ATLAS Hadronic Chronicles, Triggering Beauty and Searching for Multijet Resonances

Franchellucci, Stefano

How to cite

FRANCHELLUCCI, Stefano. ATLAS Hadronic Chronicles, Triggering Beauty and Searching for Multijet Resonances. Doctoral Thesis, 2025. doi: 10.13097/archive-ouverte/unige:185300

This publication URL: https://archive-ouverte.unige.ch/unige:185300
Publication DOI: 10.13097/archive-ouverte/unige:185300

UNIVERSITÉ DE GENÈVE
Département de physique nucléaire et corpusculaire

FACULTÉ DES SCIENCES
Professeure Anna Sfyrla

ATLAS Hadronic Chronicles

Triggering Beauty and Searching for Multijet Resonances

THÈSE

présentée à la Faculté des sciences de l'Université de Genève
pour obtenir le grade de Docteur ès sciences, mention physique

par

Stefano Franchellucci
de Roma (Italie)

Thèse N° 5903

GENÈVE
2025

UNIVERSITÉ
DE GENÈVE
FACULTÉ DES SCIENCES

DOCTORAT ÈS SCIENCES, MENTION PHYSIQUE

Thèse de Monsieur Stefano FRANCHELLUCCI

intitulée :

«ATLAS Hadronic Chronicles
Triggering Beauty and Searching for Multijet Resonances»

La Faculté des sciences, sur le préavis de

Madame A. SFYRLA, professeure associée et directrice de thèse
Département de physique nucléaire et corpusculaire

Monsieur F. SANCHEZ NIETO, professeur ordinaire
Département de physique nucléaire et corpusculaire

Madame T. AARRESTAD, docteure
Department of Physics, Institute for Particle Physics and Astrophysics, Eidgenössische Technische Hochschule Zürich (ETH Zürich), Zürich, Switzerland

Monsieur C. POLLARD, professeur
Department of Physics, Faculty of Science, University of Warwick, Coventry, United Kingdom

autorise l'impression de la présente thèse, sans exprimer d'opinion sur les propositions qui y sont énoncées.

Genève, le 28 avril 2025

Thèse - 5903 -

La Doyenne

N.B. - La thèse doit porter la déclaration précédente et remplir les conditions énumérées dans les "Informations relatives aux thèses de doctorat à l'Université de Genève".

Résumé

L'expérience ATLAS constitue l'une des plus grandes collaborations scientifiques de l'histoire des sciences naturelles, explorant les phénomènes physiques aux plus petites échelles sondées. Les particules élémentaires connues et leurs interactions fondamentales sont décrites par le Modèle Standard (MS). Cette théorie permet des prédictions extrêmement précises, mais présente également des lacunes importantes, comme l'origine de la *Matière Noire* ou le problème de la *naturalité* de la masse du boson de Higgs. ATLAS effectue des mesures de précision pour approfondir notre compréhension du MS, tout en recherchant des indices de nouvelle physique.

Le système de déclenchement (Trigger) est un élément essentiel du détecteur ATLAS. Il sélectionne les événements de collision à enregistrer en effectuant une analyse en temps réel des signaux du détecteur. Les décisions du Trigger réduisent le taux d'événements de près de 40 MHz à environ 3 kHz. Cette réduction substantielle doit préserver une efficacité maximale pour les événements de physique pertinents. Cette thèse met en œuvre une nouvelle méthodologie de déclenchement : la présélection par "fast b-tagging". Le fast b-tagging filtre les événements contenant des jets hadroniques avec des hadrons de type b dès les premières étapes du Trigger de haut niveau (HLT). Cette présélection repose sur un algorithme d'apprentissage automatique (AA) qui analyse les traces à basse résolution pour identifier la saveur des jets. Cette approche a été déployée dans le HLT pour le début du Run 3, optimisant les ressources de calcul et permettant ainsi l'utilisation de seuils plus souples pour la discrimination des événements. Une amélioration de l'efficacité de près de 50% est obtenue pour des processus physiques clés, telles que les désintégrations $\mathrm{HH} \to \mathrm{b\overline{b}b\overline{b}}$. En outre, la mise en œuvre de l'identification de la saveur des jets dans les premières phases du HLT a permis d'introduire le b-tagging pour l'analyse des données au niveau du Trigger.

La supersymétrie (SUSY) étend le MS en fournissant un cadre riche pour la nouvelle physique. Dans les théories SUSY, le problème de la naturalité du Higgs disparaît et de nouvelles particules de matière noire peuvent être introduites. Cette thèse a examiné les données d'ATLAS collectées de 2015 à 2018, à la recherche de désintégrations résonantes de particules SUSY. Les couplages violant la parité R ont été prises en compte, produisant des états finaux entièrement hadroniques. Pour traiter le bruit de fond complexe dû aux interactions fortes, deux approches ont été utilisées : une méthode améliorée de "cut-and-count" et une nouvelle technique basée sur un algorithme de AA. Les deux approches ont considérablement augmenté la sensibilité par rapport aux analyses précédentes. Les événements observés sont compatibles avec le bruit de fond attendu, donc des limites de section efficace ont été fixées pour plusieurs modèles SUSY. De plus, les résultats obtenus avec la méthode cut-and-count ont été interprétés dans le contexte de théories alternatives, contraignant une large portion de l'espace des paramètres de SUSY.

Abstract

The ATLAS experiment comprises one of the largest scientific collaborations in the history of the natural sciences, exploring physical phenomena at the smallest probed scales. The known elementary particles and their fundamental interactions are described by the Standard Model (SM). This theory makes exceptionally precise predictions, but also has compelling shortcomings, such as the origin of Dark Matter or the naturalness problem of the Higgs boson mass. ATLAS performs precision measurements to deepen our understanding of the SM, along with searches for hints of new physics.

The Trigger system is an essential component of the ATLAS detector. It selects which collision events are recorded by performing real-time analysis of detector signals. Trigger decisions reduce the event rate from almost 40 MHz to roughly 3 kHz. This significant reduction must preserve maximal efficiency for relevant physics events. In this thesis, a novel trigger methodology is implemented: the fast b-tagging preselection, filtering events containing b-jets in early stages of the High-Level Trigger (HLT). This approach relies on a Machine Learning algorithm, which analyses coarse tracks to tag jets. This preselection was deployed in HLT for the beginning of Run 3, optimising computing resources and allowing looser thresholds for event discrimination. An efficiency enhancement of nearly 50% is achieved for key physics signatures, such as $\mathrm{HH} \to \mathrm{b\bar{b}b\bar{b}}$ decays. Furthermore, implementing jet flavour tagging in early stages of the HLT allowed the introduction of b-tagging in the Trigger-level analysis data stream.

Supersymmetry (SUSY) provides a rich framework for new physics, extending the SM. In SUSY theories, the Higgs naturalness problem vanishes, and new DM candidate particles are allowed. This thesis scrutinised the full Run 2 ATLAS dataset in a search for resonant decays of heavy SUSY particles. Decay channels through R-parity violating couplings were considered, which lead to less-constrained fully-hadronic final states. To deal with the challenging background arising from strong interactions, two alternative approaches were used: an improved cut-and-count method and a new technique that leverages ML algorithms. Both approaches led to significant sensitivity improvements compared to previous analyses. No excesses were found with respect to the expected background, and cross-section limits were set on several SUSY models. Furthermore, the results obtained with the cut-and-count were interpreted in the context of alternative theories, constraining an extended portion of the SUSY parameter space.

Acknowledgments

I am extremely grateful for all the awesome people I have met through this doctoral journey. The University of Geneva, CERN, and the ATLAS collaboration allowed me to connect with great scientists, as well as with great persons. This thesis would not have been possible without many of them.

Firstly, I want to express my deepest gratitude to Anna, an amazing mentor that always provided wise guidance and strong support. Your endless energy, kindness, and motivation are always inspiring. It was really a pleasure being part of your group for the past few years. All the post-docs were also precious mentors. I would like to thank Chiara, for introducing me to ATLAS and to SUSY searches, Pantelis, for joining the effort and helping me when it was most needed (!), and Claire, for introducing me to the trigger and the fascinating data scouting world. I am grateful for all your support, patience, and willingness to share knowledge. Then, I want to thank all the other PhDs with whom I shared this journey closely, including a lot of coffees and the desperation over grid jobs: Kostas, Ondrej, Noshin, Leon, Lucas and Edward. Additionally, I would like to thank Quentin, for always having interesting discussions with great kindness.

My gratitude goes to many collaborators in ATLAS. I would like to start mentioning Chris and Dan, for taking me onboard in the b-jet trigger group, showing immense patience with my many questions and always providing very insightful feedback. I would have been still lost in some dark corner of Athena without your help. And I would like to thank in general all the great b-jet trigger folks, connecting two of the most fun areas in ATLAS, trigger and flavour tagging! Then, I want to mention some of the contributors to the multijet analysis with whom I had the pleasure of working more closely: Anthony, Javier, John, Jona, Lea, and Will. That was quite a ride! Furthermore, I thank all the awesome TLA team. Despite being a minor fraction of this manuscript, TLA has a great place in my heart, together with all the fantastic people pursuing it!

Finally, I would like to thank all my friends and my family for the endless support and affection received throughout these years. I owe you everything. Especially Domitilla, this thesis is dedicated to you.

Contents

Introduction

Understanding the mechanisms that govern physical phenomena at the subatomic scale has attracted the curiosity of a vast portion of the scientific community during the past decades. Unveiling the fundamental constituents of matter and their interactions has brought new insights into the functioning of our Universe at both microscopic and macroscopic scales. Particle physics focuses on studying such elementary components.

The recent discovery of the Higgs boson has empowered the electroweak theory and provided a mathematical description of most of the ordinary mass. The current theoretical framework that describes elementary particles and their interactions, the Standard Model (SM), makes predictions with unprecedented precision. However, the SM has compelling shortcomings. It lacks a mechanism to give mass to the elusive neutrinos and raises serious doubts about the mass of the Higgs boson. The SM cannot account for Dark Matter (DM), which appears to be the dominant mass component of our Universe based on astrophysical measurements. Furthermore, the full structure of the Higgs potential remains unknown, and it may have substantial consequences for cosmological phenomena. Hence, it is essential to pursue measurements of the SM and the Higgs boson, and to search for Beyond Standard Model (BSM) physics.

The ATLAS experiment at the Large Hadron Collider (LHC) is a general-purpose experiment. It is designed to both improve our understanding of the SM and conduct BSM searches, measuring the products of energetic particle collisions. From the Higgs boson's joint discovery with the Compact Muon Solenoid (CMS) experiment to in-depth explorations of BSM models, ATLAS has profoundly shaped the field of particle physics over the past fifteen years. This thesis was carried out within the ATLAS Collaboration, analysing data from the LHC Run 2 data acquisition campaign, and improving the experiment's physics reach for Run 3.

A core component of the ATLAS detector is its Trigger system, which is responsible for deciding which collision events are recorded. Not all events can be saved to disk, as the combination of data size and collision frequency delivers an enormous throughput that current technologies cannot sustain. The Trigger performs real-time analysis of detector signals, making prompt decisions that should be as inclusive as possible for the relevant physics processes. This thesis presents a novel trigger methodology: the *fast*

b-*tagging* preselection. This approach can efficiently select events with jets – collimated and energetic hadronic showers that deposit energy in the ATLAS calorimeter – containing b-hadrons. It leverages Machine Learning (ML) algorithms to quickly process coarse reconstructed objects, providing a powerful early rejection tool. Fast b-tagging was introduced at the start of Run 3 and allowed for looser trigger thresholds, improving the acceptance for key physics signatures, such as $\mathrm{HH} \to \mathrm{b\bar{b}b\bar{b}}$ decays, by almost 50%. Furthermore, using a similar approach, jet flavour tagging capabilities were implemented in the Trigger-Level Analysis (TLA) data stream. This stream uses exclusively real-time reconstructed objects to perform physics analyses, and flavour tagging extends its sensitivity to a wider range of processes than previously accessible.

This thesis also searches for new physics using Run 2 data. Events with high jet multiplicities are investigated to search for signs of new resonances. In some BSM scenarios, such as supersymmetry (SUSY), heavy particles could be pair-produced at the LHC and decay to an all hadronic final state, generating many energetic jets. The background arising from QCD interactions makes this signature particularly challenging to study. This thesis presents an analysis using two complementary approaches, improving previous methodologies and introducing a new technique based on ML to reconstruct the heavy resonances.

This manuscript is structured in three parts. The first part sets the stage for contemporary particle physics, and explains how collision events are studied at the ATLAS experiment. Jet flavour tagging and the ATLAS b-jet trigger are introduced in the second part before presenting the fast b-tagging method. The last part focuses on the search for new phenomena in multijet final states, explaining the analysis strategy and discussing the observed results.

Choice of units

This thesis uses the *natural units* convention that is widely adopted in the field of particle physics. In this schema, the reduced Planck constant $\hbar$ and the speed of light c are collapsed to unity

$$\hbar = \mathrm{c} = 1 \ .$$

This allows for a substantial simplification in the formulation of equations and sets the energy, momentum, mass, and inverse of time to have the same physical dimensions. The standard unit of measure for such qualities is the electronvolt (eV), which is the energy gained by an electron in an electric potential of 1 V and corresponds to:

$$1\ \mathrm{eV} = 1.602\,176\,634 \cdot 10^{-19}\ \mathrm{J} \ .$$

Personal contributions

The experiments installed on the LHC ring are conducted by scientific collaborations that are among the largest in the history of the natural sciences. The ATLAS Collaboration counts more than 170 institutions and over 3000 scientific authors. Probing the smallest known elements of nature requires the largest endeavours. In this context, the work presented in this thesis would not have been possible without the effort of a considerable amount of people. The thesis' author specific contributions are listed in the following paragraphs.

Part II The author actively participated in the b-jet trigger signature group, specifically in the preparations for the start of Run 3 and in performance studies. He was part of the team that implemented the fast b-tagging preselection, making substantial contributions that are discussed in Chapter 6. The author trained Neural Network algorithms and studied their performance in simulation, exported the models to the ATLAS trigger software framework, and validated them. He also designed and carried out studies on fast b-tagging performance using data. Furthermore, he contributed to the periodic validation of the trigger b-taggers. The most recent work of the author resulted in the validation of the new workflow for data acquisition using the Partial Event Building technique (Chapter 6.4), which enhances flavour tagging for TLA.

Part III Several members of the Collaboration participated in the analysis, which required a few years of R&D before converging on a concrete strategy. The thesis author is primarily responsible for the cut-and-count approach presented in Chapter 8. He defined and optimised the signal region selections, improved and implemented the background method together with its validation. The author evaluated the systematic errors and contributed to the statistical interpretation of the result. Regarding the Mass Resonance method (Chapter 9), the author participated mainly in strategy discussions. The development and studies for reinterpreting the results with alternative models, described in Chapter 10, are also the result of the author's work, along with the trigger efficiency estimations from Chapter 10.2.

The work of this thesis is documented in several publications. Fast b-tagging is mostly discussed in Ref. [1], with further studies presented in Ref. [2]. Its broader impact on the b-jet trigger is detailed in Ref. [3]. The search for new physics in multijet events is published in Ref. [4], while, at the time of writing, the reinterpretation document is under internal review.

I

Particle physics and the ATLAS Experiment

Chapter 1

A world of interacting particles

The foundational research question of modern particle physics is to discover all the basic building blocks of matter – the *elementary particles* – and understanding their *fundamental interactions*. Such physical entities are elementary in the sense that cannot be divided into smaller constituents, and interact through a limited number of forces that are irreducible and universal. Macroscopic physical phenomena are expected to emerge from this underlying set of elements and natural laws.

Over the past century, a comprehensive picture of the subatomic world has been developed, identifying many particles that are, to the best of our knowledge, elementary. The Standard Model (SM) encodes a significant portion of their dynamics within a coherent theory. It successfully unifies the electromagnetic and weak forces and is able to make accurate predictions that were verified experimentally with an extraordinary level of precision. However, despite its glorious success, the SM remains an effective theory, meaning that its parameters cannot be derived from first principles, and it comes with several limitations. Many pieces of the particle physics' puzzle are still missing.

This chapter gives an overview of the SM, discussing some key aspects. It also explains the shortcomings of the theory, which demand for the existence of new physics or for the formulation of new models, beyond what is currently known. Two important theoretical frameworks to extend the SM are finally discussed.

1.1 Overview of the Standard Model

Fundamental forces and particles

In total, four fundamental forces are known: the gravitational, electromagnetic, weak and strong forces. The gravitational force, while dominant at macroscopic scales, has a negligible effect at the microscopic scale. Its coupling strength is extremely weak compared to the other three. The electromagnetic force governs interactions between charged particles and is significant at both macroscopic and microscopic scales. The weak force is responsible for processes such as β-decay and plays a crucial role in nuclear reactions, but it has a very short range and has no influence at large radius. The strong force is also short ranged, and binds protons and neutrons together within atomic nuclei, overcoming the electromagnetic repulsion between similarly charged particles.

Many subatomic particles have been discovered in the past century [5–16]. Only a few are stable and can form more complex structures such as atoms and molecules. Most particles decay within fractions of a second, requiring vast scientific endeavours to produce and observe them. Amongst all the discovered particles, few of them are elementary and are commonly classified by their quantum numbers and their interactions. Figure 1.1 shows an overview of the SM content, which reports the particle masses and their main quantum numbers.

Quarks The constituents of the protons and neutrons are called *quarks*. Combinations of the *up* and *down* quark (u and d) make up most of the directly observable mass in the universe. They are fermions (spin-$\frac{1}{2}$) and have respectively an electric charge of $+\frac{2}{3}$ and $-\frac{1}{3}$. Together they form a “generation” of quarks. Two more generations of quark pairs exist, and have the same exact quantum numbers as u and d but with heavier masses: the *charm* (c) and *strange* (s), and the *top* (t) and *bottom* (b) quarks. These three families, and their corresponding antiparticles ($i.e.$, identical particles but with conjugate electric charge), are the only elementary fermions that are sensitive to the strong force.

Leptons The other class of elementary fermions is the *leptons*. Similarly to the quarks, they form three generations: the *electron* (e), *muon* (μ) and *tau* (τ) families. Each generation is composed of an electrically charged lepton ℓ and a corresponding neutrino ν_ℓ. The latter type of particle interacts only through the weak force, making it difficult to detect.

Gauge Bosons The fundamental forces are carried by the *gauge bosons*. The photon (γ) is the mediator of the electromagnetic interaction, while the $W^\pm$ and the Z bosons are responsible for the weak force. The gluons (g) mediate the strong force between quarks and *hadrons* – composite bound states of quarks – by carrying the colour charge.

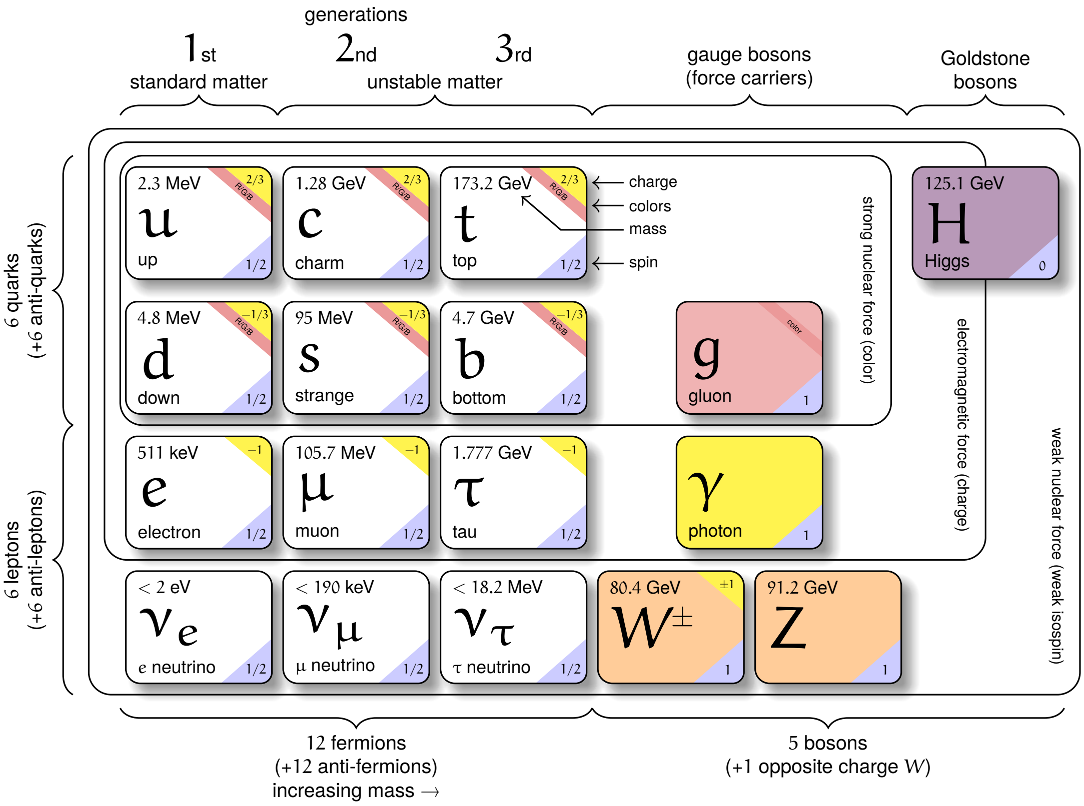

Figure 1.1: Schema of all the known elementary particles present in the SM[1]. The 12 fermions compose the matter particles, and the bosons are responsible for carrying forces.

The last remaining fundamental force, the gravitational interaction, has no known particle corresponding to it. Some extensions of the SM predict a carrier of gravity, the graviton, but it has not yet been discovered. All the gauge bosons are vector bosons, *i.e.*, they have spin 1.

Higgs Boson The only fundamental scalar (spin-0) particle is the Higgs boson. It was predicted by the (Englert-Brout-)Higgs Mechanism [17, 18], which is the means to provide mass to SM particles by breaking the electroweak symmetry. It can directly interact with the weak bosons, but has no colour or electric charge.

Theoretical framework of the SM

The SM is a relativistic Quantum Field Theory (QFT) that describes the electromagnetic, weak, and strong interactions amongst the elementary particles. It is based on a minimum number of assumptions, which are mostly symmetry principles. The importance of symmetries derives from their natural way of introducing conservation laws. According to Noether's Theorem [19] each symmetry implies the conservation of a *current*. And con-

---

[1] Original LaTeX template from: https://texample.net/tikz/examples/model-physics.

servation principles are essential tools for describing the physical world. In this section only a restricted selection of SM aspects and concepts is briefly reviewed.

Invariance principles The SM requires *Lorentz invariance*, implying that the laws of physics do not depend on observer's frame of reference. This is equivalent to imposing that the equations of motion are *covariant* with respect to changes in the observer's frame. Also, to be a consistent relativistic theory and not break *causality* ($i.e.$, allowing for causal relations only between events separated by *light*- or *time-like* distances), the SM is a *local* theory in which only contact action takes place. This has the strong physical implication that interactions propagate within a finite time, and occur through contact with a mediator of the force. For a relativistic field theory, locality implies that the Lagrangian of the system cannot depend on the distances between space-time points. It must be a function of the fields and their derivatives, and the fields must depend only on x_μ, the space-time coordinates: $L = L\left(\phi, \partial_\mu\phi\right)$, were $\phi = \phi(x_\mu)$ and $\partial_\mu = \frac{\partial}{\partial x_\mu}$.

The action $\mathcal{S} = \int dt\ L$ is a scalar, and therefore Lorentz invariant, but L is not. Instead, the equations of motion are derived from the *Lagrangian density* $\mathcal{L}(\phi, \partial_\mu\phi)$, such that $\mathcal{S} = \int dt \int d^3x\ \mathcal{L}$, ensuring that the Euler-Lagrange equations are covariant.

Imposing that $\mathcal{L}$ is invariant under local field transformations gives important properties to the theory. For example, considering a system of fermions of mass m, the Lagrangian density is:

$$\mathcal{L} = \overline{\psi}(i\gamma^\mu\partial_\mu - m)\psi\ , \tag{1.1}$$

where the ψ is the four-component spinorial field, $\overline{\psi}$ is its adjoint field, and γ^μ the *gamma* matrices introduced by Dirac [20]. The $\mathcal{L}$ invariance under a local phase $\theta(x_\mu)$ transformation:

$$\psi \rightarrow \psi' = e^{i\theta(x)}\psi\ , \tag{1.2}$$

requires the introduction of a new vectorial (spin-1) field A_μ that couples to the fermions, and has a dynamic which is does not change for *gauge* transformations:

$$A_\mu \rightarrow A'_\mu = A_\mu - \tfrac{1}{e}\partial_\mu\theta\ . \tag{1.3}$$

Then, $\mathcal{L}$ can be written in a compact form using the *covariant derivative* $D_\mu = \partial_\mu + ieA_\mu$, such that it is symmetric with respect to the $U(1)$ transformation group defined by the local phase θ:

$$\mathcal{L} = \overline{\psi}(i\gamma^\mu D_\mu - m)\psi = \overline{\psi}(i\gamma^\mu\partial_\mu - m)\psi - e\overline{\psi}\gamma^\mu\psi A_\mu\ . \tag{1.4}$$

This property is known as the *gauge invariance* of a theory. By imposing this type of symmetry condition, a gauge boson (A_μ) is introduced and its interactions with other fields in the Lagrangian are automatically fixed.

Gauge invariance constrains the additional terms that describe the kinematics of free gauge boson fields. For example, for A_μ the field's strength tensor is defined as $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$, and the following kinematic term can be added: $\mathcal{L}_A = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu}$. For the electromagnetic field, the Maxwell equations can be derived from this Lagrangian. This equation preserves gauge invariance, while adding mass terms such as $\sim m^2 A^\mu A_\mu$ would break it. Therefore, gauge bosons are introduced as massless particles.

All these considerations generalize to more elaborate transformation groups. For non-Abelian groups, such as in Yang-Mills theories [21], additional contributions are required to accommodate the non-commutative structure of the algebra.

SM forces The SM adopts the gauge principle to describe the fundamental forces, and its symmetry groups are:

$$\mathrm{SU}(3)_\mathrm{C} \otimes \mathrm{SU}(2)_\mathrm{L} \otimes \mathrm{U}(1)_\mathrm{Y} \ . \tag{1.5}$$

These groups define the basic structure of the SM Lagrangian, and in particular of all the gauge interactions.

The number of generators of each group corresponds to the number of gauge bosons. The $\mathrm{SU}(3)_\mathrm{C}$ group describes the strong interactions, and adds eight bosons, interpreted as the gluons carrying the colour charge. The gluons can interact with quark fields or with each other. The description of this part of the SM is known as Quantum Chromodynamics (QCD).

The $\mathrm{SU}(2)_\mathrm{L} \otimes \mathrm{U}(1)_\mathrm{Y}$ groups describe the electroweak sector, a unification of Quantum Electrodynamics (QED) and the weak force. The $3+1$ gauge bosons introduced are combined into the four physically observable bosons: $W^\pm$, Z, and γ. As the electric charge (Q) is associated with QED and colour charge with QCD, particles have a weak *hypercharge* (Y_W). In the electroweak unification, Y_W and Q are connected through the third component of the *weak isospin*[2](T): $T_3 = Q - \frac{1}{2}Y_W$. The weak interaction eigenstates are arranged into doublets of $T = \frac{1}{2}$ or $T = 0$ singlets, depending on their *chirality*[3]. Each generation of left(right)-handed (anti-)fermions – up-type with down-type quarks, and leptons with corresponding neutrinos – forms T doublets, determining the allowed types of charged weak currents, $i.e.$, the couplings to the $W^\pm$ bosons. While the right(left)-handed (anti-) fermions form weak isospin singlets, related to the couplings with the neutral electroweak gauge bosons.

Weak interaction universality Form observations on the phenomenology of light meson decays to leptons, Cabibbo suggested that the weak interaction coupling strength

---

[2] The *isospin* quantum number was first introduced to describe the spectrum of hadrons. It is described using the same formalism of a dimensionless spin.

[3] The *chirality* of a particle is determined by its behaviour under Lorentz transformations in the *chiral* representation (right- and left-handed spinors). For massless particles it coincides with the *helicity*, the projection of the spin along the direction of the particles motion, and is typically conserved.

is universal, requiring the quark weak interaction eigenstates to differ from the mass (or strong force) eigenstates [22]. According to the Glashow, Iliopoulos and Maiani mechanism [23], the two sets are connected by a unitary rotation matrix. Even before the c-quark discovery, the mechanism was extended to three quark generations with the Cabibbo-Kobayashi-Maskawa (CKM) matrix [24, 25]:

$$V_{CKM} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} = \begin{pmatrix} 1-\frac{1}{2}\lambda^2 & \lambda & A\lambda^3\rho e^{i\delta} \\ -\lambda & 1-\frac{1}{2}\lambda^2 & A\lambda^2 \\ A\lambda^3[1-\rho e^{-i\delta}] & -A\lambda^2 & 1 \end{pmatrix} + O(\lambda^4) \,, \tag{1.6}$$

where λ, A, and ρ are three real parameters, all smaller than 1, and δ is a complex phase. Equation 1.6 establishes the hierarchy of weak interactions across different quark generations. In particular, couplings between distant generations are suppressed by powers of the λ parameter. Furthermore, the complex phase δ introduces the violation of the charge-parity (CP) symmetry in the SM, which is observed for example in Kaon decays [26].

Regularisation and renormalisation In general, QFT predictions for physical processes are performed with a perturbative approach, employing series expansions of the *scattering matrix* (S), which represents the time evolution operator between the initial $|i\rangle$ and final $\langle f|$ state:

$$\langle f|S|i\rangle = \langle f|\mathcal{T}\left[e^{-i\int dtV_I}\right]|i\rangle = \langle f|\left(\sum_{n=0}^{\infty}\frac{(-i)^n}{n!}\int dt_0\cdots\int dt_n\mathcal{T}[V_I(t_0)\cdots V_I(t_n)]\right)|i\rangle \,. \tag{1.7}$$

The symbol $\mathcal{T}$ denotes the time-ordered product and V_I the interaction potential defined in the Lagrangian of the system. If the coupling strength of the interaction is small ($\ll 1$), the series can be truncated at the first few terms. However, depending on the structure of V_I, higher orders of the perturbative expansion may lead to divergent contributions due to poles in the integrands of Equation 1.7, resulting in inconsistent predictions. These terms appear as loops in Feynman diagrams, where virtual particles can be created with momenta that may take unconstrained values, extending up to infinity.

To give a physical meaning to such diverging integrals, the theory must be *regularised* and *renormalised*. These two steps are key aspects of QFTs. The first can be viewed as protecting the theory with some additional parameter Λ that cuts off the infinities, by setting bounds to the integrals. The second step then absorbs the dependence on Λ into the “bare” parameters of the model, leaving finite higher-order corrections.

Interactions with dimensionless coupling constants can be renormalised, as is the case for the SM. Conversely, couplings with negative mass dimensions cannot be renormalised, as the Fermi weak theory [27] with the coupling constant $G_F \simeq 10^{-5}\,\text{GeV}^{-2}$. It achieved a good description of low-energy processes such as β decays, or the decay of the muon, but becomes unusable for energy scales approaching $G_F^{-1/2}$. The renormalisability of a theory is therefore a key aspect and constrains the terms allowed in the Lagrangian.

Considering the Fermi formulation of weak interactions as an effective theory provides a physical interpretation of the regularisation cutoff Λ. The theory is valid for low energies until "new physics" becomes relevant. As the new physics' scale $\Lambda_{\text{NP}} \sim 80\,\text{GeV}$ is approached, corresponding to the mass of the mediator vector bosons, the calculations break down, and a more complete theory is required.

Radiative corrections The higher-order contributions in perturbative calculations of observable quantities are referred to as *radiative corrections*. They are a remarkable feature of QFTs and imply that SM predictions could, in principle, be made with arbitrary precision, if enough orders are considered. However, the complexity of the calculations increases rapidly with higher orders. A direct consequence of renormalisation and radiative corrections is that physically measured couplings depend on the energy scale of the process; therefore, the SM has *running couplings*. For example, an effective QED coupling constant can be defined as [28]:

$$\alpha_{em}(q^2) = \frac{\alpha_{em}(\Lambda^2)}{1 - \frac{1}{\pi}\alpha_{em}(\Lambda^2)\ln\dfrac{q^2}{\Lambda^2}}\,, \tag{1.8}$$

where q is the exchanged momentum and Λ is the energy scale at which α_{em} is determined experimentally. For $\Lambda = 1\,\text{MeV}$, $\alpha_{em}(q^2 = \Lambda^2) = 1/137$, while for increasing energies, the coupling slowly strengthens: $\alpha_{em}(m_Z) \sim 1/129$. This behaviour of QED is known as *vacuum polarisation*, and can be visualised as the fact that, at higher momenta, smaller distance scales are probed, and therefore the shielding effect of polarized virtual charges decreases. This affects also corrections to the electron gyromagnetic factor, and is one of the most precise tests of the SM [29].

Similarly to QED, considering loop corrections to QCD diagrams, the effective strong coupling constant is found to be:

$$\alpha_s(q^2) = \frac{\alpha_s(\Lambda^2)}{1 + \frac{7}{4\pi}\alpha_s(\Lambda^2)\ln\dfrac{q^2}{\Lambda^2}}\,. \tag{1.9}$$

The most relevant difference from the QED case is that the denominator has a positive sign. Consequently, the interaction strength is expected to decrease at high energy and eventually vanish, making QCD an *asymptotically free* theory. As shown by measurements and predictions in Figure 1.2, $\alpha_s \sim 0.12$ at the scale of the Z boson mass and asymptotically drops for increasing q^2. At low energies, α_s expands rapidly, going above 0.5 when $q^2 < 1\,\text{GeV}$, making the inter-quark couplings non-perturbative. Approaching the scale $\Lambda_{\text{QCD}} \sim 200\,\text{MeV}$, the coupling diverges. This behaviour is related to the *confinement* of

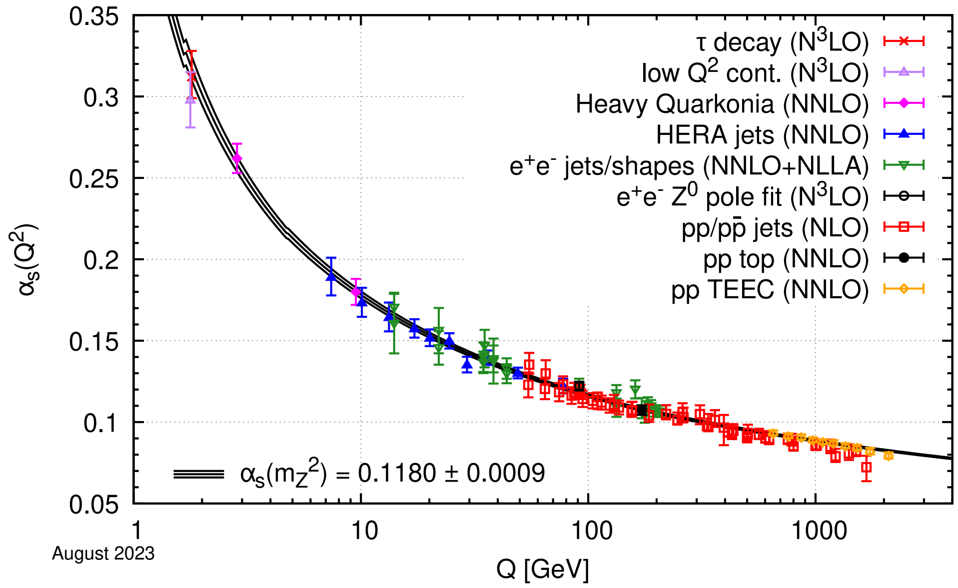

Figure 1.2: Summary of determinations of α_s as a function of the energy scale compared to the running of the coupling computed at five loops [30].

the quarks, the QCD property that causes quarks and gluons to stay within colour-neutral states (hadrons), and prevents them from being observed as free particles in nature.

Higgs mechanism To preserve gauge invariance in the SM Lagrangian under the groups from Equation 1.5, and more precisely, due to the structure of the weak interactions, fermions cannot have mass terms. Furthermore, it is observed that the $W^\pm$ and Z gauge bosons are not massless. Therefore, an additional prescription is needed to give mass to the $W^\pm$, the Z, and fermions. The Higgs Mechanism achieves this by introducing a new scalar field ϕ, with the Lagrangian:

$$\mathcal{L}_{\text{scalar}} = (\partial^\mu\phi\partial_\mu\phi) - \tfrac{1}{2}\mu^2\phi^2 - \tfrac{1}{4}\lambda\phi^4 \,, \tag{1.10}$$

where μ is the particle mass and λ a dimensionless self-coupling constant. If $\mu^2 < 0$, the potential of this scalar ($V = \frac{1}{2}\mu^2\phi^2 + \frac{1}{4}\lambda\phi^4$) can have a non-unique minimum, $i.e.$, a degenerate vacuum state: $v = \pm\sqrt{-\mu^2/\lambda}$. While the Lagrangian in Equation 1.10 is symmetric with respect to the transformation $\phi \to -\phi$, the perturbative expansion around one of the two vacua is not. Therefore, by arbitrarily selecting a specific ground state, the symmetry is *spontaneously broken*. This behaviour is often compared to a cooling ferromagnet. When above the Curie temperature, its magnetic domains are randomly oriented, and thus the system is symmetric under rotations. While at lower temperatures, a preferential orientation will emerge due to stochastic fluctuations, spontaneously breaking the rotational invariance.

With the Higgs Mechanism, adding a complex scalar doublet of $\mathrm{SU}(2)_\mathrm{L}$ to the SM – the Higgs field – and introducing the appropriate covariant derivative in Equation 1.10, the

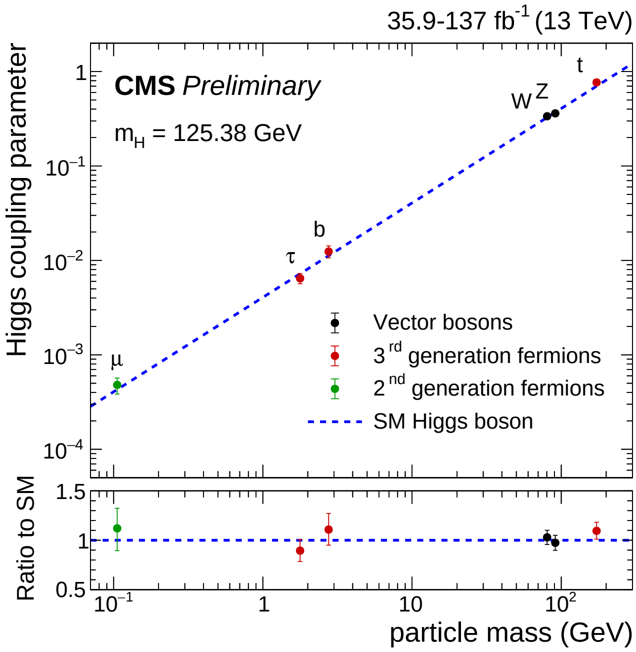

Figure 1.3: Recent measurements of the Higgs couplings as a function of the particle's mass [33], compared to the predictions from the SM.

mass terms for both the $W^{\pm}$ and Z bosons are naturally given. The fermion masses need to be manually added to the Lagrangian, with Yukawa-like couplings to the Higgs field that are proportional to the particles mass $\sim m_f \phi_H \bar{\psi}_f \psi_f$. The mechanism also predicts the existence of a new spin-0 particle, the Higgs boson. It was observed with a mass of $m_H \sim 125$ GeV in July 2012 [31, 32], and has been constantly studied since then. Many of its properties have been measured, $e.g.$, its couplings to the other particles shown in Figure 1.3, to confirm that it is indeed the boson predicted by the Higgs Mechanism. The prediction and discovery of the Higgs boson are arguably amongst the most important milestones of the SM.

1.2 Challenges and unanswered questions

As already stated, the SM is an effective theory, $i.e.$, an approximation of a more fundamental theory, and despite it's extraordinary precise predictions, the SM has clear experimental and theoretical limitations.

Dark matter and dark energy Measurements of galaxies' rotational velocity ($e.g.$, in spiral *type c* galaxies [34] or in the *Milky Way* [35]) constitute direct experimental evidence of the existence of Dark Matter (DM), according to the current understanding of gravitational laws. From these observations, large mass concentrations are present in most galaxies that do not interact via the electromagnetic force – therefore *dark* – but are only probed gravitationally. The analysis of the Cosmic Microwave Background power spectrum [36] indicates that the baryonic matter constitutes only $\sim 5\%$ of the Universe's energy content, while DM density is more than five times larger (26.5%). The remaining 68% consists of Dark Energy, which is believed to drive the accelerated expansion of the Universe. The SM framework cannot accommodate the astrophysical observations of DM and dark energy.

Neutrino masses Other experimental evidence for the limits of the SM arises from the measurements of *neutrino oscillations* [37–39]. This behaviour clearly indicates that neutrinos have a mass, despite being much lighter than other SM particles. From cosmological observations, it is expected that $\sum_\nu m_\nu < 0.082$ eV [40]. However, due to their left-handed nature and the chiral structure of the weak interactions, neutrinos cannot acquire mass through a Yukawa coupling to the Higgs, as it would violate gauge invariance. Therefore, a different mechanism is required to predict neutrino masses.

Naturalness problem When computing radiative corrections to the Higgs mass, at single-loop order, the dominant contribution is from virtual top pair production since the top quark has the largest coupling to the Higgs. In SM, the Higgs mass becomes the sum of the bare mass and the loop correction term, $m_h^2 = m_{h_0}^2 + \delta m_h^2$, with the contribution from the top loop given by [41]:

$$\delta m_h^2 \simeq -\frac{3G_F}{2\sqrt{2}\pi^2} m_t^2 \Lambda^2 \simeq -(0.2\Lambda)^2 \,, \tag{1.11}$$

where m_t is the mass of the top quark. This correction is quadratically divergent. To constrain it to the order of the observed Higgs boson mass, $\delta m_g \sim 10^2$ GeV, Λ must be close to the ~ 1 TeV scale. If the current SM is valid up to the "grand unification" scale [42] – where all the three coupling constants are supposed to coincide – $\Lambda_{\text{GUT}} \sim 10^{15}$ GeV, then, to achieve a light Higgs, the renormalisation requires an *extreme* fine-tuning of the model parameters. A cancellation between terms of the order of $\sim 10^{30}$ GeV with a precision of $\sim 10^{26}$ digits has to take place. This is known as the *naturalness problem* of the SM and provides a conceptual argument for the existence of Beyond Standard Model (BSM) physics.

Matter-antimatter asymmetry The observed Universe is predominantly composed of matter, suggesting that since the Big Bang, a mechanism must have acted to create the present asymmetry between matter and antimatter. In fact, it is usually assumed that the Big Bang produced particles and antiparticles in equal amounts. Additionally, the cosmic inflation hypothesis [43] predicts the dilution of any pre-existing baryon asymmetry. The dynamic formation of such matter-antimatter asymmetry through a *baryogenesis* process requires the Sakharov conditions [44] to be satisfied. They partially stand true in the SM, but not in the measure to explain the content of our Universe. More sources of CP violation and baryon number violation are required [45].

The preceding paragraphs outline some of the most pressing shortcomings of the SM. Many of these issues can be addressed through extensions of the theory, which predict the existence of new particles. Two prominent frameworks for BSM physics are discussed in the following section.

1.3 Beyond the Standard Model

The Beyond Standard Model (BSM) physics landscape is exceptionally vast. Numerous explanations have been proposed to overcome the limitations of the SM. Powerful ideas, such as the *supersymmetry* (SUSY), generated significant excitement within High Energy Physics (HEP) community, promising to address some of the most outstanding unresolved issues in the field. However, no experimental evidence has yet been collected to strongly favour any specific BSM model.

This section provides a brief review of SUSY, which offers a rich and well-motivated theoretical framework for BSM models. Extending the SM with SUSY can solve the Higgs mass naturalness problem, introduce DM candidate particles, and potentially incorporate a gravitational sector in the Lagrangian. These factors are arguably excellent reasons that made SUSY searches popular since the Large Electron-Positron collider (LEP) era.

The final part of the section shortly introduces the *Dark Sector*, a class of theories focusing on modelling potential DM candidates. As SUSY, also this framework provides fertile ground for BSM physics searches, encompassing various strongly motivated models. The full landscape of the Dark Sector is vast, this section discusses only its basic aspects.

Supersymmetry

One possible solution to the naturalness problem is the introduction of a new fundamental symmetry between bosons and fermions: Supersymmetry (SUSY). This symmetry postulates that the Lagrangian of the SM is invariant under transformations of bosonic fields ϕ into fermionic ones ψ, and vice versa:

$$\mathcal{Q}|\phi\rangle = |\psi\rangle \, ; \qquad \mathcal{Q}|\psi\rangle = |\phi\rangle \, . \tag{1.12}$$

Therefore, each particle must have a corresponding SUSY particle with identical quantum numbers but a spin differing by half a unit. In this way, the pairs of *super*-partners retain the same couplings, contributing equally to the Higgs mass loop corrections but with opposite signs. If the masses of particles and their corresponding “sparticles” are identical, the radiative corrections cancel out, eliminating any fine-tuning issue.

However, no SUSY particle has been observed at the scale of SM particles, implying that SUSY must be a broken symmetry. Although the SUSY scale has been pushed upwards by experimental limits [46], calculations of loop corrections in different SUSY models can still yield a precise Higgs mass compatible with the measurements [47].

The generators of SUSY are directly linked to the space-time transformations, establishing a compelling connection with general relativity. The supersymmetric generator $\mathcal{Q}$ modifies the total angular momentum by $1/2$, and following the anticommutation relation with its adjoint operator $\mathcal{Q}^\dagger$:

$$\{\mathcal{Q}, \mathcal{Q}^\dagger\} = p^\mu, \tag{1.13}$$

it is connected to the four-momentum operator p^μ, which is itself the generator of translations. The expression in Equation 1.13 shows that two consecutive SUSY transformations result in a space-time translation, thereby suggesting a deep connection with general relativity [48].

The spectrum of sparticles affects all the radiative corrections. The running of electromagnetic and strong couplings from Equations 1.8 and 1.9, along with weak coupling, is altered upon reaching the SUSY scale and beyond. If the SUSY contributions up the two-loop effects are included in the α's evolutions, the three SM couplings can unify at a scale close to $\sim 10^{16}$ GeV [49], making SUSY an interesting candidate for a grand unification.

The solution of the naturalness problem, the prospect of unifying all fundamental forces within a single theory, and the inclusion of DM candidates in many SUSY models [50] have made this framework highly attractive.

The Minimal Supersymmetric Standard Model The SM can be extended in various ways, leading to different combinations of super-partners. The Minimal Supersymmetric Standard Model (MSSM) [51] implements SUSY with a minimum set of assumptions. The particle fields are arranged in *super-multiplets*, each containing fermion-boson pairs. SM fermions are represented by chiral super-multiplets, together with their *sfermion* counterparts. Gauge bosons and their spin-$1/2$ super-partners, the *gauginos*, are described by gauge super-multiplets.

The superfield content of the MSSM is shown in Table 1.1. To preserve the electroweak spontaneous symmetry breaking, two Higgs chiral doublets are required, H_u and H_d. As in the standard Higgs Mechanism, three degrees of freedom are absorbed to provide mass to the SM gauge bosons ($W^\pm$ and Z), while one remains associated with the Higgs boson. The remaining four degrees of freedom correspond to new physical particles: two neutral and two charged Higgs bosons. These interaction eigenstates mix with other gauginos to form massive particles. The four neutral states $\widetilde{B}$, $\widetilde{W}^0$, $\widetilde{H}_d^0$, and $\widetilde{H}_u^0$ are combined into four *neutralinos*: $\widetilde{\chi}_i^0$ where $i = 1, 2, 3, 4$. They are conventionally sorted by increasing mass. The *charginos* ($\widetilde{\chi}_1^\pm$ and $\widetilde{\chi}_2^\pm$), by contrast, are the physical combinations of charged higgsinos and gauginos: $\widetilde{W}^+$, $\widetilde{H}_u^+$, $\widetilde{W}^-$, and $\widetilde{H}_d^-$. Apart from the additional Higgs doublet, an unbroken formulation of the MSSM contains no additional free parameters.

A widely adopted assumption in the MSSM is the existence of an additional $U(1)$ global

Table 1.1: Summary of chiral and gauge super-multiplets in the MSSM [49].

Names		spin-0	spin-1/2	spin-1	$SU(3)_C$, $SU(2)_L$, $U(1)_Y$
squarks, quarks	Q	$(\widetilde{u}_L\ \widetilde{d}_L)$	$(u_L\ d_L)$		$\left(3,2,\frac{1}{6}\right)$
(×3 families)	$\overline{U}$	$\widetilde{u}_R^*$	$u_R^\dagger$		$\left(\bar{3},1,-\frac{2}{3}\right)$
	$\overline{D}$	$\widetilde{d}_R^*$	$d_R^\dagger$		$\left(\bar{3},1,\frac{1}{3}\right)$
sleptons, leptons	L	$(\widetilde{\nu}\ \widetilde{e}_L)$	$(\nu\ e_L)$		$\left(1,2,-\frac{1}{2}\right)$
(×3 families)	$\overline{E}$	$\widetilde{e}_R^*$	$e_R^\dagger$		$(1,1,1)$
gluino, gluon			$\widetilde{g}$	g	$(8,2,0)$
winos, W bosons			$\widetilde{W}^\pm\ \widetilde{W}^0$	$W^\pm\ W^0$	$(1,3,0)$
bino, B boson			$\widetilde{B}$	B	$(1,1,0)$
higgsinos, Higgs	H_u	$(H_u^+\ H_u^0)$	$(\widetilde{H}_u^+\ \widetilde{H}_u^0)$		$\left(1,2,\frac{1}{2}\right)$
	H_d	$(H_d^0\ H_d^-)$	$(\widetilde{H}_d^0\ \widetilde{H}_d^-)$		$\left(1,2,-\frac{1}{2}\right)$

phase invariance, referred to as *R*-parity or *R*-symmetry. It assigns a new multiplicative quantum number to each particle:

$$R = (-1)^{3B+L+2s}\ , \tag{1.14}$$

where B and L are the baryonic and leptonic numbers, while s is the particle spin. According to this definition, all SM particles and sparticles have *R*-parity equal to $+1$ and -1, respectively. Imposing this symmetry on the Lagrangian constrains the allowed interaction terms, having significant consequences on the phenomenology of supersymmetric theories. In the so-called *R*-parity conserving (RPC) scenario, super-partners can be only pair-produced from SM particles interactions. And, most importantly, the *lightest supersymmetric particle* (LSP) is stable, implying that RPC models featuring the neutralino as the LSP naturally provide a candidate for DM.

SUSY breaking As previously mentioned, if SUSY exists in nature, it must be a broken symmetry, otherwise sparticles would have masses equal to those of SM particles. Currently, no known mechanism exists to break SUSY spontaneously in the MSSM. Therefore, the symmetry must be broken "by hand", by introducing terms into the Lagrangian that explicitly violate SUSY. However, this must be done carefully, to ensure that the benefits of radiative correction cancellations are preserved. In these context, it is said that the symmetry must be *softly* broken, which can be implemented in multiple ways. This procedure typically introduces a large number (> 100) of free parameters to the MSSM.

The standard convention is to assume the existence of a supersymmetric "hidden sector" at a significantly higher scale, effectively decoupled from the MSSM particles. The two sectors are connected uniquely through a mediator. In the hidden sector SUSY is spontaneously broken, with the messenger interaction propagating the soft breaking to the MSSM. Two prominent examples are *minimal super-gravity* (mSUGRA) [52], where the mediator interaction is gravity, and gauge-mediated SUSY breaking (GMSB) [53], though other possibilities exist.

***R*-parity violation** Building on the discussion from the previous two paragraphs, the Lagrangian of a generic SUSY theory can be factorised as:

$$\mathcal{L} = \mathcal{L}_{\text{MSSM}} + \mathcal{L}_{soft} + W_{\text{RPV}} \ . \tag{1.15}$$

The $\mathcal{L}_{\text{MSSM}}$ includes the superfields kinematic terms, along with the gauge interactions of SM and the RPC allowed interactions. The SUSY breaking part is contained in $\mathcal{L}_{soft}$. Finally, W_{RPV} describes the renormalisable super-potential for *R*-parity violating (RPV) interactions:

$$W_{\text{RPV}} = \frac{1}{2}\lambda_{ijk}L_iL_j\overline{E}_k + \lambda'_{ijk}L_iQ_j\overline{D}_k + \frac{1}{2}\lambda''_{ijk}\overline{U}_i\overline{D}_j\overline{D}_k + \mu'_iL_iH_u \ . \tag{1.16}$$

The super-multiplets in this equation are defined according to the conventions from Table 1.1. The first three terms are referred to as “trilinear” couplings, where the dimensionless coupling strengths are represented by the fully antisymmetric tensors λ_{ijk}, λ'_{ijk}, and λ''_{ijk}, with i, j and k denoting the family indexes. The last term is a “bilinear” coupling of intensity μ' between leptonic superfields and the up Higgs doublet.

In addition to violating *R*-parity, the potential in Equation 1.16 also breaks the conservation of leptonic and baryonic numbers. The λ'' interaction allows for $\Delta B = 1$ transitions, while the other interactions cause $\Delta L = 1$. There is no a priori issue with this feature, in fact, no fundamental principle enforces the conservation of B and L in the SM. They appear as accidental symmetries of the model. However, when allowing these transitions, the stability of the proton must be preserved in accordance with its experimental limits. For example, as shown in Figure 1.4, the decay $p \to e^+\pi^0$ would be allowed with the λ'' and λ' trilinear terms, but it is tightly constrained experimentally ($\tau > 1.6 \cdot 10^{34}$ years [54]). This implies that the parameter space of possible RPV coupling configurations has to account for results from low energy experiments. From measurements of proton and neutron decays, the product $|\lambda'_{i1k}\lambda''_{11k'}|$ must be smaller than $\sim 10^{-27}$ [55]. However, since such experiments primarily constrain the products of couplings, suppressing one type of interaction relaxes the constraints on the others.

A final remark on RPV theories is that the LSP is no longer stable and can decay to SM particles. At first glance, this seems problematic for the DM hypothesis, however, supersymmetric theories can still accommodate DM candidates, $e.g.$, with an unstable LSP whose lifetime exceeds the age of the Universe [56].

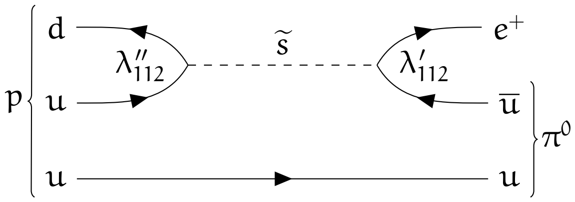

Figure 1.4: Feynman diagram of the proton decay $p \to e^+\pi^0$ through a squark mediator, allowed by the λ' and λ'' trilinear RPV interactions.

SUSY searches at colliders Since the early days of SUSY, searches for sparticles have been conducted at colliders, such as LEP [57] and Tevatron [58]. In collider experiments, RPC SUSY would leave a distinctive signature. In this scenario, interactions involving a stable LSP lead to the production of particles that do not decay within the detector volume, generating momentum imbalance in the final state of collisions. This signature is widely used in searches, as, apart from neutrinos, no SM particles can produce it. In the case of RPV scenarios, the final state is typically composed of SM particles, making it more challenging to constrain backgrounds.

No evidence of SUSY has been found to date, ruling out the lower mass scales for sparticles. To ease the analysis interpretations and provide benchmarks, a common procedure is to use simplified models where the SM Lagrangian is extended with only a few new particles and possible decays. However, this pragmatic approach represent only edge cases of broader supersymmetric theories, and reinterpretations of individual results are not always possible. The phenomenological MSSM (pMSSM) [59, 60] provides a framework for combining multiple direct searches and including external constraints from precision electroweak measurements and non-collider experiments. This methodology allows for scanning over subsets of parameters, consistently approximating extended SUSY models, and achieving a better understanding of which parts of the SUSY spectrum are effectively excluded, as shown in Figure 1.5.

The Dark Sector

A alternative approach to overcome the limitations of the SM is the definition of a Dark Sector [61], containing particles that are not charged under SM forces. Since finding DM candidates is one of its main goals, it is assumed that Dark Sector particles possess mass, and are therefore susceptible to gravity. The other key assumption is that they are *feebly interacting*, *i.e.*, weakly coupled to SM matter. A new force is responsible for this coupling, the so-called “portal” interaction. Portals are typically characterised by the spin and SM couplings of the mediator particle. A summary of the most generic renormalisable portals is reported in Table 1.2:

- One of the simplest Dark Sector models is obtained by extending the SM with right-handed *sterile neutrinos* [62]; these could be combined with the SM left-handed dou-

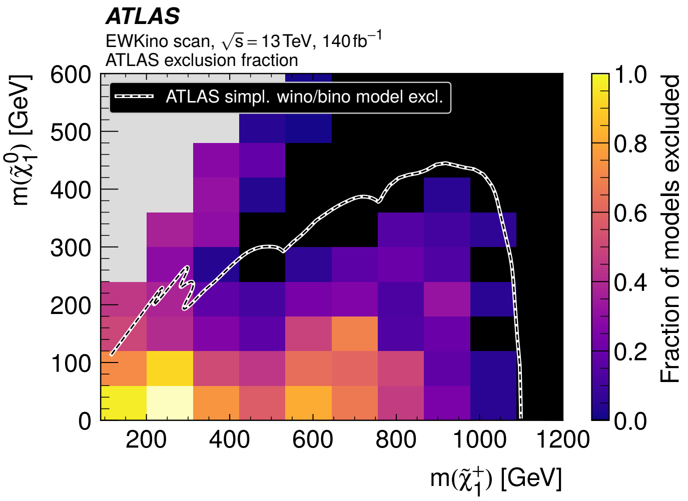

(a)

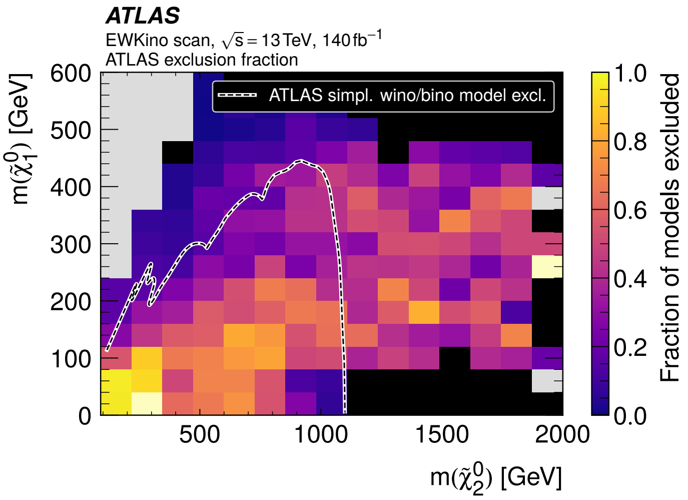

(b)

Figure 1.5: The fraction of excluded pMSSM models in the (a) $m(\tilde{\chi}_1^\pm)$–$m(\tilde{\chi}_1^0)$ and (b) $m(\tilde{\chi}_2^0)$ – $m(\tilde{\chi}_1^0)$ planes obtained using the results from eight independent searches [63]. The dashed line represents the envelope of the exclusions limits set by the analyses, considering simplified models.

Table 1.2: Main Dark Sector renormalisable portals.

Mediator	Coupling
sterile neutrino (N)	$y_N L \phi_H N$
dark photon (A'_μ)	$-\frac{\epsilon}{2\cos\theta_W} F'_{\mu\nu} B^{\mu\nu}$
dark Higgs (S)	$(\mu S + \lambda S^2)\phi_H^\dagger \phi_H$
axion-like particle (a)	$\frac{a}{f_a} F_{\mu\nu} \tilde{F}^{\mu\nu}$, $\frac{a}{f_a} G_{i,\mu\nu} \tilde{G}_i^{\mu\nu}$, $\frac{\partial_\mu a}{f_a} \overline{\psi} \gamma^\mu \gamma^5 \psi$

blet L allowing for a Yukawa coupling with the Higgs field, and thus solving the neutrino mass problem.

- A new vector mediator, the *dark photon*, could originate from a new $U(1)$ gauge symmetry group; its field strength tensor $F'_{\mu\nu}$ couples to the electroweak hypercharge field $B^{\mu\nu}$ with an intensity determined by the “kinetic mixing” parameter ϵ [64].
- The *dark Higgs*, a scalar singlet, may couple linearly and quadratically to the SM Higgs ϕ_H, affecting its possible decays [65].
- Pseudo-scalar *axion-like particles* [66] could couple to the electromagnetic or strong force mediator strength tensors, or even directly to fermion fields. They represent a generalisation of the QCD axions proposed to solve strong CP problem [67].

Given its rather generic definition, a wide spectrum of models falls within the Dark Sector. The mass range of BSM particles in these models spans from sub-eV to multi-TeV scales and beyond. Despite the limited phase space accessible, searches at high energy colliders can have unique sensitivity to some specific models of such theories.

Chapter 2

ATLAS and the energy frontier

Particle colliders are indispensable instruments for HEP research. Most known particles have been discovered by accelerating particle beams and colliding them either with fixed targets or with opposing beams. According to Einstein's equation ($E = mc^2$), the energy generated in these collisions is converted into mass, enabling the production of new unstable matter states. The primary objective of enhancing colliders power is, therefore, to access heavier particles that require higher energy thresholds for their creation

The current *energy frontier* of particle accelerators is represented by Large Hadron Collider (LHC). It hosts four large detectors that observe the collisions directly, as close as possible to the interaction points. A Toroidal LHC ApparatuS (ATLAS) and Compact Muon Solenoid (CMS) are two multi-purpose experiments, performing both precision measurements of the SM and searches new physics, forming the two largest collaborations with broad physics programmes. The other two experiments are A Large Ion Collider Experiment (ALICE) and Large Hadron Collider beatuy (LHCb). The first one is dedicated to heavy ion collisions, studying the quark-gluon plasma, while the second is designed to study CP violation and rare b-hadron decays.

The work of this thesis was carried out within the ATLAS collaboration. In this chapter the main aspects of LHC and the ATLAS detector are described, with a dedicated section on the ATLAS trigger system.

2.1 The Large Hadron Collider

The LHC is the largest and most powerful circular collider on Earth, with a circumference of 26.7 km. It collides proton beams with energies above 6 TeV and heavy ions at lower energies. The accelerator is located at the *Conseil Européen pour la Recherche Nucléaire* (CERN) laboratories, in a tunnel that was excavated previously for LEP. It is part of a large accelerator complex, as shown in Figure 2.1. The journey of protons starts in a hydrogen bottle, where they are separated from electrons by a strong electric field. They are accelerated in succession by the Linac, the Proton Synchrotron Booster, the Proton Synchrotron and the Super Proton Synchrotron, reaching an energy of 450 GeV. After this stage, they are injected into the LHC, which accelerates the proton beams to more than 99.99% of the speed of light. The highest energy achieved per beam is currently 6.8 TeV, resulting in a centre-of-mass energy of $\sqrt{s} = 13.6$ TeV, but the accelerator is designed to stretch the beams up to 7 TeV.

The LHC started operations in 2008, but the experiments' first data-taking campaign, the LHC Run 1, began in 2011 and lasted two years, until the end of 2012. This period was followed by the first long shutdown (LS1), during which maintenance and upgrades of both the LHC and its experiments took place. In 2015, Run 2 was launched and continued until

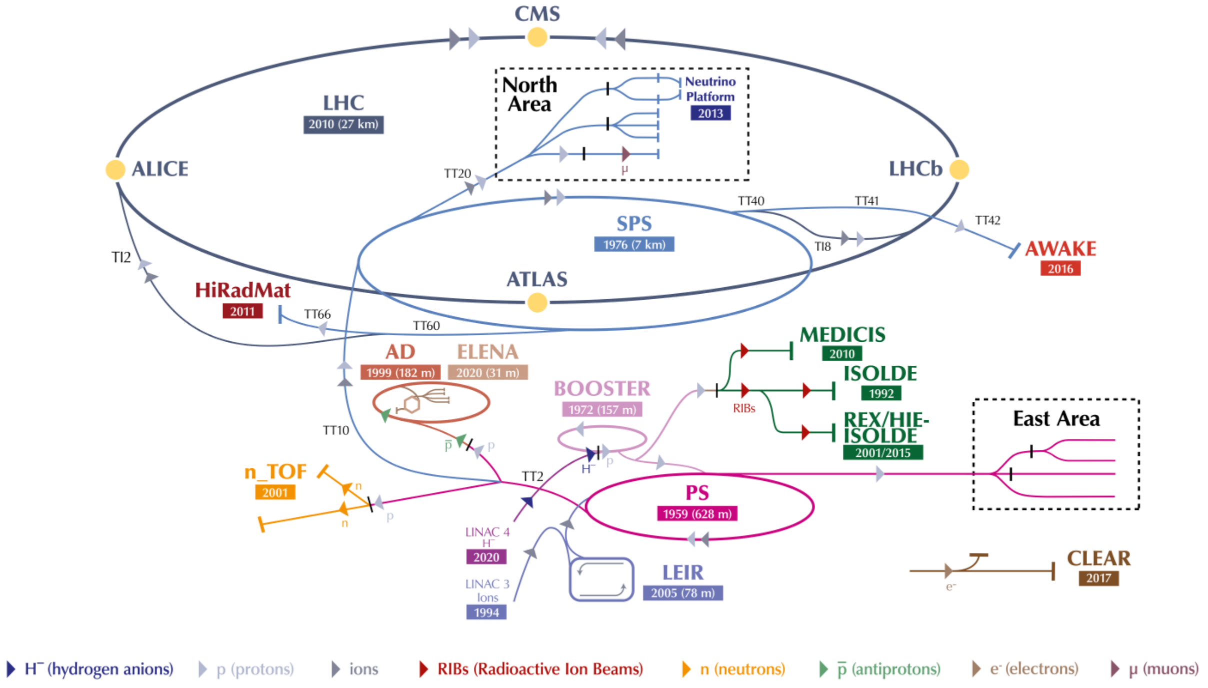

LHC - Large Hadron Collider // SPS - Super Proton Synchrotron // PS - Proton Synchrotron // AD - Antiproton Decelerator // CLEAR - CERN Linear Electron Accelerator for Research // AWAKE - Advanced WAKefield Experiment // ISOLDE - Isotope Separator OnLine // REX/HIE-ISOLDE - Radioactive EXperiment/High Intensity and Energy ISOLDE // MEDICIS // LEIR - Low Energy Ion Ring // LINAC - LINear ACcelerator // n_TOF - Neutrons Time Of Flight // HiRadMat - High-Radiation to Materials // Neutrino Platform

Figure 2.1: Layout of the CERN accelerator complex, in January 2022 [68].

2018, increasing the collisions' energy from $\sqrt{s} = 8\,\text{TeV}$ to $13\,\text{TeV}$. After the second long shut down (LS2), Run 3 began in 2022 with $0.6\,\text{TeV}$ higher center-of-mass energy, and is ongoing at the time of writing this thesis. At the end of Run 3, the accelerator and the LHC experimental projects will undergo major upgrades, in preparation for the High-Luminosity LHC (HL-LHC) [69]. In this new phase, the beam intensity is expected to increase by approximately a factor 5, introducing a few significant technical challenges. In this thesis, the data collected during Run 2 and part of Run 3 is analysed.

Principles of physics at colliders

At circular colliders, particles are accelerated by strong electric fields, to increase their energy, and are maintained in a fixed orbit through a centripetal force generated by magnetic fields. This constant force has the side effect of causing the beam energy dissipation. In fact, all charged particles emit photons when accelerated. When the source of the force is a magnetic field perpendicular to the particle's motion, as in circular colliders, the emitted energy is referred to as synchrotron radiation. The radiated power follows the Larmor equation. For relativistic particles, the emitted synchrotron power is:

$$W = \frac{2}{3}\frac{e^2}{c}\gamma^4\dot{\beta}^2 = \frac{2}{3}\frac{e^2}{c}\left(\frac{E}{mc^2}\right)^4\dot{\beta}^2\,, \tag{2.1}$$

where $\gamma = 1/\sqrt{1-\beta^2}$ is the Lorentz factor. Equation 2.1 shows that the energy loss due to beam bending is proportional to the inverse fourth power of the mass (m^{-4}). Consequently, electrons experience significantly greater energy losses compared to protons. This difference is the primary reason for the large step in energy from LEP to LHC, from $\sqrt{s} = 209\,\text{GeV}$ to above $10\,\text{TeV}$, despite both accelerators having the same diameter.

A key parameter of colliders is the luminosity [70]. It represents the proportionality constant between the cross-section σ_{process} and the event rate of a process:

$$\frac{dN_{\text{event}}}{dt} = \mathcal{L}\cdot\sigma_{\text{process}}\,. \tag{2.2}$$

This relation defines the *instantaneous* luminosity, which is typically expressed in $\text{cm}^{-2}\text{s}^{-1}$. Since protons are arranged into equidistant bunches, for the LHC beams it can be expressed as:

$$\mathcal{L} = \frac{N_1 N_2 n_b}{4\pi\sigma_x\sigma_y} f_{\text{rev}}\,, \tag{2.3}$$

where N_1 and N_2 are the numbers of particles per colliding bunch, n_b is the total number of bunches per beam and f_{rev} their revolution frequency. Finally, σ_x and σ_y correspond to the transverse dimensions of the beam. They are proportional to ϵ_b, the emittance of the beam, and β_b^*, its transverse displacement at the interaction point: $\sigma_x\sigma_y = \epsilon_b\beta_b^*$.

During each bunch-crossing, multiple proton interactions can take place, with the interactions number increasing at higher instantaneous luminosities. However, due to practical constraints not all the collisions are recorded by the experiments. Typically, only a single hard scatter from each bunch-crossing is considered interesting from a physics perspective. The other concurrent interactions are known as *pile-up*. If these interactions occur in the same bunch-crossing as the interesting collision, they are referred to as *in-time* pile-up. The expected number of such collisions is expressed as the average number of interactions per bunch-crossing:

$$\langle \mu \rangle = \frac{\mathcal{L} \cdot \sigma(pp \to X)}{n_b \cdot f_{\text{rev}}} \,, \tag{2.4}$$

where $\sigma(pp \to X)$ is the total cross-section for inelastic proton-proton (p-p) scattering. Considering that the highest luminosities achieved in Run 3 are close to $\mathcal{L} \simeq 2.1 \cdot 10^{34}/\text{cm}^2 \cdot \text{s} = 2.1 \cdot 10^{-2}/\text{pb} \cdot \text{s}$, and that $f_{rev} \simeq 1.1 \cdot 10^4$ Hz and $\sigma(pp \to X)$ is nearly $0.8 \cdot 10^{11}$ pb at $\sqrt{s} \sim 13$ TeV [71], for roughly 2400 circulating bunches, the expected average interactions are $\langle \mu \rangle \simeq \frac{2.1 \cdot 10^{-2} \cdot 0.8 \cdot 10^{11}}{2.4 \cdot 10^3 \cdot 1.1 \cdot 10^4} \simeq 63$. This number indeed corresponds closely to value of $\langle \mu \rangle$ with more recorded data in Run 3, as shown in Figure 2.2.

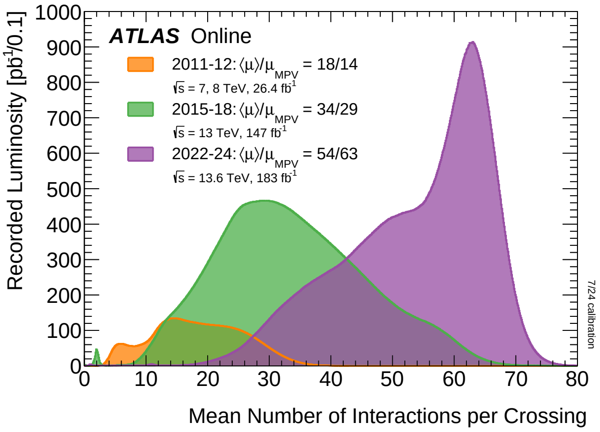

Figure 2.2: Recorded luminosity as a function of the average number of interactions, for each year of LHC data-taking campaign [72].

The Equation 2.4 shows the linear correlation between luminosity and in-time pile-up. Instead, the *out-of-time* pile-up is induced by p-p interaction from neighbouring bunches. To reconstruct collision events for measurements, it is important to mitigate biases introduced by all sources of pile-up, as discussed in Section 3.3.

