



Thèse

2025

Open Access

This version of the publication is provided by the author(s) and made available in accordance with the copyright holder(s).

Search for Top Squark Pair Production and the ATLAS Run 3 Trigger Menu

Cella, Sofia

How to cite

CELLA, Sofia. Search for Top Squark Pair Production and the ATLAS Run 3 Trigger Menu. Doctoral Thesis, 2025. doi: 10.13097/archive-ouverte/unige:187959

This publication URL: <https://archive-ouverte.unige.ch/unige:187959>

Publication DOI: [10.13097/archive-ouverte/unige:187959](https://doi.org/10.13097/archive-ouverte/unige:187959)

Search for Top Squark Pair Production and the ATLAS Run 3 Trigger Menu

THÈSE

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

par

Sofia Cella

de Varese (Italie)

Thèse N° 5936

GENÈVE
2025



**UNIVERSITÉ
DE GENÈVE**

FACULTÉ DES SCIENCES

DOCTORAT ÈS SCIENCES, MENTION PHYSIQUE

Thèse de Madame Sofia CELLA

intitulée :

«Search for Top Squark Pair Production and the ATLAS Run 3 Trigger Menu»

La Faculté des sciences, sur le préavis de

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

Monsieur S. SCHRAMM, professeur assistant
Département de physique nucléaire et corpusculaire

Monsieur B. PETERSEN, docteur
Experimental Physics Department, Conseil Européen pour la Recherche Nucléaire (CERN),
Genève, Suisse

Monsieur I. VIVARELLI, professeur
Dipartimento di Fisica e Astronomia, Università di Bologna , Bologna, Italie

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

Genève, le 3 septembre 2025

Thèse - 5936 -

La Doyenne

Resumé

Le travail présenté dans cette thèse est structuré en deux parties distinctes. La première décrit une recherche en supersymétrie, comprenant une stratégie d'analyse novatrice basée sur des techniques de machine learning. La seconde se concentre sur le système de déclenchement à haut niveau d'ATLAS, avec une attention particulière au contrôle des performances et l'optimisation des ressources dans le contexte de la récente amélioration Phase-I.

Cette thèse présente une recherche pour la production directe de paire de squark top supersymétrique dans des événements contenant deux leptons de charge opposée (électrons ou muons), des jets et de l'impulsion transverse manquante. L'analyse est basée sur des données issues de collisions protons-protons collectées avec le détecteur ATLAS au Grand Collisionneur de Hadron (LHC), correspondant à une luminosité totale intégrée de 193 fb^{-1} .

L'ensemble de données est constitué de 140 fb^{-1} à une énergie dans le centre de masse de $\sqrt{s} = 13 \text{ TeV}$, collecté entre 2015 et 2018 pendant la seconde période de prise de données (Run 2) du LHC, et 53 fb^{-1} à $\sqrt{s} = 13.6 \text{ TeV}$, enregistré en 2022 et 2023 pendant la troisième période (Run 3). La recherche cible spécifiquement les squarks top passant par une désintégration à deux corps comprenant un quark top et le plus léger des neutralinos, et elle est sensible à travers un grand intervalle de différence de masse entre les deux particules supersymétriques. En utilisant tout l'ensemble de données, l'analyse atteint une sensibilité attendue à la masse du squark top allant jusqu'à 1.1 TeV , et à la masse de neutralino jusqu'à 600 GeV . La sensibilité de l'analyse est également évaluée pour des modèles dans lesquels la différence de masse entre le squark top et le neutralino le plus léger est inférieure à la masse du quark top. Dans ce scénario, le squark top se désintègre en neutralino le plus léger ainsi qu'un boson W et un quark b par le biais d'un quark top hors de sa couche de masse. Pour ces modèles, la sensibilité de l'analyse s'étend aux masses de neutralino allant jusqu'à 450 GeV pour une masse de squark top de 600 GeV . Finalement, l'analyse est appliquée à un modèle consistant en la production associée de paires quark top-antitop avec un boson de Higgs se désintégrant invisiblement, résultant en une limite supérieure attendue sur le rapport de branchement invisible du boson de Higgs de 0.29 , basé sur l'ensemble de données complet de 193 fb^{-1} .

La seconde partie de cette thèse présente le système de déclenchement d'ATLAS dans sa configuration pour la troisième période de prise de données, suite à une importante mise à jour entre la seconde et la troisième période de prise de données du LHC. Depuis le début de la troisième période, ATLAS enregistre jusqu'à 3 kHz d'événements physiques entièrement reconstruit à partir des collisions du LHC à 40 MHz , avec un taux supplémentaire alloué à la lecture partielle d'événements. Le système de déclenchement à deux niveaux est responsable de la sélection en temps réel des événements d'intérêt pour le large programme physique d'ATLAS, tout en rejetant à un taux élevé les événements de bruit de fond. La sélection des événements, encodée dans le menu du système de déclenchement d'ATLAS, cible à la fois des signatures physiques génériques et des spécifiques pour des analyses. L'un des principaux défis du système de déclenchement d'ATLAS est de maximiser le résultat physique de l'ensemble de données enregistré tout en opérant dans des contraintes strictes sur la bande passante et les ressources des processeurs. Cette thèse fournit une vue d'ensemble détaillée des outils de surveillance de la performance du système de déclenchement à haut niveau. Un accent particulier est mis sur la stratégie d'application des facteurs de présélection pour gérer les taux de déclenchement et l'utilisation du processeur. Dans ce contexte, une contribution centrale de ce travail est le développement et l'optimisation du paquet logiciel utilisé pour générer les ensembles de facteurs de présélection pour la troisième prise de données d'ATLAS.

Abstract

The work presented in this thesis is structured in two distinct parts. The first describes a supersymmetry search, including a novel analysis strategy based on machine learning techniques. The second focuses on the ATLAS high-level trigger system, with particular attention to performance monitoring and resource optimisation in the context of the recent Phase-I upgrade.

This thesis presents a search for the direct pair production of supersymmetric top squarks in events containing two opposite-charge leptons (electrons or muons), jets, and missing transverse momentum. The analysis is based on proton-proton collision data collected with the ATLAS detector at the Large Hadron Collider (LHC), corresponding to a total integrated luminosity of 193 fb^{-1} . This dataset consists of 140 fb^{-1} at a centre-of-mass energy $\sqrt{s} = 13 \text{ TeV}$, collected between 2015 and 2018 during the LHC Run 2 data-taking period, and 53 fb^{-1} at $\sqrt{s} = 13.6 \text{ TeV}$, recorded in 2022 and 2023 during the Run 3 of the LHC. The search specifically targets top squarks undergoing a two-body decay into a top quark and the lightest neutralino, and it is sensitive across a wide range of mass differences between the two supersymmetric particles. Using the full dataset, the analysis achieves an expected sensitivity to top squark masses beyond 1.1 TeV , and neutralino masses up to 600 GeV . The sensitivity of the analysis is also evaluated for models in which the mass difference between the top squark and the lightest neutralino is smaller than the top quark mass. In these scenarios, the top squark decays to the lightest neutralino, and to a W boson and a b -quark via an off-shell top. For these models, the analysis sensitivity extends to neutralino masses up to 450 GeV , for top squark masses of 600 GeV . Finally, the analysis is applied to a model consisting of the associated production of top-antitop quark pairs with an invisibly decaying Higgs boson, resulting in an expected upper limit on the Higgs boson's invisible branching ratio of 0.29 , based on the full 193 fb^{-1} dataset.

The second part of this thesis presents the ATLAS trigger system as configured for Run 3, following a major upgrade between the Run 2 and the Run 3 data-taking periods of the LHC. Since the start of Run 3, ATLAS is recording up to 3 kHz of fully built physics events out of an LHC collision rate of 40 MHz , with additional rate allocated to partial event readout. The two-level trigger system is responsible for selecting in real time events of interest for the broad ATLAS physics programme, while rejecting a high rate of background events. The selection of events, encoded in the ATLAS trigger menu, targets both generic and analysis specific physics signatures. One of the main challenges of the ATLAS trigger system is to maximise the physics outcome of the recorded dataset while operating within strict constraints on bandwidth and CPU resources. This thesis provides a detailed overview of the high-level-trigger performance monitoring tools. A particular focus is given to the strategy for applying prescale factors to manage trigger rates and CPU usage. In this context, a central contribution of this work is the development and optimisation of the software package used to generate prescale sets for ATLAS data-taking in Run 3.

Contents

Introduction and personal contributions	1
I <i>From Physics Theory to Analysis Objects</i>	7
1 Standard Model and beyond	9
1.1 The Standard Model	9
1.1.1 Elementary particles	9
1.1.2 Fundamental interactions	10
1.2 Supersymmetry	19
1.2.1 Limitations of the Standard Model	19
1.2.2 Minimal Supersymmetric Standard Model	23
1.2.3 Top squark	30
1.2.4 Natural SUSY searches at the LHC	34
2 Experimental setup	41
2.1 The Large Hadron Collider	41
2.1.1 Accelerator complex	41
2.1.2 LHC fill cycle	43
2.1.3 Luminosity and pile-up	44
2.2 The ATLAS detector	50
2.2.1 Magnet system	53
2.2.2 Tracking: the Inner Detector	55
2.2.3 Calorimetry	57
2.2.4 Muon System	61
2.2.5 Forward Detectors	63
2.2.6 Trigger and data acquisition system	63
3 Object reconstruction in ATLAS	69
3.1 Common reconstruction algorithms	69
3.1.1 Tracks and vertices	69
3.1.2 Topological clusters	72
3.1.3 Particle-flow	73
3.2 Electrons	74
3.2.1 Electron reconstruction and calibration	74
3.2.2 Electron identification	76
3.2.3 Electron isolation	78
3.3 Muons	79

3.3.1	Muon reconstruction and calibration	79
3.3.2	Muon identification and isolation	81
3.4	Jets	82
3.4.1	Jet reconstruction	82
3.4.2	Jet energy calibration and resolution	84
3.4.3	b -tagging	86
3.5	Missing transverse momentum	88
3.5.1	Missing transverse momentum significance	89
II	<i>Search for top squark pairs</i>	93
4	Overview and datasets	95
4.1	Data sets and simulated samples	97
4.1.1	Data	97
4.1.2	Simulated samples	98
4.2	Object reconstruction	104
5	Analysis strategy and background estimation	107
5.1	Event selection	107
5.1.1	Discriminating variables	107
5.1.2	NN discriminant	115
5.2	Analysis regions	126
5.2.1	Signal region definition	127
5.2.2	Estimation of the $t\bar{t}$ background	127
5.2.3	Estimation of the $t\bar{t}Z$ background	132
6	Statistical Analysis and Results	137
6.1	Statistical Model	137
6.1.1	Likelihood function	137
6.1.2	Hypothesis test	139
6.2	Systematic uncertainties	142
6.2.1	Experimental uncertainties	143
6.2.2	Theoretical systematic uncertainties	143
6.3	Results and interpretation	148
6.3.1	Event yields	148
6.3.2	Interpretation	150
6.3.3	Invisible Higgs reinterpretation	153
III	<i>Trigger Menu and Operations</i>	157
7	ATLAS Run-3 Trigger	159
7.1	Physics motivation	159
7.2	The ATLAS HLT	163
7.3	Performance monitoring	166

7.3.1	Rates and cost analysis	166
7.3.2	Performance scaling	169
7.4	Operations	172
7.4.1	Trigger configuration	173
8	Trigger Menu	175
8.1	Menu design	175
8.1.1	System limitations	175
8.1.2	Streaming model	176
8.1.3	Menu code structure	178
8.1.4	Menu for pp physics	179
8.2	The Rulebook	182
8.2.1	Prescaling strategy	182
8.2.2	Prescale Rules	183
8.2.3	Rules processing	186
	Conclusions and outlook	195
	Appendices	
A	Nuisance Parameters	201
	Bibliography	207
	List of Acronyms	221
	List of Figures	225
	List of Tables	235
	Acknowledgments	239

Introduction and personal contributions

The Standard Model (SM) of particle physics is a unified theoretical framework that describes all known subatomic particles and their quantum interactions through the electromagnetic, weak, and strong forces. Over the past decades, its predictions have been tested with an excellent level of accuracy by a wide range of experimental results. Nevertheless, the SM leaves several fundamental questions unanswered, hinting that it may be an incomplete theory. One notable example is the absence of an explanation of the nature of dark matter and dark energy, which, according to cosmological observations, constitute approximately 96% of the content of the universe. In addition, the SM does not naturally lead to the unification of forces at high energies, a feature that characterise Grand Unification Theories, a class of theoretical models where the three fundamental gauge interactions merge into a single force described by a larger gauge group at very high energy scales. Finally, a strong hint of the SM unnaturalness is the so-called hierarchy problem: the electroweak scale (≈ 100 GeV) is unnaturally small compared to the Planck scale ($\approx 10^{19}$ GeV), leading, for example, to extreme sensitivity of the Higgs boson mass to quantum corrections without fine-tuning. These and many more theoretical and experimental shortcomings suggest the need for extensions of the SM, motivating the investigation of physics beyond the SM. Many theories have been developed across the years to offer an extension of the SM able to address as much as possible the currently open problems. Among the various proposed extensions, Supersymmetry (SUSY) stands out as one of the most promising, as it addresses simultaneously many of these open issues.

SUSY introduces a new symmetry between fermions and bosons by postulating that every SM particle has a supersymmetric partner with the same quantum numbers except for spin, which differs by half a unit. Since these particles have not been observed yet, if SUSY is realised in nature, it must be a broken symmetry, with supersymmetric partners heavier than their SM counterparts. SUSY offers natural solutions to all the open problems of the SM discussed earlier. First, it enables the unification of the gauge couplings at high energies. Under the assumption of R-parity [1] conservation, the lightest supersymmetric particle (LSP) is stable and weakly interacting, therefore it has the needed features to be the constituent of dark matter. Finally, the introduction of new particles allows the cancellation of the divergences in the calculation of the Higgs boson mass without requiring fine-tuning, thus solving the hierarchy problem. The most significant contribution to this cancellation comes primarily from the supersymmetric partner of the top quark: the top squark, or stop. In particular, a light Higgs boson mass arises naturally

in SUSY if top squark mass is close to the TeV scale. This makes the top squark a good candidate for discovery at the Large Hadron Collider (LHC) [2].

The LHC is currently the most powerful proton-proton (pp) collider in the world, and the ATLAS detector [3] is one of the two multipurpose experiments recording its collisions, aiming to study the broadest possible range of physics processes, from precise SM measurements to searches for new physics. The second data-taking period of the LHC, known as Run 2, ended in 2018. In the period between 2018 and 2022, the accelerator and the experiments were upgraded to prepare to the Run 3 data-taking period, which began in July 2022 and will continue until June 2026. The centre-of-mass energy was increased with respect to Run 2 from 13 TeV to 13.6 TeV allowing the experiments to push the discovery sensitivities for beyond the SM particles. In particular, the production cross section for top squark pair production, which depends on the assumed stop mass, increases significantly at this energy, for example by approximately 40% for a stop mass of 1.25 TeV. Top squark searches have always been a crucial component of the ATLAS SUSY programme, due to the strong theoretical motivation, and this increase in the cross section production gives a strong motivation to start new searches analysing the data collected during the first years of the Run 3.

The main search presented in this thesis targets the production of top squark pairs, in final states containing exactly two leptons with opposite electric charge, jets and large missing transverse momentum. It is based on a SUSY simplified model in which only the top squark and the lightest neutralino (assumed to be the LSP) are accessible at the LHC energies. The only decay mode considered is $\tilde{t}_1 \rightarrow t\tilde{\chi}_1^0$, restricting the kinematic phase space to regions where the mass difference $m_{\tilde{t}_1} - m_{\tilde{\chi}_1^0}$ exceeds the top quark mass. The requirement of two leptons in the final state selects events where both top quarks decay leptonically ($t \rightarrow Wb \rightarrow \ell\nu b$). Although this decay channel has a smaller branching ratio compared to the complementary semi-leptonic and fully hadronic channels, the requirement of two leptons provides a cleaner experimental signature, allowing for more effective suppression of SM backgrounds. This search will be one of the first SUSY results using partial Run 3 data in ATLAS, and benefits not only from the extended dataset and the dramatic increase in the top squark production cross section, but also from improved object reconstruction and calibration techniques developed for the processing of the ATLAS Run 3 data. In addition, the sensitivity of the analysis is enhanced thanks to the use of a novel signal-versus-background discrimination method based on a dedicated neural network, developed as part of the project of this thesis.

The LHC collides proton bunches at a rate of 40 MHz, producing a volume of data that significantly exceeds the technical capabilities for the transfer, storage and offline processing of the computing systems of the experiments. Moreover, the physics processes of interest are extremely rare, and most of the collision events produced by the LHC are characterised by the overwhelming soft QCD background arising from inelastic pp interactions. The trigger system is a crucial component of all the experiments, as it performs a real-time selection of the events of interest for physics analyses. This reduces the data rate, enabling the system to remain within the technical limitations while preserving sensitivity to a wide range of physics signatures. Trigger systems of multipurpose experiments like ATLAS have the capability of applying concurrently multiple filtering strategies, in order to collect data for multiple physics measurements at the same time. The ATLAS trigger system is composed of two levels: a hardware-based Level-1

(L1) trigger, using custom electronics to trigger on coarse granularity information from the detectors to reduce the event rate to 100 kHz, and a software-based High-Level Trigger (HLT), applying a more refined selections to further reduce the rate to about 3 kHz for permanent storage. The primary challenge in the design and operation of the trigger system is to maximise the physics potential of the experiment while operating within the constraints on the L1 rate, the HLT output bandwidth, and the processing capacity. The Run 3 LHC running conditions significantly differ with respect to Run 2. In particular, the average number of simultaneous pp interactions per bunch crossing (pile-up) has nearly doubled, reaching values of more than 60. Before the start of Run 3, the ATLAS trigger system underwent a major upgrade [4], to optimise the performance under these new, harsher conditions. The upgraded system was designed to maintain the same selection efficiency of Run 2 without exceeding the rate limitations, despite the increased event complexity and detector occupancy. The work presented in this thesis focused on the HLT, in this upgraded Run 3 configuration.

The trigger system constraints are handled through the trigger menu, which encodes the full set of L1 and HLT selections used during data-taking, and is tailored to the specific running conditions of the LHC and the operational status of the detector. It is designed to maximise the physics outcome making maximum use of the available resources based on the different requirements of the ATLAS physics programme. In particular, most of the available output bandwidth is dedicated to general-purpose triggers that can be used by different analysis groups. However, the menu also includes triggers targeting specific analyses signatures, particularly exploiting the additional bandwidth available in Run 3 thanks to the TDAQ system upgrades. The trigger menu is in constant evolution, continuously updated and is optimised to adapt to the experimental conditions. The development of the trigger menu are based on studies on the performance of the trigger system, which is constantly monitored, both in terms of physics selection and resource usage. This thesis presents the data-driven method used to estimate the HLT performance in terms of output rate and CPU usage prior to data-taking, based on a dedicated Enhanced Bias (EB) dataset, enriched with high p_T and high object multiplicity, which are more likely to be selected by the HLT. The EB event selection is invertible, by assigning a weight to each event to corrects for the sampling bias applied during the EB data-taking, restoring an effective unbiased spectrum.

Since the start of Run 3, the HLT software is capable of running both in multiprocessing and in multithreading configuration. It is therefore important to know how the performance of the HLT scales depending on the number of events processed concurrently and based on the different setup chosen to achieve parallelism. Studies on the performance scaling are used to choose the optimal software configuration to use during data-taking, and are part of the work presented in this thesis.

One of the key features of the trigger system is its ability to make maximum use of the available resources also adapting to the running conditions during data-taking. A way to do so, is the application of prescale factors, which are also part of the trigger menu definition. A trigger chain with a prescale factor of n , accepts only one out of every n events it would otherwise select. Individual prescale factors can be given to L1 items and HLT chains, and are applied to control the rate of accepted events and to manage CPU consumption at the HLT. Prescale factors can evolve during data-taking, in response to changing conditions, for example, increasing rates when more resources become available toward the end of an LHC fill. This thesis describes the methodology used to compute prescale sets adapted to a wide range of data-taking

conditions, as implemented in the Rulebook software package. The Rulebook is designed to be flexible, offering numerous configurable options to adjust prescale factors according to the evolving operational scenarios of the experiment.

Thesis outline and personal contributions

This thesis is organised into three parts.

Part I provides the theoretical and experimental foundation for the analysis presented in Part II. In particular, Chapter 1 outlines the theoretical framework, beginning with an overview of the Standard Model and its unresolved questions, and introducing Supersymmetry (SUSY) as a possible extension. The fundamental concepts of SUSY are described, with particular attention to the phenomenology of the top squark. The two following chapters introduce the experimental framework. Specifically, Chapter 2 describes the LHC and the ATLAS detector. After that, Chapter 3 then details how the particles characterising the final-state of the presented search are reconstructed and calibrated, starting from the ATLAS detector signal.

Part II is dedicated to the search for top squark pair production in final states with exactly two isolated leptons, at least two jets (including at least one b -tagged jet), and missing transverse momentum. This part is structured into three chapters. Chapter 4 provides an overview of the analysis strategy, and describes the utilised datasets and the reconstruction configurations adopted for the physics objects. Chapter 5 describes the event selection strategy, including the development of a dedicated machine learning model for signal-to-background discrimination, and the procedure used to estimate the SM backgrounds. Finally, Chapter 6 details the statistical model, discusses the sources of systematic uncertainty, and presents the results of the statistical analysis and their interpretation in terms of expected exclusion limits at 95% CL on the signal parameter space.

Part III focuses on the ATLAS High-Level Trigger and the trigger menu. Particularly, Chapter 7 discusses the requirements and challenges in designing a trigger system for a multipurpose experiment like ATLAS, aiming to maximise the physics outcome. It then presents an overview of the ATLAS HLT software in its Run 3 configuration, including a description of some key tools used for performance monitoring, and the operation of the trigger system in the context of the ATLAS detector data taking. Finally, Chapter 8 explains the trigger menu design process, describing how event selections are defined to align with the ATLAS physics programme while satisfying system limitations. A particular focus is given to the prescale strategy, with a detailed description of how the prescale sets used during data-taking are generated using the Rulebook code.

The ATLAS collaboration counts more than 170 institutions and over 3000 scientific authors. The members of the collaboration contribute to all ATLAS results working on the multiple aspects of the experiment, including the maintenance and construction of the detector, the collection and analysis of the data, the calibration of the detector, and many other tasks. The work presented in this thesis could not be performed by a single individual, and it relies on the work of many people. My main contributions are detailed below.

Part II - Search for top squark pairs: I significantly contributed to all the aspects of the analysis, being one of the two main analysers. I was responsible for the design and implementation of a

dedicated neural network for event selection (see Section 5.1.2), including the development of a complete framework for data preparation, training, and performance evaluation. I optimised the architecture of the neural network by evaluating multiple configurations to determine the optimal number of models to be used for distinct event categories within the analysis. I defined the analysis regions (control, validation, and signal regions) based on the output of the neural network (see Section 5.2). I also contributed to the implementation of the framework used for the statistical interpretation and the extraction of the final results (see Sections 6.1, 6.2, and 6.3). In addition, I served as editor of the internal supporting documentation, detailing the full analysis. This analysis is under internal review at the time of writing, and aims to be published in a paper in the next months.

Part III - Trigger Menu and Operations: I have been involved in many activities related to the ATLAS HLT software, being particularly involved in the trigger menu group. In this context, I am the main maintainer of the Rulebook package (see Section 8.2), which is used daily by ATLAS menu experts to generate the prescale sets for data-taking, as discussed in Chapter 8. My work on the Rulebook included extensive development of new functionalities, as well as the improvement and refactoring of existing code. I also provide support in case of any issues occurring during daily operations activities. In parallel, as a menu expert, I contributed to developments of the trigger menu code within the Athena framework. I am the author or co-author of all the plots shown in Chapter 8, taken from Reference [5].

In addition to my work within the menu group, I have contributed to the development and maintenance of trigger performance monitoring tools within Athena, used to estimate the trigger rates prior to collisions and to evaluate the CPU cost of trigger algorithms. Notably, I developed a new workflow to run the production of Enhanced Bias (EB) weights on the LHC computing grid. This workflow has become the standard procedure for generating the EB weights following each new EB run, as described in Section 7.3.1. Additionally, I conducted dedicated studies on the performance scaling of the HLT software as a function of the number of events processed in parallel (see Section 7.3.2). The results of these studies were published as part of Reference [4]. Beyond software development, I contributed to the smooth operation of the ATLAS trigger system by serving as both a trigger menu on-call expert and a trigger online on-call expert, covering more than 100 days of shifts. Menu on-call experts are responsible for preparing the trigger configuration used during data-taking, while the trigger online on-call expert is the primary contact for any trigger-related issues occurring during detector operations. Finally, I participated in the preparation of a white paper on the trigger systems of the LHC experiments, published as Reference [6].

Part I

From Physics Theory to Analysis Objects

Standard Model and beyond

This chapter introduces the theoretical framework that motivates the analysis presented in this thesis, which is a search for top squark pair production. Top squarks are new, still undetected particles predicted by Supersymmetry (SUSY), a proposed extension of the Standard Model (SM) of particle physics.

The SM is a quantum field theory that provides a unified description of the strong and electroweak interactions, and classifies all known elementary particles, including force carriers and matter particles. Mass terms for elementary particles are generated via the Higgs mechanism. An overview on the SM theory is provided in Section 1.1.

Although in the past decades its predictions have been experimentally verified to an excellent level of accuracy, the SM presents several theoretical and experimental limitations, suggesting that it could be an incomplete theory. Various beyond the Standard Model (BSM) theories have been proposed to address these issues, and Supersymmetry is one of the most studied and promising candidates. SUSY introduces a symmetry between bosons and fermions, predicting that each SM particle has a corresponding superpartner with a difference in spin of one-half. Section 1.2 discusses the limitations of the SM and presents the theoretical basis of SUSY, with particular focus on the Minimal Supersymmetric Standard Model (MSSM) and the top squark, including its phenomenology at hadron colliders.

1.1 The Standard Model

The SM is a gauge quantum field theory based on the symmetry group $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$, where $SU(3)_C$ corresponds to QCD, the theory of the strong interaction, and $SU(2)_L \otimes U(1)_Y$ describes the electroweak interaction. The electroweak symmetry $SU(2)_L \otimes U(1)_Y$ is spontaneously broken via the Higgs mechanism to $U(1)_{EM}$, which corresponds to the electromagnetic interactions. The SM does not include a description of gravity. However, gravitational effects are negligible at the energy scales probed by current particle colliders, being approximately 10^{38} times weaker than the strong interaction at subatomic distances.

1.1.1 Elementary particles

The elementary particles described by the SM are classified in two main categories: *fermions*, the fundamental constituents of matter, and *bosons*, which mediate the fundamental interactions. The boson sector includes the vector gauge bosons, responsible for the fundamental forces, and

the Higgs boson, responsible for mass generation mechanism.

The SM includes twelve fundamental fermions and their corresponding anti-particles, which have identical masses and opposite quantum numbers. These fermions are described by Dirac spin- $\frac{1}{2}$ fields, and can be further divided in two groups: *quarks*, and *leptons*. Both quarks and leptons are classified into three *generations* (or *families*), each with the same quantum numbers but progressively increasing masses. Leptons interact only via the electroweak force, and include both electrically charged particles (electrons e , muons μ , and taus τ) and the electrically neutral neutrinos (ν_e , ν_μ , and ν_τ). Quarks, on the other hand, interact via both the electroweak and strong forces. Due to colour confinement (see Section 1.1.2), quarks cannot be observed as free particles and can only be found in composite colour-neutral states called *hadrons*, which include baryons and mesons. In the original formulation of the SM, all the matter fields are massive, except for the neutrinos, which are considered massless. However, the observation of neutrino oscillations has confirmed that neutrinos have non-zero mass [7]. These masses are anyway extremely small relative to those of other fermions ($m_\nu < 0.8 \text{ eV}$ [8]), and can be neglected for the purposes of this thesis.

The requirement of local gauge invariance (see Section 1.1.2) in the SM leads to the introduction of spin-1 gauge fields associated with the generators of the SM gauge group. The quantum excitations of these fields are the *gauge bosons*, which are the mediators of the fundamental interactions of the SM. The *gluons* (g), mediators of the strong interaction, are massless, electrically neutral, and carry colour charge, which allows them to interact with quarks and to self-interact. The *photons* (γ), mediators of the electromagnetic interaction, are also massless and neutral, and couple only to electrically charged quarks and leptons. The massive bosons ($W^\pm$ and Z) mediate the weak interaction. The $W^\pm$ bosons are responsible for *charged-current* interactions and produce charge and flavour variations; the Z boson mediates *neutral-current* interactions. Each fundamental force is characterised by a different range of interaction: the electromagnetic force has infinite range, while the weak and strong forces act only at a subatomic level.

Finally, the Higgs boson is a neutral fundamental scalar particle, and it is responsible for the mass generation of both the weak vector bosons and the fermions (see Section 1.1.2). All SM particles and their quantum numbers are shown in Figure 1.1.

1.1.2 Fundamental interactions

The SM provides a consistent description of the strong and the electroweak interactions across a wide range of energy scales. This is achieved by imposing local gauge invariance under the already mentioned $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$ symmetry group. In gauge theories, the structure of the interactions is naturally set by the choice of the symmetry group and by the requirement that the Lagrangian be invariant under local gauge transformations. The coupling between matter and gauge fields then emerges naturally from the principle of local symmetry.

The definition of a gauge theory follows the steps described below.

- **Definition of a symmetry group.** A Lie group is selected to describe the internal symmetries of the theory. In the SM, the group $SU(3)_C$ is used to derive the QCD theory, while $SU(2)_L \otimes U(1)_Y$ governs the electroweak interaction.
- **Free Lagrangian for matter fields.** The starting point is the Lagrangian for free, massless fermions, which for a Dirac spin- $\frac{1}{2}$ field $\psi(x)$ is: $\mathcal{L}_{\text{fermion}} = i\bar{\psi}\gamma^\mu\partial_\mu\psi$. This Lagrangian




Figure 1.1 – Fundamental particles of the Standard Model [9].

is required to be invariant under *global* (i.e., space-time independent) transformations of the considered symmetry group typically represented by a unitary operator $U = e^{i\theta^a T^a}$, where θ^a are constant phases (i.e. $\partial_\mu \theta^a = 0$), and T^a are the generators of the Lie algebra associated with the group. The index a runs over the number of group generators. As regards the SM groups, $U(1)$ has only one generator, $SU(2)$ has 3 generators, and $SU(3)$ has 8 generators.

- **Gauge principle.** According to the gauge principle, the global symmetry is promoted to a *local* symmetry. This means that the constant phases θ^a are replaced by space-time dependent functions: $\theta^a \rightarrow \theta^a(x)$. Under such local transformations, the partial derivative term $\partial_\mu \psi$ in the Lagrangian introduces additional terms due to the non-zero value of $\partial_\mu \theta^a(x)$, which break the invariance of the theory. To restore invariance under local transformations, the partial derivative is replaced by a *covariant derivative*: $D_\mu = \partial_\mu + ig A_\mu^a T^a$, where A_μ^a are new spin-1 (vector) gauge fields, introduced for each generator T^a of the symmetry group, and g is the *coupling constant*. The transformation law of the vector field A_μ^a is chosen to cancel out the unwanted term containing the partial derivative of $\theta(x)$, thus restoring local gauge invariance.
- **Gauge-field kinematic term.** The Lagrangian is completed with the addition of a kinetic term describing the free propagation of the gauge fields. This term is constructed from the *field strength tensor* $F_{\mu\nu}^a = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + gf^{abc} A_\mu^b A_\nu^c$, where f^{abc} are the *structure constants* of the Lie algebra. The non-linear term involving f^{abc} , which is non-zero only for

non-Abelian¹ gauge groups, is responsible for the self-interactions of the gauge bosons. Such self-interactions are therefore absent in theories based on an Abelian gauge group, like electromagnetism, governed by the unbroken $U(1)_{\text{EM}}$ symmetry. The corresponding gauge-invariant kinetic term in the Lagrangian is $-\frac{1}{4}F_{\mu\nu}^a F^{\mu\nu a}$, which describes the propagation of the gauge bosons. In the non-Abelian case, this term also contains cubic and quartic interactions among the gauge fields, arising from the non-linear structure of $F_{\mu\nu}^a$.

This procedure is used to construct the strong and electroweak Lagrangians following the gauge principle. The theories of QCD and electroweak interactions are presented in more detail in the next paragraphs, along with an overview on the Higgs mechanism.

Quantum Chromodynamics

Quantum chromodynamics (QCD) is the quantum field theory that describes the strong interaction acting between quarks and gluons, which are the elementary constituents of the protons colliding in the LHC. An accurate understanding of QCD is therefore fundamental to describe the physics processes occurring in the LHC pp collisions. QCD is a non-Abelian gauge theory based on the $SU(3)_C$ symmetry group, where the subscript C refers to the colour charge, carried by particles taking part in strong interactions. The colour charge can take three values, labelled *red*, *green*, and *blue*, and quarks carry one of the colour charges.

Each quark field of a given flavour (u, d, c, s, t , or b) belongs to the fundamental representation of the group, transforming as a colour triplet under $SU(3)_C$. This is represented mathematically as a three-component vector, $q_f(x) = (q_f^{\text{red}}(x), q_f^{\text{green}}(x), q_f^{\text{blue}}(x))$, where each component corresponds to one colour state. Antiquarks, on the other hand, transform under the complex conjugate of the fundamental representation.

The eight generators of the $SU(3)_C$ Lie algebra correspond to the eight distinct gluon fields, which are the massless spin-1 gauge bosons that mediate the strong force. Gluons themselves carry colour charge, allowing them to interact with each other, as also confirmed by experimental evidence [10].

Quarks and gluons are never observed in isolation, but are bound into colour-neutral (singlet) states, such as *mesons*, composed of a quark and an antiquark ($q\bar{q}$), and *baryons*, composed of three quarks or three antiquarks ($qqq, \bar{q}\bar{q}\bar{q}$).

The QCD Lagrangian can be written as:

$$\mathcal{L}_{\text{QCD}} = \sum_f i\bar{q}_f^\alpha \gamma^\mu (D_\mu)_{\alpha\beta} q_f^\beta - \frac{1}{4}F_{\mu\nu}^i F^{i\mu\nu}, \quad (1.1)$$

where q_f^α represents a quark field of flavour $f = (u, d, c, s, t, b)$ and colour index $\alpha, \beta = (1, 2, 3)$ (corresponding to red, green, blue).

The local gauge invariance under $SU(3)_C$ is achieved by introducing the covariant derivative D_μ , defined as:

$$(D_\mu)_{\alpha\beta} = \partial_\mu \delta_{\alpha\beta} + ig_s G_\mu^i \lambda_{\alpha\beta}^i, \quad (1.2)$$

¹Abelian groups are defined as groups where the group operation is commutative, i.e. the result of applying the group operation to two group elements does not depend on the order in which they are written.

where g_s is the strong coupling constant, G_μ^i are the eight gluon gauge fields (with $i = 1, \dots, 8$) and λ^i are the Gell-Mann matrices, generators of the SU(3) Lie algebra in the fundamental representation.

The second term in $\mathcal{L}_{\text{QCD}}$ is the kinetic term for the gluons. The field strength tensor $F_{\mu\nu}^i$ is given by:

$$F_{\mu\nu}^i = \partial_\mu G_\nu^i - \partial_\nu G_\mu^i - g_s f^{ijk} G_\mu^j G_\nu^k, \quad (1.3)$$

where f^{ijk} ($i, j, k = 1, \dots, 8$) are the structure constants of the SU(3) group. The third term in the field strength arises from the non-Abelian nature of SU(3)_C, and generates the gluons' self-interactions.

The three quarks (or antiquarks) and quark-antiquark pairs constituting baryons and mesons, respectively, are called *valence quarks* and determine the hadron quantum numbers. However, along with the valence quarks, hadrons contain gluons, emitted and reabsorbed by quarks, and quark-antiquark pairs generated in $g \rightarrow q\bar{q} \rightarrow g$ processes, called *sea quarks*. Therefore, there is a non-zero probability of finding gluons, anti-quarks and quarks with different flavours from those of the valence quarks in hadrons. Sea quarks are more unstable than valence quarks, and they generally annihilate each other within the interior of the hadron.

The renormalisation procedure in QCD introduces a dependence of the strong coupling $\alpha_s = \frac{g_s^2}{4\pi}$ on the energy scale Q^2 (*running coupling constant*). The dependence on the energy scale Q^2 is described by the *Renormalisation Group Equation*, and for high values of Q^2 it gives:

$$\alpha_s(Q^2) \approx \frac{12\pi}{(33 - 2n_f) \cdot \ln(Q^2/\Lambda_{\text{QCD}}^2)}, \quad (1.4)$$

where n_f is the number of flavours whose mass is smaller than the considered energy scale, and the parameter $\Lambda_{\text{QCD}} \approx 200 \text{ MeV}$, called *Landau Pole*, sets the energy scale at which the coupling would diverge. At energies close to this threshold, it is necessary to take into account non-perturbative effects. The value of Λ_{QCD} is not predicted by the theory and has to be determined experimentally, from measurements of $\alpha_s(Q^2)$. The behaviour of the strong coupling constant is illustrated in Figure 1.2.

Colour confinement and **asymptotic freedom** are two properties of QCD emerging in the low- and high-energy limits of the theory, respectively.

At large momentum transfers ($Q^2 \rightarrow \infty$), the strong coupling $\alpha_s(Q^2)$ logarithmically approaches zero. This behaviour, called asymptotic freedom, implies that quarks and gluons interact weakly at short distances or high energies, allowing them to be treated as nearly free particles in high-energy processes. Colour confinement refers to the observed fact that quarks and gluons are never detected as isolated particles, but only as constituents of colour-neutral bound states such as mesons and baryons. Although colour confinement cannot be proved analytically from QCD, the phenomenon is strongly supported by experimental evidence. One strong theoretical indication supporting confinement is the behaviour of $\alpha_s(Q^2)$, diverging at low-energy scales.

These two properties have important implications in hadron-hadron collisions, such as those at the LHC: when high-energy interactions produce energetic quarks and gluons, they begin to separate. As the distance between them increases, the energy stored in the colour field grows until it becomes energetically favourable to produce quark-antiquark pairs. This pro-




Figure 1.2 – Summary of measurements of $\alpha_S(Q^2)$ [11].

cess, called *fragmentation* or *parton shower*, continues, converting part of the initial energy into quarks' masses and kinetic energies, until, eventually, many colourless quarks-aggregates (i.e., the hadrons) are formed. This phase is called *hadronisation*. As a result, quark and gluons can be observed experimentally as collimated bunches of hadrons, called *jets*. A notable exception to this behaviour is the heaviest quark, the top quark. Due to its extremely short lifetime, it decays (typically into a bottom quark and a W boson) before hadronisation can occur, allowing for direct observation of its decay products and properties.

Electroweak interactions

Electroweak interactions are described by a gauge theory based on the non-Abelian symmetry group $SU(2)_L \otimes U(1)_Y$, which provides a unified description of the electromagnetic and weak forces.

The $SU(2)_L$ group has three generators, T^i (with $i = 1, 2, 3$), which are the *weak isospin* operators. The subscript L in the group's name denotes *left-handed* (LH) fermions, indicating that only the left chiral component $\psi_L = \frac{1}{2}(1 - \gamma^5)\psi$ of quark and lepton fields couples with the three gauge fields, $W_\mu^i(x)$ (with $i = 1, 2, 3$), associated with the $SU(2)_L$ group. LH fermions are doublets under $SU(2)_L$, thus they transform non-trivially under the symmetry group and have weak isospin $T = 1/2$. The two components of each doublet have weak isospin projections $T^3 = +1/2$ (upper component) and $T^3 = -1/2$ (lower component). In each LH lepton doublet, neutrinos constitute the upper component of the doublet, paired with the corresponding

charged leptons. The original formulation of the SM does not include RH neutrinos. In the quark sector, the upper components of the doublets are the *up-type quarks* (u, c, t), while the lower components are the *down-type quarks* (d, s, b), all represented by their LH chiral projections.

In contrast, the right chiral components of fermion fields, $\psi_R = \frac{1}{2}(1 + \gamma^5)\psi(x)$, are singlets under $SU(2)_L$: they do not interact with the W^i_μ fields and are unaffected by the weak isospin transformations ($T^i\psi_R = 0$).

The full arrangement of quarks and leptons into LH doublets and RH singlets for each generation is summarized in Table 1.1.

The second part of the electroweak symmetry group, $U(1)_Y$, is associated with the generator Y ,

	Leptons	Quarks
Left-handed	$\begin{pmatrix} \nu_{e,L} \\ e_L \end{pmatrix}, \begin{pmatrix} \nu_{\mu,L} \\ \mu_L \end{pmatrix}, \begin{pmatrix} \nu_{\tau,L} \\ \tau_L \end{pmatrix}$	$\begin{pmatrix} u_L \\ d_L \end{pmatrix}, \begin{pmatrix} c_L \\ s_L \end{pmatrix}, \begin{pmatrix} t_L \\ b_L \end{pmatrix}$
Right-handed	e_R, μ_R, τ_R	$u_R, c_R, t_R, d_R, s_R, b_R$

Table 1.1 – Doublet and singlet structure for the left- and right-handed components of the fermion fields under the $SU(2)_L \otimes U(1)_Y$ symmetry group.

called *weak hypercharge* Y , which affects differently the LH and the RH components of fermion fields. Both components of the $SU(2)_L$ doublets representing the LH fermions share the same weak hypercharge. This ensures that their transformation properties under $U(1)_Y$ maintain the $SU(2)_L$ structure.

The electric charge Q of a particle is related to its weak isospin and hypercharge via the *Gell-Mann-Nishijima Relation*:

$$Q = T^3 + \frac{Y}{2}. \quad (1.5)$$

The gauge-invariant electroweak Lagrangian is given by:

$$\mathcal{L} = i\bar{\psi}_L \gamma^\mu D_\mu \psi_L + i\bar{\psi}_R \gamma^\mu D_\mu \psi_R - \frac{1}{4} \sum_{i=1}^3 W^{\mu\nu,i} W_{\mu\nu}^i - \frac{1}{4} B^{\mu\nu} B_{\mu\nu}. \quad (1.6)$$

The electroweak covariant derivative D_μ is:

$$D_\mu = \partial_\mu + ig_1 \sum_{i=1}^3 W_\mu^i(x) T^i + ig_2 B_\mu Y, \quad (1.7)$$

where g_1 and g_2 are the coupling constants associated to the groups $SU(2)_L$ and $U(1)_Y$, respectively. The last two terms in Equation 1.6 are the kinetic terms for the gauge fields associated with the $SU(2)_L$ and $U(1)_Y$ generators, $W^{\mu,i}$ and B^μ , respectively. These terms include the field strengths:

$$W^{\mu\nu,i} = \partial^\mu W^{\nu,i} - \partial^\nu W^{\mu,i} - g \sum_{j,k=1}^3 \epsilon^{ijk} W^{\mu,j} W^{\nu,k}, \quad (1.8)$$

and

$$B^{\mu\nu} = \partial^\mu B^\nu - \partial^\nu B^\mu. \quad (1.9)$$

The structure constants of the $SU(2)_L$ symmetry group appearing in Equation 1.8 are the Levi-Civita tensor components ϵ^{ijk} . The field strength $B^{\mu\nu}$ does not contain the third term describing the gauge field's self-interactions, due to the Abelian nature of the $U(1)_Y$ group.

The physical $W^\pm$ gauge bosons can be defined as linear combinations of the gauge fields $W_\mu^{1,2}$:

$$W_\mu^\pm(x) = \frac{1}{\sqrt{2}} \left[W_\mu^1(x) \mp W_\mu^2(x) \right]. \quad (1.10)$$

The $W^\pm$ bosons mediate the weak *charged current* interactions, which transform one component of a weak isospin doublet into the other, enabling transitions such as beta decay.

The electromagnetic interaction is recovered by performing a rotation in the two-dimensional space spanned by the two remaining gauge fields of the $SU(2)_L \otimes U(1)_Y$ group, W_μ^3 and B_μ :

$$\begin{pmatrix} W_\mu^3 \\ B_\mu \end{pmatrix} = \begin{pmatrix} \cos \theta_W & \sin \theta_W \\ -\sin \theta_W & \cos \theta_W \end{pmatrix} \begin{pmatrix} Z_\mu \\ A_\mu \end{pmatrix}, \quad (1.11)$$

where θ_W is the *weak mixing angle* defined by this transformation. The gauge field A_μ can be identified with the photon. This implies the following relationship between the electron charge e and the weak charges g_1 and g_2 :

$$e = g_1 \sin \theta_W = g_2 \cos \theta_W. \quad (1.12)$$

The Z boson (Z_μ) mediates weak *neutral current* interactions, which involve the exchange of a Z boson between quarks or leptons of the same flavour. These interactions do not affect the particle quantum numbers, and the coupling strength differs between LH and RH chiral components.

The $W^\pm$ and Z bosons were discovered in 1983 by the UA1 and UA2 Collaborations [12–15], using the proton-antiproton collision data from the SPS accelerator [16] at CERN. The ratio of the masses of the two bosons was found to be consistent with the SM predictions.

The full electroweak Lagrangian in Equation 1.6 thus provides a unified framework for describing both charged and neutral interactions associated with electromagnetic and weak processes. However, while it successfully predicts the photon as a massless boson, consistent with experimental observations, it does not account for the experimentally observed non-zero masses of the $W^\pm$ and Z bosons. Including in the Lagrangian explicit mass terms for these bosons of the form $m_W^2 W^{\pm,\mu} W_\mu^\mp$ and $\frac{1}{2} m_Z^2 Z^\mu Z_\mu$ would break its gauge invariance. This inconsistency is resolved by incorporating the masses in the theory via the Higgs mechanism, which dynamically generates the masses while preserving the gauge structure, as discussed in the next paragraph.

The Higgs mechanism

As discussed in Section 1.1.2, the gauge principle forbids the presence of explicit mass terms for gauge bosons in the electroweak Lagrangian, while experimental observations show that




Figure 1.3 – Shape of the potential $V(\Phi)$ in two dimensions as a function of the field components $\text{Re}(\Phi)$ and $\text{Im}(\Phi)$. The plot on the left represents the case where $\mu^2 > 0$, while the one on the right shows the case where $\mu^2 < 0$ (the so-called *Mexican Hat* potential).

the $W^\pm$ and Z bosons are massive. The *Higgs mechanism* [17–19] solves this problem maintaining the gauge principle. In this framework, massive gauge bosons acquire mass dynamically, through their interaction with a scalar field $\Phi(x)$, known as the *Higgs field*. The dynamics of the Higgs field leads to physical states which do not reflect the underlying symmetry of the theory: this phenomenon is called *spontaneous symmetry breaking*. In the SM, the spontaneous breaking of the electroweak gauge group $\text{SU}(2)_L \otimes \text{U}(1)_Y$ introduces mass terms for the $W_\mu^\pm$ and Z_μ vector fields, while leaving the photon massless and maintaining the gauge invariance of the electroweak Lagrangian.

The Higgs field is a scalar particle which is a weak isospin doublet ($T = \frac{1}{2}$) with hypercharge $Y = 1$:

$$\Phi(x) = \begin{pmatrix} \Phi^+(x) \\ \Phi^0(x) \end{pmatrix}. \quad (1.13)$$

The Lagrangian describing the dynamics of this field via a potential $V(\phi)$ is:

$$\mathcal{L}_{\text{Higgs}} = (D^\mu \Phi)^\dagger (D_\mu \Phi) - V(\Phi), \quad (1.14)$$

where D_μ is the covariant derivative of the electroweak interactions defined in Equation 1.7 and $V(\Phi)$ is the Higgs potential, defined as:

$$V(\Phi) = \mu^2 \Phi^\dagger \Phi + \lambda (\Phi^\dagger \Phi)^2. \quad (1.15)$$

The shape of the potential $V(\Phi)$ is a fundamental feature of the Higgs mechanism. While $\lambda > 0$ in Equation 1.15 guarantees the presence of global minima in $V(\phi)$, the sign of μ^2 discriminates between a unique minimum (if $\mu^2 > 0$) or a degeneracy of equivalent minima (if $\mu^2 < 0$). Figure 1.3 shows the shape taken by the potential $V(\Phi)$ in case of positive or negative μ^2 parameter. In case of $\mu^2 < 0$, there are infinite configurations of the scalar field that minimise the potential, all satisfying:

$$\Phi^\dagger \Phi = \frac{-\mu^2}{2\lambda}. \quad (1.16)$$

After choosing one of these configurations as ground (*vacuum*) state, the Lagrangian in Equation 1.14 remains gauge invariant, but the electroweak symmetry $\text{SU}(2)_L \otimes \text{U}(1)_Y$ is *sponta-*

neously broken.

The canonical choice for the potential vacuum value is given by:

$$\Phi_v = \begin{pmatrix} 0 \\ v/\sqrt{2} \end{pmatrix}, \quad (1.17)$$

with $v = \sqrt{-\mu^2/2\lambda}$. The term v is referred to as *vacuum expectation value* (VEV). After spontaneous symmetry breaking, the weak isospin doublet can be parametrised as:

$$\Phi(x) = \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix}, \quad (1.18)$$

where $H(x)$ is the Higgs field. The mass terms for the weak bosons $W^\pm$ and Z emerge when rewriting the kinetic term of $\mathcal{L}_{\text{Higgs}}$ with the parametrisation presented in Equation 1.18:

$$\begin{aligned} (D^\mu \Phi)^\dagger (D_\mu \Phi) = & \frac{1}{2} (\partial_\mu H)^2 \\ & + (v + H)^2 \left(\frac{g_1^2}{4} W^{+\mu} W_\mu^+ + \frac{g_1^2}{4} W^{+\mu} W_\mu^+ + \frac{g_1^2}{8 \cos^2 \theta_W} Z^\mu Z_\mu \right), \end{aligned} \quad (1.19)$$

where the terms proportional to v^2 correspond to the mass terms of the $W^\pm$ and Z bosons, while the terms proportional to vH and to H^2 represent the interaction terms between the Higgs and the weak bosons. The photon field A_μ remains massless, in agreement with experimental observations. The mass terms for the weak bosons take the form:

$$m_W = \frac{1}{2}vg, \quad m_Z = \frac{1}{2}v\sqrt{g_1^2 + g_2^2} = \frac{gv}{2 \cos \theta_W} = \frac{m_W}{\cos \theta_W}, \quad (1.20)$$

where m_W is the mass of the $W^\pm$ bosons, and m_Z is the mass of the Z boson.

By substituting the parametrisation of the Higgs field from Equation 1.18 into the Higgs potential of Equation 1.15, and neglecting the constant term, the potential takes the form:

$$V(\Phi) = \frac{1}{2}(-2\mu^2)H^2 + \lambda v H^3 + \frac{\lambda}{4}H^4. \quad (1.21)$$

The term proportional to H^2 indicates that the oscillations of the Higgs field correspond to a massive scalar particle, the *Higgs boson*, with mass:

$$m_H = \sqrt{-2\mu^2}. \quad (1.22)$$

The cubic and quartic terms describe the Higgs boson self-interactions.

The VEV can be measured via muon decay, yielding $v \approx 246$ GeV [20]. Once the VEV and the Higgs boson mass are known, the Higgs boson SM potential is fully determined.

The chiral nature of the electroweak gauge group also forbids explicit mass terms for fermions.

A Dirac mass term of the form $-m\bar{\psi}\psi = -m(\bar{\psi}_L\psi_R + \bar{\psi}_R\psi_L)$ is not invariant under $\text{SU}(2)_L \otimes \text{U}(1)_Y$,

since LH and RH components of fermion fields transform differently under the gauge group. However, experimental data confirm that all fermions are massive. The resolution, once again, comes from the interaction with the Higgs field and the mechanism of spontaneous symmetry breaking.

Gauge-invariant mass terms for fermions can be generated through *Yukawa interactions* between the fermion fields and the Higgs doublet $\Phi(x)$. Focusing on the first generation of quarks, the relevant Yukawa terms in the Lagrangian are:

$$\mathcal{L}_{\text{Yukawa}} = y_u (\bar{u}_L \bar{d}_L) \cdot \Phi^* u_R + y_d (\bar{u}_L \bar{d}_L) \cdot \Phi d_R + \text{h.c.}, \quad (1.23)$$

where y_u and y_d are the Yukawa coupling constants for the up-type and down-type quarks, respectively, and Φ^* denotes the complex conjugate of the Higgs doublet. After spontaneous symmetry breaking, and substituting the parametrisation of the Higgs field from Equation 1.18, the Yukawa Lagrangian becomes:

$$\begin{aligned} \mathcal{L}^{\text{Yukawa}} = & -\frac{y_u v}{\sqrt{2}} \bar{u}u - \frac{y_d v}{\sqrt{2}} \bar{d}d \\ & -\frac{y_u}{\sqrt{2}} \bar{u}uH - \frac{y_d}{\sqrt{2}} \bar{d}dH. \end{aligned} \quad (1.24)$$

The terms in the first line represent the fermion mass terms, showing that quarks acquire masses:

$$m_u \equiv \frac{y_u v}{\sqrt{2}} \quad \text{and} \quad m_d \equiv \frac{y_d v}{\sqrt{2}}, \quad (1.25)$$

for up-type and down-type quarks, respectively. The second line of Equation 1.24 describes interactions between the Higgs boson and the massive quarks, with couplings proportional to the corresponding fermion masses.

A strong experimental evidence of the Higgs mechanism was provided by the discovery of a new particle compatible with the SM Higgs boson by the ATLAS and CMS collaborations in July 2012 [21, 22].

1.2 Supersymmetry

The analysis presented in this thesis is a search for the production of top squarks, BSM particles predicted by the Supersymmetry theory. In this section, the main concepts of the Supersymmetry theory are introduced, along with the motivations for its introduction in relation to the limits of the SM theory. The phenomenology of the top squark is then described, with a focus on the production and decay processes relevant to the analysis.

1.2.1 Limitations of the Standard Model

The Standard Model (SM) has demonstrated remarkable success, with its predictions confirmed by a wide range of experimental results to high precision. However, it leaves several fundamental theoretical and experimental questions unanswered, which indicate that it may be an incomplete theory, and that a more general framework is required to achieve a unified and consistent

description of particle physics phenomenology. Among the many BSM theories proposed to address these limitations, Supersymmetry (SUSY) stands out as one of the most promising candidates. SUSY offers a coherent extension of the SM that addresses multiple open issues simultaneously. SUSY introduces a new fundamental symmetry that relates fermions and bosons: for every known particle, it predicts the existence of a superpartner differing in spin by $1/2$. This symmetry is generated by a fermionic operator Q , which acts on the quantum states as follows:

$$Q |\text{boson}\rangle = |\text{fermion}\rangle \quad Q |\text{fermion}\rangle = |\text{boson}\rangle. \quad (1.26)$$

Through this structure, SUSY provides natural solutions to several limitations of the SM. Some key examples are summarised below.

Dark Matter

Astronomical and cosmological observations indicate that only about 5% of the universe is composed of the ordinary matter described by the SM. The remaining 96% consists of *dark matter* (approximately 27%), inferred from its gravitational effects on visible matter and cosmic structures, and *dark energy* (about 68%), responsible for the accelerated expansion of the universe [23]. The SM does not contain any stable and weakly interacting particles that can be good dark matter candidates, as SM neutrinos are too light to be non-relativistic and form large galactic structures in the early universe. As a result, the SM theory cannot explain the dominant dark fraction of the universe's matter content. Since dark matter does not emit or absorb electromagnetic radiation, and is therefore invisible to direct astronomical observation, a common hypothesis is the existence of one or more unobserved stable particles, with very high mass, and interacting only via gravity or the weak force. However, the nature of such particles and the interactions related to their phenomenology is unknown. One of the most theoretically promising candidates is the *Weakly Interacting Massive Particle* (WIMP), a neutral particle with weak-scale mass and weak interactions, whose thermal relic density may naturally reproduce the observed dark matter abundance (**WIMP miracle**). SUSY provides a natural WIMP candidate: under the assumption of R -parity conservation (see Equation 1.37 in Section 1.2.2), the lightest supersymmetric particle (LSP) is stable, electrically neutral, and weakly interacting, thus matching the properties expected of dark matter.

Hierarchy problem

One of the major theoretical limitations of the SM is the large difference between the electroweak scale, characterised by the W boson mass ($m_W \approx 100, \text{GeV}$), and the Planck scale ($M_P \sim 10^{19}, \text{GeV}$), where gravitational effects become significant. This large discrepancy, also reflected in the relative strengths of the weak and gravitational forces, is considered a sign of the SM's unnaturalness. A manifestation of this issue appears in the so-called *hierarchy problem*, which regards the sensitivity of the Higgs boson mass to quantum corrections. As discussed in Section 1.1.2, the tree-level mass of the Higgs boson, $m_{H,0}$ (from Equation 1.22), is a free parameter of the SM, only depending on the shape of the Higgs potential and on the VEV of the Higgs field. However, quantum corrections introduce large contributions to the physical Higgs mass (m_H) through virtual loop effects involving all the SM particles. These corrections are quadratically divergent with Λ , which is the cut-off energy scale beyond which the SM stops to be a



Figure 1.4 – One-loop quantum corrections to the Higgs squared mass parameter m_H^2 , due to a Dirac fermion f (left), and a scalar S (right) [24].

valid theory, and can be expressed as:

$$m_H^2 = m_{H,0}^2 + \delta m_H^2 \quad \text{with} \quad \delta m_H^2 \propto \Lambda^2, \quad (1.27)$$

where δm_H^2 represents the radiative corrections. If the SM remained valid up to the Planck scale ($\Lambda = M_P$), to have the measured Higgs mass $m_H \approx 125 \text{ GeV}$, the radiative corrections to the Higgs mass should be extremely large, requiring a fine-tuned cancellation between $m_{H,0}^2$ and δm_H^2 of the order of $\frac{m_H^2}{M_P^2} = \frac{(10^2)^2}{(10^{18})^2}$. This level of fine-tuning is considered highly unnatural, and strongly suggests the presence of new physics at intermediate energy scale, which would help stabilise the Higgs mass. To obtain a more natural theory, it is necessary to introduce a cancellation mechanism between the divergent contributions to δm_H^2 .

In particular, fermions and bosons give two opposite-sign contributions to δm_H^2 . Specifically, considering a generic fermion f and a complex scalar field S that couple to the Higgs, the leading one-loop quantum corrections to δm_H^2 , illustrated in Figure 1.4, are:

$$\begin{aligned} \delta m_H^2 &\propto -\frac{|\lambda_f|^2}{8\pi^2} \Lambda^2 + \dots \quad (\text{fermion loop}), \\ \delta m_H^2 &\propto +\frac{|\lambda_S|^2}{16\pi^2} \Lambda^2 + \dots \quad (\text{boson loop}), \end{aligned} \quad (1.28)$$

where λ_f and λ_S are the couplings of the fermion and the scalar to the Higgs field, respectively, and the omitted terms grow at most logarithmically with Λ . If each of the SM fermions had a bosonic partner consisting in two complex scalars with $|\lambda_S| \approx |\lambda_f|^2$, their contributions to Λ^2 would cancel, eliminating the quadratic divergence. This is fundamental idea behind of SUSY, where coupling constants and multiplicities of the superpartners are arranged to ensure that the divergent corrections to the Higgs mass cancel exactly, giving a natural solution to the hierarchy problem.

Grand Unification Theory

The Grand Unification Theories (GUTs) are a class of theoretical models that aim to unify the three fundamental gauge interactions of the SM (strong, weak, and electromagnetic) into a single force described by a larger gauge group at very high energy scales. In these models, the differences between the three forces at low energies are the result of spontaneous symmetry




Figure 1.5 – Evolution of the coupling constants in the SM (dashed lines) and in supersymmetric models (solid lines) [24]. The masses of the SUSY particles are treated as a common threshold and varied between 750 GeV and 2.5 TeV.

breaking, and that they originate from a common interaction in the early universe. One of the primary motivations for GUTs is to reduce the large number of free parameters in the SM, such as particle masses and coupling constants, which do not emerge naturally from the theory and have to be measured experimentally. The energy dependence of the SM gauge couplings is described by the Renormalisation Group Equations. At one-loop level, these equations can be written as:

$$\frac{d\alpha_i^{-1}(t)}{dt} = -\frac{b_i}{2\pi}, \quad (1.29)$$

where α_i are the SM gauge couplings, with index $i = 1, 2, 3$ corresponding to the gauge groups $U(1)_Y$, $SU(2)_L$, and $SU(3)_C$, respectively, and $t = \ln(Q/Q_0)$ is the logarithm of the energy scale Q relative to a reference scale Q_0 . The coefficients b_i depend on the particle content of the theory and, for the SM, they take the values $(b_1, b_2, b_3) = \left(\frac{41}{10}, -\frac{19}{6}, -7\right)$. These values lead to a running of the couplings α_i such that they do not converge at high values of Q , thus preventing exact unification within the SM.

In SUSY (specifically, in the Minimal Supersymmetric Standard Model and assuming SUSY breaking at the TeV scale, see Section 1.2.2), due to the larger number of particles in the theory, the $\{b_i\}$ coefficients change to $(b_1, b_2, b_3) = \left(\frac{33}{5}, 1, -3\right)$ above the SUSY-breaking scale. As a result, the three gauge couplings evolve differently with energy and converge at a single point at an energy scale $Q \approx 10^{16}$ GeV, as shown in Figure 1.5, making unification possible.

Anomalous muon magnetic moment

The magnetic moment μ of an elementary particle with spin S , electric charge q , and mass m , is defined as:

$$\mu = g \frac{q}{2m} S, \quad (1.30)$$

where g is the gyromagnetic factor, a dimensionless constant that characterises the strength of the particle's magnetic interaction. According to quantum mechanics, a point-like spin- $\frac{1}{2}$ particle is expected to have a gyromagnetic factor $g = 2$. However, in quantum field theory, this value is modified by radiative quantum corrections, due to virtual particles in loop diagrams. The deviation from the classical value for a lepton ℓ is quantified by the *anomalous magnetic moment* $a_\ell = (g_\ell - 2)/2$. For the muon, the anomalous magnetic moment a_μ is particularly sensitive to quantum corrections due to its relatively large mass compared to the electron, which increases the sensitivity to potential contributions from new heavy particles. The most recent world average for the experimental measurements of the muon's anomalous magnetic moment is:

$$a_\mu^{exp} = 1165920715(145) \cdot 10^{-12}, \quad (1.31)$$

measured by the E821 experiment at Brookhaven and more recently by the Muon $g-2$ experiment at Fermilab [20, 25]. This result shows a discrepancy about 5 standard deviations compared to the Standard Model prediction. If this deviation is confirmed with improved theoretical and experimental precision, it could be explained by the presence of BSM massive particles acting in the quantum correction loops.

In SUSY, additional loop contributions from superpartners can increase the theoretical prediction of a_μ , and bring it into agreement with the experimental value.

1.2.2 Minimal Supersymmetric Standard Model

The model considered in this thesis is the *Minimal Supersymmetric Standard Model* (MSSM), which is the simplest supersymmetric extension of the SM, containing the minimal number of particles consistent with phenomenology. The supersymmetric generator Q , introduced earlier, is an operator that transforms bosonic states into fermionic states and vice versa, and it satisfies the following algebra:

$$\begin{aligned} \{Q, Q^\dagger\} &= P^\mu, \\ \{Q, Q\} &= \{Q^\dagger, Q^\dagger\} = 0, \\ [P^\mu, Q] &= [P^\mu, Q^\dagger] = 0 \end{aligned} \quad (1.32)$$

where P^μ is the four-momentum operator generating spacetime translations. The single-particle states of a supersymmetric theory must be irreducible representations of this algebra, and are called *supermultiplets*. Each supermultiplet contains a boson and a fermion, which are *superpartners* of each other. The supersymmetry generators commute with the generators of gauge transformations, thus superpartners must belong to the same representation of the gauge group, having the same quantum numbers, such as electric charge, weak isospin, and colour charge. Additionally, they share the number of bosonic and fermionic degrees of freedom. The two simplest types of supermultiplets that can be built starting from this principle are *chiral supermultiplets* and *gauge supermultiplets*. Chiral supermultiplets consist of a Weyl fermion (with two fermionic degrees of freedom corresponding to its helicity states), and two real scalars, that can be grouped into a single complex scalar field. Gauge supermultiplets contain a massless vector boson (with two helicity states, and thus two bosonic degrees of freedom), and a massless spin- $\frac{1}{2}$ Weyl fermion (*gaugino*). Other combinations of particles satisfying the requirement on the num-

ber of bosonic and fermionic degrees of freedom are possible, but they are always reducible to combinations of chiral and gauge supermultiplets, that are thus the only ones considered when building the particle content of the MSSM.

MSSM particle content

Only **chiral supermultiplets** can contain fermions whose LH and RH components transform differently under the gauge group, which is the case for all SM fermions (quarks and leptons). In the SM, the LH and RH components of the SM fermions are distinct two-component Weyl fermions with different gauge transformation properties. Therefore, each component must have its own complex scalar superpartner with spin 0: these are called *squarks* ($\tilde{q}_R, \tilde{q}_L$), and *sleptons* ($\tilde{\ell}_R, \tilde{\ell}_L$). The subscripts L and R of the sfermions refer to the helicity of their SM partners, since they are spin-0 scalars and therefore have no helicity. Since SM neutrinos are only LH, their scalar superpartners, the *sneutrinos*, are generally indicated simply as $\tilde{\nu}_\ell$, without any L or R subscript. The SM Higgs boson is a complex scalar, so it naturally belongs to chiral supermultiplet in the MSSM. To avoid gauge anomalies [24], the MSSM's Higgs sector is extended to contain two chiral supermultiplets with two complex $SU(2)_L$ -doublets, H_u and H_d , with hypercharge $Y = +1/2$ and $Y = -1/2$, respectively. Each Higgs doublet has a fermionic superpartner called *higgsino* ($\tilde{H}_u$ for H_u , and $\tilde{H}_d$ for H_d). Each $SU(2)_L$ doublet consists of a charged and a neutral component (H_u^+, H_u^0 for H_u , and H_d^0, H_d^- for H_d), and the observed SM Higgs boson is a linear combination of the neutral components of the two doublets. The weak-isospin components of the higgsinos are organized in the same doublet structure.

All the MSSM chiral supermultiplets, including quarks and squarks, leptons and sleptons, and Higgs and higgsinos, are listed in Table 1.2.

