.4
ZZ (EW)	6.6 $\pm$ 0.5	7 $\pm$ 5	0.8 $\pm$ 0.1	0.9 $\pm$ 0.1
WZ	400 $\pm$ 13	443 $\pm$ 12	2.4 $\pm$ 0.5	3.0 $\pm$ 1.3
Z+jets	39 $\pm$ 12	56 $\pm$ 21	0.2 $\pm$ 0.2	0.3 $\pm$ 0.3
Non-resonant-$\ell\ell$	66 $\pm$ 6	77 $\pm$ 7	0.2 $\pm$ 0.2	0.3 $\pm$ 0.2
$t\bar{t}V$, VVV	5.9 $\pm$ 0.4	5.9 $\pm$ 0.4	0.08 $\pm$ 0.02	0.04 $\pm$ 0.01
Total backgrounds	1299 $\pm$ 52	1480 $\pm$ 59	7.4 $\pm$ 0.7	8.4 $\pm$ 1.4
Observed	1280	1498	7	9

dure [276] in the asymptotic approximation.[3] It is assumed that an additional heavy scalar would be produced mainly via the ggF and VBF processes but that the ratio of the two production mechanisms might depend on the model considered. For this reason, fits for the ggF and VBF processes are done separately, and in each case the other process is allowed to float in the fit as an additional free parameter. Figure 7.13 presents the observed and expected limits at 95 % CL on the $\sigma \times B(H \to ZZ)$ of a narrow scalar resonance for the ggF (left) and VBF (right) production modes, as well as the expected limits from the $\ell^+\ell^-\ell^+\ell^-$ and $\ell^+\ell^-\nu\bar{\nu}$ searches. This result is valid for models in which the width is less than 0.5 % of m_H. When combining the two final states, the 95 % CL upper limits range from 215 fb at $m_H = 240$ GeV to 2.0 fb at $m_H = 1900$ GeV for the ggF production mode and from 87 fb at $m_H = 255$ GeV to 1.5 fb at $m_H = 1800$ GeV for the VBF production mode. Compared with the expected limits projected to the luminosity of 139 fb^{-1} from the previous publication [93], the current expected limits are decreased by a factor ranging from 20 % to 28 % for the ggF production mode and from 27 % to 43 % for the VBF production mode, depending on the mass hypothesis.

Spin-0 resonances with LWA

In the case of the LWA, upper limits on the cross section for the ggF process times branching ratio ($\sigma_{\text{ggF}} \times B(H \to ZZ)$) are set for different widths of the heavy scalar. Figure 7.14 shows the limits for a width of 1 %, 5 %, 10 % and 15 % of m_H respectively. The limits are set for masses of m_H higher than 400 GeV. The choice of 400 GeV as the lower boundary is to avoid any major instability in the parametrization of the mass spectra for the LWA signals and the interference effects, especially when the signal mass gets smaller. The interpretation has only been carried out for the ggF process, which for this signal model is assumed to have the largest production cross section, and is therefore the production mode best-suited to study the impact of a non-trivial resonance width on the search.

[3] The results from this approximation were verified to be correct within about 4 % using pseudo-experiments.

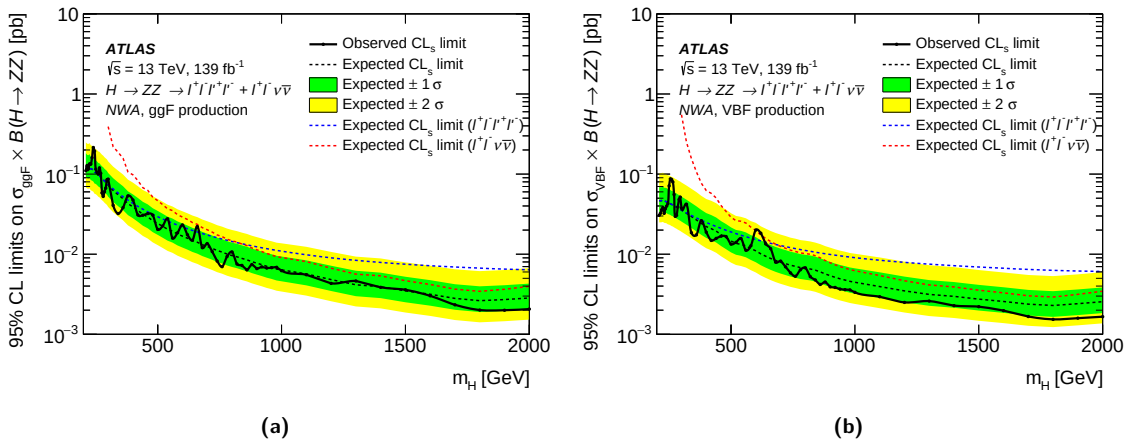

Figure 7.13: The upper limits at 95 % CL on the cross section times branching ratio as a function of the heavy resonance mass m_H for (a) the ggF production mode ($\sigma_{\mathrm{ggF}} \times B(H \to ZZ)$) and (b) for the VBF production mode ($\sigma_{\mathrm{VBF}} \times B(H \to ZZ)$) in the case of the NWA. The black line indicates the observed limit. The green and yellow bands represent the $\pm 1\sigma$ and $\pm 2\sigma$ uncertainties in the expected limits. The dashed coloured lines indicate the expected limits obtained from the individual searches.

Two-Higgs-doublet model

A search in the context of a CP-conserving 2HDM is also presented. This model has five physical Higgs bosons after electroweak symmetry breaking: two CP-even, one CP-odd, and two charged. As described in Section 2.3.1, the model considered here has seven free parameters: the Higgs boson masses, the ratio of the vacuum expectation values of the two Higgs doublets ($\tan\beta$), the mixing angle between the CP-even Higgs bosons (α), and the potential parameter m_{12}^2 that mixes the two Higgs doublets. The two Higgs doublets Φ_1 and Φ_2 can couple to leptons and up- and down-type quarks in several ways. In the Type-I model, Φ_2 couples to all quarks and leptons, whereas for Type-II, Φ_1 couples to down-type quarks and leptons and Φ_2 couples to up-type quarks. The "lepton-specific" model is similar to Type-I except for the fact that the leptons couple to Φ_1, instead of Φ_2; the "flipped" model is similar to Type-II except that the leptons couple to Φ_2, instead of Φ_1. In all these models, the coupling of the heavier CP-even Higgs boson to vector bosons is proportional to $\cos(\beta-\alpha)$. In the limit $\cos(\beta-\alpha) \to 0$, the light CP-even Higgs boson is indistinguishable from a SM Higgs boson with the same mass. In the context of $H \to ZZ$ decays there is no direct coupling of the Higgs boson to leptons, so only the Type-I and II interpretations are presented. In addition, our interpretations assume other Higgs bosons are heavy enough so that the heavy CP-even Higgs boson will not decay to them.

Figure 7.15 shows exclusion limits in the $\tan\beta$ versus $\cos(\beta-\alpha)$ plane for Type-I and Type-II 2HDMs, for a heavy Higgs boson with mass $m_H = 220\,\mathrm{GeV}$. This m_H value is chosen so that the assumption of a narrow Higgs boson is valid over most of the parameter space, and the experimental sensitivity is maximal. At this low mass, only the $\ell^+\ell^-\ell^+\ell^-$ final state contributes to this result. The range of $\cos(\beta-\alpha)$ and $\tan\beta$ explored is limited to the region where the assumption of a heavy narrow Higgs boson with negligible interference is valid. When calculating the limits at a given choice of $\cos(\beta-\alpha)$ and $\tan\beta$, the relative rates of ggF and VBF production in the fit are set to the prediction of the 2HDM for that parameter choice. Figure 7.16 shows exclusion limits as a function of the heavy Higgs boson mass m_H and the parameter $\tan\beta$ for $\cos(\beta-\alpha) = -0.1$, which is chosen

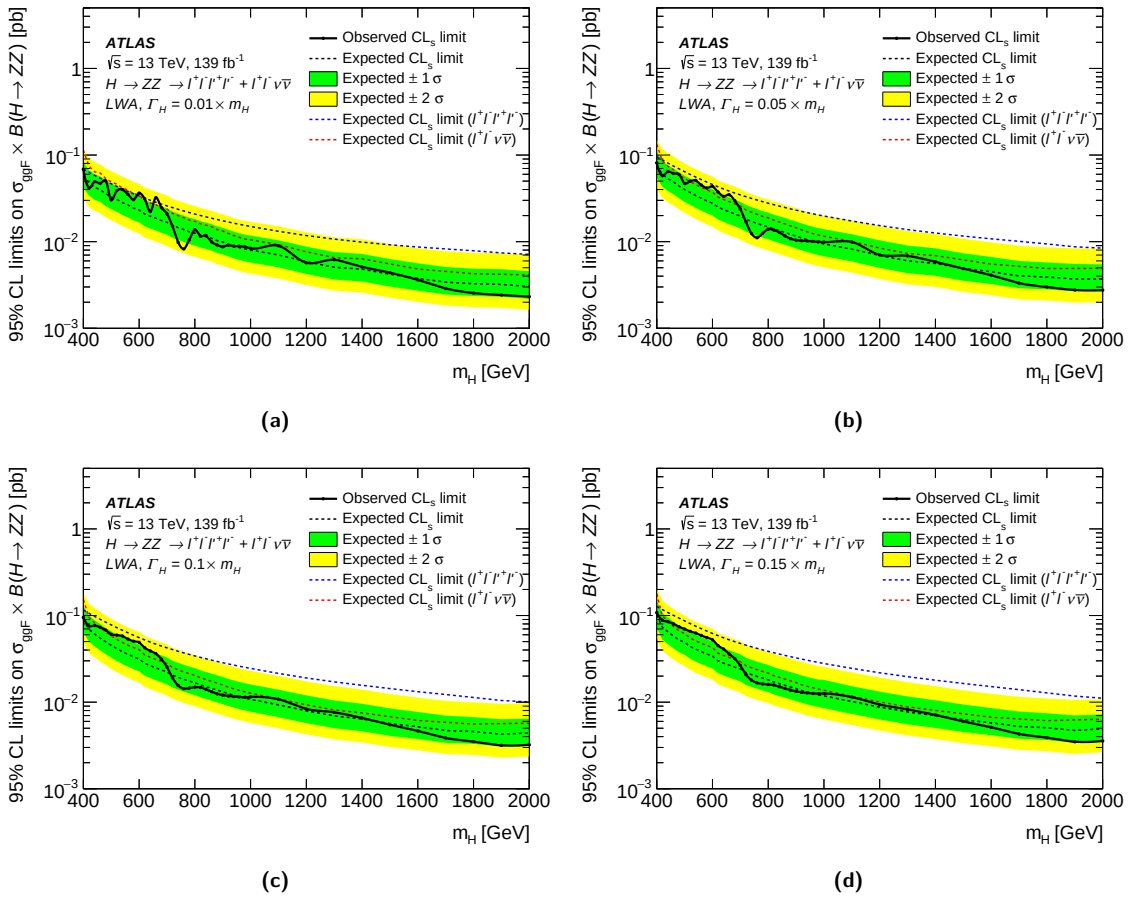

Figure 7.14: The upper limits at 95 % CL on the cross section for the ggF production mode times branching ratio ($\sigma_{\mathrm{ggF}} \times B(H \to ZZ)$) as a function of m_H for an additional heavy scalar assuming a width of (a) 1 %, (b) 5 %, (c) 10 % and (d) 15 %, of m_H. The black line indicates the observed limit. The green and yellow bands represent the $\pm 1\sigma$ and $\pm 2\sigma$ uncertainties in the expected limits. The dashed coloured lines indicate the expected limits obtained from the individual searches.

so that the light Higgs boson properties are still compatible with the recent measurements of the SM Higgs boson properties [212]. The white regions in the exclusion plots indicate regions of parameter space which are not excluded by the present analysis. In these regions the cross section predicted by the 2HDM is below the observed cross-section limit. Relative to the previous publication, the excluded regions are significantly expanded. For example, in the $\tan\beta$ versus m_H plane for the Type-II 2HDM the excluded region in $\tan\beta$ is more than 60 % larger for $200 < m_H < 400$ GeV.

7.7.2 Spin-2 resonances

The results are also interpreted as a search for a Kaluza–Klein graviton excitation, G_{KK}, in the context of the bulk RS model with $k/\overline{M}_{\mathrm{Pl}} = 1$. The limits on $\sigma \times B(G_{\mathrm{KK}} \to ZZ)$ at 95 % CL as a function of the KK graviton mass, $m(G_{\mathrm{KK}})$, are shown in Figure 7.17 together with the predicted G_{KK} cross section. A spin-2 graviton is excluded up to a mass of 1830 GeV.

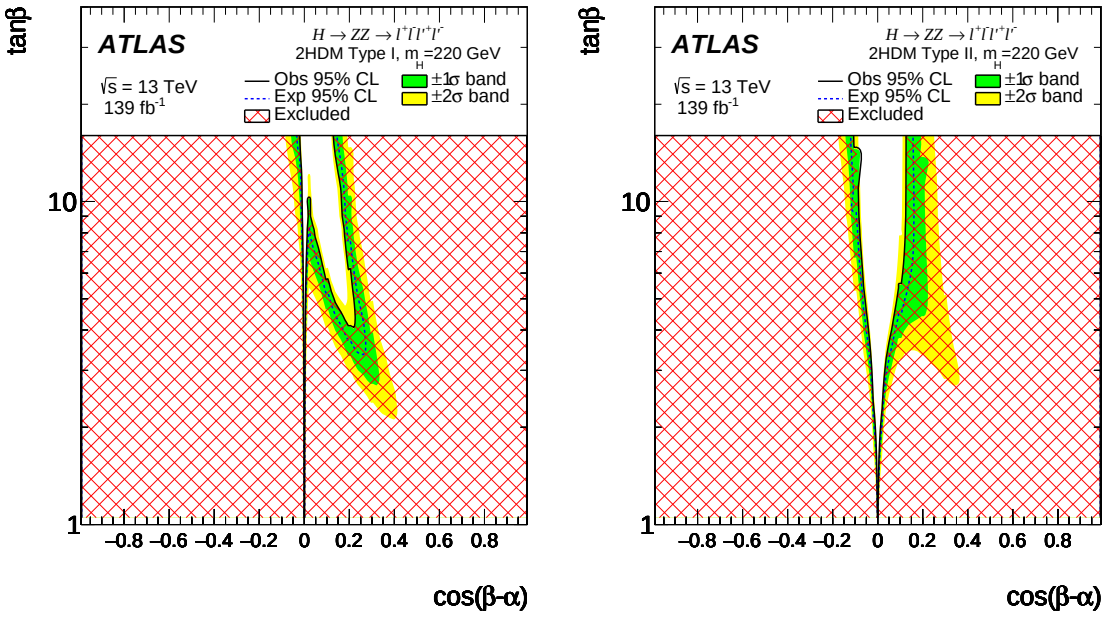

Figure 7.15: The exclusion contour in the 2HDM (a) Type-I and (b) Type-II models for $m_H = 220\,\text{GeV}$ shown as a function of the parameters $\cos(\beta-\alpha)$ and $\tan\beta$. The green and yellow bands represent the $\pm1\sigma$ and $\pm2\sigma$ uncertainties in the expected limits. The hatched area shows the observed exclusion.

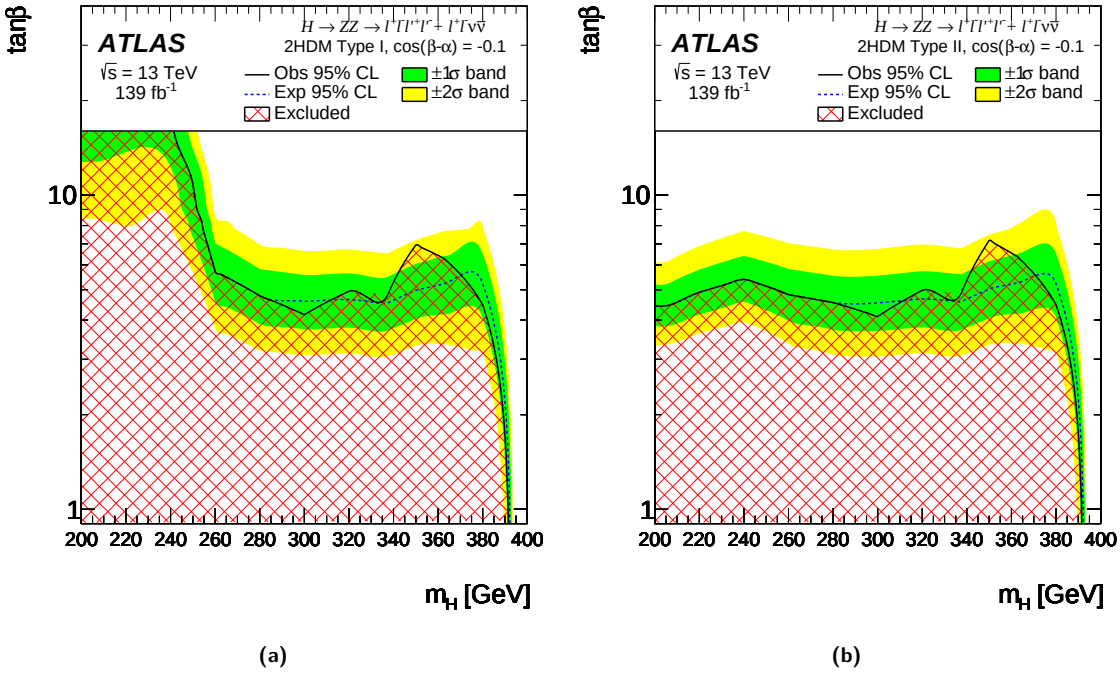

Figure 7.16: The exclusion contour in the 2HDM (a) Type-I and (b) Type-II models for $\cos(\beta-\alpha) = -0.1$, shown as a function of the heavy scalar mass m_H and the parameter $\tan\beta$. The green and yellow bands represent the $\pm1\sigma$ and $\pm2\sigma$ uncertainties in the expected limits. The hatched area shows the observed exclusion.

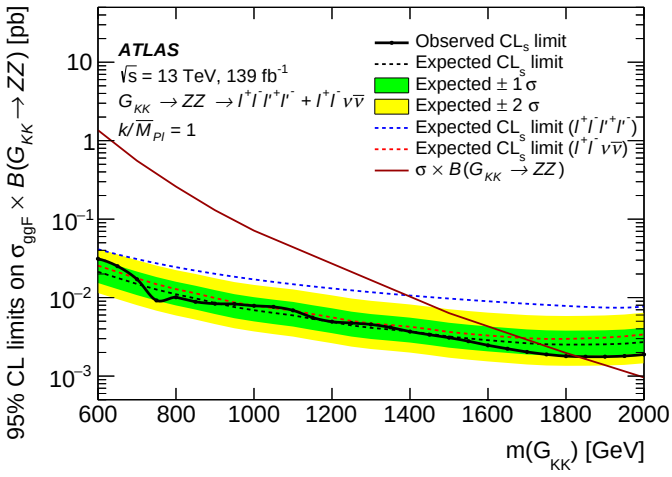

Figure 7.17: The upper limits at 95 % CL on cross section times branching ratio $\sigma \times B(G_{\mathrm{KK}} \to ZZ)$ for a KK graviton produced with $k/\overline{M}_{\mathrm{Pl}} = 1$. The black line indicates the observed limit. The green and yellow bands give the $\pm 1\sigma$ and $\pm 2\sigma$ uncertainties in the expected limits. The predicted production cross section times branching ratio as a function of the G_{KK} mass $m(G_{\mathrm{KK}})$ is shown by the red solid line.

7.8 Summary

This chapter reported on a search for heavy resonances decaying into a pair of Z bosons in the $\ell^+\ell^-\ell^+\ell^-$ and $\ell^+\ell^-\nu\bar{\nu}$ final states. The search uses proton–proton collision data collected with the ATLAS detector from 2015 to 2018 at the Large Hadron Collider at a centre-of-mass energy of 13 TeV corresponding to the full Run 2 integrated luminosity of 139 fb^{-1}. No significant excess is observed with respect to the predicted SM background; therefore, the results are interpreted as upper limits on the production cross section of spin-0 resonances or a spin-2 resonance. The mass range of the hypothetical resonances considered is between 200 GeV and 2000 GeV depending on the final state and the model considered. The spin-0 resonance is assumed to be a heavy scalar, whose dominant production modes are gluon–gluon fusion and vector-boson fusion, and it is studied in the narrow-width approximation and with the large-width assumption. In the case of the narrow-width approximation, upper limits on the production rate of a heavy scalar decaying into two Z bosons (the production cross-section times the corresponding decay branching fraction) are set separately for the ggF and VBF production modes. Combining the two final states, 95 % CL upper limits range from 215 fb at $m_H = 240$ GeV to 2.0 fb at $m_H = 1900$ GeV for the gluon–gluon fusion production mode and from 87 fb at $m_H = 255$ GeV to 1.5 fb at $m_H = 1800$ GeV for the vector-boson fusion production mode. The results are also interpreted in the context of Type-I and Type-II two-Higgs-doublet models, with exclusion contours given in the $\tan\beta$ versus $\cos(\beta-\alpha)$ (for $m_H = 220$ GeV) and $\tan\beta$ versus m_H planes. This m_H value is chosen so that the assumption of a narrow Higgs boson is valid over most of the parameter space and the experimental sensitivity is maximal. The limits on the production rate of a large-width scalar are obtained for widths of 1 %, 5 %, 10 % and 15 % of the mass of the resonance, with the interference between the heavy scalar and the SM Higgs boson as well as between the heavy scalar and the $gg \to ZZ$ continuum taken into account. In the framework of the Randall–Sundrum model with one warped extra dimension a graviton excitation spin-2 resonance with $m(G_{\mathrm{KK}}) < 1830$ GeV is excluded at 95 % CL.

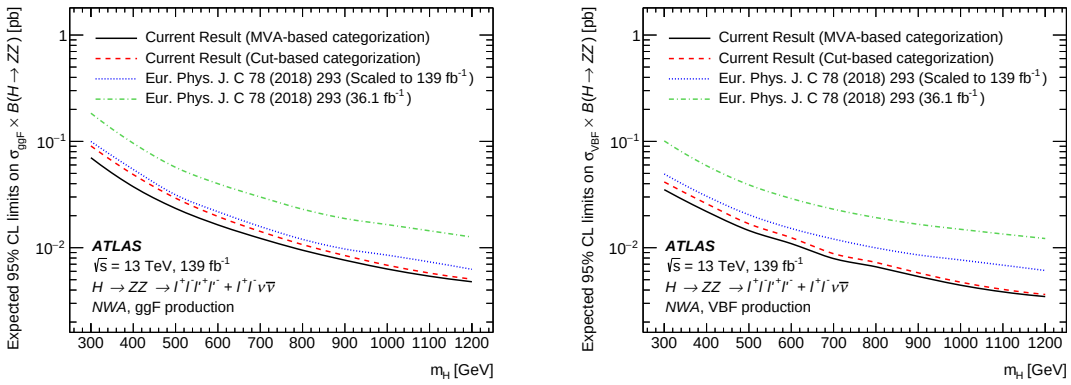

Figure 7.18: Comparisons of the expected upper limits at 95 % CL on the cross section times branching ratio as a function of the heavy resonance mass m_H for (a) the ggF production mode and for (b) the VBF production mode in the case of the NWA. The expected limits from the previous publication are shown in the green dash–dotted line and are projected to the 139 fb^{-1} as shown in the blue dotted line. In addition, the current results based on either cut-based categorization or the multivariate-based categorization are shown in dashed red and solid black lines.

The results of this analysis offer significant improvements in constraining the parameter space for the resonances and signal models studied with respect to the previous publication on this subject [93]. Figure 7.18 shows comparisons of the expected upper limits at 95 % CL on the cross section times branching ratio as a function of the heavy resonance mass m_H for the ggF and VBF production modes in the case of the NWA. Overall, these upper limits have been reduced by approximately 60–70 % over the full mass range. The largest source of improvement ($\sim 50\,\%$) comes from the full integrated luminosity of the Run 2 dataset, while the remaining improvements stem primarily from improved object-reconstruction and background-rejection methods and from the multivariate event classifiers employed in the $\ell^+\ell^-\ell^+\ell^-$ analysis.

