Measurement of Higgs couplings to top quarks and τ leptons with the ATLAS detector

by

Martina Laura Ojeda

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Department of Physics
University of Toronto

Abstract

Measurement of Higgs couplings to top quarks and τ leptons with the ATLAS detector

Martina Laura Ojeda
Doctor of Philosophy
Department of Physics
University of Toronto
2021

Since the discovery of the Higgs boson in 2012 during the first run of the Large Hadron Collider (LHC) at CERN, extensive measurements have been performed to characterize its properties. Its mass, spin, parity and couplings have been studied in detail with no significant deviations from the SM predictions observed. The second run of the LHC saw incredible progress with the addition of new measurements. This thesis highlights the measurements of Higgs couplings to third generation fermions, first with the observation of the $t\bar{t}H$ process and then through improved $pp \to H \to \tau\tau$ cross section measurements. The search for the $t\bar{t}H$ process is discussed in the context of the $t\bar{t}H \to$ multilepton analysis, with an emphasis on final states with $\tau_{\text{had-vis}}$. A $t\bar{t}H$ cross section measurement, obtained in combination with results from $t\bar{t}(H \to \gamma\gamma)$, $t\bar{t}(H \to ZZ^* \to 4\ell)$ and $t\bar{t}(H \to b\bar{b})$, is also presented. All results are found to be compatible with the Standard Model predictions. The study of Higgs decays involving τ leptons is expanded upon with the measurement of the $pp \to H \to \tau\tau$ cross section using the full Run 2 dataset, and a new method to estimate the largest background to this measurement is introduced. The improvement in the precision of $pp \to H \to \tau\tau$ measurements coupled with an increased sensitivity to broader regions of phase space allows for a deeper scrutiny of the Standard Model predictions, with the theory still providing an accurate description of the experimental results.

Acknowledgements

I have been extremely lucky to have met a lot of amazing people throughout my studies. To all who have taken the time to sit with me for a coffee in B40, to the ones whose office door was always open for me, and to those who have been on the receiving end of my messages at all times of day: thank you. While I can only name a few people here, I am eternally grateful for all the kind gestures and friendly faces I've encountered over the last couple of years.

First and foremost, j'aimerais remercier mon superviseur Pierre Savard. Thank you for the guidance and the laughs through these 5 years. Thank you for trusting me and giving me the freedom to pursue whatever projects I wished, while also providing insight and advice. Thank you for sharing your knowledge with me and always being willing to chat about physics (and maple syrup). It has been a pleasure working with and learning from you.

Thank you to Mike Luke and Pekka Sinervo for being part of my committee. You have made the graduation process feel fun even in the middle of a pandemic, and that is no easy feat.

A special thank you to the many people who have mentored me over the years: Fritz Buchinger and, particularly, John Crawford for shipping me off to TRIUMF as a young student and letting me dive into research. Thank you for believing in me from day one and giving me this amazing opportunity. Thank you also to François Corriveau and the McGill ATLAS group for doing the same thing a year later. I wouldn't be where I am today without you.

My first analysis in ATLAS was with the $t\bar{t}H$ group and I would like to thank the many people from that group who have provided me with guidance throughout my first two years as a PhD student. In particular, a big thanks goes to Tamara Vázquez Schröder for introducing me to the world of physics analysis. To Albert Casha and David Hohn, who somehow always ended up on similar tau-related projects, thanks for being so much fun to work with.

The tau CP and the upgrade physics groups have also been very welcoming. I have had the pleasure of meeting many amazing people who not only taught me a lot but also made this work so enjoyable. Thanks to everyone who has guided me and worked with me, and specifically to Pier-Olivier DeViveiros for dragging me into the world of taus.

Thanks to Pier-Olivier and Quentin Buat for then bringing me to the world of $H \to \tau\tau$ and making me a part of embedding. Working with you has been so fun and interesting and challenging (in the best possible way), and I feel lucky to have ended my PhD on such a high note. I cannot thank you enough. Thanks also to Antonio De Maria for always being only a Skype message away when I was banging my head against the wall while staring at code, and to the rest of the HLeptons team.

Being at CERN has been the highlight of my PhD, and I want to thank some of the people who have made it feel like home: Alison, Benjamin, Dan, Dirk, Heather, Mark, Sam, Stef, Tamara, and the many other people I met during my time at CERN. Thanks also to my UofT friends for making my time in Toronto (and even in the St Genis apartment) a fun one. A special thanks goes to Haider for all the help and support, and for always being enthusiastic about any and all ideas. Many thanks to Bianca for being such a great friend: I'm so glad we bonded over our fear of those terrifying stairs (and that massive milkshake). Huge thanks (again) to PO for being the greatest mentor, personal postdoc and friend. Thank you for always being around and for the sometimes serious but often not so serious chats.

I wouldn't have been able to achieve what I did without the help of my medical team, who have gone above and beyond to ensure I would be able to do what I set my mind and heart to. Thank you for everything.

Thanks also to all my non-CERN friends who have been supporting me since before this adventure began and who spent hours on video calls with me during it. A particularly big thank you to all of my Montreal friends who have gone on so many coffee, lunch and dinner dates with me over the years. Thank you many times to Devin and Jerry for housing me whenever I needed...I will miss your couch dearly. Thanks to Isabelle for many things, from providing moral support during never-ending quantum assignments in Montreal to feeding me cannelés in Toronto. Finally, thank you Bernice for being an amazing friend and supporting me no matter what, but especially for coming to visit me in Geneva and reminding me that you regret the risks you don't take.

A massive thank you to Stefan for, well, everything. Thank you for the big things, like reading this thesis basically as many times as I have and dealing with all the typesetting issues. Thank you for letting me talk to you about taus for hours on end. Thank you also for the smaller things, like convincing me that I needed to have sushi in Geneva and that I should definitely stop putting sugar in my tea. Most of all, thank you for always being there.

Y finalmente, quiero agradecer a mi familia: a los que tengo acá cerca y a los que me acompañan desde lejos. Gracias por estar siempre presentes. Gracias a mis padres Diana y Rubén, y a mis hermanos Federico y Rodrigo, por alentarme y por compartir este sueño conmigo. Esta tesis se la quiero dedicar a ustedes.

Contributions

Due to the complexity of the ATLAS detector and the elaborateness of the analysis of the data it produces, results from the ATLAS collaboration rely on the combined work of thousands of scientists across the world. Their contributions consist of, among other things, building and maintaining the detector, translating its signals into physics objects, identifying and calibrating those objects, maintaining software and computing resources, and performing the various steps required to mold and complete a physics analysis. ATLAS results, including those presented in this thesis, are therefore only possible through the sum of contributions from every member of the collaboration.

With this in mind, the elements and results to which I have contributed are outlined below.

Chapter 8: Observation of the associated production of a Higgs boson and a top quark pair ($t\bar{t}H$)

At the start of my PhD, the discovery of the $t\bar{t}H$ process was one of the highest priorities of the ATLAS collaboration. A paper highlighting evidence for this process had recently been published with the available dataset [58] and a new round of analysis was being prepared, aiming to achieve observation of this process through a combination of more advanced analysis methods, a more extensive set of final states considered and a larger dataset.

My involvement with the $t\bar{t}H$ group consisted mainly in the development and improvement of the analysis in channels involving hadronic decays of τ leptons. I implemented a new search channel with three light leptons and one hadronically-decaying τ, enabling this final state to be studied for the first time[1]. I was involved in all aspects of the analysis (in collaboration with other members of the group), including optimizing the event and object selection, determining and carrying out the background estimations and performing the statistical interpretation of the data. The $3\ell + 1\tau_{\text{had}}$ channel was added to the $t\bar{t}H$ search in the 36 fb^{-1} paper [29] and contributed to the $t\bar{t}H$ discovery [26]. I was also involved with Monte Carlo simulation validation[2], development of several techniques and cross-checks for τ background estimations, and correlation of systematic uncertainties across channels. The inclusion of τ channels improved the significance of the $t\bar{t}H$ 36 fb^{-1} result by 25%.

With the observation of the $t\bar{t}H$ process achieved, the analysis transitioned from a search to a measurement. I therefore shifted my focus to the reconfiguration of channels with hadronically-decaying τ leptons in the final state, harmonising the analysis strategy across all τ channels in $t\bar{t}H$ (object selections, single-τ fake estimates, fit strategy) and harmonising the analysis code within the various analysis frameworks to produce the first $t\bar{t}H \to$ multilepton results using the ATLAS software release intended for all publications using the full Run 2 dataset[3]. This release change introduced updates to object reconstruction and identification, affecting in particular the τ channels' background estimations. I was therefore tasked with ensuring good communication and collaboration within the analysis team and with the τ working group, as the $t\bar{t}H$ group was one of the first analysis groups scrutinising hadronically-decaying τ leptons in the new release. During that time I also maintained and developed the $t\bar{t}H$ derivation code, which takes the petabyte-scale output from the ATLAS reconstruction step and generates smaller samples, used by many groups studying processes involving both top quarks and Higgs bosons.

[1] I performed feasibility studies for $t\bar{t}H$ searches in final states with two light leptons and two hadronic τ leptons ($2\ell + 2\tau_{\text{had}}$) and three light leptons and one hadronic τ lepton ($3\ell + 1\tau_{\text{had}}$) to increase the sensitivity of the $t\bar{t}H$ search. The latter seemed the most promising and was therefore the focus of my work within the $t\bar{t}H$ group.

[2] In the context of [29], I was responsible for the validation of τ lepton modelling in the ATLAS Fast Simulation (Atlfast-II) [36] in the phase-space of the $t\bar{t}H \to$ multilepton analysis.

[3] This work is not presented in this thesis but can be found at [63].

Chapter 9: Higgs decays to τ leptons in other production modes

The SM $H \to \tau\tau$ cross section measurement depends on a reliable estimate of its largest background, $Z \to \tau\tau$. Due to the similarity between the two processes, the latter cannot be studied without looking at the $H \to \tau\tau$ signal. Alongside two contributors, I developed a new object-based embedding method that enables the use of $Z \to \ell\ell$ events as a proxy for $Z \to \tau\tau$, allowing for a thorough study of this process in a phase-space relevant to the analysis and for the determination of its normalisation using data, resulting in the improvement of both the reliability and the precision of the estimate of the $Z \to \tau\tau$ contributions to the $H \to \tau\tau$ measurement.

This method is based on the embedding method used during Run 1 [42] but aims to provide a more light-weight approach that is easier to implement within the current ATLAS workflow. A description of the new method and a thorough validation of all distributions of interest within the phase space of the $H \to \tau\tau$ analysis is provided in Chapter 9. A discussion of the method in the context of the analysis and the $H \to \tau\tau$ cross section results themselves are also presented.

Chapter 10: High-Luminosity LHC

Before analyses such as $t\bar{t}H$ or $H \to \tau\tau$ can use objects like hadronically-decaying τ leptons, combined performance groups must develop algorithms that enable the reconstruction and identification of such objects. My work was part of the τ lepton combined performance group's efforts to prepare for the detector upgrade that is scheduled for 2027. I performed and oversaw all studies, optimisations and documentation of the performance of τ lepton reconstruction, identification and calibration algorithms for the Phase-II upgrade of the ATLAS detector. Many of these results were documented in the various technical design reports for the Phase-II upgrade [30, 31, 55]. In the context of these studies I oversaw the work of two other students and was responsible for outlining the projects they worked on. This project also involved collaboration with different performance, physics analysis and detector R&D groups in order to coordinate efforts and provide the best possible τ performance estimate.

I also participated in studies aiming to assess the impact of various detector configurations and environment conditions on the quality of the identification and calibration of τ leptons, for example in the context of the luminous region studies presented in Chapter 10. These studies assess the impact of varying the pile-up rates and density and help outline the effects of introducing crab cavities on the physics potential of the HL-LHC.

Finally, I derived all needed parametrised τ efficiency and fake rate estimates to support sensitivity projections by physics analysis groups presented to the European Strategy for Particle Physics, documented in the Yellow Reports [20, 21, 22, 23, 27]. An additional document outlining the details of the object-level performance used as inputs to the Yellow Reports is available at [68].

Chapter 6: Tau lepton reconstruction, identification and calibration

Lastly, I participated in the maintenance and development of the derivation code for the τ performance group for a few years. My work consisted among other things in reducing the size of the samples to accommodate the increasing dataset size, helping with migration to a new software release and ensuring samples were produced as needed by the group so all necessary studies could be carried out.

Contents

Chapter 1

Introduction

The Standard Model (SM) of particle physics is a theory that describes fundamental particles and their interactions. It expands upon a quantum theory proposed and developed during the 20^{th} century that provided an excellent description of the electromagnetic force. Issues arose when trying to expand the theory to include the weak force, as the apparent presence of massive weak bosons couldn't be accounted for within the theory. In 1964, a new mechanism was proposed to provide an explanation for how particles acquire their mass. The introduction of this mechanism into the theory and an extension of the theory to include the strong force gave rise to the SM as it is known today.

One of the predictions of the SM was the existence of the Z boson, discovered at CERN in 1983. Another prediction was the existence of a massive Higgs boson. Successes such as the discovery of the Z boson motivated its search using particle collider experiments. Higgs searches began in the 1990s at CERN's Large Electron-Positron Collider (LEP) and continued into the 2000s at Fermilab's Tevatron, with no evidence of Higgs boson production found in either. In 2012, half a century after it was first theorised, the Higgs boson was finally discovered by the ATLAS and CMS collaborations at CERN during the first run of the Large Hadron Collider (LHC) [25, 75]. With its mass measured, all of its properties, and in particular its coupling to other particles, could be fixed within the theory. These predictions were then compared to the experimentally-obtained values, allowing for extensive searches for deviations from the SM that could indicate the presence of new physics. Since the Higgs boson was discovered mainly through its bosonic decay modes, its couplings to bosons was well understood in Run 1 and found to agree with the SM predictions within the measurements' uncertainties. Questions remained however regarding its couplings to fermions at the beginning of Run 2.

The heaviest fermion in the SM is the top quark and its coupling to the Higgs boson is therefore expected to be the strongest, making it an ideal candidate to probe fermionic couplings. While evidence of Higgs couplings to the top quark was found in Run 1 due to the fact that the Higgs' main production mode, gluon fusion (ggH), proceeds predominantly through a top quark loop, the possibility remained that contributions to the loop by processes not described by the SM affected the top quark coupling measurements derived from ggH. The observation of the $t\bar{t}H$ process, in which a Higgs boson is produced in association with a pair of top quarks, allows for a complementary test of this coupling.

While the SM provides an excellent description of particle physics, it is known to have limitations, which include its inability to describe gravity, dark matter and the matter-antimatter asymmetry, among other things, and motivate the proposal of extensions to the model. Some extensions to the SM predict deviations to the fermionic couplings, or deviations that could affect quarks and leptons differently. It is therefore important to specifically measure leptonic Higgs couplings. Leptons are generally light and therefore have small couplings to the Higgs boson. The τ lepton is the heaviest of the leptons, making its relatively large coupling the first that can be measured at the LHC. Precise measurements of the rate and kinematics of Higgs decays to τ lepton pairs provide additional handles to probe the validity of the SM.

Studies involving τ leptons face a variety of obstacles. Unlike the other members of the lepton family, τ leptons decay before they reach the detectors. This makes their identification significantly more challenging, and causes the determination of their energy to rely on a detailed understanding of the particles produced during the decay as well as their interaction with the detectors. The challenges related to the reconstruction, identification and energy calibration of τ leptons become more important as the LHC's radiation environment becomes harsher due to increased pile-up rates, with significant implications on the measurements that can be made at future colliders such as the High-Luminosity LHC (HL-LHC).

This thesis presents a study of Higgs couplings to third generation fermions through the pursuit of the observation of the $t\bar{t}H$ process using part of the Run 2 dataset and with a particular emphasis on final states involving $H \to \tau\tau$ decays. The challenges associated to the estimation of background contributions when τ leptons are part of the final state are of notable interest. The main result is the observation of the $t\bar{t}H$ process and a measurement of its cross section. Then, a more in-depth study of $H \to \tau\tau$ decays using the full Run 2 dataset becomes the focus, with the improvement of the measurement of the $pp \to H \to \tau\tau$ cross section as its main goal. In this case, a thorough understanding of the $Z \to \tau\tau$ process, which results in identical final states to those of $H \to \tau\tau$, is key.

This thesis begins with an introduction to the SM in Chapter 2, followed by an overview of the LHC accelerator complex and its environment in Chapter 3, and of the ATLAS detector in Chapter 4. A discussion of particle reconstruction, identification and calibration follows in Chapter 5, with an emphasis in Chapter 6 on the particularities associated to τ leptons. An overview of Higgs results from Run 1 is then presented in Chapter 7. The main results presented in this thesis consist of the observation of the $t\bar{t}H$ process using a partial Run 2 dataset (Chapter 8) and of the measurement of the $pp \to H \to \tau\tau$ cross section using the full Run 2 dataset (Chapter 9). A discussion of the challenges related to τ lepton reconstruction, identification and calibration at the HL-LHC is presented in Chapter 10.

Chapter 2

The Standard Model

Physics aims to understand and describe how the universe works on a fundamental level. To that end, many theories and mathematical models have been proposed. The Standard Model (SM) of Particle Physics provides the theoretical framework that constitutes what is currently accepted as the best description of particle physics, and aims to describe all elementary particles of matter as well as their interactions[1]. This chapter will provide an introduction to the SM, but more extensive descriptions can be found in [11, 85, 86, 108].

The SM proposes a structure where matter is associated to half-integer spin particles called fermions whose interactions are mediated by force-carrying integer-spin bosons. During the 20$^{\text{th}}$ century, a model that described only electromagnetic interactions was proposed but it was expanded with time to also describe the weak and strong forces. The evolution of the theory occurred thanks to both theoretical and experimental insights, alternating between particles being proposed through extensions to the theory and their existence being later confirmed by experiments, and experiments discovering particles that were then added to the theory.

At the beginning of the 21$^{\text{st}}$ century the structure of the SM, as confirmed by experiments, was as follows. The SM contained three types of bosons: the photon (γ), which mediates electromagnetic (EM) interactions, the $W^{\pm}$ and Z bosons, which mediate weak interactions[2], and the gluons (g), which mediate strong interactions. Fermions were classified according to their interactions with these bosons with quarks (q) being fermions that interact via all three forces, while leptons are those that do not interact via the strong force. Leptons themselves were categorized as either charged leptons (ℓ) or neutrinos (ν) depending on whether they interact via both forces (the former) or only via the weak force (the latter).

In the SM, Lagrangian densities (or simply Lagrangians) are used to describe the dynamics of the theory. They are functions constructed using the fields associated to the various particles as well as their derivatives. The equations of motion corresponding to the Lagrangians describe the kinematics and interactions of these particles within the theory.

[1]Except for gravity, which is described by the theory of General Relativity.
[2]The EM and weak forces are often referred to jointly as the electroweak force.

2.1 Symmetries

Within the Standard Model, interactions arise from the symmetries embedded into the theory. The SM is a non-abelian gauge theory, with the gauge group $\mathcal{G}_{\text{SM}} = SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$ describing the interaction of the fermions and bosons mentioned above. The term gauge refers to redundant degrees of freedom of a Lagrangian, with gauge (or symmetry) groups describing the set of transformations under which the Lagrangian is invariant, meaning that it describes the same fields regardless of the specific gauge chosen. These sets of transformations are called the symmetries of the theory.

2.1.1 Charges

Charges arise according to the transformation properties of the various particles under the different symmetries. States that transform non-trivially under the relevant transformation are said to be charged and interact with the associated force while the opposite is true for states that transform trivially.

- The charge of $SU(3)_C$ is called colour and can take the form of red, green or blue ($r,\ g,\ b$).

- The charge associated with $SU(2)_L$ is the weak isospin T: its third component T_3 is conserved in weak interactions. The chirality of a particle indicates how it transforms under the weak symmetry group. Left-handed chirality eigenstates interact with the weak force while right-handed eigenstates do not.

- The weak hypercharge Y_W is the charge associated to the $U(1)_Y$ symmetry group and is is a combination of the weak isospin and the electric charge Q of a particle, with $Y_W = 2(Q - T_3)$.

2.1.2 Classification of fermions

Fermions are categorized into three generations of particles with identical charges and differing masses as shown in Figure 2.1. An anti-particle exists for every particle shown, with identical mass and opposite charge.

The charges of each type of fermion are summarized in Table 2.1. The following notation is useful when referring to fermions:

$$\begin{aligned} \boldsymbol{Q}_L &= \left(\begin{pmatrix} u_L \\ d_L \end{pmatrix}, \begin{pmatrix} c_L \\ s_L \end{pmatrix}, \begin{pmatrix} t_L \\ b_L \end{pmatrix} \right) \\ \boldsymbol{u}_R &= (u_R, c_R, t_R) \\ \boldsymbol{d}_R &= (d_R, s_R, b_R) \end{aligned} \quad \text{and} \quad \begin{aligned} \boldsymbol{L}_L &= \left(\begin{pmatrix} \nu_{eL} \\ e_L \end{pmatrix}, \begin{pmatrix} \nu_{\mu L} \\ \mu_L \end{pmatrix}, \begin{pmatrix} \nu_{\tau L} \\ \tau_L \end{pmatrix} \right) \\ \boldsymbol{\ell}_R &= (e_R, \mu_R, \tau_R) \end{aligned} \tag{2.1}$$

$\boldsymbol{Q}_L$ refers to left-handed quarks while $\boldsymbol{u}_R$ and $\boldsymbol{d}_R$ refer to right-handed quarks, reflecting the fact that right-handed quark fields transform trivially as singlets under $SU(2)_L$ while left-handed quark fields transform as doublets. Similarly, left-handed leptons are denoted $\boldsymbol{L_L}$ and right-handed leptons are denoted $\boldsymbol{\ell_R}$. Right-handed neutrinos are not included here.

Particle	Q	T_3	Y_W	C
$\boldsymbol{Q_L}$: left-handed quarks	$2/3$	$1/2$	$1/3$	r, g, b
	$-1/3$	$-1/2$	$1/3$	r, g, b
$\boldsymbol{u_R}$: right-handed up-type quarks	$2/3$	0	$4/3$	r, g, b
$\boldsymbol{d_R}$: right-handed down-type quarks	$-1/3$	0	$-2/3$	r, g, b
$\boldsymbol{L_L}$: left-handed leptons	0	$1/2$	-1	none
	-1	$-1/2$	-1	none
$\boldsymbol{\ell_L}$: right-handed leptons	-1	0	-2	none

Table 2.1: Charges of fermions: electric charge (Q), weak isospin (T_3) and weak hypercharge (Y_W), and colour charge (C).

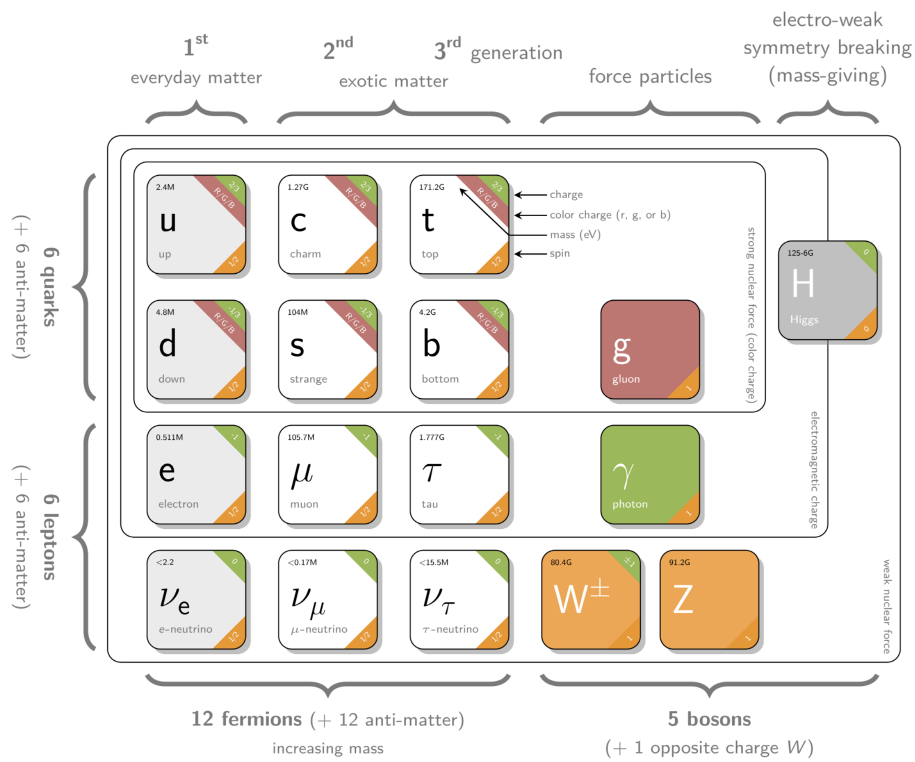

Figure 2.1: Elementary particles of the SM, categorized into half-integer spin fermions, and scalar (spin 0) and vector (spin 1) bosons. Image adapted from [18].

2.1.2.1 Fermion decays

Heavier fermions can decay into their lighter counterparts through the weak force. For example τ leptons, the heaviest of the leptons at 1.77 GeV, decay via a W boson as $\tau \to W\nu_\tau$. Because of the τ lepton's large mass, the W can decay either to a lepton and a neutrino as $W \to \ell\nu$ or to two quarks as $W \to q\bar{q}'$. This is in contrast to muons, which only decay as $\mu \to W\nu_\mu \to e\bar{\nu}_e\nu_\mu$. A similar process occurs for quarks. The heaviest member of the quark family, the top quark, characteristically decays through the W boson to a bottom quark.

2.1.3 Building a Lagrangian density

The SM Lagrangian density is created by writing down all terms that can be built from the fields associated to the fermions and bosons described above. Care must always be taken to sure that the symmetries of the theory are respected. This has implications on the characteristics of the particles and processes described by the SM.

To illustrate this, let's begin with the simple example of the Dirac Lagrangian for a fermion field ψ

$$\mathcal{L}_D = i\bar{\psi}\gamma^\mu\partial_\mu\psi - m\bar{\psi}\psi. \tag{2.2}$$

Under a local gauge transformation $\psi \to \psi' = e^{i\theta}\psi$, this Lagrangian behaves as

$$\begin{aligned}
\mathcal{L}_D \to \mathcal{L}'_D &= ie^{-i\theta}\bar{\psi}\gamma^\mu\partial_\mu(e^{i\theta}\psi) - me^{-i\theta}\bar{\psi}e^{i\theta}\psi && (2.3)\\
&= ie^{-i\theta}\bar{\psi}(e^{i\theta}\gamma^\mu\partial_\mu\psi + i\gamma^\mu(\partial_\mu\theta)e^{i\theta}\psi) - m\bar{\psi}\psi && (2.4)\\
&= i\bar{\psi}\gamma^\mu\partial_\mu\psi - \bar{\psi}\gamma^\mu(\partial_\mu\theta)\psi - m\bar{\psi}\psi && (2.5)\\
&= \mathcal{L}_D - \bar{\psi}\gamma^\mu(\partial_\mu\theta)\psi. && (2.6)
\end{aligned}$$

which is manifestly not invariant under the symmetry transformation. Invariance can however be achieved through an extension of the partial derivative ∂_μ into a covariant derivative $D_\mu \equiv \partial_\mu + iA_\mu$.

The Dirac Lagrangian becomes

$$\begin{aligned}
\mathcal{L}_D &= i\bar{\psi}\gamma^\mu(\partial_\mu + iA_\mu)\psi - m\bar{\psi}\psi && (2.7)\\
&= i\bar{\psi}\gamma^\mu\partial_\mu\psi - \bar{\psi}\gamma^\mu A_\mu\psi - m\bar{\psi}\psi && (2.8)
\end{aligned}$$

and the transformation is

$$\begin{aligned}
\mathcal{L}_D \to \mathcal{L}'_D &= ie^{i\theta}\bar{\psi}\gamma^\mu(\partial_\mu + iA_\mu)e^{-i\theta}\psi - me^{i\theta}\bar{\psi}e^{-i\theta}\psi && (2.9)\\
&= i\bar{\psi}\gamma^\mu\partial_\mu\psi - \bar{\psi}\gamma^\mu A_\mu\psi - \bar{\psi}\gamma^\mu(\partial_\mu\theta)\psi - m\bar{\psi}\psi && (2.10)\\
&= i\bar{\psi}\gamma^\mu\partial_\mu\psi - \bar{\psi}\gamma^\mu(A_\mu + \partial_\mu\theta)\psi - m\bar{\psi}\psi && (2.11)\\
&= \mathcal{L}_D - \bar{\psi}\gamma^\mu(A_\mu + \partial_\mu\theta)\psi && (2.12)
\end{aligned}$$

As long as A_μ transforms as $A_\mu \to A_\mu - \partial_\mu\theta$ the Lagrangian is invariant under the gauge transformation and one can recover Equation 2.8.

The requirement of gauge invariance therefore results in the introduction of a new field A_μ. It turns out that in this example of a theory of Dirac fields, the transformation is that of $U(1)$ and the new field that it gives rise to is equivalent to the photon field. Similar arguments can be used for the $SU(2)$ and the $SU(3)$ symmetries and give rise to $N^2 - 1$ fields, corresponding to the $2^2 - 1 = 3$ weak bosons ($W^\pm$ and Z) and the $3^2 - 1 = 8$ gluons of the strong force.

One interesting consequence of this procedure is that the fields must be massless. Indeed, a mass term of the form $\frac{1}{2}m^2A^\mu A_\mu$ transforms as

$$\frac{1}{2}m^2A^\mu A_\mu \to \frac{1}{2}m^2(A^\mu - \partial^\mu\theta)(A_\mu - \partial_\mu\theta) \tag{2.13}$$

which is not invariant. This is not a problem for the photon and gluons which are known to be massless, but it presents important consequences for the weak bosons. Another point to note is that weak transformations affect left- and right-handed states differently. The fermion mass term from Equation 2.2 becomes

$$m\bar{\psi}\psi = m(\bar{\psi}_L + \bar{\psi}_R)(\psi_L + \psi_R) \tag{2.14}$$
$$= m(\bar{\psi}_L\psi_R + \bar{\psi}_R\psi_L). \tag{2.15}$$

These terms couple left- and right-handed states and violate $SU(2)_L \times U(1)_Y$. This means that mass terms for fermions that couple to the weak force are also forbidden.

Experimental evidence shows that fermions and weak bosons have non-zero masses. This is illustrated for example by the limited range of the weak interaction. Starting with

$$\Delta x = c\Delta t \tag{2.16}$$

and using the relation $\Delta E \Delta t = \hbar$, one obtains that

$$\Delta x = c\frac{\hbar}{\Delta E}. \tag{2.17}$$

and through the relation $E = mc^2$, it can be shown that

$$\Delta x = \frac{\hbar}{mc}. \tag{2.18}$$

Therefore, massless bosons are typically[3] associated to forces with infinite ranges ($\lim_{m\to 0} \Delta x = \infty$) while short range forces imply massive bosons. Assuming a range of about 10^{-3} fm for weak interactions[4], the weak boson masses would therefore be expected to be of about

$$\begin{aligned} m &= \frac{\hbar}{\Delta x} = \frac{\hbar}{10^{-18}\ \text{m}} \\ m &\approx 100\ \text{GeV} \end{aligned} \tag{2.19}$$

which agrees with the measured values for the $W^{\pm}$ and Z bosons.

2.2 The Brout-Englert-Higgs mechanism

A mechanism was proposed by Brout, Englert, Higgs, Guralnik, Kibble and Bernstein [7, 90, 100, 102] to solve this issue through the introduction of scalar fields. A complex $SU(2)_L$ doublet Φ is introduced, with its Lagrangian given by

$$\begin{aligned} \mathcal{L}_{\text{scalar}} &= (D_\mu\Phi)^\dagger(D^\mu\Phi) - V(\Phi) \\ &= (D_\mu\Phi)^\dagger(D^\mu\Phi) - \left[-\mu^2\Phi^\dagger\Phi + \lambda(\Phi^\dagger\Phi)^2\right] \quad \text{with} \quad \mu^2, \lambda > 0. \end{aligned} \tag{2.20}$$

In this case the covariant derivative D_μ is expanded to contain terms related to all electroweak fields. $V(\Phi)$ is the Higgs self-interaction term, and is described by the most general fourth-order polynomial

[3] While mediated by massless gluons, the strong force is an exception to this rule. Its range is limited through other processes (see discussion regarding confinement in Section 2.3.2).

[4] At this range, the strength of the weak interaction is equivalent to that of the EM force.

in the Higgs field that does not violate gauge invariance. It generates a potential with an infinite set of degenerate ground states $\langle\Phi\rangle_0$ shown in Figure 2.2 and satisfying

$$\langle\Phi^\dagger\Phi\rangle_0 = \frac{\mu^2}{2\lambda} \equiv \frac{v^2}{2}, \tag{2.21}$$

where v is called the vacuum expectation value (VEV). It is possible to define a ground state of Φ such that

$$\langle\Phi\rangle_0 = \frac{1}{\sqrt{2}}\begin{pmatrix}0\\ v\end{pmatrix}. \tag{2.22}$$

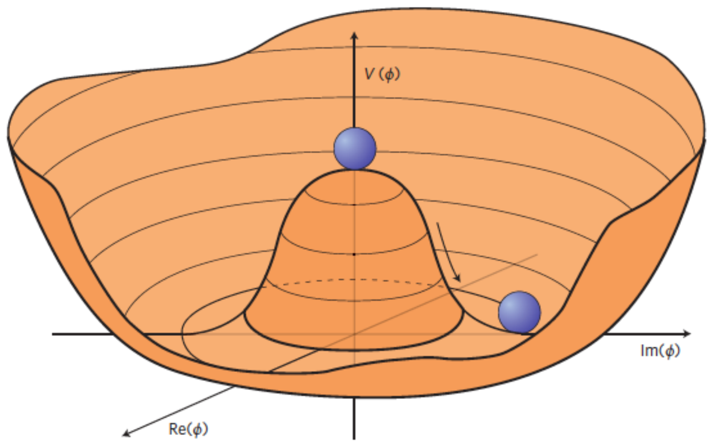

Figure 2.2: Shape of the Higgs potential with the blue ball starting at $\Phi = 0$ and rolling down to the ground state. Image from [88].

and one can expand Φ about this ground state as

$$\Phi = \frac{1}{\sqrt{2}}\exp\left(i\frac{\tau_i}{2}\theta^i(x)\right)\begin{pmatrix}0\\ v+h(x)\end{pmatrix}. \tag{2.23}$$

Two new fields are included here. The $\theta^i(x)$ fields describe scalar massless fields and the massless particles associated to it are called Goldstone bosons. In the unitary gauge, where $\theta^i(x) = 0$, they become the longitudinal polarisation of the massive weak bosons. The scalar field $h(x)$ is identified as the Higgs field. Since the covariant derivative D_μ contains terms related to the electroweak fields, the term $D_\mu\Phi$ introduces interactions between the weak bosons and the Higgs (and its VEV). The scalar Lagrangian can be written out to explicitly show these terms as follows

$$\begin{aligned}\mathcal{L}_{\text{scalar}} = &\frac{1}{2}\left(\partial^\mu h(x)\right)\left(\partial_\mu h(x)\right) - \lambda v^2 h^2(x) - \lambda v h^3(x) - \frac{\lambda}{4}h^4(x)\\ &+ \frac{1}{2}\frac{g_2^2 v^2}{4}\left(\left|W_\mu^+\right|^2 + \left|W_\mu^-\right|^2\right) + \frac{1}{2}\frac{g_2^2 v^2}{4\cos^2\theta_W}\left|Z_\mu\right|^2 + \ldots\end{aligned} \tag{2.24}$$

This structure highlights the fact that the weak boson masses are obtained through this mechanism by terms that involve both the weak bosons and the VEV. The masses are

$$m_A = 0, \quad m_{W^\pm} = \frac{1}{2}g_2 v, \quad m_Z = \frac{1}{2}\frac{g_2 v}{\cos\theta_W} = \frac{m_{W^\pm}}{\cos\theta_W}, \quad m_H = \sqrt{2\lambda v^2}. \tag{2.25}$$

Higher order terms include

$$\mathcal{L} \ni m_{W^\pm}^2 W^{+\mu} W_{-\mu} \left(1 + \frac{h(x)^2}{v^2}\right)^2 + \frac{1}{2} m_Z^2 Z^\mu Z_\mu \left(1 + \frac{h(x)^2}{v^2}\right)^2 \tag{2.26}$$

which are couplings of the Higgs field to the weak bosons. These terms highlight the fact that in the SM Higgs couplings to the weak vector bosons ($V = W, Z$) are expected to be of the form m_V^2/v^2. The vacuum expectation value of the Higgs field can be extracted from W mass and g_2 measurements and is found to be $v = \frac{2m_W}{g_2} \approx 246$ GeV, while the Higgs mass can't be predicted since the value of λ is unknown.

2.2.1 Yukawa couplings

While the masses of the weak vector bosons are a direct consequence of the introduction of a Higgs field, fermion masses are acquired via Yukawa interactions. The addition of a Yukawa term to the Lagrangian,

$$-\mathcal{L}_{\text{Yukawa}} = Y_{ij}^\ell \bar{L}_L^i \Phi \ell_R^j + \bar{Q}_L^i Y_{ij}^u \tilde{\Phi} u_R^j + Y_{ij}^d \bar{Q}_L^i \Phi d_R^j + ... \tag{2.27}$$

allows for a coupling between the fermion and Higgs fields that results in fermion masses without explicit symmetry-violating mass terms, as before via a coupling to the VEV of the Higgs field. Expanding the first term using $\Phi = \frac{1}{\sqrt{2}}(v + h(x))$ one obtains

$$Y_{ij}^\ell \bar{L}_L^i \Phi \ell_R^j = Y_{ij}^\ell \bar{L}_L^i \frac{1}{\sqrt{2}}(v + h(x)) \ell_R^j \tag{2.28}$$

$$= \frac{1}{\sqrt{2}} v Y_{ij}^\ell \bar{L}_L^i \ell_R^j + \frac{1}{\sqrt{2}} Y_{ij}^\ell \bar{L}_L^i h(x) \ell_R^j \tag{2.29}$$

from which the masses of the fermions can be read off as

$$m_{ij}^{\ell,u,d} = \frac{1}{\sqrt{2}} v Y_{ij}^{\ell,u,d}, \tag{2.30}$$

where the Yukawa coupling constants are labelled as $Y_{ij}^{\ell,u,d}$. The couplings of the Higgs field to the fermions are therefore of the form $m_{ij}^{\ell,u,d}/v$ in the SM.

2.2.2 Properties of the Higgs boson

As highlighted for both fermions and bosons, Higgs couplings are always proportional to the masses of the particles it interacts with. This means that Higgs couplings are strongest to the heaviest member of each fermion family (the top for quarks and the τ for leptons). Third generation fermions are therefore ideal candidates to probe Higgs couplings, in particular in the lepton sector where the masses tend to be very small, making the couplings difficult to measure.

Higgs couplings dictate both the rates at which Higgs bosons are produced through the interactions of bosons and fermions, and the rates at which a Higgs boson decays to bosons and fermions. A measurement of these rates therefore allows for a measurement of the couplings, and in particular a comparison of the couplings to weak vector bosons and fermions can determine whether the behaviour of these couplings as a function of mass (m^2 for bosons, m for fermions) is indeed as expected in the SM. The rate at which each process occurs is characterised by its cross section, a measure of the rate at which the process takes place. Since the cross sections themselves can be calculated from the SM Lagrangian when all parameters (such as the particles' masses) are known, it is also possible to compare the values obtained from experimental cross section measurements to the theoretical predictions to identify any deviations from the SM theory.

2.2.2.1 Higgs decays

Higgs bosons can decay in a variety of ways. Equation 2.26 contains terms involving two weak vector bosons and a Higgs boson, describing processes of the form $H \to WW$ and $H \to ZZ$. Similarly, the Yukawa Lagrangian of Equation 2.27 has terms involving two fermions and a Higgs boson, describing processes of the form $H \to q\bar{q}$ and $H \to \ell\bar{\ell}$. Other processes are also possible. In particular $H \to \gamma\gamma$ can proceed despite photons being massless (and therefore not coupling to the Higgs boson) through interactions between photons and quarks, as illustrated in Figure 2.3. Again, processes involving very massive particles proceed at higher rates than those involving lighter particles since the decay branching ratios are determined by the size of the couplings.

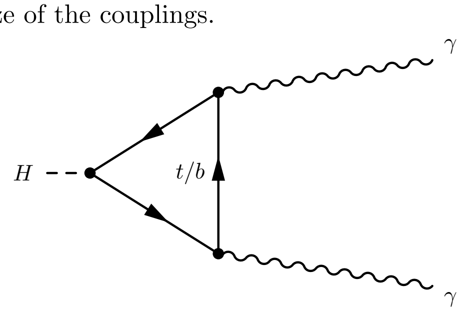

Figure 2.3: Example of a process through which $H \to \gamma\gamma$ proceeds.

2.2.2.2 Higgs production

Higgs bosons can be produced through the same processes. Therefore, Higgs production occurs as $WW \to H$, $ZZ \to H$, $q\bar{q}' \to H$ and $\ell\bar{\ell} \to H$. Production through processes involving massless particles such as $gg \to H$ are also possible through loops similar to the one shown in Figure 2.3.

To create Higgs bosons, it is therefore necessary to build machines where these processes can take place. Due to the small masses of electrons and the technical challenges associated to producing large amounts of stable muon or τ beams, it is in general easier to produce Higgs bosons in processes that involve quarks and gluons, particularly through proton collisions. Since protons are made up of quarks and gluons, both of which carry colour charge, they are subject to the strong force.

2.3 Quantum Chromodynamics

So far this chapter has focused on the electromagnetic and weak theories, since the particulars of the mathematical framework are relevant as they introduce the need for the Higgs mechanism. However, in the case of the strong force (also called Quantum Chromodynamics or QCD), it is the phenomenology of the theory that has the most important implications.

QCD is in general well-behaved. As will be discussed in Section 2.3.1 strong interactions decouple at high energies making QCD calculations and predictions relatively simple. The complications occur at lower energies, the phenomenological effects of which are discussed in Sections 2.3.2 and 2.3.3.2.

2.3.1 Running of couplings

Physics processes can in general be described to the desired accuracy by adding higher-order terms to the calculation in what is called perturbation theory. An example process in QCD would be the interaction of gluons in a 3-point vertex, where the overall interaction rate between three gluons is determined by a sum of the contributions from all possible interactions between those three gluons (including arbitrary amounts of loops) as shown in Figure 2.4. Each term is proportional to the coupling constant of the underlying interaction.

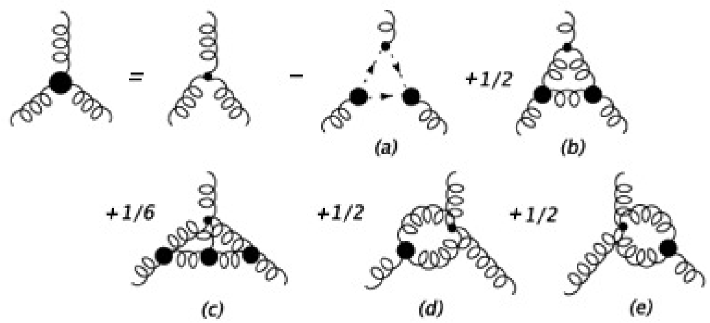

Figure 2.4: Expansion of the 3-gluon vertex, where different types of interactions between three gluons are shown to contribute to the overall interaction rate. Image from [3].

Perturbation theory calculations are in general straight-forward for processes such as the EM interaction, where the coupling constant is $\alpha_{EM} \approx 1/137$. Higher order terms appear as powers of 1/137, which become smaller and smaller as the powers increase. This results in the calculation generally being dominated by the leading-order contribution with each higher order term decreasing in size, allowing the calculation to converge smoothly.

However, interactions between charges and fields, which become apparent at higher orders, complicate this picture. These interactions are the reason why some quantities, such as the bare electron mass, cannot be measured directly. Bare electrons don't exist in the SM. In fact, they are screened by a cloud of virtual particles, and the interaction of the electron with this cloud causes the electron-cloud system to behave as if it had a different mass and charge than originally assumed for the electron alone. Probing the electron at different energies results in a different perceived mass and charge, as this is equivalent to probing the electron at different depths within the screening cloud. A procedure called renormalisation can be used to define quantities such as the electron charge and mass despite the fact that they differ at large and small distance scales. The overall effect of renormalisation is the introduction of scale-dependence into the EM coupling, leading to

$$\alpha_{EM}(Q^2) = \frac{\alpha_{EM}(\mu^2)}{1 - (\alpha_{EM}(\mu^2)/3\pi)\log(Q^2/\mu^2)} \tag{2.31}$$

where μ is called the renormalisation scale. The screening effect for electrons causes a shift in their mass and charge: for increasing momentum transfer Q of interactions that involve electrons, the strength of the EM coupling α_{EM} appears to increase as less screening is available in the smaller distances probed.

Despite this effect, the value of α_{EM} remains small within a wide range of energies and therefore higher order computations can still be performed.

In QCD, the running of the strong coupling constant with energy exhibits a different behaviour. Unlike photons, which are electrically neutral, gluons carry the strong charge and are subject to self-interactions. While only electron-positron loops are available in EM higher-order diagrams, both quark anti-quark and gluon-gluon loops exist in QCD. This feature leads to the introduction of an anti-screening effect which causes the perceived colour charge of particles to increase as the probing distance increases. Mathematically, this can be observed by the change in sign within the running of the strong coupling constant,

$$\alpha_S(Q^2) = \frac{\alpha_S(\mu^2)}{1 + (b_0 \alpha_S(\mu^2)/3\pi)\log(Q^2/\mu^2)} \tag{2.32}$$

where $b_0 = -\frac{n_f}{6\pi} + \frac{33}{12\pi}$ and n_f is the number of quark flavours for quarks with mass lower than the momentum transfer Q. Contrary to α_{EM}, α_S grows with decreasing momentum transfer. Its value can quickly become significant at low energies (large distances). As α_S grows, it becomes increasingly difficult to expand calculations to higher orders using perturbation theory since each additional term adds a significant contribution and the final computation does not converge. This has profound implications for the behaviour of colour charged objects.

2.3.2 Confinement and asymptotic freedom

The anti-screening effect due to the gluon self-coupling results in what is called confinement. When coloured particles are separated, it quickly becomes energetically favourable to produce coloured particle anti-particle pairs from the vacuum to cancel out the colour charge. The resultant coloured-neutral combination of particles are called hadrons. Similarly to the electron bare charge, bare colour charge cannot be measured: only colour-neutral hadrons can. The process via which a single colour charge from a collision evolves into a hadron is called hadronisation. Because hadronisation occurs at energies where the QCD coupling is very large, it is mostly beyond the scope of perturbation theory calculations.

Conversely, particles that are near each other experience a decrease in the strength of the colour force and behave increasingly as free particles. This behaviour is called asymptotic freedom, and processes that occur at these energies can be computed using perturbation theory.

2.3.3 Physics at proton colliders

These particularities of QCD have important implications for the physics of proton colliders since protons contain particles that interact via the strong force. There are many components to these collisions, as illustrated in Figure 2.5 and outlined below.