At LHC, the colliding projectiles are composite particles. According to the De Broglie relation[4], at energies above 200 MeV, distances shorter than the proton radius (~ 1 fm)

---

[4] De Broglie hypothesised that each massive particle has an associated wavelength $\lambda = h/p$, where h is the Planck constant and p the particle's momentum. With this hypothesis, quantum physical properties of particles are also determined by their wavelength.

are probed. Consequently, for p-p collisions, the hard-scattering process occurs between the proton constituents, described with the *parton* model [73]. In this picture, a proton is composed of: *valence* quarks, which carry most of the proton's momentum and determine its quantum numbers; gluons, which bind the quarks together; and *sea* quarks, virtual quark-antiquark pairs produced by gluons' propagation. The statistical distribution of the proton energy amongst the partons is described by the Parton Density Functions (PDFs). Since each parton carries only a fraction x_i of the total proton momentum, the underlying parton scattering develops at centre-of-mass energy $\sqrt{\hat{s}} = \sqrt{x_1 x_2 s}$, which is smaller than the colliding beams energy. Furthermore, the total scattering momentum along the beam axis is neither balanced nor deterministic.

The accelerator features

The LHC is a superconducting ring hadron accelerator. It is housed in an underground circular tunnel whose plane is slightly tilted, lying between 45 m and 170 m below the surface. The beams travel in opposite directions within two distinct pipes maintained at ultra-high vacuum. Eight radio-frequency cavities per pipe accelerate beam, each operating at a maximal voltage of 2 MV, increasing the protons' energy by 16 MeV per revolution.

To bend the beams and keep them in a circular orbit, more than a thousand dipole magnets are installed in the LHC tunnel. The structure of the dipoles' system is shown in Figure 2.3. The magnetic coils, made of niobium-titanium (Nb-Ti) fibre, require a cryogenic cooling system to operate in superconducting regime. Thereby, they are immersed in a

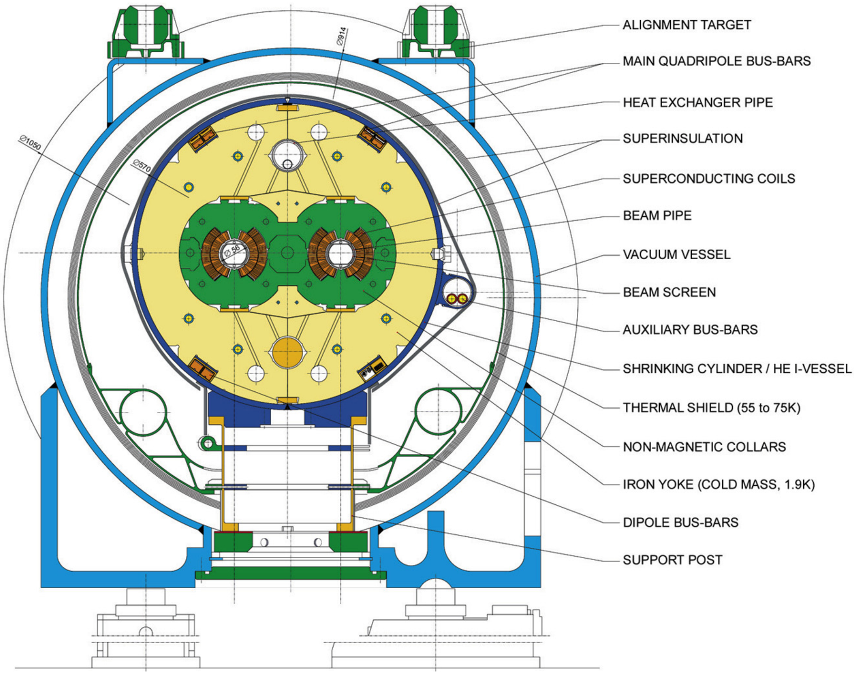

Figure 2.3: Cross section of the LHC ring in presence of a dipole magnet; lengths are reported in mm [74].

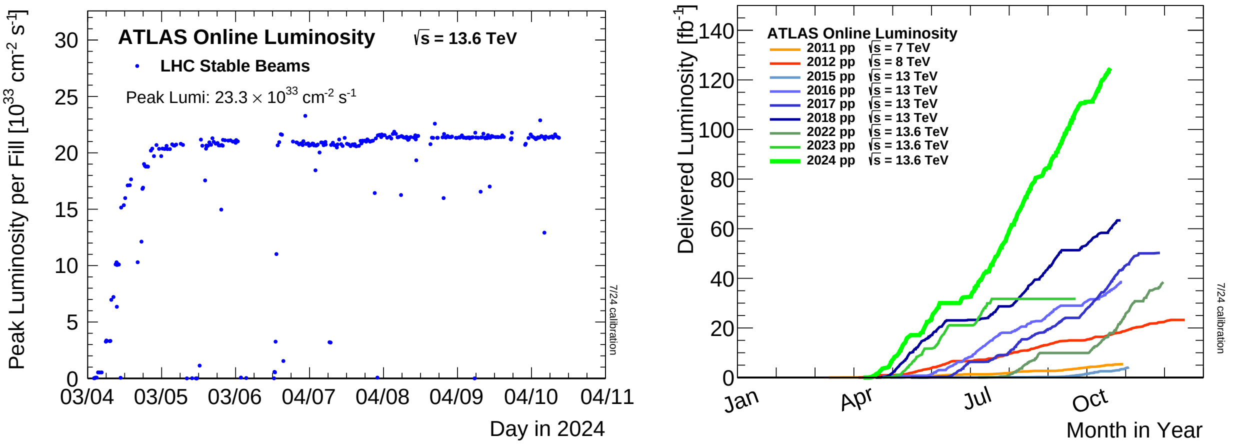

(a) **(b)**

Figure 2.4: Figure (a): peak instantaneous luminosity per filled delivered to the ATLAS Experiment during the 2024 data taking. Figure (b): integrated luminosity versus day delivered to ATLAS during stable beams and for high energy p-p collisions [72].

thermal bath of liquid helium at $1.9°$ K. These coils can generate a central magnetic field exceeding 8 T. Nearly four hundred quadrupole magnets are employed to focus the beams. Special triplets of quadrupoles apply a final focus just before the collision points, squeezing the beams to increase the probability of collisions.

At the LHC, each proton bunch contains more than 10^{11} particles. The time spacing between two bunches is 25 ns, corresponding to a maximum collision frequency of 40 MHz. Considering the almost 27 km of circumference and the fact that protons travel at the speed of light, the total number of bunches that can fit into a beam is approximately 3600. However, to keep optimal conditions and prevent excessive heat loads, LHC employs filling schemes with *bunch trains* interspaced by empty buckets. The best performance is currently achieved with ~ 2400 proton bunches per beam. With these filling strategies, and minimising β^*_b as much as possible, the luminosity at the ATLAS interaction point surpassed $2.1 \cdot 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, as shown in Figure 2.4(a).

The total delivered data is typically quantified in terms of *integrated* luminosity, corresponding to $\mathcal{L}$ integrated over time: $\mathcal{L}_{\text{int}} = \int \mathrm{d}t\,\mathcal{L}$. Figure 2.4(b) shows the delivered $\mathcal{L}_{\text{int}}$ to ATLAS for each year, since the beginning of the LHC operations. The collision production in year 2024 was remarkably abundant, with more than $120\,\mathrm{fb}^{-1}$. Considering that the gluon-gluon fusion Higgs boson production has a cross-section ~ 48.5 pb [75], close to 6 million Higgs were created in ATLAS in 2024 alone.

2.2 The ATLAS experiment

ATLAS is one of the two multi-purpose experiments installed on the LHC ring. It is designed to precisely measure the SM properties and search for BSM physics, such as SUSY or other new exotic particles at the TeV scale.

During the initial planning phase, a central objective for the experiment was the search for the SM Higgs, which was successfully discovered in 2012 [31]. Since then, its properties have been scrutinised extensively, such as its decays to bottom quarks [76] or tau leptons [77], and its coupling to the top quarks [78]. Beyond the Higgs sector, many other SM predictions have been measured, ranging from the production of same sign W bosons [79] to light by light scattering [80], and the entanglement of $t\bar{t}$ pairs [81]. In the realm of BSM physics, large portions of the SUSY parameter space have been explored [46], and a broad spectrum of exotic signatures has been investigated [82]. These results demonstrate the experiment's broad scope and versatility.

This section begins by discussing the adopted convention for the ATLAS coordinate system, followed by a description of the detector and its trigger system.

The coordinate system

Following standard practice at colliders, the ATLAS detector employs a cylindrical coordinate system to mathematically describe collision physics. The cylinder's axis corresponds to the beam line (z-axis), as shown in Figure 2.5. The origin coincides with the interaction point (IP). The azimuthal angle ϕ is measured around z, starting from the x-axis, which points horizontally towards the centre of the LHC. Finally, the polar angle θ defines the distance from the beam direction.

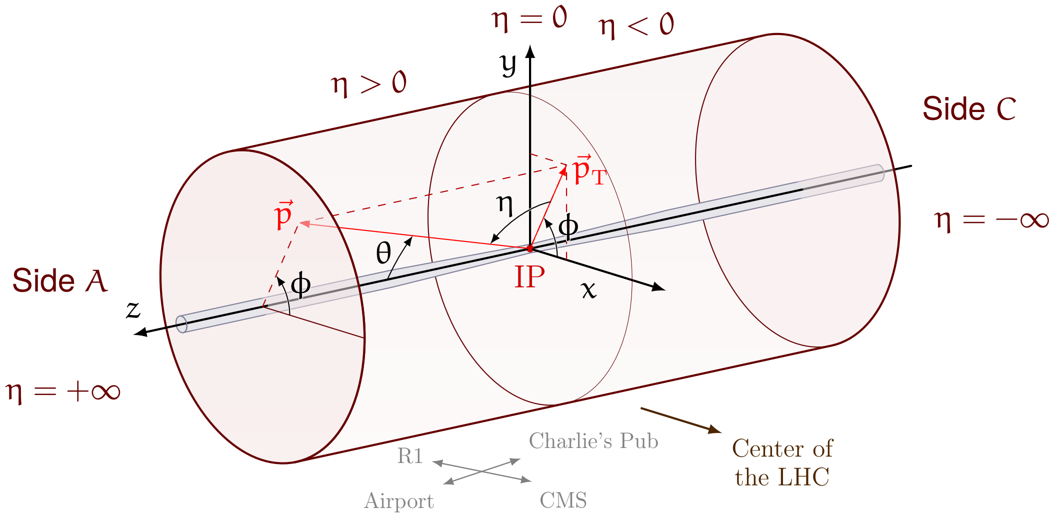

Figure 2.5: The ATLAS cylindrical coordinate system[4].

As mentioned in Section 2.1, in p-p collisions the initial momentum along the z axis is unknown, due to the probabilistic nature of PDFs. Consequently, variables invariant under Lorentz boosts in the longitudinal direction become particularly useful, such as the

---

[4] Original LaTeX template from: https://tikz.net/axis3d_cms.

transverse momentum (p_T) or the *transverse energy* ($E_T = E\sin(\theta)$). Another example of invariant variable is the difference in rapidity Δy, where the rapidity of a particle is defined as:

$$y = \frac{1}{2}\ln\left(\frac{E + p_z}{E - p_z}\right) . \tag{2.5}$$

However, since reconstructing rapidity requires precise measurement of the total energy, it is often used instead an approximate version of it, the pseudo-rapidity:

$$\eta = -\ln\left[\tan\left(\frac{\theta}{2}\right)\right] = \frac{1}{2}\ln\left(\frac{|\vec{p}| + p_z}{|\vec{p}| - p_z}\right) \simeq y . \tag{2.6}$$

In the limit of high energies or massless particles, the two quantities are equal. The main advantage of the pseudo-rapidity is its direct relation to the polar angle, making it easier to reconstruct experimentally. On the transverse plane ($\theta = 90^\circ$) η is zero. Conversely, when approaching the beam line ($\theta = 0^\circ$ or 180°), η tends to $\pm\infty$.

The ATLAS detector

To support a vast physics programme and maximise the output of the experiment, several guiding principles were considered in the design of the detector system:

- Fast and radiation-resistant sensors and electronics to cope with the harsh experimental environment.
- Large angular acceptance, extending to high pseudo-rapidities, and with close to full coverage of the azimuth angle.
- High granularity to mitigate overlapping.
- Efficient reconstruction of charged particles trajectories, with fine resolution.
- Nearly hermetic calorimetry in the transverse plane for precise measurements of momentum imbalance.
- Strong muon identification and precise measurements at high transverse momenta.
- Fast and efficient trigger with good levels of background rejections to achieve an acceptable acquisition rate.

These guidelines led to the design of the ATLAS detector [83], a multilayer barrel structure with forward-backward symmetry relative to the IP. Measuring 25 m in height and 44 m in length, the detector weighs approximately 7000 tonnes.

Figure 2.6 shows an overview of the detector, updated to the Run 3 configuration. Starting from the innermost layers, its main components are: the Inner Detector (ID) system,

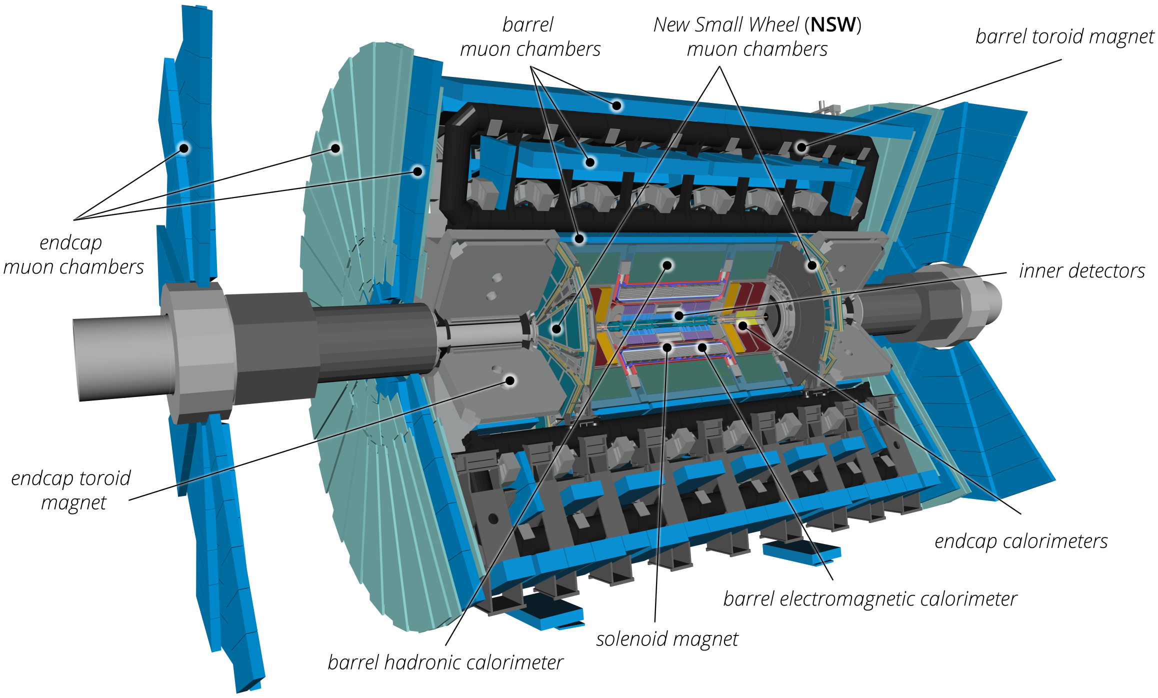

Figure 2.6: Cut-away view of the ATLAS detector system for Run 3 [84].

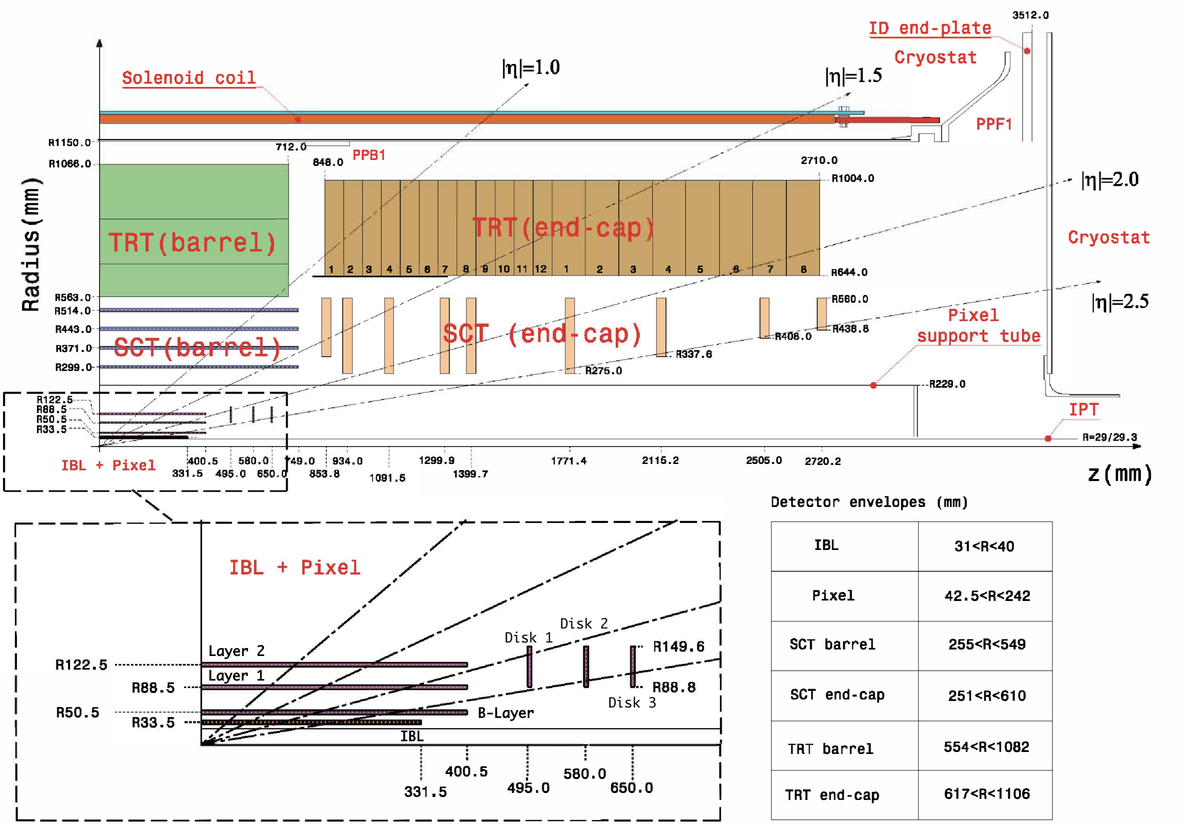

Detector envelopes (mm)

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 2.7: Cross-section of the ID system in z-r plane [84].

which covers the $|\eta| \leq 2.5$ region closest to beam line with precision trackers; the calorimeter system follows, extending up $|\eta| \leq 4.9$, composed of the high-granularity lead/Liquid Argon (LAr) calorimeter for the electromagnetic showers and the steel/scintillator-tile calorimeter for hadrons; finally, the most external part is the Muon Spectrometer (MS), characterised by several muon detection technologies. A strong magnetic field permeates the entire detector volume to distinguish charged particles and measure their momenta. A solenoid magnet encapsulates the ID, generating a 2 T axial field. Surrounding the calorimeters, superconducting barrel and end-cap toroids generate a toroidal magnetic field ranging from 2 T to 6 T, deflecting the escaping muons.

Inner Detector The ID is installed as close as possible to the beam line. It serves as the main tracking system of ATLAS, and is designed to precisely reconstruct the charged particles trajectories over a large portion of the solid angle. Its layout is illustrated in Figure 2.7, and its main parameters are summarised in Table 2.1.

The part nearest to the beam is instrumented with the Pixel Detector [85], featuring fine granularity. It consists of rectangular silicon pixels arranged into four barrel layers and three end-cap discs per side. The innermost barrel layer, know as Insertable B-Layer (IBL) [86], was added during LS1 before Run-2 to improve the vertexing capabilities of the system, and in particular to improve the resolution on measurements of the distance of closest approach to the primary vertex, a key feature to identify b-hadrons. The IBL has the finest granularity (50×250 µm pixels), and is located at a radius of 33.5 mm from the beam. The remaining three layers are almost equidistant from each other and span from a radius of 50.5 mm up to 122.5 mm. The end-cap layers ensures a good pseudorapidity coverage, with a sufficient number of hits within $|\eta| \leq 2.5$. They are a cost-efficient alternative to extending the barrel layers' lenght. Given its fine spatial resolution, the entire Pixel Detector read-out requires almost a 100 million channels.

Moving outwards from the beam, the next sub-detector is the Semiconductor Tracker (SCT) made out of silicon strips [87]. It consists of four barrels and nine discs per end-cap, and radially extends from 299 cm to 514 mm. The 6.4 cm long strips are parallel to the beam axis in the cylindrical modules, while in the end-caps they have a radial arrangement. The precision coordinate resolution is mostly determined by the pitch of the strips (80 µm). To provide a 2D measurement, each layer is the combination of two sub-layers where strips

Table 2.1: Main parameters of the ID sub-detectors.

	Dimensions	Accuracy	Channels
Pixel	(50×400) µm	10 µm $(r\phi)$, 115 µm $(z\,/\,r)$	80 M
IBL	(50×250) µm	10 µm $(r\phi)$, 72 µm $(z\,/\,r)$	12 M
SCT	length 6.4 cm, pitch 80 µm	17 µm $(r\phi)$, 580 µm $(z\,/\,r)$	6.3 M
TRT	length 114 cm / 37 cm, ⌀ 4 mm	130 µm $(r\phi)$	350 k

are tilted by a small stereo angle (40 mrad).

The Transition Radiation Tracker (TRT) is the outermost sub-detector of the ID [88], and is composed of gas-filled drift (straw) tubes of a diameter of 4 mm, separated by transition radiation material. The tubes are arranged similarly to the SCT strips: parallel to the beam in barrel region, and with radial orientation in the end-caps. On average, the TRT provides 30 points in the r-ϕ plane per track, within $|\eta| \leq 2$. When traversed by a charged particle, the inhomogeneous transition material induces the radiation of X-rays, with a probability that is proportional to the Lorentz factor γ. Therefore, electrons have higher chances of producing transition radiation, generating higher currents in the straw tubes. Exploiting this property, TRT is additionally capable of doing electron identification.

Calorimetry The calorimeter system measures the energy and position of traversing particles, both charged and neutral. It is built with azimuthal symmetry and is designed to be hermetic, absorbing most of the energy produced in the collisions, except for the fraction carried away by muons and neutrinos, and it covers $|\eta| \leq 4.9$. In ATLAS, sampling calorimeters are used, where “passive” materials that absorb energy and induce particle showering alternate with“active” regions, which sample the energies of radiation showers.

Given the different nature of electromagnetic and hadronic showers, separate systems are designed to capture the two types of radiation patterns: the *electromagnetic calorimeter* based on lead/LAr sampling [89], and the *hadronic calorimeter*, which is mostly based on steel/scintillator-tile calorimeter [90], but with LAr active material in the forward region. The composition of the two calorimeters is shown in Figure 2.8.

An important characteristic of calorimeters is their energy resolution σ_E, which is commonly parametrised as:

$$\frac{\sigma_E}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c \,. \tag{2.7}$$

The stochastic (a), noise (b), and constant (c) terms are quadratically summed. The first term accounts for random fluctuations in the showering. The second is related to electronic noise and pile-up, while the last term corresponds to fluctuations that are a constant fraction of the deposited energies, such energy left in dead material or non-uniform responses. The ATLAS detector was designed to have values of a and c close to 10% (50%) and 0.7% (3%) for the electromagnetic (hadronic) calorimeter.

The electromagnetic calorimeter is divided into three parts: one barrel (EMB) and two end-caps (EMEC). Each part is contained in a separate cryostat filled with LAr. The EMB covers the $|\eta| < 1.475$ range, with a small overlap with the EMEC, which extends in η from 1.4 to 3.2. To optimise signal extraction and geometric coverage, both the absorbers and the electrodes are accordion-shaped. Figure 2.9(a) illustrates this peculiar geometry. It also shows how the calorimeter modules are branched into three layers, with varying granularity. The innermost layer, Layer 1, has the highest η precision ($\Delta\eta = 0.003$) but its

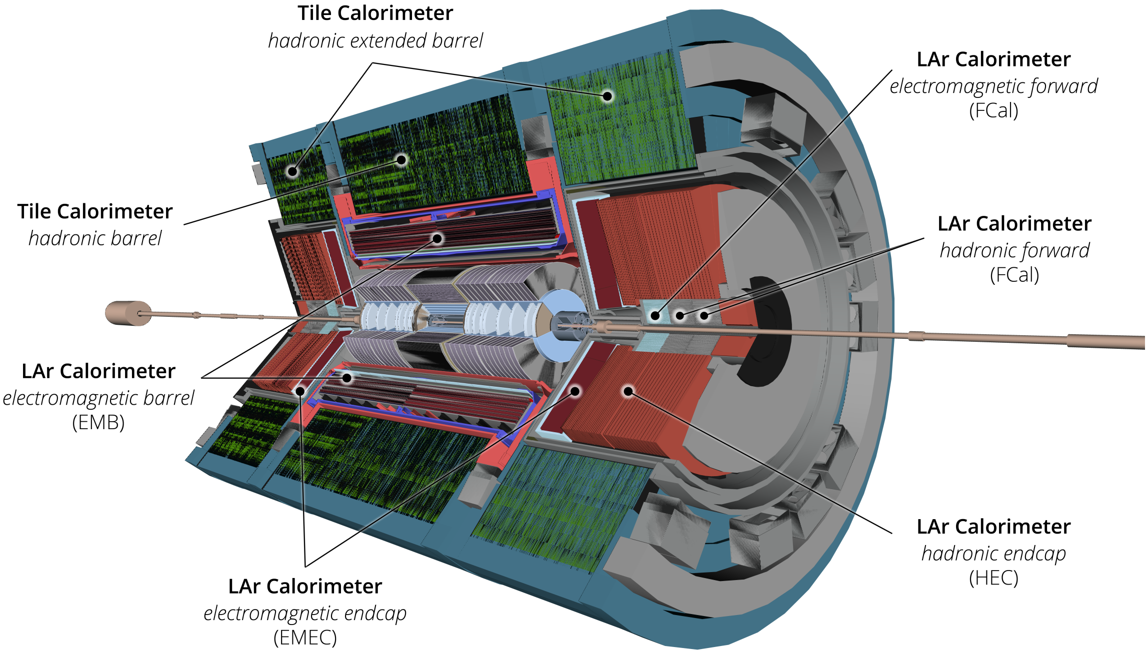

Figure 2.8: Cut-away view of the ATLAS calorimeter system [84].

cells are relatively wide in ϕ ($\Delta\phi = 0.1$). Layer 2 has cells with dimensions $\Delta\phi \times \Delta\eta = 0.0245 \times 0.025$, and cells in the last layer (Layer 3) are twice as wide in η but have the same size in ϕ as Layer 2. The same Figure also shows the layers respective *radiation lengths* (X_0), summing up to a total of $\sim 22\,X_0$. Both the granularity and the material budget vary slightly as a function of η, especially in the EMEC. In the barrel region, the calorimeter is complemented by a thin (11 mm) LAr presampler.

The hadronic calorimeter consists of a Tile calorimeter (a barrel and two extended-barrels), two LAr Hadronic End-caps (HEC), and a LAr forward calorimeter (FCal). The Tile calorimeter alternates radially oriented scintillating crystal tiles with steel ones as absorbers, as shown in Figure 2.9(b), and extends up to $|\eta| < 1.7$. It has three layers, with cells of a fixed $\Delta\phi = 0.1$ size and $\Delta\eta = 0.1\,(0.2)$ for the first two (last) layers. The HEC features similar technologies to the EMEC, with the only difference being the usage of copper absorbers instead of lead. The FCal shares the same cryostat with the HEC, covering the very forward region ($3.1 < |\eta| < 4.9$), and protects the external muon detectors from hadronic radiation. In the forward region ($2 < |\eta| < 4$), the LAr cryostat also contains Minimum Bias Trigger Scintillators (MBTS) [91], which are considered part of the hadronic calorimeter. They are employed in the trigger system and for luminosity monitoring.

The ATLAS calorimeters are *non-compensating*, meaning that their response to electromagnetic and hadronic radiation is unequal. They convert energy into digital signals more efficiently for photon and electron showers. This imposes a careful calibration procedure to correctly estimate the energy corresponding to different radiation showers.

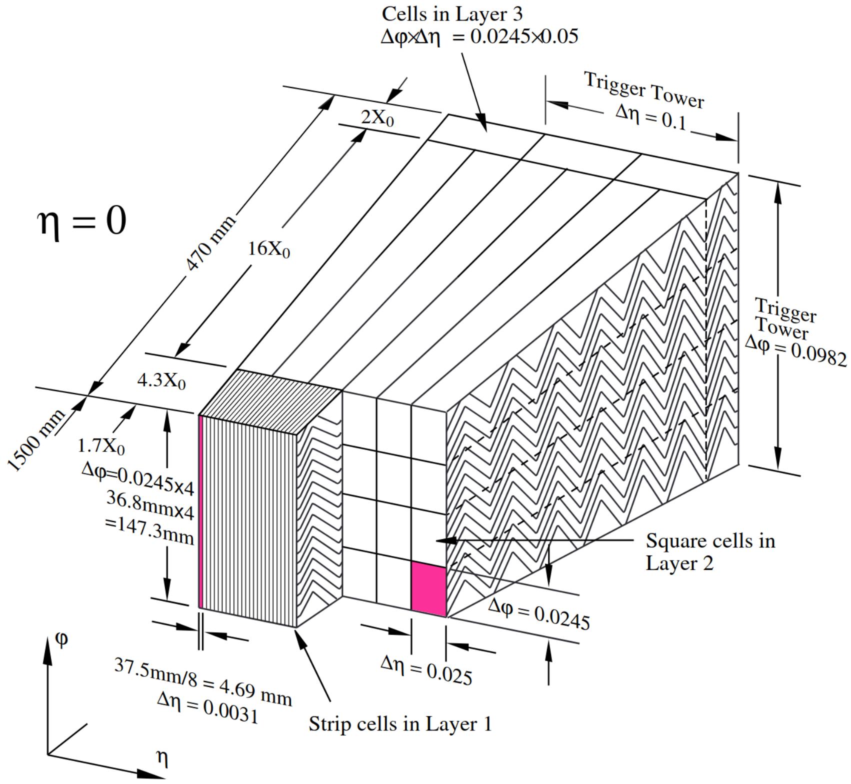

(a)

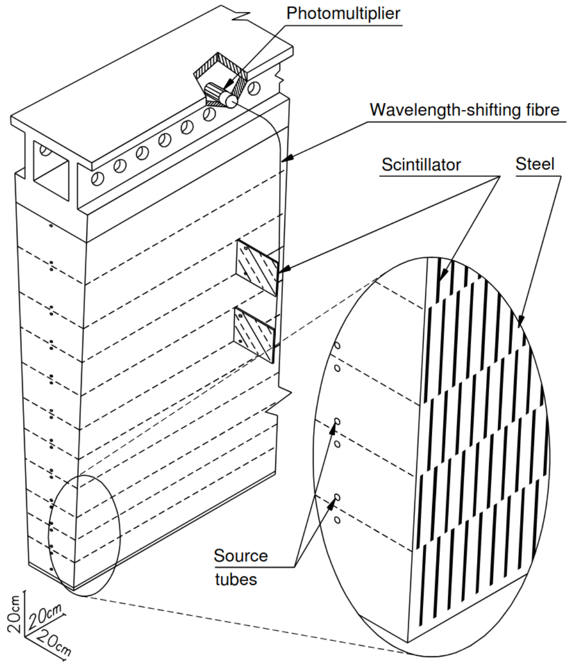

(b)

Figure 2.9: Sketches of (a) a barrel module of the electromagnetic calorimeter showing the accordion geometry and the granularity of the three different layers, and (b) the mechanical assembly and optical read-out of Tile calorimeter [83].

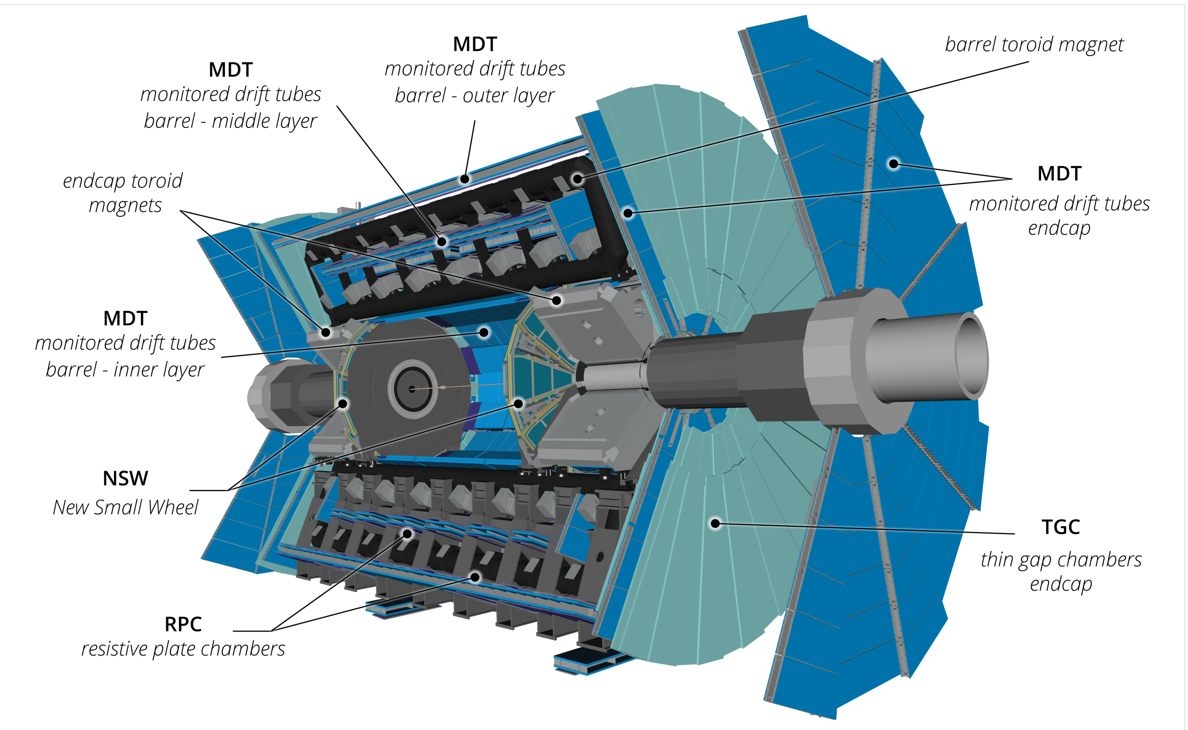

Figure 2.10: Cut-away view of the ATLAS Muon Spectrometer components [84].

Muon Spectrometer The outer part of the ATLAS detector constitutes the MS [92], which bends trajectories, measures positions and reconstructs the p_T of charged particles escaping the calorimeter volume. The magnetic field is provided by the barrel and end-caps toroids and is mostly perpendicular to the trajectories of particles originating from the IP. The MS, illustrated in Figure 2.10, primarily employs gaseous detector technologies and can measure muons up to $|\eta| < 2.7$, using three concentric barrel layers and a total of six end-cap wheel detector systems.

All the barrels are composed of Monitored Drift Tubes (MDT) for precision tracking, and the outer two also include resistive plates chambers (RPC) for triggering. The outermost wheels are equipped with only MDTs, while the middle ones have MDTs, and Thin Gap Chambers (TGC). The MDTs precisely measure the bending coordinate, while TGCs are used for azimuthal measurements and triggering. The innermost *small wheels* are positioned immediately after the end-cap calorimeters. They were upgraded for Run 3 with the New Small Wheels (NSWs), featuring small-strip TGC (sTGC) and micro-mesh gaseous structure (Micromegas) detectors. These detectors are fast and have fine resolutions, making them suitable for both triggering and 2D precision reconstruction.

Forward detectors Beyond the three main detection systems (ID, calorimeters, and MS), ATLAS is equipped with four small detector systems in the very forward region. These systems primarily serve as luminometers and employ different technologies. The LUminosity Cherenkov Integrating Detector (LUCID) system [93] sits at ± 17 m from the IP, measuring forward inelastic p-p collisions for luminosity monitoring. At ± 140 m, the Zero Degree Calorimeter (ZDC) [94], made of quartz and tungsten, is installed to measure the centrality of heavy-ion collisions. Then, at ± 210 m, the ATLAS Forward Proton (AFP) [95] silicon-based tracker studies soft, hard, and central diffractive events at low luminosities. Finally, the Absolute Luminosity for ATLAS (ALFA) Roman Pot detector [96] lies at ± 240 m from the IP. It is used in low luminosities and high β_b^* runs to provide measurements of the total p-p cross-section.

2.3 ATLAS Trigger system

As the bunch-crossing rates are close to 40 MHz, and the amount of data produced in a single collision event (or simply *event*) is greater than a MB, ATLAS cannot afford to continuously read out all the detectors and write the information to disc. The trigger and data acquisition (TDAQ) system [97, 98] is responsible for selecting collision events and saving them into permanent storage units. The trigger performs real-time reconstruction of events and decides whether to discard them or retain them if they contain processes of interest. Therefore, the trigger must be robust and efficient in retaining events for physics analysis while minimising the recording rate.

The trigger's task becomes more challenging as the luminosity – $i.e.$, pile-up – increases. However, maintaining stable and adequate triggering conditions is crucial for the experiment. For Run 3, several upgrades have been implemented [2] to handle a number of interactions per bunch-crossing ranging from 60 to 70, some of which are described in this section. Figure 2.11 illustrates the Run 3 configuration of the ATLAS TDAQ system. A two-stage approach is adopted for the trigger: a fast hardware-based Level-1 Trigger (L1) is followed by a software-based High-Level Trigger (HLT).

The L1 system is based on custom electronics that aggregate information from the calorimeter (L1Calo) and the MS (L1Muon) with coarse granularity. It also includes the L1 Topological processor (L1Topo), which combines the kinematics derived from the other two L1 processors to reconstruct the event topology. The Central Trigger Processor (CTP) combines the output of the three L1 sub-systems and is responsible for applying the event filtering. If an event passes the L1 selections, the CTP sends an "accept" signal to the Data Acquisition system (DAQ), and informs the HLT with the detector's Regions-of-Interest (RoIs) identified at L1. Furthermore, when needed, the CTP propagates "busy" signals to the DAQ, introducing *dead time* and limiting the acquisition of consecutive events to prevent overflows in the read-out buffers. The L1 system runs at a fixed latency of 2.5 μs and reduces the event rate to less than 100 kHz.

Events passing the L1 filter are processed by the HLT. The latter performs *online* event reconstruction, executing algorithms that are as close as possible to the *offline* ones, used for physics analyses. The same software suite is used for both cases. Reconstruction

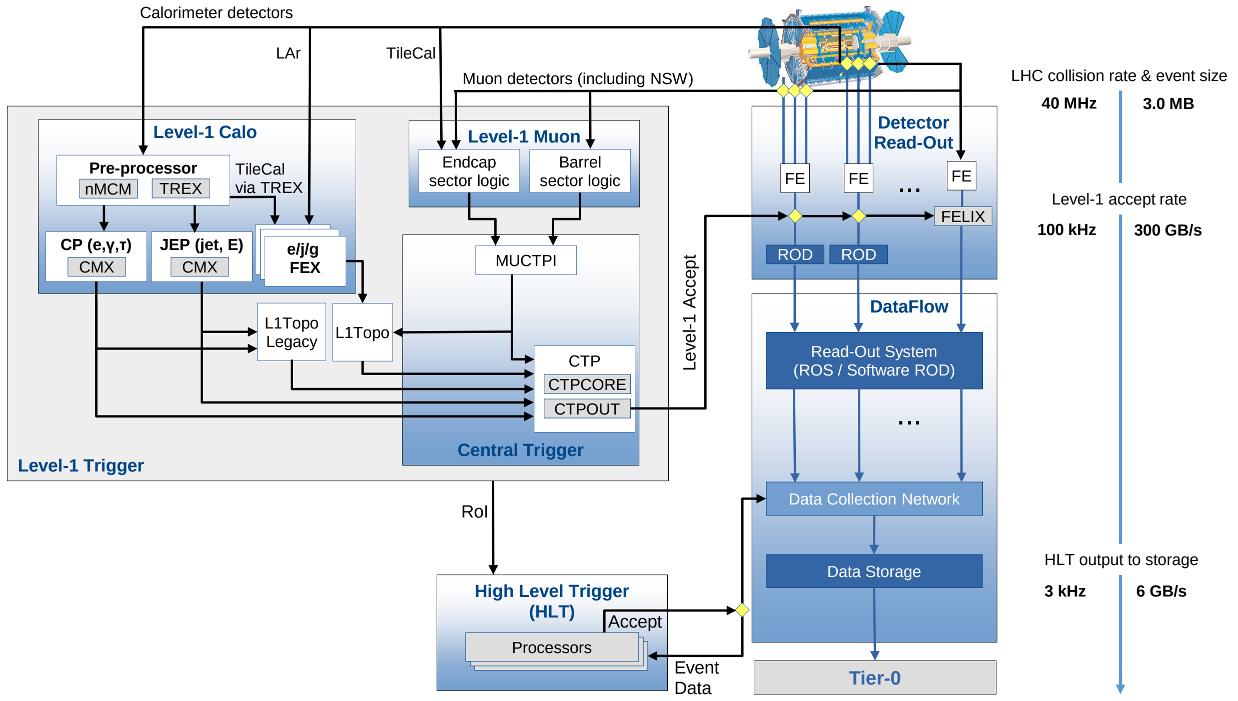

Figure 2.11: Schematics of the ATLAS TDAQ system for Run 3 [2].

steps are done mostly in RoIs to optimise computing time. However, to increase trigger efficiency, some algorithms run within the full detector acceptance, typically only after the rate is reduced by early preselections.

The L1 electronic boards are installed in the experimental cavern due to latency constraints. Conversely, the HLT computing farm is located at ground level, directly above the detector, and comprises more than $50\,000$ CPU cores. In addition, L1 accesses a limited number of sub-detectors from the calorimeters and the MS, while the HLT reads out all detectors, including the ID tracking system. The HLT output rate of interesting physics is close to $3\,$kHz. Considering that the average event size is below $2\,$MB, roughly $6\,$GB/s of data are written to disc.

Trigger reconstruction

For each calorimeter, the L1Calo system [99] stacks cells across various layers with reduced granularity ($\Delta\phi \times \Delta\eta = 0.1 \times 0.1$ in the barrel and irregularly larger in the end-caps), forming the so-called *trigger towers*. Candidate electron (e), photon (γ), or jet momenta are estimated by summing the energies in adjacent towers. For e/γ, RoIs of 2×2 towers in the LAr are identified, and isolation criteria can be imposed based on the 12 neighbouring towers or the corresponding hadronic towers, as shown in Figure 2.12(a). For jets, the granularity is reduced by using sums of 2×2 towers and forming RoIs of 4×4 or 8×8. In Run 3, the L1Calo system was upgraded for the LAr calorimeter, which now provides

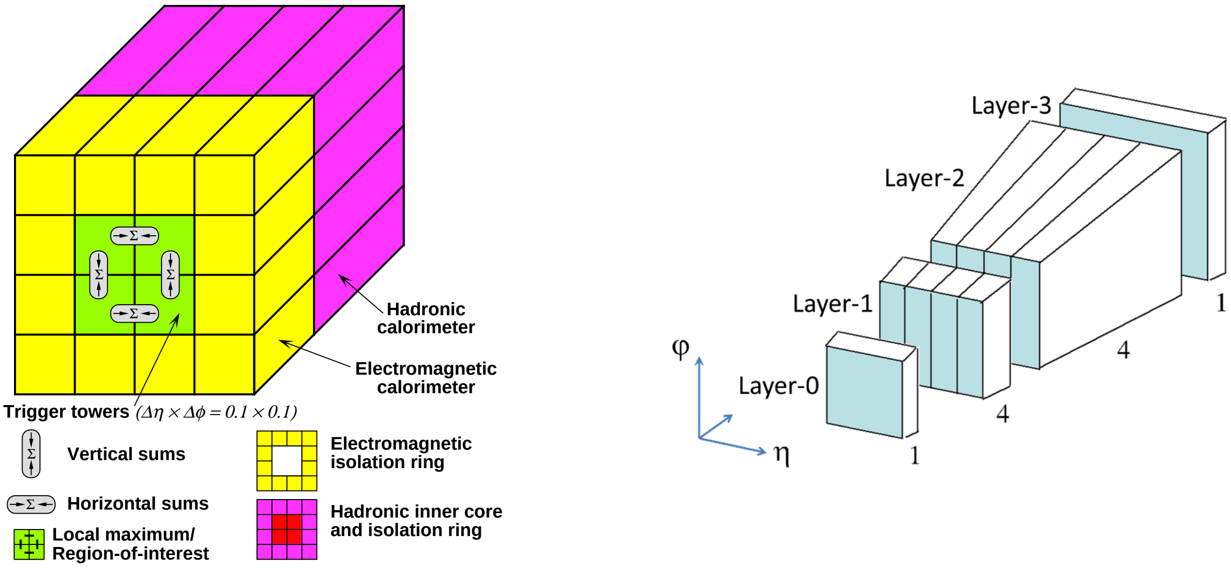

(a) (b)

Figure 2.12: Figure (a): trigger towers for L1 e/γ algorithms [97]. Figure (b): updated longitudinal granularity of LAr calorimeter SuperCells introduced in Run 3 [2].

SuperCells, as shown in Figure 2.12(b). The new system gives access to shower shape variables for e/γ identification and grants finer granularity for jets (0.1×0.1) improving the performance of multi-jet and missing transverse momentum selections. The L1Muon triggers [100] check for coincidences between the RPC stations in the two outermost barrels and TGCs in the inner and middle wheels. The hits are also used to estimate the p_T of the muons. During LS2, the NSW Micromegas and sTGC were installed to replace the innermost wheel, and further reduce rates at low p_T. The distance in the η-ϕ plane ($\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$) between L1 object pairs and their invariant mass is calculated by the L1Topo system [101], along with other event-level quantities.

The HLT event reconstruction uses the Athena software framework [102], as done for offline analysis. To overcome memory limitations and enable faster processing, Run 3 event reconstruction runs in a multithreaded fashion. It was developed following two key principles of the HLT event selection: prioritise, as much as possible, RoI reconstruction in gradual steps, and stop the event processing as soon as no trigger criteria are satisfied. These two principles required the implementation of a robust control flow mechanism, which is achieved with mostly three types of algorithms. These are the input preparation, the reconstruction (or feature extraction), and the hypothesis testing algorithms. The latter type is responsible for evaluating the predefined selections. The HLT reconstruction of physics objects emulates as much as possible the offline one, which is discussed more in depth in the next chapter (Chapter 3), while the reconstruction of trigger (b-)jets is described in Chapter 5.

Trigger Menu

The configuration of the trigger system is organised within a *trigger menu* containing a list of *trigger chains*, or simply *triggers*. Each chain specifies some L1 trigger requirements and the sequence of HLT algorithms to be executed. Triggers can be separated into two main classes. The *physics* triggers are designed for physics analyses and performance measurements. The latter are typically *prescaled*, their rate is artificially reduced by applying a *prescale* factor n, which decreases their activation probability by $1/n$. The *auxiliary* triggers are dedicated to detector calibrations and special running conditions, such as cosmic ray data acquisition or random noise studies.

The DAQ system records the accepted events into separate data streams defined in the menu by different sets of chains. The *Physics streams* comprise the *Main*, *Hadronic*, *B-physics and light states* (BLS), and *Delayed* streams. The events in the first three are reconstructed within a few days from the collection, while the latter stream is processed at a later moment, sometimes months afterward, when the computing resources are available. Another special stream for physics analysis is the Trigger-Level Analysis (TLA), described in more detail at the end of this section. The *Express* stream is a very small subset of the Physics ones, where events are promptly reconstructed for data quality monitoring. Other four auxiliary streams exist: the *Background*, *Debug*, *Calibration*, and *Monitoring* streams. Additional streams are defined for special runs.

To avoid data duplication, overlap between the streams is minimised, as shown in Figure 2.13(a) for a typical p-p run. The stacked rates of all the streams reache 10 kHz, close to the total HLT output rate of 9 kHz. Furthermore, the menu configuration is adapted based on luminosity. In Figure 2.13(a) the delivered luminosity is stable for more than five hours. After such a long time of continuous collisions, the LHC beams start to degrade, and the luminosity drops. The triggers are updated accordingly to enhance their efficiency to interesting physics. As the rate diminishes, more resources becomes available. Triggers with lower thresholds can be therefore activated, and prescale factors reduced. These changes produce the spikes in the output bandwidth visible at $8{:}30$, $11{:}00$, and $14{:}30$ in Figure 2.13(b).

Trigger-Level Analysis

The TLA stream records only a very limited number of objects reconstructed by the HLT, without including any raw detector information. This means that physics analysis of such data must rely solely on the coarser resolution of HLT reconstruction. However, this drastically reduces the event size ($e.g.$, 4.5 kB in 2022 data-taking and close to 40 kB from 2023 onwards), allowing for much higher rates compared to the standard Physics streams. In this way, TLA achieves the highest output rates while having a negligible footprint on the bandwidth, as demonstrated by the two plots in Figure 2.13.

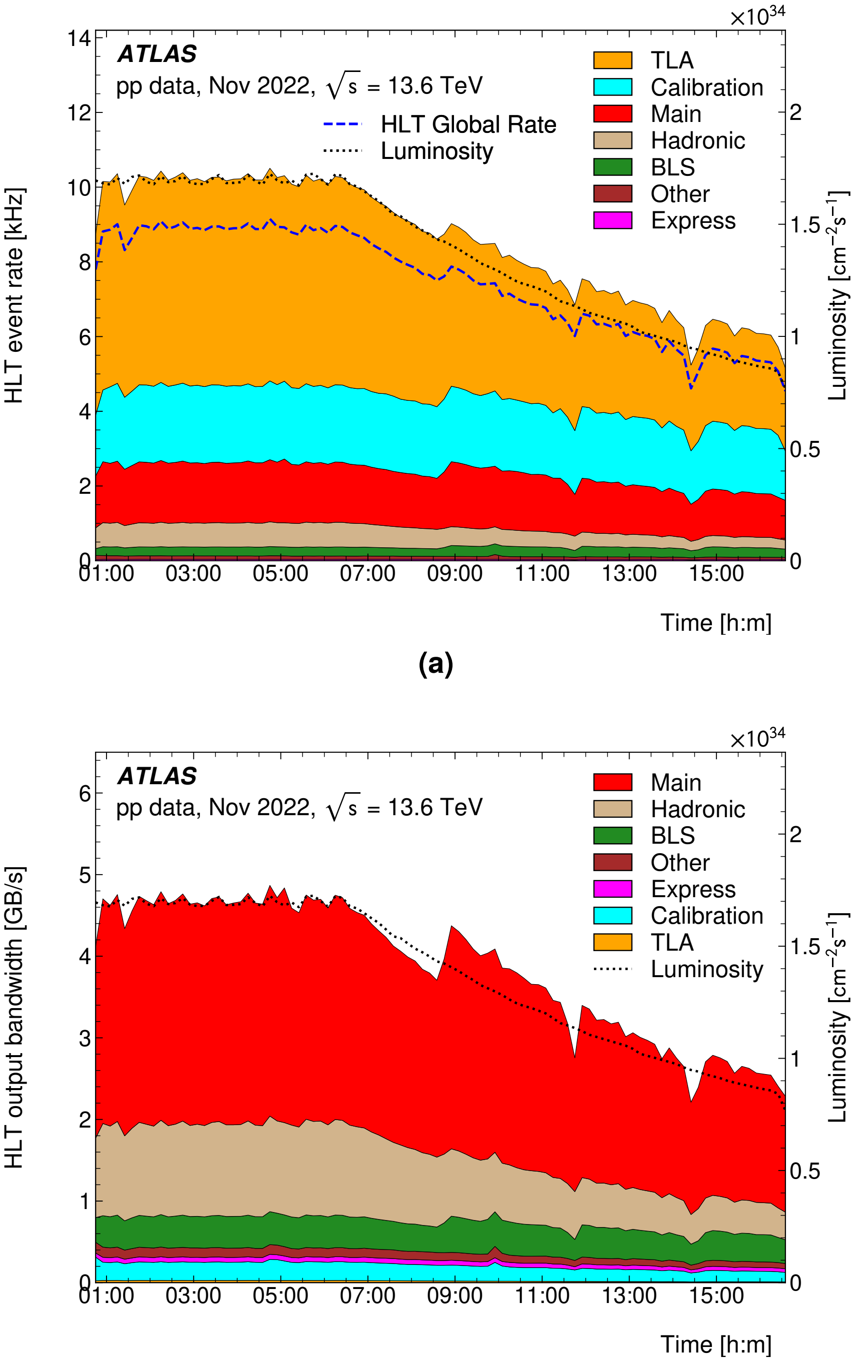

(a)

(b)

Figure 2.13: Monitored (a) event rates and (b) output bandwidths at the HLT grouped by data stream, during a typical p-p run in 2022 [2].

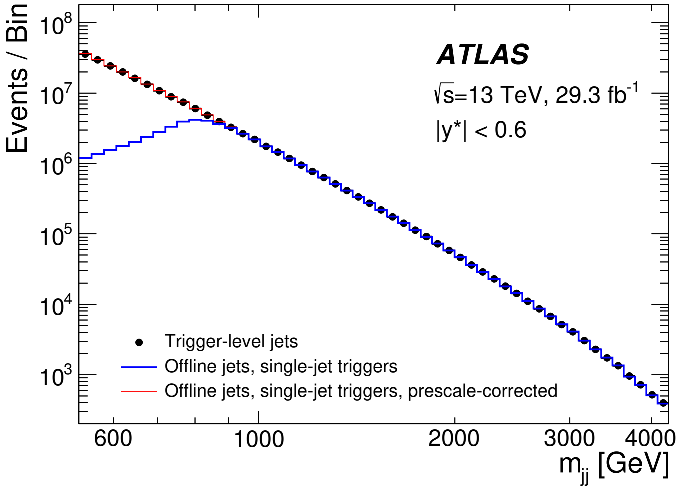

Figure 2.14: Comparison of the dijet invariant mass spectrum between: the TLA stream, the single-jet triggered, and the single-jet prescale-corrected events [103].

The TLA strategy effectively enables access to regions of phase space that would otherwise be trigger-limited. This is the case for soft dijet resonances, which are impacted by the trigger jet p_T thresholds. Figure 2.14 shows that TLA reconstructed dijet events with invariant masses smaller than 600 GeV, while the Physics stream triggers start to be heavily prescaled at 900 GeV and below. This approach was introduced at the start of Run 2, and included only a restrained set of HLT jet information, allowing for BSM searches in dijet events with unique sensitivity [103]. For Run 3, the scope of the TLA stream has been expanded, thanks in part to this thesis work, as discussed in Chapter 6.

2.4 Monte Carlo simulations

A key component of the experiment is the simulation of collision events using Monte Carlo (MC) techniques. These events are essential for many practical purposes. They are required for physical interpretation of the observed data, providing a means to compare experimental results with theory predictions. Furthermore, they are used to understand the detector performance, to design and test trigger selections prior to deployment, and to develop new reconstruction techniques.

Conventionally, simulation production is divided into two main tasks: event generation and detector simulation [104]. The first task is steered by both fundamental theory and phenomenological models, and it provides events with associated probabilities, know as *event weights*. Each event consists of a list of stable particles and their *truth-level* properties, such as kinematics, parent-particle information, and production point. The detector simulation propagates these particles through the ATLAS detector and emulates its response. Both tasks involve stochastic quantum processes; hence, they strongly rely on MC algorithms for pseudo-random number generation.

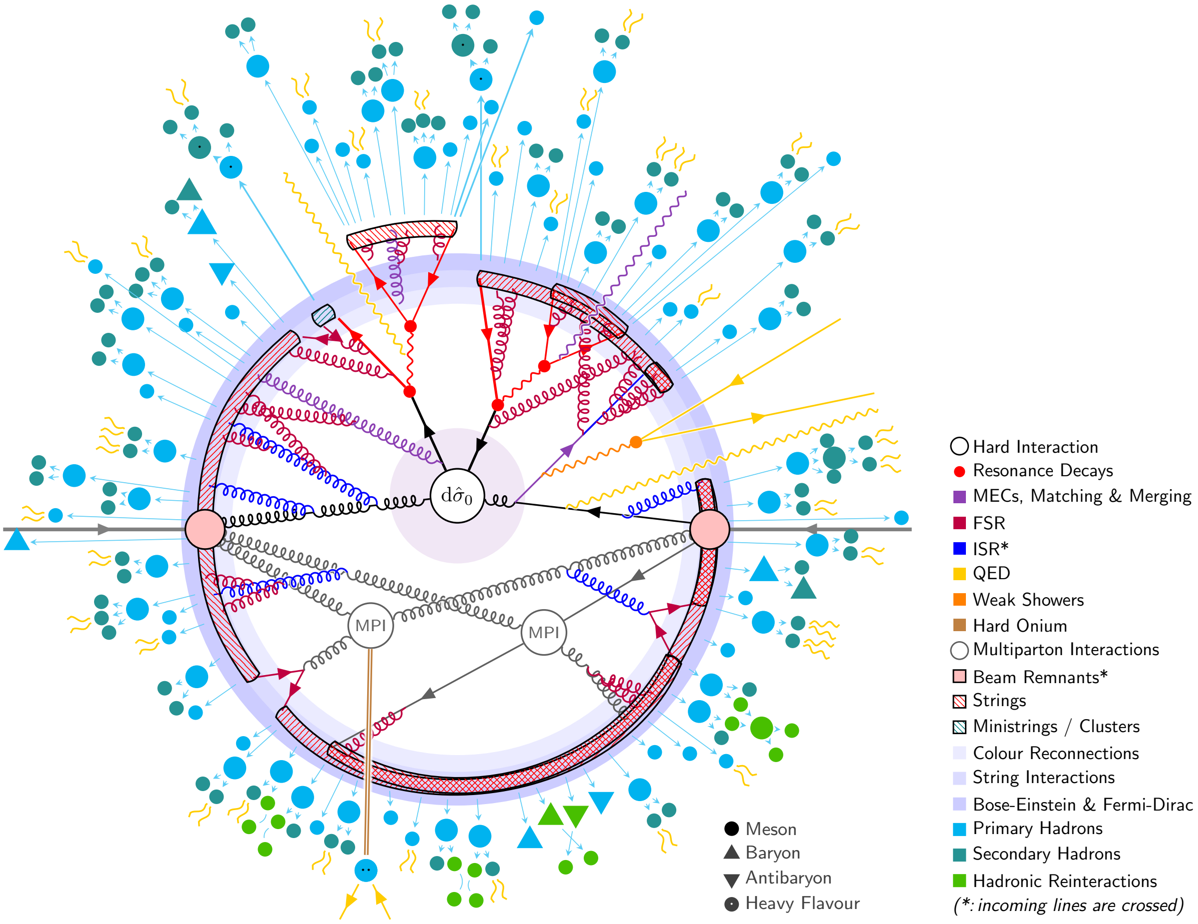

Figure 2.15: Schematics of a $pp \to t\bar{t}$ event [105].

Event generation

The collision events are extremely rich in terms of physics, as illustrated in Figure 2.15. Furthermore, as the precision of the experiment increases, more effects need to be modelled. There is no analytical form describing the full process; therefore, event generation comprises multiple steps.

A core component is the calculation of the differential *partonic cross-section* $d\hat{\sigma}_{ij\to X}$. This term represents the hard scatter – ij partons producing a final state X – which is typically the key process of interest. Along with phase space factors, $d\hat{\sigma}_{ij\to X}$ contains the probability amplitudes and, consequently, the theoretical assumptions (SM and/or BSM). This cross-section is derived using perturbation theory through *fixed-orders* expansions of the coupling constants. Typically, the dominant contributions arise from α_s terms; however, depending on the required precision, other couplings may also become relevant. This approach requires the coefficient of the expanding parameters to be sufficiently small to ensure the series convergence. In QCD, large logarithmic terms emerge due soft and collinear emissions, but their behaviour can be characterised at all orders and *resummed* into finite corrections.

As mentioned in Section 2.1, the probability of extracting a parton i with an energy fraction of x_i is determined by the PDF $f_i(x_i)$. The cross-section for p-p collisions is then expressed as:

$$d\hat{\sigma}_{pp\to X} = \sum_{ij} \int_0^1 dx_i \int_0^1 dx_j \; f_i(x_i)\, f_j(x_j)\, d\hat{\sigma}_{ij\to X} \,. \tag{2.8}$$

The evaluation of the latter equation is significantly simplified by the *factorisation* property of QCD, which allows separation of the hard- and soft-scale processes. In this way, divergent soft emissions and non-perturbative effects, below an established factorisation scale μ_F, can be removed from the cross-section calculations and treated in subsequent stages.

The QCD factorisation also allows the absorption of soft initial-state divergences into the PDFs, which are not determined from first principles. They are instead constrained by experimental data, such as deep inelastic scattering measurements at HERA [106]. These observations are done at lower energy than the LHC collisions. However, PDFs can be evolved to higher scales – specifically to the μ_F scale – using the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) equations [107].

After the hard process is determined and the cross-section is calculated above the μ_F scale, parton splitting is modelled using an algorithmic approach, the *parton showering*. These iterative algorithms rely on Sudakov factors [108] to express the splitting probabilities.

Parton showers evolve until the non-perturbative QCD regime is approached. At energy scales close to Λ_{QCD}, partons hadronise into colour-singlet states. This process requires a phenomenological description tuned to data observations. Two primary models are employed: the *Lund string* model and the *clustering* model [105]. The first one is based on the structure of the strong potential, which linearly increases with the distance between two colour charges. At sufficiently large distances, the force flux tubes can break, generating quark-antiquark pairs, as schematised in Figure 2.16, thereby leading to hadron formation. The clustering model, by contrast, is based on the *preconfinement* property of

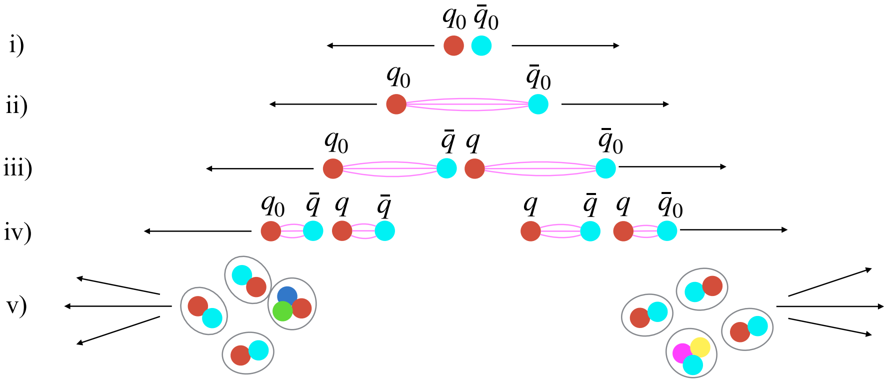

Figure 2.16: Schematics of the different steps of the hadronisation process for a dipole string between a $q\bar{q}$ pair [109].

QCD, which ensures that colour-neutral clusters can always be formed from the radiated partons following the showering stage.

Beyond the hard-scatter partons, the remaining fragments of the colliding protons must also be simulated as they contribute to the observable activity in the detector. This component of collisions is typically referred to as the *underlying event*. In some cases, if the protons have a small distance of closest approach, multi-parton interactions (MPI) can take place, as shown in Figure 2.15. These contributions correlated with the hard-scatter and must be included in the hadronisation process.

Detector simulation

The simulation of particles interactions with the multitude of ATLAS sub-detectors heavily relies on the GEANT4 toolkit [110]. GEANT4 models the passage of particles through matter, taking into account the detailed geometry and material composition of the detector. It simulates a wide range of physical processes occurring as particles traverse the detector, including electromagnetic and hadronic showers, energy loss through ionisation, multiple scattering, and particle decays. The outcome of this detector simulation stage is recorded as "hits", representing the energy deposited by particles in the detector's sensitive elements, along with their timing and spatial information.

The propagation and showering of particles through the detector is the most CPU-intensive stage in the whole MC simulation workflow. The majority of computing time is spent on the shower development within the calorimeter, which is remarkably complex due to its intricate geometry. To accelerate this stage for the production of large samples, a *fast* simulation approach is sometimes employed in ATLAS. This method replaces the *full* calorimeter simulation with a parametrised approach, and more recently, with generative ML methods [111].

After the "hits" files are produced, the *digitisation* takes place, emulating the detector readout. This step converts the energy deposits into *digits*. Depending on the specific sub-detector system, different type of information is reproduced. For example, in the SCT, sampled signal is digitised to a single-bit information to indicate where it exceeded a given threshold, whereas in the Pixel and LAr systems, the signal amplitude is also recorded. The digitised readout of all the detector channels is stored in Raw Data Objects (RDOs) files that serve as input for object reconstruction.

In addition to simulating the primary particle interaction of interest, the ATLAS experiment incorporates the pile-up effects into its MC simulations. Since Run 2, large MC datasets of *minimum-bias* events – soft inelastic p-p collisions – have been generated, which are sampled and overlaid onto the hard scatter, as shown in Figure 2.17. This pool of pile-up interactions is generated centrally once per MC production campaign, with a size nearly 500 million events per year. To optimise computing resources, the overlay is

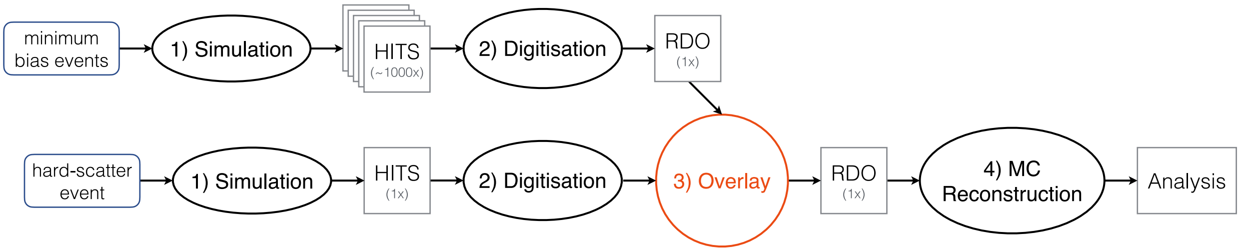

Figure 2.17: Diagram of the ATLAS MC workflow since Run 2 [112].

performed after the digitisation, using *presampled* pile-up RDO datasets [112]. For each hard scatter, numerous minimum-bias events must be sampled depending on the instantaneous luminosity. Given that the longest signal integration time is from the MDTs and spans up to 750 ns before the triggering collision and 150 ns after, $32{+}6$ additional bunch-crossings must be simulated. Therefore, for an average number of collisions $\mu = 50$, a total of $(38 + 1) \times 50 = 1950$ interactions are required. With the presampled pile-up approach, the inclusion of these supplementary collisions has a minimal computational cost that is independent of μ.

The digitisation and overlay procedures are customised for each sub-detector in order to simulate the behaviour of the different readout electronics. This applies also to the L1 trigger hardware. The L1Calo system relies on a custom digital processor. Its signal processing chain is precisely simulated, including noise contributions. In contrasts, HLT simulation runs the same software as used in data-taking.

Chapter 3

Putting the pieces together: offline event reconstruction

Many particles are generated during each p-p collision within ATLAS. These particles interact with the detector via ionisation, multiple scattering, hadronic interactions, or other mechanisms determined by the particle properties. The majority these particles lose all their energy through such processes. As a result, most of the produced radiation remains confined within the instrumented volume, where it is measured based on the energy transferred to active detection materials. Depending on the detector technology, either raw quantitative information about the energy or simply the “hit” position are digitised for further processing. Adjacent energy deposits in the calorimeter cells, as well as hits in the tracking system, are grouped using different clustering algorithms. Hit clusters are used to reconstruct the trajectories of charged particles. These tracks, along with calorimeter clusters, form the basic the building blocks for reconstructing the different physics objects (or physics *signatures*), with their associated four-vectors.

This chapter reviews the offline reconstruction procedures for physics objects, with a greater focus on hadronic signatures, which are the main objects considered in this thesis.

3.1 Tracking and vertexing

Five independent parameters are required to describe the trajectories of charged particles: $(d_0, z_0, \phi, \theta, q/p)$, where d_0 and z_0 are the transverse and longitudinal impact parameters, the angles ϕ and θ define the track direction at the point of closest approach, while the charge-momentum ratio q/p describes the track curvature. The impact parameters and the angles are defined relative to a reference point, typically the beam spot position, which represents the average position of the p-p interactions.

The track reconstruction procedure aims to precisely estimate the track parameters by connecting the tracker hits into patterns that match the particle trajectories. This task is particularly challenging in the ATLAS ID tracker considering the dense radiation environment. Incorrectly associating hits to tracks reduces reconstruction accuracy, while erroneusly combining hit patterns can produce artefact tracks (or *fake tracks*) with no corresponding real particle.

For ID tracks, a robust methodology is developed to maximise efficiency while maintaining low fake reconstruction rate [113]. In this way, tracks within the ID geometric acceptance ($|\eta| < 2.5$) and with $p_T > 500$ MeV are reconstructed. As shown in Figure 3.1, the procedure consists of two main stages: a primary inside-out tracking process followed by a backward pass. The primary tracking stage consists of the following steps:

1. **Space Point and Drift Circle Formation** Neighbouring hits in the Pixel and SCT detectors are clustered forming 3D space points, with position uncertainties determined by the detector characteristics. For TRT, 2D drift circles are formed.

2. **Pixel and Strip Seed Finding** Triplets of space-points are constructed combining compatible points in the Pixel and SCT detectors. These triplets compose the track *seeds* and are filtered based on rough estimates of their track parameters.

3. **Track Finding** To mitigate the combinatorics, for each seed, *search roads* are defined as detector regions that are expected to be compatible with the seed kinematics. Along the search roads, clusters are progressively added to a track candidate using a combinatorial Kalman Filter [114]. Finally, parameter estimation is performed by fitting the tracks to the associated clusters.

4. **Ambiguity Resolution** Track candidates are ranked by a set of quality criteria, such as the number of “holes” (number of detector layers without any cluster associated to the track), the number of shared hits with other track candidates, and the resolution of the clusters forming the track. Low-quality tracks are discarded, reducing fakes. Eventual shared hits from rejected tracks are reassigned to the retained ones. If clusters are still shared among two or more tracks, a splitting algorithm is run to determine whether the cluster has contributions from one or more charged particles [115].

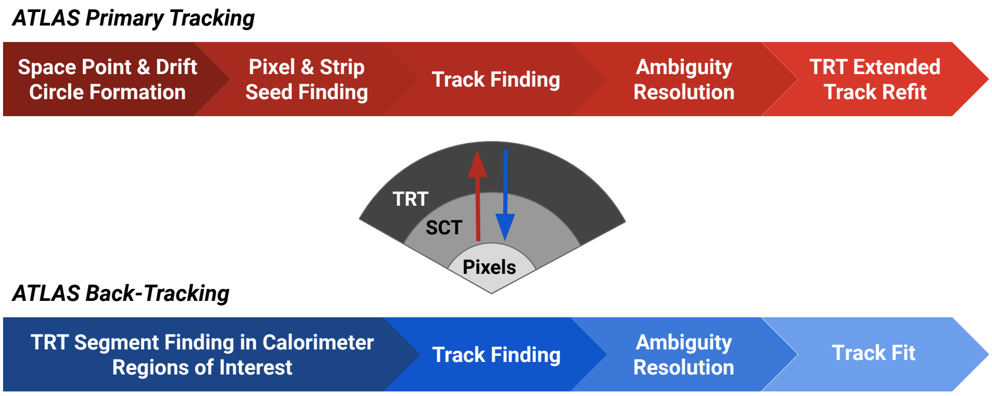

Figure 3.1: Schematics of the ID primary inside-out tracking stage and of the backward pass [116].

5. **TRT Extended and Track Refit** Tracks are re-fitted with a global χ^2 method, resulting in *precision tracks* when the fit converges, otherwise they are rejected. If compatible TRT hits are found, tracks are extended using TRT drift-circles. This can improve the momentum resolution and provide particle identification.

The inside-out approach is optimised to reconstruct particles originating from the beam spot. To improve efficiency for tracks displaced from the primary interaction point ($e.g.$, photon converting into e^+e^- pair), a backward tracking pass is performed using the remaining hits. In this step, TRT hits compatible with energy deposits in the LAr calorimeter are employed as seeds. A similar procedure to primary tracking is then applied.

For muon tracks, a two-stage procedure is also adopted, but starting with an outside-in pass followed by an inside-out pass. In this case, hits from the MS seed the primary tracking step and are then combined with information from both the calorimeters and ID systems.