Chapter 8

Conclusion

This thesis has presented a search for resonances decaying into a pair of Z bosons using $139\,\text{fb}^{-1}$ of proton–proton collision data collected with the ATLAS detector at a centre-of-mass energy of 13 TeV during LHC Run 2. This search considers two final states according to the decays of the Z bosons: $X \to ZZ \to \ell^+\ell^-\ell^+\ell^-$ and $X \to ZZ \to \ell^+\ell^-\nu\bar{\nu}$. The focus of this thesis is the analysis of the $\ell^+\ell^-\ell^+\ell^-$ channel, with a brief discussion of the $\ell^+\ell^-\nu\bar{\nu}$ analysis also given to provide context to the final combined result of the two channels. Within the search range spanning mass hypotheses from 200 GeV to 2000 GeV, no significant excess is observed with respect to the predicted Standard Model background; therefore, the results are interpreted as upper limits on the production cross section of spin-0 resonances and of a spin-2 resonance. In the case of the narrow-width approximation, and combining the two final states, 95 % confidence level upper limits on the production cross section times branching ratio range from 215 fb at $m_H = 240\,\text{GeV}$ to 2.0 fb at $m_H = 1900\,\text{GeV}$ for the gluon–gluon fusion production mode and from 87 fb at $m_H = 255\,\text{GeV}$ to 1.5 fb at $m_H = 1800\,\text{GeV}$ for the vector-boson fusion production mode. These limits represent an approximately 60–70 % improvement with respect to previous limits set using a $36.1\,\text{fb}^{-1}$ subset of the Run 2 data sample collected in the years 2015 and 2016. The results are also interpreted for large-width scalar resonances with widths of 1 %, 5 %, 10 % and 15 % of the resonance mass, and in the context of Type-I and Type-II two-Higgs-doublet models. Finally, the results are interpreted in the framework of the Randall–Sundrum model with one warped extra dimension, and a spin-2 Kaluza–Klein graviton excitation resonance with $m(G_{\text{KK}}) < 1830\,\text{GeV}$ is excluded at 95 % confidence level.

This thesis has also reported on a new luminosity-monitoring technique developed in Run 2 called *LAr energy flow*, which is based on the observation that the average amount of energy deposited in the LAr calorimeters per bunch crossing is proportional to the instantaneous luminosity. Due to the long drift time of the LAr calorimeter signal relative to the nominal 25 ns LHC bunch spacing and the bipolar shaping of the signal pulse, this method is only suitable for fills with sufficiently isolated bunches (typically $\gtrsim 450\,\text{ns}$). The LAr energy-flow algorithm was evaluated in a set of special μ-scan runs in 2017 and 2018 and compared against a track-counting luminosity algorithm in order to set limits on the relative nonlinearity of track counting in the calibration-transfer procedure of the nominal ATLAS luminosity measurement. The electromagnetic endcap (EMEC) and forward calorimeter (FCal) energy-flow measurements show the best agreement with track counting in fills

with 140 colliding bunch pairs, deviating by approximately 0.5 % over a range in μ from typical vdM conditions at $\mu = 0.5$ to typical physics data-taking conditions at $\mu = 50$. The electromagnetic barrel (EMB) and hadronic end-cap (HEC) energy-flow measurements generally result in poorer agreement with track counting, with deviations up to 2 % in similar running conditions. Fills with a greater number of colliding bunches degrade the agreement between all energy-flow and track-counting measurements, with deviations between 1 and 3 % over the same range in μ. The EMEC and FCal results can be used in a calibration transfer "uncertainty ladder" in the step from low-μ to high-μ isolated bunches to improve the preliminary calibration transfer uncertainty from 1.3 % to approximately 0.5 %. The calibration transfer is currently the largest source of uncertainty in the preliminary measurement of the ATLAS integrated luminosity, therefore these studies of LAr energy flow, along with improvements to the Tile calorimeter luminosity measurements, will offer significant improvements to the final Run 2 luminosity uncertainty. Since the uncertainty on the integrated luminosity is one that feeds into nearly all ATLAS physics analyses, these improvements will be far-reaching, particularly for those analyses that are currently limited by systematic uncertainties.

8.1 Outlook

At the time of writing, LHC Run 3 has just recently begun proton–proton collisions at a centre-of-mass energy of 13.6 TeV, with an integrated luminosity of $\sim 200\,\text{fb}^{-1}$ expected by the end of 2025. After this time, the LHC will undergo an extensive upgrade program to establish the *High-luminosity Large Hadron Collider* (HL-LHC) [277]. The HL-LHC expects to deliver peak luminosities of $5 \times 10^{34}\,\text{cm}^{-2}\,\text{s}^{-1}$ and an integrated luminosity of $\sim 3000\,\text{fb}^{-1}$ over approximately twelve years of operation beginning around the year 2030. This represents a roughly twenty-fold increase in the integrated luminosity with respect to the Run 2 data sample and will offer an excellent opportunity to revisit the searches for heavy resonances where statistical uncertainties currently limit the precision of the results. At the same time, the increase in peak luminosity at the HL-LHC will pose a significant challenge to the experiments given the associated increase in pileup (up to $\langle\mu\rangle \approx 200$). The ATLAS Collaboration is currently developing new detectors and read-out electronics to handle these higher pileup conditions with minimal loss of performance.

From its inception to the discovery of the Higgs boson in 2012 and beyond, the LHC has established itself in the rich history of particle physics as a machine of discovery and exploration at the energy frontier. Yet, questions remain after its first two successful runs concerning the completeness of our theories of Nature. Whether we will find conclusive evidence of physics beyond the Standard Model at the LHC or elsewhere remains one of the foremost and most important questions of our generation in high-energy physics.

Appendices

Appendix A

Luminosity measurements with LAr gap currents

A high-energy charged particle traversing a LAr calorimeter cell ionizes the argon within the LAr gap, producing free electrons and Ar^+ ions that drift towards the electrodes under the influence of the electric field applied across the gap. The drift of these electron-ion pairs induces a current in the high-voltage line serving that calorimeter cell. In order to keep the electric field across each LAr gap constant over time, the HV supplies are regulated such that any voltage drop induced by the particle flux through a given HV sector is counterbalanced by a continuous injection of electrical current. This current is proportional to the particle flux through the region of the calorimeter served by the HV supply and thereby provides a relative luminosity measurement using the HV line considered. This luminosity-monitoring technique is called the "LAr gap-current luminosity" and is illustrated schematically in Figure A.1.

The total pp cross section at the LHC is dominated by soft QCD events with low transverse momentum p_{T}, often referred to as "minimum-bias events" [161]. The distribution of the particles emerging from these minimum-bias events as a function of the particle pseudorapidity η is approximately uniform, therefore most of their energy is deposited in the forward region of the calorimeter system [279, 280]. Furthermore, these particles are predominantly low-p_{T} $\pi^{\pm}/\pi^0$ mesons, therefore most of the energy from minimum-bias events is deposited in the electromagnetic layers of the calorimeter system as electromagnetic showers originating from $\pi^0 \to \gamma\gamma$ decays [142, 281], where the π^0 mesons are produced both in the primary collisions and in $\pi^+ n \to \pi^0 p$ and $\pi^- p \to \pi^0 n$ "charge-exchange" nuclear interactions within the calorimeter material. The electromagnetic layers of the ATLAS end-cap colorimeters, namely the EMEC and the first layer of the FCal (FCal1), are therefore best suited for the LAr gap-current luminosity measurement.

The LAr gap-current luminosity is a *bunch-integrating* luminosity algorithm due to the relatively slow response of the HV power supplies and timing resolution of the current-readout systems. It therefore cannot distinguish individual LHC bunch crossings and is instead sensitive only to the instantaneous luminosity summed over all bunches. Consequently, no absolute luminosity calibration can be determined for LAr gap currents using the vdM method, which requires a precise bunch-by-bunch analysis due to the subtle differences in shape, size, orientation and position of each colliding-bunch pair [203]. Instead, it provides a relative luminosity measurement via a calibration

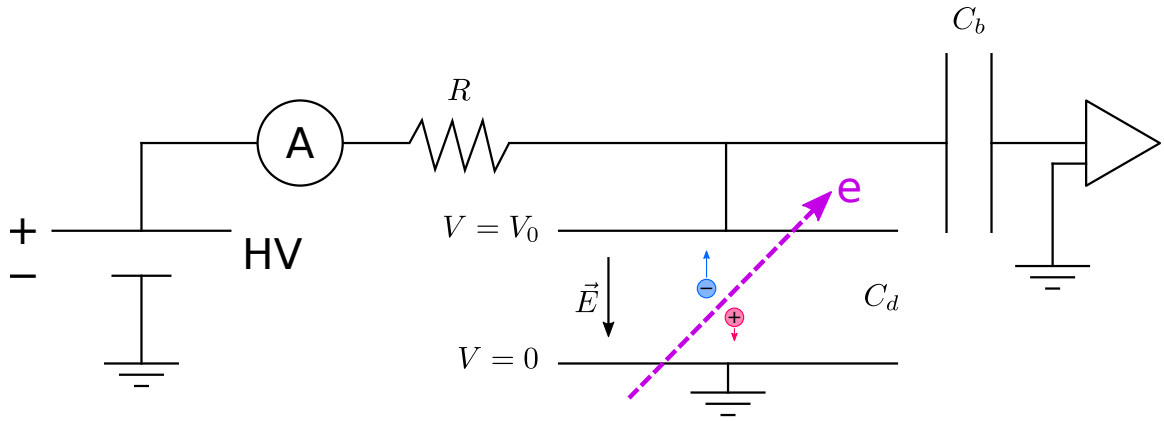

Figure A.1: Schematic diagram of the electronics chain of a single high-voltage channel used in the LAr gap-current luminosity measurement. A high-energy charged particle traversing the LAr gap ionizes the argon, producing electron-ion pairs that drift under the influence of an electric field applied across the gap. The drifting electron-ion pairs induce a current in the HV line, which is counterbalanced by the regulated HV supply to maintain a constant electric field across the gap. The current injected by the HV supply is measured by device A and is proportional to the particle flux through the region of the calorimeter served by the HV supply, thus providing a relative luminosity measurement. R represents the total resistance on the HV distribution line, C_d the detector cell capacitance and C_b the blocking capacitor that decouples the HV line from the signal readout (represented by the triangle). In practice, one LAr calorimeter HV supply serves multiple sets of electrodes, however only one is shown here for simplicity. Adapted from Ref. [278].

against another luminosity algorithm whose absolute luminosity scale is already determined, allowing independent checks of the linearity and long-term stability of these algorithms.

This section documents the ATLAS LAr gap-current luminosity measurements performed in Run 2 of the LHC, and is organized as follows. Section A.1 gives a brief history of the method and its evolution from its inception to the current methodology used in Run 2. Section A.2 describes the data samples for which the LAr gap-current luminosity has been determined. Section A.3 gives an overview of the LAr HV system and its components relevant to the luminosity measurement. The pedestal-subtraction methodology and calibration procedure are described in Sections A.4 and A.5, respectively. A new correction developed and employed in Run 2 accounting for non-linearities in the FCal measurement is presented in Section A.6. Finally, a representative set of results are given for each data sample in Section A.7.

A.1 A brief history of the LAr gap-current luminosity

The LAr gap-current luminosity-monitoring method was first proposed in Ref. [278], which includes predicted HV currents based on noble-liquid ionization-chamber theory with simplified detector models of the ATLAS EMEC and FCal1 calorimeters, as well as a discussion on the expected leading sources of systematic uncertainty and non-linearities. The leading systematics were predicted to originate from ion recombination, detector activation, space charge formation in the liquid argon, acceptance effects due to possible displacements of the mean interaction point along the beam axis, and LAr temperature variations. Additional non-linearities due to positive-ion buildup have also been studied in the context of calorimeter signal degradation [282], which may also be relevant to LAr gap-current measurements at high luminosity.

The LAr gap-current method was first studied in practice with a reproduction of a small section

of the FCal using test-beam data recorded in 2008 at the U-70 accelerator at IHEP in Protvino, Russia [283]. The HV current measured with the FCal prototype was found to be linear with the beam intensity up to 0.36 % at a beam intensity of 10^9 protons/spill, corresponding to the nominal ATLAS luminosity of $10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$.

Measurements using LAr gap currents have been used extensively since Run 1 of the LHC for both online luminosity monitoring and studies of the long-term stability of the primary ATLAS luminosity algorithms. The FCal gap-current luminosity was first included in the stability analysis of the 2010 and 2011 *pp* datasets at $\sqrt{s} = 7\,\mathrm{TeV}$ [205]. The EMEC gap-current data were also later incorporated into the stability analysis of the 2012 *pp* dataset at $\sqrt{s} = 8\,\mathrm{TeV}$ [206]. The methodology of the LAr gap-current luminosity measurements has remained mostly unchanged since its first use in Run 1 with only a few noteworthy exceptions. The first concerns the choice of HV channels to include in the final luminosity value averaged over a detector subsystem. For example, the FCal HV channel list was updated year by year to remove shorted HV channels. The selection of HV channels is discussed in Section A.3. The second is the pedestal-subtraction working point, which evolved over the course of Run 1 and Run 2 to increase its robustness against detector-activation effects and "false starts" of the LHC before declaring stable beams. These pedestal-subtraction working points are discussed further in Section A.4. Finally, a correction was developed and applied to the Run 2 FCal gap-current luminosity data to account for a non-linearity in its response. This FCal correction is discussed in Section A.6.

A.2 Data samples

Table A.1 summarizes the ATLAS Run 2 data samples for which the LAr gap-current luminosity has been determined. This table gives selected LHC parameters per year for each combination of colliding-particle species and centre-of-mass energies in Run 2. The majority of the Run 2 ATLAS physics program is concerned with the $\sqrt{s} = 13\,\mathrm{TeV}$ *pp* data samples, therefore this note will pay greater attention to this dataset, with a summary of the results given for the other datasets listed.

The LHC achieved steadily increasing luminosities over the course of Run 2. This is illustrated in Figure A.2, which shows the peak instantaneous luminosity delivered to ATLAS during stable beams as a function of time in each year of data-taking for *pp* collisions at $\sqrt{s} = 13\,\mathrm{TeV}$. These evolving LHC conditions resulted in a broad dynamic range of luminosities delivered to ATLAS, highlighting the need for high-precision and linear luminosity detectors and algorithms. The challenges imposed by these LHC conditions on the LAr gap-current luminosity, and the adopted solutions, are discussed in the sections below.

A.3 The LAr high-voltage system

The installation, commissioning and operation of the LAr high-voltage system is documented extensively in Ref. [284]. A brief overview of the primary components relevant to the LAr gap-current luminosity is given below.

The HV system consists of a set of HV power supplies located in the ATLAS service cavern (USA15), feedthrough cables that carry the high voltage to the ATLAS experimental cavern (UX15) and into the LAr cryostat, and a number of both hardware- and software-based control systems to

Table A.1: Data samples by colliding-particle species and centre-of-mass energy in Run 2 for which the LAr gap-current luminosity has been determined. Selected LHC parameters are also given per data-taking year. The "Bunch config." column indicates the number of colliding bunches and the bunch spacing. The notation "8b4e" signifies a repeating bunch train pattern of eight bunches spaced by 25 ns followed by a gap of four empty bunch slots. For heavy-ion running (Pb+Pb and p+Pb), various bunch configurations were used, typically with an individual bunch followed by alternating empty bunch-slot gaps of 2, 3, 5 or 8 bunch-slots. For proton-lead running, the dataset is subdivided into p+Pb and Pb+p periods, indicating the particle species in LHC beam 1 and beam 2, respectively. $\mathcal{L}_{\text{peak}}$ indicates the peak stable luminosity delivered per year and $\mathcal{L}$-weighted $\langle\mu\rangle$ indicates the luminosity-weighted average number of interactions per bunch crossing in a given year.

Dataset	Year	Run numbers	LHC fills	Bunch config.	$\mathcal{L}_{\text{peak}}$ ($\text{cm}^{-2}\,\text{s}^{-1}$)	$\mathcal{L}$-weighted $\langle\mu\rangle$
pp, $\sqrt{s} = 13\,\text{TeV}$	2015	266904–284484	3819–4569	≤2232, 25 ns	5.02×10^{33}	13
	2016	296939–311481	4851–5451	≤2208, 25 ns	13.8×10^{33}	25
	2017	324320–341649	5698–6417	≤2544, 25 ns / 8b4e	20.9×10^{33}	38
	2018	348197–364485	6570–7358	≤2544, 25 ns	21.0×10^{33}	36
pp, $\sqrt{s} = 5.02\,\text{TeV}$	2015	286282–286474	4634–4647	207–1813, 25 ns	0.379×10^{33}	1.1
	2017	340634–341184	6380–6399	512–1824, 8b4e	1.34×10^{33}	2.1
Pb+Pb, $\sqrt{s} = 5.02\,\text{TeV}$	2015	286665–287931	4658–4720	≤492, various	2.88×10^{27}	0.0018
	2018	365498–367384	7427–7492	≤733, various	6.17×10^{27}	0.0026
p+Pb, $\sqrt{s} = 8.16\,\text{TeV}$	2016	313062–313435	5519–5538	≤474, various	6.22×10^{29}	0.16
Pb+p, $\sqrt{s} = 8.16\,\text{TeV}$	2016	313572–314170	5545–5573	≤513, various	8.62×10^{29}	0.19

monitor and adjust all of the HV system's operational parameters.

Each HV power supply unit contains two boards of 16 channels (8 channels for the FCal). The power supplies generate high voltage for each channel and are regulated around the set value via a voltage comparator by injecting current whenever a voltage drop occurs in any of the LAr gaps fed by the corresponding HV channel. Each HV power supply board includes three primary hardware components: two potentiometers, a safety loop connector and a 32-pin connector. The potentiometers are used to set the hardware voltage and current limits over the whole module range. If either the voltage or current exceed this limit, the entire unit is automatically switched off. The safety loop connector, which is connected to an interlock system for safety controls, allows for the safe discharge of the calorimeter in case of emergency, such as cooling or power failures. The interlock loop is closed only if an HV cable is plugged into the unit, so that the power supply will ramp down to 0 V if a cable is accidentally disconnected. The 32-pin connector contains lines that carry the high voltage to the detector and the associated HV return lines, and includes module-specific interlock pins to avoid accidentally connecting cables meant for a specific detector subsystem to the wrong HV power supply board.

A set of HV cables, each approximately 120 m in length, carry the high voltage from the power supplies to the detector. These cables are connected to the three cryostats via six high-voltage feedthroughs (HVFT) [137], two for the barrel and two for each endcap cryostat. The purpose of the HVFTs is to allow the transfer of the warm HV lines into the the liquid-argon cryostats at 88 K. The HVFTs are kept at room temperature, and the HV cables pass from a warm feedthrough plate through a buffer of gaseous argon, which separates the plate from the liquid argon below, before entering the cryostat.

Dedicated components within the HV power supply boards perform the voltage and current measurements for each channel, with a resolution of 12 mV and 0.12 µA, respectively. These mea-

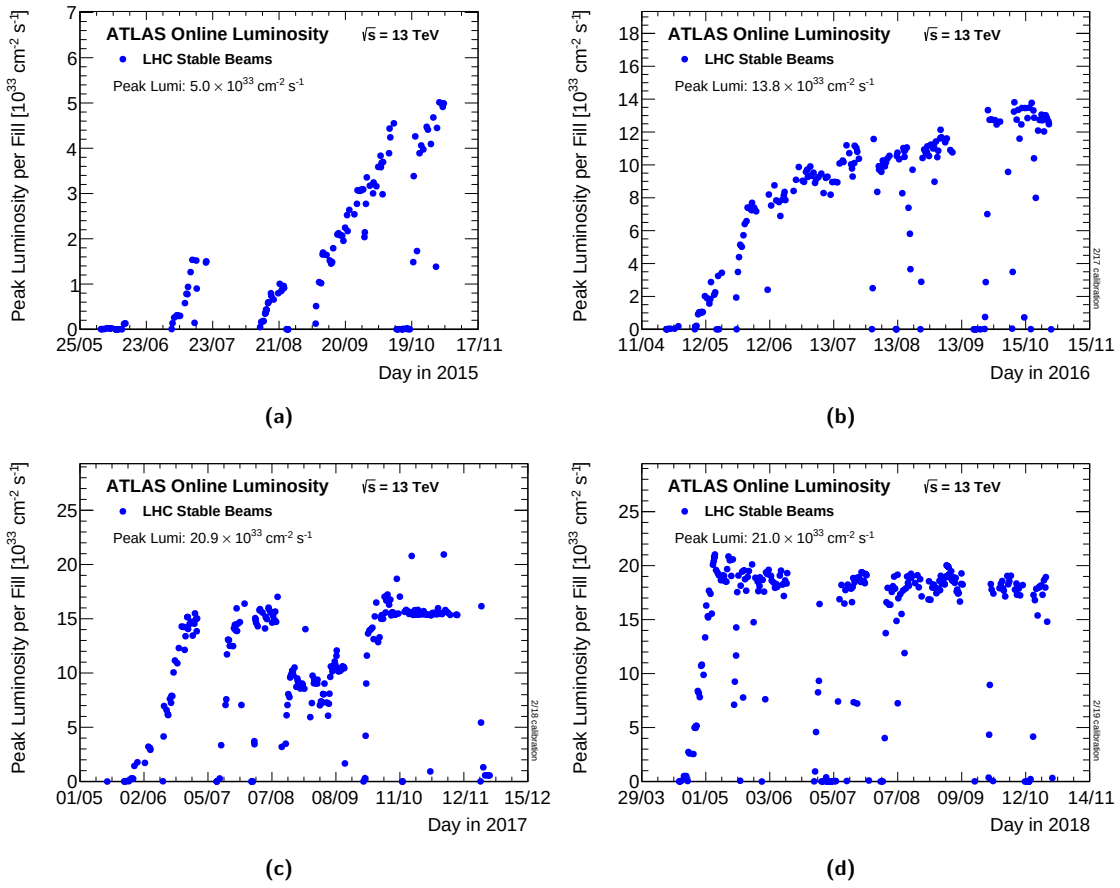

Figure A.2: Peak instantaneous luminosity delivered to ATLAS during stable beams for *pp* collisions at $\sqrt{s} = 13\,\text{TeV}$ for each LHC fill as a function of time in (a) 2015, (b) 2016, (c) 2017 and (d) 2018. The luminosity values shown here are online measurements recorded by the LUCID detector. They are based on initial estimates from van-der-Meer beam-separation scans that are updated year by year, and may differ from the final offline luminosity calibration.

surements are sent to the high-voltage control machines via a CAN-bus[1] interface where they are sent to the ATLAS Detector Control System (DCS) [285] for online monitoring and to be archived in the DCS Oracle database. The rate at which the voltages and currents are sampled and archived can be adjusted via software, however, in practice these sampling rates can only be set for an entire LAr detector subsystem. For the LAr gap-current luminosity in Run 2, the EMEC HV currents were sampled at a rate of approximately 12/min, and the FCal HV currents at a rate of approximately 20/min. A lower sampling rate was chosen for the EMEC since it has several times more HV channels than the FCal, and a higher rate would strain the DCS system.

Choice of HV channels for the luminosity measurement

In an ideal system, all HV channels serving a detector subsystem would be used to increase the statistical precision of the luminosity measurement. In practice, however, only a subset of EMEC

[1] A Controller Area Network (CAN) bus is a communication system standard that allows the transfer of data between components and for microcontrollers and devices to communicate with each other's applications. In the case of the LAr HV system, the CAN-bus interface enables the control of the HV power supplies through a PC or similar controller.

and FCal HV channels can be used, excluding channels that have a poor signal-to-noise ratio, channels that serve shorted electrodes, or channels that exhibit clearly poor linearity with a reference luminosity algorithm in a single representative LHC physics fill.

For the FCal, only the HV channels serving the FCal1 module (the "electromagnetic layer" of the FCal) are considered, given that most of the energy from minimum-bias events is deposited in this module. The FCal gap-current luminosity includes all FCal1 HV channels, except those serving tube groups containing one or more shorted tubes. A short occurs when impurities in the liquid argon connect the central rod of the tube to the outer edge of electrode. This effectively closes the circuit and the current fed into this tube by the HV system flows continuously to the ground. Since tubes are fed as a group, the current injected by the HV line includes a DC component if the group contains a shorted tube. This DC current is prone to unpredictable fluctuations due to the very nature of the shorts, hence it is prudent to exclude any of these HV lines from the luminosity analysis.

A much smaller subset of EMEC HV channels satisfy the channel-selection criteria listed above. Many EMEC HV channels are only connected to a few electrodes, resulting in a poor signal-to-noise ratio and a non-linear, sporadic response, therefore the channels connected to a greater number of electrodes (typically 24 electrodes) are preferred. Based on these criteria, eight HV channels in each of the EMEC-A and EMEC-C modules were selected. These channels serve the B3 HV sector located in the region $1.8 < |\eta| < 2.0$ (see Figure **??**). The relative position of the electrodes served by these HV channels within each EMEC module on the A- and C-sides of the detector are equivalent as viewed from the interaction point.

The complete list of HV channels included in the Run 2 EMEC and FCal gap-current luminosity measurements is given in Ref. [286].

A.4 Pedestals

Electronics pedestals in the HV system and currents induced in the HV lines due to non-collision backgrounds, collectively referred to as "pedestal" currents, must be subtracted from the recorded HV currents in order to ensure an accurate luminosity measurement. The pedestals can drift over time and must therefore be estimated separately for each LHC fill. In Run 2, the pedestals were measured using data taken at the start of each fill immediately before the beams were put into collision. The exact methodology used to compute the pedestal evolved over the course of Run 2 to improve its accuracy and its robustness against non-ideal data-taking conditions. Three pedestal-subtraction working points were developed: 1 h, 10 min and 10 min++.