2.3.3.1 Hard scatter and parton distribution functions

Protons are made of three tightly-bound valence quarks in a colour-neutral arrangement surrounded by a sea of virtual QCD particles. The composition of this sea changes with the energy of the interaction since the QCD coupling varies accordingly. The collision of two protons therefore refers not only to a collision of two point-particles, but rather the collision of two complex systems. The resulting interaction can be split into a hard perturbative component and a soft non-perturbative one. The scale at which the separation between hard and soft is made is arbitrary, but usually taken to be Q, the energy of

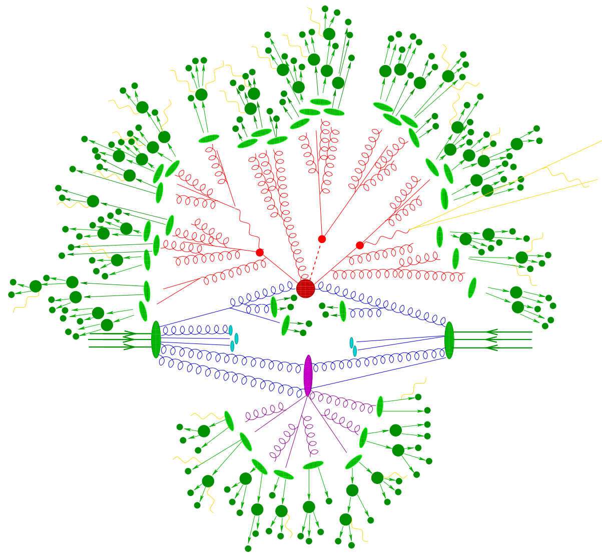

Figure 2.5: Structure of proton-proton collisions. The hard scatter interaction (red circle) originates from the quarks and gluons (blue lines) within the two incoming protons (green lines, corresponding to the three valence quarks). All end products with bare colour charges (red lines) undergo hadronisation (light and dark green circles), while those without colour charges (yellow lines) are unaffected. Image from [94].

the interaction. The first component represents the main (hard-scatter) interaction between two high-energy partons (quarks and gluons) and can be computed perturbatively using the SM Lagrangian. An example of a hard-scatter interaction is the Drell-Yan process, where a quark from one of the protons and an antiquark from the other proton annihilate into a photon or a Z boson, which then decays into a pair of oppositely-charged leptons. The second component requires a description of the structure of the proton. Parton distributions functions (PDFs) are functions that describe the probability of a certain momentum fraction to be carried by the valence and sea quarks and the gluons that make up the proton. They are not readily computable in the mathematical framework of QCD but are obtained through fits of proton-proton collision data as well as, for example, the HERA deep inelastic scattering experiments [1]. Since these PDFs depend on the energy of the interaction, they must be adapted in order to be used by experiments with different conditions. The Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) equations described in [4, 96] can be used to translate the PDFs between different energy regimes.

2.3.3.2 Parton showers

The partons involved in the proton-proton interaction can split and lose energy through radiation of gluons and quark anti-quark pairs. The resulting showers of partons need to be accounted for when describing the overall interaction. When these showers are associated to the incoming particles they are considered as initial state radiation (ISR). Showers linked to the outgoing particles are called final state radiation (FSR). The energy of these showers is generally contained around the original parton, and is described by fragmentation models that account for the non-perturbative and therefore non-computable effect of hadronisation.

Chapter 3

The Large Hadron Collider

Beams of high energy particles have proven extremely useful in the study of particle physics over the years. They have allowed the discovery of particles (to give just one example, the top quark was discovered at the Tevatron in Fermilab in 1995) and enabled a deeper understanding of their interactions. At the time of writing, the Large Hadron Collider (LHC) at CERN is the world's largest and most powerful particle accelerator and collider. Many experiments take advantage of the collisions created with the LHC for a variety of purposes, ranging from the study of quark-gluon plasma to searches for magnetic monopoles. Along with CMS [77], the ATLAS [32] experiment is one of the two general-purpose experiments at the LHC. Among the highlights of its experimental programme are the search for the Higgs boson and later the measurement of its properties, as well as precision measurements of SM parameters and searches for phenomena beyond the Standard Model (BSM). As of 2021, two experimental runs have taken place. The LHC was in operation from 2009 to 2013 in what was called Run 1. Run 2 started in 2015 and continued until 2018.

The LHC is a circular hadron-hadron collider built from 1998 to 2008 in the tunnel that used to house the Large Electron-Positron Collider (LEP) in Geneva, Switzerland. With a circumference of 27 km and a centre of mass energy of 13 TeV[1], it is approximately 4.5 times larger and 7 times more powerful than the Tevatron. It is used for both proton and heavier ion collisions. This thesis focuses on studies of proton-proton collisions.

3.1 The CERN accelerator complex

The LHC is the last element in a chain of accelerator systems (see Figure 3.1) that aim to increase the energy of hadrons before colliding them. Protons are first accelerated to 50 MeV in a linear accelerator and are then injected into the Proton Synchrotron Booster (PSB). The PSB is the first and smallest circular accelerator of the chain, with 157 m of circumference, and accelerates the protons to an energy of 1.4 GeV. Next, the protons are injected into the Proton Synchrotron (628 m circumference), accelerated to 25 GeV and arranged into bunches of approximately 10^{11} protons with a spacing of 25 ns between bunches, and passed on to the Super Proton Synchrotron (7 km circumference) until they reach an energy of 450 GeV. At this point, 2808 bunches are injected into the LHC, where the protons are accelerated to energies of 6.5 TeV.

[1] During Run 1 the centre of mass energy only reached 8 TeV; the design energy is 14 TeV.

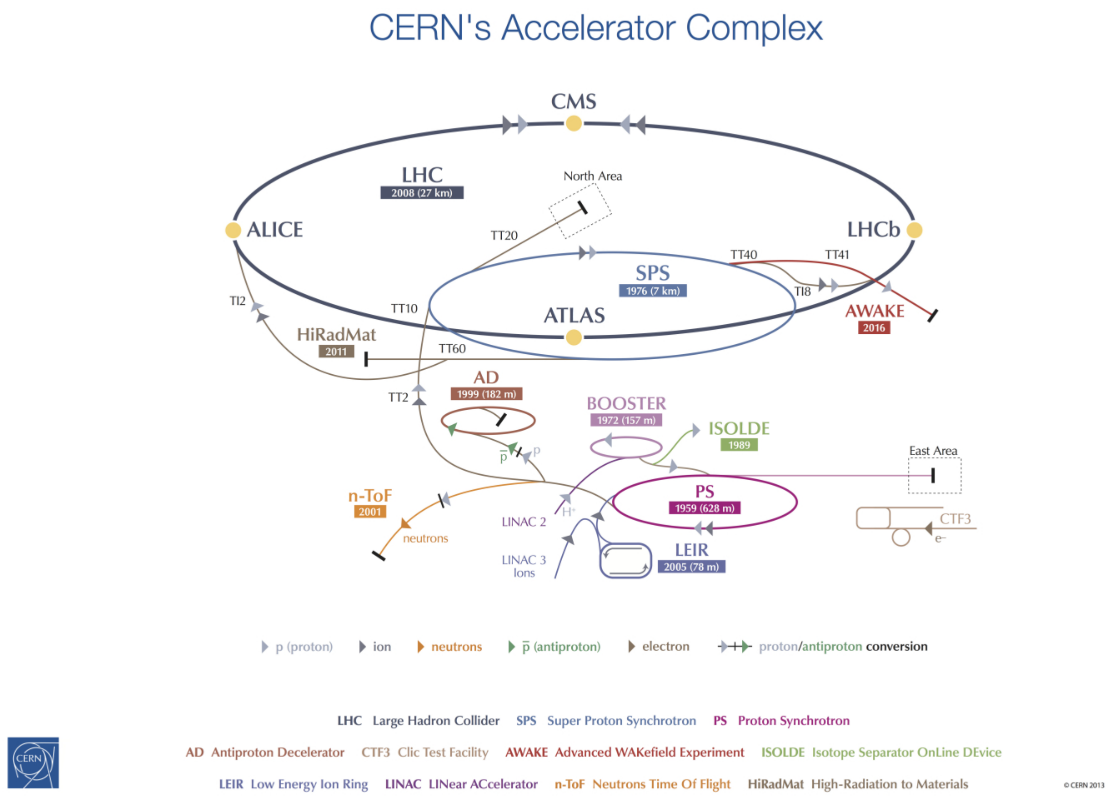

Figure 3.1: CERN accelerator complex. Image from [101].

3.2 LHC Collisions

Within the LHC, there are two beams that circulate in opposite directions. They are set on a collision course through the use of quadrupole magnets at four interaction points, where the detectors of the various LHC experiments are placed to record the products of the collisions as shown in Figure 3.2. These collisions bring the contents of the protons close together and allow interactions at high energies, which can result in the production of massive particles such as the Higgs boson. As discussed in the previous chapter, the rate at which these processes occur can be compared to theoretical predictions to determine whether the model proposed properly describes the experimental results. High energies and high collision rates allow sensitivity to processes involving particles with large masses as well as to rare processes.

3.2.1 The LHC environment

The LHC aims to deliver collisions at unprecedented energies and rates. The rate of collisions depends on the size and revolution frequency of the proton bunches in each of the colliding beams. This rate is often quantified using the measure of (instantaneous) luminosity. As a general rule, the higher the luminosity one can achieve, the more data one can accumulate to study processes of interest.

Mathematically, a luminosity relates the number of events observed in a certain amount of time to

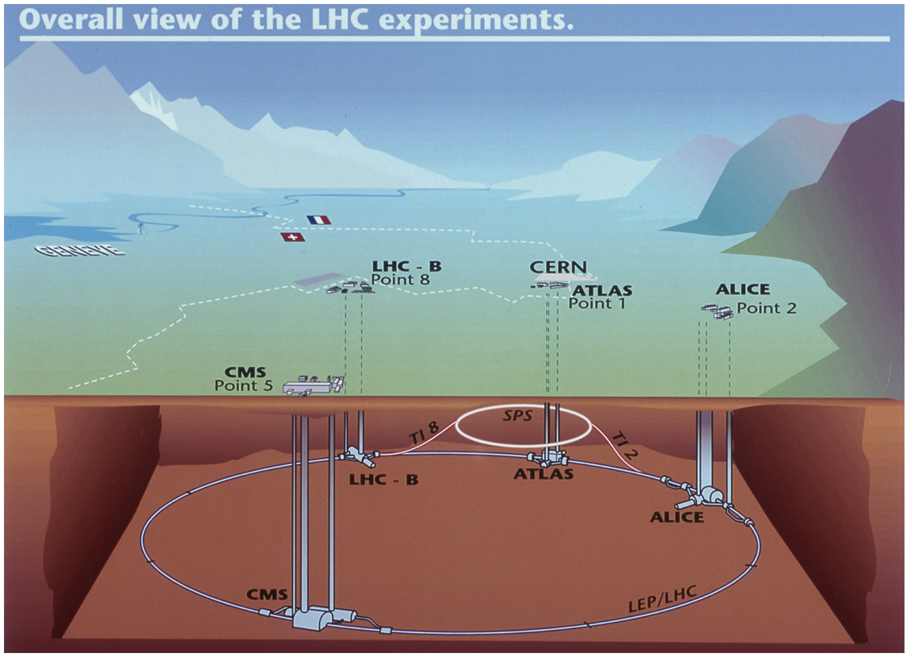

Figure 3.2: CERN LHC experiments and their location around the LHC ring. Image from [110].

the rate at which a process of interest occurs. The instantaneous luminosity L is defined as

$$L = \frac{N_b^2 n_b f_{\text{rev}}}{\mathcal{A}} F \tag{3.1}$$

where N_b is the number of protons in each bunch and n_b is the number of bunches in each beam, f_{rev} is the revolution frequency for the beams around the LHC ring, and $\mathcal{A}$ is the cross-sectional area of overlap between the beams, with F denoting the reduction in instantaneous luminosity that is due to the fact that the two beams cross at an angle.

The total number (N) of events generated in collisions in a given time interval is given by

$$N = \mathcal{L}\sigma \tag{3.2}$$

with the integrated luminosity is defined as $\mathcal{L} = \int L dt$ and where σ refers to the cross section of the process of interest. Typical units for integrated luminosities are inverse barns, with 1 barn = 10^{-28} m^2. The total integrated luminosity delivered by the LHC for the years 2015 to 2018 is shown in Figure 3.3.

3.2.1.1 Pile-up interactions

While high instantaneous luminosities lead to higher rates of interesting physics events to study, they also tend to result in an increasing number of proton-proton collisions per bunch crossing. Indeed, the more protons each bunch contains and the more area of overlap there is between beams, the more likely it is that more than one proton-proton collision will occur when two bunches collide. The main collision of interest is called the hard scatter and its location is labelled the primary vertex. The number of additional inelastic collisions, called pile-up events, can be split into two categories. In-time pile-up refers to the number of inelastic collisions that happen within a single bunch crossing. Identifying this

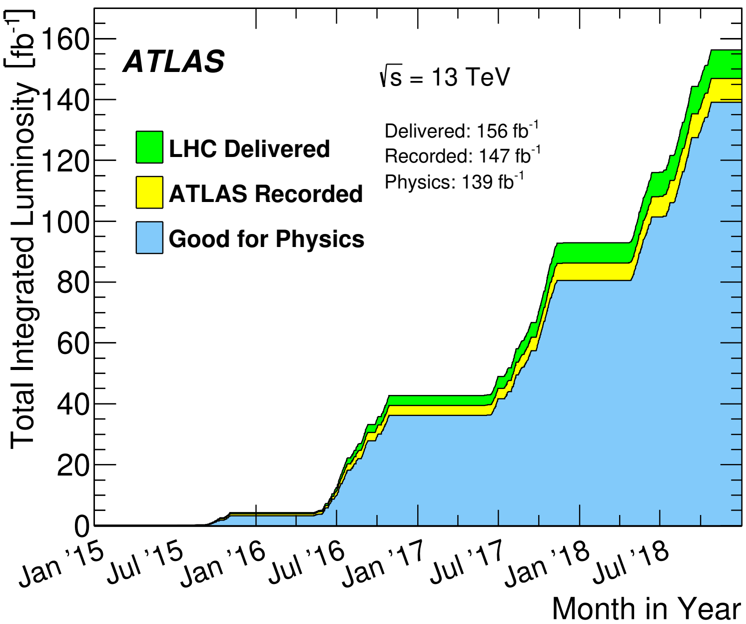

Figure 3.3: Total integrated luminosity recorded by the ATLAS detector during the years 2015 to 2018. Image from [64].

type of pile-up interaction requires the detector to be able to pinpoint the location of each collision with a high degree of accuracy and precision. Out-of-time pile-up refers instead to inelastic collisions that occur during a different bunch crossing. Due to the fact that signals within the detector take time to develop and be read out, it is possible that energy from a previous or future bunch crossing affects the signals associated to the bunch crossing of interest. This type of pile-up particularly affects parts of the detector where signals are read out at a rate slower than the bunch crossing frequency of 40 MHz (i.e. 25 ns).

Pile-up can be quantified with two measures. The number of reconstructed primary vertices (N_{PV}) is determined using the ATLAS Inner Detector (ID). Given that the ID is read out quickly, N_{PV} is mostly sensitive to in-time pile-up. The other quantity used is the average number of interactions per bunch crossing and is given by a Poisson distribution with mean μ where

$$\mu = \frac{L\sigma_{\mathrm{inel}}}{n_b f_{\mathrm{rev}}} = \frac{N_b^2}{\mathcal{A}} F \sigma_{\mathrm{inel}} \tag{3.3}$$

where L is the instantaneous luminosity from Equation 3.1 and σ_{inel} is the total inelastic cross section (its value is of approximately 80 mb [32]). The values of μ recorded by the ATLAS detector during Run 2 are shown in Figure 3.4.

The ATLAS detector was designed to handle the large pile-up rates provided by the LHC. However, upgrades to the accelerator complex over the next decade will result in even higher levels of pile-up. Some sub-systems will therefore have to be updated[2]. The impact of these changes will be discussed in detail in the context of τ-lepton reconstruction, identification and calibration in Chapter 10.

[2]The ID will be replaced by the Inner Tracker (ITk), part of which is being constructed at the University of Toronto.

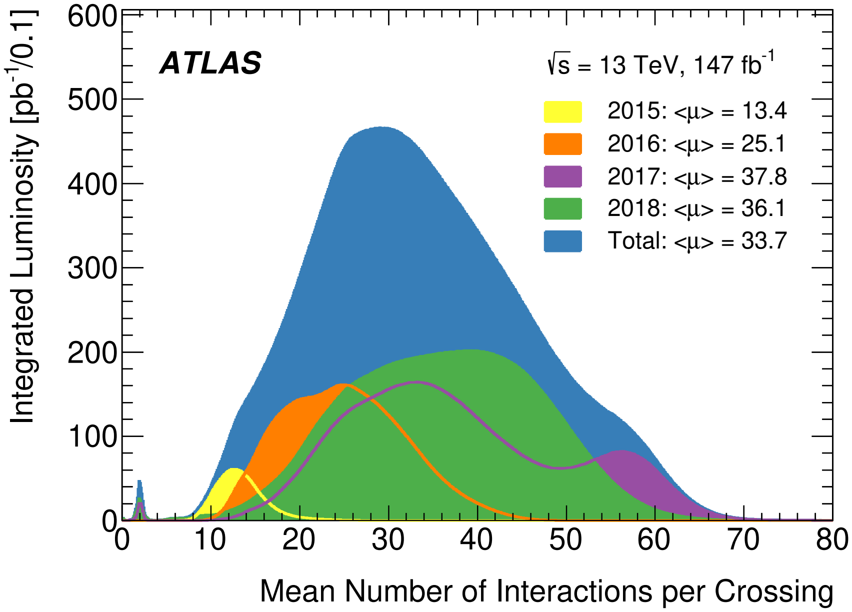

Figure 3.4: Average (mean) number of collisions per bunch crossing, as measured by the ATLAS detector. Image from [64].

3.2.1.2 Underlying event

Aside from contributions from multiple proton-proton interactions, other mechanisms can result in additional particles that interfere with the ability of the ATLAS detector to identify the products of the hard scatter of any given collision. These contributions are jointly called the underlying event and can occur through interactions of or with additional partons within the protons involved in the hard scatter.

Chapter 4

The ATLAS detector

The ATLAS detector was designed to record the collisions produced by the LHC. Its benchmark performance requirements were set in large part in preparation for the search of the Higgs boson. Given that the mass of this boson was unknown at the time, the ATLAS detector was required to be sensitive to Higgs decay products over a large mass range, and to allow for a measurement of its properties. Aside from Higgs boson physics, searches for heavy bosons such as W' and Z', and supersymmetric particles informed the design choices.

The physics programme required momentum measurements and charge identification for leptons with energies of the order of a TeV, the ability to reconstruct secondary vertices, performant calorimetry for the identification of electrons and photons, measurements of hadrons with high transverse momentum and the ability to detect hadrons produced at small angles with respect to the beamline, and finally, the ability to identify events with large imbalances of transverse energy due to the presence of undetectable particles [32]. The LHC conditions described in Section 3.2.1 placed additional requirements on the ATLAS detector, in particular the need for a high granularity as well as fast, radiation-hard electronics that could sustain the large particle flux through the detector. The final design for the ATLAS detector takes into account all of these requirements to ensure good performance in accordance for the various physics needs of the collaboration. Sections 4.2 to 4.4 describe the design of each sub-system, and outline how it addresses the requirements described here.

4.1 Overview

The ATLAS detector is a cylindrical structure that surrounds the beam pipe at one of the interaction points of the LHC. It is made up of various layers and components, each designed to measure the properties of specific particles in specific environments. ATLAS uses a right-handed coordinate system with its origin at the nominal interaction point (IP) in the centre of the detector and the z-axis along the beam pipe. The x-axis points from the IP to the centre of the LHC ring, and the y-axis points upwards. Cylindrical coordinates (r, ϕ) are used in the transverse plane, ϕ being the azimuthal angle around the z-axis. The pseudorapidity is defined in terms of the polar angle θ as $\eta = -\ln\tan(\theta/2)$. Angular distance is measured in units of $\Delta R \equiv \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$. A detailed description of the ATLAS detector can be found in [32]. A summary is presented in this chapter.

ATLAS and all of its sub-systems can be seen in Figure 4.1. The inner detector interacts with charged particles and uses the magnetic field generated by a 2 T solenoid to measure their momentum. It provides coverage up to $|\eta| = 2.5$. The next layer consists of calorimeters: electromagnetic calorimeters aim to contain showers from particles that interact through the electromagnetic force while hadronic calorimeters aim to contain showers from particles that interact through the strong force. Unlike the inner detector, the calorimeters provide coverage up to $|\eta| = 4.9$. Finally, the Muon Spectrometer surrounds the entire structure and provides measurements up to $|\eta| = 2.7$. At the LHC, muons are usually not stopped within other sub-systems and therefore require an additional layer specifically designed to measure their momentum. Its function is similar to that of the inner detector, although it relies on an air-core toroid system instead of a solenoid to provide the required magnetic field within a light and open structure. All sub-systems are segmented into barrel and end-cap components as the particle flux is larger in the forward region (end-cap, high $|\eta|$), often requiring different materials to be used.

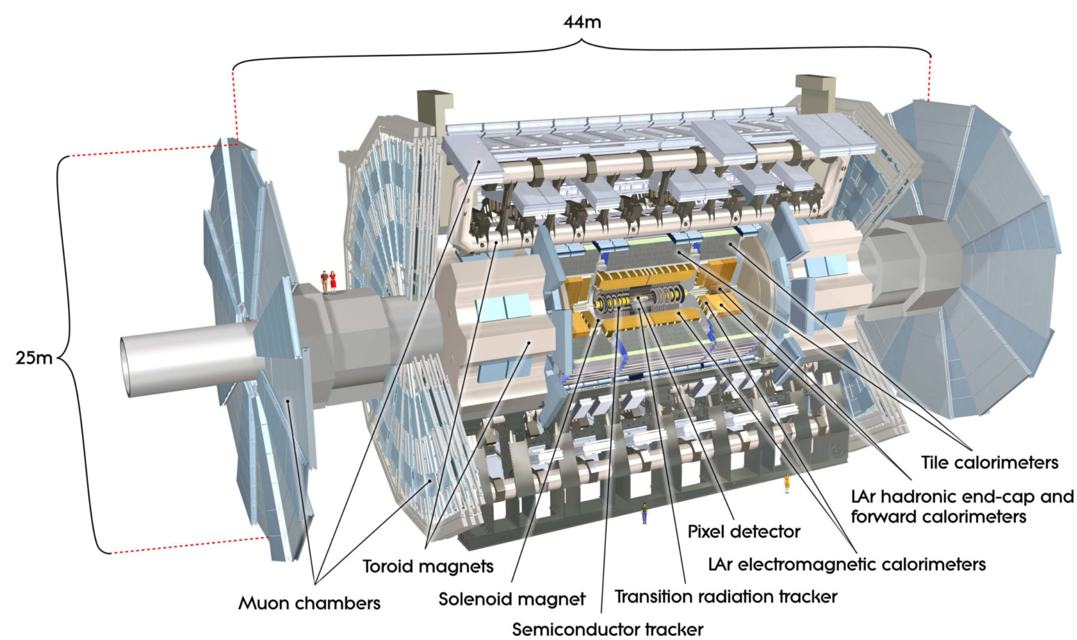

Figure 4.1: The ATLAS detector and its sub-systems: the inner detector (pixel detector, semiconductor tracker and transition radiation tracker), the electromagnetic and hadronic calorimeters and the muon chambers. Image from [32].

As discussed, the LHC was designed to provide high instantaneous luminosities (of the order of 10^{34} cm^{-2}s^{-1}). Given typical proton-proton interaction cross sections, these luminosities correspond to data flow rates of approximately 1 GHz. It is in practice not possible (or useful) to record all of these events. Current limitations set the data recording rate to approximately 1 kHz, meaning that a trigger and data acquisition system is needed to enable collision data to be filtered based on the presence of interesting physics processes. ATLAS does this through a two-level system consisting of a Level-1 Trigger (L1) and a High-Level Trigger (HLT) [74].

4.2 Inner Detector

The ATLAS Inner Detector (ID) was designed to measure the momentum of the electrically-charged particles that originate from collisions and locate their origin through the deflection of their trajectory due to the magnetic field provided by the solenoids that surround the ID. It is also used to identify the location of the various proton-proton collisions produced during each bunch crossing, in particular to enable discrimination against particles generated in pile-up events. Finally, the ID is also used to help identify specific particles. Such identification includes particles originating from displaced vertices due to decays within the detector (e.g. the decay products of b quarks), particles that can be identified through transition radiation and those that can be identified through their specific energy loss $\mathrm{d}E/\mathrm{d}x$.

The ID is sub-divided into three layers, each with a specific purpose (see Figure 4.2). The pixel detector, the layer closest to the interaction point, is designed to provide high enough granularity to locate the origin of each charged particle that crosses it. The pixel detector and the next layer, the semiconductor tracker (SCT), are used for precision momentum and position measurements. The transition radiation tracker (TRT) makes up the last layer. It plays a key role in particle identification and provides a large number of low-precision measurements at a large distance from the interaction point.

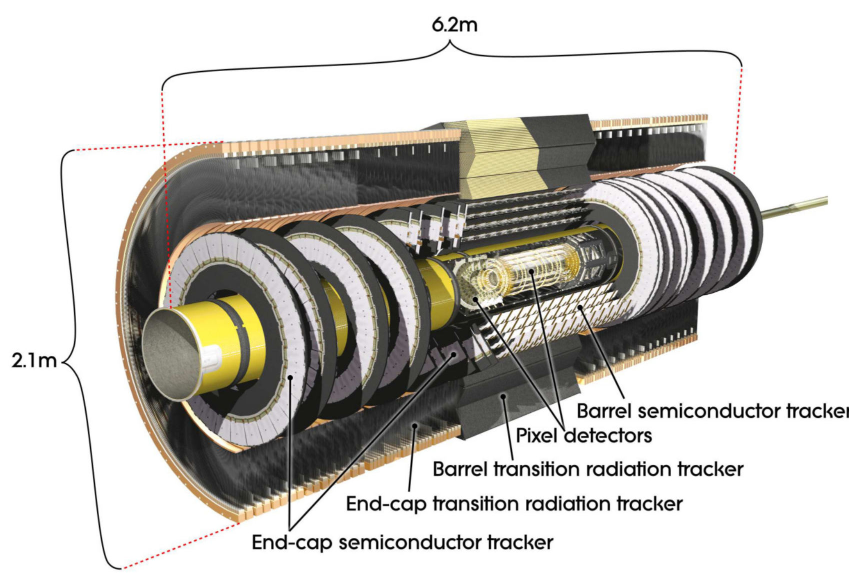

Figure 4.2: Components of the ID. Image from [32].

4.2.1 Pixel detector

The pixel detector contains three layers of silicon modules, both in the barrel and in the end-caps, with all three layers typically crossed by the charged particles that go through the detector. The layers are oriented parallel to the beam pipe in the barrel and radially in the end-caps. An additional layer called the insertable b-layer (IBL) was added to the barrel before the beginning of Run 2. Charged particles passing through the pixel detector ionize the silicon sensors and produce an electrical signal that can then be read out. The size of each sensor is therefore directly related to the resolution that can be achieved. The pixel sensors are segmented in $r-\phi$ and z with a size of 50×400 $\mu\mathrm{m}^2$ in $(r-\phi) \times z$

and are identical in the barrel and end-cap regions. This corresponds to a resolution of 10 μm in $r - \phi$ and 115 μm in z (r) for the barrel (end-cap). The pixel detector has approximately 80.4 million readout channels [32].

4.2.2 Semiconductor tracker

The SCT consists of 4 (9) layers of microstrip silicon wafers in the barrel (end-cap), resulting in at least 4 layers being crossed by each incoming particle. Each layer actually contains two sub-layers of strips that are arranged at a stereo angle of 40 mrad to provide 2-dimensional measurements. One of the strips is always oriented parallel to the beam pipe in the barrel, and radially in the end-cap. Charged particles interact with the SCT in the same way as with the pixel detector. The mean strip separation (pitch) for all strips is approximately 80 μm, resulting in a resolution of 17 μm in $r - \phi$ and 580 μm in z (r) for the barrel (end-cap). The total number of readout channels in the SCT is much lower than in the pixel detector, at approximately 6.3 million [32].

4.2.3 Transition radiation tracker

The TRT is built from a set of straw tubes consisting of a film of kapton surrounded by an aluminum and a graphite-polyimide layer, adding up to 73 tubes in the barrel and 160 in the end-cap. Incoming particles typically interact with an average of 36 tubes. In the barrel, 144 cm long straw tubes are arranged parallel to the beam pipe while in the end-cap, 37 cm long tubes are arranged radially in wheels [32]. The tubes contain a mixture of xenon (or argon), carbon dioxide and oxygen. When charged particles go through the tube, they ionize the gas and produce electrons. A negative voltage applied to the outer wall of the tube directs them towards the centre of the tube where a gold-tungsten wire is located. Once they reach the wire, an avalanche is generated and the signal can be read out. The time it takes for the avalanche to occur gives an indication of the origin of the electrons within the tube, with electrons near the outer wall taking 48 ns to reach the central wire. This information is used to estimate the location of the particle that produced the electrons upon interaction with the TRT and provides a resolution of 130 μm in $r - \phi$. The TRT has approximately 351 000 readout channels [32].

Originally, the TRT used only xenon. Unfortunately, a gas leak in the barrel was detected at the beginning of Run 2, and cost considerations motivated its replacement with argon. While argon allows for a similar performance in terms of tracking, xenon provided the TRT with electron identification power and gave the TRT the “transition radiation” part of its name. Charged particles with a large enough Lorenz boost factor emit transition radiation in the form of X-rays as they travel through the dielectric material that is used to interleave the straw tubes. The xenon gas absorbs these X-rays, and then produces a signal approximately 20 times larger than the one described above. The presence of transition radiation can therefore be inferred from the presence of such high signals [37]. This process typically only occurs for electrons (due to their smaller mass), and the presence of signals associated to X-rays can therefore be used to identify electrons in the TRT. At energies of about 150 GeV, other particles such as charged pions also begin to emit X-rays, introducing an upper limit on the electron identification power of the TRT. However, since argon does not have xenon’s ability to absorb X-rays some of the electron identification power was lost after the switch.

4.3 Calorimeters

After particles pass through the ID, they reach the calorimeters. The ATLAS calorimeters cover a wider η range than the ID and exploit different technologies in different regions to address the varying requirements of the radiation environment. A split into electromagnetic and hadronic calorimetry is also introduced. The two types of calorimeter are designed to contain either electromagnetic or hadronic showers, hence their name, but they are not meant to only measure those showers. This is because hadronic showers can begin in the electromagnetic calorimeter and also because, as will be discussed in the next section, a significant part of all hadronic showers is electromagnetic in nature. The ATLAS calorimeters are shown in Figure 4.3.

The interaction of particles with calorimeters involves some subtleties that should be discussed before a description of ATLAS's calorimetry can be carried out.

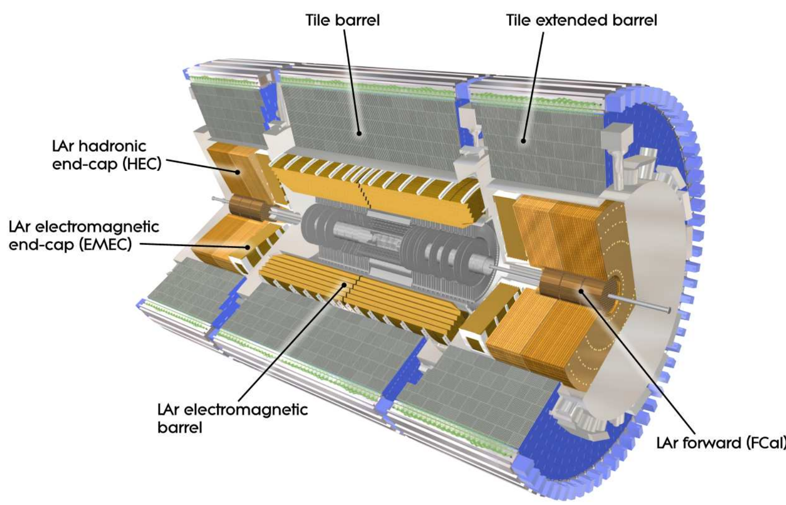

Figure 4.3: Electromagnetic and hadronic calorimeters in ATLAS. Image from [32].

4.3.1 Showers

When particles interact with the calorimeters, other particles are created in what is called a shower. Showers that occur due to electromagnetic interactions have different properties than hadronic showers (generated by strong interactions) and therefore different calorimeter technologies are often used to contain them.

Electromagnetic showers are generated by electrons, positrons and photons. As these particles pass through the calorimeters, a variety of processes take place. At high energies, electrons and positrons radiate photons via Brehmsstrahlung (at multi-GeV energies, an electron can radiate thousands of soft photons) while at lower energies they ionize the calorimeter medium. Photons can lose their energy through a variety of processes, ranging from pair production at high energies, through Compton scattering, to the production of electrons via the photoelectric effect at low energies [123]. It is important to note that this process is iterative. A high energy photon, for example, can first go through a pair production process that results in both an electron and a positron. The electron and positron would both have high energies (although lower than that of the original photon) and therefore radiate more photons. These photons would then interact with the calorimeter, and so on, until electrons with low enough energy are produced via the photoelectric effect and can be absorbed by ions in the medium [123]. The collection of photons, electrons and positrons generated in this process is what is referred to as the electromagnetic shower.

Hadronic showers are significantly more complicated. Approximately one third of the energy of hadronic interactions is associated to the production of neutral pions which almost always decay to two photons, thereby starting an electromagnetic shower. The other two thirds of the energy involve the production of:

- charged mesons (20%),
- nuclear fragments and protons (25%),
- soft neutrons and photons (15%),
- or result in energy loss through undetectable processes (40%) such as the fragmentation of nuclei or nuclear recoil.

The first two processes involve charged particles and often result in ionisation and/or spallation, where the incoming charged hadron collides with nucleons from the calorimeter material and generates intra- and internuclear cascades. In contrast, the last two result in invisible energy. For example, soft neutrons do not ionize the calorimeter medium and instead lose energy only through nuclear reactions at low rates [123]. The timescales for these processes are too slow to be used for the shower's energy measurement. This means that the fraction of the non-electromagnetic component that can be detected is lower than 1. Unless the calorimeter is built to compensate for this effect, its response to a pion and an electron of the same energy will be different. The ATLAS calorimeters in particular do not compensate for this effect.

Unlike in electromagnetic showers where only photons, electrons and positrons are present, there are many types of particles in hadronic showers. The shower's composition depends on the properties of the incoming particle. In particular, the electromagnetic fraction of hadronic showers depends on its energy. As mentioned, one third of each hadronic interaction's energy result in the production of neutral pions [123]. So, at the first interaction the electromagnetic fraction will be $\frac{1}{3}$. The remaining $\frac{2}{3}$ will result in charged hadrons that can then interact again, with $\frac{1}{3}$ of the energy of this second generation of interactions resulting in neutral pions. After this, the electromagnetic fraction is therefore $\frac{1}{3} + (\frac{1}{3} \times \frac{2}{3})$, which is of course larger than the original $\frac{1}{3}$. This process continues until the hadrons no longer have enough energy for pion production [123]. From this, it follows that the more energy the original hadron had, the more generations of interaction are possible, and the larger the electromagnetic fraction is.

These characteristics have a few consequences for hadronic showers. The main is that the hadronic response is non-linear as a function of the energy of the incoming particle, given that the response to the electromagnetic and hadronic portions of the shower is different and that the electromagnetic fraction varies with energy [123]. All of this significantly complicates the measurement of the energy of hadrons and will be of particular importance for the calibration of hadronically-decaying τ leptons, whose decay products are composed of both neutral and charged pions and therefore associated to electromagnetic and hadronic components that must be inter-calibrated with great care. This is discussed further in Chapter 6.

4.3.2 Definitions

4.3.2.1 Sampling calorimeters

In order to contain electromagnetic and hadronic showers, ATLAS employs sampling calorimeters. This means that its calorimeters contain two types of alternating material: active and passive. The active material is used to measure the energy of the shower, while the denser passive part is responsible for generating interactions between incoming particles and the material of the calorimeter. In general the denser this material, the more a particle interacts and the faster it loses its energy, therefore creating smaller showers that require less material in order to be contained. This is the main advantage of sampling calorimeters: they can remain small and so less expensive to build. The main disadvantage, on the other hand, is that some of the energy is deposited into the passive layers and therefore cannot be recorded in the active layers. This has two consequences. The first is that the energy measurement will be biased. The sampling fraction $f_{\text{samp}} = \frac{E_{\text{active}}}{E_{\text{active}} + E_{\text{passive}}}$ provides a measure of the amount of energy lost in the passive layers and can be used to correct for this effect [123]. The second is that the resolution of the energy measurement will suffer, since part of the signal is lost. This effect cannot be corrected for.

4.3.2.2 Radiation and interaction lengths

Shower size is determined by the properties of the material of the calorimeter. For EM interactions the scale is determined by the radiation length (X_0), usually of the order of a few cm. Each radiation length of material is, on average, responsible for one generation of interactions. Electromagnetic showers up to 1 TeV are typically contained within 25-30 X_0. Hadronic shower size is instead governed by the interaction length (λ_{int}), usually one order of magnitude larger than X_0 for the same material[1]. Hadronic showers at LHC energies are generally contained within approximately 10 λ_{int}.

4.3.3 Electromagnetic calorimeters

Similarly to trackers, the electromagnetic calorimeters are divided into barrel and end-cap sections. Lead has a short radiation length (i.e. particles tend to travel short distances before they interact) and was therefore chosen for the passive material. The active material is liquid argon (LAr) and is homogeneous, stable and radiation hard. The particles that come into the electromagnetic calorimeter ionize the LAr, and the resulting electrons are collected in copper electrodes.

The electromagnetic calorimeter is split into three layers. The first one is highly segmented, in particular to allow for the identification of $\pi^0 \rightarrow \gamma\gamma$. The second layer is used to absorb most of the energy of the incoming particle. Finally, the third layer has the coarsest segmentation and is used to ensure as much as possible of the electromagnetic shower's energy is contained within the electromagnetic calorimeter so that significant leaks into the hadronic calorimeter are minimized. It is also useful for the identification of electrons. In the barrel, the electromagnetic calorimeter covers the region up to $|\eta| = 1.475$ and the end-cap is within $1.375 < |\eta| < 3.2$. The region between $1.375 < |\eta| < 1.52$ is known

[1] This is a consequence of the difference between the length scales of electromagnetic and strong interactions.

as the crack region and contains inactive material, such as cabling, which reduces the performance of the calorimeter. In analyses where high precision measurements are needed, signals from this section are generally not used. In the region $|\eta| < 1.7$ a pre-sampler is included to estimate the energy loss due to the material upstream of the calorimeters [32].

The electromagnetic calorimeters provide over 22 (24) X_0 in the barrel (end-cap) and amount to approximately 1.3 λ_{int} [32].

4.3.4 Hadronic calorimeters

The hadronic calorimeters consist of a steel and scintillating tile calorimeter in the barrel, and LAr calorimeters similar to those used in the EM calorimeters in the end-cap and forward regions.

The barrel is split into barrel ($|\eta| < 0.9$) and extended barrel ($0.8 < |\eta| < 1.7$) sections. Both utilize steel as the passive material and plastic scintillating tiles as the active material. When particles go through the hadronic calorimeter, they interact with the steel plates and the resulting charged particles traverse the scintillating tiles. Photons are produced within the tiles, collected with wavelength-shifting fibers and read out with photomultiplier tubes [32]. The barrel hadronic calorimeter is segmented into three layers, all of which have coarser granularity than their electromagnetic counterparts.

The end-cap covers the range $1.5 < |\eta| < 3.2$. It uses LAr as the active medium and copper as the passive medium. It is split into four layers, with coarser granularity on the outer than on the inner ones [32].

Finally, a forward calorimeter covers the range $3.1 < |\eta| < 4.9$. It contains three layers, with the first acting as an electromagnetic calorimeter and the other two as hadronic calorimeters. All three layers use LAr as the active medium, with copper and tungsten as passive material in the electromagnetic and hadronic layers, respectively [32]. Copper is well-suited to the high particle flux (which deposits a lot of energy in the calorimeters) as it is conductive. Since the hadronic layers are placed after the electromagnetic ones, conductivity is less important as the particle flux is reduced. Tungsten was chosen as it has a very short interaction length and therefore provides better hadronic shower containment. Most particles escaping the tungsten-LAr calorimeter are stopped by a large uninstrumented brass plug to provide shielding for the Muon Spectrometer.

The number of interaction lengths of the hadronic calorimeter varies significantly throughout the detector as can be seen in Figure 4.4, but at least 9 λ_{int} are provided everywhere.

4.4 Muon Spectrometer

Most particles that pass through the ATLAS detector are stopped by the calorimeters. Particles that make it through are either non-interacting such as neutrinos, particles so energetic that they cannot be fully contained within the calorimeters, referred to as "punch-through", or muons. No attempt is made to contain non-interacting particles, while the design of the calorimeters aims to minimize punch-through. The momenta of the last set of objects, muons, are measured in the ID and the MS with the latter, illustrated in Figure 4.5, aiding in the measurement of high-momentum muons and providing a means of identifying them from other charged objects.

Like all other sub-systems, the Muon Spectrometer (MS) is split into barrel and end-cap sections. It is also divided into fast but coarse and slow but high-granularity systems. The first are used for triggering while the latter are used for high-precision measurements. The entire structure is immersed in

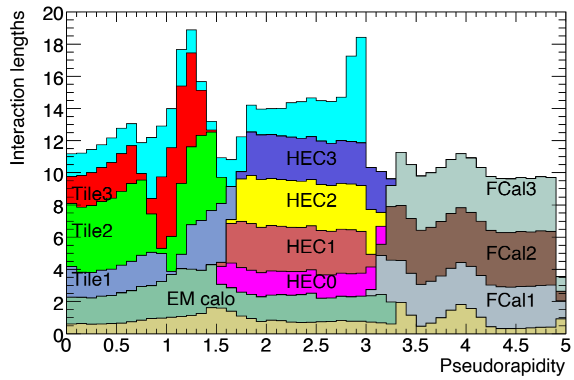

Figure 4.4: Amount of material (in λ_{int}) as a function of $|\eta|$ in the various sub-systems of the ATLAS detector. The hadronic calorimeters are labelled as Tile 1-3 for the barrel layers, HEC 0-3 for the hadronic end-caps and FCal 1-3 for the forward calorimeters. Image from [32].

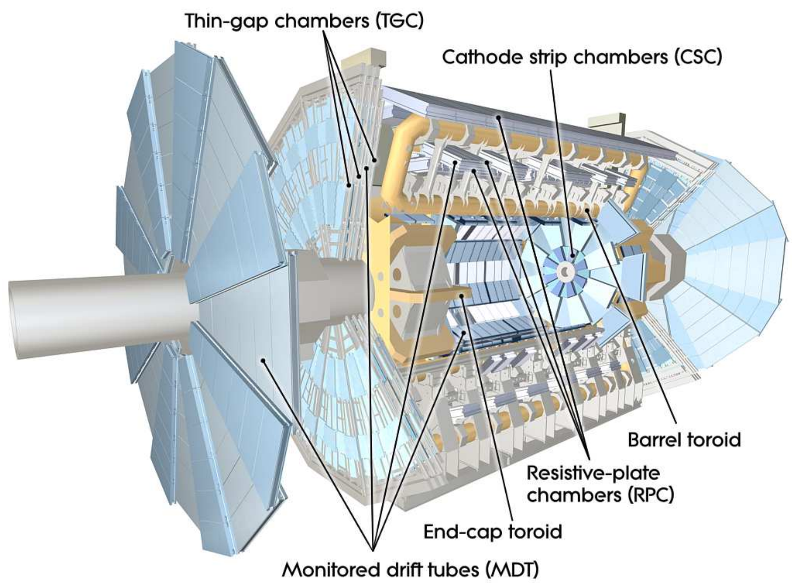

Figure 4.5: ATLAS Muon Spectrometer. Image from [32].

a magnetic field provided by toroidal magnets with a bending power of 1 (0.5) T in the central (forward) regions [33]. This field bends the trajectories of muons, allowing for momentum measurements to be carried out as in the ID. The structure is kept light and open to reduce multiple-scattering effects and to reduce the overall weight of the ATLAS detector. A gap in the coverage around $|\eta| = 0$ provides services to the solenoid magnet, the calorimeters and the ID [32].

All muon sections of the MS work in similar ways: incoming muons ionize the medium and generate electron-ion pairs. These drift in an electric field and the signal then collected and amplified. The MS is the largest ATLAS sub-system, reaching a radius of 25 meters and extending to 44 meters longitudinally. In comparison, the ID reaches a radius of 1.15 meters and extends to 3.5 meters longitudinally.

4.4.1 Triggering

The muon trigger chambers consist of Resistive Plate Chambers (RPCs) in the barrel and Thin Gap Chambers (TGCs) in the end-cap. RPCs consist of cathode and anode wires separated by glass layers, while TGCs consist of anode wires surrounded by graphite cathodes, themselves sandwiched between pickup strips. RPCs are used within $|\eta| < 1.05$ while TGCs can be used in the region covering $1.05 < |\eta| < 2.4$ for triggering (although they can provide measurements up to $|\eta| < 2.7$) [32]. The resolution achieved is of 10 mm for RPCs in z and ϕ, and of 2-6 mm and 3-7 mm for TGCs in r and ϕ respectively, with signals generated within 15-25 ns [32].

4.4.2 High-precision measurements

For high-precision measurements Monitored Drift Tubes (MDTs) and Cathode Strip Chambers (CSCs) are used. MDTs are similar to TRT straws and consist of an aluminum cathode tube filled with an argon and carbon dioxide mixture, and containing a tungsten-rhenium anode wire. MDTs are available within $|\eta| < 2.7$ and provide a resolution of 80 μm per tube and 35 μm per chamber but have a collection time of around 700 ns [32]. Their measurements are therefore 57 times more precise but 28 times slower than that of the trigger chambers. They only provide r and z measurements and so benefit from the ϕ measurements provided by the trigger chambers. Finally, CSCs are only available within $2.0 < |\eta| < 2.7$ and are used to complement the MDTs in the forward region where the increased particle flux requires chambers with higher granularity and faster collection times. CSCs consist of cathode strips and anode wires arranged into a 2D grid. They provide measurements with a resolution of 40 μm in r, and of 5 mm in ϕ.

4.5 Trigger and data acquisition system

As mentioned earlier in this chapter, while the proton-proton collision rate is approximately 40 MHz, not all of the produced events can or even should be kept. The trigger system used to decide which events should be saved operates in two stages.

A hardware-based trigger called Level 1 (L1) is used first. It makes a decision within 2.5 μs and reduces the rate to 100 kHz [74]. Since the whole detector cannot be read out that quickly, lower-granularity information from the calorimeters and the MS is used to identify regions of interest based on the presence of high-momentum objects as well as specific event quantities such as the sum of energy within the calorimeters.

The second set of triggers is software-based and is called the High-Level Trigger (HLT). The HLT first performs a fast reconstruction to quickly reject uninteresting events. The focus is on the regions of interest identified by the L1 trigger, but in this case the full detector granularity is used. Then, CPU-intensive reconstruction algorithms and calibrations are used for the remaining events for more precise and accurate decision-making. These algorithms are similar to those that will be used to reconstruct, identify and calibrate objects for physics analysis, described in Chapter 5, but slightly modified to accommodate the decision time constraints of a few hundred milliseconds[2]. The physics output of the HLT was of approximately 1.2 kHz during Run 2 [74].

A more in-depth description of the ATLAS trigger system and its performance can be found in [74].

[2]Full offline reconstruction takes approximately 10 s per event.

Chapter 5

Object reconstruction, identification and calibration

As highlighted in the previous chapter, once a hard-scatter event occurs the resulting particles go through the ATLAS detector and interact with the different sub-detectors in characteristic ways. In order to gain knowledge about the physics process associated with the hard scatter, these outgoing particles must be studied. This chapter outlines the steps required to reconstruct the passage of particles through the detector and identify the type of particle involved. The calibrations required to properly measure the energy carried by these particles are also discussed. The particulars of the reconstruction, identification and calibration of τ leptons have important consequences for the measurements presented in this thesis and will therefore be discussed separately in Section 6.

5.1 Track reconstruction

As electrically charged particles go through the inner detector (ID), they interact with the active material of the various layers. With each interaction, an electron-hole pair is created in the silicon. This is referred to as a “hit”. The collection of hits left in the ID by a charged particle is called a track, and it is reconstructed using the New Tracking (NEWT) algorithm. A complete description of the algorithm is available in [81], and a brief overview of its main steps is included below. NEWT aims to reconstruct tracks originating from the hard-scatter event (primary tracks), from the decays of long-lived particles (secondary tracks) and from interactions with the ID material (conversion tracks). It uses hits in the pixel and SCT detectors to find track candidates. A single pixel cluster provides a three-dimensional measurement while in the SCT clusters from both sides of the strip layer are required. Each three-dimensional measurement corresponds to one space point and represents the location where the charged particle interacted with the active material of the ID.

Primary tracks are tracks originating from the hard-scatter event, from particles whose lifetime is of at least 3×10^{-11} s and whose momentum is above 400 MeV. They are reconstructed using an “inside-out” algorithm that uses triplets of space points as seeds. Once the seeds are created, a combinatorial Kalman filter [92, 105] is used to propagate the seeds throughout the rest of the inner detector and full track candidates are formed. These steps are illustrated in Figure 5.1.

