Search for flavour-changing neutral-current couplings between the top quark and Higgs boson in multilepton final states with the ATLAS experiment

BERGISCHE
UNIVERSITÄT
WUPPERTAL

Fakultät für Mathematik
und Naturwissenschaften

Dissertation

Marvin Emin Geyik

Februar 2024

Abstract

This dissertation presents a search for flavour-changing neutral-current (FCNC) interactions between the top quark, the Higgs boson and either the up or the charm quark. The search is performed using a dataset corresponding to an integrated luminosity of $140\,\text{fb}^{-1}$, recorded with the ATLAS detector at the Large Hadron Collider at a centre-of-mass energy of 13 TeV. The anomalous FCNC couplings are parametrised using an effective field theory. The interactions are analysed in two distinct modes: the production of a top quark-antiquark pair, with one of the top quarks performing an FCNC decay into a Higgs boson and a different up-type quark, and the associated production of a single top quark and a Higgs boson via the new FCNC vertex. Decay modes of the Higgs boson to two W bosons, two Z bosons or two tau leptons are considered with final states containing either exactly two leptons of the same charge, or three leptons, exactly two of which possess an identical charge. No deviations from SM expectations are observed in the statistical analysis. Therefore, 95% confidence level upper exclusion limits are set on the FCNC top-quark decay branching ratios, amounting to observed (expected) limits of $\mathcal{B}(t \to Hu) < 2.8\,(3.0) \times 10^{-4}$ and $\mathcal{B}(t \to Hc) < 3.3\,(3.8) \times 10^{-4}$.

Contents

Introduction

> *The important thing is not to stop questioning. Curiosity has its own reason for existence.*
>
> - Albert Einstein, 1955

Humanity's drive to understand the world has culminated in some of the greatest achievements in history, remarkably impacting our everyday life. Our pursuit to comprehend the most remote cosmic phenomena led to the development of the theory of general relativity, essential for the high precision of the Global Positioning System (GPS). Moreover, an intricate understanding of atomic interactions has facilitated the design of modern computer processors, integral to almost every aspect of contemporary life.

In the subatomic realm, the *Standard Model of particle physics* (SM) has proven itself as the foundational theory describing all currently known elementary particles and their interactions. Experiments have demonstrated unparalleled agreement with the theory across several orders of magnitude. The *A Toroidal LHC Apparatus* (ATLAS) detector at the Large Hadron Collider (LHC) is among the most significant experiments probing the SM for its properties. A landmark achievement for ATLAS, alongside the *Compact Muon Spectrometer* (CMS) detector, was the observation of the Higgs boson in 2012, the last undetected particle predicted by the SM. All subsequent measurements at both experiments, and multiple others, have shown overwhelming agreement with the SM. Nevertheless, several notable limitations of the SM drive searches for physical phenomena beyond the SM, further fueled by the innate human proclivity for exploration.

This dissertation presents an analysis conducted using data collected with the ATLAS detector between 2015 and 2018, amounting to an integrated luminosity of $140\,\text{fb}^{-1}$. The analysis searches for *flavour-changing neutral-current* (FCNC) interactions between the top quark, the Higgs boson and either the up or the charm quark. Such processes are predicted by the SM. However, they are heavily suppressed, to the extent that they are virtually undetectable at the LHC. Therefore, any evidence of FCNC interactions would indicate the existence of physics beyond the SM.

The presented search is conducted without reliance on a specific theoretical model predicting stronger FCNC interactions. Instead, an approach utilising an effective field

theory (EFT) to model the kinematics of the signal processes is employed. The EFT is considered applicable up to a certain energy scale Λ, beyond which direct effects of the underlying BSM theories might manifest. The analysis considers FCNC couplings in two distinct channels: the decay channel, focusing in the production of top quark-antiquark pairs with one of them decaying via the new FCNC vertex, and the production channel, examining the associated production of a single-top-quark and a Higgs boson through the FCNC vertex. These processes are investigated in final states resulting in exactly two leptons of the same charge or in three leptons, exactly two of which possess an identical charge. The search is conducted separately for tHu and tHc couplings, assuming a vanishing contribution from the respective other coupling.

The dissertation begins by setting the theoretical and experimental context of the analysis. Chapter 1 presents the SM, highlighting its numerous successes and current limitations. Subsequently, Chapter 2 elaborates on the latter by detailing various BSM theories giving rise to enhanced FCNC interactions, and describes how the utilised EFT models effective couplings at LHC energies. Chapter 3 expands on the LHC and the ATLAS detector, and explains the utilised methodology of data collection. Chapters 4 and 5 explain the simulation of proton-proton collisions and the specific EFT and SM processes that are considered in the analysis. Finally, Chapter 6 outlines the methods used to reconstruct physical objects from raw detector data.

The remaining chapters of this dissertation are dedicated to the analysis strategy. Chapter 7 explains the utilised methods of **background estimation**. Chapter 8 introduces an **event selection**, focusing on phase space regions enriched either in signal or certain background processes. Chapter 9 describes the employment of specialised **reconstruction methods** to identify kinematic variables enhancing the separation between signal and background, which are subsequently combined into a single **multivariate discriminant**. **Systematic uncertainties**, expounded upon in Chapter 10, are incorporated into a **maximum-likelihood fit**, the implementation of which is explained in Chapter 11. The final results of this fit, in terms of a **determined signal strength**, are presented in Chapter 12.

1. The Standard Model of Particle Physics

In the context of contemporary particle physics, the SM holds a central position as the foundational theory that underpins the understanding of subatomic particles and their interactions. It represents the culmination of decades of work by several generations of physicists studying the most fundamental properties of the universe.

By the mid-20$^{\text{th}}$ century, several properties of subatomic particles were already known. The *electron*, as a fundamental particle, along with the existence and properties of the *electromagnetic interaction*, had been established for decades. The discovery of the neutron [1] and subsequent developments in proton-neutron interaction models led to the concept of the strong nuclear force [2–4]. Additionally, in 1933 Enrico Fermi proposed the existence of a weak nuclear force to explain observations made in β-decays [5]. This proposition also postulated the existence of *neutrinos*, which interact solely via this newly identified weak force. Following these developments, extensive research into fundamental particles led to the discovery of hundreds of subatomic particles, initially believed to be elementary constituents of the universe. This *particle zoo* was better understood, when Murray Gell-Mann and Yuval Ne'emann classified its constituents according to the *eightfold way* [6] in 1961, which later led to the development of *Quantum Chromodynamics* (QCD). This theory proposed that many of the discovered particles were composite, consisting of *quarks* interacting via the exchange of *gluons*. Shortly thereafter, in 1964, three groups led by Peter Higgs, by Robert Brout and Francois Englert and by Gerald Guralnik, C. R. Hagen, and Tom Kibble explained the fact that some fundamental particles are massive by means of the *Higgs mechanism* (also known as the Englert–Brout–Higgs–Guralnik–Hagen–Kibble mechanism) [7–9]. Finally, in 1968, Abdus Salam and Steven Weinberg, building on work by Sheldon Glashow, unified the weak force and the electromagnetic force, showing them to be manifestations of the same fundamental *electroweak interaction* [10–12]. This research, along with numerous supplementary discoveries and theoretical developments led to the final formulation of the SM, which today forms the basis of research in the field of elementary particle physics.

Today, a multitude of experiments rigorously tested the SM, largely confirming its predictions. Some notable limitations are discussed in Section 1.3. Several fundamental

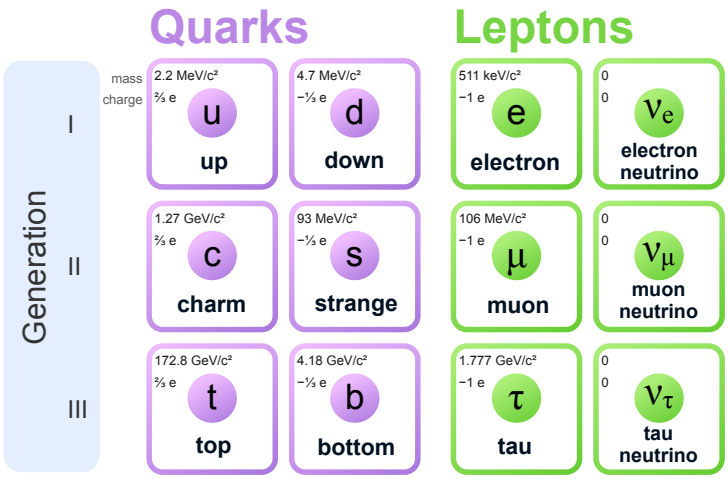

Figure 1.1.: Overview of all quarks and leptons in the SM with their electric charge and mass. The fermions are organised horizontally in their respective families. Masses taken from [13, 14]. Figure based on Ref. [15].

particles have been observed, all of which are *fermions* with a spin quantum number of $1/2$. They are categorised into two groups: The quarks, which interact via the strong and the electroweak force, and the *leptons*, which only interact via the electroweak force. Quarks are organised into up-type quarks with an electric charge of $+2/3\,e$ and down-type quarks with an electric charge of $-1/3\,e$. Additionally, they are grouped into three distinct generations, with each generation comprising one up-type and one down-type quark. The three generations of quarks are arranged in a sequence aligned with their increasing masses. In total, there are six different quarks in the SM: up (u), down (d), charm (c), strange (s), top (t) and bottom (b).

Analogous to quarks, leptons can be grouped into uncharged neutrinos and charged leptons with an electric charge of $-1\,e$, the lightest of which is the electron. In the SM, neutrinos are considered massless, while charged leptons possess a mass. In total, three charged leptons exist: The electron (e), the muon (μ) and the tau lepton (τ). Together with their corresponding neutrino (electron neutrino ν_e, muon neutrino ν_μ and tau neutrino ν_τ), they form three generations, similar to the three quark generations, and are ordered by the mass of the constituting charged lepton. An overview of all SM quarks and leptons and their properties can be seen in Figure 1.1. In addition to the listed fermions, there is an antiparticle associated to every fermion in the SM. These possess an identical mass to their nominal partners, but have inverted charges.

The three interactions described by the SM are mediated by gauge bosons, each with a spin quantum number of 1. The strong interaction is mediated by gluons (g), while the electromagnetic interaction is mediated by photons (γ). Both of these bosons are

Table 1.1.: Overview of all gauge bosons in the SM and the Higgs boson, together with their electric charge and mass. Additional information on the associated interactions are presented as well. The masses are taken from [14].

Name	Charge	Mass	Associated interaction		
			Name	Range	Rel. strength
Gluon (g)	0	0	Strong	$\sim 10^{-15}$ m	1
Photon (γ)	0	0	Electromagnetic	∞	$\alpha \approx \frac{1}{137}$
$W^{\pm}$	$\pm 1\,e$	80.38 GeV	Weak	$\sim 10^{-18}$ m	10^{-5}
Z	0	91.188 GeV			
Higgs (H)	0	125.1 GeV	-	-	-

massless. The weak interaction is mediated by massive $W^{\pm}$ and Z bosons. The SM is completed by the Higgs boson H, a scalar with a spin quantum number of 0. This boson is not a gauge boson, but emerges from the Higgs mechanism, introduced to explain the observed masses of fundamental particles. All SM bosons and their properties are listed in Table 1.1.

This chapter begins by giving an overview of the theoretical formulation of the SM, detailing both the electroweak sector in Section 1.1.1 and the QCD sector in Section 1.1.3. Special focus is placed on the possible ways in which quark flavours can change within the SM. Section 1.2 gives an overview of the most significant experimental confirmations of the SM, while Section 1.3 presents some limitations intrinsic to the theory.

1.1. Theoretical formulation of the Standard Model

The SM is a *quantum field theory* (QFT) [16]. As such, it adheres to the laws of quantum mechanics as well as special relativity and describes particles as excitations of fields. As mentioned earlier, all fundamental particles currently known are fermions with a spin quantum number of 1/2. The dynamics of such fermions are described by the *Dirac Lagrangian* [16]

$$\mathcal{L}_D = \bar{\Psi}(i\gamma^\mu \partial_\mu - m)\Psi. \tag{1.1}$$

Here, Ψ denotes the Dirac spinor describing the fermion field, γ^μ the Dirac matrices, m the fermion's mass, and $\bar{\psi} = \Psi^\dagger \gamma^0$ the Dirac-adjoint of Ψ. In general, greek indices (e.g. μ) run from 0 to 3, while latin indices (e.g. i) run from 1 to 3. Einstein's summation convention is employed throughout this dissertation, implying the summation over indices occurring twice in any given term. It should be noted that, in general, the pre-factors of terms of the order $\bar{\Psi}\Psi$ are interpreted as masses [17].

If one imposes the requirement that the Lagrangian of a given QFT should be invariant under the local transformation $\Psi \to \Psi' = \hat{U}(x)\Psi$, where $\hat{U}$ is an element of the symmetry group G, the Dirac Lagrangian needs to be extended. Group theory allows us to write the transformation as

$$\hat{U}(x) = e^{-ig\alpha_a(x)T^a} \tag{1.2}$$

where g is a real parameter, $\alpha_a(x)$ are location-dependent phases, and T^a are the generators of G [18], satisfying the commutation relation $[T^a, T^b] = if^{abc}T^c$ with the structure constants f^{abc}. Local *gauge invariance* [19] under such transformations can only be established if vector-like *gauge fields* A^a_μ exist in addition to the fermion fields of the free Dirac Lagrangian. These fields transform under the local gauge transformation as

$$A^a_\mu \to A^{a\prime}_\mu = A^a_\mu - \partial_\mu g\alpha_a(x). \tag{1.3}$$

To incorporate them into the Dirac Lagrangian, the derivative ∂_μ is replaced by the *gauge covariant derivative*

$$D_\mu = \partial_\mu - igA^a_\mu T^a. \tag{1.4}$$

This modification ensures local gauge invariance by introducing a coupling between the new gauge field and the fermion field. The parameter g parametrises the strength of this coupling.

In addition to the interaction term, the Lagrangian should include a term describing the free propagation of the gauge field. Defining the *field strength tensor* of the gauge field A^a_μ as

$$F^a_{\mu\nu} = \partial_\mu A^a_\nu - \partial_\nu A^a_\mu + gf^{abc}A^b_\mu A^c_\nu, \tag{1.5}$$

the simplest option for a gauge invariant Lagrangian of the free gauge field is

$$\mathcal{L}_{\text{dyn}} = -\frac{1}{4}F^a_{\mu\nu}F^{a\,\mu\nu}. \tag{1.6}$$

Equations (1.5) and (1.6) lead to an interesting phenomenon in the case of *non-abelian* gauge theory, which is defined by non-commuting generators T^a (i.e. $f^{abc} \neq 0$). For such a theory, the multiplication of the field strength tensor by itself leads to self-couplings of the gauge fields. The SM is an example of such a non-abelian gauge theory.

It should be noted that mass terms for the gauge fields, meaning terms of the order $A^a_\mu A^{a\,\mu}$, are not gauge invariant. Thus, unmodified gauge theories only describe fermions interacting via the exchange of massless bosons. Only after the introduction of the Higgs mechanism (cf. Section 1.1.2) the observed masses of fundamental particles are introduced into the theory.

Scattering Amplitudes and Matrix Elements in QFT

Properties of QFTs are often probed using scattering experiments. Such experiments are characterised by a given initial state $|i\rangle$ and a multitude of potential final states $|f\rangle$. Transitions to any given final state are governed by a transition matrix T constructed from the QFT Lagrangian, yielding the transition amplitude $\langle f|T|i\rangle$. This amplitude factorises into the product of a delta-distribution for four-momentum conservation, and a *matrix element* $\mathcal{M}$, describing particle interactions [16]:

$$\langle f|T|i\rangle = (2\pi)^4 \delta(p_i - p_f) \cdot \mathcal{M}. \tag{1.7}$$

The *cross-section* $\sigma_{i\to f}$, a common measure of the probability for state transitions[1], is proportional to $|\mathcal{M}|^2$. Often, the matrix element consists of multiple terms describing various ways of particle interactions resulting in the same final state. The total matrix element becomes

$$|\mathcal{M}|^2 = |\mathcal{M}_1 + \mathcal{M}_2|^2 \tag{1.8}$$

$$= |\mathcal{M}_1|^2 + |\mathcal{M}_2|^2 + \underbrace{\mathcal{M}_1^* \mathcal{M}_2 + \mathcal{M}_1 \mathcal{M}_2^*}_{\text{interference term}} \tag{1.9}$$

in the case of two contributing processes. In certain cases, the above interference term is zero, allowing for separate calculations of $\mathcal{M}_1$ and $\mathcal{M}_2$.

Feynman Diagrams

The concept of *Feynman diagrams* [20] is an immensely useful technique for visualising processes described by a QFT. Each diagram corresponds to one matrix element contributing to a given process. As explained above, the total amplitude of this process is calculated as the absolute value of the sum of all matrix elements. By utilising Feynman rules, which are derived from the given Lagrangian and general results in QFT, one can derive a complete set of matrix elements that need to be considered in the calculation. Figure 1.2 shows exemplary Feynman diagrams, representing the scattering of two identical fermions $\Psi + \Psi \to \Psi + \Psi$. In Feynman diagrams, the x-axis represents time in the sense that everything on the left-hand side of the diagram corresponds to the initial state, and everything on the right-hand side to the final state of the considered process. Particles are represented by lines, with straight lines corresponding to fermions and modified lines to (gauge) bosons. The arrows on fermion lines indicate whether a particle (arrow pointing in positive time direction) or an antiparticle (arrow pointing in

[1] A cross-section has the dimension of an area. The common unit of cross-sections in particle physics is the barn $1\,\mathrm{b} = 1 \times 10^{-24}\,\mathrm{cm}^2$.

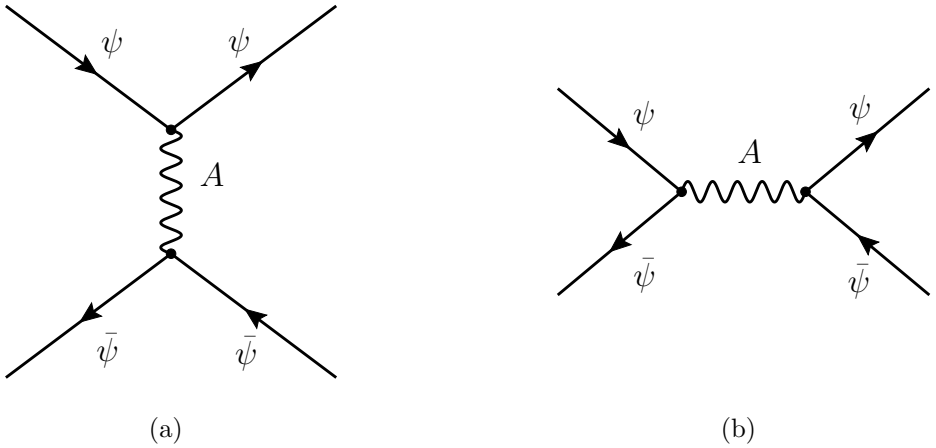

Figure 1.2.: Feynman diagrams depicting the scattering of two identical fermions Ψ (a) by exchanging a gauge boson A or (b) by annihilating into a gauge boson A, which subsequently decays into two identical fermions.

negative time direction) is involved in the process. Interactions between particles are represented by vertices at which multiple lines intersect. At each vertex, momentum must be conserved. According to the Feynman rules, vertices must correspond to couplings present in the Lagrangian. Each vertex in a diagram implies one power of the associated coupling parameter α in the corresponding matrix element.

To fully calculate the amplitude of a given process, all Feynman diagrams contributing to that process must be considered. However, the Feynman rules also allow for loops to enter the diagrams, as can be seen in Figure 1.3. This leads to an infinite number of diagrams to be considered, as loops can be added indefinitely. However, if the coupling strength α of the couplings involved in these loops is small enough, the higher order diagrams decrease in magnitude. *Perturbation theory* explains that in this case only the leading orders of diagrams need to be considered to obtain a good estimate of the total amplitude. The more diagrams are considered, the closer the calculated amplitude will approach its true value.

Regularisation and Renormalisation

The introduction of loop diagrams leads to an issue that needs to be addressed. Generally, the matrix elements corresponding to these diagrams diverge. Such results are irreconcilable with finite physical measurements. *Regularisation* and subsequent *renormalisation* [16] provide methods to absorb the infinitely large terms in the calculations into redefinitions of physical parameters. A detailed description of the formalism of

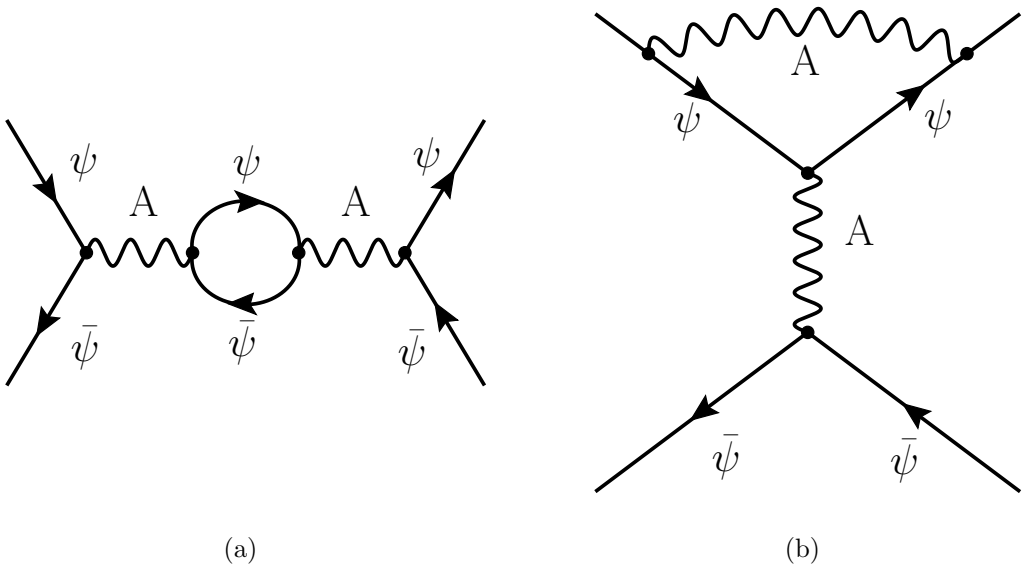

(a) (b)

Figure 1.3.: Example of loops that can be introduced to Feynman diagrams, showing (a) a fermion loop and (b) a loop introduced by a gauge boson.

this theory goes beyond the scope of this dissertation. However, the most important implications are discussed briefly. Regularisation involves introducing a boundary on the matrix-element integration to ensure the total integral is finite. Physical observables now depend on this boundary and on *bare* physical quantities as they occur in the Lagrangian, such as fermion fields, masses and coupling constants. Introducing renormalisation conditions removes the dependence on the artificial boundary by redefining bare quantities into physical fields, masses and coupling strengths, which, unlike the bare quantities, can be measured in experiments. A particularly important result of the renormalisation of the coupling strength g is its dependence on the momentum scale of a given process. To distinguish between the bare and the renormalised coupling strength, the coupling parameter

$$\alpha(Q^2) = \frac{g^2(Q^2)}{4\pi}, \tag{1.10}$$

is defined, where Q^2 represents the momentum scale and $g(Q^2)$ denotes the renormalised coupling strength.

1.1.1. The Electroweak Interaction

The electroweak interaction is described as a gauge theory, following the procedures outlined above. The symmetry group underlying this interaction is

$$G_{\mathrm{EW}} = SU(2) \otimes U(1). \tag{1.11}$$

The $SU(2)$ invariance introduces two new degrees of freedom, interpretable as the third component T_3 of a quantity called *weak isospin* T. One degree of freedom, the *weak hypercharge* Y, is obtained by the $U(1)$ symmetry.

The group $SU(2)$ possesses three distinct generators $T_i = \sigma_i/2$, $i = 1, 2, 3$ with the Pauli matrices σ_i, whereas $U(1)$ only possesses one. Thus, imposing invariance under local G_{EW} transformations yields four different gauge bosons. $W^i_\mu, i = 1, 2, 3$ are associated with the $SU(2)$ symmetry and thus couple to the weak isospin with the coupling strength g. B_μ is the gauge boson arising from local $U(1)$ invariance and couples to the weak hypercharge with the strength g'. The covariant derivative now becomes

$$D_\mu = \partial_\mu - igW^i_\mu \frac{\sigma_i}{2} - ig' B_\mu \frac{Y}{2}. \tag{1.12}$$

To fully describe the electroweak interaction, the concept of *chirality* needs to be considered. Chirality is defined as the eigenvalue of the γ^5 matrix, with $\gamma^5 = \gamma^0\gamma^1\gamma^2\gamma^3$. This matrix has two possible eigenvalues, $+1$ for so-called right-handed states and -1 for left-handed states. By employing the projection operators

$$P_L = \frac{1}{2}(1 - \gamma^5) \text{ and } P_R = \frac{1}{2}(1 + \gamma^5), \tag{1.13}$$

one can divide any given fermion field Ψ into a left-handed component $\Psi_L = P_L\Psi$ and a right-handed component $\Psi_R = P_R\Psi$. Note that these components sum up the original field $\Psi = \Psi_L + \Psi_R$.

Intriguingly, the weak interaction acts differently on the left-handed and the right-handed component of fermion fields. Left-handed fermions are doublets of the weak isospin, while right-handed fermions are singlets. By denoting the third component of the weak isospin as a two-dimensional vector, fermions interacting via the weak interaction are represented as:

$$\begin{pmatrix} U_L \\ D_L \end{pmatrix}, \begin{pmatrix} \ell_L \\ \nu_{\ell\,L} \end{pmatrix}, U_R, D_R, \ell_R, \tag{1.14}$$

where U and D refer to up-type and down-type quarks of a given generation while ℓ and ν_ℓ refer to the lepton and lepton neutrino of a given generation. There are no right-handed neutrinos in the SM.

Considering the above, the electroweak Lagrangian now has the form

$$\begin{aligned}
\mathcal{L}_{\mathrm{EW}} = &-\frac{1}{4}W^i_{\mu\nu}W^{i\,\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu} && \text{(dynamic gauge fields)}\\
&+ i\bar{\Psi}\gamma^\mu\partial_\mu\Psi && \text{(dynamic fermion fields)}\\
&+ (\bar{\Psi}_L,\, \bar{\Psi}'_L)\left(g\gamma^\mu W^i_\mu\frac{\sigma_i}{2} + g'\gamma^\mu B_\mu\frac{Y}{2}\right)\begin{pmatrix}\Psi_L\\ \Psi'_L\end{pmatrix} && \text{(left-handed interaction)}\\
&+ \bar{\Psi}_R g'\gamma^\mu B_\mu\frac{Y}{2}\Psi_R && \text{(right-handed interaction)}\\
&+ \text{h.c.}
\end{aligned}$$

with the fermion fields Ψ and Ψ' and h.c. denoting the hermitian conjugate. The field strength tensors are defined, analogous to Equation (1.5), as

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu \tag{1.15}$$

and

$$W^i_{\mu\nu} = \partial_\mu W^i_\nu - \partial_\nu W^i_\mu + g\epsilon^{ijk}W^j_\mu W^k_\nu \tag{1.16}$$

with the Levi-Cevita tensor ϵ^{ijk}. The electroweak interaction, in this form, lacks several crucial aspects necessary for alignment with experimental observation. Most strikingly, the gauge bosons and fermions possess zero mass, which is in stark contrast to the high masses of the $W^\pm$ and Z bosons, as well as fermions like the top quark. Both of these issues are resolved by the introduction of the Higgs mechanism.

1.1.2. The Higgs Mechanism

The Higgs mechanism [8] begins by postulating a scalar *Higgs field*, which possesses a hypercharge of $Y = +1$ and constitutes an $SU(2)$ doublet:

$$\phi(x) = \begin{pmatrix}\phi^+\\ \phi^0\end{pmatrix} = \frac{1}{\sqrt{2}}\begin{pmatrix}\phi_1 + i\phi_2\\ \phi_3 + i\phi_4\end{pmatrix} \tag{1.17}$$

The field is subject to the *Higgs potential*, taking the form

$$V(\phi) = \mu^2\phi^\dagger\phi + \lambda\,(\phi^\dagger\phi)^2. \tag{1.18}$$

The minimum of the potential is known as the vacuum expectation value (VEV) of the Higgs field. With positive values of μ^2 and λ, the potential has its minimum at $\phi = 0$. However, for a negative value of μ^2, the potential assumes a shape akin to a sombrero, coining the name 'mexican hat potential' often used to describe the shape. An illustration of the Higgs potential for both positive and negative values of μ^2 is shown in Figure 1.4 as a function of the real and imaginary part of a single scalar field.

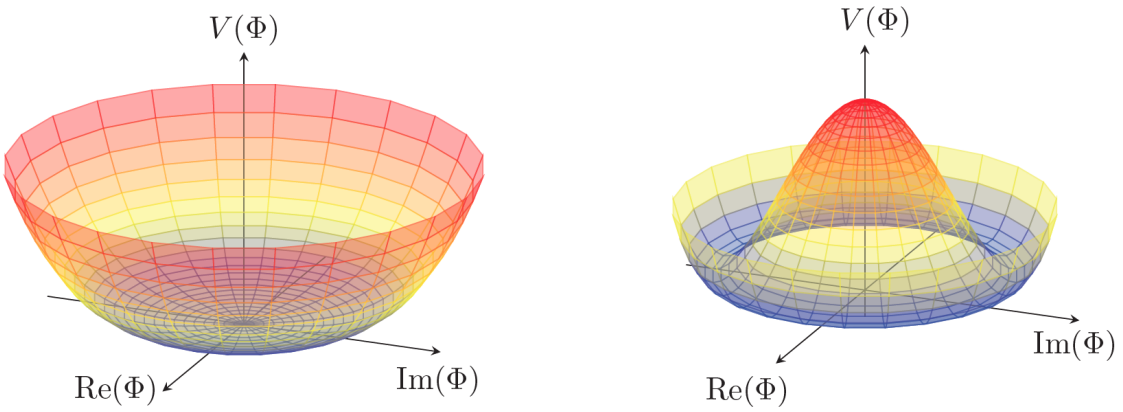

(a) $\mu^2 > 0$ (b) $\mu^2 < 0$

Figure 1.4.: Illustration of the Higgs potential for a one-dimensional scalar field Φ in the two cases (a) $\mu^2 > 0$ and (b) $\mu^2 < 0$ [21, 22].

The mexican hat potential features minima that are not equal to zero, but rather form a circle with radius $v^2 = -\mu^2/\lambda$ around the origin. No point on the circle is preferred over any other. Thus, the VEV is chosen as

$$\phi_0 = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}. \tag{1.19}$$

This non-zero VEV carries fundamental implications for the electroweak Lagrangian. Given that the Higgs field possesses both a non-zero weak hypercharge and a weak isospin, it interacts with the B_μ as well as the W^i_μ bosons. To better illustrate the arising effects, the *unitary gauge* [23] is chosen for the weak gauge bosons.[2] In this gauge, the Higgs field is transformed to

$$\phi(x) \longrightarrow \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix}. \tag{1.20}$$

Here, $H(x)$ is defined as $H(x) = \sqrt{2}\phi^0(x) - v$ and corresponds to the field of a new massive scalar boson, known as the *Higgs boson*. The transformation of the Higgs field results in a mass term with $m_H^2 = 2\lambda v^2$ and several terms of Higgs-boson self interaction.

[2] As the electroweak Lagrangian is invariant under local gauge transformations, the theory can be considered for any value of the local phase. Choosing such a value (usually with the goal of simplifying certain expressions) is known as fixing a gauge.

Considering the Lagrangian of the Higgs field and omitting all terms that include $H(x)$, one obtains [24]:

$$(D^\mu\phi)^\dagger (D_\mu\phi) = \left|\left(\partial_\mu + \frac{i}{2}gW_\mu^i\sigma^i + \frac{i}{2}g'B_\mu\right)\frac{1}{\sqrt{2}}\begin{pmatrix}0\\v\end{pmatrix}\right|^2 \tag{1.21}$$

$$= \frac{v^2}{8}\left|\left(gW_\mu^i\sigma^i + g'B_\mu\right)\frac{1}{\sqrt{2}}\begin{pmatrix}0\\1\end{pmatrix}\right|^2 \tag{1.22}$$

$$= \frac{v^2}{8}\left|\begin{pmatrix}gW_\mu^1 - igW_\mu^2\\ -gW_\mu^3 + g'B_\mu\end{pmatrix}\right|^2 \tag{1.23}$$

$$= \frac{v^2}{8}\left[g^2\left(\left(W_\mu^1\right)^2 + \left(W_\mu^2\right)^2\right)^2 + \left(g\left(W_\mu^3\right)^2 + g'\left(B_\mu\right)^2\right)^2\right]. \tag{1.24}$$

The following four vector bosons can now be defined as mixtures of the W_μ^i and B_μ bosons:

$$W_\mu^\pm \equiv \frac{1}{\sqrt{2}}(W_\mu^1 \mp iW_\mu^2) \qquad \text{with mass } m_W = \frac{gv}{2}, \tag{1.25}$$

$$Z_\mu \equiv \frac{1}{\sqrt{g^2 + g'^2}}(gW_\mu^3 - g'B_\mu) \qquad \text{with mass } m_Z = \frac{v}{2}\sqrt{g^2 + g'^2}, \tag{1.26}$$

$$\text{and } A_\mu \equiv \frac{1}{\sqrt{g^2 + g'^2}}(gW_\mu^3 + g'B_\mu) \qquad \text{with mass } m_A = 0, \tag{1.27}$$

where the masses are obtained by rewriting Equation (1.24) in terms of the newly defined bosons. The photon field A_μ now couples to the electric charge, which is related to the weak isospin and the weak hypercharge through the *Gell-Mann-Nishijima relation* [25, 26]

$$Q = T_3 + \frac{Y}{2}. \tag{1.28}$$

The $W^\pm$ bosons exclusively couple to the weak isospin, whereas the Z boson couples to a mixture of the weak isospin and the weak hypercharge. The weak and electromagnetic interactions are now described in a form that aligns with experimental observations.

Additionally, the Higgs mechanism offers a natural method for introducing mass terms for the fermions. In general, couplings between a Dirac field and a scalar field are referred to as *Yukawa-couplings* [16]. Such couplings can be introduced for the Higgs field as well, linking the left-handed $SU(2)$ doublets with the right-handed singlets as

follows (illustrated here for a quark doublet):

$$\mathcal{L}_{\text{Yukawa}} = f_d(\bar{u}_L,\, \bar{d}_L)\frac{1}{\sqrt{2}}\begin{pmatrix}0\\ v+H\end{pmatrix} d_R + f_d\,\bar{d}_R\frac{1}{\sqrt{2}}(0,\, v+\bar{H})\begin{pmatrix}u_L\\ d_L\end{pmatrix} \tag{1.29}$$

$$= \underbrace{\frac{f_d v}{\sqrt{2}}\bar{d}_L d_R + \frac{f_d v}{\sqrt{2}}\bar{d}_R d_L}_{\frac{f_d v}{\sqrt{2}}\bar{d}d} + \frac{f_d}{\sqrt{2}}\bar{d}_L H d_R + \frac{f_d}{\sqrt{2}}\bar{d}_R \bar{H} d_L. \tag{1.30}$$

Here, f_d refers to the coupling strength between the Higgs field and the down-type quark d. It can be seen that the non-zero VEV in combination with the Yukawa coupling gives rise to a mass term for the down-type quark with $m_d = -f_d v/\sqrt{2}$. Additionally, a coupling between the Higgs boson and the quark is introduced, with a coupling strength proportional to the fermion's mass. The mechanism works analogously for charged leptons.

For up-type quarks, one needs to consider the charge-conjugated version of the Higgs field

$$\tilde{\phi} = \begin{pmatrix}\phi^{0\,*}\\ \phi^-\end{pmatrix} = \frac{1}{\sqrt{2}}\begin{pmatrix}v+H\\ 0\end{pmatrix}. \tag{1.31}$$

Using the same mechanism as above, couplings between the left-handed and the right-handed up-type quark components are introduced, resulting in mass terms for the combined quark field. As there are no right-handed neutrinos in the SM, the Higgs mechanism does not give rise to neutrino masses.

Quark mixing in the Standard Model

As demonstrated earlier, the Yukawa couplings of the Higgs field introduce fermion mass terms into the electroweak Lagrangian. The fermion fields that acquire mass through this mechanism are known as the *mass* or *flavour eigenstates*. For leptons, the mass eigenstates coincide with the eigenstates of the weak interaction. However, as Nicola Cabibbo first introduced in 1963, the quark states participating in the weak interaction are a mixture of the quark mass eigenstates [27]. If one considers the up-type quarks involved in the weak interaction equal to the up-type mass eigenstates, the down-type quark eigenstates of the weak interaction d', s' and b' are linked to the mass eigenstates via the Cabibbo-Kobayashi-Maskawa (CKM) matrix [28]:

$$\begin{pmatrix}d'\\ s'\\ b'\end{pmatrix} = V_{\text{CKM}}\begin{pmatrix}d\\ s\\ b\end{pmatrix} = \begin{pmatrix}V_{ud} & V_{us} & V_{ub}\\ V_{cd} & V_{cs} & V_{cb}\\ V_{td} & V_{ts} & V_{tb}\end{pmatrix}\begin{pmatrix}d\\ s\\ b\end{pmatrix}. \tag{1.32}$$

The entries of the CKM matrix are not independent but can be reduced to four free parameters: a phase δ, introducing a violation of the charge conjugation parity (CP)

symmetry [28] and three Euler angles θ_{12}, θ_{13} and θ_{23}. Expressed as a function of these parameters, the CKM matrix assumes the following form:

$$V_{\text{CKM}} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta} \\ -s_{12}c_{23} - c_{12}s_{23}e^{-i\delta} & c_{12}c_{23} - s_{12}s_{23}s_{13} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{12}e^{-i\delta} & -c_{12}s_{23} - s_{12}c_{23}s_{12} & c_{23}c_{12} \end{pmatrix} \tag{1.33}$$

with $c_{ij} = \cos(\theta_{ij})$ and $s_{ij} = \sin(\theta_{ij})$. Owing to the CKM matrix, interactions involving the $W^{\pm}$ bosons introduce a mixing of quark mass eigenstates into the SM. Flavour mixing in connection with neutral gauge bosons is discussed in Section 1.1.4.

1.1.3. Quantum Chromodynamics

While the electroweak interaction affects all fundamental fermions, QCD is an interaction exclusive to quarks. It is fundamentally based on an $SU(3)$ gauge symmetry. Analogous to the weak isospin, three new degrees of freedom, known as *colour charges*, are introduced. These colour charges are termed red, green and blue. In contrast to the electroweak symmetry, where each fermion flavour possesses a single weak isospin and hypercharge, in QCD, each quark occurs in three distinct states, each possessing one of the three colour charges. When a QFT contains multiple particles, the Lagrangian is constructed as a sum over these particles. Consequently, the QCD Lagrangian for quarks takes the following form:

$$\mathcal{L}_{\text{QCD,quark}} = \sum_{f,c} \bar{\Psi}_{f,c} i\gamma^{\mu} D_{\mu} \Psi_{f,c}, \tag{1.34}$$

where $\Psi_{f,c}$ represents the quark field with flavour f and colour charge c.

The eight generators of $SU(3)$ can be expressed using the unitary *Gell-Mann matrices* $\lambda^a \in \mathbb{C}^{3\times 3}$ as

$$T^a = \frac{\lambda^a}{2}. \tag{1.35}$$

In this representation, the three colour charges correspond to the unit vectors $\vec{e}_1$, $\vec{e}_2$, and $\vec{e}_3$. The covariant derivative in QCD is given by

$$D_{\mu} = \partial_{\mu} - i g_S \frac{\lambda^a}{2} G^a_{\mu} \tag{1.36}$$

with the gauge fields G^a, known as *gluon* fields. Using the gluon field strength tensor

$$G^a_{\mu\nu} = \partial_{\mu} G^a_{\nu} - \partial_{\nu} G^a_{\mu} + g f^{abc} G^b_{\mu} G^c_{\nu}, \tag{1.37}$$

where f^{abc} represents the *fine-structure constants*, the free Lagrangian of the gluon fields is given as

$$\mathcal{L}_{\text{QCD},G} = -\frac{1}{4} G^a_{\mu\nu} G^{\mu\nu\, a}. \tag{1.38}$$

QCD, akin to the theory of the electroweak interaction, is a non-abelian gauge theory, introducing couplings between gluons. This results in numerous Feynman diagrams containing loops that must be considered in the renormalisation of physical parameters. Particularly, the renormalisation of the strong coupling constant g_S yields to interesting observations. As previously described, this introduces a momentum dependence into the coupling. The absolute value of the coupling strength cannot be calculate using the QFT methods. However, if known at a fixed scale μ^2, it can be interpolated to any given scale Q^2 perturbatively. The most dominant terms arising in the calculation are found to be of the order $\ln(Q^2/\mu^2)$, with the leading order being equal to

$$\alpha_{\mathrm{S}}(Q^2) = \frac{g_S^2(Q^2)}{4\pi} = \frac{\alpha_{\mathrm{S}}(\mu^2)}{1 + b_0 \frac{\alpha_{\mathrm{S}}(\mu^2)}{4\pi} \ln\left(\frac{Q^2}{\mu^2}\right)} \text{ with } b_0 = \frac{1}{3}(11N_C - 2N_f), \tag{1.39}$$

where N_C denotes the number of colour charges and N_f the number of quark flavours with $m_f^2 < Q^2$. Given that there are only six quark flavours in the SM, $\alpha_{\mathrm{S}}(Q^2)$ decreases with increasing Q^2. At low Q^2, the above approximation is no longer valid. This is characterised by the infrared cutoff scale Λ^2, defined by the vanishing of the denominator in Equation (1.39), i.e.

$$b_0 \alpha_{\mathrm{S}}^2(\mu^2) \ln\left(\frac{\Lambda^2}{\mu^2}\right) = -1. \tag{1.40}$$

Using this definition, for large values of Q^2 one can determine $\alpha_{\mathrm{S}}(Q^2)$ as

$$\alpha_{\mathrm{S}}(Q^2) = \left[\frac{1}{12\pi}(11N_C - 2N_f) \ln\left(\frac{Q^2}{\Lambda^2}\right)\right]^{-1}. \tag{1.41}$$

Figure 1.5 shows α_{S} as a function of Q^2.

The decrease in α_{S} with increasing Q^2 results in two phenomena unique to QCD. At large momenta, the coupling strength approaches zero, a phenomenon known as *asymptotic freedom*. In this regime, perturbative calculations are possible for QCD. Conversely, the increasing coupling strength for small momenta leads to the so-called *confinement*. If one were to separate two quarks from each other, the high coupling strength would result in an ever-increasing potential between the two particles. At some point the potential is high enough to produce a new quark-antiquark pair. Because of this, quarks are never observed individually in nature. Only bound, colourless states of quarks called *hadrons* can exist. A *meson* is a combination of a quark-antiquark pair with opposing colour charges, while *baryon* comprises three quarks with different colour charges.

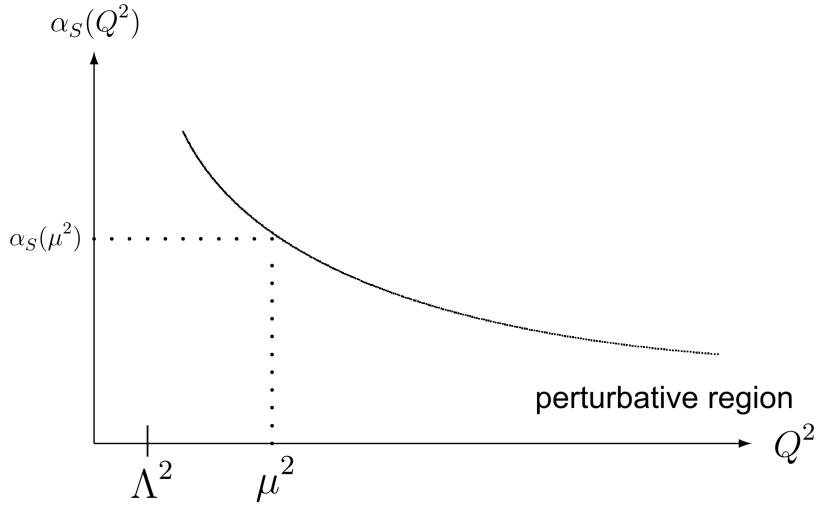

Figure 1.5.: The dependence of α_S on Q^2, with the infrared cutoff scale Λ and an arbitrary scale $\mu^2 > \Lambda^2$. Figure adapted from [23].

1.1.4. FCNC processes in the Standard Model

As demonstrated in Section 1.1.1, only quarks of different flavours can interact with the $W^\pm$ bosons. These interactions are also referred to as charged-current interactions[3], while those involving any of the neutral bosons are known as neutral-current interactions. In the SM, neutral-current interactions that change a quark's flavour, known as FCNC interactions, do not occur at leading order and are highly suppressed in higher orders in perturbation theory. The reasons for their absence vary depending on the boson potentially involved in the interaction. This section presents these reasons for each of the neutral bosons in the SM, as well as the mechanism responsible for the suppression of FCNC interactions at higher orders.

FCNC interactions involving the Higgs boson

In the electroweak interactions, the Yukawa couplings between fermions and the Higgs boson are added in an ad-hoc way, because they are gauge invariant and thus need to be considered in the Lagrangian. Notably, FCNC terms like

$$f_{ut}(\bar{u}_L,\, \bar{d}_L)\frac{1}{\sqrt{2}}\begin{pmatrix} v + H^* \\ 0 \end{pmatrix} t_R \tag{1.42}$$

[3]The *fermion current* $\bar{\Psi}\gamma^\mu\Psi$ is a conserved quantity in the SM (and generally in QFTs). Interactions can also be understood as a flow of this current, where interactions that change its charge are referred to as charged-current interactions. Analogously, all other interactions are called neutral-current interactions.

are equally gauge invariant and could be added to the Lagrangian, leading to FCNC couplings of the up quark, the top quark and the Higgs boson. However, including such terms would imply that u and t are no longer mass eigenstates. A simple rotation in flavour space would produce new mass eigenstates u_{alt} and t_{alt}, for which the original couplings to the Higgs field would hold true. Therefore, Leading-Order (LO) FCNC couplings involving the Higgs boson are absent in the SM by construction.

FCNC interactions involving the Z boson and the photon

While the Higgs boson inherently only couples to the mass eigenstates of quarks, the Z boson and the photon couple to a mixture of mass eigenstates, as determined by the CKM matrix. Exemplarily, a coupling of the form $Z_\mu \bar{d}' \gamma^\mu d'$ is considered. d' can be translated to the mass eigenstates via the corresponding CKM matrix elements, yielding

$$Z_\mu \bar{d}' \gamma^\mu d' = Z_\mu (V_{ud}^* \bar{d} + V_{us}^* \bar{s} + V_{ub}^* \bar{b}) \cdot \gamma^\mu \cdot (V_{ud} d + V_{us} s + V_{ub} b). \tag{1.43}$$

The multiplication of these terms introduces FCNC couplings among each pair of down-type quarks, for example, $Z_\mu V_{ud}^* V_{us} \bar{d} \gamma^\mu s$. However, one must also consider the two couplings $Z_\mu \bar{s}' \gamma^\mu s'$ and $Z_\mu \bar{b}' \gamma^\mu b'$. If the mass eigenstates are again substituted into these terms, two additional couplings of the form $\bar{d} \gamma^\mu s$ arise, with different elements of V_{CKM} as pre-factors. In total one obtains

$$\mathcal{L}_{Z\bar{d}s} = \underbrace{(V_{ud}^* V_{us} + V_{cd}^* V_{cs} + V_{td}^* V_{ts})}_{=0 \text{ (unitarity of } V_{\text{CKM}})} Z_\mu \bar{s} \gamma^\mu d. \tag{1.44}$$

Similarly, the unitarity of the CKM matrix ensures the cancellation of all other FCNC terms in the final Lagrangian.

An ad-hoc addition of couplings such as $Z_\mu \bar{d}' \gamma^\mu s'$ would constitute a redefinition of the weak interaction's eigenstates, akin to how cross-generational Yukawa couplings necessitate a redefinition of the mass eigenstates. Therefore, it is evident that no LO FCNC couplings can be introduced in the electroweak sector of the SM.

FCNC interactions involving gluons

Gluons act exclusively within $SU(3)$ colour space. Although the QCD Lagrangian contains a sum over all quark flavours, the specific flavour states are inconsequential for QCD processes. Only the number of distinct quark fields participating in loop diagrams is of importance. This can be interpreted as an invariance of QCD under rotations in the six-dimensional quark flavour space. A quark assumes a specific flavour only when it interacts through the electroweak interaction. Therefore, the absence of FCNC interactions involving gluons in the SM directly results from the disentanglement of $SU(3)$ colour space and $SU(2) \times U(1)$ flavour space.

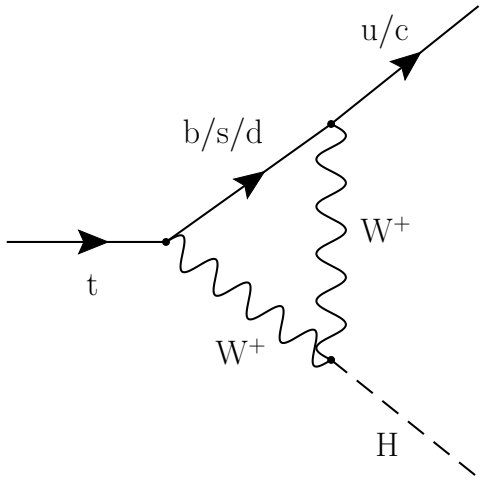

Figure 1.6.: Decay of a top quark via $t \to Hu/c$ as it occurs in the SM.

The GIM mechanism

Neither the electroweak sector nor QCD permit LO FCNC couplings. However, FCNC interactions remain feasible at higher orders through loop diagrams involving W bosons. Figure 1.6 shows a Feynman diagram of a top quark decaying into a Higgs boson and a charm quark. Processes like this one are allowed, but nevertheless suppressed in the SM due to the *Glashow-Iliopoulos-Maiani* (GIM) mechanism [29]. The mechanism is predicated on the fact that all down-type quarks contribute to the loop in the diagram. The three resulting matrix elements are identical, except for the mass of the down-type and the CKM matrix entries involved in the calculation. If all down-type quark masses were equal, the diagrams would cancel out completely, owing to the unitarity of the CKM matrix. Due to their different masses, the total amplitude of the process is approximately proportional to $m_b^2/m_W^2 \approx 2.9 \times 10^{-3}$ [30]. This causes a net suppression of top-quark related FCNC processes at higher orders.

1.2. Experimental Evidence for the Standard Model

Currently, the SM stands as one of the most rigorously tested theories in the realm of physics. Numerous experiments have studied the interactions of elementary particles across a spectrum of high and low energies, and found overwhelming support for the predictions of the SM. This section provides an overview of the most striking experimental results that have been produced in support of the SM.

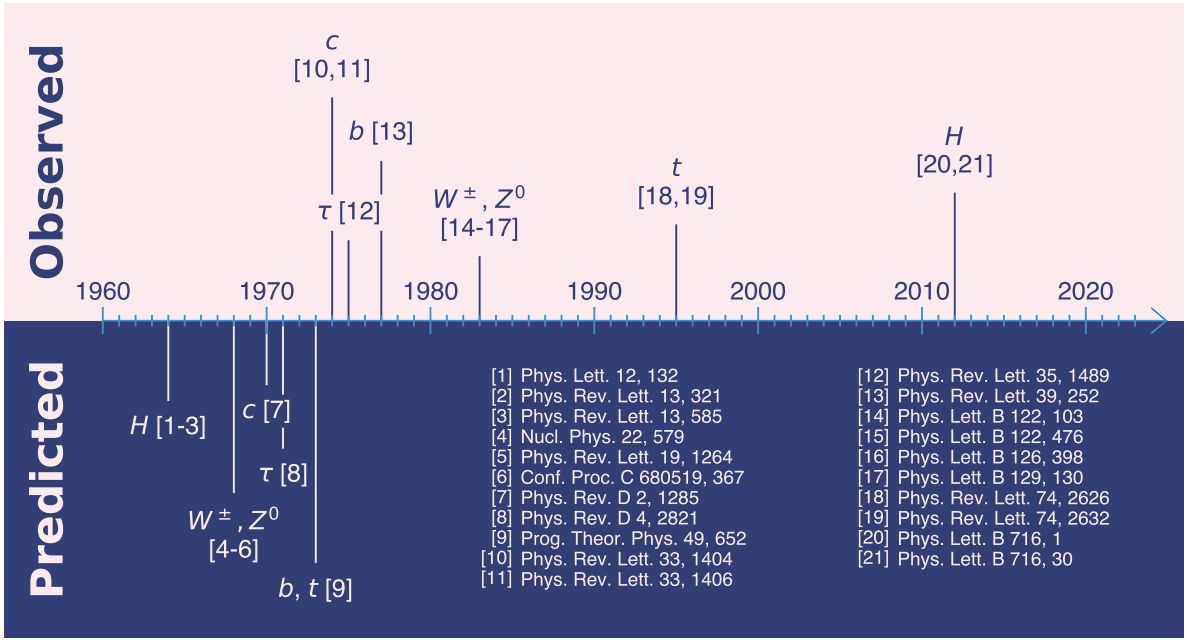

Figure 1.7.: An overview of all elementary particles that have been predicted and observed after 1960. Entries above the timeline show the experimental observations, while the lower entries show the predictions of individual particles.

Observation of Elementary Particles

As noted at the outset of this chapter, the discovery of the Higgs boson in 2012 [31, 32] marked the observation of the last undetected particle predicted by the SM. In the decades preceding this discovery, however, numerous other elementary particles were predicted by SM and subsequently detected in experiments. Figure 1.7 presents an overview of the predictions and observations of all elementary particles discovered after 1960. The fact that all of these particles were predicted prior to their experimental observation lends significant credence to the SM. The broad acceptance of a theory hinges not only on its capability to describe already observed phenomena but also on its ability to accurately predict the outcomes of new experiments. Figure 1.7 confirms that the SM has passed the latter test multiple times.

Limits on the number of fermions

Although all particles of the SM have been observed, it remains plausible that the model is incomplete, potentially omitting additional elementary particles. Experiments effectively rule out the simplest of these additions: The inclusion of further generations of fermions. Experiments conducted on the Large Electron-Positron Collider (LEP),

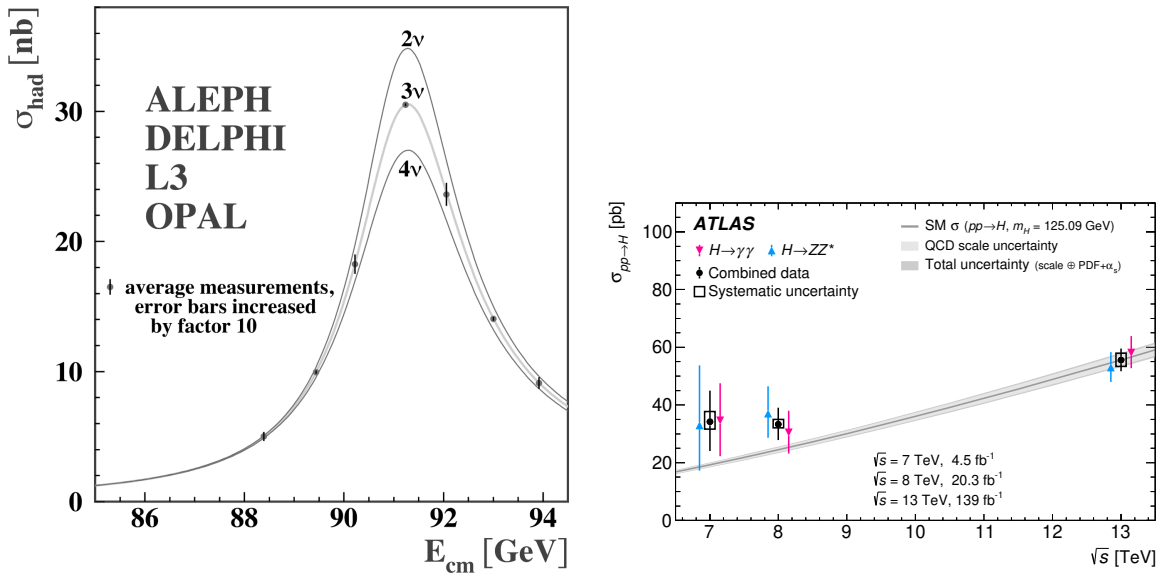

(a) (b)

Figure 1.8.: Plots depicting (a) the combination of all measurements of the Z-boson resonance performed at the four LEP experiments ALEPH, DELPHI, L3 and OPAL [33] and (b) a summary of Higgs-boson production cross-sections measured by the ATLAS collaboration at the LHC [34].

operational from 1989 to 2000, conclusively ruled out the existence of more than three neutrinos. Precision measurements of the Z-boson resonance exhibit highly sensitivity to the number of light neutrinos. This is attributed to the fact that additional neutrinos would reduce the relative contribution of visible final states to the totality of all Z-boson decays. A summary of the measurements from the four LEP experiments ALEPH, DELPHI, L3 and OPAL, depicted in Figure 1.8 (a), clearly confirms the hypothesis of exactly three light neutrinos existing.

Although the possibility of additional light neutrino-like particles has been ruled out by LEP experiments, for a long time the potential inclusion of new generations of heavy fermions in the SM remained possible. This has been ruled out by precision measurements of Higgs boson properties. As described in Section 1.1.2, the mass of a fermion is directly proportional to its coupling strength to the Higgs boson. Thus, heavy fermions have the largest contribution to loop-corrections in the calculation of Higgs-boson related observables, such as cross-sections. However, measurements of Higgs-boson cross-sections exhibit pronounced agreement with SM predictions, as illustrated in Figure 1.8 (b). This confirms non-existence of any additional fermion generations that are identical to previous generations in all respects but their mass. It should be noted that the existence of heavy fermions remains a possibility if they do not acquire their mass via the Higgs mechanism.

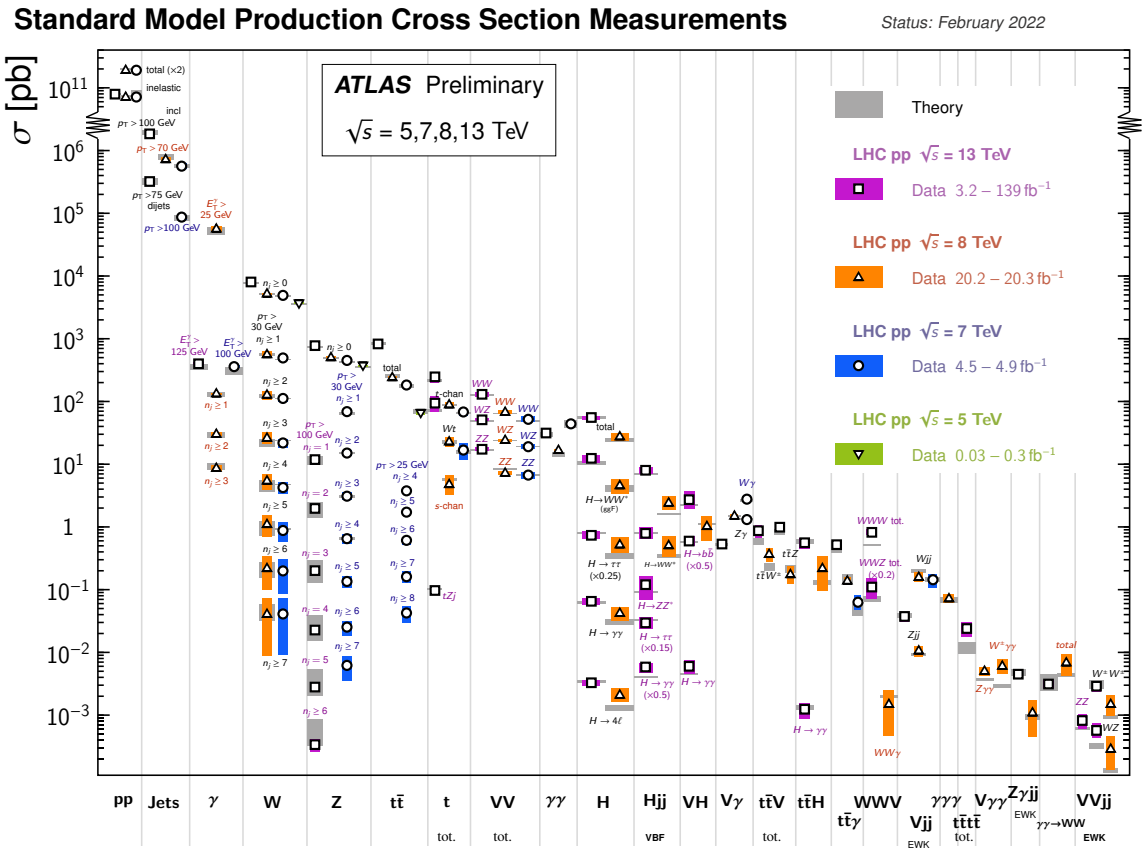

Figure 1.9.: Summary of all SM cross-sections measured by the ATLAS collaboration as of February 2022. The grey areas correspond to SM predictions, while the coloured areas represent the measurement uncertainties for various center-of-mass energies [35].

Cross-section measurements

The SM not only accurately predicts the number of particles in the SM, but also their properties. At the LHC, a vast array of processes predicted by the SM has been observed and measured to remarkable precision. A summary of all cross-section measurements published by the ATLAS collaboration until February 2022 can be found in Figure 1.9. There is an overwhelming agreement between theoretical predictions and observations. It is noteworthy that the measured cross-sections span a wide range, covering many orders of magnitude.

The electroweak fit

Finally, the SM not only predicts individual processes with accuracy but also demonstrates significant consistency among various measurements. This consistency is shown in the global fit of SM parameters to precision measurements of electroweak parameters. Such fits have successfully constrained observables, including the Higgs-boson mass, even before their direct observation [36]. With all free parameters of the SM now measured, the electroweak fit can be used to compare individual parameter measurements to the constraints set by other SM measurements. The latest electroweak fit was published in 2018 [37]. Figure 1.10 shows a comparison between the individual measurements and the fit results. A formidable agreement between the fit results and individual measurements can be observed, reaffirming the SM's ability to produce consistent and precise predictions of physical observables.

1.3. Limitations of the Standard Model

Despite its numerous successes, there are several notable observations which cannot be explained by the SM in its current form. A selection of those is presented in the following.

Neutrino Oscillations

Owing to the absence of right-handed neutrinos in the SM, they cannot acquire mass through Yukawa couplings to the Higgs field. Differences in the weak and mass eigenstates thus cannot exist in the lepton sector of the SM, because there are no mass eigenstates for neutrinos. However, the observation of neutrino oscillation [38] contradicts this. Studies on the rate of solar neutrinos of different flavours revealed that neutrinos periodically change their flavour when propagating through space. This phenomenon can also be understood as cross-generational couplings between neutrinos and charged leptons, analogous to up-type-down-type couplings in the quark-sector of the SM. To explain the observed couplings, neutrinos must have three distinct mass eigenstates. Analogously to the CKM matrix in the quark sector, one then has to introduce the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix [39, 40] translating between the mass eigenstates and the weak eigenstates of neutrinos. In its current form, the SM does not accommodate this, rendering it incapable of explaining neutrino oscillations.

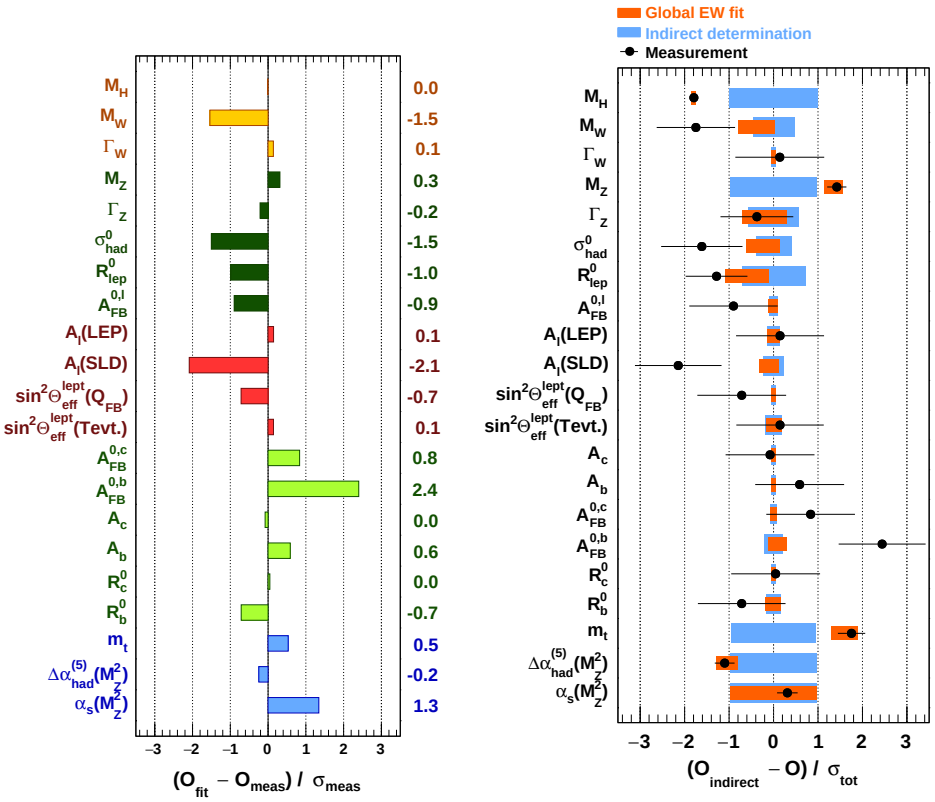

(a) (b)

Figure 1.10.: Results from the global electroweak fit. (a) Deviations of the fit results from the most precise measurements of individual parameters in units of experimental uncertainty. The fit results were obtained using all parameter measurements as input. (b) The most precise measurements of individual parameters, together with the results from the fit using all measurements as input (blue) and alternative fits, in which the respective parameter was not used as input (orange). The values are given in units of the total uncertainty [37].