The 1 h working point takes the pedestal as the average current per HV line over a 1 h period before stable beams are declared. First, the end point of the 1 h period t_{end} is determined by scanning backward in time from the time at which stable beams are declared until the current in the given HV line drops below 0.2 µA, which typically occurs a few minutes before stable beams during the ADJUST beam mode. The start point t_{start} is then determined given the sampling rate of the HV current such that $t_{\text{start}} = t_{\text{end}} - 1\,\text{h}$, which amounts to 720 samples for the EMEC and 1200 for the FCal.

As the operation of the LHC progressed, with increasing luminosities and shorter intervals between fills, it was found that residual current remained in the HV lines after the end of a fill, which

decayed exponentially over time, resulting in an artificially high pedestal value. The residual current is understood to arise from induced radioactivity, or *activation*, in the detector material due to the large flux of high-energy particles during LHC operation. The 1 h working point was therefore eventually discarded in favour of the 10 min working point, which uses the same t_{end} as before, but instead the start point is taken as $t_{\text{start}} = t_{\text{end}} - 10\,\text{min}$, or 120 samples for the EMEC and 200 for the FCal. The shorter interval of this working point improves the robustness against activation effects by averaging over a period where the HV currents are in a steady state while still giving reasonably good statistical precision.

The 10 min working point failed in a few isolated cases where collisions occurred between t_{start} and t_{end}, for example, during so-called LHC "false starts" where the procedure to bring the beams into collision was halted midway through for readjustments before bringing them back into collision and declaring stable beams. The 10 min working point in this case would include the false-start region between t_{start} and t_{end}, since the HV current would drop below 0.2 µA between the false start and the point when stable beams are declared, resulting in an artificially high pedestal value. A new working point, called 10 min++, was developed to handle cases like these by employing a more robust algorithm to locate t_{end}. The value of t_{end} is determined by traversing forward in time from the point at which the LHC enters the ADJUST beam mode up to, but not including, the luminosity block where the online ATLAS preferred luminosity rises above a certain threshold, depending on the colliding-particle species. For *pp* fills, this threshold was set at $10^{30}\,\text{cm}^{-2}\,\text{s}^{-1}$, while for heavy-ion fills, which ran at lower luminosities, the thresholds were set at $10^{27}\,\text{cm}^{-2}\,\text{s}^{-1}$ and $2 \times 10^{26}\,\text{cm}^{-2}\,\text{s}^{-1}$ for p+Pb and Pb+Pb fills, respectively. As before, the start point is taken as $t_{\text{start}} = t_{\text{end}} - 10\,\text{min}$. In this way, the pedestal is still evaluated during an interval sufficiently close in time to the beginning of stable beams, while excluding any such false starts. Figure A.3 gives an illustrative comparison of the three pedestal-subtraction working points and how the 10 min++ working point offers the most robust pedestal estimate in non-ideal data-taking conditions, and Figure A.4 shows the HV-current values selected for the pedestal estimate using this working point in a typical *pp* fill from 2018 for a few representative EMEC and FCal HV channels. The exponential decay over time of the HV current between LHC fills is also clearly visible in this figure.

The 10 min++ working point proved to provide a robust and accurate pedestal estimate and was therefore taken as the nominal pedestal-subtraction method in Run 2, with all LAr gap-current luminosity data having been reprocessed with this working point.

Figure A.5 shows the mean pedestal per detector subsystem as a function of time per LHC fill during running with *pp* collisions at $\sqrt{s} = 13\,\text{TeV}$ in each data-taking year. These figures illustrate the drift of the pedestal values over time, and that they roughly correspond to the yearly luminosity profiles shown in Figure A.2. This correlation between the LAr gap-current pedestals and the yearly luminosity profiles suggest an activation component with a half life on the order of several days is present, in addition to the activation component with a relatively short half life ($\sim 1\,\text{h}$) visible in Figure A.4.

A.5 Calibration

Since the LAr gap-current luminosity cannot be calibrated absolutely using a method such as a vdM scan, it must instead be calibrated relative to an independent luminosity algorithm whose absolute

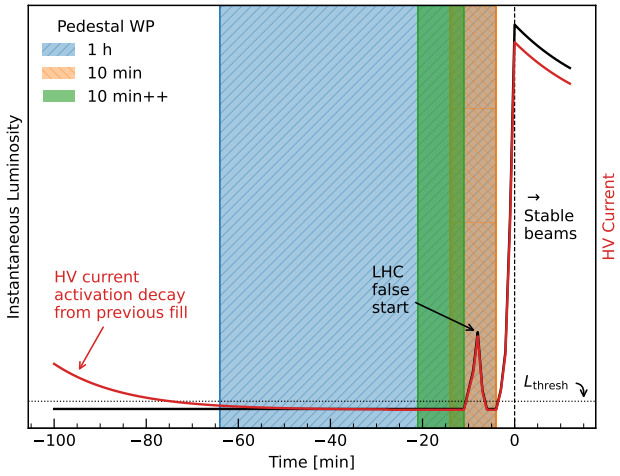

Figure A.3: Illustration of the 1 h, 10 min and 10 min++ LAr gap-current pedestal-subtraction working points in a non-ideal physics fill beginning with an LHC "false start". The beginning of stable beams is indicated by the vertical dashed line at $t = 0$. Both the 1 h (blue-shaded area) and the 10 min (orange-shaded area) working points incorrectly include the false-start region in which collisions occurred, with the 1 h working point also including the tail of the HV current activation decay from the previous fill, leading to artificially high pedestal estimates. The 10 min++ working point (green-shaded area) correctly selects the interval in which no collisions occurred immediately before the false start, and in which the HV current was in a steady state, by traversing forward in time until the online ATLAS preferred luminosity rises above a certain threshold, $\mathcal{L}_{\text{thresh}}$.

calibration has already been determined. The nominal LAr gap-current calibration assumes that the pedestal-corrected HV currents are directly proportional to the luminosity (with no constant offset).

The proportionality constant ζ_c is determined separately for each HV channel from a fit of the pedestal-corrected HV current I_c vs. a reference luminosity algorithm $\mathcal{L}_{\text{ref}}$ in a single long high-luminosity physics fill in the form:

$$I_c = \frac{1}{\zeta_c}\mathcal{L}_{\text{ref}}\,. \tag{A.1}$$

The per-channel LAr gap-current luminosity is then given by $\mathcal{L}_c = \zeta_c I_c$, and the luminosity for a given detector subsystem is taken as the average over the HV channels considered for the luminosity measurement:

$$\mathcal{L}_{\text{LAr}} = \langle \zeta_c I_c \rangle_{\text{channels}}\,. \tag{A.2}$$

In general, the offline ATLAS preferred luminosity at the time of the physics fill in question is used as the reference luminosity algorithm in the fit. The per-channel proportionality constants must also be derived for each combination of colliding-particle species and centre-of-mass energies to account for the different cross sections of these collisions.

Table A.2 summarizes the selected high-luminosity physics fills and reference luminosity algorithms used to derive the per-channel calibration constants for each combination of colliding-particle species and centre-of-mass energies in each data-taking year of Run 2. Examples of the fits to the pedestal-corrected LAr HV currents vs. instantaneous luminosity measured by LUCID BiHitOR to derive the 2015–2016 $\sqrt{s} = 13\,\text{TeV}$ pp calibration constants are shown in Figure A.6. Similarly, Figure A.7 shows the equivalent set of fits used to derive the 2018 Pb+Pb calibration constants. Table A.2 also lists the reference luminosity algorithms and runs used to extract the FCal-correction

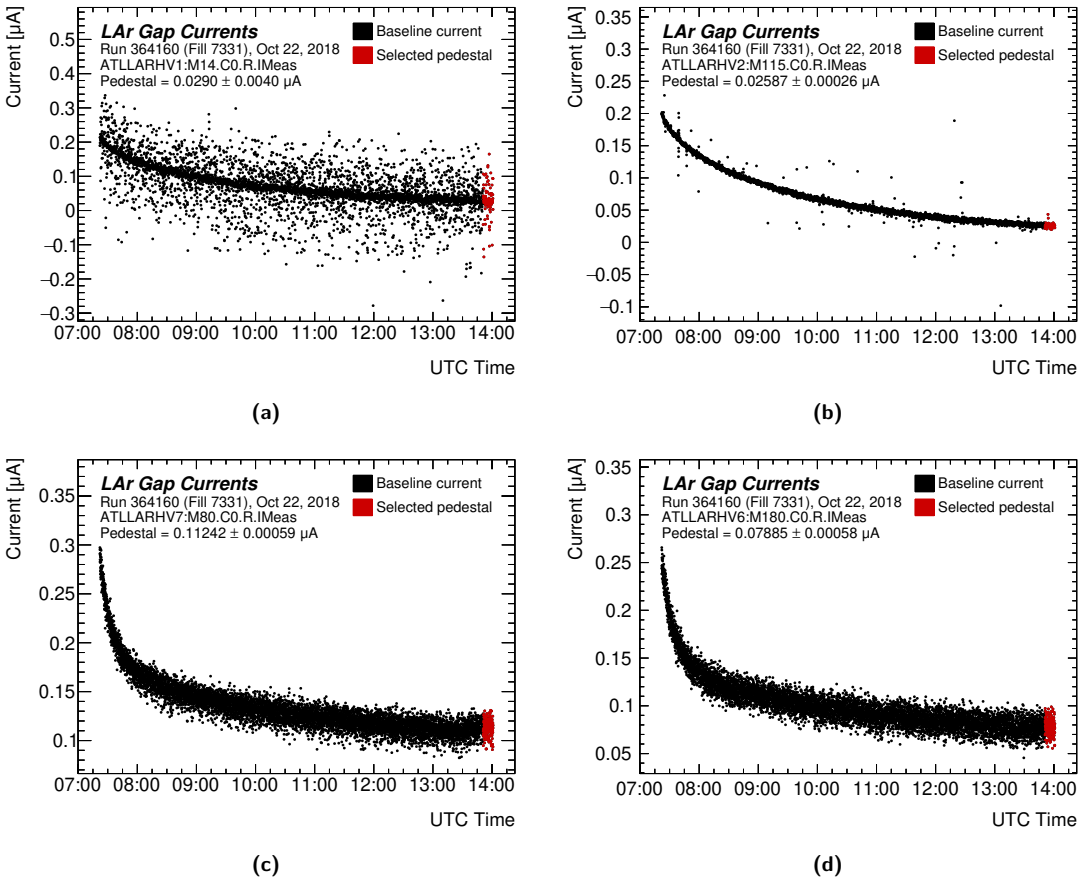

Figure A.4: HV currents as a function of time during the period without collisions immediately before the beams were put into collision in fill 7331. Shown are two representative EMEC HV channels (a) M14C0 and (b) M115C0, and two FCal channels (c) M80C0 and (d) M180C0. The red points indicate the samples used to compute the LAr gap-current pedestals for this fill using the 10 min++ working point. The exponential decay in the HV current following the previous fill is clearly visible.

parameters, which are used for a higher-order calibration of the FCal gap-current luminosity to correct for an observed non-linearity. The FCal non-linearity correction is described in detail below in Section A.6.

The per-channel calibration constants were obtained in a similar manner for the remaining datasets, with a few notable differences to account for the different running conditions and available data at the time of the analysis. First, track counting was used as the reference luminosity algorithm to derive the calibration constants for the $\sqrt{s} = 5.02\,\text{TeV}$ pp dataset. Studies of LUCID's μ-dependence in this regime were still underway at the time of the analysis and track counting was believed to serve as the most reliable reference luminosity algorithm. The run-integrated track-counting luminosity was scaled to agree with the ATLAS preferred luminosity, measured using the LUCID BiEvtOR algorithm, in the reference fill in order to set the absolute scale of the track-counting measurement. Both the 2015 and 2017 data-taking periods used the same reference fill since the detector and running conditions were comparable in both years.

The calibration constants for the p+Pb and Pb+p datasets had to be determined separately due to the asymmetric collisions, which resulted in different particle fluxes and induced HV currents

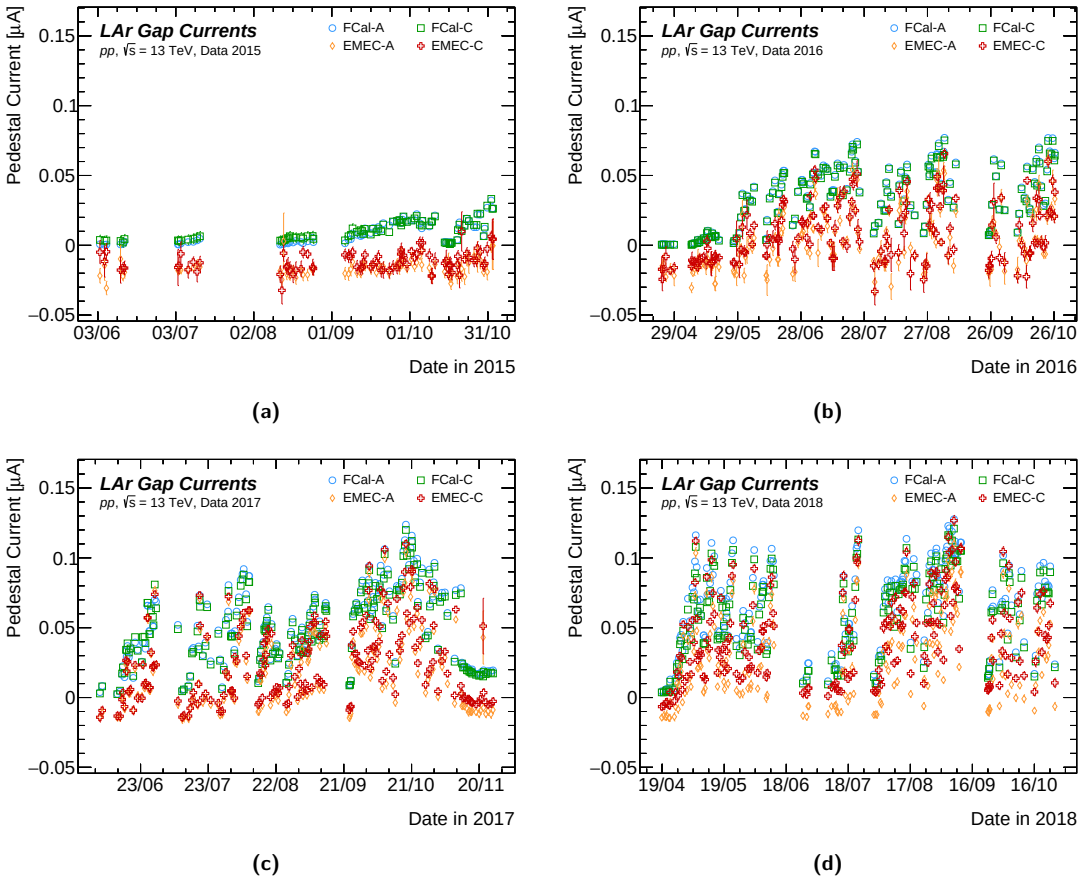

Figure A.5: LAr gap-current pedestals per LHC fill during running with pp collisions at $\sqrt{s} = 13\,\text{TeV}$ in (a) 2015, (b) 2016, (c) 2017 and (d) 2018. The pedestal currents, determined using the 10 min++ working point (see text), are averaged over all HV channels per detector subsystem.

Table A.2: Summary of the selected high-luminosity physics fills and reference luminosity algorithms $\mathcal{L}_{\text{ref}}$ used to derive the per-channel LAr gap-current luminosity calibration constants ζ_c for each combination of colliding-particle species and centre-of-mass energies in each data-taking year of Run 2. In all cases, "Tracks" refers to track counting using the 2017 track-selection working point, and "Tracks*" indicates that the run-integrated track-counting luminosity was normalized to LUCID BiEvtOR in the indicated physics fill. The reference inelastic cross sections σ_{inel} are also listed for converting between instantaneous luminosity and μ. The "FCal correction" column indicates the reference luminosity algorithms and runs used to extract the FCal non-linearity-correction parameters (see text).

Dataset	σ_{inel} [mb]	Year(s)	Run number (LHC fill)	$\mathcal{L}_{\text{ref}}$ algorithm	FCal correction
pp, $\sqrt{s} = 13\,\text{TeV}$	80.0	2015–2016	310634 (5416)	LUCID BiHitOR	Tracks in run 310634
		2017–2018	305618 (5163)	LUCID BiHitOR	Tracks in run 331085
pp, $\sqrt{s} = 5.02\,\text{TeV}$	69.0	2015, 2017	340718 (6385)	Tracks*	Tracks* in run 340718
Pb+Pb, $\sqrt{s} = 5.02\,\text{TeV}$	7650	2015	287931 (4720)	LUCID BiEvtOR	LUCID BiEvtOR in run 287931
		2018	367099 (7480)	LUCID BiEvtOR	LUCID BiEvtOR in run 367099
p+Pb, $\sqrt{s} = 8.16\,\text{TeV}$	2160	2016	313107 (5522)	LUCID BiEvtOR	LUCID BiEvtOR in run 313107
Pb+p, $\sqrt{s} = 8.16\,\text{TeV}$		2016	314157 (5571)	LUCID BiEvtOR	LUCID BiEvtOR in run 314157

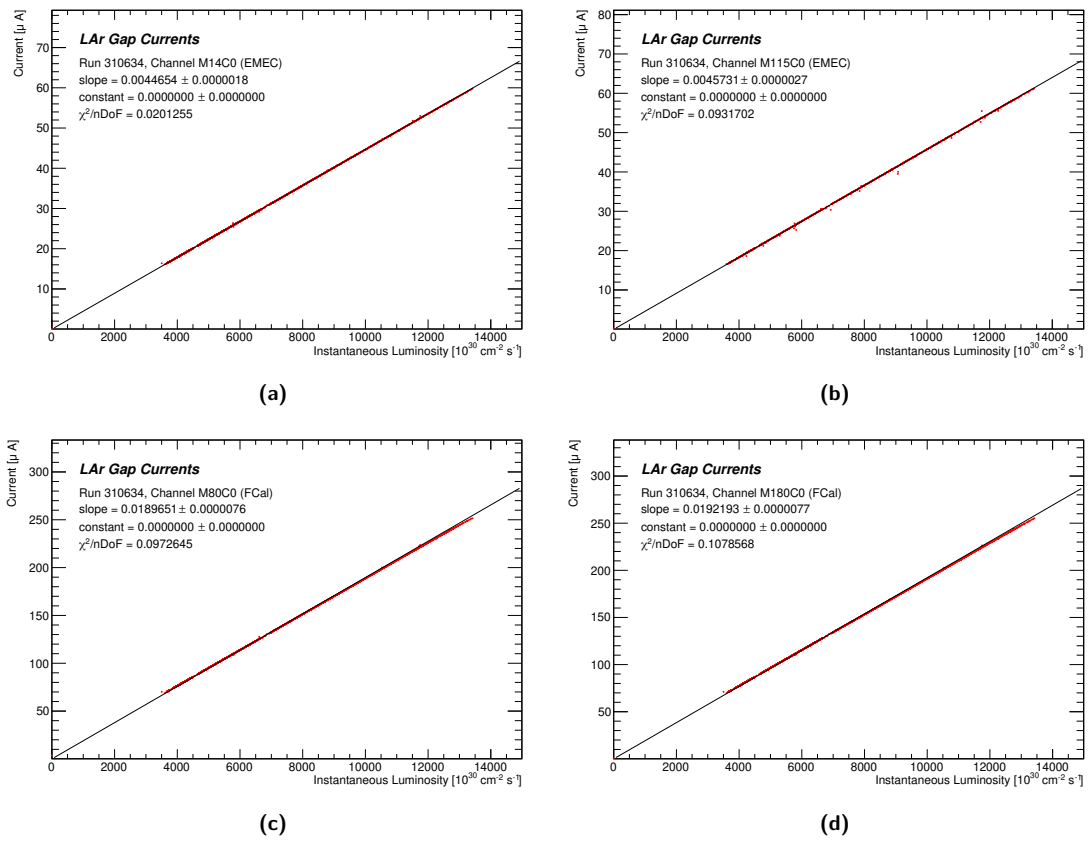

Figure A.6: Pedestal-corrected LAr HV currents vs. instantaneous luminosity measured by LUCID BiHitOR in run 310634 (from the 2016 $\sqrt{s} = 13\,\text{TeV}$ pp dataset) for two representative EMEC HV channels (a) M14C0 and (b) M115C0, and two FCal channels (c) M80C0 and (d) M180C0. The black lines show the fits used to extract the per-channel proportionality constants ζ_c, where the slope of the line is equal to $1/\zeta_c$.

on the A- and C-sides of the detector depending on the particle species in each of the two LHC beams. As with the Pb+Pb datasets, the ATLAS preferred luminosity measured using the LUCID BiEvtOR algorithm were used as the reference luminosity algorithm since the μ-dependence of LUCID is negligible over the range in μ delivered during heavy-ion running.

A.6 The FCal non-linearity correction

A non-linearity was observed in the FCal gap-current luminosity during Run 2 data-taking in which the FCal underestimates the luminosity at high instantaneous luminosities,[2] requiring a correction of $O(1.5\,\%)$ during typical physics fills. This non-linearity is clearly visible when plotting the ratio of the FCal gap-current luminosity to track counting as a function of instantaneous luminosity in a typical physics fill, as in Figure A.8.

The non-linearity is believed to arise from a combination of phenomena whose effect on the LAr

[2]Since the gap-current luminosity is a relative luminosity algorithm, it can be freely scaled, or "anchored", such that its run-integrated measurement agrees with another luminosity algorithm at some other point in time during a data-taking year (as it is in the stability plot, for example). For this reason, it is equivalent to say that the uncorrected FCal gap-current luminosity *overestimates* the luminosity at *low* instantaneous luminosities if the FCal measurement is anchored in a high-luminosity fill.

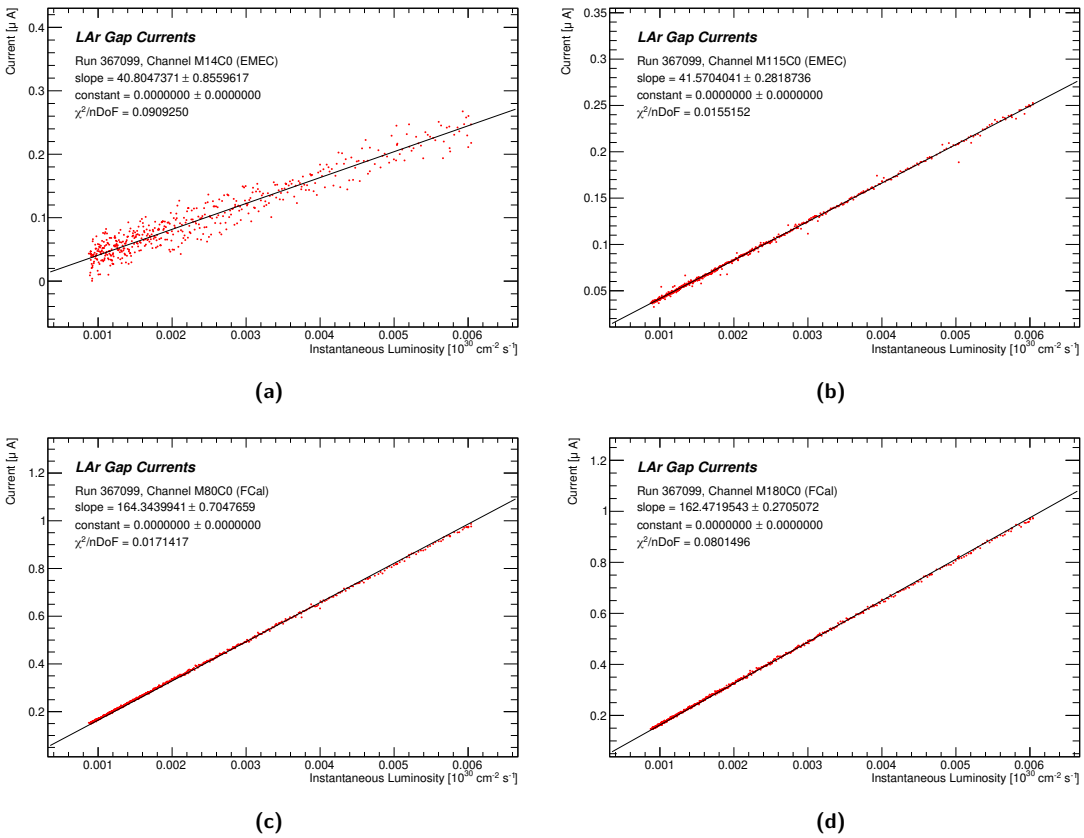

Figure A.7: Pedestal-corrected LAr HV currents vs. instantaneous luminosity measured by LUCID BiEvtOR in run 367099 (from the 2018 $\sqrt{s} = 5.02\,\text{TeV}$ Pb+Pb dataset) for two representative EMEC HV channels (a) M14C0 and (b) M115C0, and two FCal channels (c) M80C0 and (d) M180C0. The black lines show the fits used to extract the per-channel proportionality constants ζ_c, where the slope of the line is equal to $1/\zeta_c$.

gap currents becomes non-negligible at higher luminosities, mainly the buildup of space charge in the liquid argon, which distorts the electric field across the gaps, and high-voltage "sagging", in which the high voltage applied across the gaps drops due to increased current drawn over the protection resistors. However, in the absence of a specific model that accounts for this non-linearity, a correction was derived similar in principle to the LUCID calibration transfer, following the assumption that track counting is linear with luminosity and that the ratio $\mathcal{L}_{\text{FCal}}/\mathcal{L}_{\text{Tracks}}$ can be parameterized as an n^{th}-order polynomial:

$$\frac{\mathcal{L}_{\text{FCal}}^{\text{uncorr}}}{\mathcal{L}_{\text{Tracks}}} = p_0 + p_1\mathcal{L}_{\text{FCal}}^{\text{uncorr}} + \cdots + p_n(\mathcal{L}_{\text{FCal}}^{\text{uncorr}})^n\,, \tag{A.3}$$

where $\mathcal{L}_{\text{FCal}}^{\text{uncorr}}$ is the uncorrected FCal gap-current luminosity and $p_0, p_1, \ldots, p_n$ are the polynomial coefficients extracted from a fit to data in a single long high-luminosity physics fill. In practice, a second-order polynomial was consistently found to offer the best fit in all settings. The corrected FCal gap-current luminosity should satisfy $\mathcal{L}_{\text{FCal}}^{\text{corr}} = \mathcal{L}_{\text{Tracks}}$ by construction, and it follows that

$$\mathcal{L}_{\text{FCal}}^{\text{corr}} = \frac{\mathcal{L}_{\text{FCal}}^{\text{uncorr}}}{p_0 + p_1\mathcal{L}_{\text{FCal}}^{\text{uncorr}} + p_2(\mathcal{L}_{\text{FCal}}^{\text{uncorr}})^2}\,. \tag{A.4}$$

Additionally, separate fits to the FCal-A and FCal-C data were performed and found to yield sufficiently distinct fit parameters to warrant performing the non-linearity correction separately for the A- and C-side FCal modules.