During Run 2, tracking was the most CPU intensive reconstruction step due to its complexity. Considering that each inelastic p-p collision produces around 200 hits, increasing the instantaneous luminosity makes tracking even more challenging. In Run 3, for a single bunch-crossing up to $15\,000$ hits can be produced in the ID. To deal with such high density conditions, it was essential to improve the computing performance of tracking. As discussed in Section 2.3, the implementation of multi-threaded reconstruction was a crucial improvement for Run 3. Moreover, tracking specific optimisations reduced the CPU time requirement by 70% [116], including better seed definition and more rigorous usage of detector geometric bounds.

Vertex finding

After the ID tracks are built, an iterative vertex finding algorithm [117] employs them to locate all the p-p inelastic collisions. Firstly, tracks are selected according to quality criteria. Then, their z_0 distribution is used to find an initial guess of the vertex position. This estimate, together with the tracks, is iteratively fitted to precisely determine the vertex coordinates. In each iteration, the least compatible tracks are down-weighted to improve the fit. Once a vertex is found, the tracks associated are excluded and the procedure is repeated until all the tracks are processed. This method yields vertices, each with at least two tracks, achieving a resolution of $30\,\mu\mathrm{m}$ and $20\,\mu\mathrm{m}$ in the longitudinal and transverse dimensions, respectively.

In preparation for Run 3, vertexing was further improved by implementing an adaptive multi-vertex reconstruction algorithm [118] to handle higher pile-up conditions. This new method refines the initial guess for the vertex location using Gaussian seeds, and allows individual tracks to be included in multiple fits for neighbouring vertices.

Once the primary vertices (PVs) are found, the hard-scatter is defined as the one featuring the highest squared sum of tracks p_T:

$$\mathrm{PV}^{\mathrm{HS}} = \arg\max_{\mathrm{PV}_i} \sum_{\mathrm{track}\in\mathrm{PV}_i} p_T^2 \,. \tag{3.1}$$

This definition permits an efficient selection of vertices corresponding to processes of physics interest. After the $\mathrm{PV}^{\mathrm{HS}}$ is identified, the other vertices in the event are labelled as pile-up.

3.2 Calorimeter clustering

Calorimeters absorb most of the radiated particles, which typically disperse their energy in multiple cells across different calorimeter layers. To reconstruct the total energy of particles, calorimeter cells are clustered.

ATLAS employs two clustering approaches [119]: the *sliding-window* algorithm and the *topological clustering*. The former strategy is used to perform fast L1 trigger selections, as discussed in Section 2.3 on Trigger Reconstruction. In this approach, a grid of calorimeter towers is first built by summing cells at a fixed coarse granularity, radially across different layers in each calorimeter system. Next, the tower grid is scanned for single towers above a given energy threshold, which serve as *proto-clusters*, reference points for cluster formation. The last step consists in summing the proto-clusters and their neighbouring towers within a fixed size window, providing the final clusters. The main parameters of this algorithm are the proto-cluster energy threshold and the window size in $\Delta\phi \times \Delta\eta$. The latter parameter depends on the type of signature being reconstructed, and balances the amount of calorimeter noise being included with the fraction of signal leaking outside the considered area.

The sliding-window method was also used offline to build clusters associated to electrons and photons during Run 1 [120, 121]. From Run 2, these signatures use the topological clusters (or *topo-clusters*) instead, as was done for hadronic signatures.

Topological clustering

The ATLAS topological clustering approach [119, 122] is designed to group nearby cells that have recorded a significant amount of energy above the expected noise level. Its flexibility is more compatible with the irregular shapes of particle showers in the calorimeter than the sliding-window algorithm, producing clusters with varying numbers of cells.

The method relies on three threshold values for the cell significance, which is defined as:

$$\xi_{\text{cell}}^{\text{EM}} = \frac{E_{\text{cell}}^{\text{EM}}}{\sigma_{\text{noise}}^{\text{EM}}}\,, \tag{3.2}$$

where both the cell energy ($E_{\text{cell}}^{\text{EM}}$) and cell noise ($\sigma_{\text{noise}}^{\text{EM}}$) are evaluated at the calorimeter electromagnetic scale. The noise term is the combination of contributions from both electronics and pile-up: $\sigma_{\text{noise}} = \sqrt{\sigma_{\text{electronics}}^2 + \sigma_{\text{pile-up}}^2}$. The three thresholds on $\xi_{\text{cell}}^{\text{EM}}$ control the topo-cluster seeding (t_s), growth (t_g), and perimeter (t_p). The default setting is "420", where $t_s = 4$ and $t_g = 2$ to efficiently suppress the formation noise clusters, and with $t_p = 0$ for including the energy tails of showers.

The algorithm has two principal stages, the topo-cluster maker followed by a splitting stage. The first stage does run using all the calorimeter cells. It can be schematised into three steps:

1. **Seed Finding** Search for all the cells with $|\xi_{\text{cell}}^{\text{EM}}| > t_s$, creating a list of seeds (proto-clusters), as shown for example in 3.2(a).
2. **Neighbour Finding** For each seed, merge all the neighbouring cells with $|\xi_{\text{cell}}^{\text{EM}}| > t_p$, in three dimensions (usually cells can have 8 neighbours in the same detector layer, and 2 in the adjacent ones). For all the proto-clusters, the growth procedure is iteratively repeated considering all the cells with $t_g < |\xi_{\text{cell}}^{\text{EM}}| < t_s$ as new seeds. During this step, proto-clusters are merged if they share neighbouring cells.
3. **Finialise** When no more cells can be added, the clusters are formed and can be filtered by their total energy ($E_{\text{clus}}^{\text{EM}}$).

In this procedure, the absolute value of $\xi_{\text{cell}}^{\text{EM}}$ is considered, as large negative energy deposits can arise from out-of-time pile-up being negatively shaped by the LAr calorimeter read out electronics. The inclusion of negative energy cells enhances the noise suppression. However, only clusters with positive energy ($E_{\text{clus}}^{\text{EM}} > 0$) are considered for the construction of physics objects.

At higher levels of pile-up and in forward regions, clusters are more likely to merge. The splitting stage attempts to resolve distinct particle showers, reducing biases and improving the signature-specific object reconstructions. It first identifies local maxima cells with

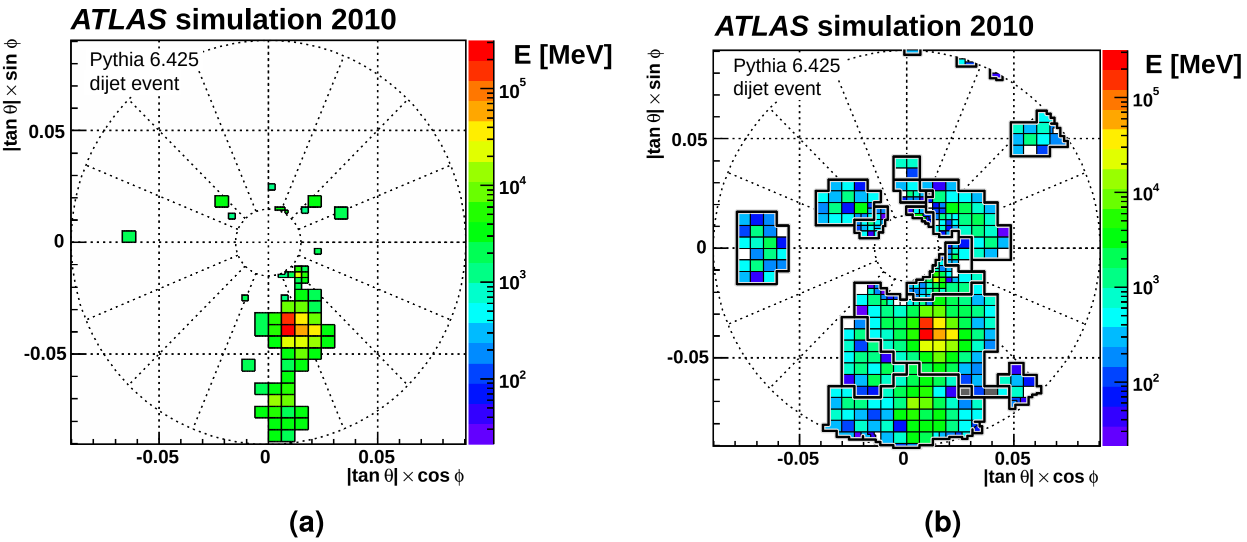

Figure 3.2: Example of (a) seeds cells passing the threshold t_s, and (b) the final topo-culsters formed after the splitting stage [122]. Only cells from the first module of the FCal system are shown for a simulated dijet event.

energies $E^{EM}_{\text{cell}} > 500$ MeV to be used as seeds. The number of local maxima determines the number of clusters formed after the splitting stage. The neighbour finding procedure is repeated from these seeds, as is done during the topo-cluster making stage. If a cell has two or more neighbouring growing clusters, it is assigned only to the two most energetic ones. Furthermore, contributions from shared cells to cluster energy are weighted based on the cluster energy and geometric distances:

$$w_{\text{cell},1} = \frac{E^{EM}_{\text{clus},1}}{E^{EM}_{\text{clus},1} + rE^{EM}_{\text{clus},2}} , \tag{3.3}$$

$$w_{\text{cell},2} = 1 - w_{\text{cell},1} , \tag{3.4}$$

$$r = \exp(d_1 - d_2) , \tag{3.5}$$

where d_1 and d_2 represent the distances from the centres of the two clusters. After this procedure is completed, the topo-clusters are formed, as shown in Figure 3.2(b).

The total energy of the topo-clusters is given by the weighted sum of the cells:

$$E^{EM}_{\text{clus}} = \sum_{i=1}^{N_{\text{cells}}} w_{\text{cell},i} E^{EM}_{\text{cell},i} . \tag{3.6}$$

The cluster directions (η and ϕ) and other kinematic observables (O) are defined by energy-weighted sums over the cells:

$$O_{\text{clus}} = \frac{\sum_{i=1}^{N_{\text{cells}}} w_{\text{cell},i} |E^{EM}_{\text{cell},i}| O_{\text{cell},i}}{\sum_{i=1}^{N_{\text{cells}}} w_{\text{cell},i} |E^{EM}_{\text{cell},i}|} . \tag{3.7}$$

3.3 Hadronic jets

At LHC, the most abundantly produced particles in p-p collisions are quarks and gluons. However, due to the confinement of the colour charge, quarks and gluons hadronise, producing *jets* – sprays of energetic colourless hadrons.

Jets are high-level objects that ideally represent the outgoing partons from the hard-scattering process, with jet and parton kinematics being equivalent. Moreover, a common jet definition is required from both the experiment and theory side, so that MC calculations can be compared with measurements. However, partons can easily branch according to perturbative QCD probabilities. For example, a parton i can split into j and k with a differential probability that is porportional to:

$$\frac{dP_{i\to jk}}{dE_k d\theta_{jk}} \propto \frac{\alpha_s}{E_k\theta_{jk}} , \tag{3.8}$$

in the limit of soft branching k and small relative angle θ_{jk}. This expression leads to divergent probabilities for soft and collinear emissions, making the definition of final-state parton ambiguous. Additionally, a wide spectrum of processes can produce partons, resulting in the existence of many jet definitions [123].

Jets are the main physics objects used in the studies presented in this thesis. Mostly two type of jet definitions were used: the EMTopo [124], and the PFlow jets [125]. Both are reconstructed with the *anti*-k_T algorithm [126] using a small R parameter, but they are made of different constituent types. The following paragraphs review the jet algorithms, constituents, and calibration procedures. Pile-up suppression is briefly discussed at the end of this section, while jet flavour tagging is addressed more thoroughly in Chapter 4.

Jet algorithms

Most of the jet definitions attempt to encompass a few important aspects. They should be simple, in line with the experimental limitations, and should be defined at any order of α_s expansion, yielding finite cross-sections that are insensitive to hadronisation [127]. Two key properties are *infrared* and *collinear safety*, which require that the reconstruction of jets in a collision event remains unchanged by QCD emissions driven by Equation 3.8. These effects are related to the hadronisation process and are difficult to precisely determine.

Sequential recombination algorithms provide a bottom-up approach to building jets, by iteratively combining particles. They are infrared and collinear safe, and can be used for precise QCD calculations, satisfying the above-mentioned properties. In this class of algorithms, the combination of constituents is based on a metric that depends on the pair angular distance $\Delta R_{ij} = \sqrt{\Delta y_{ij}^2 + \Delta\phi_{ij}^2}$, and the transverse momentum:

$$d_{ij} = \min\left(p_{T,i}^{2t}, p_{T,j}^{2t}\right) \frac{\Delta R_{ij}^2}{R^2}\,, \quad d_{iB} = p_{T,i}^{2t}\,, \tag{3.9}$$

where R and t are two tunable parameters.

Using d_{ij} and d_{iB}, it is possible to define three of the most commonly used jet algorithms at hadron colliders, which follow the same procedure but differ in the value of t. Constituents are combined according to the following steps:

1. Find the minimum values of d_{ij} and d_{iB} in the event.
2. If $d_{ij} < d_{iB}$, combine the i and j into a new single constituent k, remove i and j, and go back to step 1.
3. If $d_{iB} < d_{ij}$, define i as a jet, and return to step 1 if other particles remain.

The standard combination of particles involves summing their four-momenta. If $t = 1$, this procedure is known as the *inclusive* k_T algorithm [128]. It starts by grouping softer and collinear elements that grow into energetic jets. When $t = 0$, only the angular hierarchy between particles matters for the clustering, which is known as the *Cambridge-Aachen* algorithm [129]. Finally, $t = -1$ gives the anti-k_T algorithm [126], where jets grow around energetic seeds. In contrast with the inclusive version that produces irregular shapes, this last algorithm forms cone-shaped jets in the η-ϕ plane with a radius determined by the R parameter. This feature makes it well adapted to experimental conditions.

Figure 3.3 shows jets formed by the three sequential recombination algorithms discussed, with $R = 1$. In ATLAS, anti-k_T jets with such radius are referred to as *large-R* jets, while $R = 0.4$ is employed for the *small-R* ones.

PFlow reconstruction

Since Run 1, ATLAS has used topo-clusters at the electromagnetic scale to reconstruct jets, also called EMTopo jets. The four-momentum resolution for these jets depends on the energy resolution of the calorimeter, described by Equation 2.7, which degrades at lower energies. Moreover, soft charged particles are deflected by the magnetic field and may produce energy deposits in the calorimeter that are displaced with respect to the originating parton's direction. This energy fraction may not be clustered within the jet. Another important aspect affecting the quality of jets is pile-up. In-time pile-up interactions can produce hadrons that overlap in the calorimeter with the signal of interest, potentially leading to a misestimation of jet kinematics. To address these issues and improve jet reconstruction, the Particle Flow (PFlow) algorithm was introduced during Run 2 [125].

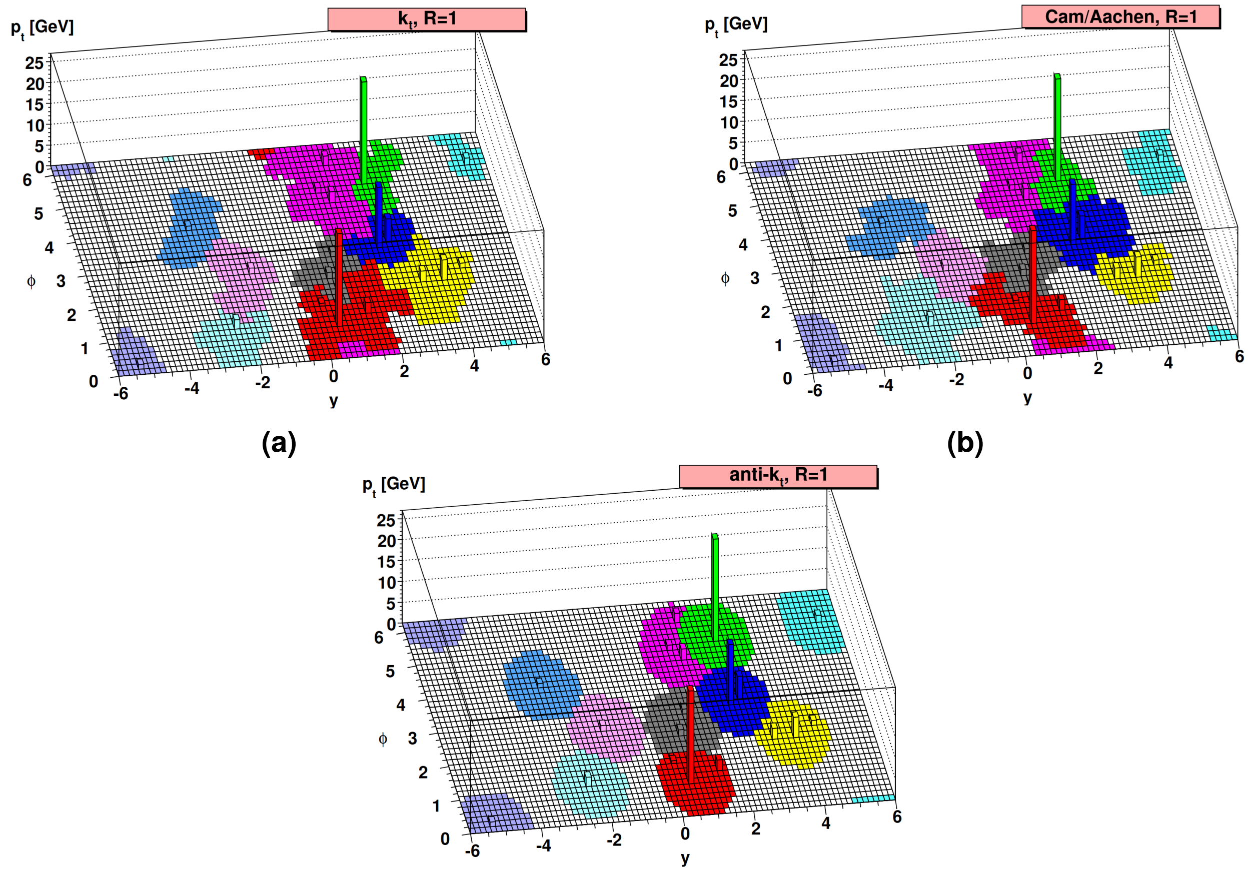

(a) (b) (c)

Figure 3.3: Jet reconstruction example with (a) inclusive k_T, (b) Cambridge-Aachen, and (c) anti-k_T algorithm [123].

The core idea of PFlow is to combine information from the ID tracker and the calorimeters. The former provides better transverse momentum resolution for soft particles, while for the latter resolution improves at high energies. Furthermore, reconstructing the trajectories of charged particles allows the recovery of the deflected ones and the verification of whether their origin is compatible with a pile-up vertex. Therefore, the main goal of PFlow is to reconstruct the "flow" of charged particles in the detector by matching tracks with the energy collected in the calorimeter.

Figure 3.4 illustrates the sequence of the PFlow algorithm. It starts by selecting tracks based on quality criteria and excluding those with $p_T^{\text{track}} > 45\,\text{GeV}$, where no advantage is gained from the tracker. Pile-up contributions are suppressed by filtering out tracks

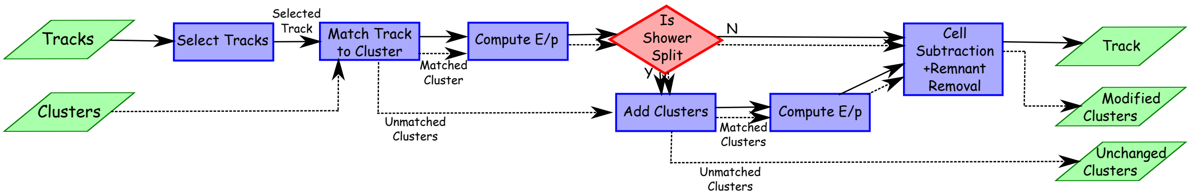

Figure 3.4: Schematics of the PFlow algorithm steps [125].

that are not compatible with the PV, *i.e.*, those with large longitudinal impact parameters. Then, tracks are matched to topo-clusters taking into account the spatial resolution of the calorimeter. The ratio of $\mathrm{E}_{\mathrm{clus}}^{\mathrm{EM}}/\mathrm{p}_{\mathrm{T}}^{\mathrm{track}}$ is computed to verify the correctness of the matching or to determine if an unrelated cluster is being considered. The energy-momentum ratio is also used to estimate if the particles' energy is expected to be spread in more than a single cluster. If so, additional compatible topo-clusters are searched for. Once the matching step is completed, the expected energy deposits in the calorimeter are extrapolated from track kinematics. The estimated energy is then subtracted from the calorimeter cluster, at the cell level. If the remaining energy is consistent with fluctuations of the calorimeter response, it is removed; otherwise, it is kept as a contribution from multiple particles overlapping in the same cluster.

The resulting output of the PFLOW sequence consists of the selected tracks, the energy-subtracted topo-clusters, and the clusters for which no matching track was found. These objects are then clustered using the anti-k_{T} algorithm to form the PFLOW jets.

Jet calibration

As previously mentioned, calorimeter jets are designed to represent the products of hard-scatter parton showers. However, many experimental effects impact the raw energy of experimentally reconstructed jets. For physics measurements, precisely determining the Jet Energy Scale (JES) through a calibration procedure for hadronic showers and measuring the relative Jet Energy Resolution (JER) are essential.

The JES calibration accounts for the calorimeter non-compensation and the presence of inactive material where deposited energy is not measured. It corrects for energy leaking outside the calorimeters or simply not contained within the jet cones. Furthermore, the calibration mitigates pile-up effects in jet reconstruction.

In ATLAS, a robust and detailed procedure is performed to calibrate jets [130–132] to the energy scale of *truth* jets – jets reconstructed in simulations using stable particles

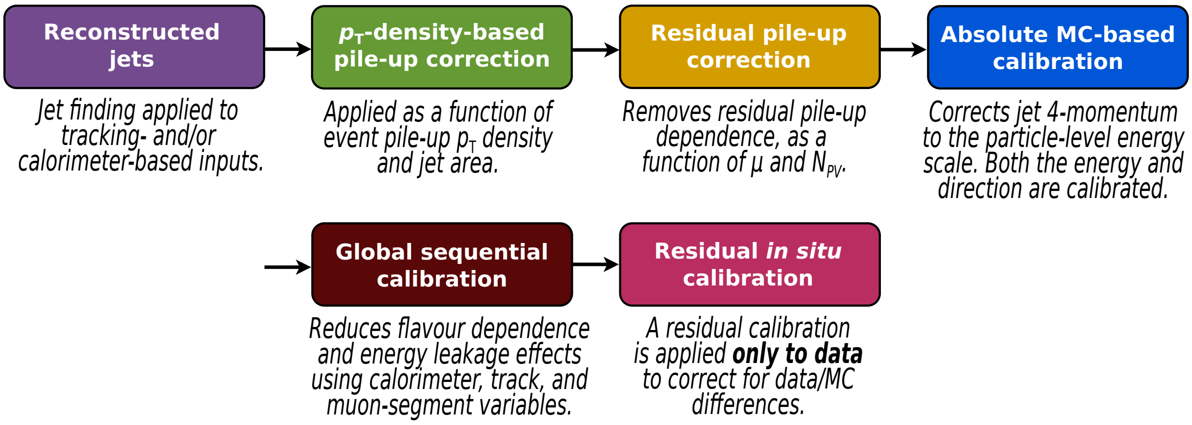

Figure 3.5: Schematics of the JES calibration sequence [130].

emerging from the hard scatter. Figure 3.5 shows a diagram of the calibration sequence applied to jets.

The first two calibration steps address pile-up corrections. The p_T density from pile-up jets is calculated in MC as a function of the p-p interactions per event (μ) and the number of reconstructed primary vertices (N_{PV}) per bunch-crossing, and is then subtracted from jets. Residual dependencies on μ and N_{PV} are then removed.

To perform the absolute JES calibration, jets are matched to truth jets. The jet response, defined as the average ratio of the energy of experimental (E^{reco}) and truth (E^{truth}) jets:

$$\mathcal{R}_E = \left\langle \frac{E^{\text{reco}}}{E^{\text{truth}}} \right\rangle , \tag{3.10}$$

is measured in bins of E^{truth} and detector η. The inverses of such measurements are taken as the JES calibration factors. A similar strategy is used to adjust the jets' η direction. These calibrations correct the non-compensating calorimeter response, out-of-cone radiation, and biases in the direction reconstruction.

The Global Sequential Calibration step uses jet variables, such as the number of tracks or the number of muon track segments associated with the jet, to factor out the dependence of $\mathcal{R}_E$ on the jet's physical properties. For example, quark-initiated jets concentrate their energy in fewer hadrons than gluon jets, or some high-p_T elongated showers can extend beyond the calorimeter, leaving signals in the muon spectrometer (punch-through effect). Correcting for these effects leads to a significant improvement of the JER.

The final calibration stage is the *in situ* correction, applied only to real data. This correction compensates the discrepancies between data and MC arising from the limitations of the simulation. It is evaluated by measuring the p_T balance relative to well-calibrated objects. In dijet events, central jets ($|\eta| < 0.8$) are used as reference to correct the response as a function of η for forward jets ($0.8 < |\eta| < 4.5$). Hadron recoils against well-measured photons or Z bosons are used to calibrate the jet's p_T. At very high p_T ($p_T \gtrsim 500$ GeV), the *in situ* correction is measured in multijet events, where hard jets are balanced by a system of calibrated low-p_T jets.

The level of accuracy achieved with the JES calibration is quantified by a broad set of uncertainties. They are defined to cover all systematic effects and statistical limitations of the calibration methods. Their impact on JES is illustrated in Figure 3.6. At very low p_T pile-up effects are dominant, while from 30 GeV to 400 GeV the largest uncertainty arises from response to different hadron flavours. Above 400 GeV the *in situ* corrections have the highest uncertainty, mostly due to limited statistics, which rapidly increase above 2 TeV, where no measurement is performed, and an extrapolated correction is applied instead.

The determination of JER obtained after calibration is essential for both SM precision measurements and BSM searches. It is parametrised using the same approach as the energy resolution of the calorimeters (Equation 2.7), with stochastic, noise and constant

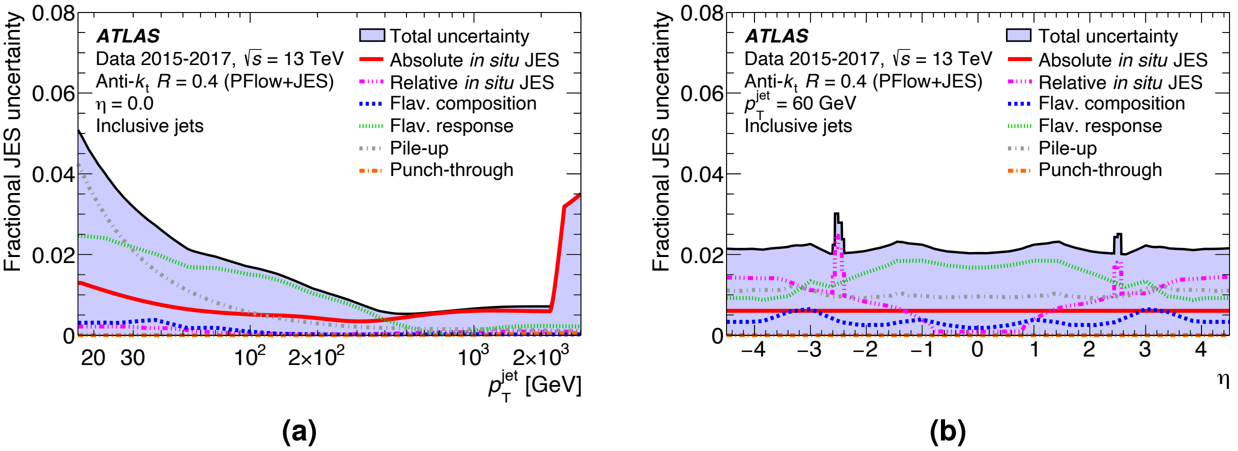

Figure 3.6: Relative JES uncertainties as a function of (a) jet's $\mathrm{p_T}$ and (b) η, for small-R PFLOW jets. The total uncertainty is determined as the quadrature sum of all the components [130].

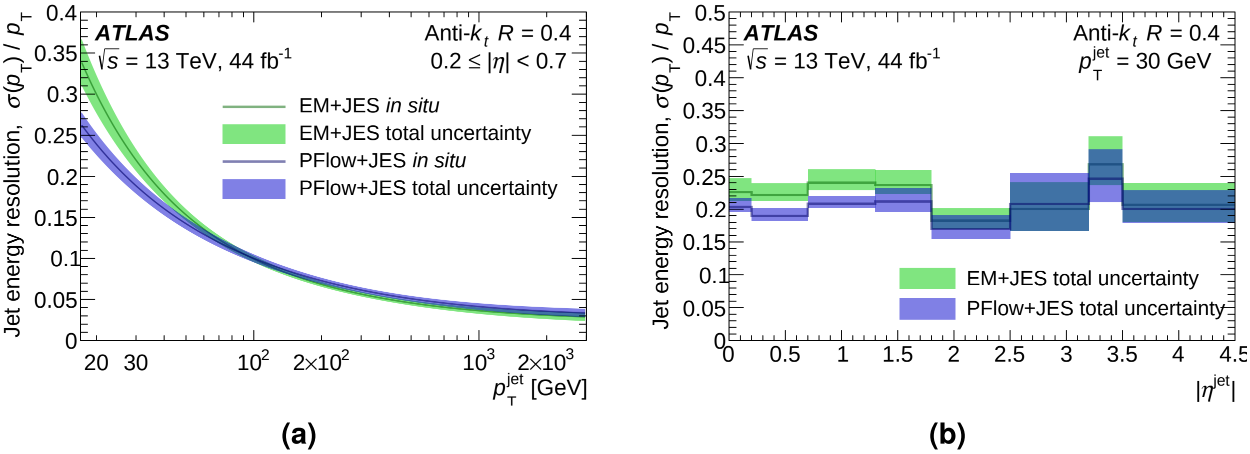

Figure 3.7: Relative JER for fully calibrated jets as a function of (a) jet's $\mathrm{p_T}$ and (b) η, for small-R PFLOW and EMTOPO jets [130].

terms. The total JER is measured in dijet events. At low $\mathrm{p_T}$, the noise contribution from pile-up dominates. However, it is constrained by measurements performed with the *random cones* method, in events selected by unbiased random triggers [130]. The overall JER obtained for fully calibrated jets is shown in Figure 3.7. As expected, the PFLOW reconstruction enhances the resolution at low $\mathrm{p_T}$ compared to EMTOPO jets.

Pile-up suppression

Pile-up suppression is crucial for accurately reconstructing the events in physics analyses. Regarding in-time pileup, “fake jets” can arise from two sources: QCD jets from pile-up vertices, or *stochastic* jets reconstructed from the random overlap of particles with different origins. The latter has a much steeper $\mathrm{p_T}$ spectrum and is therefore rarer at high $\mathrm{p_T}$ values.

The pile-up subtraction method applied in the jet calibration sequence has a strong impact on soft pile-up jets. It shifts their $\mathrm{p_T}$, bringing it below 20 GeV – the standard ATLAS jet reconstruction threshold – and effectively rejecting them.

Track-based information associated with the jets is also highly effective in pile-up mitigation. A standard variable used since Run 1 is the Jet Vertex Fraction (JVF) [133], defined as the scalar sum of p_T^{track} for tracks from the hard-scatter divided by the sum over all the tracks associated with the jet:

$$\text{JVF} = \frac{\sum_{\text{track}\in\text{jet}} p_T^{\text{track}}(PV^{HS})}{\sum_{\text{track}\in\text{jet}} p_T^{\text{track}}} . \tag{3.11}$$

For PFLOW jets, this information is inherent to the reconstruction method, hence, the fake-jet rate is reduced without any JVF requirement, as shown in Figure 3.8(a).

The expression in Equation 3.11 depends on the $\langle\mu\rangle$, as pile-up increases the denominator grows, shifting the JVF. The improved "corrJVF" variable introduces an additional term in the denominator to compensate for events with a higher number of pile-up tracks [134]. The corrJVF together with fraction of PV^{HS} tracks p_T associated to the jets

$$R_{p_T} = \frac{\sum_{\text{track}\in\text{jet}} p_T^{\text{track}}(PV^{HS})}{p_T^{\text{jet}}} , \tag{3.12}$$

are combined in a multivariate discriminant known as the Jet Vertex Tagger (JVT) [134]. This two-dimensional likelihood-based tagger successfully combines the two methods, increasing the rejection of pile-up jet, as demonstrated in Figure 3.8(b). A cut on JVT can provide a signal efficiency of 95% while accepting only 3% of pile-up jets. The JVT efficiency for hard-scatter jets is measured using data from $Z(\rightarrow \mu\mu)+$jets events.

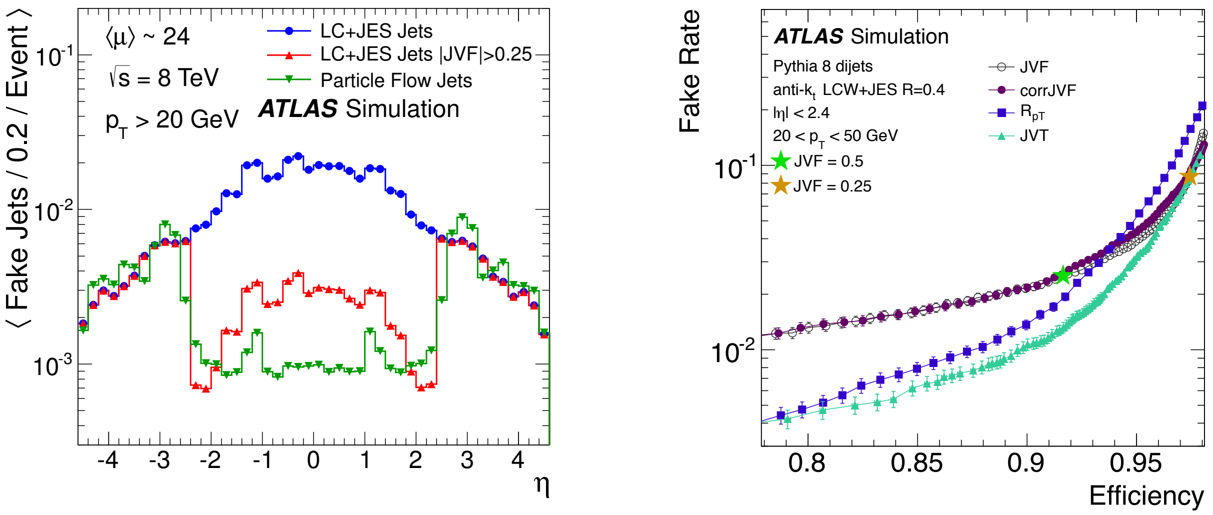

(a) **(b)**

Figure 3.8: Figure (a): jet reconstruction fake rate as a function of η [125]. Figure (b): pile-up jet fake rate versus hard-scatter jet efficiency for JVF, corrJVF, R_{p_T}, and JVT [134].

For jets beyond the ID tracker acceptance ($|\eta| > 2.5$), this approach cannot be adopted due to the lack of tracks. For hard jets with $p_T > 60$ GeV, the probability of originating from pile-up vertices becomes negligible.

3.4 Other signatures

Electrons and photons

Since Run 2, the reconstruction of e/γ signatures has used a combination of tracks and variable-sized *superclusters* [135, 136]. The dynamic shape of these clusters is better suited to recovering energy scattered from electron bremsstrahlung or photon conversions into electron-positron pairs, compared to the fixed-sized window approach.

The superclusters are formed using topo-clusters, considering only the fraction of energy deposited in the electromagnetic calorimeter cells. This fraction must be greater than half of the cluster energy and correspond to more than 400 MeV. Regional tracking is used, seeded by LAr energy deposits, with higher tolerance for the energy loss per detector layer compared to standard reconstruction. Tracks are used to identify photon conversions by locating secondary vertices compatible with massless particles, or to trace the trajectories of prompt electrons. Topo-clusters with energy above 1 GeV (1.5 GeV) and at least one (or zero) matched track(s) are used as seeds and grow into superclusters for e (γ) by adding the satellite topo-clusters within a window of 0.075×0.125 in η-ϕ around the seed barycentre, as shown in Figure 3.9. For electrons, additional satellite clusters are included in an area of 0.125×0.3 if their best-matched track coincides with that of the supercluster.

The e and γ supercluster building procedures are run independently. Therefore, to define analysis electrons and photons, superclusters compatible with both reconstructions must pass an ambiguity resolution algorithm, which primarily ranks the objects based

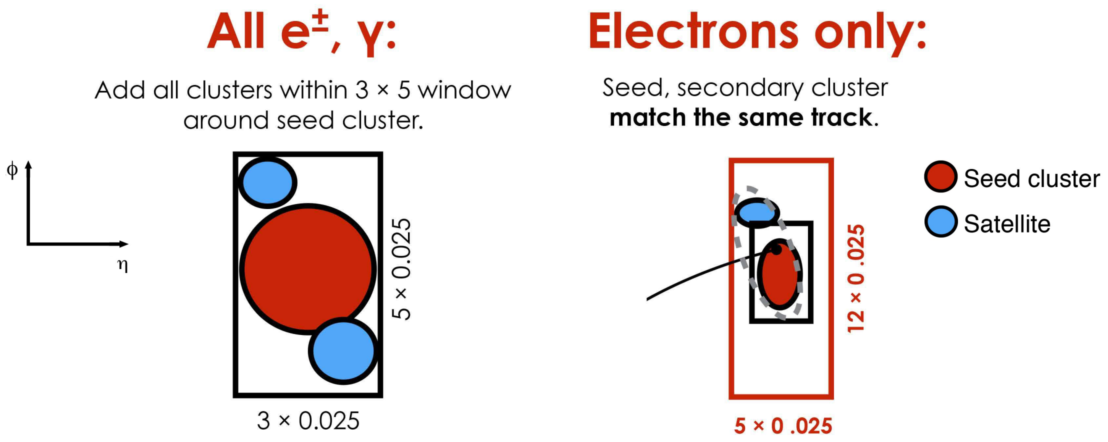

Figure 3.9: Summary of the supercluster formation for electrons and photons reconstruction [135].

on their associated tracks. Furthermore, to enhance purity and to reject deposits from hadronic jets or non-prompt e/γ, identification and isolation criteria are defined [121, 137, 138]. The identification of electrons has two standard working points (*loose* and *tight*) and is based on a likelihood approach that considers both the electromagnetic shower shape and the tracks matched to the superclusters. For photons, only rectangular cuts on the shower shapes are applied. The isolation criteria serve to quantify the activity around the e/γ candidate, and can be calculated based on calorimeter and tracker information.

To optimise the determination of the energy corresponding to electromagnetic showers, a calibration is applied to reconstructed photons and electrons [139]. The calibration also accounts for inconsistencies in the LAr calorimeter response and non-linearities due to the digitisation process. The calibration is derived using $Z \to ee$ events and validated by measuring the invariant masses of $Z \to \ell\ell\gamma$ and $J/\psi \to ee$, as shown in Figure 3.10.

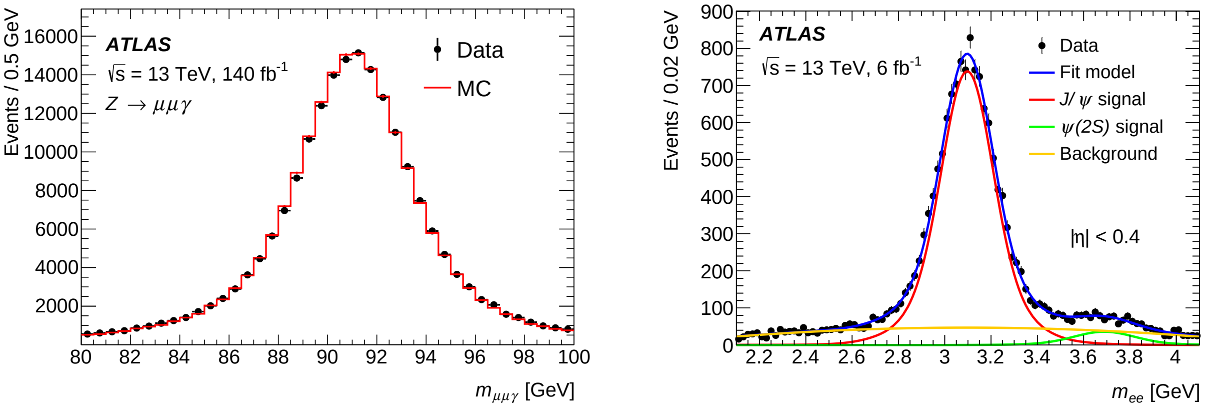

(a) **(b)**

Figure 3.10: Reconstruction of the (a) $Z \to \mu\mu\gamma$ and (b) $J/\psi \to ee$ invariant masses for the validation of the e/γ calibration procedure [139].

Muons

The passage of muons through ATLAS leaves the signature of a minimum-ionising particle across all detector systems. MS and ID are predominantly used to determine muons' trajectories, while the calorimeter information is considered to evaluate the energy loss Incurred when travelling from the ID to the MS, or to tag muons emerging from the ID independently of the MS.

Five main strategies are used in ATLAS to reconstruct muons [140, 141]. The *combined* muons are obtained by performing outside-in tracking seeded by the MS, down to the ID, and taking into account energy losses in the calorimeters during the fitting procedure. The *inside-out* muons are reconstructed by performing tracking the opposite direction, and can recover some inefficiencies due to the MS seeding. The *standalone* or *MS-extrapolated* muons are built using only muon spectrometer tracking, with the trajectory extrapolated inwards to the beam line, extending the geometric acceptance beyond the ID ($|\eta| > 2.5$). The *segment-tagged* muons use only ID information to calculate the track parameters, but require a strict angular match with detected MS segments. Finally, the *calorimeter tagged* muons are reconstructed using ID and calorimeter deposits traceable to a minimum-ionising particle.

As with electrons and photons, muon identification criteria are defined to achieve different levels of purity. Five standard working points, labelled as *Loose*, *Medium*, *Tight*, *High*-p_T, and *Low*-p_T [141], are defined. They are characterised using a few quality features: the agreement of track p_T and q/p between the estimations from ID and MS done separately, and the number of muon precision stations contributing to the track reconstruction. The various working points provide increasing levels of purity at the cost of small losses in efficiency, as shown in Figure 3.11(a). Furthermore, calorimeter isolation is used to discriminate prompt muons from those originating in hadron decays.

The muon momentum in collected data is calibrated to correct for detector misalignment [142], which leads to a charge-dependent bias. Moreover, inconsistencies in the simulations are corrected to ensure that muon performance in MC matches the observed data. After the momentum calibration, the mass resolution for dimuons is below 1.5% for soft muons and increases up to 3.4% for larger p_T, as shown in Figure 3.11(b), with a good agreement between data and MC.

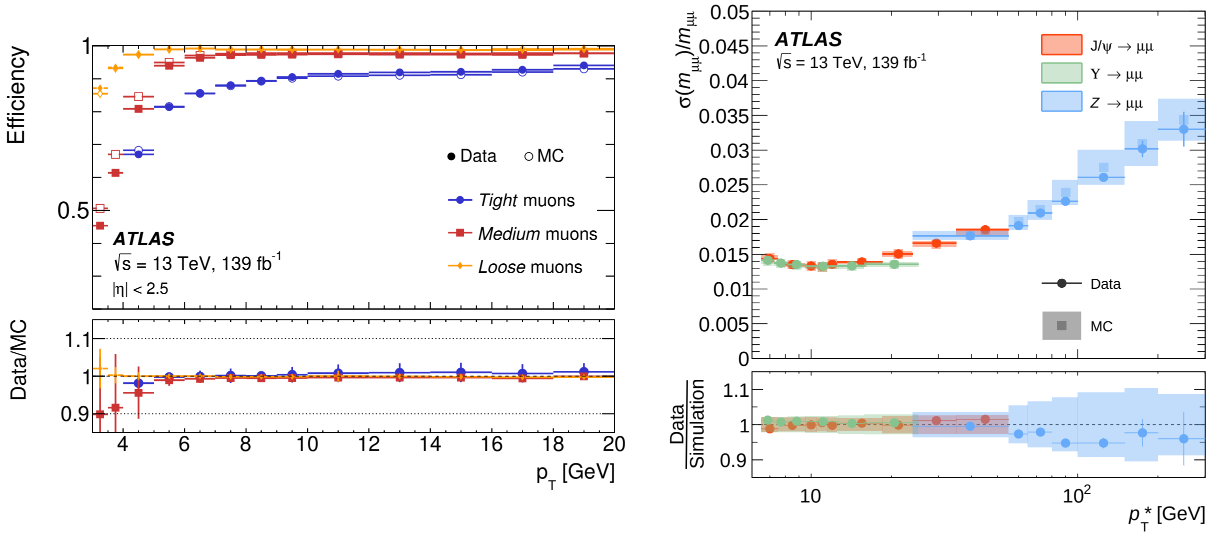

(a) **(b)**

Figure 3.11: Measurement of (a) muon reconstruction efficiency in $Z \to \mu\mu$ events [141], and (b) invariant mass resolution in $J/\psi \to \mu\mu$, $\Upsilon \to \mu\mu$, and $Z \to \mu\mu$ as a function of p_T^* (a measure of the transverse momentum independent of the mass resolution) [142].

τ-leptons

The τ-leptons have a finite lifetime, decaying with a displacement from the hard-scatter proportional to $c\tau_{\text{lifetime}} = 87.03\ \mu\text{m}$ [30], which can be resolved experimentally. For leptonic decays, corresponding to almost 35% of the τ branching ratio, they are reconstructed as muons or electrons. In the remaining 65% of cases, τ decays into hadrons, producing a signature similar to QCD jets.

In ATLAS, the visible part of hadronic τ ($\tau_{\text{had-vis}}$) is reconstructed starting from small-R anti-k_T jets [143, 144], with topo-clusters as constituents. Due to their displaced decays, they require a dedicated vertex association algorithm [145]. Moreover, $\tau_{\text{had-vis}}$ contains either one or three charged pions (or kaons with much lower probability), along with the possible presence of one or more π^0 (or K^0). To correctly reconstruct the charged prongs, a special track association and classification procedure is performed. The $\tau_{\text{had-vis}}$ candidates are then identified using multivariate tools and calibrated with a multi-step procedure [144].

Missing transverse momentum

Missing transverse momentum (p_T^{miss} or MET) is an experimental proxy used to reconstruct undetected particles produced in p-p collisions, such as neutrinos or potential DM candidates. Hence, it is a pivotal variable for both SM measurements and BSM searches. It is calculated by assuming momentum conservation and considering that, before the collisions, by definition the net momentum in the beam's transverse plane is zero.

In ATLAS, the vector of missing transverse momentum ($\mathbf{p}_T^{\text{miss}}$) is reconstructed by performing the vectorial sum over the 2D momenta in the x-y plane of the objects emerging from the hard-scatter process [146, 147]. It is composed of two classes of objects:

$$\mathbf{p}_T^{\text{miss}} = -\left(\sum \mathbf{p}_T^{\text{hard}} + \sum \mathbf{p}_T^{\text{soft}}\right) , \tag{3.13}$$

where the *hard* contribution includes all the calibrated physics signatures (γ, e, μ, $\tau_{\text{had-vis}}$, and jets), while the soft part consists of reconstructed tracks associated with the PV^{HS} that are left unmatched by the hard objects. During Run 1 and early Run 2 calorimeter deposits were also considered to account for neutrals in the soft part [146, 148], but they are strongly dependent on pile-up conditions, and therefore, no longer used. Since physics objects are reconstructed independently, a robust overlap-removal procedure is required to avoid double-counting.

From Equation 3.13, the p_T^{miss} and its direction are reconstructed as:

$$p_T^{\text{miss}} = |\mathbf{p}_T^{\text{miss}}| = \sqrt{(p_x^{\text{miss}})^2 + (p_y^{\text{miss}})^2} , \tag{3.14}$$

$$\phi^{\text{miss}} = \tan^{-1}(p_y^{\text{miss}}/p_x^{\text{miss}}) . \tag{3.15}$$

The reconstruction quality of these observables is limited by the detectors' resolution and acceptance, mis-measurements and mis-calibrations of physics objects, and the pile-up.

The strict definition of objects utilised in different analyses may vary. Therefore, for analysis consistency, different working points of $p_{\mathrm{T}}^{\mathrm{miss}}$ are defined since Run 2 [147]. The modelling and the performance of $p_{\mathrm{T}}^{\mathrm{miss}}$ are measured in both MC and data. The amount of mis-reconstruction is evaluated in $Z \to \ell\ell$ events, where no real $p_{\mathrm{T}}^{\mathrm{miss}}$ is expected. In addition to Z decays, the resolution and the response are studied using semi-leptonic $t\bar{t}$ and $W \to \mu\nu$ events, as shown in Figure 3.12. In the latter, it is possible to see how pile-up induces a resolution degradation, and how drastically it affects the response at low $p_{\mathrm{T}}^{\mathrm{miss}}$, below 50 GeV.

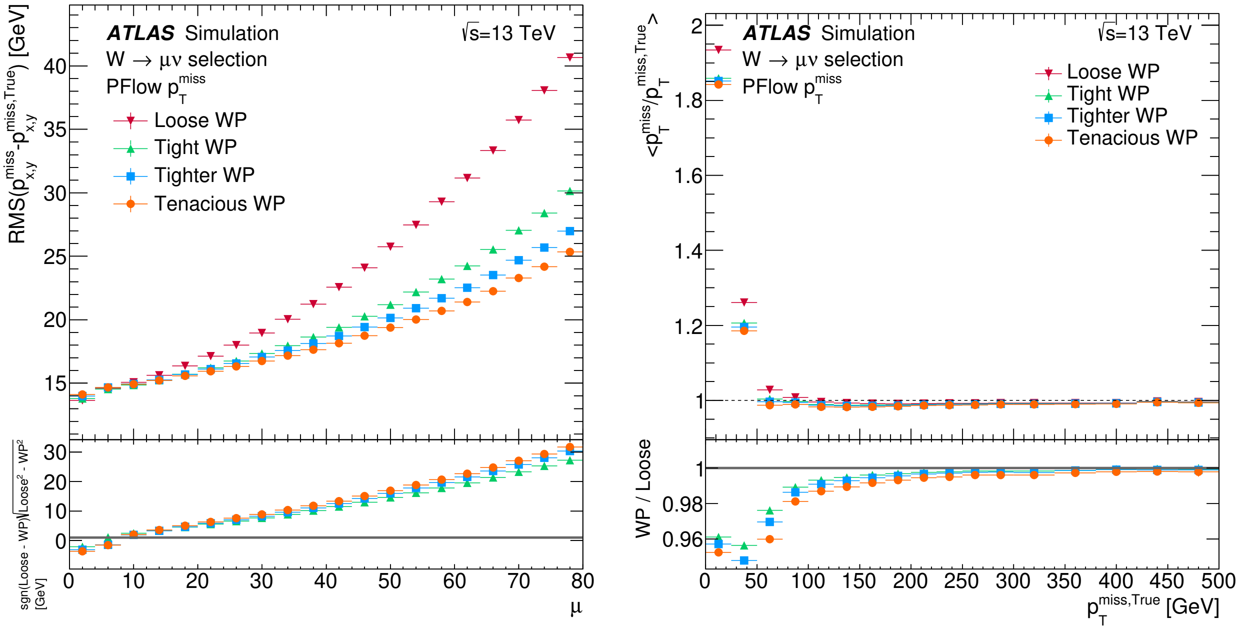

(a) **(b)**

Figure 3.12: Measurements for different working points of the $p_{\mathrm{T}}^{\mathrm{miss}}$ (a) resolution as function of pile-up (μ), and (b) response versus the true value of $p_{\mathrm{T}}^{\mathrm{miss}}$ in $W \to \mu\nu$ simulated events [147].

II

Triggering beauty jets

Chapter 4

Beauty hadrons in jets

A broad range of compelling physics processes comprise the presence beauty (or bottom) quarks. From the top quark and Higgs boson to many BSM physics models, bottom quarks are produced abundantly in decays of heavy particles. Furthermore, the properties of $\mathtt{b}$-hadrons – which count beauty quarks among its constituents – are tightly bound to many open issues in the HEP field. Hence, beauty hadrons are exceptionally powerful experimental probes.

The heavy mass and the relatively long lifetime of $\mathtt{b}$-hadrons lead to remarkable signatures in particle detectors at collider experiments. Specialised reconstruction methods were developed to identify hadronic jets containing $\mathtt{b}$-hadrons. These techniques, going under the name of $\mathtt{b}$-tagging or jet flavour tagging, heavily rely on measuring the displaced decay of such hadrons. In ATLAS, highly engaged R&D activities have profoundly improved the performances $\mathtt{b}$-tagging algorithms over the past few years.

In this chapter an overview of the physical aspects of $\mathtt{b}$-hadrons is given, including their ties with frontier HEP. The second part of the chapter reviews the jet flavour tagging methodologies in ATLAS, including some of its most recent developments.

4.1 Physics of b-hadrons

The physics of bottom quarks is extremely rich and plays a special role in HEP. Their mass is around 4.18 GeV [30], much larger than the other light quarks and Λ_{QCD} (~ 200 MeV, the confinement scale discussed in Section 1.1). This allows for precise calculations using the perturbative QCD approach or adopting a dedicated effective theory [149]. Furthermore, b-hadrons are the heaviest known hadrons. Although the top quark is heavier, it decays faster ($\sim 10^{-25}$ s) than the time required to form hadrons [150].

The bottom quark belongs to the same weak isospin doublet of top quark (as mentioned in Section 1.1), with which it shares the strongest coupling, but the top is significantly heavier. Therefore, the kinematically allowed decays of b-hadrons are suppressed by off-diagonal elements of the CKM matrix (Equation 1.6). As a result, b-hadrons have long lifetimes (~ 1.5 ps), leading to displaced decays at the LHC and other colliders, which produce secondary vertices detached from the beam spot. For reference, the lifetimes of some mesons and baryons are reported in Table 4.1, together with their corresponding decay lengths. Another relevant property of b-hadrons is their considerable branching ratios for semi-leptonic decays. These decays offer relatively clean experimental signatures for studying b-physics.

Bottom quarks are involved in many interesting processes. They provide experimental grounds for precision tests of the SM, for measurements of CKM matrix elements, and studies of CP violation. They are also excellent probes for BSM physics.

Precise measurements of B^0-$\overline{\mathrm{B}}^0$ hadron system oscillations set tight constraints on SM physics [151]. The rare decays of b-hadrons involving flavour-changing neutral currents, such as $\mathrm{B}^0 \rightarrow \mathrm{K}^{0*}\mu^+\mu^-$, are sensitive to potential new physics at higher energy scales. They are forbidden at tree level in the SM, but can arise from loop diagrams involving

Table 4.1: Summary of the several b-hadrons properties [30].

	Quarks	Mass [MeV]	Lifetime [ps]	$c\tau$ [μm]
		Mesons		
B^+	$\mathrm{u\overline{b}}$	5279.41 ± 0.07	1.638 ± 0.004	491.1
B^0	$\mathrm{d\overline{b}}$	5279.72 ± 0.08	1.517 ± 0.004	454.8
B^0_s	$\mathrm{s\overline{b}}$	5366.93 ± 0.08	1.520 ± 0.005	455.7
B^+_c	$\mathrm{c\overline{b}}$	6274.47 ± 0.32	0.510 ± 0.009	152.9
		Baryons		
Λ^0_b	udb	5619.6 ± 0.2	1.471 ± 0.009	441.0
Ξ^-_b	dsb	5797.0 ± 0.6	1.572 ± 0.040	471.3
Ξ^0_b	usb	5791.9 ± 0.5	1.480 ± 0.030	443.7
Ω^-_b	ssb	6045.8 ± 0.8	$1.64^{+0.16}_{-0.17}$	491.7

virtual boson exchanges, which may receive contributions from BSM physics [152, 153]. The semi-leptonic decay channels can test the lepton flavour universality in weak interactions, in particular measuring ratios of the $B^0 \to D^{*+}\ell^-\nu_\ell$ branching ratios for different leptons [154].

ATLAS has an extensive physics program for measuring such b-hadron properties in "light" final states [155], but it is not the focus of this thesis. This chapter and the next two focus on the reconstruction of "heavy" final states, which form hadronic jets containing b-hadrons, namely b-jets. This signature is exceptionally useful in a wide spectrum of physics analyses.

Given the intensity of the coupling, top quark decays typically produce b-jets. Hence, in ATLAS, top quark reconstruction relies on b-tagging in most cases [156].

The SM Higgs boson primarily decays into b-quarks [75], as shown in Figure 4.1, with the $b\bar{b}$ channel accounting for more than 50% of the total branching ratio. Therefore, flavour tagging is an essential tool for di-Higgs or tri-Higgs boson searches [157, 158], which are needed to study the structure of the energy potential in the Higgs sector. Understanding the Higgs potential is crucial for gaining insights into the stability of the Universe and fundamental Cosmology [159, 160].

Many BSM particles are hypothesised to have strong couplings to b-quarks. SUSY searches for *stops*, *sbottoms*, charginos, and neutralinos regularly rely on b-tagging for event reconstruction [161–163]. New light scalar particles that mix with the Higgs boson would preferentially decay to $b\bar{b}$ pairs, if kinematically allowed [164]. The theoretically motivated vector-like quarks, hypothetical colour-charged particles, would predominantly mix with third-generation quarks [165], creating b-jets if produced at LHC. These few examples of new physics searches corroborate the importance of b-tagging in ATLAS.

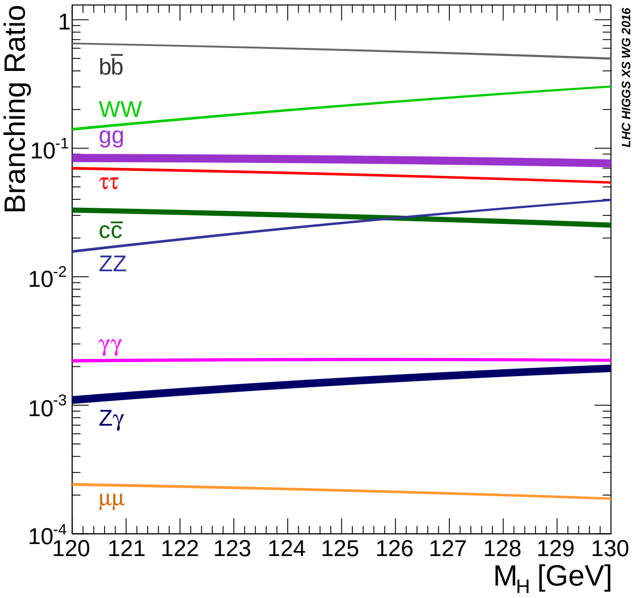

Figure 4.1: Higgs boson branching ratio as a function of its mass [75].

4.2 Jet flavour tagging in ATLAS

Jet flavour tagging provides an experimental methodology to distinguish jets originating from different partons. Commonly, beyond b-jets, two additional jet flavour classes are defined: c-jets and light-jets, which are initiated respectively by c-quarks or light partons – u, d, s quarks, or gluons. A wide range of algorithms have been developed within the ATLAS Collaboration to perform b-tagging. They rely on the identification of the peculiar traits of b-hadrons: their long lifetimes, large masses, and high decay multiplicities. Figure 4.2(a) shows the average flight length ($\beta\gamma c\tau$) of b-hadrons listed in Table 4.1 as a function of their energy. Above 20 GeV, most of them travel at least 1 mm away from the PV. Since the decay rate is exponential, for energies up to a few hundred GeV, decay occurs before reaching the IBL, as shown in Figure 4.2(b). This means that, for b-hadrons, a relatively large number of charged tracks can be reconstructed with numerous pixel hits, pointing to a vertex between the beam line and the IBL.

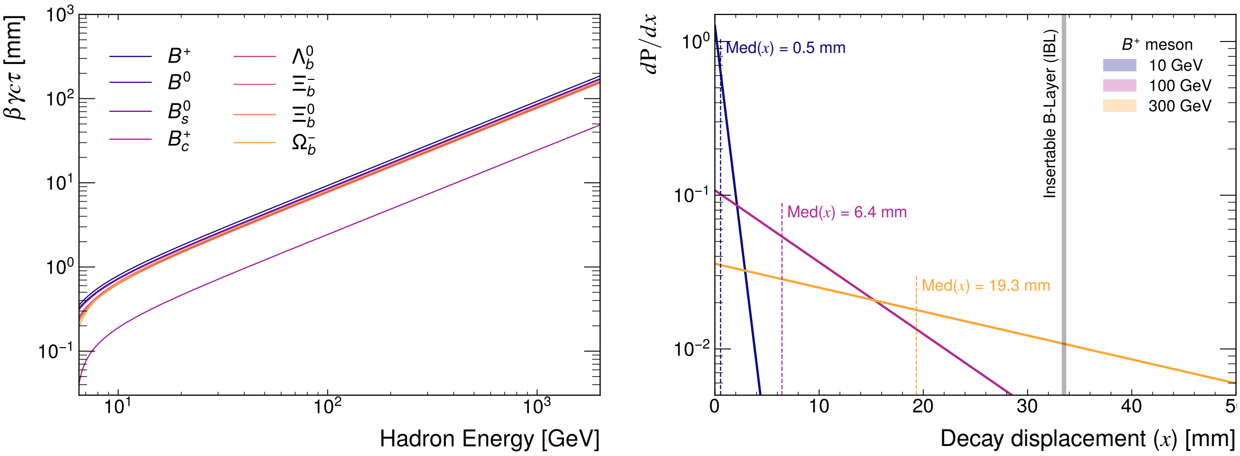

(a) **(b)**

Figure 4.2: The (a) average flight length of b-hadrons before decaying as a function of their energy, and the (b) probability distributions for the decay displacement of a B^+ meson at three different energies. In the latter plot also the median values of the decays are reported and the position of the ATLAS IBL is shown, assuming that the displacement is on the radial direction.

Since Run 1, flavour tagging in ATLAS has adopted a two-stage strategy to improve its performance. First, low-level variables, such as tracks, are used to reconstruct peculiar traits of b-hadron decays. These observables are then combined through multivariate tools to form high-level discriminants. However, recent developments have demonstrated that utilising modern Machine Learning (ML) techniques – see Appendix **A** – the best performance is achieved by skipping any further reconstruction beyond tracks and feeding them directly into large-scale ML models [166].

Two main reconstruction approaches are considered to identify b-jets in the two-stage taggers: the impact-parameter-based and the vertex-based algorithms [167–169]. The former approach exploits the fact that tracks originating with a displacement from the PV have large impact parameters. The second type of approach attempts to reconstruct the position and other physical properties of the secondary vertices corresponding to the decay of b-hadrons.

A third complementary method, known as the *soft muon tagger* [167, 170], has been developed in ATLAS. This technique uses soft muons emerging from jets. Owing to semi-leptonic and chain decays of b-hadrons, almost 20% of b-jets will contain a muon with p_T above a few GeV, while for c- and light-jets the probabilities are much lower, around 5% and 0.05% respectively. However, this methodology has limited applications due to its low intrinsic efficiency.

The following paragraphs briefly describe the elements used for b-tagging, followed by an overview of the ATLAS b-jet specific reconstruction and tagging algorithms.

Flavour tagging inputs

Jets considered for b-tagging are restricted to the geometrical acceptance of the ID tracker ($|\eta| < 2.5$), as both reconstruction approaches use ID tracks that are matched to the jets. The standard matching uses a shrinking cone in ΔR around the jet, defined by

$$\Delta R < 0.239 + \exp\left(-1.22 - \frac{0.0165}{\text{GeV}} \cdot p_T^{\text{jet}}\right) , \tag{4.1}$$

such that harder jets have tighter cones (for $20\,\text{GeV}$ $\Delta R = 0.45$, while $\Delta R = 0.25$ if p_T is $200\,\text{GeV}$). Tracks are required to pass several selections based on their kinematics and quality. Cuts on $|d_0|$ and $|z_0 \sin\theta|$ are applied to exclude tracks from pile-up vertices. Conditions on the number of hits in the Pixel and SCT detectors are imposed to reduce the fraction of fake tracks. Vertex-based algorithms additionally reject tracks that have an origin compatible with material interactions.

To develop the tagging algorithms, truth flavour labels for jets are required. These are assigned to simulated jets by matching them with truth hadrons produced in MC events. If a jet contains a b-hadron with $p_T > 5\,\text{GeV}$ within a cone of $\Delta R < 0.3$, it is labelled as b-jet. Otherwise, it is searched for c-hadrons and τ-leptons to label the jet accordingly. If none are found, the light-jet label is assigned.

Impact Parameter algorithms

For b-tagging, the track impact parameters are computed relative to the PV, rather than the beam spot. Furthermore, the impact parameters are signed based on the track's path relative to the jet direction. If a track crosses the jet axis before the PV, it has a positive sign; otherwise, it is negative. For prompt tracks, the intersection with the jet axis should ideally coincide with the PV, but it is randomised due to detector resolution. Instead, for tracks emerging from secondary vertices inside jets, the sign is mostly positive.

JetProb The first tagger implemented in ATLAS was JetProb [167, 171], adapted from the LEP experiments. It uses the signed impact parameter significance, $s_{d_0} = d_0/\sigma_{d_0}$, of tracks associated with the jet. A score is assigned to each track based on the expected resolution of s_{d_0} for light jets. Track scores are combined to derive the jet probabilities for each flavour class. This tagger is no longer used since early Run 3.

IP2D and IP3D The transverse and the longitudinal impact parameter significances (s_{d_0} and $s_{z_0} = z_0 \sin\theta/\sigma_{z_0 \sin\theta}$) together with a track categorical variable are used to build discriminant values with the IP2D and IP3D taggers [169, 172, 173]. The track categorisation is based on hit-related quality variables. The first tagger uses s_{d_0} and the track category to construct a two-dimensional likelihood template from MC simulations for estimating track weights (p_i^{track} with $i = b, c, l$). The IP3D also includes s_{z_0}, relying on a 3D template. The two algorithms assign flavour tagging scores to the jets by combining track weights with a log-likelihood ratio (LLR) method. For example, the jet discriminant for IP2D is given by:

$$\mathrm{IP2D}_{b\,\mathrm{vs}\,l} = \sum_{\text{track}\in\text{jet}} \log\left(\frac{p_b^{\text{track}}}{p_l^{\text{track}}}\right) . \tag{4.2}$$

RNNIP The IP2D and IP3D methods cannot account for correlation between tracks, as building multi-track likelihood templates becomes computationally impractical due to the high dimensionality of the problem. To capture the intrinsic correlation of impact parameters in tracks from b-hadron decays, the Recurrent Neural Network Impact Parameter (RNNIP) tagger was developed [174]. The recurrent neural network (RNN) architecture processes sequences of arbitrary lengths and learns correlations between the sequence elements. In the RNNIP tagger, the tracks associated with jets are sorted in a list by descending $|s_{d_0}|$ and passed to a *long short-term memory* (LSTM) cell [175], which preserves the correlations between distant tracks in the sequence. Once all the tracks are processed, the output of the LSTM cell is given to a feed-forward fully-connected Neural Network (NN) that generates the jet probabilities for each flavour class. The final tagging discriminant is obtained with:

$$D_b = \log\left(\frac{p_b}{p_c \cdot f_c + p_l(1 - f_c)}\right) , \tag{4.3}$$

with the *charm fraction* $0 \leq f_c \leq 1$ set according to the importance given to the c-jet background. Compared to IP3D, two additional track features are considered for RNNIP: the p_T fraction with respect to the jet p_T, and the ΔR relative to the jet axis.

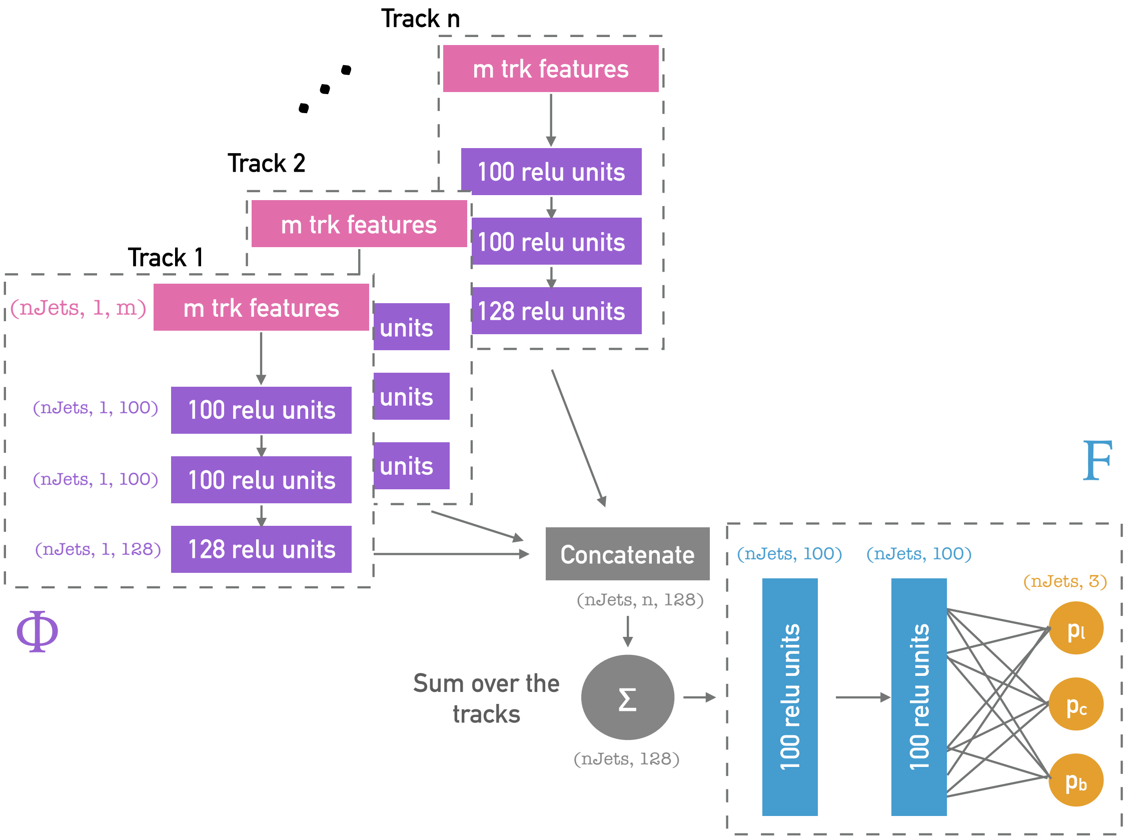

Figure 4.3: Diagram of the DIPS architecture [176].

DIPS A further improvement came with the Deep Impact Parameter Sets (DIPS) tagger [176]. It uses a *deep sets* architecture [177], in which jet tracks are treated as an unordered set of objects. Like RNNIP, this algorithm can handle an arbitrary number of tracks. Each track is processed independently by the same network, Φ, which produces a representation of the tracks in a high-dimensional latent space. These transformed tracks are aggregated by a vector sum and the result is passed to F, a fully connected NN. The output of the last layer of F consists of the posterior probabilities for each jet flavour. Figure 4.3 illustrates the DIPS architecture. For this tagger, a flavour discriminant is defined in the same way of Equation 4.3. One main advantage over RNNIP is that, due to the permutation-invariant concatenation of tracks, no artificial ordering is required. Furthermore, tracks do not have to be iteratively processed as for RNNs; instead, the Φ network can be parallelised, resulting in faster inference and shorter model training times. The input track features are similar to RNNIP, but instead of using the track category, hit-based variables are directly considered. The complete list of track variables used by impact-parameter-based algorithms is reported in Table 4.2.

Table 4.2: Summary of the track features used by the various impact parameter based taggers.

Input	**Description**	JetProb	IP2D	IP3D	RNNIP	DIPS
s_{d_0}	Transverse impact parameter significance (d_0/σ_{d_0})	×	×	×	×	×
s_{z_0}	Longitudinal impact parameter significance $(z_0 \sin\theta/\sigma_{z_0 \sin\theta})$			×	×	×
Category	Categorisation based on track quality features (mostly hits)		×	×	×	
p_T^{frac}	Fraction of the jet p_T carried by the track				×	×
ΔR(trk,jet)	Distance between track and jet axis in the η-ϕ plane				×	×
IBL hits	Number of hits in the IBL					×
PIX1 hits	Number of hits in next-to-innermost pixel layer					×
shared IBL hits	Number of shared hits in the IBL					×
split IBL hits	Number of split hits in the IBL					×
nPixHits	Number of total hits in the Pixel detector					×
shared PIX hits	Number of shared hits in the Pixel detector					×
split PIX hits	Number of split hits in the Pixel detector					×
nSCTHits	Number of total hits in the SCT detector					×
shared SCT hits	Number of shared hits in the SCT detector					×

Vertex-based algorithms

To improve b-tagging, vertex-based algorithms reconstruct the displaced decays produced by b-hadrons that travel away from the hard-scatter location. These algorithms use tracks with large impact parameters to form secondary vertices, starting from track pairs. Low-mass vertices that are compatible with light hadrons decays (K_s or Λ) or photon conversions are rejected, similarly to vertices located at detector layers, which may arise form material interactions. This is illustrated in Figure 4.4(a), where no cleaning is applied, and peaks corresponding to the ID layers are visible (at 33 mm, 50 mm, 88 mm, and 122 mm).