Names	Superfields	spin 0	spin 1/2	$SU(3)_C, SU(2)_L, U(1)_Y$
quarks and squarks ($\times 3$ generations)	Q	$(\tilde{u}_L \tilde{d}_L)$	$(u_L d_L)$	$(\mathbf{3}, \mathbf{2}, \frac{1}{6})$
	$\bar{u}$	$\tilde{u}_R$	u_R	$(\bar{\mathbf{3}}, \mathbf{1}, -\frac{2}{3})$
	$\bar{d}$	$\tilde{d}_R$	d_R	$(\bar{\mathbf{3}}, \mathbf{1}, \frac{1}{3})$
leptons and sleptons ($\times 3$ generations)	L	$(\tilde{\nu} \tilde{e}_L)$	(νe_L)	$(\mathbf{1}, \mathbf{2}, -\frac{1}{2})$
	$\bar{e}$	$\tilde{e}_R$	e_R	$(\mathbf{1}, \mathbf{1}, 1)$
Higgs and higgsinos	H_u	$(H_u^+ H_u^0)$	$(\tilde{H}_u^+ \tilde{H}_u^0)$	$(\mathbf{1}, \mathbf{2}, \frac{1}{2})$
	H_d	$(H_d^0 H_d^-)$	$(\tilde{H}_d^0 \tilde{H}_d^-)$	$(\mathbf{1}, \mathbf{2}, -\frac{1}{2})$

Table 1.2 – Chiral supermultiplets in the MSSM.

Gauge supermultiplets are constituted by the SM gauge bosons and their superpartners, the *gauginos*. The superpartners of the gluons are the *gluinos* ($\tilde{g}$). The superpartners of the $SU(2)_L$ gauge bosons ($W^+, W^-,$ and W^0) are the *winos* ($\tilde{W}^+, \tilde{W}^-, \tilde{W}^0$), and the superpartner of the

$U(1)_Y$ gauge boson (B^0) is the *bingo* ($\tilde{B}^0$). After EWSB, the neutral gauginos ($\tilde{W}^0$ and $\tilde{B}^0$) mix to form the zino ($\tilde{Z}^0$) and the photino ($\tilde{\gamma}^0$).

The gauge supermultiplets of the MSSM are listed in Table 1.3.

Names	spin 1	spin 1/2	$SU(3)_C, SU(2)_L, U(1)_Y$
gluons and gluinos	g	$\tilde{g}$	$(\mathbf{8}, \mathbf{1}, 0)$
W bosons and winos	$W^\pm, W^0$	$\tilde{W}^\pm, \tilde{W}^0$	$(\mathbf{1}, \mathbf{3}, 0)$
B boson and bino	B^0	$\tilde{B}^0$	$(\mathbf{1}, \mathbf{1}, 0)$

Table 1.3 – Gauge supermultiplets in the MSSM.

MSSM Lagrangian

The gauge interactions of the sfermion fields are the same as those of their SM fermionic partners, since they have the same gauge quantum numbers (e.g. $\tilde{u}_L$ and $\tilde{d}_L$ couple with the W boson, while $\tilde{u}_R$ and $\tilde{d}_R$ do not). The Lagrangian describing the gauge interactions of the MSSM is similar to the one of the SM, but extended to be SUSY invariant. It consists of the kinetic terms of the chiral and the gauge supermultiplets, where the ordinary derivative ∂_μ is replaced by the SM covariant derivative. This introduces the gauge interactions between the gauge bosons and the scalar (ϕ_i) and fermionic (ψ_i) components of the chiral supermultiplets, but spoils the SUSY invariance. To restore it, the Lagrangian must also include interaction terms between the chiral supermultiplets and the gauginos λ_a . Considering for simplicity only the part of the covariant derivative related to the W bosons, $D_\mu = \partial_\mu + igW_\mu^a T^a$, the gauge-interaction Lagrangian takes the form:

$$\begin{aligned} \mathcal{L}_{\text{gauge}} = & D^\mu \phi^{*i} D_\mu \phi_i + i\psi^{\dagger i} \gamma^\mu D_\mu \psi_i \\ & - \sqrt{2}g \left(\phi^{*i} T^a \phi_i \right) \lambda^a - \sqrt{2}g\lambda^{\dagger a} \left(\psi^{\dagger i} T^a \phi_i \right) - \frac{1}{2}g^2 \sum_a \left(\phi^{*i} T^a \phi_i \right)^2, \end{aligned} \quad (1.33)$$

where the first line contains the gauge-covariant kinetic terms, and the second line represents the manually added interaction terms with gauge superpartners.

In addition to gauge interactions, the MSSM also includes Yukawa-like interactions, which originate from a holomorphic function known as the *superpotential*:

$$W_{\text{MSSM}} = \bar{u} \mathbf{y}_u Q H_u - \bar{d} \mathbf{y}_d Q H_d - \bar{e} \mathbf{y}_e L H_d + \mu H_u H_d, \quad (1.34)$$

where $Q, L, \bar{u}, \bar{d}, \bar{e}, H_u$, and H_d are chiral superfields corresponding to the supermultiplets in Table 1.2. The quantity μ has the dimension of a mass and thus provides a supersymmetric contribution to the mass of the Higgs bosons. The matrices $\mathbf{y}_u, \mathbf{y}_d$, and $\mathbf{y}_e$ represent the dimensionless Yukawa coupling constants, each represented as a 3×3 matrix in the flavour space of the quark and lepton families.




Figure 1.6 – Proton decay $p \rightarrow e^+ \pi^0$ through λ' through λ'' R-parity couplings [24].

The superpotential W_{MSSM} contributes to the MSSM Lagrangian through the following terms:

$$\mathcal{L}_{\text{Yukawa}} = -\frac{1}{2} \left(W^{ij} \psi_i \psi_j + W_{ij}^* \psi^{\dagger i} \psi^{\dagger j} \right) + W^i W_i^*, \quad (1.35)$$

where ψ_i denote again the fermionic components of the chiral superfields, and W^i and W^{ij} are the partial derivatives of the superpotential with respect to the scalar components ϕ , defined as:

$$W^i = \frac{\partial W}{\partial \phi^i}, \quad W^{ij} = \frac{\partial^2 W}{\partial \phi^i \partial \phi^j}. \quad (1.36)$$

Differently from the SM, the MSSM superpotential can contain terms that violate the baryon and lepton number conservation. The violation of the baryonic and leptonic number leads potentially to proton decay, as shown in Figure 1.6, but this scenario is completely disfavoured by the experiment, with a measured limit on the decay time of the proton in a lepton-meson final state of more than 10^{34} years [26]. This issue can be solved by introducing an additional quantum number, the *R-parity*, that is assumed to be conserved (*R-parity conservation*, RPC):

$$R = (-1)^{3(B-L)+2S}, \quad (1.37)$$

where B and L are the baryon and lepton numbers, respectively, and S is the spin. It is a multiplicatively conserved quantum number, which takes the value $+1$ for all SM particles and -1 for all their supersymmetric partners. When assuming *R-parity conservation*, the supersymmetric particles can only be produced in pairs. This implies that the LSP is stable and does not decay, making it a good candidate for dark matter.

Soft SUSY breaking

If SUSY were an exact symmetry, fermions and their bosonic superpartners would be mass-degenerate. Since no supersymmetric particles have been observed so far, if SUSY is realised, it has to be a broken symmetry. The exact mechanism of SUSY breaking is unknown, therefore it is commonly modelled by adding to the Lagrangian the most general set of *soft-breaking terms*. These additional soft terms are chosen to not introduce quadratic divergences in the Lagrangian, thus preserving the capability of the SUSY theory to address the hierarchy problem. The corresponding Lagrangian, denoted as $\mathcal{L}_{\text{soft}}$, can be added to the previously defined MSSM Lagrangian, already including $\mathcal{L}_{\text{gauge}}$ and $\mathcal{L}_{\text{Yukawa}}$ (see above). The Lagrangian $\mathcal{L}_{\text{soft}}$ violates SUSY and contains all allowed terms that do not introduce quadratic divergences in the theory

and that are compatible with the SM $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$ gauge symmetry. The possible forms of the soft terms of $\mathcal{L}_{\text{soft}}$ for the MSSM are [27]:

- soft gaugino masses for each gauge group, with index $a = 1, 2, 3$ denoting $U(1)_Y$, $SU(2)_L$, and $SU(3)_C$, respectively: $\frac{1}{2}M_a\lambda^a\lambda^a$;
- soft scalar squared-mass terms for all the chiral supermultiplets: $(m^2)_i^j\phi^j\phi_i^*$;
- soft bilinear scalar interactions: $\frac{1}{2}b^{ij}\phi_i\phi_j$;
- soft trilinear scalar interactions: $\frac{1}{3!}a^{ijk}\phi_i\phi_j\phi_k$.

The most general $\mathcal{L}_{\text{soft}}$ for the MSSM can be written as:

$$\begin{aligned}\mathcal{L}_{\text{soft}} = & -\frac{1}{2}(M_1\tilde{B}\tilde{B} + M_2\tilde{W}\tilde{W} + M_3\tilde{g}\tilde{g}) + \text{h.c.} \\ & -\tilde{Q}^\dagger \mathbf{m}_{\tilde{Q}}^2 \tilde{Q} - \tilde{L}^\dagger \mathbf{m}_{\tilde{L}}^2 \tilde{L} - \tilde{u}^\dagger \mathbf{m}_{\tilde{u}}^2 \tilde{u} - \tilde{d}^\dagger \mathbf{m}_{\tilde{d}}^2 \tilde{d} - \tilde{e}^\dagger \mathbf{m}_{\tilde{e}}^2 \tilde{e} \\ & -m_{H_u}^2 H_u^* H_u - m_{H_d}^2 H_d^* H_d - (bH_u H_d + \text{h.c.}) \\ & - \left(\tilde{u} \mathbf{a}_u \tilde{Q} \cdot H_u - \tilde{d} \mathbf{a}_d \tilde{Q} \cdot H_d - \tilde{e} \mathbf{a}_e \tilde{L} \cdot H_d \right) + \text{h.c.},\end{aligned}\tag{1.38}$$

where $\tilde{Q}$, $\tilde{L}$, $\tilde{u}$, $\tilde{d}$, $\tilde{e}$, H_u , and H_d are the scalar components of the chiral superfields corresponding to the supermultiplets in Table 1.2, and summation over generations is implied. The first line contains the soft gaugino masses, the second and the third line contain soft scalar squared-mass terms and the soft bilinear scalar interactions term ($bH_u H_d$), and the last line contains the soft trilinear scalar interactions terms. The m^2 matrices are 3×3 hermitian matrices in the supermultiplets' families space, while the a matrices are 3×3 matrices in the flavour space. These terms break supersymmetry because they contribute explicitly to the masses and the interactions of the SUSY sparticles but not to their SM superpartners. The underlying supersymmetry breaking is assumed to be spontaneous and take place in a hidden sector.

The Lagrangian $\mathcal{L}_{\text{soft}}$ introduces a large number of additional free parameters defining the MSSM. In addition to the SM parameters, the MSSM contains 105 new physical parameters, including masses, phases and mixing angles, that can not be removed through phase redefinitions or changes in the flavour basis of the quark and lepton supermultiplets. This results in a total of 124 independent physical parameters in the MSSM. However, many of these parameters can be constrained by experimental data and phenomenological considerations. Requiring the suppression of flavour-changing neutral currents, the absence of new CP-violating phases, and the assumption of universality for the first two generations, the number of parameters is reduced to 19, giving a model known as the *phenomenological MSSM* (pMSSM).

After the EWSB and the soft SUSY breaking, the fields with the same quantum numbers generally mix, generating **mass eigenstates**. Mass eigenstates correspond to the physical particles that can be observed experimentally, characterising the SUSY phenomenology, and their understanding is thus fundamental for experimental physics searches.

As previously discussed, each SM fermion has two supersymmetric partners, $\tilde{f}_L$ and $\tilde{f}_R$, corresponding to the two helicity states f_R and f_L , respectively. The states $\tilde{f}_L$ and $\tilde{f}_R$ mix to form the mass eigenstates proportionally to the mass of their corresponding SM fermion. Therefore, the

mixing is expected to be small, except for the third-generation squarks, where the SM fermions masses are significantly larger. For the top squark (the *stop*) the left-right mixing becomes especially important. The stop mass-squared matrix can be approximately written as:

$$\begin{pmatrix} (\mathbf{m}_{\tilde{Q}}^2)_3 + (\frac{1}{2} + \frac{2}{3} \sin^2 \theta_W) M_Z^2 \cos 2\beta + m_t^2 & -m_t(\mathbf{a}_t + \mu \cot \beta) \\ -m_t(\mathbf{a}_t + \mu \cot \beta) & m_{\tilde{t}}^2 + \frac{2}{3} M_Z^2 \cos 2\beta \sin^2 \theta_W m_t^2 \end{pmatrix}, \quad (1.39)$$

where $(\mathbf{m}_{\tilde{Q}}^2)_3$ is the third component of the LH squarks mass matrix from Equation 1.38, $\mathbf{a}_t$ is the matrix of the trilinear interaction coefficients of the top squark, μ is the coefficient in the Higgs super-potential of Equation 1.34 that gives the mass terms for the Higgs bosons and their partners (*Higgs mass term*), $\tan \beta$ is the ratio of the VEVs of the two Higgs doublets ($\tan \beta = v_u/v_d$), θ_W is the weak mixing angle (see Section 1.1.2), and $m_{\tilde{t}}$ is the mass of the stop. Since the off diagonal terms scale with the large top quark mass, of their fermion partners, the mixing between left- and right-handed stop states is substantial. As a result, a significant mixing of the left-right gauge eigenstates can occur and determine a large mass splitting between the two mass eigenstates, $\tilde{t}_1$ and $\tilde{t}_2$, implying that the lightest of the two, $\tilde{t}_1$, might be considerably lighter than the rest of the squarks, with important implications for SUSY phenomenology and experimental searches.

The charged gauginos and higgsinos mix into four physical states, called charginos: $\tilde{\chi}_1^\pm, \tilde{\chi}_2^\pm$, assuming $m(\tilde{\chi}_1^\pm) < m(\tilde{\chi}_2^\pm)$. The mixing is described at tree-level by a 2×2 complex mass matrix:

$$\begin{pmatrix} M_2 & \sqrt{2} M_W \sin \beta \\ \sqrt{2} M_W \cos \beta & \mu \end{pmatrix}, \quad (1.40)$$

where M_2 is the soft gaugino mass term for the winos (see Equation 1.38). In the same way, the neutral gauginos and higgsinos mix into four physical states, called neutralinos: $\tilde{\chi}_1^0, \tilde{\chi}_2^0, \tilde{\chi}_3^0, \tilde{\chi}_4^0$ (index ordered in mass). The mixing is described at tree-level by a 4×4 symmetric complex mass matrix:

$$\begin{pmatrix} M_1 & 0 & -M_W \tan \theta_W \cos \beta & -M_W \tan \theta_W \sin \beta \\ & M_2 & M_W \cos \beta & M_W \sin \beta \\ & & 0 & -\mu \\ & & & 0 \end{pmatrix}, \quad (1.41)$$

where M_1 is the soft gaugino mass term for the bino (see Equation 1.38). In this thesis the $\tilde{\chi}_1^0$ is assumed to be the LSP. Since also R -parity conservation is assumed, the $\tilde{\chi}_1^0$ represents the dark matter candidate.

The mass eigenstates of the MSSM particles are listed in Table 1.4. In general, the three parameters μ, M_1 and M_2 are completely arbitrary. However, in GUTs, M_1 and M_2 are equal at the high energy scale where the gauge couplings are assumed to unify. According to the relative hierarchy between the three parameters μ, M_1 and M_2 , the mixture of higgsinos, winos, and bino contributing to the lightest chargino and neutralino differs. The phenomenology considered for this thesis assumes $M_1 < M_2 \ll |\mu|$. In this scenario, the two lightest neutralinos are mostly pure gaugino states, with masses $m(\tilde{\chi}_1^0) \approx M_1$ (*bino-like*) and $m(\tilde{\chi}_2^0) \approx M_2$, while the two heavier neutralinos are mostly pure higgsino states, with masses $m(\tilde{\chi}_3^0) \approx m(\tilde{\chi}_4^0) \approx |\mu|$.

Names	Spin	P_R	Gauge Eigenstates	Mass Eigenstates
Higgs bosons	0	+1	$H_u^0, H_d^0, H_u^+, H_d^-$	$h^0, H^0, A^0, H^\pm$
Squarks	0	-1	$\tilde{u}_L, \tilde{u}_R, \tilde{d}_L, \tilde{d}_R$	(same)
			$\tilde{s}_L, \tilde{s}_R, \tilde{c}_L, \tilde{c}_R$	(same)
			$\tilde{t}_L, \tilde{t}_R, \tilde{b}_L, \tilde{b}_R$	$\tilde{t}_1, \tilde{t}_2, \tilde{b}_1, \tilde{b}_2$
Sleptons	0	-1	$\tilde{e}_L, \tilde{e}_R, \tilde{\nu}_e$	(same)
			$\tilde{\mu}_L, \tilde{\mu}_R, \tilde{\nu}_\mu$	(same)
			$\tilde{\tau}_L, \tilde{\tau}_R, \tilde{\nu}_\tau$	$\tilde{\tau}_1, \tilde{\tau}_2, \tilde{\nu}_\tau$
Neutralinos	1/2	-1	$\tilde{B}^0, \tilde{W}^0, \tilde{H}_u^0, \tilde{H}_d^0$	$\tilde{\chi}_1^0, \tilde{\chi}_2^0, \tilde{\chi}_3^0, \tilde{\chi}_4^0$
Charginos	1/2	-1	$\tilde{W}^\pm, \tilde{H}_u^\pm, \tilde{H}_d^\pm$	$\tilde{\chi}_1^\pm, \tilde{\chi}_2^\pm$
Gluino	1/2	-1	$\tilde{g}$	(same)

Table 1.4 – MSSM particles, with sfermion mixing for the first two generations assumed to be negligible. In this table gravity (and thus its mediator, the *gravitino*, is not considered).

This hypothesis is not favoured by naturalness arguments, which require μ to be of the same scale as the mass of the Z boson. However, it still allows for a natural definition of the Higgs mass and is characterised by a phenomenology that can be easily detected at the LHC.

1.2.3 Top squark

Top squarks have an important role in the loop corrections to the Higgs boson mass, helping to address the hierarchy problem. They are also expected to be the lightest squarks, making their production potentially accessible at the LHC energies. Therefore, searches for top squarks have always been a central focus of the LHC searches programme.

This section provides an overview of the top squark phenomenology at the LHC, starting with a general introduction on physics at proton-proton (pp) colliders, followed by a more detailed discussion on the motivation behind stop searches at hadron colliders, and a description of the stop phenomenology at the LHC, including production and decay modes. A summary of some of the previous stop search results at the LHC is also presented.

Phenomenology at proton-proton colliders

Hadron accelerators have the advantage over lepton accelerators to overcome the significant energy loss experienced by accelerated leptons due to synchrotron radiation, which scales as:

$$\frac{dE}{dt} \propto \left(\frac{E}{m}\right)^4, \quad (1.42)$$

where E and m are the particle's energy and mass, respectively. Since lighter particles radiate more energy, accelerating electrons to the same energy as protons requires compensating for an energy loss approximately $(m_p/m_e)^4 \approx 10^{12}$ times larger, severely limiting the achievable centre-of-mass energy in lepton colliders.

The main drawback of colliding protons, is that they are not elementary particles: they are composed of partons (quarks and gluons), which carry only a fraction of the total proton's momentum. This produces two categories of collision events: *soft collisions* and *hard collisions*. Soft collisions involve long distance interactions, where the protons behave as elementary particles. In this kind of processes, the transferred momentum is generally low, and particles scatter at small angles. The particles produced in these interactions have high longitudinal momentum, but low transverse momentum. These are the most frequent interactions in hadron colliders and generally constitute a background in hadron collider experiments. Hard collisions, on the other hand, are rarer and contain the processes of physical interest. They are short distance collisions between the individual partons constituting the protons, and they are characterised by a high momentum exchange, so that the production of particles at large angle, high transverse momentum and high mass is possible.

Since the partonic centre-of-mass energy ($\sqrt{\hat{s}}$) depends on the fraction of momentum carried by the colliding gluons and quarks, it is lower than the pp centre-of-mass energy:

$$\sqrt{\hat{s}} = \sqrt{x_a x_b s}, \quad (1.43)$$

where $\sqrt{s}$ is the proton centre-of-mass energy and x_a, x_b are the fractions of the total proton momentum carried by scattering partons. The cross section for a hard-scattering process in pp collisions is:

$$\sigma_{pp \rightarrow X} = \sum_{a,b} \int dx_a dx_b f_a(x_a, Q^2) f_b(x_b, Q^2) \hat{\sigma}_{ab \rightarrow X}(x_a, x_b), \quad (1.44)$$




Figure 1.7 – Pair production cross section from strong-produced SUSY particles [28] at a centre-of-mass energy of 13 TeV (left) and 13.6 TeV (right), which correspond to the LHC pp collisions centre of mass energy in the Run 2 and in the Run 3 data-taking periods, respectively.

where $\hat{\sigma}_{ab \rightarrow X}$ is the cross section of the interaction between the partons a and b , x_a and x_b are the fractions of the proton momentum carried by the partons, and $f_a(x_a, Q^2)$ and $f_b(x_b, Q^2)$ are the Parton Distribution Functions (PDFs), which describe the probability that the partons carry a fraction x_a or x_b of the total proton momentum. Partons within the colliding protons that are not involved in the hard scattering undergo hadronisation or secondary multi-parton interactions, producing a large QCD background called *underlying event*.

Motivation

If SUSY is realised, the particles produced most frequently at the LHC are those with masses within the collider’s accessible energy range and that couple most strongly to the proton constituents. SUSY particles produced in strong processes, namely squarks and gluinos, typically have the largest production cross sections due to their QCD interactions. Figure 1.7 shows the production cross sections for pairs of strongly produced SUSY particles as a function of their masses. At a fixed mass, first- and second-generation squarks and gluinos have the highest cross sections. The production cross section of top squarks is an order of magnitude smaller, because, since top quarks are not among the constituents of the proton, the t -channel diagram is not available for top squark production. The centre-of-mass energy for pp collisions at the LHC increased from 13 TeV during Run 2 to 13.6 TeV in the Run 3 LHC data-taking period. This increase leads to a substantial increase in the production cross sections of strongly produced of SUSY particles, especially at higher masses. This happens because the proton PDFs decrease steeply at large momentum fractions. Even a small increase in the collider’s centre-of-mass energy significantly enhances the availability of partons with sufficient energy to produce heavy particles, resulting in a pronounced increase in the production cross section. Figure 1.8 shows as an example the ratio of the $\tilde{t}_1$ pair production cross section at 13.6 TeV (Run 3) compared to 13 TeV (Run 2), demonstrating this large increase.

In addition to having a lower production cross section compared to other squarks, stops also present experimental challenges due to their signatures resembling those of the top quark. De-




Figure 1.8 – Ratio between the $\tilde{t}_1$ pair production cross section at 13.6 TeV (Run 3 centre-of-mass energy) and at 13 TeV (Run 2 centre-of-mass energy) [29].

spite this, there are strong theoretical motivations that justify the interest in dedicated searches for direct top squark production. First, as previously discussed, the large mixing between the LH and RH top squarks ($\tilde{t}_L$ and $\tilde{t}_R$) can result in a significant mass splitting between the two physical states, $\tilde{t}_1$ and $\tilde{t}_2$, making $\tilde{t}_1$ the lightest squark, potentially lighter than the current mass limits for the first and second-generation squarks. In addition, top squarks play a crucial role in the SUSY solution to the hierarchy problem. As presented in Section 1.2.1, each supersymmetric particle contributes to the radiative corrections to the Higgs bare mass, δm_H^2 , with a sign opposite to that of its SM partner, cancelling the quadratic dependence on the high-energy cut-off Λ (see Equation 1.28). If SM particles and their SUSY superpartners were mass-degenerate (unbroken SUSY), the cancellation would be exact. This does not happen due to the spontaneous SUSY breaking described in Section 1.2.1, but the divergence is still reduced to a logarithmic one:

$$\delta m^2(H) \propto \frac{\lambda_f^2 N_c^f}{8\pi^2} (m_{\tilde{f}}^2 - m_f^2) \ln(\Lambda^2 / m_{\tilde{f}}^2). \quad (1.45)$$

According to Equation 1.45, naturalness is restored if the mass difference between the top quark and its SUSY partner is within a few hundred GeV to a few TeV.

Top squark phenomenology

The top squark **production** cross section in pp collisions depends on the stop mass, as shown in Figure 1.7. The cross sections for masses between 500 GeV and 1.5 TeV span more than three orders of magnitude. In pp collisions, the top squark is produced via strong interactions, with two possible mechanisms [30]:

- quark-antiquark annihilation ($q\bar{q} \rightarrow \tilde{t}_1 \tilde{t}_1^*$ with $\tilde{t}_1^*$ indicating the anti-particle of the $\tilde{t}_1$);
- gluon-gluon fusion ($gg \rightarrow \tilde{t}_1 \tilde{t}_1^*$)

The top squark production diagrams at tree-level are shown in Figure 1.9.




Figure 1.9 – Tree-level diagrams contributing to the quark-antiquark annihilation production channel (first row) and gluon-gluon fusion channel (second row) [31].




Figure 1.10 – Stop decay modes allowed in different kinematic regions [32].

The accessible **decay modes** of the top squark depend on the mass spectrum of the SUSY particles involved in the process in the specific model under consideration. The model considered for the results presented in this thesis is a *simplified model*, in which only the $\tilde{t}_1$ and the lightest neutralino, the $\tilde{\chi}_1^0$, are assumed to be light enough to be accessible at the LHC energies, and the stop decays into a top quark and a neutralino with a branching ratio of 100%. The masses of the stop and the neutralino are considered as free parameters. Figure 1.10 illustrates the corresponding parameter space $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$, along with the various accessible decay modes that become accessible in different mass regions. The stop decay modes consistent with the considered simplified model are summarized below.

- If $\Delta m = m_{\tilde{t}_1} - m_{\tilde{\chi}_1^0} > m_t$ (where m_t is the mass of the top quark), the top squark can decay via $\tilde{t}_1 \rightarrow t + \tilde{\chi}_1^0$, where the t further decays either hadronically or leptonically. This is called *two-body* decay.
- If $m_b + m_W < \Delta m < m_t$, the previous decay is not kinematically allowed, since the top cannot be produced on-shell. In this case, the top squark decays via an off-shell top, i.e.




Figure 1.11 – Feynman diagram representing the 2-body stop decay [33].

$\tilde{t}_1 \rightarrow b + W + \tilde{\chi}_1^0$, where b is the bottom quark and W is the W boson that further decays. This is called *three-body* decay.

- If $\Delta m < m_b + m_W$, neither the t or the W can be produced on-shell, and the stop decays via $\tilde{t}_1 \rightarrow b + f + f' + \tilde{\chi}_1^0$, where f and f' are respectively a fermion and an antifermion and they are produced according to the branching ratios (BR) of the SM W decays. This is called *four-body* decay. In this kinematic region, where the two- and the three-body decays are kinematically suppressed, the stop can also decay via $\tilde{t}_1 \rightarrow c + \tilde{\chi}_1^0$, where c is the charm quark, a very rare one-loop mediated decay.

Finally, the decay mode $\tilde{t}_1 \rightarrow b\tilde{\chi}^\pm \rightarrow bW^\pm\tilde{\chi}_1^0$ is allowed in all the previously mentioned kinematic regions, as long as $m_{\tilde{\chi}_1^0} < m_{\tilde{\chi}^\pm} < (m_{\tilde{t}_1} - m_b)$.

Figure 1.11 shows the Feynman diagram of the stop two-body decay, which is the decay mode considered in the analysis presented in this thesis.

1.2.4 Natural SUSY searches at the LHC

To preserve the naturalness of the theory, the masses of the SUSY particles are subject to theoretical constraints, directly related to the level of fine-tuning considered acceptable for the theory. Fine-tuning can be quantified as the effect that the parameters of the theory have on the value of the physical Higgs mass m_H . At tree level, m_H receives a contribution from the μ parameter, implying that naturalness favours relatively light higgsinos, as their mass is directly proportional to $|\mu|$. However, the very low production cross section of direct higgsino pair production at the LHC makes their direct detection experimentally challenging. At one-loop level, the Higgs potential is further corrected by gauge and Yukawa interactions, with the dominant contribution arising from top and stop quark loops. Given the higher production cross-section of top squark in pp collisions, experimental searches for top squarks are particularly well-motivated, as already discussed in detail in Section 1.2.3. Finally, additional constraints arise from gluinos, which contribute significantly to the radiative corrections of the stop mass through two-loop




Figure 1.12 – Expected and observed 95% CL excluded regions in the $m(\tilde{t}_1)$ - $m(\tilde{\chi}_1^0)$ plane for $\tilde{t}_1$ -pair production, assuming different $\tilde{t}_1$ decay modes with branching ratio of 100%. Signal models that lie within the contours are excluded. Uncertainty bands corresponding to the $\pm 1\sigma$ variation of the combined expected limit are also indicated. The diagonal grey dashed lines indicate the kinematic threshold of the stop decay modes. Exclusion limits shown in the plot on the left [34], the statistical combination of the semileptonic [35] and the fully-hadronic [36] analyses is shown in the plot on the right.

effects, therefore impacting indirectly the Higgs potential.

This section presents an overview of the SUSY searches performed at the LHC by the two multipurpose experiments, ATLAS and CMS, motivated by the naturalness argument. It focuses mainly on the top squark, main topic of the thesis, but it also includes a summary of searches for higgsinos and gluinos.

Top squark searches at the LHC

Extensive searches for the top squark have been performed by the ATLAS and CMS Collaborations, using the LHC data collected during the Run 1 and the Run 2 data-taking periods, corresponding to pp collisions at a centre-of-mass energy of 7-8 TeV (Run 1) and 13 TeV (Run 2). So far, no evidence for top squark production has been observed. As a result, these searches have set exclusion limits on the top squark mass, which vary depending on the assumptions of the simplified models under consideration. Figure 1.12 and 1.13 show the results obtained by the ATLAS and the CMS Collaborations, respectively, for simplified models where only the $\tilde{t}_1$ and the $\tilde{\chi}_1^0$ are assumed to be accessible at the LHC energies. The results are shown in terms of exclusion limits at 95% confidence level in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane, obtained by analyses targeting different decay channels, each dominating in specific regions of the parameter space. In these models, $\tilde{t}_1$ masses up to approximately 1.3 TeV are excluded for a very light $\tilde{\chi}_1^0$ ($m_{\tilde{\chi}_1^0} \approx 1$ GeV).

The results of searches for third-generation squarks, along with other SUSY analysis performed by ATLAS, have also been interpreted in the context of the 19-parameter pMSSM introduced in Section 1.2.2. To perform this interpretation, a set of realistic SUSY models was generated by




Figure 1.13 – Expected and observed 95% CL excluded regions in the $m(\tilde{t}_1)$ - $m(\tilde{\chi}_1^0)$ plane [37] (left) and in the $m(\tilde{t}_1)$ - $\Delta(\tilde{t}_1 - \tilde{\chi}_1^0)$ plane [38] (right) for $\tilde{t}_1$ -pair production, assuming different $\tilde{t}_1$ decay modes with branching ratio of 100%. Signal models that lie within the contours are excluded. Uncertainty bands corresponding to the $\pm 1\sigma$ variation of the combined expected limit are also indicated.


scanning the pMSSM parameter space, in a phase-space range allowed by previously existing experimental constraints and in order to ensure that the considered particles within the kinematic reach of the LHC. Figure 1.14a shows the fraction of excluded models in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane, based on results from Run 1 searches. The comparison between the black area (where all scanned models are excluded) and the white line (simplified models contour) shows that often simplified models are an over-simplification, and can overestimate the sensitivity of the searches.

Another example of an ATLAS stop search considering a less simplified SUSY model with respect to the one considered in this thesis and in the standard analyses presented in Figure 1.12 is shown in Figure 1.14b. The results are based on 36.1 fb^{-1} of ATLAS pp collision data at $\sqrt{s} = 13 \text{ TeV}$ collected during the Run 2 of the LHC. Events with 0, 1, and 2 leptons in the final state are considered. The considered model, motivated by gauge unification at the GUT scale, assumes a bino-like LSP, and a wino-like next-to-lightest SUSY particle (NLSP), with wino mass eigenstates $(\tilde{\chi}_1^\pm, \tilde{\chi}_2^0)$ with a mass two times larger than that of the LSP $(\tilde{\chi}_1^0)$. All the possible decay mode of the $\tilde{t}_1$ (and the $\tilde{b}_1$) to the LSP are considered: $\tilde{t}_1(\tilde{b}_1) \rightarrow t(b)\tilde{\chi}_1^\pm, t(b)\tilde{\chi}_{1,2}^0$. The corresponding branching ratios, whose sum is bounded to unity, vary across the mass plane, and based on the sign of the higgsino mass parameter μ .


This thesis presents the first results from top squark searches based on LHC Run 3 data.

Higgsino and gluino searches at the LHC

This paragraph provides a general overview, with some representative examples, of the searches for gluino and higgsino pair production performed by the ATLAS and CMS experiments using



(a) Fraction of pMSSM points excluded by ATLAS in the $(\tilde{t}_1, \tilde{\chi}_1^0)$ mass plane [39]. The simplified-model limit overlaid is set assuming directly produced top squark pairs, with each decaying to a top quark and neutralino with 100% BR. The colour scale indicates the fraction of pMSSM points excluded in each mass bin, with black squares indicating 100% of the generated models being excluded.



(b) Expected (dashed lines) and best observed (solid line) exclusion limits at 95% CL in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane for direct $\tilde{t}_1/\tilde{b}_1$ pair production, assuming bino-LSP and wino-NLSP with $m(\tilde{\chi}_1^\pm)/m(\tilde{\chi}_2^0) \approx 2m(\tilde{\chi}_1^0)$ [34]. All decay modes are considered, with BR varying across the mass plane and depending on the sign of the higgsino mass parameter μ .

the pp collision data collected during the LHC Run 2 data-taking period. Gluinos have a relatively large production cross section at the LHC, thus the ATLAS and CMS sensitivities extend to more than two TeV already analysing the LHC Run 2 data only. Figure 1.15 provides a broad overview of the ATLAS gluino searches using Run 2 data, considering simplified models assuming RPC. The limits are presented as a function of the gluino the lightest neutralino masses, with different curves representing different analyses with different assumptions on the gluino decay chain. All analyses assume a decay $\tilde{g} \rightarrow \tilde{q}^* q$, and the flavour of the squark together with its coupling to the electroweak sector determines the selected final state. The simplest decay chain, $\tilde{g} \rightarrow \tilde{q}^* q \rightarrow qq\tilde{\chi}_1^0$ is realised when all squarks and electroweakinos, except for $\tilde{\chi}_1^0$, are assumed to have inaccessible masses.

The results of the analysis where the squark originating from the $\tilde{g}$ decay is a third generation $\tilde{t}$ or $\tilde{b}$ are considered separately, due to their previously discussed role in the definition of the SUSY naturalness. These searches yield the most extended sensitivities, due to the presence of b -quarks in the final state, which allow to apply the powerful experimental b -tagging techniques, which provide an excellent tool to discriminate signal events from the SM background. The virtuality of the $\tilde{q}^*$ can in some cases lead to long-lived $\tilde{g}$. Dedicated analyses are designed to address this case, and are not included in Figure 1.15.

Other less simplified scenarios, especially involving longer gluino decay chains, are considered modifying the assumptions on the electroweak sector. The gluino decay chain can include the production of a vector boson (e.g. $\tilde{q}^* \rightarrow q\tilde{\chi}_1^\pm \rightarrow qW\tilde{\chi}_1^0$ or $\tilde{g} \rightarrow q\tilde{q}'W\tilde{\chi}_1^0$), leading to possible final states containing same-charge leptons or high jet multiplicity, in case of hadronic-decaying vector bosons. Another possibility is considering a scenario where the LSP is a light gravitino




Figure 1.15 – ATLAS exclusion limits at 95% CL based on 13 TeV data in the $(\tilde{g}, \tilde{\chi}_1^0)$ mass plane for different simplified models assuming different decay chains with 100% BR from the gluino to the LSP (lightest neutralino or gravitino).

($m_{\tilde{G}} \approx 1$ TeV), in which the phenomenology is determined by the NLSP. The analysis considering the decay chain $\tilde{g} \rightarrow q\bar{q}(\gamma/Z)\tilde{G}$ via $\tilde{\chi}_1^0$ achieves the maximum sensitivity for high $\tilde{\chi}_1^0$ masses, as the E_T^{miss} in the event, which primarily determines the analysis acceptance, is proportional to $m_{\tilde{\chi}_1^0} - m_{\tilde{G}}$.

Considering all the scenarios presented above, the ATLAS Run 2 analyses are able to exclude gluino masses up to more than 2.4 TeV, placing increasingly stringent constraints on natural SUSY models.

The phenomenology of scenarios with higgsino-like light electroweakinos often leads to very small mass splitting (down to a few GeV) between the lightest chargino ($\tilde{\chi}_1^\pm$) and the lightest neutralino ($\tilde{\chi}_1^0$). This results in compressed spectra and challenging experimental signatures.

An example of experimental signatures that can be exploited for these searches are those presenting two or three leptons, arising from decay chains such as $\tilde{\chi}_1^\pm \rightarrow W^\pm \tilde{\chi}_1^0$ and $\tilde{\chi}_2^0 \rightarrow Z/h\tilde{\chi}_1^0$, and *unconventional* detector signatures, such as disappearing tracks (i.e. a track in the inner detector with no associated calorimeter or muon activity, see Section 3.1.1 for more details on track reconstruction procedures in ATLAS). Disappearing tracks address scenarios with very small mass splittings ($\lesssim 500$ MeV) between $\tilde{\chi}_1^\pm$ and $\tilde{\chi}_1^0$, where the chargino becomes long-lived and leaves tracks inside the detector before decaying in a pion and a neutralino ($\tilde{\chi}_1^\pm \rightarrow \tilde{\chi}_1^0 \pi^\pm$). Figure 1.16 shows a summary of recent Run 2 CMS searches for higgsino pair production, considering the following relation between neutralinos and chargino masses: $m(\tilde{\chi}_2^0) = m(\tilde{\chi}_1^0) + 2\Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0)$. These searches fully close the sensitivity gaps for the considered higgsino simplified model in the very compressed region, which is the region favoured by naturalness arguments, where the higgsino masses are now excluded up to 140 GeV, exceeding the limits set by the LEP experiments.




Figure 1.16 – Summary plot comparing exclusion limits on chargino and neutralino masses in the $(m(\tilde{\chi}_1^\pm), \Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0))$ plane, with analyses featuring final states with disappearing tracks, a soft isolated tracks, and soft opposite-sign electron pairs.

The work presented in this thesis is based on the data collected with the ATLAS (A Toroidal LHC ApparatuS) [3] detector at the Large Hadron Collider (LHC) [2]. This chapter provides an overview of the most relevant aspects of the LHC (Section 2.1) and of the ATLAS detector (Section 2.2).

2.1 The Large Hadron Collider

The LHC, located at the CERN (Conseil Européen pour la Recherche Nucléaire) site, is a hadronic accelerator and collider, installed in a 26.7 km-long circular underground tunnel at a depth of approximately 100 m. It was designed to reach a centre-of-mass energy of 14 TeV and an instantaneous luminosity of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ in proton-proton (pp) collisions. Since it started, the LHC has undergone three operational periods: Run 1 (2009-2013), Run 2 (2015-2018), and Run 3 (2022-2026, ongoing at the time of writing). The LHC can collide protons (pp collisions) and heavy nuclei (*heavy-ion (HI) programme*). Since the results of this thesis are mostly based on pp collisions recorded during Run 2 and Run 3, this chapter will only be focused on that part of the LHC programme.

During Run 2, the LHC accelerated protons to 6.5 TeV, reaching a centre-of-mass energy ($\sqrt{s}$) of 13 TeV, which was already close to the design value. In Run 3, the centre-of-mass energy of pp collisions has been further increased to $\sqrt{s} = 13.6 \text{ TeV}$.

After the end of Run 3, the LHC will undergo a major upgrade, which will lead to the *High-Luminosity LHC (HL-LHC) phase* [40]. The HL-LHC is planned to begin operations in 2030 (start of Run 4) and continue through 2041 (end of Run 5), as shown in Figure 2.1.

2.1.1 Accelerator complex

The LHC ring is the final step of a complex accelerating chain, shown in Figure 2.2. Each accelerator injects the beam of protons into the next one, in order to increase its energy. For protons, the accelerator Run 3 chain is composed of:

- the linear accelerator **Linac 4**, accelerating protons up to 160 MeV, replacing **Linac 2**, which used to accelerate protons up to 50 MeV until the end of Run 2;
- the **Proton Synchrotron Booster**, accelerating protons up to 2 GeV (1.4 GeV in Run 2);
- the **Proton Synchrotron**, accelerating protons up to 26 GeV;




Figure 2.1 – Long term LHC schedule, as of September 2024 (most updated information at the time of the writing) [41]. The long term schedule of the LHC is shown, starting from 2021, before the start of Run 3, until 2041, when the end of Run 5 is planned.

- the **Super Proton Synchrotron**, which accelerates protons up to 450 GeV and finally injects them into the LHC.

Protons are injected into the LHC in the clockwise (*Beam 1*) and anti-clockwise (*Beam 2*) directions. They are then accelerated to the nominal energy, $\sqrt{s}/2$, using superconducting *Radio Frequency (RF) cavities*. Each RF cavity generates an oscillating voltage of 2 MV at 400 MHz and operates at a temperature of 4.5 K.

The proton beams circulating the LHC are grouped into *bunches*, which are shaped by the oscillating RF field: protons that lag behind or move ahead of the bunch centre are either accelerated or decelerated to bring them back toward the centre. The oscillation frequency of the RF cavities defines the space regions in which proton bunches can be located, called *RF buckets*. When the LHC is in operation, it has approximately 35640 RF buckets. RF buckets filled with proton bunches must have a minimal separation of 10 RF buckets, which corresponds to a minimal time between proton bunches (*bunch spacing*) of 25 ns. The LHC is designed to collide two beams made of up to 2808 proton bunches, each composed of 10^{11} protons.

The particle bending in the LHC relies on 1232 Niobium-Titanium superconducting deflecting dipole magnets, which operate at the temperature of 1.9 K and produce a field of approximately 8.33 T. A total of 392 quadrupole magnets are used to focus the beam and dedicated corrector magnets are used to correct trajectory imperfections.

In the LHC, protons collide at four different locations along the accelerator ring. These locations correspond to the position of the four main experiments:

- **ATLAS** and **CMS** (Compact Muon Solenoid) [43] are general purpose detectors, designed to cover the widest possible range of physics measurements in proton-proton and heavy ion collisions.
- **LHCb** (Large Hadron Collider beauty) [44] is focused on *B* hadron decays. The purpose of this detector is to study the properties of *b* quarks, and to measure the CKM matrix ele-




Figure 2.2 – Layout of the CERN accelerator complex [42].

ments parameters that describe CP violation and rare b quark decays, looking for physics beyond the Standard Model.

- **ALICE** (A Large Ion Collider Experiment) [45] is dedicated to heavy ion collisions, to study the strong interactions of quarks-gluon plasma.

Other smaller experiments are situated near the larger detectors: **TOTEM** (TOTal Elastic and diffractive cross-section Measurements) [46], **LHCf** (LHC forward experiment) [47], and **MoEDAL** (Monopole and Exotic Detector At the LHC) [48], and the newest ones, **FASER** (ForwArD Search ExpeRiment) [49] and **SND@LHC** (Scattering and Neutrino Detector at the LHC) [50], installed during the long shutdown between Run 2 and Run 3 (2019-2021).

2.1.2 LHC fill cycle

To deliver collisions to the experiments, the LHC operates in a cycle consisting of different phases, each corresponding to one or more beam modes [51]. The LHC *nominal cycle* is described below and illustrated in Figure 2.3.

- **Injection.** Once the current in the magnets is increased to generate the field required for injection, proton beams are injected from the accelerator chain into the LHC rings, according to a predetermined filling scheme, which defines the number of proton bunches and the bunch spacing between them.




Figure 2.3 – Example of LHC operational cycle [52]. The intensity evolution of Beam 1 (blue) and 2 (red) and the beam energy (dashed green) are shown. The vertical lines denote the beam modes changes.

- **Ramp.** During this phase, proton beams are accelerated to their collision energy. The RF systems accelerate the protons and the current in the magnets is further increased.
- **Squeeze and adjust.** In these two phases, beams are prepared for collisions. First, the beam sizes at the interaction points (IPs) are reduced (*squeeze*). Then, the beam parameters (such as position and crossing angle) are adjusted to ensure optimal conditions for collisions (*adjust*).
- **Stable beams.** This is the phase when the LHC conditions remain stable, collisions occur in the experiments, and detectors can be safely turned on to record data. Minor adjustments of beam parameters are allowed.
- **Dump and ramp down.** The beams are extracted from the rings and safely dumped. A dump may be scheduled by the LHC as part of routine operations, requested, for example by experiments in case of problems with the detector, or unplanned. Following the dump, the magnetic fields are ramped down.

The period between two consecutive stable beams phases is called *turnaround*. It includes the nominal cycle along with all the necessary actions to set up the machine for operation with beams.

2.1.3 Luminosity and pile-up

The study of rare events requires both high beam energies and high beam intensities, the former in order to discover massive particles and the latter in order to have a sufficiently large number of events. The *instantaneous luminosity* ($\mathcal{L}$) is a measurement of the number of collisions per area and time. It relates the total number of events (N) and the cross section (σ) of the considered process as:

$$\frac{dN}{dt} = \mathcal{L} \cdot \sigma. \quad (2.1)$$

For a Gaussian beam distribution, $\mathcal{L}$ can be expressed as:

$$\mathcal{L} = \frac{N_b^2 n_b f_{\text{rev}} \gamma F}{4\pi \epsilon_n \beta^*} [\text{cm}^{-2} \text{s}^{-1}], \quad (2.2)$$

where:

- N_b is the number of protons in a bunch;
- n_b is the number of colliding proton bunches in a beam;
- f_{rev} is the revolution frequency;
- γ is the relativistic Lorentz factor for the proton;
- ϵ_n is the normalised transverse beam emittance, which describes the spread of proton positions and momenta within the beam in the plane perpendicular to the beam direction;
- β^* is the beta function, which describes how the beam's transverse dimensions change with respect to its longitudinal position along the accelerator, evaluated at the IP;
- F is the geometric luminosity reduction factor, which accounts for the loss in effective luminosity when there is a non-zero crossing angle between proton bunches.

Typical values of these parameters for the LHC are shown in Table 2.1 for the different years of LHC operations. Original design values are also presented.

Parameter	Design value	2015	2016	2017	2018	2022	2023	2024
$N_b [10^{11} p]$	1.15	1.2	1.25	1.25	1.1	1.40	1.60	1.60
n_b	2808	2244	2220	2556 - 1868	2556	2462	2358	2352
$f_{\text{rev}} [\text{kHz}]$	11.245	11.245	11.245	11.245	11.245	11.245	11.245	11.245
γ_r	7462	6929	6929	6929	6929	7249	7249	7249
$\epsilon_n [\mu\text{m} \cdot \text{rad}]$	3.75	2.6-3.5	1.8-2.0	1.8-2.2	1.8-2.2	1.8-2.5	1.8-2.5	1.8-2.5
$\beta^* [\text{cm}]$	55	80	40	40 $\rightarrow$ 30	30 $\rightarrow$ 27 $\rightarrow$ 25	60 $\rightarrow$ 30	120 $\rightarrow$ 30	120 $\rightarrow$ 30

Table 2.1 – Typical parameters values at the LHC for pp collisions in Run 2 [53] and for Run 3 [54], shown together with the design values.

These operating parameters have been optimised to maximise the instantaneous luminosity delivered to the experiments. During Run 2, the peak instantaneous luminosity reached the value of $2.1 \times 10^{34} \text{ cm}^{-2} \text{s}^{-1}$, already exceeding the original LHC design target of $10^{34} \text{ cm}^{-2} \text{s}^{-1}$. In Run 3, it further increased to approximately $2.4 \times 10^{34} \text{ cm}^{-2} \text{s}^{-1}$. Figures 2.4 and 2.5 show the peak instantaneous luminosity delivered by the LHC to ATLAS during stable beams in Run 2 and Run 3, respectively.

The instantaneous luminosity in the LHC is not constant. During each LHC fill, it decreases over time, mainly due to the pp collisions occurring at the IPs that reduce the number of protons in the beam. The duration of the stable beams phase is optimised balancing the beam luminosity degradation with the amount of time needed by the accelerator chain to prepare and inject a




Figure 2.4 – Peak instantaneous luminosity delivered to ATLAS in Run 2, in the data-taking year: 2015 (upper left figure), 2016 (upper right figure), 2017 (bottom left figure), and 2018 (bottom right figure) during stable beams for pp collisions at a centre-of-mass energy of 13 TeV as a function of time [55].




Figure 2.5 – Peak instantaneous luminosity delivered to ATLAS in Run 3, in the data-taking year: 2022 (upper left figure), 2023 (upper right figure), 2024 (bottom figure) during stable beams for pp collisions at a centre-of-mass energy of 13.6 TeV as a function of time [56].




Figure 2.6 – Cumulative luminosity versus time delivered to (green) and recorded by ATLAS (yellow) during stable beams for pp collisions at 13 TeV centre-of-mass energy in Run 2 (left figure) for pp collisions at 13.6 TeV centre-of-mass energy in Run 3 (right figure) [55, 56].

new beam after beam a dump. The decrease of the instantaneous luminosity as a function of time is described by the following equation:

$$\mathcal{L}(t) = \mathcal{L}_0 \exp -t/\tau \quad \text{with} \quad \tau = \frac{N_0^{\text{tot}}}{\mathcal{L}_0 \sigma_{\text{inel}} k'} \quad (2.3)$$

where $\mathcal{L}_0$ is the initial instantaneous luminosity, N_0^{tot} is the initial beam intensity, σ_{inel} is the total inelastic pp interaction cross section, and k is a normalisation factor that takes into account the number of IPs along the accelerator ($k \simeq 2$ for the LHC, since only at the ATLAS and CMS IPs collisions occur at high-luminosity).

The *integrated luminosity* (L), defined as:

$$L = \int \mathcal{L} dt, \quad (2.4)$$

quantifies the total amount of data collected over a given period of time. The total number of collisions (N_{coll}) during this time is directly proportional to L , following the relation: $N_{\text{coll}} = \sigma_{\text{inel}} L$.

The results presented in this thesis are based on the data collected by ATLAS at a centre-of-mass energy of 13 TeV during the LHC Run 2, corresponding to an integrated luminosity of 147 fb^{-1} before data quality requirements [55], and on the data collected by ATLAS at a centre-of-mass energy of 13.6 TeV during 2022, 2023 and 2024, corresponding to an integrated luminosity of 35.7 fb^{-1} , 29.7 fb^{-1} , and 118 fb^{-1} , respectively, before data quality requirements [56].

Figure 2.6 shows the cumulative integrated luminosity delivered to and recorded by the ATLAS detector at a centre-of-mass energy of 13 TeV in Run 2, and at a centre-of-mass energy of 13.6 TeV in Run 3.

Another important parameter related to the instantaneous luminosity is the mean number

of inelastic interactions per bunch crossing (μ), defined as:

$$\mu = \mathcal{L} \frac{\sigma_{\text{inel}}}{n_b f_{\text{rev}}}. \quad (2.5)$$

This quantity is called *pile-up*, which occurs when additional soft collisions happen during the same bunch crossing as the primary hard-scattering event. Pile-up events are generally treated as background in physics analyses. Their presence deteriorates the energy resolution and makes the identification and reconstruction of interesting events more challenging. Pile-up is constituted of two main components:

- the *in-time* pile-up, constituted by the additional pp collisions occurring in the same bunch crossing as the collision of interest;
- the *out-of-time* pile-up, constituted by the additional pp collisions occurring in bunch-crossings just before and after the collision of interest that affect the signal in the collision of interest when detectors are sensitive to several bunch-crossings or their electronics integrate over more than 25 ns.

Figure 2.7 shows the distribution of the number of interactions per bunch crossing and the mean number of interactions ($\langle\mu\rangle$) per year, for the Run 2 and Run 3 (2022, 2023, -2024) data. Since




Figure 2.7 – Luminosity-weighted distribution of the mean number of interactions per crossing (μ) for Run 2 pp collision data at 13 TeV centre-of-mass energy (left), and Run 3 pp collision data at 13.6 TeV centre-of-mass energy (right). All data delivered to ATLAS during stable beams is shown, and the integrated luminosity and the mean μ value is given in the figures. The mean number of interactions per crossing corresponds to the mean of the Poisson distribution of the number of interactions per crossing calculated for each bunch [55, 56].

pile-up is directly proportional to instantaneous luminosity, it is important to find a balance between achieving high luminosity values, which correspond to larger integrated luminosity and, consequently, to larger collected datasets, and minimising pile-up, which degrades the quality of data.

In Run 3, the LHC implemented the β^* luminosity levelling technique at the ATLAS and CMS IPs during pp data-taking. The luminosity levelling allows to keep the instantaneous luminosity near its peak value for several hours during the stable beams phase, delaying the exponential




Figure 2.8 – Pileup and β^* evolution with β^* levelling during the LHC fill 8387 (November 2022) [57]. The target pile-up was $\mu = 54 \pm 2.5\%$. The final $\beta^* = 30\text{cm}$ was reached after 5.2 hours of levelling.

decay described in Equation 2.3, thus maximising the integrated luminosity. The β^* -levelling technique is executed as follows: as soon as collisions begin, the instantaneous luminosity naturally starts to decline. When it reaches the lower tolerance limit at the ATLAS and CMS IPs, the β^* value is reduced. Since, as shown in Equation 2.2, β^* is inversely proportional to the instantaneous luminosity, at each step the luminosity increases going back to its peak value. The β^* value is decreased in steps, until it reaches the final value of 30 cm, as illustrated in Figure 2.8.

2.2 The ATLAS detector

The ATLAS experiment is one of the two multipurpose detectors at the LHC, designed to explore the widest possible range of physics phenomena arising from pp collisions. Its physics program ranges from high-precision measurements of the SM parameters to the search for new physics beyond the SM (BSM). The design requirements for the ATLAS detector are defined based on the goals of the broad physics programme of the experiment, using a benchmark set of processes that cover a wide range of potential new phenomena that could be observed at the TeV scale. The identification of events containing these rare processes is particularly challenging due to the very high inelastic pp cross section, resulting in poor pile-up conditions (see Section 2.1.3), and by the high cross section of QCD jet production, which dominates over the other rarer processes in pp collisions, creating significant backgrounds. As a result, the primary objective of the detector's design is to enable the precise identification of the signatures associated with these processes of interest. The design requirements to achieve this goal are outlined below.

- **Radiation hardness and high granularity:** given the challenging experimental conditions at the LHC, the detector must have fast and radiation-hard electronics, and high granularity to handle the intense particle fluxes and minimise the effects of pile-up (see Section 2.1.3).
- **Large angular acceptance:** for accurate event reconstruction, particularly for missing trans-

verse energy (E_T^{miss}) measurements, which are crucial for new physics searches.

- Precise charged-particle momentum resolution and reconstruction efficiency, and capability to identify secondary vertices for offline τ -leptons and b -jets tagging.
- Good electromagnetic calorimetry for electron and photon identification and energy measurements, complemented by full-coverage hadronic calorimetry for precise jet and E_T^{miss} measurements.
- Good muon identification and momentum resolution over a wide range of momenta, with the ability to unambiguously determine the charge of high- p_T muons.
- Efficient triggering on low p_T objects while maintaining sufficient background rejection, to achieve an acceptable trigger rate for most physics processes of interest.

The layout of the ATLAS detector is shown in Figure 2.9. It has a cylindrical shape with a forward-backward symmetry with respect to the IP. It measures approximately 25 m in diameter and 44 m in length, with a weight of about 7000 t. The detector consists of several nested sub-detectors, that are (moving from the beam pipe outwardly in the radial direction): the **Inner Detector (ID)**, a high-granularity tracker for the reconstruction of charged particles trajectories and the identification of primary and secondary interaction vertices; the **Electromagnetic (EM) Calorimeter**, for the high-precision measurement of electrons and photons energy and direction; the **Hadronic (HAD) Calorimeter**, for the detection of charged and neutral hadrons; and the **Muon Spectrometer (MS)**, for the measurement of muons momenta. The detector also includes a magnet system, which produces the magnetic fields needed for the measurement of charged particle momenta in the Inner Detector and in the Muon Spectrometer, and the **Trigger and Data Acquisition (TDAQ) system**, which analyses collision events in real time to select the most interesting ones for storage and further analysis. Figure 2.10 shows how different types of particles interact with the ATLAS sub-detectors. The original configuration of the detector, as it was built for the start of the LHC, is described in Reference [3]. As the LHC operating conditions changed over the years, the ATLAS detector underwent substantial upgrades, to maintain excellent performance despite the harsher data-taking conditions. The major upgrade between Run 2 and Run 3, that involved mainly the EM Calorimeter, TDAQ system and the Muon Spectrometer, is called *Phase-I Upgrade*, and moves towards the requirements foreseen for the HL-LHC phase running conditions.

The work presented in this thesis is based on data collected with ATLAS during both Run 2 and Run 3. Since Run 3 is still ongoing at the time of writing, the results presented here are among the first to be based on the Run 3 dataset. Hence, this section focuses on the Run 3 configuration of the ATLAS detector [58]. For each subsystem that underwent a major upgrade, the differences from the Run 2 setup are detailed, giving a clear overview of the improvements and their impact on detector performance.

Coordinate system

A right-handed coordinate system based on the coordinates (x, y, z) is defined in order to describe the position and the direction of the particles inside the ATLAS detector. The origin of




Figure 2.9 – Cut-away view of the ATLAS detector in the Run 3 configuration [58].

the system is located at the nominal IP. The beam direction defines the z -axis and the x - y plane is transverse to the beam direction, with the positive x -axis pointing from the IP to the centre of the LHC ring, and the positive y -axis outgoing from the xz -plane. Since the momentum of partons along the beam axis is unknown (see Section 1.2.3), boost-invariant transverse quantities are defined through the projection on the xy -plane. The azimuthal angle ϕ is measured around the z -axis on the xy -plane, having $\phi = 0$ along the positive x -axis, and the polar angle θ is defined with respect to the z -axis, having $\theta = 0$ along the positive z -axis and $\theta = \pi/2$ on the xy -plane. The coordinate system described above is sketched in Figure 2.11a. The variable commonly used to describe the angle of a particle relative to the beam axis, instead of θ , is the *pseudorapidity*, η , defined as:

$$\eta = -\ln \left[\tan \left(\frac{\theta}{2} \right) \right]. \quad (2.6)$$

The pseudorapidity is obtained as the massless limit of the rapidity $y = 0.5 \ln[(E + p_z)/(E - p_z)]$, where E is the energy and p_z is the z -component of the momentum of the described particle. This variable is preferred over the polar angle θ because differences in pseudorapidity are Lorentz invariant under boosts along the longitudinal axis. Values taken by η and their relations with the polar angle θ are shown in Figure 2.11b. Large (small) values of η correspond to the forward (central) region of the detector.

The angular distance ΔR in the $\eta\phi$ -plane between two objects of coordinates (η_i, ϕ_i) and (η_j, ϕ_j) is defined as:

$$\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} = \sqrt{(\eta_i - \eta_j)^2 + (\phi_i - \phi_j)^2}. \quad (2.7)$$




Figure 2.10 – Neutral particles escape the tracker undetected. Electrons and photons release energy in the EM calorimeter, while the HAD calorimeter is designed to measure the energy of the hadrons. Muons are detected in the MS. Weakly interacting particles, such as neutrinos or BSM particles, escape all the ATLAS sub-detectors, but their presence can be deduced by a measurement of the missing transverse momentum [59]. Chapter 3 describes in detail how physics objects are reconstructed in ATLAS.




Figure 2.11 – Polar coordinates used to describe the position of a particle in the ATLAS detector (left), and relation between the pseudorapidity η and the standard polar angle θ (right) [60].

2.2.1 Magnet system

A proper magnetic field distribution is necessary to bend the trajectories of charged particles produced in collisions, allowing to measure their momenta. A charged particle moving in a




Figure 2.12 – Schematic illustration of the ATLAS magnet system.

magnetic field, in fact, follows a curved trajectory, with the curvature radius depending on both the particle's momentum and the strength of the magnetic field. Hence, by knowing the curvature radius and the strength of the magnetic field, it is possible to reconstruct the particle's momentum. For a homogeneous magnetic field of strength B , the component of the particle's momentum perpendicular to the direction of the magnetic field (p_T) is related to the curvature radius R by the equation:

$$p_T[\text{GeV}] \approx 0.3 \cdot |q| \cdot B[\text{T}] \cdot R[\text{cm}], \quad (2.8)$$

where q is the electric charge of the particle. To generate this magnetic field, ATLAS uses a superconducting magnet system, that produces a magnetic field over a volume of approximately $12\,000\text{ m}^3$, ensuring an almost complete geometric coverage of both the ID and the MS. The ATLAS magnet system consists of two parts: a **central solenoid** [61], and a **toroidal system**, composed of a barrel toroid [62] and two endcap toroids [63]. The configuration of the magnet system is shown in Figure 2.12. The central solenoid is a superconducting coil located between the ID and the EM Calorimeter. It generates an axial magnetic field of 2 T within the ID volume. The solenoid is 5.8 m long, with an inner radius of 1.23 m and an outer radius of 1.28 m. It is cooled to 4.5 K and the magnetic field is generated by an electric current of 8 kA.

The barrel toroid and the two endcap toroids, installed outside the calorimeters, produce a toroidal magnetic field of 0.5 T and 1 T in the central and endcap regions, respectively, covering the MS volume. The toroidal magnetic field is generated by an electric current of 20 kA. The barrel toroid consists of 8 separate superconducting coils, it has an inner diameter of 9.4 m, an outer diameter of 20.1 m, and is 25.3 m long. Each of the two endcap toroids consists of eight superconducting coils located inside an insulating vacuum vessel, with a diameter of 10.7 m and 5 m wide. The endcap toroids generate a magnetic field near the beam axis, in order to deflect the particles with small polar angles.




Figure 2.13 – Cut-away view of the ATLAS inner detector in the Run 3 configuration [58].

2.2.2 Tracking: the Inner Detector

The ID [64] is used for the reconstruction of the trajectories (*tracks*) of the charged particles interacting with its material, and for the identification of primary and secondary interaction vertices. The ID is designed to achieve high-precision momentum and vertex resolution despite the high track density in the detector. To handle the large number of tracks per event, high granularity detectors are needed to perform precise track measurements and accurate event vertex reconstruction. Additionally, due to its proximity to the interaction point, the ID design requires a high resistance to radiation for all the sensors, on-detector electronics, mechanical structures, and services.

The ID is composed of a barrel and of two endcaps and covers the pseudorapidity region $|\eta| < 2.5$. It is immersed in the 2 T magnetic field generated by the central solenoid (see the previous Section 2.2.1). The layout of the ID is presented in Figure 2.13.

The ID consists of three sub-detectors: the **Pixel detector**, the **Semiconductor Tracker (SCT)**, and the **Transition Radiation Tracker (TRT)**. The sub-detector structure is shown in a cut view of the ID in Figure 2.14. The Pixel Detector and the SCT are silicon detectors [65], based on a silicon p-n junction operating in reverse bias voltage [66]. Segmenting the p-n junction in one or two dimensions allows to measure the position of a particle. Devices with segments in only one direction are known as *strip detectors* (such as the SCT), while those with two-dimensional segmentation are *pixel detectors* (such as the Pixel Detector).

Pixel detector

The Pixel Detector [67, 68] is the sub-detector closest to the IP and with the best space resolution. It is composed of four cylindrical layers in the barrel area, located at radial distances of 3.3 cm, 5.1 cm, 8.9 cm and 12.3 cm from the beam line. The innermost layer is the **Insertable B-Layer**




Figure 2.14 – Barrel ID cut view in the transverse plane [58].

(IBL) [69], installed at the beginning of Run 2 to improve the tracking performance by reducing the distance between the interaction point and the first detector layer. In the two endcaps it is composed of three disks, located at ± 50 cm, 58 cm and 65.0 cm from the centre of the detector. The pixel layers are segmented in the $r\phi$ plane and along the z -axis. The pixel sensors in the endcaps and in the three outer layers in the barrel have a size of $50 \times 400 \mu\text{m}$, while they are slightly smaller in the IBL, with a size of $50 \times 250 \mu\text{m}$. The Pixel detector typically provides four measurements (*hits*) per track, with the first hit normally being in the IBL. The intrinsic spatial resolution in the three outer layers of the barrel (endcaps) is of $10 \mu\text{m}$ in the $r\phi$ -plane and $115 \mu\text{m}$ in z (r) direction.

Semi-Conductor Tracker

The SCT [70] is located at a radial distance of 30 cm to 51 cm from the beam line. It is composed of four concentric barrel layers of silicon microstrips grouped into four cylindrical structures, placed at approximate radial distances of 30 cm, 37 cm, 44 cm and 51 cm from the beam line. The endcap regions are composed of nine disks each. To allow three-dimensional position measurements, the layers are double-sided, with a 40 mrad stereo angle between the two sides, enabling the measurement of the z and the r coordinates in the barrel and in the endcap, respectively. Barrel modules consist of four rectangular silicon-strip sensors each, with strips with a constant pitch of $80 \mu\text{m}$ and approximately parallel to the magnetic field and beam axis. Endcap disks instead consist of up to three rings of modules with trapezoidal sensors. The strip direction is radial with constant azimuth and a mean pitch of $80 \mu\text{m}$. The SCT typically provides eight strip measurements per track, corresponding to four space points. The intrinsic spatial resolution per module in the barrel (endcap) is of $17 \mu\text{m}$ in the $r\phi$ -plane and $580 \mu\text{m}$ along the z (r) direction.

Transition Radiation Tracker

The TRT [71] is the outermost layer of the ID, located at a radial distance between 55 cm and 108 cm from the beam line. It consists of several layers of gas-filled straw tubes interleaved with transition radiation material, thus functioning both as a drift chamber and a transition radiation detector, contributing to particle identification. Each straw-tube has a diameter of 4 mm. In the barrel region, the straws are 140 cm long and parallel to the beam axis, while in the endcaps region they are 37 cm long and arranged radially in wheels. The tubes are filled with a Xenon-Argon gas mixture, which becomes ionised when a charged particle passes through. Xenon is the base component of the gas mixture, needed for its efficient X-ray absorption. To generate the transition radiation, the spaces between the tubes are filled with polymer fibres in the barrel region and foils in the endcaps. The transition radiation can be emitted by particles when they cross the boundary between different materials. Since the intensity of the transition radiation photons depends on the particle's relativistic Lorentz factor ($\gamma = E/m$), it can be used to discriminate between electrons and pions of similar energy, due to their mass difference ($m_e \approx 0.51 \text{ MeV}$, $m_{\pi^\pm} \approx 139.57 \text{ MeV}$). Compared to the other ID sub-detectors, the TRT has a lower spatial resolution (130 μm in the $r\phi$ -plane) and only provides $r\phi$ information. However, it can supply a high number of hits, of approximately 36 per track.