The FCal non-linearity correction was first applied to the 2017 and 2018 $\sqrt{s} = 13\,\text{TeV}$ pp datasets using the polynomial coefficients extracted from a fit to $\mathcal{L}_{\text{FCal}}^{\text{uncorr}}/\mathcal{L}_{\text{Tracks}}$ vs. $\mathcal{L}_{\text{FCal}}^{\text{uncorr}}$ in run 331085 (LHC fill 6024) recorded between July 29–30, 2017, roughly one day after the 2017 vdM session. These fits are shown in Figure A.8 separately for the FCal-A and FCal-C modules. The 2015 and 2016 $\sqrt{s} = 13\,\text{TeV}$ pp data were later reprocessed with the FCal correction applied using polynomial coefficients extracted from a fit to data in run 310634 (LHC fill 5416) recorded between October 14–15, 2016 (the same fill in which the per-channel calibration constants were determined for this period).

For the heavy-ion datasets, where track-counting luminosity measurements are challenging and have only been made for a handful of fills, a different reference luminosity algorithm was required to perform the FCal non-linearity correction. The offline ATLAS preferred luminosity (as measured by LUCID BiEvtOR for both the Pb+Pb and p+Pb datasets) was suitable for this purpose since the μ-dependence of LUCID is expected to be small over the range of luminosities delivered during heavy-ion running. The same formalism outlined in Eqs. (A.3) and (A.4) was employed for both the Pb+Pb and p+Pb datasets to perform the FCal non-linearity correction, replacing the track-counting luminosity $\mathcal{L}_{\text{Tracks}}$ with the luminosity measured by LUCID BiEvtOR. The same fills in which the per-channel calibration constants were determined were used to extract the polynomial coefficients from fits to data, as shown in Table A.2.

Figure A.8 also shows the equivalent luminosity ratios measured by the EMEC-A and EMEC-C modules to track counting in run 331085. Linear fits to these ratios as a function of instantaneous luminosity would suggest the EMEC gap-current luminosity exhibits a non-linearity as well, however the $O(0.5\,\%)$ correction would be similar in magnitude to the uncertainty on the linearity of track counting over this range of luminosities, as determined from comparisons to other luminosity detectors and algorithms. It is therefore not possible to say with certainty that this non-linearity is a feature of the EMEC or of track counting, hence no non-linearity correction is applied to the EMEC gap-current luminosity in Run 2.

A.7 Results

The linearity and long-term stability of the LAr gap-current luminosity measurements can be evaluated using so-called "stability plots" in which the fractional differences in run-integrated luminosity between the gap-current measurements and a reference luminosity algorithm are plotted as a function of time in a data-taking year or period of interest. Where both the FCal and EMEC gap-current measurements are available, the internal LAr gap-current stability is presented to evaluate the consistency of the FCal-A, FCal-C, EMEC-A and EMEC-C gap-current measurements over time. EMEC gap-current measurements were not possible during certain data-taking periods, primarily during low-luminosity running where the signal-to-noise ratio in the EMEC measurements is poor, and in these cases the stability is shown with respect to the offline ATLAS preferred luminosity as measured by LUCID. In addition to having different reference luminosity algorithms, the stability plots shown below may have subtle differences compared to those appearing in previous publications, including

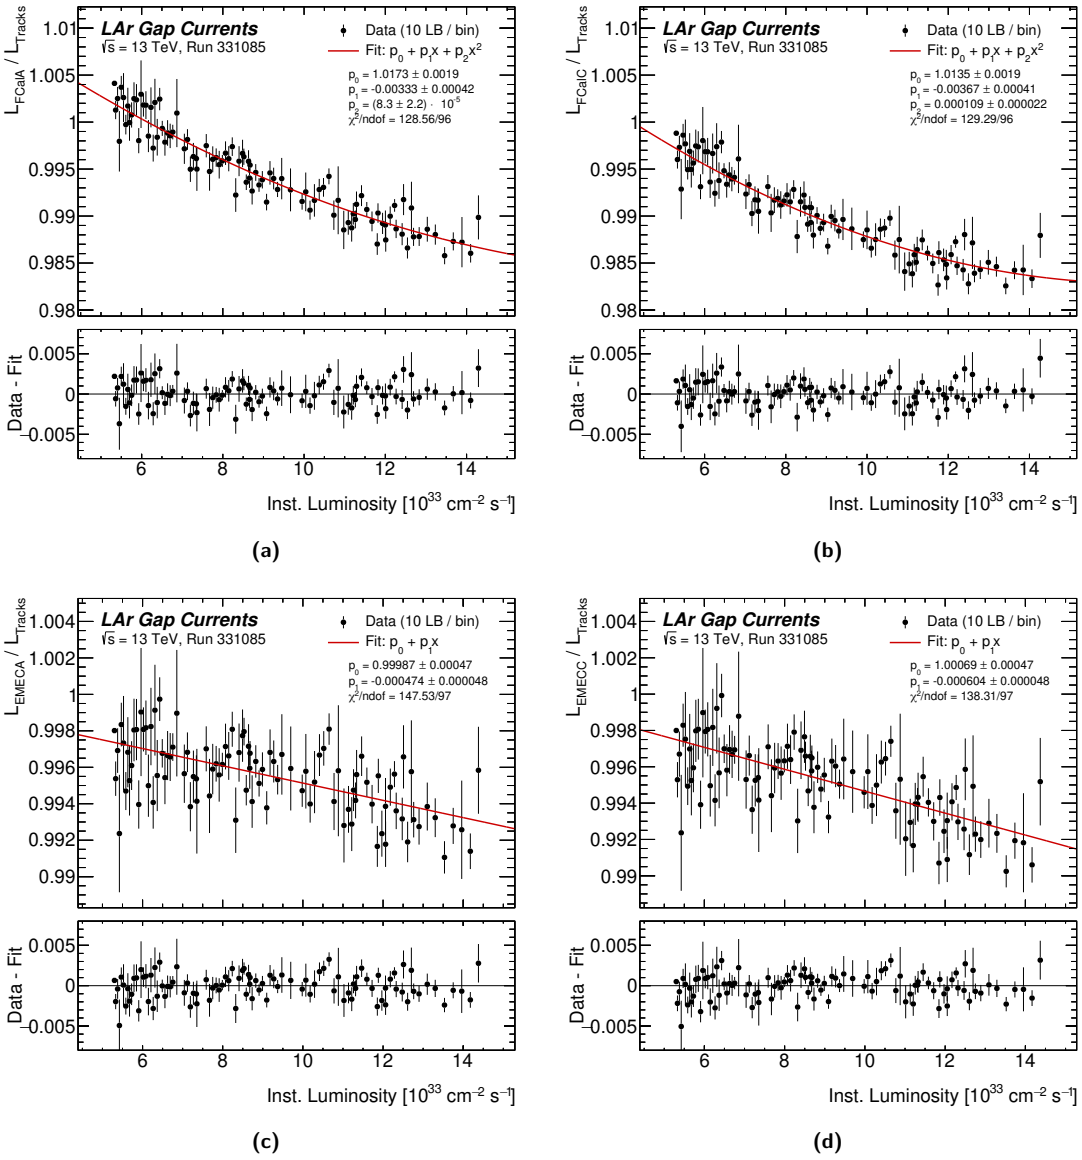

Figure A.8: Ratio of the LAr gap-current luminosity measured by the (a) FCal-A, (b) FCal-C, (c) EMEC-A and (d) EMEC-C to track counting as a function of instantaneous luminosity in run 331085 (LHC fill 6024) from the 2017 $\sqrt{s} = 13\,\text{TeV}$ pp dataset. The data are binned by ten consecutive luminosity blocks in order to estimate the statistical uncertainty on each point. The red lines show second-order polynomial fits to the FCal data and linear fits to the EMEC data. The polynomial coefficients extracted from the fits to the FCal data are used to correct the FCal gap-current luminosity non-linearity.

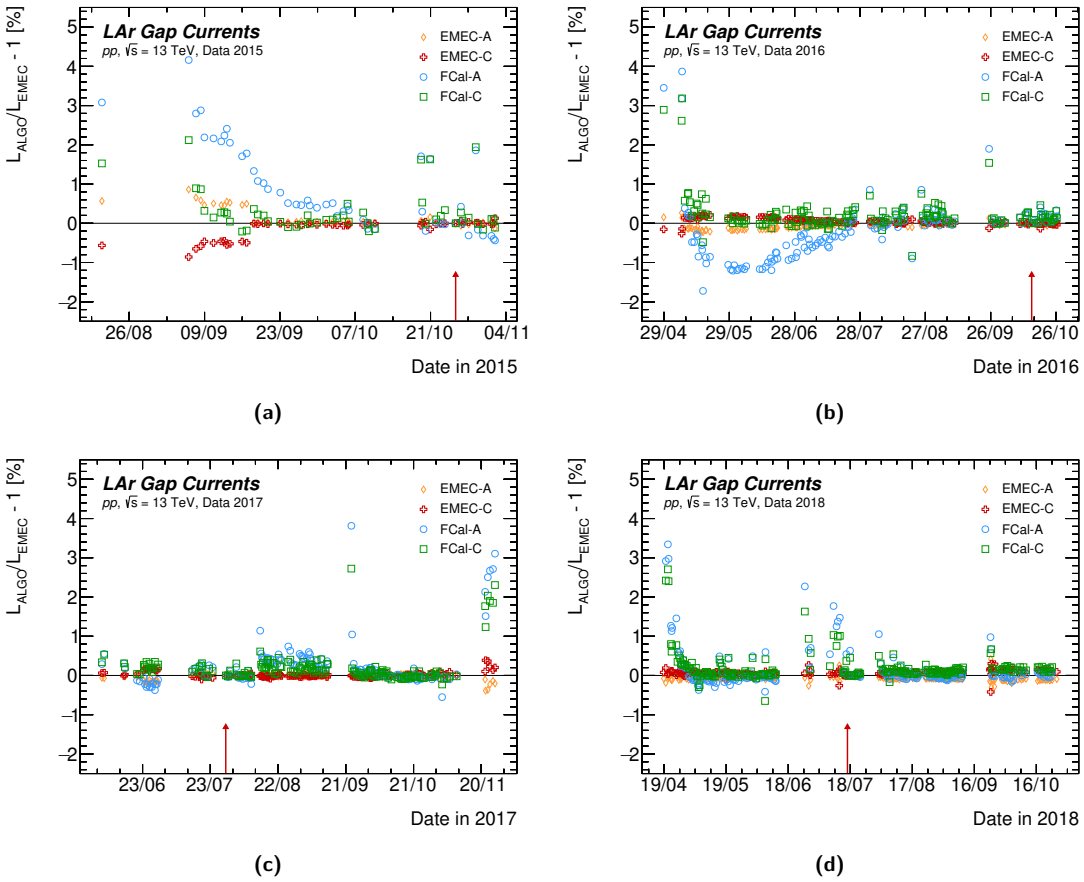

Figure A.9: Fractional differences in run-integrated luminosity between the EMEC gap currents and the FCal-A, FCal-C, EMEC-A and EMEC-C measurements for the $\sqrt{s} = 13\,\text{TeV}$ pp dataset plotted as a function of time in (a) 2015, (b) 2016, (c) 2017 and (d) 2018. The luminosity measurements from the other detectors have been normalized to that of EMEC in the reference fill indicated by the red arrows.

Refs. [115, 207], such as the choice of runs, the luminosity blocks selected within these runs, and the anchoring run in which all the luminosity algorithms are normalized to agree with the reference luminosity algorithm. In general, runs having fewer than 100 consecutive luminosity blocks with stable beams were excluded, and the luminosity blocks within the selected runs were required to have a length greater than 25 s to remove short luminosity blocks where the beams may not have been in head-on collisions (during beam scans, for example).

Figure A.9 shows the internal stability of the gap-current measurements in the $\sqrt{s} = 13\,\text{TeV}$ pp dataset as the fractional differences between the FCal-A, FCal-C, EMEC-A and EMEC-C measurements with respect to the EMEC measurement (averaged over the A- and C-side modules) in each year from 2015 to 2018. The run-to-run agreement between the gap-current measurements is generally at the level of 0.3 % or better for the bulk of the data, with a few exceptions and a number of short- and long-term trends being visible. First, many of the outliers are due to LHC fills with reduced numbers of bunches (and hence low integrated luminosity), typically in intensity ramp-up periods at the beginning of each year and after LHC stops or operational difficulties. Second, there is a set of approximately eleven low-μ fills at the end of 2017 and four fills between July 9 and 14, 2018 in which the FCal gap-current measurements were up to 3 % greater than the EMEC

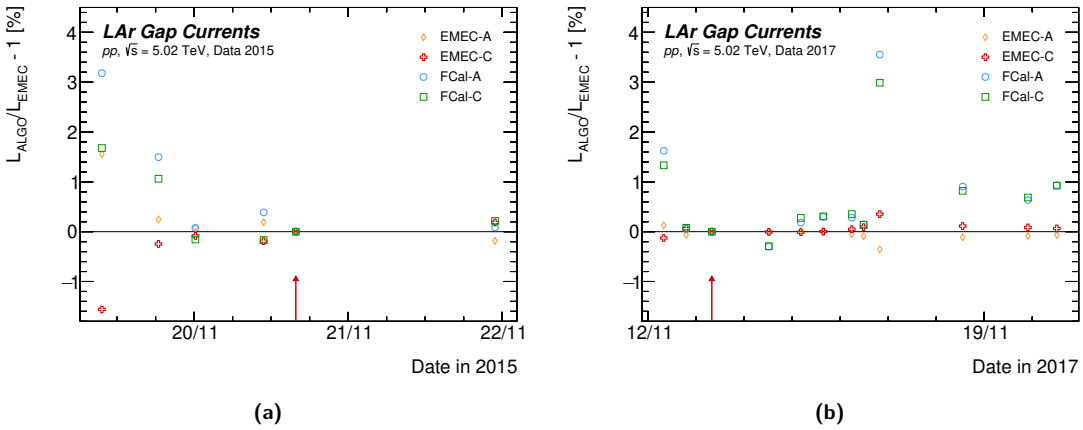

Figure A.10: Fractional differences in run-integrated luminosity between the EMEC gap currents and the FCal-A, FCal-C, EMEC-A and EMEC-C measurements for the $\sqrt{s} = 5.02\,\text{TeV}$ pp datasets plotted as a function of time in (a) 2015 and (b) 2017. The luminosity measurements from the other detectors have been normalized to that of EMEC in the reference fill indicated by the red arrows.

measurements. The luminosity in these fills was levelled to approximately $\langle\mu\rangle = 2$, as part of the precision W/Z physics program, and thus in the regime where the FCal is known to suffer from significant non-linearities. The treatment of the LAr gap-current measurements in the luminosity determination of this low-μ dataset is discussed in detail in Ref. [207]. Third, the FCal measurements were consistently $\sim 0.5\,\%$ greater than the EMEC measurements in mid-2017 between August 14 and September 13 in which the LHC ran at lower luminosities with fewer bunches (in the 8b4e configuration). As before, this lower-luminosity regime is one in which the FCal is known to suffer from non-linearities, hence the disagreement with the EMEC. Finally, there are a number of A- and C-side disagreements, most prominently in the FCal measurements at the beginning of 2015 and 2016. While the FCal-C generally agrees with the EMEC measurements, the FCal-A measurements are up to $2\,\%$ higher in the first half of 2015 followed by a $\sim 1\,\%$ drop in early 2016 before converging with the other gap-current measurements towards the end of July 2016. The possible sources of these A- and C-side disagreements have been studied in detail,[3] however, no credible explanation has been identified by the time of writing. It is therefore generally recommended to exclude the 2015 and 2016 FCal-A data from any study of the long-term stability of another luminosity algorithm in these years.

Figure A.10 shows the internal stability of the gap-current measurements in the $\sqrt{s} = 5.02\,\text{TeV}$ pp dataset as the fractional differences between the FCal-A, FCal-C, EMEC-A and EMEC-C measurements with respect to the EMEC measurement in 2015 and 2017. The run-to-run consistency is generally poorer than in the high-luminosity $\sqrt{s} = 13\,\text{TeV}$ pp datasets, with the bulk of the data having internal inconsistencies of up to $1\,\%$. This is primarily due to the difficulties associated with the FCal non-linearity and the poor signal-to-noise ratio of the EMEC measurements in this low-luminosity regime. The treatment of the LAr gap-current measurements in the luminosity determination of this dataset is discussed in detail in Ref. [207].

Figure A.11 shows the stability of the gap-current measurements in the $\sqrt{s} = 5.02\,\text{TeV}$ Pb+Pb

[3] See the presentation slides at indico.cern.ch/event/1088927/#1-fcal-stability-in-2016-the-f for a summary of the studies performed.

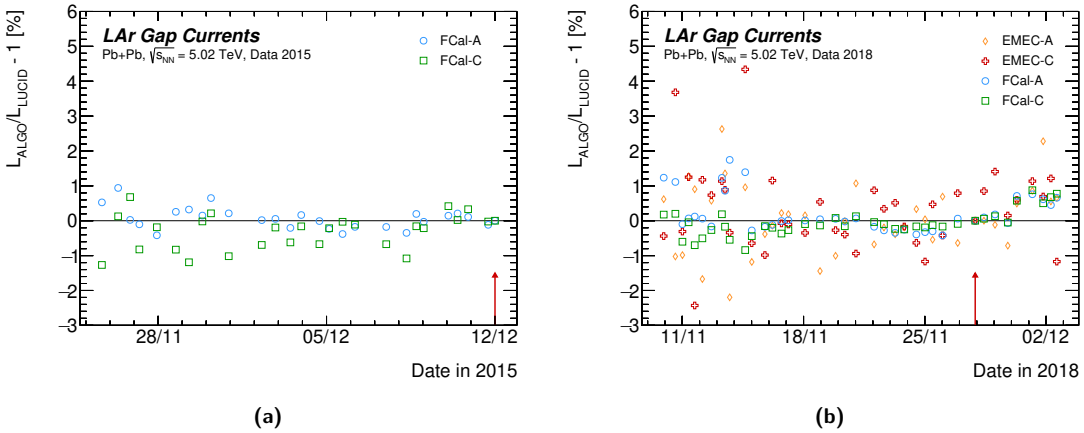

Figure A.11: Fractional differences in run-integrated luminosity between the LUCID BiEvtOR algorithm and the FCal-A, FCal-C, EMEC-A and EMEC-C gap-current measurements for the $\sqrt{s} = 5.02\,\text{TeV}$ Pb+Pb datasets plotted as a function of time in (a) 2015 and (b) 2018. The luminosity measurements from the other detectors have been normalized to that of LUCID BiEvtOR in the reference fill indicated by the red arrows. The EMEC gap-current measurements are not shown in the 2015 plot due to the poor signal-to-noise ratio at the low luminosities delivered during this period.

dataset as the fractional differences between the FCal-A, FCal-C, EMEC-A and EMEC-C measurements with respect to the offline ATLAS preferred luminosity measured by the LUCID BiEvtOR algorithm in 2015 and 2018. The EMEC gap-current measurements are not shown in the 2015 plot due to the poor signal-to-noise ratio at the low luminosities delivered during this period. The FCal measurements are generally consistent with those of LUCID at the percent level in 2015 and less so in 2018, with the EMEC showing greater run-to-run statistical fluctuations of up to $\pm 2\,\%$ in 2018.

Finally, Figure A.12 shows the stability of the gap-current measurements in the $\sqrt{s} = 8.16\,\text{TeV}$ p+Pb and Pb+p datasets as the fractional differences between the FCal-A and FCal-C measurements with respect to the offline ATLAS preferred luminosity measured by the LUCID BiEvtOR algorithm in 2016. As before, the EMEC gap-current measurements are not shown in these plots due to the poor signal-to-noise ratio at the low luminosities delivered during these periods. The bulk of the FCal data again show an agreement at the percent level or better with LUCID, with the outliers generally being a consequence of running at lower luminosity.

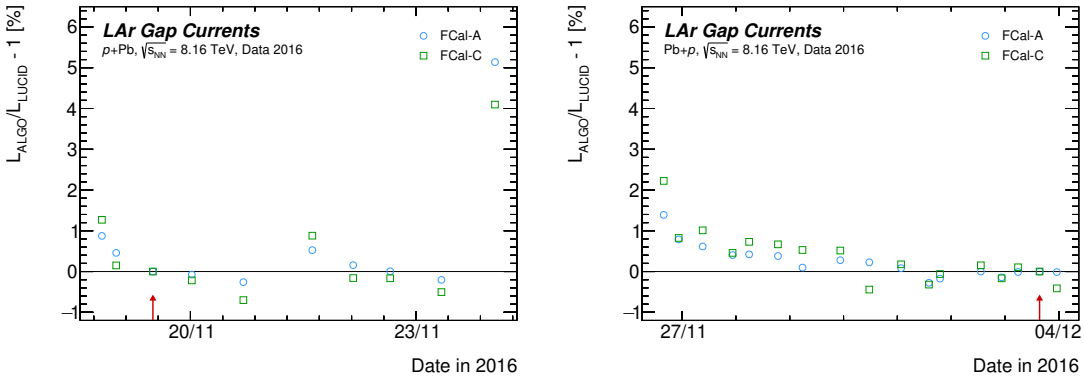

Figure A.12: Fractional differences in run-integrated luminosity between the LUCID BiEvtOR algorithm and the FCal-A and FCal-C gap-current measurements for the $\sqrt{s} = 8.16\,\text{TeV}$ (a) p+Pb and (b) Pb+p datasets plotted as a function of time in 2016. The luminosity measurements from the FCal have been normalized to that of LUCID BiEvtOR in the reference fill indicated by the red arrows. The EMEC gap-current measurements are not shown due to the poor signal-to-noise ratio at the low luminosities delivered during these periods.

Appendix B

LAr energy flow pedestal time dependence

The linearly interpolated pedestal methodology outlined in Section 6.3.2 and Eq. (6.20) explicitly assumes a linear pedestal time dependence over periods of $\sim 40\,\text{min}$ (the duration of a typical μ scan). This assumption can be verified using the data recorded during periods of beam separation by measuring the energy flow as a function of time during these periods when there are no collisions. This section examines the pedestal time dependence during the 38 min pre-scan and 12 min post-scan periods of beam separation in run 354124.