SV0, SV1 and SV2 The inclusive secondary vertices obtained from displaced tracks are used by the SV0, SV1, and SV2 taggers [167, 178]. The SV0 tagger uses only the 3D distance (L_{3D}) from the PV, divided by its uncertainty ($\sigma_{L_{3D}}$), to identify b-jets. Instead, SV1 and SV2 also consider the vertex invariant mass, calculated using all charged tracks, the ratio of the total energy of secondary vertex tracks to the jet p_T, and the number of two-track vertices found in the jet. Moreover, these two taggers define a likelihood-based discriminant, similar to the IP2D and IP3D taggers. For SV1, a 2D histogram is used for the likelihood template, while SV2 employs 3D likelihood, requiring larger MC samples to estimate the reference template. These vertex-based algorithms generally provide better background rejection compared to simple impact parameter methods (JetProb, IP2D, and IP3D), but have an intrinsic limitation in b-jet efficiency due to their reconstruction procedure.

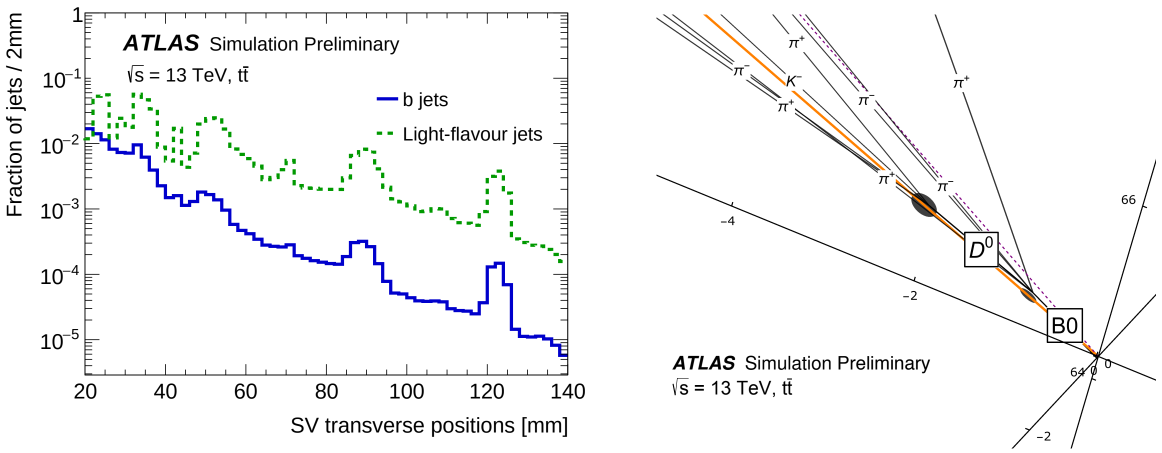

(a) **(b)**

Figure 4.4: The (a) transverse distance from the beam spot of secondary vertices before rejecting material interactions [178], and the (b) event topology for a simulated B^0 decay chain into D^0 and light hadrons [179].

JetFitter Instead of contructing inclusive secondary vertices, the JetFitter algorithm aims to trace back the entire topology of b-hadron decays [167, 179]. It assumes that sequential decays of b-hadrons into c-hadrons and subsequently into light-hadrons are aligned along the trajectory of the initial b-hadron, which is justified by their hard fragmentation, as shown in Figure 4.4(b). The b-hadron flight path is reconstructed in three dimensions using a Kalman filter, and vertices are identified at the intersections of charged tracks with this path, also considering single-track vertices. Candidate vertices and their associated tracks are then fitted to determine their exact positions and the respective uncertainties. Various features of the reconstructed topology are used to distinguish b-jets. The mass, energy fraction, and distance from the PV are calculated for secondary and, if present, tertiary vertices. The total number of vertices and tracks associated with them are also determined. During Run 1, all these variables were used as input for a small NN that provided jet flavour probabilities as output [167].

Combined high-level taggers

To enhance the performance of flavour tagging, algorithms from the two different reconstruction approaches – impact parameter and vertex-based – are combined into high-level taggers. The simplest combination is performed by a simple sum the low-level tagger scores, as was done for the IP3D+JetFitter algorithm [167]. A more powerful method involves using ML tools that can capture correlations among the various discriminants.

MV1 and MV2 The IP3D, SV1, and JetFitter discriminants were combined in a small NN known as MV1, which was widely used in Run 1 [167]. It has three input features, two hidden layers with three and two nodes, respectively, and a single output node used as the b-tagging discriminant. An alternative method was implemented with the MV2 tagger [172], which uses a Boosted Decision Tree (BDT) [180]. Many implementations of MV2 exist incorporating different combinations of input variables, and some include also a soft muon tagger.

DL1 series The usage of significantly deeper and larger networks than MV1 led to substantial improvements over the BDT approach. The first algorithm of this kind introduced in ATLAS was DL1 [172]. It takes as input the IP3D, SV1, and JetFitter reconstructed variables, along with the jet p_{T} and η. The development of RNNIP led to DL1r [169], which includes RNNIP posterior probabilities as input and has eight hidden layers with nearly $60\,000$ trainable parameters. A further improvement was achieved by replacing the RNN scores with a DIPS tagger, resulting in the DL1d algorithm, which was used in the HLT [3] during the 2022 data-taking period. All the DL1 taggers output three probabilities for jet flavours and combine them into a single discriminant, as done in Equation 4.3.

Unified taggers

In recent years, the adoption of state-of-the-art ML techniques has brought a paradigm shift in flavour tagging. The two-stage approach was surpassed by the implementation of the GN1 tagger [166]. A graph-based architecture [181] was employed to classify jets directly from track features, bypassing any intermediate low-level reconstruction. To guide the network in learning physically meaningful jet representations and improving the main classification objective, auxiliary tasks are assigned to the tagger. These tasks are designed to teach the model important features of jet structures and eliminate the need for low-level algorithms. The auxiliary training objectives include categorising tracks based on their physical origin (*e.g.*, whether a track is associated with pile-up, material interactions, a reconstruction artefact, or a given hadron) and grouping tracks emerging from common vertices.

The GN1 architecture is schematised in Figure 4.5. Each track is augmented with the

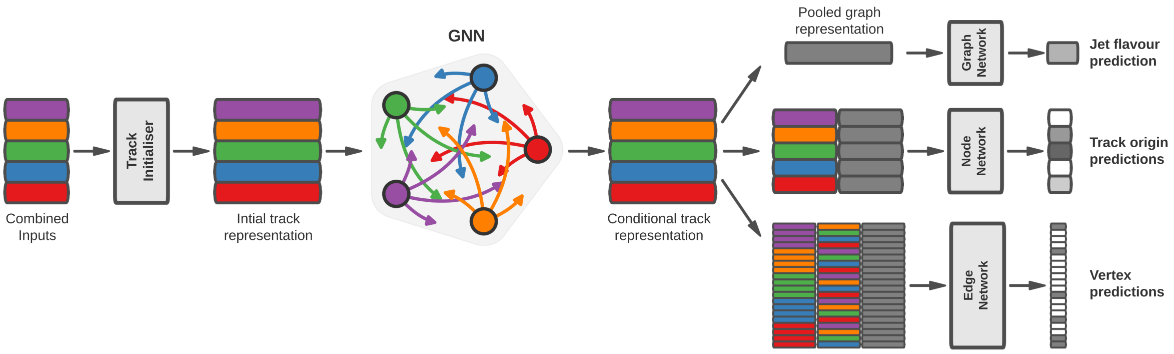

Figure 4.5: Diagram of the GN1 architecture [166].

jet p_T and η and passed through the same initialiser network, which resembles a Deep Sets model without the aggregation step. The initialised tracks are then employed in a fully connected graph neural network (GNN). Each track corresponds to a node, and the edges are formed by pairwise track relationships. After three message-passing steps, the conditional track representations are fed into three separate fully connected networks responsible for flavour tagging and the two auxiliary tasks. A further improvement is achieved with $\mathrm{GN2}$ [182] by refining the training procedure and enhancing the architecture with a transformer-like attention mechanism [183] between the graph layers.

Performance and calibrations

The performance of b-tagging algorithms is quantified using Receiver Operating Characteristic (ROC) curves, as is common for standard classifiers. The efficiency of selecting b-jets at a given threshold on the discriminant variable is compared to that of selecting background jets (c-, τ-, or light-jets). The standard practice in ATLAS is to display the b-jet efficiency on the horizontal axis and the background rejection factor (the inverse of background efficiency) on the vertical axis, as shown in Figure 4.6 for the Run 1 taggers. These ROC curves are derived from jets in $\mathrm{t\bar{t}}$ simulated events. The same convention is used to report tagger performance in the following chapters.

Figure 4.7 illustrates how the discrimination power for b-taggers has significantly improved over the first few years of Run 3, owing to the new unified taggers. With $\mathrm{GN2}$, at a 70% b-jet efficiency working point, the expected background contamination from light-jets is below 0.05%.

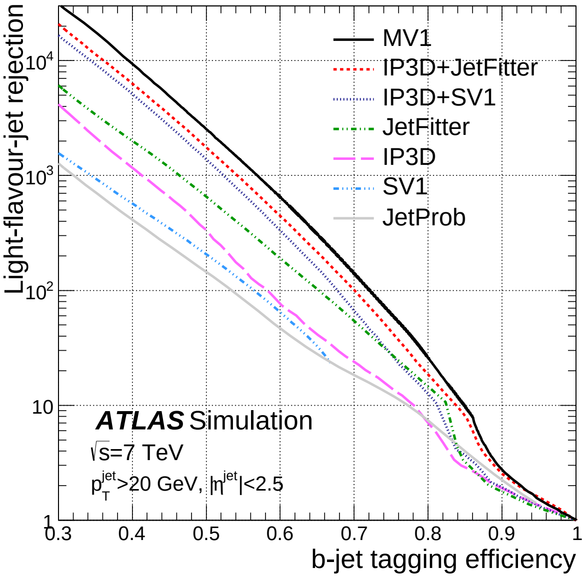

Figure 4.6: The ROC curves for b-jet efficiency versus light-jets rejection (the inverse of the miss-tag rate for light-jets) for the Run 1 taggers [167].

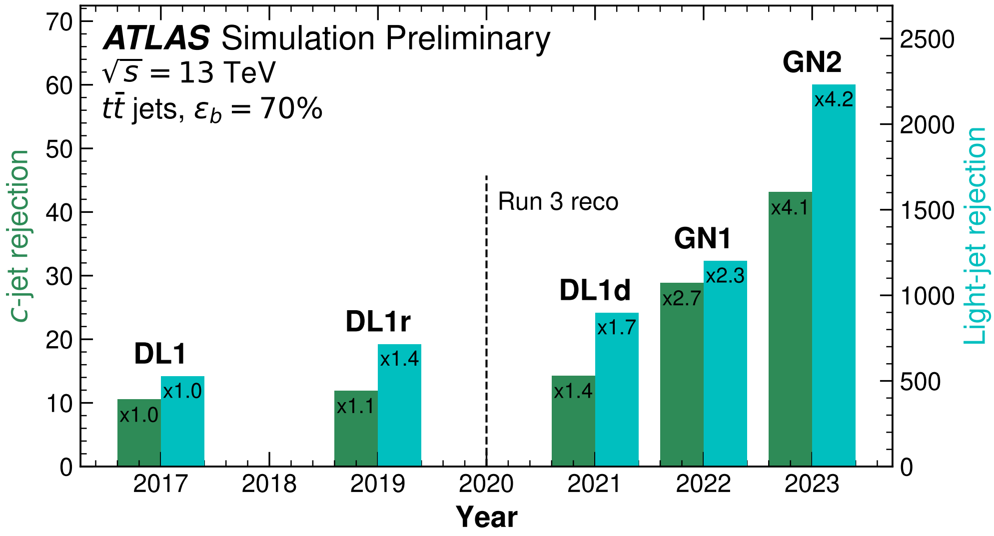

Figure 4.7: The improvement in background rejections over the years with the novel taggers at a fixed 70% b-jet efficiency [182].

To assess and correct for potential mismodelling of the low-level variables used as inputs to the algorithms, and to account for detector effects not present in MC, the taggers are calibrated using real data. The b-jet efficiency and c-jet mistagging rate are measured using $\mathrm{t\bar{t}}$ events [168, 184]. For the former measurement, both W bosons originating from top-pairs decays must produce leptons, significantly enhancing the b-jet purity in the final state. For c-jet mistagging rate measurement, one W boson must decay leptonically and the other hadronically, with the latter producing a c-jet in 33% of cases. The light-jets have an extremely low mistag rate for tighter working points leading to high statistical uncertainties. To address this challenge, a different approach named *Negative Tag* was developed [185], using Z+jets events. Since the signed impact parameter distribution is symmetric around zero for light-jets, flipping the signs of d_0 and z_0 has minimal impact on the mistag rate. Conversely, the efficiency of for heavy hadrons drops considerably. In this way, a relatively high fraction of light-jets is preserved after applying a “sign-flipped” tagger cut. For all flavours, the efficiency is measured in bins of jet $\mathrm{p_T}$, and for a limited set of b-tagging working points.

Chapter 5

The Run 3 b-jet trigger

While soft-QCD interactions have the highest rates at the LHC experiments, the QCD jets from p-p collisions are predominantly composed of fragmentation products of light-flavoured quarks or gluons. Hence, triggering events containing b-jets substantially reduces the data acquisition rates. This makes the study of relatively soft processes feasible if they decay to b-hadrons that produce collimated jets, such as the cases discussed in Section 4.1. In ATLAS, many selections use b-tagging algorithms at the HLT. This specialised sector of the trigger is known as the *b-jet trigger*.

For Run 3, the ATLAS HLT system faced several key upgrades with respect to the previous LHC Run. Some of these changes, such as the introduction of PFLOW jets, drastically shaped the design of the b-jet trigger reconstruction workflow. Furthermore, to improve the performance and maintain coherence with offline b-tagging, the tagging algorithms were updated throughout the years, following the offline developments closely.

The trigger-specific reconstruction relevant for HLT b-tagging is discussed in the first section of this chapter, excluding the fast b-tagging preselection that is the main topic of the next chapter. Then, the b-jet trigger algorithms are described along with their performance, highlighting significant improvements achieved in comparison with the Run 2 trigger configuration.

5.1 Trigger reconstruction

Tracking within the ID is not attainable at the L1-trigger stage, as discussed in Section 2.3. Therefore, b-tagging can only be performed at the HLT, after a sequence of selections that considerably reduces the rate. Typically, trigger chains with b-tagging selections are seeded by L1 jets. To meet the low-rate requirements for HLT reconstruction, either a high-energy jet ($E_T > 100\,\text{GeV}$) or several softer jets are required. For some special use cases where a high multiplicity of b-jets is expected, events are triggered by an L1Topo selection, which considers muons overlapping with calorimeter energy deposits [3]. As mentioned in Section 4.2, roughly 20% of b-hadron decay chains result in a muon, whereas this fraction is lower for lighter hadrons. At HLT, a few objects must be reconstructed before running specific algorithms for flavour tagging: tracks, vertices, and jets.

HLT tracking and vertexing

Track reconstruction at the HLT follows different strategies depending on the input rate and the physics signature it aims to select. In Run 3, a similar approach to that employed in Run 2 is retained, where tracking runs in two stages: *fast* and *precision tracking* [186, 187]. The former is performed using the Fast Track Finder (FTF) algorithm, a trigger-specialised pattern recognition algorithm. It focuses on maximising efficiency and quickly providing tracks for use in early stages of the trigger sequences. Precision tracking is always seeded by FTF tracks and refines them by running offline-like reconstruction.

The formation of triplets of space-points using hit clusters from the Pixel detectors is the first step of FTF tracking. The azimuthal angle is divided in 50 sectors, and triplets can spread across at most three of them. A conformal mapping is applied to the triplets [188], allowing for a rapid estimation of track parameters and the quick discarding of seeds below a minimum p_T requirement. Based on the obtained parameters, the seeds are grown into tracks using an alternative configuration of the standard track search procedure that focuses on enhancing the processing speed, also disregarding hits from the TRT system.

Precision tracking starts from FTF tracks and runs the same algorithms used for offline tracking. Fake tracks are rejected using the ambiguity resolution procedure described in Section 3.1. Track candidates are extended with TRT hits, improving momentum measurements and providing particle identification capabilities. Furthermore, tracks are re-fitted with a higher resolution using the offline χ^2 methodology. Therefore, precision tracking improves the purity and the quality of tracks; however, its efficiency cannot exceed that of FTF tracks, as it would require running a revised instance of the pattern recognition step.

Depending on the HLT reconstruction stage, tracking runs in RoIs or across the full ID acceptance (*full-scan* tracking). RoI tracking is usually seeded by calorimeter deposits and used to reconstruct electrons or photons. The b-jet trigger also uses precision RoI tracking

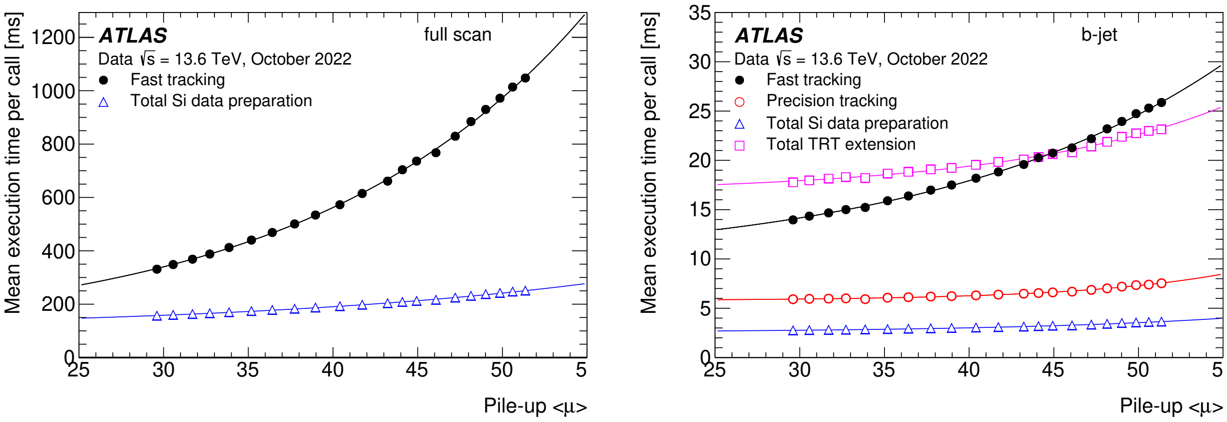

(a) **(b)**

Figure 5.1: The average processing time per algorithm call for the (a) full scan track reconstruction and (b) b-jet RoI tracking as a function of the average number of inelastic p-p collisions [2].

in detector regions corresponding to jets. Full-scan tracking was introduced in HLT for the first time in Run 3 and is necessary for reconstructing PFlow objects [2]. It runs only the FTF stage with a higher threshold on the seed p_T (1 GeV instead of 500 MeV) and is the most CPU-intensive algorithm at the HLT. As shown in Figure 5.1(a), full-scan tracking can take more than 1 second per event at higher instantaneous luminosities ($\langle\mu\rangle > 50$), when including also the input data preparation time. In contrast, at high pile-up, the two tracking stages for b-jet RoIs take less than 70 ms on average, as illustrated in Figure 5.1(b).

Similar to the standard approach for Run 3 mentioned in Section 3.1, PV finding uses Gaussian seeds from track density along the z-axis and executes an adaptive multi-vertex algorithm [118] that is robust against pile-up. This technique achieves a resolution of approximately $30\,\mu$m along the beam axis and $10\,\mu$m in the transverse plane [3].

HLT jet reconstruction

The HLT runs the same topo-clustering algorithm described in Section 3.2, always using the “420” schema for the thresholds on cell significances [2]. The topo-clusters are used to build calorimeter-only jets, known as HLT EMTopo jets, using the anti-k_T algorithm. Both $R = 0.4$ and $R = 1$ are used; however, this thesis focuses only on the small-R jets.

In most cases, the events undergo a calorimeter-based preselection, relying on EMTopo jets, before further reconstruction of PFlow candidates. This procedure can effectively reduce the event rates when the L1 threshold is significantly lower than the final HLT selection, with efficiency losses smaller than 2%. However, when the L1 and HLT jet p_T thresholds are close, this preselection is no longer sufficient and alternative methodologies are required, such as the one discussed in the next chapter.

Reconstruction of PFlow jets requires a full-scan FTF pass. After identifying fast tracks, they are matched to topo-clusters to form the PFlow objects associated with the

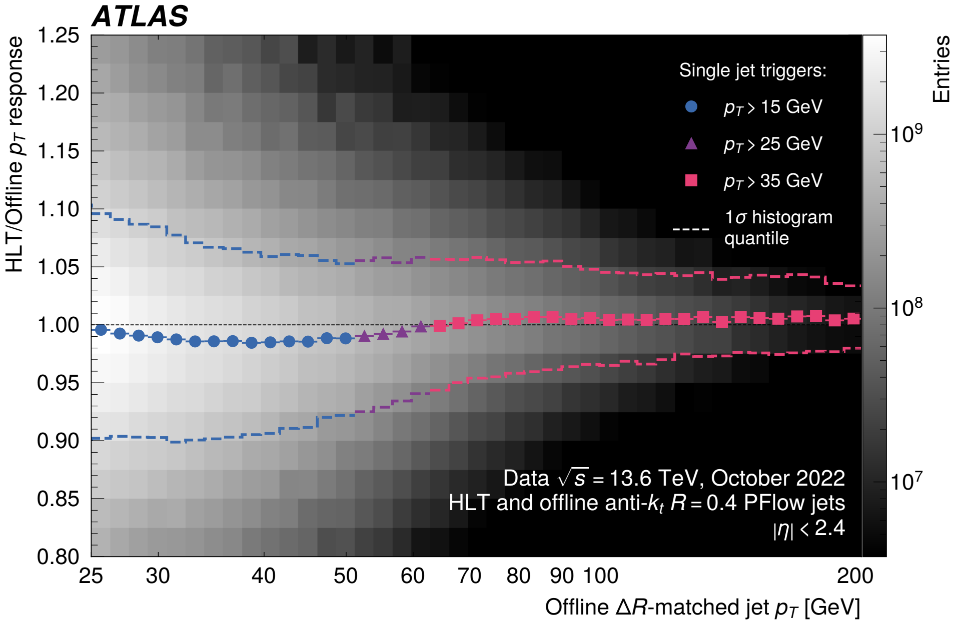

Figure 5.2: The p_T ratio for ΔR matched online and offline PFLOW jets [2]. The ratio is calculated using fully calibrated jets, with data collected during 2022 using random-seeded triggers.

PV, as described in Section 3.3. In this way, contributions from pile-up are substantially suppressed. Finally, charged and neutral PFLOW objects serves as constituents for building PFLOW jets. The latter are calibrated following a procedure identical to the offline one, but with re-derived factors. The HLT calibration yields excellent closure with offline-reconstructed jets, as demonstrated in Figure 5.2.

5.2 The b-jet trigger algorithms

The prerequisites for executing specific flavour-tagging algorithms at the HLT are jets and their respective RoI precision tracks. Once reconstructed, tracks considered for b-tagging are matched to jets using a shrinking cone association (Equation 4.1).

During the 2022 data-taking period, the two-stage reconstruction methodology was employed in the b-jet trigger. Vertex-based low-level reconstruction algorithms, described in Section 4.2 (inclusive secondary vertex finding and JETFITTER), were run together with a dedicated tuning of DIPS. The output of these algorithms provided the input features for a trigger-specific version of the DL1d algorithm, which was used as a high-level tagger to make the final HLT decisions in b-jet trigger chains. This version of DL1d consists of eight hidden layers with around 45 000 trainable parameters. Its outputs correspond to the jet probabilities for the three jet categories (b-, c-, and light-jets) and are combined into a single discriminant, as shown in Equation 4.3. The charm fraction is set to $f_c = 0.018$, a value determined by a previous optimisation study performed with offline taggers, which accounted for backgrounds from several analyses relying on b-jets [169].

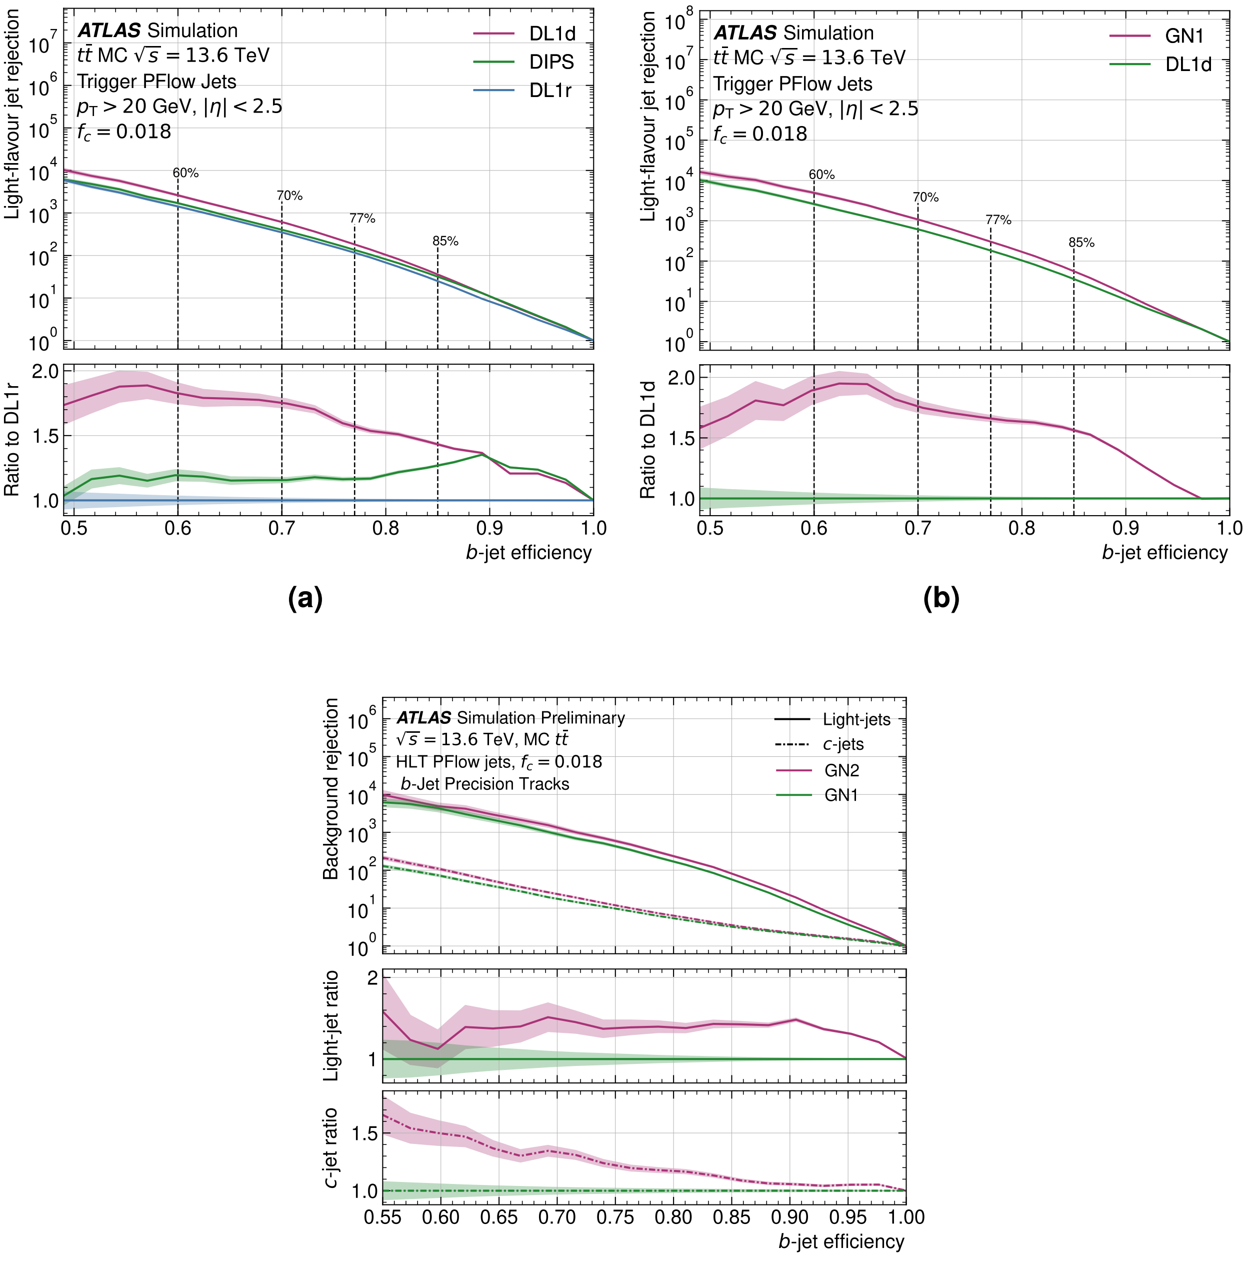

(a)

(b)

(c)

Figure 5.3: The ROC curves for several b-tagging discriminants calculated in simulated $t\bar{t}$ events for PFLOW jets [3, 189]. The (a) trigger DIPS and DL1d are compared to the DL1r algorithm, (b) the GN1 tagger introduced since 2023 is compared to DL1d, and (c) the improvement obtained with GN2 is showed in comparison with GN1.

Following the developments in offline b-tagging, a dedicated version of $\mathrm{GN1}$ was deployed in the trigger in 2023, using tracks directly as inputs. After a year of data-taking and validation of the unified tagger approach, the vertex-based algorithms were decommissioned from the trigger software in 2024. Furthermore, the $\mathrm{GN2}$ tagger was integrated into the HLT reconstruction, keeping pace with advancements in the flavour-tagging community. Both $\mathrm{GN1}$ and $\mathrm{GN2}$ provide a b-tagging discriminant in the same way as $\mathrm{DL1d}$, and in HLT, the f_c value remains identical for all the taggers.

The expected performance for $\mathrm{DL1d}$, obtained from $t\bar{t}$ simulations, is shown in Figure 5.3(a). For the standard working points used in the trigger – 85%, 77%, 70%, and 60% b-jet efficiency – improvement factors ranging from 1.5 to almost 2 are achieved compared to the $\mathrm{DL1r}$ tagger in terms of light-jet rejection, which constitutes the dominant background. A similar level of improvement was achieved by upgrading $\mathrm{DL1d}$ with $\mathrm{GN1}$, as shown in Figure 5.3(b). Additionally, $\mathrm{GN2}$ further increased light-jet rejection by 50%, as illustrated in Figure 5.3(c).

5.3 Trigger chains

The trigger menu for Run 3 includes a wide set of chains running b-jet algorithms [3]. In addition to several general-purpose trigger chains that aim to cover a broad spectrum of signatures, specialised chains are implemented to optimise acceptance for specific physics cases, such as di-Higgs decays or SUSY searches targeting final states with b-hadrons. For the various applications, the majority of trigger chains require a *fast* b*-tagging* preselection (discussed in Chapter 6) besides a cut on calorimeter-only jets.

The design of both inclusive and more model-oriented chains typically needs to find the right balance between the jet p_T thresholds and the working point used for b-jet identification. The goal is to maximise the acceptance of physics of interest while maintaining a manageable acquisition rate, which is dominated by QCD events. To assess the behaviour of trigger selections prior to data-taking, the rates are evaluated using an “enhanced bias” dataset [190], which represents a statistically significant estimate of a zero-bias data sample obtained by appropriately overweighting high p_T events.

To prepare the Run 3 menu, an important consideration was the impact of increased pile-up conditions compared to Run 2. Softer jet requirements result in rates more sensitive to the instantaneous luminosity, as shown in Figure 5.4 for the three example chains presented in Table 5.1. The two trailing jets in the di-Higgs chain $\mathrm{2b2j}_{\mathrm{HH}}^{asym}$ have a p_T threshold below $30\,$GeV, and as the pile-up increases, the rate rapidly grows compared to chains requiring smaller jet multiplicities but with higher p_T.

The advancements in the taggers and the usage of the new b-tagging preselection have greatly improved the trigger menu capabilities [3]. The two largest decay channels of Higgs pairs,$\mathrm{HH} \to b\bar{b}b\bar{b}$ and $\mathrm{HH} \to b\bar{b}\tau^+\tau^-$, are being recorded in Run 3 with significantly higher

Table 5.1: Description of three b-jet trigger chains present in the Run 3 menu. The first two have a general purpose, while the last one (`2b2j`$^{\mathrm{asym}}_{\mathrm{HH}}$) is specifically designed for di-Higgs searches. The L1 selections are applied on jets formed by calorimeter towers, as described in Section 2.3. The HLT preselection uses EMTopo jets and fast b-tagging scores (see Chapter 6), and the final HLT selections are applied on PFlow jets. The jet p_T selections are reported as "*minimum multiplicity*×*threshold*", while the flavour tagging ones "*minimum tagged-jets*@*working point*".

Chain	L1 jet	HLT preselection		HLT final thresholds	
	E_T [GeV]	p_T [GeV]	fast b-tag [ε_b]	p_T [GeV]	b-tag [ε_b]
`1b`	100	1×180	-	1×255	1 @ 70%
		1×225	-	1×300	1 @ 77%
		1×225	-	1×360	1 @ 85%
`2b1j`	1×100, 2×30	1×80, 2×45	2 @ 90%	1×150, 2×55	2 @ 70%
`2b2j`$^{\mathrm{asym}}_{\mathrm{HH}}$	1×45 $(\lvert\eta\rvert < 2.1)$, 2×15 $(\lvert\eta\rvert < 2.5)$	4×20	2 @ 85%	1×80, 1×55, 1×28, 1×20	2 @ 77%

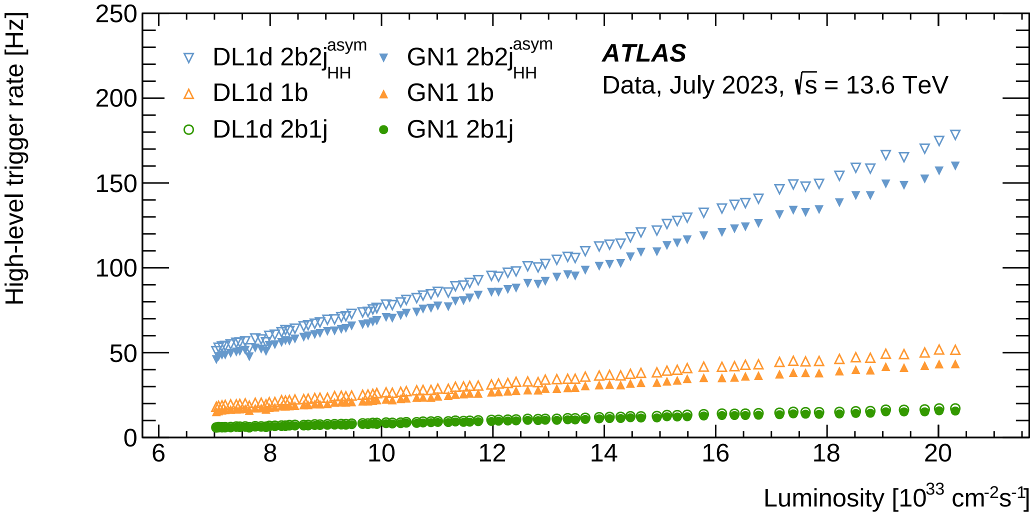

Figure 5.4: The HLT output rates observed in Run 3 for several b-jet trigger chains as a function of the instantaneous luminosity conditions [3]. Chains with identical thresholds but two different taggers (DL1d and GN1) are compared.

efficiencies than in the past. The strong background rejection from b-tagging enabled low p_{T} thresholds, such as those for the $\mathtt{2b2j}_{\mathrm{HH}}^{\mathrm{asym}}$ chain. For $\mathrm{HH} \to \mathrm{b\bar{b}b\bar{b}}$, the latter trigger alone yielded a relative increase in acceptance of more than 30% and, when combined with additional chains, the total fraction of collected di-Higgs events increased from 41% in Run 2 to nearly 60%, as shown in Figure 5.5. The most important gains are in correspondence of low invariant di-Higgs mass, which is a particularly challenging phase space to trigger at the LHC.

The selection of $\mathrm{HH} \to \mathrm{b\bar{b}}\tau^{+}\tau^{-}$ events, with hadronically decaying τ-leptons, greatly benefited from the $\mathtt{2b2j}_{\mathrm{HH}}^{\mathrm{asym}}$ chain. Figure 5.6 shows that the Run 3 τ-triggers brought an absolute gain in efficiency of 4%, while the $\mathtt{2b2j}_{\mathrm{HH}}^{\mathrm{asym}}$ chain has large unique sensitivity to this process, increasing the total efficiency by 20% more. The trigger system in the phase space within the offline fiducial region of $\mathrm{HH} \to \mathrm{b\bar{b}}\tau^{+}\tau^{-}$, with two $\tau_{\text{had-vis}}$ and two jets with p_{T} above $20\,$GeV and within the ID acceptance ($|\eta| < 2.5$), has incremented its efficiency from 54% to 78%.

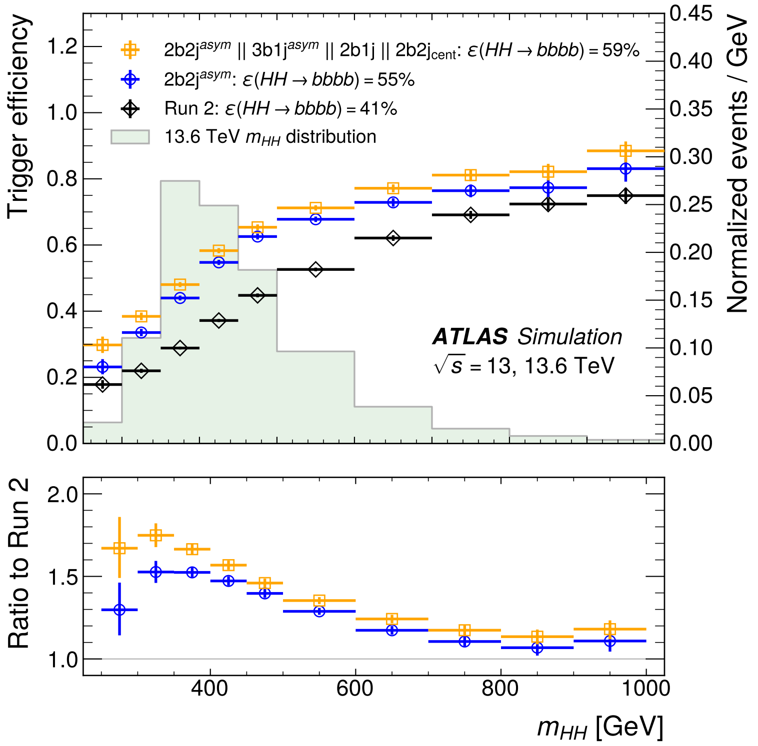

Figure 5.5: Performance of the b-jet trigger selections in $\mathrm{HH} \to \mathrm{b\bar{b}b\bar{b}}$ simulated events [3]. The reported trigger efficiency is inclusive. This plot demonstrates the improvement obtained in acceptance with respect to the Run 2 trigger configuration [191, 192].

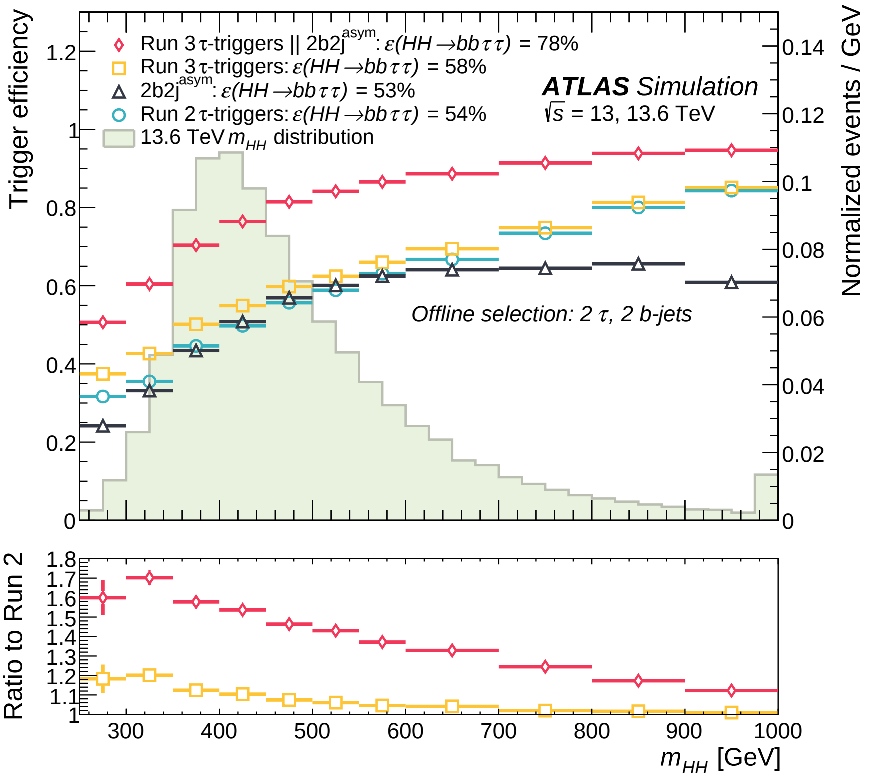

Figure 5.6: Performance of the b-jet trigger selections in $\mathrm{HH} \to \mathrm{b\bar{b}}\tau^{+}\tau^{-}$ simulated events [3]. The reported trigger efficiency is calculated with respect to fiducial cuts on offline reconstructed objects (two jets and two $\tau_{\text{had-vis}}$ with $p_{\mathrm{T}} > 20\,\mathrm{GeV}$ and $|\eta| < 2.5$). This plot demonstrates the improvement obtained in acceptance with respect to the Run 2 trigger configuration [191, 192].

Chapter 6

Fast b-tagging

The main challenge for the b-jet trigger is finding the balance between reducing the precision of reconstruction to fit within TDAQ constraints and rejecting enough background to keep a sustainable acquisition rate, while remaining highly efficient for signal events. Achieving the right balance became even more demanding at the start of Run 3, due to the introduction of more sophisticated reconstruction algorithms at HLT, such as PFlow. To tackle this challenge without compromising trigger performance, the *fast* b*-tagging* preselection was introduced. This method allowed for a substantial reduction in CPU cost of b-jet chains, enabling major improvements in HLT acceptance.

The fast b-tagging preselection relies on rapid and coarse reconstruction for the early rejection of events using a dedicated tagger, fastDIPS. Despite the reduced resolution, this procedure significantly impacts background rates while retaining high signal efficiency. This method was deployed for the first time in ATLAS in 2022 and was necessary to maintain an efficient trigger menu. Otherwise, PFlow reconstruction would have required a sizeable increase in p_T thresholds at L1. Furthermore, the integration of fast taggers at HLT enabled the inclusion of flavour-tagging information in the TLA data stream.

This chapter discusses the fast b-tagging methodology in detail, describing part of the work done for this thesis: the training of fast tagger, its expected performance in MC, and the preselection efficiency in data. Finally, improvements to the trigger menu are outlined.

6.1 The fast b-tagging preselection

Track reconstruction is highly computationally intensive. Since flavour tagging is deeply dependent on tracking, b-jet chains were among the top CPU consumers at HLT during Run 2 [190], as shown in Figure 6.1. Given their high computational cost, reducing their rates with early rejection helps conserve a large amount of resources, which can be redistributed within the HLT. Furthermore, despite tracking being heavily optimised for Run 3 (see Section 3.1), the full-scan FTF algorithm introduced to build PFLOW objects at HLT remains the most CPU-intensive process running in the HLT farm. It requires nearly 1 second per event, as shown earlier in Figure 5.1. As a result, flavour-tagging reconstruction cannot operate at the full rate of the L1 seeds (~ 10 kHz), but instead, it can only be performed after a substantial event filtering.

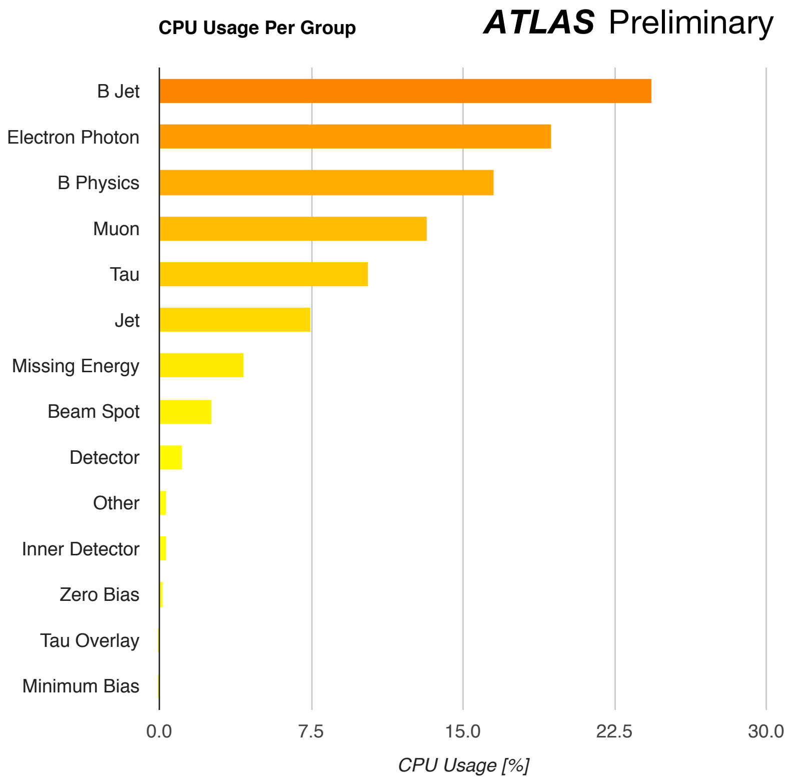

Figure 6.1: Fraction of CPU usage of trigger chain groups, based on different physics signatures, evaluated in 2016 with the Run 2 trigger menu [190].

To cope with higher rates and enhance physics acceptance, the fast b-tagging approach was introduced for Run 3. This methodology allowed the inclusion of the $2b2j^{asym}_{HH}$ chain, defined in Table 5.1, in the trigger menu, which otherwise would have required much tighter p_T cuts, compromising the efficiency for key processes such as $HH \rightarrow b\bar{b}b\bar{b}$. Fast b-tagging exploits ML algorithms to perform b-jet identification using fast tracks, *i.e.*, lower quality inputs compared to the precision taggers described in Section 5.2. Nevertheless, it achieves considerable background rejection.

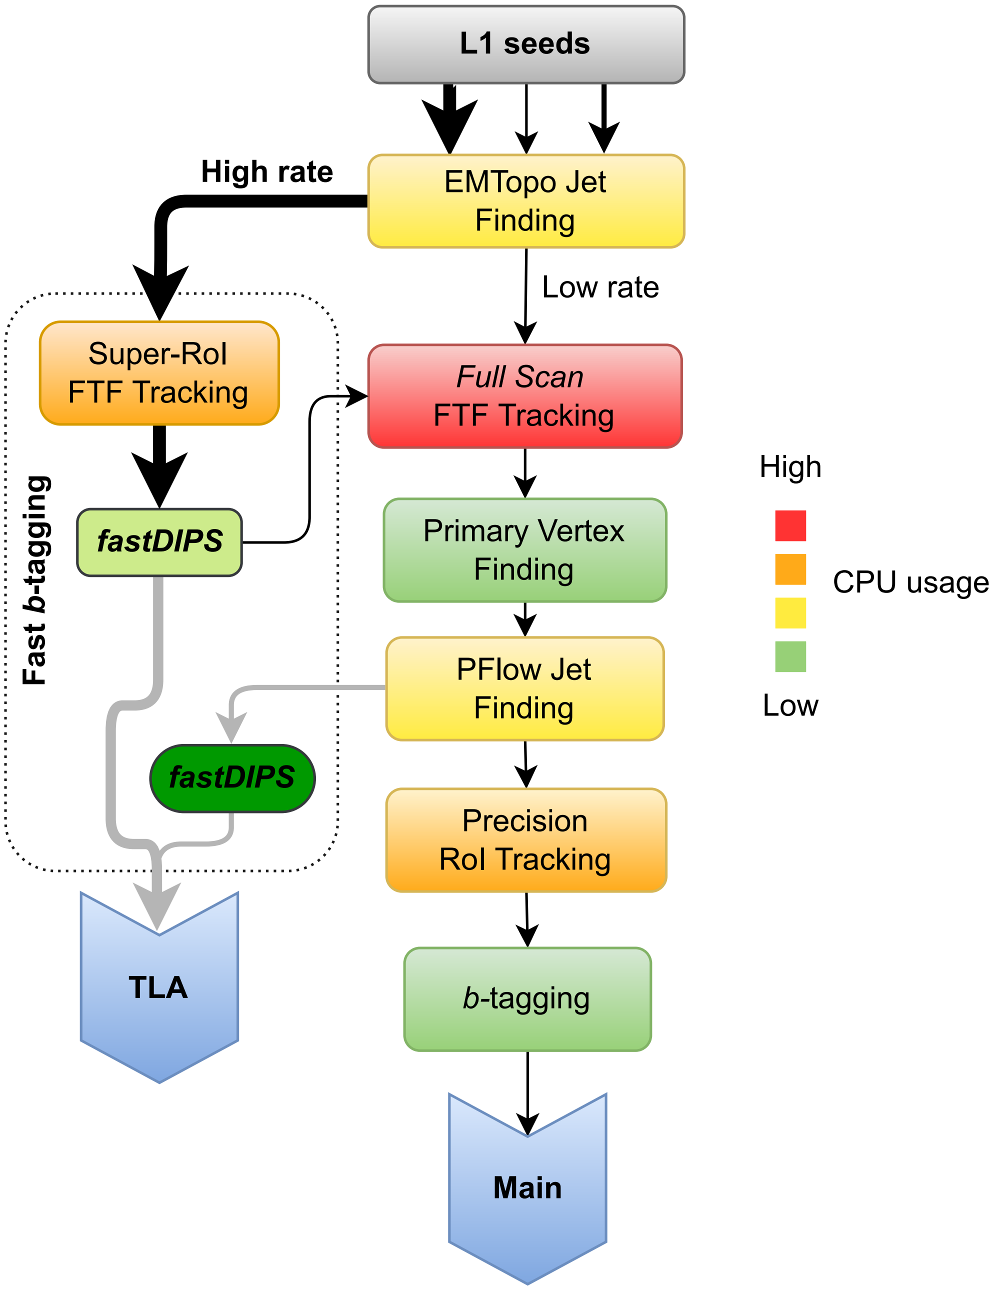

Figure 6.2: Schematic flowchart of the reconstruction steps executed for b-jet trigger selections in Run 3 HLT. The flowchart includes the novel fast b-tagging steps needed for triggering b-jet signatures with high rates. Two instances of FASTDIPS are present: the fast b-tagging preselection that tags EMTOPO jets to filter events before full scan tracking, and the PFLOW TLA tagger, which decorates jets with flavour scores for analysis purpose without applying any real-time cut. The preselection tagger score also is stored in the TLA stream for analyses using trigger EMTOPO jets, featuring higher rates. The diagram boxes are colour-coded, and each colour corresponds roughly to a different order of magnitude, ranging from green $\sim \mathrm{o}(1\,\mathrm{ms})$ to red $\sim \mathrm{o}(1\,\mathrm{s})$.

The reconstruction workflow for b-jet trigger selections is depicted in Figure 6.2. After a preselection on kinematics of EMTopo jets, full-scan tracking is run only for low rate trigger chains. Otherwise, the fast b-jet reconstruction is performed, executing fast tracking in limited RoIs around the jet axis. To avoid duplication in overlapping regions, tracking is run in a *super*-RoI, which consists of the union of all jet RoIs. Then, b-tagging inference is performed with fastDIPS to filter events seeding full scan FTF. This procedure introduces a sizeable rate reduction while maintaining signal efficiency. After tracks are reconstructed across the entire ID, the PFlow reconstruction sequence is activated, followed by a second pass of RoI precision tracking before the final b-tagging step.

The introduction of early b-tagging in the HLT workflow also enabled the augmentation of the TLA stream with flavour tagging discriminants. This is a completely new feature of the Run 3 HLT. The improved algorithms demonstrated meaningful tagging capabilities, even with low-precision inputs. Two different taggers are deployed: one for the EMTopo jets running at the highest rates, and another one for PFlow jets with slightly higher precision. The latter is used only in TLA, no HLT selections are applied based on it.

6.2 The fastDIPS algorithm

The preselection tagger uses tracks from the super-RoI associated with EMTopo jets within a shrinking cone around their axis (Equation 4.1). Similarly, the TLA tagger uses full-scan tracks associated with PFlow jets. Both track groups are reconstructed using a p_T threshold of 1 GeV for the tracking seeds (higher than the 0.5 GeV threshold for precision tracks). The major difference between the super-RoI and full-scan tracks is that the latter benefit from PV reconstruction to calculate their displacement. No PV information is available in the super-RoI, and the impact parameters are computed with respect to the beam spot, significantly reducing the resolution of the longitudinal component.

The track features fed into the fastDIPS models are listed in Table 6.1. The few differences between models for PFlow and EMTopo jets were determined by the available variables at different stages of HLT. Both taggers use the same network sizes: tracks are first embedded in a 128-dimensional latent space using a network with two hidden layers (Φ), where the tracks are summed and then passed to a deeper network (F) with four hidden layers. The F network outputs three scores (p_b, p_c, and p_l), which are combined into a discriminant D_b (Equation 4.3) to classify the jets.

All the main hyperparameters used in the algorithm training are reported in Table 6.2. The fastDIPS architecture has close to $60\,000$ trainable parameters and used up to ten million jets from simulated $t\bar{t}$ events as the training dataset. The jet flavour composition of the training sample is balanced, with roughly equal number of jets per class. Furthermore, jets were resampled in bins of p_T and η to ensure that the selection does not depend on the kinematic spectrum of the specific process used for training. The distributions for b-

Table 6.1: Summary of the track features used by the FASTDIPS taggers.

Input	Description	EMTopo	PFlow
	Track Kinematics		
d_0	Transverse impact parameter	×	×
d_0^{sign}	Signed transverse impact parameter	×	
$z_0 \sin\theta$	Longitudinal impact parameter		×
z_0^{beam}	Projection on the z-axis of the closest approach to the beamline	×	
$\log \sigma_{z_0^{\text{beam}}}$	Logarithm of uncertainty on z_0^{beam}	×	
s_{d_0}	Transverse impact parameter significance (d_0/σ_{d_0})		×
s_{z_0}	Longitudinal impact parameter significance ($z_0 \sin\theta/\sigma_{z_0 \sin\theta}$)		×
$\log p_T^{\text{frac}}$	Logarithm of the fraction of the jet p_T carried by the track	×	×
q/p	Track curvature (charge over momentum)	×	×
$\Delta\eta$	Angular separation between track and jet	×	×
$\Delta\phi$		×	×
ΔR		×	×
	Track Quality		
nPixHits	Number of total hits in the Pixel detector	×	×
nSCTHits	Number of total hits in the SCT detector	×	×
IBL hits	Number of innermost pixel layer hits	×	×
PIX1 hits	Number of hits in next-to-innermost pixel layer	×	×
n_{DOF}	Number of degrees of freedom in track fit	×	×
χ^2	$\sum_{\text{hits on track}} (r/\sigma_r)^2$ [$r \equiv$ hit residual, $\sigma_r \equiv$ residual uncertainty]	×	×

Table 6.2: Summary of the hyperparameters used to train the FASTDIPS taggers.

Track embedding network (Φ)	
Input size	15
Hidden layers size	100, 100
Output size	128
Track sum network (F)	
Input size	128
Hidden layers size	100, 100, 100, 30
Output size	3
Total trainable parameters	60 190
Training epochs	200
Learning rate	0.001
Batch size	1024
Dropout rate	0.2
Training sample size	10 000 000 jets
Validation sample size	600 000 jets

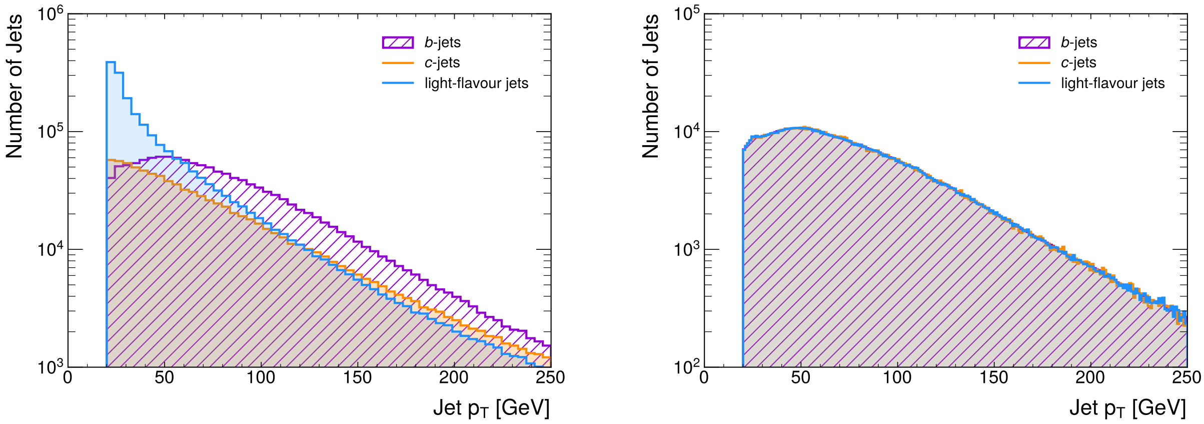

(a) **(b)**

Figure 6.3: Jet p_T distribution per flavour class (a) before and (b) after the two-dimensional resampling. The latter plot shows the more granular p_T binning used for the resampling procedure.

jets are taken as the reference, as shown in Figure 6.3 for the jet p_T, to maximise their statistical power.

Preprocessing of input datasets and network training is performed with the UMAMI [193] framework, which employs a KERAS interface to TENSORFLOW [194]. The network [195] is then exported into a format compatible with the ATLAS trigger software using the LWTNN tool [196].

6.3 Performance studies

The performance of the fast b-tagging method was initially evaluated in MC simulations using $t\bar{t}$ samples. Once deployed for data-taking, performance studies were conducted with the early Run 3 dataset. Since the main background trigger rates are due to light jets, these studies focus on light-jet rejection.

Preselection optimisation

The first step was to determine the minimal size of tracking RoIs required to achieve low CPU consumption at the preselection stage, while retaining enough information for background rejection. As the $\Delta\eta$ and $\Delta\phi$ sizes decrease, the outermost tracks from the jet axis are lost. This has a modest impact on the performance for RoI half-widths of 0.3. Further shrinking of the tracking area significantly affects b-tagging, reducing the rejection rates by a factor of almost two for 0.2 half-widths, as shown in Figure 6.4(a). The p_T threshold for track seeding was also scanned. A minor loss in performance is observed when increasing the p_T requirement from 0.5 GeV to 1 GeV, but stronger effects appear for more stringent cuts, as demonstrated for 1.5 GeV in Figure 6.4(b).

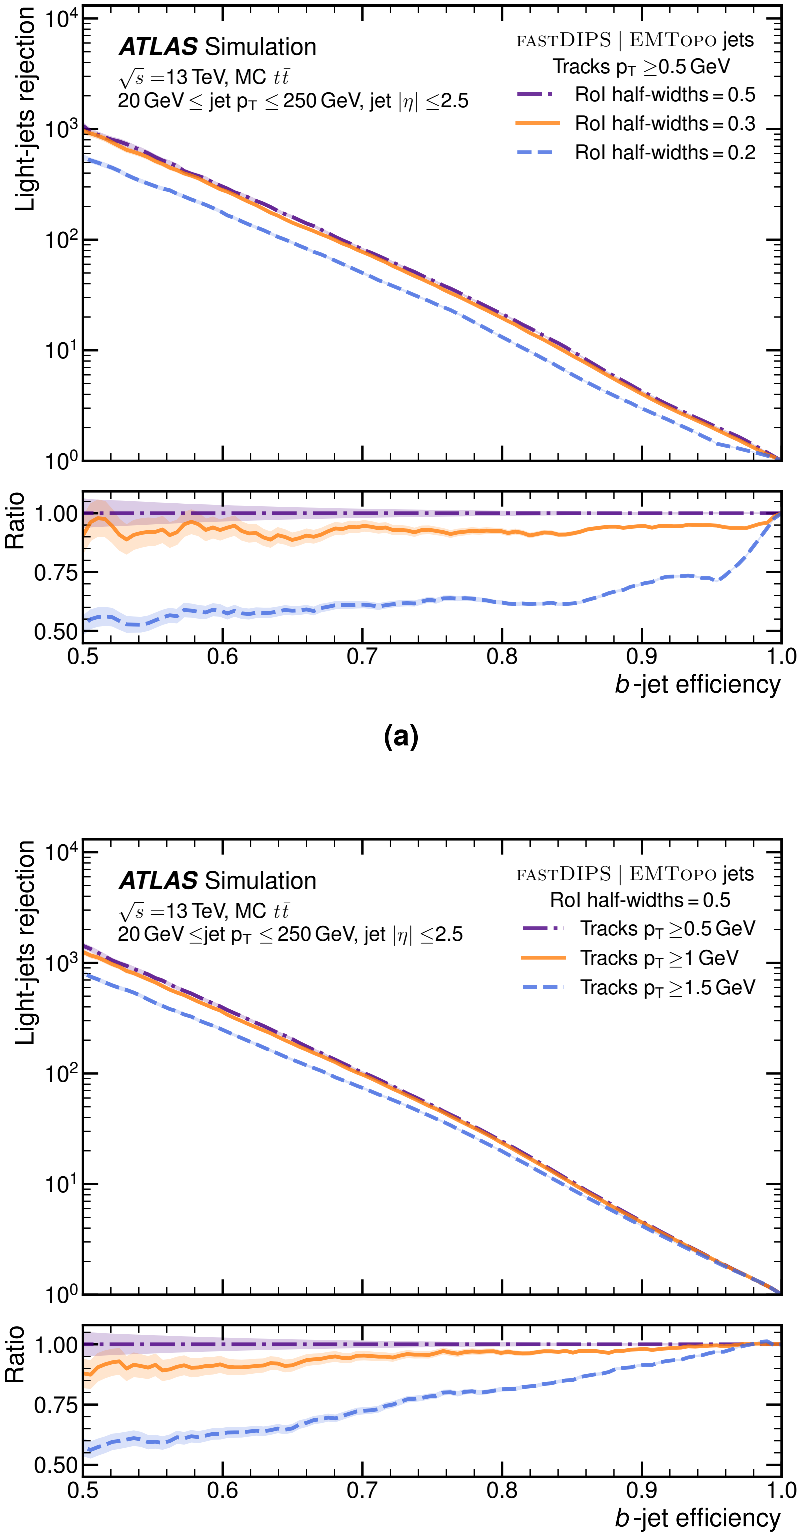

(a)

(b)

Figure 6.4: Optimisation of the RoI parameters used to run FTF tracking for the fast b-tagging preselection with EMTOPO jets. The ROC curves of different FASTDIPS versions, trained with different track sets determined by the RoI size, are reported with for a scan in (a) RoI half-widths and (b) track seeding p_T requirement. Statistical uncertainties for each ROC curve, represented with shaded regions, are computed assuming binomial efficiency errors.

Tuning these RoI parameters considerably changes the CPU time required to run the FTF algorithm. Figure 6.5 illustrates how the CPU costs drop for narrower RoIs and higher p_T thresholds. The working point used since 2022 is defined with RoI half-widths of 0.3 and a track p_T requirement of 1 GeV, resulting in an execution time four times smaller than that of the full-scan FTF algorithm. This reduction enables super-RoI tracking at a sustainable rate and has a marginal impact on b-tagging performance.

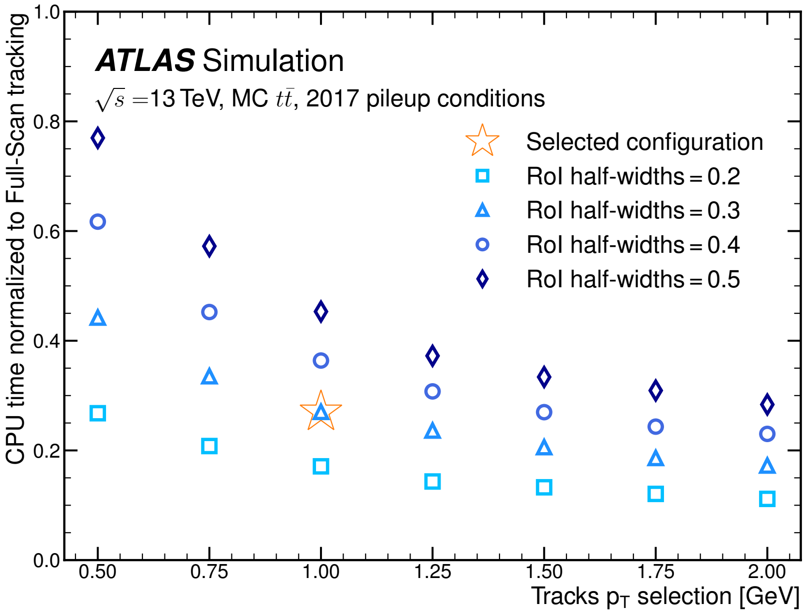

Figure 6.5: Optimisation of the RoI parameters used to run FTF tracking for the fast b-tagging preselection with EMTOPO jets. The vertical axis shows CPU time corresponding to different settings normalised to the time required to execute full scan FTF. The working point used in the ATLAS HLT is shown with an open star.

MC performance

While both network training and RoI optimisation used MC samples simulating Run 2 conditions ($\sqrt{s} = 13$ TeV and 2017 pile-up profile), the performance of fast taggers running in the trigger system was assessed using an independent $t\bar{t}$ MC sample, with $\sqrt{s} = 13.6$ TeV and the expected pile-up profile for the 2022 LHC runs.

The ROC curves for the two versions of FASTDIPS are shown in Figure 6.6(a) and compared to the trigger DL1d precision tagger, referred to as the *high-level tagger* in the following plots. The PFLOW tagger for the TLA stream achieves rejection rates comparable to those of the precision tagger, despite relying on fast tracks and lacking vertex-based reconstruction features. For preselection FASTDIPS, the most relevant region is at high b-jet efficiencies. At 85% efficiency, the expected background rejection is a factor of ten, leading to significant rate reductions.

To verify the correlation between the fast and high-level tagger, and to ensure consistency in the jets selected by both algorithms, the conditional efficiency was studied. Such metric was defined as:

$$\text{Conditional b-jet efficiency} = \frac{\text{\# true b-jets passing both fast and high-level b-tagging}}{\text{\# true b-jets passing high-level b-tagging}}\,. \tag{6.1}$$

To evaluate this metric, EMTopo and PFlow jets are geometrically matched with $\Delta R < 0.3$, and the relative p_T difference must be less than 10%. Figure 6.6(b) shows the observed overlap between the two taggers at various working points. For all working points, the

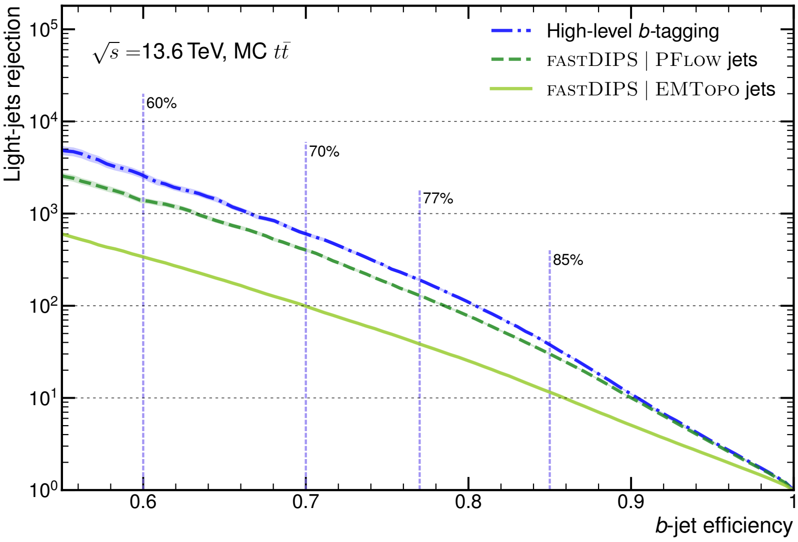

(a)

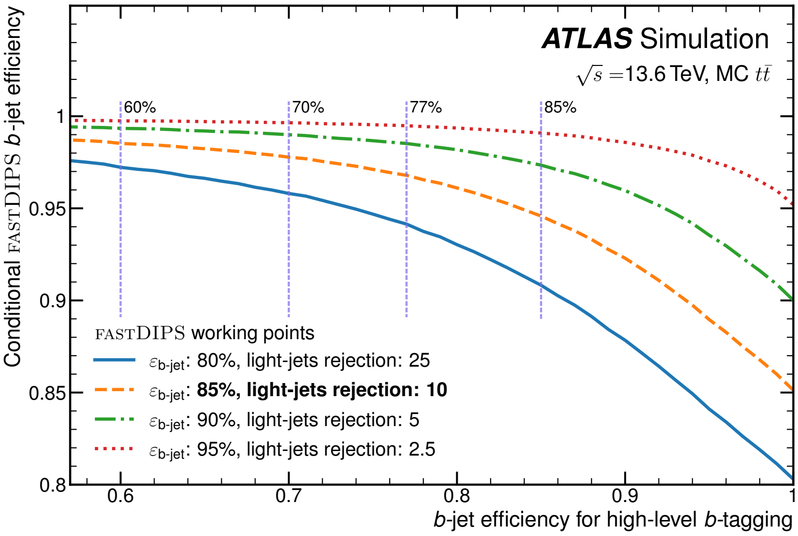

(b)

Figure 6.6: Figure (a) shows ROC curves for the three taggers deployed at the HLT in 2022: the precision DL1d tagger (dash-dotted blue line), FASTDIPS for TLA that runs on PFlow jets (dashed green line), and preselection FASTDIPS using fast tracks in EMTopo jet cones (light-green line). Statistical uncertainties for each ROC curve, represented with shaded regions, are computed assuming binomial efficiency error. The conditional efficiency between DL1d and pre-selection FASTDIPS is plotted in Figure (b), defined by Equation 6.1. The purple vertical dashed lines represent the most common working points used for b-tagging.

conditional efficiency remains above 90% and improves as the nominal b-jet efficiencies between FASTDIPS and DL1d are more distant.

Run 3 data performance

To validate the results obtained in MC, data from p-p collisions was analysed in a phase space enriched with $t\bar{t}$ pairs decaying into two leptons and jets. An L1 trigger selection requiring an isolated deposit with significant energy (> 22 GeV) in the electromagnetic calorimeter is applied. At HLT, events are selected if they contain an electron passing a *tight* identification criterion [135], with $p_T > 28$ GeV and $|\eta| < 2$, and a muon with $p_T > 25$ GeV and $|\eta| < 2$ of opposite charge. To further enhance $t\bar{t}$ purity, at least one HLT PFLOW jet with $p_T > 25$, $|\eta| < 2.5$, and JVT> 0.5, must pass the 80% DL1d working point in each event.

With this event selection applied to both data and simulation, $t\bar{t}$ becomes the dominant process. Figure 6.7 demonstrates that the MC sample for di-leptonic $t\bar{t}$ decays reasonably describes the FASTDIPS discriminant distribution observed in data, without considering additional backgrounds.

