
Search for charged-lepton-flavour violation in top-quark production and decay with the ATLAS experiment

Dissertation

zur Erlangung des akademischen Grades
Doktor der Naturwissenschaften
(Dr. rer. nat.)

vorgelegt von

Gabriel Gomes da Silva

Eingereicht bei der Naturwissenschaftlich-Technischen Fakultät
der Universität Siegen

2026

1. Gutachter: Prof. Dr. Markus Cristinziani

2. Gutachterin: PD Dr. Carmen Diez Pardos

Tag der Promotion: 29.05.2026

Abstract

In the Standard Model of particle physics with massless neutrinos, charged lepton flavour is exactly conserved in weak interactions. Although the observation of neutrino oscillations establishes that lepton flavour is not a fundamental symmetry of nature, the corresponding rates for charged lepton flavour violating processes induced by non-zero neutrino masses are suppressed to an unobservable level. Consequently, the detection of charged lepton flavour violation in processes involving charged leptons would constitute unambiguous evidence for physics beyond the Standard Model and provide direct insight into the flavour structure of new interactions.

This thesis presents a direct search for charged lepton flavour violation in top-quark processes using proton–proton collision data recorded by the ATLAS detector at the Large Hadron Collider between 2015 and 2018, corresponding to an integrated luminosity of 140 fb^{-1} at a centre-of-mass energy of $\sqrt{s} = 13 \text{ TeV}$. The analysis targets two complementary production modes: top-quark pair production ($t\bar{t}$), where one top quark decays into an electron–muon pair and an up-type quark ($t \rightarrow e\mu q$), and single top-quark production in association with an electron–muon pair. In the $t\bar{t}$ channel, the event topology is defined by requiring the Standard Model W boson from the second top quark to decay hadronically, thereby enhancing the sensitivity to the charged lepton flavour violating signature; an analogous hadronic requirement is imposed in the single top-quark case to suppress backgrounds.

Signal discrimination is achieved through the implementation of a multivariate classifier based on a boosted decision tree, optimised to exploit the kinematic and topological differences between signal and the dominant Standard Model backgrounds. A binned profile-likelihood fit to the discriminant distribution is performed to extract the signal strength, with systematic uncertainties incorporated as nuisance parameters. No significant deviation from the Standard Model expectation is observed. Upper limits at 95% confidence level are set on the branching fraction $\mathcal{B}(t \rightarrow e\mu q) < 1.12 \times 10^{-7}$, which are translated into constraints on the relevant dimension-6 Wilson coefficients within the Standard Model Effective Field Theory framework. These results significantly improve upon previous bounds and further constrain scenarios predicting flavour-violating interactions in the top-quark sector.

Keywords: charged-lepton-flavour violation, top quark, ATLAS experiment, Standard Model Effective Field Theory

Zusammenfassung

Im Standardmodell der Teilchenphysik mit masselosen Neutrinos ist die Flavourzahl geladener Leptonen in schwachen Wechselwirkungen exakt erhalten. Obwohl die Beobachtung von Neutrinooszillationen zeigt, dass Leptonflavour keine fundamentale Symmetrie der Natur darstellt, sind die entsprechenden Raten für durch nichtverschwindende Neutrinomassen induzierte Prozesse mit Flavourverletzung geladener Leptonen auf ein nicht beobachtbares Niveau unterdrückt. Folglich würde der Nachweis einer Flavourverletzung in Prozessen mit geladenen Leptonen einen eindeutigen Hinweis auf Physik jenseits des Standardmodells liefern und direkten Einblick in die Flavourstruktur neuer Wechselwirkungen ermöglichen.

Diese Dissertation präsentiert eine direkte Suche nach Flavourverletzung geladener Leptonen in Top-Quark-Prozessen unter Verwendung von Proton-Proton-Kollisionsdaten, die zwischen 2015 und 2018 mit dem ATLAS-Detektor am Large Hadron Collider aufgezeichnet wurden und einer integrierten Luminosität von 140 fb^{-1} bei einer Schwerpunktsenergie von $\sqrt{s} = 13 \text{ TeV}$ entsprechen. Die Analyse untersucht zwei komplementäre Produktionsmodi: die Top-Quark-Paarproduktion ($t\bar{t}$), bei der eines der Top-Quarks in ein Elektron-Myon-Paar und ein Up-Typ-Quark zerfällt ($t \rightarrow e\mu q$), sowie die Einzel-Top-Quark-Produktion in Assoziation mit einem Elektron-Myon-Paar. Im $t\bar{t}$ -Kanal wird die Ereignistopologie dadurch definiert, dass das Standardmodell- W -Boson des zweiten Top-Quarks hadronisch zerfallen muss, wodurch die Sensitivität gegenüber der Signatur mit Flavourverletzung geladener Leptonen erhöht wird; eine analoge hadronische Anforderung wird im Einzel-Top-Quark-Fall angewendet, um Untergrundprozesse zu unterdrücken.

Die Signaltrennung erfolgt durch die Implementierung eines multivariaten Klassifizierers auf Basis eines Boosted-Decision-Tree-Algorithmus, der darauf optimiert ist, die kinematischen und topologischen Unterschiede zwischen Signal und den dominanten Standardmodell-Untergründen auszunutzen. Zur Bestimmung der Signalstärke wird eine gebinnte Profile-Likelihood-Anpassung an die Verteilung der Diskriminante durchgeführt, wobei systematische Unsicherheiten als Nuisance-Parameter berücksichtigt werden. Es wurde keine signifikante Abweichung von der Erwartung des Standardmodells beobachtet. Es werden obere Grenzen auf dem 95%-Konfidenzniveau für das Verzweungsverhältnis $\mathcal{B}(t \rightarrow e\mu q) < 1.12 \times 10^{-7}$ gesetzt, die in Einschränkungen der relevanten Wilson-Koeffizienten der Dimension-6 im Rahmen der Standardmodell Effektiven Feldtheorie übersetzt werden. Diese Ergebnisse stellen eine deutliche Verbesserung gegenüber früheren Suchen dar und schränken Szenarien mit flavourverletzenden Wechselwirkungen im Top-Quark-Sektor weiter ein.

Schlüsselwörter: Flavourverletzung geladener Leptonen, Top-Quark, ATLAS-Experiment, Standardmodell Effektiven Feldtheorie

Contents

Abstract	iii
1 Introduction	1
2 Theoretical and Experimental Overview	3
2.1 The Standard Model	3
2.1.1 Constituents and interactions	3
2.1.2 The mathematical formalism	5
2.2 Effective field theory	6
2.3 Charged-lepton-flavour violation	7
2.3.1 Theoretical framework	7
2.3.2 EFT operators for $2\ell 2q$ contact interactions	8
2.3.3 Experimental searches	10
3 The ATLAS Experiment	17
3.1 The Large Hadron Collider	17
3.2 The ATLAS detector	19
3.2.1 The tracking system	21
3.2.2 The calorimeter system	23
3.2.3 The muon system	25
3.2.4 The trigger system	25
4 Common Elements in Top-Physics Analysis	27
4.1 Hadron-collider physics	27
4.2 Physics simulation	29
4.3 Physics object reconstruction and identification	32
4.3.1 Track reconstruction	33
4.3.2 Electrons	34
4.3.3 Muons	37
4.3.4 Jets	39
4.3.5 Missing transverse momentum	41
4.3.6 Overlap removal	41
4.4 Top-quark physics	41
4.5 Statistical analysis	44
4.5.1 The likelihood function	44

4.5.2	The test statistic	46
4.5.3	The discovery case	47
4.5.4	Setting limits	48
5	Search for Charged-Lepton-Flavour Violation in Top-Quark Processes	51
5.1	The signal processes	51
5.2	Background Monte-Carlo simulation samples	54
5.3	Object and event selection	55
5.4	Analysis strategy	58
5.5	Multivariate analysis	61
5.5.1	The boosted decision tree	61
5.5.2	Implementation	66
5.6	Systematic uncertainties	73
5.6.1	Experimental uncertainties	73
5.6.2	Theoretical uncertainties	74
5.6.3	Cross-section uncertainties	76
5.7	Results	76
5.7.1	Statistical analysis framework	76
5.7.2	Fit using Asimov data	81
5.7.3	Unblinded fit	83
5.7.4	Limits on the cLFV process	89
5.7.5	EFT interpretation	91
6	Conclusion and Outlook	97
	Bibliography	99
	List of Figures	113
	List of Tables	115
	List of Acronyms	117
	Acknowledgements	119

Introduction

The endeavour of particle physics in investigating the fundamental constituents of matter can be traced back to the early tradition of philosophical atomism. This philosophy has found its most successful application as a metaphysical thesis in the natural sciences starting in the nineteenth century with the emergence of atomic chemistry and the kinetic theory of gases. Further scientific findings on the structure of matter would force atomic models to evolve. Atoms were no longer considered indivisible: J. J. Thomson's discovery of the electron [1] revealed the existence of particles with masses much smaller than the lightest atom, and Rutherford's scattering experiment [2] proved the atom had substructure, namely, a nucleus.

The enormous development of particle physics as a field of research that took place in the 20th century is unquestionably coupled with the rapid technological developments. Most notably, the increase in the energies accessible by particle accelerators would often be followed by the discovery of new particles. This era witnessed numerous profound achievements: the detection of the positron [3], the antiparticle of the electron; the discovery of the neutron [4], a neutral constituent of the nucleus alongside the positive proton; the theoretical proposition [5, 6] and ensuing experimental evidence [7, 8] in favour of the quark model; in addition to the detection of a multitude of short-lived particles.

The current, best theory that describes the behaviour of electrons, quarks, and the other fundamental particles of the Universe is referred to as the Standard Model (SM) of particle physics. The SM is the culmination of almost a century of theoretical and experimental development and has shown itself to be both fantastically predictive and accurate throughout the years. This is best exemplified by achievements such as the SM prediction for the electron anomalous magnetic moment, which matches experiment to ten significant figures [9, 10], and the discoveries of many particles that were first predicted by the SM, such as the discovery of the W and Z bosons at SPS in the early 1980s, and more recently the discovery of the Higgs boson by the ATLAS [11] and CMS experiments [12] at the LHC in 2012.

The field of high-energy physics, however, aims to tackle the unsolved questions of the Universe, either by directly searching for particles that are not predicted by the SM or by high-precision measurements looking for deviations from SM predictions. For instance, a description of gravity is not included in the SM: it is still an open problem to seamlessly combine a quantum theory like the SM to the current theory of gravity, Einstein's General Relativity. The exclusion of gravity in the SM framework is not its only shortcoming as this becomes evident when the theory clashes with cosmological evidence. A number of astronomical observations in the last hundred years have led to

the hypothesis that the particles described by the SM constitute but a small fraction of the total mass content of the Universe. The current estimate for the fraction of the total mass that is composed of ‘ordinary matter’ is $\sim 16\%$ [13]; the remaining mass is in the form of ‘dark matter’, of which very little is known. Another gap in the model consists of the mystery behind baryogenesis, *i.e.*, the imbalance between matter and antimatter in the observed universe.

A finding that would unambiguously hint at physics beyond the Standard Model (BSM) is the observation of sizeable charged-lepton-flavour violation (cLFV). In the SM with massless neutrinos, the flavour of charged leptons is strictly conserved in weak interactions. However, the observation of neutrino oscillations [14, 15] confirms the existence and the mixing of massive neutrinos, and it also indicates that lepton flavour violation (LFV), meaning a local interaction that alters lepton flavour, is expected to occur in the charged-lepton sector. Although minimal extensions to the SM, required in order to accommodate massive neutrinos, allow for these processes, they are so highly suppressed that it would be impossible to detect them with our current experimental sensitivity. Thus, any such evidence would imply new physics.

Searches for cLFV processes involving many different particles have been going on for decades, nonetheless, it was only more recently with the advent of the LHC that such investigations in the heavy sector gained prominence. Channels involving the top quark, for example, have already shown themselves as a viable playground [16–19]. In this context, the search presented in this work represents a step towards enhancing the sensitivity for these channels.

This thesis is organised as follows. Chapter 2 presents an overview of the SM and the mechanisms by which charged-lepton-flavour violation can occur. It also outlines the theoretical motivation for the analysis and reviews the current status of cLFV searches in the top-quark sector. Chapter 3 describes the ATLAS detector and its main subcomponents. Chapter 4 provides more technical details, including the simulation of physics events, an overview of top-quark physics, and a description of the statistical analysis employed. Finally, Chapter 5 presents in detail the full analysis strategy, including the definition of the signal and control regions, the event selection criteria, and the modelling of the dominant background contributions. Particular emphasis is placed on the signal discrimination technique, where multivariate classifiers are employed to optimise the separation between the charged lepton flavour violating signal and SM backgrounds. The statistical interpretation is then described, outlining the construction of the likelihood function, the treatment of systematic uncertainties as nuisance parameters, and the fitting procedure used to extract the final results. The chapter concludes with the presentation of the fit outcomes and the corresponding constraints on the relevant dimension-6 operators within the Standard Model Effective Field Theory framework.

The work presented in this thesis is currently under review within the ATLAS collaboration.

Theoretical and Experimental Overview

This chapter provides the theoretical background and experimental context underpinning the work presented in this thesis. It begins with an outline of the Standard Model in Section 2.1, followed by a brief introduction to effective field theory in Section 2.2. Finally, Section 2.3 presents a summary of cLFV, starting with general theoretical considerations, then detailing the effective operators relevant to this study, and concluding with a review of existing cLFV searches, with a particular focus on the top quark sector.

2.1 The Standard Model

2.1.1 Constituents and interactions

The SM of particle physics is a relativistic quantum field theory constructed to respect the symmetries of the gauge group $SU(3)_C \otimes SU(2)_W \otimes U(1)_Y$. The subscripts indicate the conserved charges associated with each group: C denotes colour charge, W refers to weak isospin, and Y corresponds to weak hypercharge. This gauge structure forms the basis for a unified description of three of the four known fundamental interactions in nature: namely, the strong, weak, and electromagnetic forces. The fourth fundamental interaction, gravity, remains unincorporated within the SM, as a consistent and complete quantum theory of gravity has not yet been formulated [20].

While quantum fields are the fundamental objects in the theory, their properties are revealed through the study of their excitations: namely, the particles associated with each field. These particles are characterised by their quantum numbers and, when on shell, by a well-defined mass. Since the theory is required to be consistent with special relativity, it must be invariant under Lorentz transformations, which include spatial rotations and Lorentz boosts. The transformation properties of particles under these symmetries are determined by their intrinsic angular momentum, or spin, which leads to a classification of all fundamental particles into two broad categories: fermions and bosons.

According to the SM, the matter content of the Universe consists of twelve fundamental fermions, along with their associated antiparticles, arranged into three generations. In addition to these matter particles, the model includes four gauge bosons that mediate the strong, weak, and electromagnetic interactions. Furthermore, the SM predicts the existence of the Higgs boson, a scalar particle associated with the Higgs field, which is responsible for the mechanism of electroweak symmetry breaking and for investing particles with mass. Figure 2.1 provides an overview of all particles in the SM, including their masses, electric charges, and spin quantum numbers.




Figure 2.1: Diagram of the elementary particles of the SM and their physical properties [21].

The twelve matter particles are spin- $\frac{1}{2}$ fermions, meaning they obey Fermi-Dirac statistics and the Pauli exclusion principle. These are categorised into six quarks, which participate in the strong interaction, and six leptons, which do not. The leptons are further divided into charged leptons and neutrinos, with the latter interacting exclusively via the weak force. The column arrangement in Figure 2.1 reflects an additional classification into three generations, also referred to as ‘flavours’. While all generations carry identical quantum numbers, only the first-generation fermions contribute to stable matter in the Universe; the heavier generations appear only in high-energy processes. Each fermion also has a corresponding antiparticle with the same mass and lifetime but opposite quantum numbers.

Quarks experience the strong interaction due to their colour charge, which can take one of three values: *red*, *green*, or *blue*. The strong force is described in the SM by the $SU(3)_C$ gauge symmetry, with its dynamics governed by Quantum Chromodynamics (QCD). In this framework, the interaction is mediated by massless spin-1 bosons called gluons, which themselves carry colour charge. This self-interaction leads to a distinctive feature of QCD: the coupling strength increases with distance, resulting in *colour confinement*. Consequently, quarks and gluons are never observed in isolation but are bound into colour-neutral composite states called hadrons. These typically manifest as either mesons (quark–antiquark pairs) or baryons (three-quark combinations), although more exotic multi-quark states, such as tetraquarks and pentaquarks, have also been experimentally confirmed [22, 23].

In the SM, the electromagnetic and weak forces are unified into the electroweak (EW) theory [24–26], also known as the Glashow-Weinberg-Salam model, based on the $SU(2)_W \otimes U(1)_Y$ gauge symmetry. The electromagnetic interaction is mediated by the massless photon (γ), while the weak interaction proceeds via the massive spin-1 bosons $W^\pm$ and Z , which obey Bose-Einstein statistics. The observed mass of the weak bosons, or equivalently, the short range of the weak force, poses a theoretical challenge, as explicit mass terms for gauge bosons violate gauge invariance. This issue is resolved through the Higgs mechanism, independently proposed by Brout and Englert [27], Higgs [28], and Guralnik, Hagen, and Kibble [29]. In this framework, all particles are massless at high energies, with masses dynamically generated below the electroweak scale via spontaneous symmetry breaking of $SU(2)_W \otimes U(1)_Y$.

2.1.2 The mathematical formalism

Having outlined the particle content and interaction structure of the SM, it is now instructive to present its underlying mathematical formulation. The complete dynamics of the SM are encapsulated in its Lagrangian, which encodes the gauge symmetries, particle interactions, and mechanisms for mass generation within a single relativistically invariant framework. This Lagrangian is constructed from quantum field operators constrained by the $SU(3)_C \otimes SU(2)_W \otimes U(1)_Y$ gauge symmetry, Lorentz invariance, and renormalisability. In what follows, the various components of the SM Lagrangian are introduced, highlighting their physical interpretation and role in describing the observed phenomena of the strong, weak, and electromagnetic interactions. It is naturally organised into several distinct sectors:

Gauge sector This includes the kinetic and self-interaction terms for the gauge fields associated with the $SU(3)_C$, $SU(2)_W$, and $U(1)_Y$ symmetries. These are expressed using the corresponding field strength tensors: $G_{\mu\nu}^a$ for $SU(3)_C$, $W_{\mu\nu}^i$ for $SU(2)_W$, and $B_{\mu\nu}$ for $U(1)_Y$. The gauge sector Lagrangian is given by

$$\mathcal{L}_{\text{gauge}} = -\frac{1}{4}G_{\mu\nu}^a G_a^{\mu\nu} - \frac{1}{4}W_{\mu\nu}^i W_i^{\mu\nu} - \frac{1}{4}B_{\mu\nu} B^{\mu\nu}, \quad (2.1)$$

where the indices a and i run over the eight gluons and three weak bosons, respectively. These terms describe the propagation of the gauge bosons and, in the non-Abelian cases, their self-interactions.

Fermion sector The kinetic terms for the quarks and leptons are included, with each fermion field minimally coupled to the gauge fields through covariant derivatives that reflect their gauge charges. The fermionic part of the Lagrangian takes the form

$$\mathcal{L}_{\text{fermion}} = \sum_{\psi} \bar{\psi} i \not{D} \psi, \quad (2.2)$$

where ψ denotes each left- or right-handed fermion field, and $\not{D} = \gamma^\mu D_\mu$ contains the covariant derivative D_μ appropriate to the fermion's gauge representation. This term ensures that fermions interact with the gauge bosons according to their transformation properties under $SU(3)_C$, $SU(2)_W$, and $U(1)_Y$.

Higgs sector The scalar field responsible for electroweak symmetry breaking is introduced as an $SU(2)_W$ doublet with hypercharge $Y = +\frac{1}{2}$:

$$\Phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}. \quad (2.3)$$

The Lagrangian for the Higgs field includes a kinetic term and a potential term that induces spontaneous symmetry breaking:

$$\mathcal{L}_{\text{Higgs}} = (D_\mu \Phi)^\dagger (D^\mu \Phi) - V(\Phi), \quad V(\Phi) = \mu^2 \Phi^\dagger \Phi + \lambda (\Phi^\dagger \Phi)^2, \quad (2.4)$$

where $\mu^2 < 0$ causes the Higgs field to acquire a non-zero vacuum expectation value. This spontaneous symmetry breaking reduces the electroweak gauge symmetry $SU(2)_W \otimes U(1)_Y$ to the electromagnetic subgroup $U(1)_{\text{EM}}$, associated with Quantum Electrodynamics (QED). As a consequence, the weak interaction becomes short-ranged due to the generation of masses for the $W^\pm$ and Z bosons, while the photon remains massless and mediates a long-range force.

Yukawa sector To generate fermion masses, gauge-invariant Yukawa interactions are introduced between the Higgs field and the fermion fields. These interactions couple left-handed $SU(2)$ doublets to right-handed singlets via the Higgs field:

$$\mathcal{L}_{\text{Yukawa}} = -y_u \bar{Q}_L \tilde{\Phi} u_R - y_d \bar{Q}_L \Phi d_R - y_e \bar{L}_L \Phi e_R + \text{h.c.}, \quad (2.5)$$

where Q_L and L_L denote left-handed quark and lepton doublets, u_R , d_R , and e_R the corresponding right-handed singlets, and $\tilde{\Phi} = i\sigma_2 \Phi^*$ is the conjugate Higgs doublet. After spontaneous symmetry breaking, these terms give rise to fermion mass terms proportional to their respective Yukawa couplings.

Combining all these components, the full SM Lagrangian is given by

$$\mathcal{L}_{\text{SM}} = \mathcal{L}_{\text{gauge}} + \mathcal{L}_{\text{fermion}} + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{Yukawa}}. \quad (2.6)$$

This Lagrangian successfully describes all observed non-gravitational interactions in terms of a renormalisable, gauge-invariant quantum field theory. The phenomena of particle masses, electroweak symmetry breaking, and the unification of electromagnetic and weak forces all emerge naturally within this formalism.

2.2 Effective field theory

The concept of effective interactions dates back nearly a century, beginning with Fermi's introduction in 1933 [30] to describe nuclear β decay. Modern effective interaction theory relies on the principle of energy scale separation between different physical phenomena. At low energies, heavy mediators can be approximated as point-like, enabling their description via effective couplings. More specifically, β decay, at energies much lower than m_W , can be represented as a point-like four-fermion interaction for the process $d \rightarrow u e^- \bar{\nu}_e$. This general approach, known as effective field theory (EFT), has historically proven invaluable in simplifying calculations and achieving high precision in low-energy experiments. More recently, EFT frameworks have also been applied to processes where the relevant energy scale significantly exceeds the experimental sensitivity.

The SM itself can be regarded as an EFT, given the consensus that it is not valid up to arbitrarily high energy scales. Within this context, an EFT facilitates predictions of rates and kinematics for BSM processes without requiring detailed knowledge of the underlying high-energy theory. This perspective underpins the Standard Model Effective Field Theory (SMEFT), which systematically incorporates higher-dimensional operators built from SM fields. These operators respect the $SU(3)_C \otimes SU(2)_W \otimes U(1)_Y$ gauge symmetry and extend the SM Lagrangian as follows:

$$\mathcal{L}_{\text{SMEFT}} = \mathcal{L}_{\text{SM}}^{(4)} + \frac{1}{\Lambda} C^{(5)} O^{(5)} + \frac{1}{\Lambda^2} \sum_i C_i^{(6)} O_i^{(6)} + \dots \quad (2.7)$$

Here, $\mathcal{L}_{\text{SM}}^{(4)}$ represents the renormalisable SM Lagrangian, as denoted in Eq. 2.6, while $O^{(n)}$ denotes non-renormalisable operators of mass dimension n .¹ The Wilson coefficients (WCs) $C^{(n)}$ are complex parameters controlling the strength of new interactions and are free parameters of the theory [31], while the operators are suppressed by powers of the mass scale Λ , representing the energy scale at which new physics may emerge. Physically, Λ could correspond to the mass of a new particle mediating the interaction.

The SMEFT framework will be used to interpret the results of the search for charged-lepton-flavour-violating interactions in $e\mu q t$ discussed in Chapter 5.

2.3 Charged-lepton-flavour violation

2.3.1 Theoretical framework

As previously discussed, the leptons shown in Figure 2.1 acquire definite masses through their interaction with the Higgs boson. These masses correspond to freely propagating particles with well-defined quantum numbers, known as mass eigenstates. For charged leptons, the three flavours coincide with their mass eigenstates, ensuring that charged-lepton flavour is conserved in weak interactions.

In contrast, neutrino flavour eigenstates do not align with their mass eigenstates. As a result, neutrinos can oscillate between different flavours as they propagate through space. This phenomenon, neutrino oscillation, implies that neutrinos have mass and that individual lepton flavour numbers are not conserved. These flavour transitions (e.g., $\nu_e \leftrightarrow \nu_\mu$) violate lepton flavour in the neutral sector and provide a resolution to the solar neutrino problem [20].

In the SM, lepton flavour conservation is an accidental symmetry that emerges from its gauge structure, the minimal Higgs sector, and, notably, the absence of right-handed neutrinos. Numerous theoretical models have been proposed to explain the origin of neutrino masses and mixing, yet no single framework has been conclusively favoured by experimental evidence. While the mechanism responsible for LFV in the neutral sector remains uncertain, it is also unclear whether charged-lepton flavour is strictly conserved. The simplest extension of the SM that accommodates neutrino mass, by introducing right-handed Dirac neutrinos with a Yukawa coupling to the conjugate Higgs doublet $\tilde{\Phi}$, naturally predicts cLFV.

In this scenario, neutrino mixing is described by the Pontecorvo-Maki-Nakagawa-Sakata (PMNS) matrix, which encapsulates the mismatch between neutrino flavour and mass eigenstates. This mixing

¹ Dimension here corresponds to the mass dimension, equivalent to $\frac{1}{\text{length}}$ in natural units with $c = \hbar = 1$.

allows processes such as the loop-induced decay $\mu \rightarrow e\gamma$ via virtual neutrino exchange and a W boson, as illustrated in Figure 2.2. The corresponding branching fraction, given by Eq. 2.8 [32, 33], is of the order 10^{-54} – 10^{-55} , approximately 40 orders of magnitude below current experimental sensitivity. This extreme suppression arises from the tiny neutrino masses relative to the W boson mass, despite the sizeable PMNS matrix elements.




Figure 2.2: One-loop Feynman diagram that contributes to $\mu \rightarrow e\gamma$ in the SM extended with right-handed (Dirac) neutrinos. $V_{\mu k}, V_{ke}^*$ are PMNS matrix elements, that account for neutrino flavour mixing. Additionally, one must consider other one-loop Feynman diagrams with the photon line attached to the external leptons.

$$\mathcal{B}(\mu \rightarrow e\gamma) \approx \frac{\Gamma(\mu \rightarrow e\gamma)}{\Gamma(\mu \rightarrow e\nu\bar{\nu})} = \frac{3\alpha}{32\pi} \left| \sum_{k=1,3} \frac{(V_{\text{PMNS}})_{\mu k} (V_{\text{PMNS}})_{ek}^* m_{\nu_k}^2}{m_W^2} \right|^2 \quad (2.8)$$

Despite this suppression, it is often hypothesised that observable cLFV rates may lie just beyond current experimental limits. This expectation stems from the belief that the mechanism responsible for electroweak symmetry breaking operates at a mass scale around $O(1)$ TeV, making it accessible to present or near-future collider experiments [20]. Therefore, any experimental observation of cLFV would unambiguously signal new physics BSM and its minimal extensions, namely the inclusion of non-zero neutrino masses and lepton mixing. Many BSM models that introduce additional interactions involving lepton fields naturally predict cLFV at potentially observable rates. These include supersymmetric see-saw mechanisms [34, 35], two-Higgs-doublet models (2HDM) [36–38], and leptoquark scenarios [39].

In recent years, several measurements of B meson decays initially showed indications of lepton universality violation (LUV) [40–42]. However, some of these deviations have gradually diminished with more data and refined analyses. Despite this, models proposed to explain the original LUV anomalies often predict accompanying LFV effects [43]. Furthermore, many of the new physics scenarios that aimed to address anomalies in the b -quark sector also suggest similar signatures in the top-quark sector. For example, certain leptoquark models designed to explain the R_{D^*} anomaly predict branching fractions for processes such as $t \rightarrow \ell^+ \ell^- c$ reaching values as high as $O(10^{-6})$, where ℓ denotes a charged lepton [44].

Given this connection, searches for cLFV processes involving the top quark provide a promising avenue to explore the physics underlying the anomalies observed in B meson decays.

2.3.2 EFT operators for $2\ell 2q$ contact interactions

Following the SMEFT framework, introduced in Section 2.2, at dimension-five, the Weinberg operator [45] is the sole gauge-invariant operator in the SM. It is responsible for generating Majorana mass terms for the neutrinos by coupling the left-chiral neutrino field to the Higgs field and is therefore not of interest for cLFV. In contrast, many dimension-six operators satisfy SM gauge invariance

and can be organised into the Warsaw basis [46]. These operators are central to this thesis, as they include those facilitating cLFV four-fermion interactions involving two up-type quarks and two leptons ($2\ell 2q$) [47], detailed in Table 2.1. Here, the indices $i, j, k, l \in \{1, 2, 3\}$ represent the flavours of the fermion fields, making the operators flavour-dependent since new physics may couple differently to various fermion generations. While dimension-8 operators also induce cLFV, their contributions are suppressed by an additional factor of $1/\Lambda^2$ in the Lagrangian, as shown in Eq. 2.7. Consequently, this analysis focuses solely on dimension-6 operators.

$O_{lq}^{1(ijkl)}$	$= \left(\bar{l}_i \gamma^\mu l_j \right) \left(\bar{q}_k \gamma_\mu q_l \right)$
$O_{lq}^{3(ijkl)}$	$= \left(\bar{l}_i \gamma^\mu \sigma^I l_j \right) \left(\bar{q}_k \gamma_\mu \sigma_I q_l \right)$
$O_{eq}^{(ijkl)}$	$= \left(\bar{e}_i \gamma^\mu e_j \right) \left(\bar{q}_k \gamma_\mu q_l \right)$
$O_{lu}^{(ijkl)}$	$= \left(\bar{l}_i \gamma^\mu l_j \right) \left(\bar{u}_k \gamma_\mu u_l \right)$
$O_{eu}^{(ijkl)}$	$= \left(\bar{e}_i \gamma^\mu e_j \right) \left(\bar{u}_k \gamma_\mu u_l \right)$
$^\ddagger O_{lequ}^{1(ijkl)}$	$= \left(\bar{l}_i e_j \right) \varepsilon \left(\bar{q}_k u_l \right)$
$^\ddagger O_{lequ}^{3(ijkl)}$	$= \left(\bar{l}_i \sigma^{\mu\nu} e_j \right) \varepsilon \left(\bar{q}_k \sigma_{\mu\nu} u_l \right)$

Table 2.1: SMEFT operators capable of introducing two quark-two lepton cLFV interactions involving a top quark in the Warsaw basis. The indices ij and kl represent the lepton and quark generations respectively. The superscripts 1 and 3 are simply notations, corresponding to operators that could be combined as a linear sum (see e.g. Ref. [46]). In the convention used, l and q are the left-handed lepton and quark doublets, while u and e are the right-handed up-type quark and charged lepton singlets. σ^I are the Pauli matrices with $I \in \{1, 2, 3\}$, $\varepsilon = i\sigma^2$ and $\sigma_{\mu\nu} = \frac{i}{2} [\gamma_\mu, \gamma_\nu]$ where γ_μ are the Dirac gamma matrices. The ‡ symbol marks non-hermitian operators.

The operators can be combined into one single operator [47]:

$$O_{lq}^{-(ijkl)} \equiv O_{lq}^{1(ijkl)} - O_{lq}^{3(ijkl)}. \quad (2.9)$$

Thus, Table 2.2 displays the corresponding WC and Lorentz structure to each operator.

Operator	Wilson coefficient	Lorentz structure
$O_{lq}^{-(ijkl)}$	$c_{lq}^{-(ijkl)}$	Vector
$O_{eq}^{(ijkl)}$	$c_{eq}^{(ijkl)}$	Vector
$O_{lu}^{(ijkl)}$	$c_{lu}^{(ijkl)}$	Vector
$O_{eu}^{(ijkl)}$	$c_{eu}^{(ijkl)}$	Vector
$^\ddagger O_{lequ}^{1(ijkl)}$	$c_{lequ}^{1(ijkl)}$	Scalar
$^\ddagger O_{lequ}^{3(ijkl)}$	$c_{lequ}^{3(ijkl)}$	Tensor

Table 2.2: Relevant SMEFT operators for this thesis and the corresponding WCs and Lorentz structures.

A WC is assigned to each of the operators shown in Table 2.1. Unlike the vector operators, $O_{lequ}^{1(ijkl)}$ and $O_{lequ}^{3(ijkl)}$ are not Hermitian so their Hermitian conjugate must also be considered, with corresponding WCs $c_{lequ}^{1(jilk)}$ and $c_{lequ}^{3(jilk)}$.

The operators defined in Table 2.2 give rise to four-fermion interactions involving a top quark, either an up or charm quark, and two charged leptons of different flavours. These interactions enable rare flavour-violating decay modes of the top quark, such as $t \rightarrow \ell \ell' q$, where $\ell \neq \ell'$ and $q \in \{u, c\}$ [48]. Beyond decays, cLFV interactions can also mediate the production of single top quarks in association with a pair of charged leptons of different flavour. Representative Feynman diagrams for both the cLFV decay of a top quark in $t\bar{t}$ production and single-top production via cLFV operators alongside more top quark details are discussed in Section 4.4.

Importantly, the cross-section for cLFV single-top-quark production increases quadratically with the size of WCs. As a result, even modest non-zero values can lead to sizeable production rates, making this process the dominant contribution to the search sensitivity. Nevertheless, it is still instructive to express the total top-quark decay width Γ in terms of the EFT contributions, which can be parametrised using eight degrees of freedom [47] by:

$$\Gamma(t \rightarrow \ell^\pm \ell'^\mp u_k) = \frac{m_t}{6144\pi^3} \left(\frac{m_t}{\Lambda}\right)^4 \left[4 \left| c_{lq}^{-(jik3)} \right|^2 + 4 \left| c_{eq}^{(jik3)} \right|^2 + 4 \left| c_{lu}^{(jik3)} \right|^2 + 4 \left| c_{eu}^{(jik3)} \right|^2 \right. \\ \left. + \left| c_{lequ1}^{1(jik3)} \right|^2 + \left| c_{lequ1}^{1(ij3k)} \right|^2 + 48 \left| c_{lequ3}^{3(jik3)} \right|^2 + 48 \left| c_{lequ3}^{3(ij3k)} \right|^2 \right]. \quad (2.10)$$

Still according to Ref. [47], the WCs are assumed to be real, which implies that:

$$c_{lq}^{-(jilk)} = c_{lq}^{-(ijk l)}, c_{eq}^{(jilk)} = c_{eq}^{(ijk l)}, c_{lu}^{(jilk)} = c_{lu}^{(ijk l)}, \\ c_{eu}^{(jilk)} = c_{eu}^{(ijk l)}, c_{lequ1}^{1(jilk)} = c_{lequ1}^{1(ijk l)}, c_{lequ3}^{3(jilk)} = c_{lequ3}^{3(ijk l)}.$$

The observable final states are determined by the flavours of the charged leptons (ℓ and ℓ') as well as the decay mode of the W boson from the top decay. When the W decays hadronically, it produces two quarks that hadronise into jets.