Figure B.1 shows the distributions of the energy flow measured by the EMB, EMEC and FCal in each luminosity block for the two recorded BCIDs in this run (BCIDs 1409 and 2770). Also shown are the mean values per luminosity block (black points) and the mean value over all luminosity blocks (red line). The mean value over all luminosity blocks is the value taken as the pre-scan pedestal for this run, detector subsystem and BCID. Figures B.2 and B.3 show the same data plotted as a function of time, rather than luminosity block number, alongside the data recorded during the post-scan period of beam separation. While the variations between luminosity blocks are large, a linear fit to the pre-scan data indicates a slight time dependence that in most cases agrees with the post-scan pedestal value when extrapolated into this period. The assumption that the pedestal varies linearly with time over $\sim 1\,\text{h}$ periods (and hence during the μ scan) therefore appears valid.

Additionally, the variations in energy flow observed between luminosity blocks are partially correlated between detector subsystems and strongly correlated between BCIDs, as illustrated in Figure B.4. This result indicates a coherent, detector-wide change in response over short ($\sim 1\,\text{min}$) time scales. Whether these variations between luminosity blocks are the result of changes in the detectors themselves, or whether this is a true change in signal due to small variations in the beam conditions resulting in an environment that is not truly without collisions, should be a subject of further study.

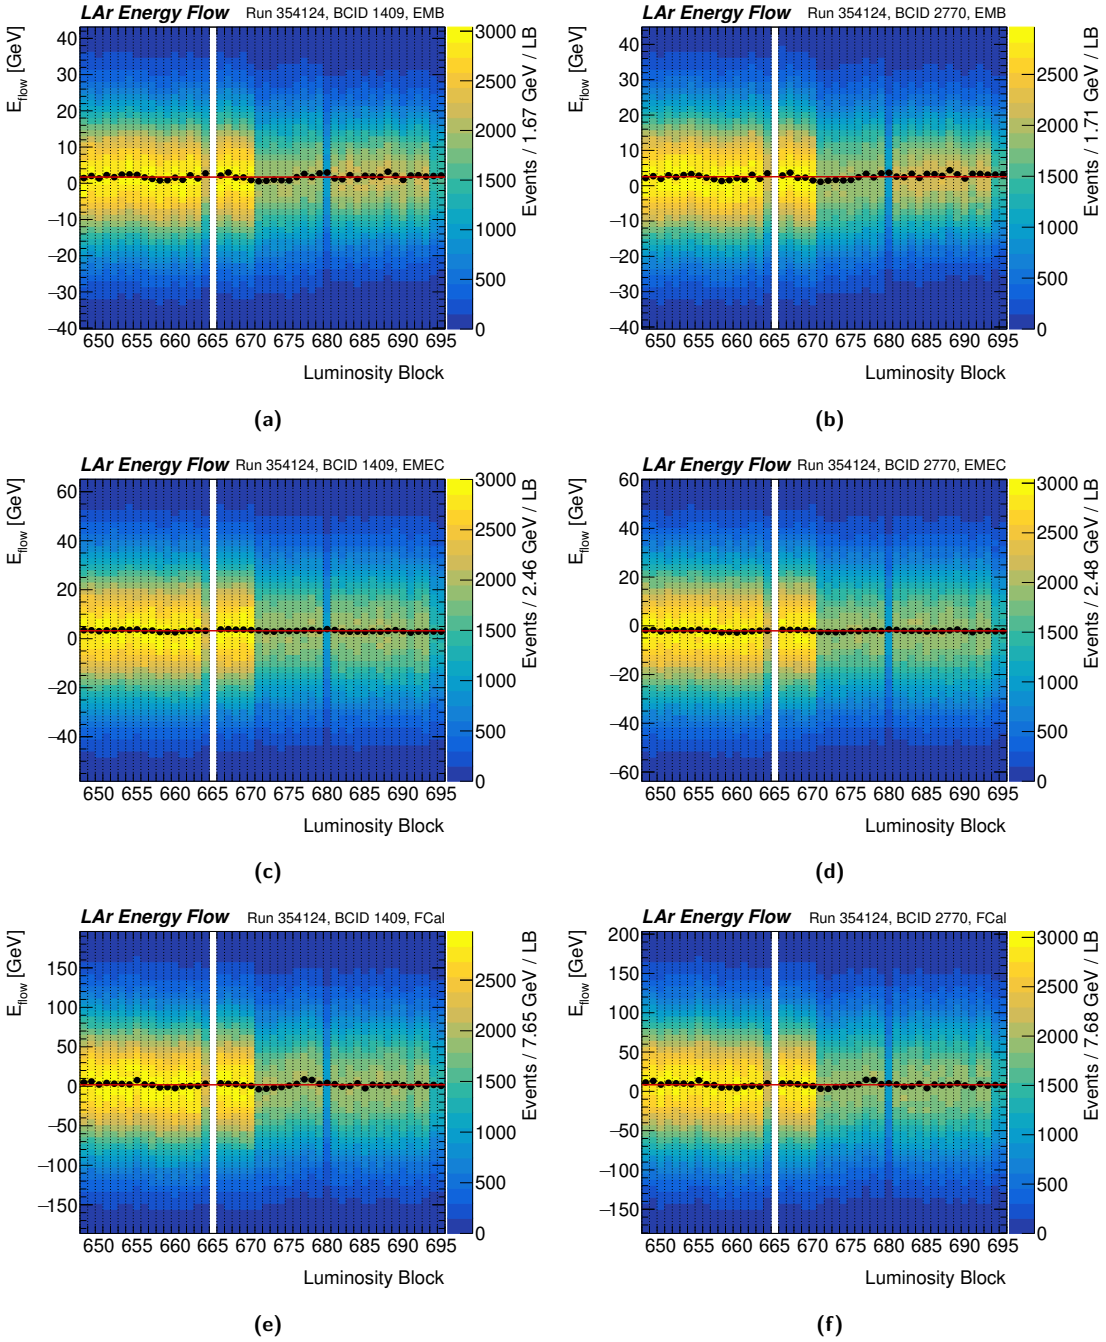

Figure B.1: From top to bottom: LAr energy-flow measurements from the EMB (a and b), EMEC (c and d) and FCal (e and f) during the pre-scan period of beam separation in run 354124. The left column shows the results for BCID 1409 and the right column shows BCID 2770. The distributions of the energy flow in each luminosity block are projected along the vertical axis. The black points indicate the mean energy flow per luminosity block, and the red line indicates the mean energy flow over all luminosity blocks. Note that the difference in total number of events recorded per luminosity block is due to differences in luminosity block duration (typically 60 s and 40 s), and that no data was recorded in luminosity block 665 due to an issue with the data-acquisition system.

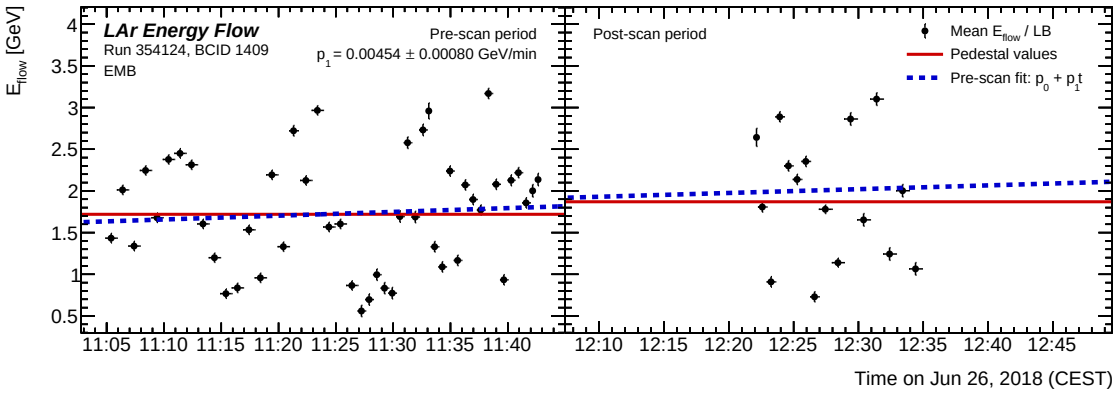

(a)

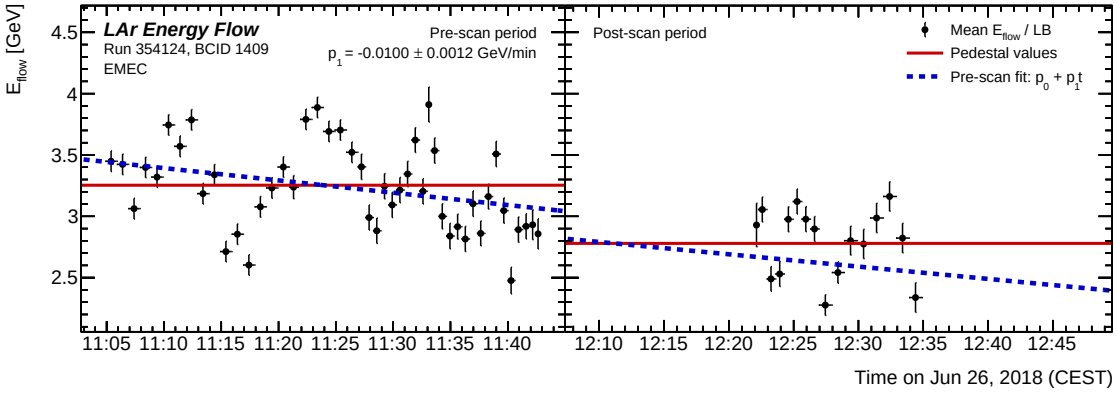

(b)

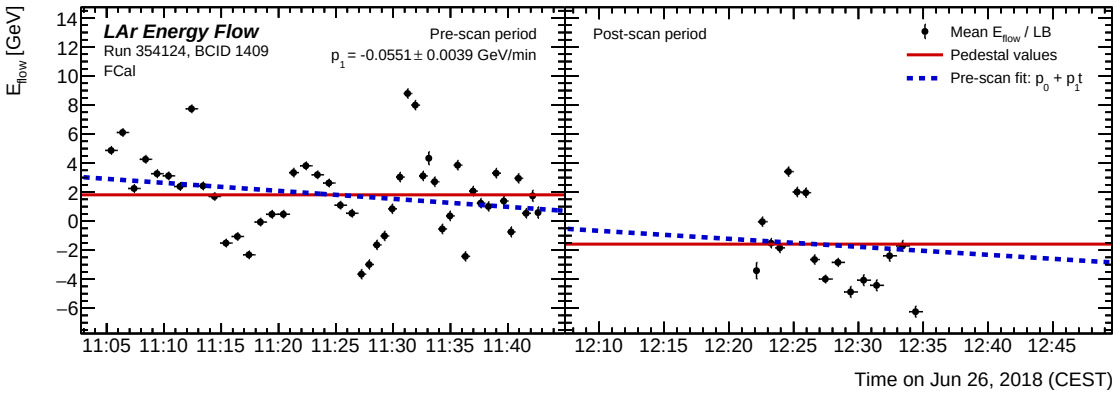

(c)

Figure B.2: LAr energy-flow measurements from the (a) EMB, (b) EMEC and (c) FCal during the pre- and post-scan periods of beam separation in run 354124 for BCID 1409. The left axes show the data recorded during the pre-scan period and the the right axes show the post-scan period. The luminosity block duration is indicated by the horizontal error bars. The pedestal value extracted from each of these periods is shown by the red horizontal lines. The blue dashed lines show the results of a linear fit to the pre-scan luminosity blocks only and extrapolated into the post-scan period.

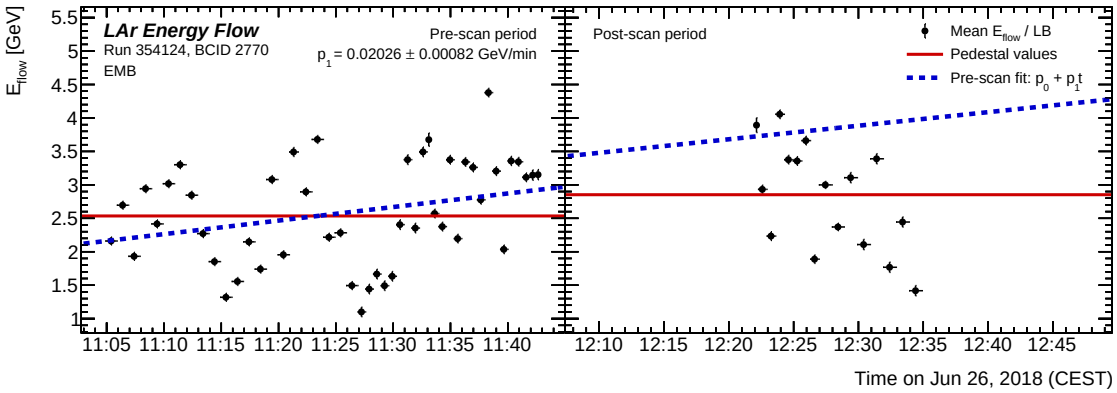

(a)

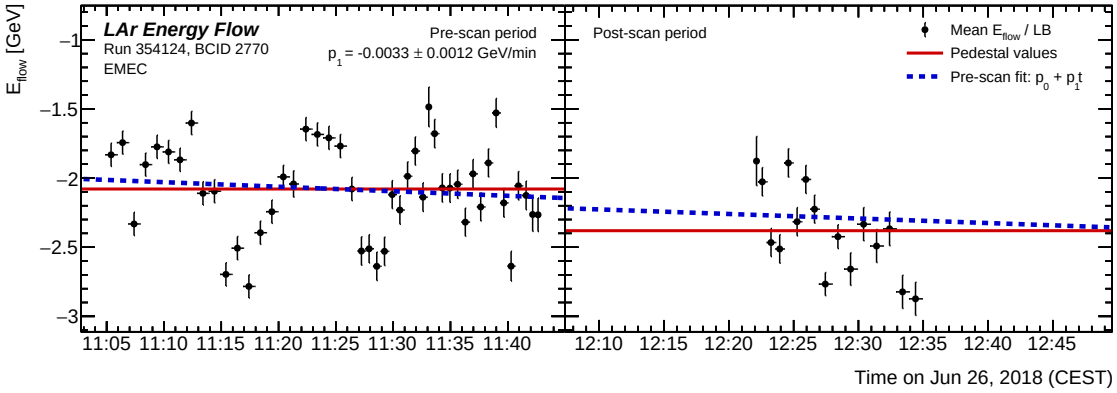

(b)

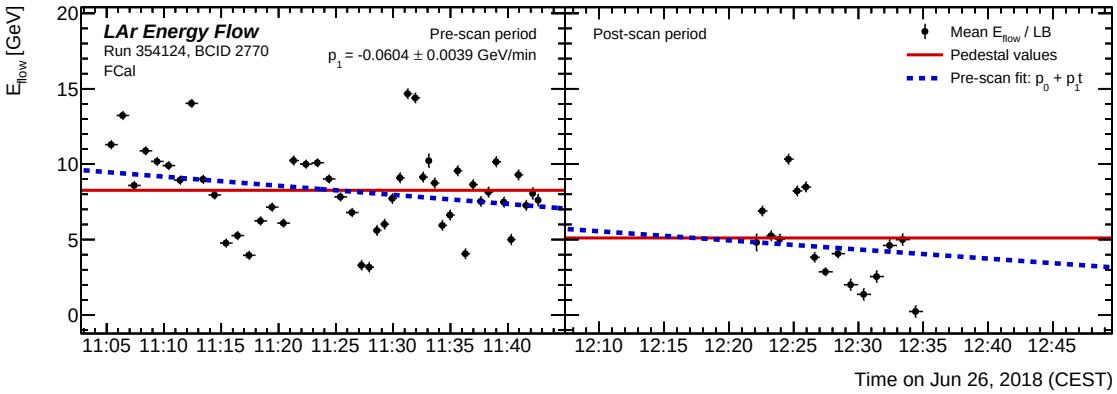

(c)

Figure B.3: LAr energy-flow measurements from the (a) EMB, (b) EMEC and (c) FCal during the pre- and post-scan periods of beam separation in run 354124 for BCID 2770. The left axes show the data recorded during the pre-scan period and the the right axes show the post-scan period. The luminosity block duration is indicated by the horizontal error bars. The pedestal value extracted from each of these periods is shown by the red horizontal lines. The blue dashed lines show the results of a linear fit to the pre-scan luminosity blocks only and extrapolated into the post-scan period.

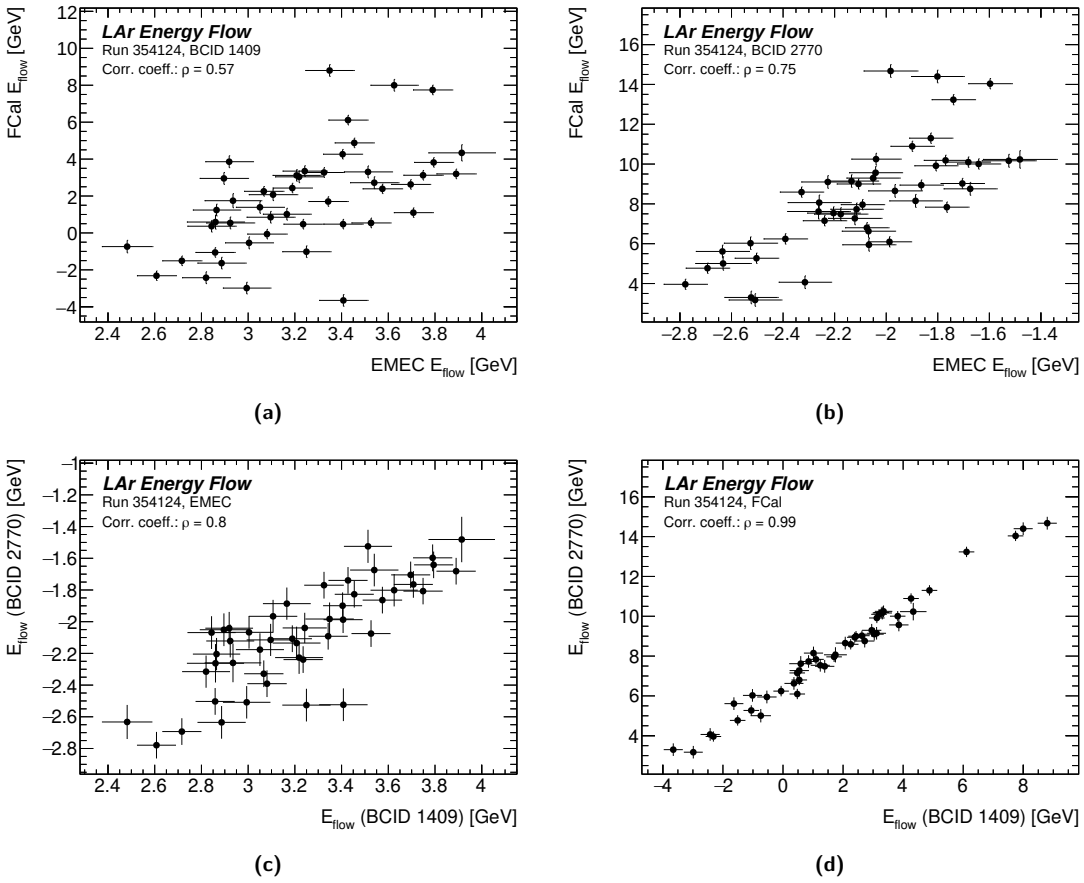

Figure B.4: Correlations between LAr energy-flow measurements per luminosity block during the pre-scan period of beam separation in run 354124. The top row shows the correlation between the simultaneous measurements from the EMEC and FCal for (a) BCID 1409 and (b) BCID 2770. The bottom row shows the correlation between the measurements from BCID 1409 and BCID 2770 for the (c) EMEC and (d) FCal. The correlation coefficients are also indicated. In all cases the uncertainties are statistical only.

Appendix C

Results of the cut-based analysis for the four-lepton final state

The results based on the cut-based categorization as described in Section 7.3.1 are presented here. The number of observed candidate events with mass above 200 GeV together with the expected background yields for each of the four categories of the $\ell^+\ell^-\ell^+\ell^-$ analysis as described in Section 7.3.1 is presented in Table C.1. The obtained ZZ normalization factors are summarized in Table C.2, and the $m_{4\ell}$ spectrum in each category is shown in Figure C.1. The upper limits at 95 % CL on the cross section times branching ratio as a function of the heavy resonance mass in the case of the NWA is presented in Figure C.2.

Table C.1: Expected and observed numbers of $\ell^+\ell^-\ell^+\ell^-$ events for $m_{4\ell} > 200$ GeV, together with their uncertainties, for the VBF-CBA-enriched and ggF-CBA-enriched categories using the cut-based categorization. The expected numbers of events, as well as their uncertainties, are obtained from a combined likelihood fit to the data under the background-only hypothesis. The uncertainties of the ZZ normalization factors, presented in Table C.2, are also taken into account.

Process	VBF-CBA-enriched	ggF-CBA-enriched		
		4μ channel	$2e2\mu$ channel	$4e$ channel
$q\bar{q} \to ZZ$	48 ± 8	860 ± 18	1360 ± 28	515 ± 11
$gg \to ZZ$	13 ± 4	114 ± 9	189 ± 14	73 ± 6
ZZ (EW)	10.9 ± 0.9	6.9 ± 0.3	11.1 ± 0.4	4.4 ± 0.2
Z+jets/$t\bar{t}/WZ$	0.3 ± 0.1	2.1 ± 0.4	6.7 ± 1.6	3.1 ± 0.4
Other backgrounds	3 ± 0.2	16.3 ± 0.4	26.8 ± 0.6	11.8 ± 0.3
Total background	75 ± 9	999 ± 20	1594 ± 31	607 ± 13
Observed	75	932	1656	612

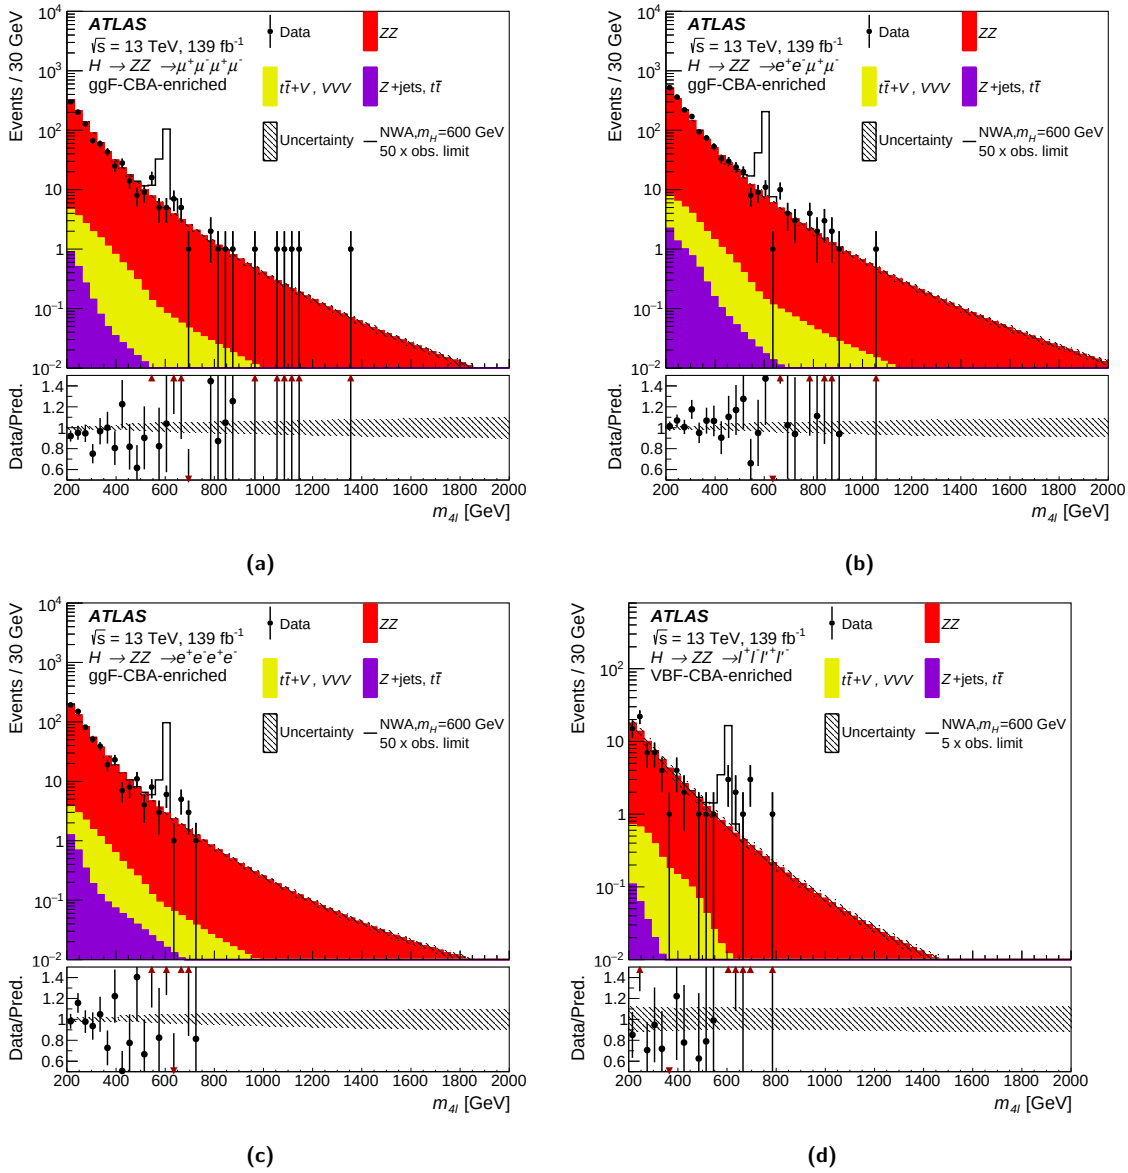

Figure C.1: Distribution of the four-lepton invariant mass $m_{4\ell}$ in the $\ell^+\ell^-\ell^+\ell^-$ search for (a), (b), (c) the ggF-CBA-enriched categories and (d) the VBF-CBA-enriched category. The backgrounds are determined from a combined likelihood fit to the data under the background-only hypothesis. The simulated $m_H = 600\,\text{GeV}$ signal is normalized to a cross section corresponding to 50 (5) times the observed limit given in Section 7.7.1 for the ggF (VBF) production mode. The error bars on the data points indicate the statistical uncertainty, while the systematic uncertainty in the prediction is shown by the hatched band. The lower panels show the ratio of data to prediction.