During this process, it is possible for hits to be shared between several track candidates, for incomplete

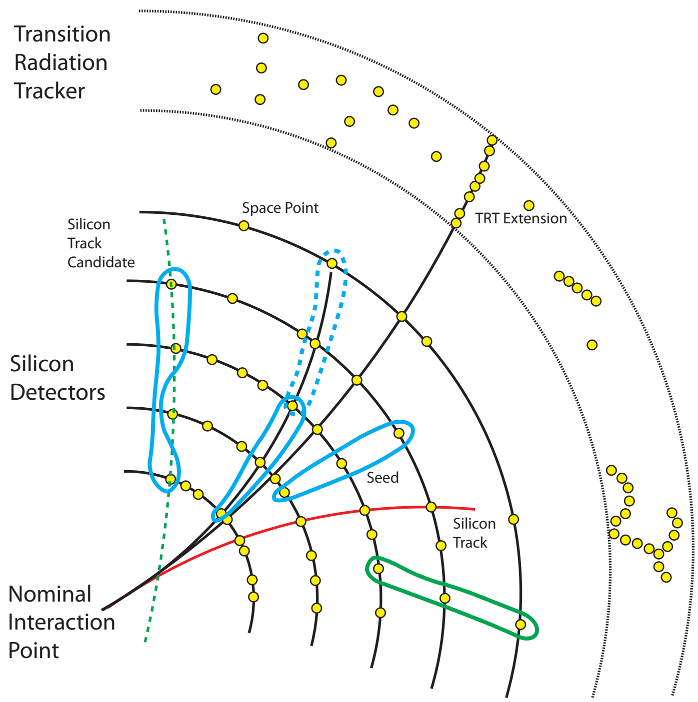

Figure 5.1: First steps of track reconstruction. Space points (yellow dots) are used to create triplet seeds (circles), and track candidates (lines). Blue circles show typical three space-point seeds, with the dotted blue circle showing a triplet that is matched to the same track candidate as another seed. The green circle shows a seed that does not point to the primary vertex and will not be used in further steps. The green dashed line shows a track candidate with the same characteristics. The red track candidate shows a fully-reconstructed track candidate with no TRT extension, while the black track candidate shows a track candidate with hits in all three sub-detectors. Both types of tracks can be used further. Image from [109].

track segments to be created, and for track candidates to originate from more than one particle. All of these scenarios can result in ambiguities regarding what track candidate a space-point belongs to. The ambiguities are solved using a weight-scoring method based on properties of the track candidates such as their momentum and the number of hits and of holes[1] associated to them. Better reconstructed track candidates, such as those with many hits from the (more precise) pixel detector, receive higher scores. Conversely, track candidates with holes receive lower scores. Track candidates that fail basic quality requirements are rejected by the ambiguity solver [53]. A neural network (NN) is used to split shared hits in the pixel detector and to assign them to one or more tracks [40]. The NN identifies and splits shared hits with an efficiency of 90% for shared hits created by two particles, and 85% for those involving three particles. The ambiguity-solving procedure is illustrated on Figure 5.2. If needed, sudden changes in curvature due to momentum loss from Bremsstrahlung interactions are also taken into account during the track reconstruction step. Finally, tracks reconstructed using pixel and SCT hits are extended to the TRT whenever possible.

Once primary tracks are reconstructed, secondary tracks from long-lived particle decays and conversion tracks from material interactions are reconstructed using an "outside-in" algorithm. Since tracks of this type originate from secondary vertices that may be located well within the inner detector, the

[1] Holes correspond to points where a hit is expected based on the track trajectory but none is seen.

search begins outside (in the TRT) and segments are then extrapolated back into the silicon detectors.

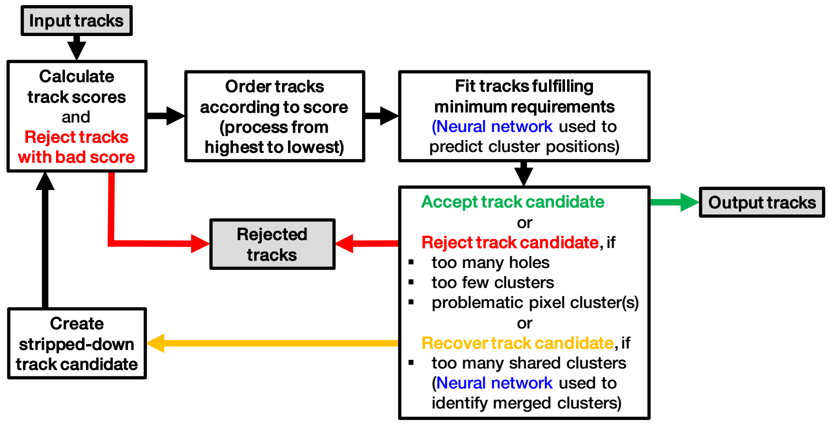

Figure 5.2: Flow of tracks through the ambiguity solver. Image from [53].

5.1.1 Vertex reconstruction

Once tracks are reconstructed, it is possible to reconstruct interaction vertices. A vertex is defined as the point from which tracks originate, and is used to identify objects originating from the hard-scatter interaction.

Vertex seeds are obtained from an iterative process [34]. In the first step, the point of highest track density in the z-direction is located and identified as a seed. Then, the seed and its nearby tracks are used to determine the true vertex location with an adaptive χ^2 fitting algorithm [93]. Any tracks that are incompatible with the reconstructed position of this vertex within 7σ, where sigma refers to the uncertainty on the position of the vertex, are used to seed further iterations of vertex-finding. The process ends when all reconstructed tracks in the event have been matched to a vertex, or when no more vertices can be found.

The primary vertex identifies the location of the hard-scatter event and is defined as the vertex seed whose associated tracks lead to the highest sum of momentum square. All other vertices are labelled as pile-up vertices.

The resolution on the position of a vertex is determined using a split-vertex method. All tracks associated to a vertex are ranked in descending order of their transverse momenta. The tracks are split into two groups: one with even- and one with odd-ranking tracks. Vertex fits are performed separately to the two groups, and the distance between the two resulting vertices gives a measurement of the resolution.

The resolution of vertex location improves with the number of tracks associated to it and with the sum of the tracks' transverse momentum. A transverse (longitudinal) vertex position resolution below 20 (30) μm is usually achieved. Single vertex reconstruction efficiency exceeds 99% in low pile-up conditions, and remains at around 98% for $\mu > 40$ [34]. The efficiency of primary vertex reconstruction decreases with increasing pile-up. Merging of adjacent vertices due to large vertex densities is the dominant effect limiting primary vertex reconstruction efficiency for samples with $\mu > 40$ [34].

5.1.2 Track fitting

Once all tracks have been found and a primary vertex has been identified, tracks are fit using the full inner detector information. The impact parameters of tracks (denoting the location of its perigee[2]) are estimated from global fits to the selected hits. Due to their curvature in magnetic fields, track candidates in the inner detector are parametrised by helices [125]. These helices (τ) are a function of the characteristic values of the track at its perigee,

$$\tau = \tau(d_0, z_0, \phi_0, \theta_0, q/p), \tag{5.1}$$

where d_0 and z_0 refer to the transverse and longitudinal impact parameters, defined as the distance in the transverse plane or z-plane of the perigee to the primary vertex, while ϕ_0 and θ_0 refer to the azimuthal and polar angles of the track candidate at its perigee. Finally, p/q is the ratio of the charge of the particle to its momentum and defines the orientation and curvature of the helix.

The resulting tracks are classified into three groups depending on their quality. The loose category is defined using basic selections requiring a minimum amount of hits in the silicon detectors and applying constraints on the number of shared hits. Tracks from the loose primary category are used to reconstruct τ hadrons and jets from heavy quark decays (called heavy-flavour jets) and are similar to those in the loose category, but are required to be associated to a larger number of hits if shared hits are present. Finally, tight primary tracks must additionally be well measured in the pixel detector. This extra selection makes it sub-optimal for use in the reconstruction of objects that do not originate from the primary vertex.

The efficiency for track reconstruction is stable as a function of μ, and it is reduced as hit requirements are tightened. Reconstruction of tracks from objects subject to hadronic interaction suffers from the hit requirements, leading to compromised reconstruction efficiencies for hadronically-decaying τ leptons and jets with very low transverse momentum (p_{T}). The transverse momentum of tracks is measured to within 2% at low p_{T} and 10% at high p_{T} in the central region, increasing to 15-50 % near $|\eta| = 2.5$ depending on the track's p_{T} [113].

Due to the the fact that there are many tracks per event and many ways to combine hits into track candidates, track reconstruction and ambiguity-solving are both computationally expensive. Event reconstruction time is in fact dominated by the inner detector reconstruction step, and it scales exponentially with pile-up. This is a limiting factor for track reconstruction in the upgrade scenario discussed in Section 10. In particular, an increase in combinatorial fake tracks is observed in events with high pile-up rates [89].

5.2 Electrons and photons

Electrons typically first pass through the inner detector, where a track can be reconstructed, and then continue to the calorimeters. Tracks used for electron reconstruction must pass the tight primary quality requirement described in Section 5.1.

While photons are not charged and therefore will not leave tracks, some photons will produce an electron-positron pair while in the inner detector. Tracks from these photon conversions can be re-

[2]The perigee is the point of closest approach to the primary vertex. The impact parameters denote its position with respect to this vertex in both the longitudinal and transverse directions. Therefore, large values of impact parameters are usually associated with tracks that do not originate from the primary vertex.

constructed up to a conversion radius of 800 mm. Since these tracks are produced inside the ID, reconstruction and identification of track segments from photon conversions are aided by the availability of hits within the TRT [41].

Since photons can emit electron-positron pairs, and electrons can produce photons via Bremsstrahlung, their showers are similar leading their reconstruction, identification and calibration to be performed in very similar ways.

5.2.1 Baseline reconstruction

Electron and photon reconstruction begins with a search for energy deposits in the EM calorimeter. Calorimeter deposits used to reconstruct seeds for both electrons and photons are chosen using a sliding-window algorithm of 3×7 (5×5) cells in $\eta \times \phi$ in the barrel (end-cap). This window size is optimized to take into account the different energy distributions in the barrel and end-cap and typically encompasses 94% of the energy deposits generated by the electron or photon. If the sum of the energy deposits contained within the window is above 2.5 GeV, the window is kept as a seed for the reconstruction [41]. The efficiency for calorimeter-based electron reconstruction ranges from 65% at $E_{\mathrm{T}} = 4.5$ GeV to more than 99% at $E_{\mathrm{T}} = 15$ GeV [67]. For photons this value is of around 70% and is dependent on μ, with the efficiency decreasing to 65% for values of μ around 60 [71].

Once calorimeter clusters are found, an attempt is made to match them to tracks in the inner detector. If no match is found, the cluster is classified as an unconverted photon candidate. If a match is found but the track points to a secondary vertex, the cluster is a converted photon candidate. Finally, electron candidates are generated if a track that points to the hard-scatter vertex is matched to the cluster. A second round of classification is performed to ensure that the candidates are properly classified and should proceed towards either electron, unconverted photon or converted photon identification. This second round of decisions is based on the presence of pixel hits in the track, secondary-vertex information, the momentum of the object and the ratio between its energy (as measured in the calorimeters) and its momentum (as measured in the ID).

5.2.2 Supercluster reconstruction

For analyses using the full Run 2 dataset, a different method for electron and photon reconstruction is available. It aims to reconstruct calorimeter clusters by following the shape of the energy deposits instead of using a fixed window size and shape. Topo-clusters, described in Section 5.4.1, are ideal for this purpose. Supercluster reconstruction [66] proceeds in two steps: first, EM topo-clusters are tested for use as seed cluster candidates. Then, since Bremsstrahlung radiation and cluster splitting can generate other clusters around the seed clusters, these additional clusters are searched for and labelled as satellite clusters.

The procedure is as follows:

- EM topo-clusters are reconstructed as either electron or photon seed candidates. Seed electron clusters are required to have $E_{\mathrm{T}} > 1$ GeV and be matched to a good quality track with at least four hits in the silicon detectors. Seed photon clusters are required to have $E_{\mathrm{T}} > 1.5$ GeV. There is no requirement on tracks or a conversion vertex matching for unconverted photon candidates. Converted photon candidates however are required to be matched to both a track and a conversion vertex.

- Next, satellite clusters are added. Satellite clusters are found sequentially for each seed, with seed clusters ranked according to their E_{T}. Satellite clusters are selected if they are found within a window of $\Delta\eta \times \Delta\phi = 0.075 \times 0.125$ around the seed cluster's barycentre. For matching to electron seeds, satellite clusters are also required to share a best-matched track with the seed cluster. For matching to converted photon seeds, the best-matched track must originate from the same conversion vertex as the seed cluster. No additional requirement is made for matching to unconverted photon seeds.

The term supercluster refers to the seed-plus-satellite system, and is illustrated in Figure 5.3 for each type of object. Since superclusters are built separately for electrons and photons, and no track requirement or veto is part of the photon reconstruction, the ambiguity-solving logic shown in Figure 5.4 was put in place to avoid having objects reconstructed as both a photon and an electron.

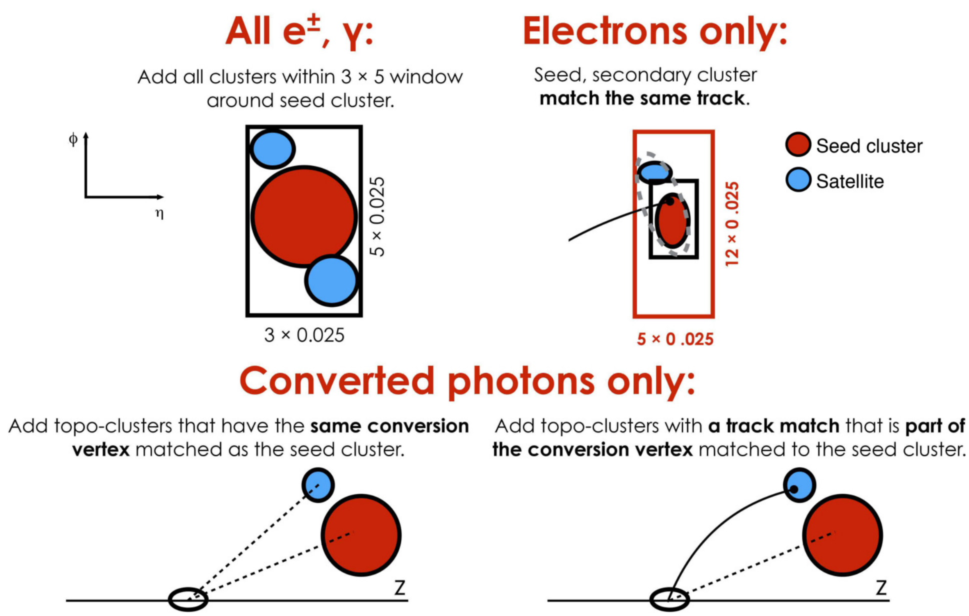

Figure 5.3: Superclustering reconstruction for electrons and photons. Image from [66].

5.2.3 Electron identification

Because electron reconstruction is similar to the reconstruction of other types of jets, discussed in Section 5.4, a likelihood function is used to specifically identify electron candidates [67], with either baseline or supercluster candidates as inputs and exploiting characteristics of electron clusters. These include the amount of activity in the hadronic calorimeters, the ratio of the energy deposited in the different layers of the EM calorimeter, the impact parameters of the tracks, the amount of transition radiation in the TRT and the quality of the matching between the tracks and the clusters.

Three working points are defined in increasing values of thresholds on the likelihood discriminant, allowing for the trade-off between electron identification efficiency and background rates to be adjusted as needed. The loose, medium and tight working points correspond to efficiencies of 93%, 88% and 80%

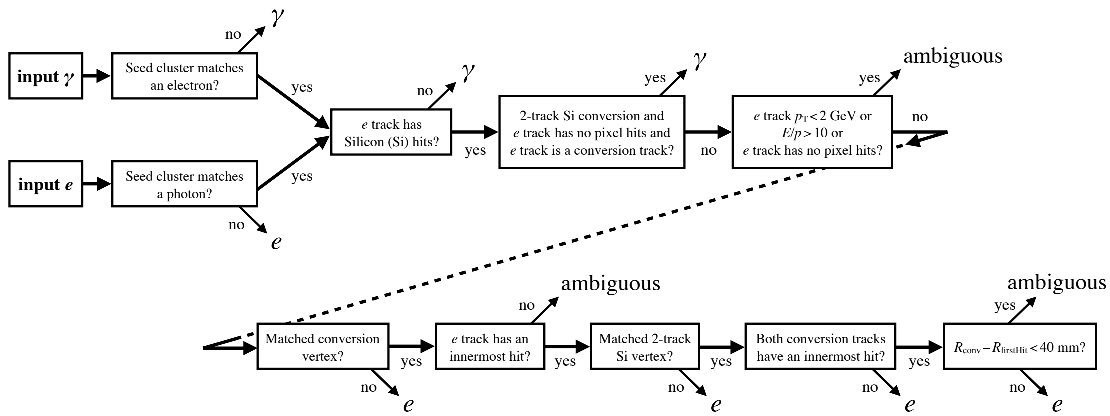

Figure 5.4: Logic of the ambiguity resolution step for particles that were reconstructed both as electron and photon candidates. Image from [66].

for prompt electrons of $E_T = 40$ GeV [67]. These working points also include explicit requirements on the number of hits within the track associated to the electron candidate. To reduce contributions from converted photons mis-reconstructed as electrons, the medium and tight working points additionally require one of the pixel hits to be in the innermost pixel layer. A more relaxed very loose working point is available for background studies. The increase in rejection factors with respect to the loose baseline, measured using multijet samples, are of 2.5 for medium electrons and of 5 for tight electrons [67].

An electron charge misidentification algorithm is used to ensure that the charge of the electron candidate is properly assigned. Misidentification can occur due to matching of an incorrect track to the electron candidate or due to mismeasurement of the curvature of the electron track. The quality of the electron track, the lateral development of the electromagnetic shower and combined tracking and calorimeter quantities are some of the variables used to identify such cases. The efficiency for correct charge reconstruction is of approximately 96% in both data and Monte Carlo for medium electrons over the entire E_T range, and varies from 90% to 99% with η [67].

Electron isolation cuts may additionally be applied, with either or both calorimeter and track information used, resulting in an efficiency between 90% and 99% depending on the isolation cuts chosen.

5.2.4 Photon identification

5.2.4.1 Unconverted photons

Similarly to electrons, unconverted photons are identified through characteristics of the shower shape of the first two EM calorimeter layers as well as the amount of leakage into the hadronic calorimeter. The high granularity of the first layer of the EM calorimeter can be exploited to distinguish photons from the hard-scatter event from those generated by neutral pion decays [41]. Both standard and supercluster reconstruction can be used as an input to the identification step.

Since it is rare for other neutral particles to decay in the EM calorimeter, no further requirements are needed. However, the selection cuts can be adjusted to produce two reference sets leading to loose and tight photon identification points. In particular, the tight selections exploit information from the finely-segmented first layer of the EM calorimeter.

5.2.4.2 Converted photons

Converted photons can be identified as calorimeter clusters with two tracks near each other in ΔR, with a secondary vertex requirement as discussed earlier. The identification procedure is performed in a similar way to unconverted photons, but with slightly different cut values to account for the generally broader lateral profile of converted photon candidates. When the photon conversion occurs with a large opening angle between the decay products, it is possible to match each converted photon cluster to a track containing no hits in the inner layer of the pixel detector [44].

5.2.4.3 General

Photon isolation criteria using tracks or calorimeter information (adjusted for pile-up and UE effects) may also be applied to distinguish photons from jets, and have a negligible effect on the identification efficiency. The overall identification efficiency ranges from 45-60% at $E_{\mathrm{T}} = 10$ GeV to 95-98% for $E_{\mathrm{T}} > 100$ GeV for both types of photons [71].

Misidentification of electrons as photons occurs infrequently, with rates of 0.015 ± 0.001 for converted photons and 0.030 ± 0.002 for unconverted photons [71]. Misidentification rates of unconverted photons as converted photons are below 5% [44].

5.2.5 Calibration

The energy recorded in the reconstructed clusters is not equivalent to the amount of energy carried by the particle studied. The fact that the reconstruction window of the baseline reconstruction does not contain all of the energy deposited by the incoming particle is one example of how the two can differ. In order to accurately measure the energy carried by an object, an energy calibration must be applied to the energy measured in the calorimeters. The energy calibration of electrons and photons is done in a four-step process [44].

For electrons and photons obtained from data, the energy scale of the first two layers of the EM calorimeter is inter-calibrated by comparing data to Monte Carlo simulation. Then, the reconstructed and truth energies of photons and electrons can be compared using Monte Carlo simulation, and a Monte Carlo based correction is derived. This correction accounts for the lateral shower leakage mentioned above. An additional correction to account for the detector uniformity and the stability of the calorimeter response is added, by comparing the momentum and energy measurements performed in the ID and the calorimeter. Finally, calibrations derived using $Z \to ee$ events are used for electrons and both types of photons to account for residual differences between data and simulation.

The uncertainty in the energy scale corrections ranges from 0.04% in the barrel to 0.2% in the endcap. The energy resolution of soft electrons is improved by up to 20% by combining ID and calorimeter measurements. The uncertainty in the resolution corrections is between 0.1% and 0.2% [66].

5.3 Muons

Unlike electrons, muons do not lose all of their energy in the calorimeters and instead reach the muon spectrometer mostly intact. Muon tracks can therefore be reconstructed both in the inner detector and in the muon spectrometer, with the latter's measurements starting to play a role in the momentum measurement of muons with $p_{\mathrm{T}} > 30$ GeV.

5.3.1 Muon reconstruction

ID track reconstruction is performed as in Section 5.1. Track reconstruction for the muon spectrometer begins with track segments created using an "outside-in" method where hits inside each muon chamber are used to form segments. Straight line segments are created from hits in the MDTs and CSCs. Once segments have been created, they are combined into track candidates, first put together using segments from the outer and middle sections of the MS, and then extrapolated back to segments reconstructed in other sections. Since segments could end up in more than one track candidate, an overlap removal is applied to select the best segment matches, with no more than two track candidates using the same segment. This procedure is performed separately in the MS and the ID. MS and ID measurements can be combined to form tracks with hits in both sub-detectors, again starting from the outside (MS) and going in (ID). Hit multiplicity and fit quality requirements are also applied independently for track segments from the inner detector and muon spectrometer [46]. The hits associated with each track candidate are fitted, and a cut on the resulting χ^2 is applied [46].

Muon candidates can be reconstructed in four ways:

- Standalone (ME): these muon candidates only have tracks in the MS. They can be used when looking at regions outside the ID acceptance.
- Combined (CB): these muon candidates use tracks created via a combination of ID and MS tracks. The combination results in better resolution in particular for muon candidates with a transverse momentum of up to 100 GeV.
- Segment-tagged (ST): when only track segments can be found in the muon spectrometer, these can be combined with full tracks from the inner detector to reconstruct ST muon candidates. This process allows for improved reconstruction efficiencies for muon candidates with $p_{\mathrm{T}} < 6$ GeV and for muon candidates that traverse regions of the detector where the acceptance is poor.
- Calorimeter-tagged (CT): when no information is available from the muon spectrometer (such as within $|\eta| < 0.1$), muon candidates can be reconstructed from a combination of ID tracks and calorimeter deposits consistent with minimally-ionizing particles.

When overlaps between muon types occur they are resolved with preference given to CB, then ST, and finally CT. Overlaps with ME muons are resolved by selecting the track with the best fit quality and largest number of hits.

5.3.2 Muon identification

5.3.2.1 Non-prompt muons

Muons produced during the decays of other particles can be identified through two characteristics: the presence of kinks, leading to a poor fit quality, and an incompatibility between the ID and MS momentum measurements.

5.3.2.2 Prompt muons

Four identification categories exist for prompt muons, based on the characteristics of the reconstructed track [46]:

- Loose: all types of muons described in Section 5.3.1 are used without restrictions.
- Medium: in order to minimize the uncertainties related to reconstruction and calibration, only CB muons are used in this category. Outside of the ID acceptance region, ME muons are considered. Requirements on the number of hits are added to ensure the selection of good quality tracks. A loose selection on the compatibility of the momentum measurement from the ID and MS is applied to suppress contributions from muons from in-flight decays and hadrons misidentified as muons.
- Tight: only CB muons from the medium category are used here. For muons with a transverse momentum below 20 GeV, more stringent selections are applied on the charge-to-momentum ratio and on the compatibility of the ID and MS momentum measurements. This category maximizes the muon selection purity.
- High-p_{T}: the muons in this category are required to have $p_{\mathrm{T}} > 100$ GeV. They are required to pass the medium requirements, be of CB type, and contain at least three hits in three different MS sections. These requirements reduce muon identification efficiency by approximately 20% but allow for good momentum resolution above 1.5 TeV.

Additional requirements may be placed on the muon isolation, with 7 working points available from a combination of cuts on the scalar sum of transverse momenta around the muon track and on the sum of transverse energy of topo-clusters within a cone of $\Delta R = 0.2$ of the muon. Further rejection rates to non-prompt muons in Monte Carlo vary between 15 and 40 depending on the working point chosen, with minimal impact on the muon efficiencies [46].

5.3.3 Calibration

Muon momentum scale and resolution are calibrated with CB muons using the J/ψ, Υ and Z mass peaks. Scale factors accounting for differences between data and simulation are derived separately for the ID and the MS, in bins of η and ϕ. The relative error on the momentum measurement is of $\Delta p_{\mathrm{T}}/p_{\mathrm{T}} \sim 3\%$ for a wide momentum range, and increases to about 10% for high-p_{T} muons [46].

5.4 Jets

Unlike the other objects discussed so far, jets are not particles from the Standard Model but rather a tool designed to represent a physical process. Before partons go through the ATLAS detector, QCD confinement causes them to undergo hadronisation, leading to a large number of particles produced in a collimated cone called jet. The methods for jet reconstruction have evolved throughout the course of this thesis, and thus two different methods need to be considered. Topo-cluster jets [39, 56] are used in analysis chapters focusing on results with partial Run 2 datasets, while particle flow jets [50, 70] are used in later analyses looking at the full Run 2 dataset. Jets associated to heavy-flavour quarks and hadronic τ decays can be identified using dedicated algorithms: heavy-flavour identification is described in Section 5.6 and τ-lepton reconstruction, identification and calibration are outlined in Chapter 6.

5.4.1 Topological clusters

When a hadron passes through the detector, clusters of hard-scatter activity in the calorimeters can be built by gathering neighboring calorimeter cells containing a significant amount of energy. A significant energy is one that is larger than what is expected from the electronic noise and pile-up levels, with the latter being the dominant contribution. Cells with an energy higher than the expected noise by at least four times its uncertainty are labelled primary and used as an input to the topo-clustering algorithm [39, 56]. The algorithm looks for neighboring secondary cells with energy levels above noise by two standard deviations in an iterative process. Neighbors of those cells with any energy above noise levels are called tertiary. This approach allows for the shape of the topo-cluster to follow that of the energy deposits in three dimensions. Clusters are kept if they have a local maximum with $E_{\mathrm{cell}} > 500$ MeV and at least four neighboring secondary or tertiary cells. If two maxima are found, it is likely that two sub-jets are present, and the cluster is split. Finally, the energies from the cells within the cluster are summed, the mass of the object is set to zero, and the angle of the object is set to match that of its energy-weighted centroid.

The resulting topo-clusters are used as inputs to an anti-k_t algorithm [15] with radius parameter $R = 0.4$ using the FASTJET 2.4.3 software package [16]. They are reconstructed at the electromagnetic (EM) energy scale, meaning that no attempt to correct for calorimeter non-compensation is made at this point. Only jets with $p_{\mathrm{T}} > 7$ GeV are kept [49]. Jets reconstructed in this way are referred to as EM topo jets.

Similarly to the procedure outlined for electrons, once a jet is reconstructed tracks from charged particles are matched to the topo-cluster [14]. To ensure a reconstruction procedure that is stable against pile-up the tracks are required to have a transverse momentum of at least 500 MeV, be within the ID acceptance and be associated with the event's hard-scatter vertex. They are also required to be associated to at least nine hits in the silicon detectors and have no holes in the pixel detector.

5.4.2 Particle flow

While EM topo jets are reconstructed based on calorimeter-only information, the particle flow algorithm aims to combine the momentum and energy measurements made in the ID and in the calorimeters respectively to improve the jet reconstruction performance, especially in regards to its stability with increasing pile-up levels. For charged particles, the tracks associated to topo-clusters sometimes allow for better determination of their energy than the calorimeter signals can. A set of modified topo-clusters can be created by replacing the energy measured in the calorimeters by that measured in the ID when this is the case, allowing for a better overall energy measurement. The full particle flow procedure is described in [50, 70] and a summary is presented below.

First, well-measured tracks are selected with requirements that are similar to those used for tracks associated to EM topo jets [70], except for the vertex requirements. The particle flow algorithm is optimized for the measurement of hadronic energy deposits and therefore tracks associated with electron and muon candidates of medium quality are not used. The selected tracks are matched to single topo-clusters in the calorimeters. Tracks are ranked using a metric that accounts for their distance to topo-cluster barycentres in both η and ϕ.

Based on a track's momentum, the expected amount of energy deposited in the calorimeter by the particle in question can be computed. The visible jet energy is mostly carried by charged pions, with another part of the energy going to neutral pions and the remainder carried by neutral hadrons [50]. Tracks, which always originate from charged particles, are therefore likely to be generated by charged pions. Monte Carlo simulation of single-pion samples can therefore be used to determine the average amount of energy deposited in the calorimeter based on the track momentum.

Lastly, a subtraction is performed. The track position is extrapolated into the calorimeter. The energy expected to be deposited in the topo-cluster by the same particle that created the track is removed from the topo-cluster cells, first from the core of the extrapolated track and then going outwards. This subtraction procedure removes the correct amount of energy on average, with a slight bias towards negative values indicating that it is more likely to remove additional energy from neutral particles rather than to miss energy from charged ones. After this step, the energy of the topo-cluster no longer contains that of particles that are best measured using ID information.

Topo-clusters with remaining energy larger than 1.5σ away from the subtracted energy are assumed to contain energy deposits from neutral particles, and so are kept [50]. Conversely, topo-clusters with no remaining energy, or energy within the error of the subtracted energy, are removed to avoid double-counting with the tracks. Since pile-up tracks are accounted for in the subtraction procedure, the jet energy after subtraction is sensitive mainly to neutral pile-up contributions.

The surviving tracks and topo-clusters are used as seeds to an anti-k_t algorithm with radius parameter $R = 0.4$. At this stage the tracks are additionally required to be matched to the hard-scatter vertex, and similarly to the EM topo reconstruction, only jets with $p_{\mathrm{T}} > 7$ GeV are kept [70].

5.4.3 Calibration

As discussed for electrons and photons, the energy measured within clusters does not always accurately describe the energy of the incoming particle. Jet calibration aims to restore the jet energy scale (JES) to that of particle-level jets for both EM topo and particle flow jets within $20 < p_{\mathrm{T}} < 1500$ GeV [50]. It follows the six-step process shown on Figure 5.5. The full calibration procedure is explained in [49].

The first step consists of a simple correction of the origin of the jet four-momentum, ensuring that it aligns with the hard-scatter vertex rather than the centre of the detector. Next comes an area-based pile-up correction. This correction removes energy associated to both in- and out-of-time pile-up contributions. The average pile-up energy density of the pile-up correction is measured in the central, lower-occupancy region of the detector. It does not fully describe the pile-up conditions in the forward region, or in the high-occupancy core of high-momentum jets, so the residual pile-up dependencies are removed by a Monte Carlo-based correction. An absolute JES correction handles the correction of the pile-up corrected four-momentum to that at the particle-level scale using reconstructed jets that can be matched to truth jets in simulation. It accounts for biases due to the transition between different calorimeter technologies and sudden changes in calorimeter granularity. The global sequential calibration aims to further improve the reconstructed energy by reducing the effects of the energy distribution within the jet as well as any flavour-dependencies. Finally, an in-situ calibration is applied only to jets in data and relies on the p_{T} balance with well-measured objects such as photons, electrons and muons from Z boson decays and well-calibrated central jets to correct for any mismodellings from the previous Monte Carlo-based calibration steps.

By the end of this procedure, the uncertainty on the jet energy scale for central jets is typically of about 1% and increases to about 5% at very low (around 20 GeV) and very high (above 2.5 TeV) momenta [70]. Another important measure is the spread of the $E_{\mathrm{truth}}/E_{\mathrm{reco}}$ distribution as it describes the fraction of jets that are calibrated to the correct scale. This spread is quantified by the jet energy resolution, with values of 17.5% for a representative EM topo jet of $p_{\mathrm{T}} = 30$ GeV and 14% for its particle flow counterpart [50].

If needed, surviving jets originating from pile-up interactions can be identified and removed using the jet vertex tagger (JVT) algorithm [47]. This algorithm uses track-based variables to reject, using its

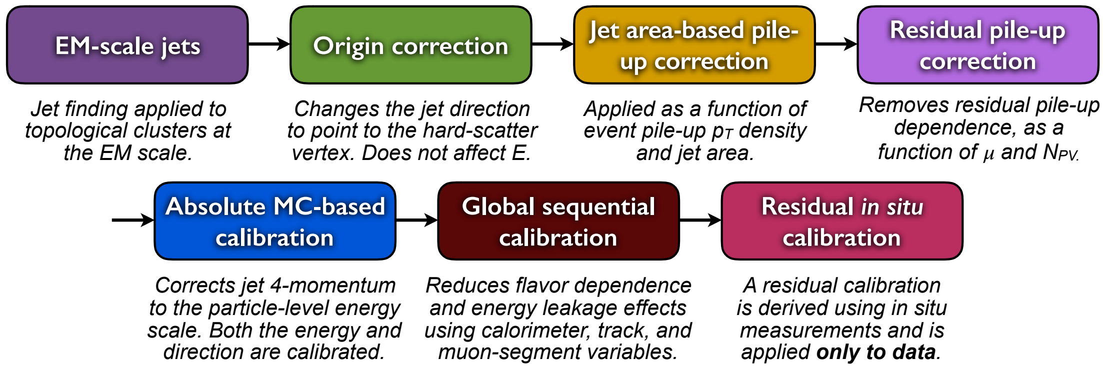

Figure 5.5: Calibration of EM-scale jets. Image from [49].

medium working point, 98% of pile-up jets while maintaining a 92% acceptance for hard-scatter jets.

5.5 Missing transverse energy

Since some particles such as neutrinos can't be detected inside of ATLAS, the amount of transverse momentum they carry is instead quantified by the missing transverse energy ($E_{\mathrm{T}}^{\mathrm{miss}}$) observable [61]. To define it, the total amount of transverse energy $E_{\mathrm{T}}^{\mathrm{tot}}$ deposited in the ATLAS detector from hard scatter objects is computed.

Since the protons that participate in the collisions are not expected to have any significant transverse energy, the partons involved in the hard-scatter interaction are also expected to have negligible amounts of transverse energy. By conservation of energy the sum of the transverse energies of all outgoing particles must then cancel out. Therefore the transverse energy of the undetected particles must balance out the detected one.

$$E_{\mathrm{T}}^{\mathrm{miss}} = -E_{\mathrm{T}}^{\mathrm{tot}} \tag{5.2}$$

Two sets of objects enter the $E_{\mathrm{T}}^{\mathrm{tot}}$ reconstruction: hard objects comprising all of the reconstructed and calibrated particles (electrons, photons, muons, hadronically-decaying τ leptons and jets), and a soft term consisting of charged particle tracks that are associated with the hard-scatter vertex but are not associated to any of the reconstructed particles. Neutral particles are not used in the soft-event signals due to the large contributions from pile-up.

$$E_{\mathrm{T}}^{\mathrm{miss}} = -\sum_{i\epsilon\{\text{hard objects}\}} p_{\mathrm{T},i} - \sum_{j\epsilon\{\text{soft term}\}} p_{\mathrm{T},j} \tag{5.3}$$

ATLAS reconstruction steps are not exclusive: while a hypothesis is made on whether objects satisfy the requirements for each kind of particle, no attempt is typically made at reconstruction level to resolve ambiguities. A single set of detector signals could therefore be reconstructed as more than one type of object, for example an electron candidate may also be a τ candidate. In order to remove double-counting of signals in the $E_{\mathrm{T}}^{\mathrm{miss}}$ reconstruction process, an explicit ambiguity resolution is introduced. The most commonly-used order for this ambiguity resolution is the following: for each hard object, electron reconstruction is given the highest priority, followed by photon, τ-lepton, and finally jet reconstruction.

Since muons are primarily reconstructed using ID and MS tracks alone, they do not have significant overlap with other reconstructed particles, except with jets.

Jets can overlap with muons for example if the jet deposit comes from pile-up but is matched to a muon track that enables it to pass the hard-scatter requirements. In this case the jet should be removed. Jets may also overlap with muons if the muon emits final state radiation that results in the jet. In this case the jet should be kept. Various variables comparing the momentum of the jet and of the muon track, as well as the shower shape of the jet, can be used to identify these cases and select the correct object.

Jets may also have fractional overlap with electrons or photons. If the jet carries more than twice the amount of energy of the electron or photon, it is kept and its energy is scaled down. Otherwise, the jet is removed but its associated tracks (excluding the one associated to the overlapping particle, if any) are kept within the soft term.

An important point to note is that different analyses might use different requirements to define the objects that they use, for example by adjusting identification working points. Therefore, a distinction is made in the $E_{\mathrm{T}}^{\mathrm{miss}}$ reconstruction between selected and accepted objects, where the accepted objects category may be re-defined by individual analyses according to their needs. This results in the following $E_{\mathrm{T}}^{\mathrm{miss}}$ definition:

$$\begin{aligned} E_{\mathrm{T}}^{\mathrm{miss}} = &- \sum_{\text{selected electrons}} \mathbf{p_T}^{e} - \sum_{\text{accepted photons}} \mathbf{p_T}^{\gamma} - \sum_{\text{accepted } \tau-\text{leptons}} \mathbf{p_T}^{\tau} \\ &- \sum_{\text{selected muons}} \mathbf{p_T}^{\mu} - \sum_{\text{accepted jets}} \mathbf{p_T}^{\text{jet}} - \sum_{\text{unused tracks}} \mathbf{p_T}^{\text{tracks}} \end{aligned} \tag{5.4}$$

The soft term aims to account for primary tracks of $p_{\mathrm{T}} > 400$ MeV. The tracks must also be reconstructed within the ID acceptance. As mentioned, they are required to be associated to the hard-scatter vertex to reduce contributions from pile-up to the soft term. Because of these requirements, very soft and neutral contributions from the hard scatter are not accounted for by the soft term. For a complete list of requirements applied to each object, see [61].

5.5.1 Validation

Events with no expected $E_{\mathrm{T}}^{\mathrm{miss}}$ such as those from $Z \to \mu\mu$ decays can be used to evaluate the performance of the $E_{\mathrm{T}}^{\mathrm{miss}}$ reconstruction. Even in these final states, where no $E_{\mathrm{T}}^{\mathrm{miss}}$ generated by undetectable particles is expected, a non-vanishing value of $E_{\mathrm{T}}^{\mathrm{miss}}$ is observed since not all of the p_{T} contributions from the hard-scatter interaction can be reconstructed, and since the resolution and acceptance of the detector is limited. It is also observed that the soft-term is inherently underestimated given that it only accounts for charged contributions that result in well-measured tracks.

Data-Monte Carlo comparisons with $Z \to \mu\mu$ events show an agreement within 20% for $E_{\mathrm{T}}^{\mathrm{miss}}$ distributions, as illustrated in Figure 5.6 [61]. The resolution[3] varies between 5% and 25% and increases as a function of $E_{\mathrm{T}}^{\mathrm{tot}}$. Is observed to be 10% worse in data than in Monte Carlo for low pile-up events ($N_{\mathrm{PV}} < 5$), and 10% better where there is high pile-up activity ($N_{\mathrm{PV}} > 20$) [61].

Other processes such as Z, W and $t\bar{t}$ production with non-vanishing expected $E_{\mathrm{T}}^{\mathrm{miss}}$ are also used to

[3]Determined by the width of the combined distribution of the differences between the measured $E_{\mathrm{T}}^{\mathrm{miss}}$ and the components of the true missing transverse momentum.

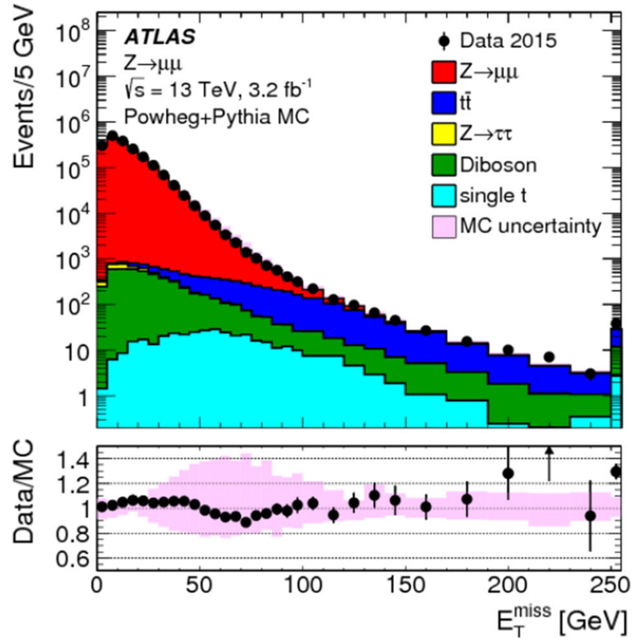

Figure 5.6: Data-Monte Carlo comparison of $E_{\mathrm{T}}^{\mathrm{miss}}$ for an inclusive sample of $Z \to \mu\mu$ events, with all relevant processes included.

validate the performance of the $E_{\mathrm{T}}^{\mathrm{miss}}$ resolution. They indicate a higher sensitivity to residual pile-up in events with higher jet multiplicities, due to increased contributions from pile-up jets to the hard term and increased fluctuations in the jet energy response due to pile-up [61].

5.6 Flavour tagging

As mentioned in Section 5.4, it is sometimes possible to identify the object that produced a given jet. In particular, flavour tagging algorithms aim to identify jets containing b and c hadrons. These jets have characteristic shower profiles, and displaced vertices on the order of 450 μm. Heavy-flavour jets can be identified using a variety of algorithms that exploit information related to the trajectories of charged particles through the ID as well those from calorimeter clusters to discern these characteristics [52]. Jets considered for flavour-tagging can be reconstructed using either the topo-cluster or the particle flow algorithms.

5.6.1 Tagging algorithms

Five low-level taggers and two high-level taggers are available for flavour tagging, with the low-level taggers serving as inputs to their high-level counterparts. All of these taggers are described in [52], and an overview is provided below.

1. **Low-level taggers**

   (a) **IP2D and IP3D**

These taggers exploit impact parameter information to identify characteristic displaced vertices, and are based on a log-likelihood ratio discriminant.

(b) **RNNIP**

This recurrent neural network is also based on impact parameter information but it accounts for track-to-track correlations. The idea is that in the decay of a b or c hadron several particles are produced with large impact parameters. In this case, if one track with a large impact parameter is found, finding additional tracks with large impact parameters is likely. In the case of a light-flavour decay, this correlation should not be observed. In general, the RNNIP algorithm is better at identifying decays with long decay distances and high track multiplicity, while the IP2D/IP3D algorithms provide better identification efficiency for decays with short distances and low track multiplicity.

(c) **Soft Muon Tagger**

This tagger exploits the presence of muons in semi-leptonic b and c decays to identify heavy-flavour jets.

(d) **SV1 and JetFitter**

These are a set of vertex taggers that aim to reconstruct the heavy-flavour decay topology.

2. **High-level taggers**

These taggers use different machine-learning algorithms to identify jets containing b and c hadrons by combining the discriminating variables produced by the low-level taggers. MV2 algorithms are based on Boosted Decision Tree discriminants while DL1 algorithms use Artificial Deep Neural Networks. For both algorithms, b-jets are considered as the signal while c- and light jets are considered background. The performance is comparable between the DL1 and MV2 algorithms and corresponds to rejection rates for light (c) jets of 30 (3) for a typical 85% efficiency working point.

5.6.2 Validation and performance

The performance of the flavour-tagging algorithms is evaluated using $t\bar{t}$ and Z samples [52]. In a data-Monte Carlo comparison, consistently good agreement for all input variables to the MV2 algorithms is observed, with discrepancies of at most 10%.

A slight decrease in performance is seen with increased pile-up, due to a shift in light-jet and c-jet variable distributions closer to the values expected for b-jets. In fact increased hadronic activity of any kind (due to nearby jets, or gluon radiation) also decreases the performance. In general, the energy resolution of heavy-flavour jets tends to be worse than that of light jets due to the presence of neutrinos in the decay [52].

Chapter 6

Tau lepton reconstruction, identification and calibration

The study of τ leptons is of special interest in the context of this thesis as the performance of the τ reconstruction, identification and calibration algorithms determines the sensitivity of the measurements that can be performed. Efficient reconstruction and identification are crucial for searches such as $t\bar{t}H$ where a lack of statistics is a primary limitation. In other cases, such as the study of $H \to \tau\tau$ decays, excellent energy calibration as well as discrimination against other objects that could be misidentified as τ leptons are key to enhancing the precision of the measurement.

The signatures of τ leptons within the ATLAS detector are complex and require an understanding of both the characteristics of τ decays as well as of the reconstruction of all types of particles mentioned in the previous chapter. As outlined in Chapter 2, τ leptons are the heaviest of the leptons at 1.77 GeV. Given their decay length of approximately 87 μm, τ leptons decay in the beam-pipe and only their decay products reach the ATLAS detector. These decays occur through a W boson as illustrated in Figure 6.1. The number of charged particles within the decay products determines the "prongness" of the τ lepton with the most common decays resulting in 1-prong (85%) and 3-prong (15%) τ leptons [107]. The most common branching ratios are summarized in Table 6.1.

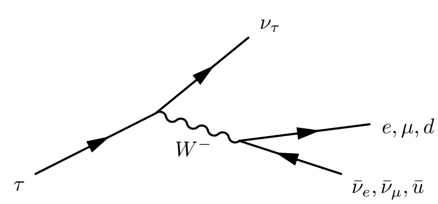

Figure 6.1: Feynman diagram illustrating a τ decay.

Leptonically-decaying τ leptons are reconstructed as electrons or muons and can only be identified in events where extra information is available. For example, $H \to \tau\tau$ decays result in a constrained system within which the characteristics of the neutrinos produced in the leptonic τ decay can be estimated, and so leptonically-decaying τ-leptons can be identified as part of a system containing an electron or a muon

Prongness	Decay	Final state	Branching ratio
1-prong	Leptonic	$\nu_\tau e \nu_e$	17.8 %
		$\nu_\tau \mu \nu_\mu$	17.4%
	Hadronic	$\pi^\pm \pi^0 \nu_\tau$	25.5 %
		$\pi^\pm \nu_\tau$	10.8 %
		$\pi^\pm \pi^0 \pi^0 \nu_\tau$	9.3 %
		$\pi^\pm \pi^0 \pi^0 \pi^0 \nu_\tau$	1.0 %
		other	3.6 %
3-prong	Hadronic	$\pi^\pm \pi^\pm \pi^\mp \nu_\tau$	9.0 %
		$\pi^\pm \pi^\pm \pi^\mp \pi^0 \nu_\tau$	2.7 %
		other	2.9 %

Table 6.1: Branching ratios of the dominant τ-lepton decays [107].

as well as missing transverse energy. Since this can be achieved through identification of light leptons (e and μ) and $E_{\mathrm{T}}^{\mathrm{miss}}$, no τ-specific reconstruction is required.

Hadronically-decaying τ leptons will be the focus of the next few sections. Since neutrinos cannot be reconstructed individually with the ATLAS detector, only the visible products of the hadronic τ decay ($\tau_{\text{had-vis}}$) are considered.

6.1 Reconstruction

Since the visible decay products of hadronically-decaying τ leptons consist of charged and neutral hadrons (predominantly pions), they are first reconstructed as any other topo jet.