For this search, the scale of new physics, Λ , is fixed at 1 TeV and the interpretation is performed in terms of the dimensionless WCs. Since the final discriminant used in the profile likelihood fit cannot distinguish between contributions from different EFT operators, the analysis considers each operator independently. In doing so, the WCs of all other operators are set to zero to isolate the contribution from the operator under study.

2.3.3 Experimental searches

Experimental searches for cLFV span a wide range of energy scales, from low-energy precision experiments to high-energy collider studies. The most stringent limits to date come from low-energy processes involving muon decays. The MEG Collaboration has set an upper limit on the branching ratio of $\mu^+ \rightarrow e^+ \gamma$ at 4.2×10^{-13} [49] in 2016, while the SINDRUM experiment reported $\mathcal{B}(\mu^+ \rightarrow e^+ e^+ e^-) < 1.0 \times 10^{-12}$ [50] in 1988, both at 90% confidence level. Figure 2.3 shows the history of cLFV searches in muons. Further constraints arise from $\mu \rightarrow e$ conversion in nuclei, as well as rare decays of mesons such as π , K , D , and B [20, 51].




Figure 2.3: The history of cLFV searches in muons (not including muonium) [52]. One sees a steady improvement in all modes and then a flattening of the rate improvement throughout the 1990s. MEG has upgrade plans for the $\mu \rightarrow e\gamma$ search. The two next generations of $\mu N \rightarrow eN$, Mu2e/COMET at FNAL and J-PARC are labelled, and possible extensions at Project X and PRIME are shown.

At high energies, cLFV searches targetting τ leptons has been very prolific. An overview of the current status of the observed upper limits on the branching fractions of cLFV τ decays is reported in Figure 2.4.

In addition, Z boson, Higgs boson, and top quark, as well as more exotic processes have been explored by high-energy experiments. ATLAS has contributed significantly to this programme through dedicated searches involving the τ lepton [53], the Z boson [54, 55], the Higgs boson [56], the top quark, and other exotic scenarios [57, 58]. CMS has also pursued similar channels.




Figure 2.4: Current status of the observed upper limits on cLFV τ decays at CLEO, BaBar, Belle, LHCb, ATLAS, and CMS [59]. In addition, projections for Belle II are reported.

For instance, cLFV decays of the Higgs boson to $e\tau$ and $\mu\tau$ have been studied at $\sqrt{s} = 13$ TeV, yielding upper limits of $\mathcal{B}(H \rightarrow e\tau) < 0.20\%$ and $\mathcal{B}(H \rightarrow \mu\tau) < 0.18\%$ at 95% CL [60]. Similar searches by CMS report comparable bounds [61]. In the case of the Z boson, ATLAS sets limits of $\mathcal{B}(Z \rightarrow e\tau) < 5.0 \times 10^{-6}$ and $\mathcal{B}(Z \rightarrow \mu\tau) < 6.5 \times 10^{-6}$ [62], with no significant excess observed.

The Belle and BaBar experiments have also placed competitive limits on $\tau \rightarrow 3\mu$ decays, at 2.1×10^{-8} and 3.3×10^{-8} , respectively [63, 64]. The current most stringent constraint is provided by the LHCb experiment, which sets $\mathcal{B}(\tau^- \rightarrow \mu^- \mu^+ \mu^-) < 4.6 \times 10^{-9}$ [65].

Searches involving the top quark are comparatively less developed but are gaining attention due to their sensitivity to flavour-violating couplings at high energies. Previous efforts include single top production via cLFV interactions at the HERA ep collider [66]. ATLAS and CMS have more recently extended this programme to top-quark decays, exploring rare modes such as $t \rightarrow \ell\ell'q$, where $\ell \neq \ell'$ and $q \in \{u, c\}$. So far, no significant excess over the SM predictions has been reported, and the observations from the ATLAS and CMS experiments are interpreted using the SMEFT framework. Brief reviews of some of these recent analyses are given below.

Preliminary $e\mu$ trilepton ATLAS analysis The ATLAS Collaboration conducted the first direct search for flavour-violating $e\mu qt$ interactions in Refs. [16, 67], using proton–proton collision data collected at $\sqrt{s} = 13$ TeV during 2015–2017, corresponding to an integrated luminosity of 79.8 fb^{-1} . The analysis focused on final states containing three leptons and at least two jets, considering only the signal topology in which the top quark undergoes a cLFV decay. The contributions from different Lorentz structures associated with dimension-6 SMEFT operators were not probed separately.

To enhance sensitivity, a boosted decision tree (BDT) classifier was trained using kinematic variables such as the transverse momentum (p_T) of the leptons. The resulting BDT discriminant distribution, shown in Figure 2.5, was used to extract limits on potential signal contributions.




Figure 2.5: Fitted BDT discriminant distribution in the signal region [16]. The left (right) histogram shows the distribution of the pre-fit (post-fit) targeting the cLFV top quark decay.

No significant deviation from SM expectations was observed. A 95% confidence level upper limit of $\mathcal{B}(t \rightarrow e\mu q) < 6.6 \times 10^{-6}$ was set using the CL_s prescription [68], improving upon a previous indirect bound by three orders of magnitude [48].

CMS $e\mu$ dilepton analysis The CMS Collaboration subsequently conducted a dedicated search for $e\mu qt$ interactions using proton–proton collision data collected between 2016 and 2018, corresponding to an integrated luminosity of 137 fb^{-1} [18]. In contrast to the earlier ATLAS study, this analysis focused on final states with two charged leptons and a hadronically decaying top quark. Both cLFV top quark production and decay processes were considered, providing broader sensitivity to new physics. Furthermore, the analysis explicitly distinguished between operators with different Lorentz structures: vector, scalar, and tensor.

To optimise signal discrimination, a BDT was trained using a range of kinematic variables. The resulting BDT output distributions, shown in Figure 2.6, were used to extract exclusion limits on the branching fractions for each operator type. The most stringent limit was obtained for the scalar operator, yielding $\mathcal{B}(t \rightarrow e\mu\mu) < 7 \times 10^{-8}$ at 95% confidence level. This result improved upon the ATLAS limit by nearly two orders of magnitude.




Figure 2.6: BDT discriminant distributions for the CMS $e\mu$ dilepton cLFV search. The left (right) histogram, taken from [18], shows the distribution of the BDT discriminant in regions with exactly (more than) one b -tagged jet. The middle (bottom) panel shows the ratio of data events and the pre-fit (post-fit) predictions.

CMS $e\mu$ trilepton analysis The CMS Collaboration later on extended its cLFV programme with a new search targeting final states containing exactly three charged leptons [69], closely resembling the earlier ATLAS analysis [16]. In addition to probing cLFV top-quark decays, this analysis also included cLFV top-quark production, a signal topology first explored by CMS in Ref. [18]. To maximise sensitivity, separate signal regions were defined for the decay and production modes, as illustrated by Figure 2.7.

Operators were analysed individually according to their Lorentz structure (scalar, vector, tensor) and quark flavour. For each SR, a dedicated binary BDT was trained to discriminate signal from

background. The contribution from non-prompt leptons was estimated using a data-driven method. No significant excess was observed, and the most stringent constraint was placed on the scalar operator, yielding an upper limit on the branching fraction of $\mathcal{B}(t \rightarrow e\mu u) < 1.2 \times 10^{-8}$ at 95% confidence level. This result improved upon the previous CMS limit [18] by nearly a factor of six.




Figure 2.7: Distributions of the post-fit BDT discriminant targeting the cLFV top-quark decay (left) and production (right) signal [69]. Contributions from the two signal modes (production and decay) are combined within each SR and are shown as the solid red line.

ATLAS $\mu\tau$ trilepton analysis The ATLAS Collaboration has also performed the first direct search for a $\mu\tau qt$ cLFV interaction using 140 fb^{-1} of proton–proton collision data collected between 2015 and 2018 [17]. The analysis targets final states containing two same-sign muons, one hadronically decaying tau lepton, and at least one jet. Both single top-quark production and top-quark decay channels are considered, with different Lorentz structures for the effective operators analysed separately.

Due to limited statistics, the observed event yields in the signal region are used directly to interpret the results, as shown in Figure 2.8. An upper limit at 95% confidence level is set on the branching fraction, yielding $\mathcal{B}(t \rightarrow \mu\tau q) < 8.7 \times 10^{-7}$.




Figure 2.8: Observed event yields in $CRt\bar{t}\mu$ (a) and the signal region (SR) (b) are shown as a function of the scalar sum of lepton and jet transverse momenta, H_T , and are compared to the post-fit expectations from simulations [17].

CMS $\mu\tau$ dilepton analysis The CMS Collaboration recently presented a dedicated search for charged-lepton-flavour violation in top-quark interactions, as reported in Ref. [70]. The analysis targets final states containing two oppositely-charged $\mu\tau$ leptons accompanied by hadronic activity consistent with top-quark production. Following the other recent searches, this study simultaneously probes flavour-violating top-quark production and decay within a unified EFT framework, thereby increasing sensitivity to a broader class of cLFV operators. Events were categorised according to lepton flavour combinations and jet multiplicities, allowing the analysis to exploit the distinct kinematic features of each signal topology.

The interpretation was performed in terms of dimension-six operators contributing to flavour-violating top interactions, considered individually to avoid cancellations among WC. A dedicated deep neural network (DNN) classifier was trained for each signal category to exploit correlations among multiple kinematic observables, significantly improving the discrimination between flavour-violating signal hypotheses and the dominant SM backgrounds. No statistically significant deviation from the SM expectation was observed. The upper limits at 95% confidence level are set at $\mathcal{B}(t \rightarrow \mu\tau u) < (0.04, 0.08, \text{ and } 0.12) \times 10^{-6}$, and $\mathcal{B}(t \rightarrow \mu\tau c) < (0.81, 1.71, \text{ and } 2.05) \times 10^{-6}$ for scalar, vector, and tensor-like operators, respectively.




Figure 2.9: Combined distributions in the DNN score after the profile likelihood fit for all data-taking periods for the vector operators with $\mu\tau u$ (left) and $\mu\tau c$ (right) couplings.

The ATLAS Experiment

This chapter details the LHC operation and specifications in Section 3.1 and an overview of the key detector systems within ATLAS in Section 3.2. Information on the LHC is mainly derived from Ref. [71], whereas concerning ATLAS it was extracted from Ref. [72].

3.1 The Large Hadron Collider

Operated by the European Organisation for Nuclear Research (CERN), the Large Hadron Collider (LHC) is the world's largest and most powerful particle accelerator, a 27 km circular collider located between 45 and 170 m deep under the Franco-Swiss border in the vicinity of Geneva, Switzerland. It is a two-ring hadron accelerator and makes use of hundreds of superconducting magnets to bend (1232 dipole magnets) and focus (392 quadrupole magnets) particles' trajectories. Moreover, it is primarily designed to collide a pair of proton beams, pp collisions, travelling at almost the speed of light at centre-of-mass energies, $\sqrt{s}$, of up to 14 TeV, as well as lead ion beams up to an energy of 2.8 TeV per nucleon. The energy is a crucial parameter for it determines the mass threshold of the particles that can be produced.

The LHC is the final step in a long chain of particle accelerators required to boost particle energies up to the TeV range. A comprehensive illustration of the accelerator complex is shown in Figure 3.1. In the first step of the chain, protons, produced by ionising hydrogen, are accelerated in a linear accelerator, LINAC4, to an energy of 160 MeV. Next up, the beam is injected into a sequence of successively larger circular accelerators: the Proton Synchrotron Booster (PSB), the Proton Synchrotron (PS), and the Super Proton Synchrotron (SPS), which raise the beam energy to 2 GeV, 26 GeV, and 450 GeV respectively. After exiting the SPS, the beam is split in two; the resulting beams are then transferred to the LHC beam pipes, in which they circle in opposite directions. The two beams afterwards accelerated up to their individual collision energies: in 2011, this energy was 3.5 TeV per beam; in 2012, 4 TeV; between 2015-2018, 6.5 TeV; finally, starting in 2022, the beam energy reached 6.8 TeV.

The LHC beams are not continuous but rather are composed of a number of proton bunches, with a designed maximum of 2808 bunches per beam, each bunch containing $O(10^{11})$ protons. The minimum spacing between bunches is 25 ns, corresponding to a bunch crossing frequency of 40 MHz. The two beams are made to collide at four points around the LHC ring where each of the four major LHC experiments are located. ATLAS and CMS (Compact Muon Solenoid) are the two independent,




Figure 3.1: The CERN accelerator complex and its experiments [73].

general-purpose detectors with wide-ranging physics programs, including Higgs, top, and electroweak physics, but also direct searches for beyond-SM physics, such as extra dimensions and dark-matter particles. ALICE (A Large Ion Collider Experiment) is dedicated to strongly-interacting matter at extreme energy densities, where a phase of matter called quark-gluon plasma is formed during heavy-ion collisions. At last, the LHCb (Large Hadron Collider beauty) experiment is focused on the study of decays of heavy-flavoured hadrons and the search for indirect evidence of non-SM physics via precision measurements of observables such as CP asymmetries, angular observables, and unitarity-triangle parameters.

Besides the beam energy, an important parameter for a collider is the instantaneous luminosity $\mathcal{L}$, related to the number of events per second generated in the collisions R by $R = \mathcal{L}\sigma_{pp}$, where σ_{pp} is the cross-section of pp collisions. For a collider using bunched beams, such as the LHC, under the assumption that the beam has a Gaussian profile, the instantaneous luminosity can be expressed as

$$\mathcal{L} = \frac{n_p^2 n_b f \gamma}{4\pi \varepsilon_n \beta^*} F, \quad (3.1)$$

where n_p is the number of protons per bunch; n_b , the number of bunches per beam; f , the revolution frequency; γ , the relativistic factor accounting for length contraction; ε_n , the normalised transverse

beam emittance; β^* , the beta function at the collision point; and F , the geometric factor arising from the beams colliding at an angle to one another. The maximum design instantaneous luminosity is $\sim 10^{34} \text{cm}^{-2}\text{s}^{-1}$, although operations to date have seen the machine achieve peak instantaneous luminosities of approximately double this number.

The main data-taking periods to date have been 2011-2012 and 2015-2018, known as Run 1 and 2, respectively, whereas Run 3, started in 2022 and is currently planned to end in 2026. The long shutdowns between Runs were necessary to prepare the machine for operation at higher energies and overall renovation work, such as the trigger system upgrade and the New Small Wheel installation, for example. A third shutdown, Long Shutdown 3, will follow Run 3 and lay the groundwork for the High-Luminosity LHC operations.

The integral of the delivered luminosity over time is called integrated luminosity. It is a measurement of the collected data set size, and it is an important value to characterise the performance of an accelerator. It is frequently expressed in terms of inverse cross-section, most commonly inverse femtobarn, fb^{-1} . Figure 3.2 shows the cumulative integrated luminosity delivered to ATLAS by LHC during Run 2.




Figure 3.2: Cumulative luminosity versus time: 156fb^{-1} delivered to ATLAS (green), 147fb^{-1} recorded by ATLAS (yellow), 140fb^{-1} and certified to be good quality data (blue) during stable beams for pp collisions at 13 TeV in LHC Run 2.

3.2 The ATLAS detector

The ATLAS experiment is a multi-purpose particle detector with a forward-backward symmetric cylindrical geometry and a nearly 4π coverage in solid angle around the interaction point. Components of the sub-detectors are divided between the barrel and end-cap regions which refer to the curved face and flat ends of the cylindrical geometry, respectively. A diagram of the ATLAS detector is shown in Figure 3.3. It consists of an inner tracking detector surrounded by a thin superconducting solenoid, electromagnetic and hadronic calorimeters, and a muon spectrometer incorporating three large superconducting toroidal magnet systems. All of the sub-detectors are cooled to maintain

stable operating temperatures for its detectors and electronics, using advanced liquid and gas-based systems to dissipate heat efficiently even under high radiation and thermal loads. Moreover, a trigger system, used for filtering the data to be kept for further processing, operates in combination with the sub-detectors. The detector has a total length of about 44 m, a diameter of about 25 m and weighs around 7,000 tonnes.




Figure 3.3: Cut-away view of the ATLAS detector [72].

The coordinate system

At this point, it is convenient to define the coordinate system used by ATLAS. It is right-handed with its origin at the nominal interaction point (IP) in the centre of the detector and the z -axis along the beam pipe. The x -axis points from the IP to the centre of the LHC ring, and the y -axis points upwards, therefore defining the xy plane as the transverse plane to the beam direction. Spherical coordinates are adopted as well: the azimuthal angle ϕ is defined on the transverse plane and by convention ranges between $-\pi$ and π ; the polar angle θ is measured from the positive direction of the z -axis, but it is conveniently expressed in terms of the pseudorapidity as

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

For a massless particle, η is equal to the rapidity y , defined as

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right), \quad (3.2)$$

where E is the particle energy and p_z its longitudinal momentum. Differences in rapidity y are Lorentz invariant under boosts along the longitudinal axis, unlike the polar angle θ and the

pseudorapidity η gives a good approximation to y for an ultra-relativistic particle. Furthermore, angular distances between two particles in the $\eta\phi$ -plane can be measured by the approximately boost-invariant variable

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

The projection of a particle's momentum vector on the transverse plane,

$$p_T = \sqrt{p_x^2 + p_y^2}, \quad (3.3)$$

is called transverse momentum. The transverse energy is then defined as $E_T^2 = p_T^2 + m^2$. Since the momentum of the incoming partons is, at first approximation, only directed along the longitudinal axis, the transverse momenta of all the outgoing particles should sum to zero unless any particles escape undetected. Therefore, the quantity

$$|\vec{p}_T^{\text{miss}}| = - \sum_i \vec{p}_{T,i} \quad (3.4)$$

called missing transverse momentum, can be used to estimate the vector sum of undetected particles, such as neutrinos. The missing transverse energy is in turn defined as $E_T^{\text{miss}} = |\vec{p}_T^{\text{miss}}|$.

3.2.1 The tracking system

The Inner Detector (ID) is the closest sub-detector to the interaction point and consists of silicon pixel, silicon micro-strip, and transition-radiation detectors, as illustrated in Figure 3.4. Silicon sensors measure particles by detecting the ionisation they cause when passing through the silicon material. The ionised electrons and holes are separated by an applied electric field, creating a measurable current or voltage signal proportional to the deposited energy. The inner detector has a geometrical coverage of $|\eta| < 2.5$ and is immersed in a 2 Tm axial magnetic field provided by the surrounding superconducting solenoid. The tracker aims to reconstruct the curved trajectories of charged particles in order to measure their momenta as well as to determine the primary and secondary vertices. These trajectories are reconstructed using position measurements (“hits”) in the three sub-detectors and the direction in which the track is bent allows the charge of the particle to be identified. The ID components, as seen in detail in Figure 3.5, are described as follows from inner to outer radii.

The innermost component is the silicon pixel detector, which provides high-granularity hits that are essential for distinguishing the numerous tracks produced by LHC collisions and identifying the primary collision vertex. At high operating luminosity, events typically contain multiple pp collisions per bunch crossing, a phenomenon known as pile-up. During Run 2, the ATLAS detector recorded an average of 33.7 interactions per bunch crossing, as illustrated in Figure 3.6. To mitigate the effects of pile-up, it is crucial to accurately measure the primary vertex position and associate particle tracks to it. In addition, the pixel detector plays an integral role in reconstructing the position of secondary vertices to detect in-flight decays of short-lived particles, such as B -mesons, which is of particular importance for identifying events involving top quark decays.

Closest to the beam pipe are four layers of 1744 silicon pixel detectors and the innermost layer, called Insertable B-Layer (IBL), has been installed before Run 2 to increase the tracking resolution. The IBL takes advantage of recent technological advancements [76], utilising both 3D and planar silicon sensors. It employs hybrid pixels with dimensions of $50 \times 250 \mu\text{m}^2$ connected to specialised front-end chips created with 130 nm CMOS technology, which are capable of enduring an ionising




Figure 3.4: Cut-away view of the ATLAS inner detector [72].




Figure 3.5: The structure of the ATLAS Inner Tracking Detector [74].




Figure 3.6: Luminosity-weighted distribution of the mean number of interactions per crossing during 2015 – 2018 [75].

dose of 250 mrad. In contrast, the three outer layers and three endcap discs use $50 \times 400 \mu\text{m}^2$ planar pixels linked to readout chips made with $0.25 \mu\text{m}$ CMOS technology. The IBL achieves a position resolution of $8 \mu\text{m}$ in the $r\phi$ -plane and $40 \mu\text{m}$ along the longitudinal direction, while the other pixel layers have resolutions of $10 \mu\text{m}$ and $115 \mu\text{m}$, respectively [77].

Surrounding the pixel detector, the semiconductor tracker (SCT) is installed. It is made up of ~ 4088 modules, each containing 12 chips with 128 ~ 12 cm-long silicon microstrip, amounting to ~ 6.2 million read-out channels in total. In the central region, the SCT consists of four barrel layers and is complemented in each endcap region by nine discs, all of them using pairs of sensors arranged back-to-back to provide two-dimensional position information. The spatial resolution is $17 \mu\text{m}$ in the $r\phi$ -plane and $580 \mu\text{m}$ in the z (r) direction for the central (endcap) region. On average, charged particles are expected to leave four (eight) hits in the pixel (SCT).

The outermost component of the tracking system is the transition radiation tracker (TRT) which covers the $|\eta| < 2.0$ region. It consists of alternating layers of drift tubes and transition radiation material. Such radiation is produced via avalanche gas ionisation, where charged particles traversing the tubes ionise the gas within, liberating electrons which drift in an electric field and are read out from a wire in the centre of the tube. In the central region, the TRT is made of 52544 144 cm long straw drift tubes collected in 73 concentric layers. In each endcap, the tubes are shorter (37 cm) and radially oriented. Each tube consists of a polyamide pipe 4 mm in diameter filled with a xenon-gas mixture. The tubes are interleaved with polypropylene fibres which serve as the transition radiation material. When a charged particle traverses the polypropylene it emits radiation proportionally to the Lorentz factor $\gamma = E/m$, particularly high for electrons. The TRT thus provides additional tracking information and contributes to the electron identification. It complements precision measurements from the pixel and SCT with in excess of 30 additional hits to radial distances of just over a meter.

3.2.2 The calorimeter system

Represented in Figure 3.7, calorimeters provide a destructive measurement of the energy of all incident particles, with the exception of muons and neutrinos.

The electromagnetic calorimeter (ECAL), positioned just outside the tracker solenoid, is designed to capture electromagnetic showers primarily from electrons and photons. Covering the region $|\eta| < 3.2$, the ECAL features active liquid argon (LAR) layers interleaved with passive lead absorbers arranged in an accordion geometry. The high-density lead absorbers promote the development of showers upon

interaction with incident particles. The ionised charge in the LAr, proportional to the energy deposited by traversing particles, is collected by electrodes within the LAr gap. Each electrode consists of three conductive copper layers separated by insulating polyamide sheets, with the outer layers at high voltage and the inner layer used for signal readout via capacitive coupling. The barrel section of the ECAL, extending to $|\eta| < 1.475$, is segmented into three layers: the first layer, with the highest granularity, resolves incoming particles; the second, thickest layer collects the majority of the energy deposit; and the third layer captures the tail of the electromagnetic shower. The endcap section comprises an outer wheel ($1.375 < |\eta| < 2.5$) and an inner wheel ($2.5 < |\eta| < 3.2$). Both the barrel and endcap calorimeters are equipped with a presampler, an additional LAr volume placed in front of the first layer, to correct for energy lost in upstream materials (ID, cryostats, coil). Depending on the η region, the ECAL components correspond to between 23 and 39 radiation lengths.¹

Particles that are not fully absorbed in the ECAL continue into the hadronic calorimeter (HCAL), which is designed to measure the energy of neutral and charged particles that interact via the strong force. This capability is crucial for reconstructing hadronic showers, known as jets, and for determining missing transverse energy. The HCAL primarily detects charged and neutral hadrons; as well as muons. It is composed of three subdetectors: the tile calorimeter, the LAr hadronic end-cap calorimeter (HEC), and the LAr forward calorimeter (FCal). The tile calorimeter, a sampling calorimeter utilizing steel as the absorber and scintillating tiles as the active medium, is situated in the region $|\eta| < 1.7$, directly surrounding the ECAL. It is segmented into three layers, each azimuthally divided into 64 modules. These modules contain the tiles, absorbers, optic fibres, photomultipliers, and front-end electronics. The front-end electronics can provide analogue channel sums, known as trigger towers, which are relayed to the trigger system. The HEC, covering the region $1.5 < |\eta| < 3.2$, is a copper/LAr sampling calorimeter positioned behind the electromagnetic endcap calorimeter, sharing its cryostat.

The FCal extends the coverage of both the ECAL and HCAL to the region $3.1 < |\eta| < 4.9$. The first layer, optimised for electromagnetic measurements, is a LAr sampling calorimeter with copper absorber plates. The second and third layers, designed for hadronic measurements, use tungsten absorbers to contain the lateral spread of hadronic showers.




Figure 3.7: Cut-away view of The ATLAS Calorimeter detector [72].

¹ A radiation length is defined as the mean distance at which a high-energy electron loses its energy due to bremsstrahlung by a factor of $1/e$.

3.2.3 The muon system

Muons, which deposit only a small fraction of their energy in the calorimeters, traverse the entire detector and reach the muon spectrometer (MS), the outermost detector system in ATLAS. As illustrated in Figure 3.8, it surrounds the calorimeters and uses large toroid magnets to create a magnetic field integral of 2.0 to 6.0 Tm, which deflects the particles, allowing their trajectories to be measured in tracking chambers. The MS consists of two main detector types: monitored drift tubes (MDTs) covering $|\eta| < 2.7$, and cathode-strip chambers (CSC) in the region $2.0 < |\eta| < 2.7$. MDT chambers are composed of three to eight layers of drift tubes, achieving an average resolution of $80\ \mu\text{m}$ per tube and approximately $35\ \mu\text{m}$ per chamber. The CSCs, located in the endcaps, are multi-wire proportional chambers with cathode planes segmented into orthogonal strips, providing a resolution of $40\ \mu\text{m}$ in the bending plane and about 5 mm in the transverse plane.

To enable trigger capabilities to the MS, additional detectors with response times under 25 ns are included. resistive plate chambers (RPCs) are used for $|\eta| < 1.05$, while thin-gap chambers (TGC) cover $1.05 < |\eta| < 2.4$. RPCs consist of small gas volumes equipped with two sets of orthogonal cathode strips, offering a resolution of about 10 mm in both the bending and non-bending directions. TGCs are multi-wire proportional chambers with a smaller wire-to-cathode distance than wire separation. They provide muon track information with a precision of 2 to 7 mm in the η coordinate and 3 to 7 mm in the ϕ coordinate.




Figure 3.8: Cut-away view of the ATLAS muon detector [72].

3.2.4 The trigger system

LHC collisions occur at the bunch crossing rate of 40 MHz with an average of 30 interactions per bunch-crossing. The kind of events typically searched for are rare and have high- p_T leptons, photons or very high- p_T jets, large E_T^{miss} , b -jet signatures and high jet multiplicity. To select events with such interesting physics signatures given the maximum event rate is determined by the transfer bandwidth

of ~ 1 GB/s and the event size of ~ 1.3 MB/event and to reduce the data-taking rate to levels which can be read-out and stored, a trigger system [78, 79] is required. The trigger system also ensures that signals in each of the detector systems are associated with the correct bunch crossings. This is a key consideration since the muon time-of-flight typically exceeds the time between bunch crossings, as does the signal persistence in the calorimeter.

A two-level trigger system is employed in the ATLAS detector. The first level, known as the Level-1 (L1) trigger, operates within a $2.5 \mu\text{s}$ decision window and generates an output event rate of 100 kHz. The L1 trigger is composed of four modules: L1Calo, L1Muon, L1Topo, and the Central Trigger Processor (CTP).

L1Calo processes energy depositions in the calorimeter towers to identify and construct jets, electron/photon candidates, and τ -lepton candidates. L1Muon identifies and classifies muon candidates based on their momentum, using hits from the RPC and TGC, with tracks required to be consistent with originating from the interaction vertex. L1Topo receives the identified objects, collectively known as Trigger Objects (TOBs), from both L1Calo and L1Muon. It then applies selection criteria based on geometric and kinematic parameters, such as estimating the angular distance between objects, removing overlaps, and calculating global quantities like total or missing transverse momentum. The CTP processes the outputs from L1Calo, L1Muon, and L1Topo to issue the final L1 trigger decision. If an event is accepted, the L1 trigger initiates detector readouts and transfers the information to the High-Level Trigger (HLT).

While the L1 trigger relies on FPGAs and custom electronics, the HLT is powered by a computer farm with approximately thirty thousand cores. The HLT reconstructs objects starting from the L1 TOBs, utilizing finer-granularity calorimeter data, tracking information, and precision measurements from the muon system. This partial reconstruction applies object identification criteria similar to those used in offline analysis. When the HLT accepts an event, the complete event data is transferred to permanent storage, with an average accepted event rate of 1 kHz.

Common Elements in Top-Physics Analysis

The LHC generates large quantities of hadronic collisions during each bunch crossing, producing a vast number of final-state particles from interactions between constituent partons. This complexity poses significant challenges for the simulation and analysis of physics events. Simulating LHC collisions involves intricate processes, as the relevant energy scales span both perturbative and non-perturbative QCD regimes and produce high particle multiplicities. In this high-multiplicity environment, advanced algorithms are required to trigger and reconstruct the processes of interest amidst pile-up noise while preventing double-counting of detector signals. This chapter provides an overview of LHC event structures and examines their impact on simulation and reconstruction, with a particular focus on reconstructing leptons and hadrons in the ATLAS detector. Finally, a discussion about top-quark physics is included as well as a description of the statistical analysis framework implemented.

4.1 Hadron-collider physics

Protons are composed of two u -quarks and one d -quark, collectively referred to as valence quarks, held together by the strong interaction. This interaction is mediated by a continuous exchange of gluons, giving rise to a sea of virtual quarks. The behaviour of these proton components lies outside the scope of QCD at energy scales below Λ_{QCD} . At LHC energy levels, a proton can be described as a collection of partons, each representing a portion of the proton's total momentum. When protons A and B collide, the hard scattering occurs between two partons, a and b , which carry momentum fractions x_a and x_b , as shown in Figure 4.1.




Figure 4.1: Schematic view of a hard scattering process. The incoming protons have momenta p_A and p_B and the interacting partons carry the momentum fractions x_a and x_b .

The centre-of-mass energy available for the hard scattering is given by $\hat{s} = x_a x_b s$, where s represents the pp centre-of-mass energy. The probability of finding a parton carrying a fraction x of the proton's momentum is characterised by a parton distribution function, $f_i(x)$. The index i identifies the various parton species within the proton, namely u, d, c, s, g , and, in some cases, b , all treated as massless [80]. The cross-section for a process $pp \rightarrow X$ can therefore be expressed as

$$\sigma_{pp \rightarrow X}(\mu_F, \mu_R) = \int_0^1 dx_a \int_0^1 dx_b \sum_{ij} f_i(x_a, \mu_F) f_j(x_b, \mu_F) \hat{\sigma}_{ij}(\hat{s}, \mu_F, \mu_R)_{ab \rightarrow X}. \quad (4.1)$$

Parton distribution functions (PDFs) exhibit distinct behaviours, as illustrated in Figure 4.2: the valence quark PDF peaks near $x \sim 1/3$, while the gluon PDF dominates at small x . PDFs are typically regarded as process-independent properties of the hadron, an assumption based on the ability to factorise the partonic interaction from the proton structure. This factorisation assumes the two are independent.




Figure 4.2: Parton distribution functions computed at next to next to leading order (NNLO) precision in QCD by the NNPDF group [81] at scales $\mu^2 = 10 \text{ GeV}^2$ (left) and $\mu = 10^4 \text{ GeV}^2$. The distributions are shown for the sea quarks and the valence up and down quarks (u_v, d_v). For s, c and b quarks, the PDFs are approximately equal for the corresponding quark and anti-quark.

Since PDFs cannot be calculated perturbatively, they are derived through global fits to experimental data, guided by assumptions about their analytical form. Key datasets include deep inelastic scattering (DIS) data from the HERA ep collider at DESY [82], which provide highly sensitive inputs across a broad phase space. Additionally, data from the LHC and the Tevatron $p\bar{p}$ collider at Fermilab [83] are crucial for constraining the gluon PDF at high x . A comprehensive review of the datasets used to construct PDFs is available in Ref. [84].

The partonic cross-section $\hat{s}$ and the PDFs also depend on the factorisation scale μ_F and the renormalisation scale μ_R . Intuitively, μ_F represents the scale that delineates long- and short-distance physics, particularly in relation to substructures like quark loops. Conversely, μ_R corresponds to the

scale of the QCD running coupling [85], reflecting how the cross-section varies with the choice of renormalisation scale. In collider cross-section calculations, these two scales are generally chosen to be equal and set to a value representative of the interaction under study. However, the choice of scales introduces theoretical uncertainties into predictions for cross-sections and other observables at the LHC, as will be elaborated in Section 5.6.2.

The remaining partons not involved in the hard scattering process give rise to the underlying event. This component of the collision is characterised by additional soft radiation, independent of the primary hard interaction, yet integral to the overall event dynamics. Both the hard scattering and underlying event produce coloured particles, such as quarks and gluons, which radiate gluons themselves. These emissions generate further QCD radiation, culminating in the development of a parton shower.

As the parton shower progresses, the energy of its constituents diminishes with each successive gluon emission. When the energy scale drops below the perturbative QCD threshold, the interaction enters the non-perturbative regime. At this stage, hadronisation occurs, a complex process where coloured partons recombine into colour-neutral hadrons, such as protons, neutrons, and mesons. These hadrons, being stable or long-lived, are detected as the observable final states of the collision. This transition from partons to hadrons is fundamental in particle physics experiments at the LHC, as it defines the measurable outcomes of proton-proton collisions.

4.2 Physics simulation

Event simulation

Analyses conducted at the LHC experiments rely extensively on simulations. Monte Carlo (MC) simulation event generators are employed to simulate physics processes, providing signal and background expectations for comparison with observed data. The use of MC generators is crucial due to the inherently stochastic nature of quantum mechanical processes. MC simulations are also indispensable for modelling the complex variations in the efficiency of the ATLAS detector and the analysis itself, which depend on factors such as angle, momentum, and other kinematic variables. In summary, MC simulations are used to design analysis strategies, estimate efficiencies, evaluate background levels, and assess the sensitivity of the analysis.