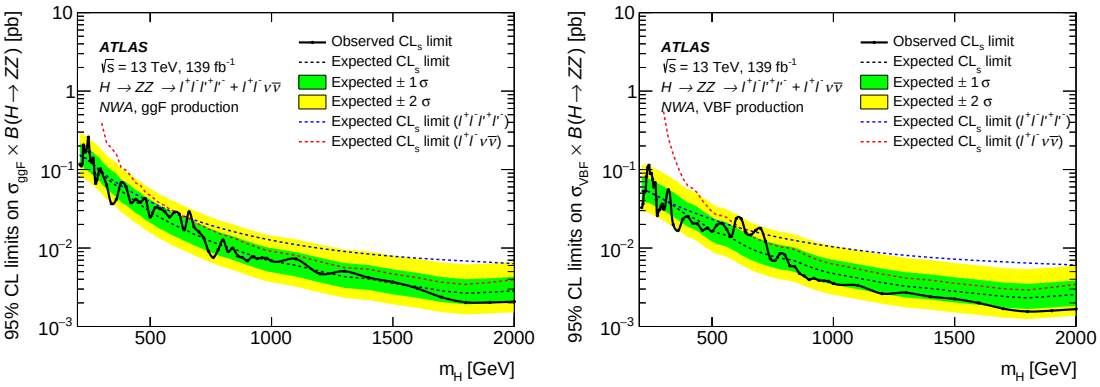

Figure C.2: The upper limits at 95 % CL on the cross section times branching ratio as a function of the heavy resonance mass m_H for (a) the ggF production mode ($\sigma_{\mathrm{ggF}} \times B(H \to ZZ)$) and (b) for the VBF production mode ($\sigma_{\mathrm{VBF}} \times B(H \to ZZ)$) in the case of the NWA. For the $\ell^+\ell^-\ell^+\ell^-$ state the cut-based categorization is used. The green and yellow bands represent the $\pm 1\sigma$ and $\pm 2\sigma$ uncertainties in the expected limits. The dashed coloured lines indicate the expected limits obtained from the individual searches.

Table C.2: The ZZ normalization factors in each category of the two final states, which scale the number of ZZ events estimated from the MC simulations, obtained from a simultaneous likelihood fit of the two final states under the background-only hypothesis. For the $\ell^+\ell^-\ell^+\ell^-$ final state, the cut-based categorization is used.

Analysis	Normalization factor	Fitted value
$\ell^+\ell^-\ell^+\ell^-$	$\mu_{ZZ}^{\text{VBF-CBA}}$	1.1 ± 0.2
	μ_{ZZ}^{ggF}	1.10 ± 0.02
$\ell^+\ell^-\nu\bar{\nu}$	μ_{ZZ}	1.04 ± 0.06

Appendix D

ATLAS Plots

Atlasplots[1] is a Python package to aid in the production of publication-ready plots and figures following the standard ATLAS style guidelines. It provides a simple interface to produce plots with a ROOT[2] backend but using matplotlib[3]-like syntax and idioms, while still giving access to the underlying ROOT objects.

Figure D.1: Atlasplots logo.

A new student of high-energy physics (HEP) has a plethora of plotting tools at their disposal to visualize their data. While ROOT is the gold standard in high-energy physics data analysis, its plotting tools have a notoriously steep learning curve and require a fair amount of "tweaking" to get your plots to look the way you want them to. Matplotlib is an attractive alternative for producing plots, thanks to its simple, intuitive syntax based on the Matlab[4] graphics libraries, and its ubiquitous usage in the Python community, but it is not always straightforward to use in existing ROOT-based analyses. There are of course good reasons to use both these tools:

Why use ROOT

- Common language for the HEP community.
- ROOT is optimized for many common computing operations in HEP, such as filling histograms (`TH1` objects) from an ntuple dataset (stored in the `TTree` format).
    - Plotting a `TH1` histogram with `TH1::Draw()` is usually much simpler than looping over its bins to convert to a matplotlib-readable format.
    - High-level functions are available to act on these objects in a way that a physicist expects, such as histogram fitting with `TH1::Fit()`.

[1] https://atlas-plots.readthedocs.io
[2] https://root.cern/
[3] https://matplotlib.org/
[4] https://www.mathworks.com/products/matlab.html

- Keep the Publication Committee and your Editorial Board happy that your plots follow the ATLAS style.

Why use matplotlib

- Industry standard for data visualization in the Python world.
- Huge community from a broad variety of disciplines.
- Simple syntax and easy to interface with many modern Python-based tools used in HEP.

Atlasplots attempts to bridge this gap between matplotlib and ROOT, and is designed with two main goals in mind:

1. Interact with ROOT objects using matplotlib-like syntax and idioms.
2. Produce plots following the standard ATLAS style guidelines *with as little typing as possible.*

Figures and Axes

Like matplotlib, **atlasplots** follows the ethos of first creating a figure with axes and then plotting data on these axes, rather than the ROOT design of having the axes be a property of the the data object itself. This design choice has a number of advantages, especially for complex plots with many datasets plotted in them, since it is inherently easier for the user to have the axes keep track of the data plotted on them rather than having the user do this and update the axes manually. The core of **atlasplots** is powered by two main classes, inspired by matplotlib: the `Figure` class and the `Axes` class.

A `Figure` object is a wrapper around a ROOT `TCanvas`, which can be accessed directly from the figure's `canvas` attribute. The `Figure` is the top-level plot "window" that holds all the plot elements.

An `Axes` object has no natural counterpart in ROOT, where axes objects are attributes of the object being plotted, such as a `TH1`, `TGraph`, `TF1`, and so on. With a `TH1` histogram object, for example, you would access its axes using `TH1::GetXaxis()` and `TH1::GetYaxis()`, which return the `TAxis` objects for its x and y axes, respectively. Instead, an `Axes` object is a wrapper around a ROOT `TPad` and a "frame", which is an empty `TH1F` histogram object drawn in the `TPad`. This frame provides the axes upon which data is plotted. Passing a drawable ROOT object to the `plot()` method of an `Axes` object draws the data on the `Axes`' frame. Doing so calls the ROOT object's `Draw()` method and automatically expands the frame's axis limits when necessary to accommodate the data being plotted. The `Axes` class provides functions with matplotlib-like syntax to change axis parameters, like `set_xlim()` and `set_xlabel()`, while still providing access to the underlying ROOT objects using the axes' `pad` and `frame` attributes.

Atlasplots provides a `subplots()` function, similar to `matplotlib.pyplot.subplots()`, to create a figure with a set of axes arranged in an $N_{\text{row}} \times N_{\text{col}}$ grid. **Atlasplots** also includes a special case of `subplots()`, called `ratio_plot()`, which produced a figure with two axes: a large upper panel to display two or more datasets and a smaller lower panel to display their ratios.

Usage

Installing

Atlasplots is available on PyPI and can be installed using `pip`:

```
$ pip install atlasplots
```

Importing atlasplots

The preferred convention is to import **atlasplots** using the alias "`aplt`":

```
import atlasplots as aplt
```

Then, to enable the ATLAS style settings, use

```
aplt.set_atlas_style()
```

You only need to call this function once; all subsequent plots in the current Python session will use the ATLAS style.

Plot formatting

Any ROOT object with a `Draw()` method can be plotted on a given set of axes using the `Axes.plot()` method, including `TH1` histograms, `TGraph` scatter plots, `TF1` functions, etc. Each of these ROOT graphics objects inherits from one or a number of graphics-attribute base classes, such as `TAttMarker`, `TAttLine`, `TAttFill`, and so on. There are two ways to change the appearance of a ROOT graphics object in **atlasplots** using the class methods inherited from these base classes:

1. Directly modifying the ROOT object prior to the call to `Axes.plot()`.
2. In the call to `Axes.plot()` itself.

The first method is the traditional ROOT way of modifying graphics objects using the object's setter methods. For example, to plot a histogram using a hatched fill pattern and red lines:

```
import ROOT as root
hist = root.TH1F(...)
hist.SetLineColor(root.kRed)
hist.SetFillColor(root.kRed)
hist.SetFillStyle(3354)
hist.Draw()
```

With **atlasplots**, this can be done all at once:

```
hist = root.TH1F(...)
ax.plot(hist, linecolor=root.kRed, fillcolor=root.kRed, fillstyle=3354)
```

The keyword argument syntax in `Axes.plot()` is the same as the equivalent ROOT attribute setter function, but in all lower case and without the "Set" prefix. For example, to set the marker style, use `markerstyle=...`, which calls `SetMarkerStyle(...)`.

A Simple Example

Below is a simple, complete example that illustrates how to use **atlasplots**. We first import the `ROOT` and `atlasplots` modules and set the plots to use the ATLAS style. We then create a histogram and fill it with Gaussian random data pulled from the distribution $\mathcal{N}(0, 1)$. Finally, we create a figure with a single set of axes using `atlasplots.subplots()`, plot the histogram on these axes, and set the x- and y-axis labels.

```
import ROOT as root
import atlasplots as aplt

aplt.set_atlas_style()                    # Use the ATLAS style
hist = root.TH1F("hist", "", 64, -4, 4)   # Create a histogram
hist.FillRandom("gaus")                   # Randomly fill it with Gaussian data
fig, ax = aplt.subplots()                 # Create a figure and axes
ax.plot(hist)                             # Plot the histogram on these axes
ax.set_xlabel("X")                        # Set the x-axis label
ax.set_ylabel("Events")                   # Set the y-axis label
```

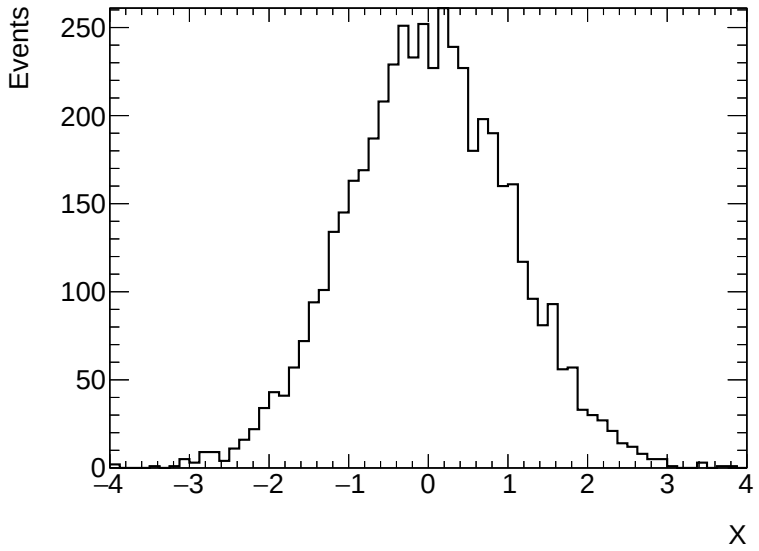

Source code and documentation

The ATLAS Plots source code is available for free under the MIT License at

https://github.com/joeycarter/atlas-plots,

and the complete package documentation is available at

https://atlas-plots.readthedocs.io.

Bibliography

[1] P. A. Zyla et al., *Review of Particle Physics*, PTEP **2020** (2020) 083C01 (cit. on pp. 1, 10, 26).

[2] SNO Collaboration, *Electron energy spectra, fluxes, and day-night asymmetries of B-8 solar neutrinos from measurements with NaCl dissolved in the heavy-water detector at the Sudbury Neutrino Observatory*, Phys. Rev. C **72** (2005) 055502, arXiv: nucl-ex/0502021 (cit. on pp. 4, 19).

[3] Super-Kamiokande Collaboration, *Evidence for oscillation of atmospheric neutrinos*, Phys. Rev. Lett. **81** (1998) 1562, arXiv: hep-ex/9807003 (cit. on pp. 4, 19).

[4] GERDA Collaboration, *Final Results of GERDA on the Search for Neutrinoless Double-β Decay*, Phys. Rev. Lett. **125** (2020) 252502, arXiv: 2009.06079 [nucl-ex] (cit. on p. 4).

[5] EXO-200 Collaboration, *Search for Neutrinoless Double-β Decay with the Complete EXO-200 Dataset*, Phys. Rev. Lett. **123** (2019) 161802, arXiv: 1906.02723 [hep-ex] (cit. on p. 4).

[6] CUORE Collaboration, *Improved Limit on Neutrinoless Double-Beta Decay in ^{130}Te with CUORE*, Phys. Rev. Lett. **124** (2020) 122501, arXiv: 1912.10966 [nucl-ex] (cit. on p. 4).

[7] KamLAND-Zen Collaboration, *Search for Majorana Neutrinos near the Inverted Mass Hierarchy Region with KamLAND-Zen*, Phys. Rev. Lett. **117** (2016) 082503, [Addendum: Phys.Rev.Lett. 117, 109903 (2016)], arXiv: 1605.02889 [hep-ex] (cit. on p. 4).

[8] K. G. Wilson, *Confinement of Quarks*, Phys. Rev. D **10** (1974) 2445 (cit. on p. 5).

[9] E. D. Bloom et al., *High-Energy inelastic e–p scattering at 6 degrees and 10 degrees*, Phys. Rev. Lett. **23** (1969) 930 (cit. on p. 5).

[10] M. Breidenbach et al., *Observed behavior of highly inelastic electron-proton scattering*, Phys. Rev. Lett. **23** (1969) 935 (cit. on p. 5).

[11] TASSO Collaboration, *Evidence for Planar Events in e^+e^- Annihilation at High-Energies*, Phys. Lett. B **86** (1979) 243 (cit. on p. 5).

[12] D. P. Barber et al., *Discovery of Three Jet Events and a Test of Quantum Chromodynamics at PETRA Energies*, Phys. Rev. Lett. **43** (1979) 830 (cit. on p. 5).

[13] PLUTO Collaboration,
Evidence for Gluon Bremsstrahlung in e^+e^- Annihilations at High-Energies,
Phys. Lett. B **86** (1979) 418 (cit. on p. 5).

[14] JADE Collaboration, *Observation of Planar Three Jet Events in e^+e^- Annihilation and Evidence for Gluon Bremsstrahlung*, Phys. Lett. B **91** (1980) 142 (cit. on p. 5).

[15] E. Fermi, *Tentativo di una teoria dei raggi β*, Nuovo Cim. **11** (1934) 1 (cit. on p. 5).

[16] S. L. Glashow, *Partial Symmetries of Weak Interactions*, Nucl. Phys. **22** (1961) 579 (cit. on pp. 5, 16).

[17] S. Weinberg, *A Model of Leptons*, Phys. Rev. Lett. **19** (1967) 1264 (cit. on pp. 5, 16).

[18] A. Salam, "Weak and Electromagnetic Interactions", *Elementary Particle Theory*, ed. by N. Svartholm, Stockholm: Almqvist & Wiksell, 1968 367 (cit. on pp. 5, 16).

[19] Gargamelle Neutrino Collaboration, *Search for Elastic ν_μ Electron Scattering*,
Phys. Lett. B **46** (1973) 121 (cit. on p. 5).

[20] Gargamelle Neutrino Collaboration, *Observation of Neutrino Like Interactions Without Muon Or Electron in the Gargamelle Neutrino Experiment*, Phys. Lett. B **46** (1973) 138 (cit. on p. 5).

[21] UA1 Collaboration, *Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540$ GeV*, Phys. Lett. B **122** (1983) 103 (cit. on p. 5).

[22] UA1 Collaboration, *Experimental Observation of Lepton Pairs of Invariant Mass Around 95-GeV/c**2 at the CERN SPS Collider*, Phys. Lett. B **126** (1983) 398 (cit. on p. 5).

[23] UA2 Collaboration, *Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN anti-p p Collider*,
Phys. Lett. B **122** (1983) 476 (cit. on p. 5).

[24] UA2 Collaboration, *Evidence for $Z^0 \to e^+e^-$ at the CERN $\bar{p}p$ Collider*,
Phys. Lett. B **129** (1983) 130 (cit. on p. 5).

[25] Wikimedia Commons, *Standard Model of Elementary Particles*,
[Online; accessed March 3, 2022], 2021, URL: https://commons.wikimedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg (cit. on p. 6).

[26] D. Griffiths, *Introduction to Elementary Particles; Second, Revised Edition*,
New York, NY, USA: Wiley, 2008, ISBN: 978-3-527-40601-2 (cit. on p. 7).

[27] M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory*,
Boulder, CO, USA: Westview Press, 1995, ISBN: 978-0-201-50397-5 (cit. on p. 7).

[28] M. D. Schwartz, *Quantum Field Theory and the Standard Model*,
Cambridge, UK: Cambridge University Press, 2014, ISBN: 978-1-107-03473-0 (cit. on p. 7).

[29] S. Tomonaga,
On a relativistically invariant formulation of the quantum theory of wave fields,
Prog. Theor. Phys. **1** (1946) 27 (cit. on p. 7).

[30] J. S. Schwinger, *On Quantum electrodynamics and the magnetic moment of the electron*, Phys. Rev. **73** (1948) 416 (cit. on p. 7).

[31] J. S. Schwinger, *Quantum electrodynamics. I. A covariant formulation*, Phys. Rev. **74** (1948) 1439 (cit. on p. 7).

[32] R. P. Feynman, *The Theory of positrons*, Phys. Rev. **76** (1949) 749 (cit. on p. 7).

[33] R. P. Feynman, *Space-time approach to quantum electrodynamics*, Phys. Rev. **76** (1949) 769 (cit. on p. 7).

[34] R. P. Feynman, *Mathematical formulation of the quantum theory of electromagnetic interaction*, Phys. Rev. **80** (1950) 440 (cit. on p. 7).

[35] F. J. Dyson, *The Radiation theories of Tomonaga, Schwinger, and Feynman*, Phys. Rev. **75** (1949) 486 (cit. on p. 7).

[36] F. J. Dyson, *The S matrix in quantum electrodynamics*, Phys. Rev. **75** (1949) 1736 (cit. on p. 7).

[37] P. A. M. Dirac, *Quantum theory of emission and absorption of radiation*, Proc. Roy. Soc. Lond. A **114** (1927) 243 (cit. on p. 7).

[38] P. A. M. Dirac, *The Quantum Theory of the Electron*, Proc. Roy. Soc. Lond. A **117** (1928) 610 (cit. on p. 7).

[39] P. A. M. Dirac, *A Theory of Electrons and Protons*, Proc. Roy. Soc. Lond. A **126** (1930) 360 (cit. on p. 7).

[40] E. Noether, *Invariant Variation Problems*, Gott. Nachr. **1918** (1918) 235, arXiv: physics/0503066 (cit. on p. 8).

[41] G. 't Hooft and M. J. G. Veltman, *Regularization and Renormalization of Gauge Fields*, Nucl. Phys. B **44** (1972) 189 (cit. on p. 8).

[42] C.-N. Yang and R. L. Mills, *Conservation of Isotopic Spin and Isotopic Gauge Invariance*, Phys. Rev. **96** (1954) 191 (cit. on p. 9).

[43] M. Gell-Mann, *The Eightfold Way: A Theory of strong interaction symmetry*, (1961) (cit. on p. 10).

[44] Y. Ne'eman, *Derivation of strong interactions from a gauge invariance*, Nucl. Phys. **26** (1961) 222 (cit. on p. 10).

[45] M. Gell-Mann, *A Schematic Model of Baryons and Mesons*, Phys. Lett. **8** (1964) 214 (cit. on p. 10).

[46] G. Zweig, *An SU_3 model for strong interaction symmetry and its breaking; Version 2*, (1964), In Developments in the Quark Theory of Hadrons, p. 22–101; version 1 is CERN preprint 8182/TH.401, URL: https://cds.cern.ch/record/570209 (cit. on p. 10).

[47] H. D. Politzer, *Reliable Perturbative Results for Strong Interactions?*, Phys. Rev. Lett. **30** (1973) 1346 (cit. on p. 10).

[48] D. J. Gross and F. Wilczek, *Ultraviolet Behavior of Nonabelian Gauge Theories*, Phys. Rev. Lett. **30** (1973) 1343 (cit. on p. 10).

[49] T. D. Lee and C.-N. Yang, *Question of Parity Conservation in Weak Interactions*, Phys. Rev. **104** (1956) 254 (cit. on p. 11).

[50] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes, and R. P. Hudson, *Experimental Test of Parity Conservation in β Decay*, Phys. Rev. **105** (1957) 1413 (cit. on p. 11).

[51] M. Kobayashi and T. Maskawa, *CP Violation in the Renormalizable Theory of Weak Interaction*, Prog. Theor. Phys. **49** (1973) 652 (cit. on p. 11).

[52] Y. Nambu, *Quasiparticles and Gauge Invariance in the Theory of Superconductivity*, Phys. Rev. **117** (1960) 648 (cit. on p. 13).

[53] J. Goldstone, *Field Theories with 'Superconductor' Solutions*, Nuovo Cim. **19** (1961) 154 (cit. on p. 13).

[54] J. Goldstone, A. Salam, and S. Weinberg, *Broken Symmetries*, Phys. Rev. **127** (1962) 965 (cit. on p. 13).

[55] P. W. Anderson, *Plasmons, Gauge Invariance, and Mass*, Phys. Rev. **130** (1963) 439 (cit. on p. 13).

[56] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, Phys. Rev. Lett. **13** (1964) 321 (cit. on p. 13).

[57] P. W. Higgs, *Broken symmetries, massless particles and gauge fields*, Phys. Lett. **12** (1964) 132 (cit. on p. 13).

[58] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, Phys. Rev. Lett. **13** (1964) 508 (cit. on p. 13).

[59] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble, *Global Conservation Laws and Massless Particles*, Phys. Rev. Lett. **13** (1964) 585 (cit. on p. 13).

[60] D. de Florian et al., *Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector*, **2/2017** (2016), arXiv: 1610.07922 [hep-ph] (cit. on pp. 20, 22, 26, 30).

[61] J. R. Ellis, M. K. Gaillard, and D. V. Nanopoulos, *A Phenomenological Profile of the Higgs Boson*, Nucl. Phys. B **106** (1976) 292 (cit. on p. 20).

[62] J. R. Andersen et al., *Handbook of LHC Higgs Cross Sections: 3. Higgs Properties*, (2013), ed. by S. Heinemeyer, C. Mariotti, G. Passarino, and R. Tanaka, arXiv: 1307.1347 [hep-ph] (cit. on pp. 22, 23).

[63] J. Ellis, M. K. Gaillard, and D. V. Nanopoulos, *A Historical Profile of the Higgs Boson*, (2012), arXiv: 1201.6045 [hep-ph] (cit. on p. 21).

[64] D. Kohler, B. A. Watson, and J. A. Becker, *Experimental Search for a Low-Mass Scalar Boson*, Phys. Rev. Lett. **33** (1974) 1628 (cit. on p. 21).

[65] R. Barbieri and T. E. O. Ericson, *Evidence Against the Existence of a Low Mass Scalar Boson from Neutron-Nucleus Scattering*, Phys. Lett. B **57** (1975) 270 (cit. on p. 21).

[66] B. W. Lee, C. Quigg, and H. B. Thacker, *Weak Interactions at Very High-Energies: The Role of the Higgs Boson Mass*, Phys. Rev. D **16** (1977) 1519 (cit. on p. 22).

[67] The LEP Collaboration (ALEPH, DELPHI, L3 and OPAL), *Search for the standard model Higgs boson at LEP*, Phys. Lett. B **565** (2003) 61, arXiv: hep-ex/0306033 (cit. on p. 22).