2.2.3 Calorimetry

The ATLAS calorimetric system consists of three calorimeters: the **EM Calorimeter**, optimised for the measurement of the energy of electrons and photon, and covering the pseudorapidity region $|\eta| < 3.2$, the **HAD Calorimeter**, for the measurement of hadrons in the pseudorapidity region $|\eta| < 3.2$, and the **Forward Calorimeter** (FCal), designed to measure both electromagnetic and hadronic showers in the forward regions, extending the pseudorapidity coverage of the ATLAS calorimetric system to $3.1 < |\eta| < 4.9$. A view of the calorimeter system is presented in Figure 2.15.

Calorimeters [72] are designed to measure the energy of incident electrons, photons and hadrons, while also providing a fundamental contribution to the measurement of the E_T^{miss} . When a particle enters a calorimeter, it interacts with its material, producing multiple new particles, which in turn interact with the detector producing other particles. This cascade of secondary particles with progressively degraded energy is called *shower*. The process ends when the particles of the shower reach the critical energy and are fully absorbed by the matter of the detector. The total energy of the secondary particles is measured by detecting the charge produced through ionisation in a gas or in a liquid, or the light produced in a scintillating material, and serves as a measurement of the energy of the initial incident particle. The size and geometry of the calorimeters are designed to contain the showers generated by incoming particles. The development of these showers depends on the particle's interactions with the detector material and varies significantly between primary electrons or photons and primary hadrons. The ATLAS calorimeters are *sampling calorimeters* [73], meaning that they consist of alternating layers of a *passive* medium of high-density material, which degrades the particle energy, and layers of *active* medium, which generates the measurable signal. Liquid Argon (LAr) or polystyrene scintillator are used as active material, lead, copper, or iron are used as passive material, depending on the calorimeter type.

A key quantity to measure a calorimeter performance is its energy resolution, which depends




Figure 2.15 – Cut-away view of the ATLAS calorimeter system in the Run 3 configuration [58]

on the initial energy of the incoming particle and can be expressed as follows:

$$\frac{\sigma(E)}{E} \approx \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c, \quad (2.9)$$

where a is the *stochastic term*, describing the contributions to the resolution from statistical fluctuations in the shower development; b is the *noise term*, including the contributions from the electronic noise induced by readout circuits and from the pile-up noise; for an individual signal in the energy range studied by ATLAS, this term is dominated by the electronic noise and generally negligible, while in ATLAS operating conditions with high luminosity, pile-up becomes the dominant source of noise and can significantly affect the resolution. Finally, c is the *constant term*, arising from factors such as detector non-uniformity, misalignments, and uncertainties in the electronic calibration. The constant term becomes the limiting factor for resolution at very high energies.

The LAr EM calorimeter

The LAr EM calorimeter [74, 75] surrounds the ID. It has a cylindrical shape, 6.65 m long and with an outer radius of 2.25 m, and it is divided into a barrel part ($|\eta| < 1.475$) and two endcap components ($1.375 < |\eta| < 3.2$). The transition areas between the barrel and endcaps are known as *crack regions* ($1.37 < |\eta| < 1.52$). Each endcap calorimeter is mechanically divided into two coaxial wheels: an outer wheel covering the region $1.375 < |\eta| < 2.5$, and an inner wheel covering the region $2.5 < |\eta| < 3.2$. The LAr detector is a sampling calorimeter made of 4 – mm thick layers of LAr as its active medium, interleaved with copper electrodes, which collect the ionisation charge generated by the EM shower, and with lead absorbers plates. The

lead layers have different thickness, ranging from 1.1 mm to 1.5 mm in the barrel, and from 1.7 mm to 2.2 mm in the end-caps. Electrodes and absorbers are arranged in an *accordion geometry* to provide a complete ϕ symmetry without azimuthal cracks due to the outgoing readout system. The LAr EM calorimeter is composed of four longitudinal layers, ensuring the precise measurement of the longitudinal development of the EM shower. The four layers are:

- **Presampler.** This layer is located in the region of $|\eta| < 1.8$ and consists of an active LAr layer of thickness 1.1 cm (0.5 cm) in the barrel (endcap) region. It is used to correct for the energy lost by electrons and photons before reaching the calorimeter.
- **Strips.** This layer has the finest segmentation along η ($\Delta\eta = 0.0031$), in order to discriminate between prompt photons and π^0 mesons decaying into two almost collinear photons.
- **Middle.** This layer is segmented in cells of size $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$. It contains most of the energy of the showers originated by photons and electrons with energy up to 50 GeV.
- **Back.** This layer is used to correct for the energy leakage of high-energy showers in the HAD calorimeter.

The nominal LAr EM calorimeter resolution is:

$$\frac{\sigma(E)}{E} \approx \frac{10\%}{\sqrt{E}} \oplus 0.7\%, \quad (2.10)$$

where E is expressed in GeV and the stochastic term depends on $|\eta|$. The noise term is negligible at the typical energy ranges studied by ATLAS, since it is not expected to contribute for energies above 0.5 GeV.

Figure 2.16 is a sketch of a barrel module of the LAr EM calorimeter. The accordion geometry of the Strips, Middle, and Back layers is visible, along with their granularity in the $\eta\phi$ -plane.

The main upgrade in the Run 3 system compared to Run 2 is the implementation of the new *digital trigger readout* [76], which provides finer granularity inputs to the upgraded trigger system (see Section 2.2.6). During Run 1 and Run 2, the calorimeter signals sent to the trigger system consisted of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ *Trigger Towers* (TTs), formed by analog summing the E_T of the calorimeter cells in the longitudinal layers of the calorimeter. The new digital trigger readout path implemented for Run 3 is based on smaller clusters called *Super Cells*, formed by introducing an additional lateral and longitudinal segmentation, splitting the TTs longitudinally in four layers, of which the middle ones are further split laterally along η into four strips each. One Super Cell can thus cover an area as small as $\Delta\eta \times \Delta\phi = 0.025 \times 0.1$, depending on which longitudinal layer it is located in, with finer granularity for the Strips and the Middle layers ($|\eta| < 2.5$), as shown in Figure 2.17. In this way, the granularity of the signals is increased by up to a factor of ten, significantly improving the both the trigger selection efficiency for events with interesting signatures and its discrimination power against background events. The *legacy* (Run 1 and Run 2) analogue trigger path continued to operate in parallel to the new digital trigger path during the Run 3 commissioning phase, and was disabled in August 2024.




Figure 2.16 – Layout of the three accordion layers in a barrel module of the LAr EM calorimeter [74].

Hadronic calorimeters

The HAD Calorimeter [77] is located after the EM calorimeter in the radial direction. It is designed to fully contain the showers generated by hadrons and to measure their energies. It has a cylindrical shape, it is 11.5 m long with an external radius of 4.25 m, and covers the pseudorapidity region $|\eta| < 3.9$. The HAD calorimeter is mainly composed of two sub-systems, which use different technologies: the **Tile Calorimeter** in the central region, and the **Hadronic Endcap Calorimeter** (HEC) in the endcaps region.

The Tile Calorimeter is placed directly outside the EM calorimeter. It is a sampling calorimeter using steel as the absorber material and scintillating tiles as the active material. Radially, the tile calorimeter extends from an inner radius of 2.28 m to an outer radius of 4.25 m. Its *barrel* covers the pseudorapidity region $|\eta| < 1.0$, and its two *extended barrels* cover the range $0.8 < |\eta| < 1.7$. The tile calorimeter is composed of cells of size $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$. The Minimum Bias Trigger Scintillators (MBTS) [78], covering the pseudorapidity range of $2.0 < |\eta| < 4.0$, are used for triggering and for luminosity monitoring. Historically, they have been considered part of the tile calorimeter as they share its readout electronics.

The HEC consists of two independent wheels in each endcap, located just behind the electromagnetic endcap calorimeter. LAr is used as the active material, and copper and tungsten plates are used as the absorber. To reduce the drop in material density at the transition between the endcap and the forward calorimeter (at around $|\eta| = 3.1$, see below), the HEC extends out to $|\eta| = 3.2$, thus overlapping with the forward calorimeter. Similarly, the HEC η range overlaps that of the tile calorimeter ($|\eta| < 1.7$) by extending up to $|\eta| = 1.5$. The HEC cells have a granularity of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ for $|\eta| < 2.5$, and $\Delta\eta \times \Delta\phi = 0.2 \times 0.2$ for $2.5 < |\eta| < 3.2$.

The nominal energy resolution for hadronic jets (combining the HAD and the EM calorimeters) is:

$$\sigma(E)/E \approx \frac{50\%}{\sqrt{E}} \oplus 3\%, \quad (2.11)$$




Figure 2.17 – Trigger granularity from each 0.1×0.1 TT after the Phase-I upgrade of the LAr calorimeter electronics. Ten E_T values are provided from $1 - 4 - 4 - 1$ longitudinal/transverse samples, each forming a SuperCell. Layer 0 corresponds to the Presampler, Layer 1 corresponds to Strips, Layer 2 corresponds to Middle and Layer 3 corresponds to Back [4].

The Forward Calorimeter

The Forward Calorimeter (FCal) [79] is designed to measure the energy of both electromagnetic and hadronic showers in the forward region. Its modules are located at a distance of 4.7 m from the IP and cover the pseudorapidity range $3.1 < |\eta| < 4.9$. The FCal consists of three modules in each endcap: the first, made of copper as the absorber material and LAr as the active medium, is optimised for electromagnetic measurements; while the other two, made of tungsten as the absorber material and LAr as the active material, primarily measure the energy of hadronic interactions. The typical energy resolution of the FCal is:

$$\frac{\sigma(E)}{E} \approx \frac{100\%}{\sqrt{E}} \oplus 10\%, \quad (2.12)$$

2.2.4 Muon System

Due to their higher mass with respect to electrons, muons experience a significantly reduced energy loss through bremsstrahlung, by a factor of approximately $m_\mu^{-2}/m_e^{-2} \approx 10^{-6}$. For this reason, muons can escape the calorimeters with minimal energy loss and the MS [80] is thus located at the outer part of the ATLAS detector. Its main goals are to detect charged particles exiting the barrel and endcap calorimeters and to measure their momentum in the pseudorapidity range $|\eta| < 2.7$. It also triggers on these particles in the region $|\eta| < 2.4$ (see Section 2.2.6). The measurement of the momentum is based on the magnetic deflection of muon trajectories caused by the toroidal magnetic field (see Section 2.2.1). A full view of the MS is shown in Figure 2.18.

The MS is composed of three parts: the central **barrel** and two **endcaps**. Each part is structured into three detector stations: an inner station, close to the calorimeters, a middle station, still inside the toroid field, and an outer station, well outside the magnetic field. In the barrel, these stations form three concentric and cylindrical layers around the beam axis. In the endcaps,




Figure 2.18 – Cut-away view of ATLAS Muon Spectrometer in the Run 3 configuration [58].

they take the shape of three discs, referred to as *Wheels*. Due to the large volume of the MS, it is impossible to provide continuous tracking; instead, tracks are typically reconstructed from straight segments formed within each station. The detectors in each station are multilayered, and provide at least six hits along the muon's trajectory, which can be reconstructed as a straight track segment with a well-defined spatial resolution. Muon tracks are then reconstructed by fitting curved paths to the three segments and matching them to ID tracks (see Section 3.1.1 in Chapter 3). The MS is designed to give a stand-alone transverse momentum resolution better than 15% for 1 TeV tracks.

In barrel, precision measurements of the bending coordinate are performed using multilayered *Monitored Drift Tube* (MDT) chambers. The middle and outer layers are also equipped with *Resistive Plate Chambers*, used for triggering and for the azimuthal coordinate of the tracks.

The endcaps employ a different set of technologies: MDT chambers are used to measure the bending coordinate in the middle and outermost layers, similar to the barrel, while *Thin Gap Chambers* (TGCs) are used for triggering and to provide the azimuthal coordinate measurements in the middle station. The innermost endcap stations are the **New Small Wheels** (NSWs), which fully replaced the **Small Wheels** used during Run 2. The NSWs use two chamber technologies: *small-strip* TGC and *Micro-Mesh Gaseous Structure* (Micromegas) detectors. Both technologies are used for triggering as well as for tracking, enhancing the resolution in both the bending and azimuthal directions compared to Run 2. The Phase-I upgrade mainly aimed to reduce the rate of the fake muons background, caused by charged particles coming from hadronic showers and increasing with luminosity, to be able to trigger on lower- p_T muons while keeping the trigger rates at a manageable level. This detector combination is designed not only to meet the requirements of the new Run 3 running conditions, but also in view of the HL-LHC phase. The Run 2 Small Wheels used *Cathode Strip Chambers* to measure the bending coordinate, and TGCs for triggering and to measure the azimuthal coordinate.

2.2.5 Forward Detectors

Four smaller detector systems cover the ATLAS forward region.

The first system (ordered according to the distance from the IP) is a Cherenkov detector called **LUCID** (LUminosity measurement using Cherenkov Integrating Detector) [81]. It is placed at approximately $z \pm 17$ m from the IP and consists of two sets of photomultiplier tubes that surround the LHC beam pipe, covering the pseudorapidity range $5.561 < |\eta| < 5.641$. It detects inelastic pp interactions in the forward direction, determining the on-line relative-luminosity delivered to ATLAS.

The second detector, the **ZDC** (Zero-Degree Calorimeter) [82], is situated at ± 140 m from the IP and is used to determine the centrality of heavy-ion collisions. It consists of a set of two sampling calorimeter modules, uses Tungsten plates as absorber material and quartz rods interspersed in the absorber as active media. It can only operate during HI or low pile-up data-taking, and is otherwise removed from the beam pipe.

The third detector is the **AFP** (ATLAS Forward Proton) [83], consists of four Roman pot stations, two on each side of the IP, equipped with silicon tracker and time-of-flight detectors. On each side, the station closest to the IP is located at $z = \pm 206$ m, while the other is located at $z = \pm 217$ m. The AFP is designed to tag and measure the properties of protons emerging intact from the pp collisions in ATLAS and travelling at very small angles with respect to the beam pipe, enabling the study of elastic or diffractive processes which are otherwise inaccessible. The last of the forward detectors, **ALFA** (Absolute Luminosity For ATLAS) [84], is located at ± 240 m from the IP. It is designed to determine the total pp cross section as well as the luminosity at the LHC, by measuring elastic proton scattering at very small angles. The system consists of four Roman Pot stations, each equipped with tracking detectors, inserted in Roman Pots which approach the LHC beams vertically, as close as 1 mm to the beam. The tracking detectors consist of multi-layer scintillating fibre structures read out by Multi-Anode-Photo-Multipliers.

2.2.6 Trigger and data acquisition system

The LHC beams collide at the IP with a frequency up to 40 MHz, producing pp collision events with an average size of 1.5 MB. This would require transferring and storing data at an unmanageable rate of almost 60 TB/s, which exceeds the ATLAS's resources. In addition, the majority of pp collisions at the LHC energies produce only soft interactions with no signatures of processes targeted by the ATLAS physics programme. Therefore, one of the key components of the ATLAS experiment is the TDAQ system, designed to efficiently select and record only potentially interesting events, reducing the output rate to a manageable level. It is a two-level system, consisting of a hardware-based **Level-1 (L1) trigger** and a software-based **High-Level Trigger (HLT)**.

The L1 trigger is mainly based on two independent systems, which use custom electronics to trigger on reduced-granularity information from the calorimeters (**L1Calo**) and the muon detectors (**L1Muon**). They identify events containing high- p_T leptons, jets, and large missing or total transverse energy, which can all indicate the presence of an interesting physics process. The L1 topological processor (**L1Topo**) takes in input Trigger Objects (TOBs) containing kinematic information from the L1Calo and L1Muon systems, and applies topological selections.

The final L1 trigger decision is then formed by the Central Trigger Processor (**CTP**), which in-

tegrates the information received by the L1Calo, L1Muon, and L1Topo systems, as well as from additional sub-detectors (MBTS, AFP, ALFA, LUCID, and ZDC). These signals are combined with information about the LHC bunch pattern (see Section 2.1.1), which is configured externally, and determines in which bunch crossings the triggers are allowed to fire. This information is combined via a logical AND operation with other trigger conditions before the L1 trigger decision is generated. Most triggers are configured to be active in bunch crossings where proton bunches are colliding, which typically have a rate of approximately 30 MHz; however, certain triggers designed to collect events for background studies are also activated in bunch crossings with only one bunch or no bunches at all. Up to 512 distinct L1 trigger items may be configured in the CTP. The CTP also applies preventive *dead time*, a mechanism to limit the number of close-by L1 accepts to be within the constraints on the detector read-out latency. It limits the minimum time between two consecutive L1 accepts (*simple dead time*), to prevent overlapping read-out windows, and restricts the number of L1 accepts allowed within a set number of bunch crossings (*complex dead time*), to prevent front-end buffers overflows [85]. Events are accepted at a rate up to the maximum detector readout rate of 100 kHz at a fixed latency below 2.5 μ s. Upon receiving an L1 trigger accept signal, events are sent to the HLT for further processing.

The HLT runs on a computing farm of 60 000 CPU cores [86] and selects up to 3 kHz of fully-built physics collision events (with additional rate for partial readout) to be sent to permanent storage, with an average processing time of approximately 600 ms. HLT algorithms reconstruct events at a progressively higher level of detail, either in the full detector volume or within *Regions-of-Interest* (RoIs), restricted regions identified by the L1 trigger as containing candidate objects. Unlike the L1 trigger, HLT algorithms use full-granularity information from the sub-detectors, to provide better energy and momentum resolution for the selections, and enabling precision tracking for particle identification. The flow of data through the Run 3 TDAQ system is illustrated in Figure 2.19.

Phase-I Upgrade

Before the start of Run 3, the trigger system underwent a major upgrade [4], to optimise the performance under the new LHC running conditions. The upgraded system was designed to handle higher levels of pile-up while keeping the same selection efficiency and lower the L1 trigger rate.

The L1Calo system was upgraded to perform on-detector digitisation of the signals from the LAr calorimeters, enabling electromagnetic cluster reconstruction with up to ten times finer granularity and jet reconstruction with four times finer granularity, as well as algorithms capable of processing data from the entire calorimeter (see Section 2.2.3). The LAr TTs are now distributed to new Feature Extraction (FEX) processors, which include new electromagnetic (eFEX) and jet (jFEX) feature extractors, as well as a global feature extractor (gFEX). The eFEX uses full SuperCell granularity for precise reconstruction of EM objects and hadronically decaying τ leptons. In the region with $|\eta| < 2.5$, the jFEX receives TTs with twice the granularity compared to the legacy system. Its performance is expected to be similar to that of the legacy L1Calo system for the single jet triggers, while improving the reconstruction for the nearby jets in the multi-jet triggers. The gFEX, has been designed with a coarser granularity, comparable




Figure 2.19 – Simplified layout of the ATLAS TDAQ hardware in Run 3 and the flow of data through the system [58]. Rates reported on the right refer to fully-built events (see Chapter 7), excluding calibration data and streams saving only part of the event information.

to the Run 2 system, so that the data from the entire calorimeter can be processed on a single module, facilitating the identification of boosted objects and global observables (for example E_T^{miss} , a crucial quantity for BSM physics searches). As already mentioned in Section 2.2.3, the legacy L1Calo system continued to operate in parallel to the Phase-I during the commissioning phase at the start of Run 3. The legacy system was finally fully decommissioned in August 2024. Figure 2.20 shows how the new Phase-I triggers provide an increased trigger efficiency while keeping the L1 trigger rate lower than the legacy triggers, which is the main goal of the Phase-I upgrade.

The L1Muon trigger system selects events containing high- p_T muons, based on the inputs from the RPCs in the barrel region, and from the TGCs and the NSWs in the endcaps, and then transmits data to the CTP via the **MUCTPI** (Muon-to-CTP Interface). For Run 3, the L1Muon system was upgraded mainly with the introduction of the NSWs and, consequently, of a new endcap trigger processors, improving the rejection of fake muons (see Section 2.2.4). To reduce the rate of fake muons in the endcaps, the Run 2 trigger already applied coincidence requirements between the outer TGC station and either the inner TGC stations or the tile calorimeter. With the replacement of the Small Wheels by the NSWs in Run 3, the system achieves a further reduction in the fake muon rates. Figure 2.21 shows how the rate of the primary single muon trigger decreases after the complete activation of the NSW coincidences during the 2024 data-taking.

The HLT software, already redesigned in Run 2 to support multiprocessing (inter-event parallelism), was further adapted for Run 3 to support multithreaded execution, introducing intra-event and intra-algorithm parallelism. This reduces the memory usage per core, allowing to operate at higher pile-up without saturating the memory of the HLT farm. Another important upgrade in the software was the speed-up of the track-reconstruction algorithms, enabling *Full Scan tracking* for hadronic signatures, which improves pile-up separation and jet energy resolution at low- p_T , and *Large Radius Tracking* (LRT), which enables high efficiency reconstruction of




Figure 2.20 – The plots compare the performance of the legacy single electron trigger L1_EM22VHI (red) and the corresponding Phase-I L1_eEM26M (blue) in data recorded by ATLAS in June 2023. On the left, the rates as a function of the instantaneous luminosity are shown. The Phase-I system improves the rate of the single electron trigger, and the ratio between the rates is approximately 80%, independent of the luminosity. On the right, trigger efficiencies as a function of the electron p_T are presented, measured using electrons from $Z \rightarrow ee$ decays [87].




Figure 2.21 – Level-1 trigger rate of the primary single muon trigger (14 GeV p_T -threshold), as a function of date. The rate for each run is normalised to an instantaneous luminosity of $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. On 17 April 2024, coincidences between the TGC in the endcap middle wheel, the Tile calorimeter, and the NSW detectors have been activated to suppress background, reducing the rate by 10 kHz. Further sectors were included on 8 May 2024, reducing the rate by 3 kHz [88].

long-lived particle trigger signatures.

A more comprehensive review of the HLT system and its upgrades is provided in the third part of this thesis (Part [III](#), *Trigger Menu and Operations*).

Object reconstruction in ATLAS

To translate the electrical signals read out by the ATLAS sub-detectors into representations of the underlying physics process, several reconstruction and identification steps are performed. The physics analysis discussed in this thesis presents a final state containing leptons (electron and muons), jets (including b -jets), and missing transverse energy (E_T^{miss}). This chapter details the methods used to reconstruct these objects from ATLAS signals. Particularly: Section 3.1 discusses the common algorithms employed for the reconstruction of multiple physics objects, including tracking, vertexing, topo-clustering and particle-flow algorithms; Section 3.2 covers the reconstruction of electrons, while Section 3.3 focuses on the reconstruction of muons. Section 3.4 describes the reconstruction of jets, with a focus on the identification of b -tagged jets in Section 3.4.3. Finally, Section 3.5 explains the reconstruction of the E_T^{miss} , including an overview of the E_T^{miss} significance quantity and its application in BSM searches.

3.1 Common reconstruction algorithms

This section provides an overview of the common reconstruction algorithms, employed in the reconstruction process of multiple physics objects. In particular, *tracking* and *vertexing* algorithms, used to reconstruct tracks and primary vertices, respectively, from the information in the ID (and in the MS for muon tracks), are described in Section 3.1.1. The *topo-clustering* algorithm, which reconstructs electromagnetic and hadronic energy deposits in the calorimeters, is described in Section 3.1.2. Finally, the *particle-flow* (Pflow) algorithm, which combines information from all sub-detectors to reconstruct particle objects, is described in Section 3.1.3. The tracks and vertices reconstructed with the tracking and vertexing algorithms serve as inputs for the reconstruction of all the physics objects described in this chapter. Topo-clusters, build by the topo-clustering algorithm, are used in the electron reconstruction process (see Section 3.2.1) and as inputs to the Pflow algorithm. Finally, Pflow objects produced by the Pflow algorithm are used as inputs to the jet reconstruction algorithm (see Section 3.4.1), and to build the isolation variables for muons (see Section 3.3.2).

3.1.1 Tracks and vertices

The reconstruction of charged-particles' tracks and primary interaction vertices is based on the information provided by the ID. Muon tracks are also reconstructed in the MS.




Figure 3.1 – Illustration of the relationship between the track parameters and the associated track. In this scenario, the hard scatter primary vertex is located at $(e_x, e_y, e_z) = (0, 0, 0)$ [89].

Tracking

Charged-particles produced in pp collisions leave signals, referred to as *hits*, in the different layers of the ID, from which their trajectories (**ID tracks**) are reconstructed. Due to the solenoidal magnetic field within the ID, charged-particle tracks follow helical trajectories in the plane transverse to the beam axis (xy -plane), which can be described by five track parameters:

$$(d_0, z_0, \phi, \theta, q/p), \quad (3.1)$$

where d_0 and z_0 are the transverse longitudinal *impact parameters*, defined as the distance between the closest track point and the reference point in the transverse and in the longitudinal plane, respectively; ϕ and θ are the azimuthal and polar coordinates of the track at the point for which d_0 and z_0 are defined, and q/p is the ratio between the particle charge and the magnitude of its momentum. The standard reference point for ATLAS tracks is the *beamspot* position, defined as the average position of the pp interactions in the ID. Figure 3.1 is an illustration of the track parameters described above.

The goal of track reconstruction is to group detector hits originating from the same charged particle and to determine the track parameters of the corresponding trajectory. The process is carried out in two stages: an *inside-out* approach, where tracks are built starting from the innermost layers of the detector, and an *outside-in* approach, where the track formation starts from the TRT. The steps of the inside-out tracking procedure are detailed below [90].

- Initial *track seeds* are formed using sets of three space points, built from hits in the Pixel detector, the SCT, or a combination of the two silicon detectors.
- Track seeds are combined using a combinatorial Kalman filter [91] to form preliminary track candidates.
- To resolve overlaps between track candidates and reject fake tracks, an ambiguity reso-

lution step is performed: track candidates are assigned a score based on a set of quality criteria, and candidates sharing a significant number of hits with higher-scoring tracks are removed, and the shared hits are assigned to the tracks that are retained.

- The remaining track candidates are refitted to obtain the final track parameter estimates.
- An extension of the track to the TRT is performed if the silicon track candidates are associated with a useful set of hits in it.

Once the inside-out tracking is complete, the outside-in tracking is performed, to increase acceptance for particles produced at a greater distance from the IP, such as electrons from photon conversions. This procedure starts from standalone TRT track segments, that are extrapolated inward to match unused silicon hits from the inside-out stage. TRT segments without associated silicon hits are also retained and used for tasks as photon conversion reconstruction. Outside-in tracking is only attempted in regions of interest determined by energy deposits in the EM calorimeter.

The set of reconstructed ID tracks is used as input to the primary **vertex reconstruction** process [92], which identifies the interaction vertices in an event and determines the Primary Vertex (PV), which is the one associated with the hard-scatter interaction. The PV reconstruction is divided in two stages: *vertex finding*, in which reconstructed tracks are associated with vertex candidates, and *vertex fitting*, which performs the reconstruction of the actual vertex position. This procedure is summarized below.

- A set of tracks satisfying the track selection criteria described above is defined.
- A vertex seed is created by estimating the most likely position of a pp interaction vertex, assumed to correspond to the maximum track density along the beam axis.
- All the tracks and the seed are fitted together to determine the best vertex position.
- After the vertex position is determined, all compatible tracks are removed from the seed-finding pool.

The procedure stops when less than two tracks remain for seed-finding, or when the algorithm can no longer form a new vertex seed. Only vertices with at least two associated tracks are considered as PVs. In the high luminosity collisions at the LHC there are typically multiple PVs associated with each pp bunch crossing. Among them, the hard-scatter vertex, where the interesting physics process is more likely to have occurred, is identified as the PV with the highest sum of squared transverse momenta of its associated tracks. Subsequent event reconstruction is performed relative to this primary hard-scatter vertex and only objects originating from it are considered in the reconstruction of physics objects. Any additional PVs are considered as pile-up vertices.

Tracks can also be reconstructed in the MS following a similar approach (**MS tracks**) [93]. The process begins with the identification of short, straight-line segments reconstructed from hits in the individual MS stations (see Section 2.2.4). The segments from different MS stations are then combined to form muon track candidates. Typically, at least two matching track segments from

different MS stations are required to form a muon track candidate. However, in the transition region between the barrel and endcap, candidates may be reconstructed using a single segment due to reduced detector coverage. Once track candidates are formed, a global fit is performed to optimise the association of hits to each muon candidate. The fit also accounts for detector effects such as possible interactions in the detector material and effects of possible misalignments between the different detector chambers. The resulting tracks are finally re-fitted with a loose constraint on the IP and taking into account the energy loss in the calorimeters, and back-extrapolated to the beam line.

3.1.2 Topological clusters

The lateral and longitudinal segmentation of the calorimeters enables a three-dimensional reconstruction of particle showers, implemented through the *topological clustering algorithm*. This algorithm is based on the selection of energy deposits measured in variable-size clusters built from topologically connected cells in the calorimeters, named *topo-clusters* [94]. The variable regulating the development of the topo-clusters is the *significance* of a calorimeter cell ($\zeta_{\text{cell}}^{\text{EM}}$), defined as:

$$\zeta_{\text{cell}}^{\text{EM}} = \left| \frac{E_{\text{cell}}^{\text{EM}}}{\sigma_{\text{noise}}^{\text{EM}}} \right|, \quad (3.2)$$

where $E_{\text{cell}}^{\text{EM}}$ is the energy deposited in the cell, and $\sigma_{\text{noise}}^{\text{EM}}$ is its expected noise, including both the electronic read-out noise and the average contribution from pile-up, which depends on the running conditions. The EM index indicates that $E_{\text{cell}}^{\text{EM}}$ and $\sigma_{\text{noise}}^{\text{EM}}$ are measured at the electromagnetic (EM) energy scale. At this scale, the energy deposited by EM particles (electrons and photons) is reconstructed correctly, while hadronic energy deposits are not corrected for the loss of signal due to the non-compensating character of the ATLAS calorimeters. This effect is compensated in later calibration steps.

The topological clustering algorithm proceeds by first identifying a seed cell; then, neighbouring cells are progressively added laterally and longitudinally to the cluster if their energy is above a certain $\zeta_{\text{cell}}^{\text{EM}}$ threshold, following the *4-2-0 scheme* described below.

- Cells with $\zeta_{\text{cell}}^{\text{EM}} \geq 4$ are identified as seed cells.
- Neighbouring cells with $\zeta_{\text{cell}}^{\text{EM}} \geq 2$ are added around the seed cells. If a cell with $\zeta_{\text{cell}}^{\text{EM}} \geq 2$ is shared between two clusters, the two clusters are merged.
- All directly adjacent cells with $\zeta_{\text{cell}}^{\text{EM}} \geq 0$ are included in the cluster.
- Clusters with multiple local maxima are split into different clusters. A local maximum is defined as a cell with $E_{\text{cell}}^{\text{EM}} > 500 \text{ MeV}$ and at least four neighbouring cells with lower energy.

Topo-clusters form the basic inputs to several reconstruction algorithms, such as the particle-flow algorithm and the algorithms for electron and photon reconstruction described in the next sections. The topological clustering algorithm employed in ATLAS, in fact, is not designed to separate energy deposits from different particles, but rather to separate continuous energy showers of different nature (electromagnetic or hadronic) and to suppress noise.




Figure 3.2 – A flow chart of the particle flow algorithm [95], starting with track selection and continuing until the energy associated with the selected tracks has been removed from the calorimeter. At the end, charged particles, topo-clusters which have not been modified by the algorithm, and remnants of topo-clusters which have had part of their energy removed remain.

3.1.3 Particle-flow

The Pflow algorithm aims to identify and reconstruct each individual particle produced in the interaction, by combining information from all the sub-detectors [95]. The resulting *Pflow objects* are a combination of charged-particle tracks in the ID and calorimeter energy deposits in the calorimeters. By combining ID and calorimeter information, the algorithm benefits from the strengths of each of these subsystems: the ID provides a better energy resolution for low-energy charged particles and allows to match the tracks to the PV, which is fundamental for pile-up rejection; the calorimeters instead provide better energy resolution at higher energies, cover a larger pseudorapidity region, and can detect neutral particles. To avoid double-counting, the energy deposited in the calorimeters by charged particles has to be identified and subtracted, and replaced by the corresponding ID tracks' momenta. It is therefore crucial to accurately identify all the energy deposited by each individual particle, so it can be subtracted without affecting the energy deposits from nearby particles. The Pflow algorithm steps are shown in Figure 3.2 and described below.

- **Track selection.** Tracks are selected following strict quality criteria: at least nine hits in the silicon detectors are required, and tracks must have no missing Pixel hits where expected. Additionally, tracks are required to have $|\eta| < 2.5$ and $0.5 < p_T < 40 \text{ GeV}$. Higher- p_T tracks are excluded as they are often not well isolated, compromising the accurate removal of the calorimeter energy associated with the track. Tracks matched to candidate electrons or muons are also excluded, as the algorithm is optimised for the subtraction of hadronic energy.
- **Tracks to topo-clusters matching.** The algorithm attempts to match each track to one topo-cluster, first selecting topo-clusters for which $E^{\text{clus}}/p^{\text{track}} > 0.1$ (where E^{clus} is the energy of the topo-cluster and p^{track} is the momentum of the track), and then selecting the topo-cluster that is closest to the track according to the metric $\Delta R = \sqrt{(\Delta\phi/\sigma_\phi)^2 + (\Delta\eta/\sigma_\eta)^2}$ (where σ_ϕ and σ_η represent the topo-cluster's angular widths).
- **Expected energy deposit.** The expected calorimeter energy deposit ($\langle E_{\text{dep}} \rangle$) from a charged particle with momentum p^{track} that formed the track is estimated as: $\langle E_{\text{dep}} \rangle = p^{\text{track}} \langle E_{\text{ref}}^{\text{clus}} / p_{\text{ref}}^{\text{track}} \rangle$, where the expectation value $\langle E_{\text{ref}}^{\text{clus}} / p_{\text{ref}}^{\text{track}} \rangle$ is determined using reference distributions derived from well-known single-pion data samples.

- **Recovering split showers.** Since a single particle may deposit energy in multiple topo-clusters, for each track/topo-cluster system, the algorithm checks if additional clusters need to be included. This decision is based on the significance of the difference between $\langle E_{\text{dep}} \rangle$ and the matched topo-cluster energy.
- **Cell-by-cell subtraction.** The expected energy deposited in the calorimeter by the particle that produced the track is subtracted cell by cell from the set of matched topo-clusters.
- **Remnant removal.** If the remaining energy in the system is consistent with expected shower fluctuations of a single particle's signal, the topo-cluster remnants are removed. This procedure is applied to tracks sorted in descending p_T -order, starting with tracks matched to a single cluster.

After these steps, the set of selected ID tracks and the remaining topo-clusters should ideally represent the reconstructed event with no double counting of energy between the sub-detectors.

3.2 Electrons

Electrons are charged particles interacting via electromagnetic force, hence their reconstruction is based on electromagnetic energy deposits in the calorimeters associated with a track in the ID. As described in the following sections, electrons are first reconstructed, together with photons, using the topo-clustering algorithm, and the reconstructed electron energy is calibrated to improve the data-to-simulation agreement using a data-driven method (Section 3.2.1). Subsequently, electron identification criteria, based on the properties of the electron's track and energy deposit in the calorimeter, are applied to distinguish prompt electrons from the background constituted of jets, converted photons, and electrons originating from hadronic decays (see Section 3.2.2). Finally, the background rejection for identified electrons is further improved by applying isolation requirements, which limit the amount of energy from other particles allowed in the vicinity of the electron (Section 3.2.3).

3.2.1 Electron reconstruction and calibration

Offline photon and electron reconstruction is based on the use of dynamic, variable-size calorimeter clusters, named *superclusters* [96]. Their flexible size allows for the recovery of energy from bremsstrahlung photons and from electrons originating from photon conversions. An electron is defined as an object consisting of a cluster built from energy deposits in the calorimeter (supercluster) matched to ID tracks. Photons are classified as converted, if their cluster is matched to a conversion vertex, and an unconverted, the cluster is not matched to either a track or a conversion vertex. Figure 3.3 illustrates the reconstruction steps of electrons and photons with $|\eta| < 2.5$, which are summarised below.

First, the algorithm prepares the inputs: topo-clusters (reconstructed as defined in Section 3.1.2), are matched to ID tracks (reconstructed as described in Section 3.1.1), which are then re-fitted accounting for bremsstrahlung energy loss. Conversion vertices are also reconstructed and matched to the corresponding clusters. Since the topo-clustering algorithm is not limited to one region of the calorimeter, topo-clusters may include cells from both the EM and the HAD




Figure 3.3 – Algorithm flow diagram for the electron and photon reconstruction [96].

calorimeters. Therefore, electromagnetic showers initiated by incoming photons or electrons are identified applying a selection based on EM energy fraction in the topo-clusters, defined as:

$$f_{\text{EM}} = \frac{E_{\text{L1}} + E_{\text{L2}} + E_{\text{L3}}}{E_{\text{cluster}}}, \quad (3.3)$$

where $E_{\text{L}x}$ is the cluster energy in the x^{th} layer of the EM calorimeter and E_{cluster} is the total cluster energy. Clusters with $f_{\text{EM}} > 0.5$ are considered as coming from an EM shower, and the hadronic cells are removed. Only topo-clusters with EM energy above 400 MeV are retained. Electron candidates within $|\eta| < 2.47$ are defined as EM clusters matched to one or more reconstructed tracks.

Then, the supercluster-building algorithm runs in parallel for electrons and photons, matching tracks to electron superclusters and conversion vertices to photon superclusters. The algorithm uses loosely matched topo-clusters and tracks as input. Different requirements are applied to distinguish electron and photon superclusters: seed topo-clusters with $E_{\text{T}} > 1 \text{ GeV}$ matched with a track with at least four hits in the silicon layers of the ID are used to reconstruct electron superclusters; seed topo-clusters with no matched track and $E_{\text{T}} > 1.5 \text{ GeV}$ are used for unconverted photon superclusters. Around each seed cluster, the algorithm searches for lower- E_{T} satellite topo-clusters within a window of $\Delta\eta \times \Delta\phi = 0.075 \times 0.125$ (3×5 cells in L2). Satellite topo-clusters can originate from bremsstrahlung or photon conversion. For electrons, additional satellite clusters are included if they lie within an extended window of




Figure 3.4 – Diagram of the superclustering algorithm for electrons and photons. Seed clusters are shown in red, satellite clusters in blue [96].

$\Delta\eta \times \Delta\phi = 0.125 \times 0.3$, have at least one matched track with hits in the silicon layers of the ID, and share the same best-matched track as the seed. For converted photons, clusters matched to the same conversion vertex as the seed are also included. Superclusters are then built by combining each seed cluster with its selected satellite clusters.

The requirements for satellite clusters applied to candidate electrons and candidate photon superclusters are summarised in Figure 3.4.

Finally, electron and photon analysis objects are built and calibrated. Additionally, discriminating variables used for electron and photon identification are calculated.

The calibration process for electrons and photons consists of different steps:

- a correction to the truth calorimeter energy is applied using a multivariate algorithm trained on Monte Carlo (MC) simulation, based on the shower properties in the EM calorimeter;
- residual corrections are applied to account for the non-uniformities in the calorimeter response and non-nominal high-voltage settings for certain calorimeter regions;
- a correction for the data-to-simulation agreement, defined based on a set of high-purity $Z \rightarrow ee$ data events, is applied to simulated events to better match the observed data.

3.2.2 Electron identification

To enhance the purity of selected electron objects, i.e. reducing the background contamination from hadronic jets, converted photons, and electrons produced in the decays of heavy-flavour

hadrons (which are hadrons containing b or c quarks), further selections are applied during *electron identification*. The identification of prompt electrons is based on a likelihood discriminant that combines various quantities measured in both the ID and the calorimeters. The electron identification process uses three primary groups of variables:

- **Properties of the primary electron track.** The primary electron track is required to fulfil specific requirements on the number of hits in the two tracking layers closest to the beam line, as well as the number of hits in the SCT. The transverse impact parameter of the track and its significance are then used to construct the likelihood discriminant. Additionally, the track's momentum resolution ($\Delta p/p$) is included, along with particle identification information from the TRT.
- **Spatial development of the EM shower in the EM calorimeter.** The lateral development of the EM shower is characterised by several variables calculated separately for the first and second layer of the EM calorimeter. To reject contributions from multiple incident particles, the total shower width is used. All lateral shower shape variables are calculated by summing energy deposits in calorimeter cells relative to the cluster's most energetic cell. The longitudinal development of the shower is addressed by dynamically selecting the number of cells contributing to the energy measurement in each layer, using the supercluster approach, which reduces the noise in the calorimeter cells. Differences between data and simulation are taken into account by optimising the electron identification procedure and applying data-to-simulation efficiency ratios in subsequent analyses.
- **Spatial compatibility of the primary electron track with the reconstructed cluster.** The spacial compatibility between the primary electron track and the reconstructed calorimeter cluster is assessed to ensure that the track and cluster originate from the same particle. This information is crucial for reducing background from non-electron processes.

A discriminant is formed from the likelihoods for a reconstructed electron to originate from signal, L_s , or background, L_b . They are calculated from probability density functions, P , which are created by smoothing histograms of the n discriminating variables with an adaptive kernel density estimator, separately for signal and background and in several $|\eta|$ and E_T bins:

$$L_{s(b)}(\mathbf{x}) = \prod_{i=1}^n P_{s(b),i}(x_i), \quad (3.4)$$

where $\mathbf{x}$ is a vector of the n discriminating variables, and $P_{s,i}(x_i)$ and $P_{b,i}(x_i)$ are the values of the probability density functions for signal and background, respectively. The likelihood discriminant d_L is defined as the natural logarithm of the ratio of L_s and L_b :

$$d_L = \ln \frac{L_s}{L_b}. \quad (3.5)$$

The probability density functions are derived with a data-driven method using $Z \rightarrow ee$ and $J/\Psi \rightarrow ee$ events for signal and dijet events for background. The electron likelihood identification imposes a selection on d_L and some additional requirements. Three Working Points (WPs) of increasing background rejection power but reduced signal efficiency are defined: *loose*,




Figure 3.5 – Electron identification efficiency for the three WPs (tight, medium, loose) for $Z \rightarrow ee$ events in 139 fb^{-1} of 2015-2018 Run 2 LHC data at 13 TeV as a function of the electron transverse energy E_T integrated over the full pseudo-rapidity range (left) and as a function of the electron pseudorapidity η (right). The error bars and uncertainty bands include the statistical and systematic uncertainties. The middle panel shows the ratios of the efficiencies measured in data over those in MC simulation. The lower panel shows the relative size of the total uncertainties for the combined efficiency measurements [97].

medium, and *tight*. All WPs are optimised in each $|\eta|$ and E_T bin. The signal efficiency increases with the electron E_T , as shown in Figures 3.5 and 3.5, that present the signal efficiency for the three WPs for Run 2 and Run 3, respectively.

The WP used in the analysis presented in this thesis is the tight WP (*TightLLH*).

3.2.3 Electron isolation

Even after the application of the identification requirements, an amount of background from non-prompt leptons originating from heavy-flavour hadronic decays and from light hadrons misidentified as electrons remains. Isolation criteria are implemented to select prompt electron suppressing the remaining background contamination. The activity around electrons is quantified by the tracks and energy deposits in the calorimeters from nearby charged particles. Electron isolation is defined using two variables: the *calorimeter isolation variable* and the *track isolation variable*. The calorimeter isolation variable is calculated from raw calorimeter isolation ($E_{T, \text{raw}}^{\text{iso}}$), which is the sum of the transverse energy of topo-clusters located within a cone of radius $R = 0.2$ around the electron cluster's barycentre. This value is then corrected to take into account the energy of the electron itself, the pile-up and the underlying event contributions. The track isolation variable is defined as the sum of the transverse momenta of the selected tracks within a cone around the electron track, excluding the tracks matched to the electron. Since for electrons produced in the decay of high-momentum heavy particles other decay products can be close to the electron's direction, the track isolation cone size is variable. Specifically, its radius decreases with increasing electron p_T , following the formula:

$$R = \min \left(\frac{10}{p_T [\text{GeV}]}, R_{\text{max}} \right), \quad (3.6)$$




Figure 3.6 – Electron identification efficiency for the three WPs (tight, medium, loose, with an extra B-Layer requirement on the latter) for $Z \rightarrow ee$ events in 3.4 fb^{-1} of early Run 3 LHC data at 13.6 TeV as a function of the electron transverse momentum p_T integrated over the full pseudo-rapidity range (left) and as a function of the electron pseudorapidity η (right). The error bars and uncertainty bands include the statistical and systematic uncertainties. The lower panel shows the ratios of the efficiencies measured in data over those in MC simulation [97].

where R_{max} is the maximum cone size, typically set to 0.2. Different isolation WPs are provided, offering different trade-offs between prompt electron identification efficiency and background rejection power [96]. The WP used in the search presented in his thesis is the *tight* WP, defined as:

$$E_T^{\text{cone20}}/p_T < 0.20, \quad p_T^{\text{varcone20}}/p_T < 0.15. \quad (3.7)$$

3.3 Muons

Muon reconstruction is based on the information from the MS and the ID [93]: as charged particles, muons leave tracks in the ID; additionally, being minimum-ionising particles, they pass through the calorimeters with minimal energy loss, so their energy is primarily inferred from the information provided by the MS (muon trajectory and momentum). As outlined in the following sections, muons are first reconstructed by combining track information from the ID and the MS, and then their momentum is calibrated using data-driven methods to correct for any effects due to detector mismodelling (Section 3.3.1). Muon identification and isolation criteria are then applied to distinguish prompt muons from those originating from hadron decays. Identification criteria focus on the quality of muon tracks to separate prompt muons from those coming from light hadrons decays. Isolation requirements instead, similarly to those for electrons, distinguish prompt muons from those coming from the decay of heavy-flavour hadrons by limiting the presence of other particles in the region around the muon (Section 3.3.2).

3.3.1 Muon reconstruction and calibration

Muon reconstruction begins with the independent reconstruction of charged-particle tracks in the ID and in the MS, which are then combined to form a complete muon track, also taking into account the information on the energy loss in the calorimeters.

Five main reconstruction strategies, each corresponding to a different reconstructed-muon *type*, are implemented. Each algorithm uses different sets of information from the ID, MS, and calorimeters.

- **Combined (CB) muons** are reconstructed by matching MS and ID tracks. Hits from both sub-detectors refitted into a single muon track, taking into account the energy loss in the calorimeters.
- **Inside-out (IO) combined muons** are reconstructed by extrapolating an ID track to the MS and matching it to at least three loosely-aligned MS hits. The ID track, the energy loss in the calorimeter, and the MS hits are then refitted into a single muon track. This algorithm does not rely on a fully reconstructed MS track, thus it is especially useful in regions of limited MS coverage, or for low- p_T muons that may not reach the middle MS station.
- **Segment-tagged (ST) muons** are identified by matching an ID track to at least one reconstructed MS segment. The muon parameters are then taken directly from the ID track fit. This method is particularly useful for muons that cross only one layer of the MS, typically due to their low- p_T or because they pass through less-instrumented regions of the MS.
- **Calorimeter-tagged (CT) muons** are reconstructed by extrapolating ID tracks to energy deposits in the calorimeters that are consistent with a minimum-ionising particle. As with ST muons, the ID track is used directly to define muon parameters. This type of muons has the lowest purity, however, it recovers acceptance in partially instrumented MS regions, particularly in the pseudorapidity region $|\eta| < 0.1$, where space is needed for cabling and services.
- **MS-Extrapolated (ME) muons** are identified when an MS track cannot be matched to an ID track. The track parameters are extrapolated back to the IP to reconstruct the muon object. This type of muons is used to extend the muon reconstruction acceptance to the forward region ($2.5 < |\eta| < 2.7$), not covered by the ID.

These different strategies ensure efficient muon reconstruction across a wide range of kinematic regions, maximizing the muon reconstruction acceptance. While CB muons offer the highest purity, other types such as ST or CT muons help recover efficiency in regions with limited detector coverage or for low- p_T muons, even if with lower purity.

After reconstruction, corrections are applied to the momentum of muon candidates to account for detector effects not well modelled in simulation, in order to improve the agreement between data and simulation. The muon momentum scale and resolution are studied using well-measured $Z \rightarrow \mu\mu$ and $J/\Psi \rightarrow$ processes, by comparing the reconstructed invariant mass distributions in data and MC simulation. Based on these studies, calibration constants are derived in different η - ϕ regions of the detector and applied to muons in simulated events. The corrections ensure that both the absolute scale and the resolution of the muon momentum are accurately reproduced in the simulation.

3.3.2 Muon identification and isolation

Muon *identification* and *isolation* criteria are applied to select prompt muons, which are muons originating from the prompt decay of SM bosons or hypothetical BSM particles, while minimising the background from non-prompt muons produced in hadron decays. Non-prompt muons from light hadrons can be distinguished from prompt muons for the lower-quality muon tracks, due to the change in trajectory deriving from the in-flight decay within the detector. These muons can be rejected by applying identification criteria to muon candidates. In contrast, non-prompt muons from heavy-flavour hadrons produce good-quality muon tracks, but tend to be less tightly associated with the primary vertex and less isolated in the ID and calorimeters with respect to prompt muons. These muons can be suppressed by applying isolation requirements. Both identification and isolation selections are available in three standard WPs: *tight*, *medium*, and *loose*, which offer increasing efficiency but decreasing purity with respect to background contamination.

Since the identification criteria target non-prompt muons coming from light hadrons, they are based on variables that assess the quality of the track fit, quantifying the consistency between the measurements in the ID and the MS:

- q/p compatibility: measures how well the charge-over-momentum (q/p) values agree between the ID and MS. It is defined as the absolute value of the difference between the ratio of the charge and momentum of the muons measured in the ID and MS divided by the sum in quadrature of the corresponding uncertainties.
- ρ' , defined as the absolute value of the difference between the transverse momentum measured in the ID and MS divided by the p_T of the combined track. It quantifies the relative discrepancy between the two sub-detectors.
- Normalised χ^2 of the combined track fit: evaluates the overall quality of the fit combining ID and MS information.

To ensure a robust momentum measurement, additional requirements are imposed on the number of track hits in both the ID and MS. The stop search presented in this thesis uses the medium identification WP for muons with $|\eta| < 2.5$, defined as follows: for $|\eta| < 2.5$, only CB and IO muons are retained, with a q/p compatibility less than 7 to ensure a loose agreement between the ID and MS measurements. In $t\bar{t}$ events, more than 98% of prompt muons passing the medium WP are found to be CB muons.

Prompt muons can be distinguished from muons originating from heavy-flavour hadron decays by measuring the amount of hadronic activity in their vicinity. Muon isolation, similarly to what is done for electrons, is defined as the transverse energy reconstructed in a cone around a muon and divided by the muon p_T . Isolation can be measured independently using only ID tracks (see Section 3.1.1), using only topo-clusters in the calorimeters (see Section 3.1.2), or using Pflow objects, which combine tracking and calorimeter information.

The analysis described in this thesis uses the *tight* isolation WP, with isolation defined based on Pflow objects (*PflowTight* WP), hence only this WP is described below.

All muon isolation WPs include at least one track-based isolation variable, with or without an additional criterion for calorimeter-based or Pflow-based isolation. Track-based isolation is

defined as the scalar sum of the transverse momenta of the ID tracks associated with the primary vertex in an η - ϕ cone of radius R around the muon, excluding the muon track itself. The size of the cone can be $R = 0.2$ (referred to as p_T^{cone20}) or $\min(10 \text{ GeV}/p_T^\mu, 0.3)$ (referred to as $p_T^{\text{varcone30}}$), where the latter is optimised for topologies where other particles are expected in the vicinity of an energetic muon. The Pflow-based isolation variable is defined as the sum of the track-based isolation $p_T^{\text{varcone30}}$ (p_T^{cone20}) for $p_T^\mu < 50 \text{ GeV}$ ($p_T^\mu > 50 \text{ GeV}$), and the transverse energy of neutral Pflow objects in a cone of radius $R = 0.2$ around the muon (E_T^{neflow20}). The E_T^{neflow20} quantity is corrected for the contribution from the energy deposit of the muon itself and for pile-up effects, and is assigned a weighting factor $w = 0.4$, optimised to maximise the rejection of heavy-flavour hadron decays. To summarise, the *PflowTight* WP is defined as:

$$\begin{aligned} (p_T^{\text{varcone30}} + 0.4 \cdot E_T^{\text{neflow20}}) &< 0.045 \cdot p_T^\mu && \text{if } p_T^\mu < 50 \text{ GeV} \\ (p_T^{\text{cone20}} + 0.4 \cdot E_T^{\text{neflow20}}) &< 0.045 \cdot p_T^\mu && \text{if } p_T^\mu > 50 \text{ GeV} \end{aligned} \quad (3.8)$$

requesting a minimum track p_T of 500 MeV for the track-based isolation variable.

3.4 Jets

Jets are collimated sprays of particles arising from the fragmentation and hadronisation of quarks and gluons produced in pp collisions, as shown in Figure 3.7 (see Section 1.1.2). Their accurate reconstruction is crucial for probing the properties of the initiating partons.

In the ATLAS experiment, jets are reconstructed by combining energy deposits measured in




Figure 3.7 – Illustration of the jet formation process, beginning with an initiating parton (quark or gluon), which hadronises to form particle jets detectable by the ATLAS ID and calorimeters [98].

the calorimeters with tracks reconstructed in the ID, used as an input for the PFlow algorithm (see Section 3.1.3), and then calibrated using MC simulated dijet events. The jet reconstruction and calibration procedures are described in the following sections.

3.4.1 Jet reconstruction

Defining a jet requires a procedure to cluster the observed particles and to combine their four-momenta to form distinct objects (the *jets*). These jets must be theoretically well-defined, meaning that their production cross section has to remain finite at all orders of perturbative QCD, and minimally sensitive to non-perturbative effects, such as the parton-to-hadron transition during hadronisation. To achieve that, jet algorithms are required to satisfy the following criteria:

- *infrared safety*: the presence of additional infinitely soft particles radiated by the primary partons should not modify the outcome of the algorithm;
- *collinear safety*: the algorithm should not be sensitive to the emission of particles at very small angles relative to the original parton direction, meaning that a configuration where a fraction of the transverse momentum is carried by a single particle should yield the same result as one where the particle splits into two collinear particles;
- *independence of the input-object*: the algorithm should be independent of the type of input provided, reconstructing the same topology at parton, particle or detector level.

The jet algorithm used by ATLAS is the *anti- k_t algorithm* [99], as implemented within the FAST-JET package [100]. As input, it uses the objects reconstructed by the ATLAS Pflow algorithm described in Section 3.1.3.

The anti- k_t algorithm is part of the class of the *sequential clustering algorithms* [101], which iteratively merge objects starting from those that are closer according to a defined distance metric. The definition of the distance d_{ij} between a pair of objects (i, j) depends on their angular separation and transverse momenta. For the anti- k_t algorithm it is defined as:

$$d_{ij} = \min \left(\frac{1}{p_{Ti}^2}, \frac{1}{p_{Tj}^2} \right) \frac{\Delta_{ij}^2}{R^2} \quad \text{and} \quad d_{iB} = \frac{1}{p_{Ti}^2}, \quad (3.9)$$

where d_{iB} is the distance between the object i and the beam axis, R is the radius parameter which determines the final size of the jet, $\Delta_{ij}^2 = (Y_i - Y_j)^2 + (\phi_i - \phi_j)^2$, and p_{Ti} , Y_i and ϕ_i are the transverse momentum, rapidity and azimuth of the i -th object, respectively.

The algorithm proceeds then with an iterative procedure, similarly to other sequential clustering algorithms, described below.

- It starts with a list of input objects (particles), each with four-momentum p_i ;
- For each object i , it calculates: the beam distance d_{iB} , the pairwise distances d_{ij} between all object pairs (i, j) . Then, it then identifies the minimum value among all the computed distances: $d_{\min} = \min(\{d_{iB}\} \cup \{d_{ij}\})$.
- If the minimum $d_{\min}$ corresponds to the beam distance (i.e. $d_{\min} = d_{iB}$), the object i is considered a final jet and removed from the list. Otherwise, if $d_{\min} = d_{ij}$, the pair (i, j) is merged into a new object with four-momentum $p_{\text{tot}} = p_i + p_j$. The new object is added to the list, and the objects i and j are removed.
- The procedure is iterated until all particles have been assigned to a jet, without requiring a predefined number of jets (*inclusive clustering*).

Since the distance metric of the anti- k_t algorithm is inversely proportional to the square of the particles' transverse momentum, it preferentially clusters harder objects first. As a result, the algorithm is less sensitive to soft radiation from the underlying event and pile-up, making it particularly suitable for the high pile-up events produced in pp collisions at the LHC. The radius parameter value used in most ATLAS analyses, including the search presented in this thesis, is $R = 0.4$. Jets of this type are referred to as *small- R jets* (as opposed to the *large- R jets* [102]), and generally represent jets initiated by quarks and gluons.

3.4.2 Jet energy calibration and resolution

The **Jet Energy Scale (JES)** calibration adjusts the jet energy to match that of *particle-level* jets, using stable final-state particles from MC simulations, without including detector effects. The procedure consists of several correction steps, which adjust the four-momentum of the jets, scaling the jet p_T , energy, and mass. The calibration steps, illustrated in Figure 3.8, can be divided in two stages: the first stage uses MC simulations to derive JES corrections that reduce the impact of pile-up, detector effects, and other parameters; the second stage is a residual *in situ* calibration, which corrects for the remaining differences between data and MC simulations, using data-driven methods based on well-measured reference objects (such as photons and Z bosons).

This simulation-based JES calibration begins with the *pile-up correction*, which removes the




Figure 3.8 – JES calibration steps [103].

contribution of in-time and out-of-time pile-up. This correction is performed in two stages: first, the per-event pile-up contribution is estimated from the median p_T density of jets and subtracted from the jet p_T . The ratio between the uncorrected and the pile-up-subtracted jet p_T is used as a correction factor and applied to the jet's four-momentum. An additional correction based on MC simulations is applied to remove the residual dependence of the jet's p_T on the pile-up as a function of η . It depends on the number of reconstructed primary vertices (sensitive to in-time pile-up), and the mean number of interactions per bunch crossing (sensitive to out-of-time pile-up).

The subsequent step is the *absolute MC-based JES and η calibration*. In this step, the reconstructed jet four-momenta are adjusted to match those of particle-level jets from dijet MC events, correcting for biases in the jet η reconstruction, which are mainly caused by the transitions between different calorimeter technologies and the sudden changes in the calorimeter granularity.

Finally, the *global sequential calibration* step corrects for the different response of jets depending on their flavour, energy distribution of constituent particles, transverse distribution, and the fluctuations of the jet development within the calorimeter. It consists of a series of multiplicative corrections to account for the differences in the calorimeter's response to different types of jets, improving the jet resolution without changing the jet energy response. The corrections are applied sequentially to the jet's four-momentum, to remove dependencies on tracking, calorimeter, and muon spectrometer observables.

The final *in-situ calibration* step is applied only to data, in order to correct for the remaining differences between data and the MC simulations. These differences are caused by the imperfect modelling of the detector response and detector material, and the physics processes involved. The differences are quantified by comparing the jet's p_T or η with those of well-measured reference objects.

The η -intercalibration corrects the JES of forward jets ($0.8 \leq |\eta| < 4.5$) to match that of well-measured central jets ($|\eta| < 0.8$) from a dijet system.

Additional in-situ calibrations are applied sequentially to correct for differences in the response of central jets compared to well-measured reference objects recoiling against a jet, each focusing on a different p_T region. The first correction uses well-calibrated photons or Z-bosons (decaying into electron or muon pairs) to measure the p_T response of recoiling jets in the central region up to a p_T of 950 GeV. The second correction uses multijet systems to calibrate central, high- p_T jets ($300 < p_T < 2000$ GeV) recoiling against a collection of well-calibrated, lower- p_T jets. Although these calibrations are derived from central jets, their corrections can also be applied to forward jets whose energy scale has been equalised by the η -intercalibration procedure.

The in situ calibration response $R_{in-situ}$ is defined as the average ratio of the reconstructed jet p_T (or η , for the η -intercalibration) to the p_T (or η) of the reference object:

$$R_{in-situ} = \frac{p_T^{\text{jet, reco}}}{p_T^{\text{ref, reco}}}. \quad (3.10)$$

The calibration factor is derived by numerically inverting the ratio:

$$C = \frac{R_{in-situ}^{\text{data}}}{R_{in-situ}^{\text{MC}}}. \quad (3.11)$$

The calibration constants from in each of these analyses are statistically combined into a final in-situ calibration covering the p_T range $20 < p_T < 2000$ GeV.

A set of 125 systematic uncertainties arises from the JES calibration process. Since this number is larger than what is required by most ATLAS analyses, the set is reduced by combining the smaller components into grouped nuisance parameter (NPs), with minimal loss of correlation information. The set of JES uncertainties used in the stop search presented in this thesis is described in Section 6.2. The total JES uncertainty is shown in Figure 3.9 as a function of the jet p_T and η . Uncertainties from the η intercalibration analysis are labelled as *relative in situ JES*, while those from other in-situ measurements are combined into the *absolute in situ JES* term. For low- p_T jets, the largest contributions come from pile-up subtraction and jet flavour response, while for high- p_T jets, the dominant uncertainty arises from the absolute in-situ calibrations.

The **Jet Energy Resolution** (JER) can be parametrised using Equation 2.9. To measure the JER, the jet momentum has to be measured precisely, and this is achieved using data samples of well-measured dijet events. Additionally, data samples collected using random, unbiased triggers are used to estimate the noise contribution from pile-up, while the contribution from the electric noise is evaluated using MC samples with $\mu = 0$ (no pile-up). These measurements are




Figure 3.9 – Fractional JES systematic uncertainty components for anti- k_t small- R jets reconstructed from Pflow objects, as a function the jet p_T at $\eta = 0$ (left), and as a function of the jet η at $p_T = 60$ GeV. The total uncertainty (shown as a filled region topped by a solid black) is determined as the quadrature sum of all components [103].

combined to determine the final JER results. To ensure that the JER in MC simulation matches that observed in data, a smearing correction is applied to simulated jets in the jet- p_T regions where the MC resolution in data is better than in data, until the average jet resolution in MC matches that in data. Uncertainties on the JER are propagated to physics analyses as systematic uncertainties (see Section 6.2).

3.4.3 b -tagging

The SUSY search presented in this thesis is characterised by final states with jets initiated by b -quarks (b -jets). Therefore, identifying these jets is fundamental to suppress the background contributions from jets initiated by charm quarks, or by lighter quarks (u , d , s) or gluons (collectively referred to as *light-jets*).

The hadronisation of bottom quarks results in the production of b -flavoured hadrons (mainly B mesons), which decay via weak interactions generating the reconstructed b -jets. Due to their relatively long lifetime of approximately 1.6 ps, B mesons can travel a measurable distance before decaying. For example, a B meson with a transverse momentum of 50 GeV typically has an average decay length of about $\langle l \rangle = \gamma v \tau \approx 3$ mm, where v is the meson's velocity and γ is the Lorentz relativistic factor. This results in the presence of a displaced secondary vertex within the jet, corresponding to the decay point of the B meson. This characteristic provides a distinctive signature to distinguish b -jets from other types of jets.

To identify b -jets, ATLAS uses advanced b -tagging algorithms, exploiting typical features of B meson decays, such as:

- **Displaced secondary vertices:** the presence of a secondary vertex, displaced by few millimetres from the primary interaction point, is a distinctive signature of a B meson decay;
- **Large impact parameters:** decay products from B mesons tend to generate tracks with large impact parameters relative to the primary vertex;
- **Soft leptons:** approximately 35% of B mesons decay semileptonically, producing low-energy, non-isolated leptons that can be found within the jet cone;

- **High particle-multiplicity:** B meson decay chains tend to produce more stable hadrons than those from charm or light hadrons, leading to a higher particle multiplicity within b -jets.

These features are used by advanced multivariate techniques that combine multiple observables into a single discriminant variable, designed to distinguish b -jets from light-flavour jets. The b -tagging algorithm used in the analysis presented in this thesis is the new transformer-based **GN2 algorithm** [104], which has been adopted as the standard in ATLAS analyses for Run 3.

The GN2 tagger is based on a single *transformer model* [105] that processes both track and jet information, and is trained using truth labels from MC simulations. While the main goal of the training is to predict jet flavour, the model also trained to reconstruct the jet's internal structure by grouping tracks that come from the same decay point and predicting the physics process from which a track originates. This additional information is incorporated into a combined loss function, which allows the model to be optimised simultaneously, giving it a flexible structure so that it can be easily retuned for different experimental conditions or specific physics goals. Figure 3.10 provides a schematic view of the GN2 transformer model.

Jets are classified using a single discriminant D_b , which combines the algorithm's output prob-




Figure 3.10 – Illustration of the GN2 algorithm. Before being passed through the model, the features of the jets are concatenated with the features of their associated tracks, allowing them to be processed simultaneously [104].

abilities of a jet being a b -jet (p_b), c -jet (p_c), τ -jet (p_τ) or a light-jet (p_u) and is defined as:

$$D_b = \log \left(\frac{p_b}{f_c p_c + f_\tau p_\tau + (1 - f_c - f_\tau) p_u} \right), \quad (3.12)$$

where the free parameters ($f_{c/\tau}$) are weights determined through an optimisation procedure aimed at maximising light-jets background rejection. A selection on the D_b discriminant determines the corresponding b -jet tagging efficiency and background rejection power. Five fixed-efficiency WPs are implemented, corresponding to inclusive b -jet selection efficiencies of 65%, 70%, 77%, 85%, and 90%. Additionally, a pseudo-continuous b -tagging option is available,

where each jet is assigned an integer score corresponding to the most stringent efficiency threshold it meets. The standard WP for ATLAS analyses is the 77% efficiency, and it is also used in the analysis presented in this thesis. Figure 3.11 shows tagger performance in terms of the c -jet and light-jet rejection as a function of b -jet tagging efficiency. Here the term light-jet refers to the jets initiated by u , d , and s quarks. The performance is evaluated on $t\bar{t}$ MC samples and is compared to the performance of the previous taggers, *DL1d* [106] and *GN1* [107].