Given a physical process of interest $pp \rightarrow X$, the MC generators must account for several key steps [86], described as follows:

- ◊ The first step involves computing the quantum mechanical matrix element (ME) for the hard scatter of partons within the incoming protons. This computation is performed to a specified order in perturbative QCD: typically leading order (LO), next to leading order (NLO), or NNLO.
- ◊ To determine the cross-section of the physical process, the PDFs of the incoming protons must be incorporated. The ME calculation is convoluted with the PDFs and integrated over the available phase space using MC integration techniques. The PDFs are factorised at a specific resolution scale, known as the factorisation scale, as briefly discussed in Section 4.1. These PDFs are then estimated at different energy scales using the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) [87] evolution equations, which can be applied at LO or NLO.
- ◊ The strong coupling constant (α_s), typically quoted at a renormalisation energy scale $Q = M_Z$, is incorporated using the Renormalisation Group Equations. These equations allow α_s to be evaluated at different energy scales relevant to the generated process.

- ◊ Uncertainties arising from missing higher-order terms in the perturbative expansion of the ME, as well as in the DGLAP evolution of the PDFs, are addressed by varying the renormalisation and factorisation scales used in the calculations. This scale variation provides an estimate of the theoretical uncertainties associated with the process.
- ◊ High-energy partons produced in the initial hard collision emit additional partons as they propagate through space-time. Quarks can emit gluons ($q \rightarrow gq$), and gluons can undergo splitting into quark-antiquark pairs ($g \rightarrow q\bar{q}$), provided the partonic energy is sufficient to prevent confinement. These processes result in the formation of cascades of coloured particles, referred to as parton showers (PSs). Parton showers are modelled using perturbative QCD, with corrections often applied via (next-to-)next-to-leading-logarithmic resummation [88, 89] to address divergences that arise from the radiation of "soft and collinear" low- p_T gluons.
- ◊ Radiation occurs in two main forms. initial state radiation (ISR) consists of emission of hard radiation from the incoming partons prior to the hard scatter, typically leading to the production of additional jets; whereas final state radiation (FSR) concerns emission from outgoing partons, which contributes to the internal structure of jets. Additionally, photons or other bosons can be radiated from either the initial or final states, further enriching the complexity of the event.
- ◊ Partons within the incoming protons that do not directly contribute to the primary physics process under investigation may engage in additional interactions, collectively referred to as the underlying event.
- ◊ As partons propagate and their energy drops below approximately 1 GeV, the strong interaction becomes confining, leading to the non-perturbative process of hadronisation. During hadronisation, partons are transformed into colour-neutral hadrons, predominantly light mesons, resulting in the formation of hadronic jets.
- ◊ The parton shower and fragmentation processes are parametrised using assigned renormalisation and factorisation scales, which influence the modelling of these phenomena.
- ◊ For unstable particles, such as heavy (b - or c -) hadrons, secondary decays may occur, producing additional particles. These decays can result in the formation of a secondary vertex within a jet composed primarily of light hadrons.

Figure 4.3 depicts the progression from the hard scatter to hadronisation, including potential secondary hadronic decays.

The following provides a brief review of the generators commonly used by the ATLAS Collaboration [91]. PYTHIA [92], a widely used general-purpose event generator, was developed over 30 years ago and remains integral to particle physics simulations. Unlike automated code generation approaches, PYTHIA relies on an extensive collection of hardcoded processes. It is capable of simulating complete events, encompassing the hard process (at LO QCD accuracy), parton showers, the underlying event, hadronisation, and particle decays. Frequently, PYTHIA is interfaced with a ME generator to focus solely on the parton shower stage. Various merging schemes exist to ensure consistency between partonic emissions from the ME and the parton shower. Unstable particle decays, particularly those of b - and c -quark hadrons, can be handled by external tools like EVTGEN [93], which specialises in heavy hadron decays. Hadronisation in PYTHIA is fully modelled using the Lund string model [94, 95].




Figure 4.3: Illustration depicting the formation of a parton shower [90]. Partons originating from incoming protons generate the hard scatter (highlighted in red). These partons subsequently emit gluons, forming a parton shower/QCD bremsstrahlung (represented in blue). At lower energies and larger distances, partons undergo hadronisation, recombining into colour-neutral hadrons (shown in green). Some hadrons may further decay (dark green), potentially emitting QED bremsstrahlung (depicted in yellow). The underlying event, initiated by the hard process, is illustrated in purple.

SHERPA is also a general-purpose event generator that can compute MEs at NLO accuracy in QCD. It can generate all SM processes and supports many BSM extensions. SHERPA implements the CKKW [96] and MENLOPS [97] merging schemes, with hadronisation based on the cluster model [98].

MADGRAPH5_aMC@NLO [99] is another versatile ME event generator, computing amplitudes on the fly from automatically generated code. Through the FEYNRULES utility [100], the SM Lagrangian can be modified, and the resulting physics supplied to MG_aMC@NLO¹ as a Universal FEYNRULES Output (UFO) [101] model. MG_aMC@NLO is capable of generating the hard process at LO or NLO, but it does not include PS algorithms; thus, it must be interfaced with a generator such as PYTHIA. The output events from MG_aMC@NLO are encoded in the Les Houches Event format [102].

POWHEG-Box [103] is also an NLO-capable ME event generator, offering features similar to MG_aMC@NLO from the user's perspective. The primary differences lie in the internal calculations and the parton shower matching/merging approach: POWHEG-Box adopts the POWHEG method [104] by default, whereas MG_aMC@NLO uses the MC@NLO scheme [105]. Technical details on these two approaches can be found in Ref. [106].

Finally, HERWIG [107, 108] is a general-purpose event generator that simulates high-energy particle collisions, focusing on the hard scattering process, parton showers, hadronisation, and particle decays. Developed as an alternative to PYTHIA, HERWIG employs a different approach to modelling QCD

¹ MADGRAPH5_aMC@NLO has been shortened to MG_aMC@NLO.

dynamics, particularly in the area of parton showers and hadronisation. Unlike PYTHIA, which uses the Lund string model for hadronisation, HERWIG relies on the cluster model, where hadrons are formed through the cluster decay of a colour-singlet state. HERWIG also includes more detailed treatments of the angular ordering in parton showers and allows for a range of processes and particle decays to be simulated, including those involving heavy quarks. Like PYTHIA, HERWIG can be interfaced with ME generators to simulate the hard scattering process, after which it handles the parton shower and hadronisation.

Pileup and underlying event modelling

To accurately simulate LHC events, numerous parameters governing parton-shower modelling are tuned based on experimental data. For PYTHIA 8, the so-called A14 tune [109] is used in conjunction with the NNPDF 2.3 LO PDF set [110] in ATLAS event simulations. Events generated with SHERPA and Herwig use their respective default tunes [111].

In addition, jets arise not only from partons produced in the hard-scatter interaction but also from multiple parton interactions and pile-up, all of which must be included to properly model the particle multiplicity in LHC events. Pile-up is accounted for by overlaying samples of inelastic pp scattering events on top of the hard-scatter process. These are generated with PYTHIA 8.186, using the NNPDF 2.3 LO PDF and the dedicated A3 tune [112]. To reflect differences in detector and machine conditions, ATLAS MC simulation production for Run 2 is split into three data-taking periods: 2015+2016, 2017, and 2018. The MC simulated samples in each period are then corrected to match the pile-up distribution observed in the corresponding data.

Detector simulation

To accurately model LHC events, the outputs of parton-shower generators (generator-level events) must be passed through a simulation of the ATLAS detector. The GEANT4 toolkit [113] is used to simulate both the interaction of particles with detector material and the digitisation of detector signals. Consequently, this simulation accounts for the detector’s geometrical acceptance and finite resolution, providing a model that can be directly compared with ATLAS data. The same particle-reconstruction techniques used on real data are then applied to the detector-simulated events. Moreover, some MC simulated samples employed for systematic uncertainty evaluations related to physics-process modelling rely on a fast-simulation of the calorimeter system [114, 115].

Information on the simulation of the samples used in this analysis is provided in Sections 5.1 and 5.2

4.3 Physics object reconstruction and identification

To analyse the events selected by the trigger menu discussed in Section 3.2.4, the signals recorded by the ATLAS detector must be reconstructed into well-identified “physics objects.” This section focuses on the objects used in the analysis presented in Chapter 5, namely electrons, muons, and jets. The reconstruction of charged tracks is introduced first, both because it is relevant to all of these objects and because it is used to identify the primary vertex. Finally, important algorithms are introduced such as overlap removal.

4.3.1 Track reconstruction

In the inner detector, signals recorded by adjacent pixel and SCT channels are assembled into clusters, which reflect energy deposits left by individual charged particles passing through [72]. Pairs of one-dimensional SCT clusters on opposite sides of a sensor module, or single pixel clusters, are subsequently converted into three-dimensional space-points, with position uncertainties determined by the detector geometry and sensor pitch. Sets of three space-points are then combined to form numerous preliminary track seeds. Each seed's quality is determined by its extrapolated distance from the principal interaction point (impact parameter), its momentum, and the pattern of hits. Using a combinatorial Kalman filter [116], these seeds are extended both outwards and inwards in radius along 'search roads', while the filter also smooths the evolving track trajectory [117].

When multiple viable space-point extensions arise on the same detector layer, the filter creates several track candidates from a single seed. Each candidate is then assigned a score determined by the cluster multiplicity, the number of holes², the track momentum, and the χ^2 of the track fit. This score serves as the primary criterion for resolving ambiguities caused by clusters shared among multiple track candidates.

All track candidates must exhibit a transverse momentum greater than 400 MeV, lie within the region $|\eta| < 2.5$, and include at least seven associated clusters, with at most two of these clusters shared with other tracks. Any candidate containing more than one hole in the pixel detector or more than two holes overall in the combined pixel plus SCT detectors is discarded. In addition, the transverse displacement from the beam-line must remain under 2 mm, and the longitudinal displacement must be less than $3 \text{ mm} / \cos \theta$. Following the initial track-building procedure, each candidate undergoes a further refinement stage to resolve cluster-to-track association ambiguities. Subsequently, several algorithms attempt to extend the track outwards to the TRT. A Kalman filter is employed once again to incorporate the TRT hits, and each track is then globally refitted with a χ^2 -based approach [118].

Candidate primary vertices are formed from reconstructed tracks that lie within the luminous interaction region, defined by the average collision point and the beam sizes of the two protons at the interaction point. The primary vertex (PV) associated with the hard-scattering process is chosen to be the one with the largest $\sum_{\text{tracks}} p_T^2$, where the sum spans all associated tracks carrying $p_T > 0.5 \text{ GeV}$ [119]. At least two tracks are required to form the PV. The vertex position is typically determined with a precision of around $20 \text{ } \mu\text{m}$ in the transverse plane and $40 \text{ } \mu\text{m}$ along the beam axis. A well-reconstructed PV is indispensable for nearly all physics analyses and performance measurements.

Each reconstructed track is described by five perigee parameters measured relative to the PV. Among these parameters, the transverse impact parameter, d_0 , and the longitudinal impact parameter, z_0 , denote the minimal transverse and longitudinal distances of the track with respect to the vertex. The remaining three parameters are: the azimuthal angle ϕ , the polar angle θ of the track momentum at the PV, and the ratio $q/|\vec{p}|$, representing the track charge divided by its total momentum [120]. Figure 4.4 provides a schematic depiction of these geometric definitions.

² Holes are defined as intersections of the reconstructed track trajectory with a sensitive detector element that does not contain a matching cluster.




Figure 4.4: The perigee parameters which characterise the track at its point of closest approach to the PV. The parameter d_0 quantifies the transverse impact parameter in the x - y plane, while z_0 captures the longitudinal displacement along z [120].

4.3.2 Electrons

Reconstruction Electron candidates are reconstructed by matching localised energy deposits in clusters of electromagnetic calorimeter cells with charged-particle tracks in the inner detector [121–123]. Such candidates are required to have $p_T > 10$ GeV and $|\eta_{\text{clu}}| < 2.47$. Any candidate found within $1.37 < |\eta_{\text{clu}}| < 1.52$ is discarded, as this interval marks the transition between the barrel and endcap regions of the ECAL. The presence of cables and other instrumentation in this zone degrades energy resolution for candidates reconstructed there [72].

If multiple tracks are associated with the same energy deposit in the calorimeter, the one closest in ΔR is selected; if no matching track is found, the electron candidate is reclassified as a photon. All electron candidates must be consistent with originating from the primary vertex. In particular, the significance of d_0 must satisfy $|d_0|/\sigma_{d_0} < 5$. Moreover, the nominal longitudinal impact parameter criterion, $|z_0 \sin(\theta)| < 0.5$ mm.

The electron reconstruction efficiency depends on both $|\eta|$ and E_T . This efficiency is determined via a tag-and-probe technique [122, 124]. Samples of electrons originating from $Z \rightarrow ee$ and $J/\psi \rightarrow ee$ decays are used, with the latter providing lower- p_T electrons. Events are selected based on the electron-positron invariant mass, designating one electron as the “tag” if it passes a strict set of criteria, while the other, the “probe”, follows looser requirements. To extract the efficiency for a particular selection, that criterion is applied to the probe electron. For electrons with $E_T > 15$ GeV, the reconstruction efficiency reaches approximately 99%.

Identification Electron candidates reconstructed in this manner often include a sizeable fraction of hadronic jets or converted photons. To ensure that only true electrons are selected, an identification algorithm is applied. This algorithm computes a multivariate likelihood ratio, drawing on 20 input variables derived from track, calorimeter, and TRT measurements [121]. For each variable x_i in the

set $\vec{x}$, separate probability density functions $P_{S,i}$ and $P_{B,i}$ describe the distributions for real-electron (signal) and non-electron (background) hypotheses, respectively. The resulting discriminant $d_{\mathcal{L}}$ assigned to each candidate is defined by:

$$d_{\mathcal{L}} = \frac{\mathcal{L}_S}{\mathcal{L}_S + \mathcal{L}_B}, \quad \text{where} \quad \mathcal{L}_{S(B)}(\vec{x}) = \prod_{i=1}^n P_{S(B),i}(x_i). \quad (4.2)$$

Different analyses place varying requirements on electron selection efficiency versus background rejection, which motivates defining multiple working points (WPs) based on different thresholds for $d_{\mathcal{L}}$. These likelihood thresholds are adjusted according to the electron candidate's p_T and $|\eta|$, ensuring that selection efficiency changes smoothly with momentum. In this study, the TightLH and LooseAndBLayerLH WPs are employed. The latter extends the LooseLH WP by also requiring a hit in the IBL, thereby reducing contamination from photon conversions [125].

Figure 4.5 illustrates the p_T and η dependence of the most frequently used WPs, derived via a tag-and-probe technique from $Z \rightarrow ee$ events. For TightLH electrons, the identification efficiency deteriorates below about 50 GeV in momentum or energy, whereas the LooseLH WP retains approximately 90% of low-energy electrons, indicating that most electrons are accepted by this looser criterion.




Figure 4.5: Identification efficiencies for various working points utilised in Run 2 [126], plotted as a function of (left) the electron transverse energy E_T and (right) the electron pseudorapidity η . In the top panels, the inner and outer error bars represent the statistical and total uncertainties, respectively. The middle panels display the ratio of the efficiencies measured in data to those obtained in MC simulation, while the bottom panels show the statistical and total uncertainties in these ratios. Efficiency is reduced in the transition region between the electromagnetic calorimeter barrel and endcap, $1.37 < |\eta| < 1.52$, in order to control background contributions.

Isolation Electrons passing the identification requirements can arise either from the decay of heavy resonances produced in the primary interaction, such as the W or Z bosons, or from secondary processes which include photon conversion, heavy hadron decays, and mis-reconstruction of light

hadrons. Electrons of the first kind, commonly referred to as prompt, are characterised by a smaller activity in their surroundings and thus can be selected by making use of isolation variables. Two discriminating variables have been designed for this purpose:

- The calorimetric isolation $E_T^{\text{cone}0.3}$ is defined as the total transverse energy of all calorimeter clusters within a cone of $\Delta R = 0.3$ surrounding the electron candidate's cluster. The transverse energy contained in a rectangular window of size $\Delta\eta \times \Delta\phi = 0.125 \times 0.175$, centred on the electron cluster barycentre, is then subtracted. Corrections are applied to account for any electron energy leakage outside this window, as well as for contributions from pile-up.
- The track-isolation variable, $p_T^{\text{varcone}0.3}$, is defined as the total transverse momentum of all tracks lying within a radius $\Delta R = \min(0.3, 10 \text{ GeV}/E_T)$ around the electron candidate.³ Only tracks that are consistent with the primary vertex are considered in this sum, excluding those associated with the electron itself (including any tracks from converted bremsstrahlung). In addition, each track must have $E_T > 1 \text{ GeV}$ and fulfil certain criteria regarding the number of hits and its longitudinal impact parameter.

A multivariate algorithm uses these variables, normalised by the lepton p_T , together with the ΔR between the lepton and the track-jet axis, the ratio of the lepton p_T to the track-jet p_T , the number of tracks within the track jet, and two b -tagging-related observables. The output of this algorithm defines the “Prompt Lepton Veto” (PLV) WPs.

The efficiency of each isolation WP depends on the lepton p_T . For prompt electrons, the PLVTight and PLVLoose WPs deliver efficiencies of around 70% and 90%, respectively, at a transverse momentum of 20 GeV, comparable to the simpler fixed-cut (FC) tight and fixed-cut loose WPs, which rely on straightforward requirements applied to the track and calorimetric isolation variables. However, for non-prompt electrons originating from light- or heavy-flavour hadrons, the PLV WPs achieve superior rejection performance compared with the FC WPs.

Efficiency corrections In the context of an analysis, the probability of selecting events that contain real electrons is governed by the product of the trigger, reconstruction, identification, and isolation efficiencies. Each of these steps may show minor discrepancies between data and simulation, which are corrected by reweighting the simulated events using scale factors defined as $w_{\text{SF}} = \epsilon_{\text{data}} / \epsilon_{\text{MC}}$. These factors, which are functions of the electron's kinematic quantities such as p_T and η , ensure that the simulation reflects the detector performance observed in data. The efficiencies in both data and simulation are measured with the same tag-and-probe method, derived from $Z/\gamma^* \rightarrow ee$ and $J/\psi \rightarrow ee$ events [122, 124].

Requirements Table 4.1 presents the two classes of electrons defined for this analysis. Tight electrons impose strict identification and isolation criteria, whereas Loose electrons are intended for the overlap removal procedure detailed in Section 4.3.6 to prevent their misidentification as other objects. Both classes of electrons place requirements on the transverse and longitudinal impact parameters, d_0 and z_0 , respectively, to ensure that these objects originate from the hard-scatter process. In addition, the condition on the transverse impact parameter accounts for the resolution on d_0 , σ_{d_0} .

³ The cone radius can vary according to the electron's transverse momentum, $\Delta R = \min(0.3, 10 \text{ GeV}/p_T^e [\text{GeV}])$. Thus, for high-energy leptons, the cone becomes narrower (hence the “var” notation) [124].

	Electron selection	
	Tight	Loose
p_T	$> 10 \text{ GeV}$	$> 10 \text{ GeV}$
$ \eta_{\text{clu}} $	< 2.47	< 2.47
Identification criteria	TightLH	LooseAndBLayerLH
Isolation criteria	PLVTight	–
$ d_0 /\sigma_{d_0}$	< 5	< 5
$ z_0 \sin(\theta) $	$< 0.5 \text{ mm}$	$< 0.5 \text{ mm}$

Table 4.1: Tight and Loose electron definitions.

4.3.3 Muons

Reconstruction Muon reconstruction is initially performed in the ID and in the MS separately, and the resulting tracks are then combined to produce the muon tracks [127, 128]. Track reconstruction in the inner detector follows the generic algorithm detailed in Section 4.3.1 and is not specific to muons. In the MS, a dedicated search algorithm identifies segments within each MDT and trigger chamber. Aligned hits in the bending plane are found via Hough transforms and are fitted linearly, whereas the perpendicular coordinate is determined from the RPCs. Track candidates are formed by simultaneously fitting hits that originate from different segments, and those compatible with the primary vertex and meeting the χ^2 selection criteria are accepted. The MS track is then extrapolated and matched to a corresponding inner detector track. Whenever a match is found, a global fit is performed, yielding the so-called combined muon. A calibration of the muon momentum scale is applied using the method described in Ref. [129].

In the region $|\eta| < 0.1$ ⁴, muon candidates can also be reconstructed from inner detector tracks matched to isolated energy deposits in the calorimeters, in a manner consistent with the behaviour of a minimum-ionising particle (calorimeter-tagged muons).

Identification The identification algorithm is designed to distinguish prompt muons from those originating in hadron decays. Among the principal variables for this discrimination are the significance of the charge-to-momentum ratios q/p measured in the ID and in the MS, the difference between the p_T values measured in the ID and in the MS, divided by the p_T of the combined track, and the reduced χ^2 of the combined track fit.

All muon tracks are required to have at least one pixel hit, at least five SCT hits, and fewer than three pixel or SCT holes. Several identification working points are defined in ATLAS. Figure 4.6 displays the p_T and η dependence of the identification working points employed in Run 2, obtained via a procedure similar to the electron tag-and-probe method, using muons from $Z \rightarrow \mu\mu$ and $J/\psi \rightarrow \mu\mu$ decays. The default working point, referred to as **Medium**, minimises systematic uncertainties related to muon reconstruction and calibration. For **Medium**, at least three hits in at least two MDT layers (for

⁴ In this region, there is a gap in the MS, essential for the passage of services to the calorimeters, the magnet solenoid and the inner detector.

$|\eta| > 0.1$) and a q/p significance smaller than seven are required. As measured in data, the **Medium** identification WP achieves an efficiency exceeding 98% for muons arising from J/ψ and Z -boson decays.




Figure 4.6: (left) Efficiency of the **Medium** identification working point used in Run 2 as a function of the muon p_T for muons with $0.1 < |\eta| < 2.5$. (right) The efficiency of the **Loose**, **Medium**, **Tight** identification WPs, measured in $Z \rightarrow \mu\mu$ events, is displayed as a function of η for muons with $p_T > 10$ GeV. The efficiency decreases in the $|\eta| < 0.1$ region, which lacks MS coverage. When not negligible, the statistical uncertainties in the measurement are indicated by the error bars. The bottom panel presents the ratio of the measured to predicted efficiencies, including both statistical and systematic uncertainties. From [128].

Isolation As in the case of electrons, muons produced in the decays of heavy particles such as the W , Z , or Higgs bosons are often isolated from other activity, whereas muons from semi-leptonic hadron decays frequently appear inside jets [127]. To distinguish between these two populations, both calorimeter- and track-based isolation variables are fed to a multivariate likelihood discriminant following the same approach described for electrons. Multiple isolation working points are defined, and in particular for this analysis, **PLVTight** and **PLVLoose** are employed [128].

Efficiency corrections Analogously to the electron case, scale factors are derived with tag-and-probe methods in $Z/\gamma^* \rightarrow \mu\mu$ and $J/\psi \rightarrow \mu\mu$ events. Corrections are made for reconstruction, identification, isolation and track-to-vertex association (TTVA) criteria.

Requirements Two categories of muons are defined, as summarised in Table 4.2. A **Tight** selection applies stringent isolation requirements to select prompt muons originating from hard-scatter interactions, whereas a **Loose** selection is specifically designed to accept non-prompt muons from heavy-flavour hadron decays by relaxing isolation criteria and inverting the transverse impact parameter requirement. Both selections mandate that muons satisfy the **Medium** identification criteria.

In addition, muon candidates must have $p_T > 10$ GeV and $|\eta| < 2.5$. The longitudinal impact parameter condition is similar to that imposed on electrons, namely $|z_0 \sin(\theta)| < 0.5$, mm. For muons, the standard (“tight”) transverse impact parameter cut is given by $|d_0|/\sigma_{d_0} < 3$.

	Muon selection	
	Tight	Loose
p_T	$> 10 \text{ GeV}$	$> 10 \text{ GeV}$
$ \eta $	< 2.5	< 2.5
Identification criteria	Medium	Medium
Isolation criteria	PLVTight	PLVLoose
$ d_0 /\sigma_{d_0}$	< 3	> 3
$ z_0 \sin(\theta) $	$< 0.5 \text{ mm}$	$< 0.5 \text{ mm}$

Table 4.2: Tight and Loose muon definitions.

4.3.4 Jets

Reconstruction Jets are formed by collimated streams of particles originating from partons that fragment after the hard-scatter process. These particles deposit energy in both the ECAL and HCAL as they traverse the detector. The jet-reconstruction procedure [130–132] begins with the creation of clusters, constructed from sets of topologically connected calorimeter cells that record a significant signal above noise. These ‘topo-clusters’ are initially built using the more accurately measured electromagnetic component of the jet and are thereafter treated as massless particles, which serve as input to a jet-finding algorithm.

The anti- k_t algorithm [133] iteratively combines pairs of particle candidates, ordered by their relative transverse momentum, into jets until a stopping criterion is reached; typically when the “distance” between adjacent jets surpasses a specified threshold. For a pair of candidates i and j , the distance may be defined as

$$d_{ij} = \min \left(p_{Ti}^{2\kappa}, p_{Tj}^{2\kappa} \right) \frac{\Delta R_{ij}^2}{R^2}, \quad (4.3)$$

where ΔR is the distance in the $r - \phi$ plane, and R and κ are arbitrary parameters. In the anti- k_t algorithm $\kappa = -1$, so that the distance d_{ij} between a soft and a hard particle is driven by the hard-particle p_T . Conversely, two soft particles with a similar separation ΔR_{ij} exhibit a larger d_{ij} . As a result, soft particles are prone to cluster with hard ones before merging among themselves. In this thesis, the radius parameter is set to $R = 0.4$. The anti- k_t algorithm is collinear and infrared safe, which means that small-angle or low-energy gluon emissions do not significantly modify its outcome.

Candidates are then calibrated using the Particle Flow (PFlow) algorithm [134], which aims to identify individual particles by combining calorimeter data with track information from the ID. This approach allows low-energy particles, deflected outside the nominal jet cone by the ID’s magnetic field, to be included in the clustering stage. In addition, contributions arising from tracks matched to pile-up vertices are excluded.

Identification Background processes, such as proton losses upstream of the interaction point, cosmic-ray showers, or calorimeter noise, can yield spurious jet candidates. Jet selection criteria are therefore employed to eliminate these background jets, while maintaining high efficiency for jets originating from pp collisions. This procedure, known as *jet cleaning*, is applied after overlap removal. Because the background-jet fraction is small, the identification requirements [135] are relatively lenient. In this thesis, the LooseBad selection is utilised.

Distinguishing jets from pile-up vertices, rather than the primary vertex, poses a greater challenge. To address this, a multivariate algorithm known as the Jet Vertex Tagger (JVT) [136] is employed for jets with $p_T < 60$ GeV and $|\eta| < 2.4$. Scale factors are then applied to correct any efficiency differences between data and simulation once the JVT requirement is applied. Reconstructed jet candidates must satisfy $p_T > 25$ GeV and $|\eta| < 2.5$.

Jet calibration A calibration of the jet energy scale (JES) is executed to correct measured jet energies to those expected at generator-level [137]. This process removes contributions from pile-up and accounts for unmeasured fractions of the hadronic shower, as well as energy that is either deposited in inactive material or lies outside the reconstruction cone. Corrections derived from simulated dijet events are applied to both simulation and data. An additional calibration is then imposed on jets in data to compensate for any residual discrepancies between data and simulation, using analyses of dijet, multi-jet, Z+jet, and γ +jet events.

jet energy resolution (JER) measurements are also performed to verify consistency between data and simulation [137]. Whenever the simulation exhibits superior resolution compared to data, a smearing procedure is applied to the former. Conversely, if the simulation's resolution is poorer than that observed in data, no smearing is employed; instead, an additional systematic uncertainty is assigned to the simulation to account for the discrepancy. This uncertainty encapsulates the entire difference in resolution for the relevant range of jet p_T and η .

Jet flavour tagging As further detailed in Section 4.4, the top quark decays almost exclusively via $t \rightarrow Wb$ in the SM, making identification of jets initiated by bottom quarks (b -tagging) a crucial tool for selecting events containing top quarks. The DL1r algorithm [138] is used to tag jets containing b -hadrons, relying on the pixel detector's vertexing capabilities to reconstruct displaced vertices associated with hadronic decays, as illustrated in Figure 4.7. Several algorithms, focused on the impact parameter of tracks, the topology of secondary decay vertices, and reconstruction of the $b \rightarrow c \rightarrow s$ decay chain, are integrated into a neural network that outputs the probabilities for the jet to be b -, c -, or light-flavoured⁵.

A requirement on this network's output, corresponding to a 70% efficiency for identifying true b -jets in top quark decays, is applied. Simulation-based efficiency corrections are then applied to jets originating from the hadronisation of b -quarks [139], c -quarks [140], or light-flavour partons [141].




Figure 4.7: Three jets emerging from the primary vertex [142]. Two originate from a light-flavour quark, while one contains a b -hadron. A secondary vertex (SV) develops when the b -hadron decays semi-leptonically, typically after travelling a few mm. The transverse impact parameter d_0 of its tracks is relatively large, as is the separation L_{xy} of the secondary vertex itself (both measured in the transverse, or xy , plane).

⁵ Light-flavour jets refers to any jet initiated by u -, d -, s -quarks or gluons.

4.3.5 Missing transverse momentum

Neutrinos produced in pp collision events are not detected by ATLAS, owing to their weak interaction with matter. Nonetheless, their presence becomes apparent through a non-conservation of the total transverse momentum in the event. Assuming all products of the hard-scatter process are reconstructed, their transverse momenta should sum to zero. Hence, any residual net transverse momentum is used to define a vector, whose inverse is the missing transverse momentum vector [143, 144], calculated as

$$\mathbf{p}_T^{\text{miss}} = - \sum_{\text{selected electrons}} \mathbf{p}_T^e - \sum_{\text{selected muons}} \mathbf{p}_T^\mu - \sum_{\text{accepted jets}} \mathbf{p}_T^{\text{jet}} - \sum_{\text{unused tracks}} \mathbf{p}_T^{\text{track}}. \quad (4.4)$$

This sum is formed from the vector sums of the transverse momenta of accepted electrons, muons, and jets, as defined in the previous Subsections, with the lepton-jet overlap removal applied according to Section 4.3.6. The final component, known as the "soft term," is derived from tracks that are not associated with any reconstructed physics object, yet remain associated to the primary vertex. It can include contributions from the hard-scatter process, the underlying event, and pile-up interactions.

The magnitude of the missing transverse momentum vector is referred to as the missing transverse energy, E_T^{miss} . In this analysis, no requirement is imposed on E_T^{miss} .

4.3.6 Overlap removal

Once reconstructed, certain physics objects may exhibit characteristics that are typical of more than one category. For example, an electron shower might be identified as both an electron and a jet, or a muon from a semi-leptonic b -hadron decay could simultaneously be assigned to a jet and counted as a stand-alone muon. To avoid such double counting, an overlap removal (OR) algorithm [145] is employed. The OR tool uses the following rules:

- Remove any calorimeter-tagged muon sharing a track with an electron within a cone of $\Delta R < 0.2$;
- Remove any electron sharing a track with a remaining muon;
- Reject any jet within a cone of $\Delta R < 0.2$ of an electron;
- Reject any electron within a cone of $\Delta R < 0.4$ of a remaining jet;
- Reject any jet within a cone of $\Delta R < 0.2$ or ghost-matched [146] to a muon, if the number of tracks associated to the jet is less than three;
- Reject any muon within a cone of $\Delta R < 0.4$ of a remaining jet.

The OR procedure is run with leptons satisfying the Loose definitions.

4.4 Top-quark physics

The top quark was the final quark to be discovered, identified in 1995 by the CDF [147] and D0 [148] experiments at Fermilab. It is the most massive particle in the SM, with a measured mass of $m_t = 172.57 \pm 0.29 \text{ GeV}$ [20].

The cross-sections for $t\bar{t}$ and other top-quark-related processes at the LHC are summarised in Figure 4.8. Top quarks are abundantly produced in $t\bar{t}$ events and at 13 TeV, its production cross-section, measured by ATLAS, is 829 ± 1 (stat.) ± 13 (syst.) ± 8 (lumi) ± 2 (beam) pb, in excellent agreement with the theoretical prediction of 832_{-29}^{+20} (scale) ± 35 (PDF + α_s) pb at NNLO+NNLL in QCD [149, 150]. Top-antitop quark pairs are predominantly produced via gluon-gluon fusion, however, quark-antiquark annihilation and quark-gluon fusion processes also contribute to $t\bar{t}$ production. The LO Feynman diagrams for $t\bar{t}$ production in pp collisions are shown in Figure 4.9.




Figure 4.8: Summary of top-quark production cross-section measurements obtained by ATLAS, compared to corresponding theoretical predictions [151]. All theoretical predictions are calculated at NLO or higher. Fiducial (fid.) cross-section measurements are obtained by considering the "final-state" particles produced in top-quark decays, typically within a restricted phase space.




Figure 4.9: LO $t\bar{t}$ production Feynman diagrams. Gluon-gluon fusion (a) and quark-antiquark annihilation (b).

In addition to $t\bar{t}$ production, top quarks are also produced via single-top processes at the LHC. The dominant contribution arises from single-top production in the t -channel ($\sim 73\%$), followed by production in association with a W boson, tW ($\sim 23\%$). LO Feynman diagrams for these processes,

as well as for single-top production in association with a Z boson (tZq) and in the s-channel, are presented in Figures 4.10 and 4.11. Additionally, example LO Feynman diagrams for $t\bar{t}$ production in association with a vector boson ($t\bar{t}W$, $t\bar{t}Z$) are shown in Figure 4.12.




Figure 4.10: LO example Feynman diagrams for single-top production in t-channel (a) and s-channel (b).




Figure 4.11: LO example Feynman diagrams for tW (a) and tZq (b) production.




Figure 4.12: LO example Feynman diagrams for $t\bar{t}W$ (a) and $t\bar{t}Z$ (b) production.

Owing to its large mass, the top quark decays extremely rapidly, with a mean lifetime of approximately 5×10^{-25} s, leaving no time to hadronise. In the SM, the top quark decays via the weak interaction into a W boson and a down-type quark, almost always a bottom quark. This preference arises because the third generation of quarks mixes minimally with other generations in weak decays ($|V_{tb}| \approx 1$).

The decay products of the top quark are determined by the subsequent decay of the W boson. The W boson decays hadronically into a qq' pair in roughly two-thirds of cases or leptonically into a lepton-neutrino pair in the remaining one-third of cases. As a result, the signature of a top quark decay in the ATLAS detector typically includes a b -tagged jet along with two additional jets or a lepton accompanied by E_T^{miss} . A Feynman diagram illustrating the possible decay channels of $t\bar{t}$ is provided in Figure 4.13.




Figure 4.13: Decay modes of $t\bar{t} \rightarrow W^+ b W^- \bar{b}$. Each of the W bosons can decay hadronically or leptonically.