The Strong CP Problem

Generally, all gauge invariant terms should be incorporated into the Lagrangian of a given QFT, even if added ad-hoc. This has been done for the Higgs-field Yukawa couplings, for example. In the context of QCD, the following CP-violating term can be added [41]:

$$\mathcal{L}_{\text{CP-viol.}} = \theta \frac{g_S^2}{4(4\pi)^2} G^a_{\mu\nu} \epsilon^{\mu\nu\sigma\rho} G^{a\,\sigma\rho}, \tag{1.45}$$

where θ is a free parameter, g_S the QCD coupling strength and G^a the gluon fields. Such a term is not relevant for the electroweak field tensors, as it can be absorbed by the CP-violating phase of the CKM matrix. In QCD, which lacks a CKM matrix, this term must be considered. However, experimental observations show no sign of CP violation in QCD processes, leading to very strong constraints on θ. There is no intrinsic reason for this particular SM parameter to possess such a low value. This apparent tension is known as the strong CP problem. Some theories propose an underlying symmetry enforcing a θ of zero. One potential method to introduce such a symmetry is through the existence of axion-like particles [42]. As of the writing of this dissertation, however, no evidence of such particles has been observed.

Gravity and the Hierarchy Problem

The SM describes three out of the four fundamental forces governing the universe. The obvious exception is gravity. The most accurate description of gravity is Einstein's theory of general relativity [43], which explains gravitational interactions in terms of the bending of space-time. Combining the SM and general relativity into a single theory poses significant challenges. Most attempts at this integration predict additional particles or interactions, none of which have yet been observed. In addition to fundamentally different approaches to describe these two sets of interactions, gravity is also phenomenologically very different from the interactions described by the SM. While all other interactions can be attractive and repulsive, gravity appears to be exclusively attractive. Furthermore, the magnitude of gravitational interactions is 24 orders of magnitude smaller than the weakest of the SM interactions. This disparity is referred to as the hierarchy problem, and as of the writing of this dissertation, there is no consensus on its origin or resolution.

The Λ-CDM Model

The SM focuses on descriptions of the interactions of fundamental interactions. At high energies and low interaction-lengths it offers predictions with unparalleled accuracy. At

cosmological scales, however, another model is used to describe the evolution of the universe: The Λ-CDM model [44]. Two assumptions coin its name. Firstly, the concept of dark energy, manifesting as a constant energy density represented by the cosmological constant Λ and responsible for the accelerated expansion of the universe. Secondly, the existence of cold dark matter (CDM), which provides explanations for cosmological observations such as higher than expected rotation speeds of galaxies. The Λ-CDM model is able to accurately describe the evolution of the universe, the observed expansion speed as well as the anisotropy spectrum of the cosmic microwave background [45]. However, the SM does not appear to be compatible with this model. Vacuum energy densities predicted by the SM, which could be the source of dark energy, are 120 orders of magnitude above the measured values [46]. Moreover, the observed matter particles of the SM are no valid candidates for CDM, and searches for additional particles have thus far been fruitless.

2. FCNC Interactions in Theories Beyond the Standard Model

This dissertation focuses on a search for FCNC couplings between the top quark and the Higgs boson. As discussed in Section 1.1.4, such interactions are highly suppressed in the SM, to the point that experimental observation is currently impossible. Nevertheless, several notable extensions of the SM predict higher contributions from FCNC interactions, to the extent that they can be constrained by current experiments. Section 2.1 presents an overview of some of the most important theoretical models beyond the SM (BSM) giving rise to FCNC interactions involving the Higgs boson. To conduct model-independent searches, *effective field theories* (EFTs) are employed. Such theories, focusing on the context of top-quark Higgs-boson FCNC couplings, are detailed in Section 2.2. Numerous theories have already been constrained by experimental findings. An overview of the latest results of FCNC searches is given in Section 2.3.

2.1. Theoretical Models of Enhanced tHq FCNC Interactions

A wide range of theoretical models give rise to FCNC interactions involving the Higgs boson. Moreover, there exist many different variations for each of these models, differing in particular details. Providing an exhaustive overview of all these models would exceed the scope of this dissertation. Instead, the dissertation presents categories of models and the manner in which each gives rise to Higgs-boson FCNC interactions.

2.1.1. Two Higgs Doublet Models

One of the most influential classes of models is made up of *Two Higgs Doublet Models* (2HDMs) [47]. They represent one of the simplest and least invasive ways to extend the SM in its current form. As the name suggests, these models assume the existence of

two $SU(2)$ doublet Higgs fields denoted Φ_1 and Φ_2. Since both fields have two complex components, there are a total of eight free parameters in the Higgs sector of 2HDMs.

Φ_1 and Φ_2 are subject to a combined mexican-hat potential. By the same logic as for the SM Higgs mechanism, two non-zero VEVs v_1 and v_2 are obtained. Three of the eight free parameters are absorbed to give mass to the $W^\pm$ and Z bosons. The remaining five manifest in the form of multiple Higgs bosons. There are two scalar Higgs bosons $H_{1,2}$, one pseudo-scalar Higgs boson H_A and two charged Higgs bosons $H^\pm$.

There exist both *flavour-violating* and *flavour-conserving* versions of 2HDMs [47]. They primarily differ in the allowed couplings of fermions to the two Higgs fields. In flavour-violating 2HDMs, fermions couple to both Higgs fields. The resulting Yukawa couplings are then of the form

$$\mathcal{L}_{\text{Yukawa}} = f^1_{ij}\bar{\Psi}^{\text{L}}_i\Phi_1\psi^R_j + f^2_{ij}\bar{\Psi}^L_i\Phi_2\Psi^{\text{R}}_j + \text{h.c.}, \tag{2.1}$$

where Ψ^{L}_i and Ψ^{R}_j denote a left-handed $SU(2)$ doublet and a right-handed $SU(2)$ singlet, and f^k_{ij} their combined coupling to the k'th Higgs field. Generally, f^1_{ij} and f^2_{ij} cannot be diagonalised simultaneously. This means that tree-level Higgs-boson FCNC couplings are part of the 2HDM Lagrangian, even if fermions are considered in their mass eigenstates.

This observation presents challenges for the development of 2HDMs, as the predicted FCNC coupling strengths are often excluded by experiments. Flavour-conserving 2HDMs address this issue by stipulating that each fermion field couples exclusively to one of the two Higgs fields [48]. However, FCNC interactions in such models are still higher than in the SM, because the new Higgs bosons contribute to loop diagrams. Specifically, the charged Higgs bosons $H^\pm$ introduce new flavour-changing charged current interactions and result in loop diagrams, exemplified by the one in Figure 2.1. Such loops are not suppressed by the GIM mechanism, because the matrix of all $H^\pm$ coupling parameters is not required to be unitary.

2.1.2. Minimal Supersymmetric Standard Models

Supersymmetry (SUSY) [49] is a concept in theoretical physics that assumes a global symmetry between fermions and bosons. This manifests as each particle having a SUSY partner with identical mass and quantum numbers. They only differ in the fact that one SUSY partner must be fermionic and the other bosonic. These unmodified SUSY models are effectively excluded by experimental findings, as they predict the existence of numerous new particles in low mass ranges. However, by introducing spontaneous symmetry breaking into the theory, it becomes feasible for SUSY partners to have different masses. These modified SUSY models offer natural solutions to several of the

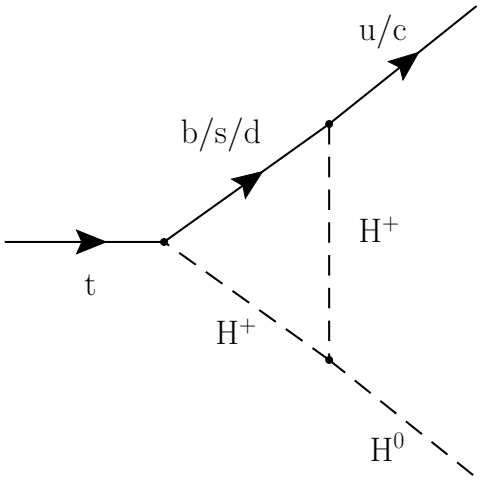

Figure 2.1.: Feynman diagram of the process $t \to Hc$ with a loop containing charged Higgs bosons from a 2HDM.

SM's current limitations. The lightest of the predicted SUSY particles could be a viable candidate for CDM, and SUSY presents a natural explanation of the hierarchy problem.

Minimal Supersymmetric Standard Models (MSSMs) [50] are a class of SUSY extensions of the SM that only add the minimum number of particles to make the SM supersymmetric. Every fermionic particle is assigned a bosonic *sparticle*. Likewise, every boson obtains a fermionic SUSY partner. All MSSMs necessitate two Higgs doublet fields, thereby introducing FCNC interactions in the same manner as discussed in the previous section. In addition to this, the new SUSY particles contribute to FCNC loop diagrams. Specifically, SUSY QCD loops involving squarks $\tilde{q}$ (quark SUSY partners) and gluinos $\tilde{g}$ (gluon SUSY partners) considerably amplify the FCNC contributions. Key SUSY QCD diagrams contributing to top-quark Higgs-boson FCNC interactions are depicted in Figure 2.2.

2.1.3. R-parity violating Supersymmetric Models

To comprehend R-parity violating MSSMs, it is essential to first consider the *baryon number*, defined as

$$B = 1/3(n_q - n_{\bar{q}}). \tag{2.2}$$

Here n_q denotes the number of quarks and $n_{\bar{q}}$ the number of antiquarks participating in a process. This number is conserved in the SM. However, no fundamental symmetry

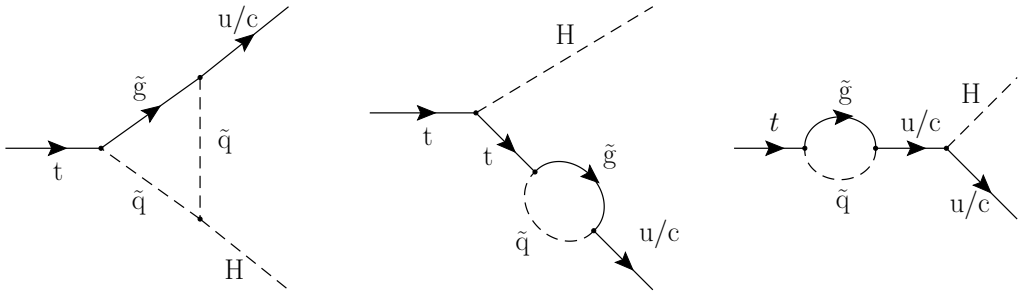

Figure 2.2.: The most important SUSY QCD loop diagrams contributing to tHq FCNC interactions in MSSMs. $\tilde{q}$ denotes involved squarks and $\tilde{g}$ gluinos. Figures based on Ref. [51].

mandates its conservation. Instead, all interactions predicted by the SM happen to conserve the number. The SUSY particles predicted by MSSMs in many cases introduce vertices that no longer conserve B. This presents a problem, as experimental observations clearly indicate the conservation of B in nature. Various models mitigate this tension by requiring the conservation of *R-parity* [49]

$$P_R = (-1)^{3B+L-2s}, \tag{2.3}$$

with the baryon number B, the lepton number $L = n_\ell - n_{\bar{\ell}}$ and the spin s.

A class of MSSMs called *R-parity violating MSSMs* [52] still manage to retain baryon number conservation without the explicit requirement of R-parity conservation. In these models, the possible R-parity violating interactions manifest as additional contributions to SUSY QCD loops. In some cases, this increases FCNC contributions further compared to bare MSSMs.

2.1.4. Randall-Sundrum Models

A distinctly different class of BSM models posits the existence of extra dimensions. Such extra dimensions are often required for the unification of the SM with a description of gravity. In most cases, these extra dimensions are limited in size due to internal curvature, sometimes referred to as warping. *Randall-Sundrum* (RS) [53] models make up a specific class of warped extra dimension models with exactly one new dimension in addition to the four dimensions of space-time. The extra dimension has two boundaries, called branes. In most RS models, our universe is situated on the *TeV-brane* at one end of the extra dimension, with the opposite being called the *Planck-brane*. It is assumed

in these models that gravity acts within the extra dimension and its strength decreases exponentially from the Planck-brane to the TeV-brane. This would solve the hierarchy problem discussed in Section 1.3.

Certain classes of RS models assume that fermion fields also have components in the extra dimension, usually to explain the fermion mass hierarchy. Due to the finite size of this dimension, the fermion wave function in this space must be a standing wave. The mode of this standing wave becomes a new degree of freedom. Fermion states with non-zero modes are called *Kaluza-Klein states* [54]. These Kaluza-Klein states contribute new mass terms to the Lagrangian, which must be accounted for alongside the mass terms arising from the Higgs mechanism. Consequently, the mass eigenstates are no longer exclusively determined by Yukawa interactions with the Higgs field, and cross-generational couplings cannot be diagonalised away. As a result, tree-level Higgs-boson FCNC interactions arise naturally within these models [55].

2.2. tHq FCNC Interactions in Effective Field Theories

As the previous section suggests, model-specific searches for FCNC interactions would require an immense amount of effort due to the multitude of models to analyse. EFTs provide a way to perform model-independent searches for phenomena, based on the assumption that new physics emerging at higher energy exerts an effective impact on experiments at lower energies. This can best be illustrated by the example of Fermi's description of the weak interaction [56].

When Fermi formulated his description, he aimed to describe observed β decays of the sort $n \to p + e^- + \bar{\nu}_e$. He assumed a four-fermion coupling between all four particles involved in the interaction, which can be seen in Figure 2.3 (a). The operator responsible for this coupling is of dimension 6. The coupling strength is given by the Fermi constant $G_\mathrm{F} = 1.1663787(6) \times 10^{-5}\,\mathrm{GeV}^{-2}$, which has to be of dimension -2 to ensure that the full Lagrangian is of dimension 4. This theory accurately described the observed interactions at the energy scales available during that period. However, the predicted cross-sections for interactions involving the four-fermion operator scale with $G_\mathrm{F}^2 E^2$, causing a divergence at high energies. This is a direct consequence of the dimension-6 operator and the dimensional coupling constant. Such a divergence was deemed to be unphysical and indicated that the theory would break down at some energy scale. This scale later turned out to be the mass-scale of the weak gauge bosons. In the SM, the β decay is described by the exchange of a W boson which couples to two fermions with a dimensionless coupling constant. This is shown in Figure 2.3 (b). The Fermi constant is found to be equal to

$$G_\mathrm{F} = \frac{\sqrt{2}}{8}\frac{g^2}{m_W^2}, \tag{2.4}$$

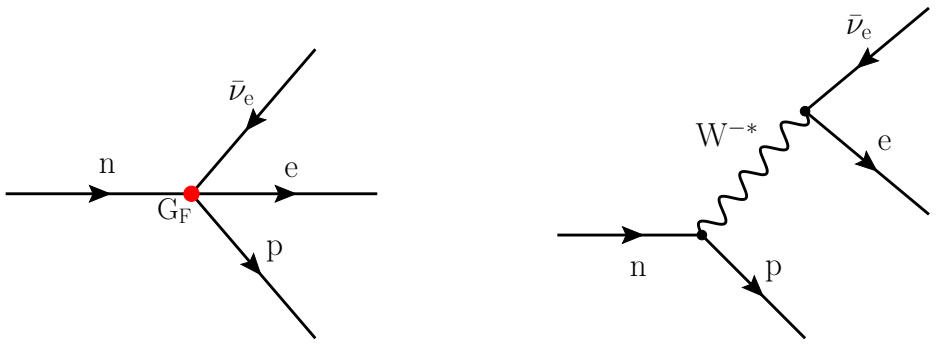

(a) (b)

Figure 2.3.: The β^- decay as described (a) by Fermi's interaction using the four-fermion coupling and (b) by the SM via the production of a virtual W boson. The four-fermion coupling is marked in red.

where g is the coupling constant of the weak interaction and m_W the mass of the W boson.

Modern approaches to EFTs, in particular EFT additions to the SM, aim to describe corrections to the SM at currently accessible energies, without knowing which high-energy phenomena give rise to these corrections. A significant number of higher-dimension operators can be constructed from SM fields, but only a fraction of these is gauge invariant while also conserving baryon and lepton numbers. The first operators fulfilling these requirements arise at dimension 6. The analysis presented in this dissertation is based on the EFT outlined in Ref. [57], which includes 59 independent dimension-6 operators, denoted $\mathcal{O}_i$. The energy scale up to which this theory is assumed to be valid is $\Lambda = 1\,\text{TeV}$. The strength of each operator is scaled by a dimensionless Wilson coefficient C_i. The full EFT Lagrangian is thus given by

$$\mathcal{L}_{\text{EFT}} = \sum_i \frac{C_i}{\Lambda^2}\mathcal{O}_i. \tag{2.5}$$

Only four operators in this theory lead to tHq FCNC couplings:

$$\mathcal{O}_{u\phi}^{13} = (\bar{u}_L\,\bar{d}_L)\tilde{\phi}t_R\left(\phi^\dagger\phi - \frac{v^2}{2}\right) \qquad \mathcal{O}_{u\phi}^{23} = (\bar{c}_L\,\bar{s}_L)\tilde{\phi}t_R\left(\phi^\dagger\phi - \frac{v^2}{2}\right) \tag{2.6}$$

$$\mathcal{O}_{u\phi}^{31} = (\bar{t}_L\,\bar{b}_L)\tilde{\phi}u_R\left(\phi^\dagger\phi - \frac{v^2}{2}\right) \qquad \mathcal{O}_{u\phi}^{32} = (\bar{t}_L\,\bar{b}_L)\tilde{\phi}c_R\left(\phi^\dagger\phi - \frac{v^2}{2}\right) \tag{2.7}$$

Here, the definitions for quark and Higgs fields are the same as those used in Chapter 1. The lower indices of the operators represent the fact, that a right-handed up-type quark and the Higgs field are involved in the coupling. The upper indices indicate the generation of the involved left-handed and right-handed quark fields. Operators in which the involved up or charm quark is left-handed (i.e. $\mathcal{O}_{u\phi}^{13}$ and $\mathcal{O}_{u\phi}^{23}$) are referred to as left-handed, the others (i.e. $\mathcal{O}_{u\phi}^{31}$ and $\mathcal{O}_{u\phi}^{32}$) as right-handed. The subtraction of $v^2/2$ is necessary to prevent the emergence of terms of order H^0, as these would constitute a redefinition of the mass eigenstates for the involved quarks. The individual vertices these operators give rise to can be obtained by evaluating the products of $SU(2)$ doublets. This is done exemplarily for $\mathcal{O}_{u\phi}^{13}$ in the following:

$$\mathcal{O}_{u\phi}^{13} = (\bar{u}_L,\, \bar{d}_L)\frac{1}{\sqrt{2}}\begin{pmatrix} v+H \\ 0 \end{pmatrix} t_R \left[\frac{1}{\sqrt{2}}(0,\, v+H)\frac{1}{\sqrt{2}}\begin{pmatrix} 0 \\ v+H \end{pmatrix} - \frac{v^2}{2}\right] \tag{2.8}$$

$$= \frac{1}{2^{3/2}}\bar{u}_L\, t_R\, (v+H)\cdot\left[2vH + H^2\right] \tag{2.9}$$

$$= \frac{1}{2^{3/2}}\bar{u}_L\, t_R \cdot \left(\underbrace{2v^2 H}_{tHu\ \text{FCNC}} + \underbrace{3vH^2 + H^3}_{\text{suppressed}}\right). \tag{2.10}$$

The first term gives rise to a LO tHu FCNC coupling. It is this coupling which is relevant in this analysis. The remaining terms correspond to vertices including multiple Higgs bosons. They don't play an important role, however, because the top-quark mass is not high enough to produce two or more Higgs bosons. Possible off-shell contributions only play a significant role at such high energies that they are barely accessible in current experiments.

2.2.1. FCNC Top-Quark Decay and Production

The search presented in this thesis is based on proton-proton collisions at the LHC. In such collisions, the new tHq FCNC vertices can play a role in two distinct ways. In $t\bar{t}$ production, the top quark or the top antiquark can decay via $t \to Hq$. Alternatively, a top quark and a Higgs boson can be produced via $qg \to Ht$. The former process is called the decay process, the latter the production process. Feynman diagrams of both processes are shown in Figure 2.4. The production process requires an up quark or a charm quark to partake in the hard-scattering collision. Since the proton contains two valence up quarks and the charm quark only contributes as a sea quark, the production signal process contributes primarily to the tHu channel.

In the previous section it has been shown that the considered EFT provides two types of FCNC operators, one parametrising the coupling of left-handed quarks and one that

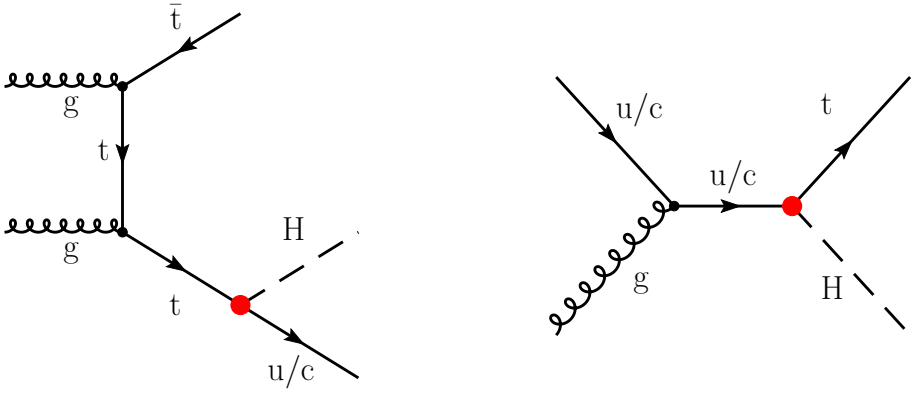

(a) $t \to Hq$ (b) $qg \to Ht$

Figure 2.4.: Feynman diagrams of (a) the FCNC $t \to Hq$ decay process and (b) the $qg \to Ht$ production process as they would arise in proton-proton collisions. The FCNC vertex is marked in red.

of their right-handed counterparts. Most BSM models giving rise to Higgs-boson related FCNC interactions do so in a non-CP violating way, meaning that in the context of these theories both operators should contribute equally. Nevertheless, to provide a treatment of these processes that is as general as possible, the operators are considered separately, taking into account potential kinematic differences between both couplings. In the case of the decay process, any sensitivity to such differences can be excluded based on theoretical considerations. The decay process is considered for a top quark-antiquark pair produced via QCD. Since gluons couple identically to left-handed and right-handed quarks, both top quarks are produced unpolarised, resulting in identical kinematics for both FCNC couplings. In principle, small effects from the spin-correlation of the two top quarks are to be expected. These effects can safely be neglected, since dedicated measurements show that they primarily impact specific, spin-related variables which are not important in this analysis [58].

For the production process, a slightly different argument is necessary. In this case, any information regarding potential polarisation of the single top quark is lost due to the scalar nature of the Higgs boson. Since its momentum is not related to the top quark's spin, the overall event kinematics are not expected to show any dependence on the handedness of the quarks involved in the FCNC couplings. This contrasts with other types of FCNC couplings, such as $tq\gamma$, where sensitivity in the coupling's handedness is observed in the production channel, owing to the photon's spin [59].

To ensure that differences between the two types of operators can truly be neglected, studies are performed with simulated samples of the left-handed and the right-handed tHu production process. Details regarding the simulations of proton-proton collisions

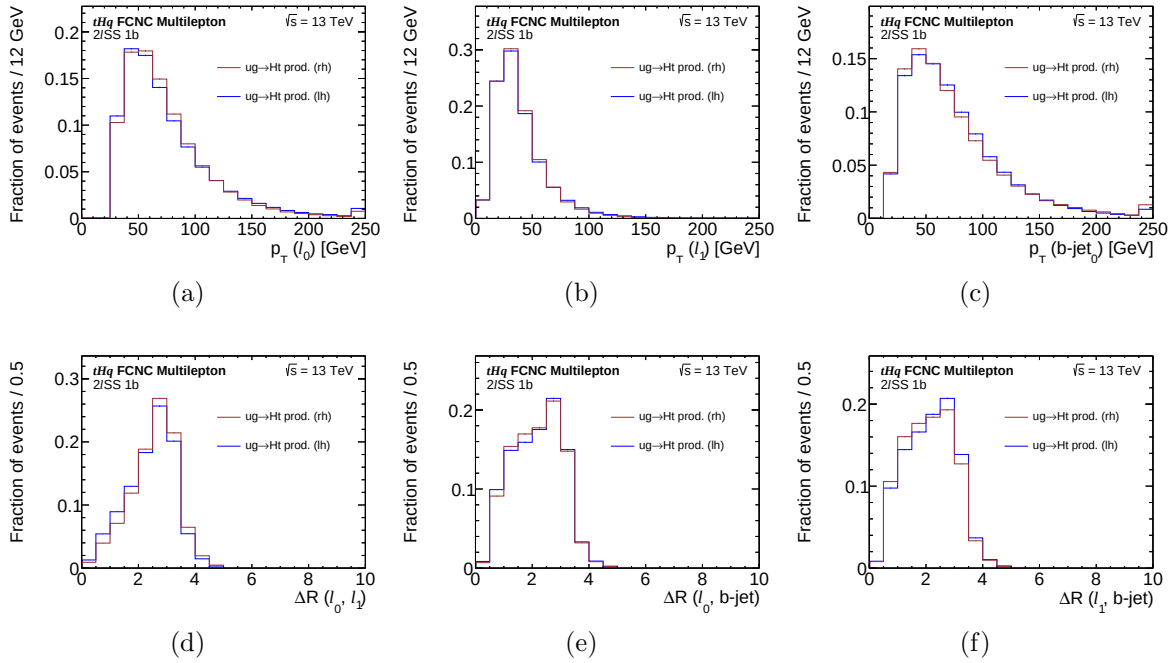

Figure 2.5.: The distribution of various kinematic variables for the left-handed (lh) and the right-handed (rh) tHu production process in the 2ℓSS final state at reconstruction-level. The vertical lines on bins depict statistical uncertainties. Transverse momenta are labelled as p_{T}, while angular separation is denoted ΔR. Leptons are ordered by their transverse momentum, with ℓ_0 being the leading-p_{T} lepton. The variables depicted include: (a) $p_{\mathrm{T}}(\ell_0)$, (b) $p_{\mathrm{T}}(\ell_1)$, (c) $p_{\mathrm{T}}(b\text{-jet})$, (d) $\Delta R(\ell_0, \ell_1)$, (e) $\Delta R(\ell_0, b\text{-jet})$ and (f) $\Delta R(\ell_1, b\text{-jet})$.

and the specific simulation samples used in this study are given in Chapters 4 and 5. The probability densities of various kinematic parameters in the 2ℓSS final state are considered and compared among the left-handed and right-handed coupling in Figure 2.5.

It can be seen that the two couplings produce effectively identical distributions, with only small differences in individual bins. These are beyond statistical fluctuations, as the statistical uncertainties in the plots indicate, and can be attributed to a slight dependence of the top quark's decay products on its handedness. Nevertheless, with the inclusion of systematic uncertainties no experimental sensitivity to these differences is expected. Identical results are obtained when performing such comparisons for the tHc coupling or in the 3ℓ final state (see Figures A.2 to A.4). Therefore, this analysis only considers the tHu and tHc processes, defined as the average of the respective left-handed and right-handed couplings. Consequently, the Wilson coefficients of individual operators are no longer subject of this analysis, being replaced by the average of the

Wilson coefficients for the left-handed and the right-handed operator:

$$C_{u\phi}^{i3,3i} := \frac{C_{u\phi}^{i3} + C_{u\phi}^{3i}}{2} \ , \ i = 1, 2. \tag{2.11}$$

2.2.2. Observables of EFT FCNC Couplings

Experiments are not directly sensitive to the EFT operators. Instead, the measurements of quantities, such as cross-sections and branching ratios, are compared to predictions, which are made as functions of the Wilson coefficients. Measurements of these quantities then allow to constrain the relevant EFT. First, it should be noted that the considered FCNC processes do not interfere with the SM, as there are no SM processes yielding identical final states at tree-level. The fact that each FCNC vertex is proportional to $C_{u\phi}$ yields the cross-section proportionality $\sigma_{\text{FCNC}} \sim |C_{u\phi}|^2$ for both the production and decay process. The decay process's cross-section can be calculated explicitly by relating it to the FCNC branching ratio $\mathcal{B}(t \to Hq)$ as follows:

$$\sigma_{(t\to Hq)(t\to Wb)} = 2\sigma_{t\bar{t}} \cdot \mathcal{B}(t \to Wb)\mathcal{B}(t \to Hq), \tag{2.12}$$

where $\sigma_{t\bar{t}}$ denotes the $t\bar{t}$ production cross-section. The two branching ratios in the equation are defined as

$$\mathcal{B}(t \to Hq) = \frac{\Gamma(t \to Hq)}{\Gamma(t \to Hq) + \Gamma(t \to Wb)} \text{ and} \tag{2.13}$$

$$\mathcal{B}(t \to Wb) = \frac{\Gamma(t \to Wb)}{\Gamma(t \to Hq) + \Gamma(t \to Wb)}, \tag{2.14}$$

where $\Gamma(t \to XX)$ denotes the *decay width* [23] of a certain top-quark decay mode. Both the EFT and SM decay widths are known from theoretical considerations, with $\Gamma(t \to Hq)$ considered at *leading order* (LO) and $\Gamma(t \to Wb)$ at *Next-to-Leading-Order* (NLO):

$$\Gamma(t \to Hq) \overset{\text{LO}}{=} \frac{|C_{u\phi}|^2}{\Lambda^4} \frac{\sqrt{2} G_{\text{F}} {m_t}^7}{8\pi} \left(1 - \frac{{m_H}^2}{{m_t}^2}\right)^2 \text{ and} \tag{2.15}$$

$$\Gamma(t \to Wb) \overset{\text{NLO}}{=} \frac{\alpha}{16 \sin^2 \theta_W} |V_{tb}|^2 \frac{{m_t}^3}{{m_W}^2} (1 - 3x^4 + 2x^6), \text{ with } x = \frac{m_W}{m_t}, \tag{2.16}$$

with the following parameters in the equations:

$\Lambda = 1.0\,\text{TeV}$	EFT energy scale
$m_t = 172.5\,\text{GeV}$	mass of the top quark
$m_H = 125.0\,\text{GeV}$	mass of the Higgs boson
$m_W = 80.385\,\text{GeV}$	mass of the W boson

$\sin^2\theta_W = 0.2342$	sine of the weak mixing angle
$G_\mathrm{F} = 1.16637 \times 10^{-5}\,\mathrm{GeV}^{-2}$	Fermi coupling constant
$\alpha = 1/137$	Fine-structure constant
$\lvert V_{tb}\rvert = 1$	tb CKM Matrix element.

These equations are used to calculate the normalisation of the decay process in the analysis. For the production process, such a simple calculation is not possible. Thus, its normalisation is determined directly from simulations, as described in Chapter 4.

2.2.3. tHq FCNC Interactions in Multilepton Final States

The FCNC interactions discussed above can be considered for any possible decay mode of the Higgs boson and the W boson from the top-quark decay $t \to Wb$. This analysis focuses on decay modes resulting in final states containing either two leptons of the same charge (2ℓSS) or three leptons with a total sum of charges equal to $\pm 1\,e$ (3ℓ). Based on the SM Higgs-boson and W-boson branching fractions, such *multilepton final states* only represent a small percentage of all possible outcomes. Figure 2.6 (a) shows the possible final states that can arise from the bare tHq processes, categorised by the number and relative charges of leptons. As can be observed, the combination of both multilepton final states comprises only approximately 1.9% of all tHq events. There are, however, only few SM processes with a similar event signature. Consequently, a high level of purity can be expected, thereby motivating a search in these final states.

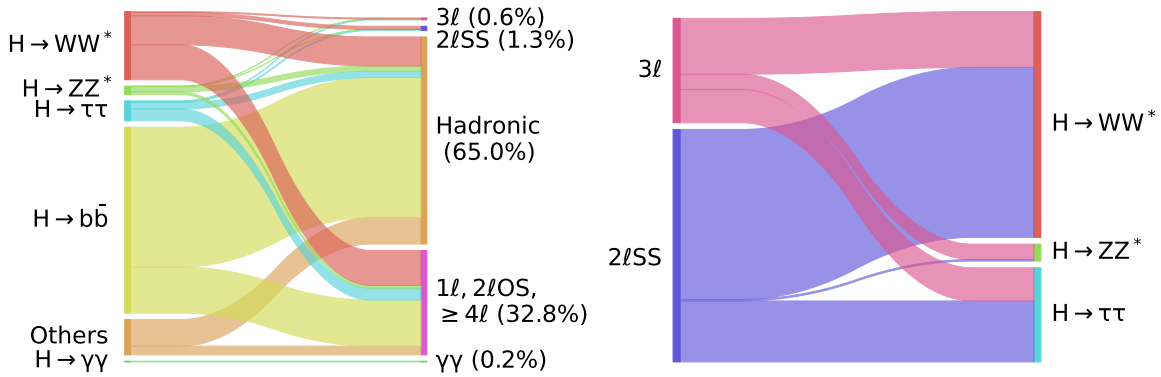

(a) (b)

Figure 2.6.: (a) The relative frequencies of possible Higgs-boson decay modes in tHq FCNC interactions and the resulting frequencies of all possible final states, and (b) the same fractions exclusively for 2ℓSS and 3ℓ final states. The graphics are based on SM branching fractions of the Higgs boson and all intermediate particles [14] and do not take into account the detector acceptance. All SM decays are considered for the top quark.

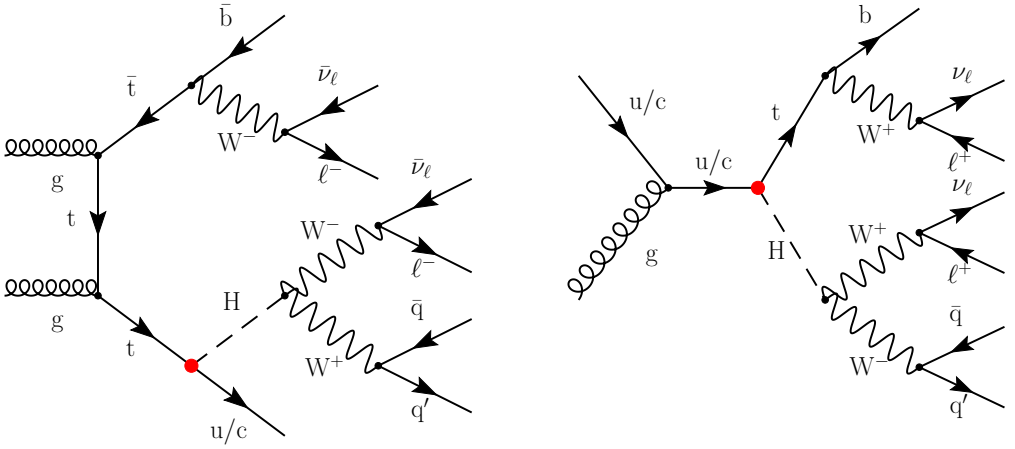

(a) $t \to Hq$ (b) $qg \to Ht$

Figure 2.7.: Feynman diagrams for tHq FCNC couplings in (a) the decay and (b) the production channel resulting in 2ℓSS final states via the $H \to WW^*$ decay mode. The FCNC vertex is marked in red.

Multiple decay modes of the Higgs boson contribute to 2ℓSS and 3ℓ final states: $H \to WW^*$, $H \to ZZ^*$ and $H \to \tau^+\tau^-$. Based on SM branching ratios, the relative proportions of the three decay modes are illustrated in Figure 2.6 (b). The $H \to WW^*$ mode is by far the most dominant in both final states, in total being responsible of 67% of all 2ℓSS and 3ℓ events. The Feynman diagrams for this mode are shown for the 2ℓSS final state in Figure 2.7, while analogous ones for the 3ℓ final state can be found in Figure A.1 The second most important decay mode is $H \to \tau^+\tau^-$, which provides 26% of all events in the 2ℓSS final state. The $H \to ZZ^*$ has a significant contribution exclusively in the 3ℓ final state, as only $Z \to \tau^+\tau^-$ decays with one leptonically decaying tau lepton can result in two leptons of the same charge. In reality, the $H \to ZZ^*$ mode will contribute slightly more to the 2ℓSS final state, as it is possible for individual leptons to not be detected in the experiment.

2.3. Searches for tHq FCNC interactions at the LHC

Searches for tHq FCNC couplings have been conducted at LHC experiments in the past. Each analysis focused on a specific Higgs-boson decay mode. In addition to decay modes resulting in 2ℓSS and 3ℓ final states, searches in the $H \to \gamma\gamma$ and the $H \to b\bar{b}$ modes, as well as the $H \to \tau^+\tau^-$ mode with one hadronically decaying tau lepton have been conducted. Orthogonality among the analyses in different channels was ensured by event selections resulting in orthogonal final states.

Table 2.1.: Summary of the results of tHq FCNC searches performed by ATLAS and CMS using Run 2 data. Analyses labelled *Combination* refer to the statistical combination of all referenced analyses with the given integrated luminosity. The observed (expected) 95% CL upper limits on the branching ratios $\mathcal{B}(t \to Hu)$ and $\mathcal{B}(t \to Hc)$ are reported. Each limit is obtained assuming the other branching ratio is zero.

Experiment/ Luminosity [fb^{-1}]	Analysis	95% CL upper limits /10^{-4} $\mathcal{B}(t \to uH)$	$\mathcal{B}(t \to cH)$
ATLAS 36.1	$H \to b\bar{b}$ [61]	52 (49)	42 (40)
	$H \to \gamma\gamma$ [62]	24 (17)	22 (16)
	$H \to \tau^+\tau^-$ [61]	17 (20)	19 (21)
	2ℓSS and 3ℓ [60]	19 (15)	16 (15)
	Combination [61]	12 (8.3)	11 (8.3)
ATLAS 140	$H \to b\bar{b}$ [66]	7.7 (8.8)	12 (7.6)
	$H \to \tau^+\tau^-$ [67]	6.9 (3.5)	9.4 (4.8)
	$H \to \gamma\gamma$ [68]	3.8 (3.9)	4.3 (4.7)
	Combination [68]	4.0 (2.4)	5.8 (3.0)
CMS 137	$H \to b\bar{b}$ [63]	7.9 (11)	9.4 (8.6)
	2ℓSS and 3ℓ [64]	7.2 (5.9)	4.3 (6.2)
	$H \to \gamma\gamma$ [65]	1.9 (3.1)	7.3 (5.1)
	Combination [64]	1.9 (2.7)	3.7 (3.5)

This section discusses three sets of searches, all performed using data obtained during Run 2 (2015-2018) with a centre-of-mass energy of $\sqrt{s} = 13\,\text{TeV}$. The first set of searches was performed with the ATLAS detector using only a partial Run 2 dataset of $36.1\,\text{fb}^{-1}$ [60–62]. The second was performed with the CMS experiment using the full Run 2 dataset of $137\,\text{fb}^{-1}$ [63–65]. Finally, searches have been conducted using the full Run 2 ATLAS dataset for all Higgs-boson decay modes [66–68], except the ones resulting in 2ℓSS and 3ℓ final states, which are the topic of this dissertation. None of the searches found evidence for the presence of tHq FCNC couplings. Consequently, 95% confidence level (CL) upper exclusion limits were set on the coefficients of the underlying effective field theory and transformed into limits on the FCNC branching ratios $\mathcal{B}(t \to Hq)$. Furthermore, the analyses in each of the three sets were statistically combined to obtain more accurate limits. The limits obtained from all searches and their respective combinations are summarised in Table 2.1.

Upon examining the expected upper limits for the $36.1\,\text{fb}^{-1}$ ATLAS searches, it can be observed that the search in 2ℓSS and 3ℓ final states is the most sensitive one. This owes to the high purity of these final states and is further motivation for the search

presented in this dissertation. The $36.1\,\text{fb}^{-1}$ ATLAS 2ℓSS and 3ℓ search did not consider the $qg \to Ht$ production signal process. Thus, improvements going beyond those from the gain in statistical precision are expected in the analysis presented here. Nevertheless, the improvements from the production signal process are expected to primarily affect the tHu signal coupling due to the aforementioned suppression of tHc production signal.

Searches for tHq FCNC couplings have been conducted by CMS in every Higgs-boson decay channel, except for the $H \to \tau^+\tau^-$ channel with one hadronically decaying tau lepton. The combination of all CMS Run 2 searches results in the strongest upper exclusion limits that have been observed for tHq FCNC branching ratios to date. They amount to $\mathcal{B}(t \to Hu) < 1.9 \times 10^{-4}$ and $\mathcal{B}(t \to Hc) < 3.7 \times 10^{-4}$.

Conversely, the combination of all previous ATLAS searches using the full Run 2 dataset of $140\,\text{fb}^{-1}$ is the most signal sensitive, as can be observed by comparing the expected upper limits. Upward fluctuations in the $H \to \tau^+\tau^-$ channel, however, cause the observed upper limit to be significantly higher than that of the CMS combination.

A summary of the strongest upper exclusion limits set by ATLAS and CMS for various top-quark related FCNC processes can be seen in Figure 2.8. The predictions of the various BSM theories discussed at the beginning of this chapter are shown. In accordance with expectations, particularly 2HDMs result in very high predicted branching ratios. As can be seen by the upper exclusion limits set by ATLAS and CMS, current experiments already manage to excluding higher ranges of flavour-conserving 2HDMs.

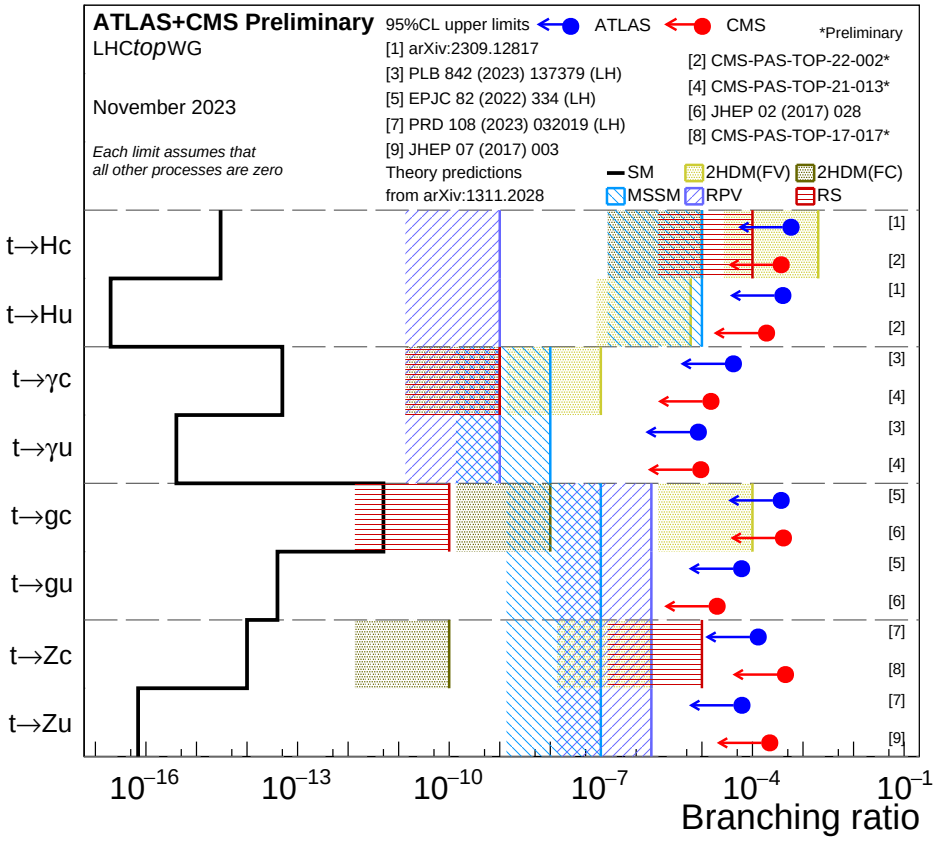

Figure 2.8.: The predicted branching ratios for various top-quark related FCNC processes, together with 95% CL upper exclusion limits set by the ATLAS and CMS collaborations [69]. The black line depicts SM predictions, while the range of predictions for multiple BSM theories are marked by textured areas. 2HDMs are separated into flavour-conserving (FC) and flavour-violating (FV) models. R-parity violating SUSY models are denoted as RPV.

3. The ATLAS Detector at the Large Hadron Collider

The analysis presented in this dissertation is based on data collected by the ATLAS detector at the LHC. This chapter presents an overview of both machines. Section 3.1 details the LHC and its properties, while Section 3.2 explains the outline of the ATLAS detector and its various subsystems.

3.1. The Large Hadron Collider

The LHC [70] is situated at the European Center for Nuclear Research (CERN) near Geneva on the Swiss-French border. With a circumference of 27 km, it is the world's largest and most powerful particle accelerator, allowing for the collision of both protons and heavy ions. To date, the maximum proton-proton centre-of-mass energy attained is $\sqrt{s} = 13.6\,\text{TeV}$ [71]. The analysis in this dissertation examines proton-proton collision data from LHC Run 2, spanning 2015 to 2018, at $\sqrt{s} = 13\,\text{TeV}$.

Four major experiments are located at the LHC. The general purpose experiments ATLAS [72] and CMS [73] aim to analyse collision data, performing precision measurements of SM parameters and searches for BSM phenomena. The *A Large Ion Collider Experiment* (ALICE) [74] focuses on heavy-ion collisions, studying quark-gluon plasma and phenomena related to the early universe. Lastly, the *Large Hadron Collider beauty* (LHCb) experiment [75] specialises in the physics of b hadrons[1], particularly in the context of CP violation.

To accelerate protons to TeV-range energies, a complex system of pre-accelerators is employed. Figure 3.1 provides an overview of the entire accelerator complex at CERN. The *Linear Accelerator 2* (LINAC 2) accelerates negatively charged hydrogen atoms to an energy of 160 MeV. These H^- atoms are then stripped of their electrons, with the resultant protons being fed into the *Proton Synchrotron Booster*. After being accelerated to 2 GeV, they enter the *Proton Synchrotron* (PS), reaching 26 GeV. Finally,

[1] b hadrons are hadrons containing at least one b quark.

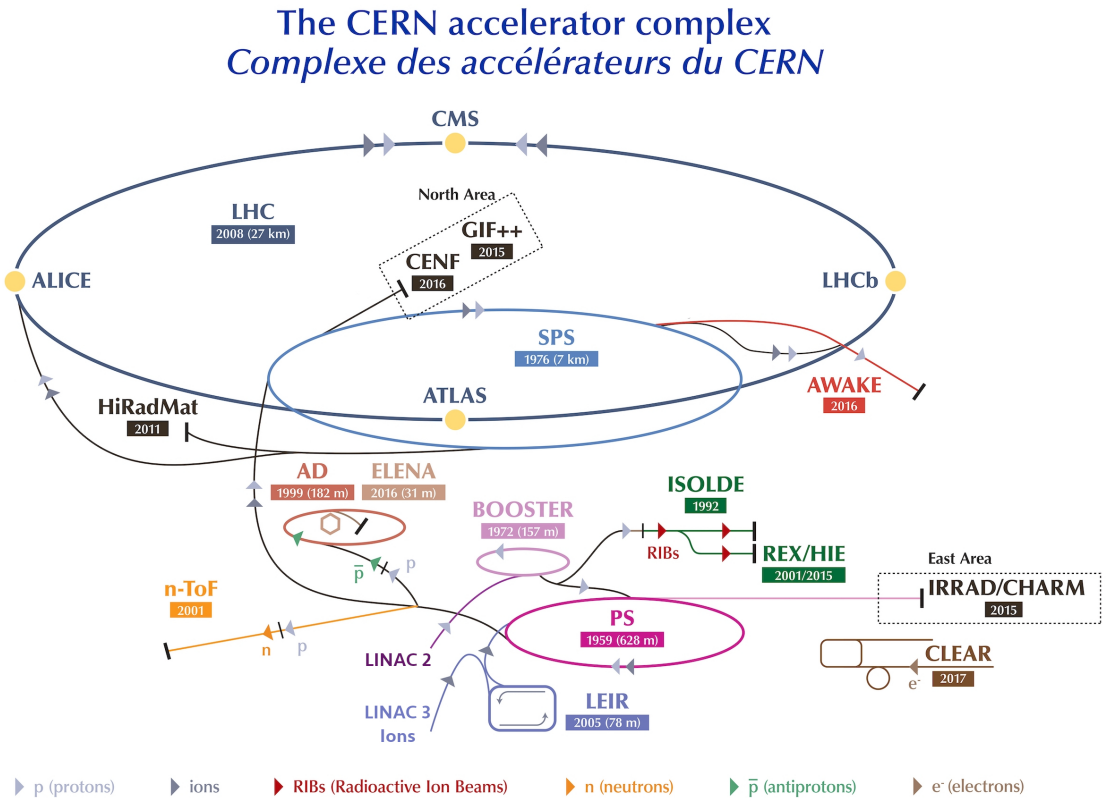

Figure 3.1.: The CERN accelerator complex, including the LHC and all its pre-accelerators. The complex is shown in the state from August 2018, the last year of Run 2 data-taking [76].

before entering the LHC, they pass through the *Super Proton Synchrotron* (SPS), which boosts them to 450 GeV. In the LHC, protons are accelerated in two parallel beam pipes, converging into a single beam pipe at interaction points, where the four major experiments are located. The process involves filling in bunches of $n_p \approx 10^{11}$ protons. There is a bunch spacing every 25 ns, theoretically allowing for up to 3564 bunches in the LHC. However, limitations in the filling and beam dumping systems permit a maximum of only 2556 bunches.

A key characteristic of the number of particle collisions in an accelerator is the *luminosity* [77]. For two beams with Gaussian profiles colliding head-on, it can be expressed as

$$\mathcal{L} = \frac{N_b n_p^2 f_R}{2\pi\sigma_x\sigma_y}, \tag{3.1}$$

where N_b represents the number of bunches, f_R the revolution frequency and $\sigma_{x,y}$ the transverse beam profile widths. As collisions in the LHC are neither strictly head-on nor feature precisely Gaussian beam profiles, real measurements require adjustments

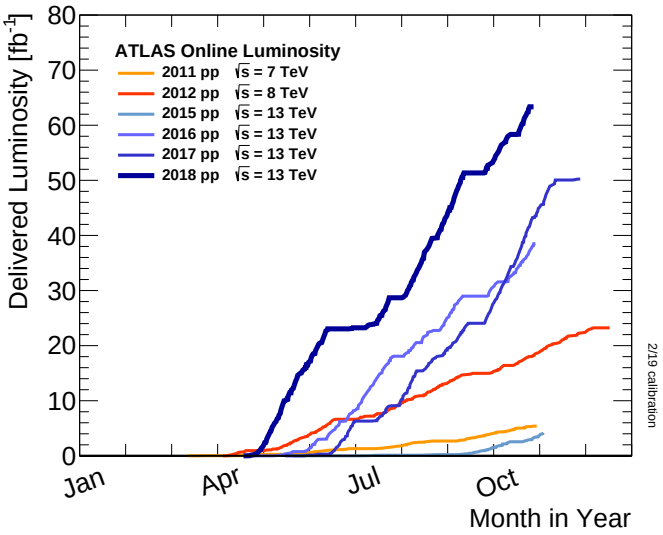

Figure 3.2.: The integrated luminosity delivered to ATLAS, shown as a function time and separated for individual years of data taking [78].

to the luminosity. The expected number of events for a given process over a specified period of time is related to the luminosity via the process's cross-section:

$$N = \sigma \int \mathcal{L} dt. \tag{3.2}$$

The time integral of the luminosity, known as integrated luminosity, often serves as a measure for the amount of data acquired at an accelerator. Figure 3.2 depicts the integrated luminosity delivered to ATLAS annually until the end of Run 2. A high instantaneous luminosity can lead to multiple proton-proton collisions within the same or neighbouring bunch-crossings, known as *pile-up*. It is characterised by the average number of collisions per bunch-crossing μ. For ATLAS in Run 2, this figure was $\mu = 33.7$.

3.2. The ATLAS detector

The ATLAS detector is the largest of the four main experiments at the LHC. It is cylindrically shaped, with a length of 46 m and a diameter of 25 m. The detector's design is optimised for the precise detection of particles resulting from hard-scattering processes, thus aiding in the deduction of fundamental interaction properties of particles. Its various detection subsystems, arranged in cylindrical layers around the detector's centre, are each tailored to measure specific properties of particles, such as their trajectory or energy. These subsystems comprise a *barrel region* with a lateral alignment around the

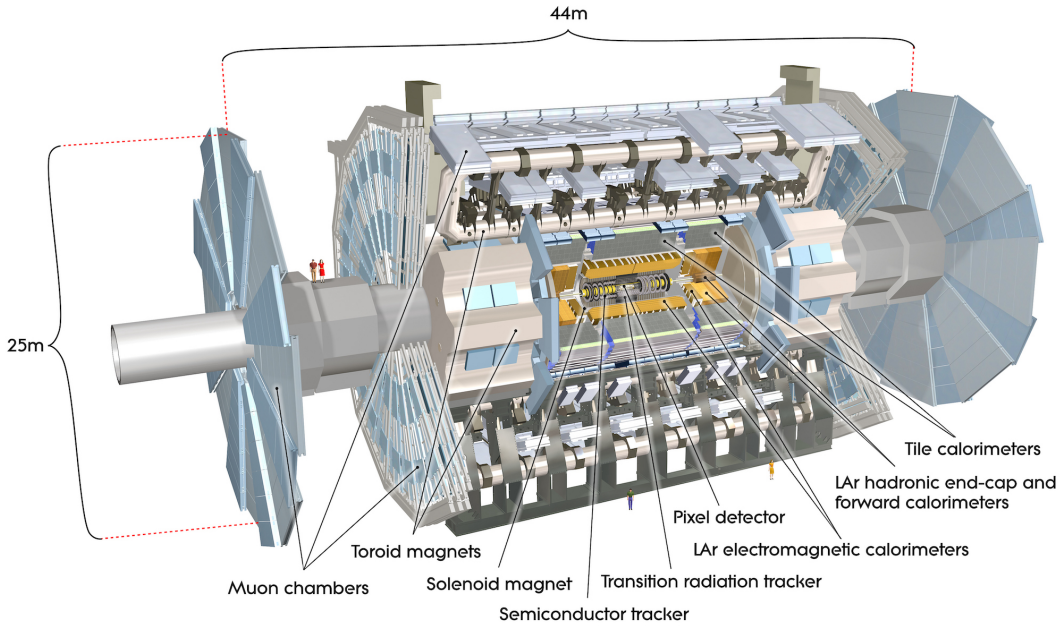

Figure 3.3.: A full schematic of the ATLAS detector, marking all sub-detectors and magnet systems [79].

beam pipe and two circular *end-caps* oriented perpendicular to the barrel. The ATLAS subsystems are described in Sections 3.2.1 to 3.2.3. A schematic of the entire ATLAS detector can be found in Figure 3.3.

ATLAS employs a right-handed coordinate system, with the origin coinciding with the centre of the detector, the interaction point. The x-axis points toward the centre of the LHC, while the y-axis points upwards. The z-axis is aligned along the beam pipe, pointing counter-clockwise when viewed from above. In describing directions of particles in the ATLAS detector, polar coordinates are often used, involving the azimuthal angle ϕ and the polar angle θ. A quantity often substituted for the information encoded in θ is the pseudorapidity

$$\eta = -\ln\left(\tan\frac{\theta}{2}\right). \tag{3.3}$$

In the relativistic limit, it approximates a particle's rapidity

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

The rapidity difference Δy between two particles is invariant under Lorentz transformations along the z-axis. Angular distances are conventionally calculated using the

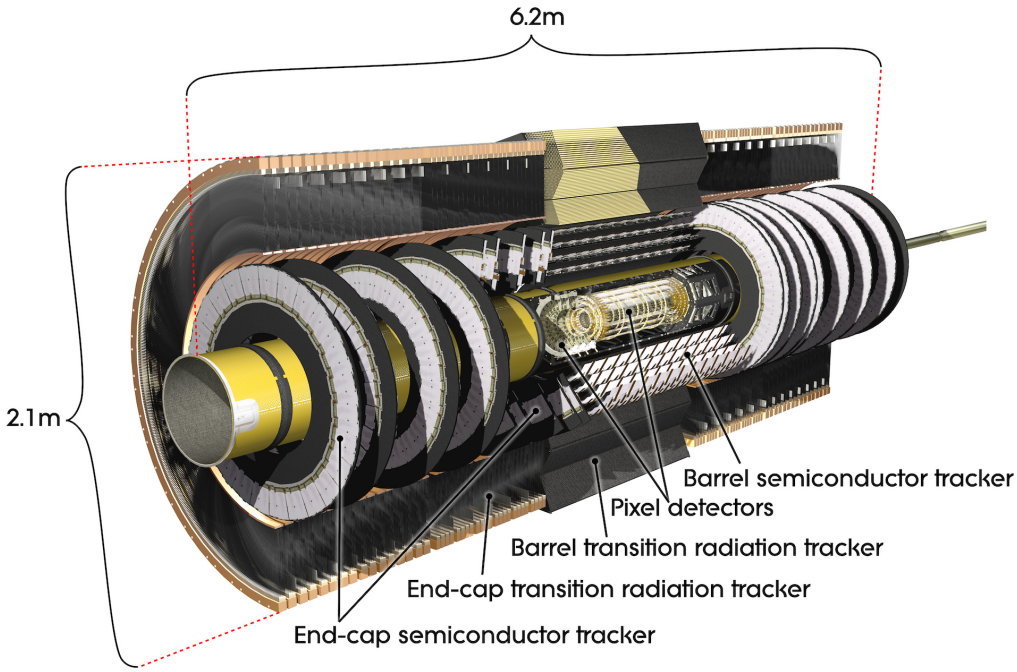

Figure 3.4.: Schematic of the ATLAS ID with all its subsystems, excluding the Pixel Detector IBL which was added after the creation of this picture [82].

pseudorapidity:

$$\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}. \tag{3.5}$$

3.2.1. The Inner Detector

The innermost part of the ATLAS detector, the Inner Detector (ID) [80, 81], is designed to identify the trajectories of ionising particles. It is immersed in a magnetic field with a magnitude of 2 T, oriented along the z-axis. This field bends the trajectories of charged particles, enabling the deduction the charge-over-momentum fraction. The ID is comprised of three complementary subsystems, each utilising different techniques to measure particle tracks. A schematic of all subsystems is shown in Figure 3.4. The total resolution in *transverse momentum* $p_\mathrm{T} = \sqrt{p_x^2 + p_y^2}$ of the ID is

$$\frac{\sigma_{\mathrm{pT}}}{p_\mathrm{T}} = 0.05\% \cdot p_\mathrm{T} \oplus 1\%. \tag{3.6}$$

The symbol $\oplus$ denotes a quadratic addition of the uncertainties.

The Pixel Detector

The innermost part of the ID is the *Pixel Detector*, covering a region up to $|\eta| = 2.5$. It comprises silicon sensors, each divided into pixels, facilitating high spatial resolution for detecting ionising particles. The original Pixel Detector consisted of three layers in both the barrel and end-cap regions, offering a resolution of 10 µm in the r-ϕ plane and 115 µm in the z direction. Before the commencement of Run 2 data taking, an additional layer known as the Insertable B-Layer (IBL) [83] was inserted in close vicinity to the beam pipe. This layer enhances the detector's ability to identify hadronic jets arising from b hadron decays, while providing an improved resolution in the z direction of 70 µm.

The Semiconductor Tracker

The *Semiconductor Tracker* (SCT), a silicon-microstrip based detector system, surrounds the Pixel Detector, covering the same $|\eta|$ range as the Pixel Detector. It adds four layers in the barrel region and nine to each end-cap. The resolution of individual sensors is coarser than that of the Pixel Detector sensors, while still maintaining an adequate double-hit resolution. Thus, the SCT still achieves accurate trajectory tracking. The overall resolution of the SCT is of 17 µm in the r-ϕ-plane and 580 µm in the z direction.

The Transition Radiation Tracker

The outermost layer of the ID is the *Transition Radiation Tracker* (TRT), which adopts a distinct method for detecting ionising particles. Unlike the preceding silicon-based sub-detectors, the TRT consists of numerous straw tubes filled with a Xenon and Argon gas mixture. These tubes are approximately 4 mm in diameter and interleaved with radiator material, which emits transition radiation in the keV range when traversed by a light ionising particle. The transition radiation and the ionising particles ionise the gas inside the straw tubes, causing a measurable change of voltage. The drift time of electrons in individual tubes is used to reconstruct the track trajectory, while the intensity of the transition radiation varies with the particle mass. For example, electrons produce more transition radiation than pions, leading to TRT hits with high charge. The number of high-charge hits along a particle trajectory is then used for particle identification, specifically to distinguish electrons from pions.

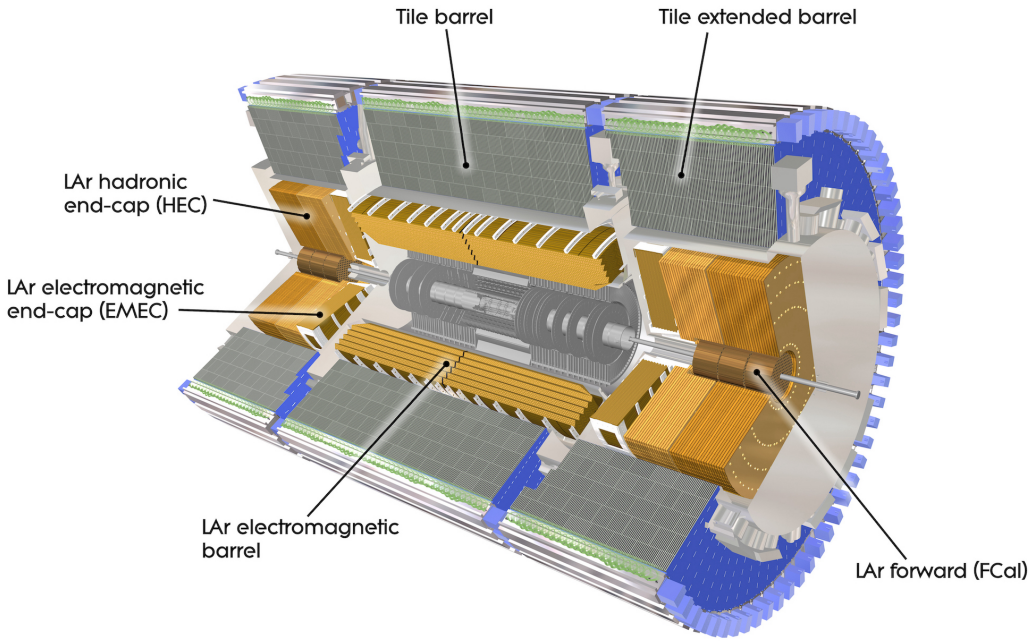

Figure 3.5.: Schematic of the ATLAS calorimeter system, depicting the individual subsystems [87].

3.2.2. The calorimeter system

The ATLAS calorimeter system [84–86] consists of three distinct parts: The *Electromagnetic Calorimeter* (ECAL), the *Hadronic Calorimeter* (HCAL) and the *Forward Calorimeter*. Their distinct designs and features are discussed in the following. A schematic of the full calorimeter system is shown in Figure 3.5.

The Electromagnetic Calorimeter

The ECAL is the next layer around the ID, measuring the energy of photons and electrons. Electrons generally lose their energy by radiating bremsstrahlung, while photons produce electron-positron pairs in the electric fields of atomic nuclei. These processes occur alternately until the bremsstrahlung photons no longer have sufficient energy to produce an electron-positron pair. The distance over which an electron on average loses all but $1/e$ of its original energy is called the radiation length X_0.

The ECAL employs a sampling technique to measure as much of this electromagnetic shower as possible. Several layers of detector material are interleaved with absorber

material to ensure that electrons and photons lose their entire energy in the ECAL. Thus, only a fraction of a particle's total energy is recorded. The correct energy is then reconstructed through precise calibration of the calorimeter. Lead and stainless steel are alternately used as absorber materials, while Liquid Argon (LAr) is used as the active detector material. The combination of all layers provides a total ECAL depth of 22-33 X_0 in the barrel region and 24-38 X_0 in the end-cap region. Layer segmentation in the η-ϕ-plane allows for spatial resolution of the measured energies, crucial for matching energy clusters to ID tracks.