Figure 3.11 – The c -jet and light-jet rejections as a function of the b -jet tagging efficiency for jets in a $t\bar{t}$ with a p_T selection of $20 < p_T < 250$ GeV. The ratio with respect to the performance of the *DL1d* algorithm is shown in the bottom panels. The shaded regions represent 68% confidence intervals calculated according to a binomial distribution [108].

3.5 Missing transverse momentum

As discussed in Section 1.2.3, protons in the LHC are assumed to carry momentum in the direction parallel to the beam-axis. Consequently, momentum conservation implies that the vector sum of the transverse momenta of all particles originating from the primary hard-scatter vertex is exactly zero. Therefore, any imbalance in the total reconstructed transverse momentum of an event can indicate the presence of invisible (weakly interacting) particles in the final state, such as neutrinos or other particles predicted by BSM theories. Thus, the precise measurement of this imbalance is fundamental for BSM searches, including the search for stop pair production, which considers final states containing neutrinos and neutralinos. This momentum imbalance is quantified by the *missing transverse energy* (E_T^{miss}), an event-level quantity whose components $E_{x(y)}^{\text{miss}}$ are calculated as the negative vectorial sum of the transverse momenta of all the reconstructed objects in each event:

$$E_{x(y)}^{\text{miss}} = - \sum_{\text{electrons}} E_{x(y)}^e - \sum_{\text{photons}} E_{x(y)}^\gamma - \sum_{\text{taus}} E_{x(y)}^\tau - \sum_{\text{muons}} E_{x(y)}^\mu - \sum_{\text{jets}} E_{x(y)}^{\text{jet}} - E_{x(y)}^{\text{soft}}, \quad (3.13)$$

where $E_{x(y)}^i$ represents the calibrated energy of the corresponding physics object projected along the x (y) axis, and $E_{x(y)}^{\text{soft}}$ refers to the contributions from tracks with $p_T > 0.5$ GeV and $|\eta| < 2.5$, originating from the primary vertex but not associated with any reconstructed object. A more detailed description of each component is provided below.

- $E_{x(y)}^{\text{miss},e} = -\sum_{\text{electrons}} E_{x(y)}^e$ is reconstructed from the calibrated energies of the clusters associated with electrons having $p_T > 10$ GeV in the pseudorapidity regions $|\eta| < 1.37$ and $1.52 < |\eta| < 2.47$ (to exclude the transition region between the barrel and endcap calorimeters).
- $E_{x(y)}^{\text{miss},\gamma} = -\sum_{\text{photons}} E_{x(y)}^\gamma$ is reconstructed from the calibrated energies of the clusters associated with photons having $p_T > 25$ GeV in the pseudorapidity regions $|\eta| < 1.37$ and $1.52 < |\eta| < 2.47$.
- $E_{x(y)}^{\text{miss},\tau} = -\sum_{\text{taus}} E_{x(y)}^\tau$ includes the contribution of hadronically-decaying τ leptons with $p_T > 20$ GeV in the pseudorapidity regions $|\eta| < 1.37$ and $1.52 < |\eta| < 2.47$.
- $E_{x(y)}^{\text{miss},\mu} = -\sum_{\text{muons}} E_{x(y)}^\mu$ includes the contribution of muons with $p_T > 10$ GeV within $|\eta| < 2.7$;
- $E_{x(y)}^{\text{miss,jets}} = -\sum_{\text{jets}} E_{x(y)}^{\text{jet}}$ is built from standard Pflow jets reconstructed using the anti- k_t algorithm with radius parameter $R = 0.4$, and with $p_T > 20$ GeV. A requirement on the jet vertex tagger variable JVT [109] ($JVT > 0.64$) is imposed to jets with $|\eta| < 2.47$ and $p_T < 50$ GeV, to remove pile-up jets. The tracks associated to jets failing this requirement are included in the $E_{x(y)}^{\text{miss,soft}}$ term.
- $E_{x(y)}^{\text{miss,soft}} = -E_{x(y)}^{\text{soft}}$ represents the *soft term*, reconstructed from tracks not associated with any of the previously selected objects. Only tracks with $p_T > 0.4$ GeV and $|\eta| < 2.5$ that are compatible with the primary vertex are included.

The E_T^{miss} is calculated from its x and y components as:

$$E_T^{\text{miss}} = \sqrt{(E_x^{\text{miss,term}})^2 + (E_y^{\text{miss,term}})^2}. \quad (3.14)$$

Since ATLAS reconstructs physics objects using independent procedures, the same calorimeter signal, for example, may be used in the reconstruction both an electron and a jet, potentially introducing double counting when calculating the E_T^{miss} . To prevent this, the contributions to the E_T^{miss} are added in the specific order defined in Equation 3.13. Lower-priority objects are fully rejected if they share their calorimeter signal with a higher-priority already included the E_T^{miss} calculation (*signal ambiguity resolution*).

3.5.1 Missing transverse momentum significance

Non-zero values of E_T^{miss} can also arise from detector mismodelling, detector noise, limited detector coverage, or miscalibration of the reconstructed objects used in its calculation. Additionally, the resolution of the E_T^{miss} measurement is sensitive to pile-up effects, and generally

degrades as the levels of pile-up increase. The **object-based E_T^{miss} significance** [110] ($\mathcal{S}$), is a quantity designed to distinguish events where the reconstructed E_T^{miss} originates from weakly interacting particles from those where it is consistent with contributions due to detector effects or reconstruction inefficiencies. It is an event-level, likelihood-based discriminant, which takes into account all the reconstructed objects in an event along with their respective resolution uncertainties. For each event, it tests the hypothesis that the total transverse momentum carried by invisible particles ($\mathbf{p}_T^{\text{inv}}$) is equal to zero against the hypothesis that $\mathbf{p}_T^{\text{inv}}$ is different from zero. Considering a likelihood function $\mathcal{L}(E_T^{\text{miss}}|\mathbf{p}_T^{\text{inv}})$ of the parameter $\mathbf{p}_T^{\text{inv}}$ for a given value of measured E_T^{miss} , $\mathcal{S}$ is defined as a log-likelihood ratio as:

$$\mathcal{S}^2 = 2 \ln \left(\frac{\max_{\mathbf{p}_T^{\text{inv}} \neq 0} \mathcal{L}(\mathbf{p}_T^{\text{miss}}|\mathbf{p}_T^{\text{inv}})}{\max_{\mathbf{p}_T^{\text{inv}} = 0} \mathcal{L}(\mathbf{p}_T^{\text{miss}}|\mathbf{p}_T^{\text{inv}})} \right). \quad (3.15)$$

The event-per-event likelihood function is computed assuming that the measurement of each reconstructed object entering the E_T^{miss} calculation is independent of the others, that the true transverse momentum absolute value associated with the hard physics process is exactly equal to that of $\mathbf{p}_T^{\text{inv}}$, and that the difference between the measured and the true transverse momentum has a Gaussian distribution. Under these assumptions, the object-based E_T^{miss} significance from Equation 3.15 can then be written as:

$$\mathcal{S} = \frac{E_T^{\text{miss}}}{\sqrt{\sigma_L^2(1 - \rho_{LT}^2)}}, \quad (3.16)$$

where σ_L is the longitudinal resolution of all objects in the event, and ρ_{LT} is the correlation factor of the longitudinal and transverse resolutions, with “longitudinal” and “transverse” referring to the direction of the E_T^{miss} vector ($\mathbf{p}_T^{\text{miss}}$). High values of $\mathcal{S}$ indicate that the E_T^{miss} observed in the event cannot be well explained by resolution inefficiencies alone, implying that the event is more likely to contain invisible objects. Figure 3.12 compares the discriminating power of the object-based E_T^{miss} significance, the E_T^{miss} , and the *event-based E_T^{miss} significance* (simply defined as $E_T^{\text{miss}}/\sqrt{\sum E_T}$) in events with jet multiplicity of two or more, that is the kind of events considered in the analysis presented in this thesis.




Figure 3.12 – Background rejection versus signal efficiency in simulated $Z \rightarrow ee$ and $ZZ \rightarrow eevv$ samples with a $Z \rightarrow ee$ selection and $E_T^{\text{miss}} > 50 \text{ GeV}$. The performance is shown for E_T^{miss} , event-based E_T^{miss} significance, and object-based E_T^{miss} significance as discriminants in events with two jets (left), and three or more jets (right). The lower panel of the figures shows the ratio of other definitions/event-based E_T^{miss} significance [110].

Part II

Search for top squark pairs

Overview and datasets

This part of the thesis describes a search for direct top squark pair production. Top squark searches have always been a crucial component of the ATLAS. They have a strong theory motivation: in fact, the top squark plays a fundamental role in the SUSY solution to the hierarchy problem if it has a mass smaller than a few TeV. Therefore, it is a candidate to be produced at the LHC. The stop decay mode considered in this analysis with 100% branching ratio is the $\tilde{t}_1 \rightarrow t\tilde{\chi}_1^0$ (two-body decay mode), which implies that the mass difference between the stop and the $\tilde{\chi}_1^0$ is larger than the mass of the top quark. The simplified model considered in this search assumes R -parity conservation, implying that the top squarks are produced in pairs and that the neutralino is the LSP, therefore it is stable and escapes the detector without interacting.

Final states presenting exactly two leptons are investigated, therefore considering only events where both the W bosons originating from the decay of the top quarks decay leptonically. The leptonic decay of the W boson has a lower branching ratio with respect to its hadronic decay (the branching ratio of $W \rightarrow \ell\nu$, where ℓ represents any lepton, is approximately 11%). Therefore, the sensitivity of this search is limited with respect to searches considering hadronic decays of the W bosons (one-lepton and zero-leptons final-states), due to the reduced event statistics. Nevertheless, the presence of two isolated, high-quality leptons in the final state provides a clean experimental signature, which benefits from low fake rates, allowing for effective suppression of reducible backgrounds. This advantage becomes particularly important in *compressed scenarios* (where the mass difference between the top squark and the neutralino is small), in which the decay products of the stop tend to have lower momentum, and the missing transverse energy is reduced compared to scenarios with a large mass splitting. The sensitivity of the zero- and one-lepton searches is reduced in these scenarios, as they typically rely on high- p_T objects and large E_T^{miss} for effective background suppression. Complementary searches with either one or zero leptons in the final state are also performed in ATLAS, and their result is planned to be combined with the result of this search, to maximise the sensitivity across the $m_{\tilde{t}_1} - m_{\tilde{\chi}_1^0}$ parameter space.

The investigated final state signature consists of two isolated leptons (electrons or muons), at least two jets, of which at least one is b -tagged, and missing transverse momentum. The targeted final state is shown in Figure 4.1.

The analysis strategy is based on the definition of different analysis regions: signal regions




Figure 4.1 – Feynman diagram representing the 2-body stop decay, with both top quarks decaying leptonically [111].

(SRs), control regions (CRs) and validation regions (VRs). **Signal regions** are defined to be enriched with signal events, through selection criteria optimised based on MC simulations (see Section 4.1.2). The aim of the selection is to extend the sensitivity of the analysis to the largest possible number of signal models. A new analysis strategy based on a Neural Network (NN) discriminant was developed for this analysis (see Section 5.1.2). This is the main difference with respect to the previous Run 2 analysis [111], whose event selection strategy was instead based on a combination of rectangular cuts on the individual kinematic distributions. The background contamination in the SR is estimated in two different ways, based on the type of SM background considered. The dominant background contribution to this search is expected to come from top-antitop pair production processes ($t\bar{t}$) and associated production of a Z boson and top-antitop pair, with the Z decaying invisibly into two neutrinos ($t\bar{t}Z \rightarrow \nu\nu$). The contribution of these backgrounds to the SR is estimated using a partially data-driven procedure, that consists in the definition of background-enriched **control regions**, required to be kinematically as close as possible to events passing the signal selections, but orthogonal to them, while maintaining high purity of the target background with low signal contamination. The normalisation of the yields of the target background is constrained with the observed data in the respective CR and extrapolated into the SRs, increasing the reliability of the background estimation. On the other hand, the background contamination of minor backgrounds is only estimated from the MC simulations. The quality of the background estimation through CRs is validated by extrapolating the obtained normalisation in dedicated **validation regions**. Once a reliable estimation of the number of signal and background events in the SRs is obtained, statistical tests are performed to assess if the observed data are better described by a background-only hypothesis, or by a background+signal hypothesis. If no excess of data with respect to the SM background is observed, the results can be interpreted in terms of exclusion limits on signal models using the CLs method (see Section 6.1).

Previous ATLAS searches for top squark pair production have set exclusion limits at 95% confidence level on the signal scenarios considered here, using pp collision data collected during

the Run 1 and Run 2 data-taking periods of the LHC, at centre-of-mass energies of $\sqrt{s} = 7\text{-}8\text{ TeV}$ and $\sqrt{s} = 13\text{ TeV}$, respectively. The analysis presented in this chapter is the first ATLAS search for direct stop production using 13.6 TeV collision data collected during Run 3. Compared to previous analyses, the results are expected to improve thanks to the larger dataset, the higher centre-of-mass energy of Run 3 collisions, advancements in reconstruction techniques (such as b -jet tagging and E_T^{miss} significance, see Chapter 3), and the development of a novel analysis strategy based on machine learning. These improvements are expected to enhance the analysis sensitivity across a broad range of signal models, including those with large stop masses and those with compressed mass spectra.

4.1 Data sets and simulated samples

This chapter describes the datasets used in the analysis, including both collision data and MC simulated samples. It summarises the characteristics of the dominant background processes and the strategy adopted for the simulation of signal events. A brief overview on the data quality requirements and MC simulation techniques employed in ATLAS is also provided.

4.1.1 Data

This analysis is based on pp collision data collected with the ATLAS detector during the Run 2 and the early phase of the Run 3 (2022-2023) LHC data-taking periods, corresponding to centre-of-mass energies of 13 TeV and 13.6 TeV, respectively. After applying standard data-quality requirements (see dedicated paragraph), the final dataset corresponds to a total integrated luminosity of 193 fb^{-1} , of which 140 fb^{-1} were collected during Run 2 and 53 fb^{-1} during Run 3.

Data quality requirements

Precise knowledge of the quality of the recorded data is essential for all physics analyses. The primary goals of the data quality (DQ) assessment are to identify and exclude data affected by detector-related problems that can compromise their reliability, and to document the relevant conditions in dedicated databases. Two types of DQ defects can be defined: *intolerable defects*, which indicate that the affected data cannot be used for physics analyses, and *tolerable defects*, which are logged only for information and bookkeeping purpose, since the affected data are still suitable for physics analyses. A *Good Runs List* (GRL) is compiled for all runs, including only the luminosity blocks¹ (LBs) not affected by intolerable defects. The average DQ efficiency of the ATLAS detector was approximately 95.6% considering the full Run 2; in Run 3, ATLAS collected data with an average DQ efficiency of 93.1% in 2022, and 96.5% in 2023 [112]. There are different versions of the GRL, based on specific analysis requirements. For this analysis, data are selected if they satisfy the GRL selection in which the LHC declared stable beams, all the sub-systems of the ATLAS detector were properly operating, and both the solenoid and toroid fields were at nominal conditions.

¹ A luminosity block, typically lasting about a minute, corresponds to a period of data-taking with stable experimental conditions (i.e. constant instantaneous luminosity, without any alteration in data-recording configuration).

Trigger Selection

Data events are selected by single-lepton triggers, accepting events containing at least one lepton, an electron or a muon. The p_T thresholds for these triggers vary across data-taking periods, with a minimum threshold of 24 GeV for electron triggers, and of 20 GeV for muon triggers. Tighter p_T requirements are applied in the lepton offline selection, to ensure that the trigger efficiency is in the *plateau region* (i.e., the p_T -region in which the trigger efficiency is maximal and stable, see Section 7.1). This approach simplifies the correction for trigger efficiency, avoiding complications due to fluctuations in the turn-on region of the trigger response. Tables 4.1 and 4.2 list the single-electron and single-muon triggers used to collect the data, broken down by data-taking year. The lowest p_T -threshold electron triggers, such as the 2022 trigger HLT_e26_lhtight_1varloose_L1EM22VHI, typically include tight identification criteria, to maintain high sample purity while also controlling the trigger rates. To increase the single electron trigger efficiency at higher lepton p_T , additional triggers with higher p_T thresholds and looser selection requirements are included. As an example, in 2022, the HLT_e60_lhmedium_L1EM22VHI and HLT_e140_lhloose_L1EM22VHI triggers are used, applying p_T thresholds of 60 GeV and 140 GeV with medium and loose identification requirements, respectively. A similar strategy is applied to both electron and muon triggers across all data-taking years, to maximise the sensitivity throughout the relevant kinematic phase-space.

This trigger strategy differs with respect to the one used in the previous Run 2 analysis [111], which relied on di-lepton triggers, selecting events with two leptons (ee , $\mu\mu$, or $e\mu$) in the final state. Di-lepton triggers allow to select lower- p_T events (e.g. the lowest p_T -threshold di-lepton trigger of 2022 has a p_T threshold of 17 GeV for both electrons), which is particularly advantageous in the compressed region of the signal parameter space, where the decay products of the stop tend to have lower momenta. However, at the time of writing, the scale factors used to correct the difference in the trigger efficiency between data and simulated samples are not yet available for di-lepton triggers, while they are available for single-lepton triggers. Different studies were performed to prove that the acceptance loss due to the use of single-lepton triggers instead of di-lepton triggers is limited, and that the sensitivity of the analysis is not significantly affected by this choice, while the MC modelling visibly improves when applying the trigger scale factors.

4.1.2 Simulated samples

Simulated event samples are used for SM background estimations and to model signal processes.

Monte Carlo samples generation

To compare theoretical predictions with experimental data, physical models must be expressed in terms of observable quantities. This requires the simulation of pp collision processes to accurately predict their signatures in the detector. Such simulations are essential in guiding the optimisation of the analysis strategy in searches for new physics, as they allow the development of reliable models of both the background-only hypothesis, including only SM processes, and the background-plus-signal hypothesis, which accounts for potential contributions from BSM physics. Monte Carlo (MC) event generators are used to produce simulated events based on the

	Single electron	Single muon
2015	HLT_e24_lhmedium_L1EM20VH HLT_e60_lhmedium HLT_e60_lhmedium_nod0 HLT_e60_medium HLT_e120_lhloose HLT_e140_lhloose_nod0	HLT_mu20_iloose_L1MU15 HLT_mu40 HLT_mu50 HLT_mu60_0eta105_msonly
2016	HLT_e24_lhmedium_L1EM20VH HLT_e24_lhtight_nod0_ivarloose HLT_e60_lhmedium HLT_e60_lhmedium_nod0 HLT_e60_medium HLT_e120_lhloose HLT_e140_lhloose_nod0 HLT_e300_etcut	HLT_mu20_iloose_L1MU15 HLT_mu26_ivarmedium HLT_mu40 HLT_mu50 HLT_mu60_0eta105_msonly
2017	HLT_e26_lhtight_nod0_ivarloose HLT_e60_lhmedium_nod0 HLT_e140_lhloose_nod0 HLT_e300_etcut	HLT_mu26_ivarmedium HLT_mu50 HLT_mu60_0eta105_msonl
2018	HLT_e26_lhtight_nod0 HLT_e26_lhtight_nod0_ivarloose HLT_e60_lhmedium_nod0 HLT_e140_lhloose_nod0 HLT_e300_etcut	HLT_mu26_ivarmedium HLT_mu50 HLT_mu60_0eta105_msonly HLT_mu80_msonly_3layersEC

Table 4.1 – Triggers used to select Run 2 data

	Single electron	Single muon
2022	HLT_e26_lhtight_ivarloose_L1EM22VHI HLT_e60_lhmedium_L1EM22VHI HLT_e140_lhloose_L1EM22VHI HLT_e140_lhloose_noringer_L1EM22VHI HLT_e300_etcut_L1EM22VHI	HLT_mu24_ivarmedium_L1MU14FCH HLT_mu50_L1MU14FCH HLT_mu60_0eta105_msonly_L1MU14FCH HLT_mu60_L1MU14FCH HLT_mu80_msonly_3layersEC_L1MU14FCH
2023	HLT_e26_lhtight_ivarloose_L1eEM26M HLT_e60_lhmedium_L1eEM26M HLT_e140_lhloose_L1eEM26M HLT_e140_lhloose_noringer_L1eEM26M HLT_e300_etcut_L1eEM26M	HLT_mu24_ivarmedium_L1MU14FCH HLT_mu50_L1MU14FCH HLT_mu60_0eta105_msonly_L1MU14FCH HLT_mu60_L1MU14FCH HLT_mu80_msonly_3layersEC_L1MU14FCH

Table 4.2 – Triggers used to select Run 3 data.

underlying theoretical calculations.

In the first stage of simulation, the MC event generators model the initial, intermediate and final states of a physics process using matrix element calculations. This stage, also referred to

as the *hard scattering*, is based on perturbative QCD. For processes occurring in pp collisions, it is necessary to include the proton PDFs. As discussed in Section 1.2.3, PDFs cannot be derived from first principles and must instead be determined from experimental data. To achieve the most accurate description, results from multiple past and current experiments are combined into various PDF sets.

The following step simulates the QCD processes of parton showering and hadronisation described in Section 1.1.2, modelling how quarks and gluons produced in the hard scattering interaction evolve into observable hadrons and jets. Initial- and final-state radiation are also simulated at this stage. All these steps can be performed within the *Athena* software framework [113], which interfaces with various third-party tools (see details in next sections). These external packages perform one or several specific steps of the simulation chain described above.

The steps described above result in a collection of four-vectors representing all stable particles after hadronisation, useful for studying physics processes at the so-called *particle-level*. However, to make predictions that can be directly compared with observed data, the interaction of these particles with the detector and its response also need to be simulated. This is achieved using a detailed GEANT4 [114] model of the ATLAS detector, which includes not only the active sensor components of each sub-detector, but also the simulation of inactive material, such as support structures, cabling, and services. The simulation also includes information on the detector alignment and run conditions, retrieved from a dedicated conditions database. Custom algorithms are developed for each of the ATLAS subdetectors to convert the energy depositions simulated by GEANT4 into detector hits. These hits are then processed by digitisation algorithms to produce simulated detector signals. At this stage, the digitised signals from minimum bias events used to model pile-up are overlaid onto those of the simulated hard-scatter event. The result of the simulation is a set of simulated digitised output signals which can be treated by the same reconstruction algorithms used for the processing of recorded data, including the trigger. This provides *reconstruction-level* information, directly comparable to that obtained from real collision data.

The final results of the simulation process allow the prediction of cross sections and distributions of the final-state observables, which can be directly compared with experimental measurements.

Background simulation

As illustrated in Figure 4.1, the final state of stop pair production is characterised by the presence of a top-antitop pair ($t\bar{t}$) and a significant amount of missing transverse momentum (E_T^{miss}) arising from the two neutralinos. This analysis specifically targets the dileptonic decay channel, in which both the W bosons produced in the decay of the top quarks decay leptonically ($t \rightarrow W + b \rightarrow \ell\nu + b$). As a result, the targeted signal events present final states characterised by two leptons with opposite electric charge (electrons or muons), two b -jets, and missing transverse momentum. This signature can also be produced by several SM processes. The expected dominant SM background contributions are outlined below.

- **Top-antitop quark pair production ($t\bar{t}$)** with both top and antitop decaying leptonically. It is the dominant background process in this search. These events can be distinguished from the signal because they typically have lower E_T^{miss} in the final state, as they do not

contain neutralinos. However, this distinction becomes less effective for signal models with small mass splitting between the stop and neutralino, which also result in reduced E_T^{miss} , making background rejection more challenging in this region of parameter space.

- $t\bar{t}Z \rightarrow \nu\nu$ consists of production of a Z boson in association with a top-antitop quark pair, with the Z decaying invisibly into two neutrinos. Despite its relatively low cross section, this process constitutes one of the dominant backgrounds due to its identical final-state signature to the signal, as the E_T^{miss} coming from the Z decay mimics the neutralinos. This process is very difficult to discriminate against, as it has exactly the same signature of the SUSY events, due to the E_T^{miss} coming from the Z decay that mimics the neutralinos. This problem makes the distinction of signal models with small mass splitting between the stop and the neutralino particularly challenging, as they tend to have lower E_T^{miss} in the final state.
- $t\bar{t}Z \rightarrow \ell\ell$ and $t\bar{t}Z \rightarrow qq$ consists of production of a Z boson in association with a top-antitop quark pair, with Z decaying leptonically or fully hadronically, respectively. These processes contribute less significantly to the analysis SM background. In the $t\bar{t}Z \rightarrow \ell\ell$ processes, the two leptons originate from the Z rather than the top quarks, and therefore they can be effectively rejected using the di-lepton invariant mass. The $t\bar{t}Z \rightarrow qq$ process is more kinematically similar to the signal, emulating $t\bar{t}$ events with additional jets, but its impact is limited by the low cross section.
- **Single-top production**, with a top produced in association with a W boson. These events contain only one b -tagged jet, and can mimic the signal events where one of the two b -jets is not reconstructed.
- **Z + jets**, with the Z boson decaying into two leptons. Similarly to the $t\bar{t}Z \rightarrow \ell\ell$ process, the two leptons originate from the Z rather than from top decays, allowing this background to be suppressed using the di-lepton invariant mass.
- **Di-boson production (VV)** gives a small contribution to the SM background in this analysis. It includes the production of pairs of W bosons both decaying leptonically, or ZZ/WZ pairs with one Z boson decaying into two leptons.
- Other minor sources of background include rare processes such as tri-boson production, three- and four-top events, associated production of $t\bar{t}$ with additional bosons (e.g. WW , WZ , W , ZZ or Higgs bosons), and of a top quark with a Z boson or a Z and a W boson (tZ , tWZ). In the following, these backgrounds will be collectively referred to as *other*.

A variety of MC generators is used to model the SM background processes described above. The generator configurations are largely consistent between the Run 2 and Run 3 MC samples, with the latter using updated generator versions for some processes. Table 4.3 provides an overview of the MC samples used for the simulation of the dominant background processes of this analysis, specifying the matrix element and parton shower generators, the perturbative order in the strong coupling constant α_s used for cross section normalisation, and the PDF sets applied.

The background contribution arising from jets incorrectly identified as leptons (*fake leptons*) and

Physics process	Generator	Parton shower	Normalisation	PDF (generator)	PDF (PS)
$t\bar{t}$	POWHEG BOX v2 [115]	PYTHIA 8.230 [116]	NNLO+NNLL	NNPDF3.0NLO [117]	NNPDF2.3LO [118]
$t\bar{t}Z$	MADGRAPH5_AMC@NLO 2.3.3 [119]	PYTHIA 8.210 [116]	NLO	NNPDF3.0NLO	NNPDF2.3LO
Single-top (Wt)	POWHEG BOX v2 [120]	PYTHIA 8.230	NLO+NNLL	NNPDF3.0NLO	NNPDF2.3LO
$Z/\gamma^* (\rightarrow \ell\ell) + \text{jets}$	SHERPA 2.2.11, 2.2.14 [121] (SHERPA 2.2.14)	SHERPA 2.2.11, 2.2.14 [122] (SHERPA 2.2.14)	NNLO	NNPDF3.0NNLO [117]	NNPDF3.0NNLO
Di-boson VV	SHERPA 2.2.11, 2.2.12 [121] (SHERPA 2.2.14, 2.2.16)	SHERPA 2.2.11, 2.2.12 [122] (SHERPA 2.2.14, 2.2.16)	LO-NLO	NNPDF3.0NNLO	NNPDF3.0NNLO

Table 4.3 – Simulated background samples with the corresponding matrix element and parton shower (PS) generators, cross section order in α_s used to normalise the event yield, the generator PDF sets used. When different, the settings used for the simulation of Run 3 samples are mentioned in parentheses. Abbreviations used are defined as: leading-order (LO), next-to-leading-order (NLO), next-to-next-to-leading-order (NNLO), next-to-leading-logarithmic (NLL), next-to-next-to-leading-logarithmic (NNLL).

from non-prompt leptons, which originate from secondary processes such as hadron decays or photon conversions occurring inside jets, is estimated using MC simulation, and collectively referred to as *MC fakes* background. The so-called lepton charge-flip background is also included in this category. Electron charge-flip occurs when a photon emitted by bremsstrahlung converts into an electron-positron pair of, resulting in the reconstruction of a track with a charge different from the original electron, or when a highly energetic electron leaves a nearly straight track that is incorrectly reconstructed. Although this background constitutes a relatively small fraction of the total SM background, it is not negligible, primarily due to the low energy of the leptons considered in the analysis. The estimation of this background is performed using the particle-level information from the MC simulation to determine the true origin of each lepton. A reconstructed lepton is classified as fake or non-prompt if it is not associated to a particle-level lepton coming from the decay of a top quark, a boson (W , Z , or Higgs), or a supersymmetric particle. To validate the modelling of this background in the simulation, a dedicated sample enriched in non-prompt and fake leptons is selected by requiring events to contain two leptons with the same electric charge, which is a configuration that is uncommon in prompt SM processes but occurs more frequently in events with fake or non-prompt leptons. In addition, the events are required to have $m_{\ell\ell}^{\ell\ell}$ greater than 50 GeV, and invariant mass of the two leptons ($m_{\ell\ell}$) greater than 50 GeV (see the definitions of $m_{\ell\ell}^{\ell\ell}$ and $m_{\ell\ell}$ in Section 5.1.1). Figure 4.2 shows the distribution of the leading lepton transverse momentum for data and MC simulated events passing this dedicated selection. The plots show good agreement between data and simulation, validating the simulation-based estimation for the fakes background.

Top squark production simulation

The SUSY signal samples for top squark pair production were generated from leading-order (LO) matrix elements with up to two extra partons using MADGRAPH5_AMC@NLO [119, 123], and the $\tilde{t}_1$ decays were simulated with MADSPIN [124]. The primary parameters defining the kinematics of the process are the masses of the $\tilde{t}_1$ and the $\tilde{\chi}_1^0$. All signal samples were generated assuming a 100% branching ratio into the specified final state. Parton-level events were generated using the NNPDF3.0NNLO [125] PDF set, interfaced to PYTHIA 8.312 [126] for two-body




Figure 4.2 – Leading lepton p_T distribution in a phase-space region enriched with fake and non-prompt leptons. The plot on the left shows the result for Run 2 samples, while the plot on the right shows the results for Run 3 samples. Only statistical uncertainties are shown.

decays and to Pythia PYTHIA 8.307 [126] for three-body decays, to model the parton showering, hadronisation, and underlying event, with parameters set according to the A14 tune [127]. Parton luminosities were evaluated using the NNPDF2.3LO PDF set. The decays of bottom and charm hadrons were performed by EVTGEN v1.7.0 [128]. Jet-parton matching was implemented following the CKKW-L prescription [129], with the matching scale set to one quarter of the $\tilde{t}_1$ mass. The generated samples were then processed through the full ATLAS detector simulation. Signal cross sections were calculated at the next-to-next-to-leading order (NNLO) in α_S , including the resummation of soft-gluon emission at next-to-next-to-leading-logarithmic (NNLO+NNLL) accuracy [29]. The nominal cross sections and the associated uncertainty were obtained using the PDF4LHC15 PDF set, following the recommendations in Reference [130]. Figure 4.3 shows the grid of signal mass points considered in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane for the simplified model under study. Both two-body and three-body $\tilde{t}_1$ decay modes are included. While the analysis strategy was optimised for two-body decay scenarios, its sensitivity was also evaluated for models with three-body decays.

As discussed in Section 1.2.2, the mass eigenstates of scalar superpartners in SUSY are formed through a mixing between their left- and right-handed interaction eigenstates. The precise mixing fraction of $\tilde{t}_R$ and $\tilde{t}_L$ in the $\tilde{t}_1$ eigenstate is an unknown parameter. Since the polarisation of the top quarks originating from the decay of the stops depends on the $\tilde{t}_L$ - $\tilde{t}_R$ mixing fraction, the consequent impact of this mixing on the event kinematics is studied. Figure 4.4 shows, as an example, the particle-level distributions of the leading lepton p_T , the E_T^{miss} , and the angular distance ΔR between the leading lepton and the leading b -jet, for stop samples considering a pp collision centre-of-mass energy of 13.6 TeV and different mixing configurations: 70% $\tilde{t}_R$ and 30% $\tilde{t}_L$, maximal mixing (50% $\tilde{t}_R$ and 50% $\tilde{t}_L$), pure $\tilde{t}_R$, and pure $\tilde{t}_L$. The decay chain $\tilde{t}_1 \rightarrow t\tilde{\chi}_1^0$ is forced in all cases. The baseline selection described in Section 4.2 is applied to the events used to produce these plots, and only events containing exactly two leptons in the final state are considered. The resulting kinematic distributions are mostly insensitive to the stop mixing configuration, except in the case of pure $\tilde{t}_L$. The mixing fraction used in the generation of the




Figure 4.3 – Generated signal points in the considered simplified model’s parameter space $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$. Blue points represent models with 2-body decay of the stop, while red points represent models with 3-body decay of the stop.




Figure 4.4 – Particle-level distribution of the leading lepton p_T (left), E_T^{miss} (centre) and the angular distance ΔR between the leading lepton and the leading b -jet (right) in stop samples with different $\tilde{t}_R$ and $\tilde{t}_L$ mixing, produced forcing the decay chain $\tilde{t}_1 \rightarrow t\tilde{\chi}_1^0$.

signal MC samples used in the analysis is 70% $\tilde{t}_R$ and 30% $\tilde{t}_L$, forcing the decay chain $\tilde{t}_1 \rightarrow t\tilde{\chi}_1^0$, to be consistent with previous stop analyses, which used this mixing percentage.

4.2 Object reconstruction

Final-state objects (electrons and muons, jets, including b -jets, and E_T^{miss}) are reconstructed following the procedures and configurations outlined in Chapter 3. This section provides a summary of the selections criteria applied in the analysis.

For electron and muon objects, two levels of selection are defined: **baseline** and **signal**. Baseline objects are required to satisfy looser selection criteria and are primarily used in the *overlap removal* procedure to resolve ambiguities between reconstructed objects. Signal objects are a subset of baseline objects that meet additional selection requirements, and are used as electron

and muon physics candidates in the following steps of the analysis. The selection for baseline electrons requires $E_T > 4.5 \text{ GeV}$, $|\eta| < 2.47$, loose likelihood-based identification with an additional requirement on B -layer hits, and impact parameter $|z_0 \cdot \sin(\theta)| < 0.5 \text{ mm}$. In addition to baseline, the signal selection for electrons includes tight likelihood-based identification, tight isolation, track significance $|d_0/\sigma(d_0)| < 5$. The baseline selection criteria for muons are: $p_T > 4 \text{ GeV}$, $|\eta| < 2.7$, medium identification, and impact parameter $|z_0 \cdot \sin(\theta)| < 0.5 \text{ mm}$. The additional requirements applied to select signal muons are the PflowTight isolation and track significance $|d_0/\sigma(d_0)| < 3$.

Jet objects are reconstructed and calibrated as described in Section 3.4. The reconstruction uses the anti- k_t algorithm with input Pflow objects and radius parameter $R = 0.4$. Jets are required to have $p_T^{\text{jet}} > 20 \text{ GeV}$ and $|\eta^{\text{jet}}| < 4.5$. To suppress pile-up collisions, jets with $p_T < 60 \text{ GeV}$ are requested to satisfy additional selections on the JVT discriminant [131], depending on their pseudorapidity, requiring a significant fraction of the tracks associated with each jet to have an origin compatible with the primary vertex. Jets resulting from the hadronisation of b -quarks are identified using the GN2 b -tagging algorithm at the 77% efficiency WP (see Section 3.4.3).

Finally, E_T^{miss} is reconstructed from calibrated electron, muon and jet, and photon objects, as well as the track soft term, as described in Section 3.5. The object-based E_T^{miss} significance variable is also reconstructed as described in Chapter 3.

Leptons and jets can produce similar signatures in the detector, and in some cases, the same detector signal may be reconstructed as both a lepton and a jet. To resolve such ambiguities, the **overlap removal algorithm** is applied to baseline electrons and muons, and to jet candidates. The algorithm follows a specific sequence of steps to remove the overlapping objects:

- jets close to leptons are removed, as they often originate from energy deposits from electron showers or muon bremsstrahlung in the calorimeter;
- calorimeter-tagged muons sharing an ID track with an electron, and electrons sharing an ID track with any remaining muons are removed;
- jets within a cone of radius $R < 0.2$ around an electron are removed if they are not b -tagged with a WP providing 85% b -tagging efficiency, or if the electron has $p_T > 100 \text{ GeV}$;
- electrons are removed in favour of jets if the distance ΔR between the two objects is $\Delta R < \min(0.4, 0.04 + 10/p_T^e)$, where p_T^e is the electron's transverse momentum;
- jets are removed in favour of muons if the distance between the two objects is $\Delta R < 0.2$, and they have less than three associated tracks;
- muons are removed in favour of jets if the distance ΔR between the two objects is $\Delta R < \min(0.4, 0.04 + 10/p_T^\mu)$, where p_T^μ is the muon's transverse momentum;
- electrons and muons close to any remaining jets are removed, to suppress non-prompt leptons originating from hadron decays.

As discussed in Section 4.1.2, MC simulated events are reconstructed using the same algorithms as real collision data. To correct for residual differences between data and simulation, such as in lepton reconstruction efficiency, energy scale, energy resolution, trigger response, b -tagging

efficiency, and pile-up conditions, event weights are applied to the simulation. These correction factors are derived from dedicated data-driven measurements (see [Chapter 3](#)).

Analysis strategy and background estimation

This chapter outlines the strategy adopted to construct a discriminant capable of separating signal and background events, with the goal of defining signal-enriched regions (signal regions) that maximise the sensitivity to the signal, across multiple signal models. A neural network is implemented for this purpose, exploiting low-level and high-level kinematic variables from the final-state objects to perform the classification task. The NN output serves as the main discriminating variable in the definition of the analysis regions. In particular, the signal regions, as well as the control and validation regions targeting the dominant $t\bar{t}$ background, are defined using selections on the NN score. For the $t\bar{t}Z$ background, control and validation regions are instead defined independently of the NN output, allowing these regions to be shared across different analysis channels. The first part of this chapter presents the input variables used for NN training, followed by a brief introduction to the machine learning concepts relevant to the model. The development, optimisation, and performance of the NN are then described. The second part of the chapter defines the signal, control, and validation regions used in the analysis.

5.1 Event selection

The event selection strategy adopted in this analysis is based on a dedicated Neural Network (NN), which is trained to discriminate signal events from background. This approach is significantly different from the analysis strategy used in the previous Run 2 analysis [111], which relied on simple rectangular cuts applied to discriminating variables. The NN model takes as input both low-level kinematic properties of the final-state objects, and a set of high-level variables that are known to provide good discriminating power between signal and background events. The definitions of these discriminating features are provided in Section 5.1.1. Section 5.1.2 then introduces some fundamental ML concepts in the context of the physics problem addressed by the NN model implemented for the analysis, followed by a description of the optimisation studies performed to reach the final NN model architecture, and an evaluation of its performance.

5.1.1 Discriminating variables

The kinematics of the signal events targeted in this analysis are influenced by several parameters, particularly the mass difference between the stop and the neutralino, $\Delta m = m_{\tilde{t}_1} - m_{\tilde{\chi}_1^0}$, in the considered signal model. Different variables can be defined to reconstruct the kinematics of the events, and these can be used as input features to the NN, to reach a good discrimination power between the signal and the SM backgrounds. This section introduces the main

discriminating variables employed in the analysis, including the mathematical definitions and a discussion on the types of background they contribute to discriminate against.

The quality of the modelling of these variables in MC simulation and their discriminating power is studied by comparing their distributions in data, and in signal and background simulated samples. All the distributions presented in this section are produced using events satisfying a loose preselection, designed to suppress part of the background contributions while preserving a very high signal efficiency. This selection requires events to have exactly two signal leptons (ee , $\mu\mu$, or $e\mu$) with opposite electric charge, with transverse momentum $p_T(\ell_1) > 27 \text{ GeV}$ and $p_T(\ell_2) > 20 \text{ GeV}$ for the leading and sub-leading leptons, respectively. Events must also contain at least two jets, of which at least one is b -tagged. Finally, the invariant mass of the two leptons $m_{\ell\ell}$ must be greater than 20 GeV , and the leptonic transverse mass, $m_{T2}^{\ell\ell}$, is required to be greater than 50 GeV (the definition of $m_{T2}^{\ell\ell}$ is provided later in this section). To illustrate the dependence of the variables' distributions on Δm , two representative signal benchmarks are included in the plots: one with $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (800, 500) \text{ GeV}$, corresponding to a small mass splitting, and another with $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (1000, 1) \text{ GeV}$, representing a large- Δm model. These plots are based on Run 2 samples only, as the distribution shapes are not significantly affected by the higher centre-of-mass energy of 13.6 TeV collisions.

E_T^{miss} and its significance

The final state signature of the top squark dileptonic decay mode targeted in this search is characterised by four invisible particles, two neutrinos and two neutralinos. These particles do not interact with the detector material, resulting in signal events characterised by the presence of large amounts of E_T^{miss} and E_T^{miss} significance (see definitions in Section 3.5). Figures 5.1a and 5.1b show the distributions of these variables in data compared to signal and background MC simulations. A clear separation between the top squark signal and the SM backgrounds can be observed, especially for signal models with large Δm values, where the stop decay products, in particular the neutralinos, carry more energy, enhancing the E_T^{miss} -related observables.

The trend observed in the data-to-MC ratio, particularly for E_T^{miss} distribution, is mainly attributed to a known mismodelling of the $t\bar{t}$ transverse momentum in the POWHEG BOX v2, which is taken into account in the $t\bar{t}$ theoretical systematic uncertainties (see Section 6.2), not included in these plots.

Particles masses

In this paragraph, several variables depending on the particles masses are presented.

The distribution of the **invariant mass** of the two leptons in the final state can be used to distinguish background events in which the two leptons originate from the decay of a Z boson, such as in the $t\bar{t}Z \rightarrow \ell\ell$ and in the $Z + \text{jets}$ processes. In such cases, the invariant mass distribution peaks at the Z mass $m_Z = 91.2 \text{ GeV}$ [20], as shown in Figure 5.2a.

The **transverse mass** from the missing transverse momentum and one of the leptons in the final state is defined as:

$$m_T = \sqrt{2E_T^{\text{miss}} p_T^\ell (1 - \cos \phi)}, \quad (5.1)$$




Figure 5.1 – Distribution of E_T^{miss} (left) and the E_T^{miss} significance (right) in the Run 2 dataset. The shaded area represents the statistical uncertainty on the total SM background. The signal normalisation is increased 250 times for visualisation purposes.

where p_T^ℓ is the transverse momentum of the lepton, which is assumed to be massless ($E \gg m$), and ϕ is the azimuthal angle between the lepton and the E_T^{miss} vectors in the transverse plane. The m_T distribution defined using a lepton and E_T^{miss} originating from the same parent particle, exhibits a characteristic kinematic end-point at the mass of the parent particle (e.g. at m_W for a $W \rightarrow \ell \nu$ decay). For SM background events, the E_T^{miss} only arises from the neutrinos, while neutralinos generate additional E_T^{miss} contributions for signal events, resulting in m_T distributions with end-points at higher values. This effect is more evident in signal models with large Δm values, where the neutralino is more energetic and the resulting E_T^{miss} is larger, as illustrated in Figure 5.2b.

The **stransverse mass** is a generalisation of the transverse mass to scenarios with two invisible particles in the final state. It is a kinematic variable designed to bound the mass of a pair of identical parent particles, each decaying into a visible and an invisible particle [132, 133]. Figure 5.3 shows the diagram of the generic process under consideration, presenting two visible particles and missing transverse momentum in the final state.




Figure 5.2 – Distribution of the invariant mass $m_{\ell\ell}$ of the leading and the sub-leading lepton (left), and distribution of the transverse mass between the E_T^{miss} and the leading lepton ℓ_1 (right) in the Run 2 dataset. The shaded area represents the statistical uncertainty on the total SM background. The signal normalisation is increased 250 times for visualisation purposes.




Figure 5.3 – Diagram of the generic process under consideration [132].

The stransverse mass is defined as:

$$m_{T,2}(\mathbf{p}_{T,1}, \mathbf{p}_{T,2}, \mathbf{p}_T^{\text{miss}}) = \min_{\mathbf{q}_{T,1} + \mathbf{q}_{T,2} = \mathbf{p}_T^{\text{miss}}} \{ \max[m_T(\mathbf{p}_{T,2}, \mathbf{q}_{T,1}), m_T(\mathbf{p}_{T,2}, \mathbf{q}_{T,2})] \}, \quad (5.2)$$

where m_T is the standard transverse mass from Equation 5.1, $\mathbf{p}_{T,1}$ and $\mathbf{p}_{T,2}$ are the transverse momentum vectors of two visible particles, and $\mathbf{q}_{T,1}$ and $\mathbf{q}_{T,2}$ represent the invisible particles' transverse momentum vectors, whose sum equals the observed missing transverse momentum: $\mathbf{p}_T^{\text{miss}} = \mathbf{q}_{T,1} + \mathbf{q}_{T,2}$. The minimisation is performed over all such possible decompositions of $\mathbf{p}_T^{\text{miss}}$ in $\mathbf{q}_{T,1}$ and $\mathbf{q}_{T,2}$. In these calculations, the invisible particles are assumed to be massless. Different variants of the stransverse mass are used as input features of the NN, computed using different pairs of visible particles, considering both leptons and jets in the final state. The most powerful version, already used as the main discriminating variable in the previous Run 2 analysis [111], uses the two leptons as the visible particles, i.e. $\mathbf{p}_{T,1}$ and $\mathbf{p}_{T,2}$ are the transverse

momentum vectors of the two leptons. This variable is denoted simply as $m_{T2}^{\ell\ell}$. For background processes like $t\bar{t}$, this quantity typically exhibits an endpoint close to the W boson mass, while in SUSY signals it is expected to reach higher values due to the presence of neutralinos. Figure 5.4a presents the distribution of $m_{T2}^{\ell\ell}$ in data and MC simulation. A clear dependence of the $m_{T2}^{\ell\ell}$ on Δm is visible when comparing the two signal benchmarks (one with $\Delta m = 300$ GeV, and one with $\Delta m = 1000$ GeV). The discrepancy between data and MC simulation observed around $m_{T2}^{\ell\ell} = 150$ GeV is due to a known MC mismodelling of the Z peak. These events are not part of the relevant analysis regions.

Other combinations of visible objects are used to compute the stransverse mass. One variant uses the two leading b -jets; others combine one or both leptons with the two leading b -jets (e.g. $m_{T2}(\mathbf{p}_T^{b_1} + \mathbf{p}_T^{\ell_1}, \mathbf{p}_T^{b_2})$ or $m_{T2}(\mathbf{p}_T^{b_1} + \mathbf{p}_T^{\ell_2}, \mathbf{p}_T^{b_2} + \mathbf{p}_T^{\ell_1})$): all possible combinations of leptons are considered. A further variant, developed to mitigate combinatorial ambiguities of the pairing of leptons and b -jets, is $m_{T2,\min}$, which corresponds to the minimum stransverse mass value obtained from a combination of leptons and b -jets when switching the two b -jets. Figures 5.4b, 5.4c, and 5.4d show the distributions of $m_{T2,\min}$ in all the three available combinations, which are:

- $m_{T2,\min}(b + \ell_1, b, E_T^{\text{miss}}) = \min[m_{T2}(b_1 + \ell_1, b_2), m_{T2}(b_2 + \ell_1, b_1)],$
- $m_{T2,\min}(b + \ell_2, b, E_T^{\text{miss}}) = \min[m_{T2}(b_1 + \ell_2, b_2), m_{T2}(b_2 + \ell_2, b_1)],$
- $m_{T2,\min}(b + \ell, b + \ell, E_T^{\text{miss}}) = \min[m_{T2}(b_1 + \ell_1, b_2 + \ell_2), m_{T2}(b_2 + \ell_1, b_1 + \ell_2)].$

In all cases, the leading light-jet is used in place of the sub-leading b -jet in events presenting only one b -jet in the final state.

The $p_{T,\text{boost}}^{\ell\ell}$ and $\Delta\phi_{\text{boost}}$ variables

Two variables are constructed to characterise the topology and the transverse momentum balance of the final-state objects [134]. The variable $\mathbf{p}_{T,\text{boost}}^{\ell\ell}$ is the vectorial sum of the missing transverse momentum vector and the leptons' transverse momentum vectors, defined as:

$$\mathbf{p}_{T,\text{boost}}^{\ell\ell} = \mathbf{p}_T^{\text{miss}} + \mathbf{p}_T(\ell_1) + \mathbf{p}_T(\ell_2) \quad (5.3)$$

Its magnitude, $p_{T,\text{boost}}^{\ell\ell}$, can also be interpreted as the magnitude of the vector sum of all the transverse hadronic activity in the event. The azimuthal angle between $\mathbf{p}_T^{\text{miss}}$ and $\mathbf{p}_{T,\text{boost}}^{\ell\ell}$ is denoted as $\Delta\phi_{\text{boost}}$. These quantities are useful for the discrimination of events where the non hadronic component (e, μ, ν and $\tilde{\chi}_1^0$) is collimated, which is the case of signal events, particularly higher- E_T^{miss} events produced by high- Δm signal models. These events tend to have higher values of $p_{T,\text{boost}}^{\ell\ell}$ and lower values of $\Delta\phi_{\text{boost}}$, as shown in Figures 5.5a and 5.5b, where the distributions of these variables are compared in data and in MC simulations.

Angular variables

The angular distance between particles in the final state has proven to be a useful quantity to discriminate the SM background processes from the top squark signal. The angular distance

ΔR between two objects a and b is defined as:

$$\Delta R = \sqrt{(\phi_a - \phi_b)^2 + (\eta_a - \eta_b)^2}, \quad (5.4)$$

as described Section 2.2. This quantity is included in the NN input features, computed between the two leptons ($\Delta R_{\ell\ell}$) and between the two b -jets, or the leading b -jet and the leading light-jet in events with only one b -tagged jet (ΔR_{bb}).

The azimuthal angular separation $\Delta\phi$ between final state objects is also considered among the discriminating variables. This quantity is computed between the two leptons and the two b -jets (or the leading b -jet and the leading light-jet in events with only one b -tagged jet), considering all the possible combinations. As an example, Figure 5.6 shows the distributions of $\Delta R_{\ell\ell}$ and $\Delta\phi$ between the two leptons, showing that these variables are particularly effective in the discrimination of background processes with two leptons originating from a Z decay (such as the $t\bar{t}Z \rightarrow \ell\ell$ and the $Z + \text{jets}$ backgrounds), where the leptons tend to be emitted at small angular separations.




Figure 5.4 – Distribution in Run 2 dataset of the stransverse mass m_{T2}^{ll} using the two leading leptons as visible particles (upper row), and distribution of $m_{T2,min}$ (bottom row) computed using the leading lepton (left), the sub-leading lepton (centre), and both leptons (right) in association with the two b -jets. The shaded area represents the statistical uncertainty on the total SM background. The signal normalisation is increased 250 times for visualisation purposes.




Figure 5.5 – Distribution of $p_{T,\text{boost}}^{\ell\ell}$ (left) and the $\Delta\phi_{\text{boost}}$ (right) in the Run 2 dataset. The shaded area represents the statistical uncertainty on the total SM background. The signal normalisation is increased 250 times for visualisation purposes.




Figure 5.6 – Distribution of the angular distance ΔR between the two leading leptons (left) and the two leading b -jets, or the leading b -jet and the leading light-jet in events with only one b -jet, (right) in the Run 2 dataset. The shaded area represents the statistical uncertainty on the total SM background. The signal normalisation is increased 250 times for visualisation purposes.

5.1.2 NN discriminant

Machine learning algorithms are powerful tools for analysing high-energy physics data. One of the problems typically addressed by ML models is the event classification, where the goal is to distinguish signal events from background processes. Neural networks, in particular, are well-suited for this task, due to their ability to model non-linear correlations among multiple input variables. Unlike traditional cut-based methods, which rely on manually defined selection criteria, NNs can automatically learn optimal threshold from data, making more efficient use of the available information. This section provides an overview of some fundamental concepts relevant to understanding the NN model implemented for the signal discrimination task in this analysis.

Neural networks for classification problems

Neural networks are a class of ML models, composed of interconnected units called *nodes* (or *neurons*), organized in *layers*: an input layer, one or more hidden layers, and an output layer. Each node in a layer receives numerical inputs from the previous layer, applies a weighted linear combination followed by a non-linear activation function, and passes the result to the next layer. The presence of a non-linear activation function allows NNs to approximate complex, non-linear functions. The NN used in this thesis is a *feed-forward* neural network, meaning that the information flows in a single direction, from the input layer, through the hidden layers, to the output layer, without forming any cycles. It is also *fully-connected*, which means that each node in a given layer is connected to all the nodes in the subsequent layer. The NN developed for this analysis is used in a *supervised learning* setting: the model is trained on a labelled dataset, where each input vector is associated with a known target output. During training, the network adjusts its internal weights to minimise a *loss function*, typically quantifying the discrepancy between the predicted and target outputs, via optimisation algorithms. The task of the NN is multiclass classification, where the goal is to assign each input event to several mutually exclusive classes based on a set of input features. The input layer consists of a set of meaningful kinematic variables, aimed at accurately describe the physics final state, while each output layer's node represents the probability of each processed event to belong to each class. Figure 5.7 provides a schematic illustration of a single neuron and the typical layered architecture of a neural network. Each input quantity x_i represents a feature associated to an event. The linear combination of the input features fed to the activation function for a given node can be expressed as:

$$z = \sum_{i=1}^N w_i \cdot x_i + b_i \quad (5.5)$$

where x_i are the input features, w_i are the corresponding weights, and b_i is the *bias* term, included to allow the model to shift the activation function independently of the input. Each input feature x_i is associated to a different weight for each node it enters. In a fully connected layer, each node receives input from all nodes in the previous layer, therefore the weights associated to the inputs of a layer with dimension m can be represented as a matrix of dimension $N \times m$, where N is the number of the input features (or the number of nodes of the previous layer).

An example of activation function is the *ReLU* (Rectified Linear Unit), which is the one used in




Figure 5.7 – Illustration of the computation performed inside a node of a neural (left) and three different types of layers (input, hidden, output) in a classical neural network.

the hidden layers of the NN designed for this analysis, and is defined as:

$$f(x) = \begin{cases} 0 & \text{for } x < 0 \\ x & \text{for } x \geq 0 \end{cases}. \quad (5.6)$$

Another example, typically used in the output layer of a multiclass classification NN, is the *softmax* function, which converts the output of the last layer into a probability distribution over the classes. It is defined as:

$$\hat{y}_c = \frac{\exp(z_c)}{\sum_{j=1}^{n_c} \exp(z_j)}, \quad (5.7)$$

where z_c is the output of the last layer for class c , and n_c is the total number of classes. The output $\hat{y}_c$ represents the probability that the input event belongs to class c , and the sum over all classes is equal to 1, ensuring that the outputs can be interpreted as probabilities.

During the training phase, a neural network learns to associate input features with target outputs by adjusting the values of weights and biases. When the final prediction is produced at the output layer, it is evaluated by quantifying its difference compared to the target label through a *loss function*. The loss function typically used for multiclass classification problems, that is the one chosen for the NN used in this analysis, is the *categorical cross-entropy*, which is defined as:

$$L(\theta, \mathbf{y}) = \sum_{i=1}^N \sum_{c=1}^{n_c} y_c^{(i)} \cdot \log(\hat{y}_c^{(i)}), \quad (5.8)$$

where θ represent the parameters of the model, that are the weights and biases; N is the total number of input samples, n_c is the number of output classes, $y_c^{(i)}$ is a binary indicator that equals 1 if the i^{th} sample belongs to class c , and 0 otherwise (called *one-hot encoding* of the true label), $\hat{y}_c^{(i)}$ is the predicted probability (output of the softmax function) for the i^{th} sample to belong to class c . Based on this discrepancy, the network updates its weights and biases through a process called *backpropagation*: the gradient of the loss function is computed with respect to each parameter, to understand how the loss function varies as a function of it, using the chain

rule of calculus, working backward from the output layer to the input layer (hence the name of the procedure). An *optimiser* uses the gradients of the loss function with respect to each parameter to determine how the parameters should be updated to decrease the value of the loss, with the size of the update determined by a value known as the *learning rate*. One commonly used optimiser is the so-called *Adam* (Adaptive Moment Estimation), which adjusts the learning rate for each parameter during the training, based on the history of gradients. This cycle is repeated over all the events in the input dataset and for many iterations, referred to as *epochs*, gradually improving the network's performance. At the end of each epoch, the reached performance is evaluated using a given metrics. A commonly used metric is the *accuracy*, that, for multiclass classifiers, is defined as the fraction of correctly classified events (i.e., those for which the class with the highest output probability matches the true label) over the total number of events in the training sample.

To properly evaluate and optimise the performance the NN, the available dataset is typically divided into three distinct subsets: a *training set*, a *validation set*, and a *test set*. The training set is used to fit the model parameters with the procedure exposed previously. During training, the validation set is used to monitor the model's performance on data on which the model is not trained, providing an estimate of its ability to generalise. Specifically, at the end of each training epoch, the model is evaluated on the validation set using the same loss function and performance metrics used during training (in the case of this analysis, categorical cross-entropy and accuracy). The validation loss is not used to update the model weights, but only serves as a tool to evaluate the model performance and guide the choice of the final architecture and hyperparameters. A well-performing model typically exhibits an increasing trend in accuracy, and decreasing trend in loss for both the training and validation sets throughout the training process. When the loss continues to decrease on the training set while beginning to increase on the validation set, the possible cause is *overfitting*. In such cases, the model begins to memorise the training data rather than learning generalisable patterns, resulting in low performance on unseen data. The hyperparameters of the model must be adjusted to achieve a configuration that does not lead to overfitting, while achieving high accuracy and low loss on both training and validation datasets. Finally, the test set is used only after the training is complete and the model architecture defined, to provide an unbiased estimate of the model's generalisation performance on unseen data.

When the available dataset has a limited size, as it is the case for some of the signal samples used for this analysis, an alternative to this splitting is the *cross-validation* technique. This approach consists in dividing the available dataset into multiple subsets, or *folds*, and to train and evaluate the model multiple times, each time using different folds for training, validation, and testing. In this analysis, a two-fold cross-validation strategy is adopted. The dataset is split into two equally sized, disjoint subsets. In each fold, one subset is further split for training and validation, while the other is used as a test set for model evaluation. The procedure is then repeated inverting the roles of the two subsets. Typically, the training set is not fed into the model all at once, but it is divided into smaller subsets called *batches*. Each batch contains a fixed number of events, and the model's weights are updated after processing each batch. The following properties, introduced earlier in this paragraph, are tuned during the training phase to maximise the performance of the network, and are referred to as *hyperparameters*.

- **Type of layers.** Depends on the type of problem addressed by the NN model. For the

classification problem considered in this thesis, a NN with fully connected (*dense*) is implemented.

- **Number of hidden layers and number of nodes per layer.** Deeper and wider networks have more parameters to optimise, therefore they can describe more complex models. However, it requires a big amount of training data to avoid overfitting, and significant computing resources. Therefore, the dimension of the network needs to be balanced considering these technical constraints.
- **Dropout rate.** Dropout is a regularisation technique in which a random fraction of nodes in the hidden layers is set to zero at each step, preventing the network from learning too much from specific nodes. Therefore, this improves the model's ability to generalise and prevents overfitting. The dropout rate defines the fraction of dropped-out nodes and must be tuned to mitigate overfitting without affecting the learning capability of the network.
- **Activation functions.** The activation functions used in this analysis, ReLU for the hidden layers and softmax for the output layer, are common choices for NNs composed of fully connected layers addressing multiclass classification problems. ReLU is commonly used in hidden layers for its computational efficiency and performance in deep networks. Softmax is used in the output layer of multiclass classifiers to ensure the outputs are interpretable as class probabilities.
- **Optimiser.** Adaptive optimisers, like the Adam optimiser used in this analysis, are preferred and commonly used for their ability to dynamically adjust learning rates for each parameter, which often leads to faster convergence and improved performance.
- **Learning rate.** A small learning rate ensures stable convergence but slows training, while a large one risks overshooting the loss function minimum. Even if the Adam optimiser is able to adjust the learning rate during the training, the initial value needs to be chosen carefully.
- **Loss function.** The choice of the loss function depends on the NN's task. For multiclass classification, categorical cross-entropy is the appropriate choice as it is able to quantify the distance between the predicted class probabilities and the true labels.
- **Number of epochs.** The number of training epochs affects the completeness of learning. It should be chosen to avoid both undertraining and overfitting, and is chosen by monitoring the validation metrics. Typically, the training process is stopped when the network's performance on the validation set stops improving, or even starts degrading, which is a sign of overfitting.
- **Batch size.** A smaller batch size leads to more frequent parameters updates, which may introduce noise in the gradient estimation. A larger batch size provides more stable gradient estimates, but requires more computational memory. Therefore, the batch size needs to be tuned based on the dataset size and the available computational resources.

The dimension of the NN architecture, dropout rates, learning rate, number of epochs and batch size chosen for the model designed for this analysis are detailed in the next section.

When the input variables of a neural network span significantly different numerical ranges, as it is the case for physical kinematic variables, their distributions are normalised to prevent features with larger numerical values from dominating the learning process. Several normalisation techniques exist, each suited to different data characteristics. In this analysis, input variables are scaled to the $[0, 1]$ interval using the so-called *min-max normalisation*, defined as:

$$x' = \frac{x - x_{\min}}{x_{\max} - x_{\min}}, \quad (5.9)$$

where x' is the normalised value of the feature, x is its original value, and $x_{\min}$ and $x_{\max}$ are the minimum and maximum values of the variable across the dataset, respectively. This approach preserves the original shape of the variable distributions, and is preferred over *standardisation* (i.e., centring to zero mean and unit variance) because the input variables in this case do not follow a Gaussian distribution.

In classification problems the training dataset may be *unbalanced*, meaning that it contains significantly more events for some classes with respect to others. In these cases, the network could be biased towards the most represented classes. To avoid this, weights can be assigned to each data point in the input dataset to have the same effective weight for all classes.

Since the NN model in this analysis is used to address a physical problem, it is important to understand its internal decision-making process, identifying which input features are most relevant for the final prediction. One commonly used approach for assessing feature importance in NNs is based on SHAP values (SHapley Additive exPlanations) [135]. This method is based on concepts from cooperative game theory, and explains the output of a model by assigning each input feature a *Shapley value* that reflects its contribution to the final decision. For each event x , the model prediction $f(x)$ for an individual class is expressed in the SHAP framework as:

$$f(x) = f_{\text{base}} + \sum_{i=1}^N \phi_i, \quad (5.10)$$

where f_{base} is the average model output for the considered class over all the training events, ϕ_i is the SHAP value of the i^{th} feature, representing how much the i^{th} feature contributed to moving the prediction away from the baseline, and N is the total number of input features. SHAP values are calculated by comparing the model's prediction when each feature is included or excluded, averaging over all the possible combinations of inputs. To evaluate which features contribute the most to the model's predictions, SHAP values are computed for all events in the dataset. The average absolute SHAP value of each feature across the full dataset then gives a measure of its importance. Input features with larger absolute mean SHAP values are interpreted as more relevant in the model's decision process.

Model optimisation

This section presents the results of the studies performed to optimise the performance of the NN model used for signal discrimination. The performance is quantified in terms of signal significance: the optimal model is capable of discriminating signal and background maximising

the signal significance at high NN score levels. The signal significance is computed following the prescription in Reference [136]:

$$Z_A = \sqrt{2 \left[(s+b) \ln \left(\frac{(s+b)(b+\sigma_b^2)}{b^2 + (s+b)\sigma_b^2} \right) - \frac{b^2}{\sigma_b^2} \ln \left(1 + \frac{\sigma_b^2 s}{b(b+\sigma_b^2)} \right) \right]} \quad (5.11)$$

where s is the number of signal events, b is the number of background events, and $\sigma_b = 0.2$ is a representative uncertainty value assigned to the number of background events to avoid overestimation of the significance value, based on the expected systematic uncertainty taken from the previous Run 2 analysis [111].

For the training of the NN model, signal events are split into three **event categories** based on Δm , since, as discussed in Section 5.1.1, the kinematics of signal events strongly depend on it. The splitting between categories is chosen to optimise the performance of the network for all the possible Δm values, and the resulting categories are the following:

- low- Δm signal models, including samples with $\Delta m = 200, 300$ GeV;
- medium- Δm signal model, including samples with $\Delta m = 300, 400, 500$ GeV;
- high- Δm signal model, including samples with $\Delta m \geq 600$ GeV;

The same background samples are used for the training of the models for all these three categories, while only signal samples with Δm in the correct range are considered for each model. Additionally, events are split based on the b -jet multiplicity: for each Δm range, one NN model is trained on events with exactly one b -tagged jet (named *1b* model), and one model is trained on events with more than one b -tagged jet (named *2b* model). For *1b* models, events are also required to have at least one light-jet in the final state.

The final number of models trained independently is six, and will be referred to as: **NN low-1b**, **NN low-2b**, **NN medium-1b**, **NN medium-2b**, **NN high-1b**, and **NN high-2b**.

A slightly different preselection is applied before training, based on the considered Δm range, due to the different kinematics of the regions. Training preselections for all models are summarised in Table 5.1.

The **input features** of the NNs include the kinematic information of the two leading leptons, E_T^{miss} , jets and b -jets in the final state. Specifically, models training on events with exactly one b -jet use the kinematic information of the b -jet and the leading light-jet, while models training on events with more than one b -jet use the kinematic information of the leading and the sub-leading b -jets. The input features are meant to provide a description of the final state kinematics as complete as possible, which the NN can exploit for the event classification. Additional event properties, mostly described by the variables introduced in Section 5.1.1, are provided to the NN to increase its discriminating power. The information provided by these features is mostly redundant kinematic information, that the NN should be able to infer from the low-level kinematic inputs in the limits of large training statistics and large network architecture. As the training statistics is limited, the addition of these features helps the NN in reaching a better classification power. The input variables are listed in Table 5.2. For models training on events

	Low Δm		Medium Δm		High Δm	
	1b	2b	1b	2b	1b	2b
$p_T(\ell_1)$ [GeV]				> 27		
$p_T(\ell_2)$ [GeV]				> 20		
$m_{\ell\ell}$ [GeV]				> 20		
$n_{b\text{-jets}}$	$= 1$	> 1	$= 1$	> 1	$= 1$	> 1
$n_{l\text{-jets}}$	> 0	≥ 0	> 0	≥ 0	> 0	≥ 0
$m_{T2}^{\ell\ell}$ [GeV]	> 80		> 90		> 100	

Table 5.1 – Preselection applied before training to the six NN models.

with only one b -tagged jet, the information of the leading light-jet is used instead of the information on the sub-leading b -jet (e.g. the angular distance $\Delta R(b_1, b_2)$ is not computed between the two leading b -jets, but between the leading b -jet and the leading light-jet). The number of b -jets is used as a discriminating variable only on models training on events with more than one b -jet.