The presence of new physics capable of inducing cLFV top quark interactions introduces novel production and decay mechanisms for the top quark. These possibilities form the basis of the search presented in Chapter 5. However, standard processes, including single top quark production and top quark pair production, serve as significant backgrounds to the cLFV search, complicating the identification of potential signals of new physics.

4.5 Statistical analysis

The principal objective of data analysis in high-energy physics is to test the prevailing model of particle interactions and, in doing so, to search for effects not anticipated by the SM. Consequently, a quantitative framework is required to define the discovery of any given process. This section provides an overview of the statistical tools used to claim a discovery or set upper limits. Specifically, we discuss the maximum likelihood fit to the data, employed to derive a limit on the signal cross-section.

As is typical in particle physics, one needs a measure (or “test statistic”) of how confidently one can reject the “signal-plus-background” (alternative) hypothesis in favour of the “background-only” (null) hypothesis. We therefore begin by describing the frequentist methodology, which uses the likelihood ratio as a test statistic.

4.5.1 The likelihood function

In a generic experiment, one measures a certain variable x , and the resulting measurements are represented in a histogram with N bins. The expectation value n_i for the bin i is

$$E[n_i] = \mu S_i(\theta) + B_i(\theta), \quad (4.5)$$

where, $S_i(B_i)$ is the expected number of signal (background) entries in the bin i . The parameter μ , known as the signal strength, governs the magnitude of the signal contribution. Here, μ serves as the parameter of interest (POI). By construction, $\mu = 0$ indicates a background-only hypothesis (no signal), whereas $\mu = 1$ corresponds to the nominal signal rate S_i . Both signal and background expectations depend on θ , a set of nuisance parameters (NPs) that encode systematic uncertainties.

While the histograms of simulated signal and background events effectively represent a distribution, the statistical model for such a binned distribution⁶ can be viewed as a series of counting experiments, each described by a Poisson distribution, for each bin. Consequently, the binned likelihood function (LF) is formed as the product of the Poisson probabilities over all bins [152].

⁶ A histogram used to model a sample is commonly referred to as a ‘template’.

Since there is a prior estimation of the systematic effects, constraints can be added to the LF for each NP. This is accomplished by introducing Gaussian constraint terms, also referred to as penalty terms, typically defined with a mean of 0 and a variance of 1 [153]. These constraint terms encode the known systematic uncertainties and prevent the corresponding nuisance parameters from drifting to significantly larger values, or shrinking away altogether, during the fit to the data, unlike fully unconstrained or “free-floating” parameters. The LF then takes the form:

$$\mathcal{L}(\mu, \theta) = \prod_{i=1}^N \mathcal{P}(n_i | \mu S_i + B_i) \prod_{j \in \text{syst}} \mathcal{G}(\theta_j), \quad (4.6)$$

where the dependence $S_i(\theta)$ and $B_i(\theta)$ have been omitted for the sake of notation, $\mathcal{P}$ represent the Poisson probabilities and $\mathcal{G}$ indicates the Gaussian constraints on the nuisance parameters.

Each NP can influence the nominal background and signal predictions by modifying either their shape, *i.e.* imposing distinct effects on each bin, or their normalisation, or both. Different interpolation methods are employed for each scenario: shape-related variations use a bin-by-bin linear interpolation, whereas normalisation shifts rely on an exponential interpolation [154]. The exponential approach guards against unphysical negative normalisation values. Moreover, when exponential interpolation is combined with a Gaussian constraint term, the normalisation follows a log-normal distribution.

A common approach is to separate the background B into multiple categories, each representing a distinct source. Certain backgrounds may be poorly modelled and thus carry large systematic uncertainties. In such situations, subsidiary measurements in control regions (CRs), where one or more background sources are enhanced and the expected signal is minimal, offer valuable insight into the background normalisation, and sometimes even its shape. In contrast, signal regions (SRs) are designed to maximise signal acceptance while suppressing background.

Similar to defining the signal strength, one can introduce normalisation factors (NFs) for specific backgrounds as additional, free-floating nuisance parameters. By integrating CRs into the LF, primarily by adding extra bins in the LF formula⁷, we gain tighter control over the nuisance parameters related to the targeted backgrounds.

Nuisance parameters that capture the impact of limited MC simulated statistics, common in reweighted simulated samples, are incorporated into the LF as so-called ‘gamma factors’ (or simply gammas). In each bin, a single “MC-stat” NP is assigned to the total prediction ($S_i + B_i$) and enters the LF via a Poisson constraint term.

From the measurement x , one obtains the best-fit values of both the POI μ and the NPs θ , including the NFs, by locating the global maximum of the LF across the parameter space (μ, θ) . Since sums and minimisation are often more convenient than products and maximisation, one typically employs the negative log-likelihood (NLL), $-\ln \mathcal{L}$. The maximum likelihood estimators (MLEs) $\hat{\mu}, \hat{\theta}$ then follow by setting the partial derivatives of the NLL functions with respect to each parameter to zero.

⁷ All regions (SRs and CRs) must be orthogonal, meaning statistically uncorrelated, so that their likelihoods can be simply multiplied.

4.5.2 The test statistic

To evaluate a particular hypothesised value of μ , one defines the profile likelihood ratio as

$$\lambda(\mu) = \begin{cases} \frac{\mathcal{L}(\mu, \hat{\theta}_\mu)}{\mathcal{L}(\hat{\mu}, \hat{\theta})}, & \hat{\mu} \geq 0 \\ \frac{\mathcal{L}(\mu, \hat{\theta}_\mu)}{\mathcal{L}(\hat{0}, \hat{\theta})}, & \hat{\mu} < 0 \end{cases} \quad (4.7)$$

In the numerator, $\hat{\theta}_\mu$ is the set of nuisance parameters that maximises the LF for a given μ , hence “profiling” over θ . The denominator contains the LF at its global maximum, obtained by the overall MLEs $\hat{\mu}$ and $\hat{\theta}$. By definition, $0 \leq \lambda \leq 1$, and values near 1 indicate that the data align well with the hypothesised value of μ . Typically, $\ln \lambda(\mu)$ exhibits a parabolic shape in the (μ, θ) space. When the constraints on the nuisance parameters are relaxed, the profile likelihood ratio becomes broader in μ , reflecting a reduction in the information available about μ due to systematic uncertainties.

In many physical scenarios, one assumes $\mu \geq 0$ because the signal rate cannot be negative. However, allowing $\mu < 0$ can sometimes be useful, for instance if $\hat{\mu} < 0$ suggests a deficit of signal-like events relative to the expected background. Unfortunately, imposing the boundary at $\mu = 0$ complicates the asymptotic distributions. An alternative procedure places a constraint directly on the test statistic, rather than on μ itself: if $\hat{\mu} < 0$, one replaces $\hat{\mu}$ with 0 and uses the conditional MLE of θ at $\mu = 0$.

Analogously to the LF treatment, it is easier to work in logarithms. Therefore, we introduce the quantity

$$q_\mu = -2 \ln \lambda(\mu). \quad (4.8)$$

To evaluate how strongly the data conflict with a given hypothesis for the signal strength μ , one employs the test statistic q_μ . Larger values of q_μ signify increased incompatibility between the data and that hypothesised value of μ .

This incompatibility is quantified by computing a p -value,

$$p_\mu = \int_{q_{\mu, \text{obs}}}^{\infty} f(q_\mu | \mu) dq_\mu \quad (4.9)$$

where $q_{\mu, \text{obs}}$ is the test-statistic value observed in data and $f(q_\mu | \mu)$ is the probability density function of q_μ under the assumption μ . In practice, the measured data provide the observed counts n_i in Eq. 4.6. Alternatively, f can then be determined by performing numerous pseudo-experiments, often called *toys*, where repeated samples of pseudo-data are generated, and the resulting values of q_μ are recorded. Small p -values indicate that the observed data are unlikely if μ were correct, prompting rejection of that hypothesis. In the asymptotic (large sample) limit, Wilks’ theorem [155] and the Wald approximation [156] imply that f follows a χ^2 distribution.

In particle physics, one often translates the p -value into a significance Z , representing the number of standard deviations a Gaussian-distributed variable would lie above its mean to yield the same upper-tail probability p . Formally, these quantities relate via

$$Z = \Phi^{-1}(1 - p), \quad (4.10)$$

where Φ^{-1} denotes the inverse of the unit Gaussian’s cumulative distribution. A standard convention sets the discovery threshold at $Z = 5(5\sigma)$, equivalent to $p = 2.87 \times 10^{-7}$. In contrast, excluding a

signal hypothesis typically employs a test size of $p = 0.05$, corresponding to a 95% confidence level and a significance of $Z = 1.64$.

The formalism set out in this section finds two primary applications: discovery and limit setting. Each of these use cases is addressed separately in the discussion that follows.

4.5.3 The discovery case

In a search for a new signal, the goal is to exclude the background-only hypothesis, which corresponds to $\mu = 0$. Here, only an excess of observed events (relative to the expectation from background processes) is regarded as evidence against $\mu = 0$. Consequently, the test statistic used is

$$q_0 = \begin{cases} -2 \ln \lambda(0), & \hat{\mu} \geq 0 \\ 0, & \hat{\mu} < 0 \end{cases} \quad (4.11)$$

If the observed event count falls below the background prediction, $\hat{\mu} < 0$, leading to $q_0 = 0$. Conversely, when more events than expected are recorded, $\hat{\mu} > 0$ and q_0 grows in proportion to the tension between the data and the $\mu = 0$ hypothesis. A strongly negative $\hat{\mu}$ is typically interpreted as indicating that the background model itself may be flawed.

The degree of incompatibility with is given by the p -value:

$$p_0 = \int_{q_{0, \text{obs}}}^{\infty} f(q_0 | 0) dq_0, \quad (4.12)$$

mirroring the structure of Eq. 4.8 but applied to q_0 . Figure 4.14 provides a graphical illustration of p_0 . One can then convert this p_0 into the corresponding Gaussian significance via

$$Z_0 = \Phi^{-1}(1 - p_0) = \sqrt{q_0}. \quad (4.13)$$




Figure 4.14: Illustration of the p -value used to reject the background-only hypothesis [67].

4.5.4 Setting limits

When setting an upper limit on the signal strength, one instead tests the signal-plus-background hypothesis. The corresponding test statistic is

$$q_\mu = \begin{cases} -2 \ln \lambda(\mu), & \hat{\mu} \leq \mu \\ 0, & \hat{\mu} > \mu \end{cases} \quad (4.14)$$

By design, an observed excess ($\hat{\mu} > \mu$) does not indicate lower compatibility with μ than the data itself. Hence, such an excess does not enter the test's rejection region. To reject the signal-plus-background hypothesis, one computes the p -value

$$p_\mu = \int_{q_{\mu,\text{obs}}}^{\infty} f(q_\mu | \mu) dq_\mu. \quad (4.15)$$

A confidence level (CL) for μ emerges by collecting all μ values whose p_μ exceeds $1 - \text{CL}$, typically using $\text{CL} = 95\%$. The upper boundary of this interval, μ_{up} , serves as the upper limit: $\mu \leq \mu_{\text{up}}$ at 95% CL implies that μ_{up} is the highest μ for which the p -value remains above 0.05.

Under the asymptotic approximation [156] the corresponding significance reduces to

$$Z_\mu = \Phi^{-1}(1 - p_\mu) = \sqrt{q_\mu}. \quad (4.16)$$

The CL_s method

Solving Eq. 4.15 directly can produce a reported upper limit that is not typically adopted. If the data exhibit a downward fluctuation, that upper limit might become arbitrarily small (though it must still include $\mu = 0$), thus overstating the experimental sensitivity. To avoid such artificially stringent limits, the modified frequentist method [68] is used. Instead of relying on p_μ , this approach utilises the quantity

$$p'_\mu = \frac{p_\mu}{1 - p_b}, \quad (4.17)$$

where p_b is the p -value derived from the same test statistic under the background-only hypothesis

$$p_b = 1 - \int_{q_{\mu,\text{obs}}}^{\infty} f(q_\mu | 0) dq_\mu. \quad (4.18)$$

By dividing the p -value by $(1 - p_b)$, the CL_s approach effectively penalises the exclusion of signals in regions with limited sensitivity. When the background-only and signal-plus-background distributions, $f(q | b)$ and $f(q | s + b)$, are clearly separated, as shown in Figure 4.15, then $1 - p_b$ remains near unity. Under such circumstances, CL_s yields an upper limit comparable to that obtained from a standard p -value. However, in regions of poor signal sensitivity where the two distributions overlap significantly, $1 - p_b$ becomes small, increasing p'_μ and thereby diminishing the degree of exclusion. In this way, the CL_s method safeguards against excluding signals that exhibit very low cross-sections, to which the experiment has negligible sensitivity.

The CL_s upper limit on μ follows from the condition $p'_\mu = 0.05$. A known caveat of this method is that it leads to over-coverage, meaning the actual confidence level exceeds 95%.




Figure 4.15: Representation of the p -values used in the computation of p'_μ . In the case of $\mu = 1$, p_μ is represented as p_{s+b} [152].

Expected upper limit

To characterise experimental sensitivity, one often examines the expected (or median) significance and upper limit derived from numerous identical experiments. Intuitively, the expected upper limit reflects what one would anticipate obtaining if the background-only hypothesis were true. However, because data can fluctuate, it is also valuable to understand how this variability affects the upper limit.

A straightforward method to find the expected limit and its uncertainty bands is to generate multiple background-only toy MC experiments, each with randomly varied nuisance parameters [154]. By recording μ_{up} for each toy experiment, one obtains a distribution whose median represents the expected limit, while additional quantiles map out the corresponding uncertainty.

An alternative, more efficient approach relies on a single MC simulated sample known as the Asimov dataset. This artificial sample uses pseudo-data set to their expected values. In notation similar to Eq. 4.5, the Asimov dataset is constructed such that $n_i = E[n_i] = \mu' S_i + B_i$, meaning the pseudo-data precisely equal their expectation for a chosen μ' . An important result [152] of this definition is that the variance σ^2 of $\hat{\mu}$ in the Asimov dataset satisfies

$$\sigma^2 = \frac{(\mu - \mu')^2}{q_\mu}. \quad (4.19)$$

where μ corresponds to the hypothesis being tested. In computing the expected limit, the Asimov dataset is meant to represent the background only, so $\mu' = \hat{\mu} = 0$. As illustrated in Figure 4.16, the expected limit $\mu_{\text{up}}^{\text{med}}$ corresponds to the median of $f(q_\mu|0)$, and, thanks to the properties of the Asimov dataset, it can be found by solving

Here μ denotes the hypothesis under investigation. To obtain the expected limit, the Asimov dataset is treated as a background-only scenario, thus setting $\mu' = \hat{\mu} = 0$. As depicted in Figure 4.16, the

expected limit $\mu_{\text{up}}^{\text{med}}$ corresponds to the median of $f(q_\mu|0)$. Exploiting the special properties of the Asimov dataset, $\mu_{\text{up}}^{\text{med}}$ is found by solving

$$\text{med} \left[Z_\mu = \Phi^{-1}(1 - 0.05) \right] = \sqrt{q_\mu}. \quad (4.20)$$

Eq. 4.20 leverages the connection between the test statistic q_μ of the Asimov dataset and the hypothesis significance Z_μ . By selecting a significance level that corresponds to a p -value of 0.05, one determines the required $\mu_{\text{up}}^{\text{med}}$. This upper-limit value can then be substituted into Eq. 4.19 to obtain σ_A^2 .




Figure 4.16: (a) Depiction of the p -value associated with the median q_μ under the assumption $\mu' = 0$ [67]. (b) Comparison of the upper limits derived from MC toy experiments (histogram) with the median and its uncertainty bands obtained using the Asimov dataset (dashed lines) [152].

Search for Charged-Lepton-Flavour Violation in Top-Quark Processes

This chapter presents a model-independent search for exotic charged-lepton-flavour-violating interactions, which involve top-quark production and decay, treated as a $e\mu qt$ interaction vertex, where $q = \{u, c\}$. The search is based on the full $\sqrt{s} = 13$ TeV proton–proton collision dataset recorded by ATLAS during Run 2 of the LHC, amounting to an integrated luminosity of 140 fb^{-1} . Details of the signal processes and the targeted final state are provided in Section 5.1. The simulated event samples used in the analysis are described in Section 5.2, while Section 5.3 summarises the object and event selection criteria. The overall analysis strategy is outlined in Section 5.4, and the multivariate analysis employed for signal discrimination is discussed in Section 5.5. Systematic uncertainties affecting both the signal and background predictions are addressed in Section 5.6. Finally, the results of the search are presented in Section 5.7.

5.1 The signal processes

In the first process considered in this search, involving the charged-lepton-flavour-violating decay of the top quark, the signal, hereafter referred to as the ‘decay signal’, consists of a $t\bar{t}$ event, where one of the top quarks decays via the cLFV mode, $t \rightarrow e^\pm \mu^\mp q$, while the other decays hadronically as $t \rightarrow W^+(\rightarrow q\bar{q}')b$ in accordance with the SM. Throughout this analysis, charge-conjugate decays of the top quark, $\bar{t} \rightarrow e^\mp \mu^\pm \bar{q}$ and $\bar{t} \rightarrow W^-\bar{b}$, are implicitly included. A representative Feynman diagram of this process is shown in Figure 5.1(a).

In the second process, which involves charged-lepton-flavour-violating top-quark production, the signal, referred to as ‘production signal’, corresponds to the single-top quark production with a cLFV vertex, $gq \rightarrow e^\pm \mu^\mp t$, where the produced top quark undergoes a hadronic decay as predicted by the SM. This process is illustrated in Figures 5.1(b) and 5.1(c). Compared to purely leptonic W decays, hadronic decays offer larger cross-sections and are therefore of particular interest in experimental searches.

In both processes, the final state of the interaction is characterised by an opposite-sign electron–muon pair and multiple jets, including at least one b -jet. The identification of hadronic jets is described in Section 4.3.4.

The signal processes are generated by extending the SM Lagrangian with the EFT operators corresponding to cLFV top-quark production and decay, as discussed in Section 2.3.2. These




Figure 5.1: Example Feynman diagrams for the EFT signal processes. Diagram (a) shows top-quark pair production in which one top undergoes the decay $t \rightarrow e\mu q$ whereas (b) and (c) show the single top-quark production t - and s -channel diagrams, respectively, contributing to the $gq \rightarrow te\mu$ process. The hatched circle represents the cLFV vertex.

interactions are implemented in a UFO model created with FEYNRULES 2.0, based on the `dim6top` model [157]. The UFO files encode the Feynman rules for all vertices allowed by the input Lagrangian.

Events are simulated at NLO in QCD using MADGRAPH 2.9.5 for the hard process, interfaced with PYTHIA 8.306 for parton showering and hadronisation. The simulation is performed in the five-flavour scheme, where all quark masses are set to zero except for the top quark. The renormalisation and factorisation scales, (μ_R, μ_F) , are defined dynamically: for decay diagrams they equal the centre-of-mass energy of the incoming partons, while for production diagrams they correspond to half the sum

of the transverse masses of all final-state particles and partons, with $m_T = \sqrt{m^2 + p_T^2}$. Top quark spin correlations are preserved in the decay products, and the top quarks are allowed to be off-shell. The NNPDF31_NLO_AS_0118 PDF set [81] is used, and PYTHIA is configured with the A14 tune [109].

All EFT coefficients are fixed to a common WC value of 1.0 and an energy scale $\Lambda = 1$ TeV. This simplifying choice reflects the absence of theoretical motivation for assigning different couplings, the lack of stringent experimental bounds, and the restriction to a single signal sample. The EFT operators affect only the charged-lepton-flavour-violating $e\mu qt$ vertices, while all other processes follow the SM. A detailed discussion of the individual operator contributions is presented in Section 5.7, where the analysis results are interpreted accordingly.

The signal MC simulated sample is generated allowing charged leptons of all three flavours (e, μ, τ) at ME level. Events are selected requiring an $e\mu$ final state at reconstruction level, such that contributions from leptonic τ decays are included in the signal acceptance. As a consequence,

the analysis is sensitive both to direct $t \rightarrow e\mu q$ transitions and to processes involving intermediate τ leptons decaying leptonically. For the interpretation in terms of WCs, events with intermediate τ -leptons are excluded, since the EFT operators carry explicit lepton-flavour indices and the analysis targets only operators involving electrons and muons.

In addition, to evaluate parton-shower uncertainties, an additional set of signal samples is produced with MADGRAPH 2.9.5 for the hard process interfaced to HERWIG 7.1.6 for showering and hadronisation. These events are also generated at NLO in QCD with the same NNPDF31_NLO_AS_0118 PDF set.

Cross-sections of the signal processes

According to Ref. [47], the decay cross-section is estimated according to:

$$\sigma_{\text{cLFV}}^{\text{decay}} = \sigma_{t\bar{t}} \times 2 \times \mathcal{B}(t \rightarrow W(q\bar{q}')b) \times \frac{\Gamma(t \rightarrow \ell\ell'q)}{\Gamma_t}. \quad (5.1)$$

The combinatorial factor 2 is due to the top-quark interchangeability, the hadronic branching ratio of the top quark is $\mathcal{B}(t \rightarrow W(q\bar{q}')b) \approx 0.67$ [158], the decay width is calculated via Eq. 2.10 and $\Gamma_t = 1.51$ GeV is the top decay width at LO. The $t\bar{t}$ production cross-section is calculated at NNLO in QCD, as discussed below in Section 5.2.

In Table 5.1, the cross-sections of the different signal processes considered in this thesis, as given by MC simulation, are presented. The cLFV in decay signal and in production signal processes are considered separately, and are further broken down by the EFT operator generating the interaction. The total cross section for each signal process (decay and production) is obtained by simply adding the separate contributions shown. For the decay signal, the cross-section value matches a calculation done in Ref. [47], and multiplying by a factor 3 to account for any lepton flavour in the cLFV decay. The cross-sections for the production signal process were taken directly from the MC generator since no analytical expression is present in the literature.

$i, j = e, \mu, \tau$	Cross-section $\sigma_{\text{scale}}^{\pm \text{scale}} \pm \text{PDF (fb)}$		
	$\mathcal{O}_{lq/eq/lu/eu}^{-(ijk3)}$	$\mathcal{O}_{lequ}^{1(ijk3)}$	$\mathcal{O}_{lequ}^{3(ijk3)}$
Production $\ell\ell'tu$	$118_{-19}^{+24} \pm 1$	$101_{-16}^{+21} \pm 1$	$2150_{-320}^{+410} \pm 20$
Production $\ell\ell'tc$	$7.9_{-1.0}^{+1.2} \pm 1.6$	$6.1_{-0.8}^{+1.0} \pm 1.5$	$153_{-18}^{+21} \pm 29$
Decay $\ell\ell'tq$	$6.9_{-1.3}^{+1.8} \pm 0.1$	$3.46_{-0.66}^{+0.90} \pm 0.03$	$166_{-32}^{+43} \pm 2$

Table 5.1: Theoretical cross-sections for single top quark production and $t\bar{t}$ decays through cLFV interactions for vector, scalar and tensor EFT operators. The cross-section for each of the vector operators is the same. These values were calculated for a top quark mass of 172.5 GeV, a top quark decay width of 1.51 GeV, a new physics scale of $\Lambda = 1$ TeV and WC values of 1.0. The scale variation and PDF uncertainties are given.

The decay contribution is significantly smaller than that arising from the production mode. In the decay channel, the contributions from up- and charm-quark interactions are comparable, as can be verified from the MC generator information of the produced particles. In contrast, the production signal is modelled using separate samples for up- and charm-quark initiated processes, with the up-quark contribution clearly dominating owing to the larger up-quark parton density in the proton. For both production and decay topologies, the tensor operator provides the leading contribution to the signal yield.

5.2 Background Monte-Carlo simulation samples

The dominant background contributions to this search arise from Standard Model top-quark production, in particular $t\bar{t}$ and single-top tW processes. In both cases, the presence of two W bosons in the final state allows for dileptonic decays, where one W decays to an electron and the other to a muon, either directly or via intermediate τ -leptons decaying leptonically. Together with the associated b -jets and missing transverse momentum from neutrinos, these events closely mimic the reconstructed final state targeted by the signal.

Subdominant backgrounds include $t\bar{t}\gamma$ and $t\bar{t}V$ ($V = H, W, Z$) production, which can yield additional leptons or photons that are not reconstructed or fail identification, resulting in an apparent $e\mu$ final state. Further smaller contributions come from diboson and triboson processes, which can produce opposite-flavour lepton pairs in association with jets and missing transverse momentum. These are hereafter collectively called “other” background.

A number of additional rare SM processes that can be reconstructed as dilepton final states are also taken into account. These include the production of three or four top quarks, as well as processes involving the production of a weak boson in association with jets or with a Higgs boson. They are collectively referred to as “rare” in figures and tables used later in this chapter. Although their cross-sections are significantly lower than those of $t\bar{t}$ and tW , these processes are nonetheless considered due to their potential to populate the signal region through similar visible final states.

$t\bar{t}$ and $t\bar{t}\gamma$ The primary background for this analysis arises from $t\bar{t}$ events, specifically those producing dilepton final states ($\ell\ell' + \text{jets}$). This process is expected to yield at most two prompt leptons. The samples were generated using the POWHEGBox v2 generator, which provides NLO matrix elements in the strong coupling constant α_S and employs the NNPDF3.0 [159] parton distribution function.

The h_{damp} ¹ parameter, which regulates the first high- p_T gluon radiation beyond the Born approximation, was set to $1.5 m_t$, where m_t is the top-quark mass. This value was optimised based on studies of observables sensitive to radiation emission from the $t\bar{t}$ system, using ATLAS data at centre-of-mass energies of $\sqrt{s} = 7, 8$, and 13 TeV [160].

The QCD renormalisation and factorisation scales are set to $\sqrt{m_t^2 + p_{T,t}^2}$, where $p_{T,t}$ is the transverse momentum of the top quark. The events are interfaced with PYTHIA 8.230 for the parton shower and hadronisation, utilising the A14 tune and the NNPDF2.3 LO PDF set.

To assess the impact of different parton shower and hadronisation models, the nominal $t\bar{t}$ sample was compared with a sample generated by the POWHEGBox v2 but interfaced with HERWIG 7.1.3 [108].

The cross-section for $t\bar{t}$ production was calculated at NNLO in QCD, including the resummation of soft-gluon emissions to next-to-next-to-leading logarithmic (NNLL) accuracy. For a top-quark mass of 172.5 GeV the cross-section is calculated using TOP++ 2.0 to be $\sigma(t\bar{t})_{\text{NNLO+NNLL}} = 832 \pm 51$ pb [161–167].

Additionally, events containing a $t\bar{t}$ pair can also involve photon radiation from an incoming charged parton, the top quarks themselves, or any charged particle in the final state. A dedicated sample of $t\bar{t}\gamma$ events was modelled at LO in QCD using MG_AMC@NLO [99] with the NNPDF2.3 PDF set, interfaced with PYTHIA 8 using the A14 tune. To avoid double-counting events, an overlap removal procedure [168] was applied to the $t\bar{t}$ and $t\bar{t}\gamma$ samples, ensuring they represent orthogonal regions of phase space based on the presence of a high- p_T photon identified through ‘truth’ information. The decays of bottom and charm hadrons were simulated using the EVTGEN 1.6.0 program.

¹ h_{damp} is related to the choice of resummation scale, separating hard emissions described by the ME calculation from soft emissions which are included instead by the parton-shower generator.

Single top The production of a single top quark alongside a W boson (tW) is anticipated to yield a maximum of two prompt leptons. This phenomenon is simulated with the POWHEGBox v2 generator, utilising the NNPDF3.0 NLO PDF set. To manage interference with $t\bar{t}$ production, a diagram removal scheme (DR) [169] is employed, which eliminates events where the W boson and an associated b -quark create a top quark. These events are interfaced to PYTHIA 8.230, using the A14 tune and the NNPDF2.3 LO PDF set. The inclusive cross-section is corrected to match the theoretical prediction computed at NLO in QCD, including NNLL soft-gluon corrections [170, 171]. Systematic uncertainties in the tW background are evaluated by comparing the diagram removal and diagram subtraction schemes; the former removes resonant $t\bar{t}$ effects at the amplitude level, whereas the latter operates at the cross-section level [172, 173]. This procedure is used to assess the corresponding theoretical uncertainty, as outlined in Section 5.6.

Additional simulations for the production of top quarks in association with bosons, specifically tZ and tWZ , are also included. In both scenarios, the Z boson is required to decay leptonically. Minor background contributions arise from single top production in the s - and t - channels, as these processes are anticipated to yield only one prompt lepton. Consequently, all single top samples, except for tW , will henceforth be designated as $t + X$.

The details described in this section are summarised in Table 5.2.

5.3 Object and event selection

The physics objects considered in this analysis are electrons, muons, and jets. Their reconstruction has been described in Section 4.3; here, we specify the identification criteria applied and the conditions defining the event selection. As illustrated in Figure 5.1, the cLFV signal processes produce an opposite-charge electron–muon pair accompanied by the products of a SM top-quark decay. This search focuses on the case where the SM top quark decays as $t \rightarrow W^+(q\bar{q}')b$, leading to a dilepton final state.

The data collected are subdivided into luminosity blocks, and only events recorded with the full detector operational are analysed. The selection is performed using the Good Run List provided by ATLAS. In addition, quality requirements are applied to reject events from periods of poor detector performance or those unlikely to originate from a hard-scatter interaction.

The presence of a reconstructed primary vertex is required. In practice, it is sufficient that at least two tracks with $p_T > 0.5$ GeV intersect in proximity of the beam crossing region [119, 174]. The vertex with the highest $\sum_{\text{tracks}} p_T^2$ is taken as primary vertex.

Furthermore, events are required to pass single-electron or single-muon trigger chains [78]. For each data-taking period, the lowest unprescaled² single-lepton triggers are used, supplemented by higher-threshold triggers with looser identification or isolation requirements to maintain efficiency at high E_T and p_T . At least one such trigger must fire for an event to be accepted. These trigger selections are applied consistently to both data and MC simulated samples, with each simulated event assigned a run number that determines the appropriate trigger and pile-up corrections. The increase in instantaneous luminosity delivered by the LHC from 2016 onwards necessitated raising the trigger thresholds to avoid prescales.

² A prescale implies that only a fraction of the events passing a given trigger are recorded. Such triggers, often referred to as *support triggers*, are typically associated with specific data-taking conditions or used for monitoring, testing low- E_T threshold triggers (with loose or no isolation), and for performance or background studies. In contrast, the *primary triggers* employed for physics measurements generally operate without prescales.

Sample	Generator	Parton shower	Accuracy (ME, PDF)
Signal	MG5_aMC	PYTHIA 8	NLO, NLO
$t\bar{t}(\rightarrow \ell\ell + \text{jets})$	POWHEGBox	PYTHIA 8	NLO, NLO
tW	POWHEGBox	PYTHIA 8	NLO, NLO
$t\bar{t}\gamma$	MG5_aMC	PYTHIA 8	LO, LO
$t\bar{t}H, H \rightarrow WW, ZZ, \tau\tau$	POWHEGBox	PYTHIA 8	NLO, NLO
$t\bar{t}Z(Z \rightarrow \ell\ell)$	MG5_aMC	PYTHIA 8	NLO, NLO
$t\bar{t}W$	SHERPA 2.2.10	–	NLO, NNLO
tWZ	MG5_aMC	PYTHIA 8	NLO, NLO
$WZ(\rightarrow 2\ell 2q)$	SHERPA 2.2.1	–	(N)LO, NNLO
$WZ(\rightarrow 3\ell\nu)$	SHERPA 2.2.2	–	(N)LO, NNLO
$ZZ(\rightarrow 2\ell 2q)$	SHERPA 2.2.1	–	(N)LO, NNLO
$Z(\rightarrow \ell\ell) + \text{jets}$	SHERPA 2.2.1	–	(N)LO, NNLO
$Z(\rightarrow \ell\ell) + \gamma$	SHERPA 2.2.4	–	(N)LO, NNLO
$ZZ(\rightarrow 4\ell)$	SHERPA 2.2.2	–	(N)LO, NNLO
VVV	SHERPA 2.2.2	–	(N)LO, NLO
$2\ell 2\nu$	SHERPA 2.2.2	–	(N)LO, NNLO
$3t$	MG5_aMC	PYTHIA 8	NLO, NLO
$4t$	MG5_aMC	PYTHIA 8	NLO, NLO
$t\bar{t}WW$	MG5_aMC	PYTHIA 8	NLO, NLO
WH, ZH	–	PYTHIA 8	LO, LO

Table 5.2: Summary of the basic generator parameters used to simulate the various processes. The last column indicates the perturbative accuracy of the ME and PDF calculations; when different approximations are used for different jet multiplicities within a given sample, this is denoted in parentheses, *e.g.* (N)LO for WZ . Only nominal samples are shown.

Electron triggers [175] select calorimeter clusters matched to inner detector tracks, with candidates required to satisfy a likelihood-based identification. In 2015, the primary electron trigger required `MediumLH` identification and $E_T > 24$ GeV. In 2016–2018, the selection was tightened to `TightLH` identification plus isolation, with $E_T > 26$ GeV. To prevent efficiency losses at high E_T , additional triggers were available throughout the data-taking period, requiring `MediumLH` electrons with $E_T > 60$ GeV or `LooseLH` electrons with $E_T > 120$ GeV (raised to 140 GeV in 2016–2018).

Muon triggers [176] rely on tracks reconstructed in both the muon spectrometer and the inner detector. In 2015, the primary trigger required a muon with $p_T > 20$ GeV and a loose isolation criterion. From 2016 onward, the threshold was increased to $p_T > 26$ GeV with tighter isolation requirements. An additional trigger without isolation requirements but with $p_T > 50$ GeV was available in all periods to preserve efficiency for high- p_T muons.

For reference, the explicit names of the triggers used in each year are listed in Table 5.3.

	Electron triggers	Muon triggers
2015	<code>HLT_e24_lhmedium_L1EM20VH</code>	<code>HLT_mu20_iloose_L1MU15</code>
	<code>HLT_e60_lhmedium</code>	<code>HLT_mu50</code>
	<code>HLT_e120_lhloose</code>	
2016–2018	<code>HLT_e26_lhtight_nod0_ivarloose</code>	<code>HLT_mu26_ivarmedium</code>
	<code>HLT_e60_lhmedium_nod0</code>	<code>HLT_mu50</code>
	<code>HLT_e140_lhloose_nod0</code>	

Table 5.3: Trigger selection for all data taking years used in this analysis.

The trigger naming convention is as follows. At the high level, `HLT` denotes a High-Level Trigger; `e (mu)` specifies an electron (muon) trigger, and the following number indicates the transverse energy threshold. The identification working point is then given (`tight`, `medium`, or `loose`), with `lh` denoting a likelihood-based identification. The suffix `nod0` indicates that no transverse impact parameter requirement is applied, while `i` specifies the isolation working point (*e.g.* `ivarloose` corresponds to a variable-cone isolation).