6.1.1 Jet seeds

Jets built with the anti-k_t algorithm with a distance parameter $R = 0.4$ are used as inputs for the $\tau_{\text{had-vis}}$ reconstruction[1] and are calibrated with what is called a local cluster weighting scheme (LC) [49]. This calibration accounts for the non-compensating nature of the ATLAS calorimeters and aims to bring electromagnetic and hadronic showers to the same scale. This step is of particular importance in determining the energy of $\tau_{\text{had-vis}}$ candidates as the neutral pions that they contain decay to photons and produce mainly electromagnetic showers while the charged pions produce hadronic ones. It is however important to note that this calibration treats not only the overall energy scale of the jet, but also the energy of each cluster within it. This improves the position resolution of $\tau_{\text{had-vis}}$ as the energy centroid of the jet can be correctly located when each cluster is properly calibrated. Jets seeding $\tau_{\text{had-vis}}$ candidates are required to have a transverse LC-calibrated energy of at least 10 GeV and to be within $|\eta| < 2.5$.

6.1.2 Track selection

Because the charged particles from the $\tau_{\text{had-vis}}$ decay leave tracks in the inner detector, it is important to identify them during the reconstruction process. These tracks are separated from other types of tracks via boosted decision trees (BDTs) [12, 91]. The BDTs used in the $\tau_{\text{had-vis}}$ reconstruction classify tracks into four categories:

[1]Clusters within $\Delta R = 0.2$ of the jet axis are considered to be part of the $\tau_{\text{had-vis}}$ candidate for the energy calibration.

- Tau tracks: tracks from the charged hadrons produced during the τ decay.
- Conversion tracks: tracks originating from photon conversions, hadronic interactions, or decays of other long-lived particles.
- Isolation tracks: tracks from the underlying event.
- Fake tracks: tracks from pile-up.

The first step is to identify the vertex that is matched to the τ. This is important since a large fraction of the energy of the τ may be contained within the neutral hadrons of the decay. These hadrons leave no tracks in the ID, resulting in an event whose hard scatter vertex has a low track momentum sum. In particular, processes such as $W \to \tau\nu$ may have a primary vertex with a track momentum sum lower than that of typical pile-up vertices, and therefore the primary vertex of the event may be misidentified. This is not in general a problem for other objects. For example, in the equivalent process $W \to e\nu$ the energy of the electron is well-described by its track and so the track momentum sum is a good quantity to use to identify the vertex of the hard scatter.

To identify the vertex that the τ originated from, a very loose selection is first applied to all the tracks within the τ jet seed to remove those that are unlikely to come from $\tau_{\text{had-vis}}$ decays. It requires that the tracks have a momentum of at least 0.5 GeV and that they be located within an angular distance of $R = 0.2$ around the seed jet axis, defined as the energy-weighted barycentre of all the topo-clusters associated with the seed jet. The surviving tracks are used to select the τ production vertex from all vertex candidates reconstructed as in Section 5.1, and all track impact parameters are re-computed with respect to this new vertex.

At this stage, a two-step selection process is used to classify the tracks. All tracks within an angular distance of $R = 0.4$ around the seed jet axis are used. The classification process consists of a first BDT to split τ and conversion tracks from isolation and fake ones, and a second BDT to define the τ and conversion track categories. Track kinematics, impact parameters and number of hits in each inner detector sub-section are some of the variables used by the BDT to discriminate between the different categories.

The prongness of the $\tau_{\text{had-vis}}$ candidate is defined by the number of tracks from the jet seed that are classified as τ tracks.

6.2 Identification

The procedure of $\tau_{\text{had-vis}}$ reconstruction is so far almost identical to that of other types of jets. An identification algorithm is therefore used in a subsequent step to distinguish $\tau_{\text{had-vis}}$ candidates from quark- and gluon-initiated jets, also called QCD jets. This is possible due to the fact that $\tau_{\text{had-vis}}$ tend to be collimated jets with low track multiplicity. Machine learning algorithms that use inputs related to track characteristics as well as shower shapes in the calorimeters are used to perform the classification. The $\tau_{\text{had-vis}}$ identification algorithm changed throughout the course of Run 2. Two algorithms were in use: a BDT-based one for analyses focusing on the partial Run 2 dataset and a Recurrent Neural Network (RNN) [95] version was later made available for results with the full Run 2 dataset.

6.2.1 Tau identification with Boosted Decision Trees

A BDT-based $\tau_{\text{had-vis}}$ identification [45] is trained separately for 1- and 3-prong $\tau_{\text{had-vis}}$ candidates, on a $\gamma^* \to \tau\tau$ signal sample for true $\tau_{\text{had-vis}}$ and an inclusive dijet sample for the QCD jet background. The jet transverse momentum distribution is re-weighted to match that of $\tau_{\text{had-vis}}$ in the $\gamma^* \to \tau\tau$ so that no bias is introduced based on the object's p_{T}. Variables such as track impact parameters, track-momentum-to-calorimeter-energy ratios, and shapes of the energy deposits in the calorimeter are used by the BDT to distinguish between $\tau_{\text{had-vis}}$ and QCD jet candidates. Once the algorithm is trained, the output score is transformed so that $\tau_{\text{had-vis}}$ identification efficiencies are constant as a function of both the object's p_{T} and the average number of interactions per bunch crossing μ of the event. Three working points are defined, with target efficiencies of 85%, 75% and 60% for loose, medium and tight for 1-prong $\tau_{\text{had-vis}}$ candidates, and 75%, 60% and 45% for 3-prong $\tau_{\text{had-vis}}$ candidates. An additional very-loose working point of 95% efficiency for 1- and 3-prong $\tau_{\text{had-vis}}$ candidates can be used to estimate the rate at which jets pass the $\tau_{\text{had-vis}}$ working point selections (also called a fake τ estimate). The corresponding rejection factors are shown in Figure 6.2 and range between 10 and 40 for 1-prong and 70 and 400 for 3-prong $\tau_{\text{had-vis}}$.

6.2.2 Tau identification with Recurrent Neural Networks

An improved rejection for same $\tau_{\text{had-vis}}$ efficiency can be achieved through the use of an RNN instead of a BDT [69]. A combination of low-level and high-level variables similar to those used for the BDT are used in the RNN. However, the RNN considers up to 10 p_{T}-ordered tracks and up to 6 E_{T}-ordered clusters per $\tau_{\text{had-vis}}$ candidate, instead of the 1 or 3 tracks and single large calorimeter topo-cluster considered by the BDT. This information is combined to gain insight on the entire $\tau_{\text{had-vis}}$ system, as shown in Figure 6.2.

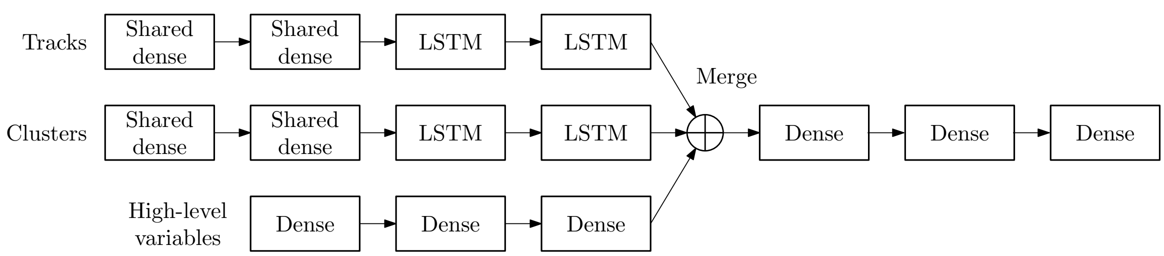

Figure 6.2: Architecture of the RNN used for tau identification. The term LSTM refers to long short-term memory architectures [103]. Image from [69].

The same three working points are defined, and a very-loose selection is also made available for fake estimates. The rejection power of the RNN is about twice that of the BDT at same efficiency, as can be seen in Figure 6.3. This improvement is reduced with increasing pile-up, particularly for 3-prong $\tau_{\text{had-vis}}$ candidates due to their wider, higher-track multiplicity signatures which are inherently harder to distinguish from QCD jets. The rejection power also varies with the quark-to-gluon composition of the background sample, as quark-initiated jets tend to be narrower and have lower track multiplicities than gluon-initiated jets, making them more difficult to differentiate from true $\tau_{\text{had-vis}}$ objects.

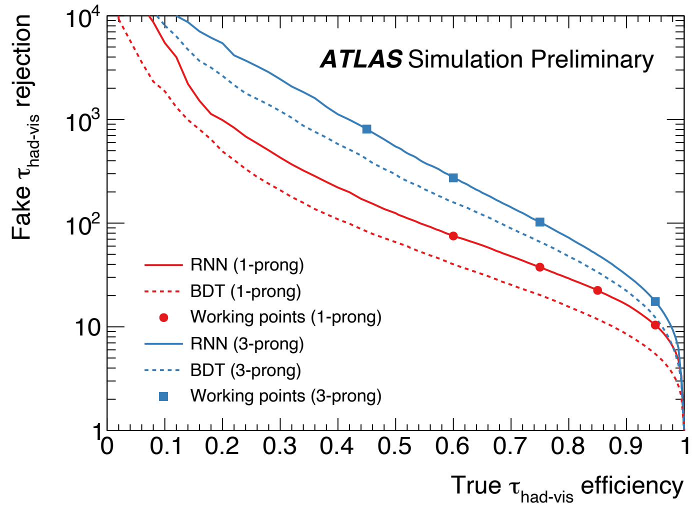

Figure 6.3: Rejection power for QCD jets misidentified as $\tau_{\text{had-vis}}$ as a function of $\tau_{\text{had-vis}}$ efficiency for both 1- (red) and 3-prong (blue) $\tau_{\text{had-vis}}$. The performance of the BDT version of the tau identification algorithm (dashed lines) can be compared to that of the RNN (solid lines). Image from [69].

6.3 Electron veto

Aside from quark- and gluon-initiated jets, $\tau_{\text{had-vis}}$ candidates can also overlap with electron candidates. In particular, 1-prong taus are composed of a single track matched to a calorimeter deposit, a signature identical to that of electrons. In order to distinguish between the two, a BDT is trained [45] with 1-prong $\tau_{\text{had-vis}}$ from a $\gamma^* \to \tau\tau$ sample and electrons from a $\gamma^* \to ee$ sample. The $\tau_{\text{had-vis}}$ candidates are required to satisfy the medium identification working point, have $p_{\text{T}} > 20$ GeV and be within $|\eta| < 2.5$. The training is performed separately in the barrel and the end-cap regions of the calorimeters, with three regions defined as follows: the barrel with $0 < |\eta| < 1.37$, endcap-1 with $1.37 < |\eta| < 2$ and endcap-2 with $2 < |\eta| < 2.5$[2]. The p_{T}-spectrum of the electron sample is re-weighted so that it matches that of the $\gamma^* \to \tau\tau$ one. The variables used include how well the leading track fits the electron hypothesis within the TRT, the shape and location of the energy deposits in the calorimeter, the location of the track with respect to the shower centre of the calorimeter cluster, and the ratio of the energy deposited in the calorimeters with respect to the track momentum. The output score is transformed this time with respect to the p_{T} and the $|\eta|$ of the object, ensuring constant efficiencies as a function of both variables. Three working points are defined, with $\tau_{\text{had-vis}}$ efficiencies of 95%, 85% and 75% for the loose, medium and tight working points respectively. The associated rejection factors are of 30-100, 50-500 and 100-1000 respectively depending on the p_{T} and η of the candidate.

6.4 Muon veto

Occasionally muons traveling through the detector will undergo Bremsstrahlung and therefore result in a track that has an associated calorimeter deposit. If this radiation is significant enough, the track in the

[2]This selection loosely follows the locations of the end of the EM barrel and the end of the hadronic extended barrel, respectively.

MS will not be matched to the one in the ID. This characteristic can be used to distinguish $\tau_{\text{had-vis}}$ from these types of muons. If a standalone MS muon candidate is identified, the $\tau_{\text{had-vis}}$ candidate matched to it is removed.

6.5 Tau decay mode classification

Before the energy of the $\tau_{\text{had-vis}}$ candidate can be properly calibrated, it is important to know its charged and neutral hadron composition. While charged hadron counting can be performed in the reconstruction step by counting the number of charged tracks in the seed jet, neutral hadron counting requires an additional step.

Tau decay mode classification [48] uses an algorithm called PanTau, also referred to as τ particle flow due to the similarities to its jet counter-part. First, tracks are used to identify charged pion deposits in the calorimeters. The energy deposition in the hadronic calorimeter and ID is used to estimate the energy deposited by charged pions in the EM calorimeter. This energy is subtracted from the matching EM calorimeter cluster. Remaining clusters are used to identify neutral particles, with information from the first EM calorimeter layer being particularly useful to distinguish the characteristic neutral pion signature of two nearby photons.

Decay modes are classified according to the number of charged pions, determined via track-counting, the number of neutral pions, determined using PanTau, and any additional neutral objects not tagged as neutral pions. The $\tau_{\text{had-vis}}$ candidates may have one or three associated charged pions (prongs, p). They may also have 0, 1, or more neutral pions (n) associated to them, and 0 or more additional neutral particles (add). This results in the following categories: 1p0n + 0 add, 1p0n + X add, 1p1n + 0 add, 1p1n + X add, 1pXn + X add, 3p0n + 0 add, 3p0n + X add, 3pXn + X add. A set of BDTs is used to distinguish between some of the ambiguous categories: the first BDT is trained on 1p0n + X vs. 1p1n + 0 add, a second is trained on 1p1n + X vs. 1pXn + X, and a third on 3p0n + X vs. 3pXn + X. Variables related to decay kinematics, neutral pion identification from PanTau and number of photons contained within the $\tau_{\text{had-vis}}$ are used in all BDTs. The output of the BDTs is both a decay mode classification and a preliminary 4-vector including the charged and neutral pion components. The diagonal efficiency, a measure of the efficiency to reconstruct all categories correctly, is 73%. Individual category efficiencies range between 40% (1pXn) and 90% (3p0n).

6.6 Energy calibration

An energy calibration is needed to correct the energy measured in the detector to the average amount of energy carried by the visible decay products at generator level. Two calibrations are performed:

- The calorimeter-based calibration aims to improve upon the LC calibration. In particular, LC was derived based on typical pion densities inside QCD jets, and isn't always appropriate for hadronically-decaying τ leptons. Corrections for pile-up effects are also not included in the LC calibration and need to be accounted for. The calorimeter-based calibration relies on comparisons of truth and reconstructed energies of $\tau_{\text{had-vis}}$ in Monte Carlo simulation (done separately for 1- and 3-prong $\tau_{\text{had-vis}}$) and affects the determination of both the energy and the position in η of the $\tau_{\text{had-vis}}$.

- The second calibration is similar to what is used for particle flow jets and is based on PanTau. Track momentum measurements are used for charged hadrons while LC clusters are used for the energy measurement of neutral hadrons. It is important to note that for high energy $\tau_{\text{had-vis}}$ associating tracks to the correct cluster becomes difficult since the clusters get collimated, and the quality of the track momentum measurement decreases. An MVA algorithm is used to determine whether the calorimeter-based or the PanTau-based calibration should be used, based on which one is most reliable or has the best resolution.

6.6.1 Calorimeter-based calibration

The calorimeter-based calibration consists of a pile-up correction ($E_{\text{pile-up}}$) to the LC-scale energy measurement (E_{LC}) as well as a correction for the detector response ($\mathcal{R}$) to $\tau_{\text{had-vis}}$ [51]. It follows

$$E_{\text{calib}} = \frac{E_{\text{LC}} - E_{\text{pile-up}}}{\mathcal{R}(E_{\text{LC}} - E_{\text{pile-up}}, |\eta|, n_P)}. \tag{6.1}$$

The first correction, $E_{\text{pile-up}}$, is derived from simulation through fitting E_{LC} as a function of the difference between the number of primary vertices in the event (N_{PV}) and the average number of primary vertices in the sample ($\langle N_{\text{PV}} \rangle$). This correction vanishes for events with N_{PV} equal to that of the simulated sample. The fit uses a linear coefficient A, and is performed in bins of η separately for 1- and 3-prong $\tau_{\text{had-vis}}$ candidates:

$$E_{\text{pile-up}}(N_{\text{PV}}, |\eta|, n_P) = A(|\eta|, n_P) \times (N_{\text{PV}} - \langle N_{\text{PV}} \rangle). \tag{6.2}$$

The detector response correction, $\mathcal{R}$ is extracted as the Gaussian mean of the $(E_{\text{LC}} - E_{\text{pile-up}})/E_{\text{true}}^{\text{vis}}$ distribution in MC, where $E_{\text{true}}^{\text{vis}}$ is the energy of the visible decay products of the $\tau_{\text{had-vis}}$ including final state radiation.

6.6.2 MVA-based calibration

The more accurate measurement of the charged pion energy that can be achieved in certain cases via the use of the ID track momentum in PanTau can be exploited to improve upon the calorimeter-based calibration.

A Boosted Regression Tree (BRT) [12] uses information from the local calibration and the PanTau object as well as additional calorimeter and tracking information to estimate the true energy of the visible τ decay products.

The BRT is trained on a sample of $\gamma^* \to \tau\tau$ events, with $\tau_{\text{had-vis}}$ candidates satisfying the medium BDT identification requirement. Two metrics are used to determine the quality of the BRT output: the resolution, defined as half of the 68% interval of the ratio of the calibrated $\tau_{\text{had-vis}}$ p_{T} to its true visible p_{T}, and the non-closure which is the offset of the most probable value of that ratio from unity [51]. Since the non-closure can be easily corrected for as needed, only the resolution is a limitation to measurements that rely on knowledge of the energy of the $\tau_{\text{had-vis}}$.

Since the resolution of the PanTau momentum measurement is better at low p_{T} and that of the corrected LC momentum measurement is better at high p_{T}, an interpolated transverse momentum is defined:

$$p_{\text{T}}^{\text{interp}} = f_x \times p_{\text{T}}^{\text{LC}} + (1 - f_x) \times p_{\text{T}}^{\text{PanTau}} \tag{6.3}$$

where f_x is a weight between 1 and 0, defined as

$$f_x(p_{\mathrm{T}}^{\mathrm{LC}}) = \frac{1}{2}\left(1 + \tanh\frac{p_{\mathrm{T}}^{\mathrm{LC}} - x}{20}\right) \tag{6.4}$$

and x is the transition point of regimes from low to high p_{T}, set at 250 GeV.

The ratios $p_{\mathrm{T}}^{\mathrm{LC}}/p_{\mathrm{T}}^{\mathrm{interp}}$ and $p_{\mathrm{T}}^{\mathrm{PanTau}}/p_{\mathrm{T}}^{\mathrm{interp}}$ are used as inputs to the BRT to provide knowledge about the $\tau_{\text{had-vis}}$ energy. The moments of the topo-clusters are used as well. The μ and N_{PV} of the event inform the BRT about the other interactions occurring in the bunch-crossing, and finally the difference in neutral-to-charged pion energies and the number of neutral pions provide additional information about the decay mode of the $\tau_{\text{had-vis}}$.

The energy resolution of the BRT exceeds that of the calorimeter-based calibration below 250 GeV, as shown in Figure 6.4, but degrades at high transverse momentum, as discussed earlier, because the track measurements deteriorate at higher p_{T} and also because matching clusters to tracks becomes more difficult as the boost of the system increases.

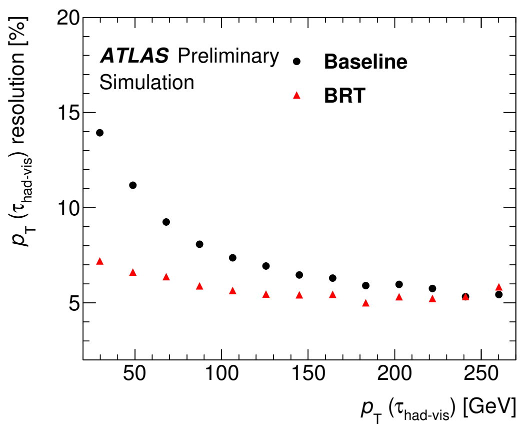

Figure 6.4: Resolution of the BRT-based and calorimeter-based (baseline) energy calibrations. Image from [51].

6.7 Data-Monte Carlo comparisons

A correction factor to account for the difference in $\tau_{\text{had-vis}}$ identification efficiencies in data and Monte Carlo is measured using a tag-and-probe method in a sample enriched in $Z \to \tau_\mu \tau_{\text{had-vis}}$ events. The relative precision of this measurement is of 5% (6%) for 1-(3-)prong $\tau_{\text{had-vis}}$ candidates using the partial Run 2 dataset [45] and is reduced to 2-3% using the full Run 2 dataset. Similar corrections are also derived for the electron veto selection. The relative precision of this measurement varies between 3% and 14% depending on the η of the $\tau_{\text{had-vis}}$ candidate [45]. Finally, the correction factors to the tau energy scale are evaluated with a relative precision of 2% (3%) for 1-(3-) prong $\tau_{\text{had-vis}}$ [45].

Chapter 7

The Higgs boson at the LHC

The search for the Higgs boson was one of the main incentives behind the upgrade of the CERN accelerator complex from the electron-positron LEP collider to the LHC (see Chapter 3). It also informed many of the design choices for the ATLAS detector, as discussed in Chapter 4, and dictated the development of reconstruction, identification and calibration algorithms presented in Chapters 5 and 6, ensuring that the searches could be carried out with the required level of precision and accuracy.

These efforts were rewarded during the first run of the LHC, which provided many physics results. Of particular importance was the announcement from the ATLAS and CMS collaborations in 2012 of the discovery of the Higgs boson, with a mass of around 125 GeV, using proton-proton collision data at center-of-mass energies of 7 and 8 TeV [25, 75]. After the long shutdown of the LHC from 2013 to 2015, a second run started in June 2015 this time at a higher center-of-mass energy of 13 TeV. During Run 2, the ATLAS collaboration continued studying the Higgs processes to improve upon the measurements of Run 1, increase sensitivity to rare processes and search for new physics. The focus in Run 2 was on to the study of the properties of the Higgs boson, such as its coupling to other fundamental particles.

7.1 Production modes

At the LHC, Higgs bosons can be produced in proton-proton collisions in a variety of ways. The production proceeds mainly through processes involving massive particles with large couplings to the Higgs boson, such as the weak vector bosons and the top (and bottom) quark. The main mechanisms for SM Higgs production at the LHC are shown in Figure 7.1. They consist of gluon fusion (ggH), vector boson fusion (VBF), associated production with a W or Z boson (VH) and associated production with a top or bottom quark pair ($t\bar{t}H/b\bar{b}H$). Rare Higgs production mechanisms exhibit much smaller cross sections and include among others the associated production of a Higgs boson and a top quark (tH).

The Higgs production rate, or cross section, for each production mode is entirely determined when the masses of all the particles involved are known. Figure 7.2 shows these production cross sections around a Higgs mass of 125 GeV.

Each production mode results in characteristic decay product kinematics and topologies, as illustrated in Figure 7.1. When analyzing LHC data, these characteristics can be exploited to identify events containing a Higgs boson, and also to determine the way in which the Higgs was produced. Gluon fusion events are generally associated to final states with low object multiplicity. VBF events, on the other

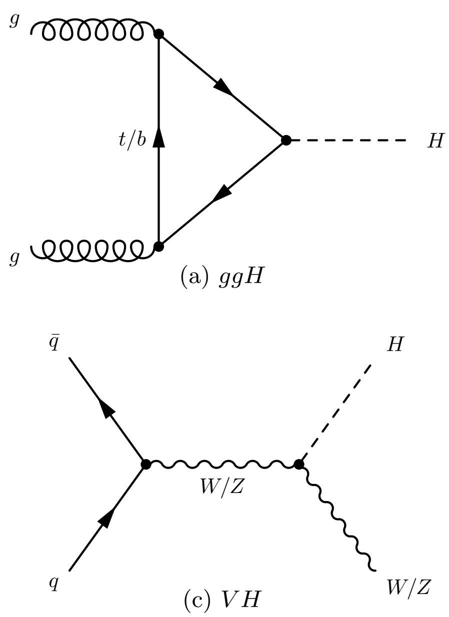

Figure 7.1: Feynman diagrams showing the main production mechanisms of Higgs bosons at the LHC.

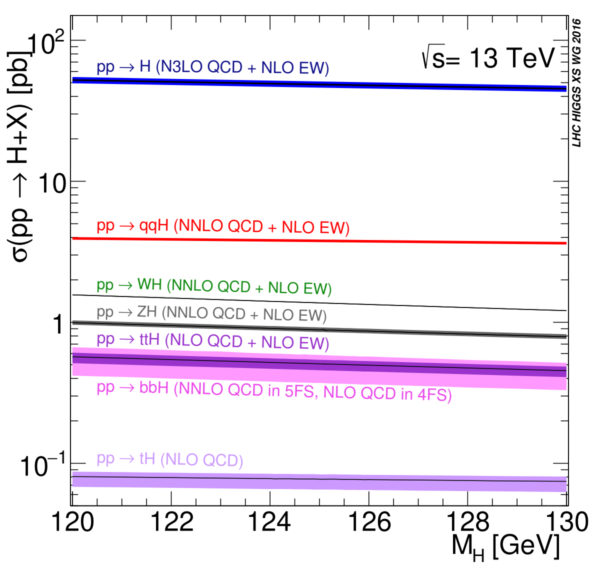

Figure 7.2: Higgs production rates around $m_H = 125$ GeV, where $pp \to H$ represents the ggH process (blue), $pp \to qqH$ represents the VBF process (red), the VH process is split into $pp \to WH$ and $pp \to ZH$ (green and grey) and $t\bar{t}H/b\bar{b}H$ are shown in purple and pink. The smaller contribution from the tH process is also shown. Image from [83].

hand, will typically contain two high-energy jets from the quarks produced alongside the Higgs boson. These jets are typically produced in the forward direction, with small angles relative to the beamline. In VH events, leptons or jets from the gauge boson decay are expected. Finally, $t\bar{t}H$ and $b\bar{b}H$ events will characteristically contain two b-quarks. In the case of $t\bar{t}H$, W bosons from the $t \to Wb$ decays contribute additional quarks or leptons to the final state.

7.2 Decay modes

As for Higgs production rates, the mass of the Higgs boson and its coupling to other particles determine the rate and branching ratios of its decay. In particular, decays to heavier particles such as the τ lepton are enhanced while decays to lighter particles like muons are much less likely to occur, as illustrated in Figure 7.3 and discussed earlier in Chapter 2. Higgs bosons can also decay to massless particles through loops, analogously to the process shown in Figure 7.1a.

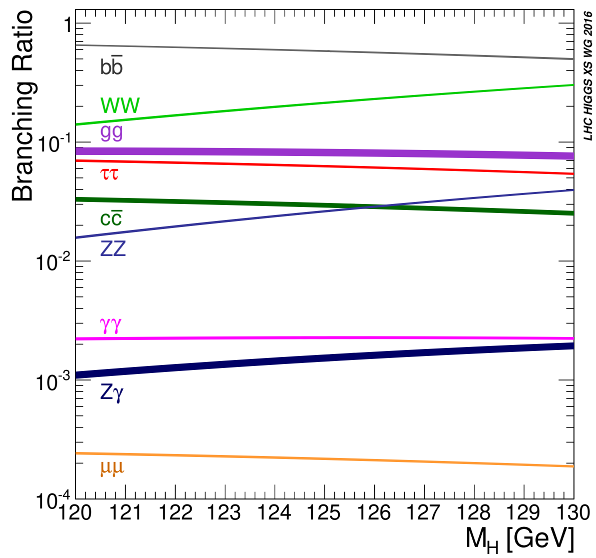

Figure 7.3: Higgs decay rates around $m_H = 125$ GeV. The $H \to \tau\tau$ decay, of particular interest for this thesis, is shown with a red line. Image from [83].

While the decays of Higgs bosons to light leptons (either directly such as $H \to \ell\ell$ or through processes such as $H \to ZZ \to 4\ell$) are subject to small branching fractions, light leptons are in general easier to identify and calibrate within the ATLAS detector as outlined in Chapter 5. Conversely, the hadronic activity at the LHC is usually quite high and makes the identification of hadronic decays of the Higgs much more difficult.

7.3 Run 1 results from ATLAS and CMS

Several measurements were made in Run 1 to characterize the newly-discovered Higgs boson. The first is the determination of its mass, with $m_H = 125.09 \pm 0.21$ (stat.) ± 0.11 (syst.) GeV measured using the ATLAS and CMS detectors in $H \to ZZ \to 4\ell$ and $H \to \gamma\gamma$ channels [5]. Upper limits were set for the total width of the Higgs boson relative to its SM prediction, with an observed value of 1.7 GeV and 4.5-7.5 GeV using the CMS and ATLAS detectors, respectively [78, 114]. The spin (J) and parity (P) of the Higgs boson were studied using events where the Higgs decays to other bosons. Hypotheses such as $J^P = 0^-$ or 2^+ were excluded in favour of the SM hypothesis of 0^+ [43]. All characteristics of the Higgs boson are similarly compared to the theoretical predictions.

The final results are the Higgs boson production cross section and decay branching ratio measurements made by the ATLAS and CMS collaborations in Run 1. They are presented in Figure 7.4. The rates of these processes can be predicted within the SM and therefore these measurements can serve as a way to test whether the particle discovered is the Higgs boson described by the SM, and to assess whether any physics beyond the SM are at play.

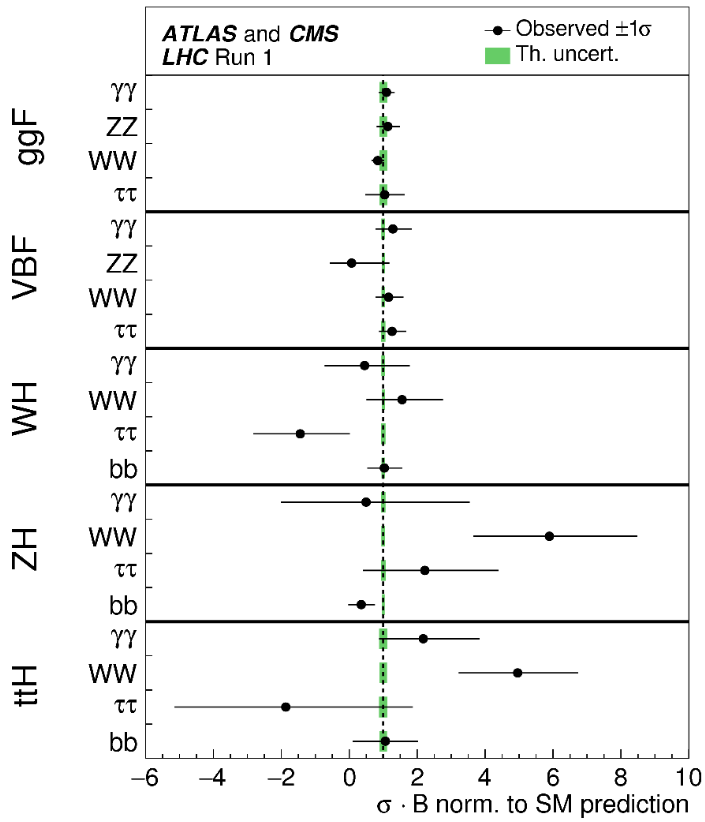

Figure 7.4: Deviation of measurements of Higgs boson production cross sections and branching ratios from theoretical predictions. Image from [5].

The various measurements of the Higgs production cross sections and decay branching ratios can be combined and interpreted as measurements of the couplings of various particles to the Higgs boson. As outlined in Chapter 2, the expected behaviour is that the couplings are proportional to m_F for fermions (F). For the weak vector bosons (V), the couplings are instead expected to follow m_V^2. Figure 7.5 illustrates the expected relationship between a particle's mass and its coupling to the Higgs boson, along with the experimentally-measured values. The results are consistent with the SM expectations in all cases within uncertainties. Since the Higgs boson was discovered in the $H \rightarrow WW$, $H \rightarrow \gamma\gamma$ and $H \rightarrow ZZ \rightarrow 4\ell$ final states given that they have large branching ratios and/or that they can be identified with relative ease. Consequently, at the end of Run 1, the couplings of the Higgs boson to the W and Z bosons were well measured. As discussed, larger uncertainties remained however on the measurement of the coupling to fermions. Increasing the precision of these measurements was one of the main physics goals of the second run of the LHC as it would help to identify any deviations of the measurements from the SM expectations.

The cross section of the ggH process is sensitive to the Yukawa coupling to the top quark since it proceeds predominantly through a top loop. This is why a relatively precise measurement of this coupling can be seen in Figure 7.5. However, this measurement relies on assumptions regarding the

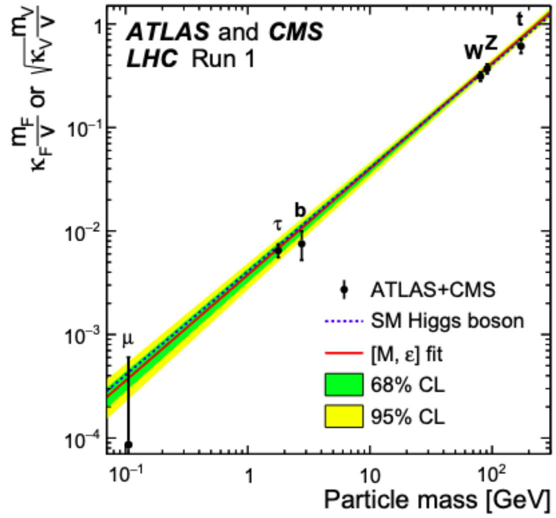

Figure 7.5: Reduced coupling strength scale factors $y_F = \kappa_F \frac{m_F}{v}$ for fermions ($F = t,\ b,\ \tau,\ \mu$) and $y_V = \sqrt{\kappa_V}\frac{m_V}{v}$ for weak gauge bosons ($V = W,\ Z$) as a function of their masses m_F and m_V, the VEV ($v = 246$ GeV) and the reduced coupling strength modifiers[1] (κ_F and κ_V) for the Run 1 dataset using the ATLAS and CMS detectors. The SM prediction is also shown with a dotted blue line. Image from [112].

other contributions to the loop. The measurement of the $t\bar{t}H$ cross section provides a complementary measurement that is independent of these assumptions. A comparison of the rates of the $t\bar{t}H$ and ggH processes therefore acts as a probe for new physics within the ggH loop. The ratio of the two is shown in the middle of Figure 7.6. The large uncertainties, due in great part to the lack of a precise $t\bar{t}H$ cross section measurement, highlight the difficulty in resolving the physics contributing to this loop using the Run 1 dataset.

In this thesis, a particular emphasis is put on the measurement of the coupling of the Higgs boson to the heaviest quark (top) and the heaviest lepton (τ) through the measurement of the $t\bar{t}H$ production cross section in events with τ leptons (see Chapters 8) as well as measurements of the $H \to \tau\tau$ decay branching ratio in events where the Higgs is produced via other modes (see Chapter 9).

Since the existence of neither the $t\bar{t}H$ nor the $H \to \tau\tau$ processes was conclusively established in single-experiment results using the Run 1 dataset, the study of these processes was particularly interesting for Run 2.

[1] The coupling strength modifiers κ_F and κ_V are measured assuming no BSM contributions to the Higgs boson decays, and assuming the SM structure of loop processes.

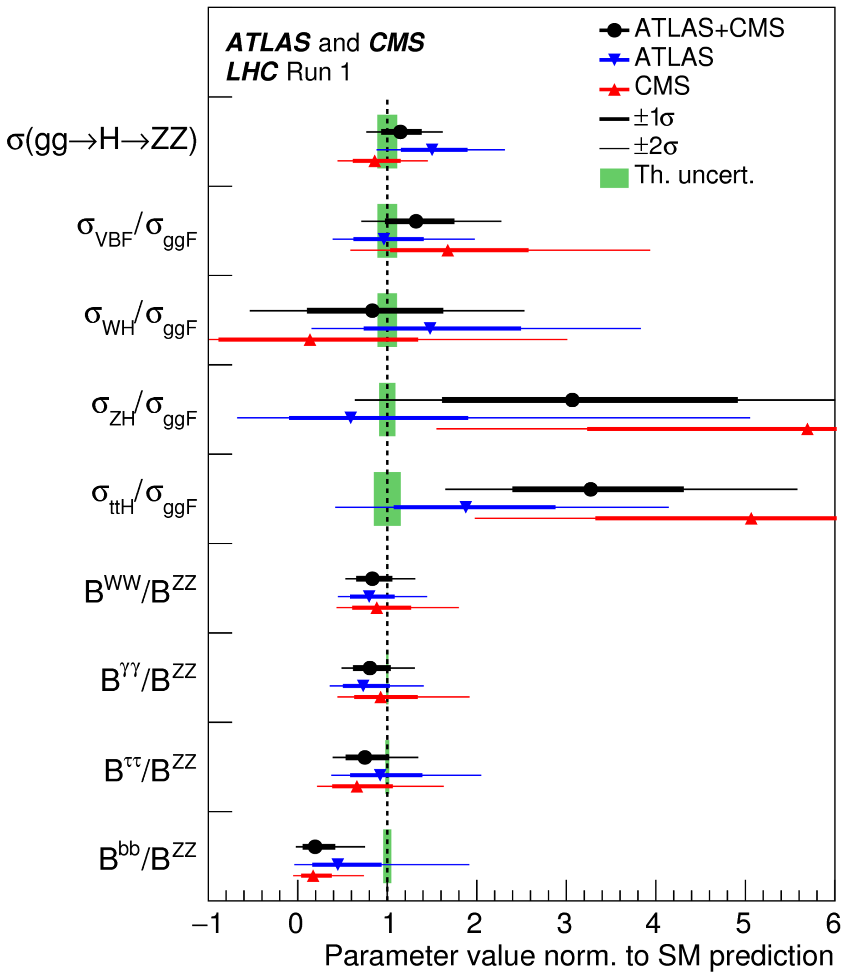

Figure 7.6: Best fit values of the $gg \to H \to ZZ$ cross section and of the ratios of cross sections (top of the image) and branching fractions (bottom of the image). Image from [5].

Chapter 8

Observation of the associated production of a Higgs boson and a top quark pair ($t\bar{t}H$)

8.1 Introduction

In the SM, the top quark has a large mass and therefore a strong coupling to the Higgs field. Measuring the top quark Yukawa coupling (Y_t) is one of the highest priorities in the physics plan for the second run of the LHC. One of the ways to extract the top quark Yukawa coupling is through the ggH production process, as it proceeds predominantly through a top quark loop in the SM. In order to extract Y_t values from this measurement, the other contributions to the loop are assumed to be those of the SM. The Y_t values obtained from ggH cross section measurements rely on the assumptions that all other Yukawa couplings are as described by the SM and that no BSM processes are involved in the loop. Therefore, while the Y_t value extracted from ggH measurements is found to be compatible with the SM, new physics hidden within the ggH loop could be biasing this result. The measurement of the cross section of the associated production of a top anti-top quark pair and a Higgs boson ($t\bar{t}H$) is an important milestone for the measurement of Y_t as it provides a complementary tree-level measurement that does not rely on these assumptions. Its comparison to measurements such as the one from ggH, for example through a ratio of the ggH to $t\bar{t}H$ cross sections, helps probe for new physics processes that may contribute to the ggH loop.

The $t\bar{t}H$ cross section is much lower than that of the ggH, VBF and VH production modes discussed in Chapter 7: it represents only about 1% of the total Higgs production. However, the distinctive final states of $t\bar{t}H$ can be used to effectively identify the relevant events within the dataset.

8.1.1 Final states

In order to measure the $t\bar{t}H$ cross section, it is necessary to collect as many $t\bar{t}H$ events as possible. For this, many potential final states must be studied. Decays of both the top quarks and the Higgs boson need to be considered. The $t\bar{t}$ system decays proceed predominantly through $t\bar{t} \to W^+W^-b\bar{b}$, with the W bosons then decaying to either an electrically-charged lepton (referred to simply as lepton in the rest of this chapter) and a neutrino, or a quark anti-quark pair. As discussed, the characteristic presence of b-quarks from the top quark decays can be used to identify events where a Higgs boson is produced via the $t\bar{t}H$ production mode. Different final states are classified according to the decays of the W and Higgs bosons. Three Higgs boson decays result in final states that can be efficiently identified within the Run 2 dataset: Higgs decays to two photons, to two b-quarks, and Higgs decays resulting in the production of leptons [58]. Analysis strategies are developed separately for each final state to achieve maximum sensitivity and to enable proper treatment of the relevant background processes and systematic uncertainties. This chapter will focus on the $t\bar{t}H \rightarrow$ multilepton analysis, which targets Higgs decays to WW, ZZ and $\tau\tau$ and requires that the Higgs final state contain at least one lepton[1], with at least one other lepton originating from the top quark decays through a leptonic decay of the W boson. This chapter presents the measurements performed in $t\bar{t}H \rightarrow$ multilepton during Run 2 with 36 fb^{-1} of data collected during the years 2015 and 2016, with some discussion of results obtained in combination with other $t\bar{t}H$ analyses (some of which include additional data collected in 2017) at the end.

8.1.2 Previous measurements at the LHC

During Run 1, the ATLAS and CMS collaborations performed $t\bar{t}H$ searches in the $t\bar{t}H \rightarrow$ multilepton, $t\bar{t}(H \rightarrow b\bar{b})$ and $t\bar{t}(H \rightarrow \gamma\gamma)$ final states. Since the dataset was relatively small at 20 fb^{-1} and the $t\bar{t}H$ production cross section $\sigma_{t\bar{t}H}$ is also expected to be small, very few $t\bar{t}H$ events were available within the dataset. It was therefore not possible to precisely measure $\sigma_{t\bar{t}H}$. Instead, the goal was two-fold: to determine whether there was evidence that the $t\bar{t}H$ process did indeed proceed and, if so, to determine how compatible its rate was with the SM expectation. In a combination with CMS results, the best-fit result of the ratio of the observed $t\bar{t}H$ production rate to the SM expectation ($\mu_{t\bar{t}H} = \frac{\sigma_{t\bar{t}H}\ \text{obs.}}{\sigma_{t\bar{t}H}\ \text{SM}}$) was found to be $2.3^{+0.7}_{-0.6}$ [111]. The individual results were $\mu_{t\bar{t}H} = 1.81 \pm 0.80$ for ATLAS [116] and $\mu_{t\bar{t}H} = 2.8 \pm 1.0$ for CMS [120]. In the multilepton final state, the ATLAS result was $\mu_{t\bar{t}H} = 2.1^{+1.4}_{-1.2}$ [115]. None of these results were able to establish the existence of the $t\bar{t}H$ process since they were also consistent with fluctuations of background processes.

The SM $t\bar{t}H$ cross section increases by a factor of 3.9 when the centre of mass energy of collisions increases from 8 TeV to 13 TeV. This implies that $t\bar{t}H$ signal events can be accumulated at a faster rate during Run 2 than they were during Run 1.

8.2 Analysis strategy

The Run 2 analysis aims to build upon the Run 1 results and establish a measurement of $\mu_{t\bar{t}H}$ that is inconsistent with background fluctuations at the level of 1 in 3.5 million. Since the focus of this thesis is on $H \rightarrow \tau\tau$ decays, only the analysis strategy particular to $t\bar{t}H \rightarrow$ multilepton will be discussed. This strategy focuses on the categorisation of possible final states of the $t\bar{t}$ and Higgs boson decays into channels with different light lepton (e, μ) multiplicities, $\tau_{\text{had-vis}}$ multiplicities and lepton charge requirements. The channels considered are shown in Figure 8.1. Final states are labelled using the objects that are searched for. A $3\ell + 1\tau_{\text{had}}$ channel therefore refers to a final state where three light leptons and one $\tau_{\text{had-vis}}$ are required.

[1]The only exception consists in Higgs decays to two Z bosons that in turn decay to four leptons. These decays result in very particular signatures in which the Higgs boson can be fully reconstructed, and the analysis strategy closely follows the one outlined for $H \rightarrow ZZ^* \rightarrow 4\ell$ measurements in other production modes and are not considered in the $t\bar{t}H \rightarrow$ multilepton analysis [59].

In general, $t\bar{t}H$ events with $H \to WW$ or $H \to \tau\tau$ decays are not distinguishable. Taking for example the process $H \to WW \to \tau\tau\nu\nu$, the final states in both cases contain only two τ as visible products. This is why it is in general not possible to reconstruct final states that specifically select one of these Higgs decay modes. However, channels without $\tau_{\text{had-vis}}$ are typically more sensitive to $H \to WW$ decays while channels with $\tau_{\text{had-vis}}$ are sensitive to $H \to \tau\tau$ decays as illustrated in Figure 8.2. Analysis of both types of channels provides some sensitivity to Higgs decay branching ratios within $t\bar{t}H \to$ multilepton.

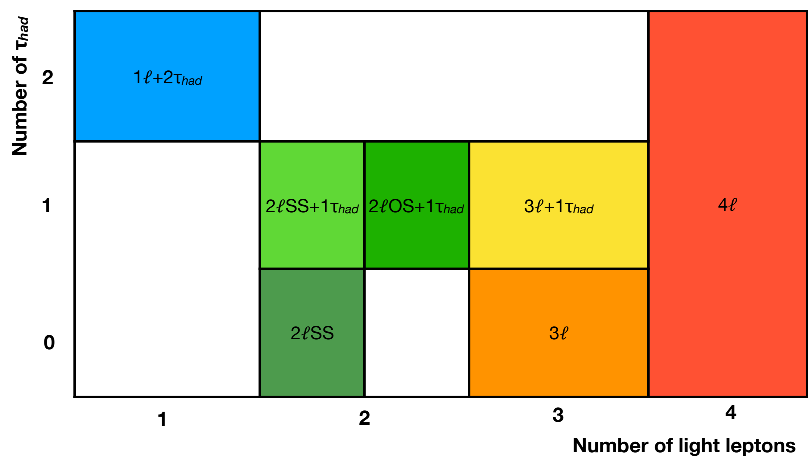

Figure 8.1: Channels of the $t\bar{t}H \to$ multilepton analysis, classified according to the light lepton and $\tau_{\text{had-vis}}$ multiplicity of the final state. Only events with same-sign light leptons or at least 3 total leptons are considered. Image from [58].

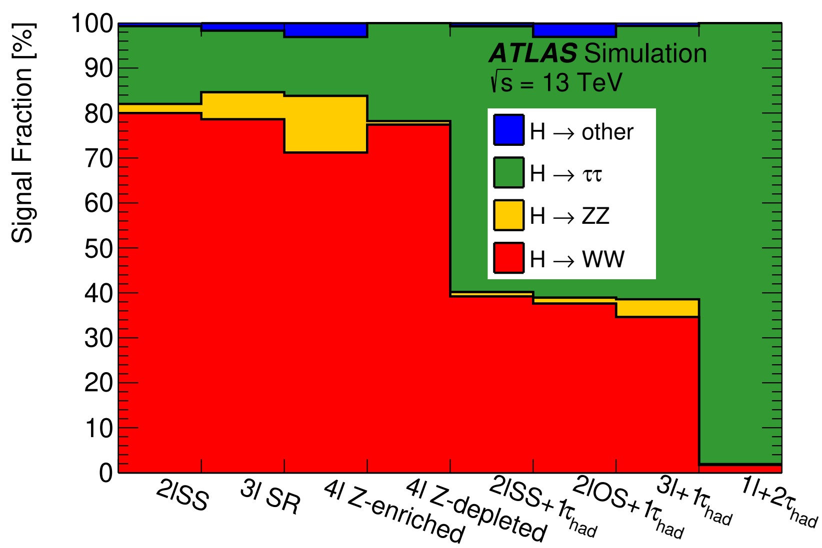

Figure 8.2: Fraction of events in each $t\bar{t}H \to$ multilepton channel with a Higgs boson decaying to WW (red), ZZ (yellow), $\tau\tau$ (green) and other decay modes such as $\mu\mu$ and $b\bar{b}$ (blue). Image from [58].