Type	Variable	Description
Lepton-related kinematic variables	p_T, η and ϕ	Transverse momentum and angular variables of the leading and sub-leading leptons.
	Same- or different-flavour	If the two leptons are same-flavour (ee or $\mu\mu$) or different-flavour ($e\mu$).
	$m_{\ell\ell}$	Invariant mass of the two leptons.
	$p_{T\text{boost}}^{\ell\ell}$ and $\Delta\phi_{\text{boost}}$	Variables defined in Section 5.1.1.
Jet-related kinematic variables	b -tag score of the two b -jets	Pseudo-continuous b -tagging score of the two selected b -jets (see Section 3.4.3).
	p_T, η , and ϕ of the two b -jets	Transverse momentum and angular variables of the two selected b -jets.
	$n_{b\text{-jets}}, n_{l\text{-jets}}$	Number of b -tagged jets and number of light-jets.
Missing transverse momentum-related variables	$E_T^{\text{miss}}, \mathcal{S}$, and ϕ^{miss}	Missing transverse momentum, its significance, and its azimuthal direction.
Transverse and transverse mass	$m_T(E_T^{\text{miss}}, \ell_1)$ and $m_T(E_T^{\text{miss}}, \ell_2)$	Transverse mass between the E_T^{miss} and the two leading leptons.
	$m_{T2}(\ell, \ell), m_{T2}(b, b), m_{T2}(b + \ell, b), m_{T2}(b + \ell, b + \ell)$	Stranverse mass, in all the variants introduced in Section 5.1.1.
Angular-distance variables	$\Delta R(\ell_1, \ell_2)$ and $\Delta R(b_1, b_2)$	Angular separation between the two leptons or the two b -jets.
	$\Delta\phi(\ell, \ell), \Delta\phi(\ell, b), \Delta\phi(E_T^{\text{miss}}, \ell)$	Angular separation in ϕ between all the combinations of leptons and b -jets.

Table 5.2 – Input variables used for training the NN in all the Δm regions, only for models requiring more than one b -tagged jet. For models training on events with only one b -tagged jet, the information of the leading light-jet is used instead of the information on the sub-leading b -jet to define the listed variables.

As discussed previously, the implemented model is a **multiclass NN**, with a class representing

the signal and multiple classes used to identify different types of backgrounds.

As detailed in Section 4.1.2, the main SM background processes considered for this analysis are: $t\bar{t}$, $t\bar{t}Z$ (with the Z decaying either in two leptons, two quarks or two neutrinos, with the latter being the dominant background in signal region), Z + jets, single top (Wt), and di-boson events (WZ , ZZ , WW). In principle, a class could be created for each background process. However, this would lead to an overcomplicated model that could not reach a satisfying performance, given the limited number of input samples. Some background processes present similar final-state kinematics. One example are the $t\bar{t}$ and single top processes, both presenting the products of top quark decays in the final state, of the Z + jets and the $t\bar{t}Z \rightarrow \ell\ell$ processes, which have a final state presenting two leptons originating from the decay of a Z boson. The approach of merging backgrounds with similar kinematics in a single class was tried, but did not achieve a satisfying performance. The best approach resulted to be the use of a small number of classes, only trained on one background process. This gives enough discriminating power to also distinguish events from background processes that are not included in the training (the NN score of, e.g., single top events, will have a similar distribution to the one of $t\bar{t}$ events). For medium and high Δm models (NN medium-1b, NN medium-2b, NN high-1b, and NN high-2b), the following classes are included in the model:

- $t\bar{t}$
- $t\bar{t}Z \rightarrow \ell\ell$
- $t\bar{t}Z \rightarrow \nu\nu$
- Signal (only signal samples with Δm in the range targeted by the considered NN model, as described earlier).

For the low Δm models (NN low-1b and NN low-2b), the two $t\bar{t}Z$ classes are merged into a single class. This is necessary because the kinematic distributions of the $t\bar{t}Z \rightarrow \nu\nu$ events and of low- Δm signal events are too similar, and the network cannot reach sufficient discriminating power to distinguish them.

The output of the NN is, for each event, a vector of probabilities, that represent the probability of the event to belong to each class:

$$\{p_i\} \quad \text{for } i = 1 \dots N, \text{ with } \sum_{i=1}^N p_i = 1, \quad (5.12)$$

with $N = 3$ for low- Δm models, and $N = 4$ for medium- and high- Δm models.

For each NN model, the input features are normalised between 0 and 1 using the min-max normalisation procedure describe earlier. The minimum and maximum values of each feature (i.e. $x_{\max}$ and $x_{\min}$ from Equation 5.9) are extracted from the full training set, including both signal and background samples. These same values are used to scale the input features when the NNs are evaluated.

The NNs are trained using the *Keras* [137] library, with the *TensorFlow* [138] backend. The 1b (2b) NNs have an input layer with 44 (45) nodes, corresponding to the number of input variables described in Table 5.2. The NN hyperparameters are optimised independently for each model, applying a manually implemented grid search. In this approach, a discrete set of values

is defined for each hyperparameter, and the performance of the model is evaluated training it using different combinations. The dimension of the NN in terms of number of hidden layers and number of nodes per layer is progressively increased, until the model starts overfitting. To mitigate overfitting, the dropout rate is also progressively increased up to a maximum value of 0.6 for the first, larger hidden layer and 0.4 for the other hidden layers. If overfitting persists under these conditions, the NN size is subsequently reduced. The number of training epochs is selected based on the behaviour of the loss function on the validation set, stopping the training once the validation loss ceases to decrease. The training batch size is chosen based on computational constraints, as discussed in the previous paragraph. The learning rate is varied between 10^{-3} and 10^{-5} , and the ideal value was found to be 10^{-4} for all models. The final architecture selected for each model is summarised in Table 5.3.

Larger architectures can be implemented for medium- and high- Δm models thanks to the

Hyperparameter	NN low-1b	NN low-2b	NN medium-1b	NN medium-2b	NN high-1b	NN high-2b
No. of hidden layers	2	2	3	3	3	2
No. of nodes per hidden layer	(1024, 32)	(512, 64)	(1024, 128, 64)	(2048, 32, 32)	(1024, 128, 32)	(1024, 128)
No. of training epochs	150	150	150	150	150	150
Learning rate	0.0001	0.0001	0.0001	0.0001	0.0001	0.0001
Training batch size	128	128	256	256	256	256
Dropout value per hidden layer	(0.6, 0.4)	(0.6, 0.4)	(0.6, 0.4, 0.4)	(0.6, 0.4, 0.4)	(0.6, 0.4, 0.4)	(0.6, 0.3)

Table 5.3 – Hyperparameters values chosen for the six NN models.

larger amount of available input signal events. The 2b models have smaller architectures than the correspondent 1b models, as they tend to converge faster, and bigger networks would lead to overfitting.

The training dataset consists of events from both the Run 2 and Run 3 MC production campaigns, including $t\bar{t}$, $t\bar{t}Z$ signal simulated samples. For the signal class, events corresponding to several benchmark points in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ mass plane (restricted to the considered Δm range) are included simultaneously. Events from each signal point are assigned a weight, to balance the contribution of each signal benchmark during training. These weights are treated as tunable hyperparameters, and are selected to optimise the network's performance in regions of the parameter space near the exclusion limits obtained in the previous Run 2 analysis [111]. Additionally, all signal and background events are weighted using the scale factors for b -tagging, JVT, electron and muon reconstruction efficiencies, and pile-up (see Section 4.2). Finally, a normalisation weight is applied to ensure that all classes have the same total weight in the loss function, even if they are unbalanced in the number of events, as discussed previously.

Figures 5.8 and 5.9 show the distributions of the NN signal score for all the six models, evaluated on the Run 2 and the Run 3 datasets, respectively. All models achieve a good performance, with a clear separation between signal and background classes.

Figure 5.10 shows the SHAP values for each of the six models, that, as explained earlier, are used to rank feature importance in the model, also showing the contribution of each feature to the different classes. As expected, $m_{T2}^{\ell\ell}$ is within the first three most relevant features for all models. The other discriminating variables presented in Section 5.1.1 also show a high relevance in the NN decision process, with E_T^{miss} significance, E_T^{miss} , m_T , and the different variants of m_{T2} scoring very high for all models. Angular variables involving the two leptons, $\Delta R_{\ell\ell}$ and




Figure 5.8 – NN signal score distributions in Run 2 MC samples for all the NN models.




Figure 5.9 – NN signal score distributions in Run 3 MC samples for all the NN models.

$\Delta\phi$ between the two leptons, along with the invariant mass of the two leptons $m_{\ell\ell}$, show, as expected, high discriminating power for the $t\bar{t}Z \rightarrow \ell\ell$ class ($t\bar{t}Z$ class for low- Δm models). The interpretation of the NNs' decision process enabled by the use of SHAP values serves to validate the NN technique, demonstrating that the network is able to exploit the input variables in a physically meaningful way, consistent with expectations based on our current understanding




Figure 5.10 – SHAP absolute values, ranking the input features' importance for each of the six models. For all models, *class 0* represents the signal class and *class 2* represents the $t\bar{t}$ class. For the low- Δm models (first line), *class 1* represents the $t\bar{t}Z$ class; for the medium- and high- Δm models (second and third lines, respectively), *class 1* represents the $t\bar{t}Z \rightarrow \nu\nu$ class, and *class 3* represents the $t\bar{t}Z \rightarrow \ell\ell$ class. Plots for the 1b and the 2b models are shown in left and right columns, respectively.

of the signal and background processes.




Figure 5.11 – Illustration of analysis regions based on NN scores.

5.2 Analysis regions

As discussed in Chapter 4, the analysis strategy is based on the definition of distinct analysis regions: signal regions (SRs), designed to be enriched in potential signal events and optimise sensitivity to the targeted process; control regions (CRs), used for the estimation of the main SM background processes; and validation regions (VRs), designed to test the reliability of the background estimation derived from the CRs.

The dominant SM background contributions to the SRs are expected to originate from $t\bar{t}$, and $t\bar{t}Z$ with invisible decay of the Z boson ($t\bar{t}Z \rightarrow \nu\nu$). Therefore, dedicated CRs are defined for these two processes.

The SRs, along with the CRs and VRs targeting the $t\bar{t}$ background, are defined using the output scores of the NN described in Section 5.1.2. For these regions, separate selections are defined for each of the six NN models targeting the different event categories presented in the previous section. Moreover, the selections are optimised independently for the Run 2 and Run 3 datasets, in order to account for the different statistics available in the two periods. This is justified by the fact that the final statistical analysis is performed separately for the two datasets (see the dedicated Section 6.1). Events with the highest NN signal score values are used to define the SRs. In order to maximise sensitivity, the SRs are further binned in NN signal score. This means the SR is subdivided into intervals (or bins) based on the value of the NN output, with higher bins representing regions of increasing signal purity. This strategy enables the statistical analysis to exploit the shape of the NN score distribution. Events with NN signal scores just below the SR threshold are assigned to the $t\bar{t}$ VRs, and those with even lower scores constitute the $t\bar{t}$ CRs. In the $t\bar{t}$ CRs, an additional selection is applied to remove events with the lowest NN $t\bar{t}$ score, which tend to be dominated by non- $t\bar{t}$ background processes. A schematic representation of how these regions are defined in the two-dimensional plane of NN signal score versus NN $t\bar{t}$ score is shown in Figure 5.11.

The definition of the SRs is detailed in Section 5.2.1, while the one of the $t\bar{t}$ CRs and VRs is

presented in Section 5.2.2.

The definition of the control and validation regions for the background from $t\bar{t}Z \rightarrow \nu\nu$ events will be shared across all three 2-body analyses (0, 1 and 2 lepton final states) when the three channels will be combined. Therefore, it is not based on the NN output, as the model is specific to this analysis. The definition of the analysis regions for the estimation of the $t\bar{t}Z \rightarrow \nu\nu$ background is discussed in Section 5.2.3.

5.2.1 Signal region definition

The SR selection and binning in NN signal score is optimised to maximise the total expected signal significance. This optimisation is performed independently for each NN model and for each data-taking period (Run 2 and Run 3), resulting in twelve distinct SR definitions. The procedure follows the same steps for each model and data-taking period. It proceeds iteratively, starting from the most signal-like region of the NN signal score, where it is close to 1, and progressively moving toward the more background-like region, at lower values of the NN signal score. At each iteration, a new bin is defined by identifying the NN score threshold that maximises the expected cumulative signal significance, integrated from that threshold up to the lower edge of the previously defined bin (which is 1 for the uppermost bin). To ensure the validity of the asymptotic approximation [139] assumed in the statistical interpretation of the results (see Section 6.1), each bin is required to contain at least two expected background events. The iteration continues until the inclusion of further bins no longer improves the total expected significance of the SR. The significance used for this optimisation process is computed using signal points close to the exclusion limits in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane from the previous Run 2 analysis [111]. The NN signal scores thresholds defining the SR bins are reported in Table 5.4. In all categories, the corresponding training preselection (Table 5.1) is applied in addition to the selection on the NN signal score. The number of bins varies across the different Δm regions, reflecting differences in the shape of the NN output distributions (see Figures 5.8 and 5.9). For simplicity, the lower boundary of the last bin is aligned across all models within the same Δm region.

5.2.2 Estimation of the $t\bar{t}$ background

Similarly to the SRs, the $t\bar{t}$ **control regions** are defined based on the NN output and therefore, also in this case, twelve distinct selections are considered: two CRs are defined for each of the six NN models, one for the Run 2 dataset and one for the Run 3 dataset. The CR definitions are designed to select events as kinematically similar as possible to the corresponding SRs, while also suppressing the signal contamination. This is achieved by applying a selection on the same NN signal score used to define the SRs, defining a region orthogonal but close to the SR. The exact threshold applied to the NN signal score is defined individually for each CR, depending on the definition of the corresponding SR. A further selection is applied for each CR to the NN $t\bar{t}$ score, to enhance the $t\bar{t}$ background purity reducing the contamination from signal and other background processes. Figures 5.12 and 5.13 show the distributions of the NN $t\bar{t}$ score from the six NN models, for the Run 2 and Run 3 datasets, respectively. The distributions only include events satisfying the criteria of the training preselection for the appropriate Δm range (see Table 5.1) and the NN signal score selection of the $t\bar{t}$ CR targeting the corresponding model

Region	Bins in NN signal score
Low-1b-Run2	[0.93, 0.948, 0.96, 0.974, 0.986, 1]
Low-2b-Run2	[0.93, 0.947, 0.962, 0.974, 0.982, 1]
Low-1b-Run3	[0.93, 0.948, 0.96, 0.968, 0.977, 1]
Low-2b-Run3	[0.93, 0.946, 0.958, 0.966, 0.979, 1]
Medium-1b-Run2	[0.53, 0.7, 0.77, 0.85, 1]
Medium-2b-Run2	[0.53, 0.66, 0.72, 0.81, 1]
Medium-1b-Run3	[0.53, 0.62, 0.7, 0.79, 1]
Medium-2b-Run3	[0.53, 0.58, 0.66, 0.74, 1]
High-1b-Run2	[0.44, 0.75, 0.87, 1]
High-2b-Run2	[0.44, 0.71, 0.83, 1]
High-1b-Run3	[0.44, 0.6, 0.76, 1]
High-2b-Run3	[0.44, 0.57, 0.71, 1]

Table 5.4 – SR binning in NN signal score. The different number of bins for different Δm regions is due to the different shape of the NN output distributions.

and data-taking period. The red arrows indicate the thresholds used to define the final CRs. The list of selections on NN signal and $t\bar{t}$ scores defined for each $t\bar{t}$ CR is reported in Table 5.5.




Figure 5.12 – NN $t\bar{t}$ score data/MC distributions for all NN models evaluated on Run 2 samples. Data in SRs are blinded.

For simplicity, the NN signal score selections are kept consistent across models targeting the same Δm region, and the selection on the NN $t\bar{t}$ score is the same for all regions. The value 0.05 is chosen to achieve a satisfying $t\bar{t}$ purity, while keeping a sufficient number of expected events in each region. As in the SRs, the training preselection specific to each Δm region, as described




Figure 5.13 – NN $t\bar{t}$ score data/MC distributions for all NN models evaluated on Run 3 samples. Data in SRs are blinded.

in Table 5.1, is applied to all CR definitions.

The expected SM yields in the $t\bar{t}$ CRs for the Run 2 and Run 3 datasets are shown in Table 5.6

Region	NN signal score interval	Selection on NN $t\bar{t}$ score
Low- Δm	[0.70, 0.85)	> 0.05
Medium- Δm	[0.20, 0.40)	> 0.05
High- Δm	[0.02, 0.20)	> 0.05

Table 5.5 – Selections on NN signal and $t\bar{t}$ scores defining the $t\bar{t}$ CRs.

for the low- Δm regions, in Table 5.7 for the medium- Δm regions, and in Table 5.8 for the high- Δm regions.

$CR_{t\bar{t}}$	Low-1b-Run2	Low-2b-Run2	$CR_{t\bar{t}}$	Low-1b-Run3	Low-2b-Run3
Observed events	232	193	Observed events	113	106
Total SM background	264.51 ± 33.02	219.73 ± 25.10	Total SM Background	111.60 ± 12.11	88.47 ± 10.62
$t\bar{t}$	233.61 ± 27.14	201.34 ± 27.12	$t\bar{t}$	98.77 ± 10.58	81.52 ± 10.56
Single-top	20.48 ± 5.77	10.76 ± 1.75	Single-top	8.67 ± 2.07	4.55 ± 0.60
Z+ jets	0.74 ± 0.66	0.71 ± 0.26	Z+ jets	0.44 ± 0.37	$0.08^{+0.09}_{-0.08}$
Di-boson	1.47 ± 0.50	$0.08^{+0.10}_{-0.08}$	Di-boson	0.61 ± 0.12	0.07 ± 0.02
$t\bar{t}Z$	3.17 ± 0.44	2.85 ± 0.31	$t\bar{t}Z$	1.20 ± 0.25	1.06 ± 0.21
Other backgrounds	3.09 ± 0.58	2.75 ± 0.38	Other backgrounds	0.96 ± 0.10	0.75 ± 0.13
MC-Fakes	1.93 ± 0.74	1.22 ± 0.47	MC-Fakes	0.95 ± 0.37	0.43 ± 0.23

Table 5.6 – MC expected yields in $CR_{t\bar{t}}$ for the low- Δm region, in Run 2 (left) and Run 3 (right) datasets. The error includes statistical and systematic uncertainties (described in Section 6.2).

$CR_{t\bar{t}}$	Medium-1b-Run2	Medium-2b-Run2	$CR_{t\bar{t}}$	Medium-1b-Run3	Medium-2b-Run3
Observed events	219	151	Observed events	107	70
Total SM background	238.47 ± 20.52	163.61 ± 13.22	Total SM background	106.31 ± 14.97	71.01 ± 13.37
$t\bar{t}$	201.62 ± 14.55	150.63 ± 14.51	$t\bar{t}$	90.92 ± 13.48	65.68 ± 13.43
Single-top	22.87 ± 4.14	7.34 ± 0.76	Single-top	9.74 ± 2.23	2.93 ± 0.43
Z+ jets	4.75 ± 1.56	0.62 ± 0.24	Z+ jets	2.36 ± 0.81	0.34 ± 0.16
Di-boson	1.65 ± 0.37	$0.00^{+0.06}_{-0.00}$	Di-boson	0.67 ± 0.17	0.04 ± 0.03
$t\bar{t}Z$	2.28 ± 0.37	1.56 ± 0.32	$t\bar{t}Z$	0.89 ± 0.21	0.69 ± 0.23
Other backgrounds	3.43 ± 0.44	2.46 ± 0.33	Other backgrounds	1.14 ± 0.15	0.70 ± 0.08
MC-Fakes	1.87 ± 0.44	0.99 ± 0.38	MC-Fakes	0.61 ± 0.48	0.63 ± 0.22

Table 5.7 – MC expected yields in $CR_{t\bar{t}}$ for the medium- Δm region, in Run 2 (left) and Run 3 (right) datasets. The error includes statistical and systematic uncertainties (described in Section 6.2).

$CR_{t\bar{t}}$	High-1b-Run2	High-2b-Run2	$CR_{t\bar{t}}$	High-1b-Run3	High-2b-Run3
Observed events	345	207	Observed events	174	116
Total SM background	362.19 ± 29.22	212.55 ± 18.36	Total SM background	170.12 ± 14.63	99.84 ± 7.97
$t\bar{t}$	283.07 ± 19.43	182.36 ± 19.37	$t\bar{t}$ events	135.52 ± 8.61	87.12 ± 8.56
Single-top	36.70 ± 4.93	15.41 ± 1.22	Single-top	15.72 ± 3.74	6.94 ± 1.04
Z+ jets	19.49 ± 3.46	3.57 ± 0.64	Z+ jets	9.24 ± 3.06	1.37 ± 0.36
Di-boson	5.13 ± 0.60	0.25 ± 0.06	Di-boson	2.56 ± 0.35	0.10 ± 0.03
$t\bar{t}Z$	6.14 ± 1.03	3.63 ± 0.38	$t\bar{t}Z$	2.64 ± 0.47	1.54 ± 0.22
Other backgrounds	7.66 ± 0.94	5.34 ± 0.28	Other backgrounds	2.56 ± 0.39	1.63 ± 0.16
MC-Fakes	4.01 ± 1.12	1.98 ± 0.33	MC-Fakes	1.90 ± 0.56	1.15 ± 0.29

Table 5.8 – MC expected yields in $CR_{t\bar{t}}$ for the high- Δm region, in Run 2 (left) and Run 3 (right) datasets. The error includes statistical and systematic uncertainties (described in Section 6.2).

The $t\bar{t}$ background estimation is validated in dedicated **validation regions**. The VRs for the $t\bar{t}$ background are defined based on the NN scores, similarly to the corresponding CRs. There are therefore twelve distinct VRs, two for each NN model, one version for Run 2 dataset, and one version for Run 3 dataset. The VRs contain events with NN signal scores between the upper limit of the corresponding $t\bar{t}$ CR and the lower limit of the corresponding SRs. Therefore, the selections on the NN signal score are the same for models targeting the same Δm region, and for the Run 2 and the Run 3 datasets. Similarly to the corresponding SRs and CRs, the training preselection specific to each Δm region described in Table 5.1 is also applied to all VRs. The selection criteria applied to define the $t\bar{t}$ VRs are listed in Table 5.9.

Region	Signal NN score interval
Low- Δm	[0.85, 0.93)
Medium- Δm	[0.40, 0.53)
High- Δm	[0.20, 0.44)

Table 5.9 – Selections applied to the NN signal scores to define the $t\bar{t}$ VRs.

The expected SM yields in the $t\bar{t}$ VRs for the Run 2 and Run 3 datasets are shown in Table 5.10 for the low- Δm regions, in Table 5.11 for the medium- Δm regions, and in Table 5.12 for the high- Δm regions.

VR $_{t\bar{t}}$	Low-1b-Run2	Low-2b-Run2	CR $_{t\bar{t}}$	Low-1b-Run3	Low-2b-Run3
Observed events	106	75	Observed events	49	40
Total SM background	121.54 $\pm$ 15.73	80.89 $\pm$ 5.23	Total SM Background	51.98 $\pm$ 7.68	32.85 $\pm$ 4.05
$t\bar{t}$	96.95 $\pm$ 12.91	65.53 $\pm$ 5.08	$t\bar{t}$	42.43 $\pm$ 6.32	26.93 $\pm$ 3.93
Single-top	10.52 $\pm$ 1.64	6.05 $\pm$ 0.69	Single-top	4.20 $\pm$ 0.79	2.43 $\pm$ 0.50
Z+ jets	0.88 $\pm$ 0.20	0.20 $\pm$ 0.12	Z+ jets	0.38 $\pm$ 0.35	0.10 ^{+0.11} _{-0.10}
Di-boson	2.62 $\pm$ 0.57	0.26 $\pm$ 0.14	Di-boson	1.04 $\pm$ 0.17	0.09 $\pm$ 0.03
$t\bar{t}Z$	5.69 $\pm$ 0.71	5.23 $\pm$ 0.28	$t\bar{t}Z$	2.35 $\pm$ 0.35	2.11 $\pm$ 0.15
Other backgrounds	3.81 $\pm$ 0.55	3.03 $\pm$ 0.27	Other backgrounds	1.14 $\pm$ 0.17	0.94 $\pm$ 0.11
MC-Fakes	1.07 $\pm$ 0.49	0.58 $\pm$ 0.34	MC-Fakes	0.45 $\pm$ 0.30	0.24 $\pm$ 0.12

Table 5.10 – MC expected yields in VR $_{t\bar{t}}$ for the low- Δm region, in Run 2 (left) and Run 3 (right) datasets. The error includes statistical and systematic uncertainties (described in Section 6.2).

VR $_{t\bar{t}}$	Medium-1b-Run2	Medium-2b-Run2	CR $_{t\bar{t}}$	Medium-1b-Run3	Medium-2b-Run3
Observed events	60	35	Observed events	23	33
Total SM background	65.00 $\pm$ 6.69	38.12 $\pm$ 4.04	Total SM background	27.32 $\pm$ 6.70	17.17 $\pm$ 3.72
$t\bar{t}$	46.77 $\pm$ 5.91	31.25 $\pm$ 3.89	$t\bar{t}$	20.01 $\pm$ 6.37	14.48 $\pm$ 3.65
Single-top	7.44 $\pm$ 0.77	1.92 $\pm$ 0.38	Single-top	3.08 $\pm$ 0.57	0.77 $\pm$ 0.22
Z+ jets	1.29 $\pm$ 0.53	0.24 $\pm$ 0.08	Z+ jets	0.69 $\pm$ 0.32	0.08 $\pm$ 0.05
Di-boson	2.46 $\pm$ 0.45	0.10 $\pm$ 0.05	Di-boson	0.89 $\pm$ 0.16	0.09 $\pm$ 0.05
$t\bar{t}Z$	3.49 $\pm$ 0.67	2.72 $\pm$ 0.20	$t\bar{t}Z$	1.48 $\pm$ 0.21	1.13 $\pm$ 0.14
Other backgrounds	2.70 $\pm$ 0.30	1.77 $\pm$ 0.17	Other backgrounds	0.79 $\pm$ 0.13	0.55 $\pm$ 0.05
MC-Fakes	0.86 $\pm$ 0.25	0.12 ^{+0.23} _{-0.12}	MC-Fakes	0.38 $\pm$ 0.16	0.06 ^{+0.18} _{-0.06}

Table 5.11 – MC expected yields in VR $_{t\bar{t}}$ for the medium- Δm region, in Run 2 (left) and Run 3 (right) datasets. The error includes statistical and systematic uncertainties (described in Section 6.2).

VR _{t$\bar{t}$}	High-1b-Run2	High-2b-Run2	CR _{t$\bar{t}$}	High-1b-Run3	High-2b-Run3
Observed events	47	24	Observed events	21	12
Total SM background	48.95 $\pm$ 4.18	31.45 $\pm$ 3.22	Total SM background	22.05 $\pm$ 5.00	14.58 $\pm$ 2.74
$t\bar{t}$	30.24 $\pm$ 2.79	22.26 $\pm$ 3.11	$t\bar{t}$ events	13.74 $\pm$ 4.64	10.90 $\pm$ 2.68
Single-top	5.85 $\pm$ 1.23	3.14 $\pm$ 0.46	Single-top	2.59 $\pm$ 0.54	1.39 $\pm$ 0.26
Z + jets	2.77 $\pm$ 0.74	0.42 $\pm$ 0.15	Z + jets	1.53 $\pm$ 0.40	0.17 $\pm$ 0.10
Di-boson	2.56 $\pm$ 0.40	0.15 $\pm$ 0.07	Di-boson	1.24 $\pm$ 0.18	0.10 $\pm$ 0.03
$t\bar{t}Z$	3.85 $\pm$ 0.37	2.89 $\pm$ 0.26	$t\bar{t}Z$	1.64 $\pm$ 0.20	1.16 $\pm$ 0.11
Other backgrounds	2.84 $\pm$ 0.37	2.04 $\pm$ 0.12	Other backgrounds	0.95 $\pm$ 0.14	0.68 $\pm$ 0.08
MC-Fakes	0.86 $\pm$ 0.19	0.54 $\pm$ 0.10	MC-Fakes	0.36 $\pm$ 0.32	0.19 $\pm$ 0.14

Table 5.12 – MC expected yields in VR_{t $\bar{t}$} for the high- Δm region, in Run 2 (left) and Run 3 (right) datasets. The error includes statistical and systematic uncertainties (described in Section 6.2).

5.2.3 Estimation of the $t\bar{t}Z$ background

The $t\bar{t}Z \rightarrow \nu\nu$ process is expected to be one of the dominant SM background contributions to the SRs, especially for events with NN signal scores, as shown in Figures 5.8 and 5.9.

Given the difficulty in selecting a sufficiently pure sample for this process, due to the high contamination from $t\bar{t}$ events, a selection strategy based on a three-lepton final state is adopted. The selection applied to define CR_{t $\bar{t}Z$} targets $t\bar{t}Z$ events, with the Z boson decaying into two leptons with same flavour (ee or $\mu\mu$) and opposite electric charge (*same-flavour, opposite-sign*, or SFOS, leptons), and the $t\bar{t}$ system decaying semileptonically. Therefore, events are selected if they have exactly three leptons in the final state, including at least one pair of SFOS leptons with an invariant mass $m_{\ell\ell}$ compatible with the Z boson mass ($|m_{\ell\ell} - m_Z| < 20$ GeV). If multiple pairs of leptons satisfying these requirements are identified, the one with $m_{\ell\ell}$ closest to m_Z is selected as the pair originating from the Z boson. The invariant mass of the pair of SFOS leptons selected with this approach is named *closest* $m_{\ell\ell}$. In order to select $t\bar{t}Z$ events whose kinematics resembles those of $t\bar{t}Z \rightarrow \nu\nu$ events in the SRs, the *corrected* E_T^{miss} variable is defined as:

$$E_{T,\text{corr}}^{\text{miss}} = |(\mathbf{p}(\ell_1^Z) + \mathbf{p}(\ell_2^Z) + \mathbf{p}_T^{\text{miss}})|, \quad (5.13)$$

where $\mathbf{p}(\ell_1^Z)$ and $\mathbf{p}(\ell_2^Z)$ are the momenta of the two leptons of the SFOS pair, and $\mathbf{p}_T^{\text{miss}}$ is the missing transverse momentum vector. In this way, the two SFOS leptons can be effectively treated as the neutrino pair coming from the Z boson decay of a $t\bar{t}Z \rightarrow \nu\nu$ process, and events with high $E_{T,\text{corr}}^{\text{miss}}$ emulate SR events with high E_T^{miss} . Events are further required to have at least four jets, including at least one b -tagged jet, to be consistent with the expected $t\bar{t}$ decay topology. The selection criteria used to define CR_{t $\bar{t}Z$} is summarised in Table 5.13, and are applied to both the Run 2 and the Run 3 datasets.

Tables 5.14 and 5.15 show the expected SM yield in CR_{t $\bar{t}Z$} from the Run 2 and the Run 3 MC samples, respectively.

Figure 5.14 shows the distribution of the closest $m_{\ell\ell}$ and the $E_{T,\text{corr}}^{\text{miss}}$ variables in CR_{t $\bar{t}Z$} . Good agreement is observed between the expected and observed distributions.

As the contribution of the minor backgrounds, labelled as “other”, is significant in CR_{t $\bar{t}Z$} , the

CR_{t$\bar{t}$Z} selection	
$p_T(\ell_1)$ [GeV]	> 27
$p_T(\ell_2)$ [GeV]	> 20
$p_T(\ell_3)$ [GeV]	> 20
$ m_{\ell\ell} - m_Z $ [GeV]	> 20 for at least one SFOS pair
$n_{b\text{-jets}}$	≥ 1 with $n_{\text{jet}} \geq 4$
$E_{T,\text{corr}}^{\text{miss}}$	> 70

Table 5.13 – Selection criteria of the control region for the $t\bar{t}Z \rightarrow \nu\nu$ background (CR_{t $\bar{t}$ Z}).

CR _{t$\bar{t}$Z} - Run 2	
Observed events	412
Total SM background	391.56 ± 17.83
$t\bar{t}$	0.00 ± 0.00
Single-top	0.00 ± 0.00
$Z + \text{jets}$	0.00 ± 0.00
Di-boson	51.96 ± 7.57
$t\bar{t}Z$ ($\nu\nu, \ell\ell, qq$)	254.17 ± 1.10
Other backgrounds	66.56 ± 7.44
MC-Fakes	18.88 ± 2.59

Table 5.14 – MC expected yields in CR_{t $\bar{t}$ Z} for the Run 2 MC samples. The error includes statistical and systematic uncertainties (described in Section 6.2).

CR _{t$\bar{t}$Z} - Run 3	
Observed events	163
Total SM background	154.13 ± 7.36
$t\bar{t}$	0.00 ± 0.00
Single-top	0.00 ± 0.00
$Z + \text{jets}$	0.00 ± 0.00
Di-boson	19.72 ± 3.39
$t\bar{t}Z$ ($\nu\nu, \ell\ell, qq$)	105.26 ± 0.69
Other backgrounds	22.11 ± 2.93
MC-Fakes	7.03 ± 1.03

Table 5.15 – MC expected yields in CR_{t $\bar{t}$ Z} for the Run 3 MC samples. The error includes statistical and systematic uncertainties (described in Section 6.2).




Figure 5.14 – Distribution of the closest $m_{\ell\ell}$ (left) and of $E_{T,\text{corr}}^{\text{miss}}$ (right) in $\text{CR}_{t\bar{t}Z}$ in Run 2 (upper row) Run 3 (bottom row) samples.

composition of this background was studied by analysing the relative contributions from its main components: $t\bar{t}W$, tZq , and tWZ . The results of this study are presented in Table 5.16.

Background process	Contribution (%)
$t\bar{t}W$	10.07
tZq	30.06
tWZ	51.90
Remaining processes	8.00

Table 5.16 – Relative contributions of the main components to the “other” background in the $\text{CR}_{t\bar{t}Z}$ control region.

The estimation of the $t\bar{t}Z$ background is validated in a dedicated **validation region** ($\text{VR}_{t\bar{t}Z}$), which is defined by selecting events containing four leptons in the final state, including at least one SFOS pair. To select events with kinematics consistent with those of signal events, a variant of the m_{T2} variable, referred to as $m_{T2}^{4\ell}$ is defined. The variable $\mathbf{p}_{T,\text{corr}}^{\text{miss}} = \mathbf{p}_T^{\text{miss}} + \mathbf{p}(\ell_1^Z) + \mathbf{p}(\ell_2^Z)$, where ℓ_1^Z and ℓ_2^Z are the leptons from the SFOS pair with $m_{\ell\ell}$ closest to m_Z , is used as the invisible object, while the momenta of the remaining two leptons are used as visible objects (see the mathematical definition of m_{T2} in Section 5.1.1).

The selection criteria used to define $\text{VR}_{t\bar{t}Z}$ are summarised in Table 5.17.

$\text{VR}_{t\bar{t}Z}$ selection	
Lepton multiplicity	4
$p_T(\ell_1)$ [GeV]	> 27
$p_T(\ell_2)$ [GeV]	> 20
$p_T(\ell_3)$ [GeV]	> 20
$p_T(\ell_4)$ [GeV]	> 20
At least one SFOS pair	
$n_{b\text{-jets}}$	> 0
$m_{T2}^{4\ell}$	> 10

Table 5.17 – Selection criteria of the validation region for the $t\bar{t}Z \rightarrow \nu\nu$ background ($\text{VR}_{t\bar{t}Z}$).

Tables 5.18 and 5.19 show the expected SM yield in $\text{VR}_{t\bar{t}Z}$ from the Run 2 and the Run 3 MC samples, respectively.

VR _{t$\bar{t}$Z} - Run 2	
Observed events	70
Total SM background	54.49 ± 1.69
$t\bar{t}$	0.00 ± 0.00
Single-top	0.00 ± 0.00
$Z + jets$	0.00 ± 0.00
Di-boson	10.70 ± 0.96
$t\bar{t}Z (vv, \ell\ell, qq)$	34.34 ± 2.06
Other backgrounds	8.19 ± 0.45
MC-Fakes	1.26 ± 0.08

Table 5.18 – MC expected yields in VR_{t $\bar{t}$ Z} for the Run 2 MC samples. The error includes statistical and systematic uncertainties (described in Section 6.2).

VR _{t$\bar{t}$Z} - Run 3	
Observed events	28
Total SM background	20.81 ± 0.93
$t\bar{t}$	0.00 ± 0.00
Single-top	0.00 ± 0.00
$Z + jets$	0.00 ± 0.00
Di-boson	4.06 ± 0.45
$t\bar{t}Z (vv, \ell\ell, qq)$	14.46 ± 1.07
Other backgrounds	1.83 ± 0.13
MC-Fakes	0.46 ± 0.04

Table 5.19 – MC expected yields in VR_{t $\bar{t}$ Z} for the Run 3 MC samples. The error includes statistical and systematic uncertainties (described in Section 6.2).

Statistical Analysis and Results

This chapter provides a detailed description of the statistical framework used to derive the results of the analysis, including the definition of a likelihood function and how it is used to define a test statistic to evaluate the compatibility of the observed data with the background-only hypothesis and to test for the presence of a potential signal (Section 6.1). The systematic uncertainties affecting the analysis results are summarised in Section 6.2, and the results of the statistical analysis and their interpretation are described in Section 6.3. This last section includes a reinterpretation of the analysis considering an alternative signal model, the $t\bar{t}$ production in association with an invisibly decaying Higgs boson, which is characterised by the same final-state signature as the targeted stop decay.

6.1 Statistical Model

The analysed data are stochastic, due to the non-deterministic nature of quantum mechanics as well as to numerous experimental effects. A well-defined statistical model is needed to be able to quantify how much a given hypothesis is likely to produce the observed data. This section introduces the statistical model used in the analysis, which is based on a likelihood function that includes the expected and observed number of events in all analysis regions, as well as the associated parameters needed to describe signal, background contributions, and systematic uncertainties. The definition of the likelihood and the statistical methods used to test hypotheses and estimate confidence levels are described in detail. The statistical analysis is implemented using the statistical analysis package HISTFITTER [140, 141] in combination with pyhf [142, 143].

6.1.1 Likelihood function

The likelihood function statistically describing the analysis model can be defined as:

$$L(\mathbf{n}|\boldsymbol{\mu}, \boldsymbol{\theta}) = \prod_{i=1}^{N^{\text{regions}}} [\text{Pois}(n_i|N_i(\boldsymbol{\mu}, \boldsymbol{\theta}))] \times G(\boldsymbol{\theta}). \quad (6.1)$$

The index i runs over the considered analysis regions, n_i is the observed number of events and $N_i(\boldsymbol{\theta})$ is expected number of signal and background events in the i^{th} region. The observed

number of events n_i follows a Poisson distribution, centred around the expected number of events $N_i(\boldsymbol{\mu}, \boldsymbol{\theta})$, which is defined as:

$$N_i(\boldsymbol{\mu}, \boldsymbol{\theta}) = \mu_{\text{sig}} \cdot N_i^{\text{signal}}(\boldsymbol{\theta}^{\text{signal}}) + \mu_{t\bar{t}} \cdot N_i^{t\bar{t}}(\boldsymbol{\theta}^{t\bar{t}}) + \mu_{t\bar{t}Z} \cdot N_i^{t\bar{t}Z}(\boldsymbol{\theta}^{t\bar{t}Z}) + N_i^{\text{other}}(\boldsymbol{\theta}^{\text{other}}), \quad (6.2)$$

where $\mu_{t\bar{t}}$ and $\mu_{t\bar{t}Z}$ are the normalisation factors (NFs) used to normalise the yields $N_i^{t\bar{t}}$ and $N_i^{t\bar{t}Z}$ of the main backgrounds, $t\bar{t}$ and $t\bar{t}Z$, in the respective CRs. The other non-dominant backgrounds yields, N_i^{other} , are simply added without applying any NF. Finally, N_i^{signal} is the signal yield, and μ_{sig} is the stop signal strength, which is the *parameter of interest* (POI) of the model; N_i^{signal} and μ_{sig} are different for each signal model in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ parameter space. A value $\mu_{\text{sig}} = 1$, corresponds the nominal event yield predicted by the respective model.

The term $G(\vec{\theta})$ is a probability density function (pdf) imposing constraints on the *nuisance parameters* of the model, $\boldsymbol{\theta}$, which take into account the systematic uncertainties. It is a product of Gaussian distributions corresponding to the auxiliary measurements describing each systematic uncertainties, typically centred in zero and with unit width.

The results of the analysis are extracted via **maximum likelihood fits**, performed simultaneously across all the relevant analysis categories. The fit calculates the values of the model parameters $(\boldsymbol{\mu}, \boldsymbol{\theta})$ that maximise the likelihood function.

As already mentioned, the yields of the two major backgrounds, $t\bar{t}$ and $t\bar{t}Z$, are constrained using dedicated NFs, while the estimation of non-dominant background processes, expected to contribute only negligibly to the SR, is based only on MC simulations. Three types of fit are performed for this search, differing by the presence of different regions in the likelihood and the assumptions made about the signal model and rate.

- **Background-only fit.** The purpose of this fit strategy is to estimate the total background in SRs and VRs, without making assumptions on any signal model. Only background samples are included in the model, and the CRs are assumed to be free of signal contamination. The fit is only performed in the CRs, and the dominant background processes, $t\bar{t}$ and $t\bar{t}Z$, are normalised to the observed number of events in the CRs to extract the NFs. The result of this fit is used to predict the number of expected background events in the SRs and VRs. As only CRs are used in the fit, these predictions are independent of the observed number of events in each SR and VR. This allows for an unbiased comparison between the predicted and observed number of events in each region and to evaluate the quality of background modelling. NFs obtained from the background-only fit of CRs are validated in the dedicated VRs. If the NFs obtained from the CRs are correct, the number of data events observed in a particular VR should be correctly described by the MC simulation after it is corrected for its NF.
- **Model-dependent fit.** This fit strategy is used with the objective of studying a specific signal model (hence the name), and is repeated for each of the signal points in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ parameter space. In absence of a significant event excess in the SRs, as concluded with the background-only fit configuration, *exclusion limits* can be set on the signal models under study. In case of excess, the model-dependent signal fit can be used to measure properties such as the signal strength. The fit is performed in the CRs and SRs simultaneously. Along

with the background samples, a signal sample is included in all regions, included the CRs, to correctly account for possible signal contamination in the CRs. The signal strength parameter μ_{sig} is assigned to the signal sample, and it is the POI of the statistical model. The different bins of the SR are included simultaneously in this fit, which is performed on their statistical combination.

- **Model-independent fit.** Used to make qualitative and model-independent statements on whether the measurements in single-bin regions are consistent with the SM predictions. All CRs and one SR are included in these fits simultaneously. No signal contributions in the CRs is assumed, but no other assumptions are made for the signal model. A signal contribution equal to the signal strength is assumed in the SR. The number of signal events in the SR is then treated as a free parameter in the fit. This fit strategy is used to set model-independent upper limits on the number of events beyond the expected number of events in each SR. This allows to exclude any signal model by estimating its number of signal events predicted in a particular SR. The SR in this fit configuration is constructed as a single-bin region, since having more bins requires assumptions on the signal spread over these bins. The model-independent signal fit strategy, fitting both the CRs and each SR, is also used to perform the background-only hypothesis test, which quantifies the significance of any observed excess of events in a SR, again in a manner that is independent of any particular signal model (*discovery fit*).

For this analysis, the background-only fit is performed separately for the Run 2 and Run 3 datasets, and independently for each of the three non-orthogonal Δm categories, combining within each fit the corresponding orthogonal event selections with different b -jet multiplicity (1b and 2b). Although $\text{CR}_{t\bar{t}Z}$ has a unique definition across the Δm regions, the resulting NFs differ slightly as a consequence of the different definition of $\text{CR}_{t\bar{t}}$. As a result, 12 NFs are obtained from the background-only fit: six for the $t\bar{t}$ and six for the $t\bar{t}Z$ background, corresponding to each Δm region (low, medium, high) and data-taking period (Run 2 and Run 3). Within a given Δm region, $\text{CR}_{t\bar{t}}$ regions differing only by b -jet multiplicity share a common NF.

6.1.2 Hypothesis test

Two exclusive hypotheses need to be defined to interpret the results of the fit, the null hypothesis, and an alternative hypothesis. A test statistic is defined, to infer which of the two hypothesis is more likely to produce the observed data. The results are obtained with the *profiled likelihood ratio* test statistic, $\Lambda(\mu_{\text{sig}})$, built from the likelihood function described in Equation 6.1 as:

$$\Lambda(\mu_{\text{sig}}) = \frac{\mathcal{L}(n|\mu_{\text{sig}}, \hat{\theta})}{\mathcal{L}(n|\hat{\mu}_{\text{sig}}, \hat{\theta})}, \quad (6.3)$$

where the signal strength μ_{sig} represents the POI, and the θ represent the NPs. In the numerator, the likelihood is maximised only with respect to the NPs, which are set to their *profiled values* $\hat{\theta}$, for a fixed value of μ_{sig} , while in the denominator the likelihood is maximised with respect to both the NPs and the POI μ_{sig} , both set to their best-fit values, $\hat{\mu}_{\text{sig}}$ and $\hat{\theta}$, respectively. Intuitively, the profile likelihood ratio measures the agreement between the hypothesised value of the POI, μ_{sig} , and the observed data: $0 \leq \Lambda(\mu_{\text{sig}}) \leq 1$, and $\Lambda(\mu_{\text{sig}})$ approaching unity implies a

good agreement between the considered model and the observed data.

Different test statistics can be derived from $\Lambda(\mu_{\text{sig}})$, depending on the objective of the hypothesis test. If an excess of events in the SR is observed from the background-only fit, a hypothesis test for discovery can be performed. In this case, the null hypothesis corresponds to the background-only scenario, and the goal is to assess whether it can be rejected in favour of the alternative hypothesis that includes a signal contribution in addition to the SM background. Conversely, if no significant excess of events is observed, a hypothesis test can be performed to set upper exclusion limits on the signal strength μ_{sig} . In this case, the null hypothesis corresponds to the presence of signal with a given signal strength μ_{sig} , and the test evaluates whether this hypothesis can be rejected in favour of the background-only alternative.

The presence of a stop signal at a specific $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ mass point can be tested using the test statistic q_0 , given by:

$$q_0 = \begin{cases} -2 \ln \Lambda(0) & \text{if } \hat{\mu}_{\text{sig}} \geq 0 \\ 0 & \text{if } \hat{\mu}_{\text{sig}} < 0 \end{cases} \quad (6.4)$$

The test statistic assumes that a signal would manifest as a positive value of $\hat{\mu}_{\text{sig}}$, indicating an excess of events over the background-only scenario ($\mu_{\text{sig}} = 0$). A q_0 value near zero corresponds to a profile likelihood ratio $\Lambda(0)$ close to 1, suggesting that the background-only hypothesis closely matches the best-fit values of the signal strength and NPs $(\hat{\mu}_{\text{sig}}, \hat{\theta})$ that jointly maximise the likelihood. In contrast, larger q_0 values indicate increasing disagreement between observed data and the background-only hypothesis, potentially indicating the presence of a signal. The level of disagreement between the observed data and the background-only null hypothesis can be quantified by the p -value, p_0 , given by:

$$p_0 = \int_{q_{0,\text{obs}}}^{\infty} f(q_0|0) dq_0, \quad (6.5)$$

where $q_{0,\text{obs}}$ is the value of the test statistic in the observed data, and $f(q_0|0)$ represents the pdf of q_0 under the background-only assumption, i.e. the expected distribution of the test statistics if the same experiment was repeated several times. The distribution of the test statistic q_0 is typically evaluated using asymptotic formulae derived under the asymptotic approximation [139], which is expected to hold unless the data contain very few events.

A smaller p_0 indicates a larger deviation from the background-only hypothesis, providing stronger confidence in the presence of a signal. The p_0 can be also expressed in terms of the *discovery significance* Z , which corresponds to the number of standard deviations above the mean of a normal Gaussian distribution with the same upper-tail probability as p_0 . The relationship between p_0 and Z is defined as:

$$Z = \Phi^{-1}(1 - p_0), \quad (6.6)$$

where Φ^{-1} is the quantile (inverse of the cumulative distribution) of the standard Gaussian distribution. To claim a discovery, the p_0 must reach a threshold of 2.87×10^{-7} [139], corresponding to a discovery significance equal to 5.0σ .

Figure 6.1 illustrates the relationship between Z and the p -value.

In absence of signal excesses, the statistical results of the search can be interpreted in terms of




Figure 6.1 – The standard Gaussian distribution $\phi(x) = (1/\sqrt{2\pi}) \cdot \exp(-x^2/2)$, showing the relation between the significance Z and the p -value [139].

upper limits on the signal strength μ_{sig} . The relevant test statistic for setting upper limits is $\tilde{q}_{\mu_{\text{sig}}}$, which differentiates the case of the signal being produced at a rate $\hat{\mu}_{\text{sig}}$ greater or smaller than μ_{sig} :

$$\tilde{q}_{\mu_{\text{sig}}} = \begin{cases} -2 \ln \frac{\mathcal{L}(\mu_{\text{sig}}, \hat{\theta}(\mu_{\text{sig}}))}{\mathcal{L}(0, \hat{\theta}(0))} & \text{if } \hat{\mu}_{\text{sig}} \leq 0 \\ -2 \ln \frac{\mathcal{L}(\mu_{\text{sig}}, \hat{\theta}(\mu_{\text{sig}}))}{\mathcal{L}(\hat{\mu}_{\text{sig}}, \hat{\theta})} & \text{if } 0 < \hat{\mu}_{\text{sig}} < \mu_{\text{sig}} \\ 0 & \text{if } \hat{\mu}_{\text{sig}} > \mu_{\text{sig}} \end{cases} \quad (6.7)$$

The reason for setting $\tilde{q}_{\mu_{\text{sig}}} = 0$ if $\hat{\mu}_{\text{sig}} > \mu_{\text{sig}}$ is that, when setting an upper limit, data with $\hat{\mu}_{\text{sig}} > \mu_{\text{sig}}$ have less compatibility with the null hypothesis than the observed data, therefore they are not part of the rejection region of the test.

For a given value of μ_{sig} , its disagreement with the observed data is quantified by the p -value, $p_{\mu_{\text{sig}}}$:

$$p_{\mu_{\text{sig}}} = \int_{\tilde{q}_{\mu_{\text{sig}}, \text{obs}}}^{\infty} f(\tilde{q}_{\mu_{\text{sig}}} | \mu_{\text{sig}}) d\tilde{q}_{\mu_{\text{sig}}}, \quad (6.8)$$

where $\tilde{q}_{\mu_{\text{sig}}, \text{obs}}$ is the test statistic evaluated on the observed data, while $f(\tilde{q}_{\mu_{\text{sig}}} | \mu_{\text{sig}})$ is the pdf of the test statistic under the null hypothesis.

As illustrated in Figure 6.2, a smaller $p_{\mu_{\text{sig}}}$ indicates a greater incompatibility between the null hypothesis of signal with signal strength μ_{sig} and the observed data, suggesting regions where the signal hypothesis can be rejected. Upper limits on the POI μ_{sig} are established using the CL_s method [144]. The quantity

$$\text{CL}_s = \frac{p_{\mu_{\text{sig}}}}{1 - p_b} \quad (6.9)$$




Figure 6.2 – Illustration of a p -value obtained from an observed value $t_{\mu,\text{obs}}$ of the test statistic t_μ [139].

is computed for various values of μ_{sig} . The numerator represents the standard p -value, while the denominator is defined as:

$$p_b = 1 - \text{CL}_b = \int_{\tilde{q}_{\mu_{\text{sig}},\text{obs}}}^{\infty} f(\tilde{q}_{\mu_{\text{sig}}} | 0), d\tilde{q}_{\mu_{\text{sig}}}, \quad (6.10)$$

where $f(\tilde{q}_{\mu_{\text{sig}}} | 0)$ is the pdf of the test statistic under the alternate background-only hypothesis. The upper limit on μ_{sig} is the largest value for which $\text{CL}_s < \alpha$, where α indicates the confidence level (CL) for exclusion. For instance, with $\alpha = 0.05$, the presence of a signal is excluded for values above the upper limit at 95% CL. Also in this case, the distributions of $\tilde{q}_{\mu_{\text{sig}}}$ under both null and alternate hypotheses are determined using asymptotic formulae derived from the asymptotic approximation [139].

6.2 Systematic uncertainties

The sensitivity of the analysis is primarily limited by statistical precision, due to the low number of expected signal events. However, it is important to identify and evaluate all potential sources of systematic uncertainty affecting the final results. These uncertainties are classified based on their origin:

- **Experimental systematic uncertainties** arise from uncertainties in auxiliary measurements that enter the statistical model (e.g., the total integrated luminosity collected by ATLAS during the considered periods of data-taking), as well as from the reconstruction of physics objects within the ATLAS detector, or the application of trigger requirements.
- **Theoretical systematic uncertainties** affect the predicted yields of both the stop signals and the main background processes (i.e. $t\bar{t}$, $t\bar{t}Z$, single top, and di-boson production), including modelling uncertainties within the event simulation for the different background and signal processes.

Typically, the impact of each source of systematic uncertainty is quantified by propagating its

effect through the full analysis chain. The impact is then expressed as the relative difference between the nominal and varied results.

6.2.1 Experimental uncertainties

The experimental systematic uncertainties considered in this analysis are summarised below. Further details on the uncertainties arising by the object reconstruction procedures are provided in Chapter 3.

Electron and muon uncertainties include uncertainties arising from energy scale and resolution calibrations, as well as from reconstruction, identification, and isolation efficiencies.

Jet-related uncertainties arise from the reconstruction and calibration of small- R jets (see Section 3.4). These include jet energy scale (JES) and jet energy resolution (JER) uncertainties. Additional contributions arise from the calibration of the GN2 algorithm used for b -tagging at the 77% working point, as well as from the Jet Vertex Tagger (JVT) algorithm [131], which is applied to suppress pile-up jets.

E_T^{miss} **uncertainties** are primarily inherited from the individual reconstructed physics objects contributing to the E_T^{miss} calculation. However, the soft term introduces additional uncertainty.

Pileup reweighting uncertainties are due to the pile-up reweighting process, used to correct for differences in the pile-up distribution between data and MC simulation, caused by the fact that the level of pile-up can vary from event to event [145].

Luminosity measurement uncertainties. An uncertainty of 0.83% on the integrated luminosity is applied for the full Run 2 dataset, while a 2.0% uncertainty is used for the combined 2022-2023 Run 3 dataset. These uncertainties are derived from dedicated measurements [146, 147] using the LUCID 2 detector (see Section 2.2.5).

All experimental systematic uncertainties are implemented as two-sided variations in the statistical model, except for the E_T^{miss} soft term and the JER uncertainties, which are treated using symmetrised variations.

6.2.2 Theoretical systematic uncertainties

Theoretical systematic uncertainties are considered for the main background processes, $t\bar{t}$, single top, $t\bar{t}Z$, and di-boson, as well as for the signal processes.

These uncertainties can be broadly classified into two main categories.

- **Parton-level modelling uncertainties**, which include uncertainties in the proton PDFs, in the value of the strong coupling constant α_s , and in the choice of the renormalisation (μ_r) and factorisation (μ_f) scales (*scale uncertainties*), associated with missing higher-order corrections in the perturbative QCD calculations. PDF uncertainties are evaluated by taking the envelope of all PDF variations in each bin of the final discriminating distribution (in

this analysis, the NN signal score). Scale uncertainties are estimated using the standard seven-point QCD variations of μ_r and μ_f , and are grouped into three independent NPs.

- **Parton shower and hadronisation uncertainties**, which arise from the description of the processes that convert partons into the hadronic final states. These uncertainties are evaluated by comparing the results obtained with MC samples produced using different parton shower modelling.

The uncertainties in the modelling of the $t\bar{t}$ background are derived from variations of μ_r and μ_f , and from changes to the amount of initial-state radiation (ISR) and final-state radiation (FSR). These ISR and FSR uncertainties are evaluated using dedicated samples in which the values of μ_f and μ_r are doubled or halved. For ISR, the showering is also varied, by modifying α_s which impacts ISR in the PYTHIA tune. The matrix element to parton shower interface uncertainty is assessed through scans of the p_T^{hard} parameter, following the prescription in [148].

The uncertainties on $t\bar{t}Z$ production include the effects of QCD scale uncertainties and ISR/FSR uncertainties, the latter evaluated by comparing the nominal sample with one generated using a PYTHIA tune that enhances the radiation.

For single top production in the Wt channel, an uncertainty is assigned to account for interference with $t\bar{t}$ production. This is estimated by comparing samples generated using the *diagram removal* (DR, nominal) and *diagram subtraction* (DS) prescriptions [149]. For $t\bar{t}$, $t\bar{t}Z$, and single top production, the parton showering and hadronisation uncertainties are estimated by comparing samples generated using the two different showering models implemented in PYTHIA and in HERWIG.

The modelling uncertainties for the di-boson background are evaluated using the seven-point variations of the renormalisation and factorisation scales. In this case, the scale variations do not take into account uncertainties on the resummation and matching scales between the matrix element and parton shower, which are included by varying the relevant scale parameters in SHERPA. For all the above processes, the PDF uncertainties are also evaluated.

For the stop signal, the uncertainty on the SUSY production cross section is derived from the envelope of the cross section predictions obtained using different PDF sets and variants of the μ_r and μ_f scales [150]. Additional uncertainties are included by varying the renormalisation, factorisation, radiation, and merging scales used in the generation of the signal samples.

Systematic uncertainties in SRs

The following tables summarise the contributions from the different sources of experimental systematic uncertainty to the total SM background predictions in the binned SRs. Tables 6.1 and 6.2 consider the low- Δm regions, for Run 2 and Run 3, respectively. Tables 6.3 and 6.4 show the medium- Δm regions, for Run 2 and Run 3, respectively. Finally, Tables 6.5 and 6.6 present the results for the high- Δm regions, for Run 2 and Run 3, respectively. The theoretical uncertainty on Wt is not presented in these tables, as it was found to be always less than 1%. The systematic uncertainty related to the JER is the dominant uncertainty in most of the regions, which can be explained by effects induced on the E_T^{miss} calculation, and therefore on $m_{T2}^{\ell\ell}$, which is the main

discriminating variable of the analysis. The variations of these uncertainties between SR bins were studied and found out to be due to fluctuations.

Uncertainty source	SR ₀ ^{low-1b}	SR ₁ ^{low-1b}	SR ₂ ^{low-1b}	SR ₃ ^{low-1b}	SR ₄ ^{low-1b}	SR ₀ ^{low-2b}	SR ₁ ^{low-2b}	SR ₂ ^{low-2b}	SR ₃ ^{low-2b}	SR ₄ ^{low-2b}
Total statistical ($\sqrt{N_{\text{exp}}}$)	51.17%	35.66%	28.40%	28.21%	20.67%	54.80%	50.77%	33.87%	29.55%	27.31%
Total background systematic	12.59%	15.27%	11.36%	18.03%	12.49%	12.09%	16.02%	10.87%	10.98%	17.06%
Jet Energy Resolution	7.1%	5.9%	8.1%	17.8%	7.4%	4.9%	6.8%	5.3%	6.2%	13.8%
Jet Energy Scale	4.9%	7.4%	2.1%	9.1%	6.5%	2.3%	11.5%	4.5%	2.6%	4.6%
$t\bar{t}$ Theoretical Uncertainties	4.0%	7.9%	5.4%	7.4%	5.4%	4.0%	3.3%	4.2%	6.3%	6.0%
$t\bar{t}$ Normalisation	3.9%	3.1%	2.1%	1.5%	0.99%	6.1%	5.2%	2.9%	2.0%	1.4%
$t\bar{t}Z$ Theoretical Uncertainties	2.5%	4.8%	0.40%	0.60%	1.2%	4.5%	2.6%	0.72%	0.32%	1.2%
Pileup Reweighting and JVT	1.4%	0.66%	0.63%	1.5%	0.95%	0.83%	1.3%	0.09%	0.47%	1.0%
Lepton Modelling	1.3%	1.0%	1.4%	1.8%	0.97%	0.56%	0.85%	1.2%	1.2%	1.5%
$t\bar{t}$ Normalisation	1.2%	1.9%	2.9%	3.2%	3.8%	0.94%	1.2%	2.6%	3.5%	4.1%
E_T^{miss} Mismodelling	1.2%	1.2%	1.5%	1.2%	1.2%	1.2%	2.1%	1.0%	1.0%	2.6%
Flavour Tagging	1.2%	0.58%	0.38%	0.42%	0.34%	0.31%	0.43%	0.18%	0.17%	0.27%
Total uncertainty (stat. + syst.)	52.69%	38.79%	30.59%	33.48%	24.15%	56.12%	53.24%	33.57%	31.52%	32.20%

Table 6.1 – Sources of systematic uncertainty in the SM background estimates, after the background fits. The values are given as relative uncertainties in the total expected background event yields in the different bins of the Run 2 low- Δm SR.

Uncertainty source	SR ₀ ^{low-1b}	SR ₁ ^{low-1b}	SR ₂ ^{low-1b}	SR ₃ ^{low-1b}	SR ₄ ^{low-1b}	SR ₀ ^{low-2b}	SR ₁ ^{low-2b}	SR ₂ ^{low-2b}	SR ₃ ^{low-2b}	SR ₄ ^{low-2b}
Total statistical ($\sqrt{N_{\text{exp}}}$)	71.43%	62.16%	43.34%	41.81%	34.48%	68.87%	68.54%	56.73%	49.27%	42.22%
Total background systematic	21.02%	21.25%	22.15%	22.19%	23.92%	21.09%	26.00%	24.35%	19.23%	23.44%
Jet Energy Resolution	13.2%	11.8%	15.5%	14.5%	20.0%	10.5%	19.3%	19.1%	11.8%	16.0%
Jet Energy Scale	10.7%	7.5%	9.9%	9.9%	11.2%	8.7%	9.0%	10.9%	8.5%	9.1%
$t\bar{t}$ Theoretical Uncertainties	8.4%	8.8%	10.4%	11.3%	11.5%	11.1%	9.8%	5.7%	5.8%	8.8%
$t\bar{t}Z$ Normalisation	5.6%	3.2%	2.4%	1.7%	1.4%	8.1%	5.2%	3.3%	2.3%	1.5%
E_T^{miss} Mismodelling	3.4%	7.5%	3.8%	1.6%	1.9%	1.9%	1.3%	1.3%	6.0%	4.9%
Lepton Modelling	2.9%	3.5%	1.7%	3.7%	2.2%	1.5%	2.7%	3.7%	3.1%	3.5%
$t\bar{t}Z$ Theoretical Uncertainties	1.9%	2.7%	2.8%	2.7%	0.73%	7.5%	5.3%	2.7%	1.0%	1.2%
$t\bar{t}$ Normalisation	1.3%	3.7%	4.5%	5.2%	5.4%	1.2%	3.0%	4.5%	5.1%	5.8%
Pileup Reweighting and JVT	0.64%	2.3%	1.8%	1.9%	1.9%	1.1%	0.76%	2.0%	0.70%	0.16%
Flavour Tagging	0.40%	0.96%	0.28%	0.24%	0.23%	0.38%	0.32%	0.34%	0.25%	0.24%
Total uncertainty (stat. + syst.)	74.46%	65.69%	48.67%	47.33%	41.96%	72.03%	73.30%	61.73%	52.89%	48.29%

Table 6.2 – Sources of systematic uncertainty in the SM background estimates, after the background fits. The values are given as relative uncertainties in the total expected background event yields in the different bins of the Run 3 low- Δm SR.

Uncertainty source	SR ₀ ^{medium-1b}	SR ₁ ^{medium-1b}	SR ₂ ^{medium-1b}	SR ₃ ^{medium-1b}	SR ₀ ^{medium-2b}	SR ₁ ^{medium-2b}	SR ₂ ^{medium-2b}	SR ₃ ^{medium-2b}
Total statistical ($\sqrt{N_{\text{exp}}}$)	64.43%	50.38%	42.49%	18.47	68.22%	51.04%	48.04%	24.09%
Total background systematic	12.73%	15.71%	17.28%	9.96%	20.18%	14.37%	22.54%	12.96%
Jet Energy Resolution	9.0%	10.9%	14.1%	8.3%	8.1%	8.8%	15.1%	8.9%
Jet Energy Scale	3.6%	2.6%	5.1%	1.6%	7.0%	5.0%	6.8%	3.7%
$t\bar{t}Z$ Normalisation	3.1%	2.7%	1.9%	1.1%	5.1%	3.4%	2.3%	1.2%
$t\bar{t}Z$ Theoretical Uncertainties	2.5%	1.0%	0.72%	0.37%	8.0%	1.9%	3.8%	0.41%
$t\bar{t}$ Theoretical Uncertainties	1.9%	7.2%	6.3%	3.4%	10.2%	7.1%	11.8%	8.0%
Flavour Tagging	1.4%	1.7%	1.4%	0.84%	1.2%	0.13%	0.41%	0.36%
E_T^{miss} Mismodelling	1.3%	1.4%	1.2%	0.78%	1.9%	2.7%	1.7%	1.1%
Pileup Reweighting and JVT	0.98%	1.1%	0.20%	0.54%	1.3%	0.83%	0.47%	0.90%
Lepton Modelling	0.89%	1.5%	2.0%	1.7%	2.1%	0.76%	2.1%	0.97%
$t\bar{t}$ Normalisation	0.72%	1.5%	2.1%	3.6%	1.3%	2.5%	3.3%	4.5%
Total uncertainty (stat. + syst.)	65.67%	52.77%	45.87%	20.98%	71.14%	53.02%	53.06	27.35%

Table 6.3 – Sources of systematic uncertainty in the SM background estimates, after the background fits. The values are given as relative uncertainties in the total expected background event yields in the different bins of the Run 2 medium- Δm SR.