[68] CDF Collaboration and D0 Collaboration, *Updated Combination of CDF and D0 Searches for Standard Model Higgs Boson Production with up to 10.0 fb^{-1} of Data*, (2012), arXiv: 1207.0449 [hep-ex] (cit. on pp. 22, 23).

[69] M. Baak et al., *Updated Status of the Global Electroweak Fit and Constraints on New Physics*, Eur. Phys. J. C **72** (2012) 2003, arXiv: 1107.0975 [hep-ph] (cit. on p. 22).

[70] ATLAS Collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, Phys. Lett. B **716** (2012) 1, arXiv: 1207.7214 [hep-ex] (cit. on pp. 22, 24, 116).

[71] CMS Collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, Phys. Lett. B **716** (2012) 30, arXiv: 1207.7235 [hep-ex] (cit. on pp. 22, 24, 116).

[72] ATLAS and CMS Collaborations, *Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV*, JHEP **08** (2016) 045, arXiv: 1606.02266 [hep-ex] (cit. on pp. 22, 30).

[73] ATLAS and CMS Collaborations, *Combined Measurement of the Higgs Boson Mass in pp Collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS Experiments*, Phys. Rev. Lett. **114** (2015) 191803, arXiv: 1503.07589 [hep-ex] (cit. on p. 23).

[74] T. Aoyama, T. Kinoshita, and M. Nio, *Revised and Improved Value of the QED Tenth-Order Electron Anomalous Magnetic Moment*, Phys. Rev. D **97** (2018) 036001, arXiv: 1712.06060 [hep-ph] (cit. on p. 23).

[75] D. Hanneke, S. Fogwell, and G. Gabrielse, *New Measurement of the Electron Magnetic Moment and the Fine Structure Constant*, Phys. Rev. Lett. **100** (2008) 120801, arXiv: 0801.1134 [physics.atom-ph] (cit. on p. 23).

[76] N. Aghanim et al., *Planck 2018 results. VI. Cosmological parameters*, Astron. Astrophys. **641** (2020) A6, [Erratum: Astron.Astrophys. 652, C4 (2021)], arXiv: 1807.06209 [astro-ph.CO] (cit. on p. 23).

[77] T. D. Lee, *A Theory of Spontaneous T Violation*, Phys. Rev. D **8** (1973) 1226 (cit. on p. 26).

[78] G. C. Branco et al., *Theory and phenomenology of two-Higgs-doublet models*, Phys. Rept. **516** (2012) 1, arXiv: 1106.0034 [hep-ph] (cit. on p. 26).

[79] G. Carrillo-Montoya et al., *Benchmarks for Heavy Higgs boson searches in the electroweak bosons decay channels in the ATLAS detector*, tech. rep., CERN, 2014, URL: https://cds.cern.ch/record/1692903 (cit. on p. 27).

[80] N. Arkani-Hamed, S. Dimopoulos, and G. R. Dvali, *The Hierarchy problem and new dimensions at a millimeter*, Phys. Lett. B **429** (1998) 263, arXiv: hep-ph/9803315 (cit. on p. 28).

[81] I. Antoniadis, N. Arkani-Hamed, S. Dimopoulos, and G. R. Dvali, *New dimensions at a millimeter to a Fermi and superstrings at a TeV*, Phys. Lett. B **436** (1998) 257, arXiv: hep-ph/9804398 (cit. on p. 28).

[82] N. Arkani-Hamed, S. Dimopoulos, and G. R. Dvali, *Phenomenology, astrophysics and cosmology of theories with submillimeter dimensions and TeV scale quantum gravity*, Phys. Rev. D **59** (1999) 086004, arXiv: hep-ph/9807344 (cit. on p. 28).

[83] E. G. Adelberger, B. R. Heckel, and A. E. Nelson, *Tests of the gravitational inverse square law*, Ann. Rev. Nucl. Part. Sci. **53** (2003) 77, arXiv: hep-ph/0307284 (cit. on p. 28).

[84] D. J. Kapner et al., *Tests of the gravitational inverse-square law below the dark-energy length scale*, Phys. Rev. Lett. **98** (2007) 021101, arXiv: hep-ph/0611184 (cit. on p. 28).

[85] L. Randall and R. Sundrum, *A Large mass hierarchy from a small extra dimension*, Phys. Rev. Lett. **83** (1999) 3370, arXiv: hep-ph/9905221 (cit. on p. 28).

[86] L. Randall and R. Sundrum, *An Alternative to compactification*, Phys. Rev. Lett. **83** (1999) 4690, arXiv: hep-th/9906064 (cit. on p. 28).

[87] W. D. Goldberger and M. B. Wise, *Modulus stabilization with bulk fields*, Phys. Rev. Lett. **83** (1999) 4922, arXiv: hep-ph/9907447 (cit. on p. 29).

[88] K. Agashe, H. Davoudiasl, G. Perez, and A. Soni, *Warped Gravitons at the LHC and Beyond*, Phys. Rev. D **76** (2007) 036006, arXiv: hep-ph/0701186 (cit. on pp. 29, 119).

[89] A. Carvalho, *Gravity particles from Warped Extra Dimensions, predictions for LHC*, (2014), arXiv: 1404.0102 [hep-ph] (cit. on pp. 29, 30).

[90] ATLAS Collaboration, *hMSSM summary plots from direct and indirect searches*, tech. rep., CERN, 2021, URL: https://cds.cern.ch/record/2777012 (cit. on p. 31).

[91] ATLAS Collaboration, *Search for an additional, heavy Higgs boson in the $H \to ZZ$ decay channel at $\sqrt{s} = 8$ TeV in pp collision data with the ATLAS detector*, Eur. Phys. J. C **76** (2016) 45, arXiv: 1507.05930 [hep-ex] (cit. on pp. 30, 122).

[92] CMS Collaboration, *Search for a Higgs boson in the mass range from 145 to 1000 GeV decaying to a pair of W or Z bosons*, JHEP **10** (2015) 144, arXiv: 1504.00936 [hep-ex] (cit. on p. 30).

[93] ATLAS Collaboration, *Search for heavy ZZ resonances in the $\ell^+\ell^-\ell^+\ell^-$ and $\ell^+\ell^-\nu\bar{\nu}$ final states using proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Eur. Phys. J. C **78** (2018) 293, arXiv: 1712.06386 [hep-ex] (cit. on pp. 30, 31, 117, 126, 139, 141, 144, 149).

[94] CMS Collaboration, *Search for a new scalar resonance decaying to a pair of Z bosons in proton–proton collisions at $\sqrt{s} = 13$ TeV*, JHEP **06** (2018) 127, arXiv: 1804.01939 [hep-ex] (cit. on p. 31).

[95] C. Llewellyn Smith, *The Large Hadron Collider: Lessons learned and summary*, Phil. Trans. Roy. Soc. Lond. A **370** (2012) 995 (cit. on p. 33).

[96] *CERN Council Gives Go-ahead for Large Hadron Collider*, CERN (1994), URL: https://home.web.cern.ch/news/press-release/cern/cern-council-gives-go-ahead-large-hadron-collider (cit. on p. 33).

[97] *Resolution concerning the Construction of the LHC and the Funding of the Organization. 106th Session of Council*, CERN (1997), URL: https://cds.cern.ch/record/34302 (cit. on p. 34).

[98] ATLAS Collaboration, *ATLAS: letter of intent for a general-purpose pp experiment at the Large Hadron Collider at CERN*, (1992), URL: https://cds.cern.ch/record/291061 (cit. on p. 34).

[99] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, JINST **3** (2008) S08003 (cit. on pp. 34, 40, 41, 47, 48, 51, 53–56, 58–60, 62).

[100] CMS Collaboration, *CMS: letter of intent by the CMS Collaboration for a general purpose detector at LHC*, (1992), URL: https://cds.cern.ch/record/290808 (cit. on p. 34).

[101] CMS Collaboration, *The CMS experiment at the CERN LHC*, JINST **3** (2008) S08004 (cit. on p. 34).

[102] ALICE Collaboration, *Letter of Intent for A Large Ion Collider Experiment [ALICE]*, (1993), URL: https://cds.cern.ch/record/290825 (cit. on p. 34).

[103] ALICE Collaboration, *The ALICE experiment at the CERN LHC*, JINST **3** (2008) S08002 (cit. on p. 34).

[104] LHCb Collaboration, *LHCb Letter of Intent*, tech. rep., CERN, 1995, URL: https://cds.cern.ch/record/691698 (cit. on p. 34).

[105] LHCb Collaboration, *The LHCb Detector at the LHC*, JINST **3** (2008) S08005 (cit. on p. 34).

[106] *LHC Design Report Vol.1: The LHC Main Ring*, CERN-2004-003-V1 (2004), ed. by O. S. Bruning et al. (cit. on pp. 34, 35).

[107] L. Evans and P. Bryant, *LHC Machine*, JINST **3** (2008) S08001 (cit. on pp. 34, 35, 38).

[108] *LHC Design Report. 2. The LHC infrastructure and general services*, CERN-2004-003-V2 (2004), ed. by O. Buning et al. (cit. on p. 34).

[109] *LHC Design Report. 3. The LHC injector chain*, CERN-2004-003-V3 (2004), ed. by M. Benedikt, P. Collier, V. Mertens, J. Poole, and K. Schindl (cit. on p. 34).

[110] E. Mobs, *The CERN accelerator complex – August 2018. Complexe des accélérateurs du CERN – Août 2018*, (2018), General Photo, URL: https://cds.cern.ch/record/2636343 (cit. on p. 37).

[111] J. Vollaire et al., *Linac4 design report*, ed. by M. Vretenar, vol. 6, CERN Yellow Reports: Monographs, Geneva: CERN, 2020, URL: https://cds.cern.ch/record/2736208 (cit. on p. 36).

[112] P. Freyermuth, *LHC Report: imaginative injectors*, CERN Bulletin **30-32** (2016), URL: https://cds.cern.ch/journal/CERNBulletin/2016/32/News%20Articles/2201549 (cit. on p. 36).

[113] J. Wenninger, *Operation and Configuration of the LHC in Run 2*, (2019), URL: https://cds.cern.ch/record/2668326 (cit. on p. 39).

[114] N. Karastathis, K. Fuchsberger, M. Hostettler, Y. Papaphilippou, and D. Pellegrini, *Crossing Angle Anti-Leveling at the LHC in 2017*, J. Phys.: Conf. Ser. **1067** (2018) MOPMF040. 5 p, URL: https://cds.cern.ch/record/2648700 (cit. on p. 39).

[115] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, ATLAS-CONF-2019-021, 2019, URL: https://cds.cern.ch/record/2677054 (cit. on pp. 39, 87, 88, 90–95, 167).

[116] L. Mether et al., "16L2: Operation, observations and physics aspects", *Proceedings of the 8th Evian Workshop on LHC beam operation*, 2017 99, URL: https://cds.cern.ch/record/2654224 (cit. on p. 40).

[117] G. Arduini and A. Milanese, *Resistive solenoid in the 16L2 interconnect*, (2017), Presented at the 319th LHC Machine Committee (LMC), URL: https://espace.cern.ch/lhc-machine-committee/Minutes/1/lmc_319.pdf (cit. on p. 40).

[118] T. G. Cornelissen et al., *Updates of the ATLAS Tracking Event Data Model (Release 13)*, ATL-SOFT-PUB-2007-003, 2007, URL: https://cds.cern.ch/record/1038095 (cit. on p. 43).

[119] ATLAS Collaboration, *ATLAS Inner Detector: Technical Design Report, Volume 1*, ATLAS-TDR-4; CERN-LHCC-97-016, 1997, URL: https://cds.cern.ch/record/331063 (cit. on p. 44).

[120] ATLAS Collaboration, *ATLAS Inner Detector: Technical Design Report, Volume 2*, ATLAS-TDR-5, CERN-LHCC-97-017, 1997, URL: https://cds.cern.ch/record/331064 (cit. on p. 44).

[121] ATLAS Collaboration, *ATLAS Central Solenoid: Magnet Project Technical Design Report, Volume 4*, ATLAS-TDR-9; CERN-LHCC-97-021, CERN, 1997, URL: https://cds.cern.ch/record/331067 (cit. on p. 44).

[122] ATLAS Collaboration,
Track Reconstruction Performance of the ATLAS Inner Detector at $\sqrt{s} = 13$ TeV,
ATL-PHYS-PUB-2015-018, 2015, URL: https://cds.cern.ch/record/2037683
(cit. on p. 45).

[123] M. Capeans et al., *ATLAS Insertable B-Layer Technical Design Report*,
(2010), Addendum: ATLAS-TDR-19-ADD-1, URL: https://cds.cern.ch/record/1291633
(cit. on p. 44).

[124] B. Abbott et al., *Production and integration of the ATLAS Insertable B-Layer*,
JINST **13** (2018) T05008, arXiv: 1803.00844 [physics.ins-det] (cit. on pp. 44, 45).

[125] ATLAS Collaboration, *ATLAS Pixel Detector: Technical Design Report*,
ATLAS-TDR-11; CERN-LHCC-98-013, 1998, URL: https://cds.cern.ch/record/381263
(cit. on p. 46).

[126] G. Aad et al., *ATLAS pixel detector electronics and sensors*, JINST **3** (2008) P07007
(cit. on p. 46).

[127] A. Ahmad et al., *The Silicon microstrip sensors of the ATLAS semiconductor tracker*,
Nucl. Instrum. Meth. A **578** (2007) 98 (cit. on p. 46).

[128] A. Abdesselam et al., *The barrel modules of the ATLAS semiconductor tracker*,
Nucl. Instrum. Meth. A **568** (2006) 642 (cit. on p. 46).

[129] A. Abdesselam et al., *The ATLAS semiconductor tracker end-cap module*,
Nucl. Instrum. Meth. A **575** (2007) 353 (cit. on p. 46).

[130] E. Abat et al., *The ATLAS Transition Radiation Tracker (TRT) proportional drift tube: Design and performance*, JINST **3** (2008) P02013 (cit. on p. 46).

[131] E. Abat et al., *The ATLAS TRT barrel detector*, JINST **3** (2008) P02014 (cit. on p. 46).

[132] E. Abat et al., *The ATLAS TRT end-cap detectors*, JINST **3** (2008) P10003 (cit. on p. 46).

[133] ATLAS Collaboration, *Performance of the ATLAS Transition Radiation Tracker in Run 1 of the LHC: tracker properties*, JINST **12** (2017) P05002, arXiv: 1702.06473 [hep-ex]
(cit. on p. 46).

[134] R. C. Fernow, *Introduction to Experimental Particle Physics*, 1983, ISBN: 978-0-521-30170-1
(cit. on p. 47).

[135] ATLAS Collaboration, *Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC*, Eur. Phys. J. C **77** (2017) 332,
arXiv: 1611.10235 [hep-ex] (cit. on pp. 47, 80).

[136] C. W. Fabjan and F. Gianotti, *Calorimetry for particle physics*,
Rev. Mod. Phys. **75** (2003) 1243 (cit. on p. 49).

[137] B. Botchev et al.,
The High Voltage Feedthroughs for the ATLAS Liquid Argon Calorimeters,
JINST **2** (2007) T10002, arXiv: 0709.0744 [physics.ins-det] (cit. on pp. 50, 156).

[138] B. Aubert et al.,
Construction, assembly and tests of the ATLAS electromagnetic barrel calorimeter,
Nucl. Instrum. Meth. A **558** (2006) 388 (cit. on p. 52).

[139] M. Aleksa et al.,
Construction, assembly and tests of the ATLAS electromagnetic end-cap calorimeters,
JINST **3** (2008) P06002 (cit. on pp. 52, 54).

[140] ATLAS Collaboration, *Drift Time Measurement in the ATLAS Liquid Argon Electromagnetic Calorimeter using Cosmic Muons*, Eur. Phys. J. C **70** (2010) 755,
arXiv: 1002.4189 [hep-ex] (cit. on p. 52).

[141] D. M. Gingrich et al.,
Construction, assembly and testing of the ATLAS hadronic end-cap calorimeter,
JINST **2** (2007) P05005 (cit. on p. 52).

[142] A. Artamonov et al., *The ATLAS Forward Calorimeter*, JINST **3** (2008) P02010
(cit. on pp. 53, 153).

[143] N. J. Buchanan et al., *ATLAS liquid argon calorimeter front end electronics*,
JINST **3** (2008) P09003 (cit. on p. 55).

[144] W. Cleland and E. Stern, *Signal processing considerations for liquid ionization calorimeters in a high rate environment*, Nucl. Instrum. Meth. A **338** (1994) 467, ISSN: 0168-9002
(cit. on p. 57).

[145] F. Ariztizabal et al., *Construction and performance of an iron scintillator hadron calorimeter with longitudinal tile configuration*, Nucl. Instrum. Meth. A **349** (1994) 384
(cit. on p. 57).

[146] M. Aharrouche et al., *Measurement of the response of the ATLAS liquid argon barrel calorimeter to electrons at the 2004 combined test-beam*,
Nucl. Instrum. Meth. A **614** (2010) 400 (cit. on p. 58).

[147] B. Dowler et al., *Performance of the ATLAS hadronic end-cap calorimeter in beam tests*,
Nucl. Instrum. Meth. A **482** (2002) 94 (cit. on p. 58).

[148] C. Cojocaru et al., *Hadronic calibration of the ATLAS liquid argon end-cap calorimeter in the pseudorapidity region $1.6 < |\eta| < 1.8$ in beam tests*,
Nucl. Instrum. Meth. A **531** (2004) 481, ed. by V. Gostilo, arXiv: physics/0407009
(cit. on p. 58).

[149] P. Strizenec and A. Minaenko,
Performance of the ATLAS liquid argon endcap calorimeter in beam tests,
J. Phys. Conf. Ser. **160** (2009) 012078, ed. by M. Fraternali, G. Gaudio, and M. Livan
(cit. on p. 58).

[150] J. P. Archambault et al., *Energy calibration of the ATLAS liquid argon forward calorimeter*,
JINST **3** (2008) P02002 (cit. on p. 58).

[151] P. Adragna et al., *Testbeam studies of production modules of the ATLAS tile calorimeter*,
Nucl. Instrum. Meth. A **606** (2009) 362 (cit. on p. 59).

[152] ATLAS Collaboration, *ATLAS Muon Spectrometer: Technical Design Report*,
ATLAS-TDR-10; CERN-LHCC-97-022, CERN, 1997,
URL: https://cds.cern.ch/record/331068 (cit. on p. 60).

[153] ATLAS Collaboration, *Performance of the ATLAS muon triggers in Run 2*, JINST **15** (2020) P09015, arXiv: 2004.13447 [hep-ex] (cit. on p. 63).

[154] I. Riu, *The ATLAS Level-1 Topological Trigger performance in Run 2*, J. Phys. Conf. Ser. **898** (2017) 032037 (cit. on p. 63).

[155] H. Bertelsen et al., *Operation of the upgraded ATLAS Central Trigger Processor during the LHC Run 2*, JINST **11** (2016) C02020 (cit. on p. 63).

[156] V. Cindro et al., *The ATLAS beam conditions monitor*, JINST **3** (2008) P02004 (cit. on p. 64).

[157] S. Abdel Khalek et al., *The ALFA Roman Pot Detectors of ATLAS*, JINST **11** (2016) P11013, arXiv: 1609.00249 [physics.ins-det] (cit. on p. 64).

[158] L. Adamczyk et al., *Technical Design Report for the ATLAS Forward Proton Detector*, CERN-LHCC-2015-009, ATLAS-TDR-024, 2015, URL: https://cds.cern.ch/record/2017378 (cit. on p. 64).

[159] P. Jenni, M. Nessi, and M. Nordberg, *Zero Degree Calorimeters for ATLAS*, CERN-LHCC-2007-001, LHCC-I-016, 2007, URL: https://cds.cern.ch/record/1009649 (cit. on p. 64).

[160] G. Avoni et al., *The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS*, JINST **13** (2018) P07017 (cit. on pp. 64, 65).

[161] ATLAS Collaboration, *Measurement of the Inelastic Proton–Proton Cross Section at $\sqrt{s} = 13$ TeV with the ATLAS Detector at the LHC*, Phys. Rev. Lett. **117** (2016) 182002, arXiv: 1606.02625 [hep-ex] (cit. on pp. 67, 153).

[162] G. Antchev et al., *First measurement of elastic, inelastic and total cross-section at $\sqrt{s} = 13$ TeV by TOTEM and overview of cross-section data at LHC energies*, Eur. Phys. J. C **79** (2019) 103, arXiv: 1712.06153 [hep-ex] (cit. on p. 67).

[163] ATLAS Collaboration, *Standard Model Summary Plots February 2022*, ATL-PHYS-PUB-2022-009, 2022, URL: https://cds.cern.ch/record/2804061 (cit. on p. 67).

[164] V. N. Gribov and L. N. Lipatov, *Deep inelastic e p scattering in perturbation theory*, Sov. J. Nucl. Phys. **15** (1972) 438 (cit. on p. 68).

[165] Y. L. Dokshitzer, *Calculation of the Structure Functions for Deep Inelastic Scattering and e+ e- Annihilation by Perturbation Theory in Quantum Chromodynamics.*, Sov. Phys. JETP **46** (1977) 641 (cit. on p. 68).

[166] G. Altarelli and G. Parisi, *Asymptotic Freedom in Parton Language*, Nucl. Phys. B **126** (1977) 298 (cit. on p. 68).

[167] L. Harland-Lang, A. Martin, P. Motylinski, and R. Thorne, *Parton distributions in the LHC era: MMHT 2014 PDFs*, Eur. Phys. J. C **75** (2015) 204, arXiv: 1412.3989 [hep-ph] (cit. on p. 68).

[168] W. A. Bardeen, A. J. Buras, D. W. Duke, and T. Muta,
Deep Inelastic Scattering Beyond the Leading Order in Asymptotically Free Gauge Theories,
Phys. Rev. D **18** (1978) 3998 (cit. on p. 69).

[169] A. Buckley et al., *General-purpose event generators for LHC physics*,
Phys. Rept. **504** (2011) 145, arXiv: 1101.2599 [hep-ph] (cit. on pp. 69, 70).

[170] T. Gleisberg et al., *Event generation with SHERPA 1.1*, JHEP **02** (2009) 007,
arXiv: 0811.4622 [hep-ph] (cit. on p. 70).

[171] B. Andersson, *The Lund model*, vol. 7, Cambridge University Press, 2005,
ISBN: 978-0-521-01734-3, 978-0-521-42094-5, 978-0-511-88149-7 (cit. on p. 71).

[172] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand,
Parton fragmentation and string dynamics, Phys. Rept. **97** (1983) 31 (cit. on p. 71).

[173] B. R. Webber, *A QCD Model for Jet Fragmentation Including Soft Gluon Interference*,
Nucl. Phys. B **238** (1984) 492 (cit. on p. 71).

[174] ATLAS Collaboration, *The ATLAS Simulation Infrastructure*,
Eur. Phys. J. C **70** (2010) 823, arXiv: 1005.4568 [physics.ins-det] (cit. on p. 72).

[175] GEANT4 Collaboration, S. Agostinelli, et al., *GEANT4 – a simulation toolkit*,
Nucl. Instrum. Meth. A **506** (2003) 250 (cit. on p. 72).

[176] ATLAS Collaboration, *Performance of the Fast ATLAS Tracking Simulation (FATRAS) and the ATLAS Fast Calorimeter Simulation (FastCaloSim) with single particles*,
ATL-SOFT-PUB-2014-001, 2014, URL: https://cds.cern.ch/record/1669341
(cit. on p. 73).

[177] ATLAS Collaboration,
ATLAS data quality operations and performance for 2015–2018 data-taking,
JINST **15** (2020) P04003, arXiv: 1911.04632 [physics.ins-det] (cit. on p. 75).

[178] J. Pequenao and P. Schaffner,
"How ATLAS detects particles: diagram of particle paths in the detector", 2013,
URL: https://cds.cern.ch/record/1505342 (cit. on p. 78).

[179] ATLAS Collaboration, *Performance of the ATLAS track reconstruction algorithms in dense environments in LHC Run 2*, Eur. Phys. J. C **77** (2017) 673, arXiv: 1704.07983 [hep-ex]
(cit. on pp. 78, 79).

[180] R. Fruhwirth, *Application of Kalman filtering to track and vertex fitting*,
Nucl. Instrum. Meth. A **262** (1987) 444 (cit. on p. 78).

[181] T. Cornelissen et al., *The new ATLAS track reconstruction (NEWT)*,
J. Phys. Conf. Ser. **119** (2008) 032014, ed. by R. Sobie, R. Tafirout, and J. Thomson
(cit. on p. 79).

[182] R. Fruhwirth, W. Waltenberger, and P. Vanlaer, *Adaptive vertex fitting*,
J. Phys. G **34** (2007) N343 (cit. on p. 80).