Other processes such as the production of a top quark pair ($t\bar{t}$), the associated production of a top quark pair and a vector boson ($t\bar{t}V$) and diboson production (VV) can result in final states similar to that of $t\bar{t}H \to$ multilepton and appear as backgrounds to the $t\bar{t}H$ measurements. Background processes can be separated into two categories. Irreducible backgrounds produce final states that are identical to or indistinguishable from those of the signal process. To suppress those contributions, only channels with high lepton multiplicity or satisfying charge requirements are chosen. Additional requirements are detailed in Section 8.4 and aim to reduce background rates by ensuring that each channel selects a final state similar to that of signal events, by enforcing selections on quantities such as jet multiplicity, invariant mass of the dilepton systems, etc. Reducible backgrounds on the other hand are those whose contribution may be significantly reduced via additional selections on the quality of the reconstructed objects. Specifically, background processes with final states that are expected to be different to the one studied, but who suffer from contributions from non-prompt and fake objects[2] that make the final state similar to that of the signal, are considered reducible backgrounds. In these cases, a more stringent object selection is expected to improve the signal-to-background ratio. In some cases, a process may be considered as a reducible or an irreducible background depending on the final state produced, with the two contributions considered separately. For example, $t\bar{t}Z$ events may enter the $3\ell + 1\tau_{\text{had}}$ channel with either a real or a fake $\tau_{\text{had-vis}}$: events with real $\tau_{\text{had-vis}}$ are considered part of the irreducible background while events with a fake $\tau_{\text{had-vis}}$ belong to the reducible background category.

In the $t\bar{t}H \to$ multilepton analysis, irreducible backgrounds are estimated with Monte Carlo simulation while reducible backgrounds are usually estimated using data-driven methods. This is in part because the types of event involved in this contribution tend to be rare, and therefore generating MC in the phase space of interest is not efficient. The other reason is that simulation of these backgrounds relies heavily on the description of the hadronisation process and interactions with the detector within the MC simulation. Any mismodellings observed in these descriptions would result in a significant misestimation of reducible background rates and kinematics. In other words, "how many times does a process X produce particle Y?" is an easier question to answer than "how many times does particle Y have specific characteristics that make it look like particle Z in the detector?" and thus the modelling of the former is in general more reliable than that of the latter.

The channels of the $t\bar{t}H \to$ multilepton analysis vary in terms of the statistics available (high in lower lepton multiplicity channels), in signal purity (high in high lepton multiplicity channels) and in background composition (very dependent on the final state, as illustrated in Figure 8.3). In high-statistics channels, the selections may be relaxed to enable the creation of sub-regions. These sub-regions correspond to signal regions (SR), where the signal process contribution is high, and control regions (CR), where the normalisation of the main background processes can be determined.

8.3 $t\bar{t}H \to$ multilepton in Run 2 with 36 fb^{-1}

Since $\tau_{\text{had-vis}}$ identification is more challenging than that of light leptons, only one channel with $\tau_{\text{had-vis}}$ was included in earlier versions of this analysis: $2\ell\text{SS} + 1\tau_{\text{had}}$, requiring two light leptons of same sign electric charge and one $\tau_{\text{had-vis}}$. In the second round of analysis, it was decided to invest time in understanding the $\tau_{\text{had-vis}}$ fake rates so more final states with $\tau_{\text{had-vis}}$ could be added. Two new channels, $2\ell\text{OS} + 1\tau_{\text{had}}$ (two light leptons of opposite sign electric charge and one $\tau_{\text{had-vis}}$) and $3\ell + 1\tau_{\text{had}}$ (three light leptons and one $\tau_{\text{had-vis}}$) were added to the existing channels. An example of a $t\bar{t}H$ event that results in the $3\ell + 1\tau_{\text{had}}$ final state is shown in Figure 8.4. The focus of this chapter will be on the

[2]Objects originating from secondary decays (e.g. an electron originating from a semi-leptonic b decay) are labelled non-prompt, while misidentified objects (such as a QCD jet that satisfies the $\tau_{\text{had-vis}}$ identification criteria) are referred to as fake.

$3\ell + 1\tau_{\text{had}}$ channel.

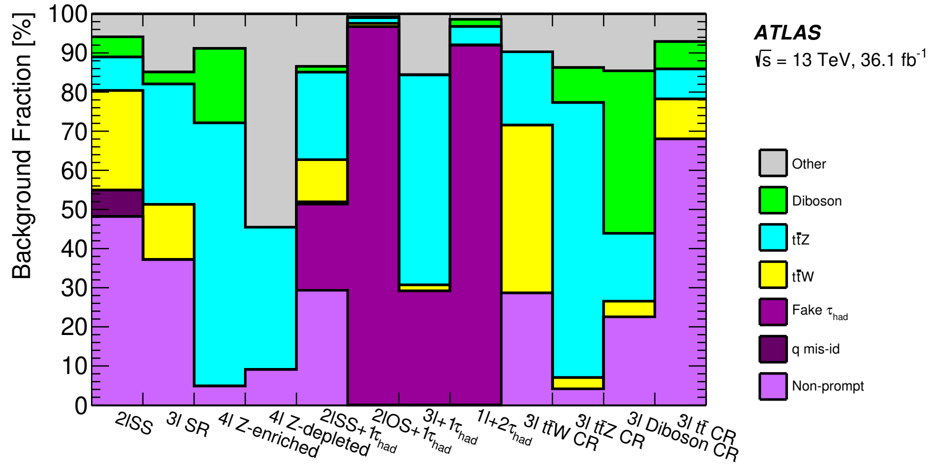

Figure 8.3: Fractional contribution of background processes to the twelve analysis categories, with all selections applied. Image from [58].

For existing channels, improvements possible due to the increase in statistics during Run 2 include the implementation of more complex signal and control region definitions that allow for better separation of signal and backgrounds through the use of multi-variate techniques[3], and more comprehensive non-prompt and fake object estimation methods.

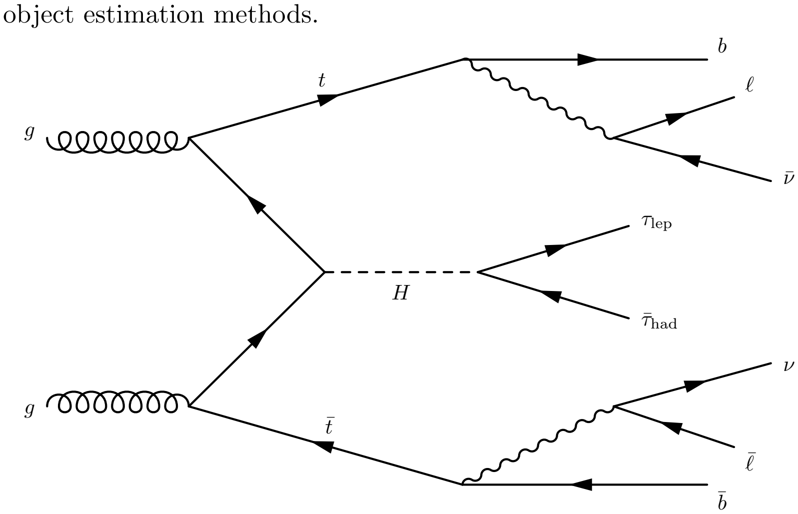

Figure 8.4: Feynman diagram of the $3\ell + 1\tau_{\text{had}}$ final state.

[3]These techniques are especially helpful for the classification of $t\bar{t}H \to$ multilepton final states since the $t \to Wb \to \ell\nu b$ decays result in the presence of neutrinos in all $t\bar{t}H \to$ multilepton final states. As discussed in Section 5.5, neutrinos cannot be reconstructed individually with the ATLAS detector and only the overall missing transverse momentum can be accounted for. While full event reconstruction is therefore not possible, MVAs can be used to gauge the kinematics and topology of the event in its entirety and provide better signal-to-background separation than would be otherwise available.

8.3.1 Data and Monte Carlo samples

The dataset used consists of data from 2015 and 2016 proton-proton collisions. After quality requirements are applied for the LHC beam conditions and the data, the integrated luminosity of the dataset is of 36.1 fb^{-1}.

The signal and many of the background contributions to each of the regions shown in Figure 8.1 are estimated using Monte Carlo simulation. The Monte Carlo samples used to model the relevant processes are produced using various generators, presented in detail within [58]. The ATLAS detector is simulated using the full Geant4 [2, 35] model, except for a few low-yield background processes where a parametrised model of the Geant4 configuration is used, with the tracking simulation taken from the full Geant4 model [36].

8.3.1.1 Triggers

Because this analysis focuses on final states with at least two light leptons, single and dilepton triggers are used to select data events, with the dilepton triggers providing additional efficiency to select events with low transverse momentum leptons. For data recorded in 2015, the single electron (single muon) trigger requires a candidate with transverse momentum of at least 24 (20) GeV while in 2016 the threshold is raised to 26 GeV. For dielectron (dimuon) triggers, the transverse momentum requirements are of 12 (17) GeV for both leptons. Finally, an electron-muon trigger is used with requirements of 17 GeV and 14 GeV respectively.

8.4 Object and event selections

The Monte Carlo samples describing the signal and background processes are used to choose the best object and event selections for each channel and maximise the sensitivity to the $t\bar{t}H$ process. While a very loose selection is used to define the light leptons and $\tau_{\text{had-vis}}$ objects that are used to categorize events into the seven analysis channels of Figure 8.1, this definition is later adjusted according to the needs of each channel. In particular, object selections are tuned to enhance signal acceptance and decrease acceptance of reducible backgrounds. Event selections are added to reject irreducible backgrounds.

8.4.1 Object selections

A variety of object selections are defined for this analysis, to adapt to the specific signal and background composition of each channel.

8.4.1.1 Light leptons

Baseline muon selection

All muon candidates are required to pass the loose identification selection described in Section 5.3. They must also satisfy $p_{\text{T}} > 10$ GeV and $|\eta| < 2.5$. To reduce contributions from non-prompt muons, the origin of the muon track is required to be compatible with the primary vertex of the event by applying selections on the significance of its transverse impact parameter $\left(\frac{|d_0|}{\sigma_{d_0}} < 3.0\right)$ and on its longitudinal impact parameter ($|z_0 \sin\theta| < 0.5$ mm)[4].

[4]Typical values for σ_{d_0} are 20 μm [57].

Baseline electron selection

Electron candidates are reconstructed using the baseline algorithm described in Section 5.2. All electron candidates are required to pass the loose likelihood identification working point, to distinguish them from contributions due to hadronic jets, photon conversions and heavy-flavour hadronic decays. Similarly to muons, a requirement of $p_{\mathrm{T}} > 10$ GeV is added. The cluster associated to the electron must also be outside of the transition region between the barrel and endcap EM calorimeters ($1.37 < |\eta_{\mathrm{cluster}}| < 1.52$) and satisfy $|\eta|_{\mathrm{cluster}} < 2.47$. To reduce contributions from non-prompt electrons, the longitudinal impact parameter selection that is applied to muons is also applied to electrons, while the transverse impact parameter selection is loosened to $\frac{|d_0|}{\sigma_{d_0}} < 5.0$.

Light lepton selection categories

Additional requirements may be added for analysis channels suffering from large non-prompt and fake background contributions.

Two isolation variables are created, using calorimeter and tracking variables. The first considers the scalar sum of transverse energies in a cone of size $\Delta R = 0.3$ around the selected light lepton. For electrons this cone excludes the reconstructed electron cluster itself, while for muons the cone excludes all clusters within $\Delta R = 0.1$ of the muon track. The sum is corrected for leakage from the electron's shower as well as for the ambient energy of the event. The track-based isolation considers the sum of transverse momenta of all tracks with $p_{\mathrm{T}} > 1$ GeV and consistent with originating from the primary vertex of the event, excluding the selected light lepton's track, and within a cone of $\Delta R = \min(0.3, 10\ \mathrm{GeV}/p_{\mathrm{T},\ell})$ around the light lepton track.

Additional rejection for non-prompt light leptons can be achieved through the use of a specifically-trained BDT that combines calorimeter and track information to identify non-prompt leptons originating from heavy-flavour decays. The output of the flavour-tagging algorithms described in Section 5.6 are also used as inputs. Its efficiency is of 70% (60%) for prompt muons (electrons) with a p_{T} of 10 GeV and goes up to 98% (96%) around 45 GeV. The rejection factor against leptons originating from b-hadron decays is 20 [58].

Since there is a small but non-zero probability for electrons to be reconstructed with the wrong electric charge, the charge misidentification BDT discussed in Section 5.2.3 is used in some channels.

The identification working point can also be adjusted, with the loose working point switched to a tight one.

The above requirements can be combined into 5 categories. Loose (L) electrons correspond to the baseline selection presented above. Loose and isolated ($\mathrm{L}^{\dagger}$) electrons are additionally required to satisfy the isolation requirements, while the last category (L^{*}) contains electrons that are loose, isolated and pass the non-prompt BDT. Tight electrons must pass the tight likelihood identification working point, be isolated, and pass the non-prompt BDT. Very tight electrons are also required to pass the charge misidentification BDT. For muons, only three categories exist. Loose (L) muons are those selected using the baseline selection. Loose and isolated ($\mathrm{L}^{\dagger}$) muons are additionally required to pass the isolation criteria. Finally, the last category consists of muons that also pass the non-prompt BDT ($\mathrm{L}^{*}/\mathrm{T}/\mathrm{T}^{*}$).

In general, the loose selection is used in channels where all light leptons are expected to be prompt, with isolation and non-prompt BDT requirements added for channels where non-prompt light leptons, in particular those originating from heavy-flavour jets, are expected. Tight categories are reserved for channels where fake electron contributions are expected, and very tight selections are used in channels that additionally suffer from backgrounds passing the selection due to the presence of electrons whose charge is misidentified. The category describes the selection applied to both electrons and muons: if loose is the light lepton working point chosen, both electrons and muons have to satisfy the corresponding requirements. When tight electron definitions are required, the muons are required to satisfy the L*/T/T* requirements.

<table>
<tr><th></th><th colspan="5">e</th><th colspan="3">μ</th></tr>
<tr><th></th><th>L</th><th>L†</th><th>L*</th><th>T</th><th>T*</th><th>L</th><th>L†</th><th>L*/T/T*</th></tr>
<tr><td>Requirement on p_{T}</td><td colspan="8">> 10 GeV</td></tr>
<tr><td>Requirement on $|\eta|$</td><td colspan="5">$1.37 < |\eta_{\mathrm{cluster}}| < 1.52$
and $|\eta_{\mathrm{cluster}}| < 2.47$</td><td colspan="3">$|\eta| < 2.5$</td></tr>
<tr><td>Passes isolation WP</td><td>No</td><td colspan="4">Yes</td><td>No</td><td colspan="2">Yes</td></tr>
<tr><td>Passes non-prompt lepton BDT WP</td><td colspan="2">No</td><td colspan="3">Yes</td><td colspan="2">No</td><td>Yes</td></tr>
<tr><td>Passes charge misidentification veto BDT</td><td colspan="4">No</td><td>Yes</td><td colspan="3">No</td></tr>
<tr><td>Identification WP</td><td colspan="3">Loose</td><td colspan="2">Tight</td><td colspan="3">Loose</td></tr>
<tr><td>Requirement on $|d_0|/\sigma_{d_0}$</td><td colspan="5">< 5</td><td colspan="3">< 3</td></tr>
<tr><td>Requirement on $|z_0 \sin\theta|$</td><td colspan="8">< 0.5 mm</td></tr>
</table>

Table 8.1: Requirements for the loose (L), loose and isolated (L†), loose, isolated an passing the non-prompt BDT (L*), tight (T) and very tight (T*) object definitions for electrons on the left and muons on the right. For muons, the L*, T and T* selections are identical.

8.4.1.2 Hadronically-decaying τ leptons

Hadronically-decaying τ are required to satisfy the loose or medium BDT identification working points (as described in Section 6.2) depending on the channel. The $\tau_{\text{had-vis}}$ is required to be matched to exactly one or three tracks so that its total charge is equal to ± 1, have $p_{\mathrm{T}} > 25$ GeV, be located within $|\eta| < 2.5$ and outside of the transition region ($1.37 < |\eta| < 1.52$), and pass the electron overlap removal BDT mentioned in Chapter 6. All $\tau_{\text{had-vis}}$ objects that overlap with a reconstructed muon candidate of $p_{\mathrm{T}} > 2$ GeV are removed. To reduce the possibility of jets originating from b-hadrons being misreconstructed as $\tau_{\text{had-vis}}$, all $\tau_{\text{had-vis}}$ candidates whose associated clusters pass the 70% identification working point of the MV2 algorithm introduced in Section 5.6 are removed. To decrease the contribution from pile-up jets reconstructed as $\tau_{\text{had-vis}}$, all selected $\tau_{\text{had-vis}}$ are required to be matched the the primary vertex of the event.

8.4.1.3 Jets and b-jets

Jets are reconstructed using the topo-clustering method presented in Section 5.4.1. They are required to satisfy $p_{\mathrm{T}} > 25$ GeV and $|\eta| < 2.5$. The JVT algorithm outlined in Section 5.4.3 is used to suppress contributions from jets originating from pile-up interactions. It is applied to jets with $p_{\mathrm{T}} < 60$ GeV and $|\eta| < 2.4$. Jets containing b-hadrons, referred to as b-jets, are identified using the MV2 algorithm presented in Section 5.6, and the working point used corresponds to an average efficiency of 70%. The expected rejection factors are of 380, 12 and 55 for jets from light quarks and gluons, for c quark jets and for $\tau_{\text{had-vis}}$, respectively.

8.4.1.4 Overlap removal

To ensure that no detector signal is used to reconstruct more than one object, an overlap removal procedure is used. If two electrons are reconstructed within $\Delta R = 0.1$, only the electron with higher p_{T} is kept. All electrons within $\Delta R = 0.1$ of a muon are removed. If a jet overlaps with an electron within $\Delta R = 0.3$, the jet is removed. Muons that overlap with remaining jets within $\Delta R = \min(0.4, 0.4 + \frac{10\ \mathrm{GeV}}{p_{\mathrm{T}}^{\mu}})$ are also removed, to reduce contributions from non-prompt muons. Then, $\tau_{\text{had-vis}}$ objects are removed if they overlap with a light lepton within $\Delta R = 0.2$. Lastly, jets overlapping with remaining $\tau_{\text{had-vis}}$ are removed if they are within $\Delta R = 0.3$. This procedure is summarized in Table 8.2.

Step	Reject	Against	Criteria
1	electron 2	electron 1	$\Delta R < 0.1$, $p_{\mathrm{T}}^{1} > p_{\mathrm{T}}^{2}$
2	electron	muon	$\Delta R < 0.1$
3	jet	electron	$\Delta R < 0.3$
4	muon	jet	$\Delta R < \min(0.4, 0.4 + \frac{10\mathrm{GeV}}{p_{\mathrm{T}}^{\mu}})$
5	$\tau_{\text{had-vis}}$	electron	$\Delta R < 0.2$
6	$\tau_{\text{had-vis}}$	muon	$\Delta R < 0.2$
7	jet	$\tau_{\text{had-vis}}$	$\Delta R < 0.3$

Table 8.2: Objects that are removed and kept during the overlap removal procedure.

8.4.2 Event selections

The event selections are chosen to reject events from irreducible backgrounds while maintaining high $t\bar{t}H$ efficiency. They are orthogonal among channels to ensure that events only contribute to one channel. The baseline object selections are used to split events into channels with different light lepton and $\tau_{\text{had-vis}}$ multiplicities, and channels without $\tau_{\text{had-vis}}$ are required to apply a $\tau_{\text{had-vis}}$ veto[5]. All events are required to contain at least one selected light lepton that is matched to one of the lepton triggers. All events must also contain at least 2 jets of which at least 1 must be b-tagged. Additional event selections are detailed in Table 8.3 for each of the 7 channels. The selections relevant to the $3\ell + 1\tau_{\text{had}}$ channel are described in more detail below.

8.4.2.1 $3\ell + 1\tau_{\text{had}}$

The events must contain exactly three reconstructed light leptons and one $\tau_{\text{had-vis}}$. The total charge of the light leptons must equal ± 1 while that of the $\tau_{\text{had-vis}}$ must be of ∓ 1.

The three light leptons are labelled with numbers, with 1 and 2 identifying the same-sign pair and 0 referring to the lepton of opposite sign. In MC lepton 0 is found to be prompt 97% of the time for both signal and background events, and therefore is only required to satisfy the L* criteria, i.e. pass a loose identification working point, be isolated and pass the non-prompt BDT selection. Due to higher likelihoods of being fake and non-prompt[6], leptons 1 and 2 are required to satisfy the T selection. Charge

[5] There is no $4\ell + 1\tau_{\text{had}}$ channel so no $\tau_{\text{had-vis}}$ veto is applied in the 4ℓ channel.

[6] In this channel, $t\bar{t}$ decays often result in a pair of opposite-sign light leptons leaving the third lepton with a higher likelihood of being fake or non-prompt.

Channel	Selection criteria
$2\ell SS + 0\tau_{had}$	Two T* light leptons with $p_T > 20$ GeV Same-charge light leptons Zero medium $\tau_{\text{had-vis}}$ $N_{jets} \geq 4$ and $N_{bjets} < 3$
$3\ell + 0\tau_{had}$	Three light leptons with sum of charges ± 1 The two same-charge leptons must be T* and have $p_T > 15$ GeV The opposite-charge lepton must be L* Zero medium $\tau_{\text{had-vis}}$ $m_{\ell^+\ell^-} > 12$ GeV and $\lvert m_{\ell^+\ell^-} - 91.2 \text{ GeV}\rvert > 10$ GeV for all SFOC pairs $\lvert m_{3\ell} - 91.2 \text{ GeV}\rvert > 10$ GeV
4ℓ	Four light leptons with sum of charges 0 Third and fourth leading leptons must be T $m_{\ell^+\ell^-} > 12$ GeV and $\lvert m_{\ell^+\ell^-} - 91.2 \text{ GeV}\rvert > 10$ GeV for all SFOC pairs $\lvert m_{4\ell} - 125 \text{ GeV}\rvert > 5$ GeV Split 2 categories: Z-depleted (0 SFOC pairs) and Z-enriched (1 or 2 SFOC pairs)
$2\ell SS + 1\tau_{had}$	Two T* light leptons with $p_T > 15$ GeV Same-charge light leptons One medium $\tau_{\text{had-vis}}$ of opposite charge to the light leptons $N_{jets} \geq 4$ $\lvert m_{ee} - 91.2\text{GeV}\rvert > 10$ GeV for ee events
$2\ell OS + 1\tau_{had}$	Two L† light leptons with $p_T > 25$, 15 GeV Opposite-charge light leptons One medium $\tau_{\text{had-vis}}$ $m_{\ell^+\ell^-} > 12$ GeV and $\lvert m_{\ell^+\ell^-} - 91.2 \text{ GeV}\rvert > 10$ GeV if SFOC pair $N_{jets} \geq 3$
$3\ell + 1\tau_{had}$	$3\ell + 0\tau_{had}$ selection, except: One medium $\tau_{\text{had-vis}}$ of opposite charge to the total charge of the light leptons The two same-charge light leptons must be T and have $p_T > 10$ GeV
$1\ell + 2\tau_{had}$	One T light lepton with $p_T > 27$ GeV Two medium $\tau_{\text{had-vis}}$ of opposite charge, with at least one being tight $N_{jets} \geq 3$

Table 8.3: Selection criteria applied in the 7 channels. "SFOC pairs" refer to pairs of light lepton with same flavor and opposite charge.

misidentification rates are small for the electrons in $3\ell + 1\tau_{had}$ events, so no requirement is placed on the charge misidentification score.

Events containing same-flavour and opposite-charge light lepton pairs with an invariant mass below 12 GeV are removed to suppress background contributions from Drell-Yan processes. Similarly, events are removed if they contain a same-flavour opposite-charge light lepton pair with an invariant mass within 10 GeV of the Z boson mass to reduce contributions from $t\bar{t}Z$.

The hadronically-decaying τ of the event is required to satisfy the medium BDT identification working point as the efficiency loss from moving to a tighter working point completely offsets the benefits from a reduction in fake $\tau_{\text{had-vis}}$ backgrounds, while a loose working point would allow too many events with fake $\tau_{\text{had-vis}}$ into the signal region.

Since the signal purity is high in this channel due to the high lepton multiplicity requirement, no further requirements are applied. After these selections, the channel is populated approximately evenly by $t\bar{t}H$ events, $t\bar{t}Z$ events and events with objects misidentified as $\tau_{\text{had-vis}}$.

8.5 Background estimations in $3\ell + 1\tau_{\text{had}}$

The two main background contributions to this channel are shown in Figure 8.3 and consist predominantly of the $t\bar{t}Z$ process with real $\tau_{\text{had-vis}}$ and events with objects misidentified as $\tau_{\text{had-vis}}$ (called fake $\tau_{\text{had-vis}}$), with smaller contributions from events with fake or non-prompt light leptons and other rare processes resulting in $3\ell + 1\tau_{\text{had}}$ final states[7].

8.5.1 Irreducible background

The rate and kinematic distributions of the irreducible backgrounds are taken from Monte Carlo simulation. Additional regions can be used to ensure proper modelling of the various physics processes and in particular to study the modelling and determine the normalisation of the $t\bar{t}Z$ process[8].

The events that make up the $3\ell + 0\tau_{\text{had}}$ channel are very similar to those populating the $3\ell + 1\tau_{\text{had}}$ region, but with jets from QCD instead of the $\tau_{\text{had-vis}}$. The $3\ell + 0\tau_{\text{had}}$ channel is divided into signal and control regions using a 5-dimensional BDT, as described in [58]. Because of the similarities between $3\ell + 0\tau_{\text{had}}$ and $3\ell + 1\tau_{\text{had}}$ events, these CRs provide a phase-space similar to that of $3\ell + 1\tau_{\text{had}}$ and allow for a study of the MC modelling of the irreducible background processes relevant to $3\ell + 1\tau_{\text{had}}$. The main advantage for the use of these CRs in $3\ell + 1\tau_{\text{had}}$ stems from the fact that the lower lepton multiplicity requirement in the $3\ell + 0\tau_{\text{had}}$ regions ensures higher statistics than those that could be obtained in $3\ell + 1\tau_{\text{had}}$ regions. The higher statistics allow for a dedicated high-purity $t\bar{t}Z$ region to be created using the 5D BDT.

Comparisons of data and MC in this region show good agreement between the two, and motivate the use of MC to describe both the rate and the kinematics of this background as illustrated in Figure 8.5 in the jet multiplicity distribution.

8.5.2 Reducible backgrounds

Thanks to stringent lepton selections, the contribution from fake and non-prompt light leptons is small, although not negligible. In contrast, the fake $\tau_{\text{had-vis}}$ contribution is much larger and its modelling in MC is seen to deviate from observed rates in data by around 35%. Since this analysis requires a more accurate determination of the fake $\tau_{\text{had-vis}}$ rate, a data-driven procedure is implemented.

Fake and non-prompt light leptons can occur independently of or in conjunction with a fake $\tau_{\text{had-vis}}$: methods that aim to estimate both contributions simultaneously, such as the ABCD method, generally do not produce reliable results due the potential correlations between the light lepton and $\tau_{\text{had-vis}}$ fake rates. For this reason, the two contributions are estimated independently. Contributions from events with a fake $\tau_{\text{had-vis}}$ are assessed using a dedicated fake estimate based on a data-driven method regardless of whether fake and non-prompt light leptons are present.

---

[7]These rare processes include processes with top quarks such as single top production, tZ, tWZ, $t\bar{t}WW$, triple top production and quadruple top production. Higgs production processes other than $t\bar{t}H$, such as tH, also contribute to the rare background category. Processes without top quarks include weak triple boson production and Drell-Yan/Z+jets: these processes require additional QCD radiation to pass the event selection criteria, and often also contribute through non-prompt or fake light leptons and $\tau_{\text{had-vis}}$.

[8]$t\bar{t}Z$ events with a fake $\tau_{\text{had-vis}}$ are considered as part of the reducible background and are studied separately.

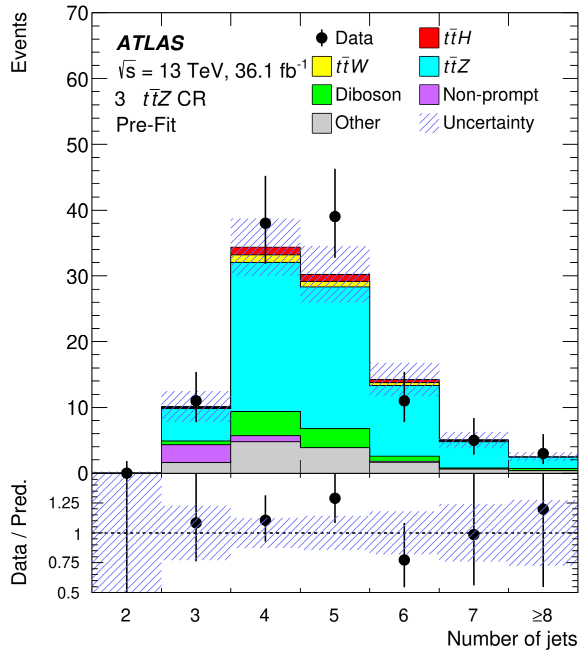

Figure 8.5: Comparison of data and MC prediction of the jet multiplicity in the $3\ell + 0\tau_{\text{had}}$ $t\bar{t}Z$-enriched control region. Image from [58].

The remaining contribution from events with real $\tau_{\text{had-vis}}$ and fake/non-prompt light leptons is extremely small and does not warrant a special treatment. Instead, the MC estimation is used and a 100% uncertainty is assigned, as studies in the $3\ell + 0\tau_{\text{had}}$ channel indicate that the MC underestimates this contribution by approximately a factor of 2.

8.5.3 Fake τ_{had}

Fake $\tau_{\text{had-vis}}$ contributions from misidentified light leptons are small. The MC $e \to \tau_{\text{had-vis}}$ rates are already scaled to those observed in data (see the tag-and-probe procedure for the electron veto BDT discussed in Section 6.3). The $\mu \to \tau_{\text{had-vis}}$ rate is extremely small and is well-modelled, so no correction is applied.

The contribution that remains to be estimated is the one from jets misidentified as $\tau_{\text{had-vis}}$. Jets from quarks and gluons produced in the hard-scatter and pile-up jets can result in detector signatures that are similar to that of a $\tau_{\text{had-vis}}$. On average, jets tend to have both higher track multiplicities and be wider than $\tau_{\text{had-vis}}$. Therefore, only very specific jets can generate fake $\tau_{\text{had-vis}}$. The probability of such a jet being produced is difficult to model. Instead of relying on the MC predictions, a data-driven method is performed to estimate this contribution. This method is described in detail in [104] and will be summarized below.

The data-driven method relies on the loosening of the $\tau_{\text{had-vis}}$ identification working point to study the rate at which fake $\tau_{\text{had-vis}}$ are selected. For $3\ell+1\tau_{\text{had}}$, a region with a relaxed $\tau_{\text{had-vis}}$ selection would overlap with the $3\ell + 0\tau_{\text{had}}$ signal region, leading to large signal contamination and making the fake estimate dependent on the measured $t\bar{t}H$ cross section. Additionally, a statistical analysis of the data would be difficult to perform as some events could contribute to both the $3\ell+0\tau_{\text{had}}$ signal region and the

$3\ell+1\tau_{\text{had}}$ region used to estimate fake $\tau_{\text{had-vis}}$ contributions. For these reasons, no dedicated data-driven estimate is performed in $3\ell+1\tau_{\text{had}}$. Instead, the MC contribution in this channel is scaled based on the mismodelling observed in $2\ell\text{OS}+1\tau_{\text{had}}$. The $2\ell\text{OS}+1\tau_{\text{had}}$ regions are ideal for studies of fake $\tau_{\text{had-vis}}$ due to the high statistics and low signal purity that they offer. The fake estimation in $2\ell\text{OS}+1\tau_{\text{had}}$ is described in Sections 8.5.3.1 and 8.5.3.2 while the scaling procedure used to bring those results to $3\ell+1\tau_{\text{had}}$ is outlined in Section 8.5.3.3. A validation of the method is discussed in Section 8.5.3.4.

8.5.3.1 Object and event selection in $2\ell\text{OS}+1\tau_{\text{had}}$

Events in the $2\ell\text{OS}+1\tau_{\text{had}}$ channel must contain exactly two reconstructed light leptons and one $\tau_{\text{had-vis}}$. The total charge of the light leptons must equal 0 while that of the $\tau_{\text{had-vis}}$ must be of ± 1.

Since the main backgrounds in this channel ($t\bar{t}$ and Z+jets) usually generate real and prompt leptons, the light leptons of the event are only required to be $\text{L}^{\dagger}$. The momentum of the leading (sub-leading) light lepton must be of at least 25 (15) GeV. When the light leptons are of same flavour, their invariant mass is required to be above 12 GeV to suppress background contributions from Drell-Yan processes. Similarly, events are removed if they contain a same-flavour light lepton pair with an invariant mass within 10 GeV of the Z boson mass to reduce contributions from Z+jets. As for the $3\ell+1\tau_{\text{had}}$ channel, the $\tau_{\text{had-vis}}$ are required to pass the medium identification BDT working point. Finally, the event must contain at least 3 jets out of which at least one must be b-tagged. The jet multiplicity requirements further reduce contributions from both $t\bar{t}$ and Z+jets.

After these selections, $t\bar{t}$ is the dominant background in the channel, mainly appearing through a decay to two prompt light leptons and a jet (often a heavy flavour jet) misidentified as a $\tau_{\text{had-vis}}$. This process accounts for 90% of the backgrounds in this channel, with the remaining 10% coming primarily from the Z+jets process. The $2\ell\text{OS}+1\tau_{\text{had}}$ regions are therefore further defined using a BDT that selects events consistent with the $t\bar{t}H$ process and reject events that are more $t\bar{t}$-like. The BDT used is described in detail in [104]. Its score is binned and only the bin with the highest score is used as a signal region. The lower-significance bins contain background-like events and are used as the nominal control region to study fake $\tau_{\text{had-vis}}$ contributions. When the two regions are considered together, the joint region is called the $2\ell\text{OS}+1\tau_{\text{had}}$ pre-selection region. The BDT distribution for $2\ell\text{OS}+1\tau_{\text{had}}$ is shown in Figure 8.6.

8.5.3.2 Fake factor estimate for τ_{had} in $2\ell\text{OS}+1\tau_{\text{had}}$

The method used to estimate the fake $\tau_{\text{had-vis}}$ background is called fake factor (FF). The FF procedure begins with the selection of a region known to be populated by processes not involving $\tau_{\text{had-vis}}$ such that any reconstructed $\tau_{\text{had-vis}}$ in this region must be a fake $\tau_{\text{had-vis}}$. Here, the number of events with fake $\tau_{\text{had-vis}}$ passing and failing an identification selection can be counted. Regions where the requirement on the $\tau_{\text{had-vis}}$ identification is reversed are called side-bands, while the ones with the identification requirement is kept are called nominal regions. The FF is the ratio of those two numbers, and provides an estimation of the number of fake $\tau_{\text{had-vis}}$ that will pass an identification working point when the number of fake $\tau_{\text{had-vis}}$ failing that working point is known. The formula for the FF is

$$\frac{\#\text{ of events with fake }\tau_{\text{had}}\text{ that pass the identification WP}}{\#\text{ of events with fake }\tau_{\text{had}}\text{ that fail the identification WP}} = \frac{N_C}{N_D} = \text{FF}_{CD}. \tag{8.1}$$

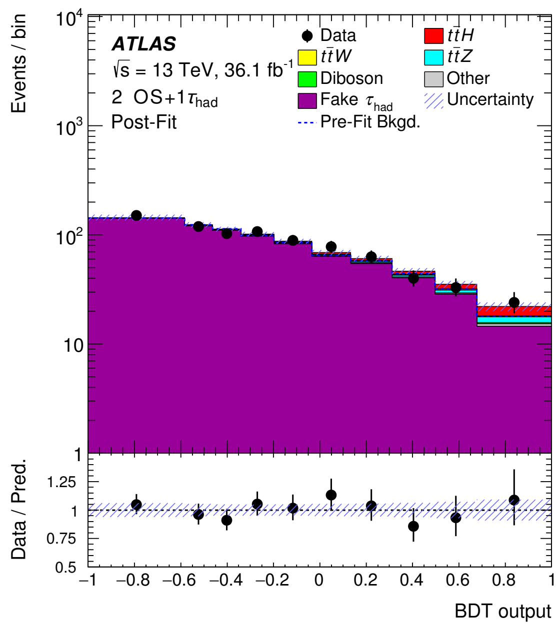

Figure 8.6: Distribution of the $2\ell\text{OS} + 1\tau_{\text{had}}$ BDT. The expected contributions are shown as coloured bands, and the observed data yields are shown as black dots with error bars indicating the statistical uncertainty of the yield measurement. Image from [58].

The FF must be derived in an extraction region (ER) that is similar to the signal region whose fake $\tau_{\text{had-vis}}$ background it aims to estimate. Only if the two regions contain events from similar processes, and with similar kinematics, is it possible to expect the FF to be valid in the SR, i.e.

$$\text{FF}_{CD} = \text{FF}. \tag{8.2}$$

In that case, one can estimate the number of events with fake $\tau_{\text{had-vis}}$ that will pass the chosen identification working point in the signal region (N_A) based on the number of events that fail the identification selection in its corresponding side-band (N_B) as follows:

$$N_A = \text{FF} \cdot N_B. \tag{8.3}$$

This procedure is illustrated in Figure 8.7.

Assumptions

In order to ensure that the FF computed in the ER can be used for the SR, the following two conditions must hold:

- The type of object faking the $\tau_{\text{had-vis}}$ is the same in the ER and the SR.
- There are no real $\tau_{\text{had-vis}}$ in the regions, or the real $\tau_{\text{had-vis}}$ contributions are small and can be subtracted using truth-matched MC.

The first point must be considered when selecting the side-band population. Figure 8.8 shows the fake $\tau_{\text{had-vis}}$ composition as a function of the $\tau_{\text{had-vis}}$ identification score. The nominal selection corresponds

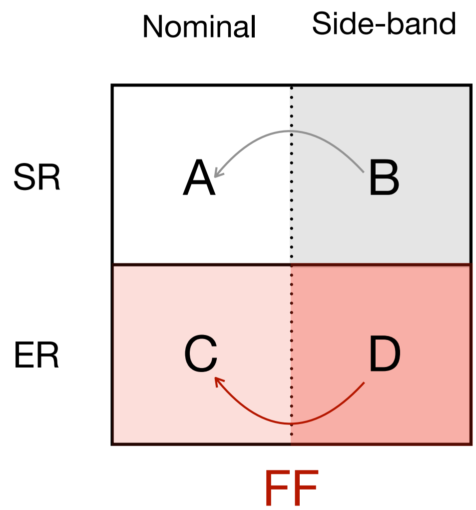

Figure 8.7: The FF method. The nominal selection is the one that requires reconstructed $\tau_{\text{had-vis}}$ to pass the identification criteria, while the side-band is a region identical to the nominal except for the requirement that the object fails the identification selection. The FF is derived in an extraction region (red) and applied to the side-band of the signal region to obtain a fake estimate (grey). The expected number of events with fake $\tau_{\text{had-vis}}$ in region A is simply $N_A = \text{FF} \cdot N_B = \frac{N_C}{N_D} \cdot N_B$.

to a medium identification working point, around a score of 0.75, and selects mostly real $\tau_{\text{had-vis}}$. At lower scores, the contributions from light quark and gluon jets vary significantly. For the side-band the selection for the score is chosen as > 0.35 to maintain a quark-to-gluon ratio that is comparable to that of the fake $\tau_{\text{had-vis}}$ in the nominal selection. Additionally, this selection ensures that the jets have similar detector signatures (that are $\tau_{\text{had-vis}}$-like), reducing the difference in FFs between the types of fake $\tau_{\text{had-vis}}$ as illustrated by Figure 8.9.

For the second point, it is important to note that no region without real $\tau_{\text{had-vis}}$ can be built. While the selections of the extraction region can be chosen to ensure that the regions are enhanced in fake $\tau_{\text{had-vis}}$, a subtraction must be performed to account for remaining contributions from events with real $\tau_{\text{had-vis}}$. MC simulation is used to estimate the number of events with real $\tau_{\text{had-vis}}$ that are expected, and Equations 8.1 and 8.3 are modified such that $N_X = N_X^{\text{data}} - N_X^{\text{MC, real}}$.

All $t\bar{t}H$ events are also subtracted since they are considered as signal regardless of whether the $\tau_{\text{had-vis}}$ is real or fake. This subtraction is performed in parallel with a measurement of $\mu_{t\bar{t}H}$ in the $2\ell\text{OS} + 1\tau_{\text{had}}$ SR, to ensure that the correct amount of signal is subtracted.

Extraction regions for $2\ell\text{OS} + 1\tau_{\text{had}}$

With the above constraints in mind, an extraction region can be chosen for $2\ell\text{OS} + 1\tau_{\text{had}}$. This region must be kinematically similar to the pre-selection region of $2\ell\text{OS} + 1\tau_{\text{had}}$, contain primarily fake $\tau_{\text{had-vis}}$, and its side-band must include a selection on the $\tau_{\text{had-vis}}$ identification score of at least 0.35.

The $2\ell\text{OS} + 1\tau_{\text{had}}$ pre-selection region is defined via three main selections: a requirement on the invariant mass of the dilepton system ($m_{\ell\ell}$), on the number of jets, and on the number of b-tagged jets. A $2\ell\text{OS} + 1\tau_{\text{had}}$ extraction region can therefore be defined by modifying one of these selections. When removing these requirements from the $2\ell\text{OS} + 1\tau_{\text{had}}$ pre-selection region definition and plotting the FFs

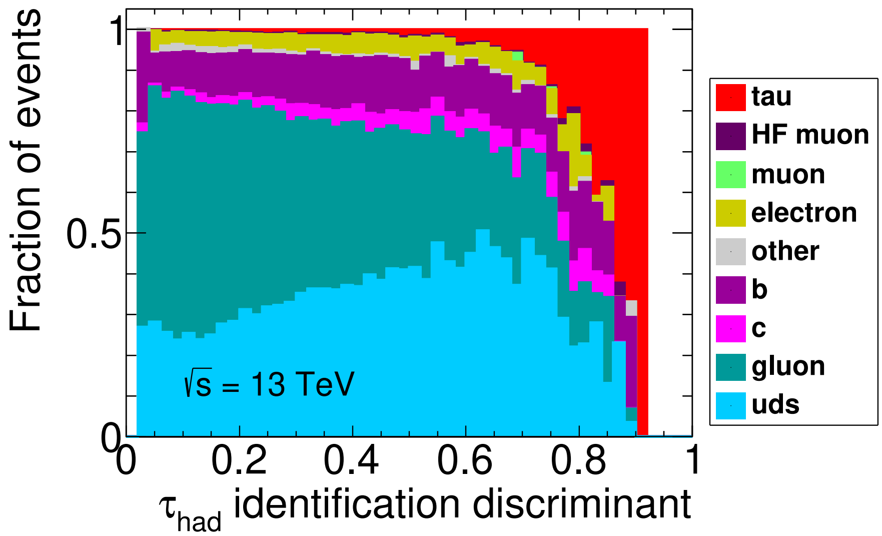

Figure 8.8: Composition of reconstructed $\tau_{\text{had-vis}}$ candidates in a $2\ell\text{OS} + 1\tau_{\text{had}}$ region where the tau identification requirement is removed. The real $\tau_{\text{had-vis}}$ contribution is shown in red, with light quark and gluon jets in light blue and dark green respectively. Image from [104].

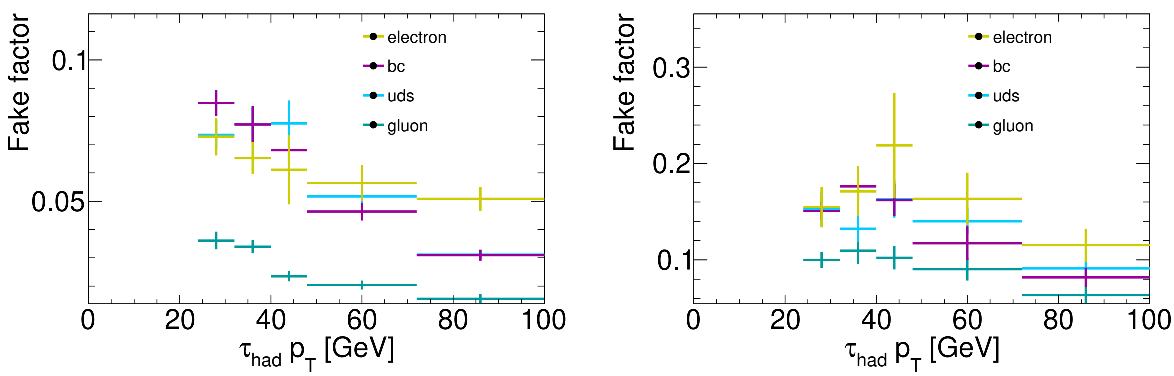

Figure 8.9: FF as a function of the transverse momentum of the $\tau_{\text{had-vis}}$ for gluons (green), light quarks (blue), heavy quarks (purple) and electrons (yellow) faking $\tau_{\text{had-vis}}$. The FF are shown with (right) and without (left) requirements of > 0.35 on the tau identification score. Without the requirement, large discrepancies between the gluon FF and the other FFs are seen (left) while when the requirement is placed, the FFs between the different types of object are similar (right). Images from [104].

in each bin, certain regions of phase space result in compatible FFs. Figure 8.10 illustrates this process for the jet multiplicity variable. Here, it can be seen that while regions with very low jet multiplicity are associated to higher FF values, all regions with at least 3 jets have similar FFs. Similar tests show that no bias is seen in the FF values as a function of $m_{\ell\ell}$ and b-jet multiplicity.

In the end, the main extraction region was selected to be identical to the signal region, but with the b-jet requirement inverted. Two other extraction regions were created:

1. The Z-enriched ER has the same selection as the main ER except that the Z boson mass veto is inverted. These events enter the $2\ell\text{OS} + 1\tau_{\text{had}}$ category with the opposite-sign light leptons from

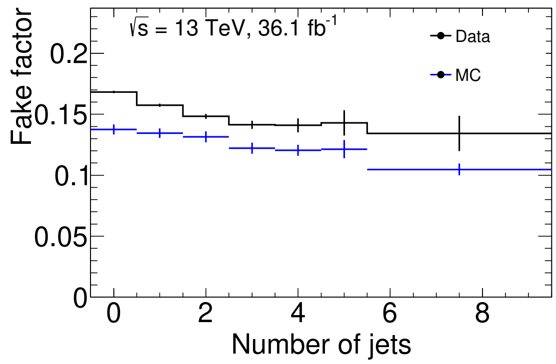

Figure 8.10: Fake factors as function of the number of jets in the event for data (black) and MC (blue). Image from [104].

the $Z \to \ell\ell$ decay, and usually contain a fake $\tau_{\text{had-vis}}$ coming from the event's light jet activity. The Z+jet contribution is almost 100% in this ER.

2. The $t\bar{t}$-enriched ER also has the same selection as the pre-selection region except that the jet requirements are changed to exactly 2 jets and at least 1 b-jet. The light leptons of opposite sign in this region are usually produced via a $t\bar{t} \to \ell\ell$ decay and a fake $\tau_{\text{had-vis}}$ is often reconstructed from a b-jet. This ER has a composition comparable to that of the SR.

The contribution from real $\tau_{\text{had-vis}}$ is found to be of at most 2% in all regions (excluding the nominal SR).

Fake factors

The FFs are derived in bins of the transverse momentum of the $\tau_{\text{had-vis}}$ as this variable is used in the BDT that defines the $2\ell\text{OS} + 1\tau_{\text{had}}$ SR and CRs and needs to be well-modelled. The dependence of the FF on other variables such as the number of vertices in the event and the prongness of the $\tau_{\text{had-vis}}$ is not taken into consideration as the BDT is not sensitive to them.

The FF value is obtained from the extraction regions. It is then applied to the side-band of the signal pre-selection region.

The total estimates for N_A, the number of events with fake $\tau_{\text{had-vis}}$ that pass the medium $\tau_{\text{had-vis}}$ identification working point in the $2\ell\text{OS} + 1\tau_{\text{had}}$ pre-selection region, and their statistical uncertainties are the following:

- When using the main ER, the estimate is $N_A = 750 \pm 40$.
- When using the Z-enriched ER, the estimate is $N_A = 714 \pm 22$. This corresponds to a 4% decrease in the expected yields with respect to the estimate using the main ER.
- When using the $t\bar{t}$-enriched ER, the estimate is $N_A = 685 \pm 35$, corresponding to an 8% reduction in yields with respect to the estimate using the main ER.

These variations give an indication of the remaining flavour dependence in the SFs, and assess how well the first assumption holds.