Uncertainty source	SR ₀ ^{medium-1b}	SR ₁ ^{medium-1b}	SR ₂ ^{medium-1b}	SR ₃ ^{medium-1b}	SR ₀ ^{medium-2b}	SR ₁ ^{medium-2b}	SR ₂ ^{medium-2b}	SR ₃ ^{medium-2b}
Total statistical ($\sqrt{N_{\text{exp}}}$)	69.56%	57.86%	44.31%	33.94%	70.50%	43.67%	52.79%	55.59%
Total background systematic	14.77%	22.90%	20.97%	17.72%	20.08%	27.87%	27.35%	23.78%
Jet Energy Resolution	6.8%	13.0%	15.2%	10.6%	7.8%	16.4%	24.9%	15.7%
$t\bar{t}$ Theoretical Uncertainties	6.5%	10.7%	9.6%	8.9%	6.5%	14.5%	7.4%	9.0%
Pileup Reweighting and JVT	4.6%	1.9%	1.5%	0.84%	0.52%	1.4%	0.62%	1.9%
Jet Energy Scale	4.2%	11.4%	8.3%	8.6%	12.1%	11.9%	9.5%	9.6%
$t\bar{t}Z$ Normalisation	3.9%	2.7%	1.5%	1.2%	6.0%	3.0%	2.0%	1.4%
$t\bar{t}Z$ Theoretical Uncertainties	2.8%	0.34%	2.1%	1.9%	0.95%	1.4%	2.1%	1.5%
$t\bar{t}$ Normalisation	2.2%	3.6%	5.1%	5.7%	2.3%	5.3%	6.3%	7.0%
Lepton Modelling	1.5%	2.6%	5.1%	2.2%	2.5%	3.8%	4.1%	4.0%
Flavour Tagging	0.98%	1.0%	0.71%	0.56%	0.27%	0.42%	0.26%	0.34%
E_T^{miss} Mismodelling	0.38%	1.1%	1.5%	4.5%	2.3%	2.3%	2.9%	2.8%
Total uncertainty (stat. + syst.)	71.11%	62.23%	49.02%	38.29%	73.30%	51.80%	59.45%	60.46%

Table 6.4 – Sources of systematic uncertainty in the SM background estimates, after the background fits. The values are given as relative uncertainties in the total expected background event yields in the different bins of the Run 3 medium- Δm SR.

Uncertainty source	SR ₀ ^{high-1b}	SR ₁ ^{high-1b}	SR ₂ ^{high-1b}	SR ₀ ^{high-2b}	SR ₁ ^{high-2b}	SR ₂ ^{high-2b}
Total statistical ($\sqrt{N_{\text{exp}}}$)	69.56%	55.38%	23.76%	69.05%	62.95%	28.58%
Total background systematic	17.46%	20.51%	9.09%	15.94%	16.27%	12.00%
Jet Energy Resolution	9.4%	16.8%	5.0%	8.1%	8.8%	8.0%
$t\bar{t}$ Theoretical Uncertainties	8.8%	5.1%	4.9%	9.1%	4.4%	5.5%
Lepton Modelling	3.9%	2.3%	2.9%	1.7%	2.5%	3.2%
Jet Energy Scale	3.7%	5.5%	2.0%	1.8%	6.5%	3.6%
Flavour Tagging	3.0%	1.8%	1.4%	0.61%	0.85%	0.73%
$t\bar{t}Z$ Normalisation	2.9%	2.6%	1.5%	3.8%	2.7%	1.4%
$t\bar{t}Z$ Theoretical Uncertainties	2.8%	2.5%	0.48%	3.8%	2.1%	0.71%
MET Mismodelling	1.1%	1.1%	1.4%	0.84%	1.5%	2.1%
$t\bar{t}$ Normalisation	0.98%	1.6%	2.4%	1.9%	2.4%	3.5%
Pileup Reweighting and JVT	0.24%	0.66%	0.62%	0.31%	1.6%	1.2%
Total uncertainty (stat. + syst.)	71.72%	59.06%	25.44%	70.87%	65.02%	31.00%

Table 6.5 – Sources of systematic uncertainty in the SM background estimates, after the background fits. The values are given as relative uncertainties in the total expected background event yields in the different bins of the Run 2 high- Δm SR.

Uncertainty source	SR ₀ ^{high-1b}	SR ₁ ^{high-1b}	SR ₂ ^{high-1b}	SR ₀ ^{high-2b}	SR ₁ ^{high-2b}	SR ₂ ^{high-2b}
Total statistical ($\sqrt{N_{\text{exp}}}$)	68.37%	58.62%	40.32%	66.37%	68.37%	57.67%
Total background systematic	20.23%	23.49%	16.93%	25.22%	24.10%	29.09%
$t\bar{t}$ Theoretical Uncertainties	12.8%	12.0%	8.6%	13.4%	7.3%	15.8%
Jet Energy Resolution	9.3%	11.9%	9.0%	15.6%	16.3%	21.4%
Jet Energy Scale	5.8%	14.0%	7.5%	9.6%	7.8%	6.5%
Lepton Modelling	5.6%	3.1%	5.3%	6.1%	11.7%	9.0%
$t\bar{t}Z$ Theoretical Uncertainties	5.1%	3.1%	3.0%	2.9%	3.6%	0.85%
Pileup Reweighting and JVT	5.0%	1.7%	0.04%	1.1%	4.6%	1.5%
$t\bar{t}Z$ Normalisation	3.6%	2.2%	1.3%	4.0%	2.3%	1.8%
$t\bar{t}$ Normalisation	2.2%	3.7%	3.8%	3.4%	4.6%	5.3%
Flavour Tagging	1.5%	1.9%	1.0%	1.2%	0.69%	0.69%
E_T^{miss} Mismodelling	0.98%	4.2%	1.7%	2.0%	0.90%	1.1%
Total uncertainty (stat. + syst.)	71.30%	63.15%	43.73%	71.00%	72.49%	64.59%

Table 6.6 – Sources of systematic uncertainty in the SM background estimates, after the background fits. The values are given as relative uncertainties in the total expected background event yields in the different bins of the Run 3 high- Δm SR.

6.3 Results and interpretation

As discussed in Section 6.1, the statistical results of the analysis are obtained through simultaneous maximum likelihood fits performed independently for each of the three Δm ranges, and separately for the Run 2 and Run 3 datasets. The expected and observed event yields obtained from the background-only fits are reported in Section 6.3.1. Since no significant excesses with respect to the SM prediction are observed, the results are interpreted in terms of model-dependent exclusion limits on all the considered stop signal models in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ parameter space.

To obtain the combined Run 2 and Run 3 result, the fits performed independently for the two datasets are combined separately within each Δm region. The final exclusion limits are then determined by selecting, for each signal point in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane, the most stringent expected limit among the three Δm regions. This procedure is applied to the Run 2, Run 3, and combined datasets. In addition, the analysis is reinterpreted using an alternative signal model that leads to the same final state signature as the stop pair production with dileptonic decay of the $t\bar{t}$ pair. The model consists of the production of a Higgs boson in association with a $t\bar{t}$ pair, with invisible decay of the Higgs boson. A more detailed description of this model and the corresponding results are presented in Section 6.3.3. All the results presented in this section include the all the systematic uncertainties described in Section 6.2.1.

6.3.1 Event yields

Figure 6.3 presents the results of the background-only fits for each analysis region on all the Δm categories, for both the Run 2 and the Run 3 datasets. Good overall agreement is observed between data and the SM expectation in the VRs, with the notable exception of the $\text{VR}_{t\bar{t}-2b}$ region for the Run 3 medium- Δm fit. This discrepancy is interpreted as a statistical fluctuation, as it appears in only one of the six fits, meaning that it cannot be interpreted as a mismodelling issue, and the fraction of $t\bar{t}Z$ events over the total number of the events in this VR is the same as in the corresponding Run 2 fit region (approximately 80%).

In the SRs, small excesses of around 2σ are observed in several bins. These too can be safely interpreted as statistical fluctuations: they do not become larger in the tightest SR bins (SR_0), where the expected signal purity is the highest. Additionally, when excesses occur in SR bins containing events with one b -tagged jet, no corresponding excess is observed in the bins with events containing more than one b -jet.

The NFs for the $t\bar{t}$ ($\mu_{t\bar{t}}$) and the $t\bar{t}Z$ ($\mu_{t\bar{t}Z}$) backgrounds, obtained from the Run 2 and the Run 3 background-only fits, are reported in Table 6.7. For both datasets, the $t\bar{t}$ NFs extracted in the different Δm categories are consistent with each other within statistical uncertainties. More plots showing the fitted values for all the NPs and the correlations between them are shown in Appendix A.




Figure 6.3 – Expected yields in the three non-orthogonal Δm categories for both the Run 2 (left column) and the Run 3 (right column) datasets. The upper panels show the comparison between the observed number of data events and the post-fit SM prediction of the background in the CRs and VRs (SRs are blinded). For CRs and VRs, the bottom panel shows the difference between data and the predicted SM background divided by the total uncertainty (σ_{tot})

NF	Value	NF	Value
Low $\mu_{t\bar{t}}$	0.86 ± 0.05	Low $\mu_{t\bar{t}}$	1.11 ± 0.08
Medium $\mu_{t\bar{t}}$	0.91 ± 0.06	Medium $\mu_{t\bar{t}}$	0.99 ± 0.09
High $\mu_{t\bar{t}}$	0.95 ± 0.05	High $\mu_{t\bar{t}}$	1.10 ± 0.09
Low $\mu_{t\bar{t}Z}$	1.08 ± 0.10	Low $\mu_{t\bar{t}Z}$	1.06 ± 0.13
Medium $\mu_{t\bar{t}Z}$	1.08 ± 0.10	Medium $\mu_{t\bar{t}Z}$	1.08 ± 0.14
High $\mu_{t\bar{t}Z}$	1.08 ± 0.10	High $\mu_{t\bar{t}Z}$	1.07 ± 0.14

Table 6.7 – NFs for $t\bar{t}$ and $t\bar{t}Z$ backgrounds obtained from the Run 2 (left) and Run 3 (right) background-only fits.

6.3.2 Interpretation

The results of the analysis are interpreted in terms of model-dependent exclusion limits in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane based on the stop simplified models considered.

Model-dependent limits are derived for the benchmark simplified stop signal models introduced in Figure 4.3. The sensitivity of the analysis is also tested on signal models where the mass difference $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0)$ lies between $m_W + m_b$ and m_t , corresponding to the kinematic regime of three-body decays. The hypothesis tests are performed including the expected signal yields and their associated uncertainties in both the CRs and SRs. To enhance statistical power, the NN signal score bins of the SRs are statistically combined. The limits are derived using the CL_s method described in Section 6.1, and all results are quoted at 95% CL. In setting exclusion limits, the event regions with one b -jet and those with more than one b -jet are combined, as they are statistically independent. The limits are evaluated separately for the three Δm regions, with the region yielding the best expected sensitivity used to define the final exclusion. Expected CL_s at 95% CL are obtained using the median of the background-only pdf as pseudo-data in the SRs, and provides an estimate of the expected exclusion sensitivity of the analysis. A model is considered excluded at 95% CL if the corresponding CL_s is less than 0.05.

Expected and observed exclusion contours in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane are shown for simplified models where the $\tilde{t}_1$ decays with 100% branching ratio into an on-shell top quark and a $\tilde{\chi}_1^0$ (stop two-body decay, targeted in the analysis optimisation), as well as simplified models where the $\tilde{t}_1$ pair decays with 100% branching ratio into a W boson and a b quark via off-shell top, and a $\tilde{\chi}_1^0$ (three-body stop decay, not targeted in the analysis optimisation). Figure 6.4 presents the expected and observed exclusion limits derived separately for Run 2 and Run 3 datasets, corresponding to integrated luminosities of 140 fb^{-1} and 53 fb^{-1} , respectively. Figure 6.5 shows the expected and observed limits obtained from the combination of both datasets, considering the full 193 fb^{-1} of integrated luminosity.

The Run 2 only limits are sensitive to models with top squark masses up to 1090 GeV for nearly massless $\tilde{\chi}_1^0$, and with $\tilde{\chi}_1^0$ masses up to 500 GeV for $m_{\tilde{t}_1} = 800 \text{ GeV}$. This represents a significant improvement over the previous Run 2 analysis, which was sensitive only to top squark masses up to 970 GeV for nearly massless $\tilde{\chi}_1^0$, and to $\tilde{\chi}_1^0$ masses up to 400 GeV for $m_{\tilde{t}_1} = 700 \text{ GeV}$. It is interesting to observe that the analysis also shows sensitivity in the three-body decay region, closing the exclusion gap previously present between the two-body and three-body regions.

The observed Run 2 limits present a discrepancy with respect to the expected limits, in particular in the high stop mass region, excluding only models with $m_{\tilde{t}_1}$ up to 1 TeV for a nearly massless $\tilde{\chi}_1^0$. This is consistent with the 2σ excesses in data relative to the SM expectation discussed in Section 6.3.1.

For the Run 3 only dataset, the expected sensitivity reaches models with top squark masses up to almost 1 TeV for nearly massless $\tilde{\chi}_1^0$, and with $\tilde{\chi}_1^0$ masses up to 500 GeV for $m_{\tilde{t}_1} = 800$ GeV. Therefore, despite the lower luminosity, the Run 3 dataset gives a valuable contribution to the analysis sensitivity. In this case, the observed limit slightly exceeds the expected one in the high stop mass region, reaching stop masses of 1030 GeV, while in the compressed region only $\tilde{\chi}_1^0$ masses up to 400 GeV are excluded for $m_{\tilde{t}_1} = 700$ GeV. The shape of the observed contour shows a step near the parameter space point $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (1000, 200)$ GeV, arising from the choice of the observed CL_s associated with the lowest expected CL_s among the three Δm regions, and indicating a shift between the high- Δm and the medium- Δm model.

Finally, the combined Run 2 and Run 3 dataset allows to extend the expected sensitivity to models with top squark masses up to 1120 GeV for nearly massless $\tilde{\chi}_1^0$, and with $m_{\tilde{\chi}_1^0}$ up to 600 GeV for $m_{\tilde{t}_1} = 900$ GeV. The observed limits, however, exclude only models with top squark masses up to 1050 GeV for nearly massless $\tilde{\chi}_1^0$, and with $m_{\tilde{\chi}_1^0}$ up to 500 GeV for $m_{\tilde{t}_1} = 800$ GeV, reflecting the same discrepancy between expected and observed limits noted for the Run 2 only results. Notably, the analysis achieves strong exclusion power even in the three-body decay region, which was not specifically targeted in the optimisation phase. In this scenario, the expected sensitivity extends to models with $m_{\tilde{\chi}_1^0}$ up to 450 GeV for $m_{\tilde{t}_1} = 600$ GeV.

These results significantly extend the reach of previous stop searches in the dileptonic final state, both at high stop masses and in the particularly challenging compressed kinematic region, demonstrating the effectiveness of the new analysis strategy. Additionally, despite the lower cross section of the dileptonic final state, the analysis on the full 193 fb^{-1} dataset achieves a sensitivity competitive with the previous Run 2 one-lepton analysis [35] when considering the compressed region, extending the expected sensitivity to $\tilde{\chi}_1^0$ masses up to 600 GeV for $m_{\tilde{t}_1} = 900$ GeV, which will give a valuable contribution when the two channels will be combined.

Figure 6.6 shows which model gives the best expected sensitivity for each point of the signal grid, and is thus used to establish the final result, for the Run 2 only dataset, the Run 3 only dataset, and combined Run 2 and Run 3 dataset. In all cases, the models tend to cover the signal points on which they were trained.




Figure 6.4 – Exclusion limit contour (95% CL) for a simplified model assuming $\tilde{t}_1$ pair production, decaying via $\tilde{t}_1 \rightarrow t\tilde{\chi}_1^0$ with 100% branching ratio, in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane, considering only the Run 2 (left) and only the Run 3 (right) dataset. The dashed lines and the shaded yellow bands represent the expected limits and their 1σ uncertainties. The red solid lines represent the observed limits. The limits are also shown in the stop three-body decay region. The grey shaded region shows the limits from the previous Run 2 analysis [111].




Figure 6.5 – Exclusion limit contour (95% CL) for a simplified model assuming $\tilde{t}_1$ pair production, decaying via $\tilde{t}_1 \rightarrow t\tilde{\chi}_1^0$ with 100% branching ratio, in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane, considering the combined Run 2 and Run 3 datasets. The dashed line and the shaded yellow band represent the expected limits and their 1σ uncertainties. The red solid line represents the observed limits. The limits are also shown in the stop three-body decay region. The grey shaded region shows the limits from the previous Run 2 analysis [111].




Figure 6.6 – Red, blue, and green points are the points for which the low, medium, and high- Δm models, respectively, give the best expected CL. The study is shown for Run 2 only (left), Run 3 only (centre) and combining Run 2 and Run 3 (right).

6.3.3 Invisible Higgs reinterpretation

One of the motivations of interest of the stop pair production searches, is that their final state signature, with $t\bar{t}$ and missing transverse momentum, is very versatile and allows for many possible reinterpretations. The analysis presented in this thesis is reinterpreted considering a model with an invisibly decaying Higgs boson produced in association with a top-antitop pair ($t\bar{t}H \rightarrow \text{inv}$), which presents the same $t\bar{t} + E_T^{\text{miss}}$ detector signature of the $\tilde{t}_1 \rightarrow t\tilde{\chi}_1^0$ model.

Many extensions of the SM predict the production of dark matter particles at the LHC. Sufficiently light dark matter particles may be produced in decays of the Higgs boson that would appear invisible to the detector. This decay channel is therefore called *invisible*. In the SM, the Higgs invisible decay happens when the Higgs boson decays into two Z bosons, that in turn decay into two neutrinos each: $HH \rightarrow ZZ \rightarrow 4\nu$ [151]. The SM branching ratio of the Higgs invisible decay is approximately 0.1% [151]. A higher branching ratio to invisible particles may imply the presence of BSM particles, and is predicted, for example, by Higgs to dark-matter portal models.

A recent ATLAS result [152], analysing the full Run 2 dataset, consists in the statistical combination of all ATLAS direct searches for invisible decays of the Higgs boson using the full Run 2 dataset. This includes the gluon-gluon fusion [153], VBF [154, 155], ZH [156] and $t\bar{t}H$ [157] production modes. The result also includes a statistical combination with the ATLAS combined Run 1 result [158], yielding the most sensitive direct constraint to invisible Higgs boson decays in ATLAS. The observed and expected upper limits at the 95% CL on the invisible decay of the Higgs boson set thanks to this combined result are 0.107 and 0.077, respectively. The contribution from the $t\bar{t}H$ channel, combining all the $t\bar{t}$ decay channels (2 leptons, 1 lepton, and 0 leptons in the final state) is $\mathcal{B}_{H \rightarrow \text{inv}} = 0.376(0.295)$ of observed (expected) upper limit at the 95% CL. The result for the dileptonic channel [157] is an observed (expected) upper limit on the Higgs boson invisible branching ratio of 0.39(0.42) at 95% CL.

The following paragraph briefly introduces the theoretical framework and the signal MC simulated samples employed to model the $t\bar{t}H \rightarrow \text{inv}$ signal. Then the results of the analysis, that is repeated without any changes with respect to the stop analysis discussed earlier, are presented.




Figure 6.7 – Example tree-level Feynman diagrams for the $pp \rightarrow t\bar{t}H$ production process.

Signal model and simulation

The production mechanisms of Higgs bosons in the LHC pp collisions are: gluon-gluon fusion (ggF), which is the production mode with the largest cross section, followed by vector boson fusion (VBF), vector boson associated production (VH, including WH and ZH), top-antitop quark pair associated production ($t\bar{t}H$), which is the production mode considered for this reinterpretation, bottom-antibottom quark pair associated production ($b\bar{b}H$), and single top quark associated production (tHW and tHq). Figure 6.7 presents tree-level Feynman diagrams contributing to the $t\bar{t}H$ production mode in pp collisions. Figure 6.8a shows the production cross sections of the different Higgs production mechanisms in pp collisions, as a function of the centre-of-mass energy of the pp collisions, assuming a Higgs mass $m_H = 125 \text{ GeV}$. The $t\bar{t}H$ production mode is one of those with the lowest cross section.

The Higgs boson has a very short lifetime (approximately $1.6 \times 10^{-22} \text{ s}$, according to the SM [151]). The main decay modes for the SM Higgs boson are $H \rightarrow b\bar{b}$, which is the leading decay mode with a branching ratio of 58.24%, followed by the $H \rightarrow WW^*$ and $H \rightarrow ZZ^*$ (with one of the weak bosons off-shell), and decays into massless gauge bosons (gluons and photons), which are also possible and occur via a loop of heavy quarks or W bosons. Finally, all the heavy fermions can constitute final states for a Higgs boson decay (except for the top quark, which is too heavy), and the corresponding branching ratio is proportional to the fermion masses. These decay modes, along with their branching ratios, are presented in Figure 6.8b as a function of the Higgs boson mass in a window between 120 GeV and 130 GeV. Considering that the branching ratio for $H \rightarrow ZZ^*$ is approximately 2.6% and the branching ratio for the $Z \rightarrow \nu\nu$ decay is around 20%, one obtains the total branching ratio for the Higgs to invisible decay of $\mathcal{B}_{H \rightarrow \text{inv}} \approx 0.1\%$.

The signal samples modelling the 125 GeV Higgs invisible decays are generated using POWHEG-BOX v2 at NLO with the NNPDF3.0NLO PDF set. The Higgs boson decays via ZZ^* to neutrinos are normalised using the total $t\bar{t}H$ cross section at NLO QCD and electroweak accuracy recommended by the LHC Higgs cross section working group [151]. Only the associated production of the Higgs boson with two top quarks is considered in the simulated samples, neglecting rarer processes such as tH and tWH .

Figure 6.9 shows a comparison of the distributions of $m_{\tau_2}^{\ell\ell}$ and E_T^{miss} significance in Run 2 MC samples for the $t\bar{t}$ background, the $t\bar{t}H \rightarrow \text{inv}$ signal, and two benchmark stop signal models with $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (700, 500) \text{ GeV}$ and $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (900, 500) \text{ GeV}$, representing a low- and a high- Δm scenario, respectively. The $t\bar{t}H \rightarrow \text{inv}$ signal distributions of these variables closely




Figure 6.8 – Cross sections for the SM Higgs boson production modes as a function of the centre-of-mass energy of pp collisions assuming a Higgs boson mass of 125 GeV (left), and branching ratio for the main decay modes of the SM Higgs boson as a function of its mass (right) [159].

resemble those of the low- Δm stop model. Since $m_{t\bar{t}}^{\ell\ell}$ and E_T^{miss} significance are among the most powerful discriminating variables in the stop analysis, the best expected sensitivity to the $t\bar{t}H \rightarrow \text{inv}$ signal is therefore anticipated to be achieved by the fit in the low- Δm regions.

Upper limits

The analysis results are interpreted for the signal model of the Higgs boson production in association with a top-antitop quark pair which decays invisibly with $\mathcal{B}_{H \rightarrow \text{inv}} = 100\%$. It should be noted that the quoted direct limit relies on the assumption that, in the presence of new physics, the cross section of the $t\bar{t}H$ process remains unchanged with respect to the SM expectation. The 95% CL upper limits are calculated using the CL_s method described in Section 6.1.2. The results are presented in terms of the signal strength, $\mu_{t\bar{t}H}$, which is the ratio of the upper limit cross section and the theoretical cross section for $\mathcal{B}_{H \rightarrow \text{inv}} = 100\%$, and thus can be interpreted as the upper limit on $\mathcal{B}_{H \rightarrow \text{inv}}$. The results do not include theoretical systematic uncertainties on signal. Figure 6.10 show the expected CL_s value as a function of the $t\bar{t}H \rightarrow \text{inv}$ signal strength $\mu_{t\bar{t}H}$. The expected upper limit at 95% CL, which corresponds to a $\text{CL}_s = 0.05$, is $\mu_{t\bar{t}H} = 0.333^{+0.148}_{-0.099}$ for the Run 2 dataset, $\mu_{t\bar{t}H} = 0.639^{+0.299}_{-0.197}$ for the Run 3 dataset, and $\mu_{t\bar{t}H} = 0.287^{+0.126}_{-0.085}$ from the combination of the two datasets. The new Run 2 result improves the previous expected limit [157] for the $t\bar{t}H \rightarrow \text{inv}$ branching ratio in the dileptonic channel, that amounted to $\mu_{t\bar{t}H} = 0.40^{+0.18}_{-0.12}$, and the new expected limit obtained using the combined datasets improves the previous $t\bar{t}H$ expected upper limit obtained combining all the $t\bar{t}$ decay channels (dileptonic, semileptonic and fully-hadronic), that corresponds to $\mu_{t\bar{t}H} = 0.30^{+0.13}_{-0.09}$ [157].

The observed upper limit is $\mu_{t\bar{t}H} = 0.60$ for the Run 2 dataset, $\mu_{t\bar{t}H} = 1.00$ for the Run 3 dataset, and $\mu_{t\bar{t}H} = 0.58$ from the combination of the two datasets. As already discussed in Section 6.3.2, the large discrepancy between the expected and observed limits is due to the 2σ excesses in data relative to the SM expectation observed in the post-fit event yields.




Figure 6.9 – Distribution of $m_{T2}^{\ell\ell}$ and E_T^{miss} significance in Run 2 MC samples comparing $\tilde{t}\tilde{t}$, $\tilde{t}\tilde{t}H$, and two stop signal points with $\Delta m = 200$ GeV (low- Δm range) and $\Delta m = 400$ GeV (medium- Δm range).




Figure 6.10 – Expected upper CL_s exclusion limits on $\mathcal{B}_{H \rightarrow \text{inv}}$ at 95% CL for the Run 2 only dataset (upper left), the Run 3 only dataset (upper right) only dataset, and combining the Run 2 and the Run 3 results (bottom). Both statistical and systematic uncertainties are included.

Part III

Trigger Menu and Operations

LHC experiments have been constructed to search for extremely rare processes. As already discussed in Section 2.2.6, the amount of data produced by the LHC collisions in the ATLAS detector significantly exceeds the technical capabilities for the transfer, storage and post-processing of the ATLAS computing systems. Staying within these limitations is possible by constraining the output data to contain only the fraction of rare events of interest, which is the goal of the trigger system. For this reason, the trigger system plays a central role in the data-taking strategy of the experiment, and the design, configuration, and operation of the trigger system are crucial for the overall physics performance of ATLAS. This chapter provides an overview of the HLT system, which operates as the final stage of the online event selection, focusing on its Run 3 implementation. In particular, Section 7.1 introduces the physics motivations that drive the definition of the trigger strategy, as well as the performance metrics that guide the design of trigger systems. After that, Section 7.2 presents an overview of the structure of the ATLAS HLT software and its execution within the *Athena* software framework, followed by a discussion on the available tools to monitor the HLT performance in Section 7.3. Finally, Section 7.4 describes how the trigger system interacts with other ATLAS subsystems in daily detector operations, including how the trigger configuration is stored and accessed during data-taking.

7.1 Physics motivation

A primary goal of the ATLAS physics programme is the exploration of extremely rare processes. The trigger system selects from the LHC collision data the events that are most relevant according to the physics priorities of the experiment. The trigger strategy needs to reflect the main physics goals of the experiment. ATLAS is a multipurpose detector, designed to address a wide spectrum of physics topics simultaneously. Therefore, its trigger system must have the capability of applying concurrently multiple selection strategies, to be able to collect data for multiple physics measurements at the same time, maximising the scientific outcome. The goals of the pp physics ATLAS programme include the studies on the EWSB through precision measurements of the properties of the Higgs boson, improved measurements of all the relevant SM parameters, including the study of rare SM processes, searches for BSM signatures and flavour physics. Table 7.1 presents an overview of the ATLAS pp physics goals, along with some representative analyses and their associated trigger signatures. General trigger signatures (e.g. the single lepton triggers used in the analysis presented previously in this thesis) are favoured in terms of

resources, since they accommodate the needs of multiple key analyses at the same time. However, the trigger system is also required to be able to select more complex signatures. More details on how resources are allocated for different physics goals are provided in Chapter 8.

The performance of a trigger system can be optimised according to three quantities: high

ATLAS Physics goals		Processes	Trigger Signatures
Precision measurements of the properties of the Higgs Boson	Couplings to fermions	$H \rightarrow \tau\tau, H \rightarrow \mu\mu, ttH, H \rightarrow bb$	single/di- e or μ / di- τ
	Couplings to bosons, cross sections	$H \rightarrow \gamma\gamma, H \rightarrow WW \rightarrow \ell\nu\ell\nu, H \rightarrow ZZ \rightarrow \ell\ell\ell\ell$	e/μ , di- γ
	Self-coupling	$HH \rightarrow bb\tau\tau / bb\gamma\gamma / 4b$	di- τ/γ , multi-jets
	Scalar Higgs boson vs. BSM composite	$H \rightarrow \ell\ell', ZH \rightarrow \ell\ell + (\text{inv})$	e/μ
Precision Standard Model Measurements	Forward/backward asymmetry	$Z \rightarrow e^+e^-, \mu^+\mu^-$	single e/μ
	Vector-boson scattering	$WWjj, WZjj$	single e/μ
	Precision top mass and cross sections	tt production	e/μ , large R-jets/multi-jets
Searches for BSM Signatures	Searches for new vector bosons	Vector Boson Fusion (VBF) $Z' \rightarrow \ell\ell$	high- p_T single e/μ
	Searches for electroweak SUSY	$\chi_1^\pm \chi_2^0 \rightarrow WH \chi_1^0 \chi_1^0$	E_T^{miss} , single/di- e, μ, τ
	SUSY top partners	$\tilde{t}_1 \rightarrow t \chi_1^0$	large-R jets/multi-jets+ E_T^{miss}
	Dark matter	$\text{ISR} + \chi_1^0 \chi_1^0$	jets+ E_T^{miss}
	New resonances, SUSY	$Z', \chi_1^0 \rightarrow jjj$	jets, large-R jets, $e/\mu, \gamma$
	Long-lived particles	$\tilde{g} \rightarrow qq \chi_1^0, \tilde{\chi}_1^\pm \rightarrow \pi^\pm \chi_1^0$	high impact parameter, E_T^{miss}
Flavour Physics	Lepton Flavour Violation	$\tau \rightarrow \mu\mu\mu$	low p_T di- μ
	Searches for FCNC in top decays	$t \rightarrow u/c + H/Z$	single e/μ
	Rare B -meson decays	$B \rightarrow \mu\mu, B_s \rightarrow J/\Psi + \Phi$	low p_T di- μ

Table 7.1 – ATLAS physics goals, with key processes and relative trigger signatures.

signal efficiency, high background rejection, and affordable output bandwidth. Moreover, this performance must be quantifiable (e.g. that trigger efficiencies must be calculable), robust and deterministic [6].

The trigger **signal efficiency** quantifies the probability of a positive trigger decision under the condition of the presence of a signal in the event. The efficiency estimator, $\epsilon_{\text{trigger}}$, is defined as:

$$\epsilon_{\text{trigger}} = \frac{N_{\text{trigger}}}{N_{\text{signal}}} \quad (7.1)$$

where N_{trigger} is the number of signal events leading to a positive trigger decisions, and N_{signal} is the number of signal events in the sample. A well-performing trigger system aims for an efficiency ideally close to 1 for the relevant signal processes, ensuring minimal signal loss. Precise knowledge of the trigger efficiency allows to correctly relate the number of recorded events to the true number of events produced in the physics process. Trigger efficiencies are typically studied as a function of key kinematic variables of the signal objects, such as the transverse momentum or pseudorapidity, as reconstructed offline (i.e., using the full detector information and calibration available after the event is stored). For an ideal trigger, the efficiency curve as a function of the offline p_T would appear as a step function, equal to zero below the p_T threshold and equal to one above. In reality, the transition is not instantaneous due to resolution inefficiencies and differences between the online and offline reconstruction. As a result, real trigger efficiency curves as a function of the offline p_T exhibit a characteristic turn-on shape, which can often be described by an *error function*. This curve is typically divided into three regions.

- **Inefficiency region:** covers p_T values under the trigger threshold, events in this region are




Figure 7.1 – Trigger efficiency curve for an ideal trigger (step-function, in blue) and a real trigger (turn-on function, in orange). The represented efficiency curves refer to a trigger with a p_T threshold of 30 GeV.

rejected by the trigger.

- **Turn-on region:** region around the trigger threshold, where the efficiency rises steeply with increasing p_T . This region needs to be modelled with precision to have good understanding of signal efficiency and avoid significant background contamination.
- **Plateau region:** region where the trigger efficiency reaches a stable maximum. Most physics analyses apply offline selections to only retain objects from the plateau region, to ensure uniform trigger performance across the dataset.

The design goal for a trigger system is to reach zero efficiency in the inefficiency region, efficiency as close as possible to one in the plateau region, and the turn-on region as narrow as possible. Figure 7.1 shows an example of different efficiency curve for an ideal trigger (step-function) and a real trigger (turn-on function). One important source of inefficiency for the ATLAS trigger system arises from the incomplete phase space coverage of the triggering detectors. In addition, due to the strict latency requirements of online event selection, trigger algorithms are simplified versions of the more precise offline reconstruction algorithms. As a result, the physical quantities used for online selection are only correlated with the corresponding offline variables. The poorer this correlation, the broader the turn-on region becomes in the efficiency curve. Another contributing factor is that online reconstruction operates using calibrations that are not fully up to date. Since calibrations evolve over time and are often derived from the same data they are meant to correct, the trigger system must rely on the best available knowledge at the time of data-taking. This introduces additional differences between online and offline variables, further broadening the efficiency turn-on. Finally, many offline reference quantities used for efficiency studies are only defined after data-taking, introducing an unavoidable offline-online mismatch.

As discussed earlier, the trigger system also needs to be able to collect auxiliary information to measure the trigger efficiency from data. Several data-driven methods are employed for this purpose. The *orthogonal trigger method* uses a sample of events where the signal for which the




Figure 7.2 – Turn-on curves of the lowest-threshold single electron trigger [160] (left) and single muon triggers [161] (right), used to collect data in 2022. The plots also show the efficiency of the L1 seed. Efficiencies are shown as a function of the offline electron transverse energy (E_T) and data compared against SM $Z \rightarrow ee$ and $Z \rightarrow \mu\mu$ MC events for the electron trigger and the muon trigger, respectively. In both cases, the tag-and-probe method is used to measure the trigger efficiency.

trigger is designed is available, selected triggering on another physics object present in the target signal signature, to avoid being biased towards the trigger under study (e.g. a sample can be created using a calorimetric E_T^{miss} trigger, and selecting events with W boson decays which are then used to study the efficiency of lepton triggers aiming at the second decay product of the W boson). The *bootstrap method*, where the efficiency of a trigger is measured using a lower-threshold trigger of the same type (e.g. a jet trigger with a 40 GeV p_T -threshold used to measure the efficiency of a jet trigger with a 60 GeV p_T -threshold). The *tag and probe method*, which measures the efficiency of single-object triggers exploiting cases where the signal appears twice in the event (e.g. $Z \rightarrow ee$ events for single-electron triggers, and $Z \rightarrow \mu\mu$ events for single-muon triggers). These type of events are selected from a sample recorded with the single-object trigger under study, and the trigger decision for the second object (probe object) is checked. The efficiency is computed as the fraction of probe objects selected by the trigger. Figure 7.2 shows the turn-on curves of two 2022 triggers used in the analysis previously presented in this thesis, measured using the tag-and-probe method.

The efficiency of single-lepton triggers in selecting stop signal events is shown in Figure 7.3 as a function of the stop and neutralino masses. The efficiencies are computed from signal MC samples simulating 2022 data, and are defined as the ratio of the number of events selected by the trigger to the total number of generated signal events. The selection considers the lowest-threshold single-electron and single-muon triggers available in 2022:

HLT_e26_lhtight_ivarloose_L1EM22VHI and HLT_mu24_ivarmedium_L1MU14FCH, respectively. An event passes the trigger selection if it is selected by at least one of these two triggers.

The HLT_e26_lhtight_ivarloose_L1EM22VHI trigger applies a 26 GeV transverse momentum threshold at the HLT level and is seeded by a L1 EM trigger with a 22 GeV threshold.

The HLT_mu24_ivarmedium_L1MU14FCH trigger instead applies a 24 GeV HLT threshold and is seeded by a 14 GeV L1 muon trigger. The loose event preselection described in Section 5.1.1 is applied before evaluating the trigger efficiency. The figure shows that the single-lepton triggers




Figure 7.3 – Trigger signal efficiency for the signal samples used in the analysis presented in this thesis in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane.

have a very high acceptance for signal events, above 97% for all the signal models in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane. The efficiency tends to decrease in the compressed mass region, where the mass difference between the stop and the neutralino approaches the top quark mass. In this region, the kinematics of the final state tend to produce lower- p_T leptons, which are more likely to fall below the trigger thresholds and escape selection.

7.2 The ATLAS HLT

The ATLAS HLT software code is part of the ATLAS software framework, *Athena* [113], which is based on the inter-experiment framework *Gaudi* [162]. Athena is used for all the steps of the event data processing in ATLAS, including detector simulation, event reconstruction and physics analysis (collectively known as *offline* processing), as well as for the real-time selections performed by the HLT (referred to as *online* processing).

The HLT selection process is organised into trigger *chains*. Each chain is seeded by a L1 decision and consists of a sequence of *steps*. Each step performs part of the event reconstruction (via *reconstruction algorithms*) and applies corresponding selection criteria (via *hypothesis algorithms*). Within each step, the control flow and data dependencies ensure that the algorithms are executed in a fixed order: first, a filter algorithm determines whether the step should run, based on the outcome of previous selections; then, an input-maker algorithm prepares the necessary input data for the reconstruction; after that, reconstruction algorithms extract the relevant physics features from the detector data; finally, hypothesis algorithms evaluate whether the reconstructed objects satisfy the selection criteria defined by the chain. Figure 7.4 shows schematically chain steps for electron and muon reconstruction and selection. Each trigger chain corresponds to one path through the graph, and the chains can be individually enabled or disabled during run time or executed on only a fraction of events. The design of the HLT soft-




Figure 7.4 – Example of control flow for menu processing. The control flow graph is created during initialisation and the steps are executed based on the available data. If a filter passes, processing continues through the next steps until all filters fail where the processing of the following steps is stopped. If the last step is reached with a chain passing all of its steps the event will be accepted. Figure adapted from Reference [163].

ware for online reconstruction has to satisfy specific constraints and optimisations due to the CPU resource limitations of the HLT farm, which require significantly shorter per-event processing times at the HLT compared to offline reconstruction. In particular, the design is based on the principles outlined below.

- **Reconstruction in RoI.** Where possible, reconstruction is performed only within the restricted RoIs identified by the L1 trigger (see Section 2.2.6), avoiding the full-detector reconstruction typical of offline processing.
- **Early chain termination.** The reconstruction process for a given trigger chain is discontinued as soon as a selection step (hypothesis algorithm) fails.
- **Early event rejection.** The processing of an event is terminated as soon as it is determined that the event failed all active trigger selections; that is, when no trigger chains remain active after a given step.

A direct consequence of early chain termination and early event rejection is that algorithms later in a chain are executed at lower rates than those earlier in the chain. This allows for a progressive allocation of computing resources: reconstruction steps occurring at lower rates are permitted longer processing times and can perform more CPU-intensive operations. To exploit this structure efficiently, when possible, trigger chains are arranged so that fast, high-rejection steps are placed early, while slower, more precise steps are deferred to later stages. Algorithmic code specifically developed for the trigger is used for some time-critical functionalities, typically executed early in the selection. Examples include fast calorimeter data preparation and clustering, and fast ID and muon tracking, which run early in chains. These fast reconstruction algorithms are followed by selection steps that significantly reduce the event rate, enabling subsequent steps to use slower, more precise reconstruction, based on offline algorithms (in some cases with a trigger-specific configuration), allowing tighter selection requirements to be applied. This operational model of fast, lower resolution reconstruction and loose selection followed by more precise reconstruction and tighter selection is an important optimisation strategy that minimises the total CPU requirements of the HLT farm. Furthermore, the use of offline code in the HLT maximises the correlation between online and offline selections, important to narrow the turn-on region of the trigger efficiency curve (see Section 7.1).

During Run 2, Athena employed a multiprocessing model, known as (*AthenaMP*). In this approach, the main process is forked after initialisation into a number of worker processes, typically equal to the number of events which should be processed in parallel. Each worker processes events independently using a single thread, sharing read-only memory with other workers. This mechanism already allowed for a significant reduction in the total memory usage compared to running fully independent processes for each event. For Run 3, the Athena framework was upgraded to support multithreaded execution (*AthenaMT*), based on the multithreaded version of Gaudi [164], itself relying on Intel's Threading Building Blocks [165]. In this approach, a single process can use multiple threads, which share both read-only and writable memory. *AthenaMT* allows for greater memory sharing across computing cores compared to *AthenaMP* and, consequently, greater flexibility and efficiency when running on hardware with limited memory per core. Consequently, an additional advantage of using *AthenaMT* is that it facilitates potential future use of compute accelerators [166], such as GPUs and FPGAs, as its

structure them to be integrated without major architectural changes. Three types of parallelism are currently supported in AthenaMT: *inter-event parallelism* (multiple events can be processed in parallel), *intra-event parallelism* (multiple algorithms can run in parallel within a single event) and *in-algorithm parallelism* (possible internal multithreading within a single algorithm). To fully benefit from the capabilities of AthenaMT, the ATLAS HLT software was largely restructured for Run 3. While in the Run 2 HLT framework all trigger algorithms relied on HLT-specific interfaces (and offline algorithms needed wrappers for integration), the Run 3 framework no longer requires such interfaces, and takes full advantage of the AthenaMT scheduler and other AthenaMT components to control the HLT event processing. These modifications simplified the integration of the offline reconstruction developments into the HLT framework, thanks to the compatibility between the offline and online software frameworks.

7.3 Performance monitoring

The efficient operation of the ATLAS trigger system requires continuous monitoring of its performance, both in terms of physics selection and resource usage. The results of these performance studies are used to make decisions related to resource allocation and to the design of the *trigger menu* (i.e., the set of enabled L1 items and HLT chains, along with the respective prescale values¹ used to control their output rates). Some aspects of the HLT performance can be monitored online, using dedicated tools that provide analytic data on high-level quantities, such as trigger rates, efficiencies, CPU usage, and the average event processing time, in real-time during data-taking. However, this section focuses on offline performance monitoring methods, which enable more detailed and flexible studies. Specifically, Section 7.3.1 describes the data-driven method used to perform offline estimations of trigger rates and computational costs. Section 7.3.2 presents the studies of the HLT software performance scaling as a function of the number of events processed in parallel, considered the different ways of achieving parallelism enabled by the Run 3 upgrade to the AthenaMT framework.

7.3.1 Rates and cost analysis

The performance of the HLT in terms of output rate and CPU usage can be estimated prior to data-taking using a data-driven method based on *Enhanced-Bias* (EB) datasets. These EB datasets are designed to overweight events with high p_T and high object multiplicity, which are more likely to be selected by the HLT. This is necessary because such events are rare, therefore estimating the HLT performance from an unbiased dataset would require an unreasonably large number of events. The EB event selection is invertible, meaning that each event can be assigned a weight that corrects for the sampling bias applied during the EB data-taking, restoring an effective unbiased spectrum. This section describes how EB data are collected and used to estimate the rates of individual HLT trigger chains, or various combinations of chains, as well as to predict their CPU usage.

¹A trigger with prescale n , has a probability of $1/n$ to be activated in the event. Prescale factors can be applied to each L1 or HLT trigger, and can take any value ≥ 1 (triggers with prescale equal to one are referred to as *unprescaled*). More details on prescales and the ATLAS prescaling strategy are provided in Chapter 8.

Enhanced Bias dataset

The EB datasets are recorded using a dedicated set of triggers that run in parallel with standard physics data-taking. This ensures that they reflect the L1 trigger configuration and the LHC beam parameters of the LHC (especially the average pile-up) at the time of their collection. Therefore, a new EB dataset must be recorded whenever there is a significant change in either the L1 trigger setup or the beam parameters. The special set of L1 triggers used to collect EB data is designed to cover all the possible signatures and kinematic regions of interest (including high- p_T electron and muons, jets, one or more low- p_T muons for B -physics). These L1 triggers are grouped in sets of approximately 10-30 items, each set seeding a corresponding EB HLT chain. Each EB HLT chain targets a specific group of physics signatures with similar expected rates. The EB HLT chains are categorised in different groups based on the p_T range of their seeding L1 items: *high* and *very-high* chains are seeded by high- p_T L1-items, *medium* and *low* chains are seeded by items in the medium and the low p_T range, respectively; finally, *random* chains are triggered randomly. Since the collection of an EB dataset occurs in parallel with standard data-taking, it must employ the same set of L1 prescales being used to take physics data. In this setup, EB HLT chains seeded by prescaled L1 items, are triggered randomly. These chains maintain an internal record of their L1 seeds, and accept the event if at least one of these items passed the raw L1 decision (before any L1 prescales were applied) in the randomly selected event. EB chains do not apply any further selection criteria at the HLT, except for the application of HLT prescales to regulate their output rates. Therefore, the EB datasets are only biased by the L1 trigger decision. EB chains are typically enabled in parallel with standard physics triggers during a data-taking period of one hour. The total output rate of the EB chains is targeted at approximately 300 Hz, resulting in a dataset of about one million events per run. The EB weights that, as mentioned earlier, are used to invert the EB selection and restore the unbiased spectrum, are calculated for each event in the sample. For each event e , the EB weight $w_{\text{EB}}(e)$ is defined as:

$$\frac{1}{w_{\text{EB}}(e)} = 1 - \prod_{j=1}^{\text{EB chains}} \left(1 - \frac{r_{je}}{p_j} \right), \quad (7.2)$$

where the product runs over the j EB chains used to select the dataset, with raw decision r_{je} (the decision before the application of any prescale) and total prescale p_j , assuming that the prescales of the EB chains were constant over the data taking periods. Since October 2023, the calculation of the EB weights is performed on the LHC computing grid [167], significantly simplifying the operational procedure with respect to the past.

Before a new version of the Athena software is deployed on the HLT farm, it is validated by using it to process an EB dataset, in a procedure called *trigger reprocessing*. This is done to validate the updates to the current software release and trigger configuration before these are deployed for actual data-taking. Trigger reprocessings are typically performed on a weekly basis, and are executed on the LHC computed grid, given the large amount of events that must be processed. The HLT performance is evaluated in terms of reconstruction efficiency and rates of the trigger chains, as well as resource usage metrics, such as run time and memory consumption. This validation process is essential to maintaining high data quality and ensuring efficient and smooth data-taking operations.

Rates and cost analysis

To extract **rate predictions** from trigger reprocessings, all L1 items and HLT chains are run unprescaled. As a result, the output of the reprocessing includes the raw trigger decision for every L1 item and all HLT chain in the considered trigger menu for each event in the EB dataset. The effects of prescales can be applied after the reprocessing via weighting factors. The rates can be estimated for individual L1 items and HLT chains, as well as more complex combinations. These include the total rate of all HLT chains defined within different physics groups, the total L1 and HLT output rates, the unique rate of a single trigger chain or group of chains (i.e., the rate of events accepted exclusively by that chain or group), and the overlap rate between one chain and the rest of the trigger menu. Event weights are applied taking into account the L1 prescales used during the EB data-taking, as well as the luminosity and beam conditions under which the dataset was collected. The computed rate can be extrapolated to different operating conditions. The rate R of a chain, or a combination of chains, and its statistical error R_{err} are computed as:

$$R = \frac{\sum_{e=1}^N w(e)}{\Delta t}, \quad R_{\text{err}} = \frac{\sqrt{\sum_{e=1}^N w(e)}}{\Delta t}, \quad (7.3)$$

where R and R_{err} are expressed in Hz, the sum runs over all the $e = 1, 2, \dots, N$ events in the EB dataset, the weight $w(e)$ is the effective number of events accepted by the chain or combination of chains in the event e , and Δt is the time period over which the EB dataset was collected. The weight $w(e)$ is expressed as

$$w(e) = w_{\text{EB}}(e) \cdot w_{\text{C}}(e) \cdot w_{\text{L}}(e), \quad (7.4)$$

where, for each event e , $w_{\text{EB}}(e)$ is the EB weight as defined in Equation 7.2 ($w_{\text{EB}}(e) \geq 1$), $w_{\text{L}}(e)$ is a luminosity extrapolation weight, and $w_{\text{C}}(e)$ is the weight of the chain, or combination of chains. The luminosity extrapolation weight ($w_{\text{L}}(e) > 0$) is applied to use an EB dataset taken at instantaneous luminosity $\mathcal{L}_{\text{EB}}$ to predict the rates for a target luminosity $\mathcal{L}_{\text{T}}$. When the rate is predicted for individual HLT chains or L1 items, w_{L} does not depend on the event properties, while for combinations an averaging procedure is required. The chain/combination weight is $w_{\text{C}}(e) = 0$ if the chain/combination does not pass the trigger selection, or $0 < w_{\text{C}}(e) \leq 1$ if the chain/combination does pass, based on the prescale value(s) being applied to the chain(s). When rates are predicted for individual chains or items, $w_{\text{C}}(e)$ is simply calculated as the ratio between the raw rate and the prescale value, while for combinations more complicated averaging procedures are applied. More details on the weighting algorithms are discussed in Reference [168].

The results of the rate predictions based on a reprocessing output are stored in a JSON file, that is used as an input for computation of the prescale values to be used during data-taking, as discussed in detail in Section 8.2 of the next chapter.

Figure 7.5 shows the HLT rate predictions compared to actual online rates for 957 physics chains, validating the rates' prediction procedure.

A procedure similar to that used for trigger rate predictions is employed to perform HLT **cost estimation** in terms of number of HLT processor cores required to run a given trigger chain, or the full menu, under specific conditions. Unlike rate predictions, where all triggers are executed unprescaled and event weights are applied offline to correct for it, cost estimations require the L1 and HLT prescales to be applied during execution of the HLT over the EB dataset in the




Figure 7.5 – Pull distribution of the rate prediction of 957 HLT chains based on an EB dataset, normalised to the combined statistical error (σ), and fitted by a Gaussian function with mean μ_{Fit} , and width σ_{Fit} , where $R_{\text{Prediction}}$ and R_{Online} refer to the rates from prediction and data-taking for the same luminosity, respectively. Error bars include statistical uncertainties only [169].

trigger reprocessing. This is necessary due to the algorithm caching mechanism in the ATLAS HLT, which introduces correlations between HLT chains, such that prescales cannot be simulated later via the application of weights. During reprocessings for cost studies, monitoring data are recorded for each executed chain, including detailed timing information. When analysing this output, event weights from the EB dataset, defined in Equation 7.2, are applied, as well as an additional weight to scale the predictions in case the prescales applied during the reprocessing are not designed for the same luminosity as the EB dataset. The number of processor cores required to execute a given HLT chain, or full menu, is obtained by normalising the total CPU usage to the EB dataset collection time. By appropriately rescaling for different luminosity scenarios, this method can be used to project the processor cores' requirements for future data-taking conditions. As prescales are applied during the HLT execution in the reprocessings, only a fraction of the EB events contributes to the CPU usage measurement, reducing the statistical power of the dataset. As a result, for prescaled chains the statistical error on the CPU utilisation is larger compared to the statistical error on their rate predictions.

7.3.2 Performance scaling

Due to the parallel structure of Athena, a fundamental insight for evaluating the performance of the HLT online applications is the study of their *performance scaling* as a function of the events processed in parallel. After the transition from AthenaMP to AthenaMT, in particular, it is useful to compare how the performance scales when using different ways for achieving parallelism. This is done by running the trigger selection algorithms in a standalone local environment on a machine identical to those used in the ATLAS HLT computing farm during data-taking. Specifically, the tests are run on a dual processor machine with 128 GB RAM using a NUMA memory architecture and two AMD EPYC 7302 CPUs, each CPU with 16 physical cores with two hyper-threads per core, giving a total number of 64 threads. The trigger configuration used in the test is identical to the one used during data-taking, and the input data sample consists of events representative of the real HLT input data. The performance is measured as a function of the

number of events processed in parallel and is evaluated in terms of application throughput in events/s, CPU usage (CPU time divided by wall time) in percent, and memory consumption in GB. The following ways of achieving the parallelism are considered:

- **Pure multiprocessing approach.** After initialisation, the main process is forked into a number of worker processes equal to the number of events requested to process in parallel. Each worker uses a single thread to process events independently, sharing read-only memory with other workers (as already defined in Section 7.2).
- **Pure multithreading approach.** A single process uses a number of threads equal to the number of events requested to process in parallel. All threads share both read-only and writable memory.
- **Hybrid approach.** A number of worker processes are forked, each using multiple threads. The total number of events requested to be processed in parallel is distributed evenly across processes and threads, i.e. $N_{\text{processes}} = N_{\text{events}} / N_{\text{threads}}$.

These performance scaling measurements are used to determine which software configuration will be used during data-taking. Figure 7.6 shows the trigger software performance scaling in terms of application throughput and memory consumption as a function of the number of events processed in parallel using 2022 and 2024 data, processed using the Athena release deployed at the time when those data were taken [170, 171]. The results demonstrate that significant memory savings can be achieved even with hybrid configurations using a small number of threads per process. However, in 2022, a pure multiprocessing configuration with 48 forks was used during data-taking. As ATLAS transition to multithreading was new for Run 3, at the time several components still made use of mutex-based locking to ensure thread safety when accessing shared resources, limiting the event throughput for multithreaded configurations. Thanks to continue software developments, also aiming at minimising this kind of bottlenecks, since the 2024 data-taking, the ATLAS HLT uses a hybrid configuration with 16 forks and 4 threads, balancing between memory efficiency and processing throughput.




Figure 7.6 – Application throughput in events/s (left column) and memory usage in GB (right column) estimated using 2022 data (upper row) [4] and 2024 data (bottom row) [172] as a function of the number of events processed in parallel. Blue squares represent a pure multiprocessing approach, pink circles represent a pure multithreading approach, while the other curves represent hybrid approaches.

7.4 Operations

The trigger system is responsible for reducing the LHC collision rate of 40 MHz to approximately 3 kHz of events selected for permanent storage, performing this selection in real-time while ensuring high efficiency across a broad range of physics signatures. As only the events accepted by the trigger are recorded, any malfunction or inefficiency in the trigger system results in irretrievable data loss. Therefore, the reliable and continuous operation of the trigger system is critical to the overall success of the ATLAS data-taking and physics programme.

The trigger system operates in coordination with all other ATLAS detector subsystems. The data-taking process is organised into *runs*, which represent periods of continuous data acquisition under stable detector conditions. During standard physics data-taking, runs typically occur in parallel with LHC filling cycles (see Section 2.1.2), which can last several hours. However, ATLAS can also perform runs outside the LHC beam time, for example to collect data for detector performance or background studies (e.g. a cosmic data-taking run). Each ATLAS run is assigned a unique identifier by the DAQ system at its start. Within a given run, each recorded event is assigned an event number, unique to that run and incremented from zero for each run. Runs are further segmented into *luminosity blocks* (LBs), typically lasting about a minute, corresponding to a period of data-taking with stable experimental conditions (i.e. constant instantaneous luminosity, without any alteration in data-recording configuration, including the trigger system and its configuration). To start a run, all software and hardware components of the ATLAS detector, including the trigger, must follow the required transitions of a shared *finite-state machine*, that ensures a synchronized and well-defined configuration state across the entire system. The transitions are the following:

- during **initialise**, all applications are being started;
- in the **configure** and **connect** transitions, the hardware and applications are configured and connections between different applications are established where necessary;
- during **start**, a run number is assigned, and the applications perform their final (run-dependent) configuration.

Physics runs in ATLAS typically begin during the injection phase of an LHC fill. In this phase, the ATLAS sub-detectors operate in *safe-mode*, to protect sensitive equipment and reduce power consumption (e.g. high-voltage systems are turned-off or run at reduced levels to mitigate the risk of damage in case of beam instabilities). In this initial phase, the trigger is configured to enable specific triggers aimed at detector calibration and monitoring. Once the LHC reaches the stable beams phase, the sub-detectors exit the safe mode, and the trigger configuration changes to one designed to record physics data. The trigger configuration is updated regularly throughout the stable-beams period, as prescale values are adjusted to maximise the resource usage in response to the LHC luminosity decline during the fill. The ATLAS prescaling strategy is discussed in detail in Section 8.2 of the following chapter. Another common reason for changing the trigger configuration during stable beams is the execution of an emittance scan. Emittance scans are short beam-separation scans performed to estimate the luminosity as a function of beam separation and are used for performance validation and calibration. They typically take place a few minutes after stable beams are declared and last only a few minutes. The dedicated

trigger configuration is used during this period to ensure that all relevant calibration data are collected. An ATLAS run usually ends after the beam dump. In the interval between the beam dump and the official end of the run, the detector, including the trigger system, returns to a standby configuration. Transitions between the standby, physics, and emittance-scan trigger configurations are handled automatically, while changes between different physics configurations during a run must be performed manually by the control-room shifter.

The next paragraph details how these trigger configurations are stored and accessed during data-taking.

7.4.1 Trigger configuration

The configuration of the trigger system is defined by a coherent set of parameters for the L1 and the HLT systems. The L1 configuration includes the list of L1 items defined in the CTP, which correspond to logical combinations of the requested calorimeter and muon multiplicities. It also defines the trigger thresholds for which the calorimeter and muon trigger hardware provide multiplicities to the CTP, the list of L1Topo algorithms with their associated parameters used in the L1Topo system, and the firmware logic implemented in the CTP. The HLT configuration includes the list of HLT chains, along with information on their configuration (e.g. the output stream in which events are recorded, see Section 8.1.2), the selection steps they execute, and their L1 seeds and thresholds. It also includes the definition of the streams themselves and the algorithms executed at each step. The specific configuration parameters of the trigger algorithms and other software components, associated to a specific Athena version, are also part of the HLT configuration. Finally, both the L1 and HLT configurations are complemented by the information on the prescale values assigned to each L1 item and HLT chain.

The trigger configuration is first developed using the offline software framework, and then exported to files in JSON format. These files are uploaded to a relational database [173], the **TriggerDB**, which stores all the configurations in a consistent, stable and accessible way. Every trigger configuration used for data-taking is saved in the TriggerDB, and subsequently used to configure the HLT during data-taking. Trigger configurations are never deleted from the TriggerDB, ensuring that all the information on the trigger parameters used during data-taking is preserved, allowing the interpretation of the trigger behaviour in offline analyses. In the TriggerDB, the different elements of the configuration are stored in separate tables. Each element is assigned a unique identifier and is linked to related elements through a hierarchical structure. The **TriggerToolWeb** (TTWeb) is a web-based tool developed to access the TriggerDB. It allows to browse and manipulate the configuration data. Its primary role is to support fast, and flexible manipulation of the trigger configuration in response to changing beam or detector conditions during data-taking.

Complete configurations are identified by the database *keys*, assigned to the top level tables which, via relational dependencies, point to the configuration of every element of the L1 and HLT system. These keys provide the unique reference to the trigger configuration for each run and can therefore be used to retrieve the configuration at a later stage. The primary key of the top-level table is the **Super Master Key** (SMK), which uniquely defines the L1 and HLT configuration. Important additional elements of the configuration are the prescale values, assigned to each L1 item and HLT chain, defined by the **Prescale Set Keys** (PSK). Separate keys exist for




Figure 7.7 – Example bunch group configurations for the first four out of the 16 possible bunch groups. The numbers in blue on the right indicate the number of bunch crossings for each group [52].

L1 and HLT, and multiple prescale sets can be associated with a single SMK. The other key that together with the SMK and the PSKs uniquely defines the trigger configuration for data-taking is the **Bunch Group Key (BGK)**, associated with the *bunch group set*, reflecting the LHC filling scheme (see Section 2.1.1).

A bunch group set is the group of 16 bunch groups that can be defined in ATLAS, each for a specific purpose. A bunch group, in turn, is a list of bunch crossings. A trigger can only accept events occurring in bunch crossings that belong to the bunch group associated with it. Bunch crossings can be *paired* (or *filled*), if they have two colliding bunches, *unpaired*, if they only have one bunch, or *empty*, if they have no bunches. A dedicated bunch group set is generated for each LHC filling scheme in advance of data-taking. This process uses information regarding the expected positions of proton bunches in each beam provided by the LHC. The resulting bunch group sets are encoded in JSON files and uploaded to the TriggerDB, where they are associated with a unique BGK. Figure 7.7 shows an example bunch group configurations for the first four bunch groups: paired, empty, CalReq and BCRVeto. The paired bunch group contains the list of paired bunch crossings, and the empty bunch group contains the list of the empty ones. The calibration requests group (*CalReq*) defines the times at which sub-detectors may request calibration triggers, typically in a long gap with no collisions, and the group of the bunch counter reset veto (*BCRVeto*) leaves a short time slice for the distribution of the LHC bunch count reset signal to the on-detector electronics. All triggers are ANDed with the BCRVeto bunch group.

The trigger configuration can be modified during data-taking via the **TriggerPanel**, a tool accessible through the DAQ graphical user interface. Currently, the control-room shifters update the configuration by setting the appropriate keys. The SMK cannot be changed after the start of the run, contrarily to PSKs and BGKs. The BGK, which defines the bunch group set reflecting the LHC filling scheme, is assigned after the end of the beam injection phase, since the filling scheme becomes fixed at that point. The PSKs are updated more frequently, often several times during stable beams, as part of the dynamic prescale strategy used to adapt to changes in instantaneous luminosity, as discussed earlier. During the 2025 data-taking period, a procedural change will be introduced: shifters will configure the trigger system using configuration aliases instead of manually setting individual keys. A configuration alias is a database entry in the TriggerDB that maps to a predefined and approved combination of SMK, L1 PSK, and HLT PSK. This mechanism, already adopted during Run 2, improves operational robustness by reducing the risk of human error in selecting incompatible or unintended combinations of keys.

The ATLAS trigger system is responsible for selecting a small subset of collision events for permanent storage, filtering out the vast majority of events in real time. Central to this process is the trigger menu, a configurable set of selection criteria that define which events are recorded based on the presence of specific physics signatures. This chapter provides an overview of the structure and design principles of the ATLAS trigger menu, also discussing how the trigger menu adapts to changing LHC conditions during data-taking periods. A particular focus is given to the *prescaling strategy* in Section 8.2, that needs to be developed balancing carefully the physics outcome and the total manageable rate during data-taking, including a detailed description of the **Trigger Menu Rulebook**, the code used to generate the prescale sets used during data-taking.

8.1 Menu design

The ATLAS Run 3 trigger menu aims to maximise the physics impact of the Run 3 dataset by exploiting the improved capabilities of the upgraded detector, the more performant HLT hardware and algorithmic advancements, while simultaneously maintaining a level of consistency with the Run 2 trigger menu to allow for combined analyses on both datasets. The menu contains 512 L1 items (maximum number configurable in the CTP, see Section 2.2.6) and more than 2000 HLT chains, and ensures that the event selection rate stays within the system's limitations, including the maximum L1 accept rate, the HLT output bandwidth, and the availability of the HLT computing farm. Not all triggers are used for physics data-taking, and the trigger menu has to allocate the available resources based on the different requirements of the ATLAS collaboration. The menu is designed to give a large portion of the output bandwidth to general-purpose triggers that can be used by different analysis groups (e.g. the single lepton triggers used in the analysis presented previously). Nevertheless, triggers targeting specific analyses (e.g. long-lived particles or Higgs boson pairs searches) are also included in the menu, particularly exploiting the additional bandwidth available in Run 3 thanks to the TDAQ system upgrades, introduced in Section 2.2.6 and in the previous chapter.

8.1.1 System limitations

The primary challenge in the design and operation of the trigger system is to respect the limitations on the L1 rate and the HLT output bandwidth and processing capacity. These constraints

are managed through the trigger menu, which is tailored to the specific running conditions of the LHC and the operational status of the detector.

The **L1 trigger rate** limit is nominally 100 kHz, due to the maximum read-out capability of the detector (see Section 2.2.6). During actual data-taking, the maximum L1 rate is typically limited to about 94–95 kHz to prevent excessive data loss from dead time. In practice, this constraint can be further reduced by specific detector subsystems requirements under specific running conditions. An example is the IBL sub-detector (see Section 2.2.2), which can have a tighter read-out rate limit depending on the number of colliding bunches and the bunch structure [174]. This is due to the wire bonds connecting the front-end chips and sensors to circuit boards, which are subject to failure if time varying currents pass through them with frequencies close to their mechanical resonance, which may happen in case of too high trigger rates that lead the trigger to fire at every bunch crossing with colliding bunches. A mechanism of protection, called *fixed frequency trigger veto* (FFTV), is in place to stop data recording when the rate exceeds the limit [174]. While the IBL limit on the L1 rate typically exceeds 100 kHz during standard physics data-taking, at lower numbers of colliding bunches it can reduce the maximum allowed L1 rate to only a few kHz, making the design of the menu particularly challenging.

The **HLT bandwidth** is limited by the storage and data reprocessing capacities of the ATLAS DAQ system. For Run 3, the maximum output bandwidth is approximately 8 GB/s. It depends on both the event rate and the size of the events and, for physics pp data-taking, it translates to an effective rate limit of around 3 kHz for physics triggers requiring full event building (with typical event sizes of 1.5 MB). Some triggers, such as those used in trigger-level-analysis (see Section 8.1.2) and several detector calibration chains, do not require the full event building. These make use of *Partial Event Building* (PEB), which reads out only specific regions of interest within the detector, producing smaller events that reduce the load on the read-out and storage systems. As a result, PEB chains can run at higher rates, and the total HLT output rate, including both full and partial events, can reach approximately 13 kHz during standard physics runs.

The **HLT computing farm** is responsible for processing all events accepted by the L1 trigger, reducing the rate from the 100 kHz to the allowed output rate, while executing increasingly complex reconstruction algorithms with an average processing time per event of approximately 600 ms. The largest percentage of processing time is consumed by ID tracking algorithms. In 2022, tracking accounted for 59% of the total HLT processing time, followed by muon reconstruction (14%), and calorimeter reconstruction (11%) [4]. Processing time increases with pile-up, due to higher event multiplicities. This impacts the overall throughput of the system and constrains the number and types of triggers that can be included in the menu. Figure 8.1 shows the effect of increasing pile-up on the HLT system's performance.