The performance of fast b-tagging is also evaluated in terms of efficiency. The impact of the preselection on the leading HLT PFLOW jet is assessed. Two different metrics are

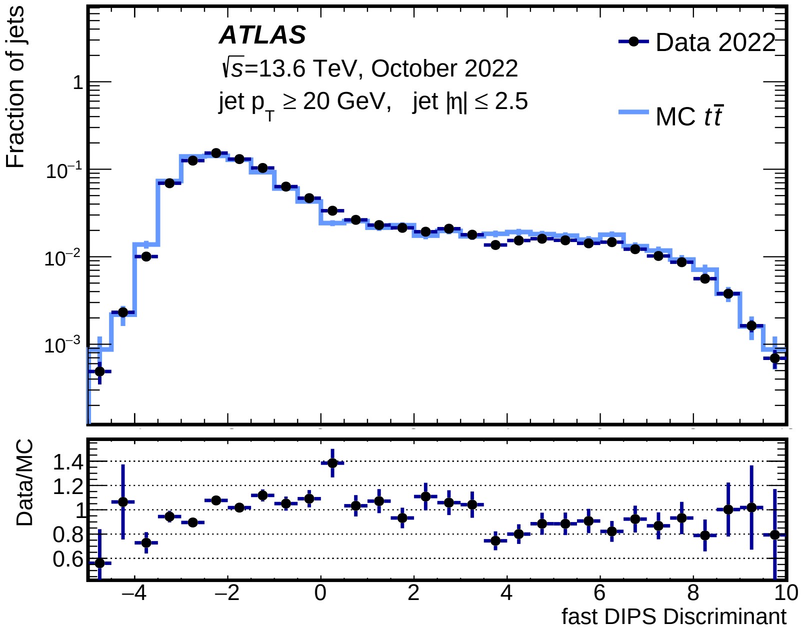

Figure 6.7: Distribution of the preselection FASTDIPS b-tagging discriminant for EMTOPO jets from selected data runs compared to MC simulations, in $t\bar{t}$ enriched events. Only statistical uncertainties are displayed for both MC and data.

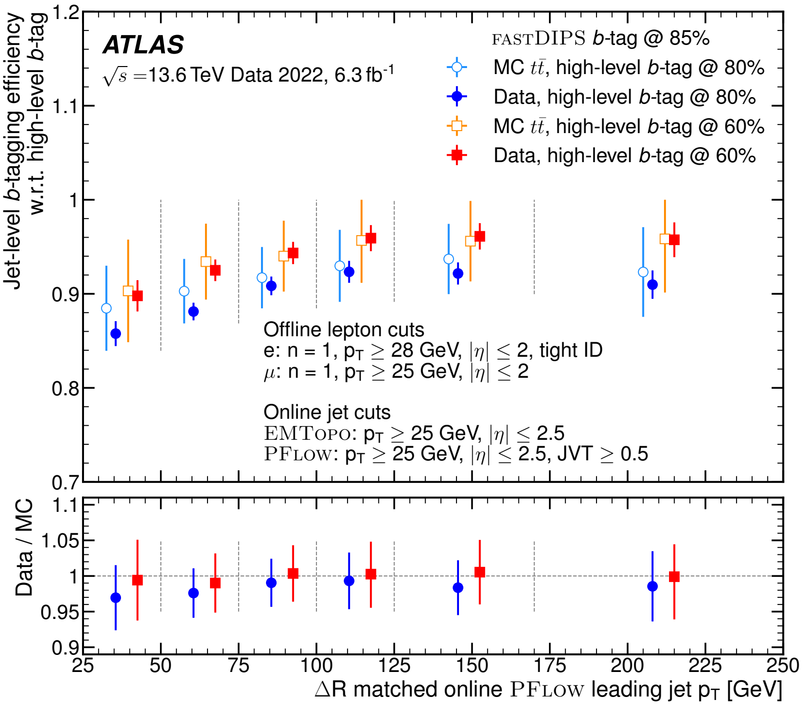

(a)

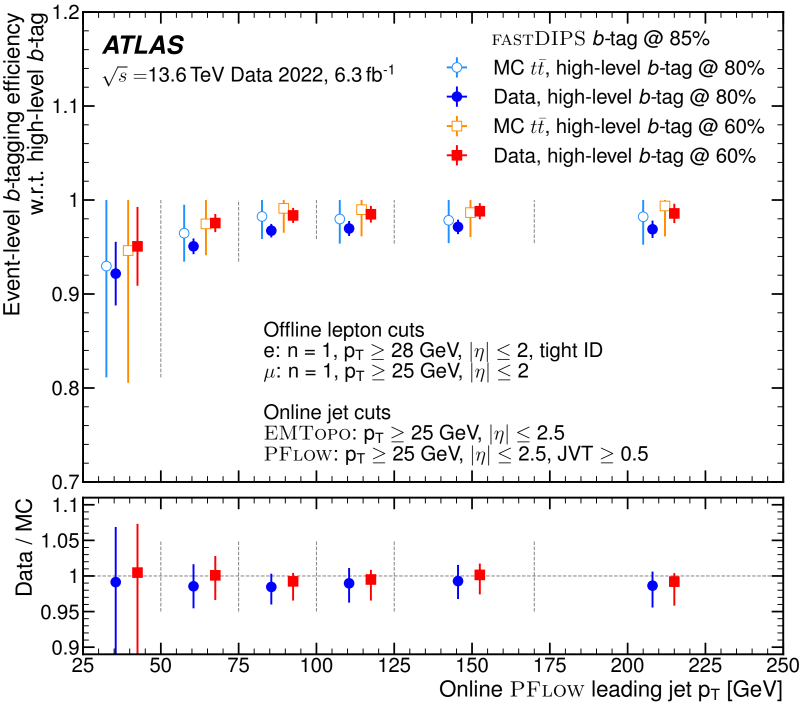

(b)

Figure 6.8: The (a) jet-level and (b) event-level efficiencies – defined by Equations 6.2 and 6.3, respectively – observed in data and simulation, in $t\bar{t}$ enriched events. The two plots have the same binning, bin edges are displayed with the vertical dashed lines. For the jet-level efficiency, the leading PFLOW jet is geometrically matched to an EMTOPO jet. For the event-level efficiency no matching is imposed. In both figures, the displayed uncertainties are statistical only. In the ratio panels, the errors are propagated as the quadratic sum of the statistical uncertainties.

introduced for this scope:

$$\text{jet-level b-tagging efficiency} = \frac{\text{\# jets passing both fast and high-level b-tagging}}{\text{\# jets passing high-level b-tagging}} \,, \tag{6.2}$$

and

$$\text{event-level b-tagging efficiency} = \frac{\text{\# events with both} \geq 1 \text{ fast and} \geq 1 \text{ high-level b-tag}}{\text{\# events with} \geq 1 \text{ high-level b-tag}} \,. \tag{6.3}$$

The former is based on jet counting and requires EMTopo and PFlow jets to be matched. The same matching procedure adopted for the conditional b-jet efficiency is applied. The event-level efficiency, by contrast, is determined by counting events, without geometric matching, and thereby more closely mimicking how the actual preselection mechanism works at HLT.

Figure 6.8 shows the efficiencies observed in data and MC. The fastDIPS working point studied is 85%, which is the tightest in the trigger menu. Looser cuts are expected to have higher efficiencies, as seen in Figure 6.6(b). The high-level b-tagger working points tested are 60% and 80%, covering the range that includes most of the cuts used in the trigger menu. Figure 6.8(a) presents the jet-level efficiencies. A slight dependence on p_T is observed, with efficiency increasing for harder jets. Figure 6.8(b) shows the event-level efficiency. In this case, a strong overlap is seen between the fastDIPS preselection and the DL1d precision tagger, with efficiencies ranging from 93% to 99%. Furthermore, for both efficiencies, excellent agreement is observed between data and simulation. The largest discrepancies are close to 3%, well below the statistical uncertainty.

Finally, the stability of the fast b-tagging preselection was studied under different pile-up conditions, using the same $\mathrm{t\bar{t}}$ phase space from previous studies. Figure 6.9 illustrates the average number of b-tagged jets per event as a function of the fastDIPS discriminant, across three different bins of $\langle\mu\rangle$ (defined in Equation 2.4). For loose cuts on fastDIPS D_b, the number of mis-tagged jets slightly increases with higher pile-up, but the b-tagged multiplicity remains consistent within $\sim 7\%$. Tighter selections rapidly reduce dependence on pile-up. At the 85% working point, the difference between the lowest and highest $\langle\mu\rangle$ bin is less than 3%, confirming the robustness of the method.

6.4 Menu improvements

Fast b-tagging has significantly impacted the b-jet trigger chains. Multijet signatures at low p_T must contend with the enormous rates of QCD events. For multi-b-jet processes, such as fully hadronic $\mathrm{t\bar{t}}$ decays or $\mathrm{HH} \to \mathrm{b\bar{b}b\bar{b}}$, flavour tagging strongly suppresses the background. However, the L1 selections are flavour-agnostic, and, therefore, the b-jet trigger reconstruction would need to run at high rates without any preselection at HLT.

The $\mathtt{2b2j}^{asym}_{\mathrm{HH}}$ trigger chain from Table 5.1 was explicitly designed for Run 3 to maximise

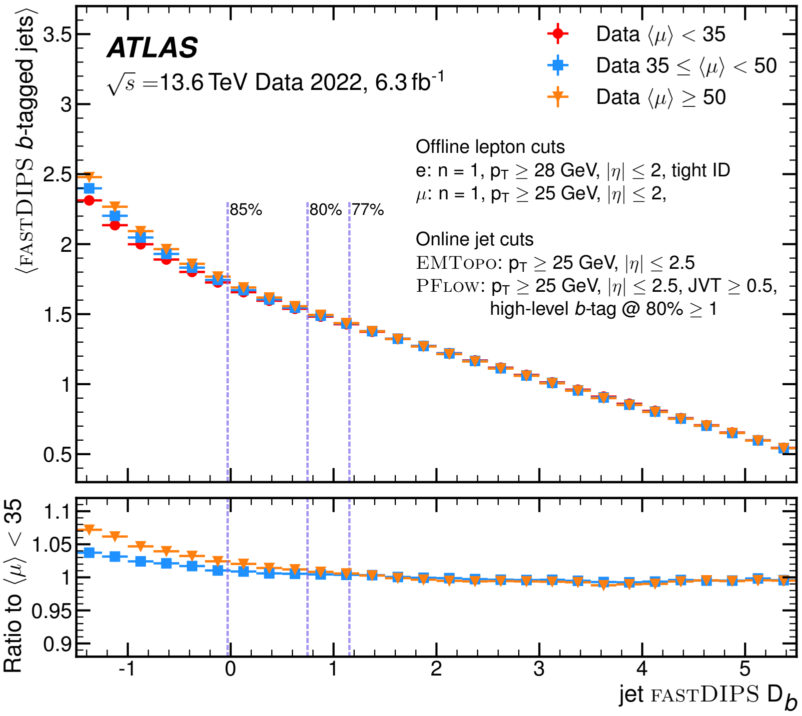

Figure 6.9: Average number of b-tagged jets per event as a function of the threshold on the preselection fastDIPS discriminant, in $t\bar{t}$ enriched data. The average is over the events in three bins of $\langle\mu\rangle$, which is the average number of p-p interactions per bunch crossing. During the considered data-taking period, $\langle\mu\rangle$ ranges roughly from 20 to 60. Statistical uncertainties, reported as vertical lines, are smaller than the marker sizes, and therefore not visible.

$HH \to b\bar{b}b\bar{b}$ acceptance, and its loose L1 selection is crucial for this purpose. However, this seed leads to a rate of roughly 8 kHz, which is not sustainable for event-wide tracking. Mitigating this burden at HLT was a key motivation for fast b-tagging.

An essential characteristic of fast b-tagging is its high signal acceptance, as it should not negate the benefit achieved from inclusive L1 triggers. This was verified by testing two relatively tight working points, using $HH \to b\bar{b}b\bar{b}$ as a benchmark. The results are reported in Table 6.3, with background rates estimated from the enhanced bias dataset [190]. The preselection can reduce the event rates considered for precision reconstruction by up to a factor of 10, largely satisfying the HLT CPU constraints. This reduction has a negligible effect on the signal, which undergoes only a 2-4% acceptance loss relative to the same

Table 6.3: Impact of the fast b-tagging preselection on background rejection and $HH \to b\bar{b}b\bar{b}$ signal acceptance. The $\mathtt{2b2j}^{asym}_{HH}$ chain from Table 5.1 is tested for two working points of fastDIPS.

Trigger selection	Preselection rejection factor on top of L1	$HH \to b\bar{b}b\bar{b}$ relative trigger acceptance
$\mathtt{2b2j}^{asym}_{HH}$ [preselection at 85%]	~ 5	0.98
$\mathtt{2b2j}^{asym}_{HH}$ [preselection at 80%]	~ 10	0.96

trigger chain with no preselection. This fraction is marginal compared to the improvement grained from the loose p_T selections, as demonstrated by the plots in Figures 5.5 and 5.6.

Overall, the preselection strategy has enabled substantial improvements for di-Higgs analyses relative to Run 2, but also for other analyses. Almost half of the b-jet trigger chains implemented in the menu use a fast b-tagging preselection [3]. Furthermore, the implementation of flavour tagging in the early stages of HLT enabled the augmentation of the TLA stream with b-tagging. All reconstructed PFlow jets at HLT are decorated with the TLA-specialised version of the fast tagger. An analysis searching for di-b-jet resonances using the TLA dataset has already been initiated, and is currently exploring a phase space that was inaccessible before Run 3 due to trigger limitations.

Outlook I – Further menu improvements

The fast b-tagging preselection has proved to be a valuable triggering strategy, which enhanced the ATLAS trigger menu for Run 3. Since its first implementation in 2022, the preselection was upgraded over the years. A PV finding algorithm is scheduled in the super-RoIs, with negligible impact on reconstruction timing but significantly refining the quality of impact parameter determination, especially in the longitudinal direction. Furthermore, the FASTDIPS architecture was upgraded with a dedicated GN2 training, which is running in the trigger since 2024. The more powerful algorithm together with the refined reconstruction greatly improved fast b-tagging, increasing the rejection factor by more than a factor of three for the $\mathtt{2b2j}_{\mathrm{HH}}^{\mathrm{asym}}$ chain preselection.

Similarly, the TLA b-taggers also benefited from moving to more expressive ML models. However, to further enhance flavour tagging in the TLA, a new data acquisition stream was set up for the beginning of 2025, the FTAGPEB stream. This data-collection technique combines TLA with Partial Event Building (PEB), meaning that raw data from partial detector is written to disk along with the HLT reconstructed objects. Trigger objects are used to identify RoIs, and only information from a list of sub-detectors of interest is readout in such RoIs. Following this principle, the event size is reduced, and offline reconstruction can be performed in limited areas of the detector, improving the precision relative to trigger objects.

The PEB strategy has been used since Run 1, but it was only considered for specific tasks such as detector calibration or noise monitoring [2, 97, 98]. It has never been used for physics analyses. The FTAGPEB data stream aims to improve flavour tagging for TLA by storing the hit information from the ID system in RoIs seeded by HLT PFLOW jets. This allows for further offline-like processing, employing CPU-intensive algorithms for tracking. Inference of the offline GN2 tagger, described in Section 4.2, is then run, achieving the state-of-the-art performance for flavour tagging in ATLAS and profiting from a centralised tagger calibration.

To validate the FTAGPEB workflow, estimate its event size, and assess its performance, a low rate (~ 10 Hz) trigger chain featuring this data-acquisition method was added to the 2024 menu. The size of events does depend on the number of saved RoIs. Considering that the test trigger chain requires at least four jets with $p_{\mathrm{T}} > 20$ GeV, and that only

RoIs for the six leading jets are stored, the observed additional 120 kB per event compared to the TLA data format implies that FTagPEB extends the event size by slightly more than 20 kB per RoI. This is still considerably smaller than the standard events (~ 2 MB); hence, as for TLA, lower trigger thresholds than the Physics Main stream are affordable.

Reconstruction of FTagPEB data was assessed by geometrically matching HLT and offline jets from a selected run, corresponding to almost 1 fb^{-1}. Figures 6.10(a) and 6.10(b) compare the standard HLT, FTagPEB, and offline tracks associated with trigger and offline PFlow jets that are geometrically matched ($\Delta R < 0.4$). It is visible that the PEB approach does benefit from the low p_T threshold for track seeding, resulting in roughly the same number of tracks per jet as offline. By contrast, TLA tracks (full scan FTF tracks) have a $p_T >$ 1 GeV cut on triplet seeds, and consequently there are slightly fewer tracks per jet. Figure 6.10(c) further confirms the higher efficiency of PEB tracking compared to the full scan fast tracks. Figure 6.10(d) demonstrates that, besides an efficiency increase, the tracking resolution is improved with PEB. When feeding these tracks into GN2, flavour tagging performance is highly correlated to the offline one, as illustrated in Figure 6.10(e). Since the PEB workflow cannot be yet emulated in MC simulations, a Gaussian smearing factor ($\mu \simeq 0$ and $\sigma \simeq 0.5$) was extracted from the distribution of the b-tagging discriminant difference between offline and FTagPEB jets, shown in Figure 6.10(f). Applying such smearing to offline jets in $t\bar{t}$ MC yields expected background rejection rates that are within 10% of the non-smeared ones and that are between a factor 1.5 and 2 better than the best TLA tagger.

After these preliminary validation and performance studies that confirmed the efficacy the PEB approach, the FTagPEB stream will start to acquire data for the first time in 2025 using a soft multi-b-jet trigger chain. This promising technique might be able to further improve the discovery potential of TLAs.

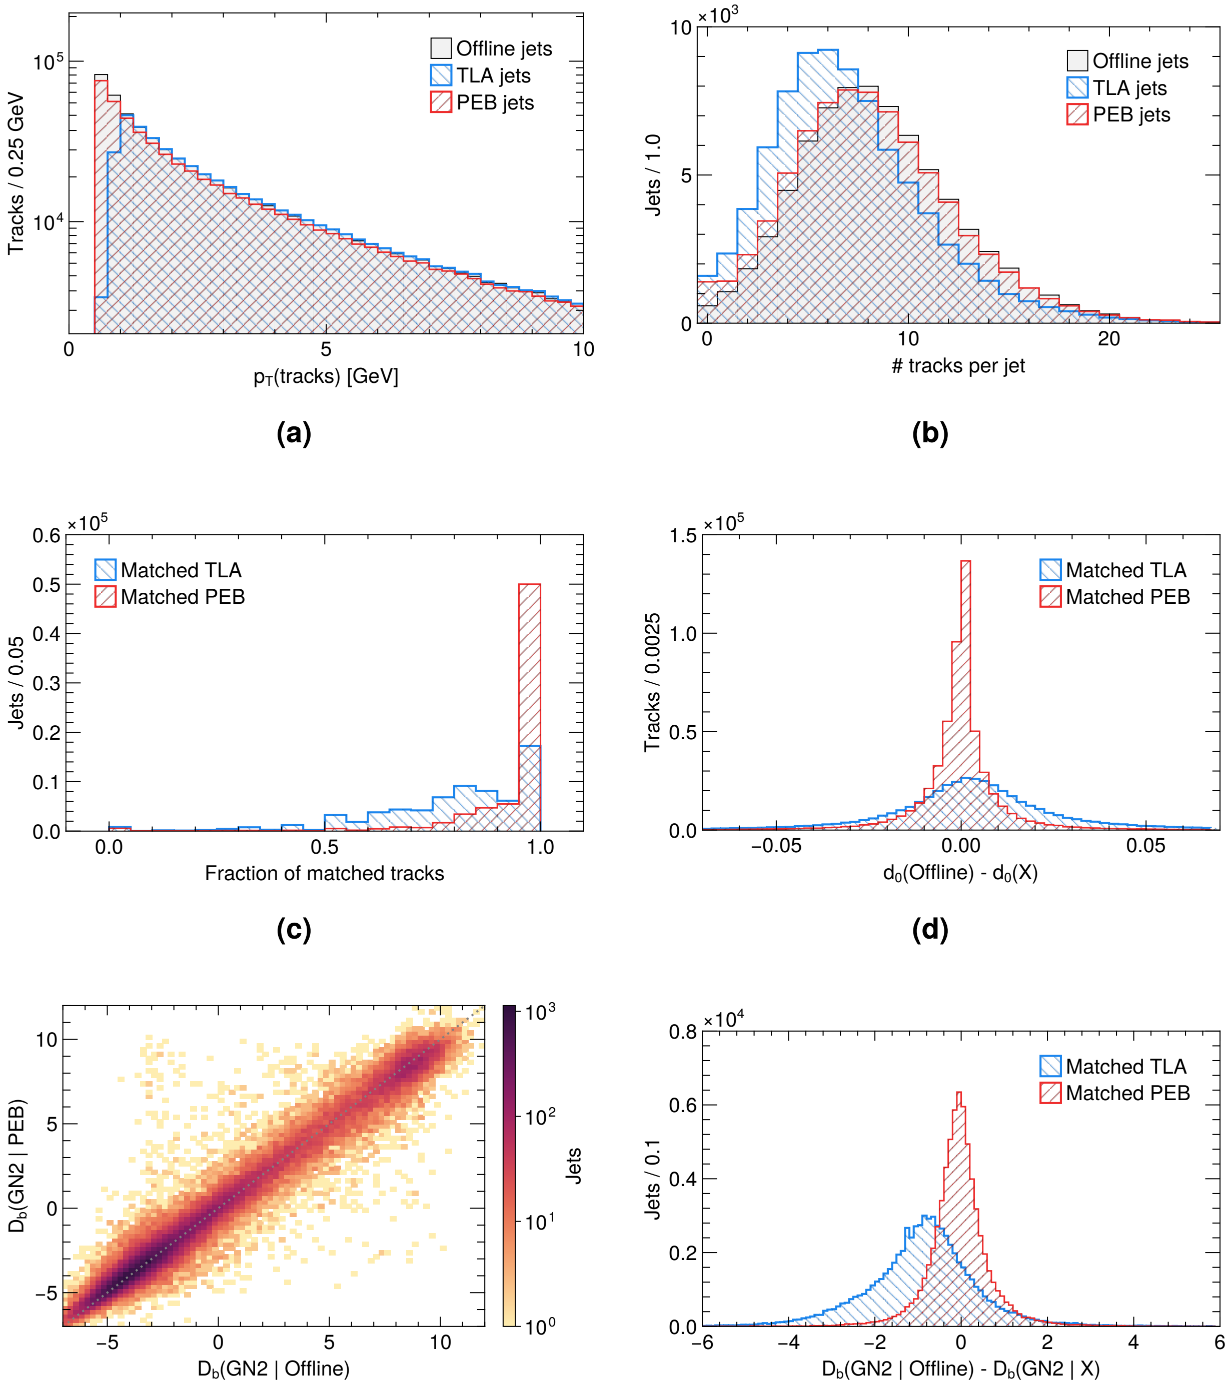

(a) (b)

(c) (d)

(e) (f)

Figure 6.10: Validation of the FTagPEB stream done using data event matched between FTagPEB, TLA, and Physics Main stream. The PFlow HLT jets are geometrically matched ($\Delta R < 0.4$) to offline ones. Figures (a) and (b) show the distribution of tracks p_T and the number of tracks associated with jets in three reconstruction procedures. Figure (c) shows the offline fraction of offline tracks per jet that are matched ($\Delta R < 0.1$) to a TLA or PEB one, while Figure (d) shows the difference in d_0 for matched tracks. The correlation of the GN2 discriminant obtained with offline and PEB tracks is evaluated in Figure (e), and Figure (f) shows the distribution of the delta between offline and PEB/TLA GN2.

III

Searching for heavy resonances with many jets

Chapter 7

Multijet analysis overview

As discussed in Chapter 1.3, RPV-SUSY offers a vast range of models with relaxed experimental constraints compared to scenarios assuming *R*-parity conservation. Interactions mediated by the λ'' coupling connect only quarks and squarks, producing fully hadronic final states. These final states, lacking any p_T^{miss} or leptons, are very hard to simulate and to constrain since they are largely dominated by QCD multijet events. This source of background has an enormous cross-section at the LHC and is remarkably complex to predict due to the intrinsic nature of QCD interactions.

In the following chapters an analysis searching for pair-produced gluinos decaying into multijet final states [4] is described. The search strategy follows a two-fold approach. The first is a Cut-and-Count Approach (CCA) also named *Jet Counting* method. The second is the *Mass Resonance* approach, and leverages Machine Learning (ML) to reconstruct the gluino decays. This chapter focuses on the general shared aspects between the two methods, details of the two strategies are described in later chapters.

7.1 Signal models

In simplified RPV frameworks, gluinos can decay to a fully hadronic decay through two different processes. Assuming that the squark masses ($m(\tilde{q})$) are much larger than the gluino mass ($m(\tilde{g})$), the gluino could decay into a quark and an off-shell squark. The squark would promptly generate two additional quarks via the λ'' coupling. This process results in a *direct* decay of the gluino into three quarks, $\tilde{g} \to qqq$. An alternative transition can take place instead if $\tilde{\chi}_1^0$ is the lightest supersymmetric particle (LSP) and therefore $m(\tilde{\chi}_1^0) < m(\tilde{g})$. In this case, gluinos can *cascade* decay via intermediate neutralinos, which would promptly decay via the RPV coupling, $\tilde{g} \to qq\tilde{\chi}_1^0(\to qqq)$.

The production of gluinos at the LHC is mediated by strong interactions that are *R*-parity conserving, implying that they must be created in pairs. Furthermore, given the high cross-section of such interactions, gluinos (and squarks) are expected to be the most abundant SUSY particles produced.

The presented search targets the topologies shown in Figure 7.1. Given the multiplicities of particles in the final states, the direct and cascade decays are also referred to as the 2×3 and 2×5 model, respectively. The high number of produced quarks directly indicates that signal events contain many energetic jets, due to quark hadronisation.

Two types of λ'' interactions can be studied with the simplified models from Figure 7.1. In fact, λ''_{ijk} represents 18 possible couplings. The i, j, k indices correspond to the generation numbers of the $\bar{U}_i\bar{D}_j\bar{D}_k$ superfields from Equation 1.16. Here only λ''_{112} and λ''_{113} are considered, which result in the production of uds and udb quark terns in the final state. For this reason, λ''_{112} and λ''_{113} are also referred to as the UDS and the UDB couplings. The conclusions derived from this search can be directly extended to other $\lambda''_{i,j,2}$ and $\lambda''_{i,j,3}$ couplings with $i = 2$ and $j = 2, 3$, since they would lead to identical detector signatures to the UDS and UDB models. If instead $i = 3$, a top quark would be produced in the

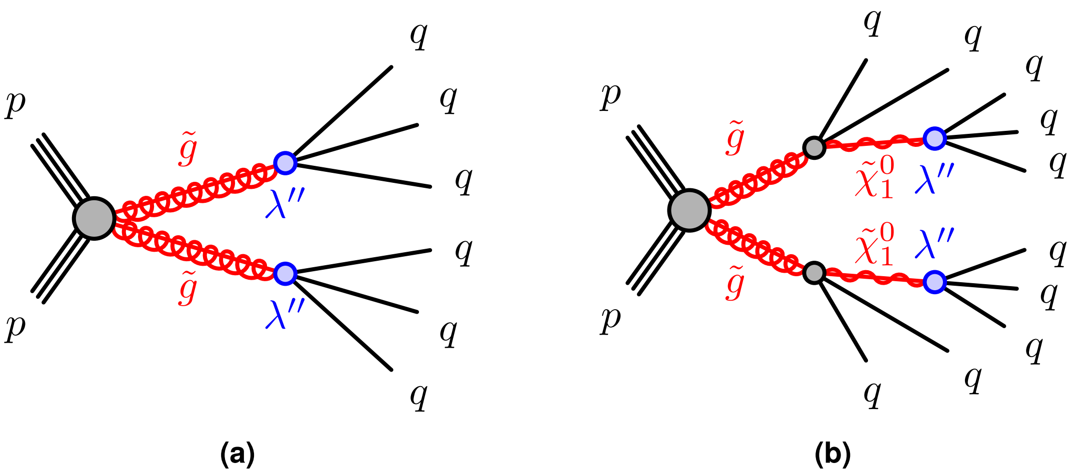

(a) (b)

Figure 7.1: Diagrams for pair-produced gluino decays via the RPV λ'' coupling. The direct gluino decay is shown in (a), while (b) is the diagram for the cascade decay through the intermediate $\tilde{\chi}_1^0$.

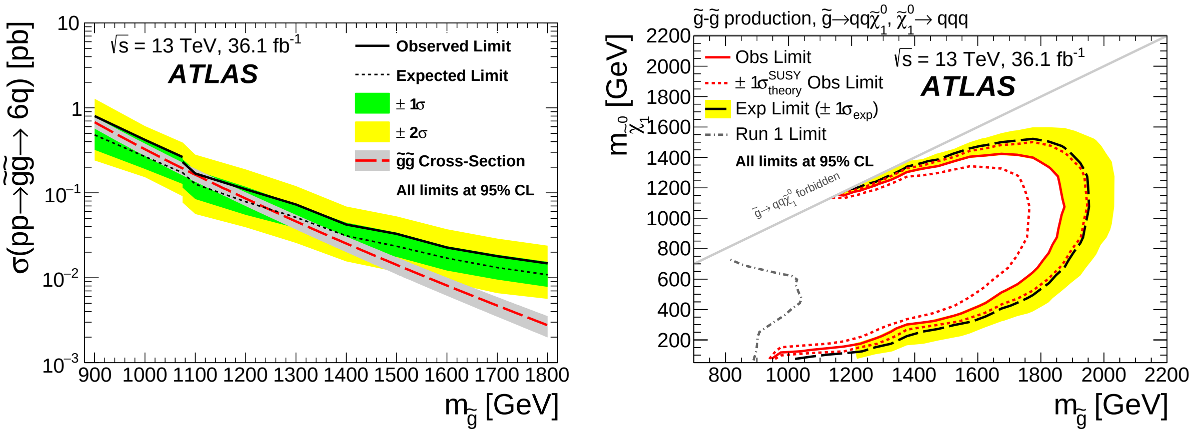

(a) **(b)**

Figure 7.2: Exclusion limits obtained by the partial Run 2 ATLAS analysis [204]. The upper limits on the cross-section for the gluino pair-production are shown in (a) for the direct decay (2×3 model). The exclusion contours in the $m(\tilde{g})$ and $m(\tilde{\chi}_1^0)$ phase space are shown in (b) for the cascade decay (2×5 model).

decay, modifying both the final state composition and kinematics. The models under consideration cannot be generalized to such conditions, requiring a separated study. From the experimental perspective, the major difference between UDS and UDB models is the formation of two b-jets in the detector. This signature provides an extra handle to deal with background rejection.

Previous searches have looked for gluino multijet decays, starting from CDF [197] to the more recent ones from CMS [198–201] and ATLAS [202–204]. The latest result from ATLAS obtained with a partial Run 2 dataset reached sensitivity up to $m(\tilde{g}) \sim 1$ TeV for the 2×3 models and $m(\tilde{g}) \sim 1.8$ TeV with $m(\tilde{\chi}_1^0) \sim 1$ TeV for 2×5 models, as shown in Figure 7.2. Gluinos in the range of 1 TeV $< m(\tilde{g}) < 1.5$ TeV are left uncovered, in the regime of light neutralinos ($m(\tilde{\chi}_1^0) < 400$ GeV). This gap makes an additional strong motivation for better investigating multijet final states. From a *natural* SUSY perspective, this mass range is of great interest. This thesis aims to improve the sensitivity to heavier gluinos in both direct and cascade decay modes, covering the gaps left by previous observations.

The direct and cascade decay models are treated separately in the above-mentioned searches. However, they act as two concurrent decay modes in extended theory frameworks. Their relative BR depends on: the intensity of the λ'' coupling, the masses of the virtual squarks mediating the RPV decay, and mass difference between $\tilde{g}$ and $\tilde{\chi}_1^0$. Figure 7.3 shows a map of points in the $\lambda'' - m(\tilde{g})$ plane transformed to values of BR($\tilde{g} \to qqq$). The calculations to derive the map are performed with SPheno [205, 206] and SARAH [207], fixing $\tilde{\chi}_1^0$ as the LSP, with $m(\tilde{\chi}_1^0) = 200$ GeV, and the squark masses at 3 TeV. Direct decays are strongly favoured when $\lambda'' \sim 1$ or when the phase space available to produce the $\tilde{\chi}_1^0$ is limited, $i.e.$, when $m(\tilde{g})$ and $m(\tilde{\chi}_1^0)$ are close. For decreasing couplings ($\lambda'' < 0.1$), the BR($\tilde{g} \to qq\tilde{\chi}_1^0(\to qqq)$) becomes largely dominant. Reducing further λ'' gives the $\tilde{\chi}_1^0$

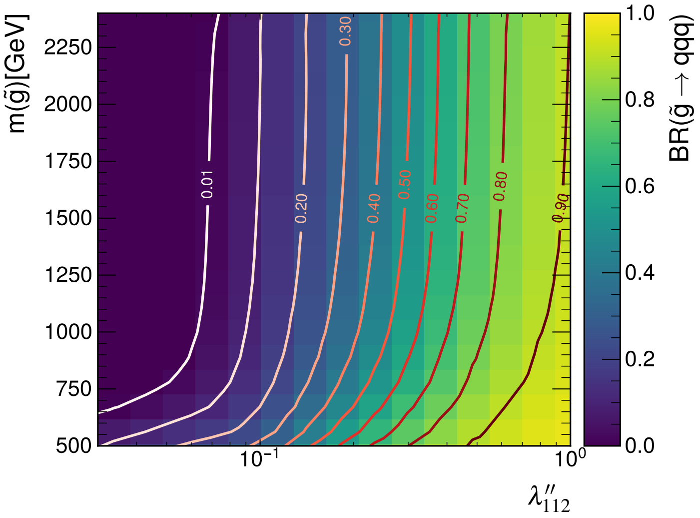

Figure 7.3: Mapping of the RPV coupling strength λ''_{112} and gluino mass to values of the gluino decay BR. It is assumed that the LSP is a bino-like $\tilde{\chi}_1^0$ with $200\,$GeV of mass, and that the other coloured sparticles have masses of $3\,$TeV.

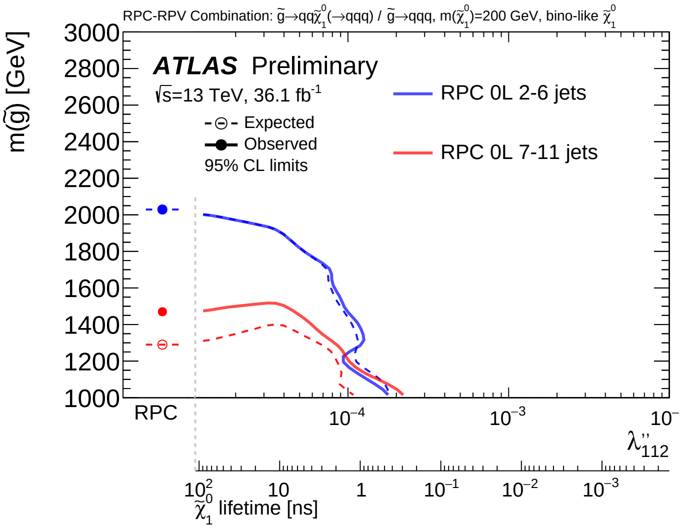

Figure 7.4: Exclusion limits for gluino as a function of λ''_{112} and $m_{\tilde{g}}$ [208]. Expected limits are shown with dashed lines, and observed as solid. The RPC-limit is shown on the leftmost part of the axes.

a lifetime (τ_{LSP}) that can be resolved experimentally as a displaced decay. The lifetime scales as $\tau_{\mathrm{LSP}} \propto 1/(\lambda'')^2$ and is correlated to $m(\tilde{\chi}_1^0)$ and $m(\tilde{q})$, but it is independent of $m(\tilde{g})$. Using the same computational tools as for the BR mapping, the lifetime is found to be nearly 1 ps for a coupling of $5 \cdot 10^{-3}$, and τ_{LSP} reaches 1 ns for $\lambda'' \sim 10^{-4}$. A regime similar to the RPC scenario appears for $\lambda'' < 10^{-5}$ ($\tau_{\mathrm{LSP}} > 100$ ns), where neutralinos decay out of the detector generating $p_{\mathrm{T}}^{\mathrm{miss}}$.

Reinterpretations of results from analyses focusing on RPC simplified models produced limits to RPV decaying gluinos [208]. As shown in Figure 7.4, signal presence is excluded for $m(\tilde{g})$ below $1.5\,$TeV with $\tau_{\mathrm{LSP}} \geq 1$ ns, and $m(\tilde{g})$ up to $2\,$TeV gluinos with $\tau_{\mathrm{LSP}} \geq 100\,$ns. These signals would produce a minimum amount of $p_{\mathrm{T}}^{\mathrm{miss}}$. However, models where $m(\tilde{g}) \geq$

1 TeV and $\tau_{\text{LSP}} < 1$ ns are left uncovered. This thesis' analysis is optimised on simplified models with prompt decays. However, its results can be reinterpreted to constrain models featuring low τ_{LSP} and high gluino masses ($m(\tilde{g}) \geq 1$ TeV).

7.2 Data and Monte Carlo samples

The analysed data was collected during the full Run 2 period of LHC activity, starting from 2015 and until the end of 2018. During this time, high intensity proton beams were colliding in ATLAS with a center-of-mass energy of 13 TeV. The instantaneous luminosity delivered by the LHC evolved during the years. Therefore, the average number of interactions per bunch crossing changed, as shown in Figure 7.5. In later years, higher pile-up was present. To cope with the changing environment, also the ATLAS trigger and detector configurations were slightly adapted during the years.

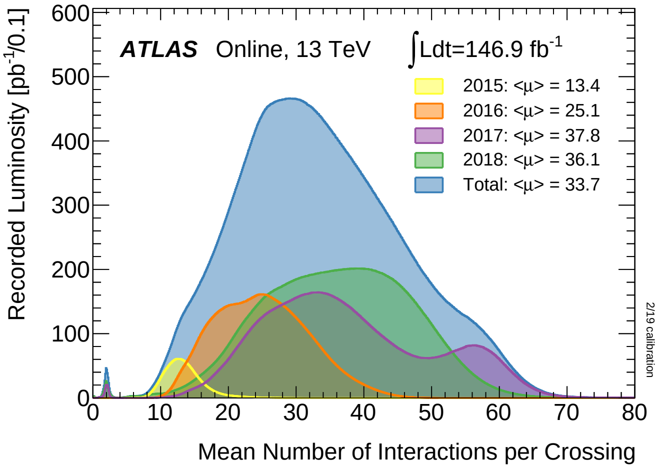

Figure 7.5: Recorded luminosity as a function of the average number of interactions, for each year of Run 2 data-taking [209].

MC techniques are utilised to replicate collision events, for both signal and background processes. These samples are needed for various tasks: the analysis R&D phase, the background estimate, and the statistical interpretation of the results. The simulations are divided into three campaigns, to match the different data-taking conditions during the years. The naming convention for each campaign and the corresponding year are reported in Table 7.1. All the MC samples were produced with these three settings.

Signal samples

For gluino production and RPV decay, the parton-level interaction is generated by computing the cross-section at Leading Order (LO) using MadGraph 5 [210], a matrix-element based event generator. To increase the precision and emulate contributions from Next-to-

Table 7.1: List of MC campaigns and corresponding years of data-taking.

MC Campaign	Data-taking Years
MC16a	2015, 2016
MC16d	2017
MC16e	2018

Leading Order (NLO), two additional partons are included in the matrix element computation. The parton showering is then simulated with Pythia 8 [211], using the A14 [212] set of tunes for the underlying event and using the NNPDF2.3LO [213] Parton Density Function (PDF) set. After the parton shower, the decay of hadrons is reprocessed with EvtGen [214] to improve the modelling of the heavy flavour decays. All previously described simulation steps take place in empty space. The last step consists in overlaying the ATLAS detector geometry and simulating the interactions with the active sensors and dead materials, and emulating the detector signals digitisation, using Geant4 [110].

For the 2×3 topology, gluinos with masses from 900 GeV and up to 2.5 TeV are generated, in steps of 100 GeV. A similar range is used for the cascade decay signal grid, but with wider steps of $\Delta m(\tilde{g}) = 200$ GeV . For each gluino mass point, neutralinos are generated spaced as well by 200 GeV, until the decay is kinematically allowed ($m(\tilde{\chi}_1^0) < m(\tilde{g})$). Two additional points are simulated at $m(\tilde{\chi}_1^0) = 50$ GeV and $m(\tilde{\chi}_1^0) = m(\tilde{g}) - 50$ GeV, to probe scenarios where the $\tilde{\chi}_1^0$ is extremely boosted, or conversely, produced almost at rest. Every mass point, for both decay modes, is generated individually for each of the two coupling schemes (UDS and UDB).

Two additional grids of MC signal samples are simulated to test models with different $\tilde{\chi}_1^0$ lifetime and variable BR between the two decay modes (direct and cascade). Pair-produced gluinos are generated in mass steps of 100 GeV, ranging from 1 TeV to 2 TeV. Each mass point is repeated for different neutralino lifetimes $\tau_{\text{LSP}} = \{0.01, 0.1, 1, 10, 100\}$ ns. For the prompt decays, samples are simulated with four values of BR: 0, 0.25, 0.75, and 1.

The inclusive cross-section estimates used for the normalization of the gluino pair-production [215] are considered with a high degree of precision, as shown in Figure 7.6. The values are derived using resummation of soft-gluon emission at Next-to-Leading Logarithmic (NLL) accuracy and with NLO SUSY QCD corrections.

Background samples

The largely dominant background for this search is multijet events emerging from QCD interactions. For UDB models, where a b-tagger is used to improve the sensitivity on signals, small contributions from top-quark pairs ($t\bar{t}$) are present, when both tops decay via hadronic channels. Other processes as γ+jets, W+jets, Z+jets, single-top quark or

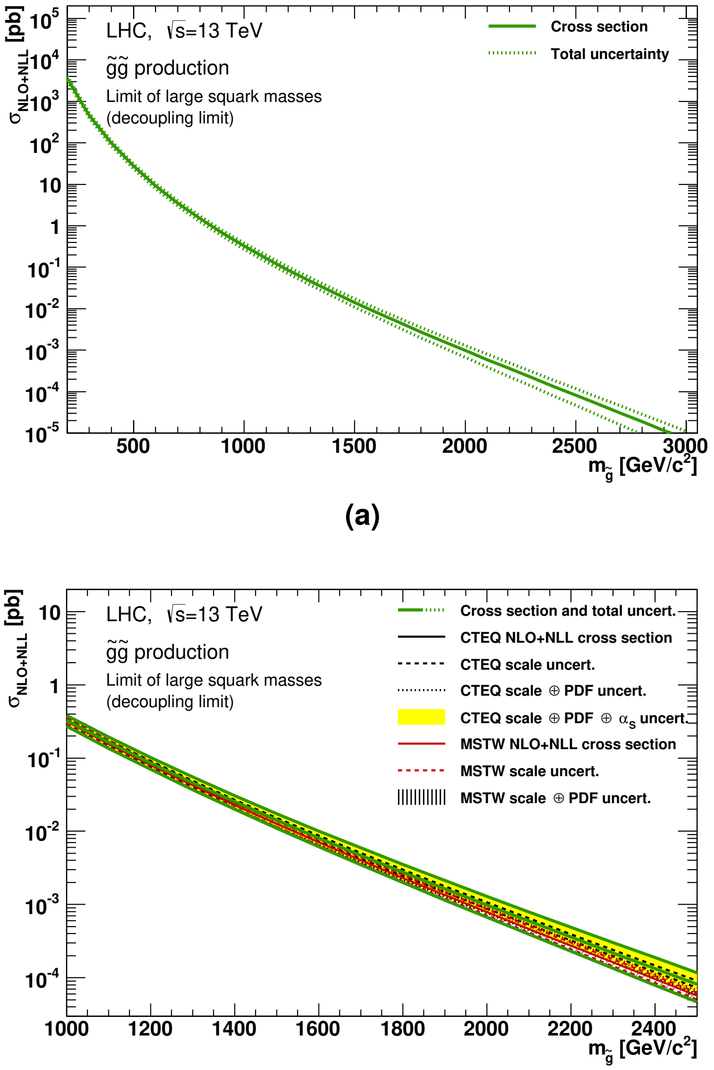

(a)

(b)

Figure 7.6: Cross-sections for gluino pair-production calculated at NLL+NLO accuracy [215]. Figure (a) shows the cross-section, and it's total uncertainty for a wide range of $m_{\tilde{g}}$. Figure (b) shows the cross-section for the mass range explored by this analysis. In the latter Figure, the estimate is broken down for the CTEQ6.6 [216] and MSTW2008 [217] PDF sets, with their respective uncertainties.

di-boson production are negligible. To design the search strategy and for part of to predict background (see Section 8.4), multijet and $t\bar{t}$ MC samples are used. These were centrally produced by the ATLAS Physics Modelling group, for the benefit of many analyses.

The QCD sample is generated at LO, as a $2 \to 2$ process [218]. PYTHIA 8 is used for both the matrix-element calculation and parton shower modelling. As for the signals, the A14 tunes and NNPDF2.3LO set of PDFs are used. The geometric mean of the squared transverse masses of the two outgoing particles

$$\hat{p}_T = \sqrt{(p_{T,1}^2 + m_1^2) \cdot (p_{T,2}^2 + m_2^2)}, \tag{7.1}$$

is used as the renormalisation scale. In order to populate the inclusive jet p_T spectrum efficiently, the leading jet p_T distribution is sliced and, within each slice, events are oversampled by a factor $(\hat{p}_T/10\,\text{GeV})^4$. The introduced bias is compensated by event weights that are inversely proportional to the sampling factor.

The $t\bar{t}$ events are produced with POWHEG BOX [219] at NLO accuracy. Additional samples created with alternative generators – MADGRAPH 5 interfaced with PYTHIA 8, and POWHEG BOX interfaced with HERWIG 7 [220] – are considered for the evaluation of systematic uncertainties associated with the MC techniques.

7.3 Analysis objects

Specialized versions of the physics objects from Chapter 3 are used for the search, as defined in the following paragraphs.

Jets Hadronic jets are reconstructed with the anti-k_T algorithm running on the PFLOW constituents – ID tracks and calorimeter clusters – and using a size parameter of $R = 0.4$. The minimum p_T is 20 GeV, while the angular acceptance is $|\eta| < 2.8$. A full JES and JER calibration chain is applied [221]. Softer jets, with $20\,\text{GeV} < p_T < 60\,\text{GeV}$ and $|\eta| < 2.5$, are considered only if they pass a "tight" working point of JVT (see Section 3.3). In this way, jets originating from pile-up interactions are suppressed.

B-tagging Jets within the ID acceptance ($|\eta| < 2.5$) are labeled as b-jets if they are successfully tagged by the DL1r [222] algorithm (see Section 4.2). This tagger is a fully connected NN, that combines as input few low-level taggers including RNNIP [223]. The used working point has a nominal efficiency of 77% on true *b*-jets and a false positive rate of 0.9% (20%) for light-jets (*c*-jets). Scale factors to correct for discrepancies in performances between data and simulations are also applied. These corrections are derived in a high purity $t\bar{t}$ data sample [224].

Leptons Events containing leptons are explicitly vetoed. Electrons are reconstructed from energy deposits in the LAr calorimeter (ECal), that are matched to an ID track [225]. Both the track and the clusters have to satisfy an isolation criteria, and shower-shape variables are used for a likelihood-based identification. Electron candidates are required to have at least $p_T > 10\,\text{GeV}$, $|\eta| < 2.47$ and have small longitudinal impact parameter $z_0 \sin\theta < 0.5$. Muon tracks are formed combining information from the ID tracker and MS. The derived track is selected as muon if it has $p_T > 10\,\text{GeV}$, $|\eta| < 2.7$ and passes "medium" identification requirement [226].

To avoid double counting of energy from calorimeter deposits and tracks, and to reduce the background from misidentified objects, an overlap removal procedure is performed. The algorithm is applied with the following order, to all the events:

- If two electron candidates share a track, the one with lower p_T is removed.
- If an electron shares its most energetic track with a muon, the former is removed.
- If an hadronic jet and an electron are within a cone of $\Delta R < 0.2$, the jet is removed.
- Electrons within $\Delta R < 0.4$ of the remaining jets are removed.
- When a jet and a muon are within $\Delta R < 0.2$ or the muon track is associated to the jet, the jet is removed if it has less than 3 tracks and its p_T is consistent with the muon energy.
- Finally, if a muon and a jet are within $\Delta R < 0.4$, the muon is removed.

The remaining objects passing these combined selections are used for the analysis.

7.4 Trigger and event preselection

Data events are selected using a trigger on high values of hadronic momentum in transverse plane (H_T). The latter is defined as:

$$H_T = \sum_{i=1}^{n} |p_{T,i}^{(\text{jet})}|, \tag{7.2}$$

where n corresponds the total number of jets in the event, with $p_T > 50$ GeV. The lowest un-prescaled H_T trigger available in the menu for each year was used, as listed in Table 7.2. For all years the same L1 requirement is used – a single L1 jet above 100 GeV. When reconstructed and calibrated offline, these jets typically correspond to double the energy in the hadronic scale. The HLT thresholds change over the years, the `ht850` selection got prescaled after 2015 due to too high rates, and therefore a tighter threshold is used.

To ensure that the trigger selection is fully efficient, a minimal analysis preselection is applied. The H_T must be above 1.1 TeV and the event should contain at least four jets with

Table 7.2: List of trigger chains used for the different years.

Years	Trigger Chain
2015	`HLT_ht850_L1J100`
2016, 2017, 2018	`HLT_ht1000_L1J100`

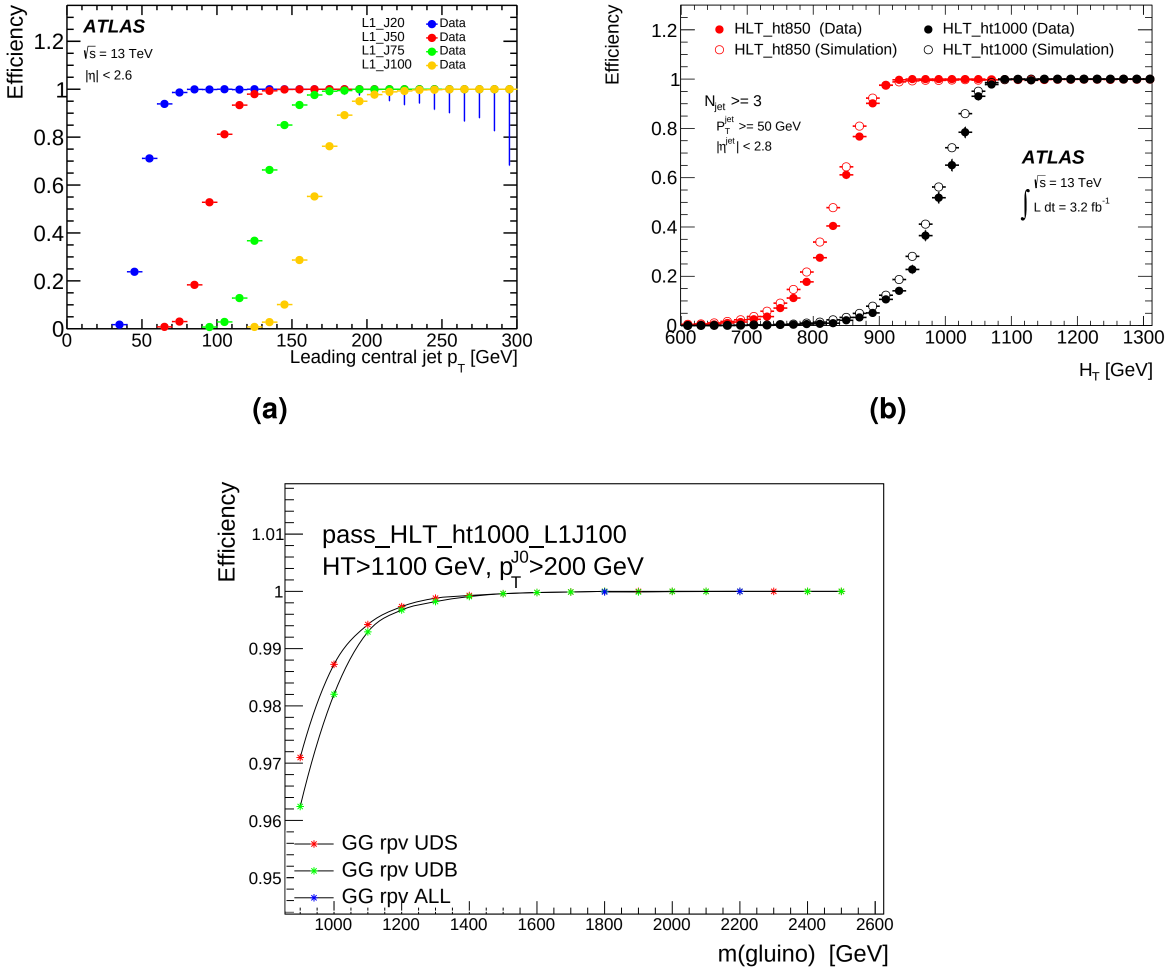

(a) **(b)**

(c)

Figure 7.7: Trigger efficiency turn-on curves (a) for L1 single-jet triggers in central region, and (b) for HLT H_T triggers as a function of the offline reconstructed H_T [227]. Figure (c): trigger efficiency on signal events, as a function of the gluino mass hypothesis. The preselections to ensure trigger plateau are applied.

$p_T > 50$ GeV, with the most energetic one having $p_T > 200$ GeV. These requirements are taken from a dedicated performance study [227], its relevant results are shown in Figures 7.7(a) and 7.7(b). The plots demonstrate that the trigger plateau regime is reached for both data and MC with the required preselection. No trigger scale factors are therefore necessary to model any trigger inefficiency.

The selected triggers have a very high acceptance for signal events. From Figure 7.7(c), the efficiency of triggering gluino decays is above 95% for low masses and increase above 99% for $m(\tilde{g}) \geq$ 1.1 TeV.

Possible detector defects, data acquisition errors or any other artefact affecting the data-taking are masked by filtering with a *Good Run List*. These lists are compiled by the ATLAS Data Quality group. The multiple causes that can corrupt data are flagged for each luminosity block. After a veto on these, the dataset used in the analysis corresponds to a total of 140.1 fb^{-1} [228].

7.5 Statistical framework

The statistical analysis follows a frequentist approach, where probabilities are computed as the fraction of occurrences. For example, the probability associated with events of type x_i is defined as:

$$p(x_i) = \lim_{N \to \infty} n_i/N \,, \tag{7.3}$$

where n_i are the occurrences of x_i over the total number of observations N. The primary statistical tool employed is the likelihood function, which encodes the compatibility between a model, parametrised by the vector $\mathbf{p}$, and the observed data $\boldsymbol{x}$:

$$L(\mathbf{p}|\boldsymbol{x}) = \prod_i p(x_i|\mathbf{p}) \,. \tag{7.4}$$

This analysis adopts a binned likelihood function, where the expected number of events is expressed as the sum of signal (s) and background (b): $E[x_i] = \mu s_i(\boldsymbol{\theta}) + b_i(\boldsymbol{\theta})$. In the latter equation the parameter μ represents the signal strength. It is equal to one for the nominal signal hypothesis (*i.e.*, gluinos produced with the nominal cross-section value), and $\mu = 0$ for the background-only hypothesis. The systematic uncertainties are modelled using Nuisance Parameters (NPs), denoted by θ. With these conventions, the complete likelihood takes the form:

$$L\left(\mu, \boldsymbol{\theta} \,|\boldsymbol{x}\right) = \prod_i \text{Poisson}\left(x_i|\lambda = \mu s_i(\boldsymbol{\theta}) + b_i(\boldsymbol{\theta})\right) \prod_j C(\alpha_j|\theta_j) \,. \tag{7.5}$$

The terms $C(\alpha_j|\theta_j)$ are the constraints on the NPs derived from the auxiliary measurements α_i (*e.g.*, calibrations or procedures for estimating the uncertainties). The vector θ typically has a high dimensionality given the large number of systematic uncertainties considered. To reduce these degrees of freedom, the “profiling” method is applied by defining the *profile likelihood ratio* as follows:

$$\lambda(\mu) = \frac{L(\mu, \hat{\hat{\theta}})}{L(\hat{\mu}, \hat{\theta})} \,. \tag{7.6}$$

The likelihood for a specific value of μ is maximised by $\hat{\hat{\theta}}$, while $\hat{\mu}$ and $\hat{\theta}$ represent the global maximum likelihood parameter values.

Several test statistics can be constructed from the profile likelihood $\lambda(\mu)$. For the discovery of a signal that increases the total event counts, the q_0 test is defined as:

$$q_0 = \begin{cases} -2\ln\lambda(0) & \text{for } \hat{\mu} \geq 0 \,, \\ 0 & \text{for } \hat{\mu} < 0 \,. \end{cases} \tag{7.7}$$

In cases of data under fluctuations, when $\hat{\mu}$ is negative, q_0 is set to 0 indicating a perfect agreement with the background-only hypothesis. In contrast, if $\hat{\mu} > 0$ the test statistic q_0 compares the likelihood of the background-only hypothesis to the maximum L value. The associated p-value is defined as:

$$p_0 = p(q_0 \geq q_0^{\text{obs}}|b) = \int_{q_0^{\text{obs}}}^{\infty} f(q_0|\mu = 0)\, dq_0 \,, \tag{7.8}$$

where $f(q_0|\mu = 0)$ is the probability density function (pdf) of the test statistic. Therefore, the p_0-value quantifies the agreement of the data with the background-only hypothesis $\mu = 0$ in terms of probability. The statistical significance of a test is often expressed in units of its equivalent value for a unitary gaussian (Z). Some commonly defined thresholds are $Z = 3\sigma$ for "evidence" and $Z = 5\sigma$ for "discovery", which correspond to $p_0 = 0.003$ and $p_0 = 5 \cdot 10^{-7}$, respectively. To define the expected sensitivity of a blinded analysis, prior to looking at the data, it is useful to quote the expected significance assuming that the data is equal to $x = s + b$. Using the asymptotic approximation [229] for Poisson distributed data with Poisson auxiliary measurements, the discovery significance is given by:

$$Z = \sqrt{2\left[x \ln\left(\frac{x(b+\sigma^2)}{b^2 + x\sigma^2}\right) - \frac{b^2}{\sigma^2}\ln\left(1 + \frac{\sigma^2(x-b)}{b(b+\sigma^2)}\right)\right]} \,, \tag{7.9}$$

where σ is an estimate of the uncertainty on b.

To set upper limits on the μ parameter, a distinct test statistic is defined:

$$q_\mu = \begin{cases} -2\ln\frac{L(\mu,\hat{\hat{\theta}})}{L(0,\hat{\hat{\theta}}(0))} & \text{for } \hat{\mu} < 0 \,, \\ -2\ln\lambda(\mu) & \text{for } 0 \leq \hat{\mu} \leq \mu \,, \\ 0 & \text{for } \hat{\mu} > \mu \,. \end{cases} \tag{7.10}$$

In this case, for downward fluctuations of data where $\hat{\mu} < 0$, the most compatible value of the signal strength that has physical meaning is $\mu = 0$. For large over fluctuations, when $\hat{\mu}$ is greater than μ, low μ values are instead compatible with data allowing the presence of a small signal. Therefore, $\mu \leq \hat{\mu}$ is not considered for rejection ($q_\mu = 0$). Using the appropriate pdf, the probability $p_\mu = p(q_\mu \geq q_\mu^{\text{obs}}|\mu s + b)$ can be computed similarly to Equation 7.8. The standard convention for limit setting is to exclude signals at a Confidence Level (CL) of 95%, which corresponds to p-values smaller than 5%. Instead of using p_μ directly (also denoted as p_{s+b}), the probability of q_μ^{obs} under the background-only hypothesis $p_b = p(q_\mu \geq q_\mu^{\text{obs}}|b)$ is also taken into account, and the threshold for exclusion is set on:

$$\text{CL}_s \equiv \frac{\text{CL}_{s+b}}{\text{CL}_b} = \frac{p_{s+b}}{p_b} \leq 0.05 \,. \tag{7.11}$$

This method is known as the CL_s approach [230]. It is important to note that $p_b \neq p_0$, since q_μ and q_0 are distinct test statistics. With this modified CL_s approach, the confidence intervals on the upper limits are more robust in cases where the discrimination between signal and background is weak, particualrly when the signal is expected to have far fewer events than the background.

Chapter 8

Jet Counting approach

The Cut-and-Count Approach (CCA) was developed to have a robust analysis with reinterpretable results. In this way, is easy to extrapolate the outcome of the analysis to alternative models. Furthermore, CCA provides a solid benchmark against which more sophisticated methods can be compared. The primary benefit of this traditional strategy lies in its conceptual simplicity. It relies on the definition of phase space Signal Regions (SRs), and an estimate of background events within each region. Data events are then counted and compared to the predicted yields to determine the final results. Therefore, the two core aspects of this analysis approach are the design of an optimal SR that maximises the sensitivity to signal events, and the elaboration of a consistent and accurate background prediction method.

Given the enormous scale of the expected multijet background compared to the signal cross-section, and the complexities in QCD modelling, designing a search that remains rather simple while being sensitive to such rare processes becomes an outstanding challenge. To address it, the Jet Counting technique is used, improving the background method and introducing the event-shape variables for this type of search.

This chapter presents the steps developed for this thesis to perform the CCA analysis, excluding the final results. The latter are discussed in Chapter 10, together with their interpretation.

8.1 Characterisation of signal and background

The starting point to define the analysis strategy is the study and characterisation of events, based on jet-related features. The distributions of simple jet kinematics quantities such as p_T or η, and more elaborated event-level observables, are examined to identify potential handles for background discrimination. For simplicity, these studies assume that the sole background source is QCD multijet events. Background composition studies verify this assumption, which holds true for most of the probed phase space. Signal hypotheses with direct decays are initially considered, and the conclusion drawn are subsequently verified for the cascade decay models. To reduce the processing time needed to produce the histograms, at this stage only the MC16e campaign is used, and scaled to the full Run 2 luminosity.

Figures 8.1(a), 8.1(b), and 8.1(c) show the p_T of the three hardest jets in the event. The effect of trigger selection bias is visible in the background distributions. The event preselection requires more than 1 TeV in the transverse plane. Given that the QCD events are mostly dijet production with additional radiation, the leading jet peaks at ~ 500 GeV

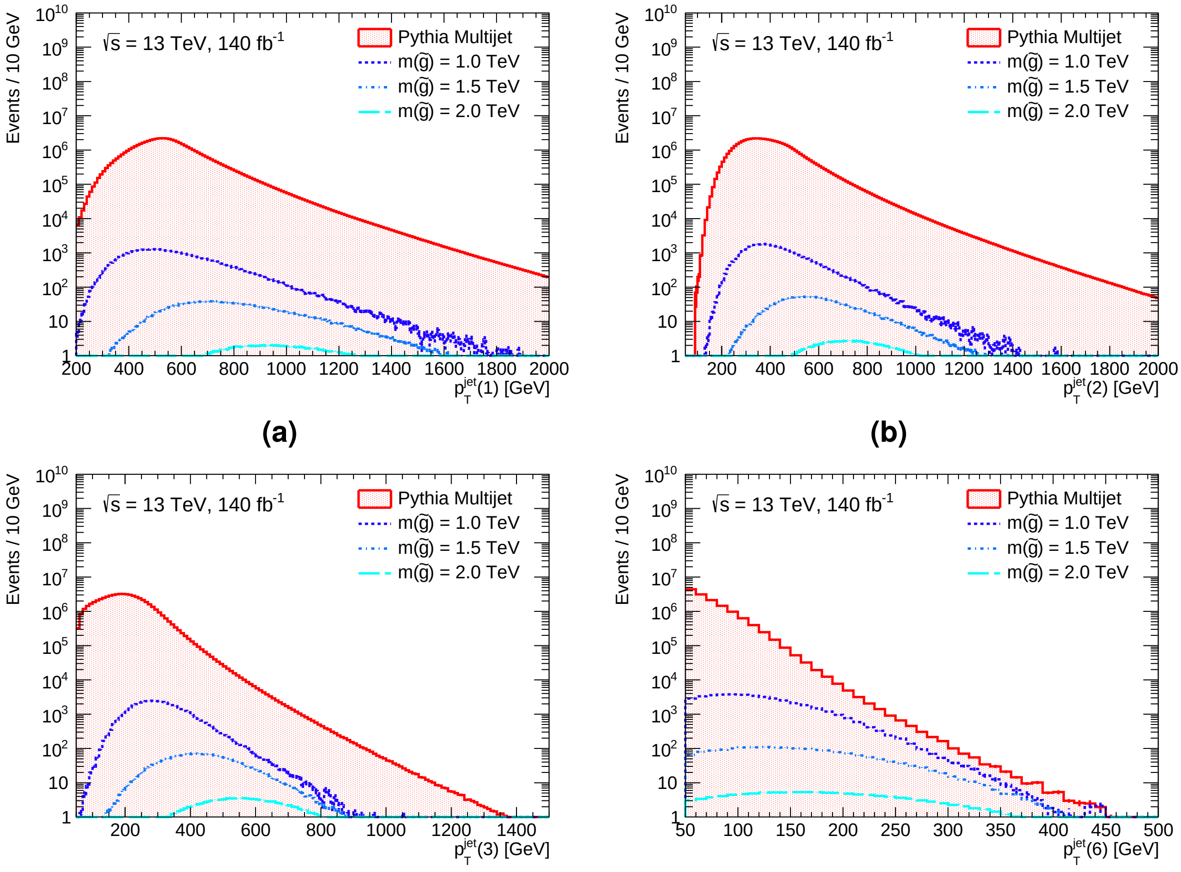

(a) **(b)**

(c) **(d)**

Figure 8.1: Distributions of p_T for the (a) first, (b) second, (c) third and (d) sixth leading jets. In red, the PYTHIA8 simulation of background events is shown. The blue lines represent three different gluino mass points. The trigger requirement and event preselection cut are applied.

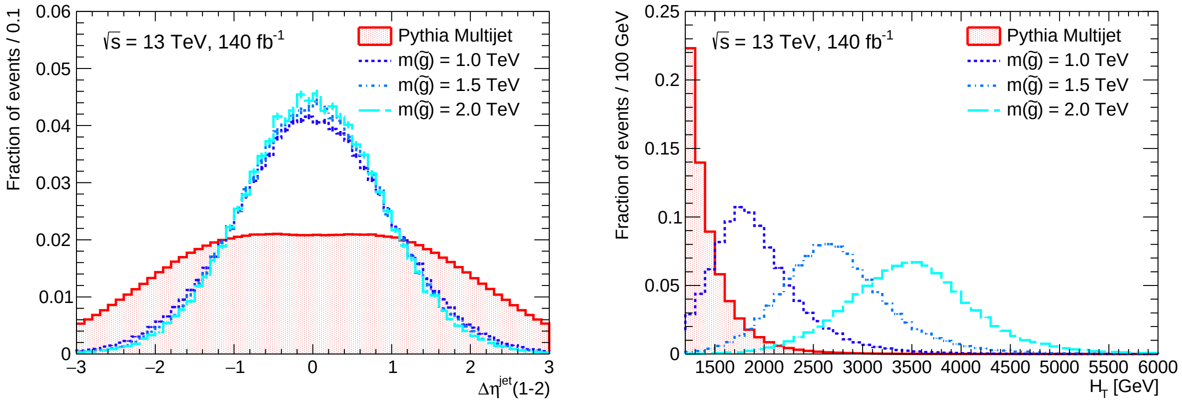

(a) **(b)**

Figure 8.2: Distributions normalized to unit area of (a) $\Delta\eta$ between the two leading jets, and (b) the total H_T in the event. In red the Pythia8 simulation of background events is shown. The blue lines represent three different gluino mass points. The trigger requirement and event preselection cut are applied.

and the subleading at ~ 400 GeV. In most cases a few outgoing particles retain the majority of energy. For signal events, the kinematic energy is more evenly distributed among the final state particles. Figure 8.1(d) illustrates the sixth-leading jet distributions. These distributions show that is more probable to observe high-p_T trailing jets in gluino decays than in the background, with a substantially larger signal-over-background ratio compared to previous plots.

Another remarkable aspect of Figure 8.1 is the difference in the event yields. The cross-section for gluino production is exceptionally smaller compared to QCD multijets. As shown in the previous chapter (figure 7.6), the probability of producing signal events decreases exponentially for higher $m(\tilde{g})$ values. To facilitate the comparison of distributions in a linear y-scale, subsequent figures in this section are scaled to unit area. Unless specified otherwise, the signal UDS samples are used, but they do generalise to UDB decays as their final state kinematics are identical.

Figure 8.2(a) reports the distribution of the distance in η between the two leading jets. For signal, jets have higher rates at low $\Delta\eta$. This is expected as the production occurs predominantly via the s-channel. QCD has instead important contributions for jet formation also in the u- and t-channel, causing a wider η distribution. The tendency for more central-jet events also results in higher transverse momentum values. In fact, large H_T is observed for signals in Figure 8.2(b). On average, almost twice the energy corresponding to the gluino mass flows in the transverse plane. For multijet, the trend is instead a decreasing exponential, which aligns with the behaviour of the individual jets shown in Figure 8.1.

High jet multiplicity is expected for signal events. Figure 8.3 displays how the number for jets per event evolves for increasing jet p_T thresholds. The preselection cut requiring four jets with $p_T \geq 50$ GeV is visible in Figure 8.3(a) in the background distribution. As transverse momenta increase, the multijet distribution rapidly shifts towards a lower av-

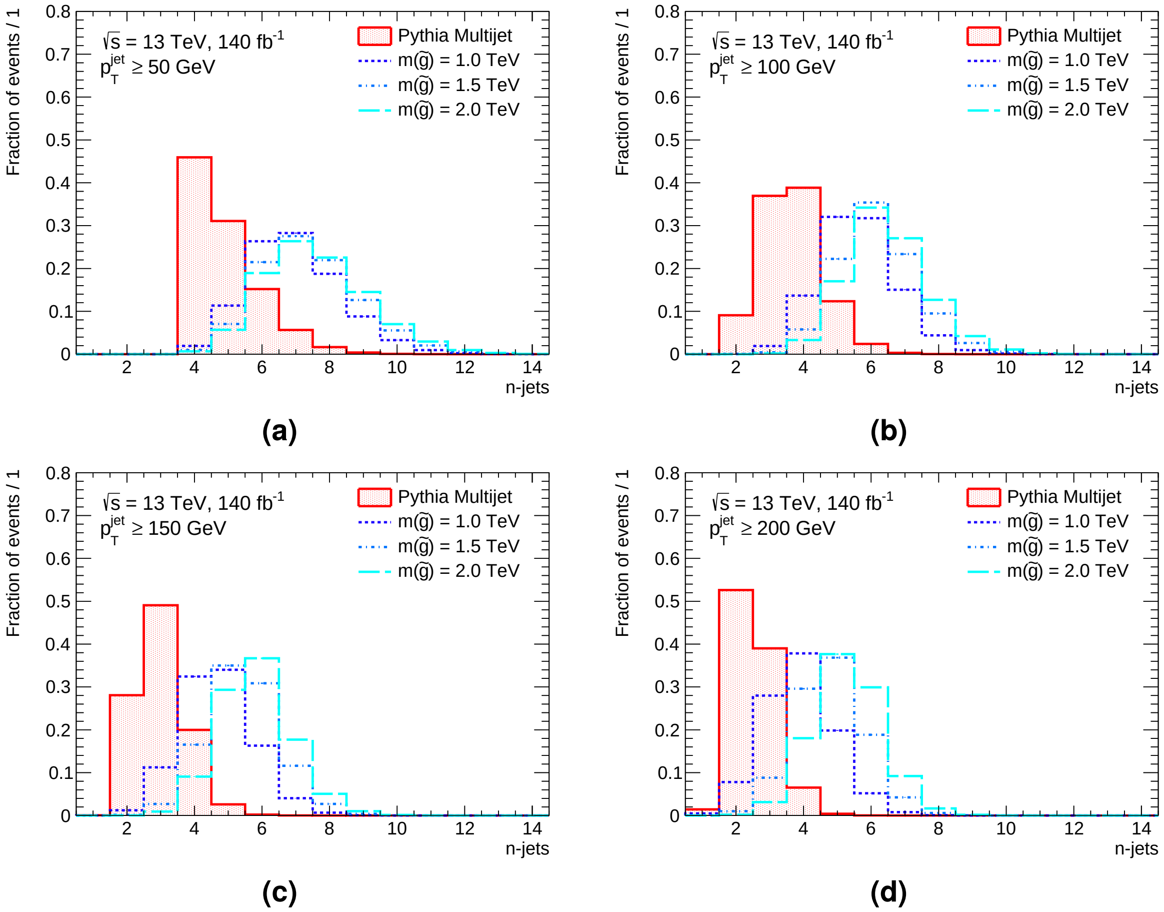

(a) (b) (c) (d)

Figure 8.3: Distributions normalized to unit area of the number of jets with p_T above (a) 50 GeV, (b) 100 GeV, (c) 150 GeV and (d) 200 GeV. In red, the PYTHIA8 simulation of background events is shown. The blue lines represent three different gluino mass points. The trigger requirement and event preselection cut are applied.

erage numbers of counted jets, becoming narrower. Conversely, gluino events maintain the average number of jets per event close to six. The relatively pronounced right tails of the distributions are associated with the generation of final state radiation, which is non-negligible given the sizeable number of hadrons involved in the hard-scatter process.