[183] ATLAS Collaboration,
Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1,
Eur. Phys. J. C **77** (2017) 490, arXiv: 1603.02934 [hep-ex] (cit. on pp. 80, 85).

[184] ATLAS Collaboration, *Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at $\sqrt{s} = 13$ TeV*, Eur. Phys. J. C **79** (2019) 639, arXiv: 1902.04655 [hep-ex] (cit. on pp. 80, 81, 138).

[185] ATLAS Collaboration, *Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run 2 data collected in 2015 and 2016*, Eur. Phys. J. C **79** (2019) 205, arXiv: 1810.05087 [hep-ex] (cit. on p. 80).

[186] ATLAS Collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton–proton collision data*, JINST **14** (2019) P12006, arXiv: 1908.00005 [hep-ex] (cit. on pp. 81, 83).

[187] ATLAS Collaboration, *Electron performance measurements with the ATLAS detector using the 2010 LHC proton–proton collision data*, Eur. Phys. J. C **72** (2012) 1909, arXiv: 1110.3174 [hep-ex] (cit. on p. 81).

[188] ATLAS Collaboration, *Electron reconstruction and identification efficiency measurements with the ATLAS detector using the 2011 LHC proton–proton collision data*, Eur. Phys. J. C **74** (2014) 2941, arXiv: 1404.2240 [hep-ex] (cit. on p. 81).

[189] ATLAS Collaboration, *Improved electron reconstruction in ATLAS using the Gaussian Sum Filter-based model for bremsstrahlung*, ATLAS-CONF-2012-047, 2012, URL: https://cds.cern.ch/record/1449796 (cit. on p. 81).

[190] ATLAS Collaboration, *Electron and photon reconstruction and performance in ATLAS using a dynamical, topological cell clustering-based approach*, ATL-PHYS-PUB-2017-022, 2017, URL: https://cds.cern.ch/record/2298955 (cit. on p. 82).

[191] ATLAS Collaboration, *Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV*, Eur. Phys. J. C **81** (2021) 578, arXiv: 2012.00578 [hep-ex] (cit. on pp. 83, 85).

[192] ATLAS Collaboration, *Jet reconstruction and performance using particle flow with the ATLAS Detector*, Eur. Phys. J. C **77** (2017) 466, arXiv: 1703.10485 [hep-ex] (cit. on p. 85).

[193] M. Cacciari, G. P. Salam, and G. Soyez, *The anti-k_t jet clustering algorithm*, JHEP **04** (2008) 063, arXiv: 0802.1189 [hep-ph] (cit. on p. 85).

[194] ATLAS Collaboration, *Jet energy scale measurements and their systematic uncertainties in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Phys. Rev. D **96** (2017) 072002, arXiv: 1703.09665 [hep-ex] (cit. on pp. 85, 138).

[195] ATLAS Collaboration, *Tagging and suppression of pileup jets with the ATLAS detector*, ATLAS-CONF-2014-018, 2014, URL: https://cds.cern.ch/record/1700870 (cit. on p. 85).

[196] ATLAS Collaboration, *Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector*, Eur. Phys. J. C **76** (2016) 581, arXiv: 1510.03823 [hep-ex] (cit. on p. 85).

[197] ATLAS Collaboration, *Performance of b-jet identification in the ATLAS experiment*, JINST **11** (2016) P04008, arXiv: 1512.01094 [hep-ex] (cit. on p. 85).

[198] ATLAS Collaboration, *Optimisation and performance studies of the ATLAS b-tagging algorithms for the 2017-18 LHC run*, ATL-PHYS-PUB-2017-013, 2017, URL: https://cds.cern.ch/record/2273281 (cit. on p. 85).

[199] ATLAS Collaboration, *Measurement of b-tagging efficiency of c-jets in $t\bar{t}$ events using a likelihood approach with the ATLAS detector*, ATLAS-CONF-2018-001, 2018, URL: https://cds.cern.ch/record/2306649 (cit. on p. 85).

[200] ATLAS Collaboration, *ATLAS b-jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV*, Eur. Phys. J. C **79** (2019) 970, arXiv: 1907.05120 [hep-ex] (cit. on p. 85).

[201] ATLAS Collaboration, *Performance of missing transverse momentum reconstruction with the ATLAS detector using proton–proton collisions at $\sqrt{s} = 13$ TeV*, Eur. Phys. J. C **78** (2018) 903, arXiv: 1802.08168 [hep-ex] (cit. on pp. 86, 138).

[202] S. van der Meer, *Calibration of the effective beam height in the ISR*, CERN-ISR-PO-68-31, 1968, URL: https://cds.cern.ch/record/296752 (cit. on pp. 87, 90).

[203] P. Grafström and W. Kozanecki, *Luminosity determination at proton colliders*, Prog. Part. Nucl. Phys. **81** (2015) 97 (cit. on pp. 87, 153).

[204] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC*, Eur. Phys. J. C **71** (2011) 1630, arXiv: 1101.2185 [hep-ex] (cit. on p. 87).

[205] ATLAS Collaboration, *Improved luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC*, Eur. Phys. J. C **73** (2013) 2518, arXiv: 1302.4393 [hep-ex] (cit. on pp. 87, 155).

[206] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector at the LHC*, Eur. Phys. J. C **76** (2016) 653, arXiv: 1608.03953 [hep-ex] (cit. on pp. 87, 91, 155).

[207] ATLAS Collaboration, *Luminosity determination for low-pileup datasets at $\sqrt{s} = 5$ and 13 TeV using the ATLAS detector at the LHC*, ATLAS-CONF-2020-023, 2020, URL: https://cds.cern.ch/record/2725195 (cit. on pp. 87, 90, 167, 168).

[208] ATLAS Collaboration, *Early Inner Detector Tracking Performance in the 2015 Data at $\sqrt{s} = 13$ TeV*, ATL-PHYS-PUB-2015-051, 2015, URL: https://cds.cern.ch/record/2110140 (cit. on p. 90).

[209] ATLAS Collaboration, *Measurements of the Higgs boson production and decay rates and coupling strengths using pp collision data at $\sqrt{s} = 7$ and 8 TeV in the ATLAS experiment*, Eur. Phys. J. C **76** (2016) 6, arXiv: 1507.04548 [hep-ex] (cit. on p. 116).

[210] CMS Collaboration, *Precise determination of the mass of the Higgs boson and tests of compatibility of its couplings with the standard model predictions using proton collisions at 7 and 8 TeV*, Eur. Phys. J. C **75** (2015) 212, arXiv: 1412.8662 [hep-ex] (cit. on p. 116).

[211] ATLAS Collaboration,
Evidence for the spin-0 nature of the Higgs boson using ATLAS data,
Phys. Lett. B **726** (2013) 120, arXiv: 1307.1432 [hep-ex] (cit. on p. 116).

[212] ATLAS Collaboration,
Combined measurements of Higgs boson production and decay using up to $80\,fb^{-1}$ *of proton–proton collision data at* $\sqrt{s} = 13$ *TeV collected with the ATLAS experiment*,
Phys. Rev. D **101** (2020) 012002, arXiv: 1909.02845 [hep-ex] (cit. on pp. 116, 146).

[213] CMS Collaboration, *Combined measurements of Higgs boson couplings in proton–proton collisions at* $\sqrt{s} = 13$ *TeV*, Eur. Phys. J. C **79** (2019) 421, arXiv: 1809.10733 [hep-ex] (cit. on p. 116).

[214] ATLAS Collaboration,
Measurements of Higgs boson production and couplings in the four-lepton channel in pp collisions at center-of-mass energies of 7 *and* 8 *TeV with the ATLAS detector*,
Phys. Rev. D **91** (2015) 012006, arXiv: 1408.5191 [hep-ex] (cit. on p. 116).

[215] ATLAS Collaboration, *Search for heavy resonances decaying into a pair of Z bosons in the* $\ell^+\ell^-\ell'^+\ell'^-$ *and* $\ell^+\ell^-\nu\bar{\nu}$ *final states using* $139\,fb^{-1}$ *of proton–proton collisions at* $\sqrt{s} = 13$ *TeV with the ATLAS detector*, Eur. Phys. J. C **81** (2020) 332, arXiv: 2009.14791 [hep-ex] (cit. on p. 117).

[216] T. Sjöstrand, S. Mrenna, and P. Skands, *A brief introduction to PYTHIA 8.1*,
Comput. Phys. Commun. **178** (2008) 852, arXiv: 0710.3820 [hep-ph] (cit. on p. 118).

[217] ATLAS Collaboration, *Summary of ATLAS Pythia 8 tunes*, ATL-PHYS-PUB-2012-003, 2012, URL: https://cds.cern.ch/record/1474107 (cit. on p. 118).

[218] A. D. Martin, W. J. Stirling, R. S. Thorne, and G. Watt, *Parton distributions for the LHC*,
Eur. Phys. J. C **63** (2009) 189, arXiv: 0901.0002 [hep-ph] (cit. on p. 118).

[219] P. Nason and G. Zanderighi, W^+W^-, WZ *and* ZZ *production in the POWHEG-BOX-V2*,
Eur. Phys. J. C **74** (2014) 2702, arXiv: 1311.1365 [hep-ph] (cit. on p. 118).

[220] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*,
Comput. Phys. Commun. **191** (2015) 159, arXiv: 1410.3012 [hep-ph] (cit. on p. 118).

[221] ATLAS Collaboration, *Measurement of the* Z/γ^* *boson transverse momentum distribution in pp collisions at* $\sqrt{s} = 7$ *TeV with the ATLAS detector*, JHEP **09** (2014) 145, arXiv: 1406.3660 [hep-ex] (cit. on p. 118).

[222] D. J. Lange, *The EvtGen particle decay simulation package*,
Nucl. Instrum. Meth. A **462** (2001) 152 (cit. on p. 118).

[223] H.-L. Lai et al., *New parton distributions for collider physics*,
Phys. Rev. D **82** (2010) 074024, arXiv: 1007.2241 [hep-ph] (cit. on p. 118).

[224] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*,
JHEP **07** (2014) 079, arXiv: 1405.0301 [hep-ph] (cit. on p. 118).

[225] ATLAS Collaboration, *ATLAS Pythia 8 tunes to* 7 *TeV data*, ATL-PHYS-PUB-2014-021, 2014, URL: https://cds.cern.ch/record/1966419 (cit. on p. 118).

[226] H. Davoudiasl, J. L. Hewett, and T. G. Rizzo,
Bulk gauge fields in the Randall–Sundrum model, Phys. Lett. B **473** (2000) 43,
arXiv: hep-ph/9911262 (cit. on p. 119).

[227] R. D. Ball et al., *Parton distributions for the LHC run II*, JHEP **04** (2015) 040,
arXiv: 1410.8849 [hep-ph] (cit. on p. 119).

[228] E. Bothmann et al., *Event generation with Sherpa 2.2*, SciPost Phys. **7** (2019) 034,
arXiv: 1905.09127 [hep-ph] (cit. on p. 119).

[229] T. Gleisberg and S. Höche, *Comix, a new matrix element generator*, JHEP **12** (2008) 039,
arXiv: 0808.3674 [hep-ph] (cit. on p. 119).

[230] F. Buccioni et al., *OpenLoops 2*, Eur. Phys. J. C **79** (2019) 866,
arXiv: 1907.13071 [hep-ph] (cit. on p. 119).

[231] F. Cascioli, P. Maierhöfer, and S. Pozzorini, *Scattering Amplitudes with Open Loops*,
Phys. Rev. Lett. **108** (2012) 111601, arXiv: 1111.5206 [hep-ph] (cit. on p. 119).

[232] A. Denner, S. Dittmaier, and L. Hofer,
Collier: A fortran-based complex one-loop library in extended regularizations,
Comput. Phys. Commun. **212** (2017) 220, arXiv: 1604.06792 [hep-ph] (cit. on p. 119).

[233] S. Schumann and F. Krauss,
A parton shower algorithm based on Catani–Seymour dipole factorisation,
JHEP **03** (2008) 038, arXiv: 0709.1027 [hep-ph] (cit. on p. 119).

[234] S. Höche, F. Krauss, M. Schönherr, and F. Siegert,
QCD matrix elements + parton showers. The NLO case, JHEP **04** (2013) 027,
arXiv: 1207.5030 [hep-ph] (cit. on p. 119).

[235] B. Biedermann, A. Denner, S. Dittmaier, L. Hofer, and B. Jäger,
Electroweak corrections to $pp \to \mu^+\mu^- e^+ e^- + X$ at the LHC: a Higgs background study,
Phys. Rev. Lett. **116** (2016) 161803, arXiv: 1601.07787 [hep-ph] (cit. on pp. 119, 139).

[236] B. Biedermann, A. Denner, S. Dittmaier, L. Hofer, and B. Jager, *Next-to-leading-order electroweak corrections to the production of four charged leptons at the LHC*,
JHEP **01** (2017) 033, arXiv: 1611.05338 [hep-ph] (cit. on pp. 119, 139).

[237] S. Kallweit, J. M. Lindert, P. Maierhöfer, S. Pozzorini, and M. Schönherr,
NLO electroweak automation and precise predictions for W+multijet production at the LHC,
JHEP **04** (2015) 012, arXiv: 1412.5157 [hep-ph] (cit. on p. 119).

[238] S. Kallweit, J. M. Lindert, S. Pozzorini, and M. Schönherr,
NLO QCD+EW predictions for $2\ell 2\nu$ diboson signatures at the LHC, JHEP **11** (2017) 120,
arXiv: 1705.00598 [hep-ph] (cit. on pp. 119, 139).

[239] M. Schönherr, *An automated subtraction of NLO EW infrared divergences*,
Eur. Phys. J. C **78** (2018) 119, arXiv: 1712.07975 [hep-ph] (cit. on p. 119).

[240] M. Grazzini, S. Kallweit, J. M. Lindert, S. Pozzorini, and M. Wiesemann, *NNLO QCD + NLO EW with Matrix+OpenLoops: precise predictions for vector-boson pair production*,
JHEP **02** (2020) 087, arXiv: 1912.00068 [hep-ph] (cit. on p. 119).

[241] F. Caola, K. Melnikov, R. Röntsch, and L. Tancredi,
QCD corrections to ZZ production in gluon fusion at the LHC,
Phys. Rev. D **92** (2015) 094028, arXiv: 1509.06734 [hep-ph] (cit. on p. 119).

[242] J. M. Campbell, R. K. Ellis, M. Czakon, and S. Kirchner,
Two loop correction to interference in $gg \to ZZ$, JHEP **08** (2016) 011,
arXiv: 1605.01380 [hep-ph] (cit. on p. 119).

[243] S. Alioli, F. Caola, G. Luisoni, and R. Röntsch,
ZZ production in gluon fusion at NLO matched to parton-shower,
Phys. Rev. D **95** (2017) 034042, arXiv: 1609.09719 [hep-ph] (cit. on p. 119).

[244] K. Melnikov and M. Dowling,
Production of two Z-bosons in gluon fusion in the heavy top quark approximation,
Phys. Lett. B **744** (2015) 43, arXiv: 1503.01274 [hep-ph] (cit. on p. 119).

[245] C. S. Li, H. T. Li, D. Y. Shao, and J. Wang,
Soft gluon resummation in the signal-background interference process of $gg(\to h^*) \to ZZ$,
JHEP **08** (2015) 065, arXiv: 1504.02388 [hep-ph] (cit. on p. 119).

[246] R. Gavin, Y. Li, F. Petriello, and S. Quackenbush,
FEWZ 2.0: A code for hadronic Z production at next-to-next-to-leading order, (2010) 2388,
arXiv: 1011.3540 [hep-ph] (cit. on p. 119).

[247] S. Frixione, V. Hirschi, D. Pagani, H. .-S. Shao, and M. Zaro, *Electroweak and QCD corrections to top-pair hadroproduction in association with heavy bosons*,
JHEP **06** (2015) 184, arXiv: 1504.03446 [hep-ph] (cit. on p. 120).

[248] J. M. Campbell, R. K. Ellis, and C. Williams, *Vector boson pair production at the LHC*,
JHEP **07** (2011) 018, arXiv: 1105.0020 [hep-ph] (cit. on p. 120).

[249] N. Kauer and G. Passarino,
Inadequacy of zero-width approximation for a light Higgs boson signal, JHEP **08** (2012) 116,
arXiv: 1206.4803 [hep-ph] (cit. on p. 120).

[250] N. Kauer,
Interference effects for $H \to WW/ZZ \to \ell\bar{\nu}_\ell\bar{\ell}\nu_\ell$ searches in gluon fusion at the LHC,
JHEP **12** (2013) 082, arXiv: 1310.7011 [hep-ph] (cit. on p. 120).

[251] ATLAS Collaboration, *2015 start-up trigger menu and initial performance assessment of the ATLAS trigger using Run-2 data*, ATL-DAQ-PUB-2016-001, 2016,
URL: https://cds.cern.ch/record/2136007 (cit. on p. 120).

[252] ATLAS Collaboration, *Trigger Menu in 2016*, ATL-DAQ-PUB-2017-001, 2017,
URL: https://cds.cern.ch/record/2242069 (cit. on p. 120).

[253] ATLAS Collaboration,
Measurement of the Higgs boson mass from the $H \to \gamma\gamma$ and $H \to ZZ^* \to 4\ell$ channels in pp collisions at center-of-mass energies of 7 and 8 TeV with the ATLAS detector,
Phys. Rev. D **90** (2014) 052004, arXiv: 1406.3827 [hep-ex] (cit. on p. 122).

[254] ATLAS Collaboration, *Higgs boson production cross-section measurements and their EFT interpretation in the 4ℓ decay channel at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Eur. Phys. J. C **80** (2020) 957, arXiv: 2004.03447 [hep-ex] (cit. on p. 122), Erratum: Eur. Phys. J. C **81** (2021) 29.

[255] F. Chollet et al., *Keras*, 2015, URL: https://github.com/fchollet/keras (cit. on p. 122).

[256] Martín Abadi et al., *TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems*, 2015, URL: https://www.tensorflow.org/ (cit. on p. 122).

[257] W. Buttinger and M. Lefebvre, *Formulae for Estimating Significance*, tech. rep. ATL-COM-GEN-2018-026, CERN, 2018, URL: https://cds.cern.ch/record/2643488 (cit. on p. 124).

[258] ATLAS Collaboration, *Measurement of inclusive and differential cross sections in the $H \to ZZ^* \to 4\ell$ decay channel in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, JHEP **10** (2017) 132, arXiv: 1708.02810 [hep-ex] (cit. on p. 126).

[259] ATLAS Collaboration, *Measurement of the Higgs boson coupling properties in the $H \to ZZ^* \to 4\ell$ decay channel at $\sqrt{s} = 13$ TeV with the ATLAS detector*, JHEP **03** (2018) 095, arXiv: 1712.02304 [hep-ex] (cit. on p. 126).

[260] M. Pivk and F. R. Le Diberder, *SPlot: A Statistical tool to unfold data distributions*, Nucl. Instrum. Meth. A **555** (2005) 356, arXiv: physics/0402083 (cit. on p. 127).

[261] M. Oreglia, *A Study of the Reactions $\psi' \to \gamma\gamma\psi$*, 1980, URL: https://www.slac.stanford.edu/cgi-wrap/getdoc/slac-r-236.pdf (cit. on p. 127).

[262] J. Gaiser, *Charmonium Spectroscopy From Radiative Decays of the J/ψ and ψ'*, 1982, URL: https://www.slac.stanford.edu/cgi-wrap/getdoc/slac-r-255.pdf (cit. on p. 127).

[263] S. Goria, G. Passarino, and D. Rosco, *The Higgs Boson Lineshape*, Nucl. Phys. B **864** (2012) 530, arXiv: 1112.5517 [hep-ph] (cit. on pp. 128–130).

[264] M. Spira, A. Djouadi, D. Graudenz, and P. M. Zerwas, *Higgs boson production at the LHC*, Nucl. Phys. B **453** (1995) 17, arXiv: hep-ph/9504378 (cit. on p. 129).

[265] R. D. Ball et al., *Parton Distribution Benchmarking with LHC Data*, JHEP **04** (2013) 125, arXiv: 1211.5142 [hep-ph] (cit. on p. 130).

[266] J. Bijnens, P. Eerola, M. Maul, A. Mansson, and T. Sjostrand, *QCD signatures of narrow graviton resonances in hadron colliders*, Phys. Lett. B **503** (2001) 341, arXiv: hep-ph/0101316 (cit. on p. 131).

[267] N. Kauer and C. O'Brien, *Heavy Higgs signalbackground interference in $gg \to VV$ in the Standard Model plus real singlet*, Eur. Phys. J. C **75** (2015) 374, arXiv: 1502.04113 [hep-ph] (cit. on p. 133).

[268] ATLAS Collaboration, *Object-based missing transverse momentum significance in the ATLAS Detector*, ATLAS-CONF-2018-038, 2018, URL: https://cds.cern.ch/record/2630948 (cit. on p. 135).

[269] ATLAS Collaboration, *Measurement of $W^{\pm}Z$ production cross sections and gauge boson polarisation in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Eur. Phys. J. C **79** (2019) 535, arXiv: 1902.05759 [hep-ex] (cit. on p. 136).

[270] ATLAS Collaboration, *Observation of electroweak $W^{\pm}Z$ boson pair production in association with two jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Phys. Lett. B **793** (2019) 469, arXiv: 1812.09740 [hep-ex] (cit. on p. 136).

[271] M. Baak, S. Gadatsch, R. Harrington, and W. Verkerke, *Interpolation between multi-dimensional histograms using a new non-linear moment morphing method*, Nucl. Instrum. Meth. A **771** (2015) 39, arXiv: 1410.7388 [physics.data-an] (cit. on p. 137).

[272] ATLAS Collaboration, *Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV*, Eur. Phys. J. C **76** (2016) 292, arXiv: 1603.05598 [hep-ex] (cit. on p. 138).

[273] J. Butterworth et al., *PDF4LHC recommendations for LHC Run II*, J. Phys. G **43** (2016) 023001, arXiv: 1510.03865 [hep-ph] (cit. on p. 138).

[274] ATLAS and CMS Collaboration, *Procedure for the LHC Higgs boson search combination in summer 2011*, ATL-PHYS-PUB-2011-011, CMS-NOTE-2011-005, 2011, URL: https://cds.cern.ch/record/1375842 (cit. on p. 139).

[275] ATLAS Collaboration, *Combined search for the Standard Model Higgs boson in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, Phys. Rev. D **86** (2012) 032003, arXiv: 1207.0319 [hep-ex] (cit. on p. 139).

[276] A. L. Read, *Presentation of search results: the CL_S technique*, J. Phys. G **28** (2002) 2693 (cit. on p. 144).

[277] O. Aberle et al., *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*, CERN Yellow Reports: Monographs, Geneva: CERN, 2020, URL: https://cds.cern.ch/record/2749422 (cit. on p. 151).

[278] W. Bonivento, *Online luminosity monitoring with liquid Argon calorimeters at ATLAS and D0*, tech. rep., CERN, 2000, URL: https://cds.cern.ch/record/684140 (cit. on p. 154).

[279] ATLAS Collaboration, *Measurements of the pseudorapidity dependence of the total transverse energy in proton–proton collisions at $\sqrt{s} = 7$ TeV with ATLAS*, JHEP **11** (2012) 033, arXiv: 1208.6256 [hep-ex] (cit. on p. 153).

[280] ATLAS Collaboration, *Charged-particle distributions at low transverse momentum in $\sqrt{s} = 13$ TeV pp interactions measured with the ATLAS detector at the LHC*, Eur. Phys. J. C **76** (2016) 502, arXiv: 1606.01133 [hep-ex] (cit. on p. 153).

[281] J. P. Archambault et al., *Performance of the ATLAS liquid argon forward calorimeter in beam tests*, JINST **8** (2013) P05006 (cit. on p. 153).

[282] J. P. Rutherfoord,
Signal degradation due to charge buildup in noble liquid ionization calorimeters,
Nucl. Instrum. Meth. A **482** (2002) 156 (cit. on p. 154).

[283] A. Afonin et al.,
Relative luminosity measurement of the LHC with the ATLAS forward calorimeter,
JINST **5** (2010) P05005, arXiv: 1005.1784 [physics.ins-det] (cit. on p. 155).

[284] S. Arfaoui, *The ATLAS liquid argon calorimeter high-voltage system: commissioning, optimisation, and LHC relative luminosity measurement*,
PhD thesis: Marseille, CPPM, 2011,
URL: https://tel.archives-ouvertes.fr/tel-00658194 (cit. on p. 155).

[285] A. Barriuso Poy et al., *The detector control system of the ATLAS experiment*,
JINST **3** (2008) P05006 (cit. on p. 157).

[286] J. W. S. Carter,
Luminosity monitoring with the ATLAS liquid-argon calorimeters in Run 2 of the LHC,
tech. rep., CERN, 2021, URL: https://cds.cern.ch/record/2776367 (cit. on p. 158).