8.1.2 Streaming model

To facilitate subsequent processing and analysis, accepted events are recorded into different datasets, known as *streams*. Each trigger chain defined in the menu is associated with one or more output streams. As a result, when a chain accepts an event, it is written to the corresponding stream(s) as specified in the menu configuration. The types of streams used during standard




Figure 8.1 – Results of the high-rate and high- μ test in October 2024. During the test, the pile-up value in ATLAS was increased from the 2024 nominal value of 64 to 69, in 1-unit steps every 10-30 minutes. During each step, the L1 trigger rate was varied from about 94 kHz to 100 kHz by increasing the low- p_T jets trigger rate. The higher number of collisions happening when increasing the pile-up value causes a heavier CPU load for the HLT reconstruction, observed as a decrease in the number of idle logical cores. The HLT output bandwidth increases as a function of μ , scaling both with the rate of accepted events in all streams, and with the increased event size due to the higher occupancy. The dead time increases with the L1 rate from about 4.5% to 8%, mainly due to the complex dead time (see Section 2.2.6) [5].

physics pp data-taking are listed below.

- **Physics streams** collect collision events relevant for physics analyses. These events contain full detector information, and dominate resource usage in terms of processing time, bandwidth and storage requirements. Three physics streams are defined for physics pp data-taking: the **Main stream**, which includes triggers for general physics analyses, the **B-physics and light states (BLS) stream**, containing triggers targeting specific to B -physics signatures, and the **Hadronic stream**, containing specialised hadronic triggers including selections for Vector Boson Fusion and Higgs boson pairs production processes. Events in the Main stream are promptly reconstructed after the first-pass calibration and data quality checks, while events in the BLS and Hadronic streams, also referred to as *delayed streams*, are processed later when resources become available.
- The **Express stream** records a small subset of events (approximately 50 Hz), selected from physics triggers, and reconstructed offline in near real-time, for prompt monitoring, detector calibration, and first-pass data quality assessment. Typically, events in the express stream are fully reconstructed offline within a day after having been recorded.
- **Background streams** include background events of interest for physics and detector performance studies.
- **Debug streams** are used to collect events for which a trigger decision could not be made due to processing problems, such as timeouts, or HLT data payloads exceeding the configured thresholds. These events need to be analysed and recovered separately, to identify and fix possible problems in the TDAQ system.

- **Calibration streams** contain events selected for specific sub-detector calibration purposes. These often use PEB techniques, significantly reducing the size of the recorded events.
- **Trigger-Level Analysis (TLA) streams** only store specific physics objects reconstructed by the HLT, with limited event information, to be used directly in the corresponding physics analyses. In 2022, the average TLA event size was about 4.5 kB, increased to approximately 25 kB in 2023 following the introduction of tracking and vertexing information.
- **Monitoring streams** are used for events that are sent to dedicated monitoring nodes for online analysis (e.g., detector monitoring), but are not saved to permanent storage.

With the exception of the Express stream, there are multiple streams of each type. Moreover, it is possible to define additional streams for special data-taking configurations, such as, for example, the EB stream, already mentioned in Section 7.3.1. Except for the debug streams, the streaming model is inclusive, meaning that a single event can be written to multiple streams if accepted by multiple chains. The streams are designed to have minimal overlap, and dedicated tools are available to assess and optimise the overlap between them. Figure 8.2 shows the evolution of HLT output rates and bandwidth as a function of time during an ATLAS run recorded in May 2024. The apparent discrepancy between the event rates and the corresponding bandwidth usage in some streams is due to the use of PEB techniques. For example, the TLA streams contribute significantly to the total HLT output rate, but their impact on the bandwidth, which is the actual limiting resource at the HLT level, is minimal, due to the smaller event sizes. The stream consuming the largest fraction of the output bandwidth is the Main stream, which collects events from generic triggers used in multiple physics analyses, in line with the ATLAS trigger menu design principles outlined earlier. While the total HLT output bandwidth corresponds to the sum of all stream output bandwidths, the total HLT output rate is lower than the sum of the individual stream rates, because some events are written to multiple streams. The periodic increase in both rate and bandwidth towards the end of the fill are due to adjustments in prescales, which are reduced as the luminosity and corresponding overall resource usage decline (see Section 8.2.1).

8.1.3 Menu code structure

The ATLAS trigger menu is implemented within the ATLAS HLT software framework in Athena, in a dedicated package called `TriggerMenuMT`.

This package contains the logic and definitions used to build trigger menus for both the L1 and HLT, for different operational purposes. Separate menus are defined depending on their intended use: some menus are used for data-taking, while others are designed for development, and testing and simulation purposes. The default physics *pp* menu is the `Physics_pp_run3_v1` menu. It defines all the triggers needed for physics analyses, along with support triggers for detector calibration, background estimation, and performance monitoring. Additional triggers needed exclusively for online data-taking, such as triggers for luminosity measurements, beam conditions monitoring, and detector-specific calibrations, are defined in a separate menu, the `P1_run3_v1`. To avoid duplication, each trigger chain is defined only once across all menus, and when a chain needs to be reused in multiple contexts, menus can be imported within one another. An example is the `PhysicsP1_pp_run3_v1` menu, used during standard *pp* data-taking, which combines both the `Physics_pp_run3_v1` and the `P1_run3_v1` menus via imports. Chains




Figure 8.2 – Output rate (left) and bandwidth (right) for the HLT streams as a function of time in a pp LHC fill taken in May 2024, with a peak instantaneous luminosity of $2.1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and a peak average number of interactions per crossing of $\langle \mu \rangle = 63$. The luminosity is levelled for a period at the beginning of the fill (see Section 2.1.3); after that, the output bandwidth decreases with the decaying luminosity [5].




Figure 8.3 – Import hierarchy for pp menus.

used in earlier data-taking campaigns, but no longer active, are archived in the `MC_pp_run3_v1` menu. This menu imports `Physics_pp_run3_v1`, and is used for the MC productions at ATLAS, as it contains all triggers ever used for data-taking during the current LHC Run. For the development of new triggers, the `Dev_pp_run3_v1` menu is employed, which itself imports the `MC_pp_run3_v1` and, consequently, the `Physics_pp_run3_v1`. Figure 8.3 illustrates the hierarchical structure of imports among the standard pp menus used in Run 3 and described above. Dedicated menus with analogous import hierarchies are defined for HI data-taking. Special menus are also available for specific data-taking conditions, including Cosmics data-taking (`Cosmic_run3_v1`), and proton-proton collisions with low pile-up (`PhysicsP1_pp_lowMu_run3_v1`), both also importing the operational `P1_run3_v1` menu. In addition, dedicated menus are being developed and tested for the upcoming Run 4 trigger configurations.

8.1.4 Menu for pp physics

The trigger menu consists of **physics triggers** and **auxiliary triggers**. These two main groups can be further divided in different subcategories.

Physics triggers can be *primary*, *support*, *alternative*, or *backup* triggers. Primary triggers cover all the signatures relevant to physics analyses and are allocated the largest fraction of the HLT output bandwidth. More details on the Run 3 pp primary triggers are discussed in the next

paragraph. Support triggers are instead used for efficiency and performance measurements, background estimates, or for monitoring. They typically run at a very low rate, of approximately 0.5-1 Hz each. Alternative triggers apply selections that are complementary to those of primary and support triggers, but use alternative reconstruction algorithms. They are often future primary triggers still in commissioning phase. Finally, backup triggers cover signatures similar to the ones of primary triggers, but with tighter selections and lower processing or output rate, so that they can replace the relevant primary triggers if their CPU usage or output rate becomes too high. These triggers select a subset of the events also selected by primary triggers, thus they do not require the use of additional computing resources or unique rate. Different prescaling strategies are applied to these different categories of physics triggers, as discussed in Section 8.2.1.

Auxiliary triggers can be classified into four main categories: *calibration*, *cosmic ray*, *beam-induced background*, and *noise* triggers. Calibration triggers are used for detector calibration and performance monitoring. They typically use PEB techniques, which allow them to run at high rates while minimising the bandwidth usage. Cosmic ray and beam-induced background triggers are used to study the *non-collision backgrounds*, i.e. detector signals not originating from pp collisions at the ATLAS IP. Specifically, cosmic ray triggers target the background resulting from the impact of extremely energetic cosmic muons on the detector. These triggers record data from empty bunch crossings during the LHC fills, or during dedicated cosmic runs in periods without beam. Beam-induced background triggers, on the other hand, are used to study the background from protons interacting inelastically with residual gas molecules near the detector, from protons with high transverse amplitudes, or from secondary particles produced when scattered protons hit the collimators. These triggers are typically active on unpaired or empty bunch crossings during the LHC fills. Noise triggers are seeded by a random L1 trigger, which selects events read-out from the detector chosen at random, independently of detector activity. These triggers are always prescaled and can be configured to sample either filled or unfilled bunch crossings. While random triggers on filled bunches can be used to mitigate potential L1 inefficiencies, those on unfilled bunches are primarily used for studying electronic noise and other detector-related backgrounds. Other dedicated auxiliary triggers are available for different scopes. One example is the *Zero Bias* trigger, which is designed to fire exactly one LHC turn (3564 bunch crossings) after a predefined electromagnetic L1 trigger (the so-called *zero-bias seed*) fires. This approach ensures that the recorded event is uncorrelated with the activity that caused the seed to fire, allowing to collect unbiased data that still reflect the instantaneous luminosity profile, which can not be achieved with random triggers. This trigger is used for dedicated detector and background studies.

Primary physics triggers

The *primary pp* menu triggers include all physics signatures relevant to the ATLAS physics programme, covering, as discussed earlier, a large range of physics goals, from SM precision measurements, such as studies of the decays of the Higgs, $W^\pm$ and Z bosons, to searches for BSM physics, including heavy resonances, supersymmetry or exotic particles. Triggers in the menu are organised into signatures, i.e., groups of related trigger chains targeting similar physics objects. The main signatures included in the physics pp menu are: electron and photons, muons,




Figure 8.4 – Example of rate given to each signature group in the Main stream in 2022, taken from Reference [4]. This chart is based on 2022 data recorded at an instantaneous luminosity of $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The label “UT” stands for *unconventional tracking signatures*, typical of, e.g., long-lived particles.

taus, jets and b -jets, E_T^{miss} , and BLS signatures. When possible, trigger thresholds at both L1 and HLT have been kept consistent with those used during Run 2, to facilitate combined analyses across data-taking periods, while benefiting from the Run 3 upgrade improvements to reduce trigger rate. The trigger menu strategy remains focused on allocating most of the system resources to inclusive triggers, which are sensitive to a wide range of processes. However, the additional bandwidth available in Run 3 compared to Run 2 has enabled the inclusion of more analysis-specific triggers, not covered by inclusive selections, the expansion of both the physics and TLA streams by lowering the trigger thresholds and including new triggers targeting previously unexplored phase space. Furthermore, the increased output capacity of the DAQ system enables a larger recording bandwidth, which is exploited to record more data in the delayed streams. Figure 8.4 shows a breakdown of the approximate HLT output rates by trigger signature. Combined chains are those applying requirements to multiple types of objects.

Data taken with electron and photon triggers are used in a wide range of physics analyses. The lowest unprescaled thresholds for single-electron and single-photon triggers are 26 GeV and 140 GeV, respectively. Lower-energy electrons and photons can be selected using multi-object triggers requiring higher multiplicity.

Muon triggers cover a broad momentum range, from few GeV, used for B -physics analyses, to several TeV, for BSM searches. The improved rejection of fake muons due to the addition of the NSW detectors (see Section 2.2.4) has allowed the isolated single-muon trigger threshold to be reduced from 26 GeV in Run 2 to 24 GeV in Run 3, without increasing the output rate.

Tau triggers target a wide spectrum of final states containing hadronically decaying τ leptons, with the lowest threshold for unprescaled single-tau triggers set to 160 GeV.

Jet triggers, including single-jet, di-jet and multi-jet selections, are used for multiple purposes,

ranging from precision physics measurements and BSM physics searches to detector performance studies and calibration (e.g. calibrations of the jet energy scale and resolution). The primary unprescaled single-jet trigger has a threshold of 420 GeV.

There are also triggers applying b -tagging algorithms to reconstructed jets, thus selecting events containing b -jets. The b -jet trigger menu includes both general triggers (single- or multi- b -jet triggers), and analysis-specific triggers, designed to target particular physics processes (e.g. a pair of Higgs bosons both decaying each into two b quarks, $HH \rightarrow b\bar{b}b\bar{b}$, which is one of the golden channel for Higgs boson pairs searches). These triggers combine different p_T thresholds and b -tagging working points, with multi- b -jet triggers requiring between one and four b -tagged jets.

Triggers targeting events containing E_T^{miss} are fundamental for BSM physics searches. The lowest threshold for triggers based on phase-I L1 seeds, used as of 2024, is 90 GeV.

A category of signatures introduced in Run 3 consists of unconventional-tracking signatures, targeting non-standard detector signatures from long-lived particles (e.g. displaced tracks or jets, or disappearing tracks). The introduction of these signatures was possible thanks to *large-radius tracking*, which uses the same algorithms of standard ID tracking, with modified configuration that allows to reconstruct tracks with large impact parameter d_0 (see Section 3.1.1).

The Run 3 TLA stream was expanded to include photon, muon, and b -jet triggers, as well as combined triggers, in addition to the jet triggers, that were used in Run 2.

8.2 The Rulebook

Prescale factors (or *prescales*) are applied to triggers to control the rate of accepted events and to manage CPU consumption at the HLT. As introduced earlier, a trigger chain with a prescale value of n has a probability of $1/n$ to be activated in the event. Individual prescale factors can be given to L1 items and HLT chains, and can have a value greater than or equal to one. The value -1 is used to disable triggers.

The *Trigger Menu Rulebook* (referred to as *Rulebook* in the following) is the software package used to generate prescale sets to be used during data-taking. It is not included in the Athena software, as it is subject to too frequent updates during data-taking.

The next paragraphs describe the prescaling strategy adopted in ATLAS and how prescale values are computed using the Rulebook based on the specific data-taking conditions.

8.2.1 Prescaling strategy

If a trigger is given a prescale value of n , a fraction $1/n$ of events that would have passed the trigger selection are actually recorded. Triggers with a prescale of exactly 1 are often referred to as *unprescaled*, and triggers with a prescale $n \geq 1$ are generally referred to as *enabled*, in contrast to *disabled* triggers, which have a prescale of -1 .

The logic used to apply prescales differs between the L1 and HLT systems.

At L1, prescales are applied after the trigger selection decision has been made by the CTP. For each L1 item, the CTP provides three types of trigger decision: the **Trigger Before Prescale** (TBP) indicates whether the item passed the selection criteria, regardless of any prescale; the **Trigger After Prescale** (TAP) indicates whether the item passed the selection and the prescale

condition; finally, the **Trigger After Veto** (TAV), corresponds to the TAP decision after applying the dead time veto (see Section 2.2.6).

On the other hand, HLT prescales are applied prior to executing the trigger algorithms. This approach ensures that computational resources are not spent processing events that will not be recorded, thus optimising the use of the HLT farm. However, as a consequence, it is not possible to determine whether a given event would have passed the HLT selection if it had not been prescaled and there is no direct equivalent of the TBP flag at HLT level.

By default, when a prescale is applied, the fraction of events that is recorded out of the total number of events that would have pass the unprescaled trigger is selected randomly for each individual trigger. However, a mechanism known as **Coherent Prescale Sets** (CPS) is implemented for defining groups of HLT chains with a common L1 seed whose prescales are applied in a coordinated way. This mechanism can be applied to group of chains that do not need to sample independent subset of events (as is generally the case). Let us consider a group of N HLT chains sharing a common L1 seed, with prescale factors $n_1 < n_2 < \dots < n_N$. If the prescales were applied without the mechanism of CPS, each chain would independently process $1/n_i$ of the L1-accepted events, for $i = 1, \dots, N$, with a probability that all chains activate simultaneously of $\prod_{i=1}^N (1/n_i)$. Consequently, the HLT reconstruction algorithms would run, on average, on $\sum_{i=1}^N (1/n_i) - \prod_{i=1}^N (1/n_i)$ of the total number of events that would pass the triggers in the absence of prescales. With CPS instead, the fraction of events selected by the triggers with prescale n_i for $i > 1$ is forced to be a subset of the $1/n_1$ events selected by the trigger with the lowest prescale. As a result, the reconstruction algorithms run only on a fraction of $1/n_1$ of the L1-accepted events, significantly reducing computational load.

Different prescaling strategies are applied to triggers based on their main purpose within the trigger menu. The *end-of-fill* strategy involves reducing prescale values or enabling certain triggers towards the end of the fill. Optimising the end-of-fill strategy is fundamental to make maximum use of available resources. Primary triggers generally run unprescaled. An exception are low- p_T TLA or BLS triggers, which are constrained by the rate limitations of their low- p_T L1 seeds and are therefore only enabled at the end of the fills, once the luminosity is significantly below its peak value. The end of fill strategy is particularly relevant to get the optimal physics outcome from these type of triggers. Triggers dedicated to background studies are also enabled, or have their prescales reduced, towards the end of the fill. Support and monitoring triggers maintain a fixed output rate throughout the fill, with their prescale values adjusted regularly to ensure stable rates. Prescale changes are clearly illustrated in Figure 8.5, which shows the L1 trigger rate as a function of time during an ATLAS run recorded in May 2024. The periodic increases in the L1 rate relative to the luminosity at the end of the fill reflect the prescale adjustments, which allow certain triggers to run at higher rates as resources become available.

8.2.2 Prescale Rules

The Rulebook processes a set of *prescale rules* and compiles them all into a consistent configuration of prescale values for both L1 and HLT, which are written in respective JSON prescale files (see Section 7.4.1). Each L1 item and HLT chain in the menu is assigned specific prescale rules that vary according to target luminosity since, as discussed previously, the prescaling strategy is adapted to the instantaneous luminosity delivered by the LHC in order to optimise the re-




Figure 8.5 – Total L1 trigger rate as a function of time in a pp LHC fill taken in May 2024, with a peak luminosity of $L = 2.1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and a peak average pile-up of $\mu = 63$, with the LHC instantaneous luminosity overlaid. The luminosity is levelled for a period at the beginning of the fill (see β^* levelling in Section 2.1.3); after that, the rate decreases with the decaying luminosity. Periodic increases in the L1 rate with respect to luminosity at the end of the fill are caused by prescale changes, as the luminosity and corresponding overall resource usage decline [5].

source usage towards the end of the fill. Prescale rules define the conditions and luminosity ranges under which a trigger is enabled, and they provide the method for computing the PS value. This can either be done by specifying the prescale value directly or by assigning a target rate. In the latter case, for HLT chains, the rule may specify either the desired rate of the HLT chain itself or the rate of its L1 seed. The target rate provided in the rule serves as the basis for calculating the final prescale value, as detailed in Section 8.2.3. When a rule specifies a target rate rather than a direct prescale value, the Rulebook requires the following inputs for the PS calculation:

- predicted online rates, taken from the rates JSON file produced during the weekly trigger reprocessings (see Section 7.3.1);
- the target luminosity, indicating the luminosity range for which the prescale set is intended;
- the expected LHC filling scheme, ensuring the prescale sets produced for online data-taking are tied to the appropriate bunch group configuration.

A distinct prescale set is produced for each considered luminosity range.

Multiple sets of prescale rules are defined to accommodate different data-taking conditions, and they can be categorised based on their purpose. A first broad distinction is between *physics rules* and *non-physics rules*. Physics rules depend on the target luminosity, and they are used to generate different prescale sets corresponding to various target luminosities. There are different sets of physics rules to define the prescale sets for standard physics data-taking, both for proton-proton physics (**physics pp rules**) and heavy-ion runs (**physics HI rules**).

Non-physics rules can be further subdivided into several categories. Two sets of rules commonly used during standard physics data-taking are the **standby rules** and the **emittance rules**. Standby rules are used to produce the prescale sets for periods with no beam between two

stable-beam phases of the LHC filling cycle. These rules enable triggers primarily aimed at detector calibration and monitoring. Emittance rules are used during emittance scans, which are short beam-separation scans performed to estimate the luminosity as a function of beam separation and are used for performance validation and calibration. The corresponding rules activate calibration-specific triggers required during these procedures. Both standby and emittance rules are implemented in two different versions specific to pp and HI data-taking. Another important category consists of the **cosmic rules**, which are optimised for cosmic data-taking. They enable dedicated triggers for the study of the cosmic-induced background and are used in runs without beam. A special category of rules is constituted of the **monitoring rules**. These rules are a common set of rules containing specific detector monitoring triggers, and are appended to all standard physics, standby, emittance and cosmic rule sets. Separate versions of monitoring rules exist for pp and HI running. Additional specialised rule sets are used for testing purposes. One example is the set of **random rules**, which consist exclusively of random triggers. These are used to test trigger infrastructure or to read out the detector without applying selection criteria. Finally, there are the **express rules**, defined for the HLT chains associated with the express stream. These rules exist separately for pp and HI configurations. Any HLT chain that selects events for the express stream is associated with two distinct prescale values: a standard prescale, which controls the fraction of events written to the regular physics or monitoring streams, and an *express prescale*, which is typically higher and determines the subset of events also sent to the express stream.

The Rulebook also includes several dedicated sets of rules designed to support the various types of ATLAS *special runs*. Special runs are recorded to perform specific measurements or to exploit special beam conditions provided by the LHC, and the trigger configuration needs to be adapted to the specific goals of the runs. Generally, dedicated sets of prescale rules are employed to generate special prescale sets enabling the triggers useful for the specific purposes. Some examples are provided below. One example is the *start-up* period, which takes place during the initial weeks of data-taking following the LHC end-of-year shutdown. During this phase, the machine operates with a reduced number of bunches and significantly lower instantaneous luminosity than nominal. The number of bunches and luminosity are progressively increased up to the nominal run conditions. A dedicated set of start-up physics rules is used to support this phase, focusing on the commissioning and validation of new triggers, reconstruction algorithms, and other updates introduced during the end-of-year shutdown. The special bunch configurations reduce significantly the IBL limit, so priority must be given to the primary chains, while supporting chains are kept at a much lower rate.

Special runs for *van der Meer (vdM) scans* [175], which typically happen twice per year (one time during pp physics and one time during HI data-taking) record events to be used to perform the ATLAS absolute luminosity calibration. During the vdM scan runs, triggers are enabled to record events based on activity in the MBTS and LUCID sub-detectors (see Section 2.2.5), in a dedicated calibration stream using PEB.

A dedicated set of physics rules is also used when the LHC delivers collisions at low pile-up, typically with average interactions per bunch crossing below one. These low pile-up runs are used for different physics goals, for example the measurement of the W boson p_T , necessary to reduce the error on the measurement of the W mass [176]. The associated prescale sets prioritise primary triggers dedicated to these specific physics analyses.

The examples above illustrate only some of the special runs recorded during each year of data-taking by ATLAS. Many others exist (some more examples are provided, for example, in Reference [177]), each requiring customised prescale rule sets. The Rulebook infrastructure provides the necessary flexibility to accommodate these diverse configurations, ensuring that trigger resources are used efficiently to meet the specific objectives of each special data-taking period.

Each set of prescale rules includes several additional options that allow further refinement of the generated prescale sets to accommodate more specific data-taking conditions. These options allow enabling or disabling specific groups of triggers or adjust their prescale values, without the need to define entirely new rule sets. For example, during the commissioning of the Phase-I system, it was necessary to test both the new Phase-I and the existing Legacy triggers, therefore specific options within the rule sets allowed selective activation of either or both systems. Similarly, these options can be used to disable or reduce the rate of specific set of chains in case of detector issues.

8.2.3 Rules processing

The code that handles the processing of the prescale rules consists of three main parts: the **RuleReader**, which processes the rules together with the inputs describing the running conditions to compute the prescale values for each L1 item and HLT chain, the **OutputChecker**, that performs consistency checks on the RuleReader output before feeding it to the **RuleWriter**, which creates the output JSON files containing the prescales sets. The output JSON files are then uploaded to the TriggerDB to create the actual trigger keys, as described in Section 7.4.1. In addition to the prescales set JSON files, the Rulebook also outputs the Alias JSON file, uploaded to the TriggerDB to generate the Alias table.

The main script in the Rulebook is `runRuleBook.py`. It takes as an input the needed rules to be processed, a target luminosity, and the desired configuration flags. Configuration flags are of two different types: general *run options*, including the SMK and the BGK for which the prescale sets are made for and the path of the rates JSON file, and the *rulebook options*, specific of a particular set of rules. There are almost one hundred configurable options in the *pp* physics rules, which allow to adjust the PS set to specific conditions. It also contains the function to gather all the needed information to build the input JSON file for the Alias panel (see Section 7.4.1).

Figure 8.6 shows the importing structure of the Rulebook code. Before starting with the configuration and the actual computation of the prescale values, a check is run on the input rules to verify that there aren't triggers with more than one rule defined, since this would create ambiguities in the prescale calculation. In case duplicate rules are found, the Rulebook breaks without generating the output prescale sets.

After this check, the *configuration* for the prescales sets is defined, and contains the information described below.

- Path to the menu files, including the L1 Menu, the HLT menu and the Monitoring groups JSON files. These files are automatically downloaded from the TriggerDB starting from the input SMK.
- Set of rules to be processed.




Figure 8.6 – Diagram showing the importing structure of the Rulebook code.

- The target luminosity, which indicates the upper limit of the luminosity range for which the generated prescale sets are valid.
- Luminosity scaling factor, which has a value greater than one, and is used to scale down the input luminosity point, corresponding to the upper edge of luminosity interval for which the generated PS sets are valid, to an intermediate value within that range, which is then used to compute the prescales.
- Target number of bunches for each bunch group: this is obtained from the Bunch group JSON file, which is automatically downloaded from the TriggerDB together with the menu files starting from the input BGK. The expected bunch pattern is communicated by the LHC in advance of collisions. Generally, the same bunch pattern is kept for the full pp physics data-taking period. During commissioning periods or during the ramp-up period of the LHC, the bunch pattern can change frequently and different sets of prescale keys need to be generated for each different bunch configuration.
- Path to the rates JSON file coming from the reprocessing of the release associated to the reference SMK.
- Luminosity for which the rates predictions in the rates JSON file were made (usually $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$)
- Number of bunches in each bunch group considered when making rates predictions. At the time of writing, Run 3 reprocessing only provides predictions for the bunch group corresponding to paired bunch crossings, and rate predictions for triggers targeting other bunch groups are not available.

After defining the configuration, prescales sets are computed and written on JSON files using the functions contained in three classes, RuleReader, OutputChecker, and RuleWriter. The details on the calculation and the formatting of these prescales is described in detail in the next paragraphs.

RuleReader

The actual computation of the prescale values is performed by RuleReader. It consists of different steps, described in detail in this section.

First, all the triggers in the menu are assigned the respective rules based on the input target luminosity. As discussed before, each trigger is assigned a set of rules containing different entries, corresponding to different target luminosities. If the set of rules assigned to the trigger contains a rule exactly matching the input target luminosity, the assignment is trivial. Otherwise, the rule corresponding to the highest luminosity that is lower than the target luminosity is assigned to the trigger. If there are no rules corresponding to luminosities lower than the target luminosity, the rule corresponding to the lowest luminosity is assigned to the trigger.

After that, the consistency between the input menu and rules is verified, checking that there are no triggers in the menu without a respective rule and vice versa. In the latter case, when a rule is defined, but the corresponding trigger is not found in the menu, the Rulebook breaks without generating the output prescale sets. When, on the other hand, a trigger is found in the menu without a respective rule, there are different cases depending on the type of trigger: if the trigger without rule is an HLT chain, then it is disabled; if instead it is a L1 item, it is disabled only if all its HLT seeded chains are also without rules; otherwise, it remains without rule, as it will get a prescale during the optimisation phase.

After these first checks, the actual computation of the prescale values begins, following a similar procedure for L1 items and HLT chains.

When the rule directly defines a **prescale value**, then the PS is set trivially. If the rule also contains a *prescale slope* value (S_{PS}), the final prescale value is corrected by a scale factor depending on the prescale slope, the target luminosity (L_{target}), and the rule's luminosity (L_{rule}):

$$PS = PS \cdot S_{PS} \frac{L_{\text{target}}}{L_{\text{rule}}}, \quad (8.1)$$

where the target luminosity L_{target} is not corrected by the luminosity scale factor previously introduced. This procedure is primarily used for HLT prescale sets, because, due to the short duration of the data-taking periods, the rate predictions from the reprocessings are less reliable. With this option the prescale is determined based on the rate measured during data-taking at a given luminosity, and then it is modified as a function of the luminosity to keep the rate constant, assuming that the rate decreases linearly with luminosity over time.

To compute the prescale value when the rule defines a **target rate**, an exact estimation of the expected rate of the considered trigger (R_{exp}) is needed. The calculation of R_{exp} is based on the predicted rate from the input rates JSON file (R_{pred}). Triggers that are not listed in the rates JSON file with a rule defining a target rate are disabled. Other types of triggers for which the rates' prediction are not reliable are triggers targeting non-paired bunch groups and multi-seeded chains. These kinds of triggers have to be defined by prescale, otherwise they are disabled. An exceptional procedure is followed for two special cases: random L1 triggers and the L1 zero-bias trigger. In these two cases, the rates are not correctly predicted, so their value in the rates JSON file is not reliable.

The L1 random triggers' rates (R_{RD}) are computed manually in the Rulebook, assuming that the 3564 bunch crossings (see Section 2.1.1) are evenly populated by the 40 million bunches per

second, and multiplying this value with the expected number of bunches (N_{bunches}) in the target bunch group for the expected LHC filling scheme:

$$R_{\text{RD}} = \frac{40 \cdot 10^6}{3564} \cdot N_{\text{bunches}} \quad (8.2)$$

If the expected number of bunches in the target bunch group is zero, then the item is disabled. By definition, the L1 zero bias trigger has the same R_{pred} of its zero bias seed (see Section 8.1), and a dedicated procedure is implemented to retrieve the information on the zero bias seed from the L1 menu configuration JSON file.

The rates obtained at this point are corrected with a *scale factor* (SF_{rate}), used to extrapolate the rate predictions from the reprocessing to the target luminosity. This scale factor is computed as the ratio between the target luminosity and the luminosity used to predict the rates for the rates JSON file. For L1 random triggers, for which the rate was already computed manually starting from the bunch structure, the value of the scale factor is manually set to 1 ($R_{\text{exp}} = R_{\text{RD}}$). In all other cases:

$$R_{\text{exp}} = R_{\text{pred}} \cdot SF_{\text{rate}}. \quad (8.3)$$

Some additional considerations are needed to evaluate R_{exp} for HLT chains. First, for HLT chains, R_{exp} needs to be further scaled by the prescale value of the L1 seed (PS_{seed}). Therefore, prescale values are calculated for L1 items before, since L1 prescales are needed as input for the calculation of the prescales of the HLT chains. Moreover, for HLT chains the rule can set a target HLT rate or a target L1 seed rate. In the first case, R_{exp} is computed starting from its predicted rate scaled by SF_{rate} , while in the second case R_{exp} corresponds to the predicted rate of its L1 seed ($R_{\text{pred,seed}}$) corrected by the same scale factor SF_{rate} . Table 8.1 summarises the different cases discussed above. Once R_{exp} is computed, the prescale value for rules defining a target

Trigger type	Rule type	R_{exp} calculation
L1 item	target L1 rate	$R_{\text{pred}} \cdot SF_{\text{rate}}$
HLT chain	target HLT rate	$R_{\text{pred}} \cdot SF_{\text{rate}} \cdot PS_{\text{seed}}$
	target L1 seed rate	$R_{\text{pred,seed}} \cdot SF_{\text{rate}} \cdot PS_{\text{seed}}$

Table 8.1 – Summary of different ways of computing the expected rate value used in the prescale computation when the rule defines a target rate. The three listed cases are: expected rate for a L1 item, for which a rule defines a target L1 rate; expected rate for an HLT chain for which the rule defines a target HLT rate, and expected rate for an HLT chain for which the rule defines a target L1 seed rate.

rate (R_{target}) is defined as:

$$PS = \begin{cases} \frac{R_{\text{exp}}}{R_{\text{target}}} & \text{if } R_{\text{exp}} > R_{\text{target}} \\ 1 & \text{otherwise} \end{cases}. \quad (8.4)$$

If the computed prescale is greater than the maximum allowed PS value, $PS_{\text{max}} = 16777215$, the prescale value is set to $PS = PS_{\text{max}} - 1$.

Figure 8.7 schematically shows the PS calculation steps described above.

Generally, a single L1 item seeds multiple HLT chains. To optimise the use of resources, prescales are configured such that the L1 TAP rates are as low as possible, avoiding accepting events at the L1 that are not used at the HLT. For this reason, prescales are applied at L1 whenever possible. To implement this efficiently, an **optimisation** procedure is performed to reduce the rate at L1 while maintaining the same overall trigger efficiency at HLT. This is done by shifting prescales from HLT to L1 in a controlled way, ensuring that each L1 item seeds at least one unprescaled HLT chain. The procedure works as follows. For each L1 item that does not have an explicit prescale rule, the minimum prescale among those of the enabled seeded HLT chains ($PS_{HLT}^{\min}$) is used to define the new L1 prescale: $PS_{L1}^{\text{new}} = PS_{HLT}^{\min}$. Finally, all prescales of the seeded HLT chains are scaled by the same factor $PS_{HLT}^{\min}$ to preserve the original total rate and selection efficiency:

$$PS_{\text{tot}}^{\text{new}} = PS_{L1}^{\text{new}} \cdot PS_{HLT}^{\text{new}} = PS_{L1}^{\text{old}} \cdot PS_{HLT}^{\min} \cdot \frac{PS_{HLT}^{\text{old}}}{PS_{HLT}^{\min}} = PS_{L1}^{\text{old}} \cdot PS_{HLT}^{\text{old}} = PS_{\text{tot}}^{\text{old}} \quad (8.5)$$

This strategy reduces the L1 rate without affecting the final trigger selection. The optimisation procedure can be applied to L1 items with an already defined rule if explicitly requested (*re-optimisation*). In this case, $PS_{HLT}^{\min}$ is multiplied to the current L1 item prescale (PS_{L1}^{old}), as:

$$PS_{L1}^{\text{new}} = PS_{L1}^{\text{old}} \cdot PS_{HLT}^{\min} \quad (8.6)$$

In this way, the final L1 prescale can only be equal or less than the initial one. This procedure is used for items whose rate needs to be kept under control, such as random triggers.

The procedure to compute the **express prescales** is similar to the one implemented to compute regular prescales. Express prescales are only computed for HLT chains defined in the express rules. Express rules are generally defined by rate ($R_{\text{target,express}}$), except for triggers not targeting colliding bunches for which, as mentioned before, there are no reliable rate predictions to compute the rule by rate. In the latter cases, the express rule directly contains an express prescale value, that is trivially assigned as express prescale to the chain. Otherwise, the express prescale is defined as follows:

$$\begin{cases} PS_{\text{express}} = \frac{R_{\text{exp,reg}}}{PS_{\text{reg}} \cdot R_{\text{target,express}}} & \text{if } R_{\text{exp,reg}} > R_{\text{target,express}} \\ 1 & \text{otherwise} \end{cases}, \quad (8.7)$$

where $R_{\text{exp,reg}}$ is the expected rate of the corresponding HLT chain, computed as in Table 8.1, and PS_{reg} its regular prescale.

OutputChecker

OutputChecker runs over the RuleReader output and performs various consistency checks in sequence, to verify that the generated prescale sets do not contain obvious inconsistencies. If issues are found, the prescale sets are created, and it is responsibility of the Menu Expert to understand if the anomaly is expected or if it needs to be followed up. The performed checks are detailed below.

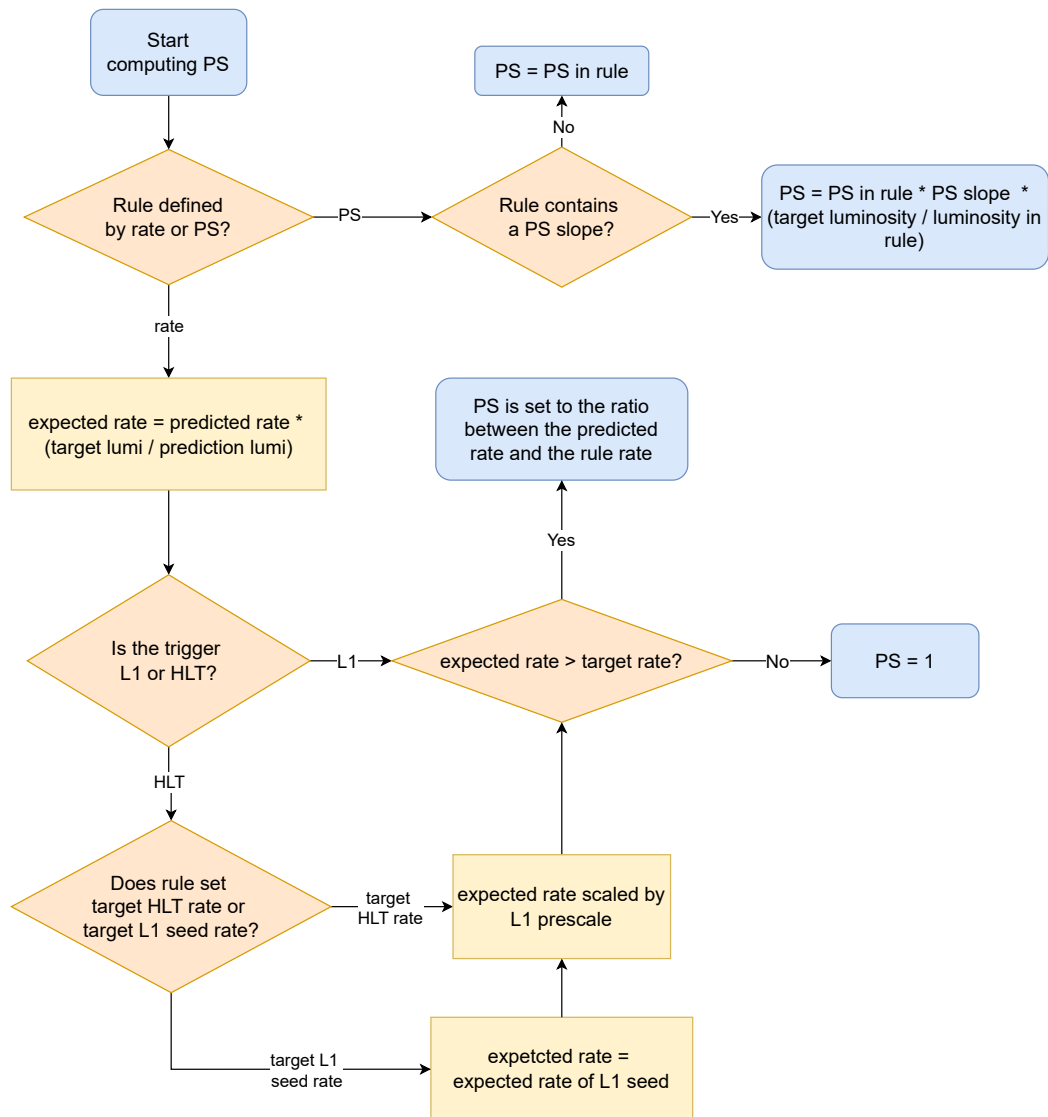


Figure 8.7 – Simplified flow chart of the PS calculation steps.

Rate request. This first check regards the HLT chains with a rule defining a target rate. The expected rate, represented by the quantity R_{exp} of Equation 8.3 scaled by the total prescale ($\text{PS}_{\text{L1}} \cdot \text{PS}_{\text{HLT}}$, where PS_{HLT} is the prescale value of the considered HLT chain and PS_{L1} is the prescale value assigned to its L1 seed) is compared with the target rate defined in the rule. The check fails if one of the following conditions is met:

- the L1 seed is disabled, resulting in negative expected rate;
- the difference between the target and the expected rate is bigger than 10%, which can happen, for example, when the second case of Equation 8.4 is verified.

In the second case, if both the HLT chain and its L1 seed are unprescaled, the inconsistency can only be solved by lowering the target rate request in the prescale rule.

Consistency of express rules and HLT menu. The second check regards the consistency between the express rules and the HLT menu. If an HLT chain has an express rule in the Rulebook, but is not defined as an express stream chain in the HLT menu JSON file, the check fails and Rulebook breaks without creating the output prescale sets. The check also fails in the opposite case, when an HLT chain is defined as an express chain in the menu but does not have an express rule in the Rulebook.

Unused L1 rate. The L1 prescales are checked to verify that there are no enabled L1 items that are not seeding any unprescaled HLT chain, and is only performed on physics prescale sets. This check is crucial to avoid having unused L1 rate. This eventuality is prevented by the L1 optimisation described in the previous paragraph for L1 items that do not have an associated rule. In other cases, the check can fail, and the amount of wasted L1 rate is also computed as

$$\left(1 - \frac{1}{\min(\text{PS}_{\text{HLT}})}\right) \cdot R_{\text{exp,L1}}, \quad (8.8)$$

where $R_{\text{exp,L1}}$ is the expected rate from Equation 8.3 of the L1 item scaled by its prescale value, and $\min(\text{PS}_{\text{HLT}})$ is the minimum prescale value assumed by its seeded HLT chains.

Allowed PS values. A very basic check is performed to verify that the prescales do not assume values that are not allowed, namely not $0 < \text{PS} < 1$ and not $\text{PS} > \text{PS}_{\text{max}}$.

PS vs expected rate consistency. A check on the consistency of the node prescale and expected rate is performed based on their **sign**. If the total prescale of the node, that is the L1 prescale for the L1 items and the product between the L1 seed's prescale and the HLT prescale for HLT chains, and the expected rate of the node have inconsistent signs (one positive and one negative), the check fails.

Triggers enabled with no rule. This check is performed both on L1 items and HLT chains. For L1 items, the check fails if they do not have HLT seeds, do not have a rule and are enabled. For HLT chains, the check fails if they have no rule but are enabled.




Figure 8.8 – Screenshot of the ATLAS TriggerPanel taken in June 2025. The names of the PS sets contain relevant information on their content.

Prescaled primaries. Only for physics prescale sets, a check is performed to verify that all primary chains are unprescaled. The check fails both if the primary chain itself is prescaled, or also if its L1 seed is prescaled. A list of primary chains allowed to be prescaled can be defined: this may be needed in commissioning phase, when chains are primary candidates but still need to be tested or tuned before being unprescaled.

Monitoring chains. A check is performed only for physics prescale sets to verify that all the chains needed for monitoring are enabled. The chains in the HLT menu file that are defined as monitoring chains in the monitoring groups JSON file are compared, and the check fails if any of these monitoring chains is disabled.

RuleWriter

RuleWriter writes on JSON files the output of RuleReader, after it is checked by OutputChecker.

First, it generates the *name* of the output prescale sets, starting from the information in the input configuration. The prescale set name contains information on the type of rules used to generate the prescales, and the target number of colliding bunches for which they were generated. For physics rules, which depend on the input luminosity, it also contains the target luminosity value. At the beginning of Run 3, the prescale set name was defined as the name of the input menu (with, therefore, only three options: physics pp , cosmic, and physics HI). The precise definition of the prescale set's name is an important feature reintroduced for Run 3 in 2023, because it is shown as key-label both in the TTWeb and in the TriggerPanel in the ATLAS Control Room, reducing the risk of using a wrong prescale set during data-taking. Figure 8.8 is a screenshot of the ATLAS TriggerPanel taken in June 2025, giving an example of meaningful PS set names: all the PSs are made for pp data-taking, and an LHC bunch pattern with 2448 colliding bunches; physics PSs are optimised for a target luminosity of $2.2 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$.

<pre> "item-name": { "cut": 5592406, "enabled": true, "info": "prescale: 1.5" }, </pre>	<pre> chain-name": { "hash": 2282772332, "prescale": 5000.0, "enabled": true, "prescale_express": 1.19, "enabled_express": true }, </pre>
(a)	(b)

Figure 8.9 – Example of an entry from a L1 (left) and an HLT (right) prescale JSON files. For the HLT file, the example comes from an entry of a chain also writing to the express stream. For chains not writing to the express stream, the last two lines are absent.

After that, it writes the computed prescales in two different JSON files, one containing the L1 items, and one for the HLT chains. The format of the two files is different and described below. The **L1 Prescales JSON file** shows the prescale values together with their *cut-off value*, which is the information needed by the CTP to make the L1 trigger decision. The L1 prescaling uses a pseudo-random binary sequence generator of 24 bit width (31 bit internal). Prescales are defined by the cut-off value (C). For each event passing the L1 trigger requirements, a random number R is generated. Both C and R take values between 0 and $2^{24} - 1 = 16\,777\,215$ (extremity values excluded). The trigger item is L1-accepted (TAP) if $R \geq C$. The PS corresponding to a cut-off C is:

$$PS = \frac{2^{24} - 1}{2^{24} - C}. \quad (8.9)$$

In this way, the cut-off value of an unprescaled trigger ($PS = 1$) is $C = 1$, which will always give $R \geq C$, meaning that the trigger will always be accepted.

In addition to the prescale value, the **HLT Prescales JSON file** also contains the information on the express prescale value for chains writing to the express stream. Figure 8.9 shows an example entry of the L1 and HLT prescales JSON files.

Conclusions and outlook

This thesis presented a search for the direct pair production of supersymmetric top squarks, based on the proton-proton collision data collected with the ATLAS detector at $\sqrt{s} = 13$ TeV and $\sqrt{s} = 13.6$ TeV, during the Run 2 and the Run 3 data-taking periods of the LHC, respectively, amounting to a total integrated luminosity of 193 fb^{-1} . The search targets simplified supersymmetric models where the lightest top squark decays into a top quark and the lightest neutralino with 100% branching ratio (two-body decay mode). Therefore, only signal scenarios with a mass splitting between the top squark and the neutralino larger than the top quark mass are considered in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ parameter space. The analysis focuses on the dileptonic decay of the top quark, selecting final states with exactly two oppositely charged leptons, jets, and missing transverse momentum. The sensitivity of the analysis was enhanced by the use of a novel machine learning technique developed to perform signal-to-background discrimination. With the full dataset, the analysis achieves sensitivity to top squark masses above 1.1 TeV for neutralino masses up to 600 GeV. This represents a significant improvement over the previous Run 2 search [111], both in the high stop mass region and for compressed scenarios, where the mass splitting between the top squark and the neutralino is close to the top quark mass. The sensitivity of the analysis was also tested in the three-body decay region, where the difference between $m_{\tilde{t}_1}$ and $m_{\tilde{\chi}_1^0}$ is smaller than the mass of the top quark, and the top squark decays directly in a W boson and a b quark, via off-shell top, and in the lightest neutralino. A good sensitivity is achieved also in this region of the parameter space, reaching neutralino masses of 450 GeV for top squark masses of 600 GeV. The analysis was also reinterpreted for an alternative signal model consisting in an invisibly decaying Higgs boson produced in association with a top-antitop quark pair, presenting the same final-state signature. An expected upper limit of $0.29^{+0.13}_{-0.09}$ is imposed on the Higgs boson invisible branching ratio when considering the full 193 fb^{-1} dataset, which equals the expected upper limit obtained combining all the $t\bar{t}H$ decay channels (dileptonic, semi-leptonic, and fully-hadronic) in the previous Run 2 analysis [157]. These results demonstrate that the new analysis strategy based on machine learning not only enhances the sensitivity to the targeted stop models, but also has the flexibility to adapt to different physics scenarios with overlapping final-state signatures.

The results of this top squark search will be combined with those of complementary analyses targeting the semileptonic and the fully-hadronic final states also considering the top squark two-body decay mode. These channels benefit from higher cross sections, due to the larger

branching ratio of the hadronic decay of the W boson ($\mathcal{B}_{W \rightarrow q\bar{q}} = 67.41 \pm 0.27\%$), and are more sensitive in the parameter space region with higher top squark masses. The dileptonic channel provides a cleaner final state with better background rejection, and employs a different trigger strategy based on lepton triggers, which enables the selection of lower- p_T leptons, improving the sensitivity to the challenging compressed region. Therefore, combining the results of these analyses is essential to fully exploit the available data and maximise the sensitivity across the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ parameter space. With Run 3 ongoing and over 206 fb^{-1} of data already collected at the time of writing, there is a unique opportunity to further push the sensitivity of top squark searches. This is being driven by novel analysis techniques, also enhanced by the use of progressively more advanced machine learning models, and improved reconstruction and identification techniques for key physics objects, especially b -tagging and missing transverse momentum reconstruction. Following the conclusion of Run 3, the High-Luminosity LHC (HL-LHC) phase will begin, aiming to deliver 3000 fb^{-1} of data at $\sqrt{s} = 14 \text{ TeV}$ by 2041. This unprecedented dataset, corresponding to 20 times the integrated luminosity achieved in Run 2, will allow for a 5σ discovery potential for top squark masses up to approximately 1.7 TeV, and a 95% CL exclusion reach extending to about 2.05 TeV under optimistic assumptions [178].

The upgraded ATLAS trigger system is successfully running, contributing to the high data-taking efficiency of the ATLAS experiment in Run 3. In 2024, the trigger system recorded physics data with a 99.7% efficiency [112]. The underlying framework of the HLT, completely rewritten in order to execute trigger algorithms within the multithreaded software framework AthenaMT, has optimised the use of the HLT computing resources, both in terms of computing power and memory consumption. The performance of the HLT software is continuously monitored, also thanks to the rates and cost analysis packages presented in this thesis, which allow to predict the trigger rates and the CPU cost of the HLT algorithms prior to data-taking. These monitoring tools provide predictions using a data-driven approach relying on Enhanced Bias datasets. The workflow used to process the data collected during the Enhanced Bias runs to correct for the sampling bias through the application of dedicated event weights was optimised as part of the work of this thesis. Thanks to this work, the Enhanced Bias weights production is executed on the LHC grid as of October 2023.

The ATLAS trigger menu defines the set of L1 and HLT selections applied during data-taking. While maintaining a level of consistency with the Run 2 trigger menu, the Run 3 trigger menu has been adapted to exploit the newly implemented detector features, more performant HLT hardware, and algorithmic advancements. The Run 2 trigger thresholds at L1 and HLT were generally preserved, benefiting from the improvements of the upgraded system to reduce the trigger rates. The additional available HLT bandwidth compared to Run 2 enabled the inclusion of new triggers targeting previously unexplored regions of phase space and triggers dedicated to analysis-specific signatures. Together with the L1 and HLT triggers used during data-taking, the trigger menu defines their prescale factors, applied to control the rate of accepted events and to manage CPU consumption at the HLT. The ATLAS trigger prescale strategy is optimised to make optimal use of the available resources and to adapt to the evolving data-taking conditions, particularly increasing the rate of specific triggers when more resources become available toward the end of an LHC fill, when the luminosity declines. In these conditions, additional bandwidth becomes available to unprescale memory-intensive triggers, such as those used in

low- p_T B -physics and trigger-level analyses. This thesis describes the methodology used to compute prescale sets under a wide range of scenarios, as implemented in the Rulebook software package, which provides a flexible and configurable framework capable of adapting to various data-taking and detector conditions. With Run 3 approaching its end, in June 2026, the preparations for Run 4 are already fully ongoing. From an operational point of view, the aim is to optimise and automate procedures as much as possible, to maximise the outcome while minimising the risk of human error. In terms of prescale strategy, efforts are already ongoing to make possible to optimise the end-of-fill rates in real time during data taking, through the development of applications capable of computing the prescale values for dedicated sets of end-of-fill physics triggers based on the actual trigger rates and data-taking conditions. This would also allow prescales to be updated automatically, without requiring manual intervention from the control room shifter. Another option to simplify the offline prescale calculations is to rely on the CTP, which has access to the bunch pattern information, to automatically scale the rate of certain triggers, such as random triggers, according to the number of colliding bunches. Depending on how the operational handling of the prescales will be developed, the Rulebook code will need to be adapted to interface with these new tools, while continuing to benefit from its flexibility and robustness, which allow it to adapt to the complexity of the many possible running conditions.

Appendices

Nuisance Parameters

This appendix presents additional studies on the fits' NPs.

Figures [A.1](#), [A.2](#), and [A.3](#) show, for each NP, the normalised difference between the fitted and the expected values (α), for the low-, medium-, and high- Δm fits, respectively. The $t\bar{t}$ and $t\bar{t}Z$ NFs (μ) are also shown in blue, together with the parameters related to the uncertainties due to limited MC statistics in individual CRs (γ).

Figures [A.4](#), [A.5](#), and [A.6](#) show the correlation matrices of the NPs for the low-, medium-, and high- Δm fits, respectively.

In both sets of figures, the plots on the left represent the results of the Run 2 fits, while the ones on the right correspond to the Run 3 fits.





Figure A.1 – NPs of the background-only low- Δm fit for Run 2 (left), and Run 3 (right).






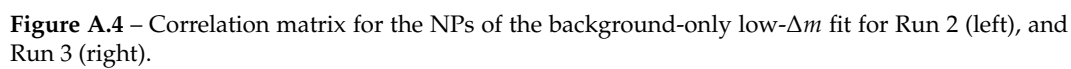


Figure A.3 – NPs of the background-only high- Δm fit for Run 2 (left), and Run 3 (right).



Bibliography

- [1] G. R. Farrar and P. Fayet, *Phenomenology of the production, decay, and detection of new hadronic states associated with supersymmetry*, *Physics Letters B* **76** (1978) 575, ISSN: 0370-2693, URL: <https://www.sciencedirect.com/science/article/pii/0370269378908584> (cit. on p. 1).
- [2] L. Evans and P. Bryant, *LHC Machine*, *Journal of Instrumentation* **3** (2008) S08001, URL: <https://dx.doi.org/10.1088/1748-0221/3/08/S08001> (cit. on pp. 2, 41).
- [3] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, *Journal of Instrumentation* **3** (2008) S08003, URL: <https://dx.doi.org/10.1088/1748-0221/3/08/S08003> (cit. on pp. 2, 41, 51).
- [4] ATLAS Collaboration, *The ATLAS trigger system for LHC Run 3 and trigger performance in 2022*, *JINST* **19** (2024) P06029, arXiv: 2401.06630 [hep-ex] (cit. on pp. 3, 5, 61, 64, 171, 176, 181).
- [5] ATLAS Collaboration, *Trigger Operation Public Results*, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TriggerOperationPublicResults> (cit. on pp. 5, 177, 179, 184).
- [6] J. Albrecht et al., *Summary of the trigger systems of the Large Hadron Collider experiments ALICE, ATLAS, CMS and LHCb*, *J. Phys. G* **52** (2025) 030501, arXiv: 2408.03881, URL: <https://cds.cern.ch/record/2924722> (cit. on pp. 5, 160).
- [7] Y. Fukuda et al., *Evidence for oscillation of atmospheric neutrinos*, *Phys. Rev. Lett.* **81** (8 1998) 1562, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.81.1562> (cit. on p. 10).
- [8] M. Aker et al., *Direct neutrino-mass measurement with sub-electronvolt sensitivity*, *Nature Physics* **18** (2022) 160, URL: <https://doi.org/10.1038/s41567-021-01463-1> (cit. on p. 10).
- [9] Wikipedia, *Standard Model*, URL: https://en.wikipedia.org/wiki/Standard_Model (cit. on p. 11).
- [10] B. R. Stella and H.-J. Meyer, *$Y(9.46\text{ GeV})$ and the gluon discovery (a critical recollection of PLUTO results)*, *Eur. Phys. J. H* **36** (2011) 203, arXiv: 1008.1869 [hep-ex] (cit. on p. 12).

- [11] T. Wengler, *Determination of the strong coupling constant at LEP*, (2006) 459, URL: <https://cds.cern.ch/record/973176> (cit. on p. 14).
- [12] UA1 Collaboration, *Experimental observation of isolated large transverse energy electrons with associated missing energy at $s=540$ GeV*, *Physics Letters B* **122** (1983) 103, ISSN: 0370-2693, URL: <https://www.sciencedirect.com/science/article/pii/0370269383911772> (cit. on p. 16).
- [13] UA2 Collaboration, *Observation of single isolated electrons of high transverse momentum in events with missing transverse energy at the CERN pp collider*, *Physics Letters B* **122** (1983) 476, ISSN: 0370-2693, URL: <https://www.sciencedirect.com/science/article/pii/0370269383916052> (cit. on p. 16).
- [14] UA1 Collaboration, *Experimental observation of lepton pairs of invariant mass around 95 GeV/c² at the CERN SPS collider*, *Physics Letters B* **126** (1983) 398, ISSN: 0370-2693, URL: <https://www.sciencedirect.com/science/article/pii/0370269383901880> (cit. on p. 16).
- [15] UA2 Collaboration, *Evidence for $Z^0 \rightarrow e^+e^-$ at the CERN pp collider*, *Physics Letters B* **129** (1983) 130, ISSN: 0370-2693, URL: <https://www.sciencedirect.com/science/article/pii/037026938390744X> (cit. on p. 16).
- [16] J. B. Adams, *The CERN 400 GeV Proton Synchrotron (CERN SPS)*, (1977), URL: <https://cds.cern.ch/record/2048898> (cit. on p. 16).
- [17] F. Englert and R. Brout, *Broken symmetry and the mass of gauge vector mesons*, *Phys. Rev. Lett.* **13** (1964) 321, ed. by J. C. Taylor (cit. on p. 17).
- [18] P. W. Higgs, *Broken symmetries and the masses of gauge bosons*, *Phys. Rev. Lett.* **13** (1964) 508, ed. by J. C. Taylor (cit. on p. 17).
- [19] G. S. Guralnik et al., *Global conservation laws and massless particles*, *Phys. Rev. Lett.* **13** (1964) 585, ed. by J. C. Taylor (cit. on p. 17).
- [20] R. L. Workman et al., *Review of Particle Physics*, *PTEP* **2022** (2022) 083C01 (cit. on pp. 18, 23, 108).
- [21] ATLAS Collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, *Physics Letters B* **716** (2012) 1, ISSN: 0370-2693, URL: <http://dx.doi.org/10.1016/j.physletb.2012.08.020> (cit. on p. 19).
- [22] CMS Collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, *Physics Letters B* **716** (2012) 30, ISSN: 0370-2693, URL: <http://dx.doi.org/10.1016/j.physletb.2012.08.021> (cit. on p. 19).
- [23] Planck Collaboration, *Planck2018 results: VI. Cosmological parameters*, *Astronomy & Astrophysics* **641** (2020) A6, ISSN: 1432-0746, URL: <http://dx.doi.org/10.1051/0004-6361/201833910> (cit. on p. 20).
- [24] S. P. Martin, *A Supersymmetry primer*, *Adv. Ser. Direct. High Energy Phys.* **18** (1998) 1, ed. by G. L. Kane, arXiv: [hep-ph/9709356](https://arxiv.org/abs/hep-ph/9709356) (cit. on pp. 21, 22, 24, 26).
- [25] The Muon g-2 Collaboration, *Measurement of the Positive Muon Anomalous Magnetic Moment to 127 ppb*, 2025, arXiv: [2506.03069 \[hep-ex\]](https://arxiv.org/abs/2506.03069), URL: <https://arxiv.org/abs/2506.03069> (cit. on p. 23).

- [26] K. Abe et al., *Search for Proton Decay via $p \rightarrow e^+ \pi^0$ and $p \rightarrow \mu^+ \pi^0$ in 0.31 megaton-years exposure of the Super-Kamiokande Water Cherenkov Detector*, *Physical Review D* **95** (2017), ISSN: 2470-0029, URL: <http://dx.doi.org/10.1103/PhysRevD.95.012004> (cit. on p. 26).
- [27] L. Girardello and M. Grisaru, *Soft breaking of supersymmetry*, *Nuclear Physics B* **194** (1982) 65, ISSN: 0550-3213, URL: <https://www.sciencedirect.com/science/article/pii/0550321382905120> (cit. on p. 27).
- [28] LHC SUSY Cross Section Working Group, URL: <https://twiki.cern.ch/twiki/bin/view/LHCPhysics/SUSYCrossSections> (cit. on p. 31).
- [29] W. Beenakker et al., *NNLL-fast 2.0: Coloured sparticle production at the LHC run 3 with $\sqrt{s} = 13.6$ TeV*, *SciPost Physics Core* **7** (2024), ISSN: 2666-9366, URL: <http://dx.doi.org/10.21468/SciPostPhysCore.7.4.072> (cit. on pp. 32, 103).
- [30] W. Beenakker et al., *Stop production at hadron colliders*, *Nuclear Physics B* **515** (1998) 3, ISSN: 0550-3213, URL: [http://dx.doi.org/10.1016/S0550-3213\(98\)00014-5](http://dx.doi.org/10.1016/S0550-3213(98)00014-5) (cit. on p. 32).
- [31] A. Broggio et al., *Approximate NNLO predictions for the stop-pair production cross section at the LHC*, *Journal of High Energy Physics* **2013** (2013), ISSN: 1029-8479, URL: [http://dx.doi.org/10.1007/JHEP07\(2013\)042](http://dx.doi.org/10.1007/JHEP07(2013)042) (cit. on p. 33).
- [32] A. Canepa, *Searches for supersymmetry at the Large Hadron Collider*, *Reviews in Physics* **4** (2019) 100033, ISSN: 2405-4283, URL: <https://www.sciencedirect.com/science/article/pii/S2405428318300091> (cit. on p. 33).
- [33] ATLAS Collaboration, *Physics Reports* **1116** (2025) 261, ISSN: 0370-1573, URL: <http://dx.doi.org/10.1016/j.physrep.2024.09.010> (cit. on p. 34).
- [34] SUSY July 2024 Summary Plot Update, tech. rep., CERN, 2024, URL: <https://cds.cern.ch/record/2904978> (cit. on pp. 35, 37).
- [35] ATLAS Collaboration, *Search for new phenomena with top-quark pairs and large missing transverse momentum using 140fb^{-1} of pp collision data at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Journal of High Energy Physics* **2024** (2024), ISSN: 1029-8479, URL: [http://dx.doi.org/10.1007/JHEP03\(2024\)139](http://dx.doi.org/10.1007/JHEP03(2024)139) (cit. on pp. 35, 151).
- [36] ATLAS Collaboration, *Search for a scalar partner of the top quark in the all-hadronic $t\bar{t}$ plus missing transverse momentum final state at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *The European Physical Journal C* **80** (2020), ISSN: 1434-6052, URL: <http://dx.doi.org/10.1140/epjc/s10052-020-8102-8> (cit. on p. 35).
- [37] CMS Collaboration, *Combined searches for the production of supersymmetric top quark partners in proton-proton collisions at 13 TeV*, *The European Physical Journal C* **81** (2021), ISSN: 1434-6052, URL: <http://dx.doi.org/10.1140/epjc/s10052-021-09721-5> (cit. on p. 36).
- [38] CMS Collaboration, *Search for top squark production in fully hadronic final states in proton-proton collisions at 13 TeV*, *Physical Review D* **104** (2021), ISSN: 2470-0029, URL: <http://dx.doi.org/10.1103/PhysRevD.104.052001> (cit. on p. 36).
- [39] ATLAS Collaboration, *Summary of the ATLAS experiment's sensitivity to supersymmetry after LHC Run 1 — interpreted in the phenomenological MSSM*, *Journal of High Energy Physics* **2015** (2015), ISSN: 1029-8479, URL: [http://dx.doi.org/10.1007/JHEP10\(2015\)134](http://dx.doi.org/10.1007/JHEP10(2015)134) (cit. on p. 37).

- [40] O. Aberle et al., *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*, CERN Yellow Reports: Monographs, Geneva: CERN, 2020, URL: <https://cds.cern.ch/record/2749422> (cit. on p. 41).
- [41] *Longer term LHC schedule*, URL: <https://lhc-commissioning.web.cern.ch/schedule/LHC-long-term.htm> (cit. on p. 42).
- [42] E. Lopienska, *The CERN accelerator complex, layout in 2022. Complexe des accélérateurs du CERN en janvier 2022*, (2022), URL: <https://cds.cern.ch/record/2800984> (cit. on p. 43).
- [43] CMS Collaboration, *The CMS experiment at the CERN LHC*, *Journal of Instrumentation* **3** (2008) S08004, URL: <https://dx.doi.org/10.1088/1748-0221/3/08/S08004> (cit. on p. 42).
- [44] LHCb Collaboration, *The LHCb Detector at the LHC*, *Journal of Instrumentation* **3** (2008) S08005, URL: <https://dx.doi.org/10.1088/1748-0221/3/08/S08005> (cit. on p. 42).
- [45] ALICE Collaboration, *The ALICE experiment at the CERN LHC*, *Journal of Instrumentation* **3** (2008) S08002, URL: <https://dx.doi.org/10.1088/1748-0221/3/08/S08002> (cit. on p. 43).
- [46] TOTEM Collaboration, *The TOTEM Experiment at the CERN Large Hadron Collider*, *Journal of Instrumentation* **3** (2008) S08007, URL: <https://dx.doi.org/10.1088/1748-0221/3/08/S08007> (cit. on p. 43).
- [47] The LHCf Collaboration, *The LHCf detector at the CERN Large Hadron Collider*, *Journal of Instrumentation* **3** (2008) S08006, URL: <https://dx.doi.org/10.1088/1748-0221/3/08/S08006> (cit. on p. 43).
- [48] V. A. Mitsou, *The MoEDAL experiment at the LHC: status and results*, *Journal of Physics: Conference Series* **873** (2017) 012010, ISSN: 1742-6596, URL: <http://dx.doi.org/10.1088/1742-6596/873/1/012010> (cit. on p. 43).
- [49] FASER Collaboration, *The FASER detector*, *Journal of Instrumentation* **19** (2024) P05066, URL: <https://dx.doi.org/10.1088/1748-0221/19/05/P05066> (cit. on p. 43).
- [50] SND@LHC Collaboration, *SND@LHC: the scattering and neutrino detector at the LHC*, *Journal of Instrumentation* **19** (2024) P05067, ISSN: 1748-0221, URL: <http://dx.doi.org/10.1088/1748-0221/19/05/P05067> (cit. on p. 43).
- [51] R. Alemany et al., *LHC modes*, LHC-OP-ES-0005-10-00, 2007 (cit. on p. 43).
- [52] ATLAS Collaboration, *Operation of the ATLAS trigger system in Run 2*, *Journal of Instrumentation* **15** (2020) P10004, ISSN: 1748-0221, URL: <http://dx.doi.org/10.1088/1748-0221/15/10/P10004> (cit. on pp. 44, 174).
- [53] B. Salvachua, *Overview of Proton-Proton Physics during Run 2*, (2019) 7, URL: <https://cds.cern.ch/record/2750272> (cit. on p. 45).
- [54] G. Arduini et al., *LHC Upgrades in preparation of Run 3*, *Journal of Instrumentation* **19** (2024) P05061, URL: <https://dx.doi.org/10.1088/1748-0221/19/05/P05061> (cit. on p. 45).