It is also possible to vary the subtracted contributions, to account for uncertainties in the expected real $\tau_{\text{had-vis}}$ and $t\bar{t}H$ rates.

- When scaling the real $\tau_{\text{had-vis}}$ MC estimate by +20%, the estimate becomes $N_A = 732 \pm 36$ while a scaling of -20% for this contribution results in $N_A = 760 \pm 36$. These variations have an impact of ±2% on the estimate.
- When scaling the $t\bar{t}H$ subtraction by +100%, the estimate is $N_A = 740 \pm 36$ while a scaling of -100% results in $N_A = 752 \pm 36$. Here, the impact is of ±1%. The $t\bar{t}H$ subtraction was varied by 100% as a conservative estimate, since its production rate was known to a precision of around 60% from previous $t\bar{t}H$ results.

This second set of variations assess the stability of the estimate with respect to the assumption regarding the real $\tau_{\text{had-vis}}$ subtraction.

8.5.3.3 Scale factors for $3\ell + 1\tau_{\text{had}}$

A comparison of the estimate using the FF and the estimate from the MC prediction in the $2\ell\text{OS}+1\tau_{\text{had}}$ low-BDT CR is performed. A fit comparing the yields generates scale factors (SFs) that indicate the scaling required to get the MC estimate to agree with the data-driven one. If the two estimates agree, the SF should be consistent with 1. A data-driven estimate twice as large as the MC one would instead result in a SF of 2. When considering the FFs from the main ER, the SF and its statistical uncertainty are 1.36 ± 0.05.

The fake $\tau_{\text{had-vis}}$ composition in the $2\ell\text{OS}+1\tau_{\text{had}}$ and $3\ell+1\tau_{\text{had}}$ regions is not expected to be identical as the main background of the former, $t\bar{t}$, results in more fake $\tau_{\text{had-vis}}$ from heavy-flavour jets compared to the main background of the latter, $t\bar{t}Z$, as illustrated in Figure 8.11. While differences in composition are observed, the effects are covered by the composition variations of the alternate extraction regions. The SF procedure is therefore repeated using the alternate FFs (Z-enriched, $t\bar{t}$-enriched) to estimate the systematic uncertainty that arises from jet flavour effects. This results in the following values for the SFs:

- Using FFs from the Z-enriched ER, the SF becomes 1.29 ± 0.04.
- With the FFs from the $t\bar{t}$-enriched ER the SF is found to be 1.23 ± 0.04.

In $3\ell + 1\tau_{\text{had}}$, the p_{T} modelling is not important so an inclusive SF is used with a final value of 1.36 ± 0.05 (stat.) $\pm\, 0.07$ (syst. from Z-enriched) $\pm\, 0.13$ (syst. from $t\bar{t}$-enriched) $= 1.36 \pm 0.16$. This SF is used to scale the MC fake $\tau_{\text{had-vis}}$ estimate for both signal and background events instead. The two types of processes have similar fake $\tau_{\text{had-vis}}$ composition (shown in Figure 8.12) so the SF is expected to be valid for both types of processes.

8.5.3.4 Validation

The estimation method is validated for $3\ell+1\tau_{\text{had}}$ events through a validation region (VR). This region is identical to the $3\ell + 1\tau_{\text{had}}$ SR except for the replacement of the b-jet requirement with a b-jet veto. It is expected to be more similar than the $2\ell\text{OS}+1\tau_{\text{had}}$ regions to the $3\ell+1\tau_{\text{had}}$ SR in terms of kinematics and event yields, but little $t\bar{t}H$ signal is expected. The validation test is performed in this VR by applying the SF to event yields from MC with fake $\tau_{\text{had-vis}}$. The fake $\tau_{\text{had-vis}}$ estimate is added to the contribution from events with real $\tau_{\text{had-vis}}$ (taken directly from MC) and the resulting total MC estimate is compared to data.

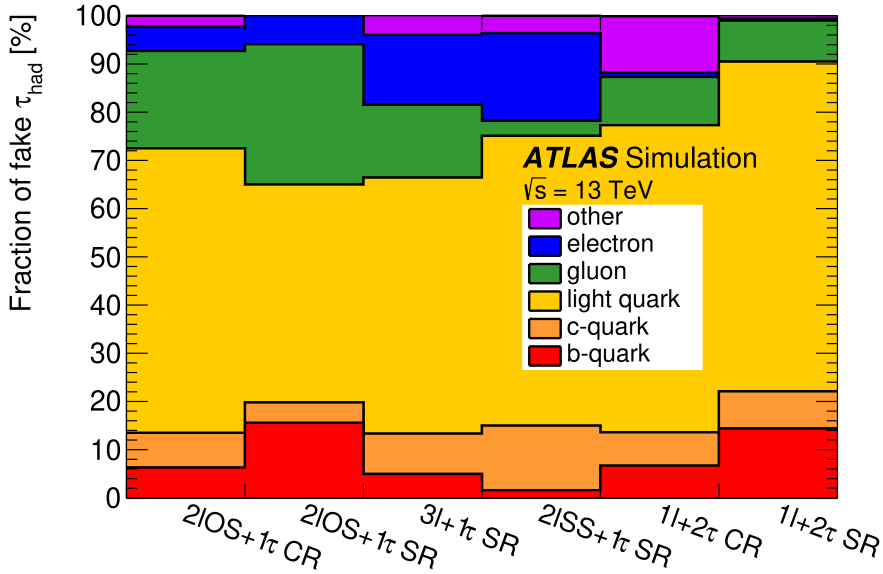

Figure 8.11: Fake $\tau_{\text{had-vis}}$ composition in the signal and control regions of the $t\bar{t}H \rightarrow$ multilepton analysis. Image from [58].

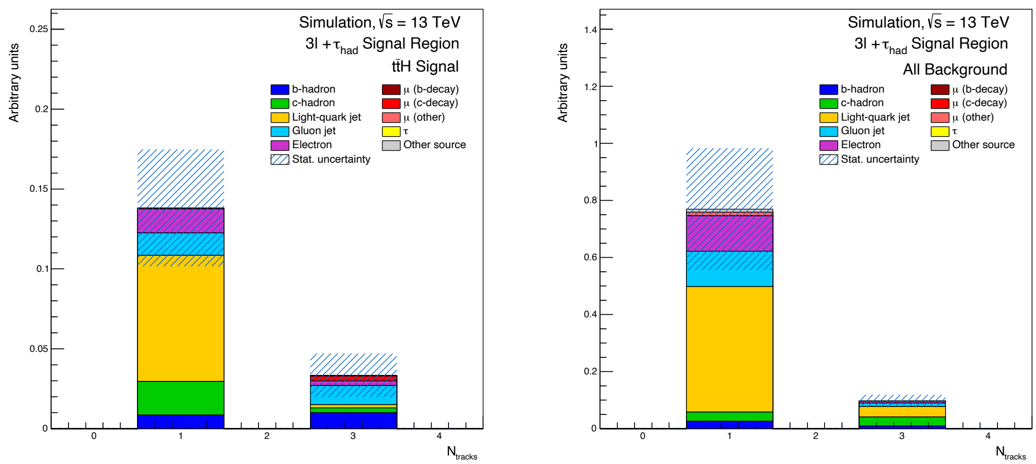

(a) $t\bar{t}H$ (b) All background

Figure 8.12: Fake $\tau_{\text{had-vis}}$ composition in the $3\ell + 1\tau_{\text{had}}$ signal region for the signal and background processes with fake $\tau_{\text{had-vis}}$.

In the VR, the MC prediction for events with real $\tau_{\text{had-vis}}$ is 1.4 ± 0.4. The MC yield for events with fake $\tau_{\text{had-vis}}$ is 1.8 ± 0.8. When the SF is applied, the estimation becomes 2.4 ± 1.1. Summing the two contributions, the expected total MC yield in this region is expected to be 3.9 ± 1.2. The data yield is 4, which is consistent with the prediction. While the statistics are very low, this test increases the confidence that the SF method is appropriate for use in the $3\ell + 1\tau_{\text{had}}$ channel.

8.5.4 Summary

With all signal and background contributions properly estimated, the total expected yields in the $3\ell + 1\tau_{\text{had}}$ signal region are computed and shown in Table 8.4.

These yields correspond to what one would expect to observe in data if the SM were to be accurate. The next step is to look at the data yield in this region, and compare it to the expected ones. A statistical analysis of the results is performed to extract information about the validity of the SM predictions, in this case as it relates to the rate of the $t\bar{t}H$ process.

Process	$\tau_{\text{had-vis}}$	Light lepton	Estimate from ...	Yield (scaled)
Signal	Real	-	MC	1.3
	Fake	-	MC with SF from $2\ell\text{OS} + 1\tau_{\text{had}}$	0.2
Background	Real	Real	MC	1.7
		Fake/n.-p.	MC with 100% syst.	0.1
	Fake (e/μ)	-	MC	0.2
	Fake (jets)	-	MC with SF from $2\ell\text{OS} + 1\tau_{\text{had}}$	0.6
			Total signal	1.5
			Total background	2.6
			Total expected	4.1

Table 8.4: Event yields for signal and background processes, with real and fake $\tau_{\text{had-vis}}$ and real, fake and non-prompt (n.-p.) light leptons. A "-" indicates that no requirement is imposed.

8.6 Fit strategy

The statistical analysis of the $t\bar{t}H \to$ multilepton results is based on the hypothesis testing procedure outlined below.

8.6.1 Counting experiments

When attempting to count the number of events where a Higgs boson is produced via the $t\bar{t}H$ process, one begins with a set of collision events N. The probability p for any event to be a $t\bar{t}H$ event is defined as follows:

$$p = \frac{\nu}{N} = \frac{L \cdot \sigma(t\bar{t}H) \cdot A \cdot \epsilon}{N} \tag{8.4}$$

where ν is the expected number of $t\bar{t}H$ events based on the collected luminosity L, the expected cross section of the $t\bar{t}H$ process $\sigma_{t\bar{t}H}$, the acceptance A defined as the fraction of $t\bar{t}H$ events that are available based on the geometry of the detector, and the efficiency ϵ corresponding to the proportion of $t\bar{t}H$ events one expects to pass the SR selections. The probability P to find n $t\bar{t}H$ events in N collision events is then

$$P(n) = \binom{N}{n} p^n (1-p)^{N-n}. \tag{8.5}$$

If there is a sufficiently large number of events N and a sufficiently small probability p, a Poisson distribution can be used to approximate $P(n)$ such that

$$P(n) \rightarrow \text{Poisson}(n, Np) = \frac{e^{-Np}(Np)^n}{n!}. \tag{8.6}$$

From Equation 8.4, Np is simply ν and so the probability to observe n $t\bar{t}H$ events depends only on the number of events one expects to find:

$$P(n) = \frac{e^{-\nu}(\nu)^n}{n!}. \tag{8.7}$$

8.6.2 Signal strength

In general, signal regions are not only composed of signal events. The total expected number of events ν depends instead on both the number of signal (s) and background (b) events as $\nu = s + b$. From Equation 8.4, the number of background events corresponds to $L\sigma_{\text{bkg}}A_{\text{bkg}}\epsilon_{\text{bkg}}$. The number of signal events s is similarly defined as $L\sigma_{\text{sig}}A_{\text{sig}}\epsilon_{\text{sig}}$. Since the purpose is to eventually measure the signal cross section, is it helpful to factorize that dependence out of s. One can define the signal strength parameter $\mu = \frac{\sigma_{\text{obs}}}{\sigma_{\text{SM}}}$, such that $\nu = \mu s_{\text{SM}} + b$ where s_{SM} is the number of signal events predicted by the SM and any deviation from the SM cross section is contained within μ.

8.6.3 Hypotheses

In order to determine the total number of events ν one expects to find in a SR, it is important to choose a hypothesis for the rate of the process of interest. In the background-only hypothesis where no $t\bar{t}H$ process is assumed to exist, one would expect μ to take on the value of 0 and therefore find that the expected event count is simply b. If instead the $t\bar{t}H$ process is hypothesized to proceed as in the SM, μ should equal 1 and $\nu = s_{\text{SM}} + b$.

8.6.4 Likelihoods

Given a set of measurements x and a hypothesis $\boldsymbol{\alpha}$, the likelihood $\mathcal{L}(x|\boldsymbol{\alpha}) = \mathcal{L}(\boldsymbol{\alpha})$ is defined as the probability of observing the data x under the hypothesis $\boldsymbol{\alpha}$. It is important to note that $\boldsymbol{\alpha}$ represents both theoretical parameters, such as the rate of the $t\bar{t}H$ process, and nuisance parameters θ, describing the impact of systematic uncertainties on both the signal and background distributions and yields through deviations of L, σ_{bkg}, A and ϵ from their assumed values. These parameters correct the probability model and account for the fact that parameters other than μ are generally not known exactly. Therefore, $\boldsymbol{\alpha}$ can be represented as $\boldsymbol{\alpha}(\mu, \boldsymbol{\theta})$ where μ refers to the parameter of interest (POI) that one aims to measure and $\boldsymbol{\theta}$ represents the nuisance parameters (NPs).

In statistical analyses, likelihoods are used to make inferences on the POIs. Often, a maximum likelihood estimator technique is used where the estimate $\hat{\mu}$ denotes the value of μ for which the likelihood $\mathcal{L}(\mu)$ is maximal. In the $t\bar{t}H \rightarrow$ multilepton case, the likelihood function is used to extract the $t\bar{t}H$ signal strength ($\mu_{t\bar{t}H}$). Other Higgs production cross sections and all branching ratios are set to their SM expectation values with appropriate theoretical uncertainties. As discussed above, in analyses where the yield of a signal region is counted, such as in $3\ell + 1\tau_{\text{had}}$, the likelihood is given by a Poisson distribution

$$\mathcal{L}(x|\boldsymbol{\alpha}) = \text{Poisson}(n|\nu(\boldsymbol{\alpha})) \tag{8.8}$$

where n corresponds to the observed count and the expected count, $\nu(\boldsymbol{\alpha})$, is dependent on $\boldsymbol{\alpha}$. In channels where a distribution is fitted, such as in $2\ell\text{OS} + 1\tau_{\text{had}}$, the distributions are binned and each bin b is similarly described by its own Poisson term leading to

$$\mathcal{L}(x|\boldsymbol{\alpha}) = \prod_{b=1}^{b_{\max}} \text{Poisson}(n_b|\nu_b(\boldsymbol{\alpha})). \tag{8.9}$$

When considering several orthogonal regions, such as when combining the results from the many $t\bar{t}H \to$ multilepton channels, the overall likelihood is simply the product of the likelihood functions in the individual channels c,

$$\mathcal{L}(x|\boldsymbol{\alpha}) = \prod_{c=1}^{c_{\max}} \text{Poisson}(n_c|\nu_c(\boldsymbol{\alpha})). \tag{8.10}$$

8.6.5 Systematic uncertainties

While Section 8.6.2 presents a direct relation between event counts and the signal strength, systematic uncertainties make the relation between the data and the POI less certain. The model must be built so that it can be regarded as correctly describing the data for some values of the parameters μ and $\boldsymbol{\theta}$, and it must be flexible enough to account for all potential sources of uncertainty.

The NPs are usually estimated from subsidiary measurements $\rho(\boldsymbol{\theta})$: in this case, for example, the overall normalisation of the fake τ_{had} contribution is measured in the $2\ell\text{OS} + 1\tau_{\text{had}}$ control regions. The likelihood model can be expanded to explicitly account for their effects by introducing the associated subsidiary measurements explicitly:

$$\mathcal{L}(x|\boldsymbol{\alpha}) = \mathcal{L}(x|\mu, \boldsymbol{\theta}) = \text{Poisson}(n_c|\nu_c(\mu, \boldsymbol{\theta})) \cdot \rho(\boldsymbol{\theta}). \tag{8.11}$$

Including these systematic uncertainties into the model using NPs instead of setting their values to those obtained from the subsidiary measurements results in a joint measurement of $\boldsymbol{\alpha}(\mu, \boldsymbol{\theta})$ which allows for stronger constraints on $\boldsymbol{\theta}$ in the analyzed phase space than would be provided by the subsidiary measurement alone. The resulting values and uncertainties on the NPs correct the original model from Equation 8.10 to provide a better description of the data in the phase space of the analysis.

Two types of subsidiary measurements can be performed: NPs that apply to object selections and calibrations come from dedicated measurements by combined performance groups, while normalisation factors (NFs), used to normalise background processes, are measured in control regions within the analysis[9]. In the latter case, the full probability model is available and so the likelihood model can be expanded to contain the information obtained in the control regions in the same way that a joint measurement is made when many signal regions are combined (see Equation 8.10). A continuous description of the model is provided for all values of θ, and the effect of simultaneous variations can be extracted. In the former case, only the maximum likelihood estimate for the NP and its error are known. In order to incorporate these terms into the likelihood model an assumption on the form of $\rho(\theta)$ must be made.

- For most systematic uncertainties, Gaussian probability density functions (pdfs) are used for $\rho(\theta)$.

[9] The use of control regions involves a trade-off between uncertainties. For example, uncertainties on the cross section of the background process do not need to be considered if one uses a CR to measure its normalisation. However, extrapolation uncertainties (that indicate the degree of confidence one has in the fact that the conclusions drawn in the CR can be used in the SR) and statistical uncertainties from the event yields in the CR must be added. A good CR is one whose associated extrapolation and statistical uncertainties are lower than the original modelling uncertainties.

The effects of θ are parametrised so that they can be described by normal distributions centered at $\tilde{\theta} = 0$ and with standard errors of 1, allowing all systematic uncertainties to be visualized on a similar scale. A parametrisation where $\theta \mapsto \ln\theta$ is often chosen and yields a log-normal pdf

$$\rho(\theta) = \frac{1}{\sqrt{2\pi}\ln\kappa_\theta}\frac{1}{\theta}\exp - \frac{(\ln(\theta/\tilde{\theta}))^2}{2(\ln\kappa_\theta)^2} \tag{8.12}$$

The width of the log-normal pdf is characterized by κ_θ. A deviation from $\tilde{\theta}$ of $\kappa_\theta = 1.1$ implies that the value will be larger or smaller than $\tilde{\theta}$ by a factor of 1.1 (each with a probability of about 16%). To explicitly evaluate the impact of the NP on the event count, the expected count ν can be described as $\nu(\theta) = \tilde{\nu} \cdot \kappa_\theta$ where $\tilde{\nu}$ is its estimate when $\theta = \tilde{\theta}$.

- For uncertainties arising from the limited number of events M of a generated Monte Carlo sample, or due to the size of a data sample in a control region that is not explicitly included in the likelihood function, the subsidiary measurement follows a Poisson distribution. The event rate ν is related to M via $\nu = aM$. Therefore, a is the luminosity ratio for an estimate from MC, and it is the ratio of the event count between a SR and a CR when data is used. The value of M is subject to statistical uncertainties, giving rise to the following pdf:

$$\rho(\nu) = \frac{1}{a}\frac{(\nu/a)^M}{M!}\exp(-\nu/a). \tag{8.13}$$

Here, the most probable value for ν is $\tilde{\nu} = aM$, the mean value is $a(M+1)$ and the dispersion is $a\sqrt{M^2+1}$. It is useful for these effects to also be parametrised to fluctuate about 1. In this case, the NP is defined as $\theta = \frac{\nu}{\tilde{\nu}}$ so that the expected event count fluctuates about its most probable value as $\nu(\theta) = \tilde{\nu} \cdot \theta$.

The robustness of the statistical analysis depends on the quality of the approximation of the effects of systematic uncertainties. A bad model will always provide unreliable results. For this reason, the behaviour of NPs must be thoroughly inspected and understood, and the assumptions made must be validated. Best-estimate values $\hat{\theta}$ of NPs that deviate significantly from 0 or errors that are much smaller or larger than 1 can often (but not always) indicate the presence inaccurate assumptions or inconsistencies. When shown, sources of uncertainty are typically ranked according to their impact on the measurement of the POI.

8.6.6 Hypothesis testing

Once the full likelihood model is built, it can be used to find maximum likelihood estimators and to test various hypotheses. In the case of $t\bar{t}H$, a maximum likelihood estimator of interest is $\hat{\mu}_{t\bar{t}H}$ and a hypothesis of interest is the background-only hypothesis.

As mentioned in Section 8.6.4, the likelihood is defined as the probability of observing the data x given a hypothesis $\boldsymbol{\alpha}$. In hypothesis testing, it is important to quantify the level of agreement between the data and the tested null hypothesis H_0, taking into account the fact that an alternate hypothesis H_1 may better fit the observed data. Test statistics are used for this purpose. The Neyman-Pearson lemma [106] states that for simple hypotheses[10] the likelihood ratio

$$\Lambda(H_0, H_1) \equiv \frac{\mathcal{L}(H_0)}{\mathcal{L}(H_1)} \tag{8.14}$$

maximizes the statistical power of a test of the null hypothesis H_0 with respect to the alternate hypothesis H_1, meaning that it maximizes the probability of rejecting the null hypothesis in the case that the alternate hypothesis is true. This makes it an ideal candidate as a test statistic.

In general, the test statistic used in Higgs analyses compares the hypothesis of the maximum likelihood estimator $\hat{\mu}$ to the hypothesis with a reference μ:

$$q_\mu = -2\ln\Lambda_\mu = -2\ln\frac{\mathcal{L}(\mu, \hat{\hat{\theta}}_\mu)}{\mathcal{L}(\hat{\mu}, \hat{\theta})} \tag{8.15}$$

where $\hat{\mu}$ and $\hat{\theta}$ are the parameters that maximize the overall likelihood, and $\hat{\hat{\theta}}_\mu$ are the parameters that maximize the likelihood for the reference value of μ. In the case of the $t\bar{t}H$ search, since the null hypothesis is that $\mu_{t\bar{t}H} = 0$, the test statistic becomes

$$q_0 = -2\ln\Lambda_0 = -2\ln\frac{\mathcal{L}(0, \hat{\hat{\theta}}_0)}{\mathcal{L}(\hat{\mu}, \hat{\theta})} \text{ if } \mu \geq 0 \text{ and } q_0 = 0 \text{ otherwise}^{11}. \tag{8.16}$$

This test statistic is quantifies the agreement of the data with a background-only hypothesis, taking into account the fact that the hypothesis $\hat{\mu}$ may better describe the data. Higher values of q_μ indicate increasing incompatibility between the data and the null hypothesis.

The value of the test statistic will depend on the data yield that one observes. Typical distributions for these values can be obtained by generating samples under the different hypotheses and determining the corresponding values that the test statistic takes, as shown in Figure 8.13. When randomly generating samples under the null hypothesis, one would expect q to fluctuate around lower values. Indeed, most of the times datasets generated using only the null hypothesis will result in event counts that are close to the null expectation value, while occasionally large fluctuations may cause the number of events to seem atypical under this hypothesis and result in higher values of q. Conversely, when generating samples under the alternate hypothesis, the value of q is expected to fluctuate around higher values while occasionally lower null-like values may occur. The probability to reject the null hypothesis when it is indeed true is proportional to the rate of the first example. This scenario is illustrated in Figure 8.13 with vertical line. When considering a random sample with a test statistic above this line one would conclude that it is unlikely, but not impossible, that the data was generated under the null hypothesis. In those cases, since the alternate hypothesis better describes the data, one might reject the null hypothesis. The probability of mistakenly rejecting the null hypothesis is defined as the integral of the null curve for values above the vertical line. This probability increases with decreasing luminosity, which is why it was not possible to reject the background-only hypothesis in $t\bar{t}H$ searches consisting of the Run 1 dataset.

The distributions of the probability density functions of the test statistics are denoted by $f(q_\mu|H)$ under

[10] Simple hypotheses are those without free parameters. Therefore likelihood ratios aren't necessarily optimal for models with important contributions from nuisance parameters.

[11] The constraint of positive μ causes an accumulation of events at $q_0 = 0$. This is because downward fluctuations indicate a deficit of events with respect to the expected background, and its incompatibility from the null hypothesis should not be considered on the same footing as an incompatibility due to the presence of signal events. For non-physical estimates of μ, i.e. those with negative values, the value $\mu = 0$ is assigned.

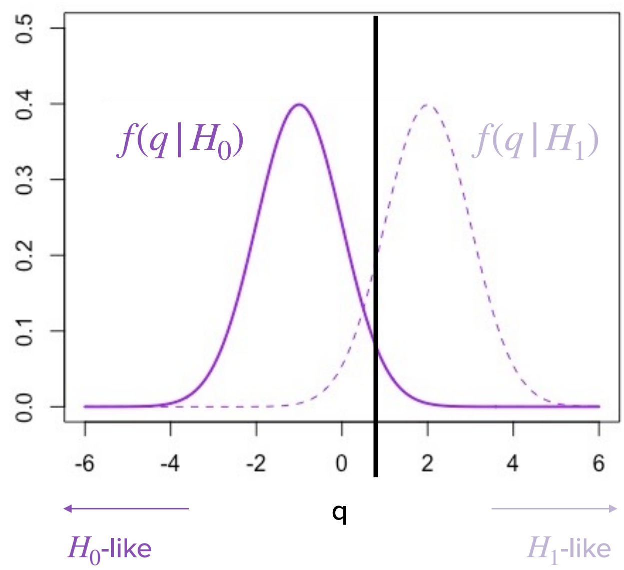

Figure 8.13: Typical distribution of test statistics for generic null (dark purple) and alternate (light purple) hypotheses: q values tend to cluster around lower values for samples generated under the null hypothesis while those generated under the alternate hypothesis tend to be associated to higher values. The value of q obtained in the experiment, such as the one shown with a black vertical line, can be used to select the hypothesis that best matches the data. Image adapted from [84].

a hypothesis H. While these pdfs can be obtained by running several experiments as mentioned above, this is often a time- and resource-consuming process. Instead, one can rely on Wald's theorem [122] to estimate the pdf. The theorem states that

$$q_\mu = \frac{(\mu - \hat{\mu})^2}{\sigma_{\hat{\mu}}^2} + \mathcal{O}\left(\frac{1}{\sqrt{N}}\right) \tag{8.17}$$

which reduces to

$$q_\mu = \frac{(\mu - \hat{\mu})^2}{\sigma_{\hat{\mu}}^2} \tag{8.18}$$

for large enough samples. Under this description, the pdf of q_μ under a generic assumption μ' is given by a non-central χ^2 with $\Lambda = \frac{(\mu-\mu')^2}{\sigma_{\mu'}^2}$:

$$f(q_\mu|H_{\mu'}) = f(q_\mu|\Lambda) = \frac{1}{2\sqrt{q_\mu}}\frac{1}{\sqrt{2\pi}}\left[\exp\left(-\frac{1}{2}(\sqrt{q_\mu}+\sqrt{\Lambda})^2\right) + \exp\left(-\frac{1}{2}(\sqrt{q_\mu}-\sqrt{\Lambda})^2\right)\right] \tag{8.19}$$

with $\hat{\mu} \sim \text{Gauss}(\mu', \sigma)$. When testing the null hypothesis, one recovers Wilks' theorem [124]

$$f(q_0|H_0) = f(q_0|\Lambda) = \frac{1}{\sqrt{2q_0\pi}}e^{-\frac{1}{2}q_0} = \chi^2 \tag{8.20}$$

since here $\mu = \mu'$ and $\Lambda = 0$. Using Wald and Wilks' theorems, one obtains the distributions shown in Figure 8.14.

To obtain the value of q for a dataset corresponding to the SM, one therefore only needs to produce one representative sample. This sample is called the Asimov dataset and it is constructed from the true

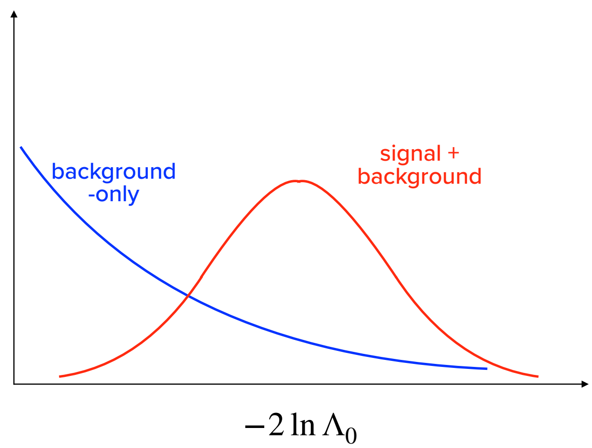

Figure 8.14: Pdfs of the test statistics for the null (background-only, blue) and alternate (signal+background, red) hypotheses. The curves correspond to the predictions from Wilks and Wald's theorems respectively.

parameter values: a set corresponding to the SM hypothesis is constructed with $\mu = 1$. The value of the test statistic associated to this dataset will correspond to the median of the $f(q_\mu|H)$ distribution and the value of $\hat{\mu}$ for this dataset will be that of the injected μ.

8.6.7 *p*-values

With the distributions of the test statistics under different hypotheses determined, one can generate an Asimov dataset and determine the value of q_0 one expects to obtain. This is done using the pre-fit yields presented in Section 8.5.4. One can then look at data and obtain the value of $\hat{\mu}$ by maximizing the likelihood with respect to all parameters. This procedure provides the best-fit value of the signal strength measurement by numerically solving the likelihood equations

$$\frac{\partial \ln \mathcal{L}(\boldsymbol{\alpha})}{\partial \alpha} = 0 \text{ for all } \alpha \in \boldsymbol{\alpha}. \tag{8.21}$$

Its uncertainty σ_μ denotes a variation of $-2\ln\Lambda$ of one unit from its minimum. The best-fit value of the signal strength in data is labelled $\hat{\mu}^{\text{obs}}$ and the corresponding test statistic value is called q_μ^{obs}.

To quantify the level of agreement between the data and a given hypothesis, the test statistic can be transformed into a p-value. The p-value quantifies the probability of finding data that is equally or less compatible with the predictions of H_0, i.e. of finding a value of q_0 that is equal to or greater than the observed value q_0^{obs}. The corresponding confidence level (CL) is $\text{CL} = 1 - p$. The p-value is given by

$$p = \int_{q_0^{\text{obs}}}^{\infty} f(q_0|H_0). \tag{8.22}$$

One can define a corresponding significance Z where Z is the point in a normal distribution that corresponds to an equivalent area under the curve, as illustrated in Figure 8.15. In the following, the significance will serve as a measure of the (in)compatibility between the observed signal strength and the hypothesis of $\mu = 0$, in terms of number of standard deviations (σ_μ), and serves to quantify the probability of incorrectly rejecting the null hypothesis when it is in fact true. Certain conventions exist in high energy physics to assess whether a null hypothesis can be conclusively rejected. Two levels of confidence in particular serve as benchmarks. "Evidence" corresponds to a significance $Z = 3\sigma$ where the background-only hypothesis H_0 is rejected at a 99.7% confidence level, corresponding to a p-value of 0.003. "Observation", on the other hand, requires a larger incompatibility of the data and the null hypothesis. It is in general accepted that a hypothesis is rejected (or, that a process is observed) when the significance of the test exceeds 5σ. This corresponds to a p-value of about 3×10^{-7}.

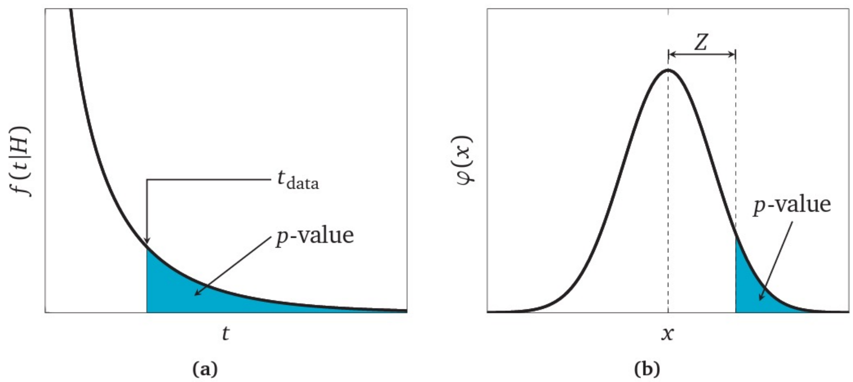

Figure 8.15: Relation between the p-value (left) and the corresponding significance Z (right) for a normal distribution. Images from [82].

8.6.8 Fit model for $t\bar{t}H \to$ multilepton

For the $t\bar{t}H \to$ multilepton analysis a total of 32 bins are fit. The bins are summarized in Table 8.5, with BDT score bins used in regions with large expected event yields, and single bins for control regions and the $3\ell + 1\tau_{\text{had}}$ and 4ℓ signal regions.

	$2\ell\text{SS} + 0\tau_{\text{had}}$	$3\ell + 0\tau_{\text{had}}$	4ℓ	$1\ell + 2\tau_{\text{had}}$	$2\ell\text{SS} + 1\tau_{\text{had}}$	$2\ell\text{OS} + 1\tau_{\text{had}}$	$3\ell + 1\tau_{\text{had}}$
BDT trained against	Fakes and $t\bar{t}V$	$t\bar{t}$,$t\bar{t}W$, $t\bar{t}Z$, VV	$t\bar{t}Z$/-	$t\bar{t}$	all	$t\bar{t}$	-
Discriminant	$2 \times$ 1D BDT	5D BDT	Event count	1D BDT	1D BDT	Event count	
Number of bins	6	5	1/1	2	2	10	1
Control regions	-	4	-	-	-	-	-

Table 8.5: Signal and control regions of the $t\bar{t}H \to$ multilepton analysis, with information about how they are used in the fit.

The fit is performed in each channel individually as well as globally for all the regions together. In Section 8.7, expected and observed values are presented for the signal strength, and the conclusions of the hypothesis test are quantified through the expected and observed values of the significance. The expected values correspond to those obtained in an Asimov dataset with $\mu = 1$. These values are presented with the data blinded in all regions with a signal-to-background ratio above 15%. The data is blinded using this criterion to ensure that knowledge of the data does not bias the design of the analysis. The observed values, on the other hand, correspond to those obtained with the data unblinded using the procedure described above.

8.6.8.1 Sources of uncertainty considered

Several NPs are considered in the $t\bar{t}H \to$ multilepton fit, with effects on the normalisation of the various contributions as well as on the distributions of the discriminating variables used to define each signal region. Aside from the uncertainties related to the data-driven estimates, discussed in detail in the case of fake $\tau_{\text{had-vis}}$ in Section 8.5, uncertainties related to experimental quantities (such as the efficiency of the reconstruction and identification of the objects considered or the calibration of their energy) as well as to theoretical predictions are also included.

Experimental uncertainties
The efficiencies of the electron, muon and $\tau_{\text{had-vis}}$ reconstruction and identification are measured in data using events containing a Z boson. The energy scale and resolution are also assessed in data using similar events as discussed in Chapter 5. The uncertainty on those measurements are included as NPs. For jets, the energy scale and resolution are also evaluated (this procedure is described in Section 5.4.3) and their uncertainties are considered. The efficiencies of the pile-up discrimination and of the b-tagging algorithms are measured in data and their uncertainties are also included. For the b-tagging algorithm, the light- and c-jet rejection rates are measured in data as well, with the c-jet rejection results assumed to also be valid for $\tau_{\text{had-vis}}$ (with an increase to the uncertainty on the measurement to account for the extrapolation). Finally, uncertainties related to the luminosity measurement in ATLAS are added.

Theoretical uncertainties
Theoretical uncertainties include effects related to the cross section calculations and to the MC event generation of the various signal and background processes. The components of the former consist of uncertainties related to the factorisation and renormalisation scales (jointly called QCD scale), the PDFs and the value of α_S used in the calculations. Uncertainties on the QCD scale are assessed by varying the scales by factors of 2 and of $\frac{1}{2}$, while those on the PDFs and on α_S are evaluated according to the recommendations of the PDF4LHC collaboration [13]. Uncertainties related to the event generation are assessed by comparing the predictions from the chosen samples to those of alternate samples. This includes samples with different parton shower simulation (for the $t\bar{t}H$ and $t\bar{t}V$ samples) as well as different matrix element computation (only for $t\bar{t}V$ contributions).

8.7 Results of the 36 fb^{-1} $t\bar{t}H \to$ multilepton analysis

8.7.1 Fit in the $3\ell + 1\tau_{\text{had}}$ channel

A blinded fit is performed in the $3\ell+1\tau_{\text{had}}$ signal region using an Asimov dataset with $\mu = 1$ to quantify the expected sensitivity of the channel (i.e. the expected significance of the associated test) and assess the impact of the many sources of uncertainty.

When only considering statistical uncertainties, the expected sensitivity of the $3\ell + 1\tau_{\text{had}}$ channel is of 0.86σ. When accounting for all relevant sources of uncertainty, the expected sensitivity of the channel falls to 0.84σ, corresponding to a signal strength $\mu_{t\bar{t}H}^{\text{expected}} = 1.0^{+1.6}_{-1.1}$ (stat.) $^{+0.5}_{-0.3}$ (syst.) $= 1.0^{+1.6}_{-1.2}$. As expected for a channel with a total of 4 expected events, the uncertainty on the $\mu_{t\bar{t}H}$ measurement in $3\ell + 1\tau_{\text{had}}$ is dominated by statistical uncertainties.

The leading sources of systematic uncertainty for the $3\ell + 1\tau_{\text{had}}$ channel are presented in Table 8.6. They include uncertainties related to τ_{had} identification and reconstruction efficiencies and to the fake

τ_{had} estimate, in particular the statistical uncertainty in the Monte Carlo simulation of the fake τ_{had} background in the $3\ell + 1\tau_{\text{had}}$ signal region, and the systematic uncertainty derived from the alternate $2\ell\text{OS}+1\tau_{\text{had}}$ extraction regions. The 100% systematic uncertainty assigned to fake and non-prompt light lepton contributions is also highly ranked. Finally, theoretical uncertainties in the calculations of the $t\bar{t}Z$ cross section and the signal cross section and parton shower are relevant, as well as the uncertainty on the luminosity measurement. All of these uncertainties have a very small impact on the measurement in this channel, as discussed.

Source of uncertainty	$\Delta\mu$	
$\tau_{\text{had-vis}}$ identification efficiency	+0.13	−0.11
$t\bar{t}Z$ cross section (QCD scale)	+0.09	−0.08
$t\bar{t}H$ cross section (QCD scale)	+0.09	−0.05
Fake $\tau_{\text{had-vis}}$ estimate (stat.)	+0.08	−0.08
$t\bar{t}H$ parton shower	+0.08	−0.08
Simulation sample size	+0.07	−0.07
Luminosity measurement	+0.06	−0.06
Fake and non-prompt light lepton estimates	+0.06	−0.06
Fake $\tau_{\text{had-vis}}$ estimate (composition)	+0.06	−0.06
$\tau_{\text{had-vis}}$ reconstruction efficiency	+0.06	−0.06

Table 8.6: Sources of systematic uncertainty with an observed impact of more than 5% on the measurement for the $3\ell + 1\tau_{\text{had}}$ channel.

With the data unblinded, a fit can be performed to obtain the best estimate of the observed value of $\mu_{t\bar{t}H}$. For the $3\ell + 1\tau_{\text{had}}$ signal region, the post-fit result is shown in Figure 8.16. The best-fit value of the signal strength in this channel is $\mu^{\text{obs}}_{t\bar{t}H} = 1.6^{+1.7}_{-1.3}\text{ (stat.) }^{+0.6}_{-0.2}\text{ (syst.) } = 1.6^{+1.8}_{-1.3}$.

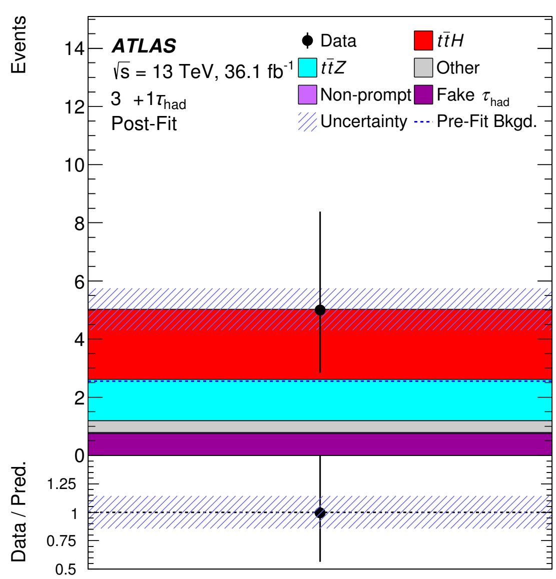

Figure 8.16: $3\ell + 1\tau_{\text{had}}$ signal region after the fit is performed. The Monte Carlo predictions are shown in colorful bands while the data yield is represented by a black dot. The $t\bar{t}H$ contribution (red) is scaled according to the fitted $\mu_{t\bar{t}H}$. Image from [58].

8.7.2 Fit in all $t\bar{t}H \to$ multilepton channels combined

All of the signal and control regions of the analysis are shown in Figure 8.17. The $t\bar{t}H$ contribution, in red, is scaled to the global best-fit result for the signal strength. The expected sensitivity is of 2.8σ corresponding to $\mu_{t\bar{t}H} = 1.00 \pm 0.3$ (stat.) ± 0.3 (syst.) $= 1.00 \pm 0.4$[12].

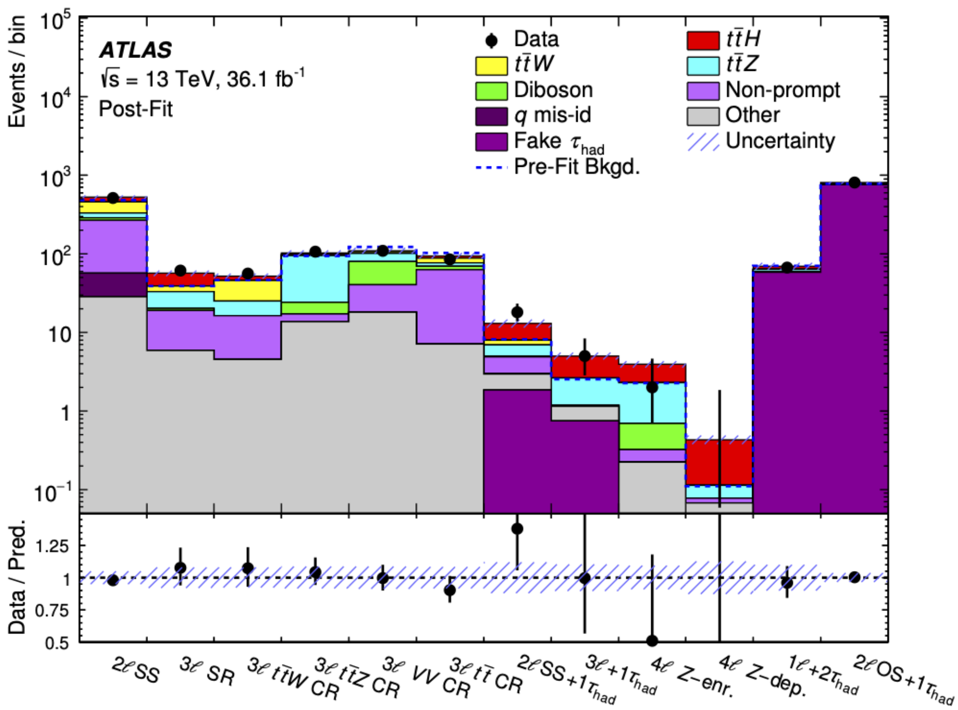

Figure 8.17: Comparison of the SM predictions to the data after the fit in the eight signal regions and four control regions of the $t\bar{t}H \to$ multilepton analysis. The $t\bar{t}H$ contribution (red) is scaled to the result of the fit. Image from [58].

The sources of uncertainty on the combined $t\bar{t}H \to$ multilepton measurement are broken down in Table 8.7. They indicate that while in the $3\ell + 1\tau_{\rm had}$ signal region uncertainties related to $\tau_{\rm had}$ identification and fake rate estimations are negligible contributions to the overall sensitivity due to the low statistics of the channel, these uncertainties have a larger impact on the overall result when considering all channels containing $\tau_{\rm had}$ as seen in Table 8.7 [13]. This observation motivates a more in-depth treatment of the fake $\tau_{\text{had-vis}}$ contribution in future rounds of analysis. Fake $\tau_{\text{had-vis}}$ estimations are indeed one of the main focuses of the studies leading to the subsequent $t\bar{t}H \to$ multilepton publication, which can be found in [63].

The most important sources of uncertainty on the combined $\mu_{t\bar{t}H}$ measurement are those that have a large impact on the measurements in the $2\ell\text{SS} + 0\tau_{\rm had}$ and $3\ell + 0\tau_{\rm had}$ channels since they are the most sensitive ones. They consist of the theoretical uncertainties on the $t\bar{t}H$ and $t\bar{t}V$ cross sections and

[12] In this case the expected sensitivity is computed assuming $\mu_{t\bar{t}H} = 1.00$ but using the unblinded fit result to assess the size of the uncertainties.

[13] These sources of uncertainty are also important within the $3\ell + 1\tau_{\rm had}$ fit when neglecting statistical uncertainties, indicating that when performing this analysis on a larger dataset where the statistical uncertainties are smaller, these sources of uncertainty will have a growing impact on the sensitivity of $3\ell + 1\tau_{\rm had}$ as well.

of uncertainties on the jet energy scale and resolution. The uncertainties related to $t\bar{t}H$ are separated into $t\bar{t}H$ cross section and $t\bar{t}H$ branching rations and parton shower. The former is only present in the $\mu_{t\bar{t}H}$ measurement and will not affect the cross section measurement discussed in Section 8.6.7. The latter however is related to the acceptance (i.e. the number of selected events in each signal region) and therefore will contribute to the $t\bar{t}H$ cross section measurement. Finally, the jet energy scale and resolution uncertainties affect both the $t\bar{t}H$ acceptance and the kinematic properties of the jets used by the BDT discriminants in the various analysis regions.

Source of uncertainty	$\Delta\mu$	
$t\bar{t}H$ cross section	+0.20	-0.09
Jet energy scale and resolution	+0.18	-0.15
Non-prompt light lepton estimates	+0.15	-0.13
Jet flavour tagging and τ_{had} identification	+0.11	-0.09
$t\bar{t}W$ theory	+0.10	-0.09
$t\bar{t}Z$ theory	+0.08	-0.07
Theory (other backgrounds)	+0.08	-0.07
Luminosity measurement	+0.08	-0.06
$t\bar{t}H$ branching ratios and parton shower	+0.08	-0.04
Fake τ_{had} estimates	+0.07	-0.07
Other experimental uncertainties	+0.05	-0.04
Simulation sample size	+0.04	-0.04
Charge mis-identification	+0.01	-0.01
Total systematic uncertainty	+0.39	-0.30

Table 8.7: Sources of uncertainty and their observed impact on the measurement of $\mu_{t\bar{t}H}$. Due to rounding effects and correlations between sources of uncertainty, the total systematic uncertainty is not identical to the sum in quadrature of the individual sources of uncertainty presented.

In the global fit to all of the regions of the analysis the signal strength measurement is $\mu^{\text{obs}}_{t\bar{t}H} = 1.6 \pm 0.3$ (stat.) $^{+0.4}_{-0.3}$ (syst.) $= 1.6^{+0.5}_{-0.4}$. This result corresponds to an observed significance of 4.1 standard deviations with respect to the background-only hypothesis, and indicates evidence for $t\bar{t}H$ production at the LHC. This result is consistent with the SM prediction.

As discussed in Section 8.6.5, NPs that deviate from their expected values can sometimes be cause for concern. In the case of this analysis, some tension is observed in the normalisation of the $t\bar{t}W$ and $t\bar{t}Z$ backgrounds between the theory expectation and the measurement carried out in the CRs of the $2\ell\text{SS}+0\tau_{\text{had}}$ and $3\ell+0\tau_{\text{had}}$ channels. Similarly, potential tensions between the $t\bar{t}H$ signal measurements in the different SRs must also be evaluated. A series of tests are made to cross-check the validity of the final $\mu_{t\bar{t}H}$ measurement in light of these comparisons:

- The agreement of the measurement with different assumptions on the normalisation of the main backgrounds is assessed. In particular, the normalisations of the $t\bar{t}W$ and $t\bar{t}Z$ backgrounds are allowed to be obtained from the control regions of the analysis instead of being set to the SM value. In this case, the NFs obtained are $\mu_{t\bar{t}W} = 0.92 \pm 0.32$ and $\mu_{t\bar{t}Z} = 1.17^{+0.25}_{-0.20}$. The use of these NFs results in a signal strength measurement of $\mu_{t\bar{t}H} = 1.57^{+0.57}_{-0.50}$ which is compatible with the original result indicating that the modelling of the normalisation of the main backgrounds is adequate in MC.