For the ECAL, the barrel region covers a range of $|\eta| < 1.475$, while the end-cap region extends from $|\eta| = 1.375$ to $|\eta| = 3.2$. Structural material in the transition region $1.37 < |\eta| < 1.52$ between the barrel and end-cap reduces the energy resolution in this part of the detector. Outside this transition region, the total energy resolution is given by

$$\frac{\sigma_{\mathrm{E}}}{E} = 10\%/\sqrt{E} \oplus 0.7\%. \tag{3.7}$$

The hadronic calorimeter

Charged hadrons, having a higher radiation length then electrons, do not lose much energy in the ECAL. Like neutral hadrons, they predominantly lose energy through hadronic showers in the HCAL. Analogous to X_0, the hadronic interaction length λ is defined as the distance after which a hadron loses all but $1/e$ of its original energy through hadronic interactions.

The HCAL is constructed similarly to the ECAL, with a scintillating material as the active medium in the barrel region and LAr in the end-caps. Stainless steel is used as the absorber material in the barrel, and copper in the end-cap region. The absorber layers in the HCAL are generally thicker than those in the ECAL, as the cross-sections for hadronic interactions are lower than those for electron bremsstrahlung or photon pair production. The combined thickness of the ECAL and HCAL at $\eta = 0$ is 9.7 λ, increasing with higher values of $|\eta|$. The total energy resolution in the HCAL is

$$\frac{\sigma_{\mathrm{E}}}{E} = 50\%/\sqrt{E} \oplus 3\%. \tag{3.8}$$

The Forward Calorimeter

The ECAL and the HCAL are extended by a Forward Calorimeter, which measures particle energies in ranges of $3.1 < |\eta| < 4.9$. This calorimeter also uses LAr as the active material, but employs different absorber materials. The innermost absorber layer is

made of copper, ensuring the containment of electromagnetic showers. Subsequent layers consist of tungsten, aimed at containing hadronic showers. The Forward Calorimeter offers an energy resolution of

$$\frac{\sigma_{\mathrm{E}}}{E} = 100\%/\sqrt{E} \oplus 10\%. \tag{3.9}$$

3.2.3. The Muon Spectrometer

The *Muon Spectrometer* (MS) [88] is the outermost component of the ATLAS detector. It is designed to accurately measure the trajectories of muons, which lose minimal energy due to bremsstrahlung in the ECAL owing to their high mass. The MS is immersed in a 2.5 T toroidal magnetic field in the barrel region, which is eponymous for the ATLAS detector. It is complemented by a 3.5 T field in the end-cap region. These fields facilitate unparalleled precision in the reconstruction of muon momenta, even at very high p_{T}. The full MS with the toroidal magnetic system is depicted in Figure 3.6. For a muon at $p_{\mathrm{T}} = 1\,\mathrm{TeV}$ the momentum resolution of the MS amounts to

$$\frac{\sigma_{\mathrm{pT}}}{p_{\mathrm{T}}} = 10\%. \tag{3.10}$$

The spectrometer is structured in three layers around the HCAL and utilises multiple detection systems. To measure muon trajectories, Monitored Drift Tube (MDT) precision chambers filled with Argon and CO_2 are employed in both the barrel and the end-cap regions, covering a region up to $|\eta| = 2.7$. The have a resolution of 35 µm. Due to a significantly higher background rate, the MDTs in the innermost end-cap layer with $2 < |\eta| < 2.7$ are replaced by Cathode Strip Chambers (CSCs). These multi-wire proportional chambers offer a shorter drift time and thus handle the higher rate of background particles more effectively.

In addition to precision tracking, the MS employs three layers of Resistive Plate Chambers (RPCs) and three to four layers of Thin Gap Chambers (TGCs) for triggering purposes, as the maximum drift time of the MDTs, 700 ns, far exceeds the LHC bunch-crossing rate. RPCs operate using gas ionised by traversing muons between two resistive plates. TGCs, being multi-wire proportional chambers constructed at a smaller scale, allow for rapid readout. Their data also provides supplementary tracking information for muon reconstruction in ATLAS.

3.2.4. Trigger system

At a bunch-crossing rate of 40 MHz, it is impossible to store data from every collision. This is not a significant issue in principle, since only a fraction of events are of interest

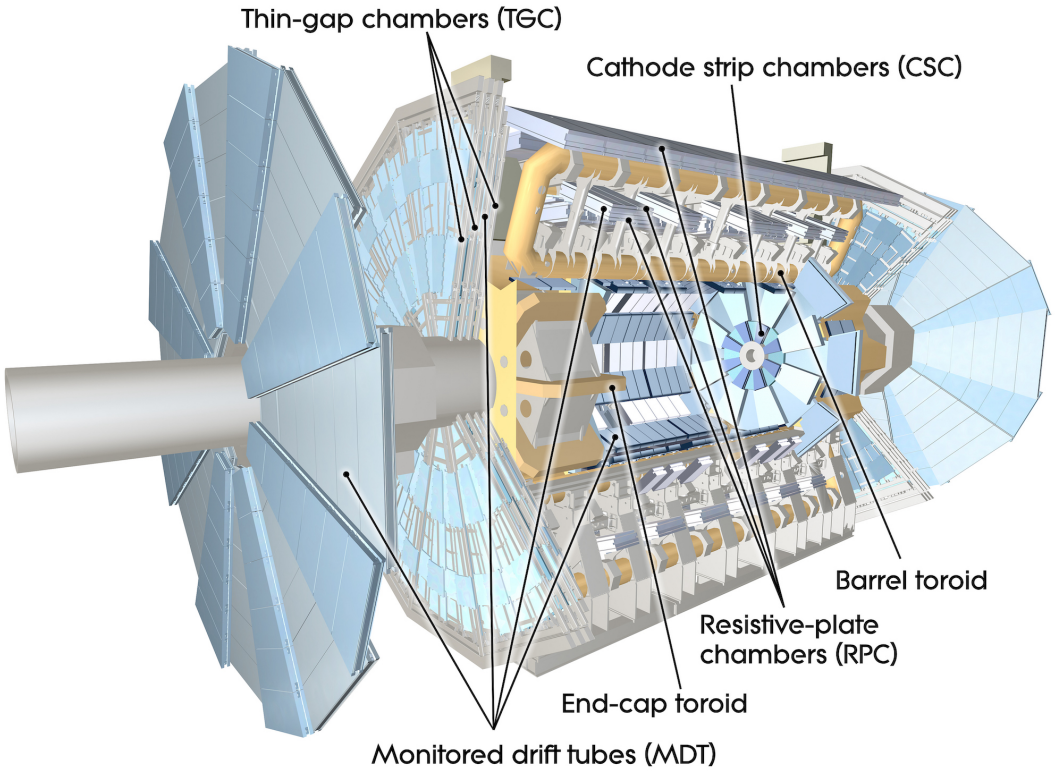

Figure 3.6.: Schematic of the ATLAS MS and the toroidal magnetic system [89].

for most physics analyses. Nevertheless, identifying these events during operation and differentiating them from background events that can be discarded remains a challenge. For this purpose, ATLAS employs a two-level trigger system [90, 91]. The *Level-1* (L1) trigger operates at the hardware level and primarily utilises information from the calorimeters and the MS, reducing the initial bunch-crossing rate to an acceptance rate of approximately 90 kHz. The *central trigger processor* evaluates this input and, if applicable, signals the readout systems of all sub-detectors to record a given event.

Events that pass the L1 trigger are then examined by the software-based *high-level trigger* (HLT). The HLT receives information on potential Regions of Interest (ROIs) from the L1 trigger and uses it to perform preliminary event reconstruction. Based on this, the HLT makes a final decision on whether or not to store the event offline. Together, both triggers reduce the rate of recorded events to approximately 1 kHz.

The analysis presented in this dissertation utilises single lepton triggers. These are designed to identify events with at least one electron or muon. A primary requirement is imposed on the transverse momentum of a candidate object. For electrons, the HLT additionally defines three different likelihood-based quality criteria: loose, medium and tight. These criteria represent increasing likelihood levels that a given ROI signal was

Table 3.1.: The HLT trigger conditions applied in this analysis. The conditions of at least one row must be met for an event to be stored. In the quality criteria, the *var* keyword represents a variable isolation cone size. If a minimum L1 p_{T} threshold is specified, the HLT only triggers if the L1 hardware trigger registers a transverse momentum exceeding this threshold.

Type	Year	Min. p_{T} [GeV]	Quality crit.	Isol. crit.	Min. L1 p_{T} [GeV]
Single electron triggers	2015	24	medium	-	20
		60	medium	-	-
		120	loose	-	-
	2016,2017 and 2018	26	tight	*var*-loose	-
		60	medium	-	-
		140	loose	-	-
Single muon triggers	2015	20	-	loose	15
		50	-	-	-
	2016, 2017 and 2018	26	-	*var*-medium	-
		50	-	-	-

caused by a true electron. For both electrons and muons, isolation criteria are established to identify electrons or muons from decays of W or Z bosons. These criteria involve cones, within which no other ROI is permitted to overlap. The sizes of these cones can vary depending on the overall event topology. The isolation criteria are again labelled loose, medium and tight, denoting increasingly stringent isolation from other ROIs. Table 3.1 displays the trigger criteria used in this analysis for each year of data-taking.

3.2.5. Luminosity measurement

A precise measurement of the luminosity is integral to obtaining accurate results in physics analyses. The integrated luminosity is measured by determining the instantaneous per-bunch luminosity

$$\mathcal{L}_b = \frac{\mu_{\mathrm{vis}} f_r}{\sigma_{\mathrm{vis}}} \tag{3.11}$$

and integrating it over time. Here, μ_{vis} is the visible interaction rate, i.e. the number of interactions per bunch-crossing that produce a measurable signal in the detector. σ_{vis} is the analogously defined visible cross-section, representing the part of the total inelastic cross-section that falls within the detector's acceptance. σ_{vis} is a calibration constant determined in *van der Meer scans* [92] during dedicated runs once each year. μ_{vis} must be measured regularly. For this purpose, the upgraded version of the Luminosity Cherenkov Integrating Detector (LUCID-2) was primarily used during Run 2, with

Table 3.2.: The Run 2 luminosity delivered by the LHC and and that recorded by ATLAS and deemed good for physics analyses [94].

Year	2015	2016	2017	2018	Total
Delivered by LHC [fb^{-1}]	4.0	39.0	50.6	63.8	157.4
Usable by ATLAS [fb^{-1}]	3.24(4)	33.40(30)	44.63(50)	58.79(64)	140.07(117)

supplementary information from the Beam Conditions Monitor [93, 94]. LUCID-2 consists of multiple photomultipliers arranged at $z = \pm 17\,\text{m}$ around the interaction point and covers a pseudorapidity of $5.561 < |\eta| < 5.641$. The quartz windows of the photomultipliers produce Cherenkov light, the intensity of which is directly related to μ_{vis}. Out of the full luminosity delivered by the LHC, only a certain fraction is recorded by ATLAS and usable for analyses due to busy readout chains or individual detector systems being offline. The delivered and usable luminosity for each year of Run 2 is shown in Table 3.2. The total luminosity recorded by ATLAS during Run 2 corresponds to $140.1(12)\,\text{fb}^{-1}$ [94].

4. Simulation of Proton-Proton Collisions

Analyses of data from the ATLAS detector predominantly involve comparisons of observed data with theoretical predictions. In the context of the analysis presented in this dissertation, such comparisons are made to ascertain the agreement of SM predictions with data and to check if EFT corrections to the SM could describe potential deviations. However, due to the complexity of the observables, predictions from both the SM and the considered EFT cannot be calculated analytically. Thus, methods of phenomenological modelling are employed to obtain accurate descriptions of physical observables. In perturbative calculations, complex integrals often arise, solved through Monte Carlo (MC) integration.

It is impossible to describe QCD phenomena at low energy scales perturbatively due to confinement. Hence, the simulation of proton-proton collisions is separated into high-energy and low-energy phenomena. Protons are composite particles, consisting of a combination of quarks and gluons, known as partons. The partonic substructure of a proton represents a low-energy phenomenon, in contrast to the hard scattering of two high-momentum partons occurring at high energies. Section 4.1 explains how the two components of the hard-scattering interaction can be separated. The partonic cross-section as a fundamental quantity used for the prediction of general observables in the interaction is calculated perturbatively. Section 4.2 explains how this calculation is performed at fixed order, while Section 4.3 details the inclusion of higher-order corrections using parton showers. These describe the evolution from the high-energy collision to low-energy partons that are subsequently confined into hadrons. Parton showers include the simulation of initial state radiation (ISR) and final state radiation (FSR), corresponding to the QCD emission of additional particles from the incoming or outgoing partons of the hard-scattering process, respectively. Most programs implementing such parton showers also simulate the underlying event (UE), the kinematics of remnants of the protons not taking part in the hard-scattering event, and multi-parton-interactions (MPIs) between these remnants. Following the showering, the hadronisation of low-energy partons, a non-perturbative phenomenon, is simulated using heuristic models. Finally, Section 4.5 explains how the response of the ATLAS detector is incorporated into the full simulation. Figure 4.1 shows a schematic overview of the most important aspects of the simulation of proton-proton collisions.

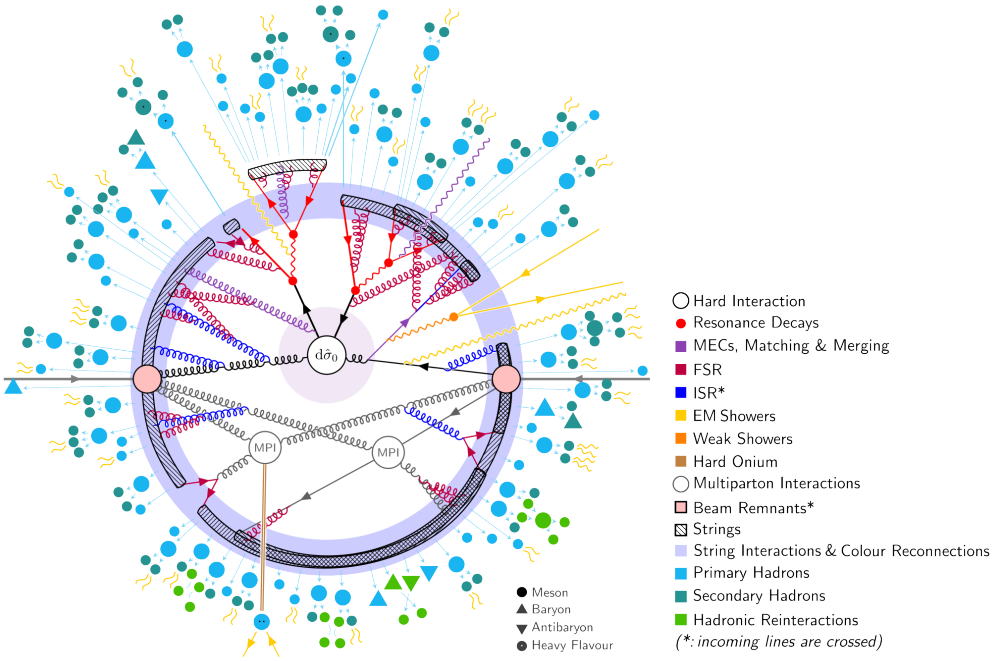

Figure 4.1.: Schematic overview of the different aspects of a proton-proton collision at the LHC. The event is separated into the hard-scattering of two partons in the proton and a complex showering process involving ISR and FSR. All labels shown are explained throughout this chapter. Figure adapted from Ref. [95].

4.1. Factorisation of PDFs and Cross-Section

Protons contain an intricate substructure of partons, most notably three *valence quarks* (two up quarks and one down quark), bound together by the strong force. The binding gluons continuously produce off-shell quark-antiquark pairs and other gluons, all of which take part in interactions with other partons. These quarks are known as *sea quarks*. When a proton is accelerated, each parton carries a fraction of the total momentum. The momentum fraction of an individual parton is given by the *Bjorken scaling variable* x [16]. The so-called *parton density functions* (PDFs) $f_a(x, Q^2)$ represent the probability to find a given parton a with a momentum fraction x when probing the proton at some energy scale Q^2. PDFs cannot be derived from first principles, but must instead be derived from experimental measurements. This can be done by means of deep inelastic scattering of electrons, muons or neutrinos on protons, by fixed-target experiments or by measurements at hadron colliders. Various groups combine the results of such experiments, employing different approaches to PDF parametrisation and varying levels of collinear parton evolution, and incorporating diverse experimental findings. The analysis presented in this dissertation uses PDFs calculated by the NNPDF group.

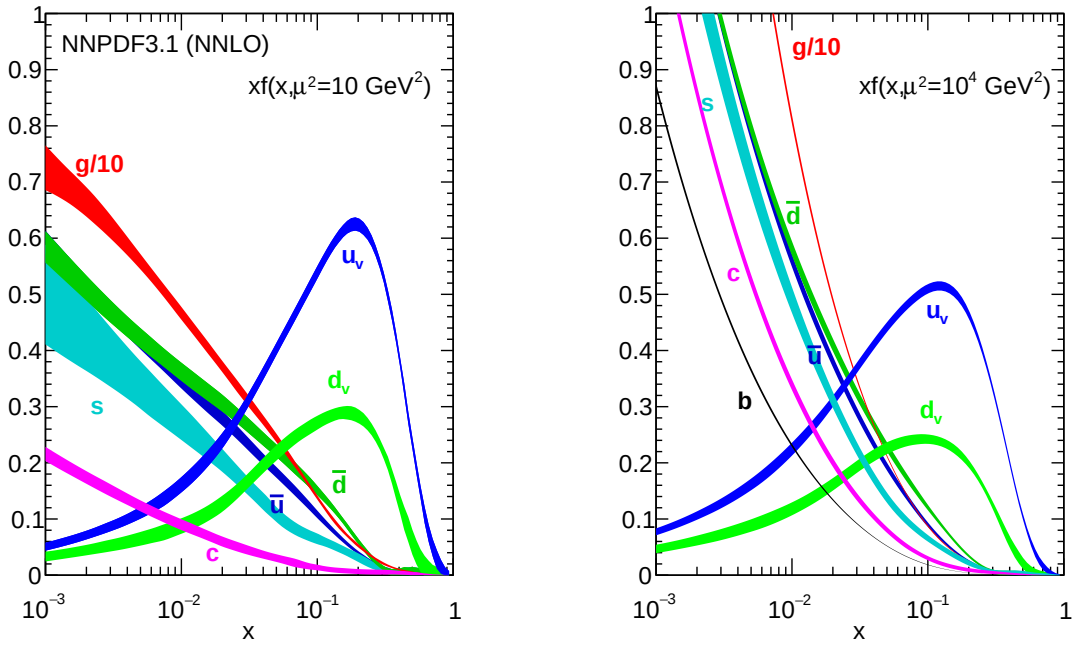

Figure 4.2.: PDFs of all partons from the NNPDF3.1NLO set, extrapolated to an energy scale of (a) $\mu^2 = 10\,\mathrm{GeV}^2$ and (b) $\mu^2 = 10\,000\,\mathrm{GeV}^2$ [96].

PDFs determined at a specific energy scale Q^2 can then be perturbatively extrapolated to other energy scales using the *DGLAP evolution equations*. They are derived based on the assumption that probing the proton at lower scales increases the sensitivity to progressively more collinear emissions. Figure 4.2 shows the NNPDF3.1NLO PDF set [96] extrapolated to scales of $\mu^2 = 10\,\mathrm{GeV}^2$ and $\mu^2 = 10\,000\,\mathrm{GeV}^2$. It can be seen that valence quarks dominate the high-x regime, while gluons peak at low momentum fractions. In high-energy collisions, gluons are by far the most dominating constituent of the proton.

The *factorisation theorem* [97] states that, at sufficiently high energy scales, the cross-section for producing a final state X in proton-proton collisions can be separated into the low-energy (*soft*) PDF contributions and the high-energy (*hard*) partonic cross-section $\sigma_{ab\to X}$:

$$\sigma_{pp\to X} = \sum_{a,b} \int dx_a dx_b f_a(x_a, \mu_\mathrm{f}^2) f_b(x_b, \mu_\mathrm{f}^2) \sigma_{ab\to X}(x_a x_b s, \mu_\mathrm{r}^2, \mu_\mathrm{f}^2). \tag{4.1}$$

The partonic cross-section depends on the centre-of-mass energy of the two colliding partons, which is related to the proton-proton centre-of-mass energy by their respective

momentum fractions. The scale μ_{f}^2 indicates the boundary beyond which partons are presumed to be treatable as free particles. Higher-order QCD loop corrections in the partonic cross-section are evaluated at the renormalisation scale μ_{r}^2. Typically, both values are set equal to a scale characteristic for the process under consideration.

4.2. Cross-sections at fixed order

Partonic cross-sections are typically calculated perturbatively up to a fixed order in $\alpha_{\mathrm{S}}(\mu_{\mathrm{r}}^2)$ [23]. The full partonic cross-section to all orders is given by

$$\sigma_{ab\to X}^{n} = \frac{1}{x_a x_b s} \sum_{k=0}^{\infty} \int d\phi_{X+k} \left| \sum_{l=0}^{\infty} \mathcal{M}_{X+k}^{(l)}(\phi_{X+k}, \mu_{\mathrm{r}}^2, \mu_{\mathrm{f}}^2) \right|^2 . \tag{4.2}$$

Here, ϕ represents the phase space and $\mathcal{M}$ the matrix element. The k in the lower index $X + k$ gives the number of additional *real emissions* in the final state, while the upper index (l) signifies the number of loops (*virtual corrections*) entering the matrix element. When considering the cross-section to a given order n, only matrix elements with $k + l \leq n$ enter the summation, with terms of the order α_{S}^n being considered in the integration. With the obtained cross-section, predictions for any observable related to the given process can be made. The corresponding expectation value can be calculated via

$$\langle \mathcal{O} \rangle_X = \sum_{k=0}^{n} \int d\phi_{X+k} \sum_{l=0}^{n-k} \frac{d\sigma_{X+k}^{(l)}}{d\phi_{X+k}} \mathcal{O}. \tag{4.3}$$

In practice, all integrations are performed using dedicated MC methods. These methods sample points in the phase space of the hard scattering process and assign a corresponding weight to each result. A large collection of these so-called *events* are expected to reproduce the distributions of all physical observables of interest. However, fixed-order calculations have certain limitations. Divergences arise in the matrix element when one or more partons in the final state are collinear to each other. This also holds true for soft radiation. These effects are handled by removing the affected regions of phase space from the integration, or by subtracting relevant parts of the integration. The soft and collinear emissions are then handled by a dedicated parton shower simulation.

4.3. Parton Shower and Hadronisation

Fixed-order calculations are capable of describing hard processes with high precision. However, in QCD such processes usually result in collimated sprays of hadrons, known

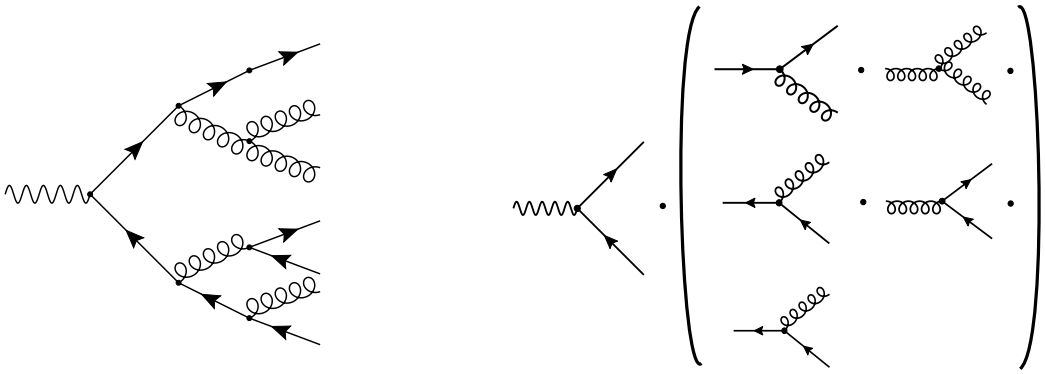

(a) (b)

Figure 4.3.: Schematic illustration of the factorisation of collinear and soft emissions. The full Feynman diagram of the process (a) is separated into the hard scattering process producing two quarks and the product of various terms of gluon emissions and quark-antiquark pair production (b). It should be noted that this factorisation presumes an ordering in a certain quantity, for example the p_{T} or the angle of the emission.

as jets. These are produced by successive radiation of bremsstrahlung and quark pair-production, with secondary contributions from electromagnetic radiation. Often this results in $\mathcal{O}(100)$ partons in the final state. Calculating matrix elements at this order is practically infeasible. Additionally, most of the radiated partons are soft or collinear to the original radiator due to logarithmic enhancements. This leads to the aforementioned singularities.

4.3.1. Parton Shower Algorithms

To account for these limitations, a number of well-separated partons are produced in the fixed-order calculation and subsequently evolved using *parton showers* [95]. These make use of the fact that in the limit of soft and collinear emissions the full matrix element factorises into the matrix element of non-collinear particles and additional terms describing the emissions. In addition to matrix elements, the phase space factorises. This full factorisation of calculations is illustrated schematically in Figure 4.3. It allows for an evolution of the shower from the scale of fixed-order calculations up to a non-perturbative scale of $Q_{\text{non-pert.}} \approx 1\,\mathrm{GeV}$, where partons are confined into hadrons.

A parton shower algorithm is based on the probability density $\mathcal{P}_a(Q^2)$ for a given parton a to split at the scale Q^2. For FSR, the differential $d\mathcal{P}(Q^2)$ is constructed using transition

probabilities $P_{a\to bc}(z)$, with $z \in [0,1]$ parametrising the momentum fraction of b. These include effects of collinear emission (usually in the form of DGLAP splitting kernels) and of soft emission. The differential branching probability $d\mathcal{P}(Q^2)$ can be calculated as

$$d\mathcal{P}(Q^2) = \frac{dQ^2}{Q^2}\frac{\alpha_{\mathrm{S}}(Q^2)}{2\pi}\sum_{b,c}\int_{z_{\min(Q^2)}}^{z_{\max(Q^2)}} P_{a\to bc}(z)dz. \tag{4.4}$$

Here, the limits of the z-integration are chosen such that energy and momentum are conserved in the splitting, as in the case of two on-shell child partons neither can obtain the parent's total momentum. Even with such boundaries, proper conservation of four-momentum cannot be achieved in the splitting of a single parton into two. Instead, an additional parton must act as a recoiler for the emission. Certain shower algorithms account for this by modelling radiation from colour dipoles rather than individual partons, thus considering transition probabilities $K_{\tilde{ij}\to ijk}(z)$ for two partons to branch into three.

To determine the probability of the parton a starting off at a scale Q_0^2 to decay at some scale $Q_1^2 < Q_0^2$, the aforementioned probability density must be multiplied by a no-transition probability $\Pi_a(Q_0^2, Q_1^2)$ known as a *Sudakov factor*. This factor can be obtained by exponentiating the probability of a decay between the two scales:

$$\Pi_a(Q_0^2, Q_1^2) = \exp\left(-\int_{Q_0^2}^{Q_1^2} d\mathcal{P}(Q^2)\right). \tag{4.5}$$

Such Sudakov factors implicitly contain approximate loop corrections to all orders. The above assumptions are used to configure the *shower veto algorithm* [95], which integrates the splitting functions into the overall MC integration of the parton shower.

For ISR, the approach is similar, but distinctively different in one way: The evolution is conducted in reverse, beginning with the high-energy partons involved in the hard-scattering interaction. Rather than calculating the probability of a parton a branching into two other partons bc, the probability that parton a originated from a splitting $x \to ay$ is considered. This is done successively, such that in each step on additional parton (y) is radiated. The evolution is performed by decreasing the scale of the radiated parton y in each step, until it either reaches the non-perturbative scale $Q_{\text{non-pert.}}$ or the total energy of the parton x surpasses that of the incoming proton. The backward evolution of ISR implicitly covers DGLAP evolution to leading order, thereby accounting for PDF effects in the simulation.

For MPIs, a different evolution type is utilised, focusing on the scale at which remnant partons of the proton interact. These interactions can give rise to additional, lower-scale parton showers that are incorporated into the simulation. In the complete event evolution, FSR, ISR and MPI each sample a scale, with the largest being selected for further evaluation. The veto algorithm subsequently determines whether a splitting occurs at this scale and the evolution continues.

4.3.2. Hadronisation Models

When the shower evolution of both ISR and FSR has reached the scale $Q_{\text{non-pert.}}$, all produced partons must be combined into confined hadrons. This is achieved using dedicated, heuristic hadronisation models. These models consider factors such as spacial separation and colour connections to produce colour-neutral composite particles. The subsequent decay of unstable hadrons is performed in this context as well, until all particles are stable. In this context, stability is defined as a lifetime that exceeds the time in which a particle traverses the detector.

Two common approaches to hadronisation exist. The first is the *Lund string model* [98]. This model posits that partons with opposing colour charges are connected via strings, with a potential proportional to the string length. Gluons are represented by kinks in these strings. Strings with a potential energy exceeding the threshold for quark-antiquark pair production can break, leading to the formation of new quarks that are incorporated into the hadronisation process. The process concludes once all strings have reached a sufficiently short length.

A second hadronisation model, known as the *cluster model* [99], begins by grouping all quarks produced in the parton shower into colour-neutral clusters. Gluons in the parton shower are decayed into quarks and also considered in the clustering. Based on their invariant mass, clusters either break down into hadrons or smaller clusters, which then undergo further subdivision. The process is repeated until only hadrons exist.

Hadronisation models, along with adjustable features of the parton shower and MPI simulation, introduce various parameters that cannot be deduced from first principles. They are instead established through comparison with experimental data. A set of parameters calibrated to align with observed data is called a *tune*.

4.3.3. Implementations of Parton Showers

Multiple programs for the simulation of high-energy collisions exist. The ones used in this analysis are Pythia [100], Sherpa [101, 102] and Herwig [103, 104]. Pythia's

default *simple shower* uses a p_{T} ordered showering together with a hadronisation model based on the Lund string ansatz. The default SHERPA *CS shower* is also ordered by p_{T}, while being based on the Catani-Seymore dipole picture [105]. The default HERWIG $\tilde{q}$ *shower* is based on angular ordering, evolving from the parton emitted at the largest angle to more and more collinear emissions. Both SHERPA and HERWIG use the cluster model for hadronisation, but provide interfaces to PYTHIA's string model.

4.4. Resolving of Ambiguities

Shower algorithms accurately describe the kinematics of soft and collinear emissions, whereas hard emissions are more effectively described by fixed-order calculations at higher orders. A combination of both approaches is thus advisable to provide an accurate description of the full phase space. At the time of writing this dissertation, fixed-order calculations at NNLO in α_{S} represent the state of the art. The combination of such higher-order calculations with the parton shower, however, results in some complications.

Considering a process with n final state partons at LO, cross-section contributions of the order $k+m$ in α_{S} are denoted by

$$\frac{d\sigma^{(k)}_{n+m}}{d\phi_{n+m}}, \tag{4.6}$$

where (k) denotes the number of loops and m the number of real emissions of the correction. Given that the parton shower approximates all real and virtual corrections to a process, showered low-order diagrams with a particularly hard emission will inevitably overlap with soft fixed-order emissions. This overlap is illustrated in Figure 4.4 exemplarily for the $e^+e^- \to q\bar{q}$ process calculated at NLO. Various techniques exist to address the potential double-counting that arises. Some of them are applied simultaneously to optimise the agreement between simulation and data. The most significant of these techniques are discussed in the following sections.

4.4.1. Matrix element corrections

Matrix element corrections (MECs) are applicable when the matrix element $|\mathcal{M}^{(0)}_{n+1}|$, which corresponds to one additional real emission, is known. In this case, a process-dependent redefinition of the splitting functions in Equation (4.4) is conducted as follows:

$$P_{a\to bc} \xrightarrow{\mathrm{MEC}} \frac{|\mathcal{M}^{(0)}_{n+1}|}{\sum\limits_{b,c} P_{a\to bc}|\mathcal{M}^{(0)}_{n}|} \cdot P_{a\to bc}. \tag{4.7}$$

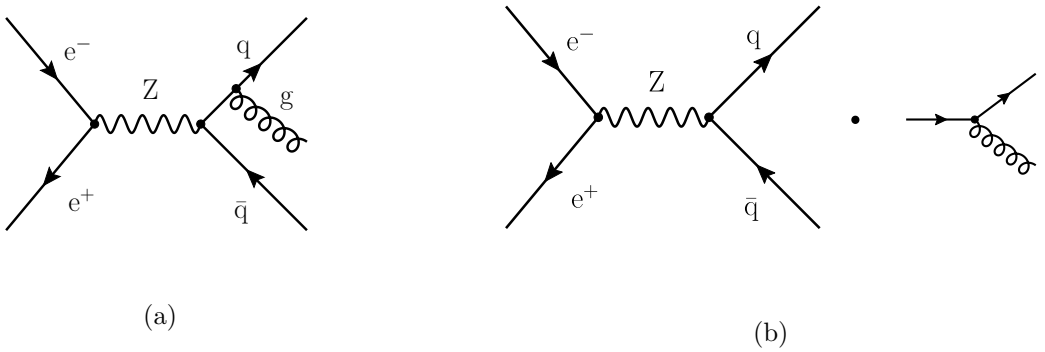

Figure 4.4.: Feynman diagrams for the $e^+e^- \to q\bar{q}$ process illustrating the overlap between fixed-order calculations and the parton-shower simulation. (a) A fixed-order diagram at NLO with one additional gluon emission and (b) a diagram at LO in fixed order with a soft or collinear gluon emission simulated by the parton shower are shown.

By construction, summing over all possible splittings then reproduces the NLO calculation $|\mathcal{M}^{(0)}_{n+1}|$, ensuring the equality

$$\frac{d\sigma^{(0)}_{n+1}}{d\phi_{n+1}} = \frac{d\sigma^{\mathrm{PS}(0)}_{n+1}}{d\phi_{n+1}}. \tag{4.8}$$

Here the superscript PS denotes the parton-shower approximation. While this process is effective for one additional real emission, devising a generalised scheme for higher orders is challenging.

4.4.2. NLO matching

Combining complete NLO calculations, including virtual corrections, with the parton shower requires the removal of all overlaps. This can be achieved by subtracting the relevant parton-shower contributions from the fixed-order calculation:

$$\frac{d\sigma^{\mathrm{NLO}}_{n}}{d\phi_{n}} \approx \frac{d\sigma^{(0)}_{n}}{d\phi_{n}} + \frac{d\sigma^{(1)}_{n}}{d\phi_{n}} - \frac{d\sigma^{\mathrm{PS}(1)}_{n}}{d\phi_{n}}, \tag{4.9}$$

$$\frac{d\sigma^{\mathrm{NLO}}_{n+1}}{d\phi_{n+1}} \approx \frac{d\sigma^{(0)}_{n+1}}{d\phi_{n+1}} - \frac{d\sigma^{\mathrm{PS}(0)}_{n+1}}{d\phi_{n+1}}. \tag{4.10}$$

This modification re-enables direct interfacing of the fixed-order calculation with the parton shower, as is typical for LO calculations. This approach corresponds to the

MC@NLO [106] matching prescription. It is automated in SHERPA and the MADGRAPH5_aMC@NLO [107] program. A modification of this approach is used in the POWHEG [108, 109] matching scheme, which is implemented in the POWHEG BOX [108–114] program. This scheme uses process-dependent MECs, ensuring the hardest real emission in a process is consistently well-behaved.

If MECs are fully used in the exponentiation for calculating parton shower Sudakov factors, the terms in Equation (4.10) vanish. In practice, agreement with data can be improved if the MEC contribution is dampened for particularly hard radiation. To achieve this, POWHEG BOX introduces a dampening factor F, which depends on a tunable parameter h_{damp} and the p_{T} of the hard radiation:

$$F = \frac{h_{\text{damp}}^2}{p_{\text{T}}^2 + h_{\text{damp}}^2}. \tag{4.11}$$

In the analysis presented in this dissertation, a nominal value of $h_{\text{damp}} = 1.5\, m_t$[1] is used in all simulations using the POWHEG scheme [115].

Another ambiguity arises in the POWHEG matching scheme, specifically when the POWHEG BOX program is interfaced to PYTHIA. Similar to PYTHIA, POWHEG BOX uses the p_{T} as an ordering parameter. Ideally, the PYTHIA shower simulation commences at the exact scale of the POWHEG emission, which, by construction, is the hardest emission of the event. However, the definition of p_{T} and hardness slightly differs between the two programs, potentially leading to double-counting or overlooking certain phase-space regions. Thus, PYTHIA implements vetoed showers with three distinct methods to calculate the initial scale of the parton shower, selectable via the $p_{\text{T}}^{\text{hard}}$ parameter. This parameter offers three possible values, each corresponding to one of the methods for estimating the initial scale of the shower:

0: Scale set externally through a dedicated Les Houches Event (LHE) file.

1: The p_{T} of the parton emitted by POWHEG relative to all other partons.

2: The minimum p_{T} of all partons with respect to all other partons.

The choice of both the h_{damp} and the $p_{\text{T}}^{\text{hard}}$ parameter introduces uncertainties in the simulation predictions, which must be accounted for when interfacing POWHEG to PYTHIA.

[1] If not specified otherwise, a value of $m_t = 172.5\,\text{GeV}$ is used throughout this dissertation.

4.4.3. Multi-jet merging

The previous two methods focused on fixed-order calculations with at most one additional emission. In certain cases, it is necessary to simulate processes with more than one hard emission. As previously outlined, these processes require fixed-order calculations, as the parton shower accuracy diminishes when producing well-separated partons. Full calculations at Next-to-Next-to-Leading-Order (NNLO) are gradually becoming feasible for parton multiplicities above 2, but are not yet applicable in this analysis. However, it is possible to calculate contributions from additional emissions at tree-level or NLO, while delegating virtual corrections at the same order in α_S to the parton shower. This approach again introduces potential overlap, addressed by slicing the phase space along a specific *merging scale* Q_{MS}.

Emissions above Q_{MS} are calculated at fixed order, whereas those below this threshold are simulated by the parton shower. To remove any overlap, all parton-shower contributions above Q_{MS} are subtracted from the fixed-order calculation. In practise, this subtraction is realised by multiplying the fixed-order cross-section with weighted Sudakov factors. The calculation of these factors introduces some ambiguity that is handled differently by various merging algorithms. For instance, the MLM merging method [116] facilitates tree-level merging by estimating the Sudakov factors by comparing the number of jets before and after the shower. Conversely, the CKKW-L method [117, 118] estimates the probabilities numerically using trial showers. A similar approach is employed by the METS method [119].

Optimal descriptions of processes are achievable by combining matching and merging methods. This combination is implemented in various ways by multiple programs. Two schemes are particularly relevant to this dissertation. METS is extended to NLO in the MEPS@NLO [120] scheme, which is used by Sherpa in combination with S-MC@NLO matching. The FxFx scheme [121] integrates the MLM merging method with aMC@NLO matching.

4.5. Simulation of Detector Interaction and Pile-up

For a comprehensive comparison between MC simulations and data, it is essential to estimate the detector response for each simulated event. This is accomplished using a detailed model of the ATLAS detector and its geometry, simulating each particle's interaction with the Geant4 toolkit [122]. Although this approach is highly accurate, it is also computationally intensive. Thus, an alternative simulation method exist, employing a parametrised calorimeter response. This method is known as AFII.

In addition to the hard-scattering event of interest, each considered SM and BSM process is subject to pile-up interactions. These pile-up interactions are incorporated into the simulation using *minimum-bias interactions*. These are overlaid on each event at the level of energy depositions, once again simulated using the GEANT4 toolkit. The minimum-bias events were simulated with the A3 set of tuned parameters [123] and the NNPDF2.3LO PDF set [124].

5. Samples of Data and Simulated Events

The data events analysed in this dissertation were recorded by the ATLAS detector between 2015 and 2018 at a centre-of-mass energy of $\sqrt{s} = 13\,\text{TeV}$. The usable data corresponds to a total integrated luminosity of $140.1(12)\,\text{fb}^{-1}$. MC processes are simulated using various combinations of matrix-element and parton-shower generators. Section 5.1 details the simulation of the BSM FCNC processes, while Section 5.2 explains the simulation of all considered SM background processes. Unless specified otherwise, samples produced with Powheg Box v2 or MadGraph5_aMC@NLO use the NNPDF3.0NLO PDF set. All samples showered with Pythia 8 use the A14 set of tuned parameters [125] and the NNPDF2.3LO [124] PDF set for showering. In each of these samples, the decays of hadrons containing b and c quarks are simulated with the EvtGen program [126]. Samples which are simulated with Sherpa employ the NNPDF3.0NNLO PDF set [127] and an internal tune for the parton shower, hadronisation and UE. The corresponding matrix-element calculations are performed using the COMIX [128] and OpenLoops [129–131] libraries.

5.1. Simulation of tHq FCNC Processes

The $t \to Hq$ decay and the $qg \to Ht$ production signal processes are modelled separately. In both cases, the FCNC EFT operators are incorporated as implemented in the Top-FCNC model [132] via the FeynRules 2.0 framework [133], allowing for calculations up to NLO in QCD.

t → Hq Decay FCNC Signal

The decay signal events are generated using Powheg Box v2 as SM $t\bar{t}$ production, prior to enforcing an FCNC decay. Both top-quark decays are modelled in MadSpin [134, 135] to preserve spin correlations, with one top quark constrained to decay via $t \to Wb$ as in the SM, and the other via the FCNC $t \to Hq$ decay mode. The produced samples are inclusive in Higgs-boson decay modes, which were modelled using Pythia 8.308

along with the parton shower, hadronisation and UE. In total, four samples are produced, modelling the tHu/tHc coupling with either the top quark or the top antiquark decaying via the left-handed FCNC vertex. As explained in Chapter 2, the resulting kinematics are expected to be identical to those of the right-handed operator. The samples are normalised to an FCNC branching ratio of $\mathcal{B}(t \to Hq) = 0.1\%$, corresponding to a Wilson coefficient of $|C_{u\phi}| = 1.33$.

qg → Ht Production FCNC Signal

The production signal process is modelled using MadGraph5_aMC@NLO 2.9.9 with the five-flavour scheme, wherein all quarks are assumed to have zero mass, except for the top quark. The decays of the top quark and the Higgs boson are modelled with MadSpin. The generated events are interfaced with Pythia 8.307. Multiple samples are produced, each with one of the four considered Wilson coefficients set to 1.0, while all others are set to zero. For each case, three samples are produced, each considering one of the three relevant Higgs-boson decay modes: $H \to WW^*$, $H \to ZZ^*$ or $H \to \tau^+\tau^-$. In the analysis, all samples are rescaled to a normalisation corresponding to a Wilson coefficient of $|C_{u\phi}| = 1.33$, consistent with the decay signal samples.

5.2. Simulation of SM background processes

Various SM processes make significant contributions to the 2ℓSS and 3ℓ final states. The approaches to simulating these processes are summarised in the following.

$t\bar{t}$ production

The production of a top quark-antiquark pair is simulated using Powheg Box v2 in the five-flavour scheme. The events are interfaced with Pythia 8.230 for the simulation of the parton shower, hadronisation and UE. They were normalised to a cross-section of $\sigma_{t\bar{t}} = 832\,\text{pb}$, as obtained from NNLO predictions by the Top++ 2.0 program [136]. Two separate samples are used: one modelling $t\bar{t}$ production with at least one charged lepton at matrix-element level, and the other requiring at least two leptons. The former is used in the 2ℓSS final state, where single-lepton events can contribute if an additional lepton is produced in the parton shower, while the latter is used in the 3ℓ final state, providing greater statistical precision for events with at least two leptons.

tW production

The production of a top quark or antiquark in association with a W boson is simulated using Powheg Box v2 interfaced to Pythia 8.230. These samples are normalised to

a cross-section of $\sigma_{tW+\bar{t}W} = 71.7\,\text{pb}$ [137]. Analogous to $t\bar{t}$ production, one-lepton inclusive samples are used in the 2ℓSS final state, while two-lepton samples are used in the 3ℓ final state, with the lepton multiplicity specified at matrix-element level.

Single-top quark s- and t-channel production

The production of a single top quark or antiquark in the s-channel or t-channel is simulated using Powheg Box v2, interfaced with Pythia 8.230. For the t-channel simulation, the alternative PDF set NNPDF2.3NLO_NF4 is used, implementing the four-flavour scheme, which sets all quark masses to zero except for those of the top quark and the bottom quark. The samples are normalised to cross-sections of $\sigma_{t\text{-channel}} = 217\,\text{pb}$ and $\sigma_{s\text{-channel}} = 10.32\,\text{pb}$, based on calculations at NLO in QCD with the Hathor 2.1 program [138, 139].

W+jets and Z+jets production

The production of a leptonically decaying W or Z boson in association with multiple jets is fully simulated using Sherpa 2.2.11. The simulation is performed in the five-flavour scheme at NLO for up to two partons and at LO for three to five partons. For the MEPS@NLO merging, a merging scale of $Q_{\text{MS}} = 20\,\text{GeV}$ is used. Separate samples are produced, categorised by the flavour of the leptons from the W- or Z-boson decay (e, μ and τ) and by the flavour of the heaviest parton emission at matrix-element level (b, c or light). The samples are normalised to the corresponding cross-sections calculated by Sherpa.

$t\bar{t}W$ production

The production of a $t\bar{t}$ pair in association with a W boson is simulated using two complementary Sherpa 2.2.10 samples. Matrix elements of the first sample are calculated at NLO for zero or one additional parton and at LO for up to two partons. Additional higher-order electroweak corrections are included by applying event-by-event correction factors, corresponding to virtual NLO electroweak corrections to $\mathcal{O}(\alpha^2\alpha_{\text{S}}^2)$ and LO corrections to $\mathcal{O}(\alpha^3)$. The second sample provides sub-leading electroweak corrections to $\mathcal{O}(\alpha^3\alpha_{\text{S}})$, as detailed in Ref. [140]. NLO QCD and NLO electroweak contributions are combined according to the methodology described in Ref. [141]. The $t\bar{t}W$ samples are normalised to a combined cross-section of $\sigma_{t\bar{t}W} = 722\,\text{fb}$, based on calculations at NLO with the inclusion of hard, non-logarithmically enhanced radiation at NLO in QCD, as performed in Ref. [141].

$t\bar{t}Z$ production

The production of a $t\bar{t}$ pair in association with a Z boson is simulated using MadGraph5_aMC@NLO 2.8.1, with the NNPDF2.3LO PDF set. Events are interfaced

with Pythia 8.244 for the parton shower, hadronisation and UE simulation. In the production of the samples, a leptonic Z-boson decay is enforced, and the samples are categorised by the flavour of the corresponding leptons (e, μ, τ). Additional samples are produced to optimally model events with an invariant mass of the Z-boson decay products of $m_{\ell\ell} < 5\,\text{GeV}$. The combined normalisation of all samples corresponds to $\sigma_{t\bar{t}Z(Z\to\ell\ell)} = 162\,\text{fb}$, in accordance with the NLO QCD and NLO electroweak accurate calculation in Ref. [142], with the supplementary inclusion of off-shell effects.

$t\bar{t}H$ production

Events of $t\bar{t}$ production in association with a Higgs boson are simulated using Powheg Box v2 and interfaced with Pythia 8.230. Two samples are produced, requiring exactly one or exactly two leptons at matrix-element level from the decay of the $t\bar{t}$ pair, respectively. The samples are normalised to a cross-section of $\sigma_t\bar{t}H = 507\,\text{fb}$ [142].

VV and VVV production

The production of two or three weak vector bosons is simulated using Sherpa 2.2.12 at NLO for up to three additional partons and at LO for four or five partons. Separate samples are produced, depending on the multiplicity of leptons after the decay of all vector bosons (2ℓSS, 3ℓ, 4ℓ). All samples are normalised to the cross-sections calculated by Sherpa.

tWZ and tZq production

The production of a single top quark in association with a Z boson and an additional quark is simulated using MadGraph5_aMC@NLO 2.9.5 interfaced with Pythia 8.230. tWZ production was simulated analogously, using the MadGraph5_aMC@NLO 2.2.2 generator interfaced with Pythia 8.212. Both samples are normalised to the values predicted by MadGraph5_aMC@NLO. The diagram removal scheme [143] was employed to handle interferences of tWZ with $t\bar{t}Z$ production.

Rare processes

Several rare processes, mainly related to top-quark production in association with other particles, can result in 2ℓSS or 3ℓ events. The processes of $t\bar{t}t$, $t\bar{t}t\bar{t}$, tHW, tHq, $t\bar{t}WW$, $t\bar{t}HH$, $t\bar{t}WH$, $t\bar{t}ZZ$ and $t\bar{t}WZ$ production are simulated using various versions of MadGraph5_aMC@NLO interfaced with Pythia 8. VH production is modelled by Powheg Box v2, also interfaced with Pythia 8. All processes are normalised to the predictions obtained from the respective matrix-element and parton-shower calculations.

6. Event Reconstruction

Precise and accurate reconstruction of physical objects is vital for measuring SM phenomena or constraining BSM physics. Intricate schemes have been devised to facilitate the identification of electrons, muons or hadronic jets, particularly those originating from b quarks. This chapter presents these schemes, along with methods used for the calibration of objects and their kinematic properties.

6.1. Tracks and Primary Vertices

The reconstruction of the ID tracks of charged particles forms the basis for the subsequent reconstruction of higher-level objects such as electrons or hadronic jets [144]. There are five fundamental quantities associated to each track. The shortest distance of the track from the collision point of an inelastic scattering process, known as the primary vertex, is called impact parameter. It can be measured in the transverse plane, denoted d_0 and in longitudinal direction, denoted z_0. The angles θ and ϕ parametrise the direction of the track at its origin, while the fraction q/p_T of the track's associated charge and its transverse momentum is determined from its curvature by measuring the sagitta. All of these quantities are depicted in Figure 6.1.

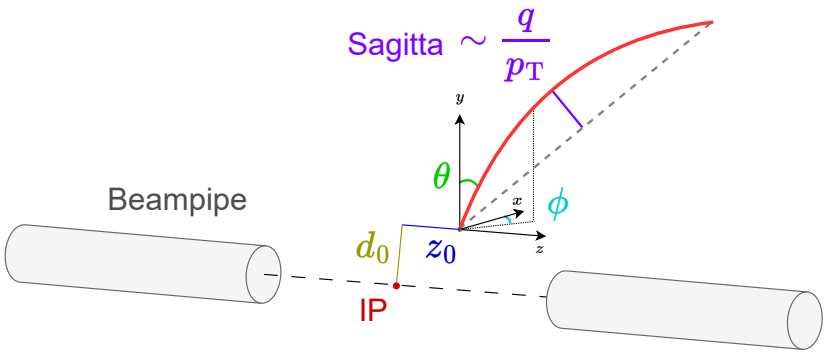

Figure 6.1.: An exemplary track in the ATLAS detector, with the defining quantities θ, ϕ, the sagitta, d_0 and z_0 indicated.

Track reconstruction begins by identifying clusters of hits in the Pixel detector and the SCT by means of *connected component analysis* [145]. The clusters serve as three-dimensional *space-points*, indicative of the traversal of one or more particles through the detector material. Numerous *track seeds* are constructed, each one consisting of three space-points. Assuming a perfectly helical trajectory in a uniform magnetic field, certain parameters, such as the charge-over-p_{T} fraction and the transverse impact parameter relative to the beam line, can be estimated. Based on these parameters, each seed is assigned a quality score. The tracks with the highest quality are input to a Kalman filter [146], which constructs track candidates from the seeds and also includes space-points not used by any seeds. Similar to the track seeds, the track candidates are assigned a multivariate track score, which measures its likelihood of being caused by a real charged particle. An ambiguity solver then removes overlaps between track candidates, giving preference to those with higher track scores. Finally, an artificial *neural network* (NN) further enhances the tracking purity, rejecting erroneous track candidates and thereby ensuring the reliability of the reconstructed tracks. Subsequently, the final track candidates are matched with TRT data to refine kinematic parameters and obtain additional information about the charged particle's mass.

Reconstructed tracks are used to identify primary vertices of hard-scattering interactions, employing an adaptive vertex finding algorithm [147]. As mentioned in Chapter 3, owing to pile-up, multiple such vertices can exist within a single bunch crossing. The algorithm establishes a seed of the first vertex based on the position of all tracks in the event. It then performs a χ^2 minimisation, considering input track parameters and assigning weights that reflect compatibility with the initial vertex estimate. The weights are then refined iteratively through further χ^2 fits, using the previously obtained weights as input, thereby ensuring that less compatible tracks exert less influence on the position calculation. Once specific convergence criteria are fulfilled, tracks that are incompatible with the obtained vertex by more than seven standard deviations are removed. The procedure is repeated with all remaining unassigned tracks to determine further vertex candidates. The vertex with the highest sum of p_{T} over all associated tracks is designated as the primary vertex, while any additional vertices in its vicinity are classified as pile-up vertices.

6.2. Electrons

Electron candidates are constructed from energy clusters in the ECAL and matching tracks in the ID [148]. Energy deposits in the ECAL are segmented into a 200×256 grid in the η-ϕ-plane, where the energy across all ECAL layers is summed up. Energy clusters are subsequently identified using a sliding-window [149] approach with a window size of 3×5. Clusters in close proximity to each other are merged into a single supercluster

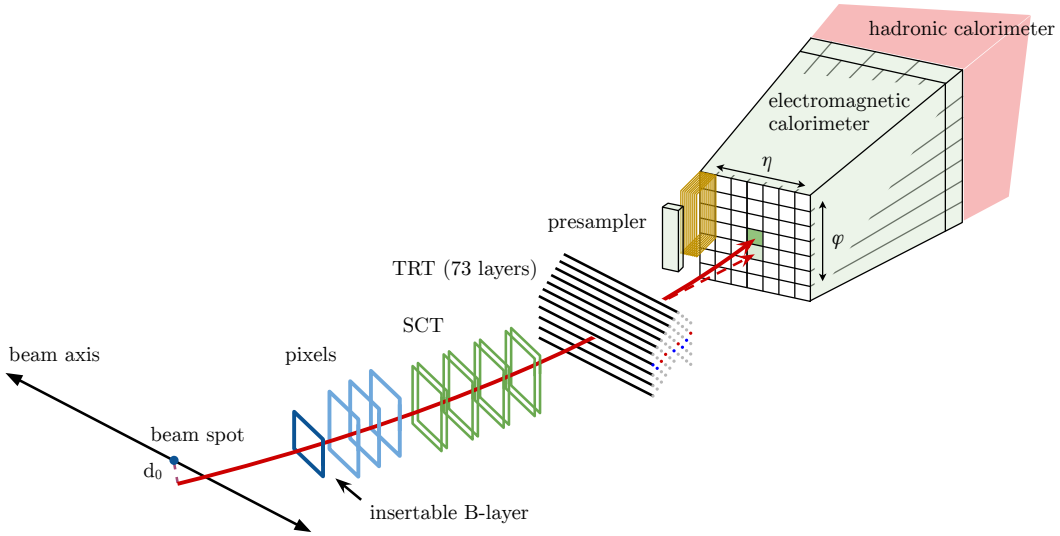

Figure 6.2.: All detector subsystems involved in the detection of an electron, together with a possible electron trajectory. After interaction with all subsystems of the ID, the electron loses its energy in the ECAL. Figure adapted from Ref. [148].

to encapsulate bremsstrahlung emitted by the electron before it enters the ECAL. The trajectory of an electron traversing all detector subsystems up to the ECAL is illustrated schematically in Figure 6.2. The electron's track is selected by employing a matching algorithm that aligns the supercluster's barycentre with the extrapolated tracks. The track is then refined using the additional data from the energy cluster. Multiple tracks may be associated to one electron, thereby accommodating the inclusion of electron-positron pairs originating from bremsstrahlung photons converting within the ID. The momentum of the electron candidate is determined using a multivariate technique, which among other things accounts for energy loss in non-active detector material. This technique is precisely calibrated using a binned likelihood fit to events in proximity to the $Z \to e^+e^-$ resonance [148].

Electron candidates are identified using a likelihood-based approach [148]. The employed likelihood function is constructed from the product of probability density functions of various quantities related to electron identification. This includes information on the signal of individual ECAL layers, track conditions, track-cluster matching and energy leakage to the HCAL. From this likelihood function, a single discriminant is derived. This analysis uses the *TightLH* working point of this discriminant [148], rejecting a substantial fraction of fake electrons, while achieving efficiencies of 65 to 90% for real electrons, increasing with electron p_{T}.

The kinematic parameters of electron candidates are evaluated to reject those erroneously identified as direct products of the hard-scattering interaction. Particular attention

is given to electrons from the decay of a b hadron. A specialised *recurrent neural network* (RNN) [150] is trained to reject these electrons, as well as those arising from $\gamma \to e^+e^-$ processes through interaction with detector material. The RNN's primary input consists of information on the electron candidates' impact parameters and the association of candidate tracks with neighbouring objects. Its output is combined with various other kinematic variables, mainly related to the angular separation between a given electron candidate and the nearest hadronic jet. This combination is executed using two dedicated boosted decision trees (BDTs) [151], one trained in the barrel region and one in the end-cap region. This elaborate procedure is called *prompt lepton improved veto* (PLIV) [152] and substantially reduces the contribution of electrons from b-hadron decays.

In the 2ℓSS final state, a significant background arises from electrons reconstructed with an inverted charge. These leptons are efficiently rejected using a specialised BDT known as *electron charge identification selection tool* (ECIDS) [153]. The ECIDS uses information on the energy, the charge, the transverse momentum, the impact parameter and the pseudorapidity of electrons, achieving a rejection rate of 82.5%.

The aforementioned reconstruction, identification and isolation methods may exhibit a different performance for MC simulations compared to actual data. Consequently, calibrations are performed using the tag-and-probe method [154] on resonant $Z \to e^+e^-$ and $J/\Psi \to e^+e^-$ decays. The method determines the efficiency for loosely selected *probe* electrons to be identified as on-resonance *tag* electrons. These efficiencies are compared between MC and data, yielding p_{T} and $|\eta|$ dependent scale factors (SFs), which are applied to MC events.

In this analysis, only electrons with a transverse impact parameter significance of $|d_0/\sigma_{d_0}| < 5$ and a longitudinal impact parameter fulfilling $|\sigma_{z_0} \sin\theta| < 0.5\,\mathrm{mm}$ are considered, ensuring all electron tracks point to the event's primary vertex (track-to-vertex association, TTVA). Moreover, a minimum transverse momentum of $p_{\mathrm{T}} > 10\,\mathrm{GeV}$ and a pseudorapidity of $|\eta| < 2.47$ are required, excluding the ECAL transition region of $1.37 < |\eta| < 1.52$ between the barrel and end-cap.

6.3. Muons

Unlike electrons, muons do not lose a lot of energy due to bremsstrahlung, owing to their high mass. Additionally, at energies in the GeV range, the energy loss through ionisation is well below the thresholds required to fully stop the muon within the ATLAS detector. Thus, the identification of muons relies predominantly on tracking data from the ID and the MS, complemented by energy deposits in the calorimeters.

The muon track reconstruction is based on local track segments, identified in individual MS layers through a Hough transform [155]. These segments are then combined into initial track candidates, incorporating the interaction point and a parabolic trajectory for a preliminary estimate. A χ^2 fit is employed to further refine these tracks, accounting for interactions with the detector and discrepancies in its alignment. The momentum scale and resolution of muon tracks are calibrated through binned likelihood fits to the mass peaks of $Z \to \mu^+\mu^-$ and $J/\Psi \to \mu^+\mu^-$ decays.

Different quality levels are defined for muon candidates in the ATLAS detector [152]. The highest quality is attributed to *combined* (CB) muons, reconstructed from a matching pair of ID and MS tracks. *Inside-out* (IO) muons rely on one ID track that is extrapolated and matched to at least three loose MS hits. So-called *segment-tagged* (ST) muons are defined analogously, but only require one matching MS hit for a given ID track. Lastly, *calorimeter-tagged* (CT) muons necessitate a muon-compatible ECAL signature at the location of an extrapolated ID track, without the requirement for any hits in the MS.

The identification requirements imposed on muons in this analysis achieve a balance between good fake-rejection and adequate real-muon efficiency. The used *medium* working point [152] exclusively includes CB and IO muons, while imposing additional criteria on the number of ID and MS hits, as well as the q/p ratio, p_{T} and η of the muon candidates.

Similar to electrons, this analysis uses a muon-specific variant of the PLIV [152]. A dedicated RNN, trained to reject muons originating from b-hadron decays, is combined with other kinematic variables in a single BDT, trained for both the barrel and the end-cap region.

TTVA requirements of $|d_0/\sigma_{d_0}| < 3$ and $|\sigma_{z_0} \sin\theta| < 0.5\,\mathrm{mm}$ are imposed on muon candidates. For the calibration of reconstruction, identification, isolation and TTVA working points, the tag-and-probe method is employed once again, utilising $Z \to \mu^+\mu^-$ and $J/\Psi \to \mu^+\mu^-$ events. Lastly, to be selected for this analysis, muon candidates must possess a minimum transverse momentum of 10 GeV and satisfy $|\eta| < 2.5$.

6.4. Jets

As explained in Chapter 1, quarks and gluons produced in the hard scattering event undergo hadronisation due to the confinement of colour charge in QCD. This process results in a cone of secondary particles, known as a hadronic jet. Hadronic jets are reconstructed using the *particle-flow* (PFlow) algorithm [156], which combines topological energy clusters in the HCAL and ECAL with tracks in the ID. The formation of clusters

is performed using a similar approach as for ECAL clustering in the identification of electrons. Tracks are subsequently matched to clusters if they align, with the constraint that one track can only be associated with one cluster. If a given track's extrapolation permits, the matched clusters are extended to encompass neighbouring clusters. The track's momentum measurement is used to estimate the expected energy deposited in the calorimeters. This estimated energy is then subtracted cell by cell, followed by a removal of any remnant deposits that fall within known fluctuations of energy measurements. The algorithm returns PFlow objects, consisting either of a track combined with a modified calorimeter cluster or of unmatched clusters, typically originating from neutral hadrons.

The PFlow objects are subsequently combined into jets using the anti-k_t algorithm [157], implemented in the FastJet software package [158]. PFlow objects with tracks are considered only if the track is associated to the event's primary vertex, satisfying the condition $|z_0 \sin\theta| < 2\,\text{mm}$. Pairwise distances between two objects i and j are calculated based on their angular separation ΔR_{ij}, and weighted by a function depending on their transverse momentum k_t and the radius parameter R:

$$d_{ij} = \min\left(\frac{1}{{k_{\text{t},i}}^2}, \frac{1}{{k_{\text{t},j}}^2}\right) \cdot \frac{{\Delta R_{ij}}^2}{R^2}. \tag{6.1}$$

Out of all input objects, the two with the smallest d_{ij} are grouped together to form a pseudo-jet, replacing the two original objects in the collection. Objects for which the smallest distance to any other object exceeds the threshold of $1/{k_t}^2$ are removed from the collection and defined as jets. This combination process is repeated until all PFlow objects become constituents of exactly one jet each. In this analysis, the anti-k_t algorithm is employed with a radius parameter of $R = 0.4$, reflecting to the expected size of jets.

The complexity of jets requires an intricate jet energy scale (JES) calibration, consisting of various stages, each designed to address specific inaccuracies in the jet modelling or detector effects [159]. The entire calibration procedure is depicted in Figure 6.3. The input to the calibration scheme are the reconstructed PFlow jets.
Initially, corrections are applied to account for the effects of pile-up. The per-event pile-up is estimated using a p_T-density function $\rho = \langle p_\text{T}/A \rangle$, which measures the pile-up contribution to the p_T of jets with a jet-area A in the y-ϕ plane. The estimated contribution is subtracted from the jet's total p_T. Subsequent corrections, depending on the average number of pile-up interactions per event, μ, and the number of reconstructed primary vertices N_PV, are applied through linear subtractions, yielding a fully pile-up corrected jet p_T of

$$p_\text{T}^\text{corr} = p_\text{T}^\text{reco} - \int \rho dA - \alpha \cdot (N_\text{PV} - 1) - \beta \cdot \mu. \tag{6.2}$$

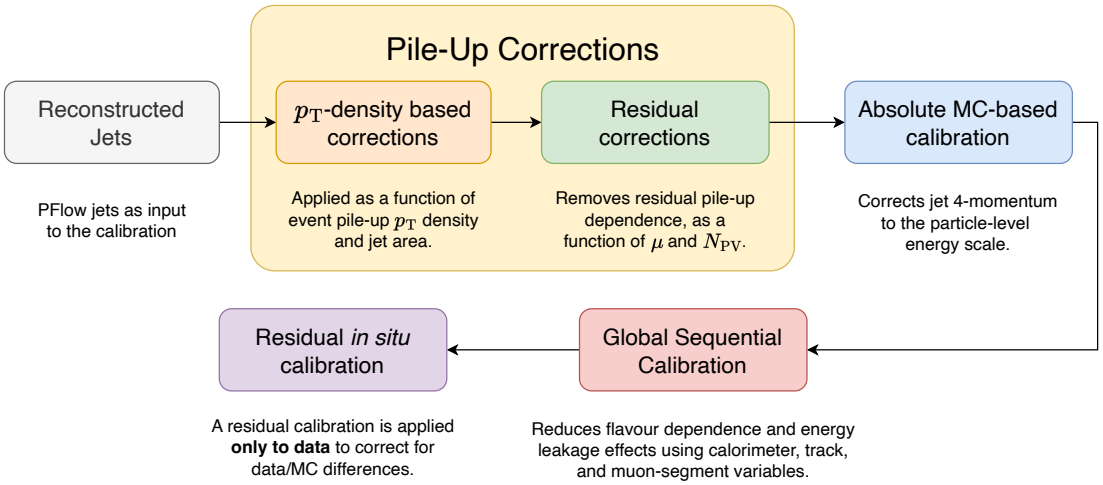

Figure 6.3.: Flow-chart depicting the different steps of the JES calibration procedure. Figure based on Ref. [159].