At Level-1, `L1` denotes the corresponding hardware trigger. Here, `EM (MU)` specifies an electromagnetic (muon) trigger, and the subsequent number indicates the transverse energy threshold. Additional letters modify the selection: `V` introduces a pseudorapidity-dependent threshold, `H` indicates hadronic isolation, and `I` denotes electromagnetic isolation. When no explicit `L1` trigger is given, the default is `L1MU20` for muons and, for electrons, an isolated electromagnetic cluster with $E_T > 22$ GeV.

In addition, each event is required to contain at least one reconstructed lepton matched to a fired trigger, a procedure known as ‘trigger matching’. In the ATLAS experiment, events are recorded only if they satisfy predefined trigger selections based on objects reconstructed online by the trigger system. Trigger matching ensures that the offline reconstructed lepton responsible for the physics analysis corresponds geometrically to the lepton candidate that initiated the trigger decision. This matching is typically performed by requiring a small angular separation between the offline lepton and the trigger-level object in the (η, ϕ) plane. The requirement guarantees that the event entered the recorded dataset because of the analysed lepton, thereby ensuring a consistent definition of trigger efficiencies and avoiding biases arising from unrelated trigger activity.

Since the reconstructed p_T and E_T used in the analysis can differ from the corresponding online quantities applied in the trigger thresholds, the matched lepton must exceed the nominal threshold by an additional margin: 1 GeV for electrons and muons [175, 176]. These requirements ensure that the trigger is fully efficient for selected lepton candidates, while residual differences between simulation and data can be corrected with trigger scale factors.

5.4 Analysis strategy

Events are required to pass a pre-selection consisting of exactly one opposite-charge electron–muon pair and at least one b -tagged jet; a veto on additional baseline leptons is not applied.. This ensures the presence of the objects necessary to accept the signal. The full set of criteria, including the object definitions listed in Section 4.3, is summarised in Table 5.4.

Events are then classified into two regions: the SR and the CR. The SR requires exactly one b -tagged jet, while the CR requires at least two, as shown in Table 5.4. The SR is designed for high signal acceptance, even at the expense of background rejection, to provide sufficiently large samples for the multivariate discriminator. The CR selection targets the dominant background, SM dileptonic $t\bar{t}$. Its normalisation is constrained in the fit using this region. At the final stage of the fit, an additional requirement on the jet multiplicity ($N_{\text{jets}} \geq 2$) is imposed in both region definitions to suppress background contamination, which proved to improve the fit sensitivity.

The distributions of H_T , the scalar sum of lepton and jet transverse momenta in both regions are shown in Figure 5.2. Although data are displayed here, it should be noted that the SR was blinded during the development of the analysis. In fact, all regions with a signal contribution exceeding 10% of the background were kept blinded; in practice, this applied only to the SR. Data in these regions were considered only after the analysis strategy was finalised, with selections and corrections fixed. This convention prevents potential bias from examining data in signal-sensitive regions while background corrections and systematic uncertainties are being determined. It can also be seen that the CR is very pure in $t\bar{t}$ background.

The description of the data by the MC prediction in the CR was deemed adequate at that stage of the analysis. The residual slope in the ratio plot observed at large H_T arises from the known limitations in the modelling of high- p_T $t\bar{t}$ production and is treated as a systematic uncertainty, as detailed in Section 5.6.

Among the signal hypotheses, cLFV single-top production provides the dominant contribution, with a cross-section significantly larger than that of cLFV $t\bar{t}$ decays. Owing to the harder lepton p_T spectrum, these events populate the high- H_T region, as illustrated in Figure 5.2, where the background contributions are suppressed. This kinematic separation is therefore exploited for signal discrimination using a multivariate approach.

The simulation and data event yields in each region, separated by physics process, are reported in Table 5.5. All EFT signal samples are accounted for inclusively and each one is normalised to a cross-section corresponding to WCs defined in Section 2.2 all set to $|C| = 1.0$.

The expected yields from simulated samples are obtained by summing the event weights provided by the MC generators, with the associated statistical uncertainty computed as the square root of the sum of the squared weights. Each event weight is constructed as the product of factors accounting for the integrated luminosity of the relevant data-taking period, the process cross-section, normalised to

Description	Value
Object definition	
Electron isolation	PLVTight
Electron ID	TightLH
Muon isolation	PLVTight
Muon ID	Medium
Baseline lepton definition	$p_T > 10 \text{ GeV}, \eta < 2.5$
Baseline jet definition	$p_T > 25 \text{ GeV}, \eta < 2.5$
Preselection cuts	
Number of leptons	$N_\ell = 2$
Lepton flavours	$1e1\mu$
Product of leptons charge	$\prod q_i = -1$
z_0 impact parameter	$ z_0 \sin \theta_{\text{lep}} < 0.5 \text{ mm}$
Number of jets	$N_{\text{jets}} \geq 1$
Number of b -jets (DL1r @ 70%)	$N_{b\text{-jets}} \geq 1$
SR cuts	
Number of b -jets	$N_{b\text{-jets}} = 1$
Number of jets*	$N_{\text{jets}} \geq 2$
CR cuts	
Number of b -jets	$N_{b\text{-jets}} > 1$
Number of jets*	$N_{\text{jets}} \geq 2$

Table 5.4: Description of the object definition, preselection cuts and region definitions used in the analysis. *The $N_{\text{jets}} \geq 2$ cut was only included for the fit.




Figure 5.2: Comparison of the H_T distributions in data and MC predictions for the SR (a) and CR (c). Corresponding plots with a logarithmic scale, (b) and (d), are also provided to enhance the visualisation. Statistical uncertainties are shown for the MC prediction (hatched bands) and for the data (error bars), with overflow events included in the final bin. The lower panel displays the ratio of data to the total MC prediction, with error bars indicating the statistical uncertainty. The hatched band in the ratio panel represents the statistical uncertainty on the MC prediction.

NLO accuracy using the appropriate k -factor³, and the finite size of the simulated sample. Additional corrections are applied to account for generator-scale and PDF variations, pile-up and jet-to-vertex association, b -tagging efficiencies, and lepton reconstruction and identification efficiencies.

After the region definitions, the analysis strategy goes as follows. A set of kinematic observables is used as input to multivariate discriminants based on boosted decision trees, which are trained in the SR to separate signal from background, as detailed in Section 5.5. The presence of a signal is then assessed via a profile likelihood fit performed simultaneously in the SR and CR, with the systematic uncertainties described in Section 5.6 incorporated as nuisance parameters. The extraction of the signal strength from a shape-based fit to the BDT output is discussed in Section 5.7.

³ The k -factors rescale the leading-order cross-sections to account for higher-order corrections in the ME calculation.

Process	SR	CR
$t\bar{t}$	390730 ± 120	240250 ± 100
tW	41340 ± 80	7557 ± 32
$t\bar{t}\gamma$	1252 ± 7	699 ± 5
Z +jets	1070 ± 50	70 ± 7
Rare	1043 ± 8	82.8 ± 1.4
$t\bar{t}W$	489.3 ± 1.9	308.7 ± 1.5
$t\bar{t}H$	232.4 ± 0.4	340.0 ± 0.4
$t\bar{t}Z$	155.0 ± 1.0	99.3 ± 0.8
$Z + \gamma$	53 ± 8	2.6 ± 1.8
W +jets	35 ± 4	3.2 ± 0.6
$t + X$	28.8 ± 1.2	4.8 ± 0.5
VVV	3.66 ± 0.06	0.193 ± 0.014
Total background	436430 ± 150	249410 ± 100
Signal (decay)	2996 ± 12	241 ± 4
Signal (production)	92910 ± 320	6020 ± 80
Total signal	95910 ± 320	6260 ± 80
S/B	$(21.97 \pm 0.07)\%$	$(2.510 \pm 0.032)\%$
Total MC	532300 ± 400	255680 ± 130
Data	425000 ± 700	242400 ± 500
Data / Bkg.	0.9738 ± 0.0015	0.9718 ± 0.0020

Table 5.5: Pre-fit event yields in the analysis regions. The signal is normalised by setting the WCs equal to 1.0. Uncertainties shown are statistical only.

5.5 Multivariate analysis

An initial signal region was defined with the variable $H_T < 3$ TeV providing reasonable separation. However, a multivariate approach was adopted instead, as it can exploit correlations among several variables in a multidimensional space. The classifier chosen is a boosted decision tree.

5.5.1 The boosted decision tree

Multivariate analysis algorithms are supervised machine-learning methods designed to predict the outcome of a process from a set of input variables. When the output is continuous, the task is referred to as a *regression* problem; when it is discrete, it becomes a *classification* problem. The input variables, also known as *input features* in pattern-recognition contexts, are used to predict the output.

Decision trees are often described in the literature as effective “out-of-the-box” classifiers, as they typically require minimal tuning to achieve good performance. This robustness stems from their simple training procedure, in which each node split corresponds to a one-dimensional optimisation. In addition, decision trees are relatively insensitive to the presence of input variables with limited discriminating power.

In the following, the set of input variables is denoted by $\vec{x}$, which in general represents a vector of p features. Each feature corresponds to a collection of N measured observations. The full dataset can therefore be represented as a matrix $X \in \mathbb{R}^{N \times p}$, where each row corresponds to one observation (or event) and each column to one input variable. For supervised learning tasks, the target variable is denoted by Y and is continuous in the case of a regression problem, or discrete in the classification scenario. The predictions returned by a trained model are typically written as $\hat{Y}$. Observed quantities are denoted using lowercase symbols; in particular, x_i refers to the i -th observed input vector and y_i to the corresponding observed value of the target variable.

To assess the performance and generalisation capability of a prediction model, the available set of observed pairs (x_i, y_i) is commonly divided into statistically independent subsets. A *training sample* is used to construct the model, while a *test* (or *validation*) *sample* is reserved for an unbiased evaluation of its predictive power on unseen data. This separation is essential to avoid overfitting, particularly for flexible models such as decision trees and their boosted variants. More generally, procedures that involve repeated or systematic partitions of the dataset into training and test subsets are referred to as *cross-validation* methods. These techniques provide robust estimates of model performance and are widely used in the optimisation of hyperparameters in boosted decision trees.

Loss function The discrepancy between the predicted output of a model and the true observed value is quantified by a *loss function*. In classification problems, the loss function provides a measure of the degree of misclassification. The specific form of the loss function determines how deviations between prediction and observation are penalised and plays a central role in the training of the model.

By averaging the loss function over an independent test sample, one obtains the test (or generalisation) error, which provides an estimate of the model’s performance on unseen data. Similarly, averaging the loss over the training sample defines the training error. As the complexity of the model increases, allowing it to capture increasingly intricate structures in the training data, the training error typically decreases monotonically. However, the test error does not necessarily follow the same behaviour. Instead, it usually reaches a minimum at an intermediate level of model complexity, beyond which it starts to increase.

A large discrepancy between the training and test errors indicates that the model has become overly sensitive to statistical fluctuations specific to the training sample rather than to the underlying data-generating process. Instead of learning systematic patterns that generalise to unseen data, the model begins to fit random noise and accidental correlations present only in the finite training dataset. Since these fluctuations are not representative of the population distribution, they do not appear in independent test samples, resulting in degraded out-of-sample performance. This phenomenon, known as *overfitting* or *overtraining*, is especially relevant for highly flexible models such as decision trees. Controlling model complexity and monitoring test performance are therefore essential for constructing robust boosted decision tree classifiers.

In many practical applications, classification problems are restricted to two classes. In this case, the task can be reformulated in terms of a quantitative output by interpreting the model prediction $\hat{y}$ as the

estimated posterior probability of one of the classes, with $\hat{y} \in [0, 1]$. A discrete class label is then assigned by applying a threshold, commonly chosen as $\hat{y} > 0.5$, although this value can be adjusted to optimise a given performance metric. This probabilistic formulation is particularly advantageous for boosted decision trees, which naturally provide continuous discriminant outputs.

The same approach can be generalised to multi-class problems with $K > 2$, where the model predicts a set of class probabilities satisfying $\sum_{k=1}^K \hat{y}_k = 1$. The predicted class label is then typically assigned to the class with the highest estimated probability. Such formulations allow for greater flexibility than the zero-one loss alone and form the basis for more sophisticated loss functions commonly employed in multi-class boosting algorithms.

The binary cross entropy (BCE) loss function is widely used in binary classification, where the target variable $y \in \{0, 1\}$. Given a predicted probability $p \in (0, 1)$ for the positive class, the BCE loss is

$$L_{\text{BCE}} = -[y \log p + (1 - y) \log(1 - p)], \quad (5.2)$$

which corresponds to the negative log-likelihood of a Bernoulli distribution. It penalises confident but incorrect predictions strongly, while giving smaller penalties to uncertain predictions.

In multi-class classification with mutually exclusive classes, the model typically outputs a vector of real-valued scores, known as *logits*, z_i , one for each class. These logits are the unnormalised outputs of the final layer and represent the model's relative preference for each class; they are not constrained to lie between 0 and 1 or to sum to unity. The logits are subsequently transformed into class probabilities using the softmax function, which normalises the scores into a probability distribution over all classes. The softmax is defined as

$$\text{softmax}(z_i) = \frac{e^{z_i}}{\sum_{j=1}^K e^{z_j}}. \quad (5.3)$$

This transformation ensures two key properties: each output probability lies in the interval (0,1), and the probabilities sum to one, so they form a valid categorical distribution. Intuitively, softmax amplifies the largest logits while suppressing smaller ones, making the most likely class stand out.

For a three-class problem ($K = 3$), the corresponding categorical cross entropy (CCE) loss is

$$L_{\text{CCE}} = - \sum_{i=1}^3 y_i \log p_i, \quad (5.4)$$

where p_i are the softmax probabilities and $y_i \in \{0, 1\}$ denotes the target indicator for class i , taking the value 1 for the correct class and 0 otherwise..

Classification trees In the following, the discussion is restricted to classification trees. A decision tree performs event classification through a sequence of binary splits, each corresponding to a node in the tree. The construction of the tree proceeds recursively: starting from a root node containing the full dataset, the events are divided into two subsets according to a splitting criterion. Each subset is then treated in the same manner, giving rise to a hierarchical structure that grows until a set of terminal nodes, or *leaves*, is reached. A schematic illustration of this procedure is shown in Figure 5.3. At each node, the algorithm selects the input variable x_i and the associated threshold that maximise the separation between signal (S) and background (B) events.




Figure 5.3: Sketch of a decision tree [177].

In the common case of a binary classification problem, the quality of a split is typically quantified using a measure of node *impurity*. A widely adopted choice is the *Gini index*, defined as

$$p(1 - p), \quad (5.5)$$

where $p = S/(S + B)$ denotes the purity of the node. The Gini index vanishes for perfectly pure nodes and reaches its maximum for an equal mixture of signal and background. During training, the optimal split is chosen by maximising the reduction in impurity, defined as the difference between the Gini index of the parent node and the weighted sum of the Gini indices of the two daughter nodes, with weights given by their relative event fractions.

Boosting Boosting extends the concept of a single decision tree to an ensemble of trees, often referred to as a *forest*. All trees in the ensemble are trained on the same underlying dataset, but with event weights that are iteratively modified during the training process. The individual trees are then combined to form a single, more powerful classifier, typically constructed as a weighted sum of the responses of the individual decision trees.

In this framework, each single decision tree acts as a *weak classifier*, meaning that its classification performance could be only marginally better than that of random guessing. The central idea of boosting is to apply these weak classifiers sequentially to a series of reweighted versions of the training sample, such that events that are misclassified by earlier trees receive increased weight in subsequent iterations. In this way, successive trees focus progressively on the more difficult-to-classify events, and the ensemble as a whole achieves significantly improved discrimination power.

In *gradient boosting*, an ensemble of decision trees is built sequentially, with each new tree improving upon the mistakes made by the previous ones. Instead of changing the weights of individual training events, as done in *adaptive boosting*, gradient boosting views learning as an optimisation problem in which the model is gradually refined to reduce a chosen loss function. At boosting each iteration m , a new decision tree is trained to predict the remaining errors of the current model, namely how its predictions should be adjusted to better match the true targets. The predictions of this new tree are

then added, with a small scaling factor, to the existing ensemble. Repeating this procedure allows each successive tree to correct the residual errors of the combined model, resulting in a progressive improvement of the classifier performance.

Formally, the boosted model can be written as an additive expansion,

$$\hat{Y} = \hat{f}(X) = \sum_{m=1}^M \alpha_m h_m(X), \quad (5.6)$$

where $h_m(X)$ denotes the individual decision trees, α_m their associated weights, and M the total number of boosting iterations, corresponding to the number of decision trees included in the ensemble.

The weights are not fixed in advance but are determined during training so that the overall model progressively improves its predictions. After a new decision tree is trained to describe the remaining errors of the current ensemble, its contribution is adjusted to achieve the best improvement in the chosen loss function. In practice, this means finding how strongly the new tree should be added to the existing model.

The updated prediction is obtained by adding the output of the new tree to the previous ensemble,

$$\hat{f}_m(X) = \hat{f}_{m-1}(X) + \nu \alpha_m h_m(X), \quad (5.7)$$

where α_m determines the importance of the new tree and $\nu \in (0, 1]$ is the learning rate. The learning rate reduces the impact of each individual tree, allowing the model to improve gradually and helping to prevent overfitting.

By repeating this procedure, each successive tree corrects the remaining errors of the ensemble, and the weights are learned adaptively so that every added tree contributes to improving the overall performance. Gradient boosting therefore provides a flexible and consistent framework for classification problems. A detailed introduction to gradient boosting can be found in Ref. [178].

Hyperparameters The performance of gradient-boosted decision trees are governed by a set of hyperparameters that control both the structure of the individual trees and the overall boosting procedure. Among the most important parameters is the number of boosting iterations (or *number of trees*), which determines the size of the ensemble. Increasing this number generally improves the expressive power of the model, but can lead to overfitting if not balanced by appropriate regularisation or early stopping.

The *learning rate* (also referred to as shrinkage) scales the contribution of each newly added tree to the ensemble. Smaller learning rates reduce the impact of individual trees, requiring a larger number of boosting iterations to achieve optimal performance, but often result in improved generalisation by making the training process more gradual and stable.

The *maximum depth* of the trees controls their structural complexity by limiting the number of successive splits from the root node to a leaf. Shallow trees act as weak learners and are commonly used in gradient boosting, as they capture simple structures while reducing sensitivity to statistical fluctuations.

The hyperparameters introduced above are primarily intended to control the model complexity and improve generalisation performance. Regularisation refers to techniques that limit the effective complexity of the model, such as reducing the learning rate, constraining the depth of individual trees, or requiring a minimum number of events in each leaf. These measures prevent individual trees from

learning statistical fluctuations specific to the training sample. Early stopping provides an additional safeguard by monitoring the model performance on an independent validation dataset and terminating the training once the validation error ceases to improve. Together, these strategies ensure that the ensemble achieves good predictive performance while maintaining robust generalisation to unseen data.

5.5.2 Implementation

The multivariate analysis was implemented using the Gradient Boosting (BDTG) model from XGBoost [179]. XGBoost provides an efficient and regularised realisation of gradient boosting, incorporating several algorithmic improvements aimed at enhancing both computational performance and robustness. In particular, it introduces explicit regularisation terms that penalise the complexity of individual trees, typically by constraining the number of leaves and the magnitude of the leaf weights. These regularisation mechanisms, together with the learning rate, subsampling of events and input variables, and early stopping criteria, play a central role in mitigating overfitting and improving generalisation, especially in high-dimensional feature spaces.

The default hyperparameter configuration already provided excellent performance without extensive tuning. The final training was therefore performed using gradient boosting [178] as the boosting algorithm, a learning rate of 0.3, an ensemble of 100 trees, and a maximum tree depth of 6.

As shown in Figure 5.2, particularly in the log-scale plots, the H_T distributions of the decay and production signals exhibit distinct features, motivating the use of a *multi-class boosted decision tree* with three categories: production signal, decay signal, and background. The BDTs were trained using MC simulated events within the SR. It was sufficient to use only $t\bar{t}$ events as the background sample during training, as this process constitutes the overwhelmingly dominant background in the analysis. Including additional minor background contributions was therefore unnecessary for the purpose of optimising the BDT performance.

Event weights derived from the Monte Carlo normalisation and correction factors were used during the BDT training. Since some simulated samples contain negative event weights originating from NLO generators, the absolute value of the event weight was supplied to the XGBoost training algorithm. The original signed event weights were retained for the evaluation of physics distributions and final statistical inference. The training samples comprise approximately 350k production-signal events, 170k decay-signal events, and 4.5M background events after the full event selection, where these numbers refer to the raw simulated event counts used for training prior to any normalisation. Each signal category accounts for all EFT samples combined inclusively.

Kinematic variables were selected as BDT inputs based on their discriminating power. The complete set of input variables used in the training is listed in Table 5.6, including their ranking of importance. This ranking is obtained by counting the number of times a given variable is used to split a node across all trees in the ensemble. Each occurrence is weighted by the squared separation gain achieved by the split and by the number of events contained in the corresponding node. This procedure provides a quantitative measure of the contribution of each input variable to the overall classification performance.

In Figure 5.4, the distributions of the input variables in the SR are given for both signal and background processes. The data–simulation agreement is generally good across all observables, with deviations mostly confined to extreme kinematic regions:

- For E_T^{miss} , the modelling reproduces the bulk of the spectrum within $\sim 10\%$, while a mild deficit is observed at large E_T^{miss} . This behaviour is commonly associated with imperfect modelling of neutrino kinematics in $t\bar{t}$ and single-top production, as well as residual uncertainties in jet

Variable	Description	Importance (%)
H_T	Scalar sum of the p_T of all jets and leptons	34.2
N_{jets}	Jet multiplicity	20.9
E_T^{miss}	Missing transverse energy	16.6
$\Delta R_{e\mu}$	ΔR between electron and muon	14.8
$p_T^{\text{bjet}} + p_T^{\text{light jet},1,2}$	p_T of the b -jet and the two leading light jets	7.0
$p_T^{\text{light jet},1}$	p_T of the leading light jet	2.3
η_e	Pseudorapidity of the electron	1.6
$p_T^{\text{jet},1}$	p_T of the leading jet	1.3
η_μ	Pseudorapidity of the muon	1.3

Table 5.6: The input variables used in the multivariate BDT analysis, listed according to the importance ranking provided by XGBoost for the average of the two trainings.

energy calibration and soft hadronic recoil, which directly affect the reconstructed missing transverse momentum.

- A similar pattern is observed for the leading-jet p_T , where agreement is good at low and intermediate values but the simulation tends to slightly underestimate the high- p_T tail. This effect is consistent with known limitations in the modelling of hard QCD radiation, including higher-order corrections, parton-shower matching, and uncertainties in initial- and final-state radiation. The same origin likely explains the small discrepancies seen in $p_T^{\text{bjet}} + p_T^{\text{light jet},1,2}$ and in the leading light-jet p_T , since these observables are strongly correlated with the overall event energy scale and additional jet activity.
- The jet multiplicity distribution N_{jets} is well reproduced at low and intermediate multiplicities, while deviations appear in the highest bin. Such differences are typically linked to the modelling of additional radiation and multi-jet emissions, where theoretical uncertainties from parton showers and scale choices become increasingly important.
- The angular observable $\Delta R_{e\mu}$ is generally well described, with small discrepancies at large separations that may arise from residual mismodelling of the top-quark boost spectrum or spin-correlation effects influencing the lepton angular topology. In contrast, the lepton pseudorapidities η_e and η_μ show very good agreement across the full range, indicating that the detector acceptance, lepton reconstruction efficiencies, and rapidity distributions of the dominant backgrounds are well modelled.

Overall, the observed mismodelling is consistent with known theoretical and experimental uncertainties affecting high-energy tails and regions dominated by additional QCD radiation, while the bulk kinematic phase space, most relevant for the BDT training, is accurately reproduced. This level of agreement validates the use of these variables as inputs to the multivariate analysis.

Figure 5.5 displays the linear correlations between the input variables. The correlation matrix shows that several kinematic observables entering the BDT are strongly related, while others provide largely independent information. Global event-activity variables exhibit the largest correlations: H_T , the leading-jet transverse momentum $p_T^{\text{jet},1}$, the scalar sum $p_T^{\text{bjet}} + p_T^{\text{light jet},1,2}$, and the leading




Figure 5.4: Comparison of input variable distributions in data and simulation within the SR used for BDT training. Statistical uncertainties are shown for the MC prediction (hatched bands) and for the data (error bars), with overflow events included in the final bin. The lower panel displays the ratio of data to the total MC prediction, with error bars indicating the statistical uncertainty. The hatched band in the ratio panel represents the statistical uncertainty in the MC prediction.

light-jet p_T are strongly correlated ($\rho \sim 0.8-0.9$), reflecting their common sensitivity to the overall hadronic energy scale. The jet multiplicity N_{jets} shows moderate correlations with these quantities ($\rho \sim 0.5-0.6$), consistent with higher jet activity in more energetic events. The missing transverse energy E_T^{miss} presents only moderate correlations with jet-based observables ($\rho \sim 0.3-0.4$), indicating partial dependence on the event kinematics. In contrast, the angular separation $\Delta R_{e\mu}$ is largely uncorrelated with the remaining variables, providing complementary geometric information to the BDT. The lepton pseudorapidities η_μ and η_e are essentially decorrelated from the global kinematic variables and show only a weak mutual correlation, demonstrating that they encode largely independent information related to the leptons' angular topology rather than the event energy scale. Overall, despite the presence of strongly correlated energy-scale observables, the inclusion of weakly correlated angular variables ensures that the BDT exploits diverse and non-redundant discriminating features; a similar correlation pattern is observed for the signal sample.


To control overtraining while making full use of the available MC simulation, a cross-validation scheme was implemented in which two separate BDTs were constructed. Specifically, the classifier referred to as 'BDT even' was trained on the even-numbered events of the signal and background samples and evaluated on the odd-numbered events, whereas 'BDT odd' was trained on the odd-numbered events and tested on the even-numbered ones. This procedure ensures that the events used for training are never used for evaluation, providing an unbiased estimate of the classifier performance.

The output of the BDT algorithm is expressed in terms of a discrete score distribution over the three output classes: decay signal, production signal, and background. For each class, the classifier provides a corresponding response, allowing the construction of distributions such as those shown in Figure 5.6, where the predicted class scores are compared for training and test samples. The good agreement observed between the two distributions indicates that no significant overtraining is present.


The performance of the multi-class BDT is evaluated using receiver operating characteristic (ROC) curves constructed by reducing the problem to a binary classification. This reduction is performed to enable a direct and well-defined comparison of the discriminating power for each signal class against the remaining event categories, following standard practice in multi-class classification problems. In particular, the signal efficiency, ε_{sig} , is defined for a given signal class by treating the corresponding signal as positive and the remaining classes as background. The ROC curve is then obtained by varying the threshold on the relevant BDT response and computing the signal efficiency and background rejection (defined as one minus the background efficiency). The signal efficiency is given by the fraction of signal events with a response above the threshold, with the background efficiency, ε_{bkg} , defined analogously. The integral of the ROC curve provides a quantitative measure of the separation power, with values closer to unity indicating better performance. The areas under the ROC curves are 0.9641 for the 'BDT odd' training and 0.9646 for the 'BDT even' training, as shown in Figure 5.7.

An additional diagnostic of the training behaviour is provided by the evolution of the loss function as a function of the boosting iteration, shown in Figure 5.8. The loss is evaluated independently on the training and testing samples and corresponds to the *multi-class softmax cross-entropy loss* minimised during the XGBoost training. The simultaneous decrease and subsequent stabilisation of the training and validation losses indicate a well-converged training, while the absence of a significant divergence between the two curves further confirms that overfitting is effectively controlled.

The output of the multi-class BDT corresponding to the background class is subsequently used as a discriminating variable in the analysis regions. This background BDT score provides a powerful summary of the multivariate information exploited by the classifier and is employed as an input observable in the fitting procedure described in Section 5.7. By construction, this variable enhances the separation between signal and background and thereby improves the sensitivity of the statistical analysis.



(a)



(b)

Figure 5.5: Correlation of the input variables for the signal (a) and background (b) classes. Signal corresponds to the combination of production and decay samples.




Figure 5.6: Overtraining check for the BDT trained on odd-numbered events. The markers show the BDT score distribution for the test sample, while the solid line corresponds to the training samples.




Figure 5.7: The ROC curve for the cross-validated BDTs.




Figure 5.8: Training and validation loss curves for the multi-class XGBoost classifier.

5.6 Systematic uncertainties

This section details the systematic uncertainties considered in the analysis. Experimental uncertainties arise from limitations in the detector measurements and collision conditions, as well as from uncertainties propagated through the calibration of physics objects. Meanwhile, theoretical uncertainties pertain to the modelling of SM background processes in the search, including cross-section calculations and MC generator configurations. Both the shape and normalisation of signal and background processes may be impacted by these uncertainties, which are incorporated as NPs in the profile likelihood fit, to be discussed in Section 5.7.

5.6.1 Experimental uncertainties

Luminosity For Run 2 of the LHC, the total integrated luminosity recorded by the ATLAS detector is determined with an uncertainty of 0.83%. This value is derived using the procedure outlined in Ref. [180] and is based on primary luminosity measurements performed with the LUCID-2 detector.

Pileup Systematic uncertainties associated with the scale factors used to reweight the pile-up distribution in simulated samples to reproduce that observed in data are taken into account. These scale factors are derived by varying the nominal pile-up distribution within its uncertainty and are applied independently to the 2015+2016, 2017, and 2018 MC simulated samples.

Lepton reconstruction efficiencies Such scale factors are obtained using tag-and-probe methods, as described in Section 4.3, and are associated with corresponding uncertainties. For electrons, separate nuisance parameters are assigned to the uncertainties in the identification, isolation, reconstruction, and trigger efficiency scale factors [122, 181]. For muons, independent statistical and systematic uncertainties are defined for the identification, isolation, track-to-vertex association, and trigger efficiencies, each implemented through dedicated nuisance parameters [127, 128].

Lepton energy/momentum scales Simulated events are corrected to account for mismodelling of lepton energy and momentum scales and resolutions with respect to data, using $Z \rightarrow \ell\ell$ events [127] and, additionally, $W \rightarrow \ell\nu$ events [122, 182]. For electrons, two nuisance parameters are introduced to describe the scale and resolution uncertainties. For muons, four nuisance parameters are defined, separating the uncertainties associated with the ID and MS track resolutions [128, 183].

Jet energy scale The JES uncertainties are derived using a combination of LHC collision data, simulated events, and test-beam measurements. An initial MC-based calibration is applied, followed by an in-situ correction using $\sqrt{s} = 13$ TeV data to constrain the remaining uncertainties [184]. Events containing a vector boson and additional jets are used to calibrate jets in the central detector region, while dijet events provide a relative calibration of forward jets with respect to central jets. The response of jets at high transverse momentum is constrained using multijet events. In analyses that are not strongly sensitive to JES, treating each of the 80 individual uncertainty source separately is unnecessary. Instead, the reduction scheme described in Ref. [137] is adopted, combining the original nuisance parameters into 30 independent components, each implemented with upward and downward variations and retaining their dependence on the reconstructed jet p_T and η .

Jet energy resolution This uncertainty is evaluated using dijet events, exploiting the width of the dijet asymmetry in bins of jet p_T and η , following procedures similar to those described in Ref. [185]. The initial fits yield approximately 120 NPs, which are subsequently condensed using a pseudo-data reduction scheme into 12 effective components plus a residual data-to-simulation term, resulting in a total of 13 NPs for the JER.

Jet vertex tagger A combined statistical and systematic uncertainty is assigned to the scale factors used to correct residual efficiency differences between data and simulation after the application of the JVT algorithm [186]. These scale factors are obtained from a tag-and-probe study in $Z \rightarrow \mu\mu$ +jets events, selecting configurations with $|\Delta\varphi(Z, \text{jet})| > 2.8$ to isolate hard-scatter jets. The systematic component covers differences in the measured efficiencies between samples generated with SHERPA and POWHEG +PYTHIA 8, as well as potential mismodelling of the $|\Delta\varphi(Z, \text{jet})|$ distribution.

Flavour-tagging uncertainties Uncertainties in the data-to-simulation correction of b -tagging efficiencies are modelled using 19 nuisance parameters. The efficiencies are measured in $t\bar{t}$ dilepton events for b -jets, in $t\bar{t} \ell$ +jets events for c -jets, and in Z +jets events for light-flavour jets [139–141]. These NPs are obtained through a reduction of the underlying systematic uncertainties, employing the same methodology as used for the JES, as described in Ref. [187]. The final set comprises nine nuisance parameters for b -jet tagging, four each for the mistagging of c -jets and light-flavour jets, and two additional parameters accounting for the extrapolation of the scale factors to high p_T .

E_T^{miss} uncertainties Systematic variations associated with the individual physics-object contributions to E_T^{miss} are propagated to the reconstructed E_T^{miss} and are therefore covered by the corresponding object-related uncertainties. In addition, dedicated scale and resolution variations are applied to the E_T^{miss} track-based soft term.

MC statistical uncertainty The limited size of the simulated MC samples introduces an additional *statistical* uncertainty in the predicted event yields. In the context of the profile likelihood fit, this effect is treated using the Barlow-Beeston approach [188] (or its light variant), in which the finite MC sample are incorporated directly into the likelihood [189, 190]. Concretely, an independent nuisance parameter is introduced for each bin of each process whose template is derived from simulation. These bin-by-bin parameters constrain the expected yield according to the corresponding MC statistical uncertainty and are typically implemented with Gaussian or Poisson constraints, depending on the chosen configuration.

5.6.2 Theoretical uncertainties

The modelling of signal and SM background processes introduces systematic uncertainties which are evaluated by the use of alternative MC simulated samples or by the application of alternative MC generator weights that correspond to variations in the nominal generator setup. These uncertainties are evaluated for signal processes and dominant background processes, namely $t\bar{t}$ and tW . Minor background contributions are instead assigned larger normalisation uncertainties, as detailed in Section 5.6.3.

The modelling of the $t\bar{t}$ process receives particular focus, due to both the tendency of the POWHEG generator to predict a harder top quark p_T distribution than observed in data [191] and its status as the primary background in the SR. Not all uncertainties are applied uniformly across all processes, as the availability of alternative simulated samples or weights may vary.

For uncertainties evaluated via alternative MC weights, the normalisation sum of weights is recalculated to ensure correct luminosity scaling, avoid double-counting cross-section uncertainties, and consequently, to account for shape and acceptance effects from modelling uncertainties. Normalisation-only cross-section uncertainties are applied separately, as discussed subsequently.

Scale variations To account for missing higher-order terms in ME calculations, variations in the renormalisation and factorisation scales, μ_R and μ_F , are applied. These parameters are independently varied by a factor of two around their nominal values for signal, $t\bar{t}$, and tW samples, producing separate NPs through internal weights without requiring alternative samples.

PDF set Uncertainties affecting the PDF set used in the generation of the sample are considered. For $t\bar{t}$ and tW , these are evaluated using the PDF4LHC prescription [84]: a set of 30 Hessian eigenvariations are considered, each of which is assigned a NP which is profiled in the fit. For signal samples, however, PDF uncertainties are estimated using sets with 100 NNPDF variations; the standard deviation of this set is used to define a symmetrised uncertainty.