To capture the peculiar relation between the topology and energy of signal events, the event-shape variables are of great interest. These variables characterise the global structure of final states. They are defined such that they tend to vanish for $2 \rightarrow 2$ processes, but instead achieve their highest values for events with spherical distributed energy across four or more jets. Gluino candidates are pair produced on-shell in p-p collisions almost at rest. This results in close to isotropic spread of the energy in the decays, $i.e.$, high event-shape values. Conversely, QCD multijets are dominated by events with two back-to-back hard jets, and several additional softer jets, leading to event shapes closer to zero.

Using the same procedure of the ATLAS measurement of QCD multijet distributions at $\sqrt{s} = 13$ TeV [231], the event shapes observables can be defined starting from the linearized sphericity tensor [232]:

$$M_{xyz} = \frac{1}{\sum_i |\vec{p}_i|} \sum_i \frac{1}{|\vec{p}_i|} \begin{pmatrix} p_{x,i}^2 & p_{x,i}p_{y,i} & p_{x,i}p_{z,i} \\ p_{y,i}p_{x,i} & p_{y,i}^2 & p_{y,i}p_{z,i} \\ p_{z,i}p_{x,i} & p_{z,i}p_{y,i} & p_{z,i}^2 \end{pmatrix} . \tag{8.1}$$

The eigenvalues of this tensor have a direct geometrical interpretation. The normalization factor of the M_{xyz} tensor ($\frac{1}{\sum_i |\vec{p}_i|}$) makes its trace to be unity, $i.e.$ the eigenvalues (ϵ_i) satisfy: $\sum_i \epsilon_i = 1$. When two of the eigenvalues are zero, the third one must be equal to one. This is case for two back-to-back jet events. For 3-jet events, one of the eigenvalues is zero, and the remaining two will be equal to $\frac{1}{2}$ for a planar and symmetric arrangement of the jets. Instead, for events with more than three jets, if the spread of the momenta in the final state is close to spherical, the three eigenvalues will have similar values between each other, close to $1/3$. Considering the set of eigenvalues ordered such that $\epsilon_1 \geq \epsilon_2 \geq \epsilon_3$, the aplanarity A and sphericity S are defined as:

$$A = \frac{3}{2}\epsilon_3 \; ; \quad S = \frac{3}{2}(\epsilon_2 + \epsilon_3) \; . \tag{8.2}$$

The quadratic and cubic combinations of ϵ_i are also typically defined as:

$$C = 3(\epsilon_1\epsilon_2 + \epsilon_1\epsilon_3 + \epsilon_2\epsilon_3) \; ; \quad D = 27(\epsilon_1\epsilon_2\epsilon_3) \; . \tag{8.3}$$

The distributions of sphericity, aplanarity, C and D are represented in Figure 8.4. For signal events, the energy tends to evenly spread, featuring a preference for high event-shape values. For QCD, a large fraction of events has instead final states concentrating in a single plane, resulting in low sphericity and aplanarity.

The shape of C (also called C-parameter) for QCD can be explained with the following considerations. Given the eigenvalues' normalization, the maximum value of C for 3-jet events is

$$C_{\text{max}} = 3(1/2 \cdot 1/2) = 3/4 = 0.75 \; , \tag{8.4}$$

since ϵ_3 is zero. Therefore, $C > 0.75$ can only occur in events with more than three jets. For such events, the population of the high-C region depends on the p_T hierarchy of the jets. If the p_T of the two leading jets is much larger than the remaining jets, also the eigenvalues are unbalanced $\epsilon_1, \epsilon_2 \gg \epsilon_3 \simeq 0$. The value $C = 1$ is achieved ideally with a perfect spherical distribution of the momentum. For example, for 4-jet events, this would occur when all four jets have an exactly equal momentum magnitude and tetrahedral relative angles. This latter configuration is extremely unlikely in QCD multijet events, consequently the C shape has a steeply falling trend after the 0.75 value.

In the case of UDB models, the production of two b-jets is expected. For the tagger's 77% b-jet efficiency working point, roughly 60% of the events should have two b-tagged

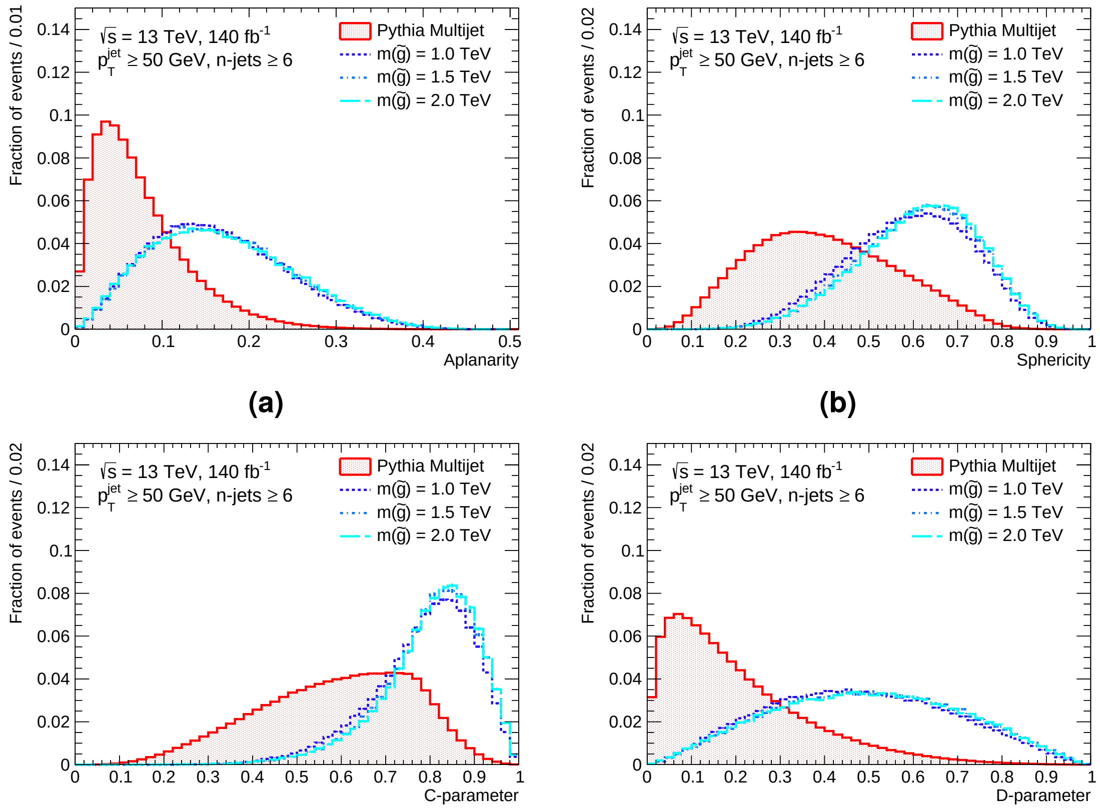

(a) **(b)**

(c) **(d)**

Figure 8.4: Distributions normalised to unit area of (a) aplanarity, (b) sphericity, (c) C and (d) D event shapes. In red, the PYTHIA8 simulation of background events is shown. The blue lines represent three different gluino mass points. The trigger selection and event preselection cut are applied, with an additional requirement of at least 6 jets per event.

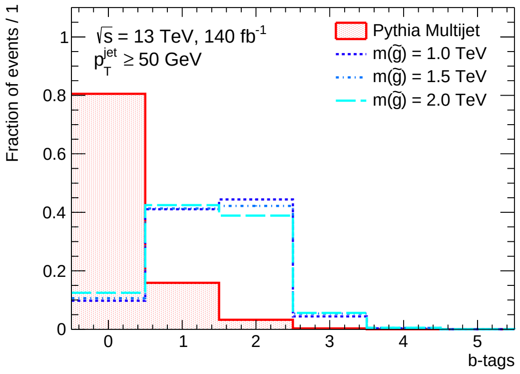

Figure 8.5: Distributions normalised to unit area of the number jets per event passing the nominal 77% b-jet efficiency DL1r cut. In red the PYTHIA8 simulation of background events is shown. The blue lines represent three different gluino mass points, decaying via the UDB coupling. The trigger requirement and event preselection cut are applied.

jets. Figure 8.5 shows that the fraction of signal events with two b-tags is slightly lower than 50%. This reflects the limitation that b-tagging acceptance is restricted to jets with $|\eta| \leq 2.5$, and the true efficiency varies slightly with jet kinematics. Conversely, the probability of b-quark production via QCD interactions is low, and thus, the distribution of b-tagged jets sharply peaks at zero for background.

The results of the studies presented in this section are used to inform the design and the strategy definition of the whole analysis, as discussed in the following.

8.2 Analysis strategy development

Various approaches were considered for the analysis strategy. All of them are based on combining roughly similar tight kinematic selections, but differ in the background method. The first approach considered is a classical $ABCD$ method, using the $\Delta\eta$ between the two leading jets and the number of jets per event. This procedure requires two uncorrelated variables with some discrimination power between signal and background. The two-dimensional plane defined by such variables is segmented in Control Regions (CRs) and a Signal Region (SR) applying rectangular cuts. The CRs are used to derive transfer factors and predict the background in the SR. However, for this analysis, the method's sensitivity is found to be highly dependent on systematic uncertainties. Moreover, high signal contamination was found in the CR, compromising the background estimate. Exploring extensions of this method using additional dimensions did not yield significant improvements.

A functional form fit of the H_T spectrum was also investigated to perform a bump hunt. The background prediction results were accurate, but the signal sensitivity was low. Further cuts improved it, but as the sensitivity increased, the prediction's reliability decreased.

Finally, it was decided to adopt a Jet Counting method, similar to previous implementations from few analyses targeting similar models [202, 203]. In this analysis the usage of a selection on the event shape is introduced, and the background method is improved. As the method's name suggests, the Jet Counting approach is based on counting energetic jets in the event. For SR definition, the events are categorised based on the total number of jets above a given p_T threshold. The background is predicted using a combination of the observed data and MC simulations, as explained in Section 8.4. The results of the search are obtained with single bin likelihood tests. For each model only the SR with maximal expected sensitivity is considered (see Chapter 10).

8.3 Signal region definition

The histograms shown in Figure 8.6 represent a series of possible signal region definitions, for the 2×3 models. The bin content in those histograms represents the number of MC

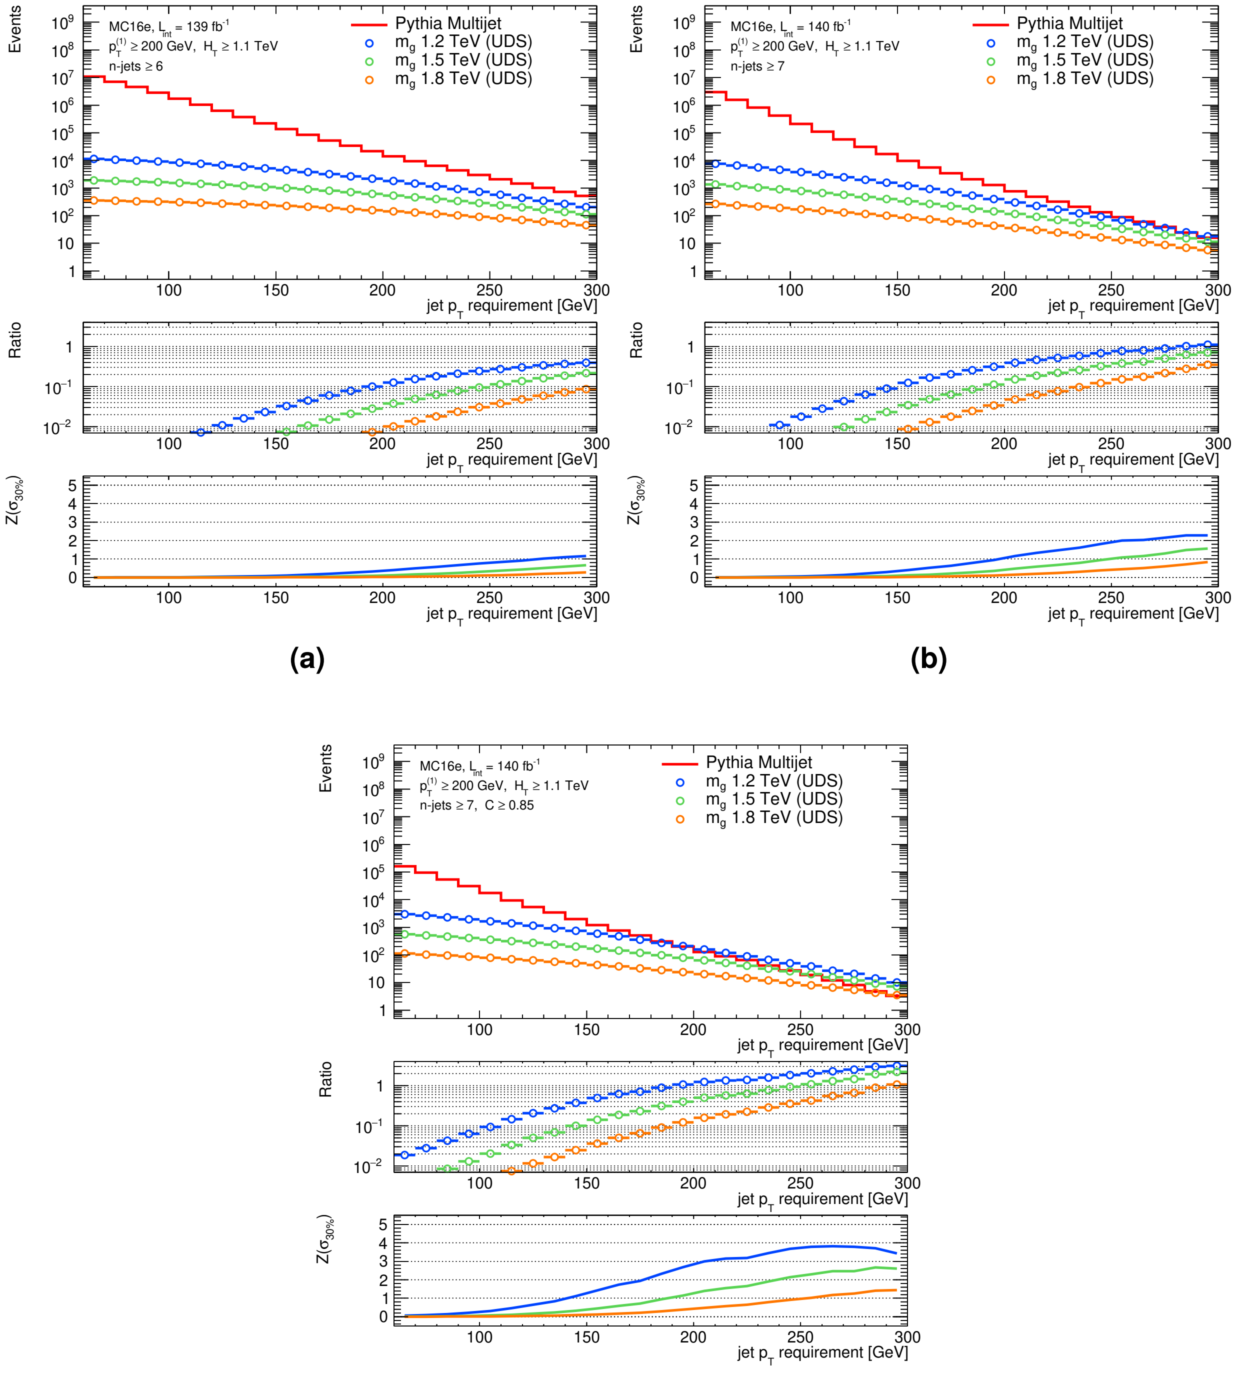

(a)

(b)

(c)

Figure 8.6: Histograms showing the number of events counted in MC samples containing at least (a) 6 or (b) 7 jets with p_T above the one indicated in the horizontal axis. An addition requirement on C is applied in (c). The middle panel shows the counts' ratio between signal and background simulations. The bottom panel shows the significance computed using the MC yields and assuming 30% systematic uncertainty on the background.

events with at least six or seven jets having a p_T above the threshold indicated in the horizontal axis. Therefore, the events in each bin are a subset of those in the left-neighbouring bin. The bottom panels display the signal-to-background ratio and the calculated statistical significance of the different signal hypotheses, assuming a flat 30% systematic uncertainty on the background. The Poisson-Poisson statistical model with asymptotic approximation is used (Equation 7.9). Although the considered signal models have six outgoing partons, requiring at least seven jets appears to suppress the background more effectively while retaining a significant fraction of the signal events. In terms of background discrimination, the situation noticeably improves with a selection on the C observable, as demonstrated by Figure 8.6(c).

Similar studies are shown in Figure 8.7 for the 2×5 models. Here, the total number of energetic jets is determined by the mass difference between the gluino and neutralino. If the neutralino is very light, it will be highly boosted, merging some of the reconstructed jets associated with its decay. Higher sensitivity is achieved by maintaining the n-jets threshold at seven, as demonstrated in Figures 8.7(a) and 8.7(b). When the gluino and neutralino masses are close, most of the event energy is contained in the three jets resulting from the neutralino decay. This leads to a reconstructed final state similar to the previous case, where a 7-jets requirement is preferred over a 6 or 8-jets threshold. Figures 8.7(c) and 8.7(d) illustrate the case where neutralinos have a large mass, yet low enough to ensure that energy evenly distributed among all final state partons. For this scenario, increasing the n-jets requirement up to eight considerably improves the sensitivity. Higher multiplicities provide no particular advantage while making the background extrapolation slightly more difficult.

Multiple SRs are needed to investigate a wide range of $m(\tilde{g})$ and $m(\tilde{\chi}_1^0)$. For this analysis, the regions definitions are optimized to maximize the discovery potential. A scan is performed over different n-jets requirements combined with all possible thresholds of p_T and C. For UDB models, the number of b-tags is also varied. In each candidate region, hypothesis test is run using MC signal and predicted background only. The selections yielding the best expected p-values are retained. The optimization procedure takes into account the full set of systematic uncertainties from Section 8.5. Table 8.1 reports the resulting SRs. Three overlapping regions are found to be optimal for direct decay models, with a 7-jet requirement at increasing values of p_T. For some cascade decays, relaxing the C cut and increasing the number of jets to eight yields maximal sensitivity. For UDB models, the preferred number of b-tags per event is ≥ 2, efficiently suppressing the background.

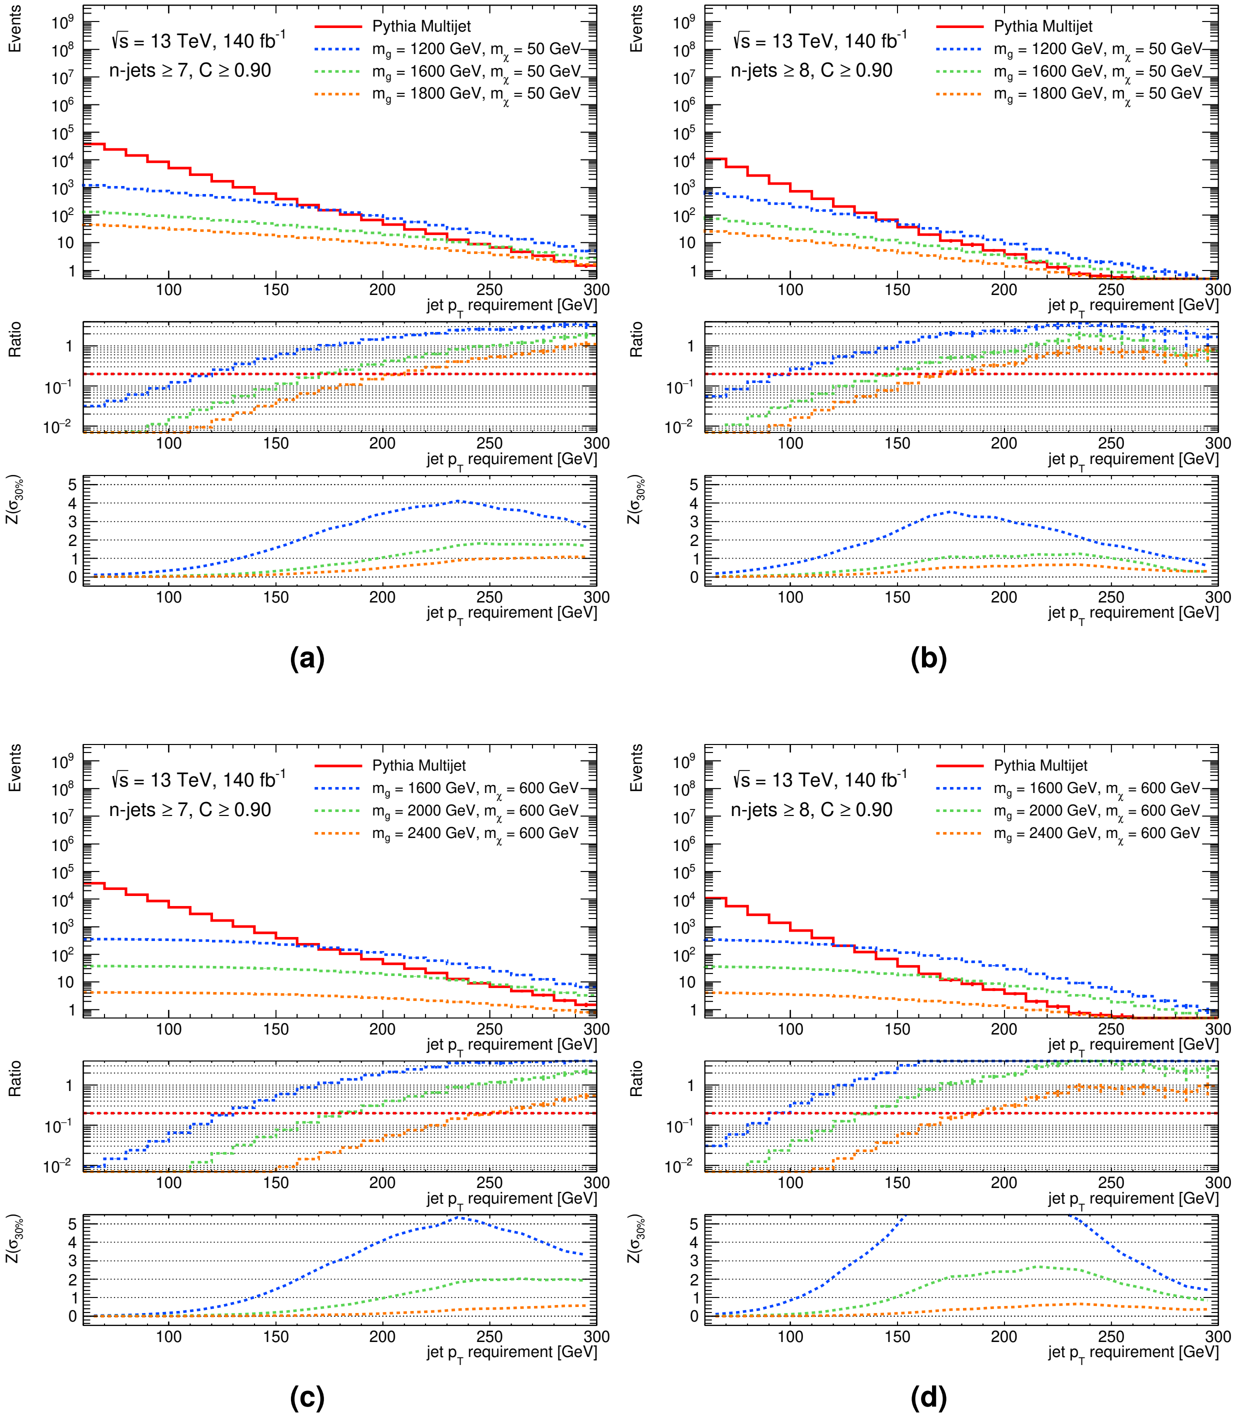

(a) **(b)**

(c) **(d)**

Figure 8.7: Histograms showing the number of events counted in MC samples for cascade decay models with $m(\tilde{\chi}_1^0) = 50\,$GeV and containing at least (a) 7 or (b) 8 jets with p_T above the one indicated in the horizontal axis. Same plots are show in (c) and (d) for a $\tilde{\chi}_1^0$ mass of $600\,$GeV A requirement of $C \geq 0.9$ is applied. The middle panel shows the counts' ratio between signal and background simulations. The bottom panel shows the significance computed using the MC yields and assuming 30% systematic uncertainty on the background.

Table 8.1: Signal regions definitions. The counting of the jets is to be considered at the given p_T threshold. The regions from SR1 to SR2 target both UDB and UDS decays, while SR1bj and SR2bj are optimized for high mass models with UDB decay.

	n-jets	jet p_T [GeV]	C cut	b-tags	Target Signature
SR1	≥ 7	≥ 180	≥ 0.90	-	(2×3) Small Masses
SR2		≥ 220		-	(2×3) Medium Masses
SR3		≥ 240		-	(2×3) High Masses
SR4	≥ 8	≥ 180	≥ 0.85	-	(2×5) Medium Masses
SR5		≥ 210		-	(2×5) High Masses
SR1bj	≥ 7	≥ 180	≥ 0.85	≥ 2	(2×3) b-jet Final States
SR2bj	≥ 8				(2×5) b-jet Final States

8.4 Background method

The background extrapolation method is based on empirical observations [202, 203], and combines MC events with collected data. Simulations are employed to compute the transfer factor across different jet multiplicities. These factors are then applied to event yields obtained from data, to extrapolate from low to high jet multiplicity regions. The number of events with n-jets with a given p_T can be extrapolated as:

$$N^{\text{Extr.}}_{n\text{-jets}} = N^{\text{Data}}_{m\text{-jets}} \cdot \frac{N^{\text{MC}}_{n\text{-jets}}}{N^{\text{MC}}_{m\text{-jets}}} \, . \tag{8.5}$$

The Equation 8.5 is used for $m < n$, and both are exclusive requirements on the number of jets. The low m-jet region is also referred to as the Normalization Region (NR). All the event yields $N^{\text{Type}}_{i\text{-jets}}$ terms assume the same p_T threshold for jet counting.

To avoid introducing bias, the data NR should ideally be signal-free. Considering that the signal is characterised by a large multiplicity in the final state, high background purity is expected in events with few jets (< 6). Figure 8.8 shows the expected yields for QCD multijet and gluino events in exclusive n-jets bins. It demonstrates that data events with up to four jets can be safely used for the background prediction. While for high-p_T 5-jet events, a significant level of signal contamination is expected.

The transfer factors are computed using only exclusive number of jets selections. In contrast, the signal regions are inclusive in n-jets. Consequently, for a SR with at least n jets, the extrapolation is obtained by summing exclusive jet multiplicities up $n + 2$:

$$N^{\text{Extr.}}_{\geq n\text{-jets}} = \sum_{i=n}^{n+2} N^{\text{Extr.}}_{i\text{-jets}} \, . \tag{8.6}$$

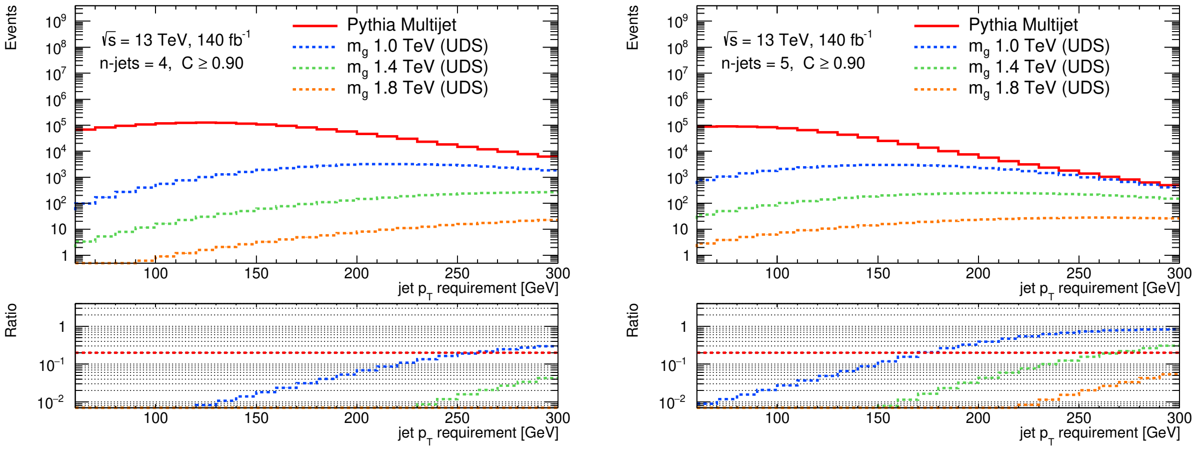

(a) **(b)**

Figure 8.8: Histograms of the number of events counted in MC samples containing different jets multiplicities with p_T above the one reported in the horizontal axis. Figures (a) and (b) correspond respectively to the exclusive n-jets equal to 4 and 5, with a C requirement of at least 0.9. The bottom panels show the ratio of the counts between signal and background simulations. The red dotted line represents the threshold of tolerable signal contamination, arbitrarily set to 20%.

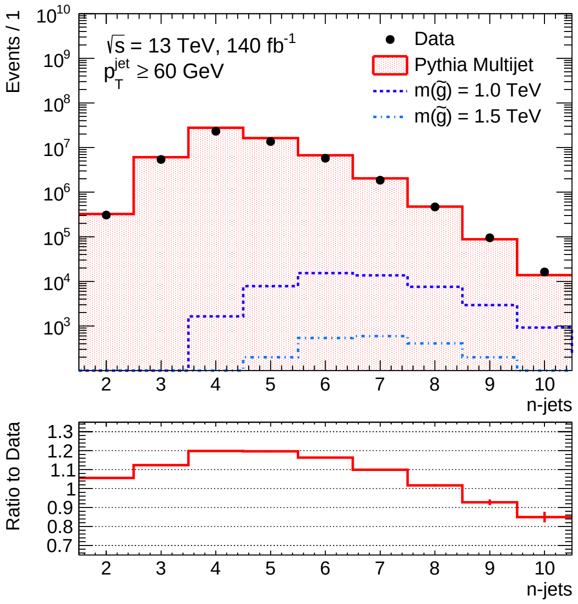

Figure 8.9: Distributions of the number of jets with $p_T \geq 60\,$GeV in data (black dots) and QCD multijet MC (red). The bottom panel shows the ratio of the MC prediction to data, after the trigger requirement and the event preselection are applied. Scale factors to correct the background extrapolation are obtained from this ratio.

The successive terms of the series are found to bring negligible contributions ($< 0.1\%$).

To address the slight bias in MC modelling of the n-jets spectrum shown in Figure 8.9 and, more generally, to adjust the prediction, scale factors are determined through direct comparisons of simulation with data in control regions (CR). Scale factors ($s_n = s_n\,|_{p_T \geq p_{T_0}}$) are derived for the MC yields from Equation 8.5, using a low transverse momentum threshold (p_{T_0}). The SR cut on C is also applied to the CR. The p_{T_0} threshold is set to be 60 GeV to ensure that the CR is as close as possible to the SR, while maintaining a negligible signal contamination.

As an example, the method to extrapolate background events in SR1 can be explicitly written as:

$$N^{\text{Extr.}}_{\geq 7\text{-jets}}\Big|_{p_T \geq 180\,\text{GeV}} = \sum_{i=7}^{9} w_{4,i} \left(N^{\text{Data}}_{4\text{-jets}} \cdot \frac{N^{\text{MC}}_{i\text{-jets}}}{N^{\text{MC}}_{4\text{-jets}}} \right)\Bigg|_{p_T \geq 180\,\text{GeV}}, \tag{8.7}$$

where:

$$w_{4,i} = \frac{s_i}{s_4}; \qquad s_i = \frac{N^{\text{Data}}_{i\text{-jets}}}{N^{\text{MC}}_{i\text{-jets}}}\Bigg|_{p_T \geq 60\,\text{GeV}}. \tag{8.8}$$

Equation 8.7 shows that the weight $w_{4,i}$ corrects the terms in the MC yields ratio, $i.e.$ each term $N^{\text{MC}}_{i-\text{jets}}$ is multiplied by the corresponding scale factor s_i.

Background composition

The primary source of background arises from QCD multijet events. Figures 8.10(a) and 8.10(b) illustrate the SM background composition. The expected fractions of all hadronic $t\bar{t}$ events are less than 1% and 5% for the NR and the SR, respectively. Thus, the transfer factors are estimated using only the simulated multijet sample. In SR1bj and SR2bj, where b-tagging is applied, the fraction of $t\bar{t}$ events increases to $\sim 20\%$, as shown in Figures 8.10(c) and 8.10(d). In this case, the secondary background component must be included in the extrapolation. To account for it, each N^{MC} term in Equation 8.7 becomes the combined sum of multijet and $t\bar{t}$ event contributions:

$$N^{\text{Extr.}}_{\geq n\text{-jets}} = \sum_{i=n}^{n+2} w_{n,i} \left(N^{\text{Data}}_{n\text{-jets}} \cdot \frac{N^{\text{Multijet}}_{i\text{-jets}} + N^{t\bar{t}}_{i\text{-jets}}}{N^{\text{Multijet}}_{n\text{-jets}} + N^{t\bar{t}}_{n\text{-jets}}} \right), \tag{8.9}$$

including the scale factors $s_i = \frac{N^{\text{Data}}_{i\text{-jets}}}{N^{\text{Multijet}}_{i\text{-jets}} + N^{t\bar{t}}_{i\text{-jets}}}\Big|_{p_T \geq 60\,\text{GeV}}$.

Signal injection tests

Eventual signal presence in regions used to extrapolate the background could lead to an overestimation of the prediction. Consequently, sensitivity to new physics would be reduced. To ensure that the method is robust against signal contamination, signal injection

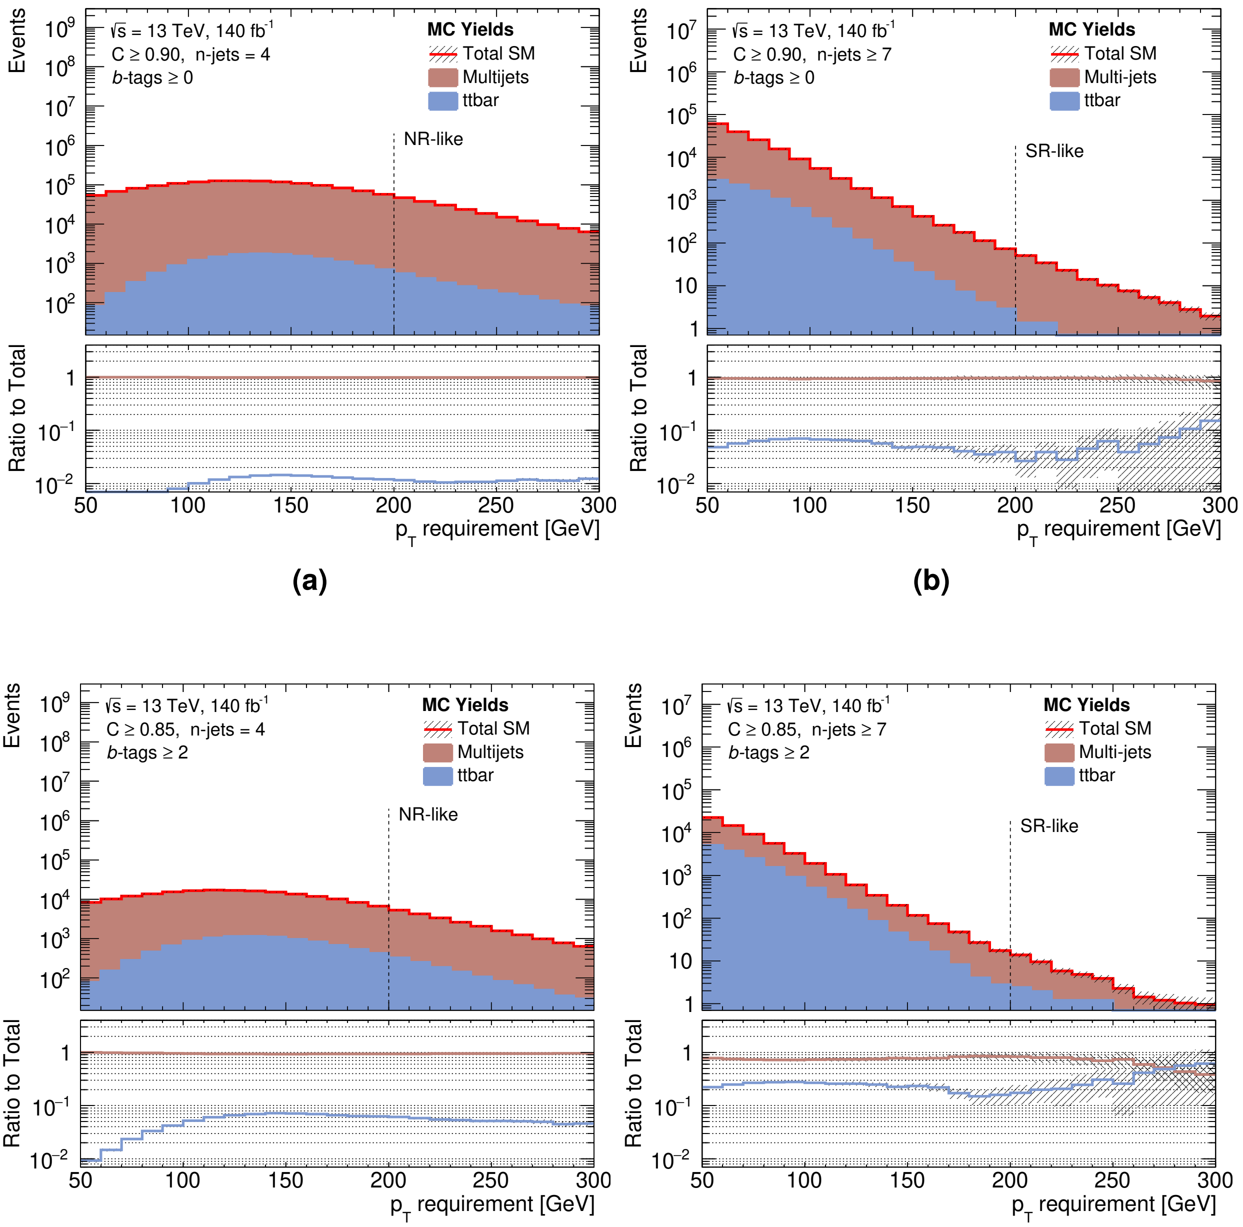

(a) **(b)**

(c) **(d)**

Figure 8.10: Total expected background yields from MC, for QCD multijet and all hadronic $t\bar{t}$ events. The (a) low and (b) high jet multiplicity regions used for the background extrapolation method are shown. Similar plots are reported in (c) and (d), with the additional requirement of at least 2 b-tagged jets. Selections on C are applied to match the corresponding SR, when b-tagging is (or not) applied. The vertical dashed lines give an approximate indication of where the SR and the NR p_T requirement are located in the horizontal axis. The hatched patterns represent the MC statistical uncertainties.

tests are performed. In these tests, the predictions of the yields from Equation 8.7 account for signal presence. All the "Data" terms are replaced with combinations of MC background and signal counts. The obtained extrapolations are compared to the pure MC background yields. Figure 8.11 displays the predicted number of events from the method, with and without signal injection. The number of signal events is also reported. Only the lighter gluinos introduce a small bias in the background estimate. Nevertheless, the expected

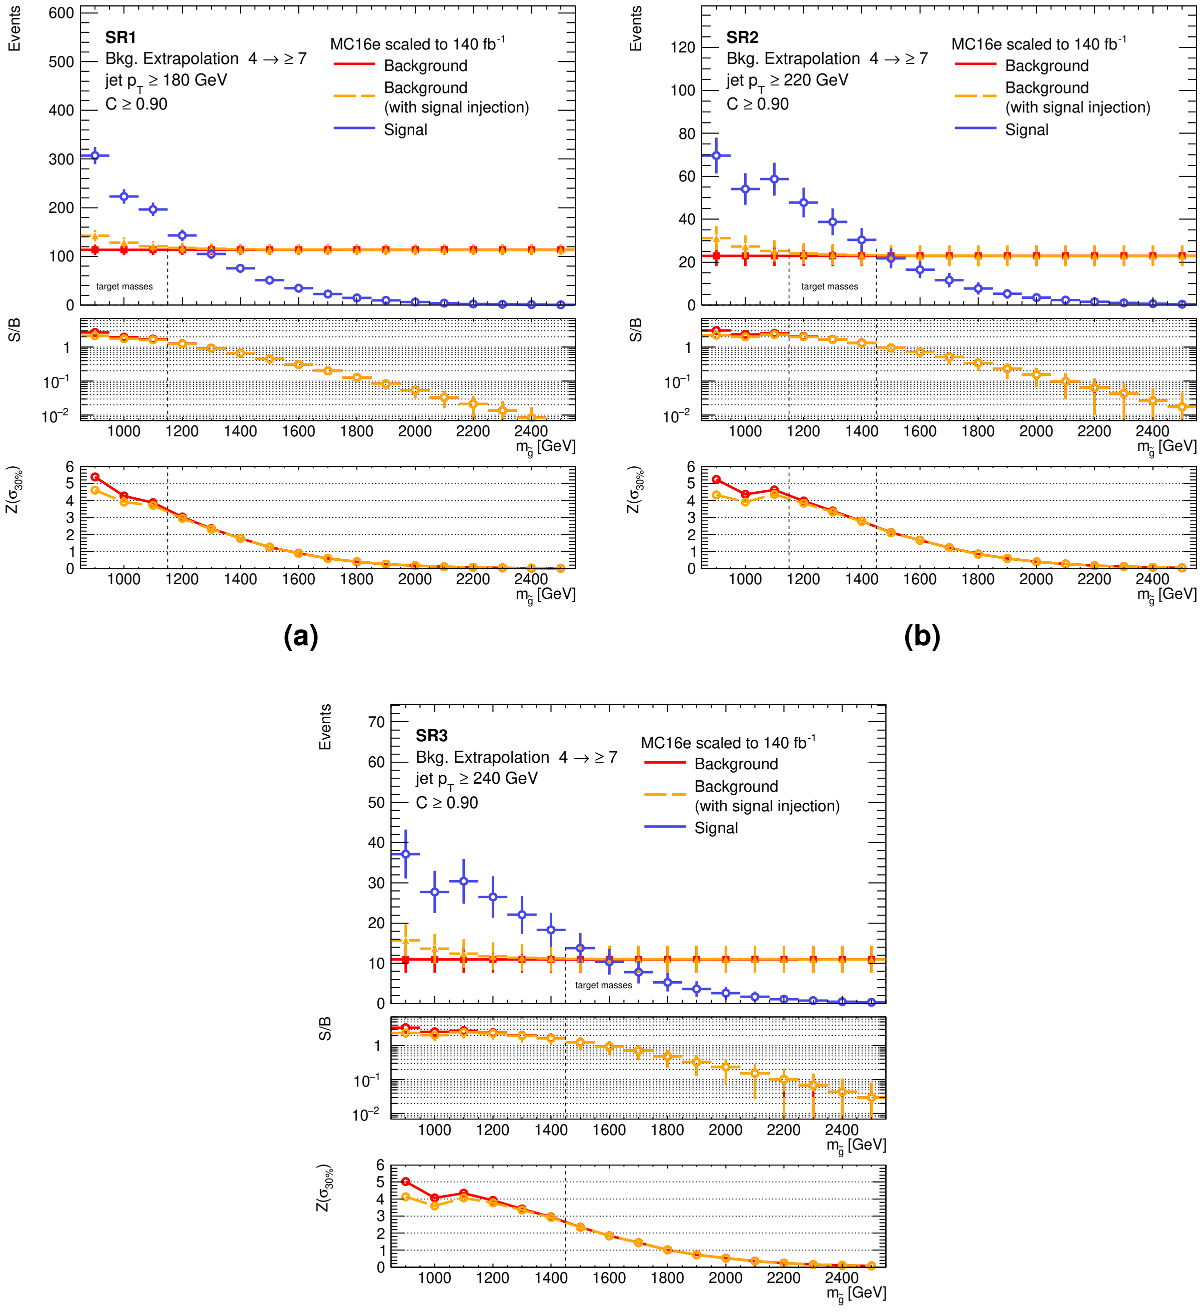

(a)

(b)

(c)

Figure 8.11: Expected event yields in (a) SR1, (b) SR2 and (c) SR3. The numbers are obtained using only MC samples. The background distribution corresponds the multijet event counts. The background with signal injection is obtained by replacing the "Data" terms with QCD multijets plus signal events, in Equation 8.7. The p_T threshold used for the CR is 60 GeV.

signal-to-noise ratio in the SRs is high enough such that sensitivity is minimally affected. If the CR is shifted towards the SR, by increasing the jet p_T threshold, the sensitivity drops rapidly for $m(\tilde{g}) \leq 1.2$ TeV.

Background validation

Tests of the background extrapolation method are performed in signal-depleted regions. Figures 8.12(a) and 8.12(b) show the comparison of background and data yields in regions with low and high number of jets, respectively. At large jet multiplicity, the selection on C is inverted to reduce signal contributions. The predictions are in good agreement with data within the uncertainties. The plots also show the small impact of the scale factors s_n. By definition, they shift the predictions to perfectly match the observed data yields in the CR, at a jet p_T requirement of 60 GeV. This leads to an overall improvement in the estimate. In Figure 8.12(c) the background method is also evaluated for b-tagged regions. The estimate well describes data when combining $t\bar{t}$ and QCD background processes.

To broadly evaluate the background modelling, specific Validation Regions (VRs) are defined. These regions are designed to fully surround the SRs while minimising signal presence. They are summarised in Table 8.2. The VR-As and VR-Bs are intended to test the method at high values of C, but with low jet multiplicities. The VR-Cs and VR-Ds validate the method in regions with many jets at high p_T, where the C selection is inverted to suppress signal. Four additional dedicated VRs are defined to validate the background estimation in the case where b-tagging is required. Figure 8.13 shows the yields of the VRs. Good agreement is found between the background expectation from the jet counting method and the observed data. A small discrepancy is observed in VR-B3, not covered by the uncertainties, which led to the definition of an additional non-closure uncertainty of nearly 5% on the expected background yields in the SRs with $C > 0.9$.

The robustness of the method across all values of the C variable is verified more thoroughly. To ensure the absence of any dangerous trend, predictions for many intervals of C are compared with data. Different jet multiplicities are considered, with minimal signal contamination in the probed regions. The outcomes are reported in Figure 8.14. Despite the known limitations in the MC modelling of event-shape variables [231], the extrapolation procedure yields accurate results across the full C spectrum.

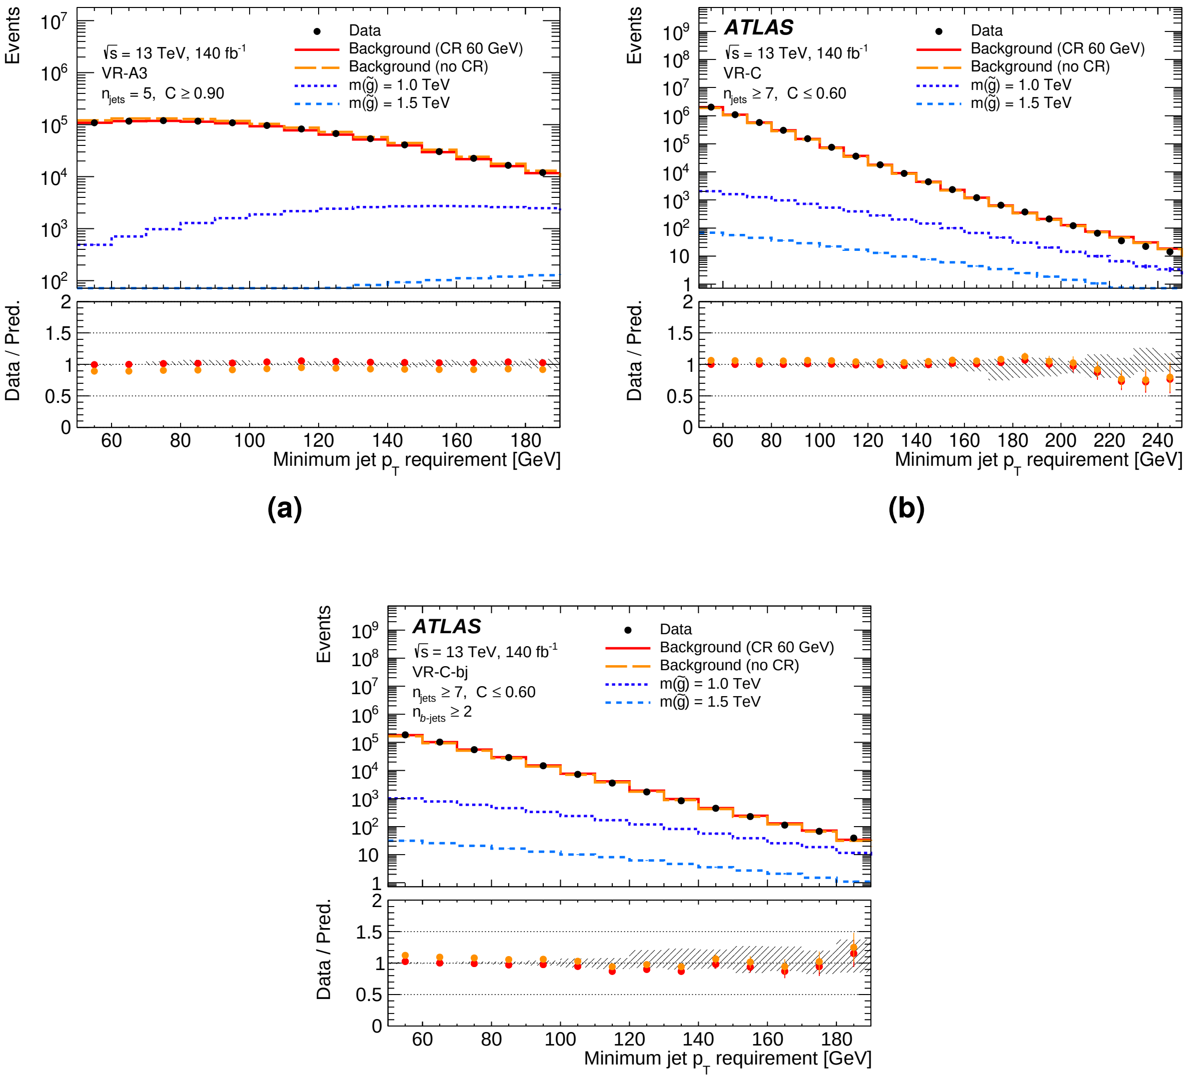

(a) **(b)**

(c)

Figure 8.12: Tests of the background estimation for different minimum jet p_T requirements, in (a) low n-jets region with high values of C and (b) high n-jets but low C region. Figure (c) does include also a b-tagging requirement. The solid red line is the prediction using correction factors estimated in the CR, while the orange line is obtained without such corrections. The bottom panels show the ratio of the data to the predicted yields, with (red dots) and without (orange dots) the CR. The hatched pattern in the bottom panels represents the total uncertainty (syst. $\bigoplus$ stat.) in the background estimate when the CR is use.

Table 8.2: Validation regions definitions. The counting of the jets is to be considered at the given p_T threshold. For all the regions, the background prediction is always normalized using scale factors computed at lower jet p_T requirements, keeping all the other selections unchanged.

	N. jets	jets p_T [GeV]	C cut	Predicted Background	Observed Data
VR-A1	5	180	≥ 0.80	73000^{+1800}_{-2400}	70184
VR-A2		160	≥ 0.85	65000^{+1800}_{-2200}	64985
VR-A3		150	≥ 0.90	30000^{+2100}_{-1000}	30360
VR-B1	6	120	≥ 0.80	80000^{+2200}_{-2800}	80271
VR-B2		110	≥ 0.85	58000^{+3900}_{-1800}	59997
VR-B3		100	≥ 0.90	28000^{+1000}_{-2000}	30212
VR-C1	≥ 7	180	≤ 0.60	350^{+37}_{-72}	372
VR-C2		220		48^{+6}_{-10}	35
VR-C3		240		18^{+4}_{-3}	14
VR-D1	≥ 8	180	≤ 0.60	23^{+5}_{-6}	16
Regions with at least 2 b-tagged jets					
VR-A-bj	5	180	≥ 0.85	2100^{+600}_{-150}	1973
VR-B-bj	6	120	≥ 0.85	3700^{+500}_{-250}	3425
VR-C-bj	≥ 7	180	≤ 0.60	34^{+13}_{-6}	39
VR-D-bj	≥ 8	160	≤ 0.60	8^{+6}_{-5}	6

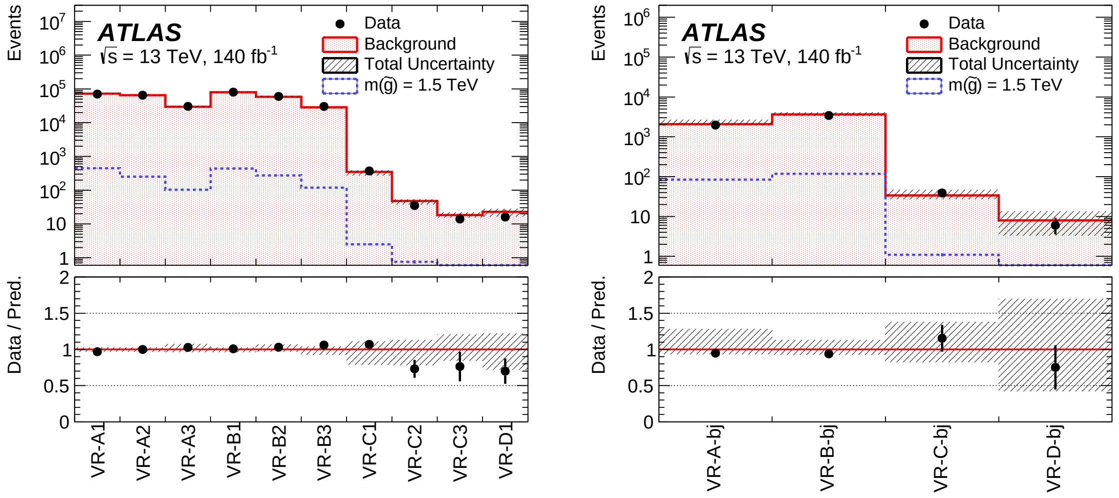

(a) **(b)**

Figure 8.13: Comparison of event yields between the observed data and the background expectation in the VRs. Figure (a): VRs containing no explicit requirement on the number of b-tagged jets. Figure (b): VRs containing at least two b-tagged jets. The bottom panel presents the ratio of data to the background prediction. The hatched pattern represents the combined statistical and systematic uncertainty in the background estimate.

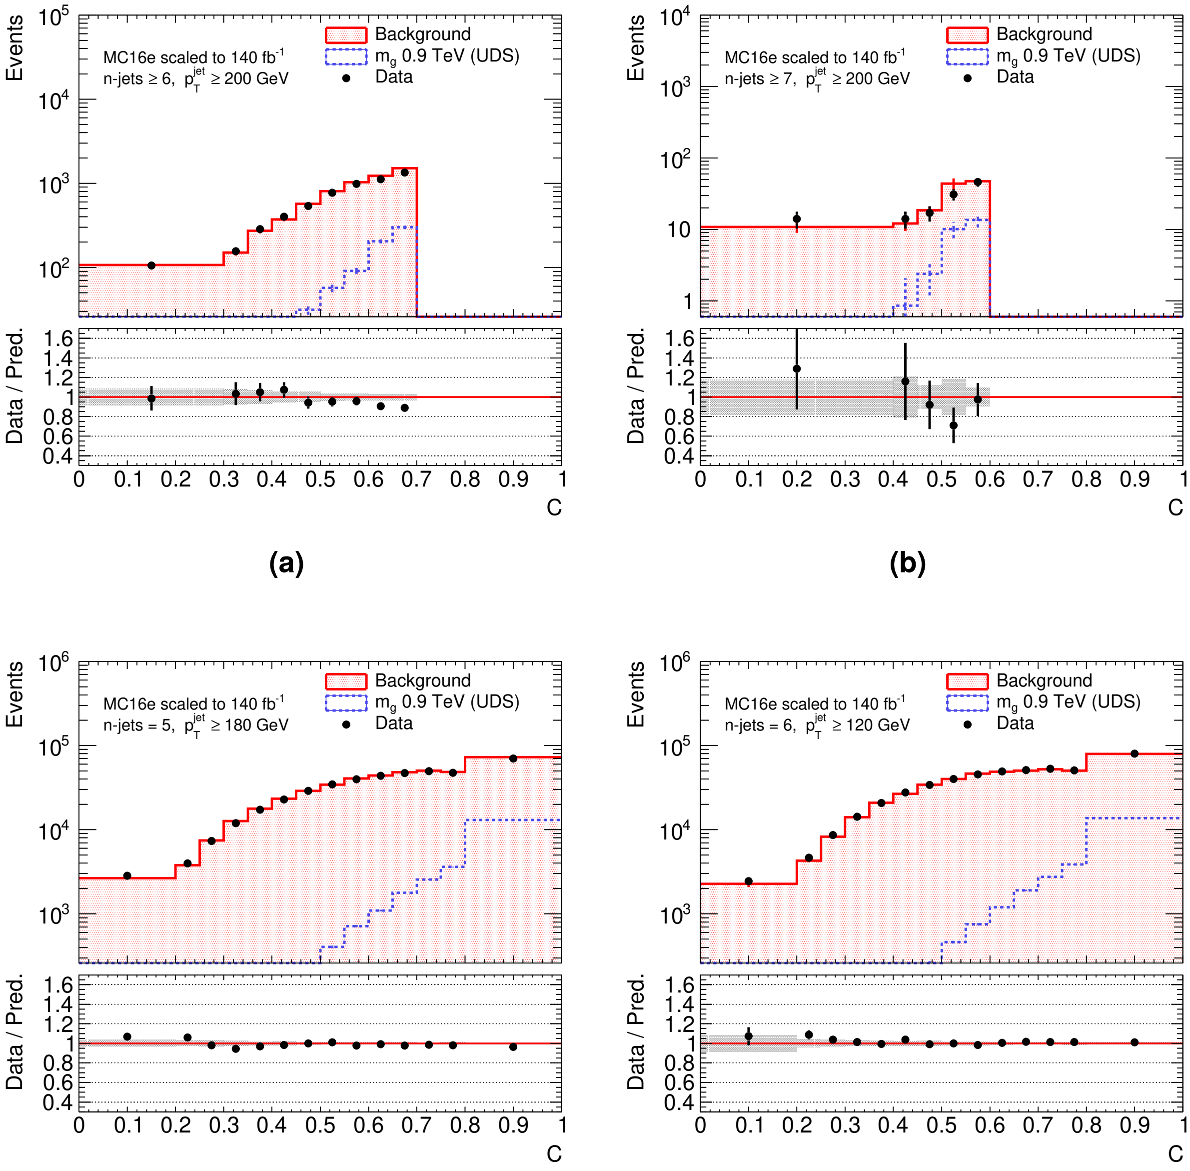

(a) **(b)**

(c) **(d)**

Figure 8.14: Closure test of the background method in binned regions of C. Figures (a) and (b) show respectively ≥ 6 and ≥ 7 jets regions, with high p_T. The high-C bins are blinded in this case. While in (c) and (d) the full C spectrum is scanned, but with only 5 and 6 jets. The background was extrapolated from a jet multiplicity of 4, with a CR at 60 GeV. The bottom panel shows the ratio of data over the predicted number of background events. In this panel, the grey-shaded region represents the MC statistical error only.

8.5 Systematic uncertainties

The interpretation of the final results accounts for multiple sources of uncertainty. They can be categorised into two distinct classes: the theory modelling and the experimental errors. The first class considers the lack of knowledge and the approximations needed for MC simulations. They are larger for multijet background, as in this case simulations are largely dependent on approximations done in QCD calculations. The second class covers uncertainties arising from objects reconstructions, estimated via dedicated calibrations. This class also includes methodological errors arising from the background method. In the statistical interpretation of the results, each error is treated as an individual Nuisance Parameter (NP) of the fit (see discussion in Section 7.5).

Modelling uncertainties

The composite nature of protons and the properties of QCD impose the usage of approximations in the computations of p-p collisions cross-sections $\sigma(pp \to X)$. As discussed in Section 2.4, the modelling of such events is usually separated in different components:

1. The hard scattering process. It involves partons from the incoming protons, and few outgoing particles. The cross-section for this process is convoluted with the proton PDFs, and can only be computed via a perturbative expansion in powers of the strong coupling constant α_s. Such calculations introduce a factorization scale (μ_F) and a renormalization scale (μ_R). The first one decouples the "soft" and "collinear" divergences from the hard process, while the second one sets the energy scale for α_s. Contributions below these energy scales are not included in the hard scattering.
2. Inital- and final-state radiation (ISR and FSR). They include additional radiated particles from the incoming scattering partons or from the reaction products.
3. The parton shower formation. It consists of an algorithmic way of dealing with divergent emissions. Different approaches are possible. Hard scatter process products, including ISR and FSR, computed with a fixed order expansion, are evolved in the non-perturbative regime with such algorithms.
4. Formation of hadrons (hadronisation). Coloured partons merge into colourless hadrons (baryons or mesons). This transition takes place in a non-perturbative regime and therefore cannot be computed analytically.

All the calculations done for these items carry some uncertainty. The most relevant ones for this analysis are arising from the missing higher orders in the perturbative expansion of the partonic cross-section, the PDF and the α_s uncertainties. The impact of these is evaluated by changing the relevant parameters in the MC simulation. This process leads to variations of the relative cross-section of events, which effectively modifies the MC weight for each event. In this way, the same MC sample can be re-weighted to evaluate different systematic effects. The parton showering and hadronisation systematic uncertainties require instead comparing the nominal MC with a different sample, generated with alternative MC techniques. For each systematic effect, the prediction of background and signal yields is compared to the nominal event counts. The difference of these two numbers is quoted as the corresponding systematic uncertainty.

QCD multijet Three leading systematic effects define three different sets of varied event weights. Within each set, combinations are built by taking the envelopes all of all the variations to estimate the corresponding uncertainty. In most of the phase space considered by this analysis, the highest systematic error is originating from variations of the μ_F and μ_R scales. In some cases, extreme fluctuations are present with huge statistical uncertainties, as shown in Figure 8.15(a). These fluctuations are caused by very few events carrying large weights. The re-weighting approach is prone to this type of features, especially when the MC has lower statistics compared to the absolute number of events it should simulate, for a given integrated luminosity. This is the case for the multijet sample, where the enormous cross-section of QCD events makes prohibitive achieving MC statis-

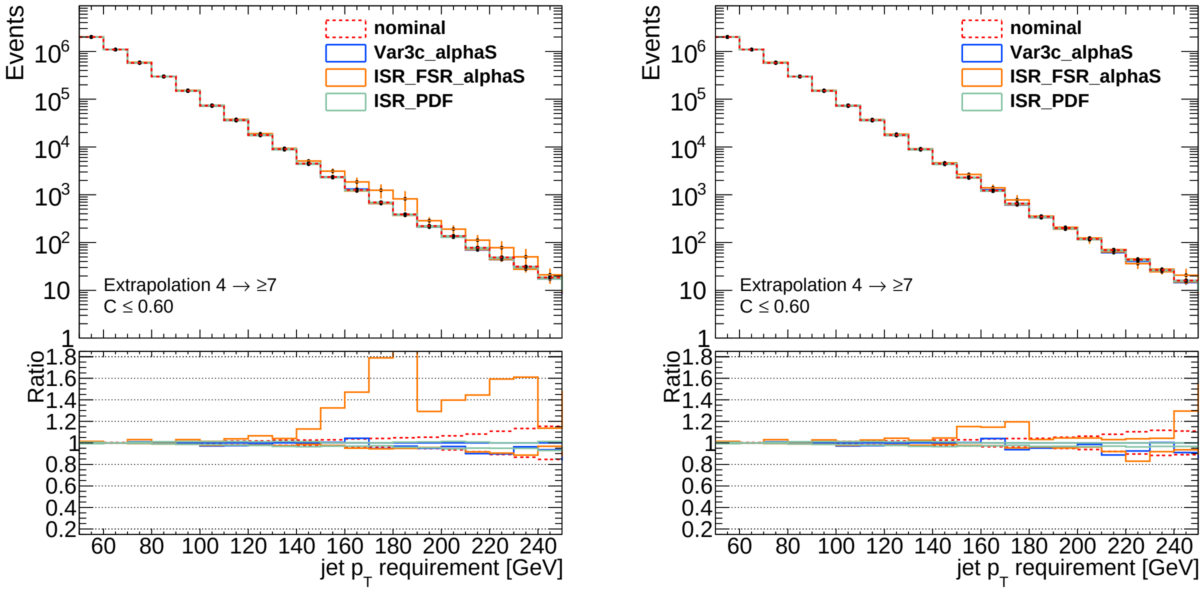

(a) (b)

Figure 8.15: Impact of the systematic variations of the MC multijet sample on the background estimate. The envelopes of the three stets are reported in (a) for all the MC events, and in (b) after the outlier removal, where the few events with very large weights are dropped. The "Var3c alphaS" are variations of the strong coupling in the initial state shower. "ISR FSR alphaS" accounts for missing higher orders in the perturbative expansion of the partonic cross-section, and for variations of the μ_F and μ_R scales. PDF uncertainties related to the ISR modelling are evaluated with the "ISR PDF" envelope.

tics larger than the expected data yields. To suppress these artificial fluctuations some events are excluded following an outlier removal procedure. For the problematic regions of the phase space, events falling in the $\sim 0.1\%$ upper tail of the event weight distributions, are removed. The results are shown in 8.15(b). For the showering and hadronisation uncertainties an alternative MC generator was considered. However, such sample was found to poorly describe data in several VRs and generally being inconsistent with observations in the high multiplicity regime. To avoid an overestimation of the error, this and other eventual contributions are extracted from data in VRs as non-closure uncertainties, as described few paragraphs below.

All hadronic $t\bar{t}$ A similar approach to the multijet one is adopted for the $t\bar{t}$ modelling uncertainties. In this case, both the re-weighting and alternative generator contributions are considered, accounting for hard-scatter, parton showering and hadronisation errors. Their impact is evaluated on the final background prediction, only for VRs and SRs that do include b-jets.

Signal $\tilde{g}\tilde{g}$ The event yields for signal are evaluated with alternative weights. As for the other MC samles, these do reflect the impact of μ_F and μ_R scales variations, the uncertainties related to PDFs, α_s, and MC tunes used.

MC statistics The modelling of all the samples is further limited by the amount of MC events generated. Usually, as a rule of thumb, the target number of MC events is three times larger than the expected total event counts. However, when the cross-section of a given process is too large, this might not be feasible. As previously mentioned, this is the case for QCD multijets. The associated MC statistical uncertainty corresponds to the sum in quadrature of the MC weights in each bin: $\sqrt{\sum w_{MC}^2}$. For signal this error is assigned directly to the event yields, for the background MCs is propagated through the formula 8.7.

Experimental uncertainties

Calibrating reconstructed quantites is a statistical procedure that does come with related uncertainties. It is often a non-trivial task and requires dedicated analyses. For most of the common physics objects, calibrations are performed centrally within the Collaboration. This is the case for all the objects used for this search.

JES and JER The calibration of jets is applied to both data and MC. As partially outlined in Section 3.3, jet calibration is a rather complex procedure composed of many steps and dependent on many parameters [221]. In total, 23 different NPs are used for JES and 8 for JER, following the central recommendations. The background extrapolation is repeated for each variation and then compared to the nominal one. An evaluation of the impact of

JES and JER on the final background estimate is reported in Figure 8.16, for n-jets≥ 7 and $C \geq 0.9$. For all the considered phase space, these systematic effects have an impact below 10%.

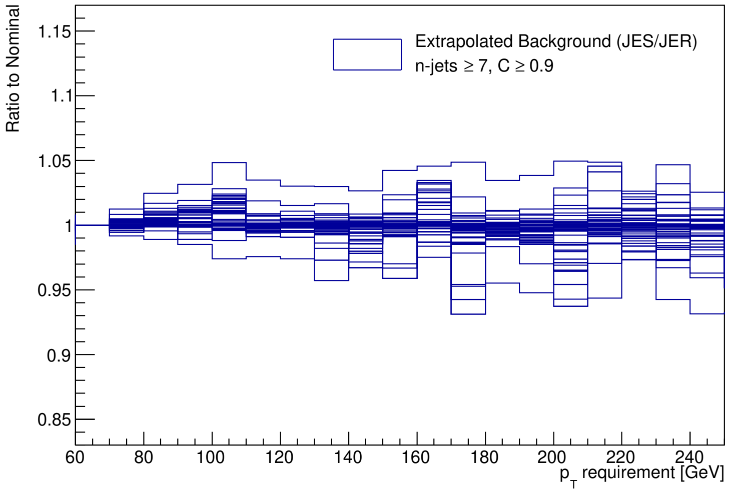

Figure 8.16: Impact of the JES and JER variations on the background extrapolation. The background is extrapolated for a total of 62 variations (31 up and 31 down) related to the JES and JER effects. The ratio to the nominal is reported on the y axis, for different p_T requirements on the horizontal axis. The event selection used in this case n-jets≥ 7 and $C \geq 0.9$.

Flavour tagging The b-taggers take as input tracks and jets, considering both high- and low-level features. Some of these are not perfectly modelled by the MC. Thus, differences in the actual b-jet efficiencies and in the l- and c-jet mis-tag rates can arise when comparing to the observed data. With a calibration of the DL1r tagger [224], multiplicative scale factors per jet can be applied in MC simulations, to compensate deviations in the estimation of the efficiencies. The scale factors and their uncertainties depend on jet p_T. In total, five independent errors related to this calibration are considered; however, they bring minor contributions ($< 5\%$) in b-tagged regions.

Luminosity A luminosity uncertainty is applied to the normalisation of the signal samples. For the combined Run 2 dataset, a luminosity uncertainty of 0.8% is used on the measured value of $140068.94\,\text{pb}^{-1}$, based on the latest assessments [228].

VR non-closure The background method has predictions compatible with data in almost all VRs within 1σ when considering all the uncertainties, including the data Poisson errors. Only in VR-B3 there is a slight underestimate (see Table 8.2 or Figure 8.13). To cover the observed discrepancy, a 4.8% systematic error is implemented on expected counts in SRs with a tight C selection ($C \geq 0.9$).

Summary

In Figure 8.17 the VR yields are shown, breaking down all the contributions to the background prediction uncertainties, except for the VR-non closure error. It is possible to see that the MC multijets errors are dominant in most of the regions. In total, more than 50 NPs are defined and included in the final SR likelihood fit. They are all summarised in Table 8.3. Some errors are applied to both signal and background. Despite having different contributions to the two populations, they are set be correlated, being described effectively by a single NP.

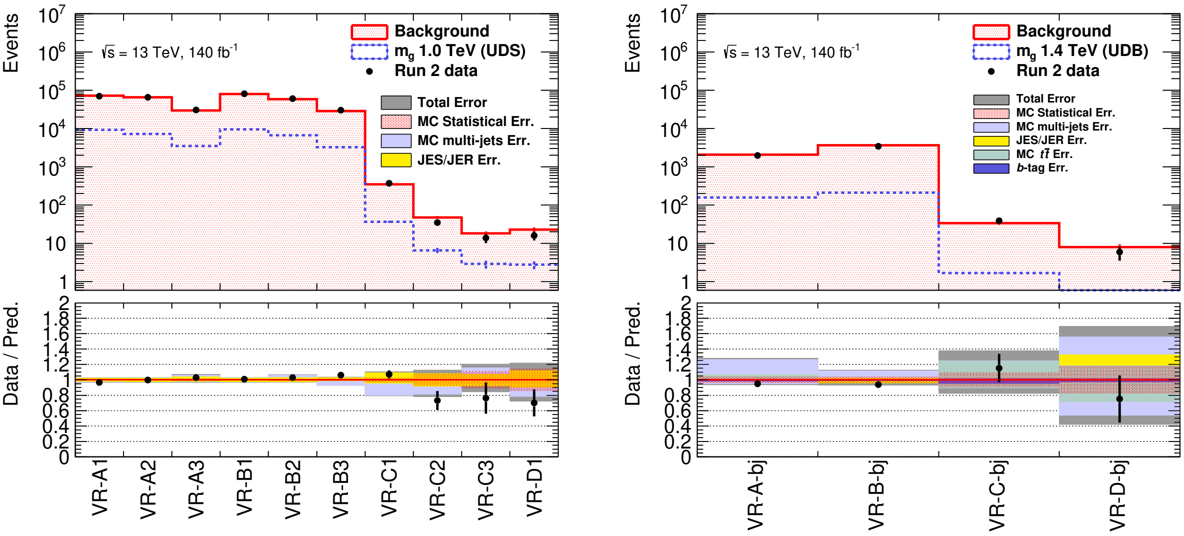

(a) (b)

Figure 8.17: Comparison of event yields between the observed data and the background expectation in the VRs. Figure (a): VRs containing no explicit requirement on the number of b-tagged jets. Figure (b): VRs containing at least two b-tagged jets. The bottom panel presents the ratio of data to the background prediction, showing the different contributions to the uncertainty.