The coefficients α and β are determined using MC information at generator level. Subsequently, the absolute energy and the pseudorapidity of a jet are calibrated using MC simulations of dijet production. These calibrations account for inefficiencies in the calorimeter response, the energy loss in passive detector material, and the presence of jet partons outside the jet cone. Simultaneously, biases in the η reconstruction, primarily arising due to transitions between calorimeter technologies and variations in calorimeter granularity, are also corrected.
Following the four-momentum calibration, the *global sequential calibration* (GSC) is applied. Even after correcting the jet's four-momentum, the JES response may still vary depending on the flavour and energy distribution of the jet's constituents. Additionally, the jet's overall composition varies according to the type of its primary particle. Another consideration is the possibility of jets not being fully contained within the HCAL, an effect known as *punch-through*. All these effects are addressed by multiplication of correction factors, derived from MC-data comparisons across six dedicated observables, resulting in an enhanced jet modelling.
As a final step, *in-situ* calibrations are employed to handle differences in the jet response between data and MC simulations. Such differences arise from inefficiencies in modelling the detector or physical processes. They are addressed by independently measuring the jet response and applying the resulting ratio as a correction factor to the data. The jet response is measured by using already calibrated objects in well-understood regions of phase space, mainly leptons from on-shell $Z \to \ell\ell$ decays, and comparing the p_{T} of the $\ell\ell$-system to that of recoiling jets. The fraction of data and MC response, as a function of jet's p_{T} and η, is used as the final correction factor. Additional corrections

are applied to synchronise forward jets with $0.8 < |\eta| < 2.4$ to the same energy scale as central jets with $|\eta| < 0.8$. The energy scale of central jets is assumed to be well modelled, while for forward jets correction factors are derived using dijet events in a process called η intercalibration. The p_{T} of both jets is expected to sum up to zero in any given event. Deviations from this for two jets with different η are used to extract efficiencies measuring the relative differences in the JES for individual $|\eta|$ regions.

The jet energy resolution (JER) is impacted by three distinct sources: the noise (N), induced by pile-up and front-end electronics, a stochastic component (S), associated with the statistical fluctuation of energy deposited in the calorimeter, and a constant term (C), related to energy deposition in passive detector material [159]. The total JER can be expressed by

$$\frac{\sigma_{p\mathrm{T}}}{p_{\mathrm{T}}} = \frac{N}{p_{\mathrm{T}}} \oplus \frac{S}{\sqrt{p_{\mathrm{T}}}} \oplus C. \tag{6.3}$$

Calibration of the noise term is conducted through random-cone measurements in zero-bias data events, which are events recorded without any trigger requirement. The remaining components of the JER are calibrated using dijet events. In phase space regions, where the JER is higher in data than in MC, jets in the simulations are smeared to align their resolution with that in data. Conversely, the data exhibit a finer resolution compared to the MC, no smearing is applied, preserving the superior data quality.

In this analysis, jet candidates are accepted if they meet the criteria of $p_{\mathrm{T}} > 20\,\mathrm{GeV}$ and $|\eta| < 2.5$. Low-p_{T} jets with $p_{\mathrm{T}} < 60\,\mathrm{GeV}$ with $|\eta| < 2.4$ must pass additional identification criteria established by the *Jet Vertex Tagger* (JVT) [160], which aims to reduce contamination from pile-up jets. The JVT employs a multivariate technique, leveraging information on the tracks constituting a given jet. The working point utilised in this analysis provides an efficiency of 92% for non-pile-up jets and succeeds in rejecting 98% of pile-up events. This working point is calibrated using the tag-and-probe method with $Z \to \mu^+\mu^-$ events, resulting in SFs that are applied to MC simulations.

6.5. Flavour Tagging

Jets originating from a b quark exhibit several distinguishing properties compared to other jets. Specifically, b quarks hadronise into b hadrons, which have a longer lifetime than other types of hadrons, at the level of one picosecond. Consequently, these hadrons can travel distances up to several millimetres before decaying, leading to a secondary vertex where the jet constituents originate. As a result, many tracks within the jet have a large impact parameter relative to the primary vertex. Additionally, the high mass of b hadrons typically allows for higher transverse momenta of jet constituents relative to

the jet's central axis, often leading to larger jet cones. The process of identifying such jets, referred to as b-jets, is known as *b-tagging*.

There are several dedicated b-tagging algorithms available. The analysis in this dissertation employs the *DL1r* tagger [161]. This tagger utilises a *deep feed-forward neural network* (DNN) to combine the inputs from several low-level b-tagging algorithms. Each of these low-level algorithms is specifically designed to exploit distinct properties of b-jets. The *IP2D* and *IP3DoubleD* algorithms [162] use information on the impact parameters of tracks associated with charged particles within the jet. Extending this approach, the *RNNIP* algorithm [163] uses an RNN trained on general track properties, thus accounting for correlations between these properties, particularly the track impact parameters. In contrast to these approaches, the *SV1* algorithm [164] aims to reconstruct the secondary vertex within a jet, indicative of the b-hadron decay point. This reconstruction is performed using a technique akin to the vertex finding algorithm used for hard scattering events, but specifically applied to the constituents of the jet. Finally, the *JetFitter* algorithm [165] expands on SV1's foundation by attempting to reconstruct the complete decay chain of hadrons containing b or c quarks within the jet.

The DL1r DNN integrates the outputs of all five algorithms to produce three output values: p_b, p_c and p_l. These correspond to the probabilities of a given jet to have been initiated by a b quark, a c quark or any other parton, respectively. The training is performed using simulated events of $t\bar{t}$ and Z' production. The probabilities are combined into a single b-tagging discriminant, calculated as

$$D_{\mathrm{DL1r}} = \ln\left(\frac{p_b}{f_c p_c + (1 - f_c) p_l}\right). \tag{6.4}$$

Here, $f_c = 0.018$ is the effective fraction of c-jets in the background sample used for the training. A working point is defined for this discriminant such that the b-tagging efficiency corresponds to 70% in simulated $t\bar{t}$ events. To compensate performance differences between MC simulations and data, p_{T}-dependent SFs are determined based on $t\bar{t}$ events, which were simulated using Pythia 8 for showering. This introduces a parton-shower dependence to these SFs, which has to be accounted for when using the DL1r algorithm with other parton shower algorithms. Therefore, MC-to-MC SFs are applied to samples simulated with Sherpa or Herwig.

6.6. Overlap Removal

The outlined object definitions can result in certain ambiguities. For instance, a single track or energy deposit might be associated with multiple object candidates. To address this, a dedicated overlap-removal procedure is employed, primarily depending

on $\Delta R_y = \sqrt{\Delta y^2 + \Delta\phi^2}$ as a measure for angular separation. Object candidates are removed from the event successively in the following order:

- Electrons if they share an ID track with a muon.
- Jets within $\Delta R_y < 0.2$ of an electron.
- Electrons within $\Delta R_y < 0.4$ of a remaining jet.
- Jets within $\Delta R_y < 0.2$ of a muon if they possess two or fewer associated tracks.
- Muons within $\Delta R_y < \min(0.4, 0.04 + 10\,\text{GeV}/p_\text{T}(\mu))$ of a remaining jet.

6.7. Missing Transverse Momentum

As outlined in Section 3.2, neutrinos (and hypothetical, weakly interacting BSM particles) are not directly detectable by the ATLAS detector. However, momentum conservation in conjunction with the fact that protons possess no momentum in the transverse plane, ensures that the transverse momenta of particles in an event must sum to zero. Leveraging this fact allows to estimate the combined transverse momenta of all invisible particles produced in the event. This is achieved by calculating the vectorial sum of all calibrated object candidates in an event. It also includes soft tracks, which are defined by not being matched to any reconstructed object, if they meet the criteria $p_\text{T} > 500\,\text{MeV}$ and $|z_0 \sin\theta| < 2\,\text{mm}$ [166]. The negative of the resulting vector is called the *missing transverse momentum* $\vec{p}_\text{T}^{\,\text{miss}}$. Its magnitude is denoted by E_T^miss.

7. Background Estimation

MC simulations typically offer a reliable baseline estimate for the majority of processes relevant to ATLAS analyses. However, for various reasons certain processes considered in this analysis are not optimally modelled by simulations and require adjustments through specialised background estimation methods. Section 7.1 defines so-called non-prompt leptons and explains the methods used to estimate events containing them. Section 7.2 discusses three SM processes whose modelling requires particular attention in this analysis.

7.1. Non-prompt Lepton Background

So-called *prompt* leptons arise from the decay of either a W, Z or H boson, from a τ-lepton decay, or from an on-shell quarkonium[1] decay in the primary hard-scattering event. Conversely, *non-prompt leptons* enter the considered final states through secondary processes. Most of them are modelled using the template method described in Section 7.1.1. Particular focus is given to the largest of these templates, which describe electrons and muons from the decay of a b hadron, in Section 7.1.2. A distinct technique, discussed in Section 7.1.3, models the background from electrons reconstructed with an inverted charge.

7.1.1. Definition of Non-prompt Lepton Templates

This analysis categorises non-prompt leptons based on MC information at the generator level, hereafter referred to as *truth* information. Multiple categories are defined to encompass all potential processes resulting in non-prompt leptons. One of the most important categories consists of leptons produced in a hadronic jet. Such leptons may be identified as isolated outside of their originating jet cone or enter the phase space because of their source jet failing the jet identification criteria. The category of leptons produced in jets is subdivided based on the initiating particle of the jet. Other origin-categories include electrons reconstructed with an inverted charge, electrons from the decay of a

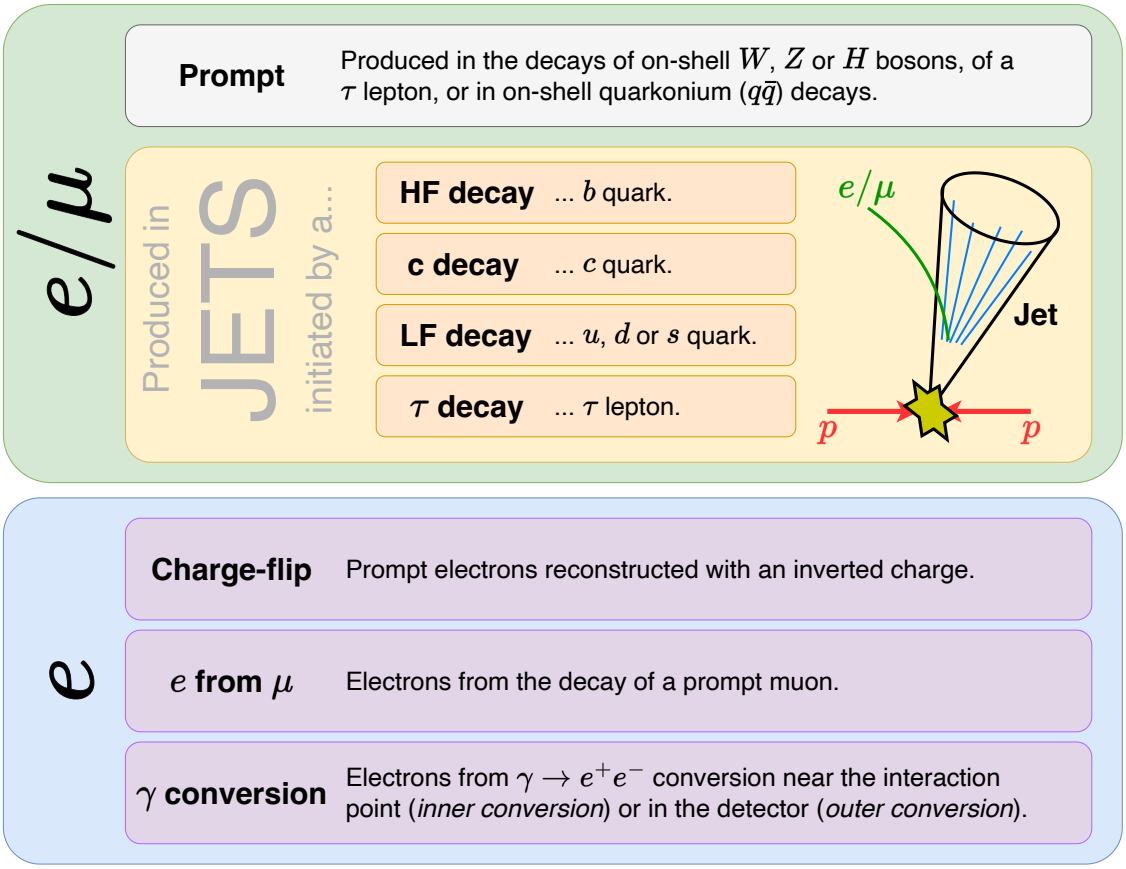

Figure 7.1.: The truth-origin classes defined for leptons in this analysis, separated into classes valid for both electrons and muons, and into classes only defined for electrons. The abbreviations HF and LF refer to heavy flavour and light flavour. An illustration of a lepton produced in a hadronic jets is shown next to the corresponding categories.

prompt muon, and electrons from photon conversion. All specified truth-origin categories are summarised in Figure 7.1.

None of these non-prompt processes can be modelled individually. Instead, non-prompt leptons may arise in any SM process. To account for this, the nominal MC samples introduced in Chapter 5 are reorganised according to the MC truth information pertaining to the leptons in a given event. Templates modelling SM processes are redefined to exclusively include events comprised solely of prompt leptons. Events containing at least one non-prompt lepton are reorganised into six distinct templates. As the background from charge-flip electrons is estimated using a purely data-driven method, detailed in Section 7.1.3, MC events containing such leptons are removed from the

[1]The term *quarkonium* refers to mesons consisting of a quark-antiquark pair with identical flavour.

analysis. In summary, the following six non-prompt templates are defined, based on the aforementioned truth-origin categories, with kinematically similar processes merged into one template:

HF-decay e	Events with exactly one electron from HF decay.
HF-decay μ	Events with exactly one muon from HF decay.
c/LF-decay	Events with exactly one lepton from the decay of an up, down, strange or charm quark.
γ-conv.	Events with exactly one electron from a photon conversion.
μ/τ-decay	Events with exactly one electron from the decay of a prompt muon or hadronic tau.
Residual non-prompt	Events with exactly one lepton that cannot be categorised due to lacking MC information or more than one non-prompt lepton.

The scarce occurrence of events with more than one non-prompt lepton within the phase space of the analysis allows for their merger into the Residual non-prompt template. HF-decay leptons constitute the largest non-prompt lepton background in this analysis. They are divided into two templates to take into account potential variances in non-prompt rejection for different lepton flavours. The estimation of this background is explained in Section 7.1.2. All other non-prompt lepton templates are assumed to be modelled adequately by their respective templates, within the assigned conservative normalisation uncertainties of 50%.

Not all SM processes contribute equally to the templates defined above. The processes leading to the production of non-prompt leptons are very rare. Thus only processes with relatively high rates significantly contribute to these templates. Almost all non-prompt lepton events originate from the production of a $t\bar{t}$ pair, a single top quark or a W/Z boson in association with hadronic jets. The relative importance of each of these processes differs by non-prompt template and final state. Figure 7.2 illustrates their relative contributions to the HF-decay e and HF-decay μ templates in the 2ℓSS and 3ℓ final state. The composition of both HF-decay templates in terms of SM processes is almost identical. The predominance of $t\bar{t}$ in is anticipated, as the process produces two b-jets. Differences between the two final states arise because the HF-decay templates require one prompt lepton in the 2ℓSS and two in the 3ℓ final state. The production of a single top quark assumes greater significance in the 2ℓSS final state, as top-quark production in the t-channel and the s-channel typically only yields one prompt lepton, whereas tW production may produce one or two prompt leptons. Similarly, W+jets

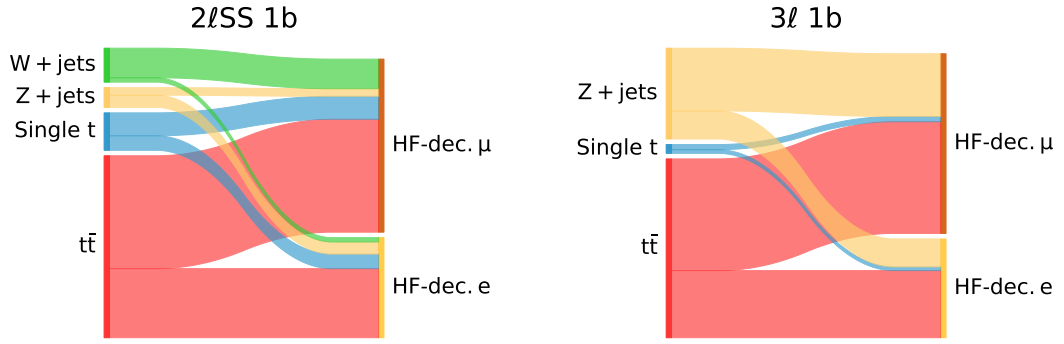

(a) (b)

Figure 7.2.: The relative share of the four main SM processes contributing to the HF-decay e and HF-decay μ templates in (a) the 2ℓSS and (b) the 3ℓ final state. A loose selection is applied, requiring exactly one b-tagged jet in addition to the respective lepton configuration.

production exclusively contributes to the 2ℓSS final state, while Z+jets contributes to both.

7.1.2. Heavy-Flavour Decay Leptons

The HF-decay lepton background is estimated using a semi-data-driven approach, employing the *template fit method* (TFM) [167]. The method posits that kinematic distributions are accurately modelled by MC simulations, necessitating only adjustments to their normalisation. Corrections are implemented through free-floating normalisation factors in the final maximum-likelihood fit to data. A fundamental premise of the TFM is the applicability of these factors independent of the event topology. Consequently, identical normalisation corrections employed in both the 2ℓSS and 3ℓ final state. To constrain the respective normalisation factors, dedicated kinematic regions enriched in either HF-decay e or HF-decay μ events are defined in Chapter 8.

7.1.3. Q-misID Electrons

The charge of an electron candidate is reconstructed using the curvature of its ID track. For electrons with exceptionally high transverse momentum, the curvature is minimal, which can lead to inaccuracies in the charge reconstruction, resulting in a charge-flip electron. Additionally, an incident electron may undergo bremsstrahlung, losing a significant fraction of its energy to photon emission within the detector material. This photon might then produce an electron-positron pair, leading to scenarios where only the

oppositely charged particle is detected, effectively mimicking a charge-flip event. The likelihood of such a *trident* interaction increases with the pseudorapidity, attributable to the larger amount of detector material encountered. The combination of the charge-flip and trident processes is referred to as Q-misID in this dissertation. It only plays a role in the 2ℓSS final state.

The Q-misID background is estimated with a fully data-driven method. This approach relies on the comparison of data events around the Z-boson mass peak in $Z \to ee$ decays between same-charge (SS) and opposite-charge (OS) final states. The objective is to derive p_{T}-$|\eta|$-dependent Q-misID efficiencies, allowing for the construction of a dedicated template for the process using OS data events. It is divided into three distinct phases, each of which is outlined in the following.

Position and Width of the Mass Peak

The initial phase involves estimating the position and width of the Z-boson mass peak. The considered events are selected by requiring zero b-tagged jets and exactly two leptons, adhering to the respective SS or OS requirement. This selection ensures the orthogonality of the efficiency-estimation with the remainder of the analysis, where at least one b-tagged jet is required. The position and width are estimated separately in the SS and OS channels due to the influence of trident electrons, which possess only a fraction of the original particle's energy, causing the mass peak to broaden and slightly shift towards lower values in SS events.

The estimation is performed using a χ^2 fit of a chosen distribution to the mass peak. Potential functions to model the peak's shape include the Breit-Wigner distribution [168], the Gaussian distribution [169] and the crystal-ball function [170]. Consistent with other ATLAS analyses using a similar Q-misID estimation technique [167, 171], a Breit-Wigner distribution is selected, described by

$$f_{\mathrm{BW}}(x|N,m,\sigma) = \frac{N}{\left(x^2 - m^2\right)^2 + \sigma^2 m^2}. \tag{7.1}$$

Here, m represents the peak's centre, σ its width and N a normalisation factor to adjust the distribution's scale. The results of the SS and the OS fit are depicted in Figure 7.3, with the fit parameters listed in Table 7.1. As expected, the SS mass peak is lower and broader compared to its OS counterpart. At the same time, the OS results agree well with previous experimental observations [14]. However, the overall reduced χ^2 values for both fits deviate significantly from the ideal value of 1, indicating that the Breit-Wigner function does not optimally describe the entire distribution. However, when focusing solely on the reduced χ^2 in the vicinity of the peak $m \pm \sigma$, based on the fitted m and σ, there is a notable improvement. Although still not optimal, this improvement suggests that the function describes the mass peak sufficiently well, enabling a reasonable estimation of its centre and width. The subsequent paragraph explains that this estimation is the only purpose of the Breit-Wigner fit.

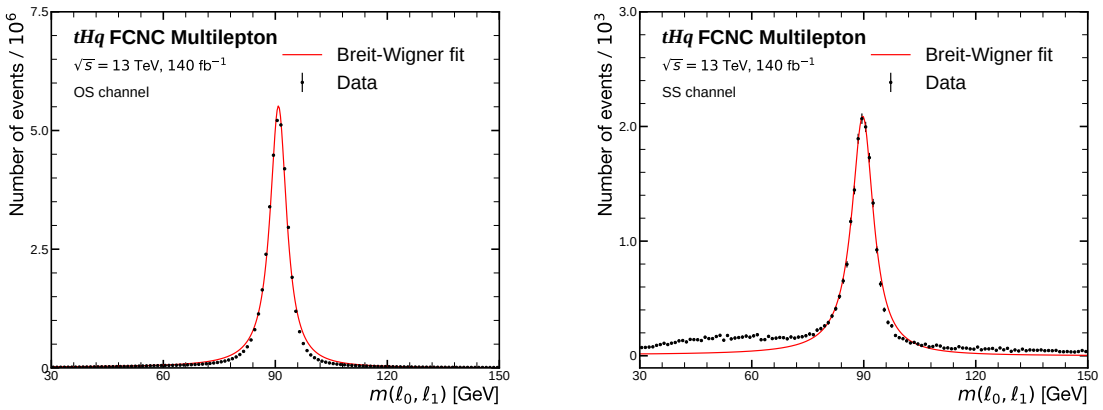

(a) OS channel (b) SS channel

Figure 7.3.: The results of the fit of a Breit-Wigner distribution to (a) OS and (b) SS data events around the Z-boson mass peak.

Table 7.1.: The resulting parameters of the Breit-Wigner fits to the Z-boson mass distribution. The reduced χ^2 values are shown once for the full spectrum of the fit and once for the on-peak region around $m \pm \sigma$.

Channel	m [GeV]	σ [GeV]	N	χ^2/n_{df}	χ^2/n_{df} $(m \pm \sigma)$
OS	90.84(3)	5.4(3)	$1.31(5) \cdot 10^{12}$	$5.1 \cdot 10^8/147$	$1.1 \cdot 10^4/8$
SS	89.7(1)	7.03(4)	$8.3(6) \cdot 10^8$	6045/147	134/11

Calculation of p_T-$|\eta|$-dependent Efficiencies

The Z-boson mass peak position and width in both channels serve as the baseline for estimating Q-misID efficiencies. The phase space is now divided into two distinct sections: the *Z-window*, containing events with $m(ee) \in [m - 4\sigma, m + 4\sigma]$, and the *side-bands*, containing events with $m(ee) \in (m \pm 4\sigma, m \pm 8\sigma]$. It is assumed that events in the Z-window are a mixture of $Z \to ee$ events and uniform background, while the side-bands only contain background events. The average number of events in both side-bands is thus subtracted from the number of events in the Z-window, yielding pure $Z \to ee$ events. This procedure is conducted separately for the OS and the SS channel, using the respective values of m and σ.

As discussed in the beginning of this section, the charge-flip background is expected to increase with p_T, while the trident background increases with $|\eta|$. Consequently, the phase space is divided into 24 electron p_T-$|\eta|$ bins. The boundaries of these bins are defined as follows:

- p_{T}-bin boundaries: $\{10, 60, 90, 130, \infty\}$ GeV and
- $|\eta|$-bin boundaries: $\{0.0, 0.6, 1.1, 1.52, 1.7, 2.3, 2.5\}$.

It is assumed that for each bin a well-defined efficiency exists, reflecting the probability of an electron in this bin becoming a charge-flip or trident electron. Under this assumption, the expected number of SS events with an electron pair falling into the bins i and j is related to the observed number of OS events via

$$\hat{N}_{\mathrm{SS}}^{i,j}(\epsilon_i, \epsilon_j) = \frac{\epsilon_i(1-\epsilon_j) + (1-\epsilon_i)\epsilon_j}{1 - \big(\epsilon_i(1-\epsilon_j) + (1-\epsilon_i)\epsilon_j\big)} N_{\mathrm{OS}}^{i,j}. \tag{7.2}$$

Here, $\epsilon_{i/j}$ are the Q-misID efficiencies for bin i/j. $\hat{N}$ denotes the number of expected events, while N signifies the number of observed events. The number of observed SS events $N_{\mathrm{SS}}^{i,j}$ is expected to follow a Poisson distribution [172], with the above expectation value:

$$f\left(N_{\mathrm{SS}}^{i,j} | \hat{N}_{\mathrm{SS}}^{i,j}(\epsilon_i, \epsilon_j)\right) = \frac{\left(\hat{N}_{\mathrm{SS}}^{i,j}\right)^{N_{\mathrm{SS}}^{i,j}} \cdot e^{\hat{N}_{\mathrm{SS}}^{i,j}}}{N_{\mathrm{SS}}^{i,j}!}. \tag{7.3}$$

Given the observed numbers in the SS channel, a likelihood function of all Q-misID efficiencies ϵ can be constructed as the product over all pairs of p_{T}-$|\eta|$ bins:

$$L(\epsilon | N_{\mathrm{SS}}) = \prod_{i,j} f\left(N_{\mathrm{SS}}^{i,j} | \hat{N}_{\mathrm{SS}}^{i,j}(\epsilon_i, \epsilon_j)\right). \tag{7.4}$$

Optimal estimates of all ϵ are then obtained by minimising $-2\ln(L)$ using the MIGRAD [173] minimiser. The resulting efficiencies for all bins are visualised in Figure 7.4.

Application of Efficiencies to Data

The final step of the Q-misID estimation involves the actual template definition. This is performed individually for every kinematic region considered in the 2ℓSS final state. A corresponding OS region is defined with identical selection criteria, except for the charge-requirement for the two leptons being inverted. Data events within this OS region are weighted by the following weights, depending on their lepton configuration:

$$\begin{aligned}
\omega_{ee} &= \frac{\epsilon_i(1-\epsilon_j) + (1-\epsilon_i)\epsilon_j}{1 - \big(\epsilon_i(1-\epsilon_j) + (1-\epsilon_i)\epsilon_j\big)} && \text{for } ee\text{-events,} \\
\omega_{e\mu} &= \frac{\epsilon_i}{1-\epsilon_i} && \text{for } e\mu/\mu e\text{-events and} \\
\omega_{\mu\mu} &= 0 && \text{for } \mu\mu\text{-events.}
\end{aligned}$$

Here ϵ_i and ϵ_j refer to the Q-misID efficiencies of the p_{T}-$|\eta|$ bins into which the respective electrons of the event fall. The weighted events form the Q-misID template in the analysis.

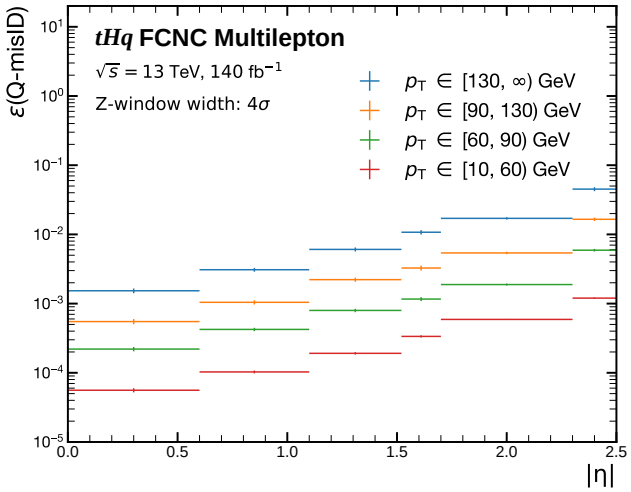

Figure 7.4.: The Q-misID efficiencies for all p_{T}-$|\eta|$ bins obtained from a maximum likelihood fit to the observed number of SS and OS events using a Z-window width of 4σ.

7.2. Prompt Lepton Background Processes

Most prompt background processes are adequately modelled by the MC simulations. In some cases, however, the simulations are known to produce imprecise results in the considered final states, while in others the predicted cross sections do not match the measured results. The affected processes include $t\bar{t}W$ and $t\bar{t}Z$ production, discussed in Section 7.2.1, and VV production in association with an additional b-jet, examined in Section 7.2.2.

7.2.1. $\mathrm{t\bar{t}V}$ Production

The production of a $t\bar{t}$ pair in association with an additional W or Z boson is one of the most significant background processes in this analysis. Recent measurements of the $t\bar{t}W$ cross-section in multilepton final states have observed a deviation from the expected value by approximately 50% [171]. Considering the uncertainties on the measurement, this discrepancy corresponds to a tension of 1.4σ. The process is particularly sensitive to higher-oder QCD and electroweak corrections, suggesting that the observed deviation may be attributable to MC modelling effects. To accommodate this, $t\bar{t}W$ is assigned a free-floating normalisation factor in the maximum-likelihood fit.

For $t\bar{t}Z$, cross-section measurements in multilepton final states are in full agreement with SM predictions, as shown in Ref. [174]. However, the referenced analysis focuses on

kinematic regions characterised by a high number of jets, with signal-enriched regions in 2ℓOS final states requiring at least six jets, while those in 3ℓ final states require a minimum of three. In contrast, the tHq FCNC signal process under investigation in this analysis is expected to exhibit a much lower jet multiplicity. It remains uncertain whether $t\bar{t}Z$ is as accurately modelled in regimes with low N_{jets} as it is in those with high N_{jets}. In light of these considerations, $t\bar{t}Z$ is also assigned a free-floating normalisation factor.

7.2.2. VV + b-jet Production

The production of two vector bosons is generally well-understood. However, the expected event signature of the tHq FCNC processes invariably includes at least one b-jet. This presents a challenge in the VV modelling, because the samples used in this analysis are generated without an additional b-jet at matrix-element level. In prompt VV events, both vector bosons must decay leptonically to contribute to 2ℓSS or 3ℓ final states. Consequently, any additional b-jet must originate from the parton shower. Given that parton shower generators are not optimised for modelling the precise kinematics of the b-jet in $VV + b$-jet production, certain inaccuracies are expected.

To address this, the prompt VV template is subdivided into multiple sub-templates based on two criteria. Firstly, the classification is based on the number of final-state leptons at matrix-element level, encompassing three scenarios: two leptons of the same charge (2ℓSS), three leptons (3ℓ) and four leptons (4ℓ). Secondly, events are categorised according to the truth-level flavour of jets present in a given event. Two distinct categories are established:

- b/c: $N_{b\text{-jets}}^{(\text{truth})} \geq 1$ or $N_{c\text{-jets}}^{(\text{truth})} \geq 1$,
- l/τ: $N_{b\text{-jets}}^{(\text{truth})} = 0$ and $N_{c\text{-jets}}^{(\text{truth})} = 0$.

Here, l stands for *light* jets, referring to jets initiated by any parton other than b or c quarks. The rationale for combining events with b and c-jets is attributed to the higher mis-tagging rate for c-jets compared to l-jets or τ-jets. Comparative studies on the differences between VV events with only b or only c-jets show very similar shapes across various kinematic distributions. Separating the VV process along these two dimensions yields the following six templates:

- $VV2\ell\text{SS} + b/c$, $VV2\ell\text{SS} + l/\tau$,
- $VV3\ell + b/c$, $VV3\ell + l/\tau$,
- $VV4\ell + b/c$ and $VV4\ell + l/\tau$.

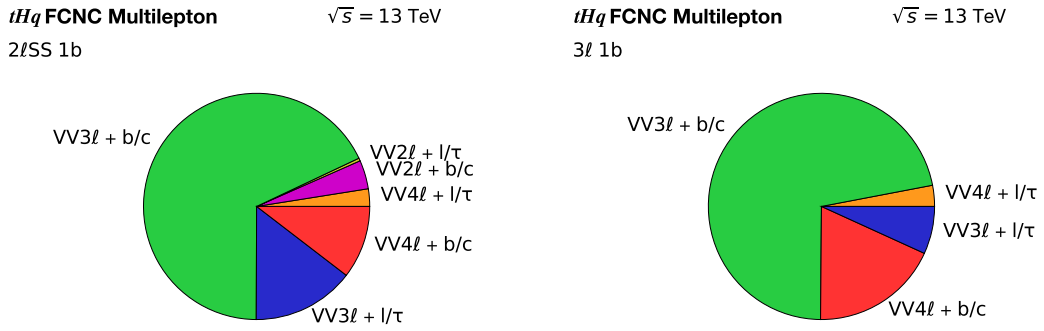

(a) 2ℓSS (b) 3ℓ

Figure 7.5.: Contribution of each VV sub-template to events with exactly one b-tagged jet in (a) the 2ℓSS and (b) the 3ℓ final state.

To determine the treatment of each template in the maximum-likelihood fit, their relative contributions to the 2ℓSS and 3ℓ final states are examined. For this purpose, events with exactly one b-tagged jet are analysed in both final states. The pie-charts in Figure 7.5 illustrate that in both final states the $VV3\ell + b/c$ template is the most significant contributor. The 3ℓ templates are likely dominated by on-shell WZ production, alongside some off-shell contributions. Their predominance is expected in the 3ℓ final state. The fact that they also contribute most significantly in the 2ℓSS final state owes to the non-detection of one of the three leptons. It also shows that $W^{\pm}W^{\pm}$ production, which is the primary process in the 2ℓSS templates, is considerably less prevalent compared to WZ production.

Given the observed dominance of the $VV3\ell + b/c$ template, it is assigned a free-floating normalisation factor in the maximum-likelihood fit. The less significant VV templates, akin to the minor non-prompt templates, are assigned uncorrelated 50% normalisation uncertainties.

8. Event Selection

This chapter presents the event selection criteria employed in this analysis. Primary quality criteria for the considered final state objects have been discussed in Chapter 6. Additionally, to reject leptons close to the single-lepton trigger p_{T}-threshold of 26 GeV, the leading-p_{T} lepton must surpass a p_{T} of 28 GeV. Moreover, this analysis exclusively considers events with at least one b-tagged jet. This requirement is driven by the event signature of the signal process, which always contains at least on b-jet from the top-quark decay. All preselection criteria are summarised in Table 8.1, along with the motivation for each criterion. The resulting phase space is further partitioned into various kinematic regions. Some of these are enriched in signal, to increase the overall sensitivity of the analysis, while others are designed to predominantly contain events from certain background processes to ensure correct modelling. The definition of all 2ℓSS regions is presented in Section 8.1, and that of 3ℓ regions in Section 8.2. Throughout this chapter and the remainder of this dissertation, leptons are labelled according to their p_{T}: the leading-p_{T} lepton is referred to as ℓ_0, followed by ℓ_1, with the third-leading-p_{T} lepton in the 3ℓ channel being denoted ℓ_2.

8.1. Kinematic Regions in the 2ℓSS Final State

The 2ℓSS final state is characterised by the presence of two leptons with an identical charge. Two signal regions are defined within this final state, both of which are

Table 8.1.: Overview of the preselection criteria for this analysis. For every cut the corresponding motivation is presented. ℓ_0 refers to the leading-p_{T} lepton of a given event.

Preselection		Motivation
$p_{\mathrm{T}}(\mathrm{jet})$	$\geq 20\,\mathrm{GeV}$	Improved quality of reconstructed objects.
$p_{\mathrm{T}}(\ell_i)$	$\geq 10\,\mathrm{GeV}$	
$p_{\mathrm{T}}(\ell_0)$	$\geq 28\,\mathrm{GeV}$	Rejection of events near trigger threshold.
N_{jets}	≥ 1	Event signature of the signal process.
$N_{b\text{-tags}}$	≥ 1	

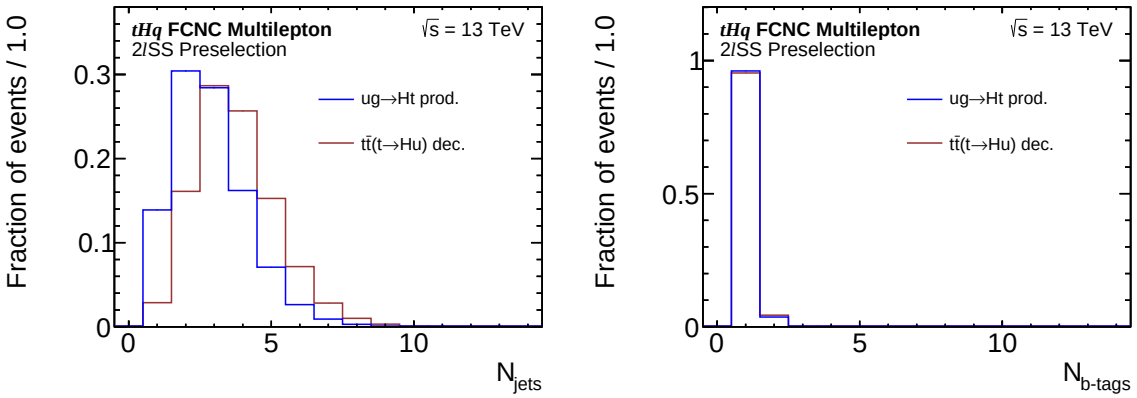

(a) N_{jets} (b) $N_{b\text{-tags}}$

Figure 8.1.: Distribution of (a) N_{jets} and (b) $N_{b\text{-tags}}$ for tHu signal events passing the 2ℓSS preselection.

motivated in Section 8.1.1. Furthermore, three control regions are defined to constrain the normalisation of the HF-decay e, HF-decay μ, $t\bar{t}W$ and $t\bar{t}Z$ templates, all detailed in Section 8.1.2. The pre-fit event yields of all 2ℓSS regions defined in the following can be found in Table B.1.

8.1.1. Signal Regions

As introduced in Chapter 2, the tHq signal coupling is considered in two distinct contexts: the production of a top quark and a Higgs boson via the new vertex and the production of a $t\bar{t}$ pair, with one of the quarks undergoing decay into a Higgs boson and either an up quark or a charm quark. At leading order, the event signature of both processes is almost identical. The sole difference lies in the presence of an additional u- or c-jet in the decay channel. In both scenarios, one top quark decays via the SM decay mode $t \to Wb$, leading to the expectation of exactly one b-tagged jet. The distributions of N_{jets} and $N_{b\text{-tags}}$ for preselected 2ℓSS tHu signal events are depicted in Figure 8.1, confirming the expectation of a peak at $N_{b\text{-tags}} = 1$ and a tendency for decay events towards higher N_{jets}. To optimally leverage the additional events provided by the production channel, two signal regions are defined. One is required to fulfil $N_{\text{jets}} \geq 4$, thereby being enriched in events from the decay channel and denoted as SR2ℓ Dec. Conversely, the second signal region, characterised by $N_{\text{jets}} \leq 3$, encompasses in a higher fraction of production signal events and is thus labelled SR2ℓ Prod.

As explained in the previous chapter, Q-misID events are responsible for a significant fraction of the total background contribution in the 2ℓSS final state. The region around

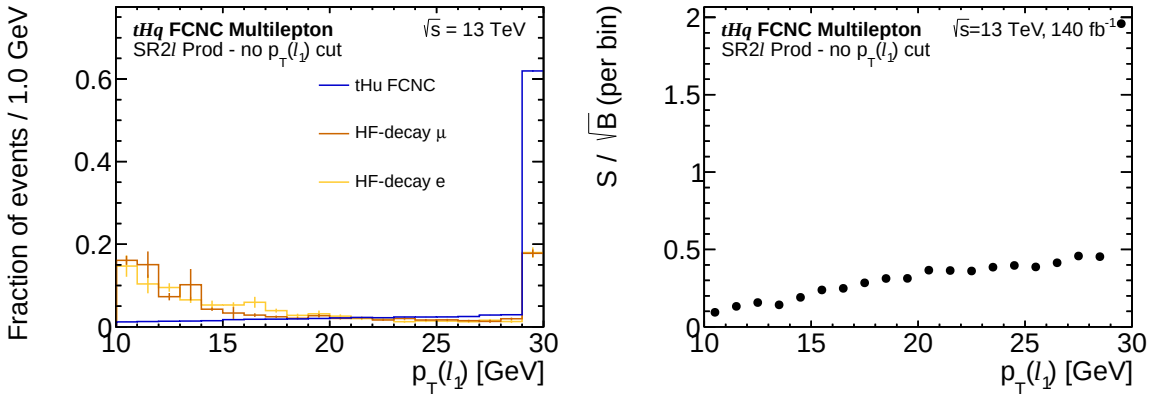

(a) (b)

Figure 8.2.: The $p_{\mathrm{T}}(\ell_1)$ variable in the SR2ℓ Prod prior to the application of any cuts on $p_{\mathrm{T}}(\ell_1)$ in the range of 10 to 30 GeV. (a) The shape of the distribution for the tHu signal and the HF-decay processes and (b) the $\mathrm{S}/\sqrt{\mathrm{B}}$ fraction for the sum of all (prompt and non-prompt) background processes are shown. For the calculation, the signal was normalised to $\mathcal{B}(t \to Hq) = 1.0$. The last bin includes all events exceeding the upper boundary of the histogram.

the Z-boson mass peak is particularly contaminated due to Z+jets production with one b-tagged jet. Consequently, events containing two electrons with an invariant mass within a window of 10 GeV around $m_Z = 91.19$ GeV are excluded from both regions.

HF-decay leptons comprise another dominant source of background. Because these leptons are produced in jets, they typically exhibit a lower p_{T} than prompt leptons. This is illustrated by the $p_{\mathrm{T}}(\ell_1)$ distribution shown in Figure 8.2 (a), where the prompt leptons from the tHu signal are compared to HF-decay leptons in the SR2ℓ Prod. A clear tendency for HF-decay e and HF-decay μ events towards p_{T} values near the identification threshold of 10 GeV is evident, particularly for events in the SR2ℓ Prod candidate. To select an optimal cut on $p_{\mathrm{T}}(\ell_1)$ with the aim of reducing the contribution of HF-decay leptons, a boundary condition for bins to be considered in the analysis of $\mathrm{S}/\sqrt{\mathrm{B}} > 0.25$ is chosen. Drawing from Figure 8.2 (b), a cut of $p_{\mathrm{T}}(\ell_1) \geq 16$ GeV is applied in the SR2ℓ Prod. Analogous distributions are considered for the SR2ℓ Dec (see Figure B.1), resulting in the requirement of $p_{\mathrm{T}}(\ell_1) \geq 12$ GeV in that region. Identical studies for the tHc process show that the above event selection is suitable for both signal processes. A summary of the final definition of both signal regions is presented in Table 8.2, with the process composition in both regions illustrated in the pie charts in Figure 8.3. Details on the tHc signal contributions can be found in Figure B.3.

Table 8.2.: Definition of the 2ℓSS signal regions. All requirements applied in addition to the event preselection are presented, together with a motivation for the specific cut.

	SR2ℓ Prod	SR2ℓ Dec	Motivation
N_{jets}	≤ 3	≥ 4	Separation of production and decay signal.
$N_{b\text{-tags}}$	$= 1$	$= 1$	Overall signal event signature.
$p_{\text{T}}(\ell_1)$	$\geq 12\,\text{GeV}$	$\geq 16\,\text{GeV}$	Reduction of HF-decay e/μ background.
$\lvert m(ee) - m_Z \rvert$	$> 10\,\text{GeV}$	$> 10\,\text{GeV}$	Reduction of Q-misID background.

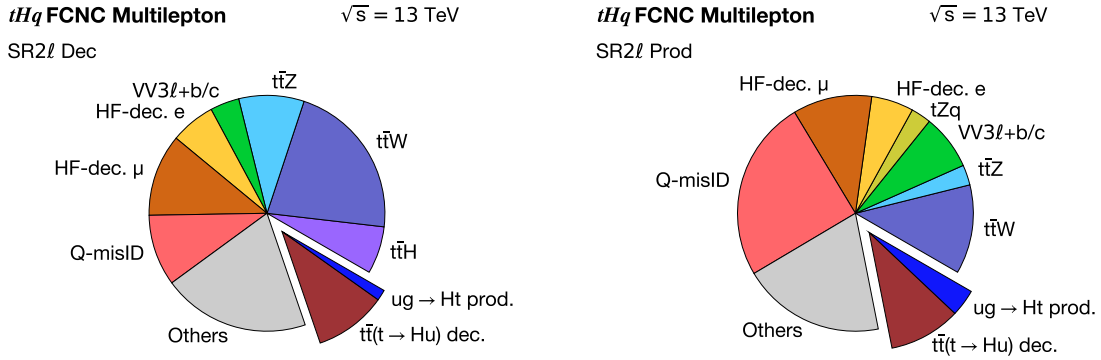

(a) (b)

Figure 8.3.: Pie charts showing the background composition in (a) the SR2ℓ Dec and (b) the SR2ℓ Prod. The tHu signal contribution for a normalisation of $\mathcal{B}(t \to Hq) = 0.1\%$ is included as well. Processes with minor contributions are collected in the *Others* category.

8.1.2. Control Regions

As outlined in Chapter 7, the normalisation of various background templates is left free-floating in the maximum-likelihood fit, allowing for adjustments based on the data. To constrain this normalisation, dedicated control regions are defined, three of them in the 2ℓSS channel of this analysis. Two specifically target the HF-decay e and HF-decay μ backgrounds, whereas the third is selected to be sensitive to $t\bar{t}W$ and $t\bar{t}Z$ production.

The previous section showed that leptons from HF decay tend towards lower p_{T} values. In defining a 2ℓSS region enriched with either HF-decay e or HF-decay μ events, selecting for subleading-p_{T} electrons or muons is an obvious choice. Additionally, specifying the leading-p_{T} lepton to be a muon significantly reduces the contribution from Q-misID events. This is attributed to an increasing difficulty to reconstruct the track curvature in the ID at higher lepton p_{T} values, thereby enhancing the charge-flip component of

Table 8.3.: Definition of the 2ℓSS control regions for the HF-decay e and HF-decay μ templates. All requirements applied in addition to the event preselection are presented, together with a motivation for the specific cut. The flavour of a given lepton is denoted by flav(ℓ).

	CR2ℓ HFe	CR2ℓ HFμ	Motivation
N_{jets}	≤ 3	≤ 3	Maximisation of HF-decay yields.
$N_{b\text{-tags}}$	≥ 1	≥ 1	
flav(ℓ_0)	μ	μ	Reduction of Q-misID contamination.
$p_{\text{T}}(\ell_1)$	$< 16\,\text{GeV}$	$< 16\,\text{GeV}$	Maximisation of HF-decay purity.
flav(ℓ_1)	e	μ	Separation of HF-decay e/ HF-decay μ.

the Q-misID background. Owing to the supplementary MS measurements, the charge-misidentification of muons is negligible, thus increasing the likelihood for a high-p_{T} muon to be prompt. The low-$p_{\text{T}}(\ell_1)$ sections that are removed from the signal regions to reduce HF-decay contributions are clearly suitable candidates for control regions. The segment removed from the SR2ℓ Prod is evaluated to be the better option, as it provides higher statistical precision, both in data and the number of simulated MC events. To further increase this precision, the requirement on the number of b-tagged jets is relaxed to $N_{b\text{-tags}} \geq 1$. The final definition of the CR2ℓ HFe and the CR2ℓ HFμ is summarised in Table 8.3. Pie charts illustrating their composition can be found in Figure B.4. The CR2ℓ HFe and CR2ℓ HFμ achieve a purity in the corresponding HF-decay processes of 57% and 61%, respectively.

In addition to non-prompt processes, the normalisation of prompt $t\bar{t}W$ and $t\bar{t}Z$ production needs to be controlled in the final fit. One notable difference from the signal signature is the presence of two top quarks in these processes. Given that both are likely to decay into a b quark and a W boson, imposing the requirement of two b-tagged jets enriches the selection in these two processes. Furthermore, the hadronic decay of the W boson from one of the top-quark decays increases the expected jet count to at least four. To reduce contamination from Q-misID events originating from $t\bar{t}$ production, which also yields two b-jets, the leading-p_{T} lepton is required to be a muon. Analogously, the subleading-p_{T} lepton must exceed a p_{T} threshold of 18 GeV to minimise contamination from HF decays. Ideally, a distinction between $t\bar{t}W$ and $t\bar{t}Z$ production would be made to separately constrain the normalisation of both processes. However, the distributions of both processes are virtually indistinguishable in the two 2ℓSS signal regions, which is due to the strong kinematic similarity between the processes, except for the resonance around the Z-boson mass peak for $t\bar{t}Z$ production, which is largely excluded in the 2ℓSS signal regions. Therefore, only one combined control region for $t\bar{t}W$ and $t\bar{t}Z$, labelled CR2ℓ $t\bar{t}V$, is defined for the 2ℓSS final state. The definition of this control region is summarised in Table 8.4, with its composition shown in Figure B.5. The region achieves a purity in $t\bar{t}W$ and $t\bar{t}Z$ of 59%.

Table 8.4.: Definition of the 2ℓSS control regions for $t\bar{t}V$ production. All requirements applied in addition to the event preselection are presented, together with a motivation for the specific cut. The flavour of a given lepton is denoted by flav(ℓ).

	CR2ℓ $t\bar{t}V$	Motivation
N_{jets}	≥ 4	Orthogonality to other control regions.
$N_{b\text{-tags}}$	$= 2$	Enrichment in $t\bar{t}V$ production.
flav(ℓ_0)	μ	Reduction of Q-misID contamination.
$p_{\text{T}}(\ell_1)$	$\geq 18\,\text{GeV}$	Reduction of HF-decay contamination.

No control region is established for the free-floating $VV3\ell{+}b/c$ template, as its potential mismodelling is expected for events containing exactly 1 b-tagged jet and relatively high lepton p_{T} values. This coincides with the defining criteria of the signal regions. However, the signal regions enter the maximum-likelihood fit with multiple bins of a discriminant distribution. The less signal-sensitive bins of this distribution can help to constrain the normalisation of this specific template.

8.1.3. Validation Region

In general, validation regions are defined to assess the applicability of the corrections derived from the maximum-likelihood fit across the entire phase space of the analysis. They do not enter the fit, as they are generally neither enriched in signal nor in any particular background process. By applying fit corrections to the MC templates within validation regions and comparing the outcomes with data, the extent to which the corrections generalise is evaluated.

Such a validation region is defined in the 2ℓSS final state as a complement to the CR2ℓ $t\bar{t}V$, requiring a minimum of four jets, exactly two of which must be b-tagged. To maintain orthogonality to the CR2ℓ $t\bar{t}V$, the leading-p_{T} lepton is required to be an electron. This region is inherently enriched in $t\bar{t}V$ and Q-misID events, thereby providing a rigorous test for the accuracy of the estimation of these background processes. Its full composition can be found in Figure B.5.

8.2. Kinematic Regions in the 3ℓ Final State

Analogous to the 2ℓSS final state, various kinematic regions are defined in the 3ℓ final state, characterised by the presence of three lepton with a total charge of $\pm 1\,e$. The

Table 8.5.: Definition of the 3ℓ signal regions. All requirements applied in addition to the event preselection are presented, together with a motivation for the specific cut.

	SR3ℓ Prod	SR3ℓ Dec	Motivation
N_{jets}	≤ 2	≥ 3	Separation of production and decay signal.
$N_{b\text{-tags}}$	$= 1$	$= 1$	Overall signal event signature.
$p_{\text{T}}(\ell_1)$	$\geq 20\,\text{GeV}$	$\geq 20\,\text{GeV}$	Reduction of HF-decay e/μ background.
$p_{\text{T}}(\ell_2)$	$\geq 16\,\text{GeV}$	$\geq 16\,\text{GeV}$	

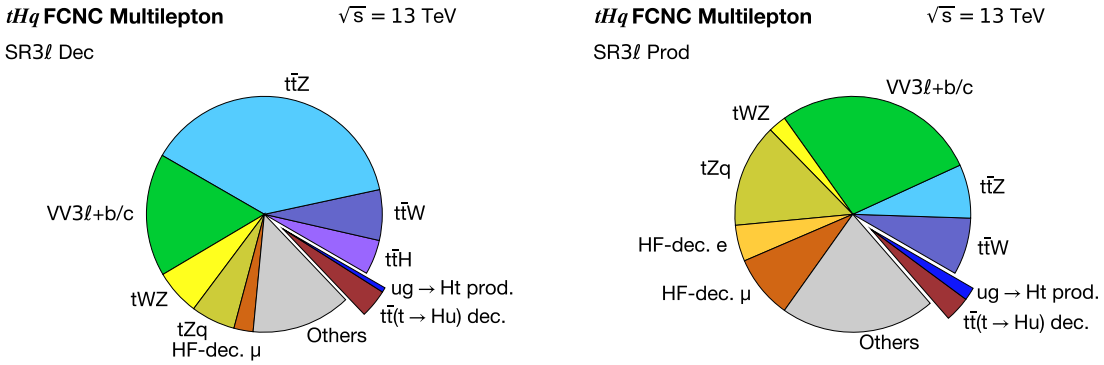

(a) (b)

Figure 8.4.: Pie charts showing the background composition in (a) the SR3ℓ Dec and (b) the SR3ℓ Prod. The tHu signal contribution for a normalisation of $\mathcal{B}(t \to Hq) = 0.1\%$ is included as well. Processes with minor contributions are collected in the *Others* category.

dissertation in Ref. [175] provides an exhaustive description on the derivation of the selection criteria. Therefore, only a brief summary of the individual region definitions is presented in this section. The pre-fit event yields of all 3ℓ regions can be found in Table B.2.

As for the 2ℓSS channel, two signal regions are defined, one enriched in the $t \to Hq$ decay signal and one with a larger fraction of $qg \to Ht$ production signal. Their selection criteria are summarised in Table 8.5. The motivation for cuts on the p_{T} of the leptons is again related to the reduction of the HF-decay background. No cuts related to the Z-boson mass peak are imposed, because a significant portion of signal events in the 3ℓ channel originate from the $H \to ZZ^*$ decay mode. The composition of both 3ℓ signal regions is shown in Figure 8.4, including the tHu signal processes. tHc signal contributions can be found in Figure B.3.

Table 8.6.: Definition of both 3ℓ HF-decay control regions. All requirements applied in addition to the event preselection are presented, together with a motivation for the specific cut. The flavour of a given lepton is denoted by flav(ℓ).

	CR3ℓ HFe	CR3ℓ HFμ	Motivation
N_{jets}	≥ 1	≥ 1	Similarity to signal regions.
$N_{b\text{-tags}}$	$= 1$	$= 1$	
$p_{\text{T}}(\ell_1)$	$\geq 20\,\text{GeV}$	$\geq 20\,\text{GeV}$	Less contamination from other backgrounds.
$p_{\text{T}}(\ell_2)$	$< 16\,\text{GeV}$	$< 16\,\text{GeV}$	Maximisation of HF-decay purity.
flav(ℓ_2)	e	μ	Separation of HF-decay e/ HF-decay μ.

Table 8.7.: Definition of both 3ℓ $t\bar{t}V$ control regions. All requirements applied in addition to the event preselection are presented, together with a motivation for the specific cut.

	CR3ℓ $t\bar{t}W$	CR3ℓ $t\bar{t}Z$	Motivation
N_{jets}	≥ 2	≥ 2	Enrichment in $t\bar{t}V$ production.
$N_{b\text{-tags}}$	$= 2$	$= 2$	
$p_{\text{T}}(\ell_1)$	$\geq 20\,\text{GeV}$	$\geq 20\,\text{GeV}$	Reduction of HF-decay contamination.
$p_{\text{T}}(\ell_2)$	$\geq 16\,\text{GeV}$	$\geq 16\,\text{GeV}$	
$\lvert m(\ell^+\ell^-) - m_Z\rvert$	$\geq 10\,\text{GeV}$	$< 10\,\text{GeV}$	Separation of $t\bar{t}W$ and $t\bar{t}Z$ production.

In the 3ℓ final state, control regions for the HF-decay e and HF-decay μ templates, as well as for the $t\bar{t}W$ and $t\bar{t}Z$ processes, are established. Similar to the 2ℓSS final state, one control region each for HF-decay e and HF-decay μ is defined. However, as the 3ℓ signal regions do not exclude the Z-boson mass peak, separate control regions for $t\bar{t}W$ and $t\bar{t}Z$ production are necessary. Consequently, a total of four control regions are defined for the 3ℓ final state. Their selection criteria are summarised in Tables 8.6 and 8.7, while pie charts depicting their composition can be found in Figures B.6 and B.7.

9. Separation of Signal and Background Processes

The signal regions, as defined in the preceding chapter, are designed to contain a high fraction of tHq FCNC signal. However, most kinematic distributions in these regions show a high resemblance among signal and background processes. The overall sensitivity to the signal process can be enhanced by constructing variables for which signal and background exhibit different behaviour, thereby generating bins with exceptionally high signal purity. To achieve this, various variables are reconstructed, some of which yield a particularly high separation power. The different reconstruction techniques employed are described in Section 9.1. The variables are subsequently combined into a single discriminant using dedicated NNs, which are presented in Section 9.2.

9.1. High-level Reconstruction Algorithms

To optimise sensitivity to the signal, several reconstruction techniques are utilised in this analysis. These techniques generally involve matching measured final-state objects to specific particles in the LO signal Feynman diagram. Kinematic parameters of these objects are then combined to reconstruct other objects in the decay chain, such as the Higgs boson or one of the top quarks. The distributions of the reconstructed variables are then expected to differ between the signal and background processes. To undertake such efforts, one must select a specific Feynman diagram on which the reconstructions are based. As discussed in Chapter 2, while three different decay modes of the Higgs boson contribute to this analysis, the $H \to WW^*$ mode is by far the most dominant in both final states. Thus, it will be the focus of all algorithms discussed in this section.

9.1.1. Combination of Basic Kinematic Parameters

One of the simplest reconstruction techniques involves defining variables obtained by performing elementary calculations using the four-momenta of final-state objects as input. Such variables include

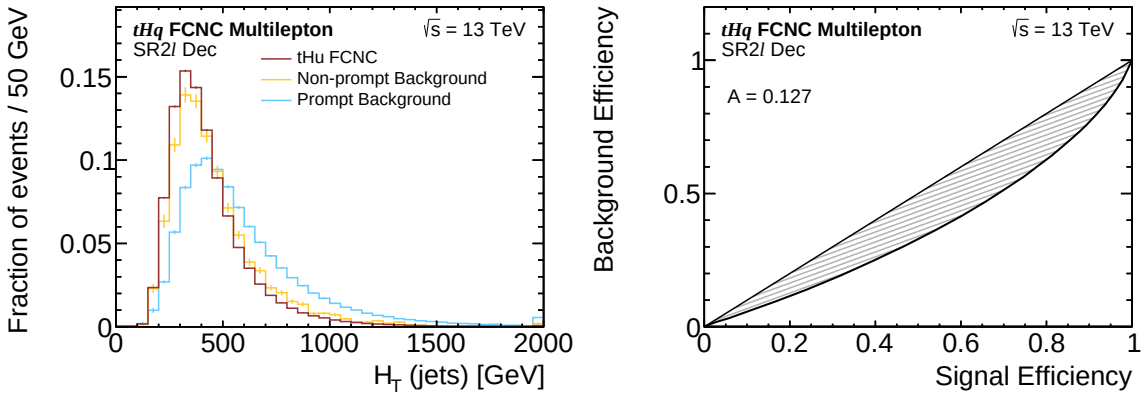

(a) (b)

Figure 9.1.: (a) The distribution of the H_T(jets) variable in the SR2ℓ Dec, shown for the tHu signal process and the combination of all prompt and non-prompt background processes, together with (b) the corresponding ROC curve. The area between the ROC curve and the diagonal is labelled A.

- the invariant mass $m(a,b) = \sqrt{(E_a + E_b)^2 - (\vec{p}_a + \vec{p}_b)^2}$ of two objects,
- the angular separation $\Delta R(a,b)$ of two objects, and
- the scalar p_{T}-sum H_T of a specified set of objects.

The invariant mass and angular separation are calculated for every possible pairing of final-state objects in each region. The scalar p_{T}-sum is computed for all visible objects, as well as for all jets and all leptons, individually. Figure 9.1 exemplarily illustrates the distribution of the scalar p_{T}-sum of all jets in the SR2ℓ Dec, separately for the tHu signal and the background. Additionally, the signal efficiency is depicted as a function of the background efficiency for events from the very left edge of the distribution to a given bin. The resulting *Receiver Operating Characteristic* (ROC) curve serves as a measure of the variable's separation power between signal and background. The greater the area between the curve and the diagonal line, the more effective the variable is at separating signal from background. The area under the ROC curve for the H_T(jets) variable is $A = 0.127$ and confirms its substantial separation power.

9.1.2. Recursive Jigsaw Reconstruction

The Recursive Jigsaw Reconstruction (RJR) technique [176] represents a more complex reconstruction effort than the mere calculation of kinematic parameters. It aims to

Table 9.1.: RJR particles and the regions in which their four-momenta are reconstructed. The specific RJR rules used for their reconstruction differ among individual regions.

Particle	SR2ℓ Prod	SR2ℓ Dec	SR3ℓ Prod	SR3ℓ Dec
t_{SM}	✓	✓	✓	✓
W_t	✓	✓	✓	✓
t_{FCNC}	✗	✓	✗	✓
H	✓	✓	✓	✓
W_{had}	✓	✓	✗	✗

reconstruct the four-vectors of the most significant particles in the signal Feynman diagram. This is accomplished by assigning an observed final-state object to each final-state particle expected from the Feynman diagram. The assignment is performed through a set of combinatoric rules, which are combined into a single likelihood discriminant. These rules are formulated based on the principle that the four-momentum is conserved at each vertex in the Feynman diagram, and that its absolute value matches the particle mass of on-shell particles. Every possible final-state objects is mapped to each matching particle in the Feynman diagram (i.e. leptons to leptons, jets to quarks, etc.). Only the combination with the highest likelihood is selected for the final reconstruction of particles in the decay chain.

Each signal region is tailored to either the decay or the production signal in the 2ℓSS or the 3ℓ final state. Consequently, the signal Feynman diagram varies slightly in each case, leading to distinct RJR rules for every region. They are exemplified for the SR2ℓ Dec in Figure 9.2. The Feynman diagram allows for the reconstruction of five particles:

t_{SM}	The top quark decaying via the SM decay mode $t_{\mathrm{SM}} \to Wb$.
W_t	The W boson from the decay of t_{SM}.
t_{FCNC}	The top quark decaying via the FCNC decay mode $t_{\mathrm{FCNC}} \to Hq$
H	The Higgs boson from the FCNC decay.
W_{had}	The W hadronically decaying W boson from the $H \to WW^*$ decay.

In other signal regions, only subsets of the above particles are reconstructed. Table 9.1 provides a summary of the particles reconstructed in each specific region. The combinatoric rules applied to derive the four-momenta of individual particles differ slightly per region, owing to differences in the considered signal Feynman diagram.