Strong coupling constant α_S For the $t\bar{t}$ process, uncertainties in α_S are evaluated using PDF sets generated with $\alpha_S(M_Z)$ values of 0.117 and 0.119, with the nominal $\alpha_S(M_Z)$ being 0.118.

Matching of NLO calculation to parton shower for $t\bar{t}$ Matching parton-level events from the ME generator to the parton shower generator beyond LO must avoid double-counting higher-order corrections. The nominal $t\bar{t}$ sample uses the POWHEG matching scheme [104, 192, 193]. The uncertainty from this matching procedure is evaluated [194] by comparing the $t\bar{t}$ sample generated with POWHEG + PYTHIA 8 to an alternative with $p_{T,\text{hard}} = 1$, which regulates the definition of the vetoed region of the showering, essential to avoid holes or overlaps in the phase space filled by POWHEG and PYTHIA.

Additional $t\bar{t}$ radiation modelling In matching parton-level events to the parton shower at NLO using the POWHEG method, the h_{damp} scale is selected to regulate the p_T of the hardest emission, which is described by the ME calculation [160]. To assess uncertainties related to h_{damp} , an alternative POWHEG + PYTHIA 8 $t\bar{t}$ sample with h_{damp} set to $3m_t$ is compared to the nominal configuration ($h_{\text{damp}} = 1.5m_t$). Intuitively, increasing values of h_{damp} result in a slightly larger p_T for the hardest additional emission beyond LO.

Interference between $t\bar{t}$ and tW As discussed in Section 5.2, the tW sample excludes interference with $t\bar{t}$ due to the overlap removal method. Sensitivity to the excluded interference is assessed by comparing the nominal tW sample to a POWHEG + PYTHIA 8 sample using a diagram subtraction scheme [160], which retains interference effects. The difference is then used to estimate the impact of interference.

Initial- and final-state radiation Uncertainties from ISR modelling are evaluated by varying the Var3c parameter in the A14 PYTHIA tune, which affects α_S in the initial-state shower [109]. FSR uncertainties are assessed by varying μ_r for QCD emissions in the final-state shower. These uncertainties are considered for both signal and background ($t\bar{t}$ and tW) processes.

Parton shower and hadronisation modelling Uncertainty in the choice of parton shower and hadronisation models, given their different frameworks, is assessed by comparing samples showered with PYTHIA to those using HERWIG. For $t\bar{t}$, the nominal PYTHIA 8 sample is compared to a HERWIG 7.1 sample with equivalent ME generator settings, where HERWIG uses its tuned parameters and the MMHT LO PDF set [195]. Nominal signal samples from MG_AMC@NLO + PYTHIA 8 are compared to those generated with MG_AMC@NLO + HERWIG 7.1.

Top p_T reweighting Earlier studies have shown that the agreement between data and simulation in $t\bar{t}$ events, in particular for the top-quark p_T spectrum, improves when higher-order calculations are considered [196]. More recently, differential $t\bar{t}$ predictions at NNLO QCD accuracy including EW corrections have become available [197]. The associated modelling uncertainty is therefore assessed by taking the full difference between results obtained with and without applying a correction that aligns the simulation with NNLO QCD and NLO EW predictions.

5.6.3 Cross-section uncertainties

Normalisation-only cross-section uncertainties are also considered for the background processes. For tW production, a theoretical cross-section uncertainty of approximately 5.3% is assigned, derived from NNLO predictions including scale, PDF, and α_S variations [198]. Additionally, a conservative 50% normalisation uncertainty is assigned to backgrounds within the ‘other SM’ category. As noted in Section 5.4, since the normalisation for the $t\bar{t}$ process is normalised by the fit, no cross-section uncertainty is assigned to it.

5.7 Results

5.7.1 Statistical analysis framework

A binned maximum-likelihood fit is performed to extract the signal strength of the cLFV top-quark production and decay processes. In the absence of a signal, exclusion limits are set on the cross-section of these processes, which are then translated into the branching ratio of the top-quark decay into the $e\mu q$ final state. These limits are derived at 95% CL using the modified frequentist approach with the asymptotic approximation.

The fit is carried out with the TREXFITTER [199] framework (v1.4.0), which constructs statistical models in the HISTFACTORY [189] format. These models are implemented using RooFIT [200], and RooSTATS [201] is used for hypothesis testing and confidence interval estimation. Both RooFIT and RooSTATS are components of the ROOT [202, 203] data analysis environment.

An essential element of the fit is the minimisation of the NLL, which is performed internally through calls to MINUIT2 [204–206]. Parameter estimation proceeds in several stages. First, the MIGRAD algorithm is employed to locate the minimum of the likelihood function. Next, the HESSE method

is applied to refine the uncertainty estimates; although more accurate than MIGRAD, it yields only symmetric uncertainties by construction. Finally, an optional step invokes the MINOS algorithm, which scans the likelihood to provide the most precise parameter uncertainties, typically asymmetric. Because MINOS substantially increases the fit runtime while altering the uncertainty estimates only marginally, it is applied exclusively to the signal strength parameter (POI).

The minimisation procedure not only locates the NLL minimum but also evaluates the correlations among the model parameters: namely the POI, other normalisation factors, and systematic-uncertainty nuisance parameters.

Systematic uncertainties treatment

Systematic uncertainties arising from instrumental effects and from the modelling of signal and background processes in the MC simulations, as detailed in Section 5.6, are incorporated into the fit as nuisance parameters.

A *pruning* procedure is employed to remove nuisance parameters with negligible pre-fit impact, thereby simplifying the fit model, reducing computational time, and improving stability. For each nuisance parameter, both shape and normalisation contributions are included in the likelihood. Pruning is performed per sample and per region, and separately for the shape and normalisation components of an uncertainty. A threshold of 0.5% is applied to both: systematic uncertainties that alter either the shape or the normalisation of a background by less than this amount are discarded.

The pruning of normalisation and shape components proceeds as follows. For each region, the $\pm 1\sigma$ (“up/down”) variations associated with a given systematic uncertainty are evaluated. For the normalisation component, the variation is retained if any up/down shift changes the total yield relative to the nominal by more than the threshold. For the shape component, each variation is first rescaled to match the nominal yield; if any bin then deviates from the nominal distribution by more than the threshold, that shape component is kept. The pruning threshold was varied between 0.1% and 1% to assess the stability of the fit with respect to the treatment of small-impact systematic uncertainties. While the 0.1% threshold retained a larger number of nuisance parameters and the 1% threshold led to a more aggressive reduction of systematics, the nominal choice of 0.5% was found to provide the best balance between preserving relevant systematic effects and maintaining fit stability and sensitivity. A table reporting the pruning procedure outcome is provided in Figure 5.9.

Subsequently, a *smoothing* procedure, implemented through the relevant TREXFITTER tools, is applied to the all systematic-uncertainty variation templates to mitigate fluctuations across adjacent bins from limited MC statistics.

All systematic uncertainties are then *symmetrised*. In this procedure, the original “up” and “down” variations provided for each nuisance parameter are first interpreted as independent shifts of the nominal template. For each bin, the symmetrised variation is constructed by taking half of the difference between the up and down yields with respect to the nominal prediction, and then defining a pair of symmetric variations about the nominal value. This symmetrisation ensures that the nuisance parameter produces equal-magnitude upward and downward shifts in the predicted yield when varied by $\pm 1\sigma$. The procedure removes potential biases associated with asymmetric fluctuations in limited MC simulated samples and enforces a linear, symmetric interpolation of the template morphing within the profile likelihood fit. As a consequence, the impact of each systematic uncertainty becomes centred on the nominal prediction, improving the numerical stability and interpretability of the fit while preserving the overall scale of the variation.




Figure 5.9: Illustration of the pruning procedure applied to nuisance parameters in the SR. Pruning of nuisance parameters associated with *modelling*, or theoretical, uncertainties (a), and with *instrumental*, or experimental, effects in lepton, jet, and b -tagging measurements (b). The NPs that are fully removed are indicated in red; those for which only the shape or normalisation component is discarded are shown in yellow or orange, respectively. Parameters highlighted in green retain both shape and normalisation components.

For cases where only a single variation is available (*e.g.*, when comparing alternative samples to the nominal to estimate the PS uncertainty), the missing downward variation is taken as the mirrored version of the upward one. Examples of the $\pm 1\sigma$ (“red–blue”) variations of modelling uncertainties affecting the $t\bar{t}$ and tW processes are shown in Figures 5.10 and 5.11.




Figure 5.10: CR distributions illustrating the impact on the $t\bar{t}$ template of $\pm 1\sigma$ variations of the $t\bar{t}$ modelling uncertainties associated with (a) NLO matching, (b) μ_F scale variation, (c) top-quark p_T reweighting, and (d) h_{damp} . As mentioned in Section 5.5.2, the analysis regions are plotted in the background multi-class BDT score. The variations are shown before (original) and after (modified) the application of smoothing and symmetrisation in TRExFITTER. Solid lines correspond to the original up/down variations (red/blue), while dashed lines indicate the modified variations. The nominal prediction is shown by the black histogram in the upper panel. Except for (b), a symmetrisation is applied, such that the down variation is obtained as a mirror of the up variation; consequently, no dashed curve is shown for the modified down variation. The hatched band represents the statistical uncertainty, normalised to the MC prediction in the ratio panel.

Fit strategy The normalisation of the dominant $t\bar{t}$ background is scaled through the introduction of a free-floating normalisation factor (NF) in the CR, denoted by $k(t\bar{t})$. This factor is determined directly from data in a simultaneous binned profile likelihood fit performed across the SR and CR. The parameter of interest is the signal-strength modifier μ_{cLFV} , which scales the expected signal yield,




Figure 5.11: CR distributions illustrating the impact on the tW template of $\pm 1\sigma$ variations of the tW modelling uncertainties associated with (a) diagram subtraction, (b) μ_R scale variation, (c) ISR, and (d) FSR.

while the $t\bar{t}$ normalisation factor enters the fit as a nuisance parameter and is correlated between the SR and CR. All other systematic uncertainties are incorporated as additional nuisance parameters and are constrained by the fit.

Limit calculation for the cLFV process

TRExFITTER provides the expected limit on the cLFV process in terms of the signal strength modifier μ_{cLFV} , which can be converted into a limit on the cLFV cross-section. This cross-section corresponds to the sum of the decay and production contributions, both of which scale with the square of the WCs responsible for the processes:

$$\sigma_{\text{cLFV}} = \sigma(qg \rightarrow t\ell\ell') + \sigma(pp \rightarrow t(\rightarrow \ell\ell'q)\bar{t}) \quad (5.8)$$

and

$$\sigma_{\text{cLFV}} \propto C^2. \quad (5.9)$$

For the initial (pre-fit) signal cross-section, all WCs introduced in Section 2.2 are assumed to take the value 1.0. In this scenario, the 95% confidence level limit on the squared WC, $(C^{95\%CL})^2$, is calculated as the ratio of the 95% CL limit on the cross-section, $\sigma^{95\%CL}$, to the starting cross-section with $|C| = 1.0$, $\sigma^{|C|=1.0}$, scaled by the square of the initial WC value:

$$\left(C^{95\%CL}\right)^2 = \frac{\sigma^{95\%CL}}{\sigma^{|C|=1.0}}. \quad (5.10)$$

The limit on the modulus of the WC is then obtained by taking the square root. From this limit, the corresponding bound on the top-quark branching ratio to the cLFV final state can be derived using Eq. 2.10 and dividing by the SM top-quark width, assuming identical couplings at the decay and production vertices.

Fit observable and binning strategy

In the analysis regions, the discriminating observable used in the fit is defined as the sum of the two signal scores calculated as $1 - (\text{multi-class BDT background score})$, which provides strong separation between signal and background. The inversion of the classifier output is chosen such that signal-like events populate the high-value end of the distribution, following the standard convention.

The binning of this observable is optimised to balance several competing requirements. A smaller number of bins improves the numerical stability of the fit, reduces computational time, and ensures adequate event counts per bin, while overly coarse binning would degrade sensitivity. In the SR, the number and bin boundaries are determined automatically using TREXFITTER by optimising the binning of the discriminant distribution to maximise the expected statistical sensitivity of the profile-likelihood fit while maintaining sufficient MC statistics per bin, resulting in 14 bins.

In the CR, a uniform binning scheme with four bins is adopted over a range in which the signal contamination in each bin remains below 5%.

5.7.2 Fit using Asimov data

To validate the analysis strategy and estimate the expected sensitivity, an initial fit is performed with the SR blinded, such that no collision data from this region are used. This blinding procedure avoids potential experimenter bias that could arise from inspecting the observed data or from optimising selection criteria in response to apparent signal-like features.

In this validation step, the fit is performed using Asimov datasets in both the SR and CR. These datasets are constructed from the expected MC simulated yields, as described in Section 4.5.4, and correspond to the background-only hypothesis, with the signal strength fixed to $\mu_{\text{cLFV}}^{\text{Asimov}} = 0$. The fitting procedure is carried out in multiple stages. First, a fit is performed in the CR alone to constrain the normalisation factors and the relevant nuisance parameters using the Asimov data. The fitted values of these parameters are then used to construct a modified Asimov dataset for the SR, in which the background predictions reflect the constraints obtained from the CR.

A final profile likelihood fit is subsequently performed simultaneously in the SR and CR using this modified dataset. This approach allows the impact of the CR constraints on the SR prediction to be assessed in a controlled environment prior to unblinding.

The resulting pre-fit and post-fit distributions for all analysis regions are shown in Figure 5.12. The normalisation factor obtained by the fit is $k(t\bar{t}) = 1.00 \pm 0.05$. Pre- and post-fit yields are collected in Tables 5.7 and 5.8.




Figure 5.12: Pre-fit (a,b) and post-fit (c,d) plots for the fit using Asimov data in SR and CR. The uncertainty band includes both the statistical and systematic uncertainties.

Process	SR	CR
Signal	$79\,000 \pm 8\,000$	720 ± 250
$t\bar{t}$	$161\,000 \pm 19\,000$	$132\,000 \pm 13\,000$
tW	$8\,000 \pm 1\,500$	$3\,600 \pm 600$
Other	$1\,900 \pm 900$	$1\,100 \pm 600$
Total Background	$171\,000 \pm 19\,000$	$137\,000 \pm 13\,000$

Table 5.7: Pre-fit event yields in the analysis regions for the Asimov fit. The signal is normalised by setting the WCs equal to 1.0. The quoted uncertainties include the combined statistical and systematic contributions, summed in quadrature.

Process	SR	CR
Signal (μ_{cLFV} set to 0)	0 ± 40	0.00 ± 0.34
$t\bar{t}$	$161\,500 \pm 1\,800$	$132\,400 \pm 900$
tW	$8\,000 \pm 1\,500$	$3\,600 \pm 600$
Other	$1\,900 \pm 900$	$1\,100 \pm 600$
Total Background	$171\,400 \pm 2\,500$	$137\,100 \pm 1\,200$

Table 5.8: Post-fit event yields in the analysis regions for the Asimov fit. The signal is normalised by setting the WCs equal to 1.0. The quoted uncertainties include the combined statistical and systematic contributions, summed in quadrature.

5.7.3 Unblinded fit

After confirming the stability of the Asimov fit, an additional validation was performed using real data in the CR while keeping the SR blinded with Asimov data. This step also showed stability, confirming that the analysis strategy was robust and could be applied to the full dataset with an unblinding⁴ of the SR. Unblinding allows the inclusion of SR data in the fit, enabling a direct comparison with background predictions and the extraction of the signal strength.

The pre-fit and post-fit distributions for all analysis regions are presented in Figure 5.13. The fit returns a normalisation factor for the dominant $t\bar{t}$ background of $k(t\bar{t}) = 0.95 \pm 0.05$, and a fitted signal strength of $\mu_{\text{cLFV}} = (-0.04 \pm 0.5) \times 10^{-3}$, consistent with the background-only hypothesis. This indicates no significant evidence for the presence of the cLFV signal in the SR.

⁴ This is an unblinding for the purposes of this thesis, the results are not yet approved by the ATLAS Collaboration.




Figure 5.13: Pre-fit (a,b) and post-fit (c,d) plots for the unblinded fit. The uncertainty band includes both the statistical and systematic uncertainties. The lower panel plots the ratio between data and background for each bin.

The pre-fit event yields for each process are summarised in Table 5.9, with the corresponding post-fit yields provided in Table 5.10. These comparisons illustrate the effect of the fit in constraining background contributions and demonstrate the consistency of the observed data with the expectations from the background model.

Process	SR	CR
Signal	$79\,000 \pm 8\,000$	720 ± 250
$t\bar{t}$	$161\,000 \pm 19\,000$	$132\,000 \pm 13\,000$
tW	$8\,000 \pm 1\,500$	$3\,600 \pm 600$
Other	$1\,900 \pm 900$	$1\,100 \pm 600$
Total Background	$171\,000 \pm 19\,000$	$137\,000 \pm 13\,000$
Data	166 149	133 224

Table 5.9: Pre-fit event yields in the analysis regions for the unblinded fit.

Process	SR	CR
Signal	0 ± 40	0.00 ± 0.34
$t\bar{t}$	$156\,700 \pm 1\,800$	$128\,600 \pm 900$
tW	$7\,700 \pm 1\,400$	$3\,500 \pm 600$
Other	$1\,800 \pm 900$	$1\,100 \pm 600$
Total Background	$166\,200 \pm 2\,200$	$133\,200 \pm 1\,100$
Data	166 149	133 224

Table 5.10: Post-fit event yields in the analysis regions for the unblinded fit.

Correlations, constraints and pulls of nuisance parameters

The correlation coefficients between the NPs and the POI are computed from the correlation matrix obtained during the minimisation of the NLL function. These correlations are shown in Figure 5.14. NPs associated with background sources that have similar shapes tend to be strongly correlated, as variations in these parameters affect the predicted distributions in comparable ways. Understanding these correlations is important, since highly correlated NPs can influence the precision with which the POI is determined and can indicate potential degeneracies in the fit.

Systematic NPs, which are constrained by Gaussian priors, and NFs, which are typically free-floating, are introduced to account for differences between data and simulation predictions. A correlation




Figure 5.14: Correlation matrix (in %) of the nuisance parameters obtained from the full fit. Only NPs with at least one correlation larger than 15% are drawn to improve readability.

between a systematic-uncertainty NP and a NF indicates that variations in the NP strongly influence both the value and uncertainty of the NF, while the NF can, in turn, reduce the post-fit uncertainty of the NP. As shown in Figure 5.14, the systematic NPs associated with backgrounds normalised via NFs often exhibit large correlations or anti-correlations with their corresponding NFs.

For example, in the case of the $t\bar{t}$ background, variations in top-quark p_T reweighting, NLO matching, final-state radiation, and parton-shower modelling are adjusted during the fit to correct the data/simulation differences observed in the CR. The anti-correlation between the top-quark p_T reweighting NP and the $t\bar{t}$ NF reflects the fact that an upward variation of the NP increases the predicted $t\bar{t}$ yield in the CR, which the NF compensates for by decreasing its value, and vice versa. Similar reasoning applies to other modelling NPs that influence the total predicted yield.

A significant correlation between a NP and the parameter of interest, μ_{cLFV} , indicates that variations in that NP have a substantial effect on the expected signal strength. Identifying these correlations is important, as they highlight which systematic uncertainties most strongly impact the measurement and can inform future improvements in modelling or control-region design.

The *pull* of a NP quantifies the deviation of its fitted value from its pre-fit expectation, relative to the pre-fit uncertainty. It is defined as

$$\text{pull}(\theta) = \frac{\hat{\theta} - \theta_0}{\Delta\theta}, \quad (5.11)$$

where $\hat{\theta}$ is the post-fit value of the NP, θ_0 is its nominal (pre-fit) value, and $\Delta\theta$ is the pre-fit uncertainty.

The pulls of the non-pruned systematic uncertainties are shown in Figure 5.15, with the error bars indicating the post-fit uncertainty on each NP. Note that the plot only includes NPs that enter the likelihood function with a defined pre-fit constraint; free-floating normalisation factors, the parameter of interest, and gamma parameters are not represented in this figure.




Figure 5.15: Nuisance parameter pull plot obtained for the full fit, showing how the values of the nuisance parameters and their uncertainties changed relative to those before the fit. The error bars correspond to the post-fit uncertainty, $\Delta\hat{\theta}$, which is to be compared to the pre-fit $\Delta\theta = 1$ uncertainty.

The pull distributions for the instrumental experimental and theory modelling NPs show that most NPs are well within one standard deviation of their pre-fit values, indicating good agreement between the background model and the observed data. In the instrumental plot (left), the majority of NPs, covering lepton efficiencies, pile-up corrections, jet energy scale and resolution, and b -tagging uncertainties, display small pulls, typically below 1, and remain well within the post-fit uncertainties (green bands). This shows that the fit did not require large adjustments to these systematic effects, reflecting both the robustness of the detector calibration and the stability of the fit in the control and signal regions.

In the theory NPs (right), covering top-quark modelling, $t\bar{t}$ and tW cross-section variations, PDFs, parton-shower and ISR/FSR modelling, as well as signal modelling uncertainties, the pulls similarly remain close to zero for most parameters. This reflects that the observed data in the CR are consistent with the MC predictions within the pre-fit systematic uncertainties. Some NPs, notably top-quark p_T

reweighting, NLO matching scales, and parton-shower variations, show more noticeable pulls, which is expected given their direct impact on the $t\bar{t}$ prediction in the CR. Corresponding uncertainties are often reduced post-fit, reflecting constraints provided by the data. The correlations between these modelling NPs and the $t\bar{t}$ normalisation factor further allow the fit to adjust the background predictions in a coherent way. Overall, no NP exhibits extreme pulls or an unreasonably large post-fit shift, indicating that the fit is stable, the nuisance parameters are properly constrained, and systematic uncertainties are under control.

Ranking of nuisance parameters

The ranking of NPs is determined by their impact on the estimated value of the POI, μ_{CLFV} . This impact is quantified as the change in the POI when a given NP is fixed to a value offset x from its best-fit estimate, $\hat{\theta}$:

$$\text{Impact} = \mu_{\text{CLFV}}(\hat{\theta} \pm x) - \mu_{\text{CLFV}}(\hat{\theta}). \quad (5.12)$$

Two types of impact estimates are commonly considered:

- **Pre-fit impact:** $x = \Delta\theta$, corresponding to a one-standard-deviation variation of the NP relative to its pre-fit uncertainty. For unconstrained NPs, the pre-fit impact is not defined.
- **Post-fit impact:** $x = \Delta\hat{\theta}$, where $\Delta\hat{\theta}$ is the uncertainty on the post-fit value of the NP. Since the fit may constrain NPs using the observed data, the post-fit impact is often smaller than the pre-fit estimate.

Figure 5.16 shows the ranking of nuisance parameters according to their impact on the fitted signal strength μ_{CLFV} , together with their post-fit pulls. The dominant contributors are theory-related uncertainties associated with the modelling and normalisation of the tW and $t\bar{t}$ backgrounds. In particular, the tW DR modelling uncertainty and the $t\bar{t}$ NLO matching uncertainty exhibit the largest pre-fit impacts, indicating a strong sensitivity of the signal extraction to the theoretical description of these backgrounds. While their post-fit impacts are reduced, they remain among the leading sources of uncertainty, reflecting only partial constraint by the data.

Several additional NPs related to background normalisation and modelling, such as the free-floating normalisation term and the $t\bar{t}$ renormalisation and factorisation scale variations, also show sizeable impacts. The reduction from pre-fit to post-fit impact observed for many of these parameters demonstrates the constraining power of the control regions, which help to mitigate modelling uncertainties through the simultaneous fit. Experimental uncertainties, including jet flavour composition, pile-up modelling, and top p_T reweighting, contribute at a subleading but non-negligible level, consistent with their moderate correlations with μ_{CLFV} .

Signal-related theoretical uncertainties appear near the bottom of the ranking, with limited impact on μ_{CLFV} , confirming that the sensitivity of the analysis is primarily driven by background modelling rather than signal theory uncertainties. Overall, the observed ranking and impact pattern is consistent with expectations for a background-dominated search and supports the robustness of the statistical interpretation.

The ranking obtained from the unblinded fit is consistent with the relative magnitudes of the correlations between each NP and μ_{CLFV} , providing an independent check that the most influential

systematic uncertainties are correctly identified. This ranking is particularly useful for understanding which uncertainties dominate the sensitivity of the analysis and for guiding potential improvements in modelling or control-region design.




Figure 5.16: Ranking of nuisance parameters NPs from the unblinded fit. The NPs shown correspond to the systematic uncertainties with the largest influence on the fitted signal strength μ_{cLFV} . The points along the bottom axis indicate the deviation of each fitted NP from its nominal value, zero for systematic-uncertainty NPs and one for NFs, expressed in units of the pre-fit standard deviation, $\Delta\theta$. Error bars represent the post-fit uncertainties, $\Delta\hat{\theta}$; values close to one indicate little constraint from the data, while smaller values reflect a reduction compared to the pre-fit uncertainty. The hollow blue bars show the pre-fit impact of each NP on μ_{cLFV} , whereas the filled bars represent the corresponding post-fit impact. NPs are ordered according to their post-fit impact on the signal strength, as indicated by the top axis.

5.7.4 Limits on the cLFV process

The expected upper limit on the signal strength for the cLFV processes is computed using the CL_s method, based on the Asimov dataset in the SR under the background-only hypothesis, employing the fit described in section 5.7.2. The resulting expected upper limit on the signal strength at 95% confidence level is

$$\mu_{\text{cLFV}}^{\text{exp.}} < 9.5 \times 10^{-4}. \quad (5.13)$$

Eq. 5.13 translates into an expected upper limit on the cross-section of the combined production and decay signal processes of

$$\sigma_{\text{cLFV}}^{\text{exp.}} < 1.98 \text{ fb} \quad (95\% \text{ CL}). \quad (5.14)$$

Following the prescription outlined in Section 5.7.1, this cross-section limit is further translated into an expected upper limit on the branching ratio of the cLFV top-quark decay:

$$\mathcal{B}(t \rightarrow e\mu q)_{\text{exp.}}^{\text{stat.}+\text{syst.}} < 1.16 \times 10^{-7} \quad (95\% \text{ CL}). \quad (5.15)$$

The sensitivity to μ_{cLFV} is driven predominantly by the production signal, which provides the dominant statistical contribution compared to the decay signal.

Results from the unblinded fit yield an observed upper limit on the signal strength of

$$\mu_{\text{cLFV}}^{\text{obs.}} < 9.2 \times 10^{-4}, \quad (5.16)$$

which is slightly more stringent than the expected limit. This indicates good agreement between the observed data and the background-only hypothesis.

The corresponding observed upper limit on the cross-section of the combined production and decay signal processes is

$$\sigma_{\text{cLFV}}^{\text{obs.}} < 1.92 \text{ fb} \quad (95\% \text{ CL}), \quad (5.17)$$

which again closely follows the expected sensitivity of the analysis.

Using the same translation procedure as for the expected limit, the observed upper limit on the branching ratio of the cLFV top-quark decay, including both statistical and systematic uncertainties, is found to be

$$\mathcal{B}(t \rightarrow e\mu q)_{\text{obs.}}^{\text{stat.}+\text{syst.}} < 1.12 \times 10^{-7} \quad (95\% \text{ CL}). \quad (5.18)$$

This observed limit is marginally stronger than the expected one, consistent with the slight downward fluctuation observed in the fitted signal strength.

To assess the impact of systematic uncertainties, the observed limit is also evaluated considering statistical uncertainties only, yielding

$$\mathcal{B}(t \rightarrow e\mu q)_{\text{obs.}}^{\text{stat.}} < 8.73 \times 10^{-8} \quad (95\% \text{ CL}). \quad (5.19)$$

Comparing the statistical-only and statistical-plus-systematic results, the inclusion of systematic uncertainties weakens the observed branching ratio limit by approximately 28%. This indicates that, while the analysis remains predominantly statistically limited, systematic effects introduce a degradation of the sensitivity.

The observed upper limit on the branching ratio obtained in this analysis, Eq. 5.18, represents a substantial improvement over the previous ATLAS result reported in Ref. [16], which set an observed limit of

$$\mathcal{B}(t \rightarrow e\mu q)_{\text{obs.}}^{\text{stat.} + \text{syst.}} < 6.6 \times 10^{-5} \quad (95\% \text{ CL}). \quad (5.20)$$

This corresponds to an improvement of nearly three orders of magnitude in sensitivity. The enhanced reach of the present analysis can be attributed primarily to the inclusion of the cLFV production signal, which dominates the statistical sensitivity, as well as to the use of the full Run 2 dataset and an improved analysis strategy. In contrast, the earlier result relied on a smaller integrated luminosity of 80 fb^{-1} and considered only the decay-signal contribution. Together, these factors lead to a significantly stronger constraint on the branching ratio of the cLFV top-quark decay.

5.7.5 EFT interpretation

The exclusion limit on the cross-section of the $e\mu tq$ cLFV process can be interpreted within the SMEFT framework, where each effective operator is considered independently. In this approach, limits are derived under the assumption that a single dimension-6 operator is present at a time, while all others are set to zero, allowing for a straightforward mapping between the experimental constraint and the corresponding WC.

Although differences in the kinematic properties of the cLFV signal $gq \rightarrow te\mu$ are expected due to the distinct parton distribution functions of the u - and c -quark initial states, these contributions cannot be experimentally resolved with the current dataset. They are therefore treated independently in the interpretation, with separate limits derived for the u - and c -quark-initiated processes. This procedure ensures a consistent and conservative interpretation of the results within the SMEFT framework.

EFT interpretation procedure

The EFT interpretation is performed using the profile likelihood fit described in Section 4.5. Sensitivity to different new-physics structures is probed by replacing the inclusive signal sample, comprising contributions from all SMEFT operators listed in Table 2.1, with samples corresponding to individual operators. As in the fit to the inclusive sample, the initial value of each WC is fixed to unity for a reference scale of $\Lambda = 1 \text{ TeV}$. This procedure results in twelve independent fits, corresponding to the six EFT operators in Table 2.1, each considered separately for the two up-type quark flavours.

In addition to the single-operator fits, a scenario in which all vector operators contribute equally is also considered. This is achieved by setting the corresponding WCs to the same value, reflecting the assumption that the underlying new-physics interaction couples with equal strength to left- and right-handed fermion fields. This benchmark provides a complementary interpretation that is less sensitive to the chiral structure of the interaction [207]. Such a construction is redundant for the scalar and tensor cases, as each corresponds to a single operator and therefore admits no analogous equal-contribution scenario.

For the cLFV signal in $t\bar{t}$ decay, six simulated samples are used, as described in Section 2.3.2. These samples are inclusive with respect to the final-state quark flavour. The decays $t \rightarrow e\mu u$ and $t \rightarrow e\mu c$ are therefore generated with equal probabilities, with only minor differences expected due to the non-zero mass of the charm quark. As a result, no attempt is made to distinguish between the two decay modes in the reconstruction or statistical analysis.

EFT fit results

The behaviour of NPs in each individual fit is observed to be consistent with the results presented in Section 5.7.3, with the pulls on NPs showing very close agreement. This is expected, given that the

underlying data and background models are identical across the fits. Some differences, however, are noted in the relative impact of various signal-modelling systematic uncertainties between the EFT fits. In particular, for the $e\mu tc$ samples, the PDF uncertainty associated with the $gq \rightarrow te\mu$ signal process has a more significant effect than the parton shower and hadronisation uncertainty. Furthermore, for charm quark-initiated processes, the PDF uncertainties are substantially larger than those for up quark-initiated processes. Since the overall sensitivity of the fit is primarily limited by the size of the dataset, a detailed discussion of the ranking of systematic NPs is not presented.

Table 5.11 summarises the expected and observed 95% CL exclusion limits on the WCs for up- and charm-quark-initiated processes. The same expected and observed limits are shown graphically in Figure 5.17. Overall, significantly stronger constraints are obtained for the u -initiated case than for the c -initiated case. The expected limits for the u scenario are at the level of $\mathcal{O}(0.14)$ for the four-fermion operators and $c_{lequ}^{1(ijk3)}$, with a notably tighter bound of 0.036 for $c_{lequ}^{3(ijk3)}$ due to the larger cross-section of this process. In contrast, the expected limits in the c scenario are weaker by approximately a factor of four to five, lying around $\mathcal{O}(0.62\text{--}0.64)$ for most operators and 0.151 for $c_{lequ}^{3(ijk3)}$.

Exclusion limits on SMEFT WCs at 95% CL						
$i, j = e, \mu$	$c_{lq}^{-(ijk3)}$	$c_{eq}^{(ijk3)}$	$c_{lu}^{(ijk3)}$	$c_{eu}^{(ijk3)}$	$c_{lequ}^{1(ijk3)}$	$c_{lequ}^{3(ijk3)}$
u (Exp.)	0.142	0.144	0.145	0.145	0.145	0.036
u (Obs.)	0.138	0.139	0.140	0.141	0.147	0.036
c (Exp.)	0.62	0.62	0.64	0.63	0.63	0.151
c (Obs.)	0.67	0.62	0.68	0.63	0.61	0.148

Table 5.11: Observed and expected exclusion limits at 95% CL on each of the WCs probed by the analysis.




Figure 5.17: Observed and expected exclusion limits at 95% CL on the WCs probed in this analysis, corresponding to the SMEFT operators listed in Table 2.1. A new-physics scale of $\Lambda = 1$ TeV is assumed. Each limit is derived under the assumption that the WC under consideration is non-zero while all other operators are set to zero.

The observed limits are generally in good agreement with the expectations in both scenarios, indicating no significant deviation from the background-only hypothesis. Small fluctuations are visible: in the u case, the observed limits are slightly stronger than expected for most operators, while in the c case the observed limits are in some instances weaker (e.g. $c_{lq}^{-(ijk3)}$ and $c_{lu}^{(ijk3)}$) and in others slightly stronger (e.g. $c_{lequ}^{1(ijk3)}$ and $c_{lequ}^{3(ijk3)}$). The consistently tighter constraints in the u -initiated case reflect the greater sensitivity of the analysis to up-quark-induced signal production, whereas the reduced sensitivity in the charm-quark case leads to correspondingly weaker bounds across all operators.

Limits obtained for vector-, scalar- and tensor-like interactions are summarised in Table 5.12. Corresponding branching ratio and WC limits are displayed in Figure 5.18 comparing them between the quark involved in the interaction, up or charm.

Vertex	Int. type	$C_{e\mu tq}/\Lambda^2 [\text{TeV}^{-2}]$		$\mathcal{B}(10^{-6})$	
		Exp.	Obs.	Exp.	Obs.
$e\mu tu$	Vector	0.072	0.069	0.044	0.040
	Scalar	0.145	0.147	0.022	0.023
	Tensor	0.036	0.036	0.068	0.065
$e\mu tc$	Vector	0.31	0.33	0.84	0.95
	Scalar	0.63	0.61	0.42	0.39
	Tensor	0.151	0.148	1.17	1.11

Table 5.12: Expected and observed 95% CL upper limits on the cLFV WCs, and top quark cLFV branching fractions.