Table 8.3: Summary of the uncertainties considered for the CCA.

Uncertainty type	# NP	Short description
MC multijets	3	MC theory uncertainties, deriving from PDF, scale, ISR and FSR modeling uncertainties
MC $t\bar{t}$	4	Similar to multijets, includes hard scatter, ISR and FSR modeling uncertainties and considers additionally parton showering and hadronisation errors
MC background stats	1	MC and data Poisson errors from CR, NR and transfer factors, propagated for Equation 8.7
MC signal	7	Signal modeling systematics, mostly scale and PDF uncertainties (excluding errors on the cross-section)
MC signal stats	1	Weighted Poisson error, computed as the sum of MC weights $\sqrt{\sum_i w_i^2}$
JES/JER set	31	Jet energy scale and resolution errors, for the Global Reduction JES and Simple JER sets; this uncertainty is applied to both signal and background (correlated NPs)
Flavour tagging	5	Uncertainties on *b*-tagging scale factors; this uncertainty is applied to both signal and background (correlated NPs)
Luminosity	1	Uncertainty on the total Run 2 luminosity (0.8%), for signal
VR non-closure	1	Error to account discrepancies above 1σ in VR, only for background with $C \geq 0.9$

Chapter 9

Mass Resonance approach

The Mass Resonance approach complements the CCA and aims to improve the sensitivity using more elaborate analysis tools. In this approach, the main goal is to search for evidence of resonances in the invariant mass spectrum. The key assumption is that, for background, the mass spectrum is a smooth and steeply falling distribution, as the mass increases. While for signal, a localized peak should appear in the spectrum at the value corresponding to the new particle's mass. However, to build the invariant mass spectrum, the four momenta of the decaying particles must be reconstructed. This step becomes particularly challenging for gluino RPV decays. In fact, given the high multiplicity in the final state, it is non-trivial to define the correct combination of jets originating from a common parent. The space of possible combinations exhibits factorial scaling, with the increase in number of jets. To deal with such combinatorial background, an ML-based method is developed to predict the assignments of the jets to the corresponding gluino decay. The method also needs to ensure that the assumption for the background remains valid, thus avoiding the artificial sculpting of the reconstructed mass spectrum for non-resonant QCD multijet decays.

The content of this chapter is in large part the fruit of the work done by the other members of the Analysis Team. The analysis strategy is discussed together with the background method and the systematic uncertainties.

9.1 Analysis strategy

Figure 9.1 shows a schematic of the approach's strategy. In each event, the hardest eight jets are passed through a NN, which uses only the jets four-momenta to predict their origin. The NN outputs three scores: the probability of a jet originating from either of the two gluinos or from another source, such as ISR or pile-up jets. Each jet is assigned to the class that has the highest probability. The gluino candidates' four-momenta are obtained by summing together the jets within the corresponding classes. The average mass spectrum $m_{avg} = 1/2 \cdot \left(m_{\tilde{g}}^{(1)} + m_{\tilde{g}}^{(2)}\right)$ is then used to run a bump hunt [233], with profile likelihood ratio tests. The background distribution is estimated in a data-driven manner, using a functional form fit. To increase sensitivity, the SRs apply cuts similar to those in the CCA. Events must include at least six jets with $p_T \geq 100$ GeV and the selection criterion $C \geq 0.9$ is applied. For models featuring UDB decays, an additional requirement of at least one b-tagged jet per event is included.

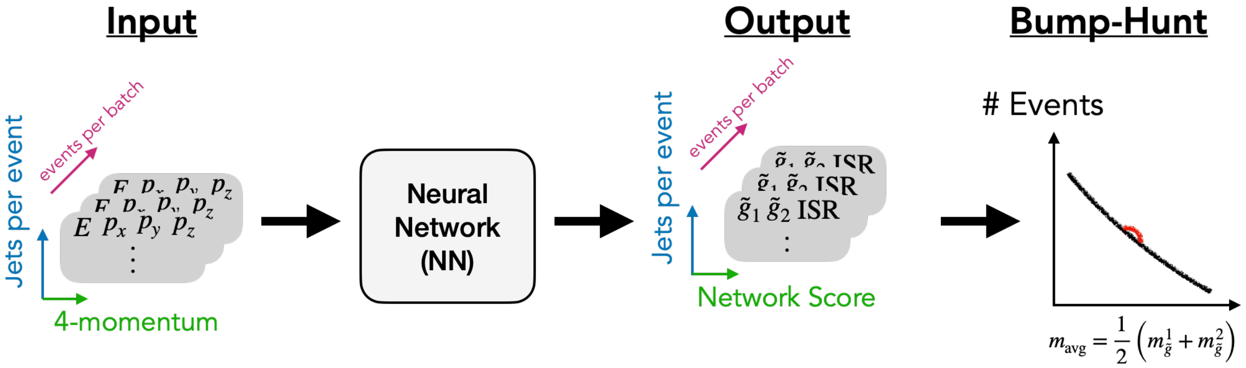

Figure 9.1: Overview of the analysis strategy. In each event, the jets 4-momenta are given as inputs to a NN that returns the jet-to-gluino assignments in the form of probability scores. Using these outputs, gluino candidates are reconstructed summing the corresponding jet 4-momenta, and the obtained invariant mass spectrum is used to perform a bump hunt.

9.2 The neural network

The combinatorial problem in event reconstruction is common to many physics analyses at colliders. For pair-produced resonances with few objects in the final state, traditional methods typically minimise or maximise a quantity dependent on the pairings, such as the mass asymmetry ($A_m = |m_1 - m_2|$). In recent years, several applications of advanced ML algorithms have been developed to solve same problem. These studies primarily focus on reconstructing SM processes, such as $t\bar{t}$ decays [234, 235]. In the context of this search, there is the additional challenge of not knowing the resonance mass. Furthermore, the phenomenology studies [236] usually do not account for background modelling. This analysis represents one of the first results employing such a technique.

The model

The implemented NN model is inspired by the Transformer architecture [183], based on the attention mechanism. Figure 9.2(a) illustrates the two basic building blocks. The `Embedding Block` consists of a fully connected feedforward network that employs Rectified Linear Unit (ReLU) activation functions (see Appendix A). Between each linear transformation and the application of the ReLU, Layer Normalisation is applied [237] to stabilize the network. In the input `Embedding Blocks` ReLU is not applied to the final Linear layer such that the output is not constrained to be positive. The `Attention Block` begins with multihead attention, which implements scaled dot-product attention. This method maps input queries (Q) to an output weighted sum of values (V), based on the compatibility between Q and keys (K). The latter is evaluated by the dot-product between Q and K:

$$\text{Attention}(Q, K, V, \Lambda) = \text{SoftMax}\left(\frac{QK^T}{\sqrt{d_k}} + \Lambda\right) V \,. \tag{9.1}$$

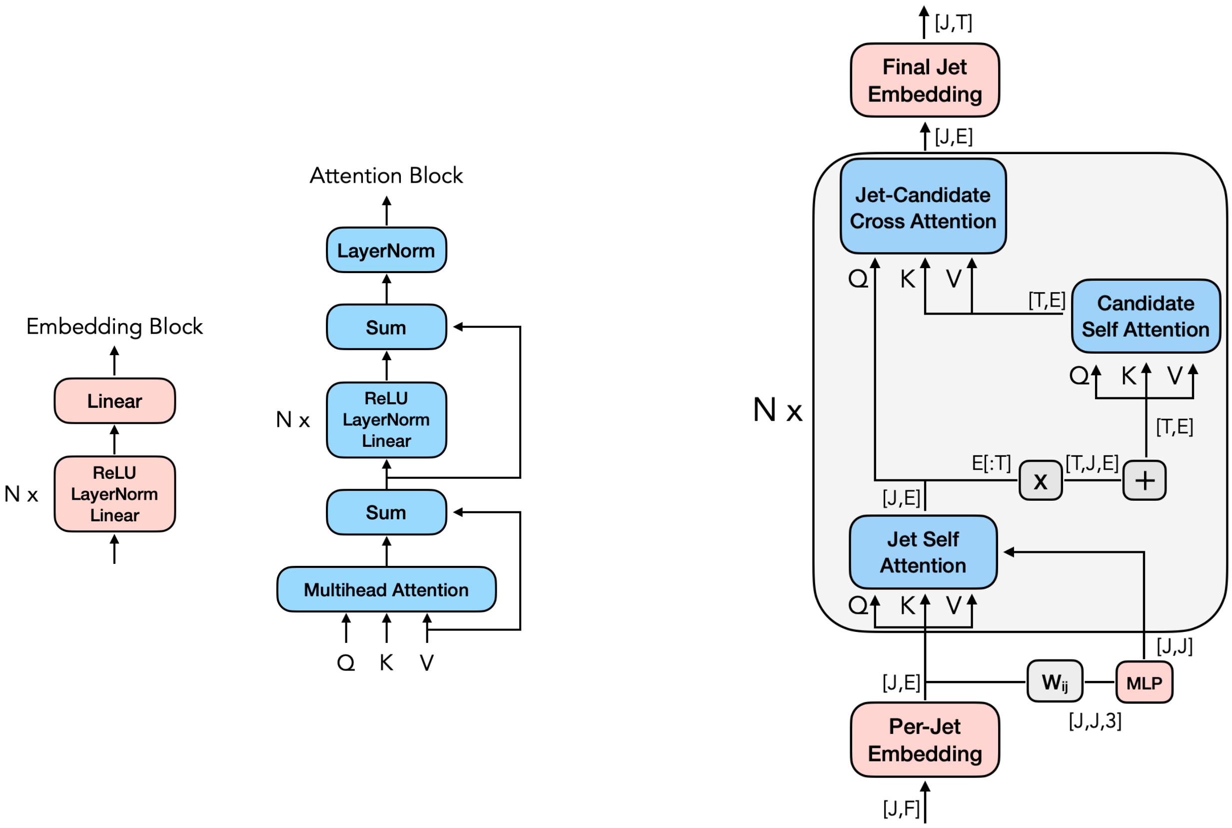

(a) **(b)**

Figure 9.2: Schematics of (a) the basic blocks of (b) the model architecture. In each block of the model, the input X and output Y dimensions are reported in square brackets [X,Y]. The values J and F corresponds the number of jets and features (8 and 5), E is the dimension of the embedded representation (32), and T the size of the targets space (3).

The scaling by the square root of the size of K ($1/\sqrt{d_K}$) controls excessive growths of the dot-product. Bias term Λ in the SoftMax argument can be optionally included. When the dimensions are compatible, residual connections with V are summed to the output. The final steps of this block include an additional feedforward network.

Figure 9.2(b) shows the architecture built by combining the basic building blocks in different fashions. The input to the model consists of a matrix containing the jets four-momenta stacked in rows:

$$x = \begin{bmatrix} \log p_T^{(0)} & \eta^{(0)} & \cos\phi^{(0)} & \sin\phi^{(0)} & \log E^{(0)} \\ \log p_T^{(1)} & \eta^{(1)} & \cos\phi^{(1)} & \sin\phi^{(1)} & \log E^{(1)} \\ \dots & \dots & \dots & \dots & \dots \end{bmatrix} \in \mathbb{R}^{J\times 5} \,. \tag{9.2}$$

This matrix passes through a first embedding step, where per-jet information is learned

$$x' = \text{PerJetEmbedding}(x) \in \mathbb{R}^{J\times E} \,. \tag{9.3}$$

An additional tensor is constructed out of the input jets:

$$W_{ij} = (\Delta\eta^{(ij)},\ \Delta\cos\phi^{(ij)},\ \Delta\sin\phi^{(ij)}) \in \mathbb{R}^{J\times J\times 3} \,, \tag{9.4}$$

containing pair-wise angular relationships across all the jets. The latter goes through a different `Embedding Block` and is transformed into a matrix $W' \in \mathbb{R}^{J\times J}$. The Jet Self-Attention step uses this angular tensor as a bias mask $\Lambda = W'_{ij}$, while the embedded jets x' are employed as Q,K, and V:

$$x'' = \text{AttentionBlock}(x', x', x', W') \in \mathbb{R}^{J\times E} \,. \tag{9.5}$$

The first three output coordinates of the E-dimensional latent space are arbitrarily chosen to represent the probabilities of the per-jet assignment scores for the three target classes: $\tilde{g}^{(1)}$, $\tilde{g}^{(2)}$, and ISR (the latter class also includes pile-up jets). Intermediate output candidates are formed using these three scores The jets assigned to each class are summed, resulting in a candidate tensor $c \in \mathbb{R}^{3\times E}$. This tensor is then passed through a Candidate Self-Attention block:

$$c' = \text{AttentionBlock}(c, c, c) \in \mathbb{R}^{3\times E} \,. \tag{9.6}$$

To learn relationships between them, the embedded jets and intermediate candidates Cross-Attention is computed, after their respective Self-Attentions

$$x''' = \text{AttentionBlock}(x'', c', c') \in \mathbb{R}^{J\times E} \,. \tag{9.7}$$

The entire sequence of attention blocks is repeated six times. After the last pass, a final `Embedding Block` calculates the per jet probability of belonging to each output class. The two sets of jets associated with gluino decays are formed and the corresponding four-momenta are computed. No constrains are imposed on the number of jets per output class.

The training

The model is trained exclusively on signal MC. To be consistent with the SR, only events with $C \geq 0.8$ and containing at least six hard jets ($p_T \geq 100\,\text{GeV}$) are considered. Target labels are assigned in each event by matching jets to truth partons with a threshold of $\Delta R \leq 0.4$. Jets with no corresponding parton are labelled as ISR. The NN takes exactly eight jets as input; events with fewer jets are zero-padded. To prevent introducing bias towards any specific mass value and sculpting the background, all the mass points are trained on simultaneously. A total of 1.7 million events from the MC16e campaign are used for the training, only 10% of them is reserved for validation.

The cross-entropy loss function is used (see Appendix A). To symmetrise it, the loss for swapped jet groupings between the two gluinos is also computed in each event, and only the assignments with minimal loss are considered:

$$L_M = \min_{\tilde{g}^{(0)} \leftrightarrow \tilde{g}^{(1)}} \text{CrossEntropyLoss}(\hat{y}, y) \ , \tag{9.8}$$

where $\hat{y}$ and y are the predicted and the true label. Auxiliary losses $L_M^{(i)}$ per repeated block are computed for the intermediate candidates. Their contribution is averaged and scaled by a factor $\lambda_{\text{aux}} = 0.25$. The total loss per training step is:

$$L = L_M + \frac{\lambda_{\text{aux}}}{6} \sum_{i=1}^{6} L_M^{(i)} \ . \tag{9.9}$$

Results and performance

The MC16d campaign is used as an independent test set to evaluate the performance of the network. Figure 9.3(a) shows the inferred confusion matrix. As a consequence of the symmetric loss, the two gluino labels can be interchanged. It is also notable that less than 20% of the jets from the gluinos are mislabelled as ISR. Roughly 30% of the ISR jets are incorrectly assinged to the gluinos.

The total fraction of correct labels across all mass points was evaluated. More than 95% of the events have at least four correctly predicted labels, between 50% and 60% of the events have at least six correct labels depending on the models' mass. Only in $\sim 15\%$ of the events all eight jet assignments matching the true labels.

The target and predicted m_{avg} spectra are shown in Figure 9.3(b). The former are produced using jets that are truth matched to the gluinos. It is possible to observe a significant skew and a downshift in the predicted spectra relative to the actual $m_{\tilde{g}}$. These effects are in large portion consequences of inefficiencies in the matching, FSR, and energy losses ($e.g.$, out-of-cone radiation). The predicted spectra are in good agreement with the target ones. They are slightly smeared, but the peak position remains unchanged.

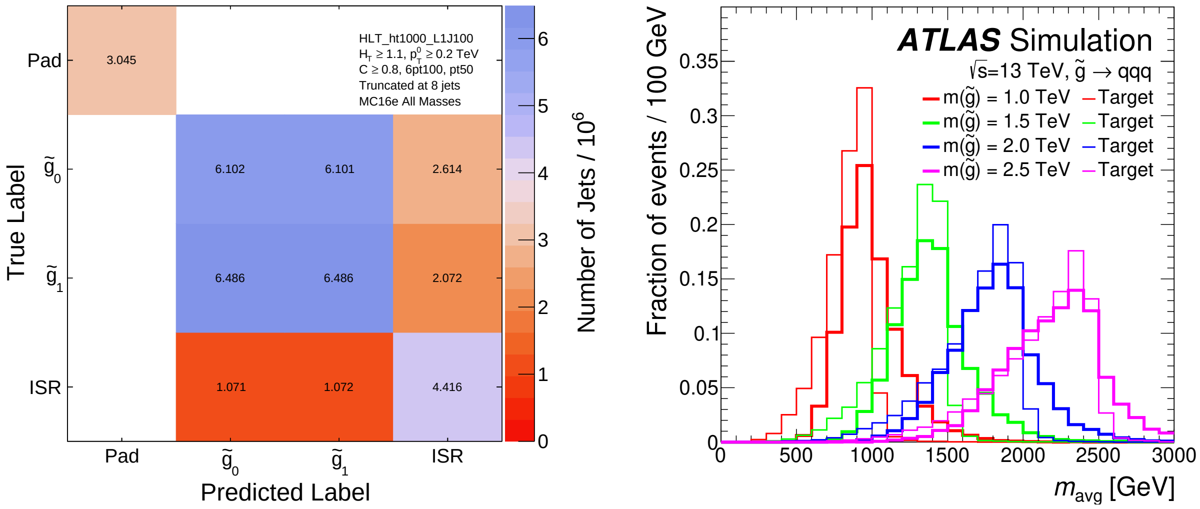

(a) **(b)**

Figure 9.3: Figure (a): confusion matrix for the output classes of the model evaluated on the independent test set. Besides the two gluinos and the ISR classes also padded jets are shown to verify that the padding mask is correctly propagated. Figure (b): normalised m_{avg} spectrum comparing the shapes of the reconstructed (solid) and target (light) distributions for different masses. The reconstructed distribution is produced using the NN assignments, whereas the target distribution is calculated from gluino jets identified by truth parton matching.

9.3 Background modelling

As with the CCA, the largest SM background for this approach is composed of non-resonant QCD events. The corresponding m_{avg} spectrum is expected to be a smoothly falling distribution. The total background is obtained through a fully data-driven method. A functional form is used to directly fit the observed data distribution, employing the PDF-inspired family of parametric functions:

$$f(x) = p_1(1-x)^{p_2}\, x^{p_3+p_4 \ln(x)+p_5 \ln^2(x)+\ldots} \tag{9.10}$$

where $x = m_{avg}/\sqrt{s}$ and $p_{i=1,2,3,\ldots}$ are the freely floating fit parameters. This function has been widely used to model QCD backgrounds in resonance searches by CDF, CMS and ATLAS experiments [238–240].

Initially, a family of exponential functions was also tested. To decide which function to use, the following criteria were considered. The function must perform well under background-only fits, based on the values and probabilities from χ^2 tests, and must produce low *spurious signal* (see section below). The function that satisfies the criteria and has the fewest parameters is preferred. These considerations lead to the decision of using the function from Equation 9.10, with only the first three parameters.

The fit uses a binned maximum likelihood method, implemented within the `HistFitter` framework [241]. The background-only fits have the signal strength parameter (μ) fixed at zero, while in signal-plus-background fits μ is left floating. The fit region starts at 0.7 TeV and extends up to 3 TeV, in bins of 100 GeV.

Background validation

Figure 9.4 shows the tests of the background estimation in data and MC. The method is tested using a small portion of the dataset, in a region with looser cuts that achieve statistical power comparable to that of the SR, avoiding any premature unblinding. The validation dataset corresponds to 3.2 fb^{-1} and loosens the selections to $C \geq 0.7$ and minimum jet p_T to 70 GeV. The validation in MC is performed in the SR, scaling the sample to the full luminosity of the available dataset (140 fb^{-1}). Given the different selections, the data and MC distributions are not expected to perfectly match. In both cases, the function with the best fit parameters has an acceptable agreement with the tested sample.

The modelling bias of the fitting function is estimated by running *spurious signal* tests. If the function being evaluated as a candidate for the background method has N parameters,

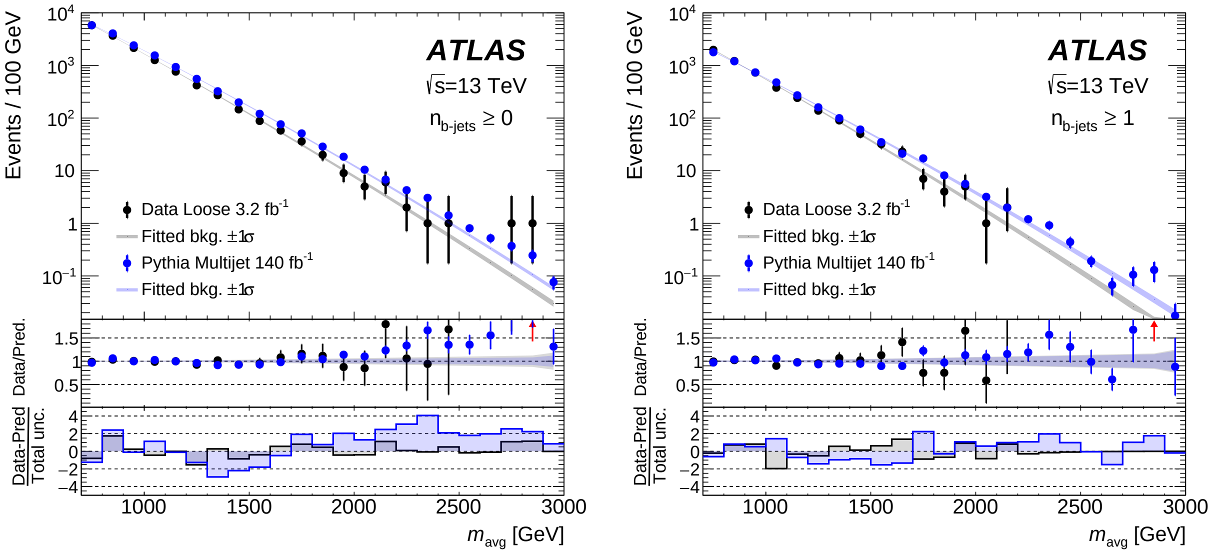

(a) **(b)**

Figure 9.4: Validation of the background method in the (a) nominal and (b) b-tagged regions. The method is tested using the Pythia multijet sample, shown in blue, and a loose selection of data, in black, corresponding to 3.2 fb^{-1}. The cuts applied to the latter are tuned such that the event yields roughly correspond to the ones with SR selections using the full Run 2 luminosity. The grey and blue bands present the combined statistical and systematic uncertainty in the background estimate for the data and MC fit functions respectively. The red arrow denotes points which lie above the range of the ratio plot.

a background-only template is obtained by fitting the data (or MC) m_{avg} distribution with an $N + 1$ parameter function. Many pseudo-experiments are generated by fluctuating the background template. For each experiment, a signal-plus-background fit is performed with the N-parameter function, allowing the signal contribution to be either positive or negative. The average extracted signal events across all experiments does correspond to the spurious signal s_{Sp} that the function can create. Since in this case the true value of signal events is zero, s_{Sp} quantifies the amount of signal that can be artificially generated or absorbed by a biased model of the background. The s_{Sp} together with the standard deviation of the extracted signal events σ_{Sp} are used to determine whether the function is suitable for the background estimate. The requirements to be satisfied are $s_{\mathrm{Sp}} \leq s \cdot 20\%$ and $s_{\mathrm{Sp}}/\sigma_{\mathrm{Sp}} \leq 20\%$, where only gluinos up to $m(\tilde{g}) \leq 1.5$ TeV are considered for signal events (s). The choice of these thresholds includes a degree of arbitrariness, but it generally ensures that s_{Sp} has a reasonably small impact and introduces minimal uncertainty to the background estimate.

Signal injection tests are performed by running both the background-only and signal-plus-background fitting procedures over a sample that contains both signal and background templates. These checks ensure that background function cannot absorb signal, and that the signal can be properly extracted. The linear relationship between the amount of injected signal events and the extracted ones is verified.

9.4 Systematic uncertainties

As done for the Jet Counting method, two categories of systematic uncertainties can be identified: the physics modelling and experimental uncertainties. In the following paragraphs, the errors included as NP in the statistical interpretation of the results are discussed.

Modelling uncertainties

Given that the background estimate is derived using only data, no errors associated with the multijet and $t\bar{t}$ simulations are considered. The uncertainties affecting the signal models are the same as those discussed in Section 8.5: the renormalisation and factorisation scales in the partonic cross-section, PDF, and α_s uncertainties. For each of these, variations of the m_{avg} spectrum are compared to the spectrum obtained from nominal signal simulation. The discrepancies are used as the corresponding systematic errors. The MC statistical uncertainty is included as a NP affecting the normalisation of signal events.

Experimental uncertainties

Similarly to the CCA the impacts of JES and JER are accounted for by propagating the total 31 related uncertainties to the m_{avg} spectrum. For this analysis approach, they only affect the signal yields. The same applies to the b-tagging scale factor errors, in the SR with at least one b-jet. The error on the luminosity is included for scaling the expected signal. No additional uncertainty is applied to the ML algorithm performance, as it simply creates combinations of jets based on their four-momentum. It may result in a suboptimal reconstruction of the event, but does not require any further systematic error.

Background method Two types of dedicated systematic errors are introduced for the background method. The first propagates the uncertainties on the background function parameters obtained through the fit minimisation. The second uncertainty pertains to the choice of the functional form, and how it can bias the background modelling. It is quantified by the spurious signal. A relative uncertainty is included with a Gaussian prior for $\mu_{\text{Sp}} = s_{\text{Sp}}/s$, and it is symmetrised around zero. Figure 9.5 shows the size of this uncertainty, estimated for each reconstructed gluino mass value.

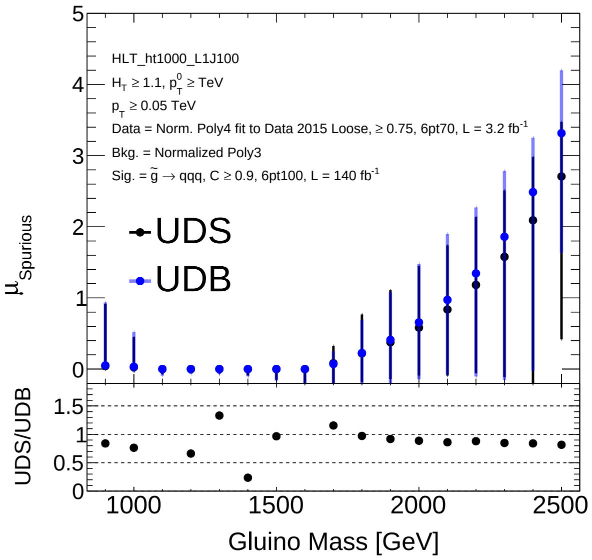

Figure 9.5: The spurious signal uncertainty (μ_{Sp}) per mass point derived from fitting the 3-parameter background function to the pseudo-experiments generated with a background-only template. The latter is obtained by fitting a 4-parameter function to the loose data VR.

Summary

Table 9.1 summarises all the systematic errors included for the Mass Resonance approach. In total 49 different NP are considered for the final interpretation of the results.

Table 9.1: Summary of the uncertainties considered for the Mass Resonance approach.

Uncertainty type	# NP	Short description
MC signal	7	Signal modeling systematics, mostly scale and PDF uncertainties (excluding errors on the cross-section)
MC signal stats	1	Weighted Poisson error, computed as the sum of MC weights $\sqrt{\sum_i w_i^2}$
JES/JER set	31	Jet energy scale and resolution errors, for the Global Reduction JES and Simple JER sets; this uncertainty is applied to signal only
Flavour tagging	5	Uncertainties on *b*-tagging scale factors; this uncertainty is applied to signal only
Luminosity	1	Uncertainty on the total Run 2 luminosity (0.8%), for signal
Fitted parameters	3	Uncertainties on the fitted parameters for the background modelling
Spurious signal	1	Uncertainty related to choice of the functional form, and its ability to fake or absorb signal

Chapter 10

Results and interpretation

The results obtained with the two approaches described in Chapters 8 and 9 are discussed here. After unblinding the full Run 2 p-p collisions dataset, the interpretation of the results uses the statistical tools from Section 7.5. No excesses are observed with respect to the predicted backgrounds, and therefore upper limits on cross-sections are set.

The improvement in sensitivity of the Mass Resonance approach compared to the Jet Counting approach is visible in some signal regions, especially where the additional handles, such as b-tagging, cannot be used to reject the backgrounds. However, owing to its simplicity, the results obtained with the CCA are generalised and reinterpreted with alternative models. The simplified SUSY models initially considered have a maximal intensity for the λ'' RPV couplings, thus the $\tilde{\chi}_1^0$ decays in the 2×5 models are prompt. If the value of the coupling λ'' is decreased, the cascade decay starts to have a finite lifetime. The sensitivity to such models is explored, and new limits are set using the same results obtained from the prompt Jet Counting search.

10.1 Unblinded results

The unblinding procedure occured in two steps. Firstly the regions without the b-tag requirement were unveiled. In the second step, after excluding the lowest $\tilde{g}$ mass points considered and ensuring that no signal contamination was expected in the CR and NR with b-tagging, the remaining regions were unblinded.

For the Jet Counting method, the predicted and counted events are illustrated in Figure 10.1. No significant excesses are observed above the predicted background in any of the signal regions. The data in the inclusive regions (SR1-SR5) is well in agreement with the background. In the b-tagged regions (SR1bj and SR2bj) a small under fluctuation is observed, in agreement with predictions within $\sim 1.6\,\sigma$. Figure 10.2 presents the results for the Mass Resonance method. Consistent with the CCA, no significant data excesses are seen. An under fluctuation of data is observed at $m_{avg} \sim 1.5$ TeV, which is slightly more pronounced in the b-tagged regions – 30% less than the expected background.

Model-independent limits are set for the *visible cross-section*, defined as the product of cross-section, acceptance, and efficiency ($\langle \epsilon\sigma \rangle^{95}_{obs}$). These upper limits are set using a profile likelihood test statistic and the CL$_s$ method (equation 7.11). The expected limits are obtained by setting the data equal to the background estimate. After having ensured that the likelihood is well in agreement with a Gaussian shape, the p-values are calculated using an asymptotic approximation [229] for the pdf of the test statistic q_μ from Equation 7.10. An exception is made for regions featuring few events in the very high-mass part

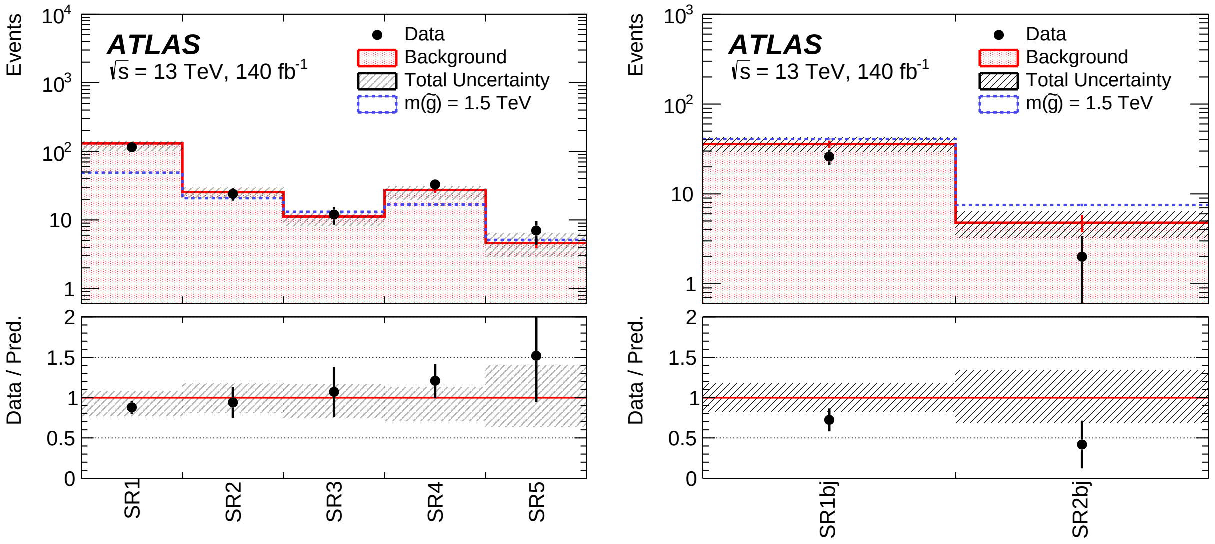

(a) (b)

Figure 10.1: Observed and predicted yields in the signal regions of the Jet Counting method for the (a) b-tagging inclusive and (b) ≥ 2 b-tags regions. The bottom panel shows the ratio of the data to the background prediction. The hatched pattern represents total uncertainty on the background estimate, combining the statistical and systematic errors.

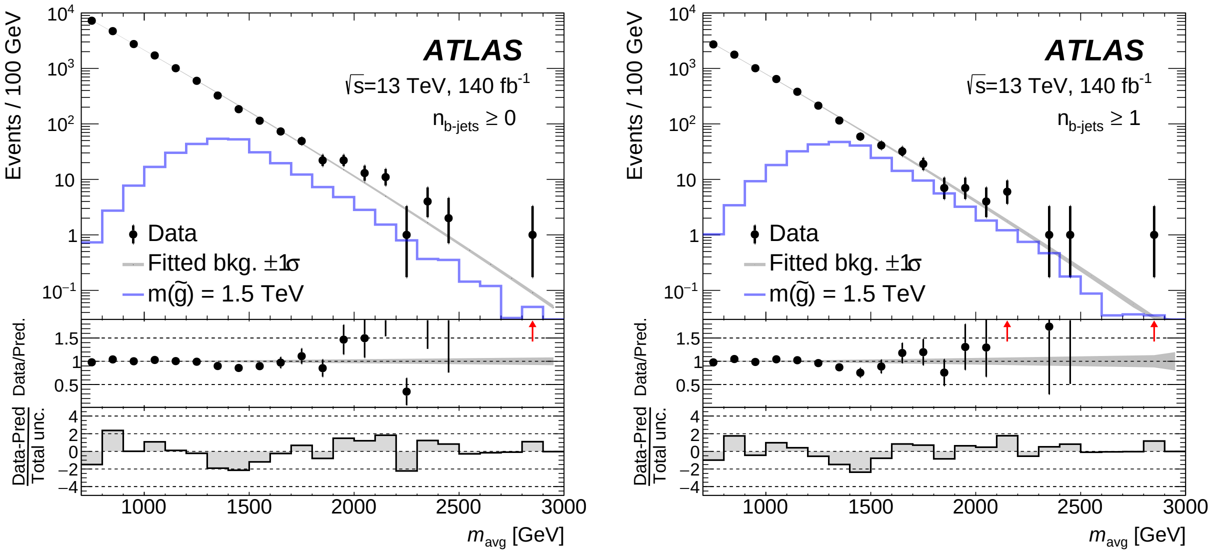

(a) **(b)**

Figure 10.2: Observed data distribution and the background-only fit of the reconstructed average mass spectrum for the (a) b-tagging inclusive and (b) ≥ 1 b-tags regions. The grey bands include both the statistical and systematic uncertainties. The red arrows denote points that lie above the range of the ratio plot.

Table 10.1: The upper limit table for the signal regions for the jet counting method. Left to right: 95% CL upper limits on the visible cross-section ($\langle\epsilon\sigma\rangle^{95}_{\text{obs}}$) and on the number of signal events (S^{95}_{obs}). The third column (S^{95}_{exp}) shows the 95% CL upper limit on the number of signal events, given the expected number (and $\pm 1\sigma$ excursions on the expectation) of background events. The last two columns indicate the $1 - \text{CL}_\text{b}$ value, *i.e.* one minus the confidence level observed for the background-only hypothesis, and the discovery p-value (p_0), with its corresponding Gaussian significance (Z). The p_0 measures the compatibility of the observed data with the background-only (zero signal strength) hypothesis, relative to fluctuations of the background. Larger values indicate greater relative compatibility. In signal regions with a deficit relative to the nominal background prediction, the p_0 value is capped at 0.50.

Signal region	$\langle\epsilon\sigma\rangle^{95}_{\text{obs}}$[fb]	S^{95}_{obs}	S^{95}_{exp}	$1 - \text{CL}_\text{b}$	p_0 (Z)
SR1	0.32	45	57^{+18}_{-14}	0.49	0.50 (0.00)
SR2	0.09	13	$14.1^{+5.7}_{-4.1}$	0.44	0.50 (0.00)
SR3	0.07	10	$9.5^{+4.1}_{-2.7}$	0.48	0.42 (0.20)
SR4	0.16	22	$17.4^{+6.5}_{-4.7}$	0.74	0.21 (0.79)
SR5	0.07	9.4	$7.4^{+3.6}_{-2.4}$	0.58	0.32 (0.46)
SR1bj	0.08	11	$17.0^{+6.9}_{-4.8}$	0.45	0.50 (0.00)
SR2bj	0.03	4.4	$6.6^{+2.9}_{-1.9}$	0.34	0.50 (0.00)

Table 10.2: The upper limit table for the ≥ 0 b-tagged jets region of the mass resonance method. Left to right: 95% CL upper limits on the visible cross-section ($\langle\epsilon\sigma\rangle^{95}_{\mathrm{obs}}$) and on the number of signal events (S^{95}_{obs}). The third column (S^{95}_{exp}) shows the 95% CL upper limit on the number of signal events, given the expected number (and $\pm 1\sigma$ excursions on the expectation) of background events. The last two columns indicate the $1 - \mathrm{CL}_b$ value, *i.e.* one minus the confidence level observed for the background-only hypothesis, and the discovery p-value (p_0), with its corresponding Gaussian significance (Z). The p_0 measures the compatibility of the observed data with the background-only (zero signal strength) hypothesis, relative to fluctuations of the background. Larger values indicate greater relative compatibility. In signal regions with a deficit relative to the nominal background prediction, the p_0 value is capped at 0.50.

m_{avg} range [GeV]	$\langle\epsilon\sigma\rangle^{95}_{\mathrm{obs}}$[fb]	S^{95}_{obs}	S^{95}_{exp}	$1-\mathrm{CL}_b$	p_0 (Z)
$700-1000$	7.30	1000	1300^{+460}_{-300}	0.22	$0.50\ (0.00)$
$800-1100$	5.70	800	360^{+150}_{-49}	0.99	$0.01\ (2.50)$
$900-1200$	2.10	290	210^{+88}_{-25}	0.81	$0.18\ (0.91)$
$1000-1300$	1.50	210	160^{+50}_{-34}	0.80	$0.18\ (0.90)$
$1100-1400$	0.54	76	120^{+45}_{-30}	0.09	$0.50\ (0.00)$
$1200-1500$	0.27	37	85^{+33}_{-24}	0.00	$0.50\ (0.00)$
$1300-1600$	0.16	23	63^{+37}_{-18}	0.00	$0.50\ (0.00)$
$1400-1700$	0.16	22	47^{+19}_{-13}	0.00	$0.50\ (0.00)$
$1500-1800$	0.24	33	39^{+16}_{-10}	0.25	$0.50\ (0.00)$
$1600-1900$	0.26	37	38^{+15}_{-10}	0.47	$0.50\ (0.00)$
$1700-2000$	0.30	42	34^{+12}_{-7}	0.71	$0.29\ (0.55)$
$1800-2100$	0.25	35	28^{+12}_{-8}	0.72	$0.28\ (0.57)$
$1900-2200$	0.29	41	25^{+11}_{-4}	0.93	$0.06\ (1.50)$
$2000-2300$	0.19	27	$21.5^{+7.6}_{-4.4}$	0.78	$0.19\ (0.89)$
$2100-2400$	0.15	21	$15.5^{+6.2}_{-2.3}$	0.74	$0.20\ (0.84)$
$2200-2500$	0.08	11	$10.5^{+3.2}_{-1.9}$	0.57	$0.40\ (0.26)$
$2300-2600$	0.08	11	$9.2^{+3.9}_{-1.2}$	0.66	$0.27\ (0.61)$
$2400-2700$	0.05	6.9	$6.8^{+2.1}_{-1.4}$	0.51	$0.48\ (0.05)$
$2500-2800$	0.02	2.3	$3.1^{+2.1}_{-1.2}$	0.26	$0.50\ (0.01)$
$2600-2900$	0.04	5.3	$5.2^{+2.2}_{-1.3}$	0.52	$0.46\ (0.10)$
$2700-3000$	0.06	8.3	$8.2^{+0.4}_{-0.7}$	0.53	$0.44\ (0.16)$

Table 10.3: The upper limit table for the ≥ 1 b-tagged jets region of the mass resonance method. Left to right: 95% CL upper limits on the visible cross-section ($\langle \epsilon\sigma \rangle^{95}_{\text{obs}}$) and on the number of signal events (S^{95}_{obs}). The third column (S^{95}_{exp}) shows the 95% CL upper limit on the number of signal events, given the expected number (and $\pm 1\sigma$ excursions on the expectation) of background events. The last two columns indicate the $1 - \text{CL}_b$ value, *i.e.* one minus the confidence level observed for the background-only hypothesis, and the discovery p-value (p_0), with its corresponding Gaussian significance (Z). The p_0 measures the compatibility of the observed data with the background-only (zero signal strength) hypothesis, relative to fluctuations of the background. Larger values indicate greater relative compatibility. In signal regions with a deficit relative to the nominal background prediction, the p_0 value is capped at 0.50.

m_{avg} range [GeV]	$\langle \epsilon\sigma \rangle^{95}_{\text{obs}}$[fb]	S^{95}_{obs}	S^{95}_{exp}	$1 - \text{CL}_b$	p_0 (Z)
700 − 1000	5.70	800	960^{+330}_{-240}	0.31	0.50 (0.00)
800 − 1100	3.30	460	320^{+100}_{-65}	0.89	0.11 (1.20)
900 − 1200	1.10	150	130^{+38}_{-31}	0.74	0.24 (0.71)
1000 − 1300	0.92	130	92^{+40}_{-13}	0.81	0.18 (0.91)
1100 − 1400	0.36	51	70^{+27}_{-20}	0.17	0.50 (0.00)
1200 − 1500	0.16	23	52^{+21}_{-15}	0.00	0.50 (0.00)
1300 − 1600	0.11	16	39^{+15}_{-11}	0.00	0.50 (0.00)
1400 − 1700	0.12	17	29^{+12}_{-8}	0.04	0.50 (0.00)
1500 − 1800	0.20	27	25^{+25}_{-7}	0.61	0.38 (0.29)
1600 − 1900	0.25	35	30^{+10}_{-7}	0.68	0.45 (0.13)
1700 − 2000	0.21	30	28^{+10}_{-8}	0.58	0.42 (0.20)
1800 − 2100	0.17	24	$24.0^{+5.9}_{-6.2}$	0.51	0.49 (0.03)
1900 − 2200	0.18	25	$21.6^{+5.9}_{-5.8}$	0.71	0.26 (0.65)
2000 − 2300	0.13	18	$17.1^{+5.3}_{-2.1}$	0.63	0.32 (0.47)
2100 − 2400	0.10	13	$12.4^{+3.3}_{-2.6}$	0.63	0.30 (0.51)
2200 − 2500	0.05	6.4	$6.4^{+2.5}_{-1.5}$	0.50	0.50 (0.00)
2300 − 2600	0.05	6.8	$6.7^{+2.6}_{-0.8}$	0.54	0.42 (0.20)
2400 − 2700	0.03	4.0	$3.9^{+2.2}_{-1.2}$	0.52	0.45 (0.14)
2500 − 2800	0.01	2.0	$2.1^{+1.8}_{-0.9}$	0.47	0.49 (0.02)
2600 − 2900	0.04	5.4	$5.3^{+2.2}_{-1.3}$	0.53	0.43 (0.19)
2700 − 3000	0.04	6.1	$6.0^{+2.3}_{-0.6}$	0.53	0.42 (0.20)

of the m_{avg} spectrum. In these cases the assumptions of the approximation might break, and pseudo-experiments are needed to evaluate the distributions of q_{μ}. By the same means, the asymptotic approximation is validated for other regions with moderately small yields. The resulting limits are reported in Tables 10.1, 10.2 and 10.3. For the CCA, all the observations are compatible with the background-only hypothesis within $Z \leq 1\sigma$. With the Mass Resonance method, limits ranging from 7.3 fb down to 0.01 fb are set for the visible cross-section.

Model-dependent limits are shown in Figure 10.3 for direct decays as a function of the gluino mass. For both methods, the final upper limits are evaluated for each signal model only in the regions that yield the best expected limit. A combination of the two methods is performed by taking the more sensitive between the two approaches for each model, and the corresponding expected and observed limit are quoted as the result of the search. On the one hand, it is evident that the Mass Resonance method extends the reach to higher gluino masses, by excluding gluinos with masses up to 1730 GeV and 1800 GeV for the UDS and UDB, respectively. On the other hand, the CCA has competitive sensitivity and can set tighter constraints for some lighter models ($m(\tilde{g}) < 1.5$ TeV), particularly for the b-tagged regions. Figure 10.4 shows the exclusion contours for the cascade decay models. Only the Jet Counting method is used. As in the previous case, hypothesis testing with data is performed for each model only in the region with the best expected limits. A wide region of the $m(\tilde{g})$-$m(\tilde{\chi}_1^0)$ plane is explored and gluinos with masses up to 2230 (2340) GeV are excluded for a neutralino with 1250 GeV mass and UDS (UDB) coupling.

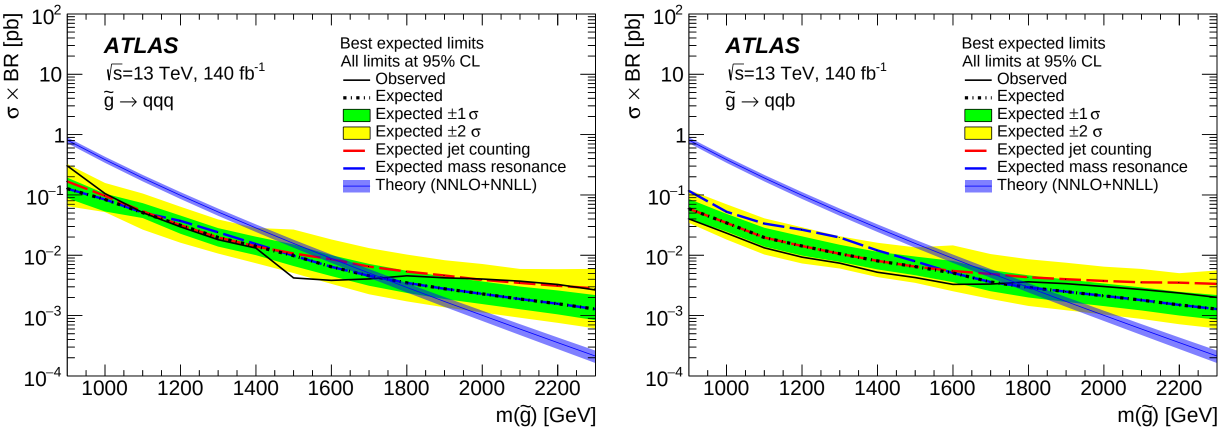

(a) **(b)**

Figure 10.3: Observed and expected limits on the signal cross-section production times branching ratio ($\sigma\times$BR) as a function of the gluino mass for the gluino direct decay model with (a) UDS and (b) UDB decays. The expected limits for the Jet Counting and Mass Resonance methods are shown in red and blue, respectively. The best expected limit per mass point between the methods is chosen (dashed black) and corresponding observed limit reported (solid black). The green and yellow bands around the expected limit correspond to the $\pm1\sigma$ and $\pm2\sigma$ variations including both systematic and statistical uncertainties. The theoretical prediction with its uncertainty is shown as blue line with a coloured band.

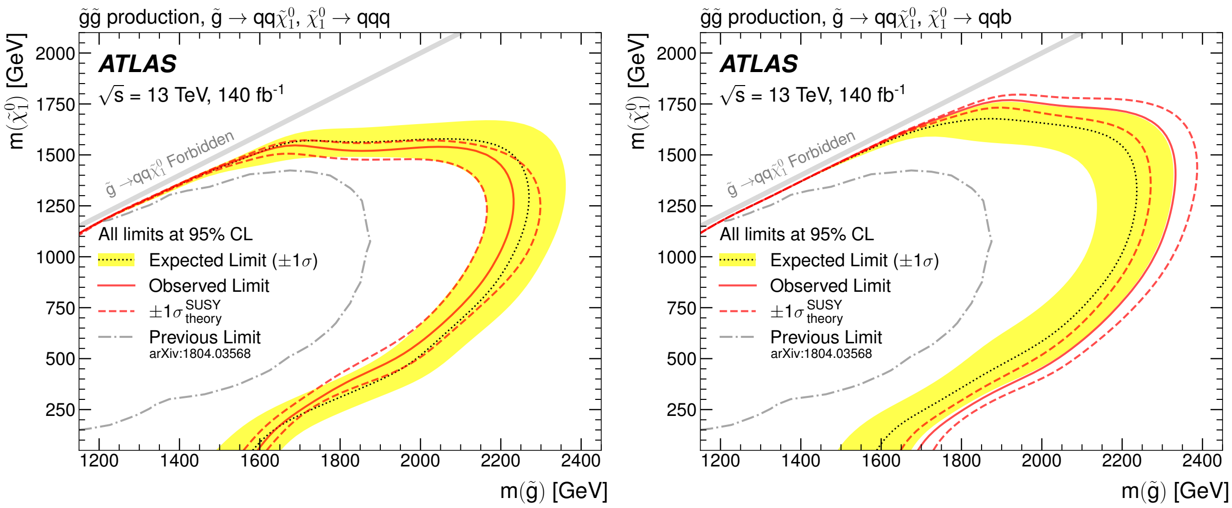

(a) **(b)**

Figure 10.4: Observed and expected exclusion contours for the gluino cascade decay model with (a) UDS and (b) UDB decays. The contours of the band around the expected limit correspond to the $\pm 1\sigma$ variations including both systematic and statistical uncertainties. The dotted lines around the observed limit illustrate the change in the observed limit as the nominal signal cross-section is scaled up and down by the theoretical uncertainty. The diagonal line indicates the kinematic limit for the decay of the gluino.

Comparing the results with the previous ATLAS search (Figure 7.2), the implemented analysis strategy has largely improved the signal sensitivity. The increase in significance expected by the larger dataset, $Z^{\text{Full Run-2}}/Z^{\text{Par. Run-2}} \sim \sqrt{140\,\text{fb}^{-1}/36\,\text{fb}^{-1}} \simeq 2$, was out-scaled, and, in some regions, cross-sections close to a factor 10 smaller were probed.

10.2 RPC-RPV reinterpretation using RECAST

In general, for practical reasons, most of the SUSY searches focus on specific simplified models. A few decay channels are considered at a time, neglecting several features of more complete theories. For example, if the lightest supersymmetric particle (LSP) is a neutralino $\tilde{\chi}_1^0$ and RPV couplings are set to be very small ($\ll 1$), the LSP can have a finite lifetime depending on the other sparticles masses. In some cases simplified traditional searches are sensitive to these signatures. The results obtained with the Jet Counting method are reinterpreted in more elaborate models as part of a larger effort to evaluate the coverage of RPC-RPV intermediate scenarios. In this section, only models with $\tilde{g}$ decaying via different values of λ''_{112} coupling are discussed. As mentioned earlier in Section 7.1, the results extend to cases where the generation indexes of λ''_{ijk} are $i = 2$ and $j, k = 2, 3$.

In the benchmark theories taken into account for the reinterpretation, the LSP is assumed to a pure bino-like $\tilde{\chi}_1^0$ with a mass of $m(\tilde{\chi}_1^0) = 200\,\text{GeV}$. The mass is selected to allow decays to top quarks, and the bino-like requirement is imposed to simplify the spectrum of the neutralino decays.

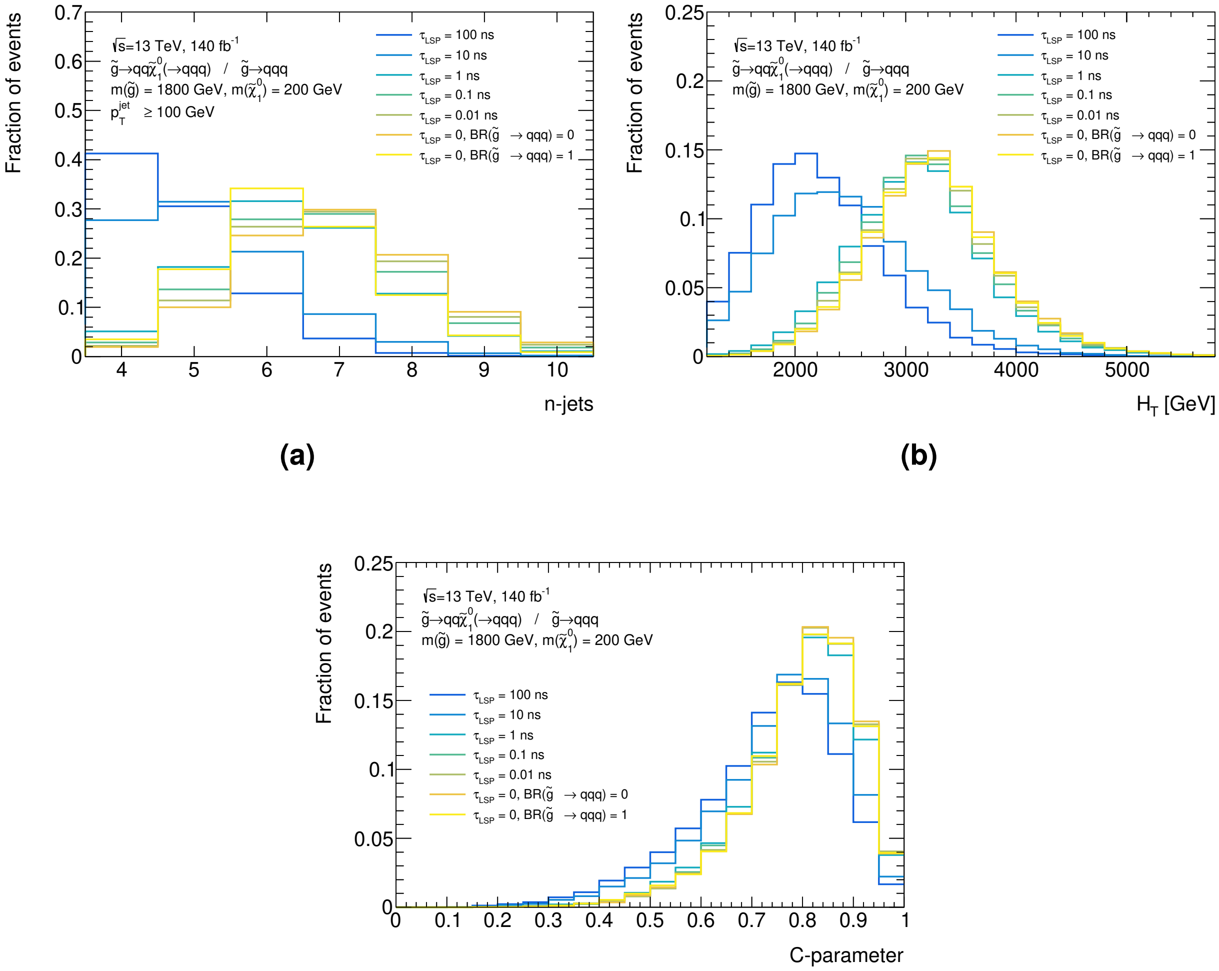

(a)

(b)

(c)

Figure 10.5: Distributions normalised to unit area of (a) the number of jets above 100 GeV, (b) the H_T, and (c) the event-shape C, for gluino and neutralino masses of 1.8 TeV and 200 GeV respectively. Models with different lifetime of the neutralino are shown with different colour and all of them feature a cascade decay BR of 100%, except for the prompt decay shown in bright yellow.

Figure 10.5 shows some kinematic distributions for various lifetimes and BRs. When τ_{LSP} increases, the neutralinos gain a macroscopic decay length, which reduces the number of energetic jets that are formed in the detector (Figure 10.5(a)), and consequently lowers the total hadronic energy deposited in the transverse plane (Figure 10.5(b)). The event-shapes are also affected, the detectable decay products become less balanced, reducing the value of C (Figure 10.5(c)). These changes become evident when $\tau_{LSP} > 1$ ns. Above this threshold the expected displacement reaches the order of the meters ($\beta\gamma c\tau_{LSP} \gtrsim 1$ m), assuming $\beta\gamma \sim 1$. Thus, the neutralinos can easily escape the calorimeters volume.

To run the statistical inference on these alternative samples, a RECAST [242] workflow for the Jet Counting method is implemented using the `Yadage` and `Packtivity` toolkits [243]. The full event selection process is preserved, along with the observed data and background expectation values for each analysis' SR. The framework for the statistical interpretation is also integrated in the workflow. An additional systematic uncertainty is included to account for changes in the JES that depend on the displacement of $\tilde{\chi}_1^0$ decay (L_{xy}). It is estimated [244] that the p_T response per jet increases by up to 30% for decays occurring far from the beam axis. A per event error proportional to the L_{xy} of the two neutralinos is included, but becomes relevant only for samples with $\tau_{\text{LSP}} \geq 1$ ns.

Executing the RECAST workflow on the MC samples yields the CL_s values as output for all the SRs. For each signal model, the cross-section upper limit is derived by considering only the outcome of hypothesis tests in the region with maximum expected sensitivity. Figure 10.6 displays the resulting exclusion contours in green, labelled as “RPV Multijet”. Two other analyses, the “RPC 0L” [245] and the “DV+Jets” [246], are reinterpreted and compared in the same plane. The first is a search originally designed for RPC signatures, with many energetic jets and large p_T^{miss} in the final state. The second focuses on long-lived particles decaying into hadrons. It reconstructs displaced vertices by leveraging large radius tracking [247] up to a distance of 300 mm from the beam axis.

The three searches are complementary and cover different regions of the phase space. The RPC 0L sets the most competitive limits at very small λ'' values, near the RPC regime. The DV+Jets has the best coverage for scenarios where the $\tilde{\chi}_1^0$ has a large decay length that fits within the inner tracker. For higher λ'', the RPV Multijet has unique sensitivity and tests signal hypotheses in regions of the phase space not accessible to other searches. Overall, gluino masses from ~ 1.7 TeV and up to ~ 2.6 TeV are excluded for various intensities of the RPV coupling.

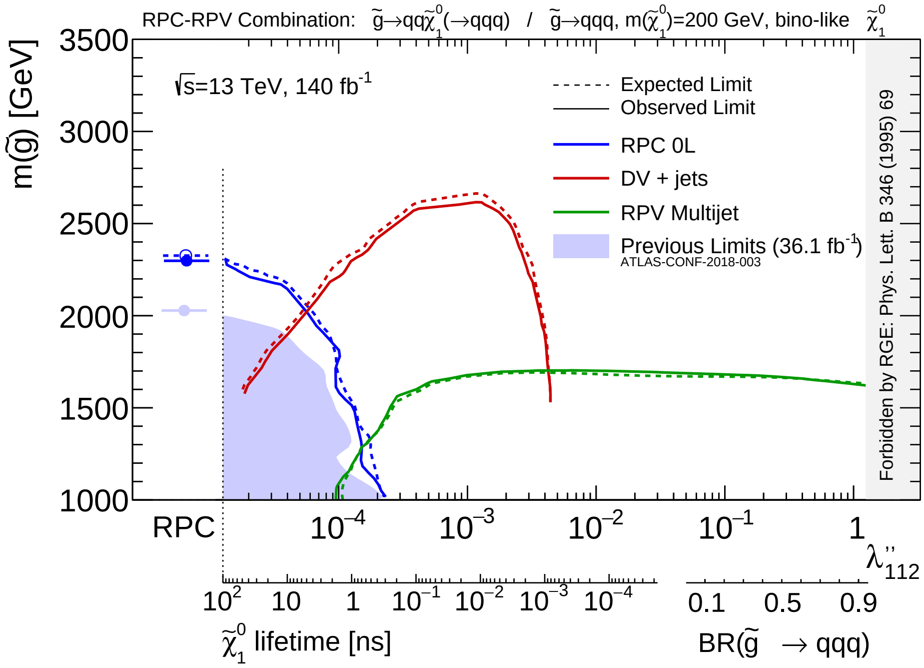

Figure 10.6: Exclusion contours for the gluino decaying into quarks models as a function of λ''_{112} and $m_{\tilde{g}}$. Expected and observed limits are shown with dashed and solid lines. The reinterpretations of RPC 0L [245], the DV+Jets [246], and the RPV Multijet (this thesis) analyses are respectively shown in blue, red, and green. The shaded light-blue area represents the combined total area excluded by Partial Run 2 analyses [208]. The RPC-limit is shown on the leftmost part of the axes.

Outlook II – Future searches

The presented search extended the discovery potential of ATLAS by introducing a novel analysis method that leverages ML techniques, and by improving the traditional CCA approach with enhanced background modelling and better-designed search regions. With the Mass Resonance analysis the expected exclusion increases from 1.6 TeV to nearly 1.75 TeV gluinos for the direct decay modes, meaning that it was able to extend the sensitivity to cross-sections smaller than a factor of 2. In b-tagged regions the improvement is similar, but slightly smaller. A relevant advantage of the Mass Resonance approach resides in its background method, which avoids relying on complex QCD simulations. In contrast, the optimisation of Jet Counting to different models such as gluino cascade decays is straight forward, making it versatile. The strength of the latter method is further demonstrated through the reinterpretation of the results considering more sophisticated models.

Reaching sensitivity to heavier gluinos will require even more powerful background discrimination methods and much higher statistics, as the production cross-section decreases significantly for large masses. At the lower end of the mass spectrum, searches in multijet final states become instead limited by the trigger selection.

Soft gluinos have been excluded by previous experiments [248]. However, the SM superpartners that are not produced via the strong coupling, such as winos or higgsinos, feature far lower cross-sections – three orders of magnitude smaller. The RPV-coupled decays for such particles have not yet been probed by ATLAS. These models produce rather soft hadronic signatures that are barely captured by the trigger system due to tight thresholds required to handle the overwhelming background rates. To overcome this limitation, the TLA paradigm can be adopted, trading resolution in the analysis objects reconstruction to access softer final states that otherwise could not be recorded.

Owing to the Trigger Menu developments for Run 3, including the methodology from Chapter 6, several chains have been implemented in the TLA stream for data-taking, aiming to record low mass multi(-b-)jet resonances. These chains are listed in Table 10.4. For the H_T chain, the benefit of the TLA approach is evident, as the trigger threshold is halved. For the multijet trigger selections the gain might seem less striking. However, since the jet p_T spectra are roughly exponential, even small threshold shifts can have a significant impact. Furthermore, for 2025 data-taking, the TLA version of the "4-jets $(2b)$" chain is being moved to the FTAGPEB stream, introduced in Chapter 6.4, and will thereby benefit from PEB reconstruction for flavour tagging.

Figure 10.7 compares the trigger acceptance of the standard triggers and the TLA stream to light SUSY states decaying into hadrons. Gluinos were used as a benchmark since their final state kinematics would be identical to the higgsino one. From this simplified study, the TLA approach can significantly increase trigger acceptance of benchmark models, becoming a promising avenue to search for such new physics candidates. The TLA and PEB approaches open new possibilities for BSM searches in multijets final states that are beyond the capacity of the standard triggering strategies.

Table 10.4: Comparison of the analysis level thresholds between Physics Main and the corresponding TLA chains, for three trigger chains targeting hadronic final states. The standard thresholds are on offline reconstructed objects, while for TLA stream the thresholds are on HLT objects.

Signature	Analysis Thresholds [GeV] Physics Main Stream	 TLA Stream
4-jets $(2b)$	$p_T \geq 100, 60, 60, 30$ $+\ 2\ (3)$ tight (medium) b-tags	$p_T \geq 100, 60, 60, 20$ $+2$ loose b-tags
6-jets	$p_T \geq 50, 50, 50, 50, 40, 40$	$p_T \geq 50, 50, 50, 50, 25, 25$
H_T	$H_T \geq 1100$	$H_T \geq 550$

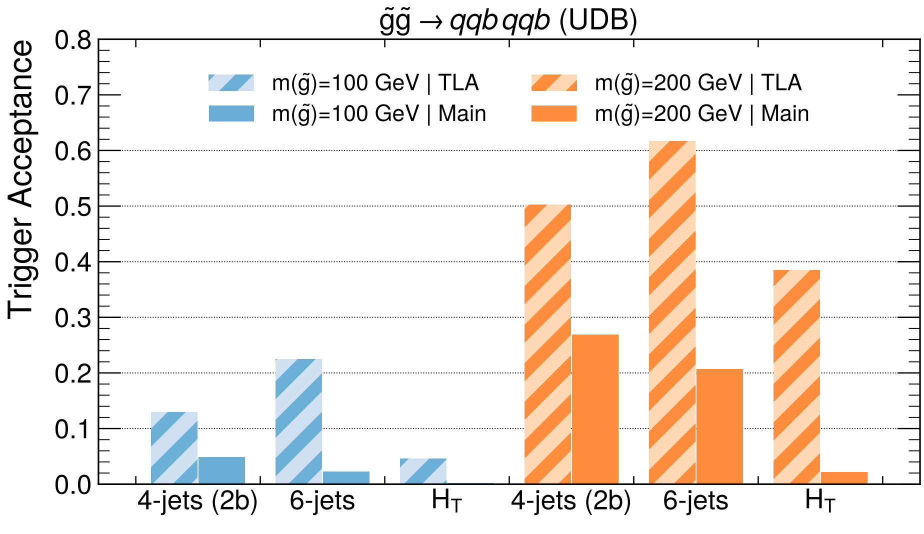

Figure 10.7: Fraction of accepted signal events by the trigger chains from Table 10.4. Pair produced gluinos with direct decays to hadrons of masses of $100\,\text{GeV}$ and $200\,\text{GeV}$ are shown respectively in blue and orange. A UDB coupling is assumed. The hatched pattern bars represent the TLA Stream trigger selection acceptances, while the homogeneously filled bars represent the Physics Main Stream ones.

Conclusions

The SM remains the most precisely tested theory of particle physics, despite knowingly providing an incomplete picture. Some of its key parameters, such as the Higgs self-coupling strength, are yet to be determined and may have strong implications for our understanding of nature. Extensions of the SM are required to explain experimental observations, such as DM or neutralino masses. Many BSM theories predict new particles that could be observed at the LHC; therefore, it is crucial to “leave no stone unturned”.

This work contributed to enhancing ATLAS potential to further expand the boundaries of SM knowledge and discover new physics. The novel fast b-tagging methodology has substantially improved the ATLAS trigger system. It enabled looser selection thresholds, resulting in great acceptance gains for multijet signatures. Using this approach, the $\mathrm{HH} \to \mathrm{b\bar{b}b\bar{b}}$ acceptance in Run 3 increased by nearly 50% relative to the Run 2 trigger configuration. Moreover, the Run 3 TLA stream incorporates flavour tagging information, which may further improve background rejection using the PEB technique in 2025 data-taking. This enables sensitivity to new phase space regions that were previously limited by the trigger.

This thesis scrutinised the entire Run 2 dataset ($140\,\mathrm{fb}^{-1}$) to search for new physics, looking for excesses over background expectations in inclusive regions or localised bumps. Events with many energetic jets in a nearly spherical distribution were selected, as they could contain heavy resonances, such as pair-produced gluinos decaying via RPV couplings. Two alternative approaches were used to analyse the data: an improved cut-and-count workflow and a new event-reconstruction technique that leverages ML algorithms. Neither method produced a significant excess compared to the expected background events, leading to stringent cross-section limits for the production of new physics in this channel. Furthermore, the cut-and-count results were utilised to constrain additional supersymmetric models covering a broad range of SUSY parameter space. The analysis methods significantly improved sensitivity compared to previous similar searches. Further extending the discovery potential at high masses will require greater statistical powers, while improvements in the trigger may allow probing soft resonances at considerably smaller cross-sections in the Run 3 TLA dataset.

Appendix A

Machine Learning

Machine learning (ML) algorithms are excellent tools for analysing HEP data, which are vast and complex, typically with high dimensionality. Moreover, large amounts of labelled data are accessible through MC simulations. Over the past decade, ML techniques have become increasingly present in HEP [249], partly due to their versatility. For instance, ML is employed for event classification to distinguish between signal events and background processes. It is also widely employed for object reconstruction and identification. Furthermore, machine learning aids in sophisticated regression tasks, such as background modelling and object calibration. This appendix gives a brief introduction to the basic concepts of ML.

Key ingredients

The typical application of machine learning in particle physics follows a structured workflow, as illustrated in Figure A.1. Firstly, a large dataset is gathered, usually from simulations. In the *supervised learning* approach, this data must be labelled according to the *true* characteristics that the ML model aims to learn. The labels can be, for example, the type of particle or the type of event (signal vs. background). If the dataset is only partially

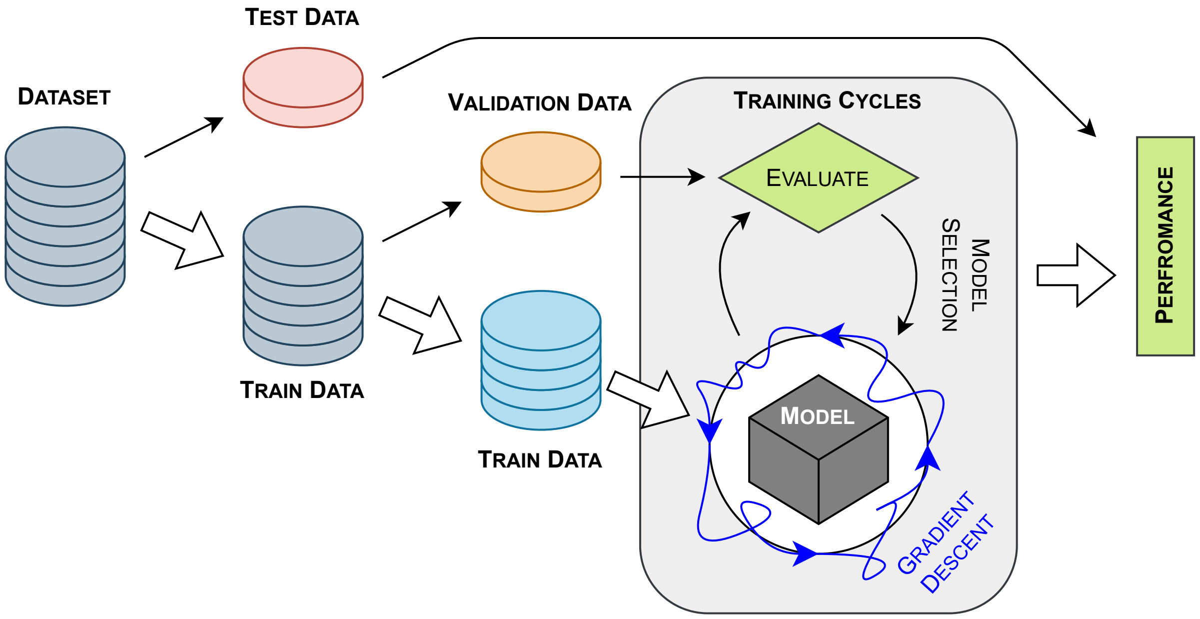

Figure A.1: Diagram of a classical workflow to train ML algorithms.

labelled, or not labelled at all, *semi-supervised* or *unsupervised learning* techniques can be used for example to cluster data entities based on complex correlations.

Data partitioning In supervised frameworks, the dataset is separated into *training* and *test* samples. The test sample is put aside and used only after the training is completed to obtain an unbiased estimate of the final model's performance. Often, a small *validation* sample is partitioned from the training data. While the bulk of the training sample is used to optimise the parameters of the ML model, the validation sample is used to evaluate the performance during the training, allowing for the tuning of several hyperparameters. For example, the accuracy on the validation data can be used to select different model variations and for *early stopping* of the training if the model starts to overfit the input data.