The four-momenta derived through RJR are employed to construct fundamental kinematic parameters, such as the invariant mass or the angular separation between particles. The separation power of these RJR-based parameters is, in some cases, significantly greater than that achieved using only final-state objects. Figure 9.3 exemplarily shows

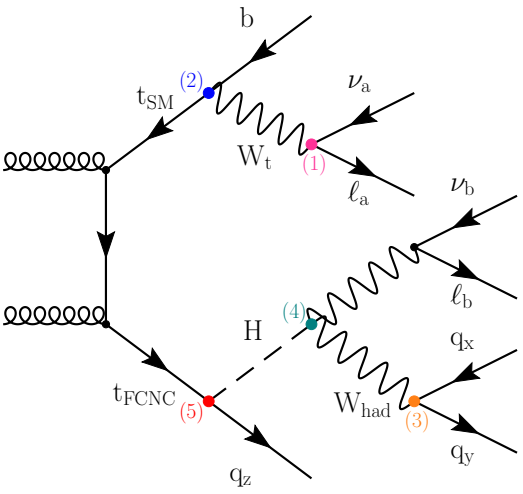

$$\vec{p}_{\mathrm{T}}^{\,\mathrm{miss}} = \vec{p}_{\mathrm{T}}(\nu_a) + \vec{p}_{\mathrm{T}}(\nu_b)$$

(1) $p(\ell_a) + p(\nu_a) = p(W_t)$, $\quad m(W_t) = m_W$

(2) $p(t_{\mathrm{SM}}) = p(b) + p(W_t)$, $\quad m(t_{\mathrm{SM}}) = m_t$

(3) $p(W_{\mathrm{had}}) = p(q_x) + p(q_y)$

(4) $p(H) = p(W_{\mathrm{had}}) + p(\ell_b) + p(\nu_b)$, $\quad m(H) = m_H$

(5) $p(t_{\mathrm{FCNC}}) = p(H) + p(q_z)$, $\quad m(t_{\mathrm{FCNC}}) = m_t$

Figure 9.2.: The combinatoric RJR rules applied in the SR2ℓ Dec, together with the vertex in the signal Feynman diagram motivating the rule. An object's four-momentum is denoted p.

the distribution of $m(\ell_1, t_{\mathrm{FCNC}})$, where the top quark was reconstructed using RJR in the SR2ℓ Dec. The signal and background shapes together with the ROC curve show a substantial separation power.

9.1.3. Dedicated Algorithms in the 3ℓ Final State

Additional reconstruction algorithms are employed exclusively within the 3ℓ final state. Details on their development are provided in Ref. [175]. This dissertation only presents a concise summary of their most important attributes.

The *Neutrino Independent Combinatorics Estimator* (NICE) reconstruction assigns the three leptons of any given event to either the Higgs-boson decay or the top-quark decay. The two leptons associated with the Higgs-boson decay, $\ell_{H,0}$ and $\ell_{H,1}$, (ordered in p_{T}) are defined as the lepton pair with the minimal angular separation. The remaining lepton is designated as ℓ_t. Specific combinatorial rules facilitate the reconstruction of fundamental kinematic parameters, such as invariant masses, with quality criteria labelled *NICE* and *ReallyNICE*.

An alternative method for relabelling leptons, which demonstrates notable separation power, is based not on angular separation, but on the charge of the leptons. In any given event, two leptons must have an identical charge and are designated as $\ell_{\mathrm{SS},0}$ and $\ell_{\mathrm{SS},1}$ in order of increasing p_{T}. The third lepton is labelled ℓ_{OS}. Analogous to the NICE reconstruction, the calculation of simple kinematic variables in combination with other reconstructed objects yields a multitude of variables with substantial separation capability.

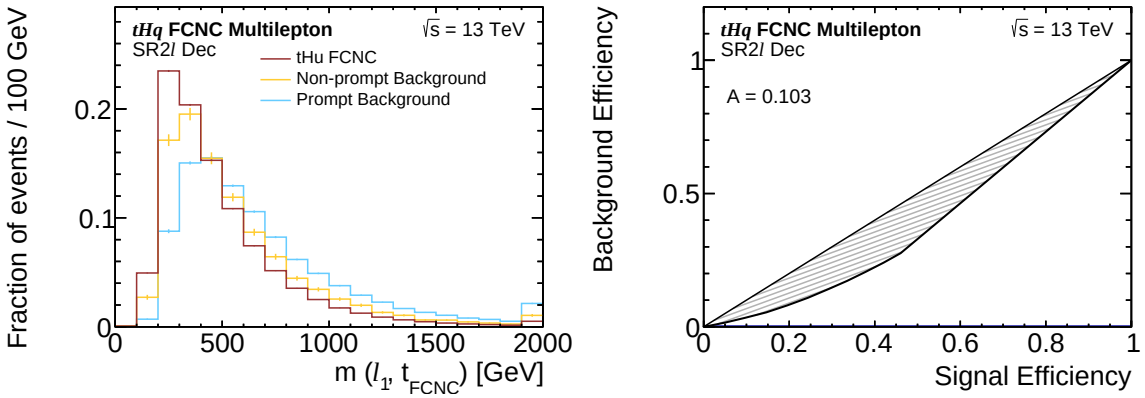

(a) (b)

Figure 9.3.: (a) The distribution of the $m(\ell_1, t_{\mathrm{FCNC}})$ variable in the SR2ℓ Dec for the tHu signal process and the combination of all prompt and non-prompt background processes, together with (b) the corresponding ROC curve. The area between the ROC curve and the diagonal is labelled A.

9.2. Multivariate Analysis using Artificial Neural Networks

The previous section provided an overview of a variety of reconstructed variables, each capable of separating signal from background processes. This chapter explains, how artificial feed-forward NNs are used to combine the separation power of the most significant variables into a single discriminant.

9.2.1. Feed-Forward Neural Networks

NNs, a category of machine learning methods, have gained substantial recognition in recent decades. They are adaptable to a wide range of applications. Prior chapters already explained their use in the identification and isolation of physical objects. In this analysis, NNs are harnessed to achieve an optimal separation between signal and background processes.

NNs are categorised based on their internal structure. This dissertation uses *feed-forward NNs* [177], the simplest and simultaneously the most comprehensively understood type. NNs consist of units called nodes, organised into several disjoint groups called layers. The initial layer, the input layer, functions as the interface for external data, encoding the raw information for network processing. The intermediate layers, referred to as

hidden layers, introduce transformations to the input data, characterised by parameters called *weights and biases*, which are learned during the training process. For a given layer i with $n^{(i)}$ nodes labelled $a_j^{(i)}$, the values of nodes in the subsequent layer $i+1$ are calculated as

$$a_k^{(i+1)} = \sigma\left(\sum_{j=1}^{n^{(i)}} \omega_{j,k}^{(i+1)} \cdot a_j^{(i)} + \mu_k^{(i+1)}\right), \tag{9.1}$$

where $\omega_{j,k}^{(i+1)}$ are the weights connecting the j'th node in layer i to the k'th node in layer $i+1$, and $\mu_k^{(i+1)}$ denotes the bias for the k'th node. The function σ is a non-linear *activation function*, mapping the summation's result to a certain interval. Mathematically, the activation function introduces a non-linearity, allowing the NN to model patterns that linear transformations alone would fail to capture.

By design, the propagation of information through the NN occurs in a feed-forward manner. That is, the input data sequentially traverses each layer, with the output of one layer forming the input for the next. This process concludes at the output layer, which provides the final outcome of the NN's computation.

The optimal values of all weights and biases in the NN are determined in a dedicated training procedure. In this dissertation, the training of an NN is performed using MC events labelled as either signal or background. A loss function measures the NN's precision in evaluating the classification of an event based on the input data. The NNs parameters are optimised in this respect through a procedure called *backpropagation* [177], which entails the reverse transmission of an error signal through the network, starting from the output layer. By calculating the derivative of the loss function with respect to each weight, the algorithm effectively measures the impact of each weight on the overall loss. During the training, weights are adjusted in the direction opposite to this gradient, gradually reducing the error.

However, training NNs carries the inherent risk of overfitting. Overfitting occurs when a network becomes overly specialised at modelling the training data, losing its ability to generalise to new, unseen data. This is typically the result of an overly complex architecture, with too many parameters relative to the amount of training data, or of excessively long training periods. To counteract overfitting, various strategies are employed, such as the introduction of regularisation techniques that penalise complexity in the network, using dropout layers that randomly deactivate a subset of neurons during training, or the adoption of early stopping, wherein training is halted once the performance on a validation set ceases to improve. Addressing overfitting is crucial for developing NNs that are robust and perform well on real-world data.

9.2.2. The NeuroBayes Framework

All NNs presented in this dissertation are trained using the NeuroBayes framework [178, 179]. Developed specifically for machine learning applications in particle physics, NeuroBayes takes a set of variables as input and produces a single output value, designed to approach -1 for background events and 1 for signal events. The guiding principle of NeuroBayes is that the features required to be learned by the NN are reduced to a minimum and yet optimal results are achieved. To this end, NeuroBayes incorporates an extensive preprocessing procedure, ensuring that only the most relevant variables are included in the training. This allows for a compact NN architecture with a single hidden layer.

The preprocessing consists of two main parts: the variable selection and the single-variable preprocessing. The variable selection ranks all potential input variables according to the gain in correlation they provide to the *target variable*, which is defined as -1 for background and $+1$ for signal events. This is accomplished by calculating the total correlation of the set of all N input variables to the target variable via

$$C_N^2 = \vec{\rho}^{\intercal} \cdot R^{-1} \cdot \vec{\rho}. \tag{9.2}$$

Here, $R \in [-1,1]^{N\times N}$ represents the correlation matrix of all input variables, while $\vec{\rho} \in [0,1]^N$ is the vector of correlation coefficients between each input variable and the target. Afterwards, each variable is removed from the total set of variables sequentially, and the new correlation coefficient C_{N-1}^2 is computed for the $N-1$ remaining variables. The variable whose removal results in the smallest loss of correlation, $\Delta C^2 = C_N^2 - C_{N-1}^2$, is identified as the least significant variable for the training and is excluded. This iterative process continues until only one variable remains. A tunable threshold determines, how many of the most significant variables are used as input to the NN.

For each selected variable, a single-variable preprocessing is applied. Initially, a binning is determined such that each bin contains an identical number of events. The purity of each bin is calculated as $S/(S+B)$, where S and B represent the signal and background contribution, respectively. A spline fit is conducted through the resultant distribution, facilitating the transformation of the input variable to the interval $[0,1]$. This spline fit also aims to eliminate statistical fluctuations in regions of nearly constant purity. Subsequently, each variable is normalised to a mean of 0 and a standard deviation of 1. As a final step before commencing the actual training procedure, all selected and preprocessed input variables are decorrelated.

The training is performed using the memory-efficient L-BFGS-B[1] algorithm [180] for backpropagation. The algorithm employs a logarithmic loss function

$$E = \sum_n \ln\left(\frac{1 + t_n p_n + \epsilon}{2}\right) + \alpha \cdot \sum_i \omega_i^2, \tag{9.3}$$

where n indexes the training events, with t_n and p_n denoting the target and the prediction of the n'th event. The parameter ϵ ensures that singularities at $t_n \cdot p_n = -1$ can be handled. Such singularities usually only occur for the first training iterations, so ϵ is set to zero for later iterations. The summation over the square of all weights ω_i in the network ensures an overall preference for smaller weights, thereby reducing the risk of overfitting, which is often characterised by individual large weights. The impact of this so-called regularisation term can be tuned through the parameter α. NeuroBayes utilises a sigmoid activation function, which maps the interval $(-\infty, \infty)$ to the finite interval $[-1, 1]$ and is defined by

$$\sigma(x) = \frac{2}{1 + e^{-x}} - 1. \tag{9.4}$$

The structure of an NN, as well as details in the training procedure can be adapted using dedicated *hyperparameters*. For every NN trained in this analysis, these hyperparameters are optimised using a grid search, which involves setting a range of possible values for each parameter and training an individual NN for every permutation. The optimised NN is selected as the one with the highest area under the ROC curve. An summary of all available hyperparameters and the values considered in the grid search can be found in Table C.1.

The likelihood of overfitting with NNs trained using NeuroBayes is inherently low due to the compact nature of the NN architecture. The risk is further mitigated by adjusting weights to values close to zero based on Bayesian inference from prior training iterations, further reducing the effective number of parameters in the training. Furthermore, NeuroBayes divides the training data into an 80% training set and a 20% validation set. An early-stopping technique, as described in the preceding section, is implemented based on the performance on both datasets.

Finally, the analysis is performed in a way that completely ensures overfitting effects do not compromise the final agreement between MC and data. MC events are divided based on the parity of the unique and pseudorandom event number. For every NN to be trained in the analysis, two models are produced subsequent to the input variable selection and the establishment of the architecture. One model is trained exclusively on events with an odd event number, and the other on events with an even event number.

[1] The acronym L-BFGS-B stands for *Low-memory Broyden-Fletcher-Goldfarb-Shanno Bound* constraints.

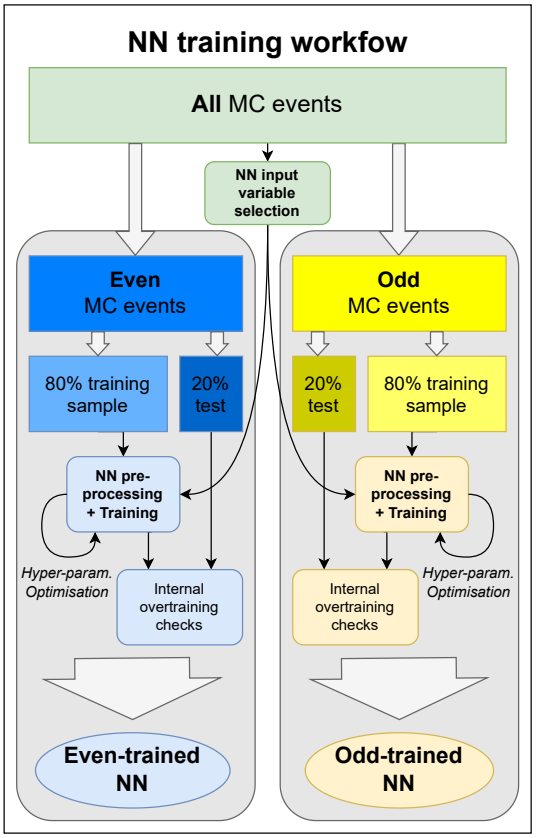

Figure 9.4.: The NN training process employed in this analysis. Variables are selected using the standard NeuroBayes procedure and subsequently divided according to event-number parity. Subsequently, the entire NeuroBayes training process, including variable preprocessing and overfitting checks, is executed, yielding two fully trained NN models.

In the final analysis, the model trained on even-numbered events is applied exclusively to events with an odd event number, and vice versa. This strategy does not eliminate the danger of overfitting – it remains possible for an NN to only perform well on the MC events used in its training. However, when such a model is then applied to previously unseen MC events and real data, any deficiency in performance will be consistent across both, thus preserving overall MC-data agreement. The entire training process of a single NN, split by even and odd event numbers, is illustrated in Figure 9.4.

9.2.3. Neural Networks trained in the Analysis

In this analysis, separate NNs are trained for each signal process across every signal region, yielding a total of eight unique NN architectures. Each training considers the sum of the decay and the production tHq FCNC process as signal. The input variables

selected by the NeuroBayes preprocessing for each of the architectures are summarised in Tables C.2 to C.5. Hyperparameters for each architecture were fine-tuned through a grid search, the result of which can be found in Table C.6. The distributions of NN output variable D_{NN} in every signal region for the tHu signal and the sum of all prompt and non-prompt background process are shown in Figure 9.5. They are depicted as they are incorporated into the statistical analysis, applying even-trained NNs to MC events with odd event numbers, and vice versa, and subsequently adding the two NN distributions. Furthermore, in the analysis, the NeuroBayes output is scaled from the interval $[-1, 1]$ to $[0, 1]$. Analogous distributions of the tHc signal are available in Figure C.1. The overall separation power strikingly surpasses that of the individual reconstructed variables discussed in the previous section. It is also notable that the separation is more pronounced for prompt backgrounds compared to non-prompt background processes. This discrepancy is attributed to the inherent challenge in finding distinguishing properties for non-prompt lepton events with NNs, given their random occurrence in a multitude of different SM processes. This inherent randomness results in a fraction of non-prompt leptons with high kinematic similarity to leptons from signal events.

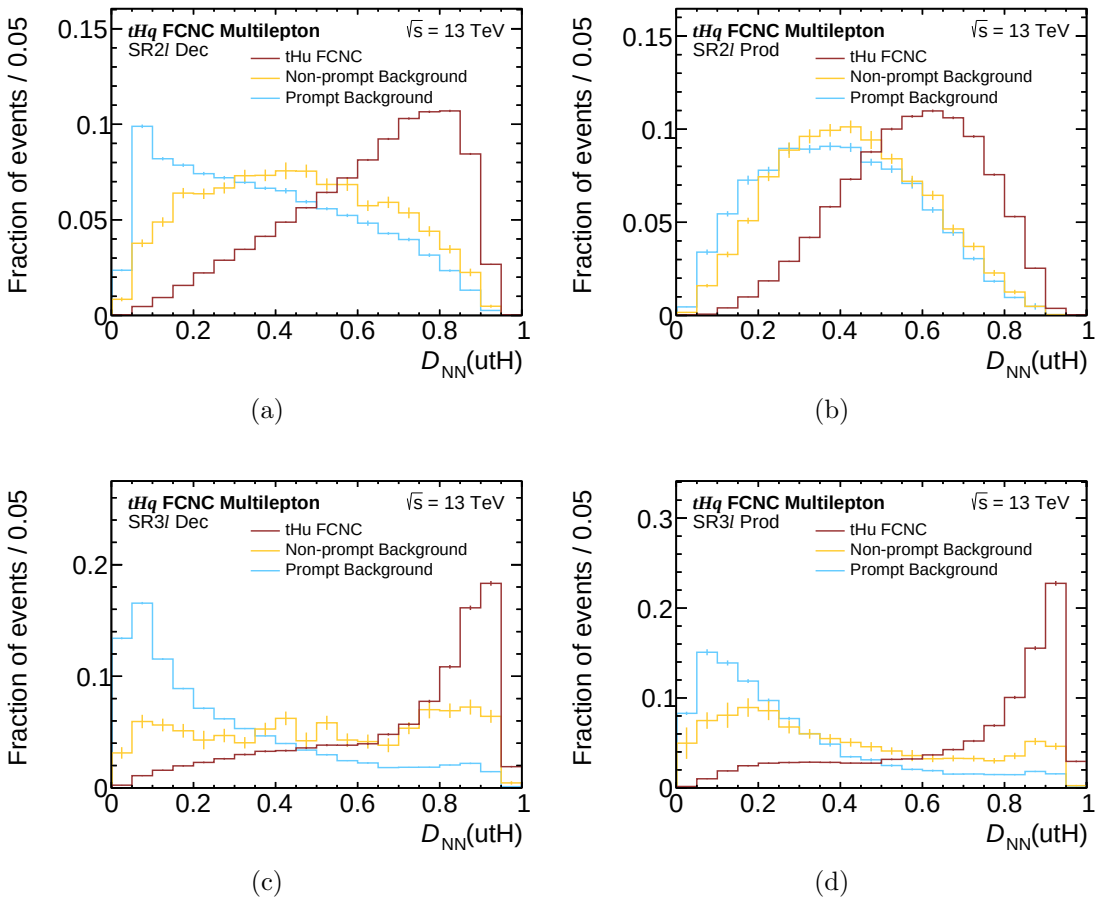

Figure 9.5.: The distribution of the tHu-trained NN output variable D_{NN} in (a) the SR2ℓ Dec, (b) the SR2ℓ Prod, (c) the SR3ℓ Dec and (d) the SR3ℓ Prod for the tHu signal process and the combination of all prompt and non-prompt background processes. The distributions are obtained by applying the even-trained NN to events with an odd event number, and vice versa, and adding the two resulting distributions.

10. Systematic Uncertainties

This chapter presents the estimation of systematic uncertainties inherent to analyses conducted at the ATLAS experiment. Systematic uncertainties, distinct from statistical uncertainties, do not originate from random fluctuations in data but from potential biases and errors in the measurement process and analysis methodology. These may arise from various sources, including the detector calibration, the modelling of physical processes in simulations, and the theoretical premises of the analysis.

Sources of uncertainty can be categorised into two main groups. The first consists of *experimental uncertainties*, presented in Section 10.1, which pertain to the imprecise reconstruction of physical objects. *Modelling uncertainties*, associated with the MC simulation or data-driven estimates of physical background processes, comprise the second group and are discussed in Section 10.2. Before entering the statistical analysis, systematic uncertainties undergo an intricate preprocessing procedure, which is detailed in Section 10.3.

10.1. Experimental Uncertainties

Each source of systematic uncertainty is evaluated through dedicated analyses, which aim to quantify the potential impact of each source on the measurements. The result of this quantification is typically expressed as an upward and a downward variation by one standard deviation. These variations are then propagated through the analysis in order to estimate their effect on the considered physical observables. The uncertainties presented in the following are relevant for all MC simulated samples and assumed to be fully correlated across different processes.

Electron and Muon Calibration

To account for differences in the efficiencies of lepton triggers, as well as lepton isolation, identification and reconstruction, between MC simulations and data, SFs are derived and applied to MC events. For muons, additional SFs associated to the TTVA are included. Variations on all SFs are taken into account in the analysis. For muons, the

uncertainties are split into a statistical component, related to the size of the calibration samples, and a systematic component, stemming from biases in the SF determination.

Moreover, calibration uncertainties related to the energy and momentum of leptons must be addressed. These uncertainties arise from various sources, including but no limited to, noise in the detector signal, fluctuations in the shower shape, and misalignment between the MS and the ID. By leveraging known correlations among these different sources, they are combined into single uncertainties for the scale and resolution of measured electron energies and muon momenta. An additional uncertainty pertains to the energy scale of AFII samples. A specific uncertainty accounts for biases in the sagitta reconstruction of muons, induced by minor deviations from the cylindrical geometry of the ATLAS detector.

Jet Calibration

Uncertainties on the JES originate from more than 100 different sources. Given the correlations among some of them, a simplified scheme comprising 31 independent uncertainties, of which 30 are non-zero depending on the calorimeter simulation employed. They can be separated into 10 categories, covering:

- Limited statistical precision in the calibration samples (6 uncertainties).
- Jet modelling (4 uncertainties).
- Detector modelling (2 uncertainties).
- A mixture of minor uncertainties from various sources (3 uncertainties).
- Biases in the jet η intercalibration (6 uncertainties).
- Jet flavour composition and response (2 uncertainties).
- Jet p_{T} corrections due to pile-up (4 uncertainties).
- Single-particle response of jets (1 uncertainty).
- Differences between b-jets and non-b-jets (1 uncertainty).
- The punch-through effect of jets (2 uncertainties, one applied to samples with a full detector simulation and the other to AFII samples).

The primary uncertainties related to the JER stem from the calibration process. Twelve orthogonal components are identified and propagated through the analysis by smearing jet energies with a Gaussian distribution of a specified width. This width is chosen such that the smearing of jets in addition to the nominal JER reproduce the varied resolution. As established in Chapter 6, the JER varies between MC and data across different phase-space regions. In regions where an uncertainty increases the JER, smearing is applied to the sample with the initially coarser resolution. Conversely, when an uncertainty leads to a reduced JER, the sample with the finer resolution is subject to smearing. In all scenarios, variations are applied as uncertainties on the MC template. Moreover, to

avoid limitations by the size of the dataset, smearing is applied to pseudo-data wherever data should be smeared.

An additional uncertainty is defined for the JER to account for phase-space regions where the data resolution is finer compared to MC simulation. In these cases, pseudo-data is smeared to align with the MC resolution, with the adjustment being subsequently applied as a one-sided uncertainty to the MC samples. Two distinct uncertainties are specified: one for samples using a full detector simulation and one for AFII samples. Lastly, an uncertainty related to the JVT SF is implemented.

$\mathbf{E_T^{miss}}$ Soft Term

The missing transverse momentum is calculated as a function of the reconstructed p_{T} of all detected objects, thereby ensuring the propagation of all uncertainties related to these objects to $E_{\mathrm{T}}^{\mathrm{miss}}$. The only $E_{\mathrm{T}}^{\mathrm{miss}}$-specific uncertainties are those associated with the scale and resolution of the soft term [166]. The resolution uncertainty is divided into two components, addressing the parallel and perpendicular components, respectively.

Flavour-Tagging Efficiencies

b-tagging rates are matched to data using p_{T}-dependent SFs. Uncertainties on these SFs arise from variations in the b-tagging efficiency and the mistagging rates of c-jets and light jets. A total of 45 orthogonal components are introduced for the b-tagging efficiency [181], while 20 components each account for the mistagging uncertainties of c-jets and light jets, respectively [182, 183].

Other Experimental and Instrumental Uncertainties

The uncertainty associated with the measurement of the integrated luminosity is accounted for as a normalisation uncertainty affecting the sum of all MC samples. Concluding the set of experimental uncertainties, an uncertainty on the SF for pile-up reweighting is applied across all samples.

10.2. Modelling Uncertainties

Modelling uncertainties arise from innate inaccuracies in simulating physical processes. They are treated analogously to experimental uncertainties, with variations representing one standard deviation from the nominal sample. In some instances, modelling uncertainties are estimated by comparing the nominal sample to an alternative sample, resulting in just one variation. Unless otherwise specified, the impact of such variations

on kinematic distributions is mirrored around the nominal sample. Additionally, all of these samples are produced using the AFII scheme. To exclude effects from differences in the detector simulation, additional samples, identical to the nominal ones with the full detector simulation being replaced by the AFII parametrisations, are produced for all relevant processes. The difference between the AFII samples is then applied to the nominal sample, defining the final uncertainty.

10.2.1. Modelling of MC Samples

The MC samples employed in this analysis are subject to a multitude of potential modelling discrepancies. The following paragraphs present these sources, the affected samples, and the definition of specific uncertainties in the analysis.

Normalisation of Templates

MC templates are normalised to the predicted cross-sections presented in Chapter 5. These normalisations carry uncertainties, which are addressed differently depending on the MC template. No uncertainties are defined for templates left free-floating in the maximum-likelihood fit, including $t\bar{t}W$, $t\bar{t}Z$ and $VV3\ell + b/c$, as well as $t\bar{t}$, single-t and V+jets production, which only contribute to 2ℓSS and 3ℓ final states by means of non-prompt leptons. Minor non-prompt templates, which are not left free-floating, receive a 50% normalisation uncertainty, as do minor VV templates. The $t\bar{t}H$ production template is subject to an uncertainty of $^{+6.9\%}_{-9.9\%}$, based on variations of renormalisation and factorisation scales, alongside PDF and α_S uncertainties [142]. Templates modelling tZq and tWZ production are assigned 30% normalisation uncertainties, based on the measurement in Ref. [184]. Remaining templates, representing rare top-quark processes, such as $t\bar{t}$ production in association with two heavy bosons, have minimal contribution to the analysis and are assigned a 50% normalisation uncertainty.

Parton Shower and Hadronisation Modelling

Uncertainties on the choice of parton shower and hadronisation models are considered for the tHq FCNC signal, as well as the $t\bar{t}$, $t\bar{t}H$, $t\bar{t}Z$ and $t\bar{t}W$ background processes. For all processes excluding $t\bar{t}W$, these uncertainties are assessed by interfacing the respective nominal matrix-element generators with different versions of Herwig 7 instead of Pythia 8, followed by a comparison between the alternative and the nominal samples. An alternative approach is adopted for $t\bar{t}W$ production, because the nominal Sherpa 2.2.10 samples cannot easily be interfaced with other parton-shower generators. Instead, two distinct samples modelling $t\bar{t}W$ production are generated using Powheg Box v2, each interfaced with either Pythia 8 or Herwig 7. The uncertainties for $t\bar{t}W$ production are then derived from the relative differences between these two samples and applied to the nominal sample.

Parton Shower Matching and Merging

The simulation of $t\bar{t}$ production involves an uncertainty related to the choice of the h_{damp} parameter, as outlined in Chapter 4. This uncertainty is estimated using an alternative sample wherein the parameter is set to $3 \times m_t$. Furthermore, to account for a potential overlap between hard radiation generated by POWHEG and PYTHIA, an alternative sample is produced with the $p_{\text{T}}^{\text{hard}}$ parameter set to 1 instead of its default of 0, following the recommendation in Ref. [185].

For $t\bar{t}W$ production, a global generator uncertainty is defined using an alternative sample generated with MADGRAPH5_AMC@NLO interfaced with PYTHIA 8. This sample employs the FXFX prescription for matching and merging, with a merging scale of $Q_{\text{MS}} = 30\,\text{GeV}$. The uncertainty is referred to as the $t\bar{t}W$ *FxFx* uncertainty.

Factorisation and Renormalisation Scale

The calculations of matrix elements depend on μ_{r} and μ_{f}. Systematic uncertainties associated with these scales are accounted for by varying each scale upwards by a factor of 2 and downwards by a factor of 0.5. These variations are applied for $t\bar{t}$, $t\bar{t}W$, $t\bar{t}Z$, $t\bar{t}H$, VV and VVV production, and considered decorrelated among these different processes. The implementation of these variations involves applying additional weights to selected events.

ISR and FSR

Uncertainties related to ISR are considered for $t\bar{t}$, $t\bar{t}Z$ and $t\bar{t}H$ production, by varying the Var3c parameter of the A14 tune. This corresponds to a variation of α_{S} specifically for ISR [125]. An uncertainty pertaining to FSR is introduced by varying the scale μ_{r} at which α_{S} is evaluated in the FSR part of the parton shower. For the upward variation, μ_{r} is increased by a factor of 2, while for the downward variation it is decreased by a factor of 0.5. This uncertainty is considered for $t\bar{t}$ and $t\bar{t}H$ only. All above variations are implemented as additional event weights, except for the $t\bar{t}$ FSR α_{S} variation, which is estimated using two dedicated samples corresponding to each of the two variations. All ISR and FSR uncertainties are treated as uncorrelated between different processes.

PDFs

Correlated PDF uncertainties on $t\bar{t}$ and $t\bar{t}H$ production are evaluated using the PDF4LHC prescription [186]. The method encompasses 30 variations, obtained by first reweighting all considered samples to a new central value of the PDF4LHC21 PDF set and then calculating the relative difference to one of 30 eigenvectors.

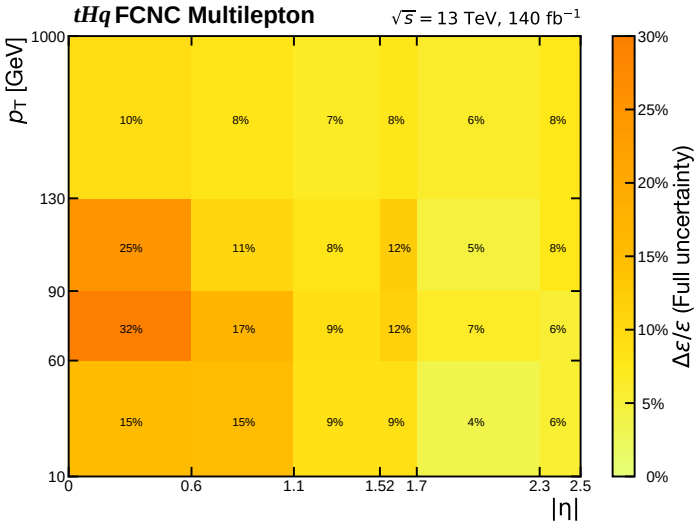

Figure 10.1.: Total relative uncertainty on the Q-misID efficiencies, obtained by considering all sources of systematic uncertainty on the Q-misID background.

10.2.2. Modelling of the data-driven Q-misID Background

As the Q-misID background is modelled using data events, uncertainties related to the calibration of MC simulations to data do not apply. However, during the estimation of Q-misID efficiencies several ambiguities arise, which have to be accounted for as systematic uncertainties. For each source of uncertainty, alternative Q-misID templates are produced, using the respective upward or downward variation. The total uncertainty on the Q-misID efficiencies in each of the considered p_{T}-$|\eta|$ bins, depicted in Figure 10.1, is calculated with the individual uncertainties added in quadrature. Additional figures showing the individual uncertainties are available in Figures D.1 to D.3. All sources of Q-misID efficiency uncertainty are explained in the following.

Q-misID Efficiencies

The Q-misID efficiencies are determined using a maximum likelihood fit. As fit parameters, they possess a statistical uncertainty, derived from the likelihood function's shape.

Choice of the Z-window

The estimation of the Q-misID background is significantly influenced by the definition of the Z-window, expected to be predominantly populated by Q-misID events in the SS channel. The width of this window is chosen arbitrarily. Hence, the nominal value of 4σ

is varied to 3σ and 5σ. For each variation, new Q-misID efficiencies are calculated. The largest deviation from the nominal efficiency in each p_{T}-$|\eta|$ bin is symmetrised around the nominal value and defined as the Z-window uncertainty.

Non-closure with Z → ee MC Samples

A last uncertainty is established by assessing the Q-misID efficiencies using $Z \to ee$ MC events simulated with Sherpa 2.2.11. The method presumes an overall adequate MC modelling of Q-misID events within the respective MC samples. Upon completing the procedure on the samples, the symmetrised difference between the MC Q-misID efficiencies and the nominal efficiencies is defined as an uncertainty in the corresponding p_{T}-$|\eta|$ bin.

10.3. Preprocessing of Systematic Uncertainties

Prior to their incorporation into the statistical analysis, systematic uncertainties are refined to ensure stability in the maximum-likelihood fit while preserving the integrity of the results. The modifications are separated into four distinct stages, detailed in the following.

Rescaling

This procedure is tailored specifically to the $t\bar{t}W$ FxFx uncertainty due to its unique effects in the analysis, as elaborated in Chapter 12. The alternative FxFx sample used to define the uncertainty is rescaled to match the event count across all defined kinematic regions with that of the nominal $t\bar{t}W$ sample. Subsequently, the derived uncertainty is divided into two parts: a *shape* component, measuring the shape differences within regions, and a *region migration* component, measuring the inter-regional differences across regions induced by the uncertainty.

Smoothing

This phase is implemented for all systematic uncertainties estimated using an alternative MC sample. Statistical fluctuations in alternative samples can distort the impact of individual uncertainties on various templates. Such discrepancies may cause unphysical effects or instabilities within the maximum-likelihood fit. To mitigate this, a dedicated smoothing algorithm is employed, which preserves significant features of systematic variations while adapting the values in outlier bins.

Symmetrisation

A first step of symmetrisation has already been introduced: the mirroring of systematics which are defined by just one variation, referred to as one-sided symmetrisation. Additionally, certain uncertainties may exhibit up and a down variations of different magnitudes, resulting in an asymmetry that could lead to instabilities in the maximum-likelihood fit. To counter this, two-sided symmetrisation is employed, adjusting the variations bin by bin to position the nominal value at their centre.

Pruning

Finally, the overall impact of systematic uncertainties on the considered kinematic distributions is evaluated, separately for each region and process. The normalisation component of an uncertainty is omitted if its impact is less than 0.5%. At the same time, the shape component is removed if it modifies a distributions by less than 0.5% in every bin. This pruning procedure significantly stabilises the maximum-likelihood fit, ensuring a more robust behaviour without altering the overall fit results.

11. Methods of Statistical Analysis

This chapter presents the methods of the statistical analysis used in this analysis. Details on the maximum-likelihood fit and the profiling technique it employs are discussed in Section 11.1. Section 11.2 elaborates on methods of hypothesis testing applied, particularly how limits are set on the normalisation of the signal process should the fit show it to be compatible with zero. Finally, Section 11.3 shows how these limits are transformed into limits on physical observables related to the underlying EFT.

11.1. The Profile Likelihood Fit

The statistical analysis in this dissertation is conducted using a specific subset of binned maximum likelihood fits. They are designed to incorporate systematic uncertainties as nuisance parameters (NPs) θ through a technique known as profiling, leading to the name *profile likelihood fit*. NPs are continuous variables that interpolate between up and down variations of a given uncertainty, with a value of 0 corresponding to the nominal value, and ± 1 representing the systematic variations. Interpolations for the shape and the normalisation components of an uncertainty are performed separately. The former is estimated through a linear function between both variations, and the latter via an exponential function, ensuring that normalisation effects cannot lead to negative event yields. The expected number of events in bin i is then a function of the vector of all NPs $\vec{\theta}$:

$$\hat{n}_i = \mu \cdot S_i\left(\vec{\theta}\right) + B_i\left(\vec{\beta}, \vec{\theta}\right), \tag{11.1}$$

where S_i and B_i represent the contributions of signal and background, respectively. The signal normalisation μ is the parameter of interest in this analysis. B_i additionally depends on the normalisation factors $\vec{\beta}$ for certain background templates, and can be expressed as

$$B_i = \sum_{k=1}^{T} \beta_k B_i^{(k)}\left(\vec{\theta}\right) + B_i^{(\text{rest})}\left(\vec{\theta}\right), \tag{11.2}$$

with $B_i^{(k)}$ denoting the contribution from one of T free-floating background templates to bin i, and $B_i^{(\text{rest})}$ aggregating the contributions from all other templates.

A likelihood depending on the parameter of interest, NPs and the free-floating normalisation factors can now be formulated. The event count in each bin follows a Poisson distribution with the expectation value given in Equation (11.1), while Gaussian constraint terms

$$\rho(\theta_j) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{\theta_j^2}{2}\right) \tag{11.3}$$

are included for the NPs. Hence, the likelihood can be expressed as

$$L\left(\mu, \vec{\theta}, \vec{\beta}\right) = \prod_{i=1}^{N} \frac{(\mu \cdot S_i + B_i)^{n_i}}{n_i!} e^{-(\mu \cdot S_i + B_i)} \prod_{j=1}^{M} \rho(\theta_j), \tag{11.4}$$

where N and M represent the total number of bins and constrained NPs, respectively. A deviation of an NP from its pre-fit value of 0, termed a *pull*, signifies a particular sensitivity of the analysis to the corresponding uncertainty. A reduction in the post-fit variance of an NP, known as a *constraint*, might suggest an overestimation of the initial systematic uncertainty. Given that the uncertainties used in this analysis are largely derived from detailed studies, such constraints are generally unexpected and may hint at fit instabilities. Both, pulls and constraints of NPs, thus require meticulous attention in the evaluation of fit results.

Statistical uncertainties on the number of MC events have not yet been addressed. They are included in the fit using Poisson-constrained γ-factors as additional NPs. In principle, one such factor should be defined per MC template and bin. This analysis employs a simplified approach with a single γ-factor per bin, modifying the bin-content's Poisson term in the likelihood to

$$\mathcal{P}(n_i; \hat{n}_i) \to \mathcal{P}(n_i; \gamma_i \hat{n}_i) \cdot \mathcal{P}(x_i; \gamma_i x_i) \tag{11.5}$$

where x_i denotes the total MC event count in bin i.

To ascertain the statistical model's validity, the goodness of fit is evaluated using a saturated fit model, designed to include as many free parameters as there are bins in the fit. The difference between the corresponding maximum likelihood value of this model and the nominal fit model is expected to follow a χ^2 distribution [172] with $n_{\mathrm{df}} = n_{\mathrm{bins}} - n_{\mathrm{NF}}$ degrees of freedom, where n_{NF} is the number of free-floating normalisation factors in the nominal fit model.

11.2. Hypothesis Testing

The obtained fit results can be used to evaluate hypotheses regarding the validity of certain theoretical models. In this analysis, the background-only hypothesis H_0,

presuming the SM is entirely valid, and the signal hypothesis H_μ, positing a signal strength $\mu > 0$, are compared. While some analyses might also consider signal hypotheses for $\mu < 0$, in the context of tHq FCNC couplings, there is no interference between the SM and the considered EFT process, indicating the EFT can only cause an increase in event yields.

To decide whether or not H_0 is rejected, a discriminant λ is constructed, exhibiting distinct distributions under each hypothesis. Such a discriminant is sometimes also referred to as a *test statistic*. Based on the experimental observed λ^{obs}, a *p-value* for H_0 is calculated. This p-value represents the likelihood of obtaining a result exhibiting less compatibility with H_0 than λ^{obs}, assuming H_0 is valid. In particle physics, often the significance $Z = \phi^{-1}(1-p)$ is reported instead of the p-value itself, where ϕ^{-1} is the inverse of the Gaussian cumulative distribution function. The significance threshold for declaring the observation of new physics is set at 5σ, corresponding to a p-value of $2.87 \cdot 10^{-7}$.

To optimally harness the information contained in experimental data, the aforementioned test statistic must be chosen to maximise the separation power between H_0 and H_μ. According to the Neyman-Pearson lemma [187], the most effective discriminant is the likelihood ratio

$$\lambda_\mu = \frac{L(\mu, \hat{\hat{\Theta}})}{L(\hat{\mu}, \hat{\Theta})}. \tag{11.6}$$

Here, $\hat{\mu}$ is the best-fit value for the parameter of interest. Θ represents all parameters in the likelihood besides μ, with $\hat{\Theta}$ being the best-fit values in the nominal fit and $\hat{\hat{\Theta}}$ maximising the likelihood for a fixed μ.

In this analysis, the likelihood ratio undergoes two modifications. First, it is transformed to

$$t_\mu = -2\ln(\lambda_\mu) \tag{11.7}$$

which maps the likelihood ratio to the interval $[0, \infty]$, simplifying further calculations. Secondly, given that negative signal normalisations are unphysical in the context of this analysis, the likelihood ratio for $\hat{\mu} = 0$ is considered when $\hat{\mu} < 0$, leading to a revised test statistic:

$$\tilde{\lambda}_\mu = \begin{cases} \lambda_\mu, & \hat{\mu} \geq 0 \\ \frac{L(\mu, \hat{\hat{\Theta}})}{L\left(0, \hat{\hat{\Theta}}(0)\right)}, & \hat{\mu} < 0. \end{cases} \tag{11.8}$$

Should the fit result indicate a signal normalisation compatible with zero, corresponding to an acceptance of the background-only hypothesis, upper exclusion limits are established for the signal normalisation. Such limits are typically set at a 95% CL,

corresponding to the signal strength $\mu^{(\mathrm{lim})}$ which yields a p-value of 0.05 for the given experimental observation. To obtain the relevant $\mu^{(\mathrm{lim})}$, the test statistic is recalculated for varying μ values until the p-value aligns with the specified threshold. In the context of limit-setting, further modifications to $\tilde{\lambda}_\mu$ are warranted, to avert limits on μ below the best-fit value $\hat{\mu}$. The new test statistic $\tilde{q}_\mu$ is defined as zero in such cases, yielding:

$$\tilde{q}_\mu = \begin{cases} -2\ln \frac{L\left(\mu,\hat{\hat{\Theta}}(\mu)\right)}{L(0,\hat{\hat{\Theta}}(0))}, & \hat{\mu} < 0 \\ -2\ln \frac{L\left(\mu,\hat{\hat{\Theta}}(\mu)\right)}{L(\hat{\mu},\hat{\Theta})}, & 0 \leq \hat{\mu} \leq \mu \\ 0, & \hat{\mu} > \mu. \end{cases} \tag{11.9}$$

This formulation allows for the calculation of the p-value for a given hypothesis H_μ via

$$p_\mu = \int_{\tilde{q}_\mu^{(\mathrm{obs})}}^{\infty} f(\tilde{q}_\mu|\mu)\, d\tilde{q}_\mu, \tag{11.10}$$

where $\tilde{q}_\mu^{(\mathrm{obs})}$ denotes the experimentally observed value of the test statistic. The test statistic's probability density $f(\tilde{q}_\mu|\mu)$ is typically determined through pseudo-experiments with toy data generated by varying the predicted event yield within its uncertainties. Since this induces high computational demands, this analysis employs a parametrisation of q_μ, valid in the limit of sufficiently high event count [188]. This parametrisation presumes a Gaussian distribution for $\hat{\mu}$ with a standard deviation σ, derived from the covariance matrix of the fit result. The parametrised test statistic can be expressed as

$$\tilde{q}_\mu \approx \begin{cases} \frac{\mu^2 - 2\mu\hat{\mu}}{\sigma^2}, & \hat{\mu} < 0 \\ \frac{(\mu-\hat{\mu})^2}{\sigma^2}, & 0 \leq \hat{\mu} \leq \mu \\ 0, & \hat{\mu} > \mu. \end{cases} \tag{11.11}$$

The described methodology for setting upper exclusion limits focuses solely on different signal hypotheses H_μ, without considering the data's consistency with the background-only hypothesis H_0. In scenarios with a significant downward fluctuation in data, the limits on the signal strength might become overly stringent. The $\mathrm{CL_S}$ method [189] aims to provide a more balanced estimation of the upper exclusion limit by factoring in the p-value of the background-only hypothesis H_0. This approach modifies the p-value calculation as follows:

$$p_\mathrm{S} = \frac{p_\mathrm{S+B}}{1 - p_\mathrm{B}}. \tag{11.12}$$

Here, $p_\mathrm{S+B}$ represents the p-value of the signal hypothesis and p_B that of the background-only hypothesis. Should the test-statistic distributions for both hypotheses be sufficiently well separated, p_S is approximately equal to $p_\mathrm{S+B}$. However, if there is considerable overlap between the two distributions, p_S will increase as $1 - p_\mathrm{B}$ decreases. This means the threshold of 0.05 is reached earlier, mitigating the issue of too stringent limits.

11.3. Transformation of tHq FCNC Upper Limits

Following the outlined procedure, limits are derived for the parameter of interest μ, which then must be transformed into constraints on physical parameters to assess restrictions on the underlying EFT. The signal normalisation μ is proportional to the total signal cross-section $\sigma_{\text{FCNC}} = \sigma_{\text{dec}} + \sigma_{\text{prod}}$, encompassing both the decay and the production signal component. This allows for the deduction of

$$\sigma_{FCNC}^{(\text{lim})} = \frac{\mu^{(\text{lim})}}{\mu^{(\text{MC})}} \sigma_{FCNC}^{(\text{MC})}, \tag{11.13}$$

where the upper indices (lim) and (MC) denote the limit and the pre-fit normalisation, respectively. To impose constraints on the EFT Wilson coefficients, these limits must be further transformed. As outlined above, the signal does not interfere with any SM processes, implying that

$$\sigma_{\text{prod}} \propto |C_{u\phi}|^2 \text{ and } \sigma_{\text{dec}} \propto |C_{u\phi}|^2 \tag{11.14}$$

$$\Rightarrow \sigma_{\text{FCNC}} \propto |C_{u\phi}|^2. \tag{11.15}$$

This relation, combined with Equation (11.13), facilitates the calculation of Wilson coefficient limits as

$$|C_{u\phi}^{(\text{lim})}| = \sqrt{\frac{\mu^{(\text{lim})}}{\mu^{(\text{MC})}}} |C_{u\phi}^{(\text{MC})}|. \tag{11.16}$$

Finally, it is beneficial to derive limits on the $\mathcal{B}(t \to Hq)$ branching ratio to enable comparison with other analyses. Given the proportional relationship between the FCNC branching ratio and the square of the Wilson coefficient, as indicated in Equation (2.13), the limit on the branching ratio can be expressed as

$$\mathcal{B}(t \to Hq)^{(\text{lim})} = \frac{\mu^{(\text{lim})}}{\mu^{(\text{MC})}} \mathcal{B}(t \to Hq)^{(\text{MC})}. \tag{11.17}$$

It is important to note that these limits are, strictly speaking, still only applicable to the considered EFT, since all signal kinematics were simulated using the corresponding operators.

12. Results of the Statistical Analysis

This chapter presents the results of the profile-likelihood fit performed to obtain an estimate for the signal normalisation μ. Details on the fit's explicit configuration, particularly the distributions entering the fit in each region, are discussed in Section 12.1. Prior to performing a full fit across the entire phase space, a background-only fit is conducted in bins with a low signal sensitivity to ascertain a consistent modelling of background processes across all kinematic regions. The background-only fit and its results are detailed in Section 12.2. Following this evaluation, Section 12.3 discusses the results of the full fit to data, including the impact of specific systematic uncertainties. Finally, Section 12.4 presents the upper exclusion limits set on the signal process based on the obtained fit results.

12.1. Selection of Fitted Distributions

As explained in Chapter 9, significant improvement in the overall sensitivity of the analysis can be achieved by identifying kinematic distributions in the signal regions with a high discriminatory power between signal and background. Analogously, it is advantageous to identify distributions in control regions that effectively separate the background processes to be controlled from other processes, as this considerably enhances the fit's ability to constrain free-floating normalisation factors. In each control region, numerous distributions were evaluated for their separation power regarding the relevant process. The ones ultimately selected are summarised in Table 12.1.

In the CR2ℓ HFe and CR2ℓ HFμ, only the total number of events are considered in the fit, owing to the limited number of MC events modelling the HF-decay processes. This limitation is less stringent in the CR3ℓ HFe and CR3ℓ HFμ, where sufficiently large numbers of HF-decay MC events are present. The discrepancy among the two final states is related to the use of different MC samples to model $t\bar{t}$ production in each of them, outlined in Chapter 5. As explained in Chapter 7, $t\bar{t}$ is responsible for the majority of all HF-decay events in this analysis. Thus, this process primarily dictates MC-statistical precision of the HF-decay e and HF-decay μ templates. The dedicated $t\bar{t}$ sample with two leptons at matrix-element level used in the 3ℓ final state provides a

Table 12.1.: The kinematic distributions entering the profile-likelihood fit for each region. A description of the respective variable is presented alongside its notation.

Region	Variable	Description
SR2ℓ Dec SR2ℓ Prod SR3ℓ Dec SR3ℓ Prod	$D_{\mathrm{NN}}(tHq)$	The NN discriminant of each region.
CR2ℓ HFe CR2ℓ HFμ	-	Only event yields used.
CR2ℓ $t\bar{t}V$	$p_{\mathrm{T}}(\ell_1)$	The transverse momentum of the subleading-p_{T} lepton.
CR3ℓ HFe CR3ℓ HFμ	$p_{\mathrm{T}}(\ell_2)$	The transverse momentum of the third-leading-p_{T} lepton.
CR3ℓ $t\bar{t}W$ CR3ℓ $t\bar{t}Z$	$p_{\mathrm{T}}(b\text{-jet}_0)$	The transverse momentum of the leading-p_{T} b-tagged jet.

notable increase in the available MC events, leading to higher statistical accuracy in this final state and allowing for the preservation of some shape information in the CR3ℓ HFe and CR3ℓ HFμ.

As a final step prior to performing the profile-likelihood fit, the binning of each distribution entering the fit is optimised. The optimisation of the number of bins in each region is performed with respect to MC-statistical uncertainties, requiring a maximum uncertainty of 12.5% for each individual process in every bin. After determining the number of bins, the bin boundaries are optimised using a binning optimisation algorithm. The algorithm initiates with a finely-binned histogram of $\mathcal{O}(100)$ bins. Starting from the bin with the highest overall content, it merges adjacent bins until the parameter

$$Z = \alpha_{\mathrm{B}} \frac{n_{\mathrm{B}}}{N_{\mathrm{B}}} + \alpha_{\mathrm{S}} \frac{n_{\mathrm{S}}}{N_{\mathrm{S}}} \tag{12.1}$$

reaches a value of one. Here, n_{X} and N_{X} represent the number of signal or background events in a given bin and in the entire region, respectively. The coefficients α_{X} are adjustable parameters of the algorithm, indicating the extent to which signal and background are considered in the selection of bin boundaries, relative to each other. The algorithm yields asymmetric bin boundaries, ensuring that $\alpha_{\mathrm{B}} n_{\mathrm{B}} + \alpha_{\mathrm{S}} n_{\mathrm{S}}$ is approximately constant across all bins. In this analysis, $\alpha_{\mathrm{B}} = \alpha_{\mathrm{S}}$ is chosen in all signal regions, resulting in a linear increase of the signal purity S/B. For control regions, α_{S} is assigned a value of zero, reflecting the negligible signal contribution.

The binning of the signal regions is crucial to the analysis, as an increased number of bins generally enhances the overall sensitivity, provided statistical uncertainties remain

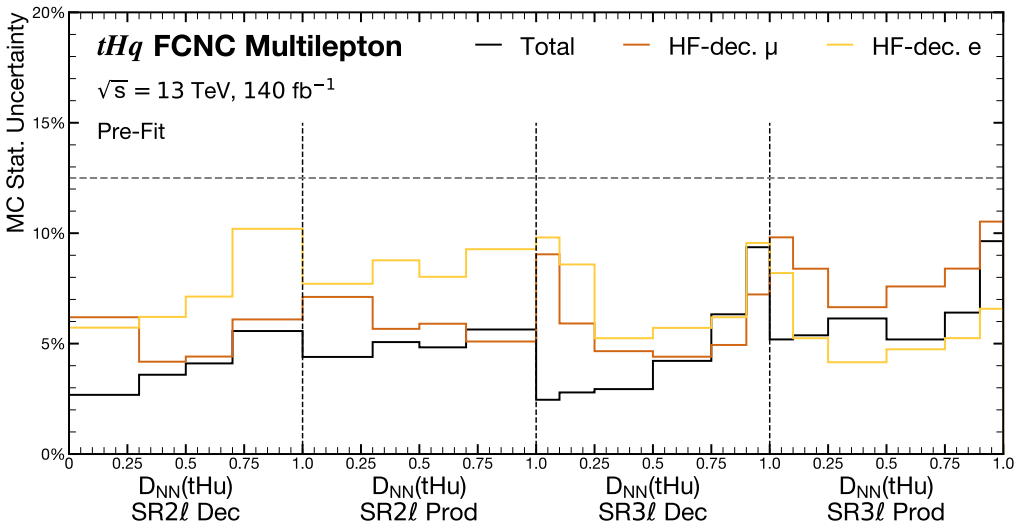

Figure 12.1.: The MC-statistical uncertainties for the HF-decay templates and the combined uncertainty for all processes in the four signal regions are depicted for the optimised binning of the respective $D_{\mathrm{NN}}(tHu)$ distributions.

sufficiently low. For the 2ℓSS signal regions, a total of four bins are included in the fit, while six bins are used for each 3ℓ signal region. The limitation across all regions is primarily imposed by the quantity of available HF-decay events, which is lower in the 2ℓSS final state for this analysis. Figure 12.1 illustrates this by displaying the total MC statistical uncertainties and those for the HF-decay e and HF-decay μ templates in each signal region. The HF-decay templates are the ones with the highest statistical uncertainty and thus constitute the primary constraint on the number of bins. As indicated in Chapter 9, these templates are also the predominant background in the signal-sensitive high-D_{NN} bins. A further increase in the number of bins, leading to an increased MC-statistical uncertainty for HF-decay templates, would significantly impair the overall signal sensitivity.

12.2. Background-Only Fit

Prior to conducting the full fit to data, it is crucial to verify the adequate modelling of the considered background processes, particularly for all variables considered in the fit. This verification is relatively straightforward for the control regions, as they only consider a single kinematic variable each. Conversely, the D_{NN} distributions in the signal regions contain information from $\mathcal{O}(20)$ different variables, all of which require accurate modelling. Ideally, this modelling must be validated prior to exposing the most

signal-sensitive bins to the fit to avoid biases in the analysis. To test the modelling in all regions, a *background-only fit* is performed, wherein only background processes are fitted to data, excluding the most signal-sensitive bins. In each signal region, only bins with $D_{\mathrm{NN}} < 0.5$ are considered. This section primarily presents the results obtained with the $D_{\mathrm{NN}}(tHu)$ distributions. Comparable results were observed for the NNs trained on tHc signal.

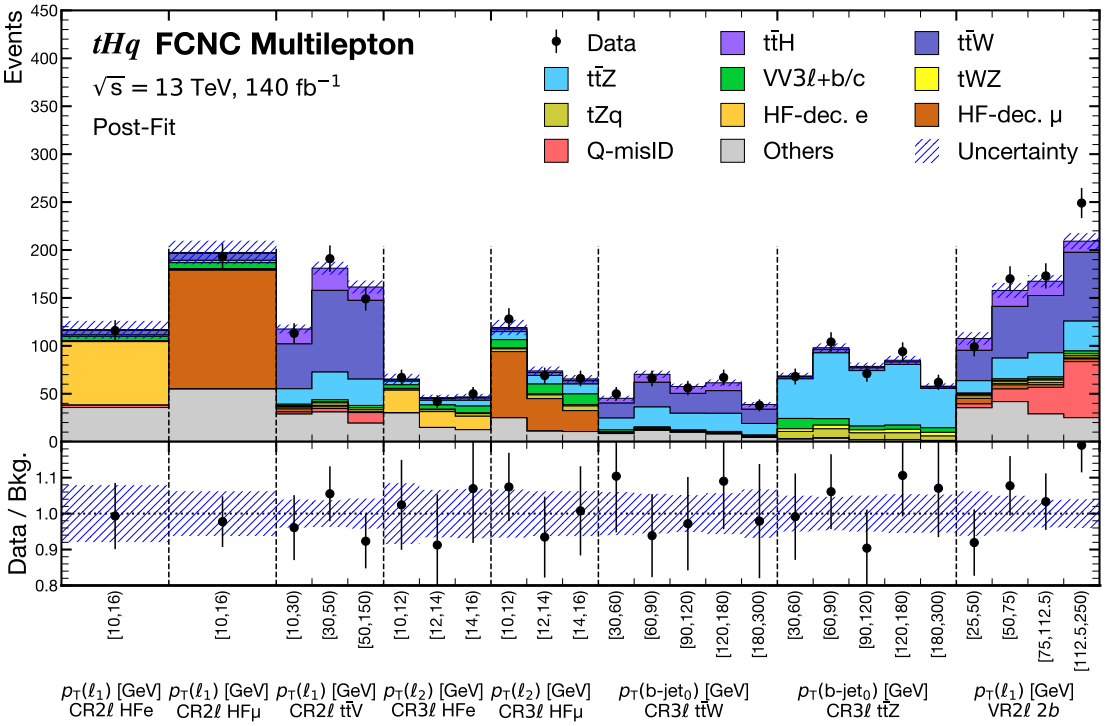

Figure 12.2.: The post-fit agreement between the MC-based model and data for the background-only fit is depicted for all control regions. The shaded area represents all MC-statistical and systematic uncertainties.

The post-fit distributions for all control regions and the validation region are summarised in Figure 12.2, while the low-$D_{\mathrm{NN}}(tHu)$ signal regions are presented in Figure 12.4. An equivalent figure for the $D_{\mathrm{NN}}(tHc)$ distributions is available in Figure E.3. Pre-fit distributions and individual plots for each control and validation region are accessible in Figures E.1 and E.2. A notable agreement between the fit-corrected MC simulations and data is observed across all regions and bins. This consistency, particularly in the VR2ℓ 2b, which was not included in the fit, lends support to the reliability of the background estimation methodology employed in this analysis. The required adjustments of significant background templates is indicated by the fitted free-floating normalisation factors, shown in Figure 12.3. Their overall scale aligns with the expectations expressed

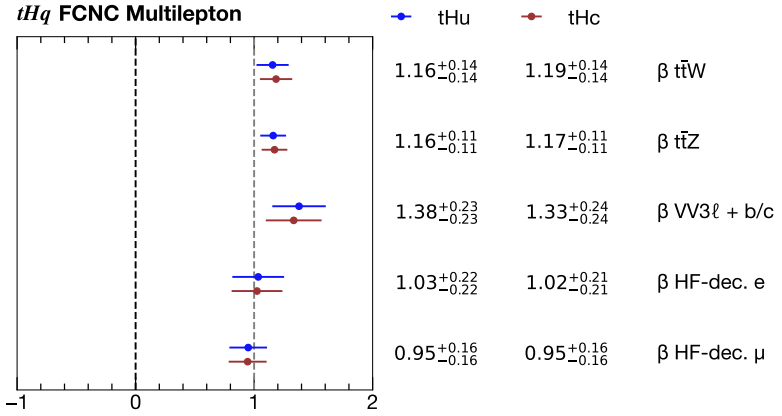

Figure 12.3.: The post-fit values of the free-floating normalisation factors, as determined by the background-only fit.

The NPs obtained from the background-only fit are subsequently applied to the NN input variables for each of the signal regions. The resultant distributions for the highest-ranking input variables of each signal region's NN are shown in Figure 12.5. The evident MC-data agreement not only validates the generalisability of the fit's corrections, but also reinforces the premise that the NN output discriminants are accurately modelled. This suggests that the blinded high-D_{NN} bins are likely to be accurately described through the employed methods.

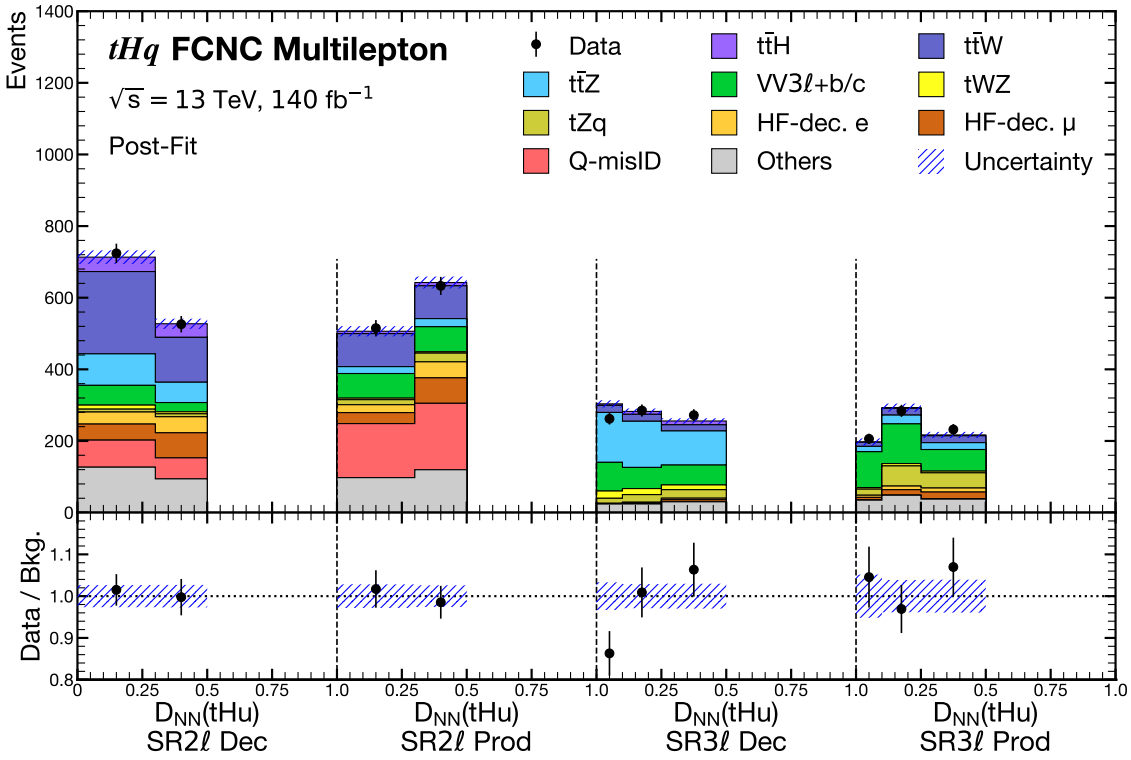

Figure 12.4.: The post-fit agreement between MC and data for the background-only fit across the low-D_{NN} bins of all signal region, using the $D_{\mathrm{NN}}(tHu)$ distributions in the fit. The shaded area encompasses all MC-statistical and systematic uncertainties.

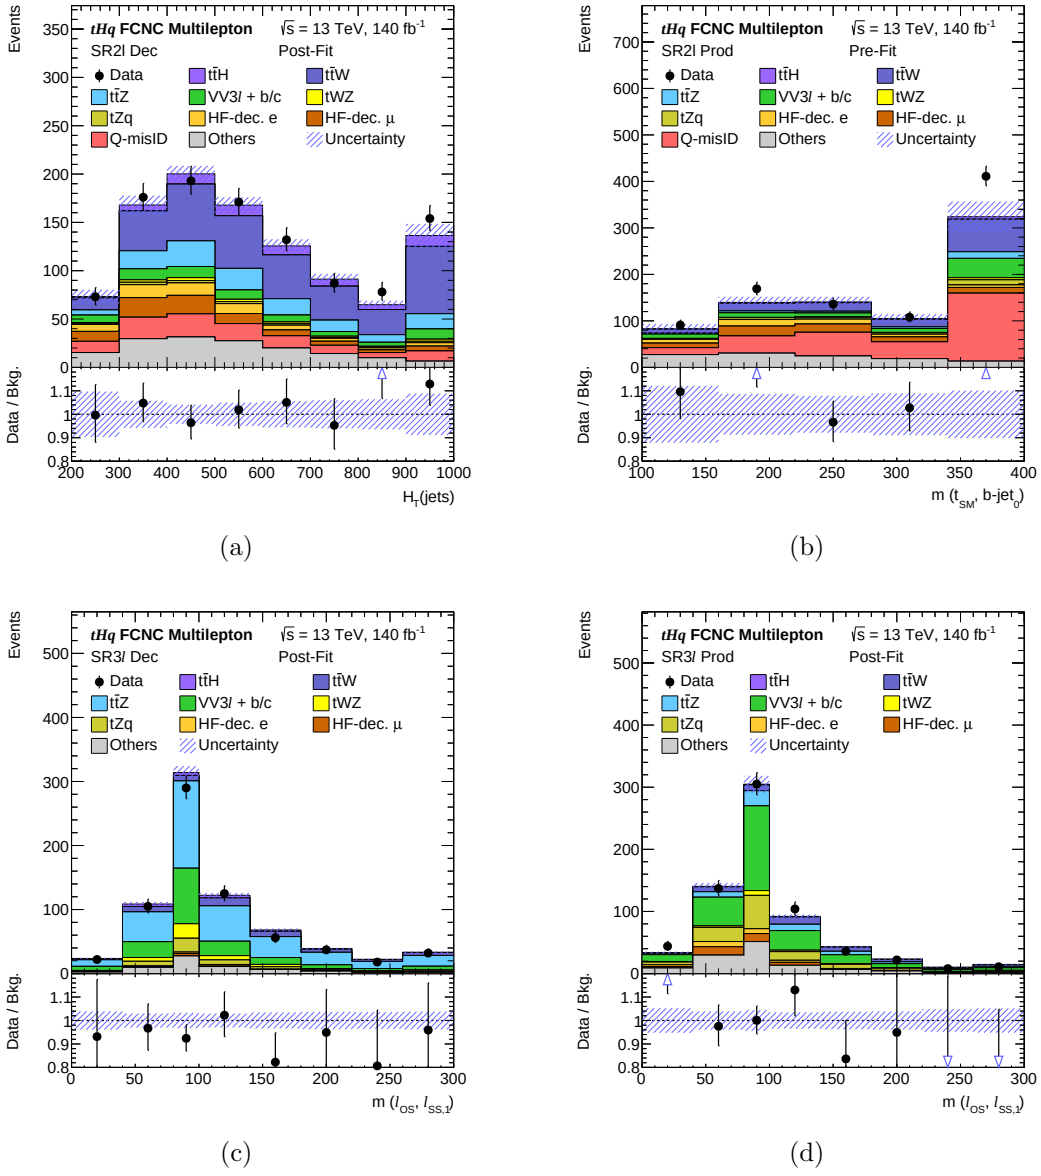

Figure 12.5.: Distributions of the highest-ranking NN input variables of (a) the SR2ℓ Dec, (b) the SR2ℓ Prod, (c) the SR3ℓ Dec and (d) the SR3ℓ Prod after applying the NP values obtained from the background-only fit. The distributions were not used as input to the fit.