Figure 5.18: The observed 95% CL exclusion limits on the $\mathcal{B}(t \rightarrow e\mu c)$ as a function of $\mathcal{B}(t \rightarrow e\mu u)$ (left) and $e\mu tc$ vs. $e\mu tu$ Wilson coefficient (right) for the vector-, scalar-, and tensor-like cLFV interactions. The hatched bands indicate the regions containing 68% of the distribution of limits expected under the background-only hypothesis.

The results obtained by this analysis with ATLAS data can be directly compared with those reported by the CMS Collaboration [70], which performed a complementary search for charged-lepton-flavour violation in top-quark interactions using a similar EFT interpretation. Although the two analyses differ in event selection, signal modelling, and multivariate techniques, both experiments probe the same set of dimension-six operators and present limits in terms of WCs and corresponding branching fractions. This enables a meaningful comparison of their sensitivities and provides an important cross-check of the experimental reach of the LHC programme in constraining flavour-violating top-quark interactions.

Limits obtained for vector-, scalar- and tensor-like interactions are summarised in Table 5.13. Corresponding branching ratio and WC limits are displayed in Figure 5.19 comparing them between the quark involved in the interaction, up or charm.

Vertex	Int. type	$C_{e\mu tq}/\Lambda^2 [\text{TeV}^{-2}]$		$\mathcal{B}(10^{-6})$	
		Exp.	Obs.	Exp.	Obs.
$e\mu tu$	Vector	0.12	0.12	0.14	0.13
	Scalar	0.23	0.24	0.06	0.07
	Tensor	0.07	0.06	0.27	0.25
$e\mu tc$	Vector	0.39	0.37	1.49	1.31
	Scalar	0.87	0.86	0.91	0.89
	Tensor	0.24	0.21	3.16	2.59

Table 5.13: Expected and observed 95% CL upper limits on the cLFV WCs, and top quark cLFV branching fractions reported by CMS.




Figure 5.19: The observed 95% CL exclusion limits on the $\mathcal{B}(t \rightarrow e\mu c)$ as a function of $\mathcal{B}(t \rightarrow e\mu u)$ (left) and $e\mu tc$ vs. $e\mu tu$ Wilson coefficient (right) for the vector-, scalar-, and tensor-like cLFV interactions reported by CMS. The hatched bands indicate the regions containing 68% of the distribution of limits expected under the background-only hypothesis.

A comparison between the ATLAS and CMS results shows a consistent qualitative behaviour across operators and quark flavours, while highlighting a clear difference in overall sensitivity. In both experiments, limits on interactions involving the up quark are significantly stronger than those involving the charm quark, reflecting the larger up-quark parton density in the proton and consequently higher signal acceptance for $e\mu tu$ vertices. Furthermore, both analyses exhibit the same hierarchy among Lorentz structures: tensor operators yield the strongest constraints on the Wilson coefficients, followed by vector interactions, while scalar operators lead to comparatively weaker limits. Quantitatively, however, the ATLAS analysis achieves substantially tighter bounds, improving the limits on $C_{e\mu tq}/\Lambda^2$ by roughly a factor of two relative to CMS across most channels, which translates into branching-fraction limits typically smaller by factors of three to five. The observed limits are generally consistent with expectations in both experiments, indicating no significant deviation from the Standard Model prediction and demonstrating good agreement between the two independent measurements despite differences in the analyses.

Conclusion and Outlook

The ATLAS experiment has as one of its key missions to explore new physics beyond the Standard Model. This thesis investigates the potential violation of charged-lepton flavour in top-quark processes using the full Run 2 dataset of pp collisions at a centre-of-mass energy of $\sqrt{s} = 13$ TeV, with an integrated luminosity of 140 fb^{-1} . The analysis highlights the capability of the current ATLAS dataset to constrain the effective coupling strengths of BSM operators within the EFT framework. The significantly larger datasets anticipated in the future will enhance the precision of SM parameter measurements and enable the probing of even smaller effective coupling strengths.

This search is among the first to investigate cLFV in the top-quark sector and represents the first ATLAS search focusing on final states with a hadronically decaying top-originated W boson from charged-lepton flavour violating top-quark decay and production processes. These include $t\bar{t} \rightarrow (\ell\ell'q)(W \rightarrow qq)b$ and $qg \rightarrow t\ell\ell'$, where ℓ and ℓ' have distinct flavours. The leptons considered are electrons and muons, $\ell = \{e, \mu\}$, also including those from τ leptonic decays. The target signal topology is characterised by an isolated, opposite-charge electron-muon pair and multiple jets, with at least one being a b -jet. The signal region is optimised to enhance signal acceptance by selecting events containing exactly one b -jet, while a dedicated control region requiring more than one b -jet is used to constrain the dominant $t\bar{t}$ background. Boosted decision trees were employed to distinguish the signal, with events generated using a dedicated setup. The multivariate discriminant shape was used as input for a binned profile likelihood fit to test for the presence of signal events.

The search is also conducted within a model-independent EFT framework, offering sensitivity to a variety of new physics models capable of inducing such processes. The fitted observables align with the SM hypothesis, with no evidence of a cLFV signal detected. Exclusion limits are set on the Wilson coefficients of dimension-6 SMEFT operators that can induce cLFV $e\mu tq$ interactions, focusing on the leading contributions to these processes. While these limits do not directly reveal the nature of the underlying new physics, they constrain the parameter space of theoretical models. Importantly, these bounds serve as valuable inputs for global EFT fits, aiming to offer a comprehensive view of potential new physics in the 1–10 TeV energy range.

The presented results complement previous searches for four-fermion $e\mu tq$ interactions by ATLAS [16, 17] and CMS [18, 19, 70]. While combining exclusion limits for $e\mu tq$ and $\mu\tau tq$ operators is not relevant within the SMEFT framework, it may become significant in models like leptoquark scenarios that predict correlations between their effective coupling strengths. A combination with CMS $e\mu tq$ searches could yield valuable insights but is complicated by the unknown correlations

in systematic uncertainties between the two collaborations.

Looking ahead, this analysis provides a solid foundation for future investigations of charged lepton flavour violation in top-quark processes within the Standard Model Effective Field Theory framework. As the integrated luminosity delivered by the CERN Large Hadron Collider continues to increase, and with the advent of the High-Luminosity Large Hadron Collider, significantly improved sensitivity to rare top-quark decays will become achievable. This will allow the tightening of constraints on dimension-6 operators inducing charged lepton flavour violation, and potentially enable the first indirect hints of physics beyond the Standard Model in the top sector. On the theoretical side, further refinements, including higher-order corrections and global SMEFT fits incorporating complementary low-energy observables, will be essential to fully exploit the experimental precision. In this context, the methodology and results presented in this thesis can be extended to larger datasets and more comprehensive operator bases, thereby contributing to a coherent and systematically improvable programme of searches for flavour-violating new physics.

From the perspective of analysis techniques, several foreseeable improvements could substantially enhance the sensitivity to charged lepton flavour violation in top-quark processes. The adoption of more sophisticated multivariate strategies, such as deep neural networks or graph-based architectures tailored to event topology, would improve signal–background discrimination beyond boosted decision trees. In addition, refined object reconstruction, including improved flavour tagging, lepton identification, and pile-up mitigation, particularly relevant in the high-occupancy environment of the HL-LHC, will directly reduce systematic uncertainties and increase effective signal efficiency. The use of advanced background modelling techniques, data-driven constraint strategies, and simultaneous multi-channel fits within a global SMEFT interpretation could further stabilise the extraction of limits and strengthen their robustness against modelling assumptions. Finally, the incorporation of differential information and optimised observables sensitive to the Lorentz structure of dimension-6 operators would not only improve the overall reach but also enhance the capability to disentangle operator contributions in the event of a deviation from the Standard Model prediction.

The measurements presented in this thesis constitute a result being scrutinised by the ATLAS Top Group before publication. The analysis has been performed within the Universität Siegen ATLAS Top Physics Group, and is carried out by a research team with colleagues from other institutions, such as Universities of Birmingham and Sussex.

Bibliography

- [1] J. J. Thomson. “Cathode Rays”. In: *Lond.Edinb.Dubl.Phil.Mag.* 44 (1897). doi: [10.1080/14786449708621070](https://doi.org/10.1080/14786449708621070).
- [2] E. Rutherford. “LXXIX. The scattering of α and β particles by matter and the structure of the atom”. In: *Lond.Edinb.Dubl.Phil.Mag.* 21 (1911). doi: [10.1080/14786440508637080](https://doi.org/10.1080/14786440508637080).
- [3] C. D. Anderson. “The Positive Electron”. In: *Phys. Rev.* 43 (1933). doi: [10.1103/PhysRev.43.491](https://doi.org/10.1103/PhysRev.43.491).
- [4] J. Chadwick. “Possible Existence of a Neutron”. In: *Nature* 129 (1932). doi: [10.1038/129312a0](https://doi.org/10.1038/129312a0).
- [5] M. Gell-Mann. “A schematic model of baryons and mesons”. In: *Phys. Lett.* 8 (1964). doi: [https://doi.org/10.1016/S0031-9163\(64\)92001-3](https://doi.org/10.1016/S0031-9163(64)92001-3).
- [6] G. Zweig. *An SU_3 model for strong interaction symmetry and its breaking; Version 2*. 1964. doi: [10.17181/CERN-TH-412](https://doi.org/10.17181/CERN-TH-412).
- [7] E. D. Bloom et al. “High-Energy Inelastic $e - p$ Scattering at 6° and 10° ”. In: *Phys. Rev. Lett.* 23 (1969). doi: [10.1103/PhysRevLett.23.930](https://doi.org/10.1103/PhysRevLett.23.930).
- [8] M. Breidenbach et al. “Observed Behavior of Highly Inelastic Electron–Proton Scattering”. In: *Phys. Rev. Lett.* 23 (1969). doi: [10.1103/PhysRevLett.23.935](https://doi.org/10.1103/PhysRevLett.23.935).
- [9] D. Hanneke, S. Fogwell and G. Gabrielse. “New Measurement of the Electron Magnetic Moment and the Fine Structure Constant”. In: *Phys. Rev. Lett.* 100 (2008). doi: [10.1103/PhysRevLett.100.120801](https://doi.org/10.1103/PhysRevLett.100.120801).
- [10] T. Aoyama et al. “Tenth-order electron anomalous magnetic moment: Contribution of diagrams without closed lepton loops”. In: *Phys. Rev. D* 91 (2015). doi: [10.1103/PhysRevD.91.033006](https://doi.org/10.1103/PhysRevD.91.033006).
- [11] ATLAS Collaboration. “Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC”. In: *Phys. Lett. B* 716 (2012). doi: [10.1016/j.physletb.2012.08.020](https://doi.org/10.1016/j.physletb.2012.08.020).
- [12] CMS Collaboration. “Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC”. In: *Phys. Lett. B* 716 (2012). doi: [10.1016/j.physletb.2012.08.021](https://doi.org/10.1016/j.physletb.2012.08.021).
- [13] Planck Collaboration. “Planck 2018 results - VI. Cosmological parameters (Corrigendum)”. In: *AA* 652 (2021). doi: [10.1051/0004-6361/201833910e](https://doi.org/10.1051/0004-6361/201833910e).
- [14] Super-Kamiokande Collaboration. “Evidence for Oscillation of Atmospheric Neutrinos”. In: *Phys. Rev. Lett.* 81 (1998). doi: [10.1103/PhysRevLett.81.1562](https://doi.org/10.1103/PhysRevLett.81.1562).

- [15] SNO Collaboration. “Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory”. In: *Phys. Rev. Lett.* 89 (2002). DOI: [10.1103/PhysRevLett.89.011301](https://doi.org/10.1103/PhysRevLett.89.011301).
- [16] ATLAS Collaboration. *Search for charged lepton-flavour violation in top-quark decays at the LHC with the ATLAS detector*. Geneva, 2018. URL: <https://cds.cern.ch/record/2638305>.
- [17] ATLAS Collaboration. “Search for charged-lepton-flavour violating $\mu\tau qt$ interactions in top-quark production and decay in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector at the LHC”. In: *Phys. Rev. D* 110 (2024). DOI: [10.1103/physrevd.110.012014](https://doi.org/10.1103/physrevd.110.012014).
- [18] CMS Collaboration. “Search for charged-lepton flavor violation in top quark production and decay in pp collisions at $\sqrt{s} = 13$ TeV”. In: *JHEP* 2022 (2022). DOI: [10.1007/jhep06\(2022\)082](https://doi.org/10.1007/jhep06(2022)082).
- [19] CMS Collaboration. “Search for Charged-Lepton Flavor Violation in the Production and Decay of Top Quarks using Trilepton Final States in Proton-Proton Collisions at $\sqrt{s} = 13$ TeV”. In: *Phys. Rev. D* 111 (2025). DOI: [10.1103/PhysRevD.111.012009](https://doi.org/10.1103/PhysRevD.111.012009).
- [20] Particle Data Group. “Review of particle physics”. In: *Phys. Rev. D* 110 (2024). DOI: [10.1103/PhysRevD.110.030001](https://doi.org/10.1103/PhysRevD.110.030001).
- [21] D. Galbraith and C. Burgard. *Diagram of the Standard Model of particle physics*. URL: <https://texample.net/tikz/examples/model-physics/>. (accessed: 21.07.2025).
- [22] LHCb Collaboration. “Observation of the Resonant Character of the $Z(4430)^-$ State”. In: *Phys. Rev. Lett.* 112 (2014). DOI: [10.1103/PhysRevLett.112.222002](https://doi.org/10.1103/PhysRevLett.112.222002).
- [23] LHCb Collaboration. “Observation of $J/\psi p$ Resonances Consistent with Pentaquark States in $\Lambda_b^0 \rightarrow J/\psi K^- p$ Decays”. In: *Phys. Rev. Lett.* 115 (2015). DOI: [10.1103/PhysRevLett.115.072001](https://doi.org/10.1103/PhysRevLett.115.072001).
- [24] S. L. Glashow. “Partial-symmetries of weak interactions”. In: *Nucl. Phys.* 22 (1961). DOI: [https://doi.org/10.1016/0029-5582\(61\)90469-2](https://doi.org/10.1016/0029-5582(61)90469-2).
- [25] S. Weinberg. “A Model of Leptons”. In: *Phys. Rev. Lett.* 19 (1967). DOI: [10.1103/PhysRevLett.19.1264](https://doi.org/10.1103/PhysRevLett.19.1264).
- [26] A. Salam. “Weak and Electromagnetic Interactions”. In: *Conf. Proc. C* 680519 (1968). DOI: [10.1142/97898127959150034](https://doi.org/10.1142/97898127959150034).
- [27] F. Englert and R. Brout. “Broken Symmetry and the Mass of Gauge Vector Mesons”. In: *Phys. Rev. Lett.* 13 (1964). DOI: [10.1103/PhysRevLett.13.321](https://doi.org/10.1103/PhysRevLett.13.321).
- [28] P. W. Higgs. “Broken Symmetries and the Masses of Gauge Bosons”. In: *Phys. Rev. Lett.* 13 (1964). DOI: [10.1103/PhysRevLett.13.508](https://doi.org/10.1103/PhysRevLett.13.508).
- [29] G. S. Guralnik, C. R. Hagen and T. W. B. Kibble. “Global Conservation Laws and Massless Particles”. In: *Phys. Rev. Lett.* 13 (1964). DOI: [10.1103/PhysRevLett.13.585](https://doi.org/10.1103/PhysRevLett.13.585).
- [30] E. Fermi. “An attempt of a theory of beta radiation. 1.” In: *Z. Phys.* 88 (1934). DOI: [10.1007/BF01351864](https://doi.org/10.1007/BF01351864).
- [31] B. Henning, X. Lu and H. Murayama. *How to use the Standard Model effective field theory*. 2015. arXiv: [1412.1837](https://arxiv.org/abs/1412.1837) [hep-ph].
- [32] S. Bilenky, S. Petcov and B. Pontecorvo. “Lepton mixing, $\mu \rightarrow e + \gamma$ decay and neutrino oscillations”. In: *Phys. Lett. B* 67 (1977). DOI: [https://doi.org/10.1016/0370-2693\(77\)90379-3](https://doi.org/10.1016/0370-2693(77)90379-3).

-
- [33] “Charged lepton flavour violation: An experimental and theoretical introduction”. In: *La Rivista del Nuovo Cimento* 41 (2018). DOI: [10.1393/ncr/i2018-10144-0](https://doi.org/10.1393/ncr/i2018-10144-0).
 - [34] J. Hisano et al. “Lepton-flavor violation via right-handed neutrino Yukawa couplings in the supersymmetric standard model”. In: *Phys. Rev. D* 53 (1996). DOI: [10.1103/physrevd.53.2442](https://doi.org/10.1103/physrevd.53.2442).
 - [35] A. Ilakovac, A. Pilaftsis and L. Popov. “Charged lepton flavor violation in supersymmetric low-scale seesaw models”. In: *Phys. Rev. D* 87 (2013). DOI: [10.1103/physrevd.87.053014](https://doi.org/10.1103/physrevd.87.053014).
 - [36] G. Branco et al. “Theory and phenomenology of two-Higgs-doublet models”. In: *Phys. Rep.* 516 (2012). DOI: [10.1016/j.physrep.2012.02.002](https://doi.org/10.1016/j.physrep.2012.02.002).
 - [37] A. Crivellin, C. Greub and A. Kokulu. “Flavor-phenomenology of two-Higgs-doublet models with generic Yukawa structure”. In: *Phys. Rev. D* 87 (2013). DOI: [10.1103/physrevd.87.094031](https://doi.org/10.1103/physrevd.87.094031).
 - [38] A. Vicente. *Lepton flavor violation beyond the MSSM*. 2015. arXiv: [1503.08622](https://arxiv.org/abs/1503.08622) [[hep-ph](#)].
 - [39] I. Doršner et al. “Physics of leptoquarks in precision experiments and at particle colliders”. In: *Phys. Rep.* 641 (2016). DOI: [10.1016/j.physrep.2016.06.001](https://doi.org/10.1016/j.physrep.2016.06.001).
 - [40] LHCb Collaboration. “Test of Lepton Universality Using $B^+ \rightarrow K^+ \ell^+ \ell^-$ Decays”. In: *Phys. Rev. Lett.* 113 (2014). DOI: [10.1103/physrevlett.113.151601](https://doi.org/10.1103/physrevlett.113.151601).
 - [41] LHCb Collaboration. “Measurement of the Ratio of Branching Fractions $\mathcal{B}(\bar{B}^0 \rightarrow D^{*+} \tau^- \bar{\nu}_\tau) / \mathcal{B}(\bar{B}^0 \rightarrow D^{*+} \mu^- \bar{\nu}_\mu)$ ”. In: *Phys. Rev. Lett.* 115 (2015). DOI: [10.1103/PhysRevLett.115.111803](https://doi.org/10.1103/PhysRevLett.115.111803).
 - [42] LHCb Collaboration. “Test of lepton universality with $B^0 \rightarrow K^{*0} \ell^+ \ell^-$ decays”. In: *JHEP* 2017 (2017). DOI: [10.1007/jhep08\(2017\)055](https://doi.org/10.1007/jhep08(2017)055).
 - [43] S. L. Glashow, D. Guadagnoli and K. Lane. “Lepton Flavor Violation in B Decays?” In: *Phys. Rev. Lett.* 114 (2015). DOI: [10.1103/PhysRevLett.114.091801](https://doi.org/10.1103/PhysRevLett.114.091801).
 - [44] T. J. Kim et al. “Correlation between $R_{D^{(*)}}$ and top quark FCNC decays in leptoquark models”. In: *JHEP* 2019 (2019). DOI: [10.1007/jhep07\(2019\)025](https://doi.org/10.1007/jhep07(2019)025).
 - [45] S. Weinberg. “Baryon- and Lepton-Nonconserving Processes”. In: *Phys. Rev. Lett.* 43 (1979). DOI: [10.1103/PhysRevLett.43.1566](https://doi.org/10.1103/PhysRevLett.43.1566).
 - [46] B. Grzadkowski et al. “Dimension-six terms in the Standard Model Lagrangian”. In: *JHEP* 2010 (2010). DOI: [10.1007/jhep10\(2010\)085](https://doi.org/10.1007/jhep10(2010)085).
 - [47] M. Chala, J. Santiago and M. Spannowsky. “Constraining four-fermion operators using rare top decays”. In: *JHEP* 2019 (2019). DOI: [10.1007/jhep04\(2019\)014](https://doi.org/10.1007/jhep04(2019)014).
 - [48] S. Davidson et al. “Lepton flavour violating top decays at the LHC”. In: *EPJ C* 75 (2015). DOI: [10.1140/epjc/s10052-015-3649-5](https://doi.org/10.1140/epjc/s10052-015-3649-5).
 - [49] MEG Collaboration. *Search for the Lepton Flavour Violating Decay $\mu^+ \rightarrow e^+ \gamma$ with the Full Dataset of the MEG Experiment*. 2016. arXiv: [1605.05081](https://arxiv.org/abs/1605.05081) [[hep-ex](#)].
 - [50] U. Bellgardt et al. “Search for the Decay $\mu^+ \rightarrow e^+ e^+ e^-$ ”. In: *Nucl. Phys. B* 299 (1988). DOI: [10.1016/0550-3213\(88\)90462-2](https://doi.org/10.1016/0550-3213(88)90462-2).
 - [51] “Charged Lepton Flavour Violation: An Experimental and Theoretical Introduction”. In: *Riv. Nuovo Cimento* 41 (2018). DOI: [10.1393/ncr/i2018-10144-0](https://doi.org/10.1393/ncr/i2018-10144-0).
 - [52] R. Bernstein and P. S. Cooper. “Charged lepton flavor violation: An experimenter’s guide”. In: *Phys. Rep.* 532.2 (2013). DOI: [10.1016/j.physrep.2013.07.002](https://doi.org/10.1016/j.physrep.2013.07.002).

- [53] ATLAS Collaboration. “Probing lepton flavour violation via neutrinoless $\tau \rightarrow 3\mu$ decays with the ATLAS detector”. In: *EPJ C* 76 (2016). DOI: [10.1140/epjc/s10052-016-4041-9](https://doi.org/10.1140/epjc/s10052-016-4041-9).
- [54] ATLAS Collaboration. “Search for the lepton flavor violating decay $Z \rightarrow e\mu$ in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector”. In: *Phys. Rev. D* 90 (2014). DOI: [10.1103/PhysRevD.90.072010](https://doi.org/10.1103/PhysRevD.90.072010).
- [55] ATLAS Collaboration. “Search for lepton-flavor-violating decays of the Z boson into a τ lepton and a light lepton with the ATLAS detector”. In: *Phys. Rev. D* 98 (2018). DOI: [10.1103/PhysRevD.98.092010](https://doi.org/10.1103/PhysRevD.98.092010).
- [56] ATLAS Collaboration. “Search for lepton-flavour-violating decays of the Higgs and Z bosons with the ATLAS detector”. In: *EPJ C* 77 (2017). DOI: [10.1140/epjc/s10052-017-4624-0](https://doi.org/10.1140/epjc/s10052-017-4624-0).
- [57] ATLAS Collaboration. “Search for lepton-flavor violation in different-flavor, high-mass final states in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *Phys. Rev. D* 98 (2018). DOI: [10.1103/PhysRevD.98.092008](https://doi.org/10.1103/PhysRevD.98.092008).
- [58] ATLAS Collaboration. “Search for lepton flavour violation in the $e\mu$ continuum with the ATLAS detector in $\sqrt{s} = 7$ TeV pp collisions at the LHC”. In: *EPJ C* 72 (2012). DOI: [10.1140/epjc/s10052-012-2040-z](https://doi.org/10.1140/epjc/s10052-012-2040-z).
- [59] S. Banerjee et al. *Snowmass 2021 White Paper: Charged lepton flavor violation in the tau sector*. 2022. arXiv: [2203.14919](https://arxiv.org/abs/2203.14919) [[hep-ph](#)].
- [60] ATLAS Collaboration. “Searches for lepton-flavour-violating decays of the Higgs boson into $e\tau$ and $\mu\tau$ in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector”. In: *JHEP* 2023 (2023). DOI: [10.1007/JHEP07\(2023\)166](https://doi.org/10.1007/JHEP07(2023)166).
- [61] CMS Collaboration. “Search for lepton-flavor violating decays of the Higgs boson in the $\mu\tau$ and $e\tau$ final states in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Phys. Rev. D* 104 (2021). DOI: [10.1103/PhysRevD.104.032013](https://doi.org/10.1103/PhysRevD.104.032013).
- [62] ATLAS Collaboration. “Search for Lepton-Flavor Violation in Z -Boson Decays with τ Leptons with the ATLAS Detector”. In: *Phys. Rev. Lett.* 127 (2021). DOI: [10.1103/PhysRevLett.127.271801](https://doi.org/10.1103/PhysRevLett.127.271801).
- [63] Belle Collaboration. “Search for lepton-flavor-violating tau decays into three leptons”. In: *Phys. Lett. B* (2010). DOI: [10.1016/j.physletb.2010.03.037](https://doi.org/10.1016/j.physletb.2010.03.037).
- [64] BaBar Collaboration. “Limits on tau lepton-flavor violating decays in three charged leptons”. In: *Phys. Rev. D* (2010). DOI: [10.1103/PhysRevD.81.111101](https://doi.org/10.1103/PhysRevD.81.111101).
- [65] LHCb Collaboration. “Search for the lepton flavour violating decay $\tau^- \rightarrow \mu^- \mu^+ \mu^-$ ”. In: *JHEP* 2015 (2015). DOI: [10.1007/JHEP02\(2015\)121](https://doi.org/10.1007/JHEP02(2015)121).
- [66] H1 Collaboration. “Search for single top quark production at HERA”. In: *Phys. Lett. B* 678 (2009). DOI: [10.1016/j.physletb.2009.06.057](https://doi.org/10.1016/j.physletb.2009.06.057).
- [67] C. Gottardo. “Search for charged lepton-flavour violation in top-quark decays at the LHC with the ATLAS detector”. PhD thesis. Universität Bonn, 2019. URL: <https://cds.cern.ch/record/2676841>.
- [68] A. L. Read. “Presentation of search results: the CL_s technique”. In: *J. Phys. G: Nucl. Part. Phys* 28 (2002). DOI: [10.1088/0954-3899/28/10/313](https://doi.org/10.1088/0954-3899/28/10/313).

-
- [69] CMS Collaboration. “Search for charged-lepton flavor violation in the production and decay of top quarks using trilepton final states in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *Phys. Rev. D* 111 (2025). DOI: [10.1103/PhysRevD.111.012009](https://doi.org/10.1103/PhysRevD.111.012009).
- [70] CMS Collaboration. “Search for charged-lepton flavour violation in top quark interactions with an up-type quark, a muon, and a τ lepton in proton-proton collisions at $\sqrt{s} = 13$ TeV”. In: *JHEP* (2025). DOI: [10.1007/jhep12\(2025\)065](https://doi.org/10.1007/jhep12(2025)065).
- [71] L. Evans and P. Bryant. “LHC Machine”. In: *JINST* 3 (2008). DOI: [10.1088/1748-0221/3/08/s08001](https://doi.org/10.1088/1748-0221/3/08/s08001).
- [72] ATLAS Collaboration. “The ATLAS Experiment at the CERN Large Hadron Collider”. In: *JINST* 3 (2008). DOI: [10.1088/1748-0221/3/08/S08003](https://doi.org/10.1088/1748-0221/3/08/S08003).
- [73] E. Lopienska. “The CERN accelerator complex, layout in 2022. Complexe des accélérateurs du CERN en janvier 2022”. In: (2022). General Photo. URL: <https://cds.cern.ch/record/2800984>.
- [74] ATLAS Collaboration. “Experiment Briefing: Keeping the ATLAS Inner Detector in perfect alignment”. General Photo, 2020. URL: <https://cds.cern.ch/record/2723878>.
- [75] ATLAS Collaboration. *Public ATLAS Luminosity Results for Run-2 of the LHC*. URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2>. (accessed: 05.07.2024).
- [76] M. Capeans et al. *ATLAS Insertable B-Layer Technical Design Report*. Tech. rep. 2010. URL: <https://cds.cern.ch/record/1291633>.
- [77] G. Ripellino. *The alignment of the ATLAS Inner Detector in Run 2*. Tech. rep. Geneva: CERN, 2016. DOI: [10.22323/1.276.0196](https://doi.org/10.22323/1.276.0196).
- [78] The ATLAS TDAQ Collaboration. “The ATLAS Data Acquisition and High Level Trigger system”. In: *JINST* 11 (2016). DOI: [10.1088/1748-0221/11/06/P06008](https://doi.org/10.1088/1748-0221/11/06/P06008).
- [79] ATLAS Collaboration. “Operation of the ATLAS trigger system in Run 2”. In: *JINST* 15 (2020). DOI: [10.1088/1748-0221/15/10/P10004](https://doi.org/10.1088/1748-0221/15/10/P10004).
- [80] T. Plehn. *Lectures on LHC Physics*. Springer Berlin Heidelberg, 2012. DOI: [10.1007/978-3-642-24040-9](https://doi.org/10.1007/978-3-642-24040-9).
- [81] R. D. Ball et al. “Parton distributions for the LHC run II”. In: *JHEP* 2015 (2015). DOI: [10.1007/jhep04\(2015\)040](https://doi.org/10.1007/jhep04(2015)040).
- [82] HERA Study Group. *HERA - A Proposal for a Large Electron Proton Colliding Beam Facility at DESY*. Tech. rep. 1981. DOI: [10.3204/PUBDB-2017-01301](https://doi.org/10.3204/PUBDB-2017-01301).
- [83] R. R. Wilson. “The Tevatron”. In: *Physics Today* (1977). URL: <https://doi.org/10.1063/1.3037746>.
- [84] J. Butterworth et al. “PDF4LHC recommendations for LHC Run II”. In: *J. Phys. G: Nucl. Part. Phys* 43 (2016). DOI: [10.1088/0954-3899/43/2/023001](https://doi.org/10.1088/0954-3899/43/2/023001).
- [85] J. M. Campbell, J. W. Huston and W. J. Stirling. “Hard interactions of quarks and gluons: a primer for LHC physics”. In: *Rep. Prog. Phys* 70 (2006). DOI: [10.1088/0034-4885/70/1/r02](https://doi.org/10.1088/0034-4885/70/1/r02).
- [86] T. Sjöstrand. *Monte Carlo Generators*. 2007. URL: <https://cds.cern.ch/record/999717>.

- [87] G. Altarelli and G. Parisi. “Asymptotic freedom in parton language”. In: *Nucl. Phys. B* 126 (1977). DOI: [https://doi.org/10.1016/0550-3213\(77\)90384-4](https://doi.org/10.1016/0550-3213(77)90384-4).
- [88] D. de Florian and M. Grazzini. “Next-to-Next-to-Leading-Order Logarithmic Corrections at Small Transverse Momentum in Hadronic Collisions”. In: *Phys. Rev. Lett.* 85 (2000). DOI: [10.1103/physrevlett.85.4678](https://doi.org/10.1103/physrevlett.85.4678).
- [89] Z. Nagy and D. E. Soper. “What is a parton shower?”. In: *Phys. Rev. D* 98 (2018). DOI: [10.1103/physrevd.98.014034](https://doi.org/10.1103/physrevd.98.014034).
- [90] E. Bothmann et al. “Event generation with Sherpa 2.2”. In: *SciPost Phys.* (2019). DOI: [10.21468/scipostphys.7.3.034](https://doi.org/10.21468/scipostphys.7.3.034).
- [91] C. Ay et al. “Monte Carlo generators in ATLAS software”. In: *J.Phys.: Conference Series* 219 (2010). DOI: [10.1088/1742-6596/219/3/032001](https://doi.org/10.1088/1742-6596/219/3/032001).
- [92] T. Sjöstrand, S. Mrenna and P. Skands. “A brief introduction to PYTHIA 8.1”. In: *CPC* 178 (2008). DOI: [10.1016/j.cpc.2008.01.036](https://doi.org/10.1016/j.cpc.2008.01.036).
- [93] D. J. Lange. “The EvtGen particle decay simulation package”. In: *Nucl. Instrum. Methods Phys. Res. A* 462 (2001). DOI: [https://doi.org/10.1016/S0168-9002\(01\)00089-4](https://doi.org/10.1016/S0168-9002(01)00089-4).
- [94] B. Andersson et al. “Parton fragmentation and string dynamics”. In: *Phys. Rep.* 97 (1983). DOI: [https://doi.org/10.1016/0370-1573\(83\)90080-7](https://doi.org/10.1016/0370-1573(83)90080-7).
- [95] B. Andersson. “The Lund Model”. In: *Nucl. Phys. A* 461 (1987). DOI: [10.1016/0375-9474\(87\)90510-0](https://doi.org/10.1016/0375-9474(87)90510-0).
- [96] S. Catani et al. “QCD Matrix Elements + Parton Showers”. In: *JHEP* 2001 (2001). DOI: [10.1088/1126-6708/2001/11/063](https://doi.org/10.1088/1126-6708/2001/11/063).
- [97] K. Hamilton and P. Nason. “Improving NLO-parton shower matched simulations with higher order matrix elements”. In: *JHEP* 2010 (2010). DOI: [10.1007/jhep06\(2010\)039](https://doi.org/10.1007/jhep06(2010)039).
- [98] G. C. Fox and S. Wolfram. “A model for parton showers in QCD”. In: *Nucl. Phys. B* 168 (1980). DOI: [https://doi.org/10.1016/0550-3213\(80\)90111-X](https://doi.org/10.1016/0550-3213(80)90111-X).
- [99] 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”. In: *JHEP* 2014 (2014). DOI: [10.1007/jhep07\(2014\)079](https://doi.org/10.1007/jhep07(2014)079).
- [100] A. Alloul et al. “FeynRules 2.0 — A complete toolbox for tree-level phenomenology”. In: *CPC* 185 (2014). DOI: [10.1016/j.cpc.2014.04.012](https://doi.org/10.1016/j.cpc.2014.04.012).
- [101] C. Degrande et al. “UFO – The Universal FeynRules Output”. In: *CPC* 183 (2012). DOI: [10.1016/j.cpc.2012.01.022](https://doi.org/10.1016/j.cpc.2012.01.022).
- [102] J. Alwall et al. “A standard format for Les Houches Event Files”. In: *CPC* 176 (2007). DOI: [10.1016/j.cpc.2006.11.010](https://doi.org/10.1016/j.cpc.2006.11.010).
- [103] S. Alioli et al. “A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX”. In: *JHEP* 2010 (2010). DOI: [10.1007/jhep06\(2010\)043](https://doi.org/10.1007/jhep06(2010)043).
- [104] S. Frixione, P. Nason and C. Oleari. “Matching NLO QCD computations with parton shower simulations: the POWHEG method”. In: *JHEP* 2007 (2007). DOI: [10.1088/1126-6708/2007/11/070](https://doi.org/10.1088/1126-6708/2007/11/070).