- [55] ATLAS Collaboration, *Public ATLAS Luminosity Results for Run-2 of the LHC*, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2> (cit. on pp. 46, 48, 49).
- [56] ATLAS Collaboration, *Public ATLAS Luminosity Results for Run-3 of the LHC*, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun3> (cit. on pp. 47–49).
- [57] M. Hostettler et al., *Operational β^* levelling at the LHC in 2022 and beyond*, *JACoW IPAC 2023 (2023) MOPL045*, URL: <https://cds.cern.ch/record/2884758> (cit. on p. 50).
- [58] ATLAS Collaboration, *The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run 3*, *JINST 19 (2024) P05063*, arXiv: 2305.16623 [physics.ins-det] (cit. on pp. 51, 52, 55, 56, 58, 62, 65).
- [59] S. Mehlhase, *ATLAS detector slice (and particle visualisations)*, (2021), URL: <https://cds.cern.ch/record/2770815> (cit. on p. 53).
- [60] I. Neutelings, *Images from TikZ.net*, URL: <https://tikz.net/> (cit. on p. 53).
- [61] ATLAS Collaboration, *ATLAS central solenoid: Technical design report*, (1997) (cit. on p. 54).
- [62] ATLAS Collaboration, *ATLAS barrel toroid: Technical design report*, (1997) (cit. on p. 54).
- [63] ATLAS Collaboration, *ATLAS end-cap toroids: Technical design report*, (1997) (cit. on p. 54).
- [64] ATLAS Collaboration, *ATLAS inner detector: Technical Design Report, 1*, Technical design report. ATLAS, Geneva: CERN, 1997, URL: <https://cds.cern.ch/record/331063> (cit. on p. 55).
- [65] F. Hartmann, “Basic Principles of a Silicon Detector”, *Evolution of Silicon Sensor Technology in Particle Physics: Basics and Applications*, Cham: Springer Nature Switzerland, 2024, ISBN: 978-3-031-59720-6, URL: https://doi.org/10.1007/978-3-031-59720-6_1 (cit. on p. 55).
- [66] G. Aad et al., *ATLAS pixel detector electronics and sensors*, *JINST 3 (2008) P07007*, URL: <https://cds.cern.ch/record/1119279> (cit. on p. 55).
- [67] N. Wermes and G. Hallewel, *ATLAS pixel detector: Technical Design Report*, Technical design report. ATLAS, Geneva: CERN, 1998, URL: <https://cds.cern.ch/record/381263> (cit. on p. 55).
- [68] ATLAS Collaboration, *Performance of ATLAS Pixel Detector and Track Reconstruction at the start of Run 3 in LHC Collisions at $\sqrt{s} = 900$ GeV*, tech. rep., CERN, 2022, URL: <https://cds.cern.ch/record/2814766> (cit. on p. 55).
- [69] M. Capeans et al., *ATLAS Insertable B-Layer Technical Design Report*, (2010) (cit. on p. 56).
- [70] ATLAS Collaboration, *Operation and performance of the ATLAS semiconductor tracker*, *JINST 9 (2014) P08009*, arXiv: 1404.7473, URL: <https://cds.cern.ch/record/1698966> (cit. on p. 56).
- [71] A. Vogel, *ATLAS Transition Radiation Tracker (TRT): Straw Tube Gaseous Detectors at High Rates*, tech. rep., CERN, 2013, URL: <https://cds.cern.ch/record/1537991> (cit. on p. 57).

- [72] C. W. Fabjan and F. Gianotti, *Calorimetry for particle physics*, *Rev. Mod. Phys.* **75** (2003) 1243 (cit. on p. 57).
- [73] R. Wigmans, *Sampling calorimetry*, *Nucl. Instrum. Meth. A* **494** (2002) 277, ed. by G. V. Fedotovitch and B. Khazin (cit. on p. 57).
- [74] ATLAS Collaboration, *ATLAS liquid argon calorimeter: Technical design report*, (1996) (cit. on pp. 58, 60).
- [75] M. Aleksa et al., *ATLAS Liquid Argon Calorimeter Phase-I Upgrade Technical Design Report*, (2013) (cit. on p. 58).
- [76] G. Aad et al., *The Phase-I trigger readout electronics upgrade of the ATLAS Liquid Argon calorimeters*, *Journal of Instrumentation* **17** (2022) P05024, URL: <https://dx.doi.org/10.1088/1748-0221/17/05/P05024> (cit. on p. 59).
- [77] ATLAS Collaboration, *ATLAS tile calorimeter: Technical Design Report*, Technical design report. ATLAS, Geneva: CERN, 1996, URL: <https://cds.cern.ch/record/331062> (cit. on p. 60).
- [78] A. Sidoti, *Minimum Bias Trigger Scintillators in ATLAS Run II*, *Journal of Instrumentation* **9** (2014) C10020, URL: <https://dx.doi.org/10.1088/1748-0221/9/10/C10020> (cit. on p. 60).
- [79] A. Artamonov et al., *The ATLAS Forward Calorimeter*, *JINST* **3** (2008) P02010, URL: <https://cds.cern.ch/record/1094547> (cit. on p. 61).
- [80] ATLAS Collaboration, *ATLAS muon spectrometer: Technical design report*, (1997) (cit. on p. 61).
- [81] G. Avoni et al., *The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS*, *Journal of Instrumentation* **13** (2018) P07017, URL: <https://dx.doi.org/10.1088/1748-0221/13/07/P07017> (cit. on p. 63).
- [82] P. Jenni et al., *Zero Degree Calorimeters for ATLAS*, tech. rep., CERN, 2007, URL: <https://cds.cern.ch/record/1009649> (cit. on p. 63).
- [83] L. Adamczyk et al., *Technical Design Report for the ATLAS Forward Proton Detector*, tech. rep., 2015, URL: <https://cds.cern.ch/record/2017378> (cit. on p. 63).
- [84] S. A. Khalek et al., *The ALFA Roman Pot detectors of ATLAS*, *Journal of Instrumentation* **11** (2016) P11013, URL: <https://dx.doi.org/10.1088/1748-0221/11/11/P11013> (cit. on p. 63).
- [85] H. Bertelsen et al., *Operation of the upgraded ATLAS Central Trigger Processor during the LHC Run 2*, *Journal of Instrumentation* **11** (2016) C02020, URL: <https://dx.doi.org/10.1088/1748-0221/11/02/C02020> (cit. on p. 64).
- [86] S. Cella, *The ATLAS Run-3 Trigger Menu*, *PoS ICHEP2024* (2024) 943 (cit. on p. 64).
- [87] ATLAS Collaboration, *L1 Calorimeter Trigger Public Results*, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/L1CaloTriggerPublicResults> (cit. on p. 66).
- [88] ATLAS Collaboration, *L1 Muon Trigger Public Results*, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/L1MuonTriggerPublicResults> (cit. on p. 66).

- [89] M. Mironova, “Experimental Particle Physics with the ATLAS Detector”, *Search for Higgs Boson Decays to Charm Quarks with the ATLAS Experiment and Development of Novel Silicon Pixel Detectors*, Cham: Springer Nature Switzerland, 2023 23, ISBN: 978-3-031-36220-0, URL: https://doi.org/10.1007/978-3-031-36220-0_3 (cit. on p. 70).
- [90] ATLAS Collaboration, *Software Performance of the ATLAS Track Reconstruction for LHC Run 3*, *Computing and Software for Big Science* **8** (2024), ISSN: 2510-2044, URL: <http://dx.doi.org/10.1007/s41781-023-00111-y> (cit. on p. 70).
- [91] R. Frühwirth, *Application of Kalman filtering to track and vertex fitting*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **262** (1987) 444, ISSN: 0168-9002, URL: <https://www.sciencedirect.com/science/article/pii/0168900287908874> (cit. on p. 70).
- [92] ATLAS Collaboration, *Development of ATLAS Primary Vertex Reconstruction for LHC Run 3*, tech. rep., CERN, 2019, URL: <https://cds.cern.ch/record/2670380> (cit. on p. 71).
- [93] ATLAS Collaboration, *Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **81** (2021) 578, arXiv: 2012.00578 [hep-ex] (cit. on pp. 71, 79).
- [94] ATLAS Collaboration, *Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1*, *Eur. Phys. J. C* **77** (2017) 490, arXiv: 1603.02934, URL: <https://cds.cern.ch/record/2138166> (cit. on p. 72).
- [95] ATLAS Collaboration, *Jet reconstruction and performance using particle flow with the ATLAS Detector*, *Eur. Phys. J. C* **77** (2017) 466, arXiv: 1703.10485 [hep-ex] (cit. on p. 73).
- [96] ATLAS Collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data*, *JINST* **14** (2019) P12006, arXiv: 1908.00005 [hep-ex] (cit. on pp. 74–76, 79).
- [97] ATLAS Collaboration, *Electron and Photon Performance Public Plots*, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ElectronGammaPublicCollisionResults> (cit. on pp. 78, 79).
- [98] CMS Collaboration, *Jets at CMS and the determination of their energy scale*, URL: <https://cms.cern/news/jets-cms-and-determination-their-energy-scale> (cit. on p. 82).
- [99] M. Cacciari et al., *The anti- k_t jet clustering algorithm*, *Journal of High Energy Physics* **2008** (2008) 063, URL: <https://dx.doi.org/10.1088/1126-6708/2008/04/063> (cit. on p. 83).
- [100] M. Cacciari et al., *FastJet user manual: (for version 3.0.2)*, *The European Physical Journal C* **72** (2012), ISSN: 1434-6052, URL: <http://dx.doi.org/10.1140/epjc/s10052-012-1896-2> (cit. on p. 83).
- [101] S. Catani et al., *Longitudinally-invariant $k_{\perp}$ -clustering algorithms for hadron-hadron collisions*, *Nuclear Physics B* **406** (1993) 187, ISSN: 0550-3213, URL: <https://www.sciencedirect.com/science/article/pii/055032139390166M> (cit. on p. 83).
- [102] ATLAS Collaboration, *Optimisation of large-radius jet reconstruction for the ATLAS detector in 13 TeV proton-proton collisions*, *The European Physical Journal C* **81** (2021), ISSN: 1434-6052, URL: <http://dx.doi.org/10.1140/epjc/s10052-021-09054-3> (cit. on p. 83).

- [103] ATLAS Collaboration, *Jet energy scale and resolution measured in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *The European Physical Journal C* **81** (2021), ISSN: 1434-6052, URL: <http://dx.doi.org/10.1140/epjc/s10052-021-09402-3> (cit. on pp. 84, 86).
- [104] A. Collaboration, *Transforming jet flavour tagging at ATLAS*, 2025, arXiv: 2505.19689 [hep-ex], URL: <https://arxiv.org/abs/2505.19689> (cit. on p. 87).
- [105] A. Vaswani et al., *Attention Is All You Need*, 2023, arXiv: 1706.03762 [cs.CL], URL: <https://arxiv.org/abs/1706.03762> (cit. on p. 87).
- [106] ATLAS Collaboration, *ATLAS flavour-tagging algorithms for the LHC Run 2 pp collision dataset*, *The European Physical Journal C* **83** (2023), ISSN: 1434-6052, URL: <http://dx.doi.org/10.1140/epjc/s10052-023-11699-1> (cit. on p. 88).
- [107] ATLAS Collaboration, *Graph Neural Network Jet Flavour Tagging with the ATLAS Detector*, tech. rep., CERN, 2022, URL: <https://cds.cern.ch/record/2811135> (cit. on p. 88).
- [108] ATLAS Collaboration, *Flavour Tagging Performance Public Results*, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/FlavourTaggingPublicResultsCollisionData> (cit. on p. 88).
- [109] ATLAS Collaboration, *Tagging and suppression of pileup jets with the ATLAS detector*, tech. rep., CERN, 2014, URL: <https://cds.cern.ch/record/1700870> (cit. on p. 89).
- [110] *Object-based missing transverse momentum significance in the ATLAS detector*, tech. rep., CERN, 2018, URL: <https://cds.cern.ch/record/2630948> (cit. on pp. 90, 91).
- [111] ATLAS Collaboration, *Search for new phenomena in events with two opposite-charge leptons, jets and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Journal of High Energy Physics* **2021** (2021), ISSN: 1029-8479, URL: [http://dx.doi.org/10.1007/JHEP04\(2021\)165](http://dx.doi.org/10.1007/JHEP04(2021)165) (cit. on pp. 96, 98, 107, 110, 120, 123, 127, 152, 195).
- [112] ATLAS Collaboration, *Data Quality Efficiency Plots*, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/DataQualityResults> (cit. on pp. 97, 196).
- [113] ATLAS Collaboration, *Athena*, 2024, URL: <https://gitlab.cern.ch/atlas/athena/> (cit. on pp. 100, 163).
- [114] S. Agostinelli et al., *GEANT4—a simulation toolkit*. *GEANT4. A Simulation toolkit*, *Nucl. Instrum. Methods Phys. Res., A* **506** (2003) 250, URL: <https://cds.cern.ch/record/602040> (cit. on p. 100).
- [115] J. M. Campbell et al., *Top-Pair Production and Decay at NLO Matched with Parton Showers*, *JHEP* **04** (2015) 114, arXiv: 1412.1828 [hep-ph] (cit. on p. 102).
- [116] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, *Computer Physics Communications* **191** (2015) 159, ISSN: 0010-4655, URL: <http://dx.doi.org/10.1016/j.cpc.2015.01.024> (cit. on p. 102).
- [117] R. D. Ball et al., *Parton distributions for the LHC run II*, *Journal of High Energy Physics* **2015** (2015), ISSN: 1029-8479, URL: [http://dx.doi.org/10.1007/JHEP04\(2015\)040](http://dx.doi.org/10.1007/JHEP04(2015)040) (cit. on p. 102).

- [118] R. D. Ball et al., *Parton distributions with LHC data*, *Nuclear Physics B* **867** (2013) 244, ISSN: 0550-3213, URL: <http://dx.doi.org/10.1016/j.nuclphysb.2012.10.003> (cit. on p. 102).
- [119] J. Alwall et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *Journal of High Energy Physics* **2014** (2014), ISSN: 1029-8479, URL: [http://dx.doi.org/10.1007/JHEP07\(2014\)079](http://dx.doi.org/10.1007/JHEP07(2014)079) (cit. on p. 102).
- [120] E. Re, *Single-top Wt-channel production matched with parton showers using the POWHEG method*, *Eur. Phys. J. C* **71** (2011) 1547, arXiv: 1009.2450 [hep-ph] (cit. on p. 102).
- [121] E. Bothmann et al., *Event generation with Sherpa 2.2*, *SciPost Physics* **7** (2019), ISSN: 2542-4653, URL: <http://dx.doi.org/10.21468/SciPostPhys.7.3.034> (cit. on p. 102).
- [122] S. Schumann and F. Krauss, *A parton shower algorithm based on Catani-Seymour dipole factorisation*, *Journal of High Energy Physics* **2008** (2008) 038, ISSN: 1029-8479, URL: <http://dx.doi.org/10.1088/1126-6708/2008/03/038> (cit. on p. 102).
- [123] R. Frederix et al., *The automation of next-to-leading order electroweak calculations*, *Journal of High Energy Physics* **2018** (2018), ISSN: 1029-8479, URL: [http://dx.doi.org/10.1007/JHEP07\(2018\)185](http://dx.doi.org/10.1007/JHEP07(2018)185) (cit. on p. 102).
- [124] P. Artoisenet et al., *Automatic spin-entangled decays of heavy resonances in Monte Carlo simulations*, *Journal of High Energy Physics* **2013** (2013), ISSN: 1029-8479, URL: [http://dx.doi.org/10.1007/JHEP03\(2013\)015](http://dx.doi.org/10.1007/JHEP03(2013)015) (cit. on p. 102).
- [125] R. D. Ball et al., *Parton distributions for the LHC run II*, *Journal of High Energy Physics* **2015** (2015), ISSN: 1029-8479, URL: [http://dx.doi.org/10.1007/JHEP04\(2015\)040](http://dx.doi.org/10.1007/JHEP04(2015)040) (cit. on p. 102).
- [126] C. Bierlich et al., *A comprehensive guide to the physics and usage of PYTHIA 8.3*, 2022, arXiv: 2203.11601 [hep-ph], URL: <https://arxiv.org/abs/2203.11601> (cit. on pp. 102, 103).
- [127] *ATLAS Pythia 8 tunes to 7 TeV data*, tech. rep., CERN, 2014, URL: <https://cds.cern.ch/record/1966419> (cit. on p. 103).
- [128] D. J. Lange, *The EvtGen particle decay simulation package*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **462** (2001) 152, ISSN: 0168-9002, URL: <https://www.sciencedirect.com/science/article/pii/S0168900201000894> (cit. on p. 103).
- [129] L. Lönnblad and S. Prestel, *Merging multi-leg NLO matrix elements with parton showers*, *Journal of High Energy Physics* **2013** (2013), ISSN: 1029-8479, URL: [http://dx.doi.org/10.1007/JHEP03\(2013\)166](http://dx.doi.org/10.1007/JHEP03(2013)166) (cit. on p. 103).
- [130] J. Butterworth et al., *PDF4LHC recommendations for LHC Run II*, *Journal of Physics G: Nuclear and Particle Physics* **43** (2016) 023001, ISSN: 1361-6471, URL: <http://dx.doi.org/10.1088/0954-3899/43/2/023001> (cit. on p. 103).
- [131] *Tagging and suppression of pileup jets with the ATLAS detector*, tech. rep., CERN, 2014, URL: <https://cds.cern.ch/record/1700870> (cit. on pp. 105, 143).

- [132] C. Lester and D. Summers, *Measuring masses of semi-invisibly decaying particle pairs produced at hadron colliders*, *Physics Letters B* **463** (1999) 99, ISSN: 0370-2693, URL: [http://dx.doi.org/10.1016/S0370-2693\(99\)00945-4](http://dx.doi.org/10.1016/S0370-2693(99)00945-4) (cit. on pp. 109, 110).
- [133] A. Barr et al., *A variable for measuring masses at hadron colliders when missing energy is expected; mT2: the truth behind the glamour*, *Journal of Physics G: Nuclear and Particle Physics* **29** (2003) 2343, ISSN: 1361-6471, URL: <http://dx.doi.org/10.1088/0954-3899/29/10/304> (cit. on p. 109).
- [134] G. Polesello and D. R. Tovey, *Supersymmetric particle mass measurement with the boost-corrected contranverse mass*, *Journal of High Energy Physics* **2010** (2010), ISSN: 1029-8479, URL: [http://dx.doi.org/10.1007/JHEP03\(2010\)030](http://dx.doi.org/10.1007/JHEP03(2010)030) (cit. on p. 111).
- [135] S. Lundberg and S.-I. Lee, *A Unified Approach to Interpreting Model Predictions*, 2017, arXiv: 1705.07874 [cs.AI], URL: <https://arxiv.org/abs/1705.07874> (cit. on p. 119).
- [136] G. Cowan, *Discovery sensitivity for a counting experiment with background uncertainty*, tech. rep., Royal Holloway, 2012, URL: <https://www.pp.rhul.ac.uk/~cowan/stat/medsig/medsigNote.pdf> (cit. on p. 120).
- [137] F. Chollet et al., *Keras*, <https://keras.io>, 2015 (cit. on p. 122).
- [138] M. Abadi et al., *TensorFlow: A system for large-scale machine learning*, 2016, arXiv: 1605.08695 [cs.DC], URL: <https://arxiv.org/abs/1605.08695> (cit. on p. 122).
- [139] G. Cowan et al., *Asymptotic formulae for likelihood-based tests of new physics*, *The European Physical Journal C* **71** (2011) 1554, ISSN: 1434-6052, URL: <https://doi.org/10.1140/epjc/s10052-011-1554-0> (cit. on pp. 127, 140–142).
- [140] M. Baak et al., *HistFitter software framework for statistical data analysis*, *The European Physical Journal C* **75** (2015), ISSN: 1434-6052, URL: <http://dx.doi.org/10.1140/epjc/s10052-015-3327-7> (cit. on p. 137).
- [141] G. J. Besjes et al., *HistFitter: a flexible framework for statistical data analysis*, *J. Phys.: Conf. Ser.* **664** (2015) 072004, URL: <https://cds.cern.ch/record/2134618> (cit. on p. 137).
- [142] L. Heinrich et al., *pyhf: v0.7.6*, version 0.7.6, URL: <https://doi.org/10.5281/zenodo.1169739> (cit. on p. 137).
- [143] L. Heinrich et al., *pyhf: pure-Python implementation of HistFactory statistical models*, *Journal of Open Source Software* **6** (2021) 2823, URL: <https://doi.org/10.21105/joss.02823> (cit. on p. 137).
- [144] A. L. Read, *Presentation of search results: The CL_s technique*, *J. Phys. G* **28** (2002) 2693, ed. by M. R. Whalley and L. Lyons (cit. on p. 141).
- [145] ATLAS Collaboration, *Measurement of the Inelastic Proton-Proton Cross Section at 13 TeV with the ATLAS Detector at the LHC*, *Physical Review Letters* **117** (2016), ISSN: 1079-7114, URL: <http://dx.doi.org/10.1103/PhysRevLett.117.182002> (cit. on p. 143).
- [146] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, *The European Physical Journal C* **83** (2023), ISSN: 1434-6052, URL: <http://dx.doi.org/10.1140/epjc/s10052-023-11747-w> (cit. on p. 143).
- [147] *Preliminary analysis of the luminosity calibration for the ATLAS 13.6 TeV data recorded in 2023*, tech. rep., CERN, 2024, URL: <https://cds.cern.ch/record/2900949> (cit. on p. 143).

- [148] S. Höche et al., *A Study of QCD Radiation in VBF Higgs Production with Vincia and Pythia*, *SciPost Phys.* **12** (2022) 010, URL: <https://scipost.org/10.21468/SciPostPhys.12.1.010> (cit. on p. 144).
- [149] S. Frixione et al., *Single-top hadroproduction in association with a W boson*, *Journal of High Energy Physics* **2008** (2008) 029, ISSN: 1029-8479, URL: <http://dx.doi.org/10.1088/1126-6708/2008/07/029> (cit. on p. 144).
- [150] C. Borschensky et al., *Squark and gluino production cross sections in pp collisions at $\sqrt{s} = 13, 14, 33$ and 100 TeV*, *The European Physical Journal C* **74** (2014), ISSN: 1434-6052, URL: <http://dx.doi.org/10.1140/epjc/s10052-014-3174-y> (cit. on p. 144).
- [151] D. de Florian et al., *CERN Yellow Reports: Monographs, Vol 2 (2017): Handbook of LHC Higgs cross sections: 4. Deciphering the nature of the Higgs sector*, 2017, URL: <https://e-publishing.cern.ch/index.php/CYRM/issue/view/32> (cit. on pp. 153, 154).
- [152] ATLAS Collaboration, *Combination of searches for invisible decays of the Higgs boson using 139fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment*, *Physics Letters B* **842** (2023) 137963, ISSN: 0370-2693, URL: <http://dx.doi.org/10.1016/j.physletb.2023.137963> (cit. on p. 153).
- [153] ATLAS Collaboration, *Search for new phenomena in events with an energetic jet and missing transverse momentum in pp collisions at $\sqrt{s} = 13$ with the ATLAS detector*, *Physical Review D* **103** (2021), ISSN: 2470-0029, URL: <http://dx.doi.org/10.1103/PhysRevD.103.112006> (cit. on p. 153).
- [154] ATLAS Collaboration, *Search for invisible Higgs-boson decays in events with vector-boson fusion signatures using 139fb^{-1} of proton-proton data recorded by the ATLAS experiment*, *Journal of High Energy Physics* **2022** (2022), ISSN: 1029-8479, URL: [http://dx.doi.org/10.1007/JHEP08\(2022\)104](http://dx.doi.org/10.1007/JHEP08(2022)104) (cit. on p. 153).
- [155] ATLAS Collaboration, *Observation of electroweak production of two jets in association with an isolated photon and missing transverse momentum, and search for a Higgs boson decaying into invisible particles at 13 TeV with the ATLAS detector*, *The European Physical Journal C* **82** (2022), ISSN: 1434-6052, URL: <http://dx.doi.org/10.1140/epjc/s10052-021-09878-z> (cit. on p. 153).
- [156] ATLAS Collaboration, *Search for associated production of a Z boson with an invisibly decaying Higgs boson or dark matter candidates at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Physics Letters B* **829** (2022) 137066, ISSN: 0370-2693, URL: <http://dx.doi.org/10.1016/j.physletb.2022.137066> (cit. on p. 153).
- [157] ATLAS Collaboration, *Constraints on spin-0 dark matter mediators and invisible Higgs decays using ATLAS 13 TeV pp collision data with two top quarks and missing energy in the final state*, tech. rep., CERN, 2022, URL: <https://cds.cern.ch/record/2805211> (cit. on pp. 153, 155, 195).
- [158] ATLAS Collaboration, *Constraints on new phenomena via Higgs boson couplings and invisible decays with the ATLAS detector*, *Journal of High Energy Physics* **2015** (2015), ISSN: 1029-8479, URL: [http://dx.doi.org/10.1007/JHEP11\(2015\)206](http://dx.doi.org/10.1007/JHEP11(2015)206) (cit. on p. 153).

- [159] S. Heinemeyer et al., *Handbook of LHC Higgs Cross Sections: 3. Higgs Properties: Report of the LHC Higgs Cross Section Working Group*, ed. by S. Heinemeyer, CERN Yellow Reports: Monographs, 2013, URL: <https://cds.cern.ch/record/1559921> (cit. on p. 155).
- [160] ATLAS Collaboration, *Public Egamma Trigger Plots for Collision Data*, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/EgammaTriggerPublicResults> (cit. on p. 162).
- [161] ATLAS Collaboration, *Muon Trigger Public Results*, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/MuonTriggerPublicResults> (cit. on p. 162).
- [162] G. Barrand et al., *GAUDI — A software architecture and framework for building HEP data processing applications*, *Computer Physics Communications* **140** (2001) 45, ISSN: 0010-4655, URL: <https://www.sciencedirect.com/science/article/pii/S0010465501002545> (cit. on p. 163).
- [163] R. Bielski and A. Collaboration, *Implementation of the ATLAS trigger within the multi-threaded AthenaMT framework*, (2019), URL: <https://cds.cern.ch/record/2692714> (cit. on p. 164).
- [164] M. Clemencic et al., *Introducing concurrency in the Gaudi data processing framework*, *Journal of Physics: Conference Series* **513** (2014) 022013, URL: <https://dx.doi.org/10.1088/1742-6596/513/2/022013> (cit. on p. 165).
- [165] *Intel Threading Building Blocks*, URL: <https://github.com/uxlfoundation/oneTBB> (cit. on p. 165).
- [166] *Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System*, tech. rep., CERN, 2017, URL: <https://cds.cern.ch/record/2285584> (cit. on p. 165).
- [167] M. Borodin et al., *Scaling up ATLAS production system for the LHC Run 2 and beyond: project ProdSys2*, *Journal of Physics: Conference Series* **664** (2015) 062005, URL: <https://dx.doi.org/10.1088/1742-6596/664/6/062005> (cit. on p. 167).
- [168] V. Lendermann et al., *Combining Triggers in HEP data analysis*, *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* **604** (2009) 707, ISSN: 0168-9002, URL: <http://dx.doi.org/10.1016/j.nima.2009.03.173> (cit. on p. 168).
- [169] *Trigger monitoring and rate predictions using Enhanced Bias data from the ATLAS Detector at the LHC*, tech. rep., CERN, 2016, URL: <https://cds.cern.ch/record/2223498> (cit. on p. 169).
- [170] ATLAS Collaboration, *Athena release 22.0.102*, 2022, URL: <https://gitlab.cern.ch/atlas/athena/-/releases/release/22.0.102> (cit. on p. 170).
- [171] ATLAS Collaboration, *Athena release 24.0.40*, 2024, URL: <https://gitlab.cern.ch/atlas/athena/-/releases/release/24.0.40> (cit. on p. 170).
- [172] ATLAS Collaboration, *Trigger Core Software Public Results*, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TriggerCoreSWPublicResults> (cit. on p. 171).
- [173] E. F. Codd, *A relational model of data for large shared data banks*, *Commun. ACM* **13** (1970) 377, ISSN: 0001-0782, URL: <https://doi.org/10.1145/362384.362685> (cit. on p. 173).

- [174] D. Alvarez Feito et al., *Studies of IBL wire bonds operation in a ATLAS-like magnetic field.*, tech. rep., CERN, 2015, URL: <https://cds.cern.ch/record/2010249> (cit. on p. 176).
- [175] S. van der Meer, *Calibration of the effective beam height in the ISR*, tech. rep., CERN, 1968, URL: <https://cds.cern.ch/record/296752> (cit. on p. 185).
- [176] ATLAS Collaboration, *Measurement of the W-boson mass and width with the ATLAS detector using proton-proton collisions at $\sqrt{s} = 7$ TeV*, *The European Physical Journal C* **84** (2024), ISSN: 1434-6052, URL: <http://dx.doi.org/10.1140/epjc/s10052-024-13190-x> (cit. on p. 185).
- [177] *Trigger menu in 2018*, tech. rep., CERN, 2019, URL: <https://cds.cern.ch/record/2693402> (cit. on p. 186).
- [178] H. Baer et al., *Prospects for supersymmetry at high luminosity LHC*, 2025, arXiv: 2502.10879 [hep-ph], URL: <https://arxiv.org/abs/2502.10879> (cit. on p. 196).

Nomenclature

pp	proton-proton
Adam	Adaptive Moment Estimation
AFP	ATLAS Forward Proton
ALFA	Absolute Luminosity For ATLAS
ATLAS	A Toroidal LHC ApparatuS
BLS	B -physics and Light States
BSM	Beyond the Standard Model
CB	ComBined
CERN	Conseil Européen pour la Recherche Nucléaire
CL	Confidence Level
CPS	Coherent Prescale Sets
DQ	Data Quality
EB	Enhanced Bias
EM	Electromagnetic
EWSB	Electroweak Symmetry Breaking
FCal	Forward Calorimeter
FEX	Feature EXtraction
FFTV	Fixed-Frequency Trigger Veto
FPGA	Field Programmable Gate Array
GPU	Graphics Processing Unit

GRL	Good Runs List
GUT	Grand Unification Theory
HAD	Hadronic
HEC	Hadronic Endcap Calorimeter
HI	Heavy-ions
HL	High-Luminosity
HLT	High-Level-Trigger
IBL	Insertable <i>B</i> -Layer
ID	Inner Detector
IO	Inside-Out
IP	Interaction Point
JER	Jet Energy Resolution
JES	Jet Energy Scale
JVT	Jet Vertex Tagger variable
L1	Level-1
LAr	Liquid Argon
LB	Luminosity Block
LH	left-handed
LHC	Large Hadron Collider
LSP	Lightest Supersymmetric Particle
LUCID	LUminosity measurement using Cherenkov Integrating Detector
MBTS	Minimum Bias Trigger Scintillators
MC	Monte Carlo
MDT	Monitored Drift Tube
ME	MS-Extrapolated
MLE	Maximum Likelihood Estimate
MS	Muon Spectrometer
MSSM	Minimal Supersymmetric Standard Model

NF	Normalisation Factor
NLSP	Next-to-lightest SUSY particle
NP	Nuisance Parameter
NSW	New Small Wheel
PEB	Partial Event Building
Pflow	Particle-flow
pMSSM	phenomenological Minimal Supersymmetric Standard Model
POI	Parameter Of Interest
PS	PreScale
PV	Primary Vertex
QCD	Quantum Chromodynamics
ReLU	Rectified Linear Unit
RF	Radio Frequency
RH	right-handed
RPC	R-Parity Conservation
SCT	Semi-Conductor Tracker
SM	Standard Model
ST	Segment-tagged
SUSY	SUperSYmmetry
TAP	Trigger After Prescale
TAV	Trigger After Veto
TBP	Trigger Before Prescale
TDAQ	Trigger and Data AcQuisition
TGC	Thin Gap Chamber
TLA	Trigger-Level Analysis
TOB	Trigger OBject
TRT	Transition Radiation Tracker
TT	Trigger Tower

TTWeb TriggerToolWeb

vdM van der Meer

VEV Vacuum Expectation Value

WIMP Weakly Interacting Massive Particle

WP Working Point

ZDC Zero-Degree Calorimeter

NN Neural Network

SHAP SHapley Additive exPlanations

List of Figures

1.1	Fundamental particles of the Standard Model [9].	11
1.2	Summary of measurements of $\alpha_s(Q^2)$ [11].	14
1.3	Shape of the potential $V(\Phi)$ in two dimensions as a function of the field components $\text{Re}(\Phi)$ and $\text{Im}(\Phi)$. The plot on the left represents the case where $\mu^2 > 0$, while the one on the right shows the case where $\mu^2 < 0$ (the so-called <i>Mexican Hat</i> potential).	17
1.4	One-loop quantum corrections to the Higgs squared mass parameter m_H^2 , due to a Dirac fermion f (left), and a scalar S (right) [24].	21
1.5	Evolution of the coupling constants in the SM (dashed lines) and in supersymmetric models (solid lines) [24]. The masses of the SUSY particles are treated as a common threshold and varied between 750 GeV and 2.5 TeV.	22
1.6	Proton decay $p \rightarrow e^+ \pi^0$ through λ' through λ'' R-parity couplings [24].	26
1.7	Pair production cross section from strong-produced SUSY particles [28] at a centre-of-mass energy of 13 TeV (left) and 13.6 TeV (right), which correspond to the LHC pp collisions centre of mass energy in the Run 2 and in the Run 3 data-taking periods, respectively.	31
1.8	Ratio between the $\tilde{t}_1$ pair production cross section at 13.6 TeV (Run 3 centre-of-mass energy) and at 13 TeV (Run 2 centre-of-mass energy) [29].	32
1.9	Tree-level diagrams contributing to the quark-antiquark annihilation production channel (first row) and gluon-gluon fusion channel (second row) [31].	33
1.10	Stop decay modes allowed in different kinematic regions [32].	33
1.11	Feynman diagram representing the 2-body stop decay [33].	34
1.12	Expected and observed 95% CL excluded regions in the $m(\tilde{t}_1)$ - $m(\tilde{\chi}_1^0)$ plane for $\tilde{t}_1$ -pair production, assuming different $\tilde{t}_1$ decay modes with branching ratio of 100%. Signal models that lie within the contours are excluded. Uncertainty bands corresponding to the $\pm 1\sigma$ variation of the combined expected limit are also indicated. The diagonal grey dashed lines indicate the kinematic threshold of the stop decay modes. Exclusion limits shown in the plot on the left [34], the statistical combination of the semileptonic [35] and the fully-hadronic [36] analyses is shown in the plot on the right.	35

- 1.13 Expected and observed 95% CL excluded regions in the $m(\tilde{t}_1)$ - $m(\tilde{\chi}_1^0)$ plane [37] (left) and in the $m(\tilde{t}_1)$ - $\Delta(\tilde{t}_1 - \tilde{\chi}_1^0)$ plane [38] (right) for $\tilde{t}_1$ -pair production, assuming different $\tilde{t}_1$ decay modes with branching ratio of 100%. Signal models that lie within the contours are excluded. Uncertainty bands corresponding to the $\pm 1\sigma$ variation of the combined expected limit are also indicated. 36
- 1.15 ATLAS exclusion limits at 95% CL based on 13 TeV data in the $(\tilde{g}, \tilde{\chi}_1^0)$ mass plane for different simplified models assuming different decay chains with 100% BR from the gluino to the LSP (lightest neutralino or gravitino). 38
- 1.16 Summary plot comparing exclusion limits on chargino and neutralino masses in the $(m(\tilde{\chi}_1^\pm), \Delta m(\tilde{\chi}_1^\pm, \tilde{\chi}_1^0))$ plane, with analyses featuring final states with disappearing tracks, a soft isolated tracks, and soft opposite-sign electron pairs. 39
- 2.1 Long term LHC schedule, as of September 2024 (most updated information at the time of the writing) [41]. The long term schedule of the LHC is shown, starting from 2021, before the start of Run 3, until 2041, when the end of Run 5 is planned. 42
- 2.2 Layout of the CERN accelerator complex [42]. 43
- 2.3 Example of LHC operational cycle [52]. The intensity evolution of Beam 1 (blue) and 2 (red) and the beam energy (dashed green) are shown. The vertical lines denote the beam modes changes. 44
- 2.4 Peak instantaneous luminosity delivered to ATLAS in Run 2, in the data-taking year: 2015 (upper left figure), 2016 (upper right figure), 2017 (bottom left figure), and 2018 (bottom right figure) during stable beams for pp collisions at a centre-of-mass energy of 13 TeV as a function of time [55]. 46
- 2.5 Peak instantaneous luminosity delivered to ATLAS in Run 3, in the data-taking year: 2022 (upper left figure), 2023 (upper right figure), 2024 (bottom figure) during stable beams for pp collisions at a centre-of-mass energy of 13.6 TeV as a function of time [56]. 47
- 2.6 Cumulative luminosity versus time delivered to (green) and recorded by ATLAS (yellow) during stable beams for pp collisions at 13 TeV centre-of-mass energy in Run 2 (left figure) for pp collisions at 13.6 TeV centre-of-mass energy in Run 3 (right figure) [55, 56]. 48
- 2.7 Luminosity-weighted distribution of the mean number of interactions per crossing (μ) for Run 2 pp collision data at 13 TeV centre-of-mass energy (left), and Run 3 pp collision data at 13.6 TeV centre-of-mass energy (right). All data delivered to ATLAS during stable beams is shown, and the integrated luminosity and the mean μ value is given in the figures. The mean number of interactions per crossing corresponds to the mean of the Poisson distribution of the number of interactions per crossing calculated for each bunch [55, 56]. 49
- 2.8 Pileup and β^* evolution with β^* levelling during the LHC fill 8387 (November 2022) [57]. The target pile-up was $\mu = 54 \pm 2.5\%$. The final $\beta^* = 30\text{cm}$ was reached after 5.2 hours of levelling. 50
- 2.9 Cut-away view of the ATLAS detector in the Run 3 configuration [58]. 52

- 2.10 Neutral particles escape the tracker undetected. Electrons and photons release energy in the EM calorimeter, while the HAD calorimeter is designed to measure the energy of the hadrons. Muons are detected in the MS. Weakly interacting particles, such as neutrinos or BSM particles, escape all the ATLAS sub-detectors, but their presence can be deduced by a measurement of the missing transverse momentum [59]. Chapter 3 describes in detail how physics objects are reconstructed in ATLAS. 53
- 2.11 Polar coordinates used to describe the position of a particle in the ATLAS detector (left), and relation between the pseudorapidity η and the standard polar angle θ (right) [60]. 53
- 2.12 Schematic illustration of the ATLAS magnet system. 54
- 2.13 Cut-away view of the ATLAS inner detector in the Run 3 configuration [58]. 55
- 2.14 Barrel ID cut view in the transverse plane [58]. 56
- 2.15 Cut-away view of the ATLAS calorimeter system in the Run 3 configuration [58] 58
- 2.16 Layout of the three accordion layers in a barrel module of the LAr EM calorimeter [74]. 60
- 2.17 Trigger granularity from each 0.1×0.1 TT after the Phase-I upgrade of the LAr calorimeter electronics. Ten E_T values are provided from $1 - 4 - 4 - 1$ longitudinal/transverse samples, each forming a SuperCell. Layer 0 corresponds to the Presampler, Layer 1 corresponds to Strips, Layer 2 corresponds to Middle and Layer 3 corresponds to Back [4]. 61
- 2.18 Cut-away view of ATLAS Muon Spectrometer in the Run 3 configuration [58]. 62
- 2.19 Simplified layout of the ATLAS TDAQ hardware in Run 3 and the flow of data through the system [58]. Rates reported on the right refer to fully-built events (see Chapter 7), excluding calibration data and streams saving only part of the event information. 65
- 2.20 The plots compare the performance of the legacy single electron trigger L1_EM22VHI (red) and the corresponding Phase-I L1_eEM26M (blue) in data recorded by ATLAS in June 2023. On the left, the rates as a function of the instantaneous luminosity are shown. The Phase-I system improves the rate of the single electron trigger, and the ratio between the rates is approximately 80%, independent of the luminosity. On the right, trigger efficiencies as a function of the electron p_T are presented, measured using electrons from $Z \rightarrow ee$ decays [87]. 66
- 2.21 Level-1 trigger rate of the primary single muon trigger (14 GeV p_T -threshold), as a function of date. The rate for each run is normalised to an instantaneous luminosity of $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. On 17 April 2024, coincidences between the TGC in the endcap middle wheel, the Tile calorimeter, and the NSW detectors have been activated to suppress background, reducing the rate by 10 kHz. Further sectors were included on 8 May 2024, reducing the rate by 3 kHz [88]. 66
- 3.1 Illustration of the relationship between the track parameters and the associated track. In this scenario, the hard scatter primary vertex is located at $(e_x, e_y, e_z) = (0, 0, 0)$ [89]. 70

- 3.2 A flow chart of the particle flow algorithm [95], starting with track selection and continuing until the energy associated with the selected tracks has been removed from the calorimeter. At the end, charged particles, topo-clusters which have not been modified by the algorithm, and remnants of topo-clusters which have had part of their energy removed remain. 73
- 3.3 Algorithm flow diagram for the electron and photon reconstruction [96]. 75
- 3.4 Diagram of the superclustering algorithm for electrons and photons. Seed clusters are shown in red, satellite clusters in blue [96]. 76
- 3.5 Electron identification efficiency for the three WPs (tight, medium, loose) for $Z \rightarrow ee$ events in 139 fb^{-1} of 2015-2018 Run 2 LHC data at 13 TeV as a function of the electron transverse energy E_T integrated over the full pseudo-rapidity range (left) and as a function of the electron pseudorapidity η (right). The error bars and uncertainty bands include the statistical and systematic uncertainties. The middle panel shows the ratios of the efficiencies measured in data over those in MC simulation. The lower panel shows the relative size of the total uncertainties for the combined efficiency measurements [97]. 78
- 3.6 Electron identification efficiency for the three WPs (tight, medium, loose, with an extra B-Layer requirement on the latter) for $Z \rightarrow ee$ events in 3.4 fb^{-1} of early Run 3 LHC data at 13.6 TeV as a function of the electron transverse momentum p_T integrated over the full pseudo-rapidity range (left) and as a function of the electron pseudorapidity η (right). The error bars and uncertainty bands include the statistical and systematic uncertainties. The lower panel shows the ratios of the efficiencies measured in data over those in MC simulation [97]. 79
- 3.7 Illustration of the jet formation process, beginning with an initiating parton (quark or gluon), which hadronises to form particle jets detectable by the ATLAS ID and calorimeters [98]. 82
- 3.8 JES calibration steps [103]. 84
- 3.9 Fractional JES systematic uncertainty components for anti- k_t small-R jets reconstructed from Pflow objects, as a function the jet p_T at $\eta = 0$ (left), and as a function of the jet η at $p_T = 60 \text{ GeV}$. The total uncertainty (shown as a filled region topped by a solid black) is determined as the quadrature sum of all components [103]. 86
- 3.10 Illustration of the GN2 algorithm. Before being passed through the model, the features of the jets are concatenated with the features of their associated tracks, allowing them to be processed simultaneously [104]. 87
- 3.11 The c -jet and light-jet rejections as a function of the b -jet tagging efficiency for jets in a $t\bar{t}$ with a p_T selection of $20 < p_T < 250 \text{ GeV}$. The ratio with respect to the performance of the DL1d algorithm is shown in the bottom panels. The shaded regions represent 68% confidence intervals calculated according to a binomial distribution [108]. 88

- 3.12 Background rejection versus signal efficiency in simulated $Z \rightarrow ee$ and $ZZ \rightarrow eev\nu$ samples with a $Z \rightarrow ee$ selection and $E_T^{\text{miss}} > 50 \text{ GeV}$. The performance is shown for E_T^{miss} , event-based E_T^{miss} significance, and object-based E_T^{miss} significance as discriminants in events with two jets (left), and three or more jets (right). The lower panel of the figures shows the ratio of other definitions/event-based E_T^{miss} significance [110]. 91
- 4.1 Feynman diagram representing the 2-body stop decay, whit both top quarks decaying leptonically [111]. 96
- 4.2 Leading lepton p_T distribution in a phase-space region enriched with fake and non-prompt leptons. The plot on the left shows the result for Run 2 samples, while the plot on the right shows the results for Run 3 samples. Only statistical uncertainties are shown. 103
- 4.3 Generated signal points in the considered simplified model's parameter space $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$. Blue points represent models with 2-body decay of the stop, while red points represent models with 3-body decay of the stop. 104
- 4.4 Particle-level distribution of the leading lepton p_T (left), E_T^{miss} (centre) and the angular distance ΔR between the leading lepton and the leading b -jet (right) in stop samples with different $\tilde{t}_R$ and $\tilde{t}_L$ mixing, produced forcing the decay chain $\tilde{t}_1 \rightarrow t\tilde{\chi}_1^0$. 104
- 5.1 Distribution of E_T^{miss} (left) and the E_T^{miss} significance (right) in the Run 2 dataset. The shaded area represents the statistical uncertainty on the total SM background. The signal normalisation is increased 250 times for visualisation purposes. 109
- 5.2 Distribution of the invariant mass $m_{\ell\ell}$ of the leading and the sub-leading lepton (left), and distribution of the transverse mass between the E_T^{miss} and the leading lepton ℓ_1 (right) in the Run 2 dataset. The shaded area represents the statistical uncertainty on the total SM background. The signal normalisation is increased 250 times for visualisation purposes. 110
- 5.3 Diagram of the generic process under consideration [132]. 110
- 5.4 Distribution in Run 2 dataset of the stransverse mass $m_{T2}^{\ell\ell}$ using the two leading leptons as visible particles (upper row), and distribution of $m_{T2,\text{min}}$ (bottom row) computed using the leading lepton (left), the sub-leading lepton (centre), and both leptons (right) in association with the two b -jets. The shaded area represents the statistical uncertainty on the total SM background. The signal normalisation is increased 250 times for visualisation purposes. 113
- 5.5 Distribution of $\mathbf{p}_{T,\text{boost}}^{\ell\ell}$ (left) and the $\Delta\phi_{\text{boost}}$ (right) in the Run 2 dataset. The shaded area represents the statistical uncertainty on the total SM background. The signal normalisation is increased 250 times for visualisation purposes. 114
- 5.6 Distribution of the angular distance ΔR between the two leading leptons (left) and the two leading b -jets, or the leading b -jet and the leading light-jet in events with only one b -jet, (right) in the Run 2 dataset. The shaded area represents the statistical uncertainty on the total SM background. The signal normalisation is increased 250 times for visualisation purposes. 114

- 5.7 Illustration of the computation performed inside a node of a neural (left) and three different types of layers (input, hidden, output) in a classical neural network. 116
- 5.8 NN signal score distributions in Run 2 MC samples for all the NN models. 124
- 5.9 NN signal score distributions in Run 3 MC samples for all the NN models. 124
- 5.10 SHAP absolute values, ranking the input features' importance for each of the six models. For all models, *class 0* represents the signal class and *class 2* represents the $t\bar{t}$ class. For the low- Δm models (first line), *class 1* represents the $t\bar{t}Z$ class; for the medium- and high- Δm models (second and third lines, respectively), *class 1* represents the $t\bar{t}Z \rightarrow \nu\nu$ class, and *class 3* represents the $t\bar{t}Z \rightarrow \ell\ell$ class. Plots for the 1b and the 2b models are shown in left and right columns, respectively. 125
- 5.11 Illustration of analysis regions based on NN scores. 126
- 5.12 NN $t\bar{t}$ score data/MC distributions for all NN models evaluated on Run 2 samples. Data in SRs are blinded. 128
- 5.13 NN $t\bar{t}$ score data/MC distributions for all NN models evaluated on Run 3 samples. Data in SRs are blinded. 129
- 5.14 Distribution of the closest $m_{\ell\ell}$ (left) and of $E_{T,\text{corr}}^{\text{miss}}$ (right) in $\text{CR}_{t\bar{t}Z}$ in Run 2 (upper row) Run 3 (bottom row) samples. 134
- 6.1 The standard Gaussian distribution $\phi(x) = (1/\sqrt{2\pi}) \cdot \exp(-x^2/2)$, showing the relation between the significance Z and the p -value [139]. 141
- 6.2 Illustration of a p -value obtained from an observed value $t_{\mu,\text{obs}}$ of the test statistic t_μ [139]. 142
- 6.3 Expected yields in the three non-orthogonal Δm categories for both the Run 2 (left column) and the Run 3 (right column) datasets. The upper panels show the comparison between the observed number of data events and the post-fit SM prediction of the background in the CRs and VRs (SRs are blinded). For CRs and VRs, the bottom panel shows the difference between data and the predicted SM background divided by the total uncertainty (σ_{tot}) 149
- 6.4 Exclusion limit contour (95% CL) for a simplified model assuming $\tilde{t}_1$ pair production, decaying via $\tilde{t}_1 \rightarrow t\tilde{\chi}_1^0$ with 100% branching ratio, in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane, considering only the Run 2 (left) and only the Run 3 (right) dataset. The dashed lines and the shaded yellow bands represent the expected limits and their 1σ uncertainties. The red solid lines represent the observed limits. The limits are also shown in the stop three-body decay region. The grey shaded region shows the limits from the previous Run 2 analysis [111]. 152
- 6.5 Exclusion limit contour (95% CL) for a simplified model assuming $\tilde{t}_1$ pair production, decaying via $\tilde{t}_1 \rightarrow t\tilde{\chi}_1^0$ with 100% branching ratio, in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane, considering the combined Run 2 and Run 3 datasets. The dashed line and the shaded yellow band represent the expected limits and their 1σ uncertainties. The red solid line represents the observed limits. The limits are also shown in the stop three-body decay region. The grey shaded region shows the limits from the previous Run 2 analysis [111]. 152

- 6.6 Red, blue, and green points are the points for which the low, medium, and high- Δm models, respectively, give the best expected CL. The study is shown for Run 2 only (left), Run 3 only (centre) and combining Run 2 and Run 3 (right). 153
- 6.7 Example tree-level Feynman diagrams for the $pp \rightarrow t\bar{t}H$ production process. 154
- 6.8 Cross sections for the SM Higgs boson production modes as a function of the centre-of-mass energy of pp collisions assuming a Higgs boson mass of 125 GeV (left), and branching ratio for the main decay modes of the SM Higgs boson as a function of its mass (right) [159]. 155
- 6.9 Distribution of $m_{\ell\ell}^{\text{miss}}$ and E_T^{miss} significance in Run 2 MC samples comparing $t\bar{t}$, $t\bar{t}H$, and two stop signal points with $\Delta m = 200$ GeV (low- Δm range) and $\Delta m = 400$ GeV (medium- Δm range). 156
- 6.10 Expected upper CL_s exclusion limits on $\mathcal{B}_{H \rightarrow \text{inv}}$ at 95% CL for the Run 2 only dataset (upper left), the Run 3 only dataset (upper right) only dataset, and combining the Run 2 and the Run 3 results (bottom). Both statistical and systematic uncertainties are included. 156
- 7.1 Trigger efficiency curve for an ideal trigger (step-function, in blue) and a real trigger (turn-on function, in orange). The represented efficiency curves refer to a trigger with a p_T threshold of 30 GeV. 161
- 7.2 Turn-on curves of the lowest-threshold single electron trigger [160] (left) and single muon triggers [161] (right), used to collect data in 2022. The plot on the also shows the efficiency of the L1 seed. Efficiencies are shown as a function of the offline electron transverse energy (E_T) and data compared against SM $Z \rightarrow ee$ and $Z \rightarrow \mu\mu$ MC events for the electron trigger and the muon trigger, respectively. In both cases, the tag-and-probe method is used to measure the trigger efficiency. 162
- 7.3 Trigger signal efficiency for the signal samples used in the analysis presented in this thesis in the $(m_{\ell_1}, m_{\chi_1^0})$ plane. 163
- 7.4 Example of control flow for menu processing. The control flow graph is created during initialisation and the steps are executed based on the available data. If a filter passes, processing continues through the next steps until all filters fail where the processing of the following steps is stopped. If the last step is reached with a chain passing all of its steps the event will be accepted. Figure adapted from Reference [163]. 164
- 7.5 Pull distribution of the rate prediction of 957 HLT chains based on an EB dataset, normalised to the combined statistical error (σ), and fitted by a Gaussian function with mean μ_{Fit} and width σ_{Fit} , where $R_{\text{Prediction}}$ and R_{Online} refer to the rates from prediction and data-taking for the same luminosity, respectively. Error bars include statistical uncertainties only [169]. 169
- 7.6 Application throughput in events/s (left column) and memory usage in GB (right column) estimated using 2022 data (upper row) [4] and 2024 data (bottom row) [172] as a function of the number of events processed in parallel. Blue squares represent a pure multiprocessing approach, pink circles represent a pure multithreading approach, while the other curves represent hybrid approaches. 171

7.7	Example bunch group configurations for the first four out of the 16 possible bunch groups. The numbers in blue on the right indicate the number of bunch crossings for each group [52].	174
8.1	Results of the high-rate and high- μ test in October 2024. During the test, the pile-up value in ATLAS was increased from the 2024 nominal value of 64 to 69, in 1-unit steps every 10-30 minutes. During each step, the L1 trigger rate was varied from about 94 kHz to 100 kHz by increasing the low- p_T jets trigger rate. The higher number of collisions happening when increasing the pile-up value causes a heavier CPU load for the HLT reconstruction, observed as a decrease in the number of idle logical cores. The HLT output bandwidth increases as a function of μ , scaling both with the rate of accepted events in all streams, and with the increased event size due to the higher occupancy. The dead time increases with the L1 rate from about 4.5% to 8%, mainly due to the complex dead time (see Section 2.2.6) [5].	177
8.2	Output rate (left) and bandwidth (right) for the HLT streams as a function of time in a pp LHC fill taken in May 2024, with a peak instantaneous luminosity of $2.1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and a peak average number of interactions per crossing of $\langle \mu \rangle = 63$. The luminosity is levelled for a period at the beginning of the fill (see Section 2.1.3); after that, the output bandwidth decreases with the decaying luminosity [5].	179
8.3	Import hierarchy for pp menus.	179
8.4	Example of rate given to each signature group in the Main stream in 2022, taken from Reference [4]. This chart is based on 2022 data recorded at an instantaneous luminosity of $1.8 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. The label “UT” stands for <i>unconventional tracking</i> signatures, typical of, e.g., long-lived particles.	181
8.5	Total L1 trigger rate as a function of time in a pp LHC fill taken in May 2024, with a peak luminosity of $L = 2.1 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ and a peak average pile-up of $\mu = 63$, with the LHC instantaneous luminosity overlaid. The luminosity is levelled for a period at the beginning of the fill (see β^* levelling in Section 2.1.3); after that, the rate decreases with the decaying luminosity. Periodic increases in the L1 rate with respect to luminosity at the end of the fill are caused by prescale changes, as the luminosity and corresponding overall resource usage decline [5].	184
8.6	Diagram showing the importing structure of the Rulebook code.	187
8.7	Simplified flow chart of the PS calculation steps.	191
8.8	Screenshot of the ATLAS TriggerPanel taken in June 2025. The names of the PS sets contain relevant information on their content.	193
8.9	Example of an entry from a L1 (left) and an HLT (right) prescale JSON files. For the HLT file, the example comes from an entry of a chain also writing to the express stream. For chains not writing to the express stream, the last two lines are absent.	194
A.1	NPs of the background-only low- Δm fit for Run 2 (left), and Run 3 (right).	202
A.2	NPs of the background-only medium- Δm fit for Run 2 (left), and Run 3 (right).	203
A.3	NPs of the background-only high- Δm fit for Run 2 (left), and Run 3 (right).	204

-
- A.4 Correlation matrix for the NPs of the background-only low- Δm fit for Run 2 (left), and Run 3 (right). 205
- A.5 Correlation matrix for the NPs of the background-only medium- Δm fit for Run 2 (left), and Run 3 (right). 205
- A.6 Correlation matrix for the NPs of the background-only high- Δm fit for Run 2 (left), and Run 3 (right). 205

List of Tables

1.1	Doublet and singlet structure for the left- and right-handed components of the fermion fields under the $SU(2)_L \otimes U(1)_Y$ symmetry group.	15
1.2	Chiral supermultiplets in the MSSM.	24
1.3	Gauge supermultiplets in the MSSM.	25
1.4	MSSM particles, with sfermion mixing for the first two generations assumed to be negligible. In this table gravity (and thus its mediator, the <i>gravitino</i> , is not considered).	29
2.1	Typical parameters values at the LHC for pp collisions in Run 2 [53] and for Run 3 [54], shown together with the design values.	45
4.1	Triggers used to select Run 2 data	99
4.2	Triggers used to select Run 3 data.	99
4.3	Simulated background samples with the corresponding matrix element and parton shower (PS) generators, cross section order in α_s used to normalise the event yield, the generator PDF sets used. When different, the settings used for the simulation of Run 3 samples are mentioned in parentheses. Abbreviations used are defined as: leading-order (LO), next-to-leading-order (NLO), next-to-next-to-leading-order (NNLO), next-to-leading-logarithmic (NLL), next-to-next-to-leading-logarithmic (NNLL).	102
5.1	Preselection applied before training to the six NN models.	121
5.2	Input variables used for training the NN in all the Δm regions, only for models requiring more than one b -tagged jet. For models training on events with only one b -tagged jet, the information of the leading light-jet is used instead of the information on the sub-leading b -jet to define the listed variables.	121
5.3	Hyperparameters values chosen for the six NN models.	123
5.4	SR binning in NN signal score. The different number of bins for different Δm regions is due to the different shape of the NN output distributions.	128
5.5	Selections on NN signal and $t\bar{t}$ scores defining the $t\bar{t}$ CRs.	129

5.6	MC expected yields in $CR_{t\bar{t}}$ for the low- Δm region, in Run 2 (left) and Run 3 (right) datasets. The error includes statistical and systematic uncertainties (described in Section 6.2).	130
5.7	MC expected yields in $CR_{t\bar{t}}$ for the medium- Δm region, in Run 2 (left) and Run 3 (right) datasets. The error includes statistical and systematic uncertainties (described in Section 6.2).	130
5.8	MC expected yields in $CR_{t\bar{t}}$ for the high- Δm region, in Run 2 (left) and Run 3 (right) datasets. The error includes statistical and systematic uncertainties (described in Section 6.2).	130
5.9	Selections applied to the NN signal scores to define the $t\bar{t}$ VRs.	131
5.10	MC expected yields in $VR_{t\bar{t}}$ for the low- Δm region, in Run 2 (left) and Run 3 (right) datasets. The error includes statistical and systematic uncertainties (described in Section 6.2).	131
5.11	MC expected yields in $VR_{t\bar{t}}$ for the medium- Δm region, in Run 2 (left) and Run 3 (right) datasets. The error includes statistical and systematic uncertainties (described in Section 6.2).	131
5.12	MC expected yields in $VR_{t\bar{t}}$ for the high- Δm region, in Run 2 (left) and Run 3 (right) datasets. The error includes statistical and systematic uncertainties (described in Section 6.2).	132
5.13	Selection criteria of the control region for the $t\bar{t}Z \rightarrow \nu\nu$ background ($CR_{t\bar{t}Z}$).	133
5.14	MC expected yields in $CR_{t\bar{t}Z}$ for the Run 2 MC samples. The error includes statistical and systematic uncertainties (described in Section 6.2).	133
5.15	MC expected yields in $CR_{t\bar{t}Z}$ for the Run 3 MC samples. The error includes statistical and systematic uncertainties (described in Section 6.2).	133
5.16	Relative contributions of the main components to the “other” background in the $CR_{t\bar{t}Z}$ control region.	135
5.17	Selection criteria of the validation region for the $t\bar{t}Z \rightarrow \nu\nu$ background ($VR_{t\bar{t}Z}$).	135
5.18	MC expected yields in $VR_{t\bar{t}Z}$ for the Run 2 MC samples. The error includes statistical and systematic uncertainties (described in Section 6.2).	136
5.19	MC expected yields in $VR_{t\bar{t}Z}$ for the Run 3 MC samples. The error includes statistical and systematic uncertainties (described in Section 6.2).	136
6.1	Sources of systematic uncertainty in the SM background estimates, after the background fits. The values are given as relative uncertainties in the total expected background event yields in the different bins of the Run 2 low- Δm SR.	145
6.2	Sources of systematic uncertainty in the SM background estimates, after the background fits. The values are given as relative uncertainties in the total expected background event yields in the different bins of the Run 3 low- Δm SR.	145
6.3	Sources of systematic uncertainty in the SM background estimates, after the background fits. The values are given as relative uncertainties in the total expected background event yields in the different bins of the Run 2 medium- Δm SR.	146
6.4	Sources of systematic uncertainty in the SM background estimates, after the background fits. The values are given as relative uncertainties in the total expected background event yields in the different bins of the Run 3 medium- Δm SR.	146

6.5	Sources of systematic uncertainty in the SM background estimates, after the background fits. The values are given as relative uncertainties in the total expected background event yields in the different bins of the Run 2 high- Δm SR.	147
6.6	Sources of systematic uncertainty in the SM background estimates, after the background fits. The values are given as relative uncertainties in the total expected background event yields in the different bins of the Run 3 high- Δm SR.	147
6.7	NFs for $t\bar{t}$ and $t\bar{t}Z$ backgrounds obtained from the Run 2 (left) and Run 3 (right) background-only fits.	150
7.1	ATLAS physics goals, with key processes and relative trigger signatures.	160
8.1	Summary of different ways of computing the expected rate value used in the prescale computation when the rule defines a target rate. The three listed cases are: expected rate for a L1 item, for which a rule defines a target L1 rate; expected rate for an HLT chain for which the rule defines a target HLT rate, and expected rate for an HLT chain for which the rule defines a target L1 seed rate.	189

Acknowledgments

And so this is the end! If I've made it this far, it's thanks to the many people who have walked alongside me, and, at times, pulled me forward, throughout this journey.

I'm first of all grateful to my supervisors. Thank you Brian for your daily guidance and support. Thank you for everything you've taught me, and for the many opportunities you gave me. Your door was always open when I had questions or just needed some encouragement. A big thanks to Anna, for always being available for help and suggestions. Conversations with you are always enlightening, and your enthusiasm is an inspiration. I also want to thank Mark. You left CERN before this thesis was done, but you accompanied me in my first steps in the trigger world, passing on the enthusiasm that has stayed with me ever since. Finally, I'll always be deeply grateful to Leonardo. I wasn't even sure I wanted a Ph.D. when I started my Master's thesis, but you showed me how exciting and rewarding research can be and convinced me I had what it takes: I wouldn't have even started this without you.

I was so lucky to have the chance to pursue my Ph.D. within the SMARTHEP network, which gave me many opportunities and resources I am truly grateful for. A special thanks to all the SMARTHEP ECRs, particularly Leon, my fellow Geneva ECR and secondment mate. Our (not-so) short journey in the world of finance would have been much less productive without you.

I was so lucky to be part of the CERN ATLAS Trigger team during my Ph.D. Thanks to all the members of the group, for always being ready to give a hand and answer questions. The traditional line closing the presentations at the yearly group meetings, "the doors on the fifth floor are always open for help and discussion", is really true! In particular, I thank Stefanie, for being an endless source of knowledge on everything about the trigger menu and operations; and Marco, for guiding me through the analysis. I am very grateful for everything I learned from you and for the patience you had in answering my many (many) questions. Without all your support, especially during these last crazy months, I would have never pushed the analysis to this point: this thesis is much more interesting thanks to you! A big thanks also to John and Michael, who often got stuck with my SUSY questions over the years, and to my office mates, Aleksandra, Eleni and Viktoriia, for all the days spent together in the "Students' Office" sharing good and bad moments of student life.

I was so lucky to have many great people supervising my work and guiding me during this journey, but I was as lucky to have the opportunity to be a supervisor myself. So thanks a lot to

Clarisse: I hope you learned from me as much as I learned from you.

I also want to thank the Geneva group. Even if I didn't spend much time at the university, I always felt welcome and included, and I hope to cross my path with many of you again!

A big part of my Ph.D. work was in the Trigger Menu and Operations group. I met so many great people there, and I hope I'll still have the chance to work with many of you. In particular, I want to thank TJ, Antonia, Stefanie, Takuya, Lidija and Claudia, for their patience in guiding me through my first steps in trigger operations. And especially thank you to Sara, not only for teaching me everything I know during my shadows (and even afterwards), but for being someone I feel I can trust and rely on completely, inside and outside of trigger and physics.

I was very happy to carry on my analysis in the SUSY group, where I always felt welcome and supported. In particular, a big thanks to all the Run 3 stop team: it was a pleasure to work with all of you, and I'm excited to continue! And especially thanks to Marco and Claudia (again!), and Jonas, for their continuous support and patience in keeping everything on track. To Claudia, I was so lucky to work with you both in trigger and in SUSY. Thanks for everything I learned from your explanations, but even more importantly for your help and support when it was most needed.

Finally, I want to thank my friends and family.

I'll start with my CERN friends, who lived this Ph.D. with me day by day: my life would have been much more difficult, and less fun, without all of you. Elena, you were always there for me, to live all the fun and the nights out, but also to share the daily life and struggles, and to cheer me up when I was at my lowest – I will always be grateful for this. Laura and Federica, you welcomed me at CERN: I had my first b40 coffees and OB dinners with you, and we've never stopped since! And thank you Chiara for being with me since the Welcome Day: I'm so grateful to have you when I need an honest opinion (and possum memes). I also want to thank Alberto, and Alberto (Pe. and Ple.), the lunch crew – Mattia, Alessandro, Mathis (et merci encore pour ton aide à la traduction du résumé de cette thèse!) –, Camilla, Chiara M., Leonardo, Simone and everyone who shared at least a small part of this journey with me. A big thank you also to Giacomo, for welcoming me at CERN from day one. It was so nice to see a familiar face when I first moved and everything was new. Thanks also to Giulio, Sabrina, Stefano and all the others, I look forward to watching many more football matches with all of you!

And lastly, to those who make every trip back to Italy truly feel like coming home.

Thank you to my friends in Italy, in particular Roby, Miriam, Silvia, and Elena. Even if we don't talk every day as we once did, I know that you are there for me and as soon as we get together it feels as if no time has passed. And Ilaria, thanks for being my mate in physics struggles, from the Bachelor's to the Ph.D.!

Infine, ringrazio la mia famiglia. Grazie nonna, per avermi sempre sostenuto e cercato di stare dietro alle mie ossessioni, da quando ho cercato di spiegarti il DNA, il latino e l'inglese, e infine i segreti dell'LHC. Thanks to my brother Pietro: I have no doubt you are one of my biggest supporters, ever since ten years ago you explained to me the Stevin's law at midnight for that physics test I had forgotten about. This Ph.D. is also a bit yours.

And I conclude this work by thanking my parents: for always teaching me to have big dreams and to fight for them, and for never making me feel alone while doing so. This thesis is dedicated to you.