12.3. Full Fit to Data

Given that the validation checks support the background modelling chosen for this analysis, two full fits to data, including the most signal-sensitive high-D_{NN} bins, are conducted. One fit includes the tHu signal, the other the tHc signal. Both fits use the respective signal's D_{NN} distributions in the signal regions. The resulting MC-data agreement in control and validation regions mirrors that of the background-only fit. Illustrations of the pre-fit and post-fit distributions in these regions are available in Figures E.15 to E.18. Further details on the post-fit event yields of individual processes are provided in Tables E.1 and E.2. The post-fit D_{NN} distributions in all signal regions for the tHu fit are depicted in Figure 12.6. Analogous distributions for the tHc signal, alongside pre-fit distributions for both signals, are available in Figures E.19 to E.21. In both fits, the corrected MC background closely matches data, indicating no substantial evidence of signal presence. This information is confirmed by examining the post-fit values of the free-floating normalisation factors, shown in Figure 12.7. While those pertaining to background processes align with the results from the background-only fit, the newly added signal normalisation factor is clearly consistent with zero. Prior to determining upper exclusion limits on the signal strength, the overall fit quality is further scrutinised. The goodness-of-fit, as delineated in Chapter 11, is reported in Table 12.2 for both fits, affirming the overall validity of the statistical model. The post-fit values of all NPs are in agreement with the background-only fit, reinforcing the statistical model's applicability to signal-sensitive high-D_{NN} bins as well as those with minimal signal influence. Still, to ascertain the accurate estimation of systematic uncertainties, the most significant NPs are analysed in the following.

Table 12.2.: The goodness of fit, measured by the χ^2/n_{df} fraction together with the corresponding p-value. The χ^2-value is calculated by comparing the maximum likelihood of the nominal fit with that of a saturated fit model.

Signal	χ^2/n_{df}	p-value
tHu	$34.6/35 = 0.99$	0.49
tHc	$37.4/35 = 1.07$	0.36

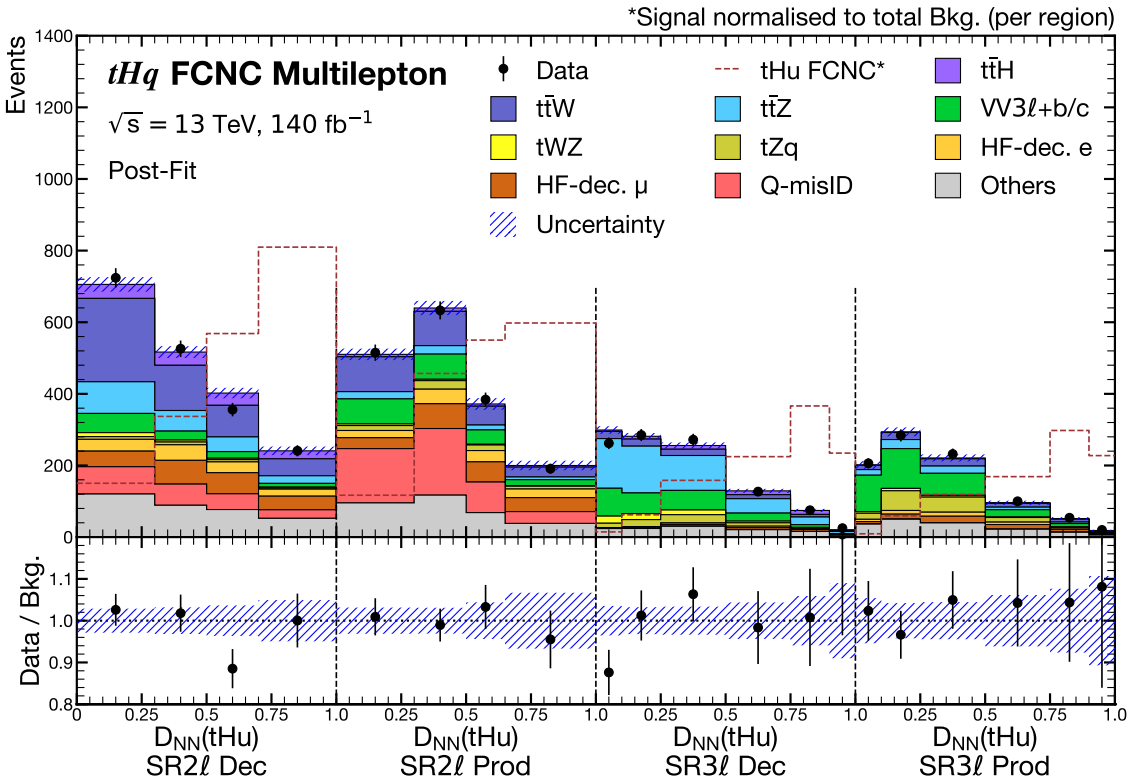

Figure 12.6.: The fitted D_{NN} distributions of the tHu fit. The shape of the respective signal process is shown as a dashed line, scaled to the total number of background events per region.

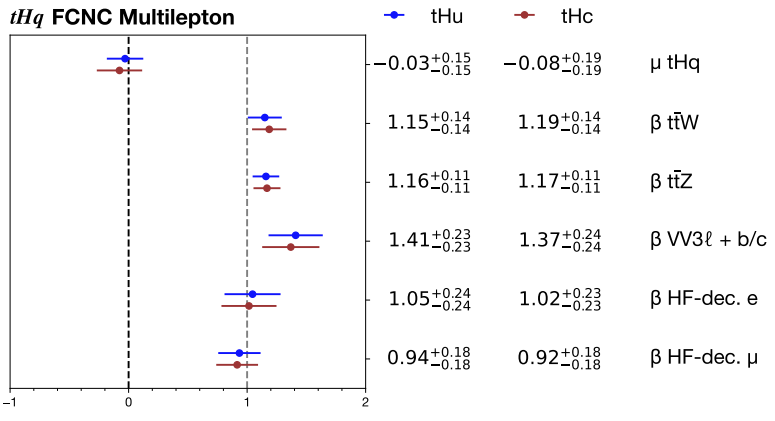

Figure 12.7.: The post-fit values of the free-floating normalisation factors, as determined by the full fit to data. The signal normalisation $\mu\,tHq$ as the parameter of interest is shown in the top row.

12.3.1. Impact of Nuisance Parameters

As presented in Chapter 10, a wide array of systematic uncertainties is incorporated into the fit. The NPs related to each of them may be pulled from their nominal value, while their uncertainties can be constrained. The former is also true for γ-factors, modelling the MC-statistical uncertainty. Understanding the origins of these pulls and constraints is vital to ensure the treatment of uncertainties is correct. The same is true for strong correlations among NPs. Both, strongly pulled and highly correlated NPs are discussed in the following.

Correlations of Nuisance Parameters

The correlation matrix, showing all NPs correlated with at least one other NP by more than 0.25, is presented in Figure 12.8 for the tHu fit, with an analogous plot for the tHc fit available in Figure E.22. Analysis of the correlation matrix reveals that most of the correlations are driven by the free-floating normalisation factors. The NPs to which these normalisation factors are correlated are strongly related to the respective process being normalised. The normalisation factors for HF-decay e and HF-decay μ exhibit strong anti-correlations with the normalisation of other non-prompt processes, as well as NPs related to the electron and muon identification and isolation. The former arise because the same SM processes ($t\bar{t}$, single-t and V+jets) underpin all non-prompt processes, resulting in some kinematic similarities. The latter attest to the significant effect of lepton identification and isolation inefficiencies on the likelihood of misidentifying HF-decay leptons as real ones. Moreover, both HF-decay e and HF-decay μ normalisations are anti-correlated with the signal normalisation $\mu\, tHu$, emphasising that non-prompt background processes have the highest impact in the most signal-sensitive high-D_{NN} bins.

The normalisation factor for $VV3\ell + b/c$ is strongly anti-correlated to the normalisation NP of another VV template, which is understandable given their shared origins. It exhibits an even stronger anti-correlation with the NP for the tZq cross-section uncertainty, which can be traced back to similar shapes of both processes in the 3ℓ signal regions.

The $t\bar{t}W$ normalisation factor shows a strong anti-correlation with the NP for the γ-conv. normalisation uncertainty. This may be an artifact of the non-prompt process's somewhat random behaviour. A physically meaningful correlation exists with the NP for the principal b-tagging uncertainty for b-jets, attributable to $t\bar{t}W$ necessitating one b-jet to go untagged in order to enter the signal regions.

The same correlation with the principal b-tagging uncertainty is observed for $t\bar{t}Z$. Additionally, the normalisation factor exhibits substantial anti-correlations with the $t\bar{t}Z$

shower NP and the NP related to the uncertainty on the pile-up ρ-topology. The former shows that the shower uncertainty has a strong normalisation impact, while the latter pertains to the reason for $t\bar{t}Z$'s normalisation to be left unconstrained in the fit: the modelling quality of the process for low N_{jets} is uncertain, and the kinematics of jets are affected by the systematic uncertainty in question.

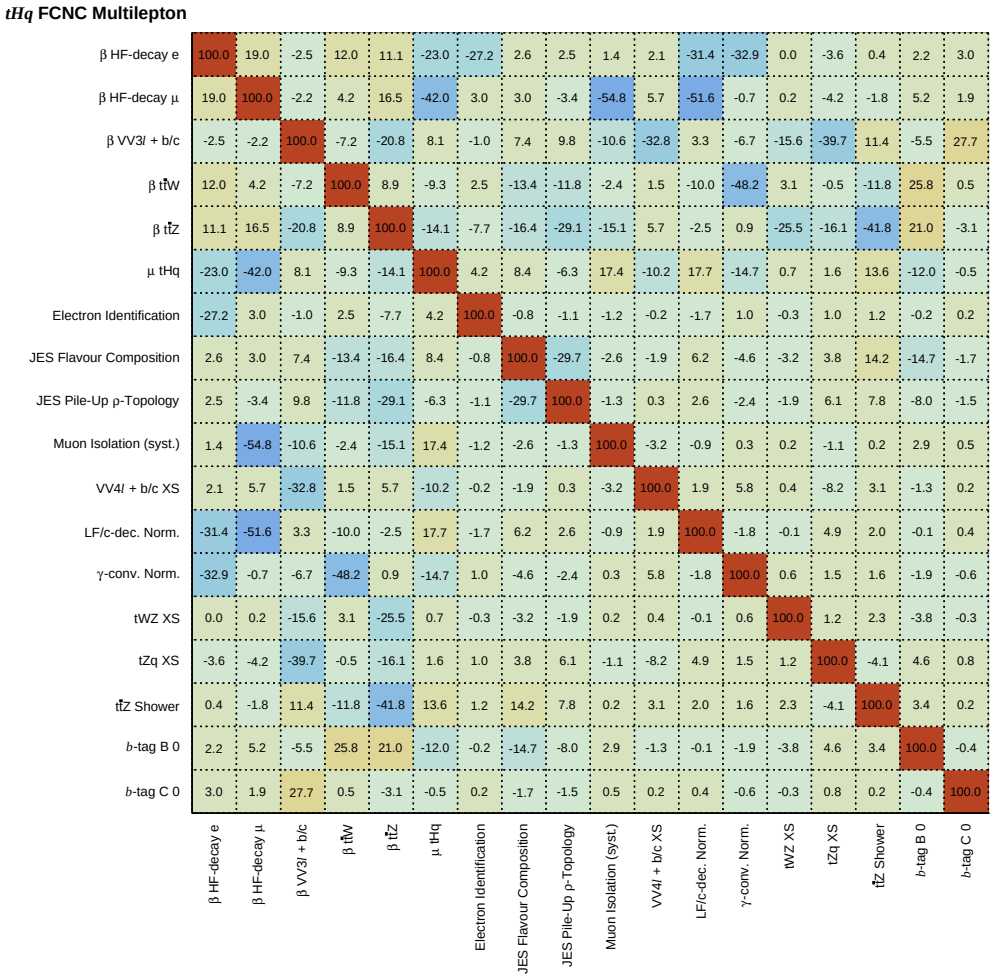

Figure 12.8.: Correlations of NPs for the full tHu fit. Only NPs with a correlation above 0.25 to at least one other NP are shown in the correlation matrix.

Pulls and Constraints of Nuisance Parameters

Figure 12.9 presents NPs that are significantly constrained or pulled beyond a threshold of 0.3σ for the tHu fit, with an analogous plot for the tHc plot available in Figure E.23. A

complete set of all post-fit NP values is provided in Figures E.24 to E.33. Many of the NPs exhibiting strong pulls are also highly correlated with at least one normalisation factor. Given that the free-floating processes are among the most significant in the analysis, and that dedicated control regions are established to constrain their normalisation, a certain sensitivity to NPs closely related to this normalisation is anticipated.

Two NPs with notable pulls pertain to the modelling of pile-up, which significantly influences the distribution of N_{jets}. The required number of jets is one of the key distinctions between decay and production signal regions in both final states, rendering this analysis particularly sensitive to NPs affecting this variable. A similar rationale applies to the two b-tagging related NPs: one corresponding to the already mentioned principal uncertainty for the tagging of b-jets, the other representing the principal component for the mistagging of light jets. The region definition of this analysis also relies on $N_{b\text{-tags}}$, making it susceptible to uncertainties on the b-tagging methodology.

Another category of heavily influenced NPs includes those related to modelling uncertainties. NPs associated with parton-shower uncertainties in $t\bar{t}$, $t\bar{t}W$ and $t\bar{t}Z$ are pulled significantly. For $t\bar{t}Z$, this is additionally true for the NP related to the Var3c variation of the parton-shower tune. The parton shower plays a crucial role in the modelling of jets. For $t\bar{t}$, this strongly impacts the likelihood of lepton production in jets and their subsequent detection as isolated, resulting in notable variations in the HF-decay e and HF-decay μ templates. The steps taken to control both processes ensures a high sensitivity to NPs altering their kinematics. In contrast, the most significant impact of $t\bar{t}W$ and $t\bar{t}Z$ shower uncertainties relates to the modelling of b-jets and the consequent impact on $p_{\text{T}}(b\text{-jet}_0)$, the variable with which the CR3ℓ $t\bar{t}W$ and the CR3ℓ $t\bar{t}Z$ enter the fit.

Lastly, the NP related to the region migration component of the $t\bar{t}W$ FxFx variation is pulled to almost -0.5σ. At the same time, it exhibits the strongest constraint in the entire fit, of approximately 25%. This warrants particular attention to this parameter. As explained in Chapter 10, the NP quantifies the difference in normalisation across different regions caused by the $t\bar{t}W$ FxFx variation. When assessing the resulting effect among regions with a non-negligible $t\bar{t}W$ content, an general trend of reduced normalisation in 2ℓSS regions and an increase in 3ℓ regions is noted. Moreover, the increase is particularly pronounced in the CR3ℓ $t\bar{t}W$, amounting to approximately 9% compared to a maximum of 4% in other regions. The strong pull thus enables the fit to correct the stronger MC-data discrepancies in this region, without negatively affecting regions with better agreement.

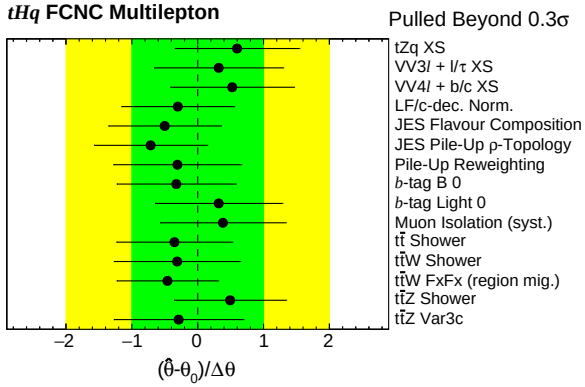

Figure 12.9.: Pulls and constraints of NPs from the full tHu fit, surpassing a pull threshold of 0.3σ. Among the depicted NPs are also the ones exhibiting the strongest constraints in the fit.

Impact of Nuisance Parameters on the Signal Normalisation

The NPs discussed above represent systematic uncertainties that are important in the description of the 2ℓSS and 3ℓ phase space considered in the analysis. However, this analysis's parameter of interest is the signal normalisation, while the precise values of individual NPs are of lesser concern. To assess the impact of NPs on the signal normalisation, each one is fixed to its nominal value plus or minus its post-fit uncertainty. The fit is then rerun, keeping the fixed NP at the varied value, and the altered signal normalisation is recorded. NPs are ranked in descending order based on the deviation of the obtained signal normalisation from the nominal fit's result. Figure 12.10 presents the ensuing NP ranking for both fits.

It is evident that the NPs with the highest impact are predominantly related to the non-prompt background modelling. This assessment holds for the HF-decay normalisation factors, as well as many of the constrained NPs shown, for reasons discussed above. It owes to the fact that non-prompt background processes, particularly HF decay, predominantly populate the most signal-sensitive high-D_{NN} bins due to the NN's inability to distinctly differentiate these processes from signal. Other high-ranking NPs are related to $t\bar{t}W$ and $t\bar{t}Z$, which constitute the largest prompt background processes. An outlier is the γ-factor for the highest D_{NN} bin in the SR2ℓ Dec, the high ranking of which owes to the aforementioned limited statistical precision of the HF-decay processes in the 2ℓSS final state. Predominantly, the ranked NPs are neither significantly pulled nor constrained, indicating that the signal normalisation is largely influenced by uncertainties not strongly affected by the fit, thus ensuring overall stability.

The backgrounds from $t\bar{t}W$ and $t\bar{t}Z$ are inherently challenging to reduce in this analysis, since these processes enter the 2ℓSS and 3ℓ final states due to their innate kinematics.

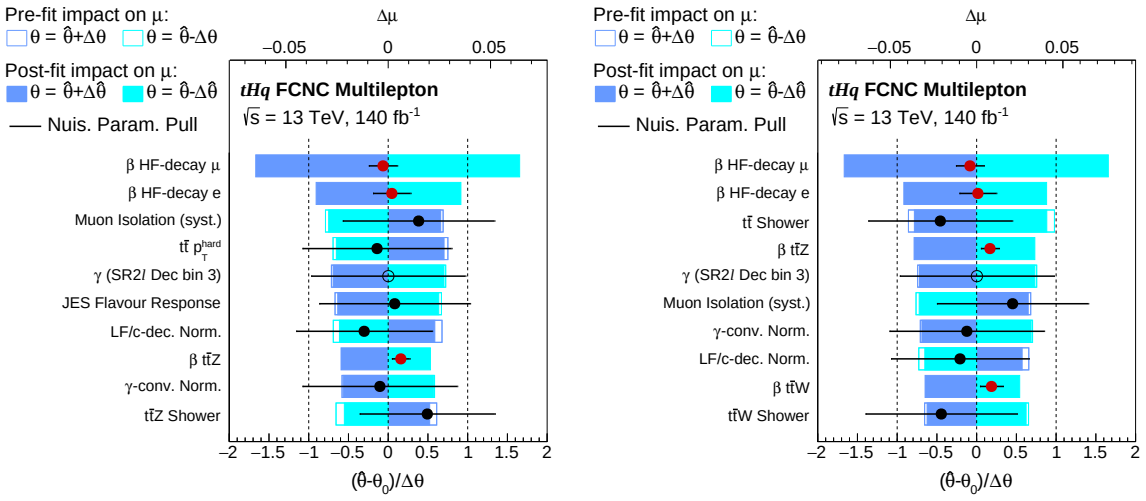

(a) (b)

Figure 12.10.: NPs ranked according to their impact on the signal normalisation μ. For each NP, the fit is rerun with the NP fixed to its post-fit up or down variation. The resulting difference $\Delta\mu$ between the signal normalisation of the nominal and the modified fit is represented by the solid bars and relates to the upper y-axis. Similar fits are performed using the pre-fit uncertainties, with the results indicated by the empty bars. The pulls and constraints of each NP are illustrated by the black data points and errorbars, associated with the lower y-axis. Red data points indicate that the related NP is a free-floating normalisation factor, while empty data points represent γ factors.

Conversely, non-prompt processes could potentially be mitigated using more sophisticated identification techniques. Focusing on enhancing such techniques would likely be the most effective strategy for minimising the impact of systematic uncertainties in this analysis.

12.4. Upper Exclusion Limits on tHq FCNC Couplings

Given that the full fit to data results in a signal normalisation compatible with zero, upper exclusion limits are set on both the tHu and tHc signal processes. The 95% CL limits on the signal normalisation μ are computed using the $\mathrm{CL_S}$ method and subsequently converted into limits on the branching ratios $\mathcal{B}(t \to Hq)$ and the EFT Wilson coefficients $C_{u\phi}$, as explained in Chapter 11. Prior to determining the observed upper exclusion limits, various studies on the sensitivity of this analysis are conducted in Section 12.4.1. Following these studies, the observed upper exclusion limits are presented in Section 12.4.2 and compared with other Run 2 searches for tHq FCNC couplings.

Table 12.3.: The 95% CL expected upper exclusion limits on the branching ratio $\mathcal{B}(t \to Hq)$, obtained by including either only 2ℓSS or only 3ℓ final state regions in the fit. The limits of the nominal analysis, including both final states, are labelled *2ℓSS and* 3ℓ. Limits are shown for both signal processes.

Final State	$\mathcal{B}(t \to Hq)^{(\mathrm{lim})}$ /10^{-4}									
	-2σ	-1σ	tHu	$+1\sigma$	$+2\sigma$	-2σ	-1σ	tHc	$+1\sigma$	$+2\sigma$
3ℓ	3.4	4.6	6.5	9.3	12.7	4.7	6.3	8.9	12.6	17.3
2ℓSS	1.9	2.6	3.6	5.1	7.0	2.3	3.1	4.3	6.2	8.5
2ℓSS and 3ℓ	1.6	2.2	3.0	4.2	5.7	2.0	2.7	3.8	5.3	7.1

12.4.1. Sensitivity of the Analysis

The sensitivity of an analysis searching for a process whose post-fit normalisation is compatible with zero is estimated using the expected upper exclusion limit. This expected limit is determined based on a realistic Asimov dataset, constructed by applying all fitted background corrections to background MC samples and summing up the predictions for each bin. The limits obtained from a fit to this Asimov data are defined as the expected limits. Variations of the Asimov dataset allow for the computation of 1σ and 2σ confidence intervals around each expected limit. All limits discussed in the following are transformed into limits on the branching ratio $\mathcal{B}(t \to Hq)$.

Comparison of Final States

As an initial step, the expected upper exclusion limits of the complete analysis for both the tHu and tHc signal are determined and listed in Table 12.3. Additionally, the limits for each individual final state are shown, obtained by excluding all regions of the other final state from the fit. Figure 12.11 illustrates all expected exclusion limits. A comparison of the three sets of limits reveals that the 2ℓSS final state predominantly contributes to the sensitivity of the analysis for both signal processes. However, the addition of the 3ℓ final state offers a discernible benefit, enhancing the upper limit by 20% and 13% for the tHu and the tHc signal, respectively.

Impact of Systematic Uncertainties

Systematic uncertainties and the influence of their associated NPs on the fit result were discussed in detail in the preceding section. In general, additional uncertainties deteriorate the sensitivity of an analysis. To understand, how strongly this analysis is impacted by systematic uncertainties, an additional fit including only statistical and cross-section uncertainties is performed. Furthermore, two fits are performed, one of

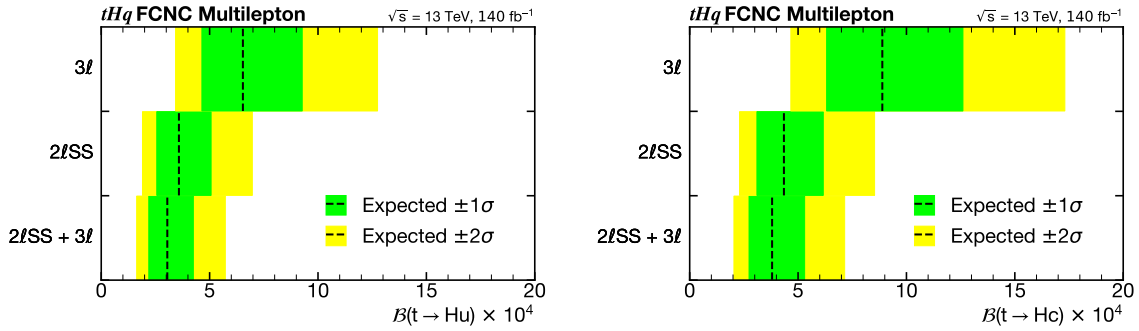

(a) (b)

Figure 12.11.: The 95% CL expected upper exclusion limits on the branching ratio (a) $\mathcal{B}(t \to Hu)$ and (b) $\mathcal{B}(t \to Hc)$, obtained by including only 2ℓSS or only 3ℓ final state regions in the fit, together with the limits of the nominal analysis.

which includes exclusively experimental uncertainties, while the other only incorporates modelling uncertainties. They allow to estimate the extent to which each category impacts the final result. The expected upper limits derived from these fits are listed in Table 12.4 and visualised in Figure 12.12.

A comparison of the limits obtained using only statistical uncertainties with those from the nominal analysis reveals an increase of approximately 13% for both signal processes. Moreover, no clear dominance of experimental versus modelling uncertainties is observed, as both appear to equally increase the upper limits. This suggests that the analysis is predominantly limited by statistical uncertainties, and an enlarged dataset would be the most effective way of achieving a higher sensitivity.

Table 12.4.: The 95% CL expected upper exclusion limits on the branching ratio $\mathcal{B}(t \to Hq)$, obtained by including only statistical uncertainties, or statistical and either experimental or modelling uncertainties in the fit. The limits of the nominal analysis, including all systematic uncertainties, are also shown and labelled *full syst.*. Limits are shown for both signal processes.

Uncertainties	$\mathcal{B}(t \to Hq)^{(\mathrm{lim})}$ $/10^{-4}$									
	-2σ	-1σ	tHu	$+1\sigma$	$+2\sigma$	-2σ	-1σ	tHc	$+1\sigma$	$+2\sigma$
Stat. only	1.4	1.9	2.6	3.7	4.9	1.7	2.3	3.3	4.5	6.1
Exp. only	1.5	2.0	2.8	3.9	5.3	1.9	2.5	3.5	4.9	6.6
Mod. only	1.5	2.0	2.8	4.0	5.4	1.9	2.6	3.6	5.0	6.7
Full syst.	1.6	2.2	3.0	4.2	5.7	2.0	2.7	3.8	5.3	7.1

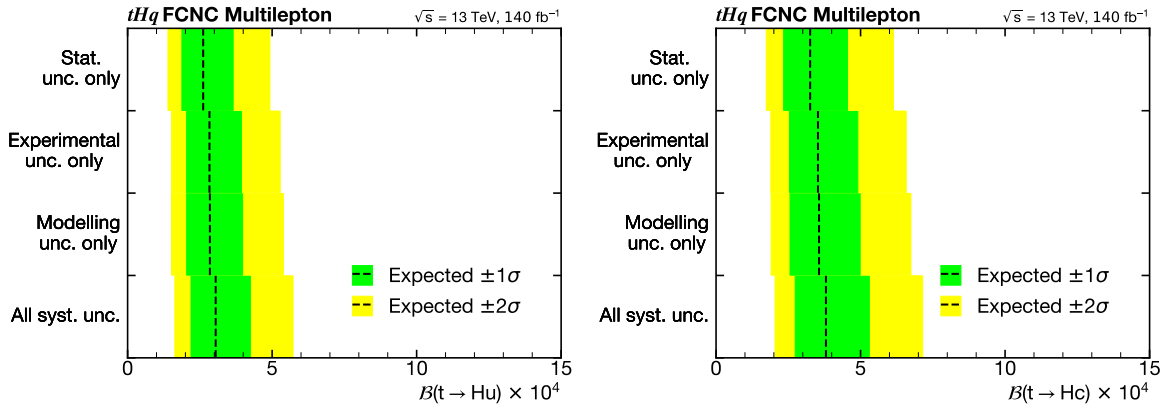

(a) (b)

Figure 12.12.: The 95% CL expected upper exclusion limits on the branching ratio (a) $\mathcal{B}(t \to Hu)$ and (b) $\mathcal{B}(t \to Hc)$, obtained by including only statistical uncertainties, or statistical and either experimental or modelling uncertainties in the fit, together with the limits of the nominal analysis, including the full set of systematic uncertainties.

Comparison to 36.1 fb^{-1} tHq FCNC search

In Chapter 2, various previous analyses searching for tHq FCNC interactions were mentioned, including another ATLAS analysis considering 2ℓSS and 3ℓ final states with a partial Run 2 dataset of 36.1 fb^{-1}, published in 2018 [60]. The results of this analysis are clearly improved upon by this analysis, owing to both higher statistical precision and differences in analysis methodologies. To disentangle both effects, the realistic Asimov dataset used to determine expected upper limits is rescaled to an integrated luminosity of 36.1 fb^{-1}. The resulting expected upper limits, alongside those obtained by the 2018 analysis, are listed in Table 12.5 and graphically depicted in Figure 12.13.

The expected limit for the tHc signal is improved by a factor of 2.3. This notable gain in sensitivity owes to refined analysis techniques, especially an enhanced rejection of non-prompt leptons, facilitated by the usage of the PLIV isolation method, discussed in Chapter 6. Whereas non-prompt leptons from HF decay constitute 44% of the entire background in the 2ℓSS signal regions in the 2018 analysis, this proportion is reduced to 18% in this analysis. Given that this background presents the most significant challenge for NN discrimination, its effective exclusion is crucial for a heightened signal sensitivity.

Regarding the tHu signal, an even more pronounced improvement by a factor of 2.8 can be observed. The additional enhancement is attributable to the incorporation of the $qg \to Ht$ production signal process into the analysis, which was not considered in the 2018 analysis. This process only contributes significantly to the tHu signal process, thus amplifying the improvement for this signal.

Table 12.5.: The 95% CL expected upper exclusion limits on the branching ratio $\mathcal{B}(t \to Hq)$, obtained by performing a fit to Asimov data with the luminosity scaled to $36.1\,\text{fb}^{-1}$. For comparison, the expected limits from the 2018 ATLAS analysis searching for tHq FCNC couplings in 2ℓSS and 3ℓ final states with a partial Run 2 dataset corresponding to that luminosity are also shown [60]. Limits are shown for both signal processes.

Analysis	$\mathcal{B}(t \to Hq)^{(\text{lim})}\,/10^{-4}$ $(\int \mathcal{L} = 36.1\,\text{fb}^{-1})$									
	-2σ	-1σ	tHu	$+1\sigma$	$+2\sigma$	-2σ	-1σ	tHc	$+1\sigma$	$+2\sigma$
2018 Analysis	7.8	10.5	14.8	21.2	29.9	8.2	11.0	15.5	22.2	31.1
This Analysis	2.9	3.9	5.3	7.4	10.0	3.6	4.9	6.7	9.4	12.6

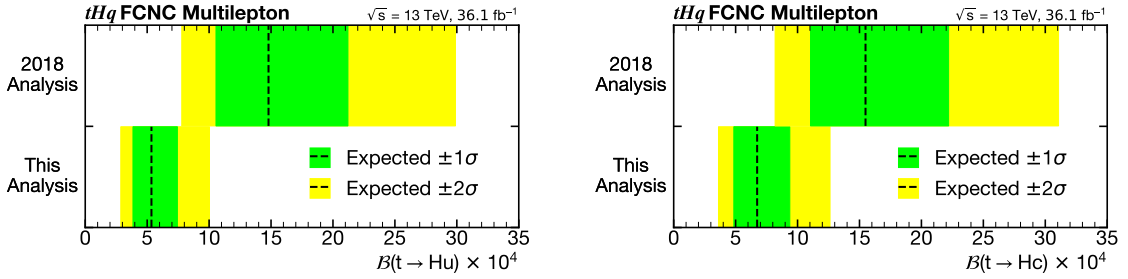

(a) (b)

Figure 12.13.: The 95% CL expected upper exclusion limits on the branching ratio (a) $\mathcal{B}(t \to Hu)$ and (b) $\mathcal{B}(t \to Hc)$, obtained by performing a fit with the luminosity scaled to $36.1\,\text{fb}^{-1}$. The limits from the 2018 ATLAS analysis searching for tHq FCNC couplings in 2ℓSS and 3ℓ final states with a partial Run 2 dataset are also shown [60].

12.4.2. Observed Upper Exclusion Limits

As the final stage of the analysis presented in this dissertation, the observed upper exclusion limits on the considered signal processes are determined. The observed (expected) exclusion limits on the signal strength are found to be $\mu(tHu) < 0.28(0.30)$ and $\mu(tHc) < 0.33(0.38)$. Based on a pre-fit signal normalisation corresponding to $\mathcal{B}(t \to Hq) = 0.1\%$, this corresponds to limits on the FCNC branching ratios of $\mathcal{B}(t \to Hu) < 2.8\,(3.0) \times 10^{-4}$ and $\mathcal{B}(t \to Hc) < 3.3\,(3.8) \times 10^{-4}$. When converted to limits on the EFT Wilson coefficient, observed (expected) values of $|C_{u\phi}^{13,31}| < 0.71\,(0.73)$ and $|C_{u\phi}^{23,32}| < 0.76\,(0.82)$ are derived. All limits are summarised in Table 12.6.

This analysis obtains the strongest expected upper limits amongst all analyses searching for tHq FCNC couplings in individual final states, rendering this the most signal-sensitive

Table 12.6.: The observed (expected) 95% CL upper exclusion limits on the signal strength $\mu(tHq)$, the branching ratio $\mathcal{B}(t \to Hq)$ and the EFT Wilson coefficient $C_{u\phi}^{i3,3i}$ obtained by this analysis for both signal processes. The limit on the signal strength is based on a pre-fit signal normalisation corresponding to $\mathcal{B}(t \to Hq) = 0.1\%$.

Signal	$\mu^{(\text{lim})}(tHq)$	$\mathcal{B}(t \to Hq)^{(\text{lim})}$	$C_{u\phi}^{i3,3i\,(\text{lim})}$
tHu	0.28 (0.30)	$2.8\,(3.0) \times 10^{-4}$	0.71 (0.73)
tHc	0.33 (0.38)	$3.3\,(3.8) \times 10^{-4}$	0.76 (0.82)

single-channel analysis. A direct comparison with other ATLAS analyses, searching for identical couplings in different final states using the full Run 2 dataset, is depicted in Figure 12.14, confirming the high sensitivity of this analysis. The observed upper exclusion limits slightly exceed expectations, but lie well within the 1σ confidence interval of the expected limits. For the tHu signal, a CMS analysis searching for tHq FCNC interactions in the $H \to \gamma\gamma$ decay mode observed a pronounced downward fluctuation [65], making the limits at hand the second-most stringent observed limits for this signal in a single-channel analysis. For tHc, the limits obtained here surpass those of any other analysis, including statistical combinations of single-channel analyses.

A direct comparison of this search with its CMS counterpart [64], depicted in Figure 12.15, demonstrates a significantly higher sensitivity of this analysis. It is generally challenging to make detailed comparisons between two analyses without comprehensive knowledge of both. However, Ref. [64] provides some indications to possible sources of the different sensitivities. Firstly, the CMS analysis imposes much stricter preselection criteria, particularly regarding the minimum lepton p_{T} and 2ℓSS cuts around the Z-boson mass peak. Secondly, while this analysis defines dedicated signal regions, focusing on specific signal processes within a final state, the CMS analysis merely considers the collection of all preselected events and employs a single BDT across all of them. This approach, which heavily relies on machine learning, appears to compromise the overall analysis sensitivity. This observation affirms the analysis strategy adopted in this dissertation, demonstrating that specialised event selection and reconstruction algorithms. applied prior to a multivariate analysis, yield optimal results.

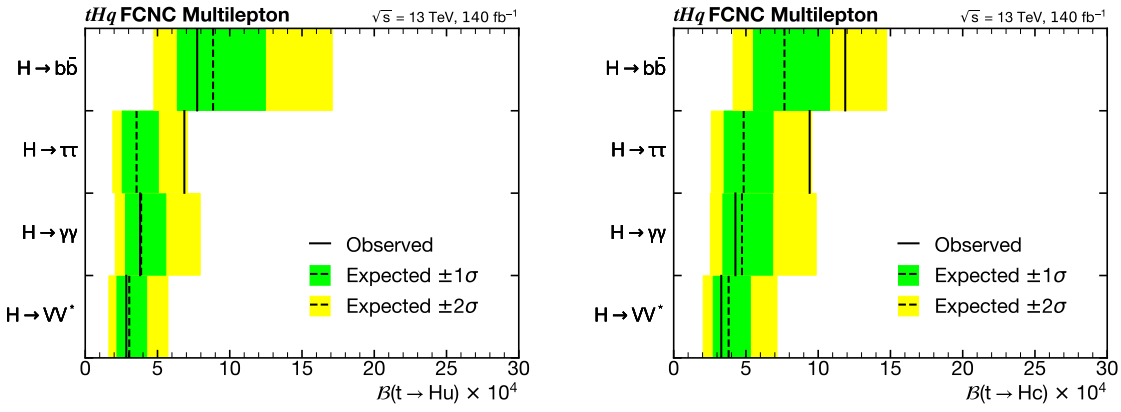

(a) (b)

Figure 12.14.: The expected and observed 95% CL upper exclusion limits on the branching ratio (a) $\mathcal{B}(t \to Hu)$ and (b) $\mathcal{B}(t \to Hc)$ of this analysis and other ATLAS tHq FCNC searches in different final states, all using the full Run 2 dataset. The results of other analyses are obtained from Refs. [66–68].

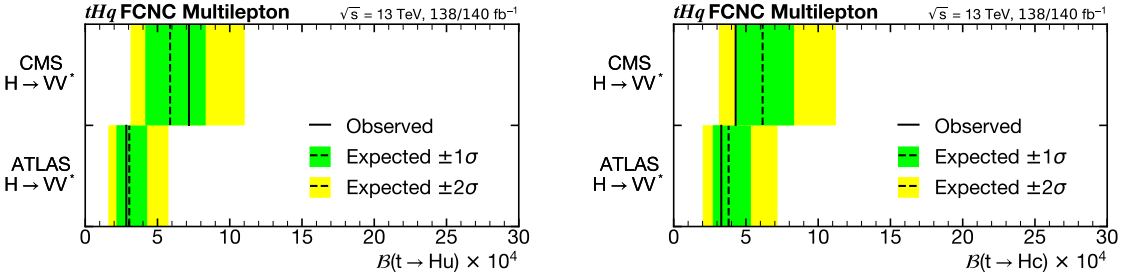

(a) (b)

Figure 12.15.: The expected and observed 95% CL upper exclusion limits on the branching ratio (a) $\mathcal{B}(t \to Hu)$ and (b) $\mathcal{B}(t \to Hc)$ of this analysis and a CMS analysis searching for the same couplings in identical final states, also using the full Run 2 dataset. The results of the CMS analysis are obtained from Ref. [64].

Conclusion

This dissertation presents an analysis searching for tHq flavour-changing neutral-current (FCNC) couplings in final states with exactly two leptons of the same charge (2ℓSS) and with three leptons, exactly two of which possess an identical charge (3ℓ), conducted with the ATLAS detector across the full Run 2 dataset of $140\,\text{fb}^{-1}$. Both, the $t\bar{t}(t \to Hq)$ decay and the $qg \to Ht$ production channel are considered. Signal processes are parametrised using an effective field theory (EFT), independent of any specific theoretical model.

A variety of methods for background estimation are employed to control specific processes. Non-prompt leptons from the decay of a b hadron (HF-decay leptons) are estimated using the template fit method, leaving the normalisation of their Monte-Carlo (MC) templates unconstrained in the final maximum-likelihood fit. In contrast, the background from electrons reconstructed with an inverted charge, relevant in the 2ℓSS final state, is estimated in a purely data-driven way, considering an expected p_{T}-$|\eta|$ dependence of the process. Moreover, the normalisations of $t\bar{t}W$ production is left free-floating to account for known MC-data discrepancies. The same is done for $t\bar{t}Z$ production, due to unknown quality of the $t\bar{t}Z$ modelling in low-N_{jets} regions. Lastly, the VV background is separated into six templates, one of which is left free-floating, owing to suboptimal b-jet modelling within the used MC samples. To constrain the normalisation of free-floating background processes, various control regions are established.

Additionally, based on kinematics of simulated events, four signal regions are defined, two per final state. Each region is aimed at either the decay or the production signal, with the primary distinction being the number of jets in an event. In the signal regions, basic kinematic variables are used as input for several event reconstruction algorithms, employed to enhance the sensitivity to the signal processes. The most significant reconstructed variables are combined into a single neural network discriminant, subsequently used as the input distribution for the maximum-likelihood fits. The fits incorporate systematic uncertainties, related to an imprecise reconstruction of physical objects and to the modelling of physical processes, through constrained nuisance parameters. A background-only fit, excluding the most signal-sensitive bins, is performed to validate the background estimation methods, revealing an overall immaculate MC-data agreement. A second, final fit considering also the signal sensitive bins of the analysis, shows similar agreement between MC predictions and data. A detailed evaluation of

the post-fit normalisation factors and the impact of systematic uncertainties on the fit results is performed, validating the efficacy of the employed statistical model.

The FCNC signal strength obtained by the full fit to data is found to be compatible with zero for both signal processes, showing no deviation from SM predictions. Consequently, 95% confidence level upper exclusion limits are imposed on the signal strength, which are subsequently converted into limits on the FCNC branching ratio and the EFT Wilson coefficient, the latter of which parametrises the strength of the FCNC coupling. The observed (expected) limits on the branching ratio amount to

$$\mathcal{B}(t \to Hu) < 2.8\,(3.0) \times 10^{-4} \text{ and } \mathcal{B}(t \to Hc) < 3.3\,(3.8) \times 10^{-4},$$

while for the Wilson coefficient

$$|C_{u\phi}^{13,31}| < 0.71\,(0.73) \text{ and } |C_{u\phi}^{23,32}| < 0.76\,(0.82)$$

are obtained. The expected upper limits of this analysis are the lowest obtained by any previous tHq FCNC analysis conducted focusing on individual Higgs-boson decay channels. Moreover, the observed limits on the tHc coupling are the strongest across all analyses to date. Thus, this analysis reaffirms the SM's validity in the domain of FCNC couplings between the top quark and the Higgs boson to a previously unattained degree.

The evaluation of the limiting factors of this analysis provides valuable insights for potential improvements, both to this and to other analyses in 2ℓSS and 3ℓ final states. The inclusion of systematic uncertainties is observed to increase the expected upper exclusion limits by merely 16%. Thus, the most straightforward improvement is the accumulation of a larger dataset to increase the overall statistical precision of the analysis. When specifically targeting the impact of systematic uncertainties on the analysis sensitivity, a clear recommendation to aim to reduce the fraction of non-prompt leptons in the analysis, mainly those from HF decay, can be expressed. In more concrete terms, the further improvement of tools such as the PLIV is paramount for minimising the impact of non-prompt leptons.

In the spirit of Albert Einstein's profound reminder that "Curiosity has its own reason for existence", this dissertation searched for fundamental interactions, whose existence is not yet indicated by any experimental findings. Although the investigation did not culminate in the discovery of tHq FCNC interactions, the stringent limits set on their branching ratios contribute valuable insights into the fundamental properties of elementary particles. Moreover, the analysis underscored the significance of 2ℓSS and 3ℓ final states as highly sensitive sectors in the search for phenomena beyond the SM. This dissertation highlights the principle that the merit of scientific exploration lies not solely

in the discovery of new phenomena but in the relentless pursuit of knowledge itself. It emphasises that even in the absence of groundbreaking revelations, the continuous quest for understanding continues to be essential to scientific progress.

> *We are driven by the usual insatiable curiosity of the scientist, and our work is a delightful game.*
>
> Murray Gell-Mann, 1969

Appendices

A. The tHq FCNC Signal Process

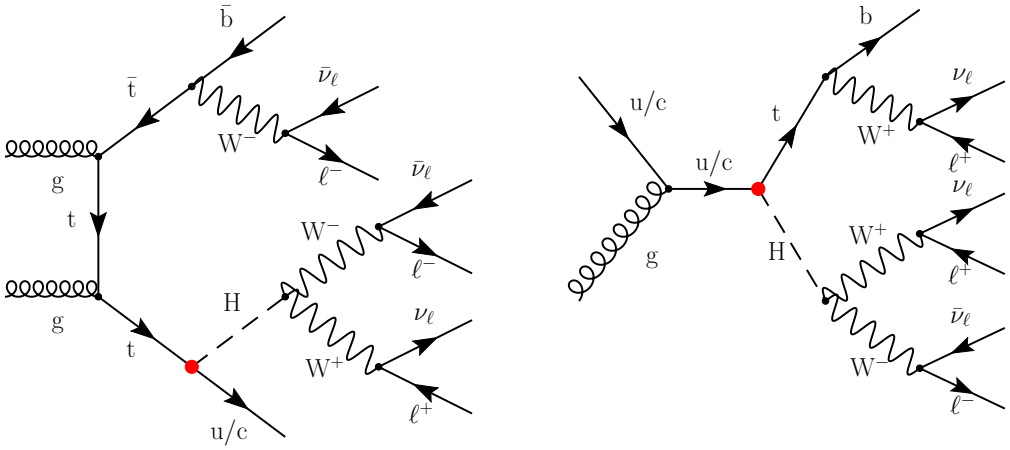

(a) $t \to Hq$ (b) $qg \to Ht$

Figure A.1.: Feynman diagrams for tHq FCNC couplings in (a) the decay and (b) the production channel resulting in 3ℓ final states via the $H \to WW^*$ decay mode. The FCNC vertex is marked in red.

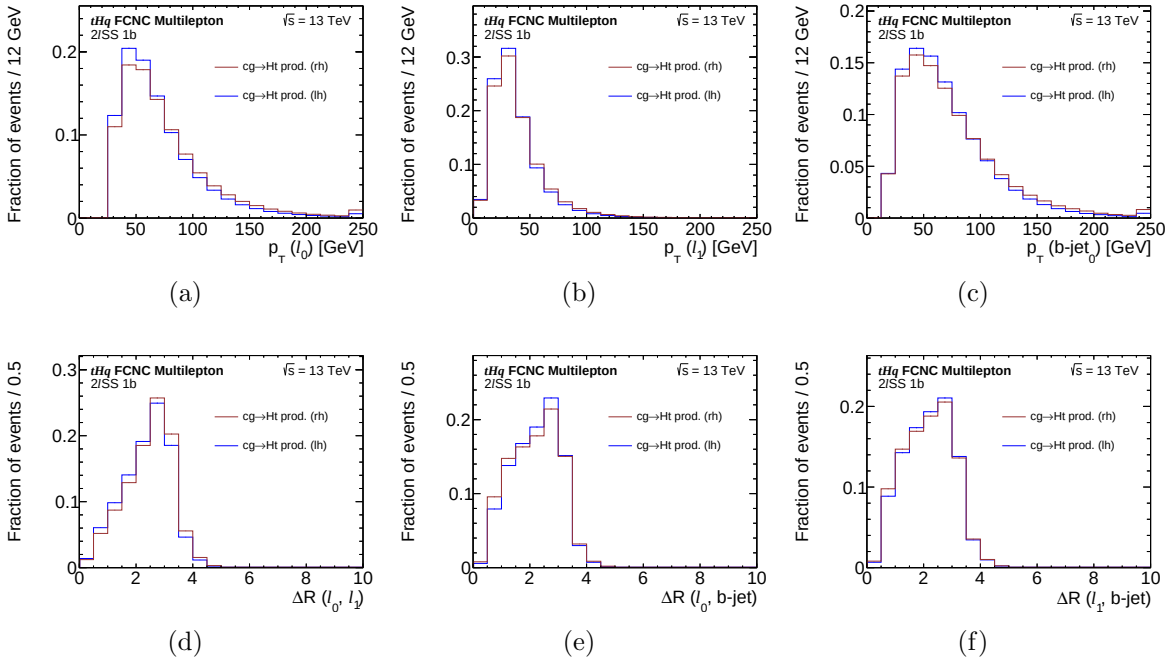

Figure A.2.: The distribution of various kinematic variables for the left-handed (lh) and the right-handed (rh) tHc production process in the 2ℓSS final state. The vertical lines on bins depict statistical uncertainties. Transverse momenta are labelled as p_{T}, while angular separation is denoted ΔR. Leptons are ordered by their transverse momentum, with ℓ_0 being the leading-p_{T} lepton. The variables depicted include: (a) $p_{\mathrm{T}}(\ell_0)$, (b) $p_{\mathrm{T}}(\ell_1)$, (c) $p_{\mathrm{T}}(b\text{-jet})$, (d) $\Delta R(\ell_0, \ell_1)$, (e) $\Delta R(\ell_0, b\text{-jet})$ and (f) $\Delta R(\ell_1, b\text{-jet})$.

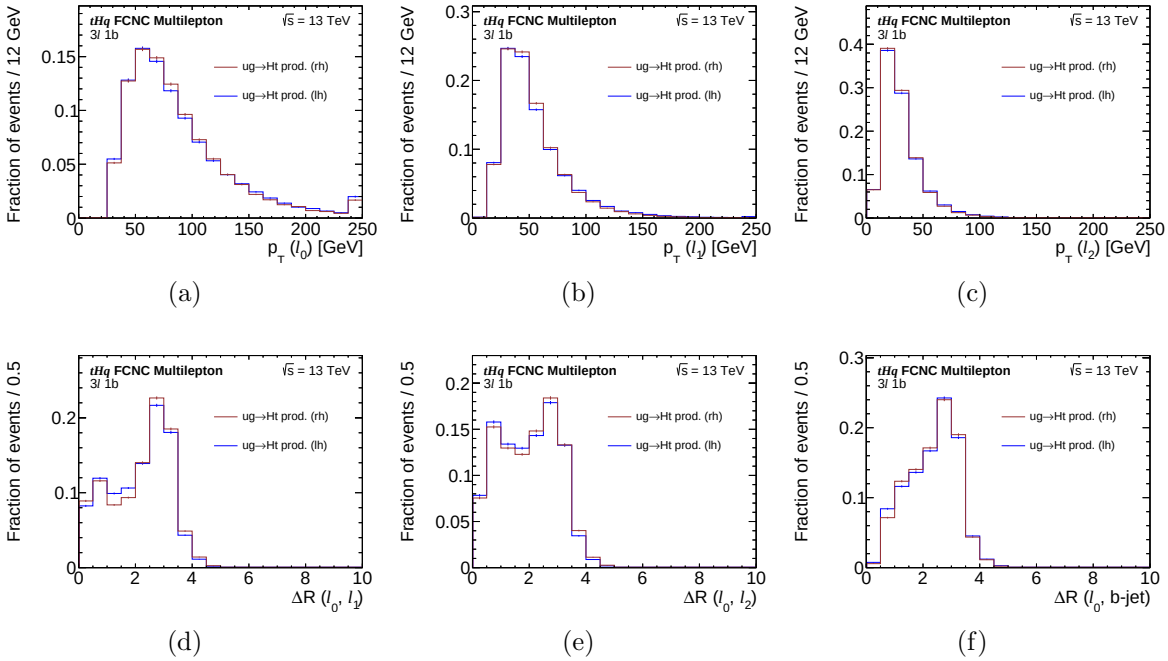

Figure A.3.: The distribution of various kinematic variables for the left-handed (lh) and the right-handed (rh) tHu production process in the 3ℓ final state. The vertical lines on bins depict statistical uncertainties. Transverse momenta are labelled as p_{T}, while angular separation is denoted ΔR. Leptons are ordered by their transverse momentum, with ℓ_0 being the leading-p_{T} lepton. The variables depicted include: (a) $p_{\mathrm{T}}(\ell_0)$, (b) $p_{\mathrm{T}}(\ell_1)$, (c) $p_{\mathrm{T}}(b\text{-jet})$, (d) $\Delta R(\ell_0, \ell_1)$, (e) $\Delta R(\ell_0, b\text{-jet})$ and (f) $\Delta R(\ell_1, b\text{-jet})$.

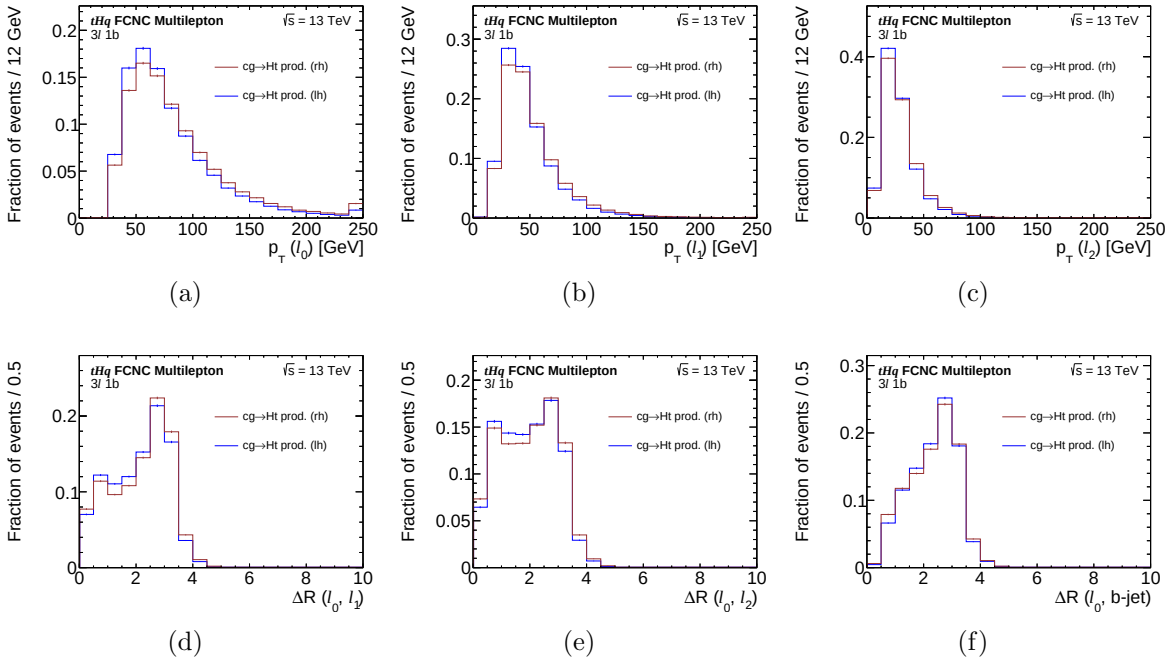

Figure A.4.: The distribution of various kinematic variables for the left-handed (lh) and the right-handed (rh) tHc production process in the 3ℓ final state. The vertical lines on bins depict statistical uncertainties. Transverse momenta are labelled as p_{T}, while angular separation is denoted ΔR. Leptons are ordered by their transverse momentum, with ℓ_0 being the leading-p_{T} lepton. The variables depicted include: (a) $p_{\mathrm{T}}(\ell_0)$, (b) $p_{\mathrm{T}}(\ell_1)$, (c) $p_{\mathrm{T}}(b\text{-jet})$, (d) $\Delta R(\ell_0, \ell_1)$, (e) $\Delta R(\ell_0, b\text{-jet})$ and (f) $\Delta R(\ell_1, b\text{-jet})$.

B. Event Selection

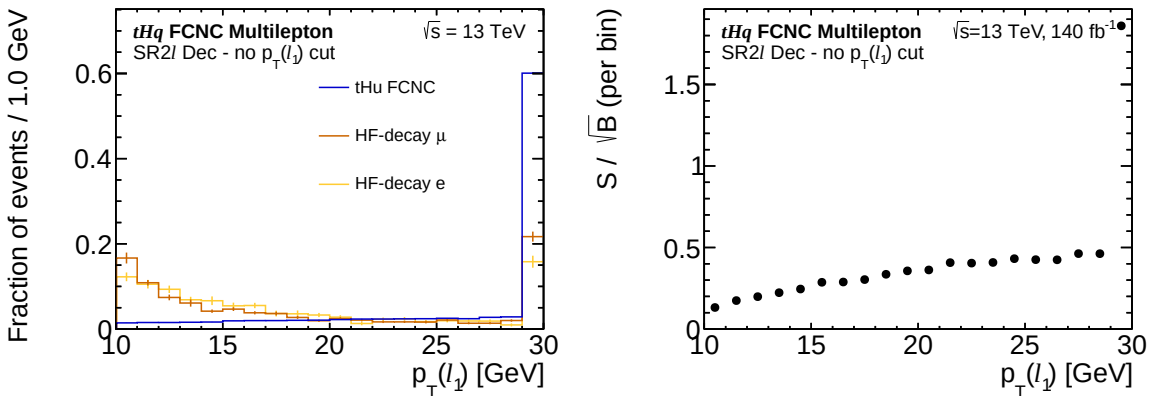

(a) (b)

Figure B.1.: The $p_{\mathrm{T}}(\ell_1)$ variable in the SR2ℓ Dec prior to the application of any cuts on $p_{\mathrm{T}}(\ell_1)$ in the range of 10 to 30 GeV. (a) The shape of the distribution for the tHu signal and the HF-decay processes and (b) the $\mathrm{S}/\sqrt{\mathrm{B}}$ fraction for the sum of all background processes are shown. For the calculation, the signal was normalised to $\mathcal{B}(t \to Hq) = 1.0$. The last bin includes all events exceeding the upper boundary of the histogram.

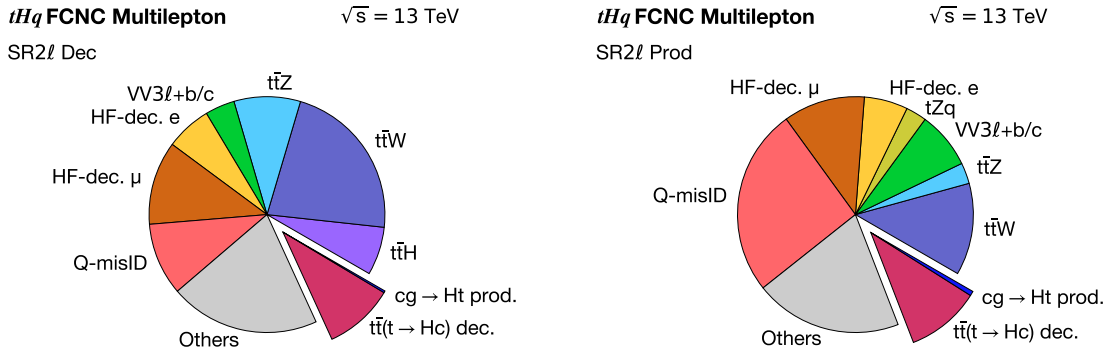

(a) (b)

Figure B.2.: Pie charts showing the background composition in (a) the SR2ℓ Dec and (b) the SR2ℓ Prod. The tHc signal contribution for a normalisation of $\mathcal{B}(t \to Hq) = 0.1\%$ is included as well. Processes with minor contributions are collected in the *Others* category.

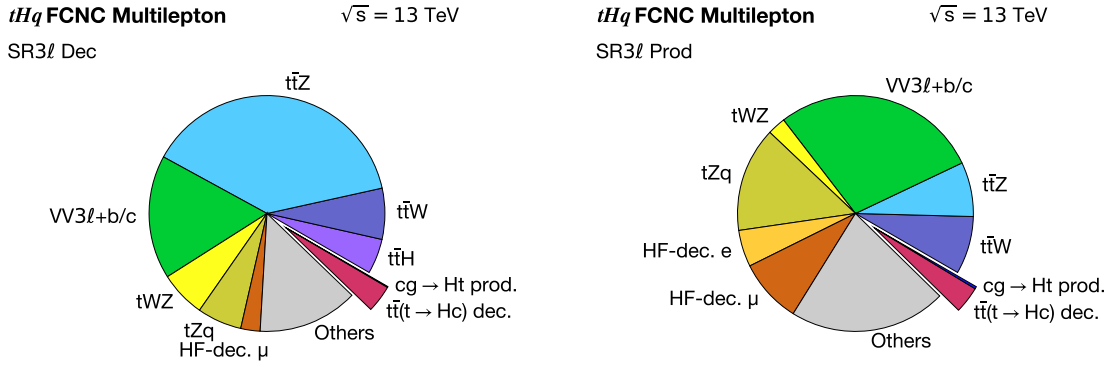

(a) (b)

Figure B.3.: Pie charts showing the background composition in (a) the SR3ℓ Dec and (b) the SR3ℓ Prod. The tHc signal contribution for a normalisation of $\mathcal{B}(t \to Hq) = 0.1\%$ is included as well. Processes with minor contributions are collected in the *Others* category.

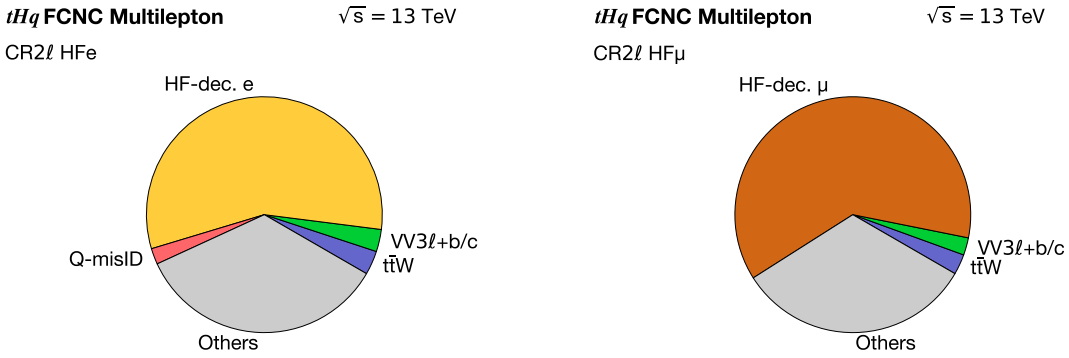

(a) (b)

Figure B.4.: Pie charts showing the background composition in (a) the CR2ℓ HFe and (b) the CR2ℓ HFμ. Processes with minor contributions are collected in the *Others* category.

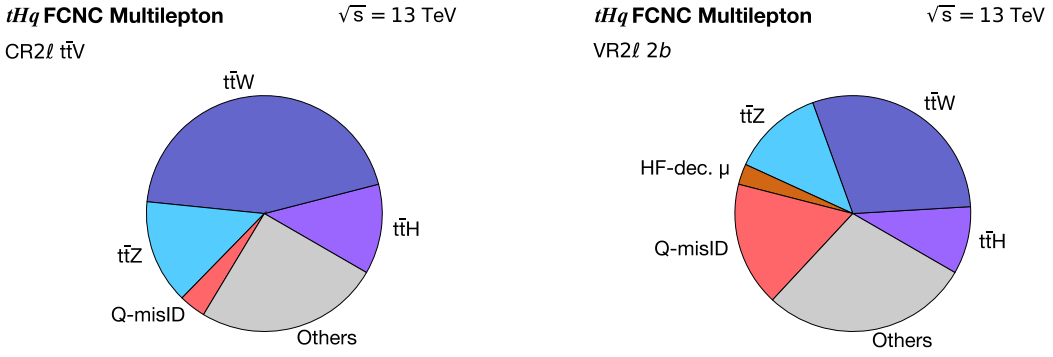

(a) (b)

Figure B.5.: Pie charts showing the background composition in (a) the CR2ℓ $t\bar{t}V$ and (b) the VR2ℓ 2b. Processes with minor contributions are collected in the *Others* category.