-
- [105] S. Frixione and B. R. Webber. “Matching NLO QCD computations and parton shower simulations”. In: *JHEP* 2002 (2002). DOI: [10.1088/1126-6708/2002/06/029](https://doi.org/10.1088/1126-6708/2002/06/029).
- [106] A. Buckley et al. “General-purpose event generators for LHC physics”. In: *Phys. Rep.* 504 (2011). DOI: [10.1016/j.physrep.2011.03.005](https://doi.org/10.1016/j.physrep.2011.03.005).
- [107] J. Bellm et al. “Herwig 7.0/Herwig++ 3.0 release note”. In: *EPJ C* 76 (2016). DOI: [10.1140/epjc/s10052-016-4018-8](https://doi.org/10.1140/epjc/s10052-016-4018-8).
- [108] J. Bellm et al. *Herwig 7.1 Release Note*. 2017. arXiv: [1705.06919](https://arxiv.org/abs/1705.06919) [[hep-ph](#)].
- [109] ATLAS Collaboration. *ATLAS Pythia 8 tunes to 7 TeV data*. Tech. rep. Geneva: CERN, 2014. URL: <https://cds.cern.ch/record/1966419>.
- [110] R. D. Ball et al. “Parton distributions with LHC data”. In: *Nucl. Phys. B* 867 (2013). DOI: [10.1016/j.nuclphysb.2012.10.003](https://doi.org/10.1016/j.nuclphysb.2012.10.003).
- [111] S. Schumann and F. Krauss. “A parton shower algorithm based on Catani-Seymour dipole factorisation”. In: *JHEP* 2008 (2008). DOI: [10.1088/1126-6708/2008/03/038](https://doi.org/10.1088/1126-6708/2008/03/038).
- [112] ATLAS Collaboration. *The Pythia 8 A3 tune description of ATLAS minimum bias and inelastic measurements incorporating the Donnachie-Landshoff diffractive model*. Tech. rep. Geneva: CERN, 2016. URL: <https://cds.cern.ch/record/2206965>.
- [113] S. Agostinelli et al. “Geant4—a simulation toolkit”. In: *Nucl. Instrum. Methods Phys. Res. A* 506 (2003). DOI: [https://doi.org/10.1016/S0168-9002\(03\)01368-8](https://doi.org/10.1016/S0168-9002(03)01368-8).
- [114] ATLAS Collaboration et al. *The simulation principle and performance of the ATLAS fast calorimeter simulation FastCaloSim*. Geneva, 2010. URL: <https://cds.cern.ch/record/1300517>.
- [115] W. Lukas. “Fast Simulation for ATLAS: Atfast-II and ISF”. In: *J.Phys.: Conference Series* 396 (2012). DOI: [10.1088/1742-6596/396/2/022031](https://doi.org/10.1088/1742-6596/396/2/022031).
- [116] T. Cornelissen et al. “The new ATLAS track reconstruction (NEWT)”. In: *J.Phys.: Conference Series* 119 (2008). DOI: [10.1088/1742-6596/119/3/032014](https://doi.org/10.1088/1742-6596/119/3/032014).
- [117] V. Kartvelishvili. “Electron bremsstrahlung recovery in ATLAS”. In: *Nucl. Phys. B - Proceedings Supplements* 172 (2007). DOI: <https://doi.org/10.1016/j.nuclphysbps.2007.08.137>.
- [118] ATLAS Collaboration. *Characterization of Interaction-Point Beam Parameters Using the pp Event-Vertex Distribution Reconstructed in the ATLAS Detector at the LHC*. Tech. rep. Geneva: CERN, 2010. URL: <https://cds.cern.ch/record/1277659>.
- [119] ATLAS Collaboration. *Vertex Reconstruction Performance of the ATLAS Detector at $\sqrt{s} = 13$ TeV*. Geneva, 2015. URL: <https://cds.cern.ch/record/2037717>.
- [120] ATLAS Collaboration. *ATLAS Track Reconstruction – General Overview*. <https://atlassoftwaredocs.web.cern.ch/internal-links/tracking-tutorial/idoverview/>. Accessed: 2025-01-07. 2024.
- [121] ATLAS Collaboration. *Electron efficiency measurements with the ATLAS detector using the 2015 LHC proton-proton collision data*. Tech. rep. Geneva: CERN, 2016. URL: <https://cds.cern.ch/record/2157687>.
- [122] ATLAS Collaboration. “Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data”. In: *JINST* 14 (2019). DOI: [10.1088/1748-0221/14/12/p12006](https://doi.org/10.1088/1748-0221/14/12/p12006).

- [123] F. He. “High Precision Electron and Muon Reconstruction Performance with ATLAS at LHC Run-2”. In: *PoS ICHEP2022* (2022). DOI: [10.22323/1.414.0675](https://doi.org/10.22323/1.414.0675).
- [124] ATLAS Collaboration. “Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at $\sqrt{s} = 13$ TeV”. In: *EPJ C* 79 (2019). DOI: [10.1140/epjc/s10052-019-7140-6](https://doi.org/10.1140/epjc/s10052-019-7140-6).
- [125] ATLAS Collaboration. “Tools for estimating fake/non-prompt lepton backgrounds with the ATLAS detector at the LHC”. In: *JINST* 18 (2023). DOI: [10.1088/1748-0221/18/11/t11004](https://doi.org/10.1088/1748-0221/18/11/t11004).
- [126] ATLAS Collaboration. “Electron and photon efficiencies in LHC Run 2 with the ATLAS experiment”. In: *JHEP* 2405 (2024). DOI: [10.1007/JHEP05\(2024\)162](https://doi.org/10.1007/JHEP05(2024)162).
- [127] ATLAS Collaboration. “Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV”. In: *EPJ C* 76 (2016). DOI: [10.1140/epjc/s10052-016-4120-y](https://doi.org/10.1140/epjc/s10052-016-4120-y).
- [128] ATLAS Collaboration. “Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV”. In: *EPJ C* 81 (2021). DOI: [10.1140/epjc/s10052-021-09233-2](https://doi.org/10.1140/epjc/s10052-021-09233-2).
- [129] ATLAS Collaboration. “Studies of the muon momentum calibration and performance of the ATLAS detector with pp collisions at $\sqrt{s} = 13$ TeV”. In: *EPJ C* 83 (2023). DOI: [10.1140/epjc/s10052-023-11584-x](https://doi.org/10.1140/epjc/s10052-023-11584-x).
- [130] W. Lampl et al. *Calorimeter Clustering Algorithms: Description and Performance*. Tech. rep. Geneva: CERN, 2008. URL: <https://cds.cern.ch/record/1099735>.
- [131] ATLAS Collaboration. “Jet energy measurement and its systematic uncertainty in proton–proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector”. In: *EPJ C* 75 (2015). DOI: [10.1140/epjc/s10052-014-3190-y](https://doi.org/10.1140/epjc/s10052-014-3190-y).
- [132] ATLAS Collaboration. “Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1”. In: *EPJ C* 77 (2017). DOI: [10.1140/epjc/s10052-017-5004-5](https://doi.org/10.1140/epjc/s10052-017-5004-5).
- [133] M. Cacciari, G. P. Salam and G. Soyez. “The anti- k_t jet clustering algorithm”. In: *JHEP* 2008 (2008). DOI: [10.1088/1126-6708/2008/04/063](https://doi.org/10.1088/1126-6708/2008/04/063).
- [134] ATLAS Collaboration. “Jet reconstruction and performance using particle flow with the ATLAS Detector”. In: *EPJ C* 77 (2017). DOI: [10.1140/epjc/s10052-017-5031-2](https://doi.org/10.1140/epjc/s10052-017-5031-2).
- [135] ATLAS Collaboration. *Selection of jets produced in 13TeV proton-proton collisions with the ATLAS detector*. Tech. rep. Geneva: CERN, 2015. URL: <https://cds.cern.ch/record/2037702>.
- [136] ATLAS Collaboration. *Tagging and suppression of pileup jets with the ATLAS detector*. Tech. rep. Geneva: CERN, 2014. URL: <https://cds.cern.ch/record/1700870>.
- [137] ATLAS Collaboration. “Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *EPJ C* 81 (2021). DOI: [10.1140/epjc/s10052-021-09402-3](https://doi.org/10.1140/epjc/s10052-021-09402-3).
- [138] ATLAS Collaboration. “ATLAS flavour-tagging algorithms for the LHC Run 2 pp collision dataset”. In: *EPJ C* 83 (2023). DOI: [10.1140/epjc/s10052-023-11699-1](https://doi.org/10.1140/epjc/s10052-023-11699-1).

-
- [139] ATLAS Collaboration. “ATLAS b-jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV”. In: *EPJ C* 79 (2019). DOI: [10.1140/epjc/s10052-019-7450-8](https://doi.org/10.1140/epjc/s10052-019-7450-8).
 - [140] ATLAS Collaboration. “Measurement of the c-jet mistagging efficiency in $t\bar{t}$ events using pp collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS detector”. In: *EPJ C* 82 (2022). DOI: [10.1140/epjc/s10052-021-09843-w](https://doi.org/10.1140/epjc/s10052-021-09843-w).
 - [141] ATLAS Collaboration. “Calibration of the light-flavour jet mistagging efficiency of the b-tagging algorithms with Z+jets events using 139 fb^{-1} of ATLAS proton–proton collision data at $\sqrt{s} = 13$ TeV”. In: *EPJ C* 83 (2023). DOI: [10.1140/epjc/s10052-023-11736-z](https://doi.org/10.1140/epjc/s10052-023-11736-z).
 - [142] DØ Collaboration. *Observation of single top quark production*. 2009. URL: https://www-d0.fnal.gov/Run2Physics/top/singletop_observation/singletop_%20observation_updated.html.
 - [143] ATLAS Collaboration. E_T^{miss} performance in the ATLAS detector using 2015–2016 LHC pp collisions. Tech. rep. Geneva: CERN, 2018. URL: <https://cds.cern.ch/record/2625233>.
 - [144] ATLAS Collaboration. “Performance of missing transverse momentum reconstruction with the ATLAS detector using proton–proton collisions at $\sqrt{s} = 13$ TeV”. In: *EPJ C* 78 (2018). DOI: [10.1140/epjc/s10052-018-6288-9](https://doi.org/10.1140/epjc/s10052-018-6288-9).
 - [145] D. Adams et al. *Recommendations of the Physics Objects and Analysis Harmonisation Study Groups 2014*. Tech. rep. Geneva: CERN, 2014. URL: <https://cds.cern.ch/record/1743654>.
 - [146] M. Cacciari, G. P. Salam and G. Soyez. “The catchment area of jets”. In: *JHEP* 2008 (2008). DOI: [10.1088/1126-6708/2008/04/005](https://doi.org/10.1088/1126-6708/2008/04/005).
 - [147] CDF Collaboration. “Observation of Top Quark Production in $\bar{p}p$ Collisions with the Collider Detector at Fermilab”. In: *Phys. Rev. Lett.* 74 (1995). DOI: [10.1103/PhysRevLett.74.2626](https://doi.org/10.1103/PhysRevLett.74.2626).
 - [148] DØ Collaboration. “Observation of the Top Quark”. In: *Phys. Rev. Lett.* 74 (1995). DOI: [10.1103/PhysRevLett.74.2632](https://doi.org/10.1103/PhysRevLett.74.2632).
 - [149] A. Collaboration. “Inclusive and differential cross-sections for dilepton $t\bar{t}$ production measured in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector”. In: *JHEP* 07 (2023). DOI: [10.1007/JHEP07\(2023\)141](https://doi.org/10.1007/JHEP07(2023)141).
 - [150] M. Czakon, D. Heymes and A. Mitov. “High-precision differential predictions for top-quark pairs at the LHC”. In: *Phys. Rev. Lett.* 118 (2017). DOI: [10.1103/PhysRevLett.118.082003](https://doi.org/10.1103/PhysRevLett.118.082003).
 - [151] ATLAS Collaboration. *Top cross section summary plots - April 2024*. Tech. rep. Geneva: CERN, 2024. URL: <https://cds.cern.ch/record/2896104>.
 - [152] G. Cowan et al. “Asymptotic formulae for likelihood-based tests of new physics”. In: *EPJ C* 71 (2011). DOI: [10.1140/epjc/s10052-011-1554-0](https://doi.org/10.1140/epjc/s10052-011-1554-0).
 - [153] L.-G. Xia. “Study of constraint and impact of a nuisance parameter in maximum likelihood method”. In: *J. Phys. G: Nucl. Part. Phys* 46 (2019). DOI: [10.1088/1361-6471/ab02c0](https://doi.org/10.1088/1361-6471/ab02c0).
 - [154] K. Cranmer. *Practical Statistics for the LHC*. 2015. URL: <https://arxiv.org/abs/1503.07622>.
 - [155] S. S. Wilks. “The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses”. In: *Ann. Math. Statist.* 9 (1938). DOI: [10.1214/aoms/1177732360](https://doi.org/10.1214/aoms/1177732360).
 - [156] A. Wald. “On Cumulative Sums of Random Variables”. In: *Ann. Math. Statist.* 15 (1944). DOI: [10.1214/aoms/1177731235](https://doi.org/10.1214/aoms/1177731235).

- [157] J. A. A. Saavedra et al. *Interpreting top-quark LHC measurements in the Standard-Model Effective Field Theory*. 2018. arXiv: [1802.07237 \[hep-ph\]](#).
- [158] Particle Data Group Collaboration. “Review of Particle Physics”. In: *Phys. Rev. D* 110 (3 2024). DOI: [10.1103/PhysRevD.110.030001](#). URL: <https://link.aps.org/doi/10.1103/PhysRevD.110.030001>.
- [159] R. D. Ball et al. “Parton distributions for the LHC run II”. In: *JHEP* 2015 (2015). DOI: [10.1007/jhep04\(2015\)040](#).
- [160] ATLAS Collaboration. *Studies on top-quark Monte Carlo modelling for Top2016*. Tech. rep. Geneva: CERN, 2016. URL: <https://cds.cern.ch/record/2216168>.
- [161] M. Beneke et al. “Hadronic top-quark pair production with NNLL threshold resummation”. In: *Nucl. Phys. B* 855 (2012). DOI: [10.1016/j.nuclphysb.2011.10.021](#).
- [162] M. Cacciari et al. “Top-pair production at hadron colliders with next-to-next-to-leading logarithmic soft-gluon resummation”. In: *Phys. Lett. B* 710 (2012). DOI: [10.1016/j.physletb.2012.03.013](#).
- [163] P. Bärnreuther, M. Czakon and A. Mitov. “Percent-Level-Precision Physics at the Tevatron: Next-to-Next-to-Leading Order QCD Corrections to $q\bar{q} \rightarrow t\bar{t} + X$ ”. In: *Phys. Rev. Lett.* 109 (2012). DOI: [10.1103/physrevlett.109.132001](#).
- [164] M. Czakon and A. Mitov. “NNLO corrections to top-pair production at hadron colliders: the all-fermionic scattering channels”. In: *JHEP* 2012 (2012). DOI: [10.1007/jhep12\(2012\)054](#).
- [165] M. Czakon and A. Mitov. “NNLO corrections to top pair production at hadron colliders: the quark-gluon reaction”. In: *JHEP* 2013 (2013). DOI: [10.1007/jhep01\(2013\)080](#).
- [166] M. Czakon, P. Fiedler and A. Mitov. “Total Top-Quark Pair-Production Cross Section at Hadron Colliders Through $\mathcal{O}(\alpha_s^4)$ ”. In: *Phys. Rev. Lett.* 110 (2013). DOI: [10.1103/physrevlett.110.252004](#).
- [167] M. Czakon and A. Mitov. “Top++: A program for the calculation of the top-pair cross-section at hadron colliders”. In: *CPC* 185 (2014). DOI: [10.1016/j.cpc.2014.06.021](#).
- [168] ATLAS Collaboration. “Measurement of the charge asymmetry in top-quark pair production in association with a photon with the ATLAS experiment”. In: *Phys. Lett. B* 843 (2023). DOI: [10.1016/j.physletb.2023.137848](#).
- [169] S. Frixione et al. “Single-top hadroproduction in association with a W boson”. In: *JHEP* 2008 (2008). DOI: [10.1088/1126-6708/2008/07/029](#).
- [170] N. Kidonakis. “Two-loop soft anomalous dimensions for single top quark associated production with a W^- or H^- ”. In: *Phys. Rev. D* 82 (2010). DOI: [10.1103/physrevd.82.054018](#).
- [171] N. Kidonakis. *Top Quark Production*. 2013. arXiv: [1311.0283 \[hep-ph\]](#).
- [172] A. Giammanco. “Single top quark production at the LHC”. In: *Reviews in Physics* 1 (2016). DOI: <https://doi.org/10.1016/j.revip.2015.12.001>.
- [173] C. D. White et al. “Isolating Wt production at the LHC”. In: *JHEP* 2009 (2009). DOI: [10.1088/1126-6708/2009/11/074](#). URL: <https://dx.doi.org/10.1088/1126-6708/2009/11/074>.

-
- [174] ATLAS Collaboration. *Comparison of vertex reconstruction quantities in ATLAS software releases 21 and 20.7 for 2016 proton-proton data*. Tech. rep. Geneva: CERN, 2017. URL: <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/IDTR-2017-006/>.
- [175] ATLAS Collaboration. *Electron trigger recommendations*. URL: <https://twiki.cern.ch/twiki/bin/view/Atlas/TrigEgammaRecommendations2015>.
- [176] ATLAS Collaboration. *Muon trigger recommendations*. URL: <https://twiki.cern.ch/twiki/bin/view/Atlas/MuonTriggerPhysicsRecommendationsRel212017>.
- [177] A. Hoecker et al. *TMVA - Toolkit for Multivariate Data Analysis*. 2009. arXiv: [physics/0703039](https://arxiv.org/abs/physics/0703039) [[physics.data-an](https://arxiv.org/abs/physics/0703039)].
- [178] J. H. Friedman. “Greedy Function Approximation: A Gradient Boosting Machine”. In: *The Annals of Statistics* 29 (2001). DOI: [10.1214/aos/1013203451](https://doi.org/10.1214/aos/1013203451).
- [179] T. Chen and C. Guestrin. “XGBoost: A Scalable Tree Boosting System”. In: *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. KDD ’16. San Francisco, California, USA: Association for Computing Machinery, 2016. DOI: [10.1145/2939672.2939785](https://doi.org/10.1145/2939672.2939785).
- [180] ATLAS Collaboration. “Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC”. In: *EPJ C* 83 (2023). DOI: [10.1140/epjc/s10052-023-11747-w](https://doi.org/10.1140/epjc/s10052-023-11747-w).
- [181] ATLAS Collaboration. “Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton-proton collision data at $\sqrt{s} = 13$ TeV”. In: *EPJ C* 79 (2019). DOI: [10.1140/epjc/s10052-019-7140-6](https://doi.org/10.1140/epjc/s10052-019-7140-6).
- [182] ATLAS Collaboration. “Electron and photon energy calibration with the ATLAS detector using 2015-2016 LHC proton-proton collision data”. In: *JINST* 14 (2019). DOI: [10.1088/1748-0221/14/03/p03017](https://doi.org/10.1088/1748-0221/14/03/p03017).
- [183] ATLAS Collaboration. “Studies of the muon momentum calibration and performance of the ATLAS detector with pp collisions at $\sqrt{s} = 13$ TeV”. In: *EPJ C* 83 (2023). DOI: [10.1140/epjc/s10052-023-11584-x](https://doi.org/10.1140/epjc/s10052-023-11584-x).
- [184] ATLAS Collaboration. “Jet energy scale measurements and their systematic uncertainties in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”. In: *Phys. Rev. D* 96 (2017). DOI: [10.1103/physrevd.96.072002](https://doi.org/10.1103/physrevd.96.072002).
- [185] ATLAS Collaboration. “Jet energy resolution in proton-proton collisions at $\sqrt{s} = 7$ TeV recorded in 2010 with the ATLAS detector”. In: *EPJ C* 73 (2013). DOI: [10.1140/epjc/s10052-013-2306-0](https://doi.org/10.1140/epjc/s10052-013-2306-0).
- [186] ATLAS Collaboration. “Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector”. In: *EPJ C* 76 (2016). DOI: [10.1140/epjc/s10052-016-4395-z](https://doi.org/10.1140/epjc/s10052-016-4395-z).
- [187] ATLAS Collaboration. “Measurements of b -jet tagging efficiency with the ATLAS detector using $t\bar{t}$ events at $\sqrt{s} = 13$ TeV”. In: *JHEP* 2018 (2018). DOI: [10.1007/jhep08\(2018\)089](https://doi.org/10.1007/jhep08(2018)089).
- [188] R. J. Barlow and C. Beeston. “Fitting using finite Monte Carlo samples”. In: *CPC* (1993). DOI: [10.1016/0010-4655\(93\)90005-W](https://doi.org/10.1016/0010-4655(93)90005-W).

- [189] K. Cranmer et al. *HistFactory: A tool for creating statistical models for use with RooFit and RooStats*. Tech. rep. New York: New York U., 2012. doi: [10.17181/CERN-OPEN-2012-016](https://doi.org/10.17181/CERN-OPEN-2012-016).
- [190] J. S. Conway. “Incorporating nuisance parameters in likelihoods for multisource spectra”. In: *PHYSTAT 2011*. 2011. doi: [10.5170/CERN-2011-006.115](https://doi.org/10.5170/CERN-2011-006.115).
- [191] ATLAS Collaboration. “Measurement of the $t\bar{t}$ production cross-section in the lepton+jets channel at $\sqrt{s} = 13$ TeV with the ATLAS experiment”. In: *Phys. Lett. B* 810 (2020). doi: [10.1016/j.physletb.2020.135797](https://doi.org/10.1016/j.physletb.2020.135797).
- [192] S. Frixione, G. Ridolfi and P. Nason. “A positive-weight next-to-leading-order Monte Carlo for heavy flavour hadroproduction”. In: *JHEP* 2007 (2007). doi: [10.1088/1126-6708/2007/09/126](https://doi.org/10.1088/1126-6708/2007/09/126).
- [193] P. Nason. “A New Method for Combining NLO QCD with Shower Monte Carlo Algorithms”. In: *JHEP* 2004 (2004). doi: [10.1088/1126-6708/2004/11/040](https://doi.org/10.1088/1126-6708/2004/11/040).
- [194] S. Höche et al. “A Study of QCD Radiation in VBF Higgs Production with Vincia and Pythia”. In: *SciPost Phys.* 12 (2022). doi: [10.21468/SciPostPhys.12.1.010](https://doi.org/10.21468/SciPostPhys.12.1.010).
- [195] L. A. Harland-Lang et al. “Parton distributions in the LHC era: MMHT 2014 PDFs”. In: *EPJ C* 75 (2015). doi: [10.1140/epjc/s10052-015-3397-6](https://doi.org/10.1140/epjc/s10052-015-3397-6).
- [196] ATLAS Collaboration. “Measurements of top-quark pair differential cross-sections in the jets channel in pp collisions at $\sqrt{s} = 8$ TeV s = 8 TeV using the ATLAS detector”. In: *EPJ C* 76 (2016). doi: [10.1140/epjc/s10052-016-4366-4](https://doi.org/10.1140/epjc/s10052-016-4366-4).
- [197] M. Czakon et al. “Top-pair production at the LHC through NNLO QCD and NLO EW”. In: *JHEP* 2017 (2017). doi: [10.1007/jhep10\(2017\)186](https://doi.org/10.1007/jhep10(2017)186).
- [198] N. Kidonakis. “Single-top transverse-momentum distributions at approximate NNLO”. In: *Phys. Rev. D* (2016). doi: [10.1103/PhysRevD.93.054022](https://doi.org/10.1103/PhysRevD.93.054022).
- [199] *TRExFitter documentation*. 2025. URL: <https://trexfitter-docs.web.cern.ch/trexfitter-docs/>.
- [200] W. Verkerke and D. Kirkby. *The RooFit toolkit for data modeling*. 2003. arXiv: [physics/0306116](https://arxiv.org/abs/physics/0306116) [[physics.data-an](https://arxiv.org/abs/physics/0306116)].
- [201] L. Moneta et al. *The RooStats Project*. 2011. arXiv: [1009.1003](https://arxiv.org/abs/1009.1003) [[physics.data-an](https://arxiv.org/abs/1009.1003)].
- [202] R. Brun and F. Rademakers. “ROOT — An object oriented data analysis framework”. In: *Nucl. Instrum. Methods Phys. Res. A* 389 (1997). New Computing Techniques in Physics Research V. doi: [https://doi.org/10.1016/S0168-9002\(97\)00048-X](https://doi.org/10.1016/S0168-9002(97)00048-X).
- [203] I. Antcheva et al. “ROOT — A C++ framework for petabyte data storage, statistical analysis and visualization”. In: *CPC* 180 (2009). doi: [10.1016/j.cpc.2009.08.005](https://doi.org/10.1016/j.cpc.2009.08.005).
- [204] F. James and M. Roos. “Minuit - a system for function minimization and analysis of the parameter errors and correlations”. In: *CPC* 10 (1975). doi: [https://doi.org/10.1016/0010-4655\(75\)90039-9](https://doi.org/10.1016/0010-4655(75)90039-9).
- [205] F. James. “MINUIT Function Minimization and Error Analysis: Reference Manual Version 94.1”. In: (1994). URL: <https://inspirehep.net/literature/1258343>.
- [206] F. James and M. Winkler. *MINUIT User’s Guide*. 2004. URL: <https://inspirehep.net/literature/1258345>.

-
- [207] I. Brivio and M. Trott. “The Standard Model as an Effective Field Theory”. In: *Phys. Rep.* (2019). doi: [10.1016/j.physrep.2018.11.002](https://doi.org/10.1016/j.physrep.2018.11.002).

List of Figures

2.1	Diagram of the elementary particles of the SM and their physical properties.	4
2.2	One-loop Feynman diagram that contributes to $\mu \rightarrow e\gamma$ in the SM extended with right-handed (Dirac) neutrinos.	8
2.3	The history of cLFV searches in muons.	11
2.4	Current status of the observed upper limits on cLFV τ decays at several experiments.	11
2.5	BDT discriminant distributions for the preliminary $e\mu$ trilepton ATLAS cLFV search.	12
2.6	BDT discriminant distributions for the CMS $e\mu$ dilepton cLFV search.	13
2.7	BDT discriminant distributions for the CMS $e\mu$ trilepton cLFV search.	14
2.8	H_T distributions for the ATLAS $\mu\tau$ trilepton cLFV search.	14
2.9	DNN distributions for the CMS $\mu\tau$ dilepton cLFV search.	15
3.1	The CERN accelerator complex and its experiments.	18
3.2	Cumulative luminosity versus time.	19
3.3	Cut-away view of the ATLAS detector.	20
3.4	Cut-away view of the ATLAS inner detector.	22
3.5	The structure of the ATLAS Inner Tracking Detector	22
3.6	Luminosity-weighted distribution of the mean number of interactions per crossing during 2015 – 2018.	23
3.7	Cut-away view of The ATLAS Calorimeter system.	24
3.8	Cut-away view of the ATLAS muon system.	25
4.1	Schematic view of a hard scattering process.	27
4.2	Parton distribution functions computed at NNLO precision in QCD by the NNPDF group.	28
4.3	Illustration depicting the formation of a parton shower.	31
4.4	The perigee parameters which characterise the track at its point of closest approach to the PV.	34
4.5	Identification efficiencies for various working points utilised in Run 2.	35
4.6	Efficiency of the the Loose, Medium, Tight identification working points used in Run 2.	38
4.7	Illustration of the PV and SV.	40
4.8	Summary of top-quark production cross-section measurements obtained by ATLAS, compared to corresponding theoretical predictions.	42
4.9	LO $t\bar{t}$ production Feynman diagrams.	42
4.10	LO example Feynman diagrams for single-top production.	43
4.11	LO example Feynman diagrams for tW and tZq production.	43

4.12	LO example Feynman diagrams for $t\bar{t}W$ and $t\bar{t}Z$ production.	43
4.13	Decay modes of $t\bar{t} \rightarrow W^+bW^-\bar{b}$	44
4.14	Illustration of the p -value used to reject the background-only hypothesis.	47
4.15	Representation of the p -values used in the computation of p'_μ	49
4.16	p -value of the median test statistic q_μ under the background-only hypothesis ($\mu' = 0$); comparison of upper limits from MC toy experiments with the median expectation and uncertainty bands from the Asimov dataset.	50
5.1	Example Feynman diagrams for the EFT signal processes.	52
5.2	Comparison of the H_T distributions in data and MC predictions for the SR and CR.	60
5.3	Sketch of a decision tree.	64
5.4	Comparison of input variable distributions in data and simulation within the SR used for BDT training.	68
5.5	Correlation of the input variables for the signal and background classes.	70
5.6	Overtraining check for the BDT trained on odd-numbered events.	71
5.7	The ROC curve for the cross-validated BDTs.	72
5.8	Training and validation loss curves for the multi-class XGBoost classifier.	72
5.9	Illustration of the pruning procedure applied to nuisance parameters in the SR.	78
5.10	CR distributions illustrating the impact on the $t\bar{t}$ template of $\pm 1\sigma$ variations of some of the $t\bar{t}$ modelling uncertainties.	79
5.11	CR distributions illustrating the impact on the tW template of $\pm 1\sigma$ variations of some of the tW modelling uncertainties.	80
5.12	Pre-fit and post-fit plots for the fit using Asimov data in SR and CR.	82
5.13	Pre-fit and post-fit plots for the unblinded fit.	84
5.14	Correlation matrix of the nuisance parameters obtained from the full fit.	86
5.15	Nuisance parameter pull plot obtained for the full fit.	87
5.16	Ranking of nuisance parameters NPs from the unblinded fit.	89
5.17	Observed and expected exclusion limits at 95% CL on the WCs probed in this analysis.	92
5.18	ATLAS: Observed and expected exclusion limits at 95% CL on the branching ratio of the $t \rightarrow e\mu q$ decay, and on the corresponding WCs for scalar, vector, and tensor couplings.	93
5.19	CMS: Observed and expected exclusion limits at 95% CL on the branching ratio of the $t \rightarrow e\mu q$ decay, and on the corresponding WCs for scalar, vector, and tensor couplings.	94

List of Tables

2.1	SMEFT operators capable of introducing two quark-two lepton cLFV interactions involving a top quark in the Warsaw basis.	9
2.2	Relevant SMEFT operators for this thesis and the corresponding WCs and Lorentz structures.	9
4.1	Tight and Loose electron definitions.	37
4.2	Tight and Loose muon definitions.	39
5.1	Theoretical cross-sections for single top quark production and $t\bar{t}$ decays through cLFV interactions for vector, scalar and tensor EFT operators. The cross-section for each of the vector operators is the same. These values were calculated for a top quark mass of 172.5 GeV, a top quark decay width of 1.51 GeV, a new physics scale of $\Lambda = 1$ TeV and WC values of 1.0. The scale variation and PDF uncertainties are given.	53
5.2	Summary of the basic generator parameters used to simulate the various processes.	56
5.3	Trigger selection for all data taking years used in this analysis.	57
5.4	Description of the object definition, preselection cuts and region definitions used in the analysis.	59
5.5	Pre-fit event yields in the analysis regions.	61
5.6	The input variables used in the multivariate BDT analysis, listed according to the importance ranking provided by XGBoost.	67
5.7	Pre-fit event yields in the analysis regions for the Asimov fit.	83
5.8	Post-fit event yields in the analysis regions for the Asimov fit.	83
5.9	Pre-fit event yields in the analysis regions for the unblinded fit.	85
5.10	Post-fit event yields in the analysis regions for the unblinded fit.	85
5.11	Observed and expected exclusion limits at 95% CL on each of the WCs probed by the analysis.	92
5.12	ATLAS: Expected and observed 95% CL upper limits on the cLFV WCs, and top quark cLFV branching fractions.	93
5.13	CMS: Expected and observed 95% CL upper limits on the cLFV WCs, and top quark cLFV branching fractions.	94

List of Acronyms

CERN European Organisation for Nuclear Research.

LHC Large Hadron Collider.

BDT boosted decision tree.

BSM beyond the Standard Model.

CL confidence level.

cLFV charged-lepton-flavour violation.

CR control region.

ECAL electromagnetic calorimeter.

EFT effective field theory.

EW electroweak.

FSR final state radiation.

HCAL hadronic calorimeter.

HLT High-Level Trigger.

ID Inner Detector.

ISR initial state radiation.

JER jet energy resolution.

JES jet energy scale.

JVT Jet Vertex Tagger.

LFV lepton flavour violation.

LO	leading order.
LUV	lepton universality violation.
MC	Monte Carlo.
MDT	monitored drift tube.
ME	matrix element.
MS	muon spectrometer.
NF	normalisation factor.
NLL	negative log-likelihood.
NLO	next to leading order.
NNLL	next-to-next-to-leading logarithmic.
NNLO	next to next to leading order.
NP	nuisance parameter.
PDF	parton distribution function.
POI	parameter of interest.
PS	parton shower.
QCD	Quantum Chromodynamics.
QED	Quantum Electrodynamics.
RPC	resistive plate chamber.
SCT	semiconductor tracker.
SM	Standard Model.
SMEFT	Standard Model Effective Field Theory.
SR	signal region.
TRT	transition radiation tracker.
UFO	Universal FEYNRULES Output.
WC	Wilson coefficient.
WP	working point.

Acknowledgements

The research work presented in this thesis would not have been conceived without the assistance and support of several people. I want to acknowledge and express my deepest gratitude:

To my supervisor, Prof. Dr. Markus Cristinziani, for his constant encouragement, reliability in addressing challenges, and invaluable guidance, all of which were essential to the success of this research.

To Dr. Diptaparna Biswas, for his availability, patience, and generosity in sharing his expertise, particularly during the final stages of the analysis.

To Dr. Wolfgang Walkowiak, for his support with all computing-related matters and for kindly providing feedback on this thesis.

To Dr. Vadim Kostyukhin, for his technical support and insightful feedback, especially regarding the BDT and the fit procedure.

To PD Dr. Carmen Diez Pardos, for her valuable feedback on the analysis, particularly during her term as Top+X subgroup convener.

To Dr. Qader Dorosti, for his considerate aid with the masterlab course.

To Dr. Elisabeth Schopf, for her attentive feedback and for generously agreeing to review part of this thesis in detail.

To Stefanie Grebe, for her administrative assistance and the pleasant conversations.

To my fellow first-generation PhD students, Katharina, Beatrice, and Tongbin, whose companionship and support gave me the strength to persevere until the very end.

To Oğul Öncel, whose advice was fundamental during my early, and remote, PhD days.

To my other fellow PhD students, Adam, Nils, Arpan, and many others, for their collaboration, camaraderie, and friendship.

To Prof. Dr. Miriam Watson, Dr. Jacob Kempster, Alexios Stampekis and William George, for all the collaboration in the ATLAS cLFV in top-quark processes analysis team.

To all members of the Experimental Particle and Astroparticle Physics Group, for fostering such a welcoming and delightful work environment.

To my friends in Siegen outside the university, for their constant encouragement.

A Márcia Ribeiro, que me ajudou a manter a calma, a lucidez e a não desistir deste desafio.

E, finalmente, à minha mãe, Lourdes, fonte inesgotável de devoção e apoio incondicional.