- The agreement of the measurement within signal regions is investigated. For this purpose the signal strength is fitted in the seven analysis channels separately. The measurements of the individual channels are found to be compatible with one another with a probability of 34%. The breakdown of these measurements is presented in Figure 8.18.

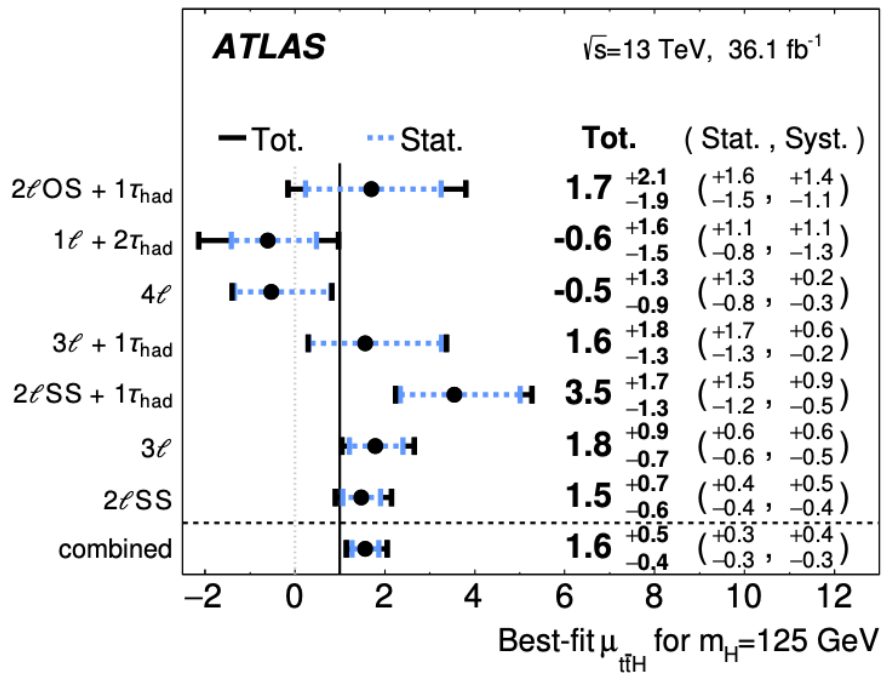

Figure 8.18: Observed $t\bar{t}H$ signal strength values for each channel and in the combination. The black vertical line represents a $t\bar{t}H$ production as predicted by the SM ($\mu_{t\bar{t}H} = 1$) while the dashed grey vertical line represents the background-only hypothesis ($\mu_{t\bar{t}H} = 0$). Image from [58].

For reference, the corresponding $t\bar{t}H \to$ multilepton analysis performed using 36 fb^{-1} of data from the run years 2015 and 2016 collected with the CMS detector found $\mu_{t\bar{t}H} = 1.5 \pm 0.5$ [121].

8.8 Observation of the $t\bar{t}H$ process

As discussed in Section 8.6.7, to claim that the $t\bar{t}H$ process is observed the background-only hypothesis must be rejected with a significance of at least 5σ. Projections were made to assess when and how this goal could be achieved. As mentioned at the beginning of this chapter, other analyses aim to study the $t\bar{t}H$ process with the Higgs decaying to objects such as b-quark pairs and two photons. Combining all $t\bar{t}H$ measurements performed with 36 fb^{-1} of data, as illustrated in Figure 8.19, results in a sensitivity of approximately 4σ. The 5σ threshold can only be reached when more data is added. By the end of 2017, the size of the available dataset had increased to 80 fb^{-1}. As shown in Figure 8.19, the observation of the $t\bar{t}H$ process could be achieved with a dataset of this size. Not all analyses could be easily updated to include the larger dataset, so in the end a combination of results with 36 fb^{-1} and 80 fb^{-1} was used.

When combining the results presented in Section 8.7 with the results for $t\bar{t}(H \to b\bar{b})$ using the same 36 fb^{-1} dataset [60], and $t\bar{t}(H \to \gamma\gamma)$ and $t\bar{t}(H \to ZZ^* \to 4\ell)$ using a dataset corresponding to an integrated luminosity of 80 fb^{-1} [54, 59], it is possible to measure the $t\bar{t}H$ production cross section with an observed (expected) significance of 5.8 (4.9) standard deviations with respect to the background-only

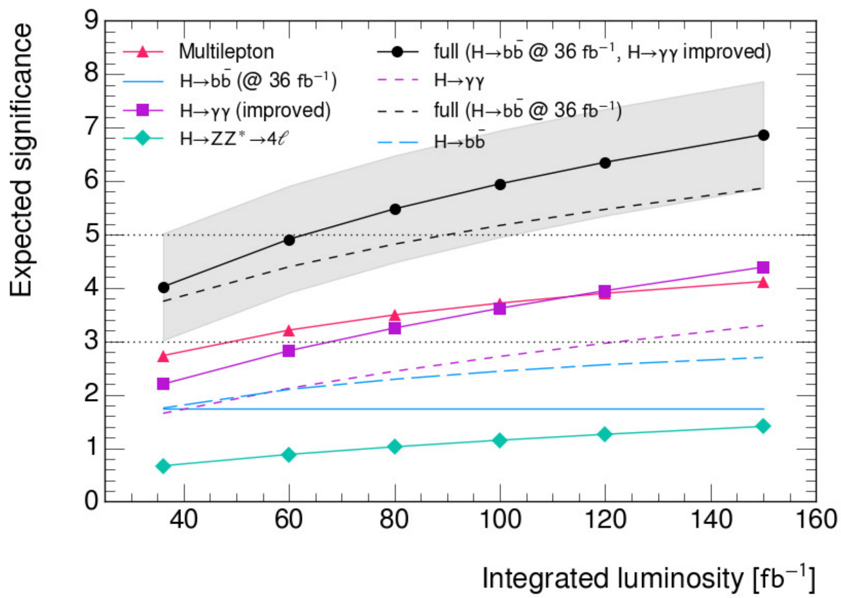

Figure 8.19: Expected significance for the $t\bar{t}H$ process as a function of integrated luminosity. The contribution from the $t\bar{t}H \to$ multilepton measurement is shown in red.

hypothesis [26]. When also combining this result with the $t\bar{t}H$ measurements from Run 1, at 7 and 8 TeV, the observed (expected) significance is of 6.3 (5.1) standard deviations [26]. With this result, it is possible to reject the background-only hypothesis in favour of the hypothesis where $t\bar{t}H$ production proceeds. The result can be interpreted as the observation of the $t\bar{t}H$ process.

The CMS collaboration also produced a set of results with data from the years 2015-2017. Both collaborations individually claimed observation of the $t\bar{t}H$ process [26, 76], with a probability of the measurement resulting from a background fluctuation of under 1 in 3.5 million.

Having established the existence of the $t\bar{t}H$ process, a $t\bar{t}H$ cross section measurement is also provided with a value of $\sigma_{t\bar{t}H} = 670 \pm 90$ (stat.)$^{+110}_{-100}$ (syst.) fb [26] in agreement with the SM prediction of 508 ± 45 (PDF+α_S) $^{+29}_{-47}$ (QCD scale) fb [98].

Using the new $t\bar{t}H$ measurement, the ratio of the $t\bar{t}H + tH$ cross section to the ggH cross section are again computed as illustrated in Figure 8.20. This new result is significantly more precise than its Run 1 counterpart, presented in Figure 7.6, and provides a much more sensitive probe of the contents of the ggH loop, which are observed to be consistent with the SM predictions.

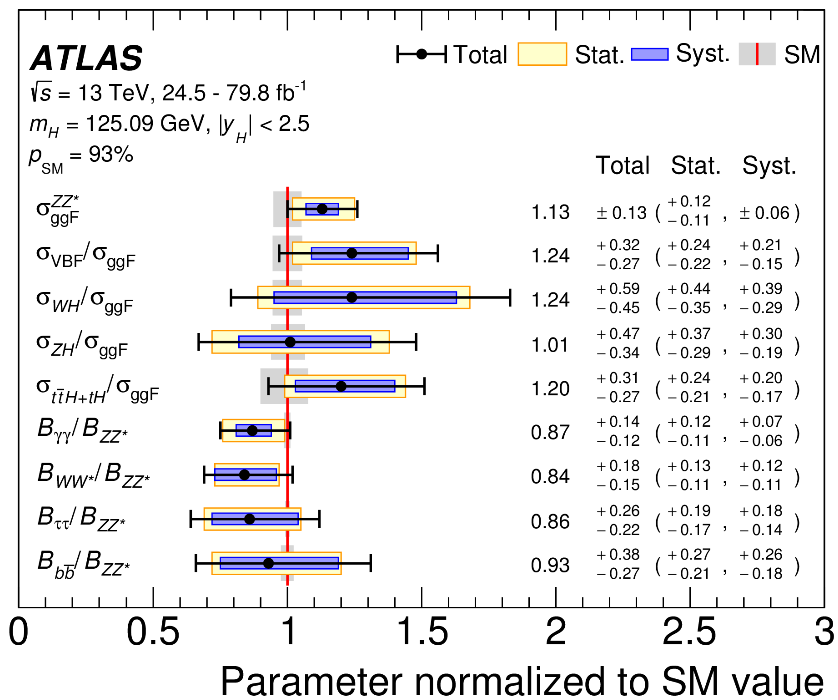

Figure 8.20: Ratios of cross sections (top of the image) and branching fractions (bottom of the image) with the Run 2 dataset. Image from [72].

Chapter 9

Higgs decays to τ leptons in other production modes

Couplings of the Higgs boson to leptons are in general difficult to measure due to the small masses of the leptons. The only coupling that can be precisely measured using the full Run 2 dataset is the one to the τ lepton, due to this lepton's relatively large mass. Sensitivity to this coupling is obtained from measurements of the $pp \rightarrow H \rightarrow \tau\tau$ cross section. The study of events with $H \rightarrow \tau\tau$ decays also provides sensitivity to the rates of the various Higgs production modes. In particular, cross section measurements obtained in the $H \rightarrow \tau\tau$ analysis can be used to study the rates of the ggH, VBF, VH and $t\bar{t}H$ processes.

As discussed in Section 6, τ leptons decay before they reach the ATLAS detector. Unlike other leptons, they can decay to both leptons and quarks due to their large mass. Leptonic decays, labelled τ_{lep}, refer to the process $\tau \rightarrow W\nu_\tau \rightarrow \ell + \bar{\nu}_\ell + \nu_\tau$. They are reconstructed as electrons or muons. Hadronic decays on the other hand are labelled τ_{had} and correspond to the process $\tau \rightarrow W\nu_\tau \rightarrow \text{hadron(s)} + \nu_\tau$. The visible decay products are reconstructed as $\tau_{\text{had-vis}}$[1]. In both cases neutrinos are present, so $H \rightarrow \tau\tau$ events are in general expected to be associated to non-negligible amounts of missing transverse energy.

Measurements of the $pp \rightarrow H \rightarrow \tau\tau$ cross section are made in three different final states: $H \rightarrow \tau_{\text{lep}}\tau_{\text{lep}}$, $H \rightarrow \tau_{\text{lep}}\tau_{\text{had}}$ and $H \rightarrow \tau_{\text{had}}\tau_{\text{had}}$. These final states are reconstructed respectively as $\ell\ell + E_{\text{T}}^{\text{miss}}$, $\ell + \tau_{\text{had-vis}} + E_{\text{T}}^{\text{miss}}$ and $\tau_{\text{had-vis}}\tau_{\text{had-vis}} + E_{\text{T}}^{\text{miss}}$. The production mode associated to the event adds additional information to the final state: ggH events are associated to clean final states and chosen when they produce boosted Higgs bosons. VBF events are associated to two jets in opposite hemispheres of the detector, while VH events are characterized by the presence of jets with $m_{jj} \approx m_V$ from the $V \rightarrow q\bar{q}$ decay. VH events with leptonically-decaying vector bosons are considered in a separate analysis, so in the following VH will only refer to events hadronically-decaying vector bosons. Finally, $t\bar{t}H$ events are identified through the presence of complex final states containing b-jets.

9.1 Previous measurements at the LHC

The Higgs discovery, achieved in the $H \rightarrow \gamma\gamma$, $H \rightarrow ZZ \rightarrow 4\ell$ and $H \rightarrow WW$ final states in ATLAS, was closely followed by direct evidence of Higgs boson couplings to fermions established in the $H \rightarrow \tau\tau$ channel. In combination with the equivalent CMS result, the $H \rightarrow \tau\tau$ process was observed with a

[1] These were the only τ considered in the $t\bar{t}H \rightarrow$ multilepton analysis.

significance of 5.5 σ [112], and was the first observation of Higgs couplings to leptons. Evidence for this process was re-established in Run 2 using a dataset corresponding to 36 fb^{-1} and its cross section was measured to be $3.77^{+0.60}_{-0.59}$ (stat.) $^{+0.87}_{-0.74}$ (syst.) pb [65], in agreement with the SM expectation of 3.46 ± 0.13 pb [97], as shown in Figure 9.1. A combination with the Run 1 ATLAS result allowed for single-experiment observation of $H \to \tau\tau$ with a significance of 6.4σ. These measurements were sensitive to different Higgs production modes (ggH and VBF) as well as to the Higgs branching ratio to $\tau\tau$, and therefore contributed to measurements of various Higgs coupling parameters.

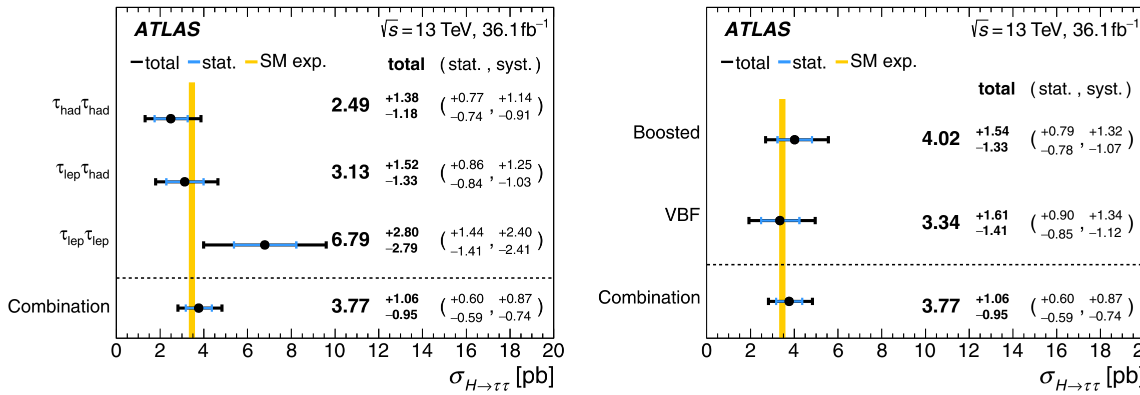

Figure 9.1: Measured cross section for $H \to \tau\tau$, split into the different τ decay modes on the left and into ggH (boosted) and VBF on the right. Images from [118].

As illustrated in Figure 9.2 the measurements were performed using the $m_{\tau\tau}^{\text{MMC}}$ [87], a di-τ mass estimator that aims to reconstruct the energy and direction of the neutrinos from τ decays using the kinematics of both the visible decay products of the τ leptons and the $E_{\text{T}}^{\text{miss}}$ of the event. This variable provides some separation from the dominant background, $Z \to \tau\tau$.

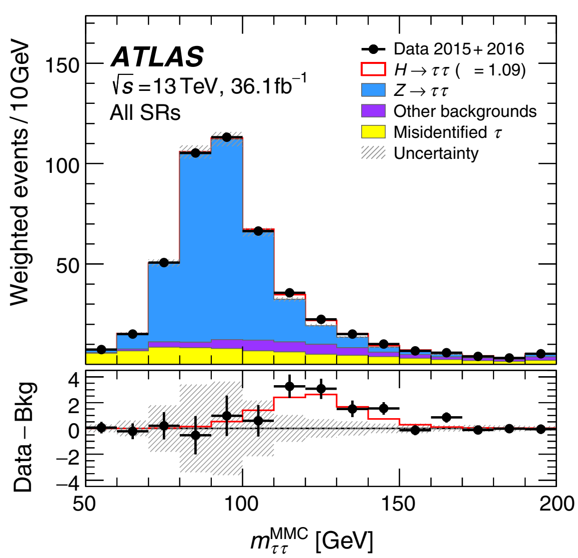

Figure 9.2: Reconstructed di-τ $m_{\tau\tau}^{\text{MMC}}$ mass in the $H \to \tau\tau$ signal regions. The $H \to \tau\tau$ contribution is shown in red. Image from [118].

Measurements of the $pp \to H \to \tau\tau$ cross section were also performed using the simplified template cross section (STXS) framework [99], as shown in Table 9.1. STXS measurements provide essential information about the behaviour of the cross section as a function of various kinematic and topological variables, and allow for more comprehensive tests of the SM as they provide sensitivity to processes that may have a small impact on the total cross section while significantly modifying the variable distributions. The STXS framework is specifically designed to provide cross section measurements for specific production modes in orthogonal regions of phase space defined by the event kinematics. While a fine binning of the phase space (such as the one in differential cross section measurements) provides excellent sensitivity to potential BSM effects, it can make the combination of results from channels targeting different Higgs decay modes challenging since the optimal binning might be different for each channel. The STXS binning is designed to enable easier combination while still isolating regions of phase space where BSM effects might be notable. The goal being to increase the number of bins to increase the sensitivity to BSM effects while at the same time maintaining good experimental sensitivity, the binning is adjusted as more data is gathered. Different binning configurations are defined in stages, with finer binning corresponding to later stages.

Process	Particle-level selection	σ [pb]	σ^{SM} [pb]
ggH	$N_{\text{jets}} \geq 1$, $60 < p_T^H < 120$ GeV, $\lvert y_H \rvert < 2.5$	$1.79 \pm 0.53(\text{stat}) \pm 0.74(\text{syst})$	0.40 ± 0.05
ggH	$N_{\text{jets}} \geq 1$, $p_T^H > 120$ GeV, $\lvert y_H \rvert < 2.5$	$0.12 \pm 0.05(\text{stat}) \pm 0.05(\text{syst})$	0.14 ± 0.03
VBF	$\lvert y_H \rvert < 2.5$	$0.25 \pm 0.08(\text{stat}) \pm 0.08(\text{syst})$	0.22 ± 0.01

Table 9.1: $pp \to H \to \tau\tau$ cross section measurements based on Stage 1 of the STXS framework. Table from [118].

The measurements performed by the CMS collaboration include the total $pp \to H \to \tau\tau$ cross section [119] as well as differential cross sections using the STXS framework [79] with 36 fb^{-1} and 80 fb^{-1} of data, respectively.

9.2 Analysis strategy for measurements with the full Run 2 dataset

Using the full Run 2 dataset, it is possible to expand upon the previous measurements and obtain:

1. A measurement of the total cross section of the $pp \to H \to \tau\tau$ process, which can later be re-interpreted as a measurement of the Yukawa coupling of the τ lepton.

2. A simultaneous measurement of the cross sections of the four main production modes (ggH, VBF, VH and $t\bar{t}H$) of the Higgs boson.

3. A measurement of the cross section using the STXS framework.

The larger dataset and improved particle reconstruction, identification and calibration techniques achieved at the end of Run 2 are extremely beneficial to this analysis. They allow for increased precision of all three types of measurement presented above and provide sensitivity to larger regions of the phase space, in particular regions associated to events where the Higgs boson is produced with high transverse momentum.

The three final states ($H \to \tau_{\text{lep}}\tau_{\text{lep}}$, $H \to \tau_{\text{lep}}\tau_{\text{had}}$ and $H \to \tau_{\text{had}}\tau_{\text{had}}$) are organized into three channels referred to as $\tau_\ell\tau_\ell$, $\tau_\ell\tau_{\text{had}}$ and $\tau_{\text{had}}\tau_{\text{had}}$, and are outlined in Table 9.2. This is because processes involving the Z boson contribute as the dominant background. In particular, the $Z \to \tau\tau$ process results in final states that are almost identical to those of $H \to \tau\tau$. The mass difference between the Z and Higgs bosons is especially difficult to resolve in this analysis due to the presence of neutrinos in their decays and so the two contributions overlap significantly. Since $Z \to \tau\tau$ is the largest irreducible background, an accurate estimate of its contribution is key for the $H \to \tau\tau$ measurements to be possible. The challenges related to its estimation are presented in detail in Section 9.5 and will be the focus of this chapter.

Previous rounds of analysis showed that the $Z \to \ell\ell$ process can also contribute as background. Events from $Z \to \ell\ell$ with large missing transverse energy due to jet miscalibrations contribute to the $\tau_\ell\tau_\ell$ signal region with rates roughly 15 times larger than those expected for the signal. These events are difficult to simulate given that detailed descriptions of jet miscalibrations and $E_{\text{T}}^{\text{miss}}$ would be necessary to properly estimate this contribution, and because the size of the available MC sample is very limited in the signal regions. These combined effects lead to uncertainties on the order of 20% in the $Z \to \ell\ell$ estimations. Consequently, it was decided to only study events with opposite flavour light leptons in the $\tau_\ell\tau_\ell$ signal region to reduce contributions from $Z \to ee$ and $Z \to \mu\mu$. The $\tau_\ell\tau_\ell$ channel is therefore labelled $\tau_e\tau_\mu$ in the rest of the text to indicate that only opposite flavour events are used.

	$\tau_e\tau_\mu$	$\tau_\ell\tau_{\text{had}}$	$\tau_{\text{had}}\tau_{\text{had}}$
$H \to \tau\tau$ branching ratio	6%	45.5%	42.5%
Number of neutrinos	4	3	2

Table 9.2: Characteristics of the analysis channels.

Each channel is further classified into one of three categories: boost (targeting ggH events), VBF and VH. The $\tau_{\text{had}}\tau_{\text{had}}$ channel additionally contains an all-hadronic $t\bar{t}H$ category, which is orthogonal to the categories presented in the $t\bar{t}H \to$ multilepton analysis of Chapter 8 as no light leptons are selected. The selections corresponding to each of these categories are outlined in Section 9.4 below.

9.3 Data and Monte Carlo samples

The dataset used consists of the full Run 2 dataset of proton-proton collisions at a centre of mass energy of 13 TeV. After quality requirements are applied for the LHC beam conditions and the data, the integrated luminosity of the dataset is of 139 fb^{-1}.

The various processes relevant to this analysis are simulated as described in Table 9.3.

9.3.1 Triggers

Light lepton triggers are used to select data events in the $\tau_e\tau_\mu$ and $\tau_\ell\tau_{\text{had}}$ channels while di-τ triggers are used for $\tau_{\text{had}}\tau_{\text{had}}$. In particular, $\tau_e\tau_\mu$ events are required to pass a single electron, a single muon or an electron-muon trigger while only single electron and single muon triggers are used in $\tau_\ell\tau_{\text{had}}$. These triggers are the same as those described for the $t\bar{t}H \to$ multilepton analysis. In $\tau_{\text{had}}\tau_{\text{had}}$, the di-$\tau$ trigger selects events with two $\tau_{\text{had-vis}}$ of at least 25 and 35 GeV. It also includes a requirement of a jet of

Process	Generator ME	Generator PS	PDF set ME	PDF set PS	Tune
$H \to \tau\tau/WW$					
ggF	Powheg	Pythia8	PDF4LHC15	CTEQ6L1	AZNLO
VBF	Powheg	Pythia8	PDF4LHC15	CTEQ6L1	AZNLO
VH	Powheg	Pythia8	PDF4LHC15	CTEQ6L1	AZNLO
$t\bar{t}H$	Powheg	Pythia8	NNPDF2.3		A14
Background					
$V+$jets	Sherpa 2.2.1		NNPDF30		Sherpa
$t\bar{t}$	Powheg	Pythia8	NNPDF2.3		A14
Single top	Powheg	Pythia8	NNPDF2.3		A14
Di-Boson	Sherpa 2.2.1		NNPDF30		Sherpa

Table 9.3: Overview of the MC generators used for the main signal and background samples.

at least 70 GeV[2] that is within $|\eta| < 3.2$, to ensure that low-momentum $\tau_{\text{had-vis}}$ can be selected while maintaining reasonable trigger rates. The addition of the jet requirement only results in an efficiency loss of under 0.5% for $H \to \tau\tau$ events.

9.4 Object and event selection

The Monte Carlo samples describing the signal and background processes are used to choose the best object and event selections for each channel and maximize the sensitivity to the $H \to \tau\tau$ process. Categories are additionally built to provide sensitivity to the various production modes considered.

9.4.1 Object selections

As in Chapter 8, a variety of object selections are defined for this analysis, to adapt to the specific signal and background composition of each channel and category. As before, object selections aim to decrease contributions from reducible backgrounds to the signal regions.

9.4.1.1 Electrons

Electron candidates in this analysis are reconstructed using the supercluster procedure outlined in Section 5.2.2. They are required to pass the loose identification criteria, have $p_{\text{T}} > 15$ GeV, $|\eta| < 2.47$ and be outside of the crack region ($1.37 < |\eta| < 1.52$). These requirements define the baseline electron candidates used in the overlap removal procedure described in Section 9.4.1.5.

Electrons surviving the procedure are additionally required to pass the medium identification and a set of loose calorimeter- and track-based isolation requirements to reduce contributions from jets misidentified as electrons. In events where trigger chains containing electrons are used, the p_{T} requirements are adjusted to ensure that triggers are used in a phase-space where they are maximally efficient, shown in Table 9.4 for the different data-taking periods.

[2] The requirement is of a 25 GeV jet at L1.

9.4.1.2 Muons

Muons are required to pass the loose identification criteria, have $p_{\mathrm{T}} > 10$ GeV and $|\eta| < 2.47$. The muons that satisfy these requirements are called baseline muons and are used in the overlap removal procedure.

Surviving muons are required to pass a tight track-based isolation selection. Similarly to electrons, muons are required to pass more stringent p_{T} requirements in channels where muons are used in the trigger chains (see Table 9.4).

9.4.1.3 Hadronically-decaying τ leptons

Hadronically-decaying τ leptons are required to have one or three associated tracks, $p_{\mathrm{T}} > 20$ GeV and $|\eta| < 2.47$ excluding the crack region, and pass a very loose selection on the $\tau_{\text{had-vis}}$ identification RNN. A muon overlap removal selection is also applied. In the $\tau_\ell\tau_{\text{had}}$ channel, only the highest p_{T} object passing these selections is considered as a potential $\tau_{\text{had-vis}}$ candidate to simplify the fake estimation procedure. This baseline $\tau_{\text{had-vis}}$ definition is used for the overlap removal procedure.

For $\tau_{\text{had-vis}}$ that survive this procedure, additional requirements are placed. The $\tau_{\text{had-vis}}$ identification RNN criteria are increased to those of the medium working point to reduce contributions from jets misidentified as $\tau_{\text{had-vis}}$. In the $\tau_e\tau_{\text{had}}$ channel, $\tau_{\text{had-vis}}$ objects are also required to pass the loose working point of the electron veto BDT discussed in Section 6.3 to further reject $Z \to ee$ events, and must satisfy $p_{\mathrm{T}} > 30$ GeV. In the $\tau_{\text{had}}\tau_{\text{had}}$ channel, the p_{T} selections for the $\tau_{\text{had-vis}}$ are adjusted as in Table 9.4 due to their use in the trigger chains.

Trigger chain	Data period	HLT p_{T} threshold	Offline p_{T} threshold
Single electron	2015 2016-2018	p_{T} (e) > 24 GeV p_{T} (e) > 26 GeV	p_{T} (e) > 25 GeV p_{T} (e) > 27 GeV
Single muon	2015 2016-2018	p_{T} (μ) > 20 GeV p_{T} (μ) > 26 GeV	p_{T} (μ) > 21 GeV p_{T} (μ) > 27.3 GeV
Electron-muon	2015-2018	p_{T} (e) > 17 p_{T} (μ) > 14 GeV	p_{T} (e) > 18 p_{T} (μ) > 14.7 GeV
Ditau	2015-2018	p_{T} (leading $\tau_{\text{had-vis}}$) > 35 GeV p_{T} (subleading $\tau_{\text{had-vis}}$) > 25GeV	p_{T} (leading $\tau_{\text{had-vis}}$) > 40 GeV p_{T} (subleading $\tau_{\text{had-vis}}$) > 30 GeV

Table 9.4: The p_{T} thresholds applied at trigger level and offline for the selected electrons, muons and $\tau_{\text{had-vis}}$.

9.4.1.4 Jets and b-jets

Jets in this analysis are reconstructed using the particle flow algorithm and required to have $p_{\mathrm{T}} > 20$ GeV and $|\eta| < 4.5$.

If they survive the overlap procedure, the jets must additionally satisfy a medium requirement on the JVT score to reduce contributions from pile-up jets as the presence of pile-up jets decreases the performance of the $m_{\tau\tau}^{\text{MMC}}$[3]. Jets from decays of heavy-flavour quarks are identified using the DL1

[3] Pile-up jets selected in the analysis are included in the computation of the missing transverse energy. This introduces

tagger presented in Section 5.6, with working points of 85% efficiency in the $\tau_e\tau_\mu$ and $\tau_\ell\tau_{\text{had}}$ channels and 70% efficiency in the $\tau_{\text{had}}\tau_{\text{had}}$ channel corresponding to rejections of approximately 30 (3) and 300 (10) for light (c) jets [52].

9.4.1.5 Overlap removal

As in Chapter 8, an overlap removal procedure is performed to ensure that tracks and calorimeter clusters are not used to reconstruct more than one object. Table 9.5 presents the selection that is applied for this purpose.

Step	Reject	Against	Criteria
1	electron	electron	Shared track, $p_T^1 < p_T^2$
2	$\tau_{\text{had-vis}}$	electron	$\Delta R < 0.2$
3	$\tau_{\text{had-vis}}$	muon	$\Delta R < 0.2$
4	muon	electron	Is calo muon and shared ID track
5	electron	muon	Shared ID track
6	jet	electron	$\Delta R < 0.2$
7	electron	jet	$\Delta R < 0.4$
8	jet	muon	Number of tracks < 3 and ghost-associated or $\Delta R < 0.2$
9	muon	jet	$\Delta R < 0.4$
10	jet	$\tau_{\text{had-vis}}$	$\Delta R < 0.2$

Table 9.5: Overlap removal selection. The selection criteria are applied sequentially as ordered in this table.

9.4.1.6 Missing transverse energy

The missing transverse energy of each event is computed using the baseline electrons, muons and $\tau_{\text{had-vis}}$, and all of the jets defined above. An additional selection is added to ensure that the transverse momentum of forward jets (i.e. jets with $|\eta| > 2.5$) considered is above 30 GeV. The purpose of this selection is also to reduce contributions from pile-up jets, which the $m_{\tau\tau}^{\text{MMC}}$ is sensitive to.

9.4.2 Event selections

The event selections for the three channels and the categories targeting ggH, VBF, VH and $t\bar{t}H$ final states are shown in Tables 9.6 and 9.7. The purpose of these selections is to ensure optimal sensitivity to the Higgs boson signal as well as a reliable estimate of its mass via the reconstruction of the invariant mass of the di-τ system using the $m_{\tau\tau}^{\text{MMC}}$, and harmonisation with the STXS framework binning.

significant biases to the $E_{\text{T}}^{\text{miss}}$ calculation and results in incorrect neutrino kinematic estimates in the $m_{\tau\tau}^{\text{MMC}}$ leading to an incorrectly-estimated mass.

	$\tau_e\tau_\mu$	$\tau_\ell\tau_{\text{had}}$		$\tau_{\text{had}}\tau_{\text{had}}$
		$e\tau_{\text{had}}$	$\mu\tau_{\text{had}}$	
Preselection				
Object counting	1 e, 1 μ, 0 $\tau_{\text{had-vis}}$	1 e or μ, 1 $\tau_{\text{had-vis}}$		0 e and 0 μ, 2 $\tau_{\text{had-vis}}$
p_{T} cut	e/μ : 10 GeV to 27.3 GeV	e/μ : 21 GeV to 27.3 GeV, $\tau_{\text{had-vis}}$: 30 GeV		$\tau_{\text{had-vis}}$: 40 ,30 GeV
Identification, isolation, and eveto	e/μ: medium e: loose, μ: track-based tight	e/μ: medium, $\tau_{\text{had-vis}}$: medium e: loose, μ: track-based tight 1-prong $\tau_{\text{had-vis}}$: eleBDT e-veto		$\tau_{\text{had-vis}}$: medium
Charge product	Opposite charge	Opposite charge		Opposite charge
Kinematics	$m_{\tau\tau}^{\text{coll}} > m_Z - 25$ GeV $30 < m_{e\mu} < 100$ GeV	$m_T > 70$ GeV		
b-veto	0 b-jets DL1-based, 85% efficiency working point (wp)	0 b-jets DL1-based, 85% efficiency wp		0 b-jets DL1-based, 70% efficiency wp not applied in tt(0L)$H \to \tau_{\text{had}}\tau_{\text{had}}$
$E_{\text{T}}^{\text{miss}}$	$E_{\text{T}}^{\text{miss}} > 20$ GeV	$E_{\text{T}}^{\text{miss}} > 20$ GeV		$E_{\text{T}}^{\text{miss}} > 20$ GeV
Leading jet	$p_{\text{T}} > 40$ GeV	$p_{\text{T}} > 40$ GeV		$p_{\text{T}} > 70$ GeV, $\|\eta\| < 3.2$
Angular	$\Delta R_{e\mu} < 2.0$, $\|\Delta\eta_{e\mu}\| < 1.5$	$\Delta R_{\ell\tau} < 2.5$, $\|\Delta\eta_{\ell\tau}\| < 1.5$		$0.6 < \Delta R_{\tau\tau} < 2.5$ $\|\Delta\eta_{\tau\tau}\| < 1.5$
Coll. app. x_1/x_2	$0.1 < x_1 < 1.0$, $0.1 < x_2 < 1.0$	$0.1 < x_1 < 1.4$, $0.1 < x_2 < 1.2$		$0.1 < x_1 < 1.4$, $0.1 < x_2 < 1.4$
specific cut for tt(0L)$H \to \tau_{\text{had}}\tau_{\text{had}}$				At least 6 jets and at least 1 b-jet or at least 5 jets and at least 2 b-jets

Table 9.6: Summary of the event selection for all sub-channels.

Category	Selection
VBF inclusive	sub-leading jet $p_{\text{T}} > 30$ GeV $m_{jj} > 350$ GeV, $\|\Delta\eta_{jj}\| > 3$ $\eta(j_0) \times \eta(j_1) < 0$ lepton centrality: visible decay products of the τ leptons between VBF jets
VH inclusive	$60\text{ GeV} < m_{jj} < 120$ GeV sub-leading jet $p_{\text{T}} > 30$ GeV
Boost inclusive	Not VBF inclusive Not VH inclusive $p_{\text{T}}\ H > 100$ GeV

Table 9.7: Summary of the event selection for the VBF, VH and boost categories.

9.4.2.1 Decay channels

The three channels, designed to separate events containing leptonic and hadronic τ decays, are defined as follows:

$\tau_e\tau_\mu$: Events must contain exactly one reconstructed electron and one reconstructed muon, their electric charges must be opposite, and they must have an invariant mass between 30 and 100 and a collinear mass[4] higher than $m_Z - 25$ GeV to reduce contributions from processes involving vector bosons and top quarks. Events containing b-tagged jets at the 85% efficiency working point are rejected to further reduce contributions from top processes. Finally the two light leptons are required to have separations of $\Delta R < 2.0$ and $|\Delta\eta| < 1.5$ to reduce contributions from multi-jet processes. To improve the $m_{\tau\tau}^{\mathrm{MMC}}$ mass estimation, the missing transverse energy of the event is required to be of at least 20 GeV and the momentum fraction carried by each $\tau_{\text{had-vis}}$[5] is required to be between 0.1 and 1.0.

$\tau_\ell\tau_{\text{had}}$: Events must contain exactly one reconstructed electron or muon and exactly one $\tau_{\text{had-vis}}$, again requiring their electric charges to be opposite. To suppress processes with W bosons, the transverse mass of the lepton-plus-$\tau_{\text{had-vis}}$ system must be at least 70 GeV. Events containing b-tagged jets at an 85% efficiency working point are again rejected to reduce contributions from top processes, and the lepton and $\tau_{\text{had-vis}}$ are required to have separations of $\Delta R < 2.5$ and $|\Delta\eta| < 1.5$. Due to the $Z \to ee$ contributions to the $e\tau_{\text{had}}$ channel, τ candidates are required to pass an electron veto BDT selection to suppress contributions from electrons misreconstructed as a $\tau_{\text{had-vis}}$. As in the $\tau_e\tau_\mu$ channel, the missing transverse energy of the event is required to be of at least 20 GeV, while the momentum fraction carried out by each $\tau_{\text{had-vis}}$ is required to be between 0.1 and 1.4 and that carried out by the light lepton is required to be between 0.1 and 1.2.

$\tau_{\text{had}}\tau_{\text{had}}$: Events must contain exactly two $\tau_{\text{had-vis}}$ with opposite electric charges. Events containing b-tagged jets at the 70% efficiency[6] working point are rejected to reduce contributions from top processes (except for the $t\bar{t}H$ regions where at least 5 (or 6) jets and at least 2 (or 1) b-tagged jets are required), and the two leptons are required to have separations of $0.6 < \Delta R < 2.5$ and $|\Delta\eta| < 1.5$. To improve the $m_{\tau\tau}^{\mathrm{MMC}}$ mass estimation, the missing transverse energy of the event is required to be of at least 20 GeV as before and the momentum fraction carried out by each $\tau_{\text{had-vis}}$ is required to be between 0.1 and 1.4. Finally, the presence of a jet with a transverse momentum of at least 70 GeV is required to ensure full efficiency of the trigger chain used in this channel.

9.4.2.2 Production categories

For each channel, three categories are defined to gain sensitivity to the Higgs production modes. The VBF and VH categories are described below, while the boost category simply aims to gather events from the ggH process via a requirement on the transverse momentum of the Higgs boson. The boost region is explicitly required to be orthogonal to the VBF and VH regions.

VBF : In VBF events, Higgs bosons are produced alongside two jets from the hard scatter of the event. These jets are produced with high energies and a large rapidity gap, leading to a high invariant mass of

[4]The collinear mass is an invariant mass computed under the assumptions that the τ decay products travel in the same direction as the original object and that the missing transverse energy of the event comes only from the neutrinos produced in τ decays. This assumption generally holds for boosted Higgs and Z decays, but fails in the case of large angular separation between the two τ. Since it is simple to compute the collinear mass is used at this stage to loosely select events, but it is replaced by the $m_{\tau\tau}^{\mathrm{MMC}}$ in all other stages of the analysis, as the latter provides a more reliable invariant mass reconstruction.

[5]The momentum fraction is defined as $x_i = \frac{p_{\mathrm{T}}^i}{p_{\mathrm{T}}^i + p_{\mathrm{T}}^{\mathrm{miss}}}$ where i represents the visible decay products. It is computed under the collinear assumption. Values larger than 1 are associated to events where this assumption fails.

[6]This working point used is looser in the $\tau_e\tau_\mu$ and $\tau_\ell\tau_{\text{had}}$ channels since there is no $t\bar{t}H$ category in those channels.

the di-jet system[7]. Events in this category are required to have $m_{jj} > 350$ GeV, and the subleading jet must have a transverse momentum of at least 30 GeV. The product of the pseudorapidity of the two leading jets is required to be negative to ensure that the two jets are located in different hemispheres of the detector. Additionally, a pseudorapidity gap of $|\Delta\eta| > 3.0$ is required between them. Finally, the visible τ decay products are required to be within this gap.

VH : Events where the vector boson decays leptonically can be efficiently selected and are the subject of a dedicated analysis. Events where the vector bosons decay as $V \to q\bar{q}$ are investigated here instead. Events in this category are required to contain a di-jet system satisfying $60 > m_{jj} > 100$ GeV to target the $V \to q\bar{q}$ decay, with the subleading jet required to have a transverse momentum of at least 30 GeV.

To increase the sensitivity of the measurement to the Higgs production mode, both categories are further sub-divided using dedicated BDT discriminants (also referred to as taggers) based on the jet activity of the event as well as the characteristics of the Higgs system. Examples of variables used include the separation in η and ϕ between the two leading jets, the ratio of the transverse momentum of the Higgs boson to the leading dijet system and the transverse momentum of the Higgs-plus-dijet system. Two regions are created in each case, with VBF 1 corresponding to the more VBF-like region and VBF 0 being slightly more consistent with ggH events (and similarly for VH).

The $t\bar{t}H$ categories are only defined for the $\tau_{\text{had}}\tau_{\text{had}}$ channel to avoid any overlap with the signal regions of the $t\bar{t}H \to$ multilepton analysis. Because top and Z processes are important backgrounds to the $t\bar{t}H$ category, it is sub-divided using two BDT discriminants to distinguish the Higgs signal from that of $t\bar{t}$ and $Z \to \tau\tau$. Two categories (ttH 0 and ttH 1) are created by applying selections on both discriminants, with ttH 1 being purer in signal contributions.

The boost category, orthogonal to the VBF, VH and $t\bar{t}H$ categories described above, is also sub-divided to follow the recommended STXS binning. The reconstructed transverse momentum of the Higgs boson is defined as the vector sum of the di-τ system and the $E_{\text{T}}^{\text{miss}}$, and events in the boost category are classified according to its magnitude, as well as according to the number of jets of at least 30 GeV that are present in the event as shown in Table 9.8.

$N_{\text{jets}} > 30$GeV	$p_{\text{T}}(H)$ bins in GeV 100, 120	120, 200	200, 300	300, ∞
Exactly 1	boost_0_1J	boost_1_1J	boost_2_1J	boost_3_1J
At least 2	boost_0_ge2J	boost_1_ge2J	boost_2_ge2J	boost_3_ge2J

Table 9.8: Definition of the eight categories in the boost phase space.

Each of the three channels is therefore described by 12 categories (8 for boost, 2 for VBF and 2 for VH). The $\tau_{\text{had}}\tau_{\text{had}}$ channel contains an additional 2 categories for the $t\bar{t}H$ selection. In total, 38 categories are considered for the cross section measurement.

[7] Since $m_{j1,j2} \propto p_{\text{T,j1}} p_{\text{T,j2}}(\cosh(\eta_{j1} - \eta_{j2}) - \cos(\phi_{j1} - \phi_{j2}))$ and thus $m_{j1,j2} \propto \eta_{j1} - \eta_{j2}$.

9.5 Estimation of the $Z \to \tau\tau$ background

The $Z \to \tau\tau$ process is the dominant background for most channels of the $H \to \tau\tau$ cross section measurement, representing 70-90% of background contributions depending on the signal region. The sensitivity of the $H \to \tau\tau$ measurements therefore relies on a detailed and accurate modelling of the $Z \to \tau\tau$ process. However, the similarity between the two processes makes it difficult to create a region where the modelling and normalisation of $Z \to \tau\tau$ can be scrutinised. For this reason, in Run 1 an embedding method had been used to estimate $Z \to \tau\tau$ distributions using $Z \to \ell\ell$ events [42]. This method consisted in selecting $Z \to \mu\mu$ events from data and replacing the reconstructed muons in those events with the reconstructed tracks and calorimeter deposits of τ leptons from simulated Z decays. The original muons were used as a proxy for the undecayed τ leptons, and the kinematics of the τ leptons were matched to those of the muons. The event was then re-reconstructed to account for the new objects and enable proper treatment of quantities such as the missing transverse energy. This method allowed for the use of simulation for processes that are well understood and modelled, such as Z boson and τ lepton decays, while relying on data for the description of other observables, such as the kinematics of the jets in the event. The latter in particular has been observed to be poorly modelled in the $H \to \tau\tau$ analysis as well as in dedicated Z production cross section measurements (see for example [117]).

The embedding as it was performed in Run 1 met significant hurdles in being deployed for Run 2. This was the unfortunate side effect of mandatory changes to the reconstruction software and data management to handle the unprecedented flow of data.

In the $H \to \tau\tau$ analysis using the 2015-2016 dataset [118], the modelling of important variables for the analysis was investigated in a validation region (VR) pure in $Z \to \ell\ell$ events, where ℓ refers to either an electron or a muon, and given lepton universality the conclusions regarding the overall normalisation and the jet activity of the event were assumed to also be valid for events with leptonically- and hadronically-decaying τ leptons. However, while all the decay products of the $Z \to \ell\ell$ events can be reconstructed, the decays of the tau leptons from $Z \to \tau\tau$ events result in the production of visible decay products (typically an electron, a muon, or one or three charged pions and up to two neutral pions) as well as neutrinos, which cannot be detected. Since region definitions in the analysis were only based on the properties of the visible decay products, differences were expected in the distributions of $Z \to \ell\ell$ and $Z \to \tau\tau$ events resulting in mismatched phase-spaces between the signal regions and the $Z \to \ell\ell$ VRs and raising questions about the reliability of the extrapolations into the signal regions.

For the analysis of the full Run 2 dataset, a slightly different approach based on the Run 1 embedding method has been developed to address this issue. This method, referred to as object-level embedding, aims to correct for the differences between $Z \to \ell\ell$ and $Z \to \tau\tau$ events by using the light leptons of $Z \to \ell\ell$ events[8] to estimate the kinematics of the truth τ before they are decayed, and then scaling their momentum down by the average energy lost to neutrinos in a τ decay using parametrised functions. In this way, the $Z \to \ell\ell$ events are transformed to be kinematically similar to $Z \to \tau\tau$ events without having to replace the light leptons by fully simulated τ lepton decays. While it is possible to also correct the direction of the objects, the choice to correct only the momentum was made for simplicity. Object-level embedding also accounts for effects related to the efficiency of the reconstruction and identification of τ decay products. Quantities that depend on the momenta of the leptons in the event, such as the missing transverse energy, are also re-scaled to account for the momentum corrections.

[8]The triggering, reconstruction and identification efficiencies of these light leptons are unfolded to ensure that they are as similar as possible to truth leptons.

Object-level embedding has certain advantages over the original embedding method:

- The cell-by-cell subtraction process in the Run 1 approach prevents the usage of $Z \to ee$ as it is very challenging to accurately remove the electron deposits in the calorimeters. This is however much simpler with the object-level embedding since no such subtraction is performed: the electron energy is simply rescaled. Considering $Z \to ee$ events in addition to $Z \to \mu\mu$ events increases the size of the embedded sample by approximately a factor two.

- The implementation can be performed at the latest stage of the event reconstruction workflow, minimizing the computing time and allowing for efficient reprocessing if needed.

Drawbacks of the object-level embedding method include the fact that the momentum correction parametrisations are subject to limited statistics, and that it is impossible to correctly apply object overlap removal criteria to the corrected objects since the event is not re-reconstructed: while the object-level embedding procedure can correct the momenta of muons from $Z \to \mu\mu$ events to match those of τ leptons from $Z \to \tau\tau$ events, the events will have muon-related overlap removal selections that cannot be removed and replaced by τ-appropriate ones.

Using the object-level embedding procedure, the distributions from the $Z \to \ell\ell$ control regions are corrected to match those of the $Z \to \tau\tau$ background in the signal regions. The data and MC in the control regions are then compared to assess the modelling of input variables to the multi-variate discriminants used to define signal regions as before, and to constrain the normalisation of the $Z \to \tau\tau$ background in the $H \to \tau\tau$ fits.

Object-level embedding will simply be referred to as embedding in the rest of this thesis for simplicity, with the understanding that the name refers to the new method and not its Run 1 counterpart.

9.5.1 Selections (pre-embedding)

Since this method aims to compare the $Z \to \ell\ell$ distributions in the control regions to those of the $Z \to \tau\tau$ background present in the various $H \to \tau\tau$ signal regions, all events used for this procedure are required to be kinematically similar to the events from those signal regions. Three types of regions are defined: signal regions, source regions and control regions. The signal regions are the ones from the $H \to \tau\tau$ analysis and are used as reference throughout this chapter to illustrate the $Z \to \tau\tau$ distributions. The source regions refer to the regions rich in $Z \to \ell\ell$ events that will be used for the embedding procedure, while the control regions refer to those same regions after the embedding procedure is applied. Each signal region is matched (and compared) to a control region that aims to replicate its phase space.