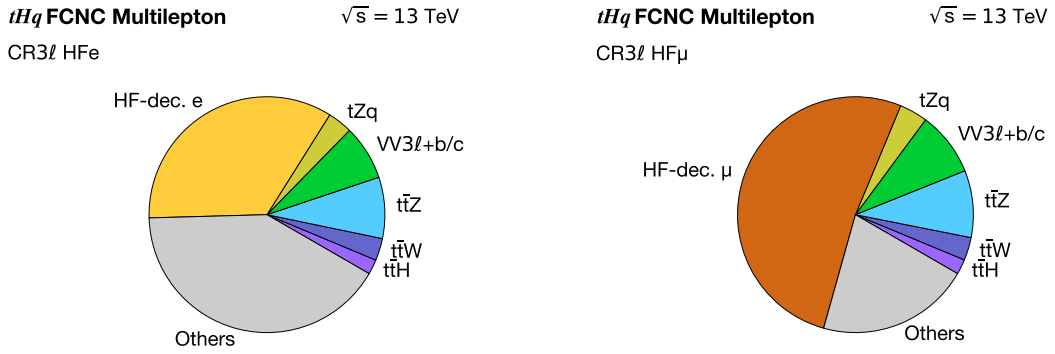

(a) (b)

Figure B.6.: Pie charts showing the background composition in (a) the CR3ℓ HFe and (b) the CR3ℓ HFμ. Processes with minor contributions are collected in the *Others* category.

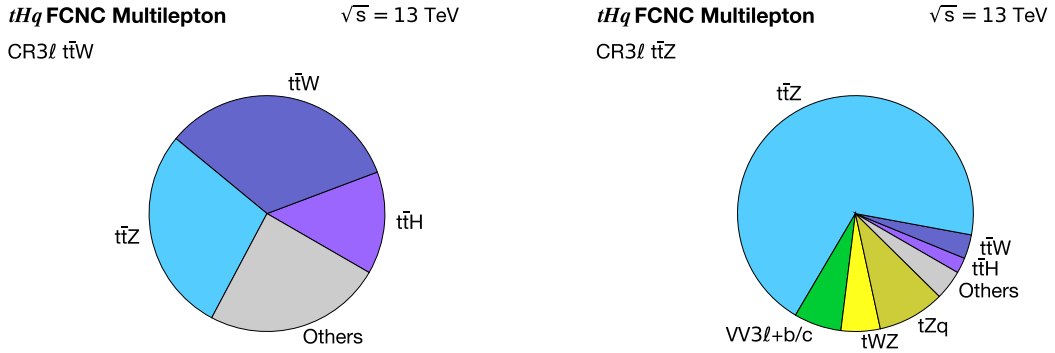

(a) (b)

Figure B.7.: Pie charts showing the background composition in (a) the CR3ℓ $t\bar{t}W$ and (b) the CR3ℓ $t\bar{t}Z$. Processes with minor contributions are collected in the *Others* category.

Table B.1.: The predicted and observed pre-fit yields in all 2ℓSS regions of the analysis. The pre-fit predictions for the two signal components is presented as well, scaled to a branching ratio $\mathcal{B}(t \to Hq) = 0.1\%$ and separated into the decay and production process. Minor MC templates are combined in the *Others* category. The uncertainties on MC yields reflect the quadratic sum of all systematic and statistical uncertainties.

Process	SR2ℓ Dec	SR2ℓ Prod	CR2ℓ HFe	CR2ℓ HFμ	CR2ℓ $t\bar{t}V$	VR2ℓ 2b
HF-decay e	128 ± 10	106 ± 6	63 ± 6	-	2.6 ± 0.6	7.3 ± 1.3
HF-decay μ	235 ± 21	199 ± 10	0.1 ± 0.02	130 ± 18	6.7 ± 1.0	16.6 ± 1.5
Q-misID	204 ± 18	456 ± 40	2.4 ± 0.3	-	15.6 ± 1.6	100 ± 11
$t\bar{t}H$	135 ± 21	24 ± 5	0.5 ± 0.1	0.9 ± 0.2	52 ± 9	54 ± 9
$t\bar{t}W$	453 ± 29	224 ± 19	3.6 ± 0.6	5.9 ± 0.8	188 ± 10	174 ± 12
$t\bar{t}Z$	187 ± 8	51 ± 5	1.1 ± 0.2	1.8 ± 0.3	61 ± 5	74 ± 5
$VV3\ell + b/c$	83 ± 9	139 ± 6	3.4 ± 0.4	5.0 ± 0.7	4.7 ± 0.5	5.4 ± 0.5
tWZ	25 ± 8	12 ± 4	0.11 ± 0.04	0.2 ± 0.1	3.9 ± 1.2	4.6 ± 1.4
tZq	23 ± 7	51 ± 15	0.6 ± 0.2	0.9 ± 0.3	5.2 ± 1.6	7 ± 2
Others	372 ± 77	323 ± 50	36 ± 9	65 ± 27	84 ± 16	144 ± 36
Total BG	1845 ± 91	1585 ± 70	111 ± 11	210 ± 32	424 ± 22	587 ± 41
$t\bar{t}(t \to Hu)$	207 ± 22	181 ± 10	3.4 ± 0.3	5.4 ± 0.7	6.8 ± 0.6	7.0 ± 0.7
$ug \to Ht$	31 ± 4	68 ± 2	1.2 ± 0.1	2.1 ± 0.2	1.1 ± 0.1	1.1 ± 0.1
$t\bar{t}(t \to Hc)$	196 ± 22	180 ± 10	3.5 ± 0.4	5.9 ± 0.7	13.4 ± 1.5	13.8 ± 1.7
$cg \to Ht$	5 ± 1	11 ± 1	0.2 ± 0.1	0.4 ± 0.1	0.2 ± 0.1	0.23 ± 0.05
Data	1847	1723	116	193	443	647

Table B.2.: The predicted and observed pre-fit yields in all 3ℓ regions of the analysis. The pre-fit predictions for the two signal components is presented as well, scaled to a branching ratio $\mathcal{B}(t \to Hq) = 0.1\%$ and separated into the decay and production process. Minor MC templates are combined in the *Others* category. The uncertainties on MC yields reflect the quadratic sum of all systematic and statistical uncertainties.

Process	SR3ℓ Dec	SR3ℓ Prod	CR3ℓ HFe	CR3ℓ HFμ	CR3ℓ $t\bar{t}W$	CR3ℓ $t\bar{t}Z$
HF-decay e	14.2 ± 1.5	37 ± 5	51 ± 5	-	1.3 ± 0.1	0.29 ± 0.06
HF-decay μ	26 ± 3	65 ± 4	0.2 ± 0.1	128 ± 11	1.8 ± 0.1	0.41 ± 0.06
Q-misID	-	-	-	-	-	-
$t\bar{t}H$	47 ± 7	8.5 ± 1.6	3.0 ± 0.6	5.1 ± 0.9	32 ± 5	6.7 ± 1.1
$t\bar{t}W$	68 ± 5	59 ± 7	4.5 ± 0.7	7.7 ± 0.7	76 ± 13	10.2 ± 0.8
$t\bar{t}Z$	378 ± 11	55 ± 8	12.4 ± 1.1	23 ± 2	64 ± 5	216 ± 16
$VV3\ell + b/c$	166 ± 13	210 ± 11	11.2 ± 1.1	22 ± 2	3.6 ± 0.3	20.3 ± 1.2
tWZ	62 ± 19	18 ± 6	2.0 ± 0.6	3.7 ± 1.1	3.0 ± 0.9	17 ± 5
tZq	60 ± 18	106 ± 32	5.0 ± 1.6	10 ± 3	2.9 ± 0.9	29 ± 9
Others	120 ± 26	151 ± 31	59 ± 10	48 ± 10	43 ± 8	11.9 ± 1.7
Total BG	941 ± 42	710 ± 48	148 ± 11	248 ± 16	228 ± 17	312 ± 19
$t\bar{t}(t \to Hu)$	39 ± 3	26.2 ± 1.9	4.6 ± 0.5	8.0 ± 0.8	1.2 ± 0.2	0.7 ± 0.1
$ug \to Ht$	6.5 ± 0.7	13.4 ± 0.7	1.2 ± 0.1	2.1 ± 0.2	0.36 ± 0.03	0.16 ± 0.02
$t\bar{t}(t \to Hc)$	37 ± 3	27 ± 2	4.4 ± 0.5	7.8 ± 0.7	4.3 ± 0.4	1.5 ± 0.1
$cg \to Ht$	1.1 ± 0.2	2.4 ± 0.3	0.2 ± 0.1	0.4 ± 0.1	0.08 ± 0.01	0.04 ± 0.01
Data	1046	896	159	263	268	381

C. Separation of Signal and Background

Table C.1.: An overview of the different NeuroBayes hyperparameters that were varied in order to optimise the NN training. A short explanation of each hyperparameter is provided, together with the values the parameter was set to in the grid search.

Parameter	Description	Grid-search values
`Speed`	A multiplicative factor modifying the magnitude of the weight update in each training step.	0.5, 4, 8, 25, 50
`Momentum`	An additional term added to the weight update, proportional to the previous training step's update.	0.25, 0.5, 0.75, 0.8, 0.85, 0.9
`MaxLearn`	The maximum learning value by which a weight is updated.	1, 3, 5
$\texttt{N}_{\texttt{HiddenNodes}}$	The number of nodes in the hidden layer of the NN.	10, 20, 25, 30

Table C.2.: List of input variables to the NN in the SR3ℓProd, approximately ordered by the increase in significance provided by each variable. The exact order differs among the two signal processes.

Variable	Description
$m(\ell_{\mathrm{OS}}, \ell_{\mathrm{SS},1})$	Invariant mass of the opposite-charge and the subleading-p_{T} same-charge lepton
$m(\ell_{\mathrm{OS}}, \ell_{\mathrm{SS},0})$	Invariant mass of the opposite-charge and the leading-p_{T} same-charge lepton
$m(\ell_t, b_t)$	Invariant mass of the b-tagged jet and the lepton assigned to the top-quark decay
N_{jets}	The number of jets
$H_T(\mathrm{jets})$	Scalar p_{T}-sum of all jets
$m(t_{\mathrm{SM}}, H)$	Invariant mass of the RJR top quark decaying via $t \to Wb$ and the Higgs boson
$\Delta R(\ell_{\mathrm{SS},0}, \ell_{\mathrm{SS},1})$	Angular separation between the leading and subleading-p_{T} same-charge lepton
$m(\ell_{H,0}, \ell_{H,1})$	Invariant mass of the two leptons assigned to the Higgs-boson decay
$m(b\text{-jet}, \ell_{\mathrm{SS},0})$	Invariant mass of the b-tagged jet and the leading-p_{T} same-charge lepton
$\Delta R(\ell_t, b_t)$	Angular separation between the b-tagged jet and the lepton assigned to the top-quark decay
$p_{\mathrm{T}}(t_{\mathrm{SM}})$	Transverse momentum of the RJR top quark decaying via $t \to Wb$
$p_{\mathrm{T}}(b\text{-jet})$	Transverse momentum of the b-tagged jet
$\eta(\ell_{\mathrm{SS},1})$	Pseudorapidity of the subleading-p_{T} same-charge lepton
$p_{\mathrm{T}}(\ell_{\mathrm{SS},1})$	Transverse momentum of the subleading-p_{T} same-charge lepton
$m(H, \ell_{\mathrm{SS},1})$	Invariant mass of the RJR Higgs boson and the subleading-p_{T} same-charge lepton
$\Delta R(t_{\mathrm{SM}}, \ell_{\mathrm{OS}})$	Angular separation between the RJR top quark decaying via $t \to Wb$ and the opposite-charge lepton
$\Delta R(H, \ell_{\mathrm{OS}})$	Angular separation between the RJR Higgs boson and the opposite-charge lepton
$\Delta R(\ell_{\mathrm{OS}, \ell_{\mathrm{SS},1}})$	Angular separation between the opposite-charge and the subleading-p_{T} same-charge lepton

Table C.3.: List of input variables to the NN in the SR3ℓ Dec, approximately ordered by the increase in significance provided by each variable. The exact order differs among the two signal processes. Variables labelled NICE were reconstructed with a fulfilled ReallyNICE Reco condition.

Variable	Description
$m(\ell_{\text{OS}}, \ell_{\text{SS},1})$	Invariant mass of the opposite-charge and the subleading-p_{T} same-charge lepton
$m(\ell_{\text{OS}}, \ell_{\text{SS},0})$	Invariant mass of the opposite-charge and the leading-p_{T} same-charge lepton
NICE $m(\ell_t, b_t)$	Invariant mass of the b-tagged jet and the lepton assigned to the top-quark decay with a fulfilled NICE Reco condition
$H_T(\text{jets})$	Scalar p_{T}-sum of all jets
$m(b\text{-jet}, \ell_{\text{SS},0})$	Invariant mass of the b-tagged jet and the leading-p_{T} same-charge lepton
$m(t_{\text{SM}}, H)$	Invariant mass of the RJR top quark decaying via $t \to Wb$ and the RJR Higgs boson
$m(\ell_{H,0}, \ell_{H,1})$	Invariant mass of the two leptons assigned to the Higgs-boson decay
$m(H, \ell_{\text{SS},1})$	Invariant mass of the RJR Higgs boson and the subleading-p_{T} same-charge lepton
$\Delta R(b\text{-jet}, t_{\text{SM}})$	Angular separation between the b-tagged jet and the RJR top quark decaying via $t \to Wb$
$m(\ell_0, t_{\text{SM}})$	Invariant mass of the leading-p_{T} lepton and the RJR top quark decaying via $t \to Wb$
$p_{\text{T}}(t_{\text{SM}})$	Transverse momentum of the RJR top quark decaying via $t \to Wb$
$m(t_{\text{SM}}, \ell_{\text{SS},1})$	Invariant mass of the RJR top quark decaying via $t \to Wb$ and the subleading-p_{T} same-charge lepton
$\Delta R(\ell_{\text{OS}}, \ell_{\text{SS},0})$	Angular separation between the opposite-charge and the leading-p_{T} same-charge lepton
$p_{\text{T}}(\ell_{\text{OS}})$	Transverse momentum of the opposite-charge lepton
$m(b\text{-jet}, \ell_{\text{OS}})$	Invariant mass of the b-tagged jet and the opposite-charge lepton
$m(b\text{-jet}, H)$	Invariant mass of the b-tagged jet and the RJR Higgs boson
$p_{\text{T}}(\ell_2)$	Transverse momentum of the third-leading-p_{T} lepton
$\eta(\ell_0)$	Pseudorapidity of the leading-p_{T} lepton
$m(W_t)$	Mass of the RJR W boson from the top-quark decay
$m(\ell_t, b_t)$	Invariant mass of the b-tagged jet and the lepton assigned to the top-quark decay

Table C.4.: List of input variables to the NN in the SR2ℓProd, approximately ordered by the increase in significance provided by each variable. The exact order differs among the various signal processes.

Variable	Description
$m(\ell_1, H)$	Invariant mass of the subleading-p_{T} lepton and the RJR Higgs boson
N_{jets}	The number of jets
$m(b\text{-jet}, t_{\mathrm{SM}})$	Invariant mass of the b-tagged jet and the RJR top quark decaying via $t \to Wb$
$m(H, b\text{-jet})$	Invariant mass of the RJR Higgs boson and the b-tagged jet
$p_{\mathrm{T}}(W_{\mathrm{had}})$	Transverse momentum of the hadronically decaying RJR W boson
$\Delta R(\ell_1, H)$	Angular separation between the subleading-p_{T} lepton and the RJR Higgs boson
$m(W_{\mathrm{had}})$	Mass of the hadronically decaying RJR W boson
$p_{\mathrm{T}}(\ell_1)$	Transverse momentum of the subleading-p_{T} lepton
$\eta(\ell_1)$	Pseudorapidity of the subleading-p_{T} lepton
$\Delta R(H, W_t)$	Angular separation between the RJR Higgs boson and the RJR W boson from the top-quark decay
$\Delta R(\ell_0, \ell_1)$	Angular separation between leading and subleading-p_{T} lepton
$m(\ell_1, b\text{-jet})$	Invariant mass of the subleading-p_{T} lepton and the b-tagged jet
$\eta(b\text{-jet})$	Pseudorapidity of the b-tagged jet
$\Delta R(\ell_0, t_{\mathrm{SM}})$	Angular separation between the leading-p_{T} lepton and the RJR top quark decaying via $t \to Wb$
$E_{\mathrm{T}}^{\mathrm{miss}}$	Missing transverse momentum
fl.(ℓ_0)	Flavour of the leading-p_{T} lepton
$\eta(\ell_0)$	Pseudorapidity of the leading-p_{T} lepton
$p_{\mathrm{T}}(\ell_0)$	Transverse momentum of the leading-p_{T} lepton
$\Delta R(\ell_1, t_{\mathrm{SM}})$	Angular separation between the subleading-p_{T} lepton and the RJR top quark decaying via $t \to Wb$
$m(H, W_t)$	Invariant mass of the RJR Higgs boson and the RJR W boson from the top-quark decay
$\Delta R(\ell_1, W_t)$	Angular separation between the subleading-p_{T} lepton and the RJR W boson from the top-quark decay
$m(\ell_0, H)$	Invariant mass of the leading-p_{T} lepton and the RJR Higgs boson
$p_{\mathrm{T}}(b\text{-jet})$	Transverse momentum of the b-tagged jet

Table C.5.: List of input variables to the NN in the SR2ℓ Dec, approximately ordered by the increase in significance provided by each variable. The exact order differs among the various signal processes.

Variable	Description
H_T(jets)	Scalar p_{T}-sum of all jets
$m(\ell_0, b\text{-jet})$	Invariant mass of the leading-p_{T} lepton and the b-tagged jet
$\Delta R(\ell_1, H)$	Angular separation between the subleading-p_{T} lepton and the RJR Higgs boson
$p_{\mathrm{T}}(\ell_1)$	Transverse momentum of the subleading-p_{T} lepton
$m(\mathrm{jets}_{\mathrm{min}\Delta R})$	Invariant mass of the two non-b-tagged jets with the smallest ΔR
$m(t_{\mathrm{SM}}, l\text{-jet}_0)$	Invariant mass of the RJR top quark decaying via $t \to Wb$ and the leading-p_{T} non-b-tagged jet
$\eta(\ell_1)$	Pseudorapidity of the subleading-p_{T} lepton
$\Delta R(\ell_0, l\text{-jet}_1)$	Angular separation between the leading-p_{T} lepton and the subleading-p_{T} non-b-tagged jet
$m(\ell_1, l\text{-jet}_0)$	Invariant mass of the subleading-p_{T} lepton and the leading-p_{T} non-b-tagged jet
$m(\ell_0, l\text{-jet}_0)$	Invariant mass of the leading-p_{T} lepton and the leading-p_{T} non-b-tagged jet
$\Delta R(\ell_0, l\text{-jet}_2)$	Angular separation between the leading-p_{T} lepton and the third-leading-p_{T} non-b-tagged jet
$\Delta R(\ell_1, l\text{-jet}_2)$	Angular separation between the subleading-p_{T} lepton and the third-leading-p_{T} non-b-tagged jet
$m(t_{\mathrm{FCNC}}, l\text{-jet}_0)$	Invariant mass of the RJR top quark decaying via $t \to Hq$ and the leading-p_{T} non-b-tagged jet
$m(\ell_1, l\text{-jet}_1)$	Invariant mass of the subleading-p_{T} lepton and the subleading-p_{T} non-b-tagged jet
$m(\ell_1, t_{\mathrm{FCNC}})$	Invariant mass of the subleading-p_{T} lepton and the RJR top quark decaying via $t \to Hq$
$m(W_t, W_{\mathrm{had}})$	Invariant mass of the RJR W boson from the top-quark decay and the hadronically decaying RJR W boson
$\Delta R(\ell_0, l\text{-jet}_0)$	Angular separation between the leading-p_{T} lepton and the leading-p_{T} non-b-tagged jet
$m(\ell_1, b\text{-jet})$	Invariant mass of the subleading-p_{T} lepton and the b-tagged jet
N_{jets}	The number of jets
$m(H, b\text{-jet})$	Invariant mass of the RJR Higgs boson and the b-tagged jet
$H_T(\ell_0, \ell_1)$	Scalar p_{T}-sum of all leptons
$p_{\mathrm{T}}(\ell_0)$	Transverse momentum of the leading-p_{T} lepton
$m(W_t, t_{\mathrm{FCNC}})$	Invariant mass of the RJR W boson from the top-quark decay and the RJR top quark decaying via $t \to Hq$

Table C.6.: The optimised hyperparameter values obtained through the grid search, shown individually for the NNs trained in each region and for each signal.

Region	Signal	`Speed`	`Momentum`	`MaxLearn`	$\texttt{N}_{\texttt{HiddenNodes}}$
SR2ℓ Dec	tHu	50	0.85	3	30
	tHc	50	0.85	3	25
SR2ℓ Prod	tHu	25	0.85	1	30
	tHc	50	0.85	3	25
SR3ℓ Dec	tHu	50	0.8	3	30
	tHc	50	0.9	1	30
SR3ℓ Prod	tHu	50	0.85	3	30
	tHc	25	0.8	3	20

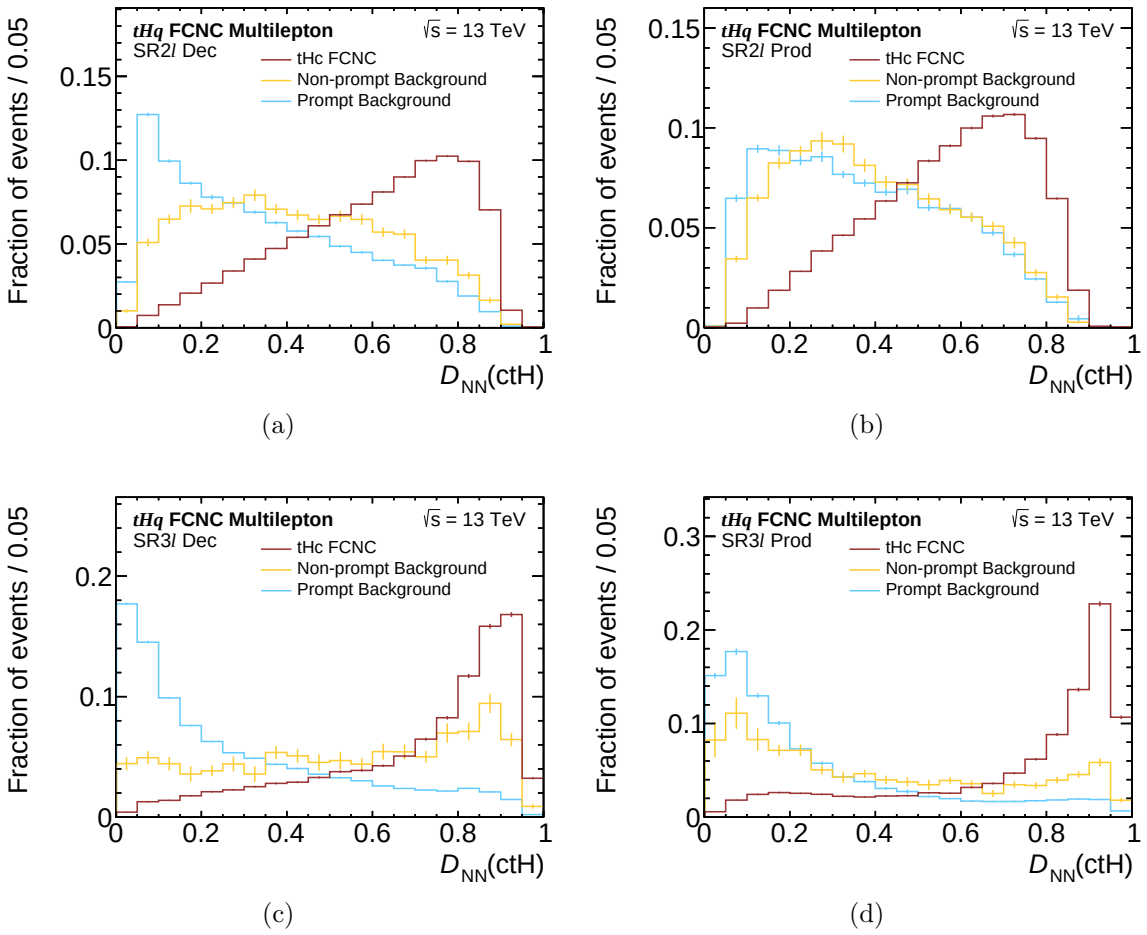

Figure C.1.: The distribution of the tHc-trained NN output variable D_{NN} in (a) the SR2ℓ Dec, (b) the SR2ℓ Prod, (c) the SR3ℓ Dec and (d) the SR3ℓ Prod for the tHc signal process and the combination of all prompt and non-prompt background processes. The distributions are obtained by applying the even-trained NN to events with an odd event number and vice versa, and adding the two resulting distributions.

D. Systematic Uncertainties

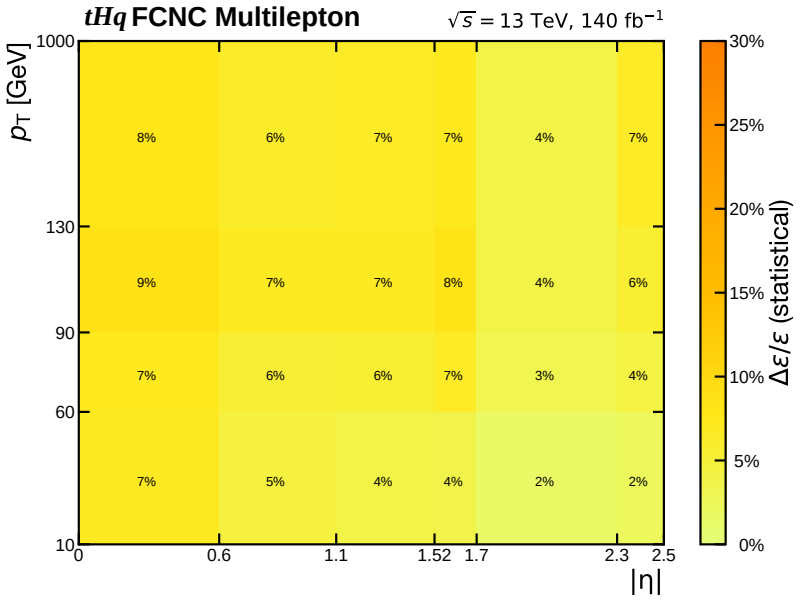

Figure D.1.: Statistical uncertainty on the Q-misID efficiencies, obtained by from the maximum-likelihood fit for their determination.

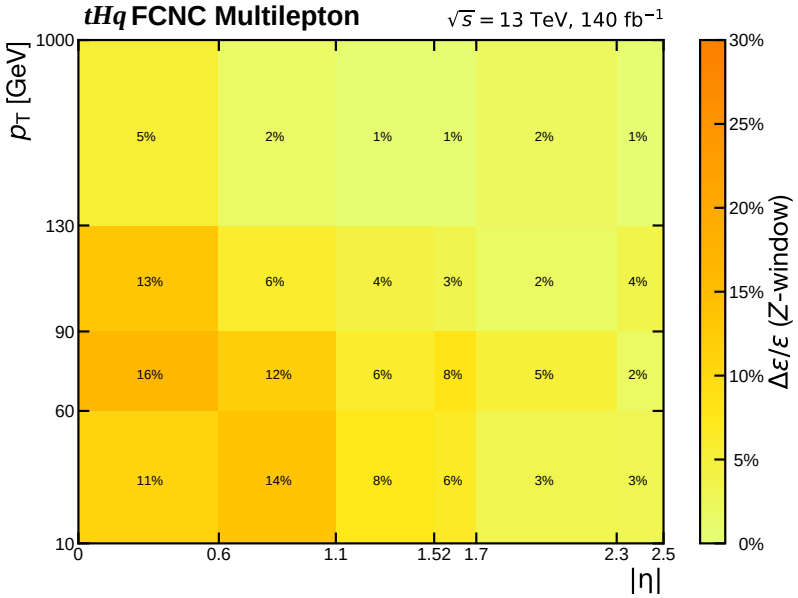

Figure D.2.: Uncertainty on the Q-misID efficiencies related to choice of the Z-boson mass window, obtained by varying the width of this window by $\pm 1\sigma$.

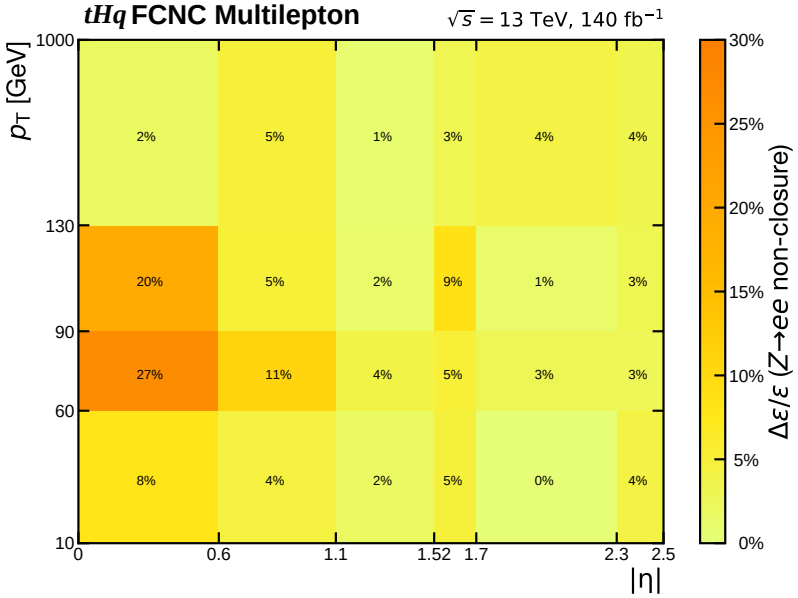

Figure D.3.: Non-closure uncertainty on the Q-misID efficiencies, obtained by comparing the nominal Q-misID efficiencies obtained with data to alternative ones obtained using $Z \to e^+e^-$ MC events.

E. Results of the Statistical Analysis

Results of the Background-Only Fit

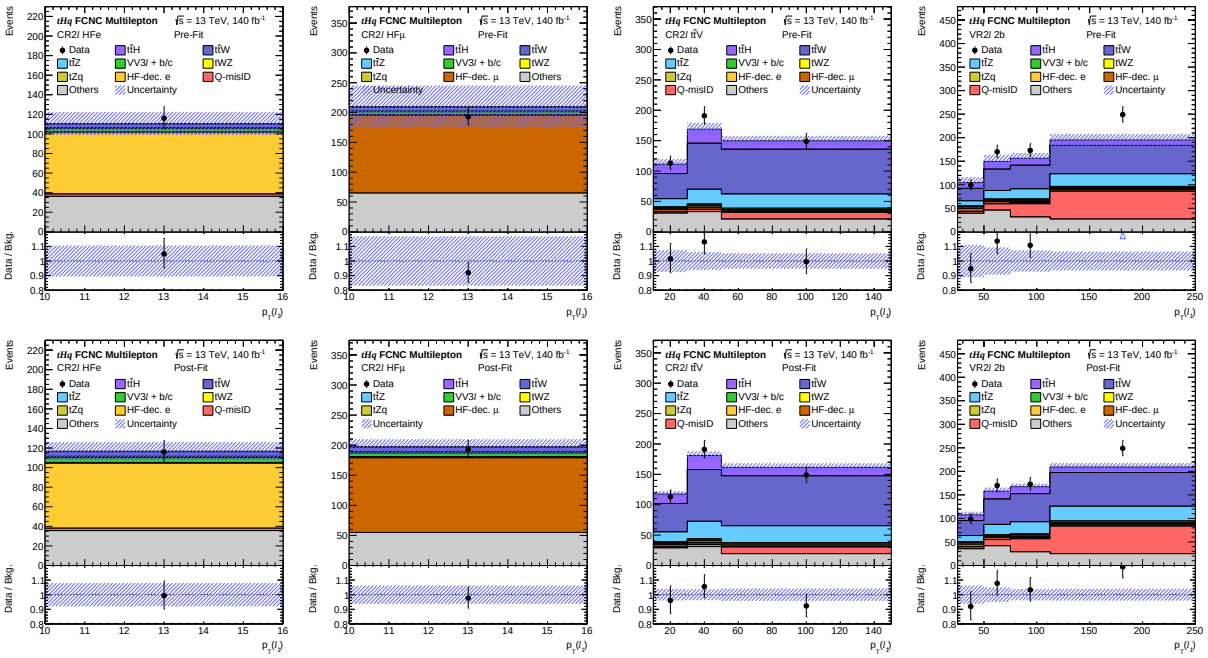

Figure E.1.: Pre-fit distributions (upper row) and post-fit distributions (lower row) of the three control regions and the validation region in the 2ℓSS final state for the tHu background-only fit to data.

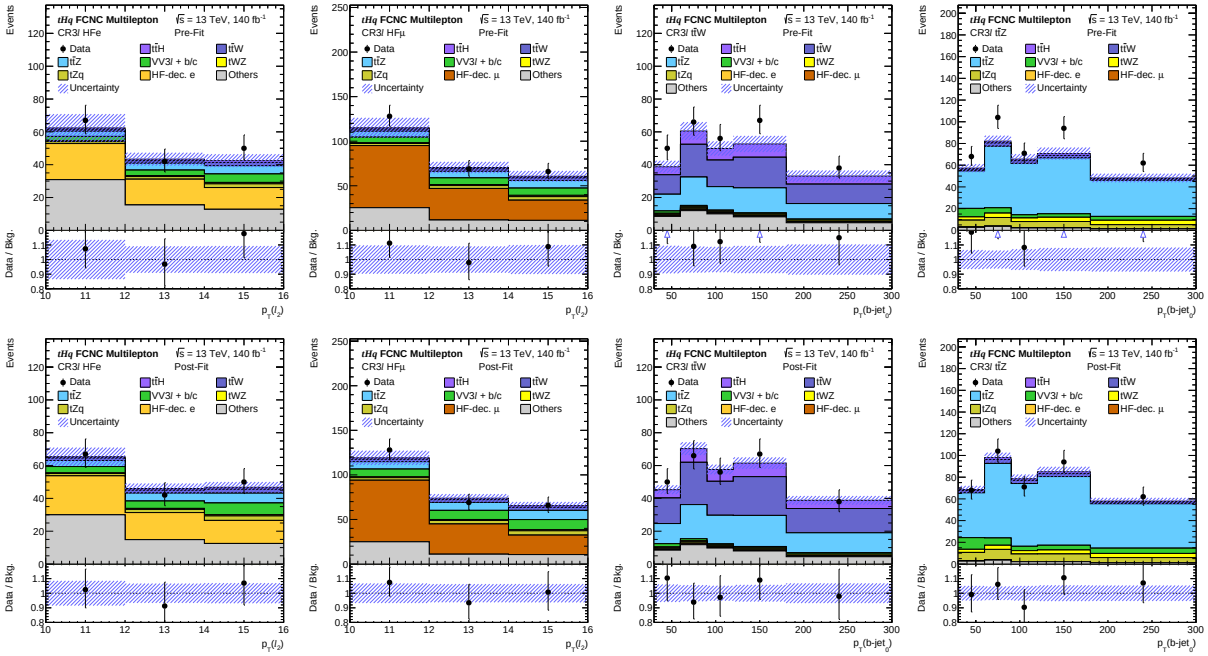

Figure E.2.: Pre-fit distributions (upper row) and post-fit distributions (lower row) of the control regions in the 3ℓ final state for the tHu background-only fit to data.

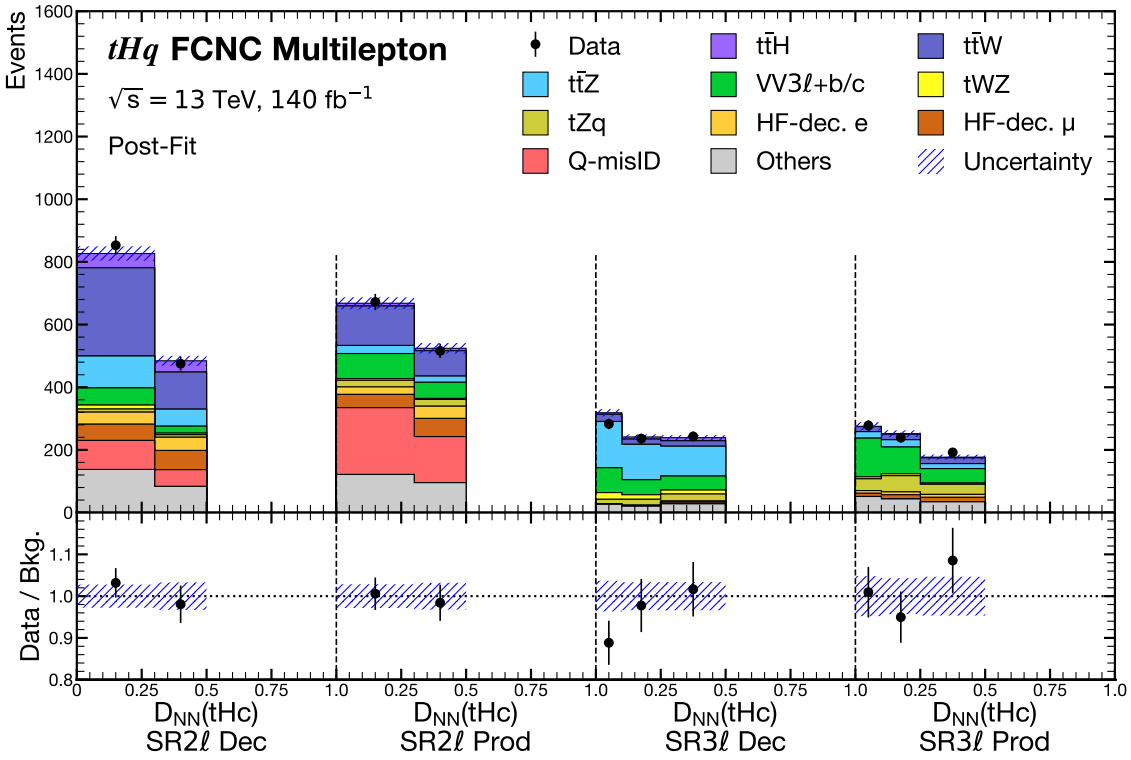

Figure E.3.: The post-fit agreement between MC and data for the background-only for the low-D_{NN} bins of all signal regions. The $D_{\mathrm{NN}}(tHc)$ distributions were used in the fit. The shaded area marks all MC-statistical and systematic uncertainties.

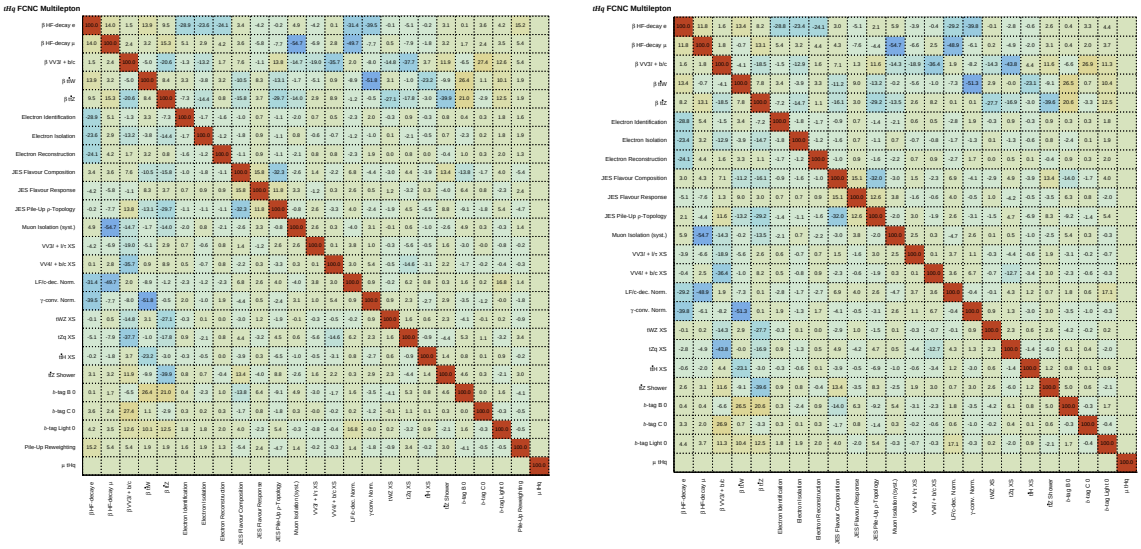

(a) (b)

Figure E.4.: NP correlation matrices for (a) the tHu and (b) the tHc background-only fit. Only NPs with a minimum correlation of 0.25 to at least one other NP are shown.

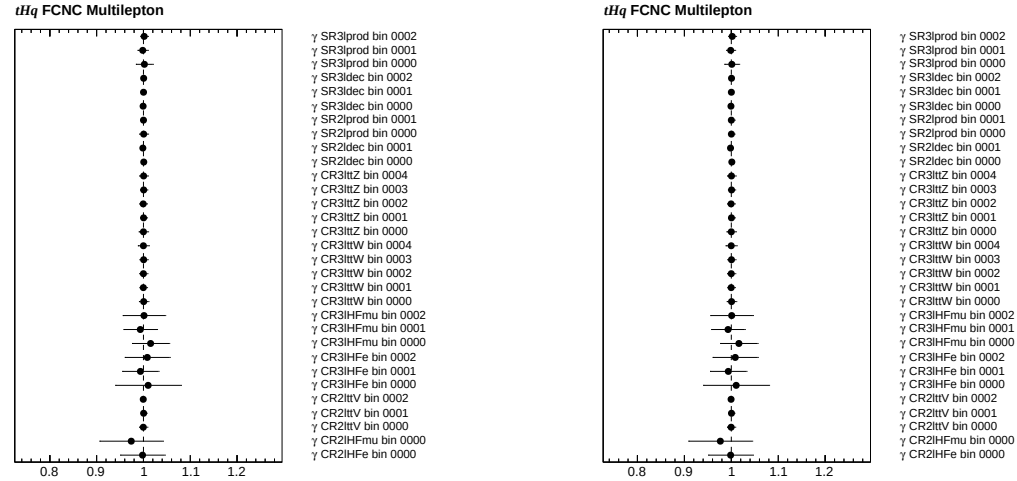

(a) (b)

Figure E.5.: γ-factors for (a) the tHu and (b) the tHc background-only fit.

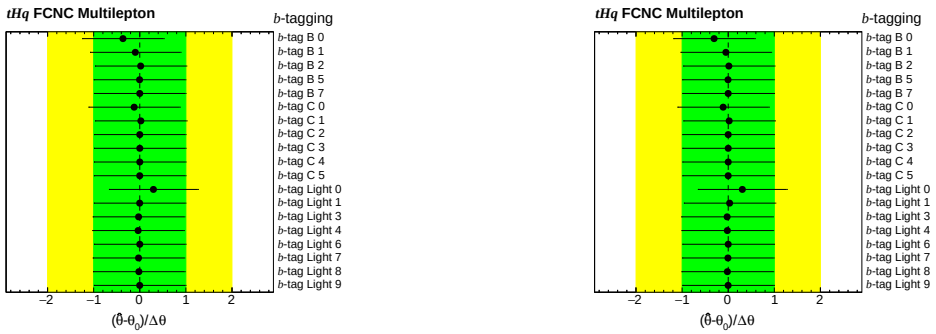

(a) (b)

Figure E.6.: Post-fit values of b-tagging related NPs for (a) the tHu and (b) the tHc background-only fit. Pruned NPs are excluded from the plots.

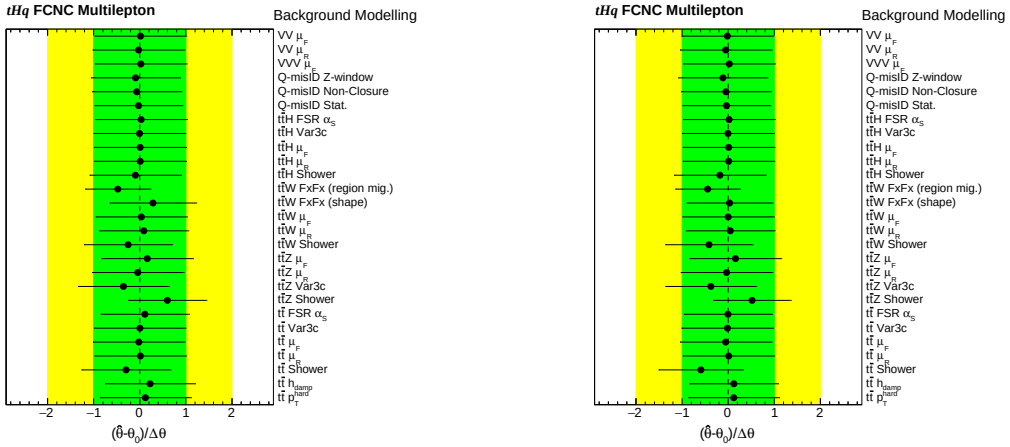

(a) (b)

Figure E.7.: Post-fit values of background-modelling related NPs for (a) the tHu and (b) the tHc background-only fit. Pruned NPs are excluded from the plots.

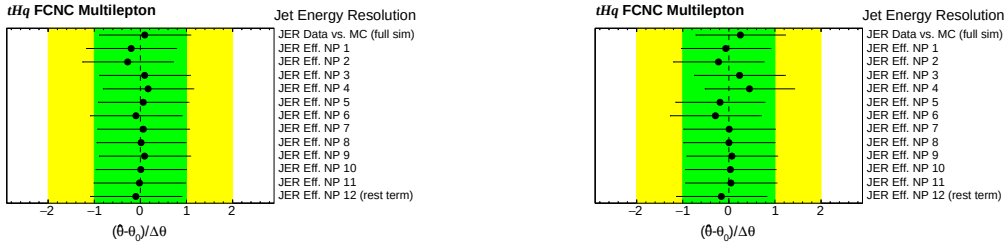

(a) (b)

Figure E.8.: Post-fit values of JER related NPs for (a) the tHu and (b) the tHc background-only fit. Pruned NPs are excluded from the plots.

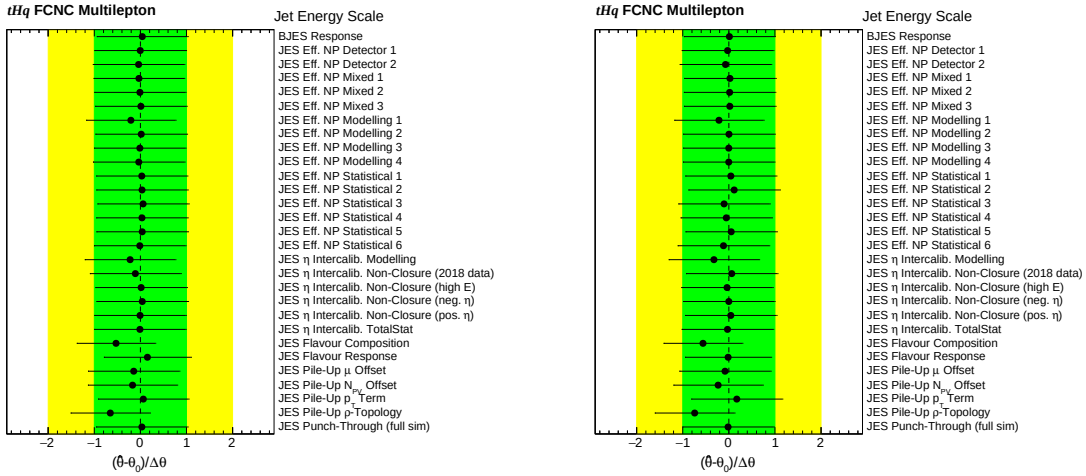

(a) (b)

Figure E.9.: Post-fit values of JES related NPs for (a) the tHu and (b) the tHc background-only fit. Pruned NPs are excluded from the plots.

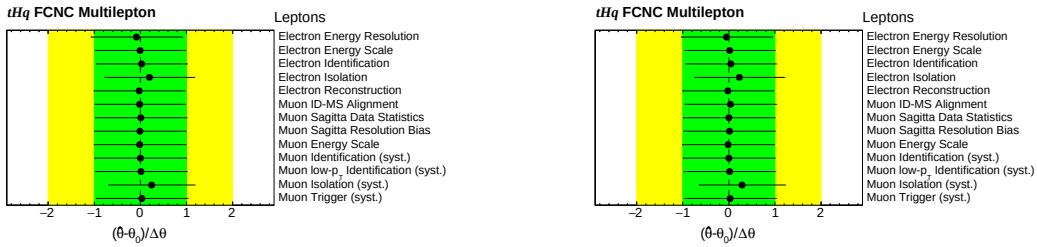

(a) (b)

Figure E.10.: Post-fit values of lepton related NPs for (a) the tHu and (b) the tHc background-only fit. Pruned NPs are excluded from the plots.

(a) (b)

Figure E.11.: Post-fit values of $E_{\mathrm{T}}^{\mathrm{miss}}$ related NPs for (a) the tHu and (b) the tHc background-only fit. Pruned NPs are excluded from the plots.

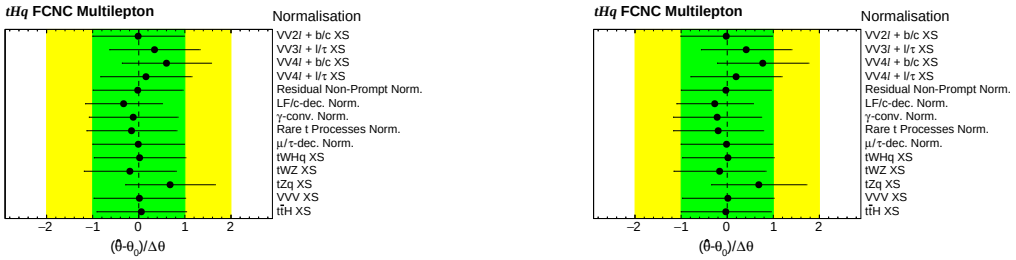

(a) (b)

Figure E.12.: Post-fit values of normalisation related NPs for (a) the tHu and (b) the tHc background-only fit. Pruned NPs are excluded from the plots.

(a) (b)

Figure E.13.: Post-fit values of additional minor NPs for (a) the tHu and (b) the tHc background-only fit. Pruned NPs are excluded from the plots.

(a) (b)

Figure E.14.: Post-fit values of PDF related NPs for (a) the tHu and (b) the tHc background-only fit. Pruned NPs are excluded from the plots.

Results of the full Fit to Data

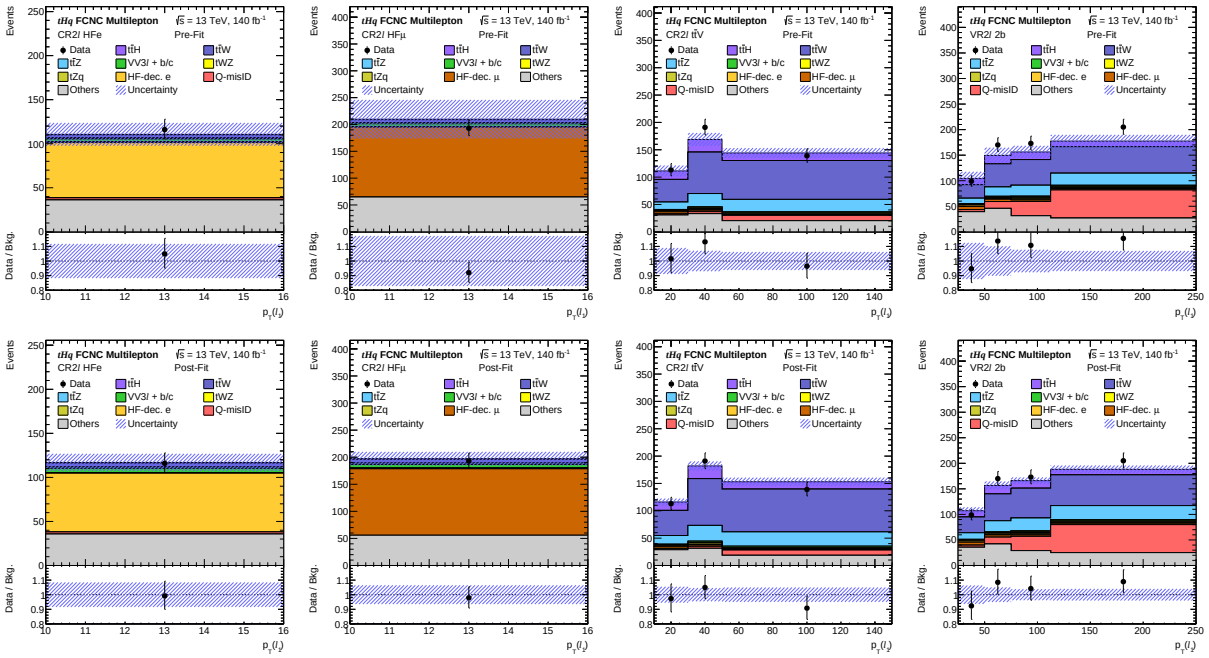

Figure E.15.: Pre-fit distributions (upper row) and post-fit distributions (lower row) of the three control regions and the validation region in the 2ℓSS final state for for the tHu full fit to data.

Table E.1.: The post-fit predicted and observed yields in all 2ℓSS regions of the analysis, obtained from the full fit to data. Minor MC templates are combined in the *Others* category. Post-fit signal normalisations are compatible with zero and thus not shown. The uncertainties on MC yields reflect the quadratic sum of all post-fit systematic and statistical uncertainties.

Process	SR2ℓ Dec	SR2ℓ Prod	CR2ℓ HFe	CR2ℓ HFμ	CR2ℓ $t\bar{t}V$	VR2ℓ $2b$
HF-decay e	122 ± 27	113 ± 25	66 ± 13	-	2.9 ± 0.9	6.8 ± 1.9
HF-decay μ	201 ± 36	192 ± 35	0.1 ± 0.02	120 ± 22	5.6 ± 1.2	15 ± 3
Q-misID	204 ± 16	457 ± 35	2.4 ± 0.2	-	15.5 ± 1.4	100 ± 10
$t\bar{t}H$	132 ± 20	27 ± 5	0.6 ± 0.1	1.0 ± 0.2	51 ± 8	54 ± 8
$t\bar{t}W$	512 ± 61	285 ± 42	4.8 ± 0.9	7.5 ± 1.4	216 ± 24	208 ± 25
$t\bar{t}Z$	210 ± 21	66 ± 9	1.5 ± 0.2	2.4 ± 0.4	70 ± 6	87 ± 7
$VV3\ell + b/c$	104 ± 20	192 ± 32	4.7 ± 1.0	6.6 ± 1.4	6.0 ± 1.2	7.1 ± 1.4
tWZ	23 ± 7	12 ± 4	0.11 ± 0.04	0.17 ± 0.06	3.6 ± 1.1	4.3 ± 1.3
tZq	26 ± 8	63 ± 18	0.7 ± 0.2	1.1 ± 0.3	5.8 ± 1.7	9 ± 2
Others	340 ± 64	322 ± 46	36 ± 8	59 ± 20	79 ± 14	132 ± 31
Total BG	1874 ± 107	1729 ± 92	117 ± 15	198 ± 30	455 ± 30	623 ± 43
Data	1847	1723	116	193	443	647

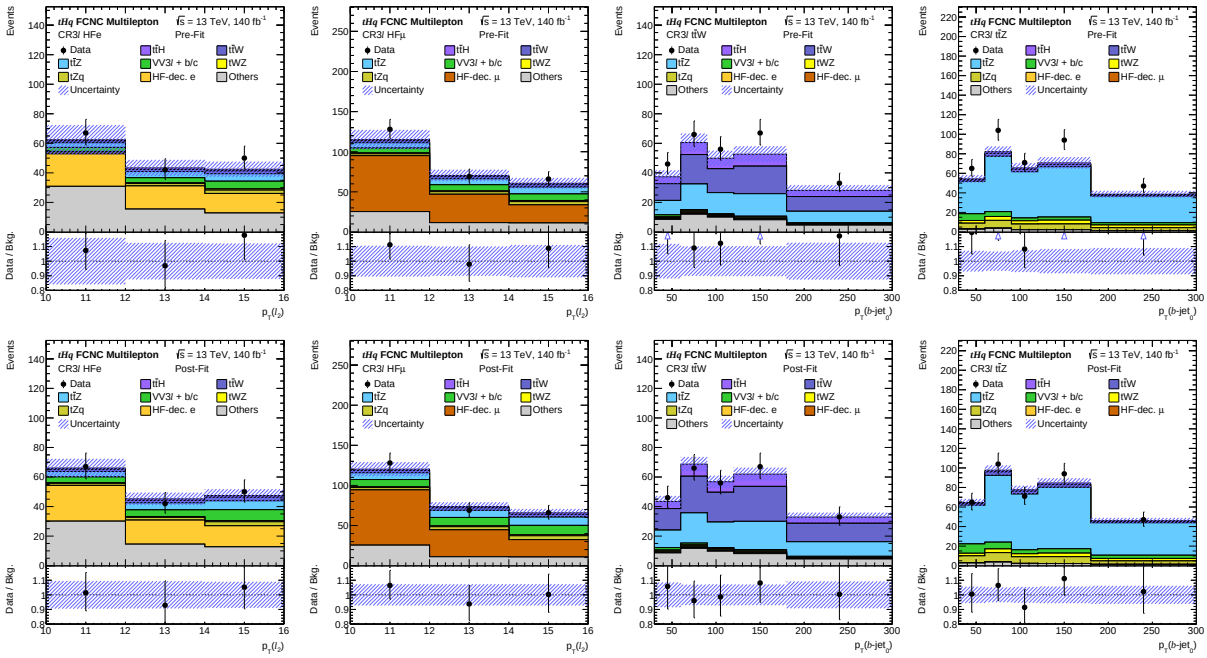

Figure E.16.: Pre-fit distributions (upper row) and post-fit distributions (lower row) of the control regions in the 3ℓ final state for the tHu full fit to data.

Table E.2.: The post-fit predicted and observed yields in all 3ℓ regions of the analysis, obtained from the full fit to data. Minor MC templates are combined in the *Others* category. Post-fit signal normalisations are compatible with zero and thus not shown. The uncertainties on MC yields reflect the quadratic sum of all post-fit systematic and statistical uncertainties.

Process	SR3ℓ Dec	SR3ℓ Prod	CR3ℓ HFe	CR3ℓ HFμ	CR3ℓ $t\bar{t}W$	CR3ℓ $t\bar{t}Z$
HF-decay e	14 ± 3	38 ± 9	53 ± 11	-	1.3 ± 0.3	0.28 ± 0.09
HF-decay μ	22 ± 4	63 ± 11	0.2 ± 0.1	122 ± 19	1.6 ± 0.3	0.37 ± 0.08
Q-misID	-	-	-	-	-	-
$t\bar{t}H$	47 ± 7	10 ± 2	3.0 ± 0.5	5.2 ± 0.9	32 ± 5	6.7 ± 1.1
$t\bar{t}W$	80 ± 10	77 ± 12	5.8 ± 1.0	9.5 ± 1.4	98 ± 16	12.5 ± 1.6
$t\bar{t}Z$	438 ± 40	75 ± 11	14.7 ± 1.8	28 ± 3	78 ± 7	261 ± 20
$VV3\ell + b/c$	215 ± 39	296 ± 49	15 ± 3	30 ± 5	4.8 ± 0.9	27 ± 5
tWZ	57 ± 18	19 ± 6	1.9 ± 0.6	3.6 ± 1.1	2.9 ± 0.9	16 ± 5
tZq	69 ± 20	134 ± 38	6.1 ± 1.8	12 ± 3	3.5 ± 1.0	35 ± 10
Others	119 ± 23	171 ± 32	59 ± 8	48 ± 9	43 ± 7	11.7 ± 1.5
Total BG	1061 ± 67	882 ± 73	159 ± 14	258 ± 22	265 ± 20	371 ± 24
Data	1046	896	159	263	268	381

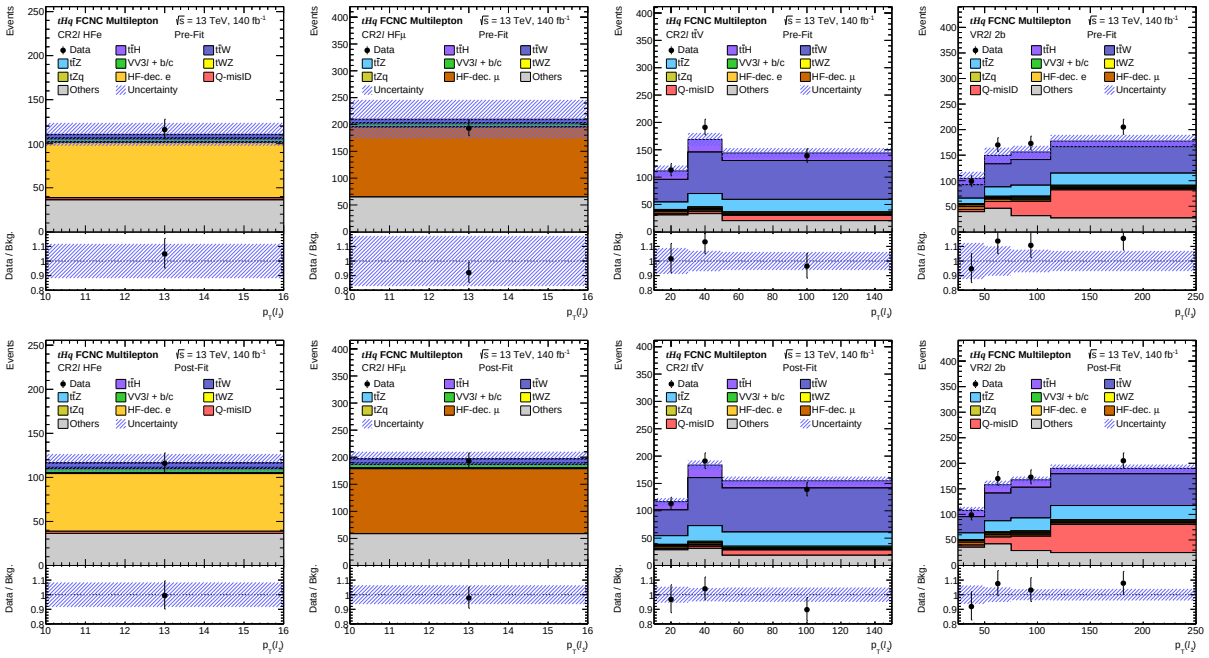

Figure E.17.: Pre-fit distributions (upper row) and post-fit distributions (lower row) of the three control regions and the validation region in the 2ℓSS final state for for the tHc full fit to data.

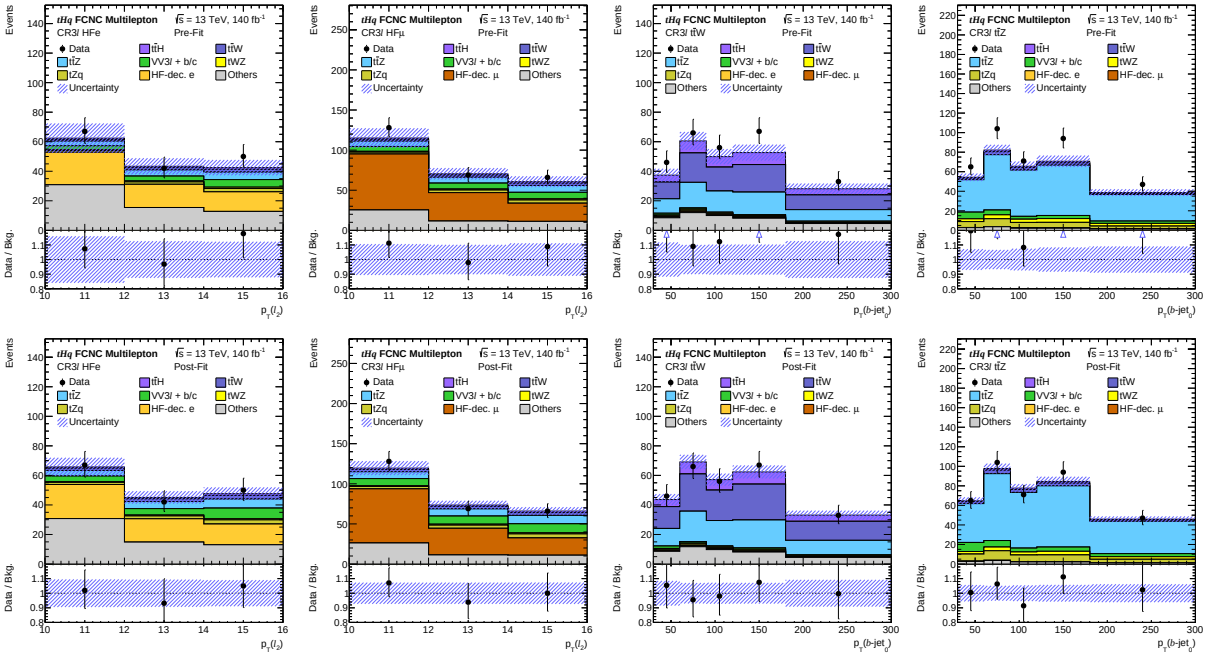

Figure E.18.: Pre-fit distributions (upper row) and post-fit distributions (lower row) of the control regions in the 3ℓ final state for the tHc full fit to data.