Loss functions The data usually consist of a set of vectors $x^{(i)}$, where the components are meaningful physical quantities relevant to a given task. For each element in the dataset (X) the associated label $y^{(i)}$ can be either a continuous or a discrete (categorical) variable. ML models are usually described by a large number of parameters (often referred to as trainable parameters), here represented by the vector θ, which are also referred to as *trainable* parameters. The goal of the training is to find the optimal values for the parameters θ of a given model M, such that its predictions $\hat{y}^{(i)} = M(x^{(i)}; \theta)$ are as close as possible to the true labels $y^{(i)}$. The deviation of $\hat{y}^{(i)}$ from $y^{(i)}$ is quantified by a *loss function* L. For regressions, it is common to use the mean square error (MSE):

$$L_{\mathsf{MSE}}(\theta;\mathbf{X}) = \frac{1}{N}\sum_{i=1}^{N}\left(y^{(i)} - \hat{y}^{(i)}\right)^2 . \tag{A.1}$$

Generally, in classification problems, the evaluation of the model M on the data returns the probability of belonging to a given class. For the simplest scenario, a binary classifier, $\hat{y}^{(i)}$ is a scalar value between 0 and 1, and the *cross-entropy* loss is most often used:

$$L_{\text{cross-entropy}}(\theta;\mathbf{X}) = -\frac{1}{N}\sum_{i=1}^{N}\left[y^{(i)} \cdot \log\left(\hat{y}^{(i)}\right) + \left(1 - y^{(i)}\right) \cdot \log\left(1 - \hat{y}^{(i)}\right)\right] , \tag{A.2}$$

which represents the negative log-likelihood of a Bernoulli distribution. In the case of multiple classes, the output of $M(\mathbf{x}^{(i)};\theta) = \hat{\mathbf{y}}^{(i)}$ is a vector, where each component $\hat{y}_c^{(i)}$ represent the predicted probability for class c. In this case, the cross-entropy loss becomes:

$$L_{\text{cross-entropy}}(\theta;\mathbf{X}) = -\frac{1}{N}\sum_{i=1}^{N}\sum_{c=1}^{n_c} y_c^{(i)} \cdot \log\left(\hat{y}_c^{(i)}\right) . \tag{A.3}$$

Gradient descent Once a loss is the defined, the training of an algorithm becomes a minimisation task. This problem is solved numerically using the *gradient descent* approach, an iterative process that searches for minima following the gradient of the loss. The parameters are updated after each iteration i according to:

$$\theta_{i+1} = \theta_i - \alpha_{\mathsf{LR}}\nabla_\theta L(\theta_i;\mathbf{X}) . \tag{A.4}$$

In this equation, α_{LR} is the *learning rate*, in other words, it is the fraction of the gradient used as a step in the parameter space at each iteration. If this rate is too large, the convergence to the minimum might be compromised; conversely, if it is too small, the training procedure might converge only to a local minimum.

Equation A.4 represents the standard gradient descent, where each parameter update requires the computation over the full dataset. The stochastic gradient descent (SGD) is a more computationally efficient approach. In this case, the gradient is computed over one or few data instances (*mini-bacthes*) at each iteration. One of the most used methods, the *Adam* optimiser [250], combines SGD with adaptive learning, adjusting the learning rate at each step based on the moments of the gradient distribution.

The gradient-based methods require that the loss function $L(\theta;x)$ is sub-differentiable with respect to θ. In fact, the gradient term usually cannot be calculated analytically. However, models usually consist of nested mathematical operations, allowing the chain rule to be used:

$$\frac{dx_1}{dx_2}\frac{dx_2}{dx_3}\cdots\frac{dx_{n-1}}{dx_n} = \frac{dx_1}{dx_n} . \tag{A.5}$$

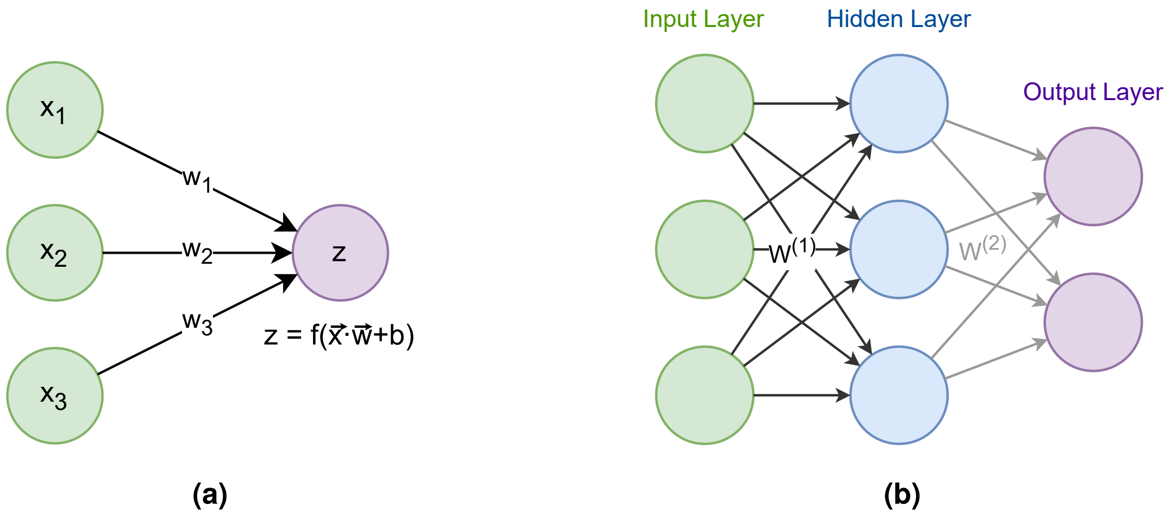

Figure A.2: Diagrams of a single (a) neuron and (b) a simple fully connected NN.

Backpropagation [251] is the standard algorithm used to compute these long chains of nested gradients efficiently.

Neural Networks The majority of the ML architectures comprise different arrangements of Neural Networks (NNs). The wide usage of NNs is due to their versatility; in fact, they are universal approximators [252]. In principle, with enough parameters, they can describe any arbitrary non-linear function.

The basic units of NNs are the *neurons*, also known as *perceptrons*. As shown in Figure A.2(a), an artificial neuron applies an *activation function* f to the scalar product of the input features vectors x with the *weights* vector w, plus a *bias* term b, applying the mapping

$$x \to f(x \cdot w + b)\ . \tag{A.6}$$

A NN is simply a combination of multiple neurons (or nodes), arranged in layers, as shown in Figure A.2(b). In this case, the action of each layer involves a linear transformation of each layer, upgrading the vector w to a matrix W. Taking as an example the *feed-forward* NN from Figure A.2(b), the mapping in this case is:

$$x \to g\left(\mathbf{b}^{(2)} + W^{(2)} f\left(\mathbf{b}^{(1)} + W^{(1)} x\right)\right)\ , \tag{A.7}$$

where $W^{(1)} \in \mathbb{R}^{3\times3}$ and $W^{(2)} \in \mathbb{R}^{2\times3}$, and the two bias terms $\mathbf{b}^{(1)}$ and $\mathbf{b}^{(2)}$ are vectors of dimension three and two, respectively. Furthermore, the functions f and g are to be applied element-wise to the vectors produced by the transformations, assuming that all neurons within a given layer have the same activation function type.

Activation functions The non-linearity in NNs is introduced by the activation functions. Many options exist, and better choices of functions may depend on the specific task. One of the most widely used is the Rectified Linear Unit (ReLU), also thanks to the computational efficiency of its derivative:

$$\text{ReLU}(x) = \begin{cases} 0 & \text{for } x < 0 \\ x & \text{for } x \geq 0 \end{cases}, \qquad \text{ReLU}'(x) = \begin{cases} 0 & \text{for } x < 0 \\ 1 & \text{for } x \geq 0 \end{cases}. \tag{A.8}$$

Different functions are usually used for output layers. For classification, it is desirable that the outputs of the NN represent probabilities, and therefore their values are between 0 and 1. In the case of binary classification a single output node with the Sigmoid function is typically used:

$$\text{Sigmoid}(x) = \frac{1}{1 + e^{-x}} \, . \tag{A.9}$$

The Sigmoid is generalised to the multi-class case by the SoftMax function:

$$\text{SoftMax}_i(\boldsymbol{x}) = \frac{e^{x_i}}{\sum_{c=1}^{n_c} e^{x_c}} \, , \tag{A.10}$$

where n_c is the number of target classes and corresponds to the number of output nodes. Following its definition, the SoftMax function satisfies the condition that the sum of its output vector components is 1: $\sum_{i=1}^{n_c} \text{SoftMax}_i(\boldsymbol{x}) = 1$.

List of Acronyms

AFP ATLAS Forward Proton
ALFA Absolute Luminosity for ATLAS
ALICE A Large Ion Collider Experiment
ATLAS A Toroidal LHC ApparatuS

BDT Boosted Decision Tree
BSM Beyond Standard Model

CCA Cut-and-Count Approach
CERN *Conseil Européen pour la Recherche Nucléaire*
CKM Cabibbo-Kobayashi-Maskawa
CL Confidence Level
CMS Compact Muon Solenoid
CP charge-parity
CR Control Region
CTP Central Trigger Processor

DAQ Data Acquisition system
DIPS Deep Impact Parameter Sets
DM Dark Matter

FTF Fast Track Finder

GNN graph neural network

HEP High Energy Physics
HL-LHC High-Luminosity LHC
HLT High-Level Trigger

IBL Insertable B-Layer
ID Inner Detector
IP interaction point

JER Jet Energy Resolution
JES Jet Energy Scale
JVF Jet Vertex Fraction
JVT Jet Vertex Tagger

L1 Level-1 Trigger
L1Calo L1 Calorimeter processor
L1Muon L1 Muon processor
L1Topo L1 Topological processor
LAr Liquid Argon
LEP Large Electron-Positron collider
LHC Large Hadron Collider
LHCb Large Hadron Collider beatuy
LLR log-likelihood ratio
LO Leading Order
LSP lightest supersymmetric particle
LSTM long short-term memory
LUCID LUminosity Cherenkov Integrating Detector

MBTS Minimum Bias Trigger Scintillators
MC Monte Carlo
MDT Monitored Drift Tubes
MET Missing transverse momentum
Micromegas micro-mesh gaseous structure
ML Machine Learning
MPI multi-parton interactions
MS Muon Spectrometer
MSSM Minimal Supersymmetric Standard Model

NLL Next-to-Leading Logarithmic
NLO Next-to-Leading Order
NN Neural Network
NP Nuisance Parameter
NR Normalization Region
NSW New Small Wheel

PDF Parton Density Function
pdf probability density function
PEB Partial Event Building
PFlow Particle Flow
PV Primary Vertex

QCD Quantum Chromodynamics
QED Quantum Electrodynamics
QFT Quantum Field Theory

RDO Raw Data Object
ReLU Rectified Linear Unit
RNN recurrent neural network
RNNIP Recurrent Neural Network Impact Parameter
ROC Receiver Operating Characteristic
RoI Region-of-Interest
RPC *R*-parity conserving
RPV *R*-parity violating

SCT Semiconductor Tracker
SGD stochastic gradient descent
SM Standard Model
SR Signal Region
sTGC small-strip TGC
SUSY Supersymmetry

TDAQ trigger and data acquisition
TGC Thin Gap Chambers
TLA Trigger-Level Analysis
TRT Transition Radiation Tracker

VR Validation Region

ZDC Zero Degree Calorimeter

Bibliography

[1] Georges Aad et al. “Fast b-tagging at the high-level trigger of the ATLAS experiment in LHC Run 3”. In: *JINST* 18.11 (2023), P11006. DOI: 10.1088/1748-0221/18/11/P11006. arXiv: 2306.09738 [hep-ex].

[2] Georges Aad et al. “The ATLAS trigger system for LHC Run 3 and trigger performance in 2022”. In: *JINST* 19.06 (2024), P06029. DOI: 10.1088/1748-0221/19/06/P06029. arXiv: 2401.06630 [hep-ex].

[3] Georges Aad et al. “Configuration, Performance, and Commissioning of the ATLAS b-jet Triggers for the 2022 and 2023 LHC data-taking periods”. In: (Jan. 2025). arXiv: 2501.11420 [hep-ex].

[4] ATLAS Collaboration. “A search for R-parity-violating supersymmetry in final states containing many jets in pp collisions at $\sqrt{s}$ = 13 TeV with the ATLAS detector”. In: *Journal of High Energy Physics* 2024.5 (May 2024). ISSN: 1029-8479. DOI: 10.1007/jhep05(2024)003.

[5] S. H. Neddermeyer and C. D. Anderson. “Note on the Nature of Cosmic Ray Particles”. In: *Phys. Rev.* 51 (1937), pp. 884–886. DOI: 10.1103/PhysRev.51.884.

[6] C. M. G. Lattes et al. “Process involving Charged Mesons”. In: *Nature* 159 (1947), pp. 694–697. DOI: 10.1038/159694a0.

[7] G. D. Rochester and C. C. Butler. “Evidence for the Existence of New Unstable Elementary Particles”. In: *Nature* 160 (1947), pp. 855–857. DOI: 10.1038/160855a0.

[8] W. D. Walker. “$\Lambda^0 - \theta^0$ Production in $\pi^- - P$ Collisions at 1 Bev”. In: *Phys. Rev.* 98 (5 June 1955), pp. 1407–1410. DOI: 10.1103/PhysRev.98.1407. URL: https://link.aps.org/doi/10.1103/PhysRev.98.1407.

[9] A. Pevsner et al. “Evidence for a Three Pion Resonance Near 550-MeV”. In: *Phys. Rev. Lett.* 7 (1961), pp. 421–423. DOI: 10.1103/PhysRevLett.7.421.

[10] J. J. Aubert et al. “Experimental Observation of a Heavy Particle J”. In: *Phys. Rev. Lett.* 33 (1974), pp. 1404–1406. DOI: 10.1103/PhysRevLett.33.1404.

[11] J. E. Augustin et al. “Discovery of a Narrow Resonance in e^+e^- Annihilation”. In: *Phys. Rev. Lett.* 33 (1974), pp. 1406–1408. DOI: 10.1103/PhysRevLett.33.1406.

[12] Martin L. Perl et al. “Evidence for Anomalous Lepton Production in e+ - e- Annihilation”. In: *Phys. Rev. Lett.* 35 (1975), pp. 1489–1492. DOI: 10.1103/PhysRevLett.35.1489.

[13] S. W. Herb et al. “Observation of a Dimuon Resonance at 9.5-GeV in 400-GeV Proton-Nucleus Collisions”. In: *Phys. Rev. Lett.* 39 (1977), pp. 252–255. DOI: 10.1103/PhysRevLett.39.252.

[14] G. Arnison et al. “Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540$ GeV”. In: *Phys. Lett. B* 122 (1983), pp. 103–116. DOI: 10.1016/0370-2693(83)91177-2.

[15] P. Bagnaia et al. “Evidence for $Z^0 \to e^+e^-$ at the CERN $\bar{p}p$ Collider”. In: *Phys. Lett. B* 129 (1983), pp. 130–140. DOI: 10.1016/0370-2693(83)90744-X.

[16] J. Konigsberg. “Observation of top quark production in $\bar{p}p$ collisions with the Collider Detector at Fermilab”. In: *30th Rencontres de Moriond: Electroweak Interactions and Unified Theories*. 1995, pp. 103–112.

[17] Peter W. Higgs. “Broken symmetries, massless particles and gauge fields”. In: *Phys. Lett.* 12 (1964), pp. 132–133. DOI: 10.1016/0031-9163(64)91136-9.

[18] F. Englert and R. Brout. “Broken Symmetry and the Mass of Gauge Vector Mesons”. In: *Phys. Rev. Lett.* 13 (9 Aug. 1964), pp. 321–323. DOI: 10.1103/PhysRevLett.13.321. URL: https://link.aps.org/doi/10.1103/PhysRevLett.13.321.

[19] E. Noether. “Invariante Variationsprobleme”. In: *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse* (1918), pp. 235–257.

[20] Paul A. M. Dirac. “The quantum theory of the electron”. In: *Proc. Roy. Soc. Lond. A* 117 (1928), pp. 610–624. DOI: 10.1098/rspa.1928.0023.

[21] C. N. Yang and R. L. Mills. “Conservation of Isotopic Spin and Isotopic Gauge Invariance”. In: *Phys. Rev.* 96 (1 Oct. 1954), pp. 191–195. DOI: 10.1103/PhysRev.96.191. URL: https://link.aps.org/doi/10.1103/PhysRev.96.191.

[22] Nicola Cabibbo. “Unitary Symmetry and Leptonic Decays”. In: *Phys. Rev. Lett.* 10 (1963), pp. 531–533. DOI: 10.1103/PhysRevLett.10.531.

[23] S. L. Glashow, J. Iliopoulos, and L. Maiani. “Weak Interactions with Lepton-Hadron Symmetry”. In: *Phys. Rev. D* 2 (1970), pp. 1285–1292. DOI: 10.1103/PhysRevD.2.1285.

[24] Makoto Kobayashi and Toshihide Maskawa. “CP Violation in the Renormalizable Theory of Weak Interaction”. In: *Prog. Theor. Phys.* 49 (1973), pp. 652–657. DOI: 10.1143/PTP.49.652.

[25] Lincoln Wolfenstein. “Parametrization of the Kobayashi-Maskawa Matrix”. In: *Phys. Rev. Lett.* 51 (1983), p. 1945. DOI: 10.1103/PhysRevLett.51.1945.

[26] J. H. Christenson et al. “Evidence for the 2π Decay of the K_2^0 Meson”. In: *Phys. Rev. Lett.* 13 (1964), pp. 138–140. DOI: 10.1103/PhysRevLett.13.138.

[27] Enrico Fermi. “Tentativo di una teoria dell’emissione dei raggi beta”. In: *Ric. Sci.* 4 (1933), pp. 491–495.

[28] D. H. Perkins. *Introduction to high energy physics*. 2000. ISBN: 978-0-521-62196-0. DOI: 10.1017/CBO9780511809040.

[29] X. Fan et al. “Measurement of the Electron Magnetic Moment”. In: *Phys. Rev. Lett.* 130.7 (2023), p. 071801. DOI: 10.1103/PhysRevLett.130.071801. arXiv: 2209.13084.

[30] S. Navas et al. “Review of particle physics”. In: *Phys. Rev. D* 110.3 (2024), p. 030001. DOI: 10.1103/PhysRevD.110.030001.

[31] Georges Aad et al. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Phys. Lett. B* 716 (2012), pp. 1–29. DOI: 10.1016/j.physletb.2012.08.020. arXiv: 1207.7214 [hep-ex].

[32] Serguei Chatrchyan et al. “Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC”. In: *Phys. Lett. B* 716 (2012), pp. 30–61. DOI: 10.1016/j.physletb.2012.08.021. arXiv: 1207.7235 [hep-ex].

[33] “Measurement of Higgs boson decay to a pair of muons in proton-proton collisions at $\sqrt{s} = 13\,\mathrm{TeV}$”. In: (2020).

[34] V. C. Rubin, N. Thonnard, and W. K. Ford Jr. “Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 /R = 4kpc/ to UGC 2885 /R = 122 kpc/”. In: *Astrophys. J.* 238 (1980), p. 471. DOI: 10.1086/158003.

[35] Markus Weber and Wim de Boer. “Determination of the Local Dark Matter Density in our Galaxy”. In: *Astron. Astrophys.* 509 (2010), A25. DOI: 10.1051/0004-6361/200913381. arXiv: 0910.4272 [astro-ph.CO].

[36] N. Aghanim et al. “Planck 2018 results. VI. Cosmological parameters”. In: *Astron. Astrophys.* 641 (2020). [Erratum: Astron.Astrophys. 652, C4 (2021)], A6. DOI: 10.1051/0004-6361/201833910. arXiv: 1807.06209 [astro-ph.CO].

[37] Y. Fukuda et al. “Evidence for oscillation of atmospheric neutrinos”. In: *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562. arXiv: hep-ex/9807003.

[38] Q. R. Ahmad et al. “Measurement of the rate of $\nu_e + d \to p + p + e^-$ interactions produced by ^{8}B solar neutrinos at the Sudbury Neutrino Observatory”. In: *Phys. Rev. Lett.* 87 (2001), p. 071301. DOI: 10.1103/PhysRevLett.87.071301. arXiv: nucl-ex/0106015.

[39] B. Aharmim et al. “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”. In: *Phys. Rev. C* 72 (2005), p. 055502. DOI: 10.1103/PhysRevC.72.055502. arXiv: nucl-ex/0502021.

[40] Samuel Brieden, Héctor Gil-Marín, and Licia Verde. “Model-agnostic interpretation of 10 billion years of cosmic evolution traced by BOSS and eBOSS data”. In: *JCAP* 08.08 (2022), p. 024. DOI: 10.1088/1475-7516/2022/08/024. arXiv: 2204.11868 [astro-ph.CO].

[41] Guido Altarelli. “The Higgs: so simple yet so unnatural”. In: *Phys. Scripta T* 158 (2013). Ed. by Tord Ekelöf, p. 014011. DOI: 10.1088/0031-8949/2013/T158/014011. arXiv: 1308.0545 [hep-ph].

[42] H. Georgi and S. L. Glashow. “Unity of All Elementary Particle Forces”. In: *Phys. Rev. Lett.* 32 (1974), pp. 438–441. DOI: 10.1103/PhysRevLett.32.438.

[43] Alexei A. Starobinsky. “A New Type of Isotropic Cosmological Models Without Singularity”. In: *Phys. Lett. B* 91 (1980). Ed. by I. M. Khalatnikov and V. P. Mineev, pp. 99–102. DOI: 10.1016/0370-2693(80)90670-X.

[44] A. D. Sakharov. “Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe”. In: *Pisma Zh. Eksp. Teor. Fiz.* 5 (1967), pp. 32–35.

[45] Laurent Canetti, Marco Drewes, and Mikhail Shaposhnikov. “Matter and Antimatter in the Universe”. In: *New J. Phys.* 14 (2012), p. 095012. DOI: 10.1088/1367-2630/14/9/095012. arXiv: 1204.4186 [hep-ph].

[46] Georges Aad et al. “The quest to discover supersymmetry at the ATLAS experiment”. In: (Mar. 2024). arXiv: 2403.02455 [hep-ex].

[47] Emanuele Bagnaschi et al. “Higgs Mass and Unnatural Supersymmetry”. In: *JHEP* 09 (2014), p. 092. DOI: 10.1007/JHEP09(2014)092. arXiv: 1407.4081 [hep-ph].

[48] Pran Nath and Richard L. Arnowitt. “Generalized Supergauge Symmetry as a New Framework for Unified Gauge Theories”. In: *Phys. Lett. B* 56 (1975), pp. 177–180. DOI: 10.1016/0370-2693(75)90297-X.

[49] Stephen P. Martin. “A Supersymmetry primer”. In: *Adv. Ser. Direct. High Energy Phys.* 18 (1998). Ed. by Gordon L. Kane, pp. 1–98. DOI: 10.1142/9789812839657_0001. arXiv: hep-ph/9709356.

[50] Gerard Jungman, Marc Kamionkowski, and Kim Griest. “Supersymmetric dark matter”. In: *Phys. Rept.* 267 (1996), pp. 195–373. DOI: 10.1016/0370-1573(95)00058-5. arXiv: hep-ph/9506380.

[51] Howard E. Haber. “Introductory low-energy supersymmetry”. In: *Theoretical Advanced Study Institute (TASI 92): From Black Holes and Strings to Particles*. Apr. 1993, pp. 589–686. arXiv: hep-ph/9306207.

[52] John R. Ellis et al. “Likelihood analysis of the CMSSM parameter space”. In: *Phys. Rev. D* 69 (2004), p. 095004. DOI: 10.1103/PhysRevD.69.095004. arXiv: hep-ph/0310356.

[53] G. F. Giudice and R. Rattazzi. “Theories with gauge mediated supersymmetry breaking”. In: *Phys. Rept.* 322 (1999), pp. 419–499. DOI: 10.1016/S0370-1573(99)00042-3. arXiv: hep-ph/9801271.

[54] K. Abe et al. “Search for proton decay via $p \to e^+\pi^0$ and $p \to \mu^+\pi^0$ in 0.31 megaton·years exposure of the Super-Kamiokande water Cherenkov detector”. In: *Phys. Rev. D* 95.1 (2017), p. 012004. DOI: 10.1103/PhysRevD.95.012004. arXiv: 1610.03597 [hep-ex].

[55] R. Barbier et al. “R-parity violating supersymmetry”. In: *Phys. Rept.* 420 (2005), pp. 1–202. DOI: 10.1016/j.physrep.2005.08.006. arXiv: hep-ph/0406039.

[56] Wilfried Buchmuller et al. “Gravitino Dark Matter in R-Parity Breaking Vacua”. In: *JHEP* 03 (2007), p. 037. DOI: 10.1088/1126-6708/2007/03/037. arXiv: hep-ph/0702184.

[57] S. Ask. “A Review of the supersymmetry searches at LEP”. In: *38th Rencontres de Moriond on Electroweak Interactions and Unified Theories*. May 2003. arXiv: hep-ex/0305007.

[58] Philippe Gris. “SUSY Searches at the Tevatron”. In: *46th Rencontres de Moriond on QCD and High Energy Interactions*. June 2011, pp. 41–44. arXiv: 1106.1032 [hep-ex].

[59] A. Djouadi et al. “The Minimal supersymmetric standard model: Group summary report”. In: *GDR (Groupement De Recherche) - Supersymetrie*. Dec. 1998. arXiv: hep-ph/9901246.

[60] Carola F. Berger et al. “Supersymmetry Without Prejudice”. In: *JHEP* 02 (2009), p. 023. DOI: 10.1088/1126-6708/2009/02/023. arXiv: 0812.0980 [hep-ph].

[61] Jim Alexander et al. “Dark Sectors 2016 Workshop: Community Report”. In: Aug. 2016. arXiv: 1608.08632 [hep-ph].

[62] Takehiko Asaka and Mikhail Shaposhnikov. “The νMSM, dark matter and baryon asymmetry of the universe”. In: *Phys. Lett. B* 620 (2005), pp. 17–26. DOI: 10.1016/j.physletb.2005.06.020. arXiv: hep-ph/0505013.

[63] Georges Aad et al. “ATLAS Run 2 searches for electroweak production of supersymmetric particles interpreted within the pMSSM”. In: *JHEP* 05 (2024), p. 106. DOI: 10.1007/JHEP05(2024)106. arXiv: 2402.01392 [hep-ex].

[64] Bob Holdom. “Two U(1)’s and Epsilon Charge Shifts”. In: *Phys. Lett. B* 166 (1986), pp. 196–198. DOI: 10.1016/0370-2693(86)91377-8.

[65] David Curtin et al. “Exotic decays of the 125 GeV Higgs boson”. In: *Phys. Rev. D* 90.7 (2014), p. 075004. DOI: 10.1103/PhysRevD.90.075004. arXiv: 1312.4992 [hep-ph].

[66] Joerg Jaeckel and Andreas Ringwald. “The Low-Energy Frontier of Particle Physics”. In: *Ann. Rev. Nucl. Part. Sci.* 60 (2010), pp. 405–437. DOI: 10.1146/annurev.nucl.012809.104433. arXiv: 1002.0329 [hep-ph].

[67] R. D. Peccei and Helen R. Quinn. “CP Conservation in the Presence of Instantons”. In: *Phys. Rev. Lett.* 38 (1977), pp. 1440–1443. DOI: 10.1103/PhysRevLett.38.1440.

[68] E. Lopienska. *The CERN accelerator complex, layout in 2022*. CERN-GRAPHICS-2022-001. URL: https://cds.cern.ch/record/2800984.

[69] O. et al. Aberle. *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2020. DOI: 10.23731/CYRM-2020-0010. URL: https://cds.cern.ch/record/2749422.

[70] W. Herr and B. Muratori. “Concept of luminosity”. In: *CERN Accelerator School and DESY Zeuthen: Accelerator Physics*. Sept. 2003, pp. 361–377.

[71] Jan Kretzschmar. “Standard Model physics at the LHC”. In: *From My Vast Repertoire ...: Guido Altarelli’s Legacy*. Ed. by Aharon Levy, Stefano Forte, and Giovanni Ridolfi. 2019, pp. 153–171. DOI: 10.1142/9789813238053_0009. arXiv: 1803.10800 [hep-ex].

[72] ATLAS Collaboration. *Public ATLAS Luminosity Results for Run-3 of the LHC.* Public Twiki. URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun3.

[73] Tung-Mow Yan and Sidney D. Drell. “The Parton Model and its Applications”. In: *Int. J. Mod. Phys. A* 29 (2014), p. 0071. DOI: 10.1142/S0217751X14300713. arXiv: 1409.0051 [hep-ph].

[74] AC Team. *Diagram showing the cross-section of an LHC dipole magnet with cold mass and vacuum chamber.* CERN-DI-9906025. URL: https://cds.cern.ch/record/40524.

[75] D. de Florian et al. “Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector”. In: 2/2017 (Oct. 2016). DOI: 10.23731/CYRM-2017-002. arXiv: 1610.07922 [hep-ph].

[76] M. Aaboud et al. “Evidence for the $\mathrm{H} \to \mathrm{b\bar{b}}$ decay with the ATLAS detector”. In: *JHEP* 12 (2017), p. 024. DOI: 10.1007/JHEP12(2017)024. arXiv: 1708.03299 [hep-ex].

[77] Georges Aad et al. “Measurements of Higgs boson production cross-sections in the $\mathrm{H} \to \tau^+\tau^-$ decay channel in pp collisions at $\sqrt{s}$ = 13 TeV with the ATLAS detector”. In: *JHEP* 08 (2022), p. 175. DOI: 10.1007/JHEP08(2022)175. arXiv: 2201.08269 [hep-ex].

[78] M. Aaboud et al. “Observation of Higgs boson production in association with a top quark pair at the LHC with the ATLAS detector”. In: *Phys. Lett. B* 784 (2018), pp. 173–191. DOI: 10.1016/j.physletb.2018.07.035. arXiv: 1806.00425 [hep-ex].

[79] Georges Aad et al. “Evidence for Electroweak Production of $W^{\pm}W^{\pm}jj$ in pp Collisions at $\sqrt{s} = 8$ TeV with the ATLAS Detector”. In: *Phys. Rev. Lett.* 113.14 (2014), p. 141803. DOI: 10.1103/PhysRevLett.113.141803. arXiv: 1405.6241 [hep-ex].

[80] Morad Aaboud et al. “Evidence for light-by-light scattering in heavy-ion collisions with the ATLAS detector at the LHC”. In: *Nature Phys.* 13.9 (2017), pp. 852–858. DOI: 10.1038/nphys4208. arXiv: 1702.01625 [hep-ex].

[81] Georges Aad et al. “Observation of quantum entanglement with top quarks at the ATLAS detector”. In: *Nature* 633.8030 (2024), pp. 542–547. DOI: 10.1038/s41586-024-07824-z. arXiv: 2311.07288 [hep-ex].

[82] Georges Aad et al. “Exploration at the high-energy frontier: ATLAS Run 2 searches investigating the exotic jungle beyond the Standard Model”. In: (Mar. 2024). arXiv: 2403.09292 [hep-ex].

[83] G. Aad et al. “The ATLAS Experiment at the CERN Large Hadron Collider”. In: *JINST* 3 (2008), S08003. DOI: 10.1088/1748-0221/3/08/S08003.

[84] Georges Aad et al. “The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run 3”. In: *JINST* 19.05 (2024), P05063. DOI: 10.1088/1748-0221/19/05/P05063. arXiv: 2305.16623 [physics.ins-det].

[85] G. Aad et al. “ATLAS pixel detector electronics and sensors”. In: *JINST* 3 (2008), P07007. DOI: 10.1088/1748-0221/3/07/P07007.

[86] M. Capeans et al. “ATLAS Insertable B-Layer Technical Design Report”. In: (Sept. 2010).

[87] Georges Aad et al. “Operation and performance of the ATLAS semiconductor tracker”. In: *JINST* 9 (2014), P08009. DOI: `10.1088/1748-0221/9/08/P08009`. arXiv: `1404.7473 [hep-ex]`.

[88] Morad Aaboud et al. “Performance of the ATLAS Transition Radiation Tracker in Run 1 of the LHC: tracker properties”. In: *JINST* 12.05 (2017), P05002. DOI: `10.1088/1748-0221/12/05/P05002`. arXiv: `1702.06473 [hep-ex]`.

[89] *ATLAS liquid-argon calorimeter: Technical Design Report*. Technical design report. ATLAS. Geneva: CERN, 1996. DOI: `10.17181/CERN.FWRW.FOOQ`. URL: `https://cds.cern.ch/record/331061`.

[90] *ATLAS tile calorimeter: Technical Design Report*. Technical design report. ATLAS. Geneva: CERN, 1996. DOI: `10.17181/CERN.JRBJ.7028`. URL: `https://cds.cern.ch/record/331062`.

[91] Regina Kwee. “Minimum Bias Triggers at ATLAS, LHC”. In: *34th International Conference on High Energy Physics*. Dec. 2008. arXiv: `0812.0613 [hep-ex]`.

[92] *ATLAS muon spectrometer: Technical Design Report*. Technical design report. ATLAS. Geneva: CERN, 1997. URL: `https://cds.cern.ch/record/331068`.

[93] A Sbrizzi. *A Cherenkov Detector for Monitoring ATLAS Luminosity*. Tech. rep. 1. Geneva: CERN, 2011. DOI: `10.1016/j.nuclphysbps.2011.04.017`. URL: `https://cds.cern.ch/record/1291970`.

[94] “Zero degree calorimeters for ATLAS”. In: (Jan. 2007).

[95] L. Adamczyk et al. “Technical Design Report for the ATLAS Forward Proton Detector”. In: (May 2015).

[96] S. Franz and P. Barrillon. “ATLAS ALFA-measuring absolute luminosity with scintillating fibres”. In: *Nucl. Instrum. Meth. A* 610 (2009). Ed. by Philippe Bourgeois et al., pp. 35–40. DOI: `10.1016/j.nima.2009.05.148`.

[97] Georges Aad et al. “Performance of the ATLAS Trigger System in 2010”. In: *Eur. Phys. J. C* 72 (2012), p. 1849. DOI: `10.1140/epjc/s10052-011-1849-1`. arXiv: `1110.1530 [hep-ex]`.

[98] Morad Aaboud et al. “Performance of the ATLAS Trigger System in 2015”. In: *Eur. Phys. J. C* 77.5 (2017), p. 317. DOI: `10.1140/epjc/s10052-017-4852-3`. arXiv: `1611.09661 [hep-ex]`.

[99] R. Achenbach et al. “The ATLAS level-1 calorimeter trigger”. In: *JINST* 3 (2008), P03001. DOI: `10.1088/1748-0221/3/03/P03001`.

[100] Georges Aad et al. “Performance of the ATLAS muon triggers in Run 2”. In: *JINST* 15.09 (2020), P09015. DOI: `10.1088/1748-0221/15/09/p09015`. arXiv: `2004.13447`.

[101] Georges Aad et al. “Performance of the ATLAS Level-1 topological trigger in Run 2”. In: *Eur. Phys. J. C* 82.1 (2022), p. 7. DOI: `10.1140/epjc/s10052-021-09807-0`. arXiv: `2105.01416 [hep-ex]`.

[102] ATLAS Collaboration. *Athena*. Version 22.0.1. Apr. 2019. DOI: `10.5281/zenodo.2641997`. URL: `https://doi.org/10.5281/zenodo.2641997`.

[103] M. Aaboud et al. “Search for low-mass dijet resonances using trigger-level jets with the ATLAS detector in pp collisions at $\sqrt{s} = 13$ TeV”. In: *Phys. Rev. Lett.* 121.8 (2018), p. 081801. DOI: `10.1103/PhysRevLett.121.081801`. arXiv: `1804.03496 [hep-ex]`.

[104] Georges Aad et al. “Software and computing for Run 3 of the ATLAS experiment at the LHC”. In: *Eur. Phys. J. C* 85.3 (2025), p. 234. DOI: `10.1140/epjc/s10052-024-13701-w`. arXiv: `2404.06335 [hep-ex]`.

[105] Christian Bierlich et al. “A comprehensive guide to the physics and usage of PYTHIA 8.3”. In: *SciPost Phys. Codeb.* 2022 (2022), p. 8. DOI: 10.21468/SciPostPhysCodeb.8. arXiv: 2203.11601 [hep-ph].

[106] Vladinir Chekellan. “Proton structure and PDFs at HERA”. In: *PoS* Photon2013 (2013), p. 004.

[107] Guido Altarelli and G. Parisi. “Asymptotic Freedom in Parton Language”. In: *Nucl. Phys. B* 126 (1977), pp. 298–318. DOI: 10.1016/0550-3213(77)90384-4.

[108] John C. Collins. “Sudakov form-factors”. In: *Adv. Ser. Direct. High Energy Phys.* 5 (1989), pp. 573–614. DOI: 10.1142/9789814503266_0006. arXiv: hep-ph/0312336.

[109] Javira Altmann and Peter Skands. “String junctions revisited”. In: *JHEP* 07 (2024), p. 238. DOI: 10.1007/JHEP07(2024)238. arXiv: 2404.12040 [hep-ph].

[110] S. Agostinelli et al. “GEANT4–a simulation toolkit”. In: *Nucl. Instrum. Meth. A* 506 (2003), pp. 250–303. DOI: 10.1016/S0168-9002(03)01368-8.

[111] Georges Aad et al. “AtlFast3: The Next Generation of Fast Simulation in ATLAS”. In: *Comput. Softw. Big Sci.* 6.1 (2022), p. 7. DOI: 10.1007/s41781-021-00079-7. arXiv: 2109.02551 [hep-ex].

[112] Georges Aad et al. “Emulating the impact of additional proton–proton interactions in the ATLAS simulation by presampling sets of inelastic Monte Carlo events”. In: *Comput. Softw. Big Sci.* 6.1 (2022), p. 3. DOI: 10.1007/s41781-021-00062-2. arXiv: 2102.09495 [hep-ex].

[113] M. Aaboud et al. “Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2”. In: *Eur. Phys. J. C* 77.10 (2017), p. 673. DOI: 10.1140/epjc/s10052-017-5225-7. arXiv: 1704.07983 [hep-ex].

[114] R. Fruhwirth. “Application of Kalman filtering to track and vertex fitting”. In: *Nucl. Instrum. Meth. A* 262 (1987), pp. 444–450. DOI: 10.1016/0168-9002(87)90887-4.

[115] *Training and validation of the ATLAS pixel clustering neural networks*. Tech. rep. Geneva: CERN, 2018. URL: https://cds.cern.ch/record/2309474.

[116] Georges Aad et al. “Software Performance of the ATLAS Track Reconstruction for LHC Run 3”. In: *Comput. Softw. Big Sci.* 8.1 (2024), p. 9. DOI: 10.1007/s41781-023-00111-y. arXiv: 2308.09471 [hep-ex].

[117] Morad Aaboud et al. “Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC”. In: *Eur. Phys. J. C* 77.5 (2017), p. 332. DOI: 10.1140/epjc/s10052-017-4887-5. arXiv: 1611.10235 [physics.ins-det].

[118] *Development of ATLAS Primary Vertex Reconstruction for LHC Run 3*. Tech. rep. Geneva: CERN, 2019. URL: http://cds.cern.ch/record/2670380.

[119] W Lampl et al. *Calorimeter Clustering Algorithms: Description and Performance*. Tech. rep. Geneva: CERN, 2008. URL: https://cds.cern.ch/record/1099735.

[120] Georges Aad et al. “Electron reconstruction and identification efficiency measurements with the ATLAS detector using the 2011 LHC proton-proton collision data”. In: *Eur. Phys. J. C* 74.7 (2014), p. 2941. DOI: 10.1140/epjc/s10052-014-2941-0. arXiv: 1404.2240 [hep-ex].

[121] Morad Aaboud et al. “Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run-1 data”. In: *Eur. Phys. J. C* 76.12 (2016), p. 666. DOI: 10.1140/epjc/s10052-016-4507-9. arXiv: 1606.01813 [hep-ex].

[122] Georges Aad et al. “Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1”. In: *Eur. Phys. J. C* 77 (2017), p. 490. DOI: 10.1140/epjc/s10052-017-5004-5. arXiv: 1603.02934 [hep-ex].

[123] Gavin P. Salam. “Towards Jetography”. In: *Eur. Phys. J. C* 67 (2010), pp. 637–686. DOI: `10.1140/epjc/s10052-010-1314-6`. arXiv: `0906.1833 [hep-ph]`.

[124] Georges Aad et al. “Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{s} = 7$ TeV”. In: *Eur. Phys. J. C* 73.3 (2013), p. 2304. DOI: `10.1140/epjc/s10052-013-2304-2`. arXiv: `1112.6426 [hep-ex]`.

[125] Morad Aaboud et al. “Jet reconstruction and performance using particle flow with the ATLAS Detector”. In: *Eur. Phys. J. C* 77.7 (2017), p. 466. DOI: `10.1140/epjc/s10052-017-5031-2`. arXiv: `1703.10485 [hep-ex]`.

[126] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. “The anti-k_t jet clustering algorithm”. In: *JHEP* 04 (2008), p. 063. DOI: `10.1088/1126-6708/2008/04/063`. arXiv: `0802.1189 [hep-ph]`.

[127] John E. Huth et al. “Toward a standardization of jet definitions”. In: *1990 DPF Summer Study on High-energy Physics: Research Directions for the Decade (Snowmass 90)*. Dec. 1990, pp. 0134–136.

[128] Stephen D. Ellis and Davison E. Soper. “Successive combination jet algorithm for hadron collisions”. In: *Phys. Rev. D* 48 (1993), pp. 3160–3166. DOI: `10.1103/PhysRevD.48.3160`. arXiv: `hep-ph/9305266`.

[129] Yuri L. Dokshitzer et al. “Better jet clustering algorithms”. In: *JHEP* 08 (1997), p. 001. DOI: `10.1088/1126-6708/1997/08/001`. arXiv: `hep-ph/9707323`.

[130] Georges Aad et al. “Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *Eur. Phys. J. C* 81.8 (2021), p. 689. DOI: `10.1140/epjc/s10052-021-09402-3`. arXiv: `2007.02645 [hep-ex]`.

[131] Morad Aaboud et al. “Determination of jet calibration and energy resolution in proton-proton collisions at $\sqrt{s}$ = 8 TeV using the ATLAS detector”. In: *Eur. Phys. J. C* 80.12 (2020), p. 1104. DOI: `10.1140/epjc/s10052-020-08477-8`. arXiv: `1910.04482 [hep-ex]`.

[132] Georges Aad et al. “New techniques for jet calibration with the ATLAS detector”. In: *Eur. Phys. J. C* 83.8 (2023), p. 761. DOI: `10.1140/epjc/s10052-023-11837-9`. arXiv: `2303.17312 [hep-ex]`.

[133] *Pile-up subtraction and suppression for jets in ATLAS*. Tech. rep. Geneva: CERN, 2013. URL: `https://cds.cern.ch/record/1570994`.

[134] Georges Aad et al. “Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector”. In: *Eur. Phys. J. C* 76.11 (2016), p. 581. DOI: `10.1140/epjc/s10052-016-4395-z`. arXiv: `1510.03823 [hep-ex]`.

[135] Georges Aad et al. “Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data”. In: *JINST* 14.12 (2019), P12006. DOI: `10.1088/1748-0221/14/12/P12006`. arXiv: `1908.00005 [hep-ex]`.

[136] Georges Aad et al. “Electron and photon efficiencies in LHC Run 2 with the ATLAS experiment”. In: *JHEP* 05 (2024), p. 162. DOI: `10.1007/JHEP05(2024)162`. arXiv: `2308.13362 [hep-ex]`.

[137] Morad Aaboud et al. “Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton-proton collision data at $\sqrt{s}$ = 13 TeV”. In: *Eur. Phys. J. C* 79.8 (2019), p. 639. DOI: `10.1140/epjc/s10052-019-7140-6`. arXiv: `1902.04655 [physics.ins-det]`.

[138] Georges Aad et al. “Tools for estimating fake/non-prompt lepton backgrounds with the ATLAS detector at the LHC”. In: *JINST* 18.11 (2023), T11004. DOI: `10.1088/1748-0221/18/11/T11004`. arXiv: `2211.16178 [hep-ex]`.

[139] Georges Aad et al. “Electron and photon energy calibration with the ATLAS detector using LHC Run 2 data”. In: *JINST* 19.02 (2024), P02009. DOI: 10.1088/1748-0221/19/02/P02009. arXiv: 2309.05471 [hep-ex].

[140] Georges Aad et al. “Muon reconstruction performance of the ATLAS detector in proton-proton collision data at $\sqrt{s}$ =13 TeV”. In: *Eur. Phys. J. C* 76.5 (2016), p. 292. DOI: 10.1140/epjc/s10052-016-4120-y. arXiv: 1603.05598 [hep-ex].

[141] Georges Aad et al. “Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV”. In: *Eur. Phys. J. C* 81.7 (2021), p. 578. DOI: 10.1140/epjc/s10052-021-09233-2. arXiv: 2012.00578 [hep-ex].

[142] Georges Aad et al. “Studies of the muon momentum calibration and performance of the ATLAS detector with pp collisions at $\sqrt{s} = 13$ TeV”. In: *Eur. Phys. J. C* 83.8 (2023), p. 686. DOI: 10.1140/epjc/s10052-023-11584-x. arXiv: 2212.07338 [hep-ex].

[143] Georges Aad et al. “Reconstruction of hadronic decay products of tau leptons with the ATLAS experiment”. In: *Eur. Phys. J. C* 76.5 (2016), p. 295. DOI: 10.1140/epjc/s10052-016-4110-0. arXiv: 1512.05955 [hep-ex].

[144] *Reconstruction, Identification, and Calibration of hadronically decaying tau leptons with the ATLAS detector for the LHC Run 3 and reprocessed Run 2 data*. Tech. rep. Geneva: CERN, 2022. URL: http://cds.cern.ch/record/2827111.

[145] Georges Aad et al. “Identification and energy calibration of hadronically decaying tau leptons with the ATLAS experiment in pp collisions at $\sqrt{s}$=8 TeV”. In: *Eur. Phys. J. C* 75.7 (2015), p. 303. DOI: 10.1140/epjc/s10052-015-3500-z. arXiv: 1412.7086 [hep-ex].

[146] Morad Aaboud et al. “Performance of missing transverse momentum reconstruction with the ATLAS detector using proton-proton collisions at $\sqrt{s}$ = 13 TeV”. In: *Eur. Phys. J. C* 78.11 (2018), p. 903. DOI: 10.1140/epjc/s10052-018-6288-9. arXiv: 1802.08168 [hep-ex].

[147] Georges Aad et al. “The performance of missing transverse momentum reconstruction and its significance with the ATLAS detector using 140 fb^{-1} of $\sqrt{s} = 13$ TeV pp collisions”. In: (Feb. 2024). arXiv: 2402.05858 [hep-ex].

[148] G. Aad et al. “Performance of algorithms that reconstruct missing transverse momentum in $\sqrt{s} = 8$ TeV proton-proton collisions in the ATLAS detector”. In: *Eur. Phys. J. C* 77.4 (2017), p. 241. DOI: 10.1140/epjc/s10052-017-4780-2. arXiv: 1609.09324 [hep-ex].

[149] Matthias Neubert. “Heavy quark effective theory”. In: *Subnucl. Ser.* 34 (1997). Ed. by A. Zichichi, pp. 98–165. arXiv: hep-ph/9610266.

[150] Arnulf Quadt. “Top quark physics at hadron colliders”. In: *Eur. Phys. J. C* 48 (2006), pp. 835–1000. DOI: 10.1140/epjc/s2006-02631-6.

[151] R. Aaij et al. “Precise determination of the B_s^0–$\overline{B}_s^0$ oscillation frequency”. In: *Nature Phys.* 18.1 (2022), pp. 1–5. DOI: 10.1038/s41567-021-01394-x. arXiv: 2104.04421 [hep-ex].

[152] Wolfgang Altmannshofer and David M. Straub. “New Physics in $B \to K^*\mu\mu$?” In: *Eur. Phys. J. C* 73 (2013), p. 2646. DOI: 10.1140/epjc/s10052-013-2646-9. arXiv: 1308.1501 [hep-ph].

[153] Marco Ciuchini et al. “Constraints on lepton universality violation from rare B decays”. In: *Phys. Rev. D* 107.5 (2023), p. 055036. DOI: 10.1103/PhysRevD.107.055036. arXiv: 2212.10516 [hep-ph].

[154] Syuhei Iguro, Teppei Kitahara, and Ryoutaro Watanabe. “Global fit to $b \to c\tau\nu$ anomalies as of Spring 2024”. In: *Phys. Rev. D* 110.7 (2024), p. 075005. DOI: 10.1103/PhysRevD.110.075005. arXiv: 2405.06062 [hep-ph].

[155] Aaron Taylor. "A Selection of Three ATLAS B-Physics Results: A Search Beyond the Standard Model, A Precision Measurement, and the Discovery of a New Heavy Meson". In: *Meeting of the APS Division of Particles and Fields*. Oct. 2015. arXiv: `1510.01212 [hep-ex]`.

[156] Georges Aad et al. "Climbing to the Top of the ATLAS 13 TeV data". In: (Apr. 2024). arXiv: `2404.10674 [hep-ex]`.

[157] Georges Aad et al. "Characterising the Higgs boson with ATLAS data from Run 2 of the LHC". In: *Phys. Rept.* 11 (2024), p. 001. DOI: `10.1016/j.physrep.2024.11.001`. arXiv: `2404.05498 [hep-ex]`.

[158] Georges Aad et al. "A search for triple Higgs boson production in the $6\mathrm{b}$ final state using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: (Nov. 2024). arXiv: `2411.02040 [hep-ex]`.

[159] Dario Buttazzo et al. "Investigating the near-criticality of the Higgs boson". In: *JHEP* 12 (2013), p. 089. DOI: `10.1007/JHEP12(2013)089`. arXiv: `1307.3536 [hep-ph]`.

[160] Valerio De Luca, Alex Kehagias, and Antonio Riotto. "On the cosmological stability of the Higgs instability". In: *JCAP* 09 (2022), p. 055. DOI: `10.1088/1475-7516/2022/09/055`. arXiv: `2205.10240 [hep-ph]`.

[161] Georges Aad et al. "Search for R-parity violating supersymmetric decays of the top squark to a b-jet and a lepton in s=13 TeV pp collisions with the ATLAS detector". In: *Phys. Rev. D* 110.9 (2024), p. 092004. DOI: `10.1103/PhysRevD.110.092004`. arXiv: `2406.18367 [hep-ex]`.

[162] Georges Aad et al. "Search for new phenomena in final states with b-jets and missing transverse momentum in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector". In: *JHEP* 05 (2021), p. 093. DOI: `10.1007/JHEP05(2021)093`. arXiv: `2101.12527 [hep-ex]`.

[163] Georges Aad et al. "Search for direct production of electroweakinos in final states with one lepton, missing transverse momentum and a Higgs boson decaying into two b-jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: *Eur. Phys. J. C* 80.8 (2020), p. 691. DOI: `10.1140/epjc/s10052-020-8050-3`. arXiv: `1909.09226 [hep-ex]`.

[164] Yuri Gershtein, Simon Knapen, and Diego Redigolo. "Probing naturally light singlets with a displaced vertex trigger". In: *Phys. Lett. B* 823 (2021), p. 136758. DOI: `10.1016/j.physletb.2021.136758`. arXiv: `2012.07864 [hep-ph]`.

[165] J. A. Aguilar-Saavedra et al. "Handbook of vectorlike quarks: Mixing and single production". In: *Phys. Rev. D* 88.9 (2013), p. 094010. DOI: `10.1103/PhysRevD.88.094010`. arXiv: `1306.0572 [hep-ph]`.

[166] "Graph Neural Network Jet Flavour Tagging with the ATLAS Detector". In: (2022).

[167] Georges Aad et al. "Performance of b-Jet Identification in the ATLAS Experiment". In: *JINST* 11.04 (2016), P04008. DOI: `10.1088/1748-0221/11/04/P04008`. arXiv: `1512.01094 [hep-ex]`.

[168] Georges Aad et al. "ATLAS b-jet identification performance and efficiency measurement with $\mathrm{t}\bar{\mathrm{t}}$ events in pp collisions at $\sqrt{s} = 13$ TeV". In: *Eur. Phys. J. C* 79.11 (2019), p. 970. DOI: `10.1140/epjc/s10052-019-7450-8`. arXiv: `1907.05120 [hep-ex]`.

[169] Georges Aad et al. "ATLAS flavour-tagging algorithms for the LHC Run 2 pp collision dataset". In: *Eur. Phys. J. C* 83.7 (2023), p. 681. DOI: `10.1140/epjc/s10052-023-11699-1`. arXiv: `2211.16345 [physics.data-an]`.

[170] Andrea Sciandra. "Development of a new Soft Muon Tagger for the identification of b-jets in ATLAS". In: *PoS* EPS-HEP2017 (2017). Ed. by Paolo Checchia et al., p. 768. DOI: `10.22323/1.314.0768`.

[171] "First look at the JetProb b-tagging algorithm in the 900 GeV collision data with the ATLAS detector". In: (June 2010).

[172] "Optimisation and performance studies of the ATLAS b-tagging algorithms for the 2017-18 LHC run". In: (July 2017).

[173] "Expected performance of the ATLAS b-tagging algorithms in Run-2". In: (July 2015).

[174] "Identification of Jets Containing b-Hadrons with Recurrent Neural Networks at the ATLAS Experiment". In: (Mar. 2017).

[175] Sepp Hochreiter and Jürgen Schmidhuber. "Long Short-Term Memory". In: *Neural Computation* 9.8 (Nov. 1997), pp. 1735–1780. ISSN: 0899-7667. DOI: 10.1162/neco.1997.9.8.1735. eprint: https://direct.mit.edu/neco/article-pdf/9/8/1735/813796/neco.1997.9.8.1735.pdf. URL: https://doi.org/10.1162/neco.1997.9.8.1735.

[176] "Deep Sets based Neural Networks for Impact Parameter Flavour Tagging in ATLAS". In: (2020).

[177] Manzil Zaheer et al. *Deep Sets*. 2018. arXiv: 1703.06114 [cs.LG]. URL: https://arxiv.org/abs/1703.06114.

[178] "Secondary vertex finding for jet flavour identification with the ATLAS detector". In: (June 2017).

[179] "Topological b-hadron decay reconstruction and identification of b-jets with the JetFitter package in the ATLAS experiment at the LHC". In: (2018).

[180] Yoav Freund and Robert E. Schapire. "A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting". In: *J. Comput. Syst. Sci.* 55.1 (1997), pp. 119–139. DOI: 10.1006/jcss.1997.1504.

[181] Peter W. Battaglia et al. "Relational inductive biases, deep learning, and graph networks". In: (June 2018). arXiv: 1806.01261 [cs.LG].

[182] *Flavour Tagging with Graph Neural Network*. 2023. URL: https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/FTAG-2023-01.

[183] Ashish Vaswani et al. "Attention Is All You Need". In: *31st International Conference on Neural Information Processing Systems*. June 2017. arXiv: 1706.03762 [cs.CL].

[184] Georges Aad et al. "Measurement of the c-jet mistagging efficiency in $t\bar{t}$ events using pp collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS detector". In: *Eur. Phys. J. C* 82.1 (2022), p. 95. DOI: 10.1140/epjc/s10052-021-09843-w. arXiv: 2109.10627 [hep-ex].

[185] Georges Aad et al. "Calibration of the light-flavour jet mistagging efficiency of the b-tagging algorithms with Z+jets events using 139 fb^{-1} of ATLAS proton–proton collision data at $\sqrt{s} = 13$ TeV". In: *Eur. Phys. J. C* 83.8 (2023), p. 728. DOI: 10.1140/epjc/s10052-023-11736-z. arXiv: 2301.06319 [hep-ex].

[186] *The ATLAS Inner Detector Trigger performance in pp collisions at 900 GeV and 13.6 TeV for LHC Run 3 operation during 2022*. Tech. rep. Geneva: CERN, 2023. URL: http://cds.cern.ch/record/2881679.

[187] Georges Aad et al. "The ATLAS inner detector trigger performance in pp collisions at 13 TeV during LHC Run 2". In: *Eur. Phys. J. C* 82.3 (2022), p. 206. DOI: 10.1140/epjc/s10052-021-09920-0. arXiv: 2107.02485 [hep-ex].

[188] Pablo Yepes. "A fast track pattern recognition". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 380.3 (1996), pp. 582–585. ISSN: 0168-9002. DOI: https://doi.org/10.1016/0168-9002(96)00726-7. URL: https://www.sciencedirect.com/science/article/pii/0168900296007267.

[189] *Public b-Jet Trigger Plots for Collision Data*. URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/BJetTriggerPublicResults.

[190] *Trigger monitoring and rate predictions using Enhanced Bias data from the ATLAS Detector at the LHC*. Tech. rep. Geneva: CERN, 2016. URL: https://cds.cern.ch/record/2223498.

[191] Georges Aad et al. "Search for nonresonant pair production of Higgs bosons in the bbbb final state in pp collisions at $\sqrt{s}$=13 TeV with the ATLAS detector". In: *Phys. Rev. D* 108.5 (2023), p. 052003. DOI: 10.1103/PhysRevD.108.052003. arXiv: 2301.03212 [hep-ex].

[192] Georges Aad et al. "Search for the nonresonant production of Higgs boson pairs via gluon fusion and vector-boson fusion in the $bb\tau\tau$ final state in proton-proton collisions at $\sqrt{s}$=13TeV with the ATLAS detector". In: *Phys. Rev. D* 110.3 (2024), p. 032012. DOI: 10.1103/PhysRevD.110.032012. arXiv: 2404.12660 [hep-ex].

[193] Jackson Barr et al. "Umami: A Python toolkit for jet flavour tagging". In: *Journal of Open Source Software* 9.102 (2024), p. 5833. DOI: 10.21105/joss.05833. URL: https://doi.org/10.21105/joss.05833.

[194] TensorFlow Developers. *TensorFlow*. Version v2.19.0-rc0. Feb. 2025. DOI: 10.5281/zenodo.14919893. URL: https://doi.org/10.5281/zenodo.14919893.

[195] ATLAS Collaboration. *Fast* b*-tagging Neural Network Configuration*. https://atlas-groupdata.web.cern.ch/atlas-groupdata/BTagging/20220211trig/fastDips/antikt4empflow/network.json. (Visited on 02/25/2025).

[196] Daniel Hay Guest et al. *lwtnn/lwtnn: Version 2.12.1*. Version v2.12.1. July 2021. DOI: 10.5281/zenodo.5082190. URL: https://doi.org/10.5281/zenodo.5082190.

[197] CDF Collaboration. "First Search for Multijet Resonances in $\sqrt{s} = 1.96\ \text{TeV}\ p\overline{p}$ Collisions". In: *Phys. Rev. Lett.* 107 (4 July 2011), p. 042001. DOI: 10.1103/PhysRevLett.107.042001.

[198] CMS Collaboration. "Search for new phenomena in events with high jet multiplicity and low missing transverse momentum in proton–proton collisions at $\sqrt{s} = 8$ TeV". In: *Physics Letters B* 770 (2017), pp. 257–267. ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2017.01.073.

[199] CMS Collaboration. "Searches for light- and heavy-flavour three-jet resonances in pp collisions at $\sqrt{s} = 8$TeV". In: *Physics Letters B* 730 (2014), pp. 193–214. ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2014.01.049.

[200] CMS Collaboration. "Search for Pair-Produced Resonances Each Decaying into at Least Four Quarks in Proton-Proton Collisions at $\sqrt{s} = 13\ \text{TeV}$". In: *Phys. Rev. Lett.* 121 (14 Oct. 2018), p. 141802. DOI: 10.1103/PhysRevLett.121.141802.

[201] CMS Collaboration. "Search for pair-produced three-jet resonances in proton-proton collisions at $\sqrt{s} = 13\ \text{TeV}$". In: *Phys. Rev. D* 99 (1 Jan. 2019), p. 012010. DOI: 10.1103/PhysRevD.99.012010.

[202] ATLAS Collaboration. "Search for pair production of massive particles decaying into three quarks with the ATLAS detector in $\sqrt{s} = 7\ \text{TeV}$ pp collisions at the LHC". In: *Journal of High Energy Physics* 2012.12 (Dec. 2012). ISSN: 1029-8479. DOI: 10.1007/jhep12(2012)086.

[203] ATLAS Collaboration. "Search for massive supersymmetric particles decaying to many jets using the ATLAS detector in pp collisions at $\sqrt{s} = 8\ \text{TeV}$". In: *Phys. Rev. D* 91 (11 June 2015), p. 112016. DOI: 10.1103/PhysRevD.91.112016.

[204] ATLAS Collaboration. "Search for R-parity-violating supersymmetric particles in multi-jet final states produced in p–p collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC". In: *Physics Letters B* 785 (2018), pp. 136–158. ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2018.08.021.

[205] Werner Porod. “SPheno, a program for calculating supersymmetric spectra, SUSY particle decays and SUSY particle production at e+ e- colliders”. In: *Comput. Phys. Commun.* 153 (2003), pp. 275–315. DOI: `10.1016/S0010-4655(03)00222-4`. arXiv: `hep-ph/0301101`.

[206] W. Porod and F. Staub. “SPheno 3.1: Extensions including flavour, CP-phases and models beyond the MSSM”. In: *Comput. Phys. Commun.* 183 (2012), pp. 2458–2469. DOI: `10.1016/j.cpc.2012.05.021`. arXiv: `1104.1573 [hep-ph]`.

[207] Florian Staub. “SARAH 4 : A tool for (not only SUSY) model builders”. In: *Comput. Phys. Commun.* 185 (2014), pp. 1773–1790. DOI: `10.1016/j.cpc.2014.02.018`. arXiv: `1309.7223 [hep-ph]`.

[208] “Reinterpretation of searches for supersymmetry in models with variable R-parity-violating coupling strength and long-lived R-hadrons”. In: (Mar. 2018).

[209] Georges Aad et al. “The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run 3”. In: *JINST* 19.05 (2024), P05063. DOI: `10.1088/1748-0221/19/05/P05063`. arXiv: `2305.16623 [physics.ins-det]`.

[210] Johan Alwall et al. “MadGraph 5: going beyond”. In: *Journal of High Energy Physics* 2011.6 (June 2011). ISSN: 1029-8479. DOI: `10.1007/jhep06(2011)128`. URL: `http://dx.doi.org/10.1007/JHEP06(2011)128`.

[211] S. Mrenna and P. Skands. “Automated parton-shower variations in pythia 8”. In: *Physical Review D* 94.7 (Oct. 2016). ISSN: 2470-0029. DOI: `10.1103/physrevd.94.074005`.

[212] P. Skands, S. Carrazza, and J. Rojo. “Tuning PYTHIA 8.1: the Monash 2013 tune”. In: *The European Physical Journal C* 74.8 (Aug. 2014). ISSN: 1434-6052. DOI: `10.1140/epjc/s10052-014-3024-y`.

[213] Richard D. Ball et al. “Parton distributions with LHC data”. In: *Nuclear Physics B* 867.2 (Feb. 2013), pp. 244–289. ISSN: 0550-3213. DOI: `10.1016/j.nuclphysb.2012.10.003`.

[214] D. J. Lange. “The EvtGen particle decay simulation package”. In: *Nucl. Instrum. Meth. A* 462 (2001). Ed. by S. Erhan, P. Schlein, and Y. Rozen, pp. 152–155. DOI: `10.1016/S0168-9002(01)00089-4`.

[215] Christoph Borschensky et al. “Squark and gluino production cross sections in pp collisions at $\sqrt{s}$ = 13, 14, 33 and 100 TeV”. In: *Eur. Phys. J. C* 74.12 (2014), p. 3174. DOI: `10.1140/epjc/s10052-014-3174-y`. arXiv: `1407.5066 [hep-ph]`.

[216] Pavel M. Nadolsky et al. “Implications of CTEQ global analysis for collider observables”. In: *Phys. Rev. D* 78 (2008), p. 013004. DOI: `10.1103/PhysRevD.78.013004`. arXiv: `0802.0007 [hep-ph]`.

[217] A. D. Martin et al. “Parton distributions for the LHC”. In: *Eur. Phys. J. C* 63 (2009), pp. 189–285. DOI: `10.1140/epjc/s10052-009-1072-5`. arXiv: `0901.0002 [hep-ph]`.

[218] “Multijet simulation for 13 TeV ATLAS Analyses”. In: (2019).

[219] Simone Alioli et al. “A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX”. In: *JHEP* 06 (2010), p. 043. DOI: `10.1007/JHEP06(2010)043`. arXiv: `1002.2581 [hep-ph]`.

[220] M. Bahr et al. “Herwig++ Physics and Manual”. In: *Eur. Phys. J. C* 58 (2008), pp. 639–707. DOI: `10.1140/epjc/s10052-008-0798-9`. arXiv: `0803.0883 [hep-ph]`.

[221] Georges Aad et al. “Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *Eur. Phys. J. C* 81.8 (2021), p. 689. DOI: `10.1140/epjc/s10052-021-09402-3`. arXiv: `2007.02645 [hep-ex]`.

[222] Georges Aad et al. “ATLAS flavour-tagging algorithms for the LHC Run 2 pp collision dataset”. In: *Eur. Phys. J. C* 83.7 (2023), p. 681. DOI: 10.1140/epjc/s10052-023-11699-1. arXiv: 2211.16345 [physics.data-an].

[223] “Identification of Jets Containing b-Hadrons with Recurrent Neural Networks at the ATLAS Experiment”. In: (Mar. 2017).

[224] Georges Aad et al. “ATLAS b-jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV”. In: *Eur. Phys. J. C* 79.11 (2019), p. 970. DOI: 10.1140/epjc/s10052-019-7450-8. arXiv: 1907.05120 [hep-ex].

[225] Georges Aad et al. “Electron and photon performance measurements with the ATLAS detector using the 2015-2017 LHC proton-proton collision data”. In: *JINST* 14.12 (2019), P12006. DOI: 10.1088/1748-0221/14/12/P12006. arXiv: 1908.00005 [hep-ex].

[226] Georges Aad et al. “Muon reconstruction performance of the ATLAS detector in proton-proton collision data at $\sqrt{s}$ =13 TeV”. In: *Eur. Phys. J. C* 76.5 (2016), p. 292. DOI: 10.1140/epjc/s10052-016-4120-y. arXiv: 1603.05598 [hep-ex].

[227] Morad Aaboud et al. “Performance of the ATLAS Trigger System in 2015”. In: *Eur. Phys. J. C* 77.5 (2017), p. 317. DOI: 10.1140/epjc/s10052-017-4852-3. arXiv: 1611.09661 [hep-ex].

[228] ATLAS Collaboration. “Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC”. In: *The European Physical Journal C* 83.10 (Oct. 2023). ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-023-11747-w. URL: http://dx.doi.org/10.1140/epjc/s10052-023-11747-w.

[229] Robert D. Cousins, James T. Linnemann, and Jordan Tucker. “Evaluation of three methods for calculating statistical significance when incorporating a systematic uncertainty into a test of the background-only hypothesis for a Poisson process”. In: *Nucl. Instrum. Meth. A* 595.2 (2008), pp. 480–501. DOI: 10.1016/j.nima.2008.07.086. arXiv: physics/0702156.

[230] Thomas Junk. “Confidence level computation for combining searches with small statistics”. In: *Nucl. Instrum. Meth. A* 434 (1999), pp. 435–443. DOI: 10.1016/S0168-9002(99)00498-2. arXiv: hep-ex/9902006.

[231] Georges Aad et al. “Measurement of hadronic event shapes in high-p_T multijet final states at $\sqrt{s}$ = 13 TeV with the ATLAS detector”. In: *JHEP* 01 (2021). [Erratum: JHEP 12, 053 (2021)], p. 188. DOI: 10.1007/JHEP01(2021)188. arXiv: 2007.12600 [hep-ex].

[232] G. Parisi. “Super Inclusive Cross-Sections”. In: *Phys. Lett. B* 74 (1978), pp. 65–67. DOI: 10.1016/0370-2693(78)90061-8.

[233] Georgios Choudalakis. “On hypothesis testing, trials factor, hypertests and the BumpHunter”. In: *PHYSTAT 2011*. Jan. 2011. arXiv: 1101.0390 [physics.data-an].

[234] Alexander Shmakov et al. “SPANet: Generalized permutationless set assignment for particle physics using symmetry preserving attention”. In: *SciPost Phys.* 12.5 (2022), p. 178. DOI: 10.21468/SciPostPhys.12.5.178. arXiv: 2106.03898 [hep-ex].

[235] Lukas Ehrke et al. “Topological reconstruction of particle physics processes using graph neural networks”. In: *Phys. Rev. D* 107.11 (2023), p. 116019. DOI: 10.1103/PhysRevD.107.116019. arXiv: 2303.13937 [hep-ph].

[236] Anthony Badea et al. “Solving combinatorial problems at particle colliders using machine learning”. In: *Phys. Rev. D* 106.1 (2022), p. 016001. DOI: 10.1103/PhysRevD.106.016001. arXiv: 2201.02205 [hep-ph].

[237] Jimmy Lei Ba, Jamie Ryan Kiros, and Geoffrey E. Hinton. “Layer Normalization”. In: (July 2016). arXiv: 1607.06450 [stat.ML].

[238] T. Aaltonen et al. “Search for new particles decaying into dijets in proton-antiproton collisions at s**(1/2) = 1.96-TeV”. In: *Phys. Rev. D* 79 (2009), p. 112002. DOI: `10.1103/PhysRevD.79.112002`. arXiv: `0812.4036 [hep-ex]`.

[239] G. Aad et al. “Search for New Particles in Two-Jet Final States in 7 TeV Proton-Proton Collisions with the ATLAS Detector at the LHC”. In: *Phys. Rev. Lett.* 105 (2010), p. 161801. DOI: `10.1103/PhysRevLett.105.161801`. arXiv: `1008.2461 [hep-ex]`.

[240] Serguei Chatrchyan et al. “Search for Resonances in the Dijet Mass Spectrum from 7 TeV pp Collisions at CMS”. In: *Phys. Lett. B* 704 (2011), pp. 123–142. DOI: `10.1016/j.physletb.2011.09.015`. arXiv: `1107.4771 [hep-ex]`.

[241] M. Baak et al. “HistFitter software framework for statistical data analysis”. In: *Eur. Phys. J. C* 75 (2015), p. 153. DOI: `10.1140/epjc/s10052-015-3327-7`. arXiv: `1410.1280 [hep-ex]`.

[242] Kyle Cranmer and Itay Yavin. “RECAST: Extending the Impact of Existing Analyses”. In: *JHEP* 04 (2011), p. 038. DOI: `10.1007/JHEP04(2011)038`. arXiv: `1010.2506 [hep-ex]`.

[243] Kyle Cranmer and Lukas Heinrich. “Yadage and Packtivity - analysis preservation using parametrized workflows”. In: *J. Phys. Conf. Ser.* 898.10 (2017). Ed. by Richard Mount and Craig Tull, p. 102019. DOI: `10.1088/1742-6596/898/10/102019`. arXiv: `1706.01878 [physics.data-an]`.

[244] “Transverse momentum response and reconstruction efficiency for jets from displaced decays in the ATLAS detector”. In: (2019).

[245] Morad Aaboud et al. “Search for squarks and gluinos in final states with jets and missing transverse momentum using 36 fb^{-1} of $\sqrt{s} = 13$ TeV pp collision data with the ATLAS detector”. In: *Phys. Rev. D* 97.11 (2018), p. 112001. DOI: `10.1103/PhysRevD.97.112001`. arXiv: `1712.02332 [hep-ex]`.

[246] Georges Aad et al. “Search for long-lived, massive particles in events with displaced vertices and multiple jets in pp collisions at $\sqrt{s}$ = 13 TeV with the ATLAS detector”. In: *JHEP* 06 (2023), p. 200. DOI: `10.1007/JHEP06(2023)200`. arXiv: `2301.13866 [hep-ex]`.

[247] “Performance of the reconstruction of large impact parameter tracks in the ATLAS inner detector”. In: (July 2017).

[248] T. Aaltonen et al. “First Search for Multijet Resonances in $\sqrt{s} = 1.96$ TeV $p\bar{p}$ Collisions”. In: *Phys. Rev. Lett.* 107 (2011), p. 042001. DOI: `10.1103/PhysRevLett.107.042001`. arXiv: `1105.2815 [hep-ex]`.

[249] HEP ML Community. *A Living Review of Machine Learning for Particle Physics*. URL: `https://iml-wg.github.io/HEPML-LivingReview/`.

[250] Diederik P. Kingma and Jimmy Ba. *Adam: A Method for Stochastic Optimization*. 2017. arXiv: `1412.6980 [cs.LG]`. URL: `https://arxiv.org/abs/1412.6980`.

[251] David E. Rumelhart, Geoffrey E. Hinton, and Ronald J. Williams. “Learning representations by back-propagating errors”. In: *Nature* 323.6088 (1986), pp. 533–536. DOI: `10.1038/323533a0`.

[252] Kurt Hornik, Maxwell Stinchcombe, and Halbert White. “Multilayer feedforward networks are universal approximators”. In: *Neural Networks* 2.5 (1989), pp. 359–366. DOI: `10.1016/0893-6080(89)90020-8`.