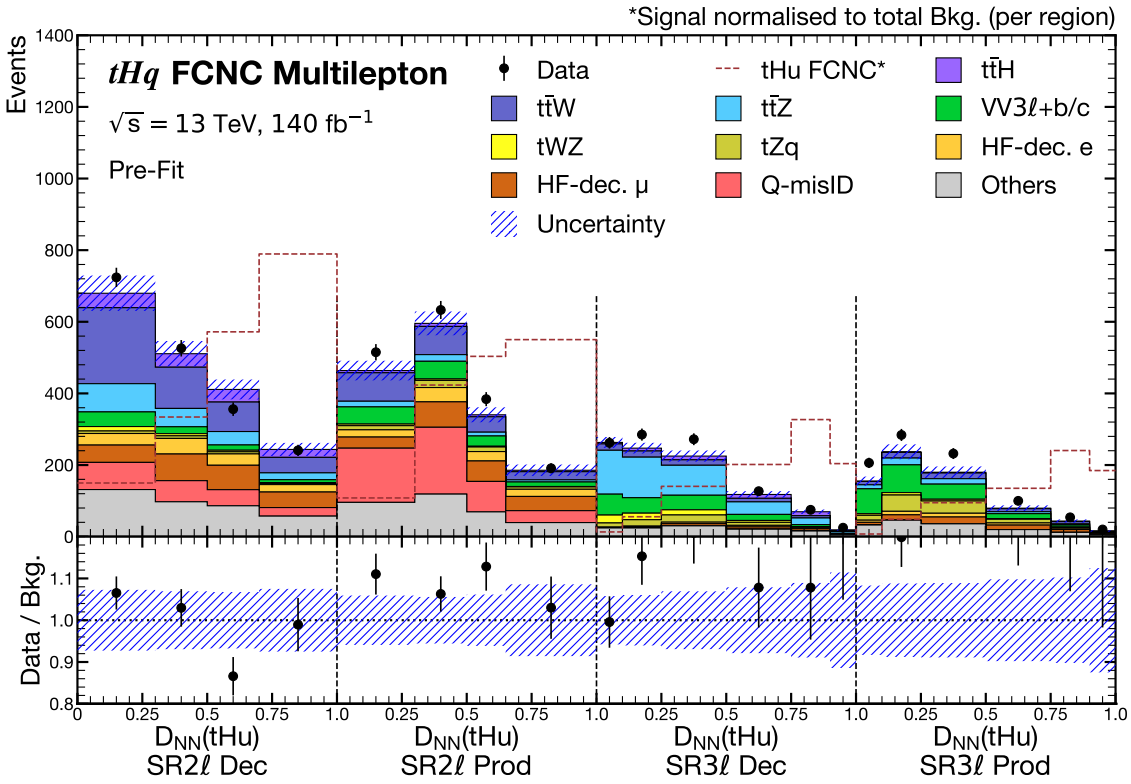

Figure E.19.: The pre-fit $D_{\mathrm{NN}}(tHu)$ distributions. The shape of the respective signal process is shown as a dashed line, scaled to the total number of background events per region.

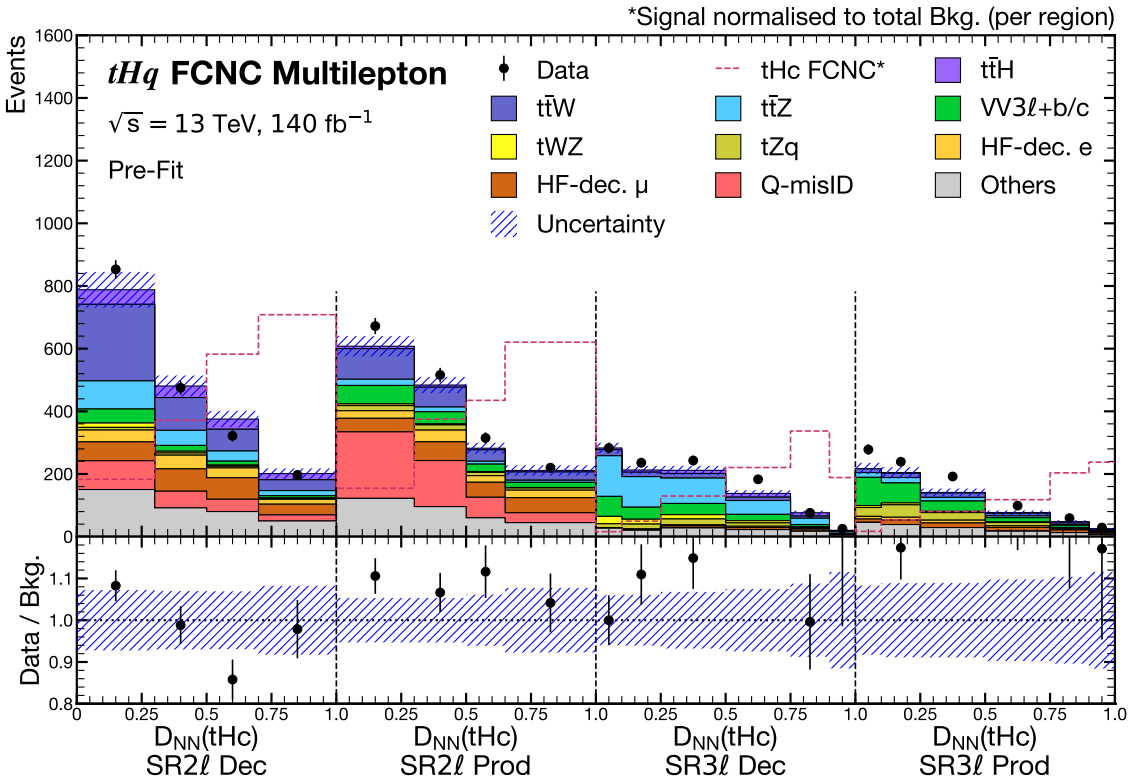

Figure E.20.: The pre-fit $D_{\mathrm{NN}}(tHc)$ distributions. The shape of the respective signal process is shown as a dashed line, scaled to the total number of background events per region.

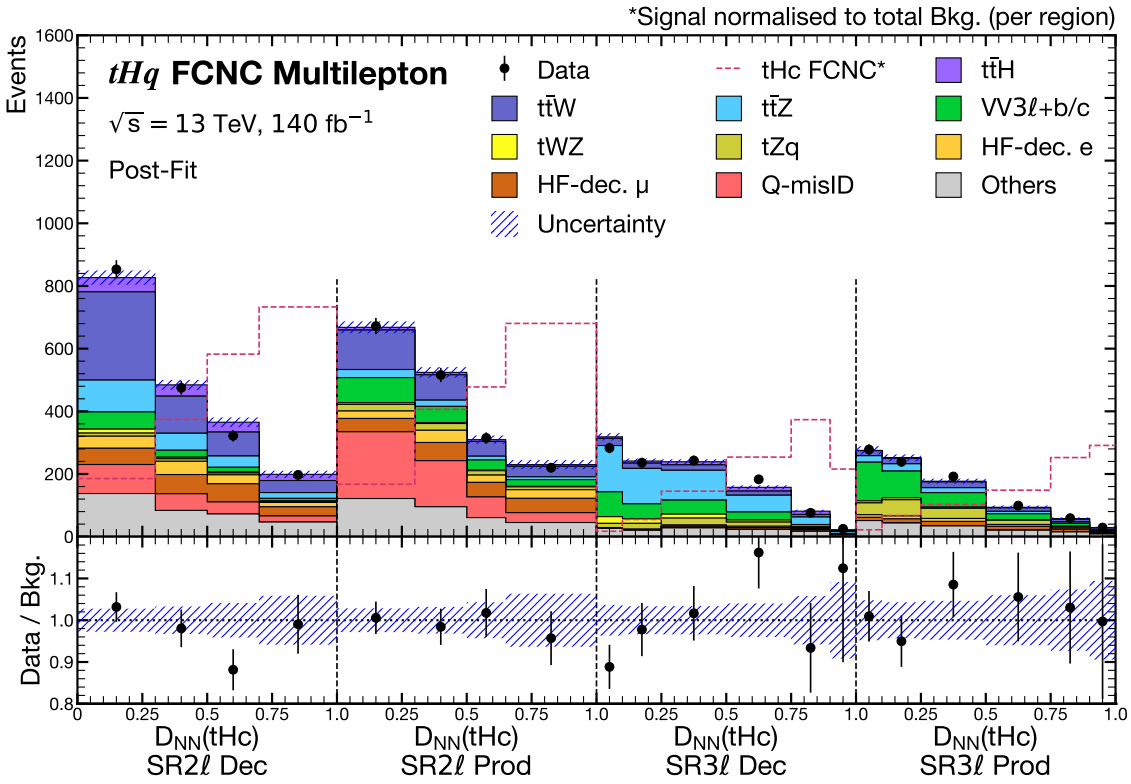

Figure E.21.: The fitted D_{NN} distributions of the tHc fit. The shape of the respective signal process is shown as a dashed line, scaled to the total number of background events per region.

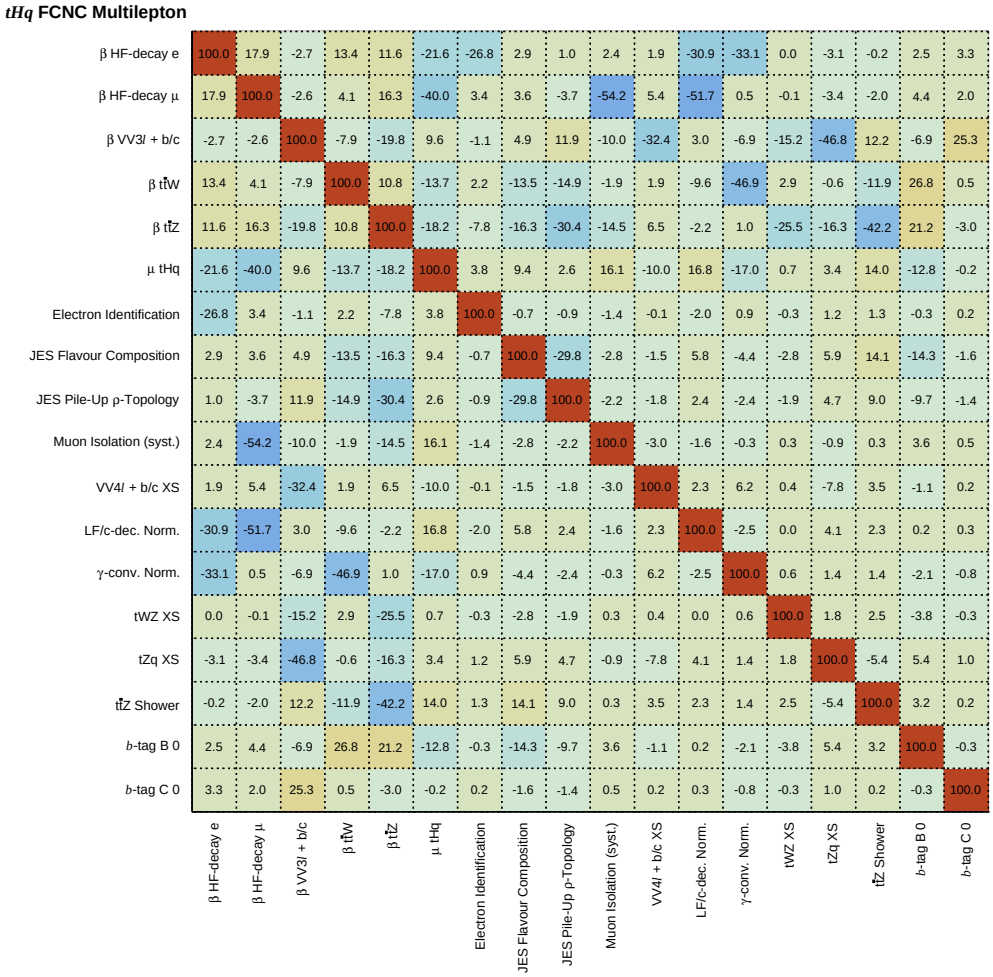

Figure E.22.: Correlations of NPs for the full tHc fit. Only NPs with a correlation above 0.25 to at least one other NP are shown in the correlation matrix.

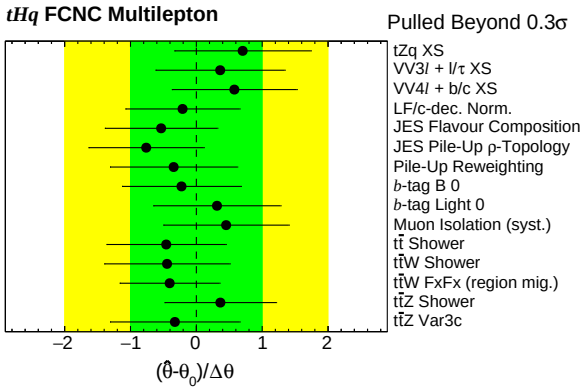

Figure E.23.: Pulls and constraints of NPs from the full tHc fit, surpassing a pull threshold of 0.3σ. Among the depicted NPs are also the ones exhibiting the strongest constraints in the fit.

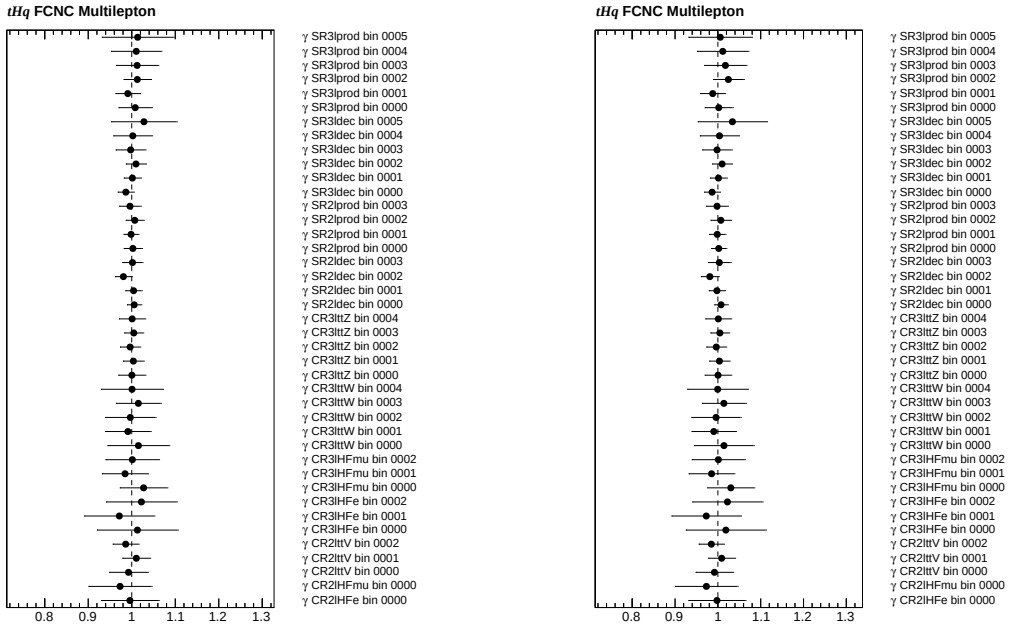

(a) (b)

Figure E.24.: γ-factors for (a) the tHu and (b) the tHc full fit to data.

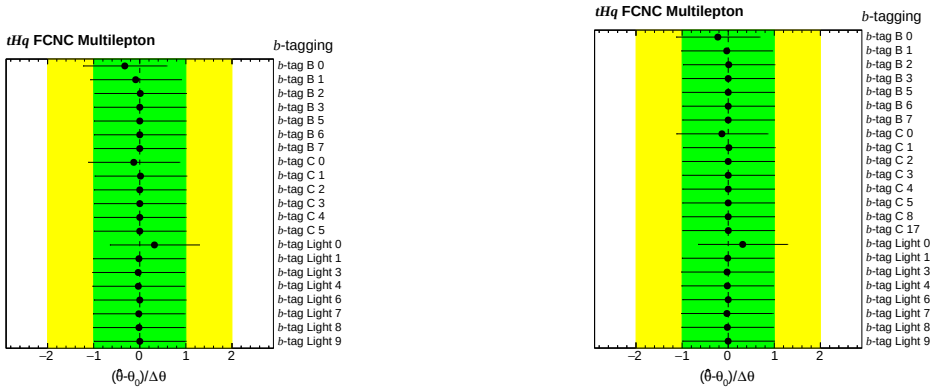

(a) (b)

Figure E.25.: Post-fit values of b-tagging related NPs for (a) the tHu and (b) the tHc full fit to data. Pruned NPs are excluded from the plots.

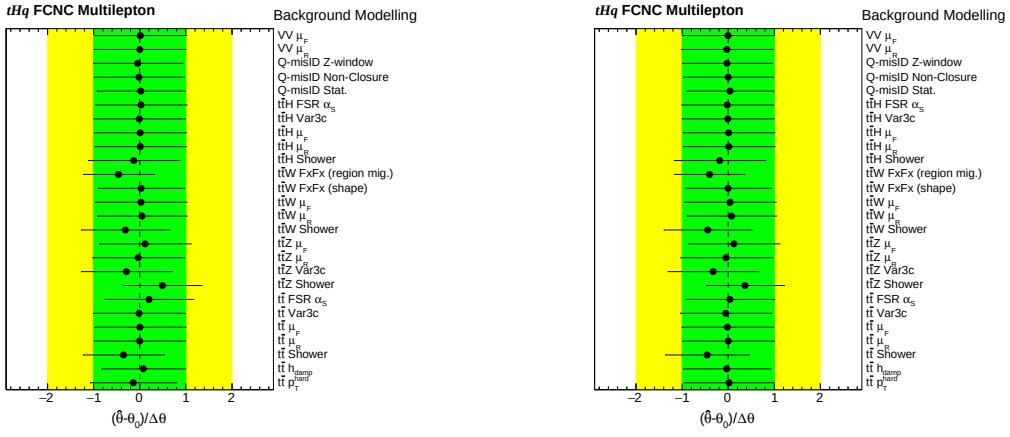

(a) (b)

Figure E.26.: Post-fit values of background-modelling related NPs for (a) the tHu and (b) the tHc full fit to data. Pruned NPs are excluded from the plots.

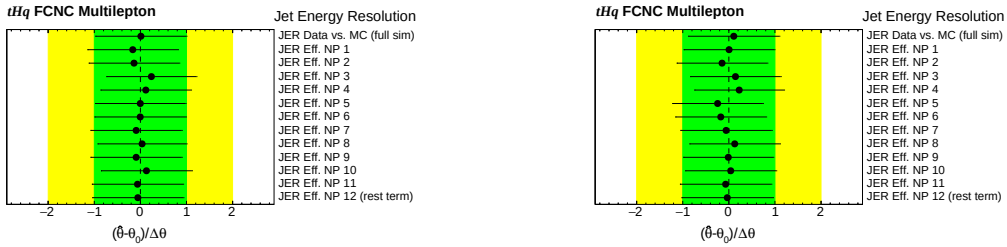

(a) (b)

Figure E.27.: Post-fit values of JER related NPs for (a) the tHu and (b) the tHc full fit to data. Pruned NPs are excluded from the plots.

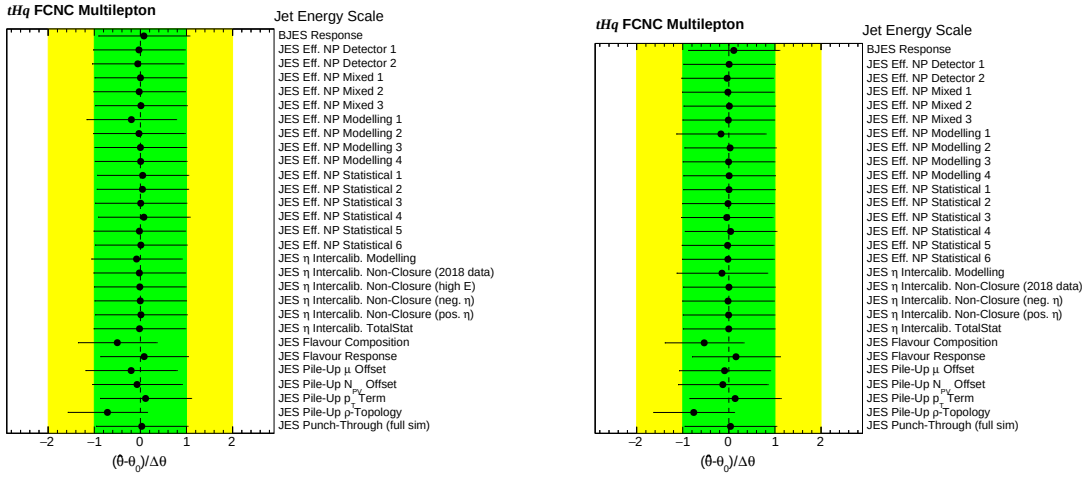

(a) (b)

Figure E.28.: Post-fit values of JES related NPs for (a) the tHu and (b) the tHc full fit to data. Pruned NPs are excluded from the plots.

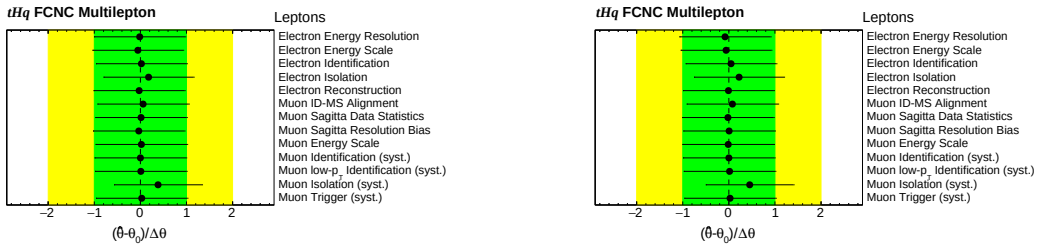

(a) (b)

Figure E.29.: Post-fit values of lepton related NPs for (a) the tHu and (b) the tHc full fit to data. Pruned NPs are excluded from the plots.

(a) (b)

Figure E.30.: Post-fit values of $E_{\mathrm{T}}^{\mathrm{miss}}$ related NPs for (a) the tHu and (b) the tHc full fit to data. Pruned NPs are excluded from the plots.

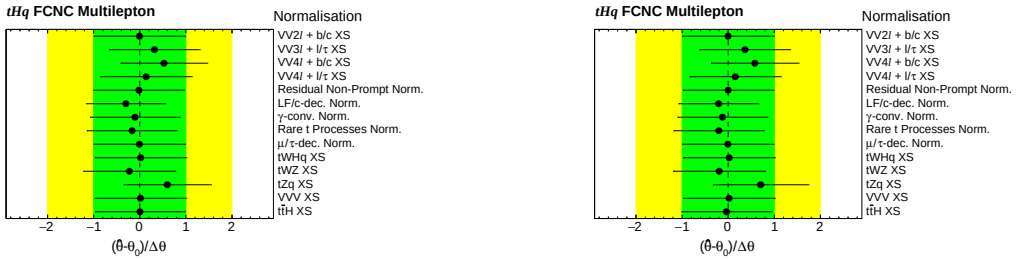

(a) (b)

Figure E.31.: Post-fit values of normalisation related NPs for (a) the tHu and (b) the tHc full fit to data. Pruned NPs are excluded from the plots.

(a) (b)

Figure E.32.: Post-fit values of additional minor NPs for (a) the tHu and (b) the tHc full fit to data. Pruned NPs are excluded from the plots.

(a) (b)

Figure E.33.: Post-fit values of PDF related NPs for (a) the tHu and (b) the tHc full fit to data. Pruned NPs are excluded from the plots.

Bibliography

[1] J. Chadwick, *The Existence of a Neutron*,
Proc. Roy. Soc. Lond. A **136** (1932) 692 (cit. on p. 3).

[2] W. Heisenberg, *Über den Bau der Atomkerne. I*,
Zeitschrift fur Physik **77** (1932) 1 (cit. on p. 3).

[3] W. Heisenberg, *Über den Bau der Atomkerne. II*,
Zeitschrift fur Physik **78** (1932) 156 (cit. on p. 3).

[4] W. Heisenberg, *Über den Bau der Atomkerne. III*,
Zeitschrift fur Physik **80** (1933) 587 (cit. on p. 3).

[5] E. Fermi, *Tentativo di una Teoria Dei Raggi β*,
Il Nuovo Cimento (1924-1942) **11** (1934) 1 (cit. on p. 3).

[6] M. Gell-Mann, *The Eightfold Way: A Theory of strong interaction symmetry*,
(1961) (cit. on p. 3).

[7] F. Englert and R. Brout,
Broken Symmetry and the Mass of Gauge Vector Mesons,
Phys. Rev. Lett. **13** (9 1964) 321 (cit. on p. 3).

[8] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*,
Phys. Rev. Lett. **13** (16 1964) 508 (cit. on pp. 3, 11).

[9] G. S. Guralnik, C. R. Hagen and T. W. B. Kibble,
Global Conservation Laws and Massless Particles,
Phys. Rev. Lett. **13** (20 1964) 585 (cit. on p. 3).

[10] S. L. Glashow, *The renormalizability of vector meson interactions*,
Nuclear Physics **10** (1959) 107 (cit. on p. 3).

[11] A. Salam, *Weak and Electromagnetic Interactions*,
Conf. Proc. C **680519** (1968) 367 (cit. on p. 3).

[12] S. Weinberg, *A Model of Leptons*, Phys. Rev. Lett. **19** (1967) 1264 (cit. on p. 3).

[13] Particle Data Group, P. Zyla et al., *Review of Particle Physics*,
PTEP **2022** (2022) 083C01 (cit. on p. 4).

[14] Particle Data Group, P. Zyla et al., *Review of Particle Physics*,
Phys. Rev. D **98** (3 2018) 030001 (cit. on pp. 4, 5, 37, 85).

[15] W. Commons, *Standard Model of Elementary Particles*, 2023, URL: https://en.wikipedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg (cit. on p. 4).

[16] M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory*, Reading, USA: Addison-Wesley (1995) 842 p, Westview Press, 1995 (cit. on pp. 5, 7, 8, 13, 56).

[17] J. D. Bjorken and S. D. Drell, *Relativistische Quantenmechanik*, BI - Hochschultaschenbücher, Berlin, Germany: Bibliographisches Institut, 1990 (cit. on p. 5).

[18] R. N. Cahn, *Semi-simple lie algebras and their representations*, Dover Books on Mathematics, Mineola, NY: Dover Publications, 2006 (cit. on p. 6).

[19] T.-P. Cheng and L.-F. Li, *Gauge theory of elementary particle physics*, en, Oxford science publications, London, England: Oxford University Press, 1984 (cit. on p. 6).

[20] R. P. Feynman, *Space-Time Approach to Quantum Electrodynamics*, Phys. Rev. **76** (6 1949) 769 (cit. on p. 7).

[21] J. Roggel, *Search for vector-like partners of the top and bottom quarks with the ATLAS experiment*, Presented 27 Jan 2023, Wuppertal U., 2022, URL: https://cds.cern.ch/record/2850965 (cit. on p. 12).

[22] V. A. Austrup, *Search for scalar and vector leptoquarks decaying into quarks and leptons of different generations*, Presented 01 Dec 2022, Bergische Universitaet Wuppertal, 2022, URL: https://cds.cern.ch/record/2847970 (cit. on p. 12).

[23] G. Münster, *Von der Quantenfeldtheorie zum Standardmodell*, de, de Gruyter Studium, de Gruyter, 2019 (cit. on pp. 12, 17, 36, 58).

[24] K. Nguyen, *The Higgs Mechanism*, (2009), URL: https://www.theorie.physik.uni-muenchen.de/lsfrey/teaching/archiv/sose_09/rng/higgs_mechanism.pdf (cit. on p. 13).

[25] K. Nishijima, *Charge Independence Theory of V Particles**, Progress of Theoretical Physics **13** (1955) 285 (cit. on p. 13).

[26] M. Gell-Mann, *The interpretation of the new particles as displaced charge multiplets*, Nuovo Cim. **4** (1956) 848 (cit. on p. 13).

[27] N. Cabibbo, *Unitary Symmetry and Leptonic Decays*, Phys. Rev. Lett. **10** (12 1963) 531 (cit. on p. 14).

[28] M. Kobayashi and T. Maskawa,
CP-Violation in the Renormalizable Theory of Weak Interaction,
Progress of Theoretical Physics **49** (1973) 652 (cit. on pp. 14, 15).

[29] S. L. Glashow, J. Iliopoulos and L. Maiani,
Weak Interactions with Lepton-Hadron Symmetry, Phys. Rev. D **2** (7 1970) 1285 (cit. on p. 19).

[30] J. Iliopoulos, *Glashow-Iliopoulos-Maiani mechanism*,
Scholarpedia **5** (2010) 7125, revision #137391 (cit. on p. 19).

[31] ATLAS Collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*,
Phys. Lett. B **716** (2012) 1, arXiv: 1207.7214 [hep-ex] (cit. on p. 20).

[32] CMS Collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, Phys. Lett. B **716** (2012) 30,
arXiv: 1207.7235 [hep-ex] (cit. on p. 20).

[33] A. Blondel, *The Number of Neutrinos and the Z Line Shape*, (2016),
URL: https://cds.cern.ch/record/2217139 (cit. on p. 21).

[34] ATLAS Collaboration, *Measurement of the total and differential Higgs boson production cross-sections at $\sqrt{s} = 13$ TeV with the ATLAS detector by combining the $H \to ZZ^* \to 4\ell$ and $H \to \gamma\gamma$ decay channels*, JHEP **05** (2022) 028,
arXiv: 2207.08615 [hep-ex] (cit. on p. 21).

[35] ATLAS Collaboration, *Standard Model Summary Plots February 2022*,
tech. rep., CERN, 2022, URL: http://cds.cern.ch/record/2804061
(cit. on p. 22).

[36] M. Baak et al.,
Updated status of the global electroweak fit and constraints on new physics,
The European Physical Journal C **72** (2012) (cit. on p. 23).

[37] J. Haller et al.,
Update of the global electroweak fit and constraints on two-Higgs-doublet models,
The European Physical Journal C **78** (2018) (cit. on pp. 23, 24).

[38] F. P. An et al.,
Observation of Electron-Antineutrino Disappearance at Daya Bay,
Physical Review Letters **108** (2012) (cit. on p. 23).

[39] B. Pontecorvo, *Inverse beta processes and nonconservation of lepton charge*,
Zh. Eksp. Teor. Fiz. **34** (1957) 247 (cit. on p. 23).

[40] Z. Maki, M. Nakagawa and S. Sakata,
Remarks on the Unified Model of Elementary Particles,
Progress of Theoretical Physics **28** (1962) 870 (cit. on p. 23).

[41] D.-d. Wu, *A Brief Introduction to the Strong CP Problem*, Z. Naturforsch. A **52** (1997) 179, ed. by B. G. Englert and G. Sussmann (cit. on p. 25).

[42] R. D. Peccei and H. R. Quinn, CP *Conservation in the Presence of Pseudoparticles*, Phys. Rev. Lett. **38** (25 1977) 1440 (cit. on p. 25).

[43] A. Einstein, *Die Feldgleichungen der Gravitation*, Sitzungsberichte der Königlich Preussischen Akademie der Wissenschaften (1915) 844 (cit. on p. 25).

[44] S. Weinberg, *Cosmology*, London, England: Oxford University Press, 2008 (cit. on p. 26).

[45] E. Komatsu et al., *Seven-Year Wilkinson Microwave Anisotropy Probe (WMAP) Observations: Cosmological Interpretation*, The Astrophysical Journal Supplement Series **192** (2011) 18 (cit. on p. 26).

[46] R. J. Adler, B. Casey and O. C. Jacob, *Vacuum catastrophe: An elementary exposition of the cosmological constant problem*, American Journal of Physics **63** (1995) 620 (cit. on p. 26).

[47] G. Branco et al., *Theory and phenomenology of two-Higgs-doublet models*, Physics Reports **516** (2012) 1 (cit. on pp. 27, 28).

[48] S. Bejar, *Flavor changing neutral decay effects in models with two Higgs boson doublets: Applications to LHC Physics*, 2006, arXiv: `hep-ph/0606138 [hep-ph]` (cit. on p. 28).

[49] M. E. Peskin, *Supersymmetry in Elementary Particle Physics*, 2008, arXiv: `0801.1928 [hep-ph]` (cit. on pp. 28, 30).

[50] H. E. Haber, *The status of the minimal supersymmetric standard model and beyond*, Nuclear Physics B - Proceedings Supplements **62** (1998) 469 (cit. on p. 29).

[51] J. Cao, C. Han, L. Wu, J. M. Yang and M. Zhang, *SUSY induced top quark FCNC decay $t \to ch$ after Run I of LHC*, The European Physical Journal C **74** (2014) (cit. on p. 30).

[52] G. Eilam, A. Gemintern, T. Han, J. Yang and X. Zhang, *Top-quark rare decay $t \to ch$ in R-parity-violating SUSY*, Physics Letters B **510** (2001) 227 (cit. on p. 30).

[53] A. Azatov, M. Toharia and L. Zhu, *Higgs mediated flavor changing neutral currents in warped extra dimensions*, Physical Review D **80** (2009) (cit. on p. 30).

[54] E. Accomando and K. Benakli, *Kaluza-Klein states at the LHC in models with localized fermions*, 2005, arXiv: `hep-ph/0512243 [hep-ph]` (cit. on p. 31).

[55] A. Azatov, M. Toharia and L. Zhu,
Higgs mediated flavor changing neutral currents in warped extra dimensions,
Phys. Rev. D **80** (2009) 035016, arXiv: 0906.1990 [hep-ph] (cit. on p. 31).

[56] E. Fermi, *Tentativo di una Teoria Dei Raggi β*, Il Nuovo Cimento **11** (1934) 1 (cit. on p. 31).

[57] G. Durieux, F. Maltoni and C. Zhang,
Global approach to top-quark flavor-changing interactions,
Phys. Rev. D **91** (7 2015) 074017 (cit. on p. 32).

[58] ATLAS Collaboration, *Measurements of top-quark pair spin correlations in the $e\mu$ channel at $\sqrt{s} = 13$ TeV using pp collisions in the ATLAS detector*,
Eur. Phys. J. C **80** (2020) 754, arXiv: 1903.07570 [hep-ex] (cit. on p. 34).

[59] ATLAS Collaboration, *Search for flavour-changing neutral currents in processes with one top quark and a photon using $81\,fb^{-1}$ of pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS experiment*, Phys. Lett. B **800** (2020) 135082,
arXiv: 1908.08461 [hep-ex] (cit. on p. 34).

[60] ATLAS Collaboration, *Search for flavor-changing neutral currents in top quark decays $t \to Hc$ and $t \to Hu$ in multilepton final states in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*,
Phys. Rev. D **98** (2018) 032002, arXiv: 1805.03483 [hep-ex]
(cit. on pp. 39, 141, 142).

[61] ATLAS Collaboration, *Search for top-quark decays $t \to Hq$ with $36\,fb^{-1}$ of pp collision data at $\sqrt{s} = 13$ TeV with the ATLAS detector*, JHEP **05** (2019) 123,
arXiv: 1812.11568 [hep-ex] (cit. on p. 39).

[62] ATLAS Collaboration, *Search for top quark decays $t \to qH$, with $H \to \gamma\gamma$, in $\sqrt{s} = 13$ TeV pp collisions using the ATLAS detector*, JHEP **10** (2017) 129,
arXiv: 1707.01404 [hep-ex] (cit. on p. 39).

[63] CMS Collaboration,
Search for flavor-changing neutral current interactions of the top quark and the Higgs boson decaying to a bottom quark-antiquark pair at $\sqrt{s}$ =13 TeV,
JHEP **02** (2021) 169, arXiv: 2112.09734 [hep-ex] (cit. on p. 39).

[64] CMS Collaboration, *Search for flavor-changing neutral current interactions of the top quark and Higgs boson in proton-proton collisions at $\sqrt{s} = 13$ TeV*,
tech. rep., CERN, 2023, URL: https://cds.cern.ch/record/2872677
(cit. on pp. 39, 143, 144).

[65] CMS Collaboration, *Search for flavor-changing neutral current interactions of the top quark and Higgs boson in final states with two photons in proton–proton collisions at $\sqrt{s} = 13$ TeV*, (2021), arXiv: 2111.02219 [hep-ex]
(cit. on pp. 39, 143).

[66] ATLAS Collaboration, *Search for a new scalar resonance in flavour-changing neutral-current top-quark decays $t \to qX$ $(q = u, c)$, with $X \to b\bar{b}$, in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, JHEP **07** (2023) 199, arXiv: 2301.03902 [hep-ex] (cit. on pp. 39, 144).

[67] ATLAS Collaboration, *Search for flavour-changing neutral current interactions of the top quark and the Higgs boson in events with a pair of τ-leptons in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, JHEP **06** (2022) 155, arXiv: 2208.11415 [hep-ex] (cit. on pp. 39, 144).

[68] ATLAS Collaboration, *Search for flavor-changing neutral tqH interactions with $H \to \gamma\gamma$ in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector*, (2023), arXiv: 2309.12817 [hep-ex] (cit. on pp. 39, 144).

[69] ATLAS and CMS Collaborations, *Summary of the current 95 % confidence level observed limits on the branching ratios of the top quark decays via flavour changing neutral currents (FCNC) by the ATLAS and CMS Collaborations compared to several new physics models.*, (2023), URL: https://twiki.cern.ch/twiki/pub/LHCPhysics/FCNCHistory/fcnc_summarybsm_nov23.pdf (cit. on p. 41).

[70] L. Evans and P. Bryant, *LHC Machine*, JINST **3** (2008) S08001 (cit. on p. 43).

[71] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider: A Description of the Detector Configuration for Run 3*, (2023), arXiv: 2305.16623 [physics.ins-det] (cit. on p. 43).

[72] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, JINST **3** (2008) S08003 (cit. on p. 43).

[73] CMS Collaboration, *The CMS Experiment at the CERN LHC*, JINST **3** (2008) S08004 (cit. on p. 43).

[74] ALICE Collaboration, *The ALICE experiment at the CERN LHC. A Large Ion Collider Experiment*, JINST **3** (2008) S08002, Also published by CERN Geneva in 2010 (cit. on p. 43).

[75] LHCb Collaboration, *The LHCb Detector at the LHC*, JINST **3** (2008) S08005, Also published by CERN Geneva in 2010 (cit. on p. 43).

[76] E. Mobs, *The CERN accelerator complex - August 2018. Complexe des accélérateurs du CERN - Août 2018*, (2018), General Photo, URL: https://cds.cern.ch/record/2636343 (cit. on p. 44).

[77] S.-Y. Lee, *Accelerator physics (fourth edition)*, en, 2018 (cit. on p. 44).

[78] ATLAS Collaboration, *ATLAS Luminosity - Public Results*, 2018, URL: https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2 (visited on 08/01/2023) (cit. on p. 45).

[79] J. Pequenao, *Computer generated image of the whole ATLAS detector*, (2008), URL: https://cds.cern.ch/record/1095924 (cit. on p. 46).

[80] ATLAS Collaboration, *ATLAS Inner Detector: Technical Design Report, Volume 1*, ATLAS-TDR-4; CERN-LHCC-97-016, 1997, URL: https://cds.cern.ch/record/331063 (cit. on p. 47).

[81] ATLAS Collaboration, *ATLAS Inner Detector: Technical Design Report, Volume 2*, ATLAS-TDR-5, CERN-LHCC-97-017, 1997, URL: https://cds.cern.ch/record/331064 (cit. on p. 47).

[82] J. Pequenao, *Computer generated image of the ATLAS inner detector*, (2008), URL: https://cds.cern.ch/record/1095926 (cit. on p. 47).

[83] B. Abbott et al., *Production and integration of the ATLAS Insertable B-Layer*, JINST **13** (2018) T05008, arXiv: 1803.00844 [physics.ins-det] (cit. on p. 48).

[84] ATLAS Collaboration, *ATLAS Calorimeter Performance: Technical Design Report*, ATLAS-TDR-1; CERN-LHCC-96-040, 1996, URL: https://cds.cern.ch/record/331059 (cit. on p. 49).

[85] ATLAS Collaboration, *ATLAS Liquid Argon Calorimeter: Technical Design Report*, ATLAS-TDR-2; CERN-LHCC-96-041, 1996, URL: https://cds.cern.ch/record/331061 (cit. on p. 49).

[86] ATLAS Collaboration, *ATLAS Tile Calorimeter: Technical Design Report*, ATLAS-TDR-3; CERN-LHCC-96-042, 1996, URL: https://cds.cern.ch/record/331062 (cit. on p. 49).

[87] J. Pequenao, *Computer Generated image of the ATLAS calorimeter*, (2008), URL: https://cds.cern.ch/record/1095927 (cit. on p. 49).

[88] ATLAS Collaboration, *ATLAS Muon Spectrometer: Technical Design Report*, ATLAS-TDR-10; CERN-LHCC-97-022, CERN, 1997, URL: https://cds.cern.ch/record/331068 (cit. on p. 51).

[89] J. Pequenao, *Computer generated image of the ATLAS Muons subsystem*, (2008), URL: https://cds.cern.ch/record/1095929 (cit. on p. 52).

[90] ATLAS Collaboration, *Performance of the ATLAS trigger system in 2015*, Eur. Phys. J. C **77** (2017) 317, arXiv: 1611.09661 [hep-ex] (cit. on p. 52).

[91] ATLAS Collaboration, *Operation of the ATLAS trigger system in Run 2*, JINST **15** (2020) P10004, arXiv: 2007.12539 [hep-ex] (cit. on p. 52).

[92] S. van der Meer, *Calibration of the effective beam height in the ISR*, tech. rep., CERN, 1968 (cit. on p. 53).

[93] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector at the LHC*, Eur. Phys. J. C **76** (2016) 653, arXiv: 1608.03953 [hep-ex] (cit. on p. 54).

[94] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, Eur. Phys. J. C **83** (2023) 982, arXiv: 2212.09379 [hep-ex] (cit. on p. 54).

[95] C. Bierlich et al., *A comprehensive guide to the physics and usage of PYTHIA 8.3*, 2022, arXiv: 2203.11601 [hep-ph] (cit. on pp. 56, 59, 60).

[96] The NNPDF Collaboration, *Parton distributions from high-precision collider data: NNPDF Collaboration*, The European Physical Journal C **77** (2017) (cit. on p. 57).

[97] J. C. Collins, D. E. Soper and G. Sterman, *Factorization of Hard Processes in QCD*, 2004, arXiv: hep-ph/0409313 [hep-ph] (cit. on p. 57).

[98] T. Sjöstrand, *Jet fragmentation of multiparton configurations in a string framework*, Nucl. Phys. B **248** (1984) 469 (cit. on p. 61).

[99] A. Kupco, *Cluster Hadronization in HERWIG 5.9*, 1999, arXiv: hep-ph/9906412 [hep-ph] (cit. on p. 61).

[100] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, Comput. Phys. Commun. **191** (2015) 159, arXiv: 1410.3012 [hep-ph] (cit. on p. 61).

[101] E. Bothmann et al., *Event generation with Sherpa 2.2*, SciPost Phys. **7** (2019) 034, arXiv: 1905.09127 [hep-ph] (cit. on p. 61).

[102] T. Gleisberg et al., *Event generation with SHERPA 1.1*, JHEP **02** (2009) 007, arXiv: 0811.4622 [hep-ph] (cit. on p. 61).

[103] M. Bähr et al., *Herwig++ physics and manual*, Eur. Phys. J. C **58** (2008) 639, arXiv: 0803.0883 [hep-ph] (cit. on p. 61).

[104] J. Bellm et al., *Herwig 7.0/Herwig++ 3.0 release note*, Eur. Phys. J. C **76** (2016) 196, arXiv: 1512.01178 [hep-ph] (cit. on p. 61).

[105] S. Schumann and F. Krauss,
A parton shower algorithm based on Catani–Seymour dipole factorisation,
JHEP **03** (2008) 038, arXiv: 0709.1027 [hep-ph] (cit. on p. 62).

[106] S. Frixione and B. R. Webber,
Matching NLO QCD computations and parton shower simulations,
JHEP **06** (2002) 029, arXiv: hep-ph/0204244 (cit. on p. 64).

[107] J. Alwall et al.,
The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations,
JHEP **07** (2014) 079, arXiv: 1405.0301 [hep-ph] (cit. on p. 64).

[108] P. Nason,
A new method for combining NLO QCD with shower Monte Carlo algorithms,
JHEP **11** (2004) 040, arXiv: hep-ph/0409146 (cit. on p. 64).

[109] S. Frixione, P. Nason and C. Oleari, *Matching NLO QCD computations with parton shower simulations: the POWHEG method*, JHEP **11** (2007) 070,
arXiv: 0709.2092 [hep-ph] (cit. on p. 64).

[110] S. Frixione, G. Ridolfi and P. Nason, *A positive-weight next-to-leading-order Monte Carlo for heavy flavour hadroproduction*, JHEP **09** (2007) 126,
arXiv: 0707.3088 [hep-ph] (cit. on p. 64).

[111] S. Alioli, P. Nason, C. Oleari and E. Re, *NLO single-top production matched with shower in POWHEG: s- and t-channel contributions*, JHEP **09** (2009) 111,
arXiv: 0907.4076 [hep-ph] (cit. on p. 64), Erratum: JHEP **02** (2010) 011.

[112] S. Alioli, P. Nason, C. Oleari and E. Re, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*,
JHEP **06** (2010) 043, arXiv: 1002.2581 [hep-ph] (cit. on p. 64).

[113] E. Re, *Single-top Wt-channel production matched with parton showers using the POWHEG method*, Eur. Phys. J. C **71** (2011) 1547, arXiv: 1009.2450 [hep-ph]
(cit. on p. 64).

[114] R. Frederix, E. Re and P. Torrielli, *Single-top t-channel hadroproduction in the four-flavour scheme with POWHEG and aMC@NLO*, JHEP **09** (2012) 130,
arXiv: 1207.5391 [hep-ph] (cit. on p. 64).

[115] ATLAS Collaboration, *Studies on top-quark Monte Carlo modelling for Top2016*,
ATL-PHYS-PUB-2016-020, 2016,
URL: https://cds.cern.ch/record/2216168 (cit. on p. 64).

[116] M. L. Mangano, M. Moretti, F. Piccinini and M. Treccani, *Matching matrix elements and shower evolution for top-pair production in hadronic collisions*,
Journal of High Energy Physics **2007** (2007) 013 (cit. on p. 65).

[117] L. Lönnblad,
Correcting the Colour-Dipole Cascade Model with Fixed Order Matrix Elements,
JHEP **05** (2002) 046, arXiv: hep-ph/0112284 (cit. on p. 65).

[118] L. Lönnblad and S. Prestel,
Matching tree-level matrix elements with interleaved showers,
JHEP **03** (2012) 019, arXiv: 1109.4829 [hep-ph] (cit. on p. 65).

[119] S. Höche, F. Krauss, S. Schumann and F. Siegert,
QCD matrix elements and truncated showers,
Journal of High Energy Physics **2009** (2009) 053 (cit. on p. 65).

[120] S. Höche, F. Krauss, M. Schönherr and F. Siegert,
QCD matrix elements + parton showers. The NLO case, JHEP **04** (2013) 027,
arXiv: 1207.5030 [hep-ph] (cit. on p. 65).

[121] R. Frederix and I. Tsinikos, *On improving NLO merging for* $t\bar{t}W$ *production*,
JHEP **11** (2021) 029, arXiv: 2108.07826 [hep-ph] (cit. on p. 65).

[122] S. Agostinelli et al., *Geant4 – a simulation toolkit*,
Nucl. Instrum. Meth. A **506** (2003) 250 (cit. on p. 65).

[123] ATLAS Collaboration,
The Pythia 8 A3 tune description of ATLAS minimum bias and inelastic measurements incorporating the Donnachie–Landshoff diffractive model,
ATL-PHYS-PUB-2016-017, 2016,
URL: https://cds.cern.ch/record/2206965 (cit. on p. 66).

[124] NNPDF Collaboration, R. D. Ball et al., *Parton distributions with LHC data*,
Nucl. Phys. B **867** (2013) 244, arXiv: 1207.1303 [hep-ph] (cit. on pp. 66, 67).

[125] ATLAS Collaboration, *ATLAS Pythia 8 tunes to* 7 *TeV data*,
ATL-PHYS-PUB-2014-021, 2014,
URL: https://cds.cern.ch/record/1966419 (cit. on pp. 67, 115).

[126] D. J. Lange, *The EvtGen particle decay simulation package*,
Nucl. Instrum. Meth. A **462** (2001) 152 (cit. on p. 67).

[127] The NNPDF Collaboration, R. D. Ball et al.,
Parton distributions for the LHC run II, JHEP **04** (2015) 040,
arXiv: 1410.8849 [hep-ph] (cit. on p. 67).

[128] T. Gleisberg and S. Höche, *Comix, a new matrix element generator*,
JHEP **12** (2008) 039, arXiv: 0808.3674 [hep-ph] (cit. on p. 67).

[129] F. Buccioni et al., *OpenLoops 2*, Eur. Phys. J. C **79** (2019) 866,
arXiv: 1907.13071 [hep-ph] (cit. on p. 67).

[130] F. Cascioli, P. Maierhöfer and S. Pozzorini,
Scattering Amplitudes with Open Loops, Phys. Rev. Lett. **108** (2012) 111601,
arXiv: 1111.5206 [hep-ph] (cit. on p. 67).

[131] A. Denner, S. Dittmaier and L. Hofer,
Collier: A fortran-based complex one-loop library in extended regularizations,
Comput. Phys. Commun. **212** (2017) 220, arXiv: 1604.06792 [hep-ph]
(cit. on p. 67).

[132] C. Degrande, F. Maltoni, J. Wang and C. Zhang, *Automatic computations at next-to-leading order in QCD for top-quark flavor-changing neutral processes*,
Phys. Rev. D **91** (3 2015) 034024 (cit. on p. 67).

[133] A. Alloul, N. D. Christensen, C. Degrande, C. Duhr and B. Fuks,
FeynRules 2.0 — A complete toolbox for tree-level phenomenology,
Computer Physics Communications **185** (2014) 2250 (cit. on p. 67).

[134] S. Frixione, E. Laenen, P. Motylinski and B. R. Webber, *Angular correlations of lepton pairs from vector boson and top quark decays in Monte Carlo simulations*,
JHEP **04** (2007) 081, arXiv: hep-ph/0702198 (cit. on p. 67).

[135] P. Artoisenet, R. Frederix, O. Mattelaer and R. Rietkerk,
Automatic spin-entangled decays of heavy resonances in Monte Carlo simulations,
JHEP **03** (2013) 015, arXiv: 1212.3460 [hep-ph] (cit. on p. 67).

[136] M. Czakon and A. Mitov, *Top++: A program for the calculation of the top-pair cross-section at hadron colliders*, Comput. Phys. Commun. **185** (2014) 2930,
arXiv: 1112.5675 [hep-ph] (cit. on p. 68).

[137] N. Kidonakis, *Theoretical results for electroweak-boson and single-top production*,
PoS **DIS2015** (2015) 170, arXiv: 1506.04072 [hep-ph] (cit. on p. 69).

[138] M. Aliev et al.,
HATHOR – HAdronic Top and Heavy quarks crOss section calculatoR,
Comput. Phys. Commun. **182** (2011) 1034, arXiv: 1007.1327 [hep-ph]
(cit. on p. 69).

[139] P. Kant et al., *HatHor for single top-quark production: Updated predictions and uncertainty estimates for single top-quark production in hadronic collisions*,
Comput. Phys. Commun. **191** (2015) 74, arXiv: 1406.4403 [hep-ph]
(cit. on p. 69).

[140] R. Frederix, D. Pagani and M. Zaro, *Large NLO corrections in $t\bar{t}W^{\pm}$ and $t\bar{t}t\bar{t}$ hadroproduction from supposedly subleading EW contributions*,
JHEP **02** (2018) 031, arXiv: 1711.02116 [hep-ph] (cit. on p. 69).

[141] R. Frederix and I. Tsinikos, *On improving NLO merging for* $\mathrm{t\bar{t}W}$ *production*,
JHEP **11** (2021) 029, arXiv: 2108.07826 [hep-ph] (cit. on p. 69).

[142] D. de Florian et al., *Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector*, (2016), arXiv: 1610.07922 [hep-ph]
(cit. on pp. 70, 114).

[143] O. B. Bylund, *Modelling Wt and tWZ production at NLO for ATLAS analyses*, 2016, arXiv: 1612.00440 [hep-ph] (cit. on p. 70).

[144] ATLAS Collaboration, *Performance of the ATLAS track reconstruction algorithms in dense environments in LHC Run 2*, Eur. Phys. J. C **77** (2017) 673, arXiv: 1704.07983 [hep-ex] (cit. on p. 71).

[145] A. Rosenfeld and J. L. Pfaltz, *Sequential Operations in Digital Picture Processing*, J. ACM **13** (1966) 471 (cit. on p. 72).

[146] R. Frühwirth, *Application of Kalman filtering to track and vertex fitting*, Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment **262** (1987) 444 (cit. on p. 72).

[147] ATLAS Collaboration, *Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC*, Eur. Phys. J. C **77** (2017) 332, arXiv: 1611.10235 [hep-ex] (cit. on p. 72).

[148] ATLAS Collaboration, *Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at* $\sqrt{s} = 13$ *TeV*, Eur. Phys. J. C **79** (2019) 639, arXiv: 1902.04655 [physics.ins-det] (cit. on pp. 72, 73).

[149] W. Lampl et al., *Calorimeter Clustering Algorithms: Description and Performance*, ATL-LARG-PUB-2008-002, 2008, URL: https://cds.cern.ch/record/1099735 (cit. on p. 72).

[150] E. N. Sanchez, *Discrete-Time Recurrent Neural Control (Automation and Control Engineering)*, English, Paperback, CRC Press, 2022 292, ISBN: 978-1032338965 (cit. on p. 74).

[151] Y. Coadou, 'Boosted Decision Trees', *Artificial Intelligence for High Energy Physics*, WORLD SCIENTIFIC, 2022 9 (cit. on p. 74).

[152] ATLAS Collaboration, *Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at* $\sqrt{s} = 13$ *TeV*, Eur. Phys. J. C **81** (2021) 578, arXiv: 2012.00578 [hep-ex] (cit. on pp. 74, 75).

[153] ATLAS Collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton–proton collision data*, JINST **14** (2019) P12006, arXiv: 1908.00005 [hep-ex] (cit. on p. 74).

[154] ATLAS Collaboration, *Electron efficiency measurements with the ATLAS detector using 2012 LHC proton–proton collision data*, Eur. Phys. J. C **77** (2017) 195, arXiv: 1612.01456 [hep-ex] (cit. on p. 74).

[155] J. Illingworth and J. Kittler, *A survey of the hough transform*, Computer Vision, Graphics, and Image Processing **44** (1988) 87 (cit. on p. 75).

[156] ATLAS Collaboration, *Jet reconstruction and performance using particle flow with the ATLAS Detector*, Eur. Phys. J. C **77** (2017) 466, arXiv: 1703.10485 [hep-ex] (cit. on p. 75).

[157] M. Cacciari, G. P. Salam and G. Soyez, *The anti-k_t jet clustering algorithm*, JHEP **04** (2008) 063, arXiv: 0802.1189 [hep-ph] (cit. on p. 76).

[158] M. Cacciari, G. P. Salam and G. Soyez, *FastJet user manual*, Eur. Phys. J. C **72** (2012) 1896, arXiv: 1111.6097 [hep-ph] (cit. on p. 76).

[159] ATLAS Collaboration, *Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Eur. Phys. J. C **81** (2020) 689, arXiv: 2007.02645 [hep-ex] (cit. on pp. 76–78).

[160] ATLAS Collaboration, *Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector*, Eur. Phys. J. C **76** (2016) 581, arXiv: 1510.03823 [hep-ex] (cit. on p. 78).

[161] ATLAS Collaboration, *ATLAS flavour-tagging algorithms for the LHC Run 2 pp collision dataset*, Eur. Phys. J. C **83** (2022) 681, arXiv: 2211.16345 [physics.data-an] (cit. on p. 79).

[162] ATLAS Collaboration, *Optimisation and performance studies of the ATLAS b-tagging algorithms for the 2017-18 LHC run*, ATL-PHYS-PUB-2017-013, 2017, URL: https://cds.cern.ch/record/2273281 (cit. on p. 79).

[163] ATLAS Collaboration, *Identification of Jets Containing b-Hadrons with Recurrent Neural Networks at the ATLAS Experiment*, ATL-PHYS-PUB-2017-003, 2017, URL: https://cds.cern.ch/record/2255226 (cit. on p. 79).

[164] ATLAS Collaboration, *Secondary vertex finding for jet flavour identification with the ATLAS detector*, ATL-PHYS-PUB-2017-011, 2017, URL: https://cds.cern.ch/record/2270366 (cit. on p. 79).

[165] ATLAS Collaboration,
Topological b-hadron decay reconstruction and identification of b-jets with the JetFitter package in the ATLAS experiment at the LHC,
ATL-PHYS-PUB-2018-025, 2018,
URL: https://cds.cern.ch/record/2645405 (cit. on p. 79).

[166] ATLAS Collaboration,
Performance of missing transverse momentum reconstruction with the ATLAS detector using proton–proton collisions at $\sqrt{s} = 13$ *TeV*,
Eur. Phys. J. C **78** (2018) 903, arXiv: 1802.08168 [hep-ex]
(cit. on pp. 80, 113).

[167] ATLAS Collaboration, *Observation of four-top-quark production in the multilepton final state with the ATLAS detector*, Eur. Phys. J. C **83** (2023) 496, arXiv: 2303.15061 [hep-ex] (cit. on pp. 84, 85).

[168] G. Breit and E. Wigner, *Capture of Slow Neutrons*, Phys. Rev. **49** (7 1936) 519 (cit. on p. 85).

[169] C. Gauss,
Theoria motus corporum coelestium in sectionibus conicis solem ambientium,
Carl Friedrich Gauss Werke, Sumtibus F. Perthes et I.H. Besser, 1809
(cit. on p. 85).

[170] M. Oreglia, *A Study of the Reactions* $\psi' \to \gamma\gamma\psi$, PhD thesis, 1980
(cit. on p. 85).

[171] ATLAS Collaboration, *Measurement of the total and differential cross-sections of* $t\bar{t}W$ *production in pp collisions at* $\sqrt{s} = 13$ *TeV with the ATLAS detector*, 2024, arXiv: 2401.05299 [hep-ex] (cit. on pp. 85, 88).

[172] G. R. Iversen and M. Gergen, *Statistics*, Springer New York, 1997,
ISBN: 9781461222446 (cit. on pp. 87, 120).

[173] F. James and M. Roos, *Minuit - a system for function minimization and analysis of the parameter errors and correlations*,
Computer Physics Communications **10** (1975) 343 (cit. on p. 87).

[174] ATLAS Collaboration, *Inclusive and differential cross-section measurements of* $t\bar{t}Z$ *production in pp collisions at* $\sqrt{s} = 13$ *TeV with the ATLAS detector, including EFT and spin-correlation interpretations*, 2023,
arXiv: 2312.04450 [hep-ex] (cit. on p. 88).

[175] O. Thielmann, *Search for flavour-changing neutral current interactions in the top-quark Higgs boson sector in multi-lepton final states with the ATLAS detector at the LHC at* $\sqrt{s} = 13$ *TeV*, Bergische Universität Wuppertal, 2023,
URL: https://cds.cern.ch/record/2884767 (cit. on pp. 97, 102).

[176] P. Jackson and C. Rogan, *Recursive jigsaw reconstruction: HEP event analysis in the presence of kinematic and combinatoric ambiguities*, Physical Review D **96** (2017) (cit. on p. 100).

[177] J. Chapmann, *Neural networks*, North Charleston, SC: Createspace Independent Publishing Platform, 2017 (cit. on pp. 103, 104).

[178] M. Feindt, *A Neural Bayesian Estimator for Conditional Probability Densities*, 2004, arXiv: physics/0402093 [physics.data-an] (cit. on p. 105).

[179] M. Feindt and U. Kerzel, *The NeuroBayes neural network package*, Nucl. Instrum. Meth. A **559** (2006) 190 (cit. on p. 105).

[180] R. H. Byrd, P. Lu, J. Nocedal and C. Zhu, *A Limited Memory Algorithm for Bound Constrained Optimization*, SIAM Journal on Scientific Computing **16** (1995) 1190 (cit. on p. 106).

[181] ATLAS Collaboration, *ATLAS b-jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV*, Eur. Phys. J. C **79** (2019) 970, arXiv: 1907.05120 [hep-ex] (cit. on p. 113).

[182] 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*, Eur. Phys. J. C **82** (2022) 95, arXiv: 2109.10627 [hep-ex] (cit. on p. 113).

[183] 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*, Eur. Phys. J. C **83** (2023) 728, arXiv: 2301.06319 [hep-ex] (cit. on p. 113).

[184] ATLAS Collaboration, *Measurement of the production cross-section of a single top quark in association with a Z boson in proton–proton collisions at 13 TeV with the ATLAS detector*, Phys. Lett. B **780** (2018) 557, arXiv: 1710.03659 [hep-ex] (cit. on p. 114).

[185] S. Höche, S. Mrenna, S. Payne, C. T. Preuss and P. Skands, *A Study of QCD Radiation in VBF Higgs Production with Vincia and Pythia*, SciPost Phys. **12** (2022) 010, arXiv: 2106.10987 [hep-ph] (cit. on p. 115).

[186] J. Butterworth et al., *PDF4LHC recommendations for LHC Run II*, J. Phys. G **43** (2016) 023001, arXiv: 1510.03865 [hep-ph] (cit. on p. 115).

[187] J. Neyman and E. Pearson, *On the Problem of the Most Efficient Tests of Statistical Hypotheses*, Philosophical Transactions of the Royal Society, A **231** (1933) 289 (cit. on p. 121).

[188] G. Cowan, K. Cranmer, E. Gross and O. Vitells,
Asymptotic formulae for likelihood-based tests of new physics,
The European Physical Journal C **71** (2011) (cit. on p. 122).

[189] A. L. Read, *Presentation of search results: the CL_S technique*,
J. Phys. G **28** (2002) 2693 (cit. on p. 122).

Acknowledgements

This dissertation would not exist without the immense practical, professional and emotional support of many wonderful people. First, I would like to thank Prof. Dr. Wolfgang Wagner for being the best supervisor one could wish for. You were always available when I was in need of guidance and continuously provided helpful and constructive feedback to bring the analysis forward. Special thanks go also to Dr. Oliver Thielmann for the invaluable collaboration on this analysis. I will forever cherish the great moments we had in our senior executive office, and your continuous triumphs in Kicker, Magic, Skat, or whatever competitive pass-time activities you could think of were simultaneously fun and humbling. To Rukije: I will be eternally grateful for your support at every step of this process. Particularly the last two months, in which you effectively managed my entire life in addition to your own, were extraordinarily stressful. I admire your resilience during this period, and I am anxiously looking ahead to the time when you have to write your dissertation.

I also wish to acknowledge the support of Dr. Dominic Hirschbühl and PD Dr. Frank Ellinghaus, as you continuously offered your assistance in all technical, physical and statistical manners. I am also very grateful to Dr. Christian Preuss for all your private theory lectures during lunch. Moreover, I am deeply indebted to the people who, in all their kindness, proof-read chapters of this dissertation: Anna, Christian, Dominic, Johanna, Lukas, Marc, Maren, Oliver, Shayma, Tim and Tobias – Thank you very much! On a personal note, I would like to thank all the members of the WHEP group, who were a source of both immense professional support and delightful entertainment throughout my time here. Specifically, I would like to thank Dr. Carsten Dülsen for getting me into the group in the first place. And special thanks to Anna and Johanna for putting up with me as an office mate for the last few months and accepting all of my rubbish into your immaculate kingdom.

I would also like to express my gratitude to the rest of the tHq FCNC Multilepton analysis team. The last four years of collaboration were very fruitful and enjoyable. Additionally, it would not have been possible to conduct most of the analysis the way we did without a functioning computing cluster. The cluster team continuously worked hard to keep it up and running. Thank you very much for this.

I would also like to thank Dr. Carmen Diez Pardos for agreeing to be the second referee of this dissertation, and Prof. Dr. Dirk Lützenkirchen-Hecht and Prof. Dr. Christian

Hölbling for being part of the examination committee. Reading such a long document, preparing for the examination and the examination itself take up a lot of time from your already busy schedules. I genuinely appreciate that, nonetheless, you agreed to occupy these positions.

Finally I want to express deep gratitude to my family for their support, not just for this dissertation, but throughout my entire life. Mom and dad, you kept me fed for many years and taught me how important it is to work hard towards your goals. Jordan and Jason, you are among the most dependable people in my life and I am proud to see you having grown into such proud young men. Thank you all.