9.5.2 $Z \to \ell\ell$ source region

A $Z \to \ell\ell$ source region (i.e. the $Z \to \ell\ell$ region before embedding) must first be defined through a selection that increases the statistics of the final control regions and that ensures that the required phase-space is available after the embedding procedure. Consequently, looser selections must be used for any variable that is transformed by the embedding procedure. Once the procedure is performed, a new set of selections is applied to the $Z \to \ell\ell$ events to ensure that the final selections match the ones applied to $Z \to \tau\tau$ in the signal regions.

The selections for the source region are summarised in Table 9.9. Only events with exactly two muons or two electrons of opposite charge are selected. They must pass the single lepton triggers from

Table 9.4. The requirements on their transverse momenta are of at least 15 GeV for electrons and 10 GeV for muons, with quality, isolation and identification requirements identical to those in the $\tau_e\tau_\mu$ channel. To select contributions from $Z \to \ell\ell$ the invariant mass of the two light leptons is required to be greater than 80 GeV[9]. Events in the source region are additionally required to contain at least one jet with $p_{\mathrm{T}} > 40$ GeV. No requirement is placed on the number of b-tagged jets, the missing transverse energy of the event or the momentum fractions at this stage since the region is already very pure in $Z \to \ell\ell$.

	ee	$\mu\mu$
Object counting	2 electrons	2 muons
Transverse momentum	15 GeV to 27 GeV depending on whether the e is trigger-matched	10 GeV to 27.3 GeV depending on whether the μ is trigger-matched
Identification Isolation	medium loose	medium track-based tight
Trigger	At least one e satisfies the single-electron trigger	At least one μ satisfies the single muon trigger
Charge product	opposite charge	opposite charge
Kinematics	$m_{ee} > 80$ GeV	$m_{\mu\mu} > 80$ GeV
Leading jet	> 40 GeV	> 40 GeV
Angular selections	$\Delta R_{ee} < 2.5$ $\lvert\Delta\eta_{ee}\rvert < 1.5$	$\Delta R_{\mu\mu} < 2.5$ $\lvert\Delta\eta_{\mu\mu}\rvert < 1.5$

Table 9.9: Summary of the event selection for $Z \to \ell\ell$ events before embedding.

Finally, a random number is generated uniformly between 0 and 1 and assigned to each event. It is used to split the events from the source region into the three channels as follows: a third of the events are assigned to the $\tau_e\tau_\mu$ channel, a third of the events to the $\tau_\ell\tau_{\mathrm{had}}$ channel (half to $e\tau_{\mathrm{had}}$ and half to $\mu\tau_{\mathrm{had}}$) and the remaining third is assigned to the $\tau_{\mathrm{had}}\tau_{\mathrm{had}}$ channel[10]. The events assigned to each channel are then re-weighted according to the appropriate branching ratios. In the $\tau_e\tau_\mu$ and $\tau_{\mathrm{had}}\tau_{\mathrm{had}}$ channels, leading lepton refers to the lepton with the highest transverse momentum. In the $\tau_\ell\tau_{\mathrm{had}}$ channel, the τ_{had} of the event is always labelled as the leading lepton while the subleading lepton always refers to the light lepton, regardless of which object has the highest transverse momentum. This convention was chosen to ensure that τ_{had} and light lepton kinematics do not appear mixed in plots.

[9] This selection introduces a bias since in the signal regions, the $Z \to \tau\tau$ contribution has no such selection, but the impact is small (it is measured to be approximately 2-3%) and it is therefore ignored in this procedure.

[10] This was seen to be sub-optimal as many more $\tau_e\tau_\mu$ events survive the embedding procedure, so this choice should be re-evaluated in future iterations.

9.5.3 Object-level embedding procedure

In order to estimate $Z \to \tau\tau$ distributions it is necessary to obtain information about the two main processes at hand: the (truth) kinematics of the Z boson decay and the reconstruction, identification and triggering effects related to τ decays. The former is difficult to achieve using $Z \to \tau\tau$ events as the leptons produced in τ decays do not allow for a full reconstruction of the Z boson decay: some of the information required is lost through the production of neutrinos in τ decays. However, the study can be achieved through the use of $Z \to \ell\ell$ events as the reconstructed objects in $Z \to \ell\ell$ offer a direct measurements of the truth leptons from the Z boson decay[11]. The latter requires studies of both leptonic and hadronic τ decays, in particular with regards to the energy that is lost to neutrinos during the decay process, and can be performed using $Z \to \tau\tau$ events from Monte Carlo[12]. It is therefore possible to estimate $Z \to \tau\tau$ distributions by using $Z \to \ell\ell$ events as a proxy for the truth Z decay, and then applying an embedding procedure to those events to emulate the τ decays once the relevant effects are well understood.

The embedding procedure is performed as follows:

1. In a first step, to estimate the kinematics of the truth decay, effects related to the reconstruction, identification and triggering of the light leptons are compensated for in reconstructed $Z \to \ell\ell$ events. This step is called unfolding.

2. Then, the embedding procedure itself is applied. This procedure consists of parametrising the amount of energy that an average τ loses to neutrinos. These parametrisations are derived separately for τ leptons decaying leptonically to e, leptonically to μ, and hadronically as $\tau_{\text{had-vis}}$, as a function of the total transverse momentum of the truth τ. The embedding parametrisation functions are computed using $Z \to \tau\tau$ Monte Carlo simulation by matching truth τ to their associated visible decay products and assessing the fraction of the original truth tau momentum (with neutrinos) that is carried by the reconstructed visible decay products (e, μ, $\tau_{\text{had-vis}}$). The functions also account for the reconstruction and identification efficiencies associated to the decay products. In this step all variables that depend on the momentum of the leptons are also corrected.

3. In the next step, referred to as folding, trigger efficiencies are estimated and applied to the final objects. The efficiencies are derived for both single and dilepton triggers, and therefore this set of corrections must be applied to whole events after the two relevant leptons have passed through the previous steps.

4. Finally, the signal region selections are applied to obtain distributions equivalent to those from $Z \to \tau\tau$ in the phase space of interest.

This procedure is outlined in Figure 9.3 and is discussed in more detail below.

9.5.3.1 Efficiency unfolding

To ensure a good agreement between the reconstructed and truth $Z \to \ell\ell$ distributions, the selection efficiencies for the light leptons need to be taken into account. The triggering, reconstruction and identification efficiencies are measured in $Z \to ee$ and $Z \to \mu\mu$ events using POWHEG PYTHIA8 samples,

[11] This is what the VRs of the $H \to \tau\tau$ cross section measurement with the 2015-2016 dataset aimed to study.
[12] This is the component that was missing in the 2015-2016 measurement.

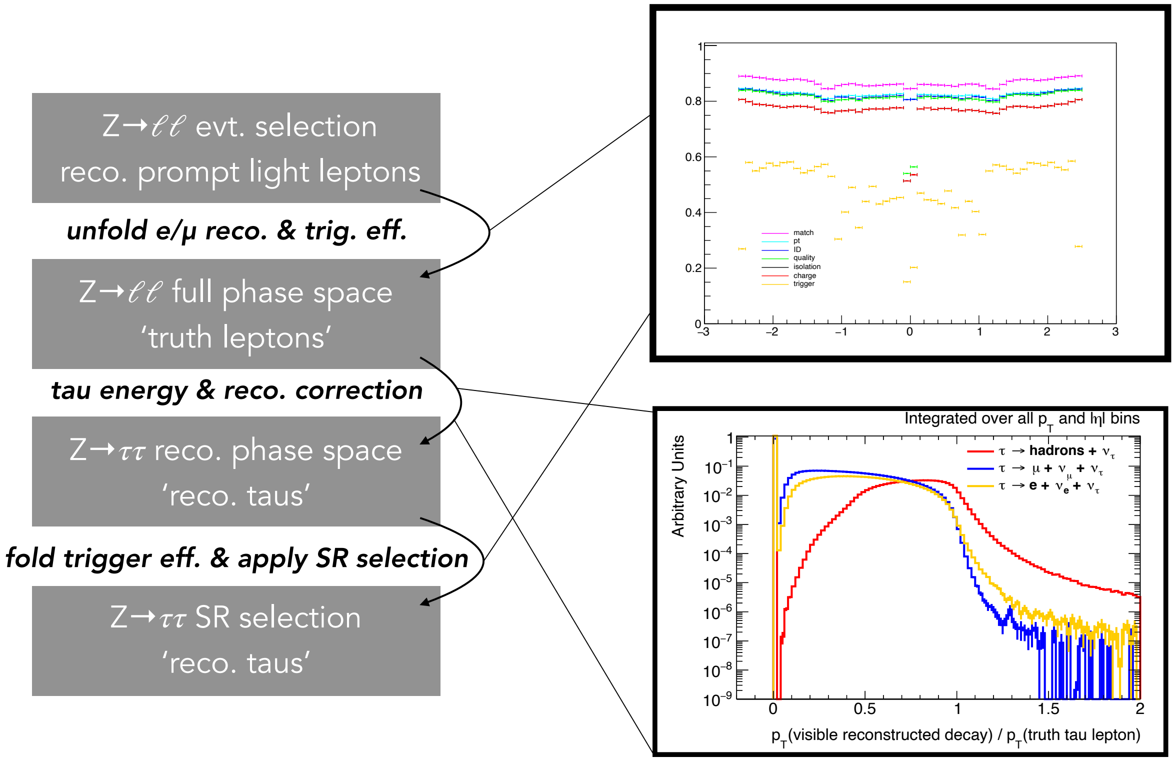

Figure 9.3: Embedding procedure. The unfolding step compensates for reconstruction effects in $Z \to \ell\ell$ events so that they may be used to estimate truth Z decay distributions. Then, the embedding procedure is applied to these events to approximate the effects of τ decays. Next, trigger efficiencies are folded in and finally the signal region selection is applied. At this stage, the events can be used to estimate $Z \to \tau\tau$ distributions in the signal regions of the analysis.

in two-dimensional bins of transverse momentum and η. Energy correction effects and data-MC scale factors for the triggering, reconstruction and identification efficiencies are taken into account, and their associated systematic uncertainties are propagated as needed. Figure 9.4 shows an example of the efficiencies for electrons and muons as a function of each variable individually, integrated over the other variable for illustration purposes. Once the efficiency of passing all three types of selections is computed, its inverse is used to scale the $Z \to \ell\ell$ event weights.

9.5.3.2 Embedding parametrisations

The next step is to determine the characteristics of visible τ decay products. As in the previous step, one can determine the reconstruction and identification efficiencies for reconstructed τ decay products by starting from a sample of truth τ leptons and determining the proportion that is matched to reconstructed and identified objects. These objects may be electrons, muons or $\tau_{\text{had-vis}}$ depending on the decay mode of the τ, as determined by its truth record in Monte Carlo, which will have experienced some change to its momentum during the decay process. The embedding parametrisations derived must therefore account for both the reconstruction and identification efficiencies, and the momentum scaling. Energy correction effects and scale factors are also accounted for in this part of the procedure. The systematic uncertainties associated to all of these effects are propagated through the procedure accordingly, and

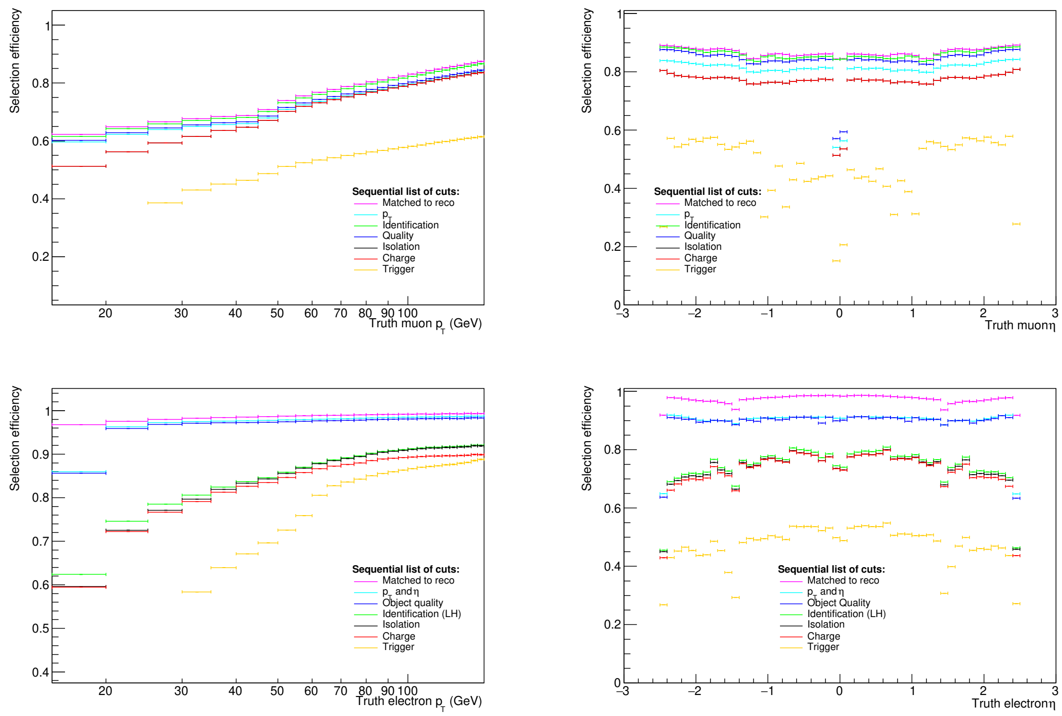

Figure 9.4: Efficiency as a function of p_{T} (left) and η (right) for muons and electrons from $Z \to \mu\mu$ (top) and $Z \to ee$ (bottom) events respectively. In the procedure, the efficiency parametrisation used is binned in 2D.

their impact is discussed in Section 9.5.6.2.

The embedding parametrisations are derived using a POWHEG PYTHIA8 $Z \to \tau\tau$ sample. While a SHERPA sample is used to model this process in the analysis, it cannot be used in this procedure as it is filtered based on the visible lepton p_{T}, which could bias the embedding parametrisation functions. The use of $Z \to \tau\tau$ (instead of a single-τ simulation) allows the parametrisations to account, on average, for spin effects and correlations[13].

The first step in the calculation is to assess whether a reconstructed τ decay product can be matched to the truth τ object. Reconstructed lepton candidates are selected from the electron, muon and $\tau_{\text{had-vis}}$ candidates described in Section 9.5.2 by imposing a proximity requirement to a truth τ object of $\Delta R < 0.2$. Reconstructed leptons are matched to the true τ according to its truth record, e.g. a truth τ whose truth record indicates that it decayed leptonically to a muon can only be matched to a reconstructed muon.

The embedding parametrisations are computed as the ratio of the momentum of the reconstructed lepton candidate to that of the truth τ object it is matched to. This ratio is computed in a 2-dimensional parametrisation of truth τ p_{T} and truth τ $|\eta|$ with bins as follows:

- p_{T} bin edges: [0, 30, 40, 50, 60, 70, 80, 100, 150, 50000]
- $|\eta|$ bin edges: [0, 0.4, 0.8, 1.37, 1.52, 2.5]

[13]The fact that spin correlations are not taken into account explicitly is not observed to significantly bias the distributions within the phase-space considered.

The embedding parametrisations are not seen to depend on any other variable. In particular, they are stable against pile-up. For each bin, a distribution of momentum ratios is created using all the events in the sample, as shown in Figure 9.5. This distribution can later be sampled to obtain individual embedding factors to scale the energy of the light leptons from the $Z \to \ell\ell$ source region.

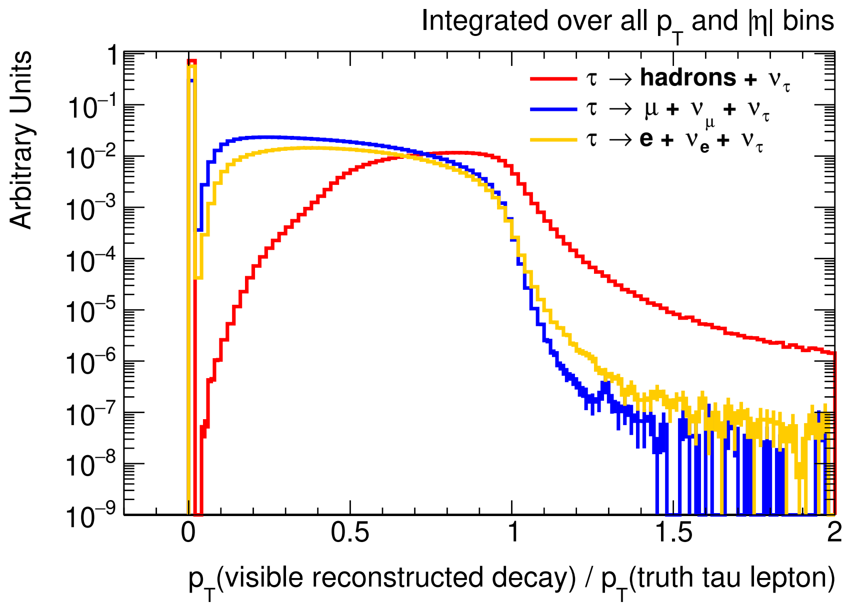

Figure 9.5: Embedding parametrisations, defined as reconstructed τ momentum divided by truth τ momentum for τ leptons decaying leptonically to electrons (yellow) and muons (blue) and for τ leptons decaying hadronically (red). The various truth momentum and η bins are integrated to make this plot easier to read.

It is possible for truth τ objects to have no matching reconstructed object that passes the selections described in Section 9.5.3.2 due to the reconstructed object failing either the reconstruction or the identification requirements. This type of object is added to a bin at zero in the parametrisation curve plots. This approach ensures that the probability that a true τ will fail to be reconstructed as either an electron, a muon, or a $\tau_{\text{had-vis}}$ once it has decayed is properly taken into account by the embedding procedure. The result is that when the distributions from Figure 9.5 are sampled, some leptons will be assigned an embedding factor from the zero-bin and therefore obtain an embedded momentum of zero. This will cause them to fail the final object selections, and ultimately be discarded from the background estimation.

9.5.3.3 Variable re-calculation

Assuming that after unfolding the reconstructed light leptons from $Z \to \ell\ell$ events are an appropriate substitute for truth leptons, one can apply the embedding parametrisations to the reconstructed leptons of the $Z \to \ell\ell$ events.

To estimate the $Z \to \tau\tau$ contribution in the various signal regions, the random number used to split $Z \to \ell\ell$ events into $\tau_\ell\tau_\ell$, $\tau_\ell\tau_{\text{had}}$ and $\tau_{\text{had}}\tau_{\text{had}}$ is used to choose between curves associated to either a reconstructed $\tau_{\text{had-vis}}$, a reconstructed muon or a reconstructed electron as the final product, ensuring that events that will be used to model the $Z \to \tau\tau$ background in the $\tau_{\text{had}}\tau_{\text{had}}$ channel are assigned embedding factors from the $\tau_{\text{had-vis}}$ curves, and similarly for the different channels. Embedding for the momenta of the $Z \to \ell\ell$ leptons is applied as follows: the reconstructed momentum of a light lepton to be embedded is used to choose a corresponding p_{T}-$|\eta|$ bin curve (such as the one illustrated in Figure 9.5). Embedding factors for the light lepton are obtained by using new random numbers to sample from the chosen parametrisation curve. The embedding factors are then applied to the original reconstructed light lepton momentum. Finally, the embedded light lepton momentum, now equivalent to that of a reconstructed τ decay object, is used to replace its unembedded counterpart in all calculations. Variables such as missing transverse energy[14] and momentum fractions are re-computed using the embedded momenta as inputs.

To ensure reproducibility, all random numbers used in the embedding procedure are generated with seeds based on the event number. Various offsets are added to the event number so different seeds can be produced for the different random numbers that are needed. In particular, the random number used to split the $Z \to \ell\ell$ events into $\tau_\ell\tau_\ell$, $\tau_\ell\tau_{\text{had}}$ and $\tau_{\text{had}}\tau_{\text{had}}$ differs from the one used to select an embedding factor through an offset that is added to the seed.

9.5.3.4 Oversampling

In practice, a significant fraction of truth τ do not have corresponding reconstructed electrons, muons or $\tau_{\text{had-vis}}$. As shown in Figure 9.5, the zero-bin is usually quite large. For example, in the bin with truth p_{T} between 0 and 30 GeV and $0 < |\eta| < 0.4$, the zero-bin contains 45% of the events. In the bin with truth p_{T} below 30 GeV and $0 < |\eta| < 0.4$, it contains 92% of events. In order to recover the statistics that would otherwise be lost to the zero-bin, an oversampling procedure is implemented. Instead of assigning a single set of embedding factors to each event, the embedding procedure is repeated 10 times for each seed event using different random numbers and producing pairs of leptons with different p_{T} values. An arbitrary lepton pair that satisfies the p_{T} criteria defined in the analysis is to be kept as the embedded event. The first di-lepton system that satisfies the loosest p_{T} criteria was chosen for this purpose. The criteria are: $p_{\mathrm{T}} > 40$ (30) GeV for the leading (subleading) lepton in the $\tau_{\text{had}}\tau_{\text{had}}$ channel, $p_{\mathrm{T}} > 30$ (20) GeV in the $\tau_\ell\tau_{\text{had}}$ channel and $p_{\mathrm{T}} > 17$ (10) GeV in the $\tau_\ell\tau_\ell$ channel. The embedded event is then weighted by the probability that the seed event used yields a pair of leptons satisfying the p_{T} criteria, calculated using the 10 different generated lepton pairs. An oversampling rate of 10 was deemed high enough to recover the statistics needed for the analysis without significantly increasing the run time of the procedure.

9.5.3.5 Trigger efficiency folding

The only efficiency for the embedded object that is not accounted for in the embedding parametrisations is the one related to the triggers. This is because the embedding factors are applied to each individual lepton, while trigger selections need to be applied on a per-event basis. So, as a final step, the rate at which the embedded leptons would pass the trigger selections is computed and used to re-weight the events. Efficiency maps identical to those described in Section 9.5.3.1, but derived using the $Z \to \tau\tau$ samples from Section 9.5.3.2 with events where the τ leptons decay to muons and electrons, are used

[14] Assuming that the contribution to the soft term (which uses energy deposits outside of the original objects) from both light leptons and embedded leptons is minimal, it is possible to estimate the corrected missing transverse energy contribution by simply subtracting away the original light lepton vector, and adding back the reconstructed embedded lepton vector.

for events with electrons and muons. This ensures that the trigger efficiencies are computed using non-prompt electrons and muons from τ decays. Events with $\tau_{\text{had-vis}}$ are also assigned weights based on the di-τ trigger efficiencies computed using the samples from Section 9.5.3.2. Efficiencies related to jet component of this trigger are not taken into account. The trigger chains considered are those presented in Table 9.4. Their efficiencies are slightly different in the various run years so a weighted-average is used in this step.

Weights are applied to the whole event, accounting for the possibility of either of the leptons to have fired the relevant (leg of the) event's trigger. The efficiency for events to pass the di-lepton trigger selections is taken into account with care taken to avoid any double-counting with single-lepton trigger efficiencies. At this stage, the embedded $Z \to \ell\ell$ leptons are equivalent to the τ decay objects from $Z \to \tau\tau$ events.

9.5.3.6 Selections post-embedding

Once the embedding procedure is performed, a new event selection is applied. While the original $Z \to \ell\ell$ source region was defined using selections on the variables before embedding, chosen to increase $Z \to \ell\ell$ statistics, a new selection is required on the embedded $Z \to \ell\ell$ events in order to obtain events that are as similar as possible to the $Z \to \tau\tau$ events that contribute as background to the signal regions. Each of these new control regions matches one of the signal regions of the analysis. For this purpose, the nominal event selection described in Table 9.6 is applied to the embedded objects and corrected quantities. The final distributions are then used to estimate those from the $Z \to \tau\tau$ process. An example of a comparison between a $Z \to \ell\ell$ distribution and its $Z \to \tau\tau$ equivalent is shown for the boosted region of the $\tau_{\text{had}}\tau_{\text{had}}$ channel in Figure 9.6. The agreement is significantly improved after embedding.

9.5.3.7 Comparison of $Z \to ee$ and $Z \to \mu\mu$ before and after embedding

Figure 9.7 illustrates the effect of the embedding on the missing transverse momentum distributions in MC. Before the embedding procedure, differences in the reconstruction, identification and trigger efficiencies between electrons and muons can be seen to bias the distributions. During embedding, these biases are corrected through the unfolding step leading $Z \to ee$ and $Z \to \mu\mu$ to result in identical distributions (within a few percent). This provides further validation of the embedding method, as one would indeed expect the embedded distributions to be independent of which $Z \to \ell\ell$ sample was used to generate them, i.e. no difference should be observed between $Z \to ee$ and $Z \to \mu\mu$ if all trigger, reconstruction and identification efficiencies are properly unfolded.

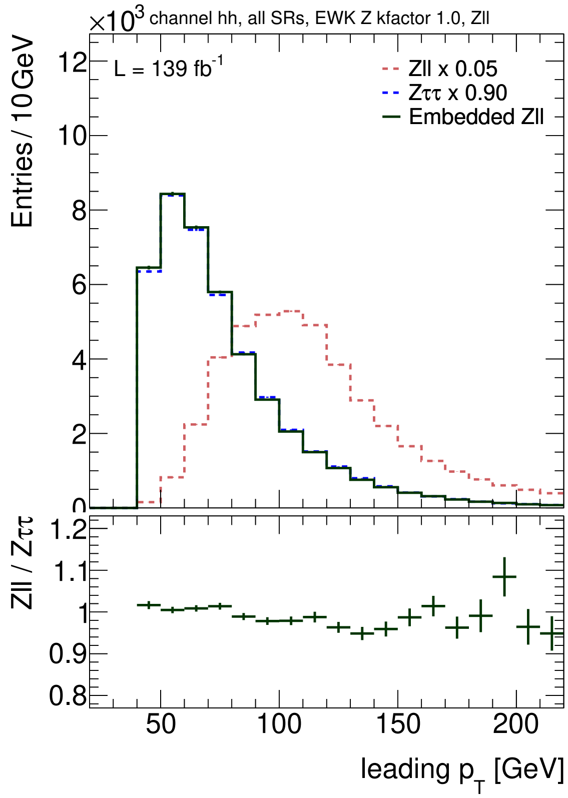

Figure 9.6: Leading tau p_{T} distributions for $Z \to \ell\ell$ in the source region (red, equivalent to the validation region used in the 2015-2016 analysis), embedded $Z \to \ell\ell$ control region (dark green) and $Z \to \tau\tau$ in the signal region (blue). The overall normalisation of the unembedded $Z \to \ell\ell$ and of $Z \to \tau\tau$ events are required to be identical to that of the embedded $Z \to \ell\ell$ events. The scaling required is shown at the top of this figure and allows for a comparison of the normalisation of the three distributions. The shape agreement of the embedded $Z \to \ell\ell$ distributions with $Z \to \tau\tau$ is shown in the ratio plot. A significant improvement in the $Z \to \ell\ell/Z \to \tau\tau$ agreement can be seen after the embedding procedure is performed.

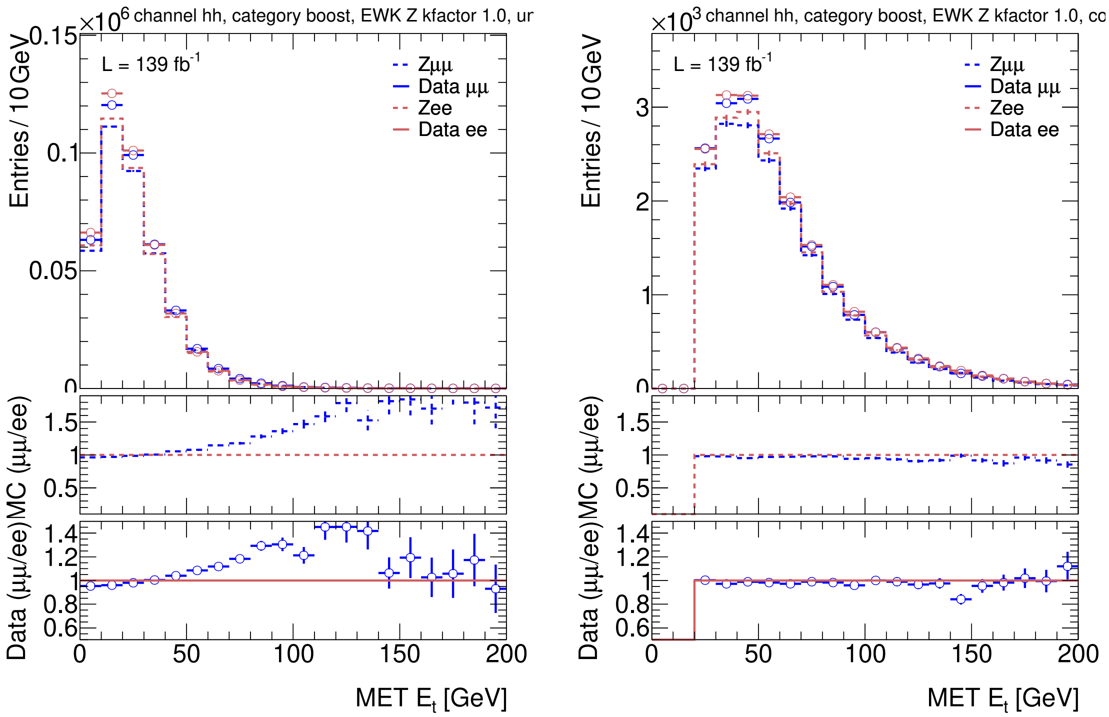

Figure 9.7: Missing transverse energy in $Z \to \mu\mu$ (blue) and $Z \to ee$ (red) events, for data (dots) and Monte Carlo (dashed lines) in the control region matching the boosted signal region of the $\tau_{\text{had}}\tau_{\text{had}}$ channel. The distributions are shown both before (left) and after (right) embedding. Before embedding, significant differences can be seen between $Z \to \mu\mu$ and $Z \to ee$, while after embedding the two distributions are similar.

9.5.4 Validation of the embedding method

While only the transverse momentum distributions are directly affected by the embedding procedure, many other variable distributions are also transformed through either a re-calculation for variables that depend on the p_{T} of the leptons in the event, or through changes to the event selection once it is updated after embedding. Agreement between the MC of the original $Z \to \ell\ell$ distributions in the relevant source regions, their counterparts in the embedding CRs and $Z \to \tau\tau$ in the signal regions are shown in Figures 9.8-9.11. The level of agreement indicates the level of similarity between the phase-spaces of the embedded $Z \to \ell\ell$ and of $Z \to \tau\tau$ and can be used to assess the level of non-closure of the method. In general, the closure is within a few percent except for the ΔR between the two leptons (Figure 9.11) where the non-closure reaches 10% and the lepton momenta in the $\tau_\ell\tau_\ell$ channel (Figure 9.9) where it also reaches 10%. Both effects are understood to be caused by simplifications in the embedding parametrisations, in particular those related to ignoring spin-correlation effects for ΔR and the coarseness of the functions' p_{T} binning for the lepton momenta. Overall normalisation differences of approximately 10% are seen between the embedded $Z \to \ell\ell$ and the $Z \to \tau\tau$ distributions in the $\tau_\ell\tau_{\mathrm{had}}$ and $\tau_{\mathrm{had}}\tau_{\mathrm{had}}$ channels and indicate residual phase-space mismatching, potentially also due to these simplifications.

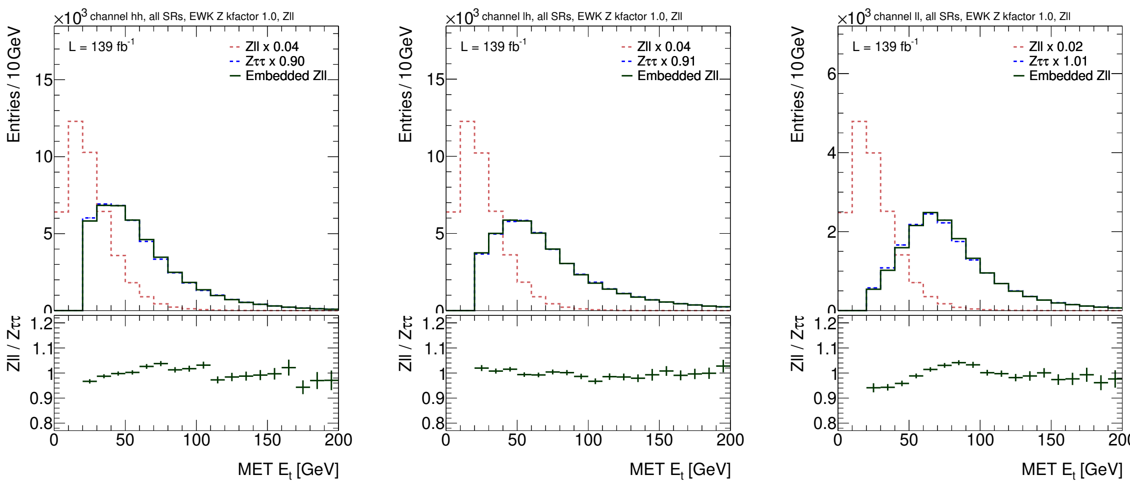

Figure 9.8: Missing transverse energy in the $\tau_{\mathrm{had}}\tau_{\mathrm{had}}$ (left), $\tau_\ell\tau_{\mathrm{had}}$ (middle) and $\tau_\ell\tau_\ell$ (right) channels. As before, $Z \to \ell\ell$ in the source region (red), embedded $Z \to \ell\ell$ (dark green) and $Z \to \tau\tau$ (blue) can be compared. The difference in the overall normalisations can be found at the top of the figure, while the ratio plot at the bottom allows for a shape comparison of the distributions of embedded $Z \to \ell\ell$ and of $Z \to \tau\tau$.

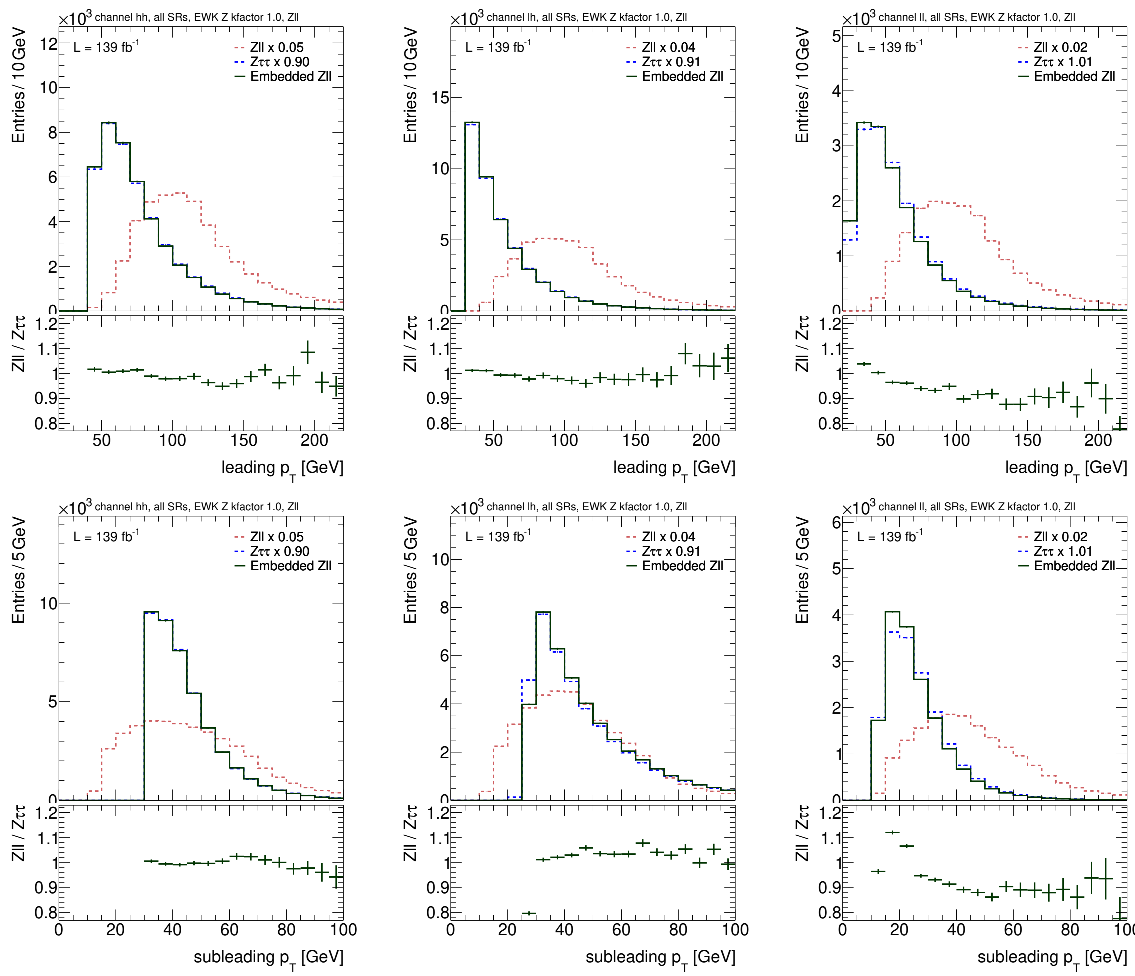

Figure 9.9: Leading (top) and subleading (bottom) lepton p_{T} in the $\tau_{\mathrm{had}}\tau_{\mathrm{had}}$ (left), $\tau_{\ell}\tau_{\mathrm{had}}$ (middle) and $\tau_{\ell}\tau_{\ell}$ (right) channels.

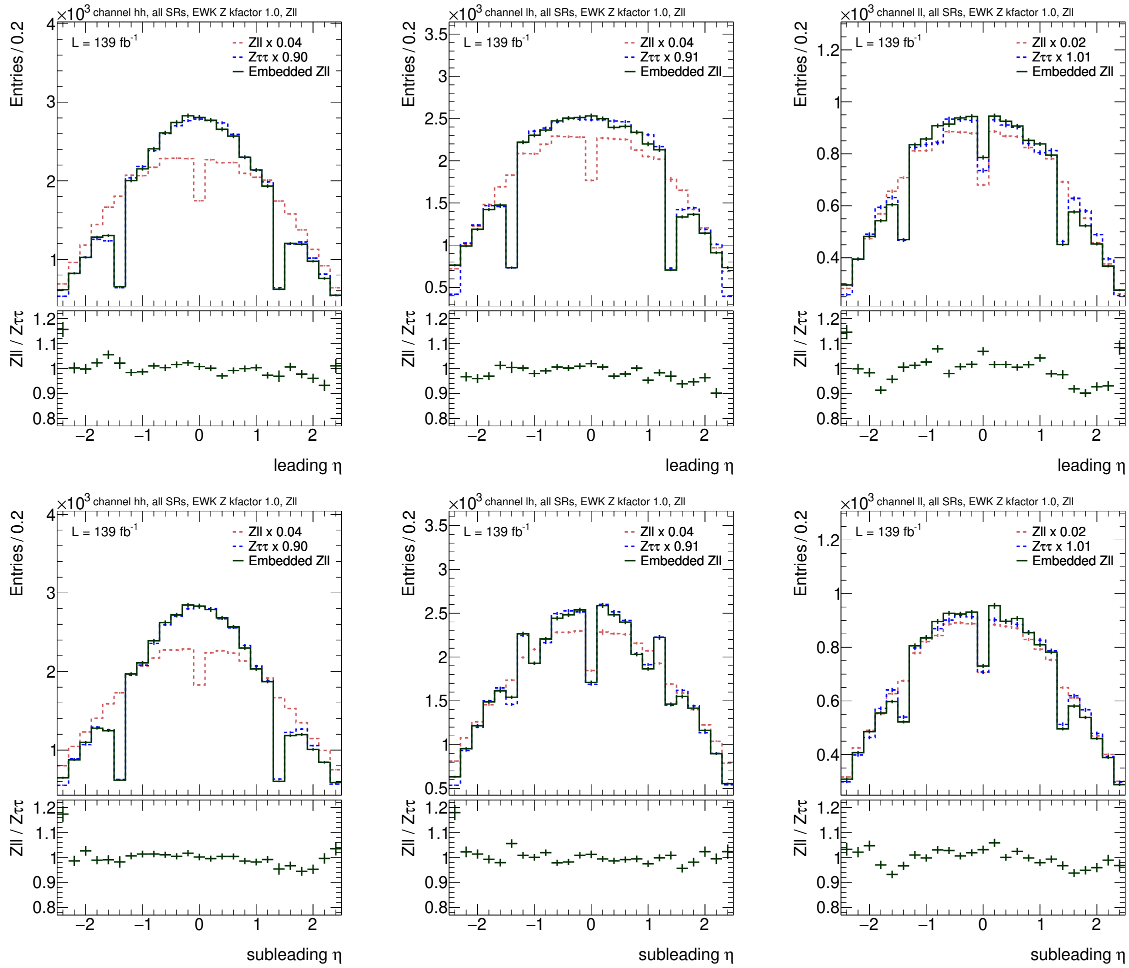

Figure 9.10: Leading (top) and subleading (bottom) lepton η in the $\tau_{\text{had}}\tau_{\text{had}}$ (left), $\tau_{\ell}\tau_{\text{had}}$ (middle) and $\tau_{\ell}\tau_{\ell}$ (right) channels.

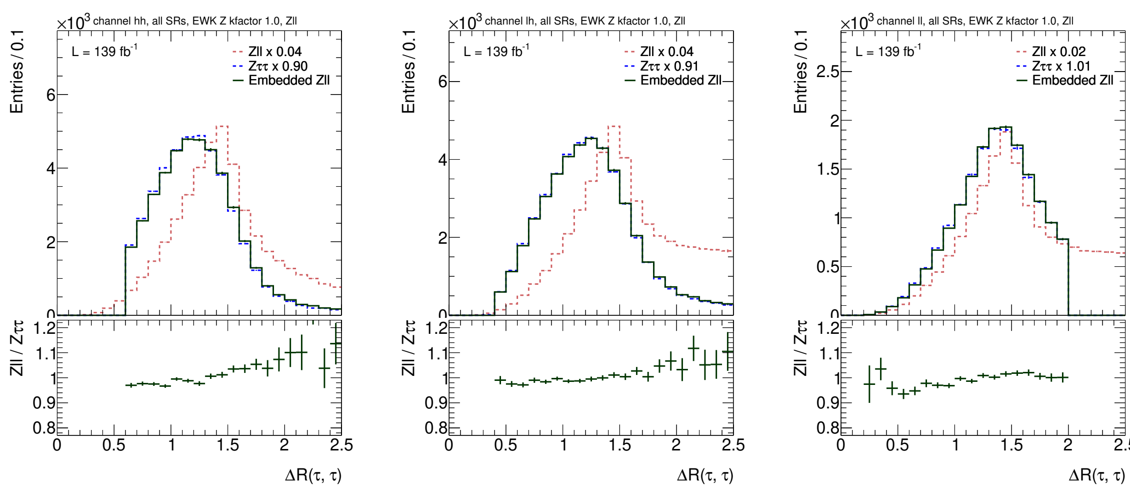

Figure 9.11: ΔR between the two leptons in the $\tau_{\text{had}}\tau_{\text{had}}$ (left), $\tau_{\ell}\tau_{\text{had}}$ (middle) and $\tau_{\ell}\tau_{\ell}$ (right) channels.

9.5.5 Comparison to the Run 1 method

Similar plots are presented in Figures 9.12, 9.13 and 9.14 for comparison to the Run 1 method, for the $\tau_\ell\tau_{\text{had}}$ channel. The lepton naming convention has changed since Run 1: the light lepton in $\tau_\ell\tau_{\text{had}}$ is now called the subleading τ. Note also that the phase-spaces considered are slightly different (in particular the Run 1 plots include events with large ΔR between the two leptons). These figures highlight that the quality of the embedding was not significantly affected by the simplifications within the object-level procedure.

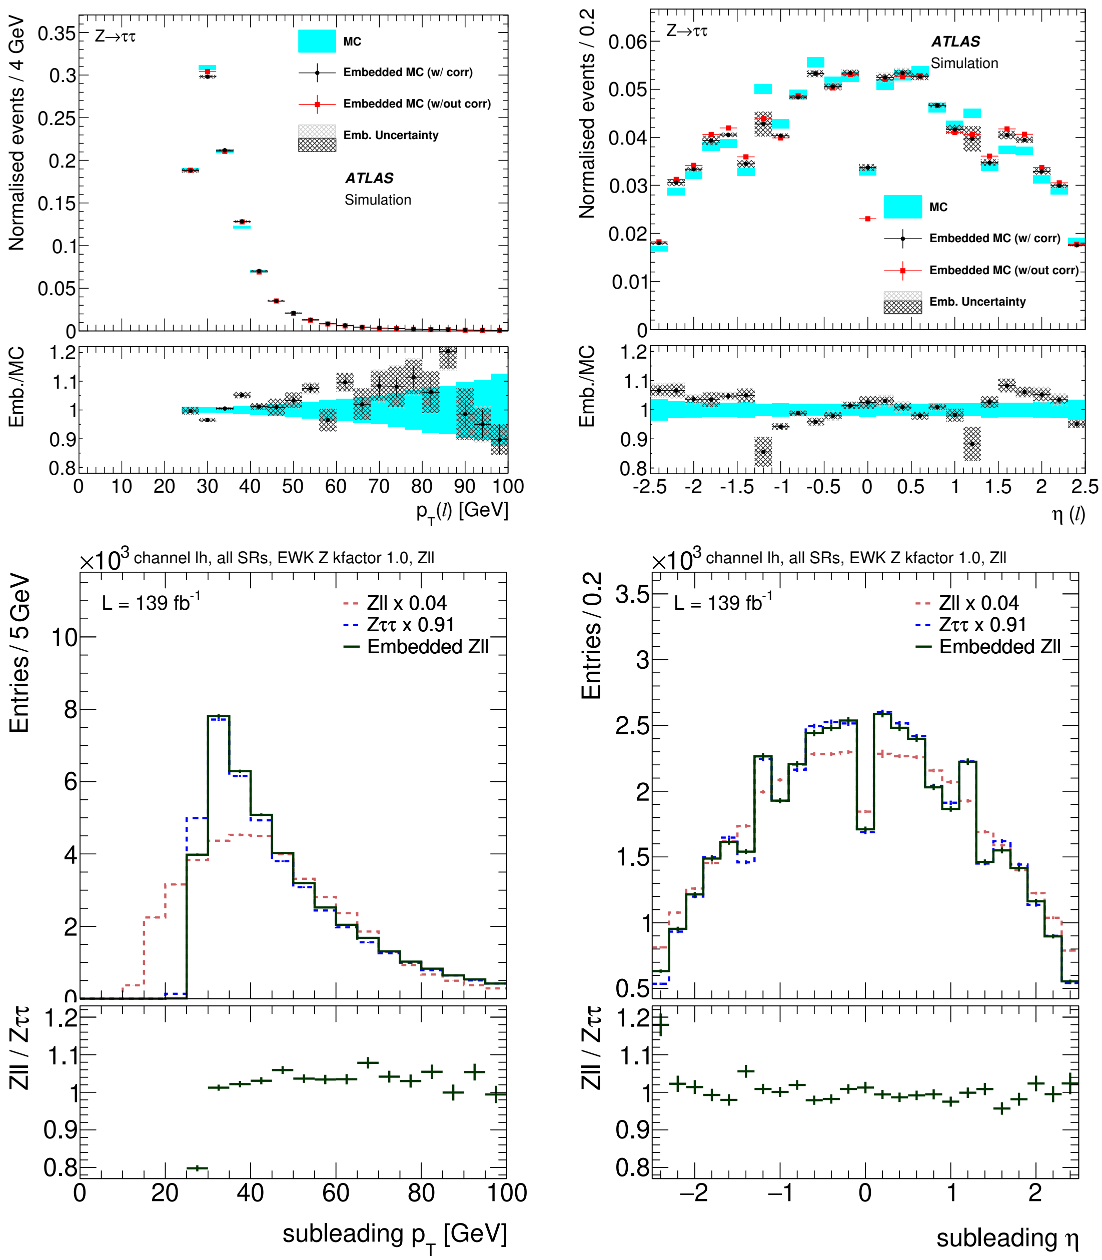

Figure 9.12: Comparison between the Run 1 embedding (top) and the object-level embedding (bottom) in the $\tau_\ell\tau_{\text{had}}$ channel for the light lepton p_{T} (left) and η (right).

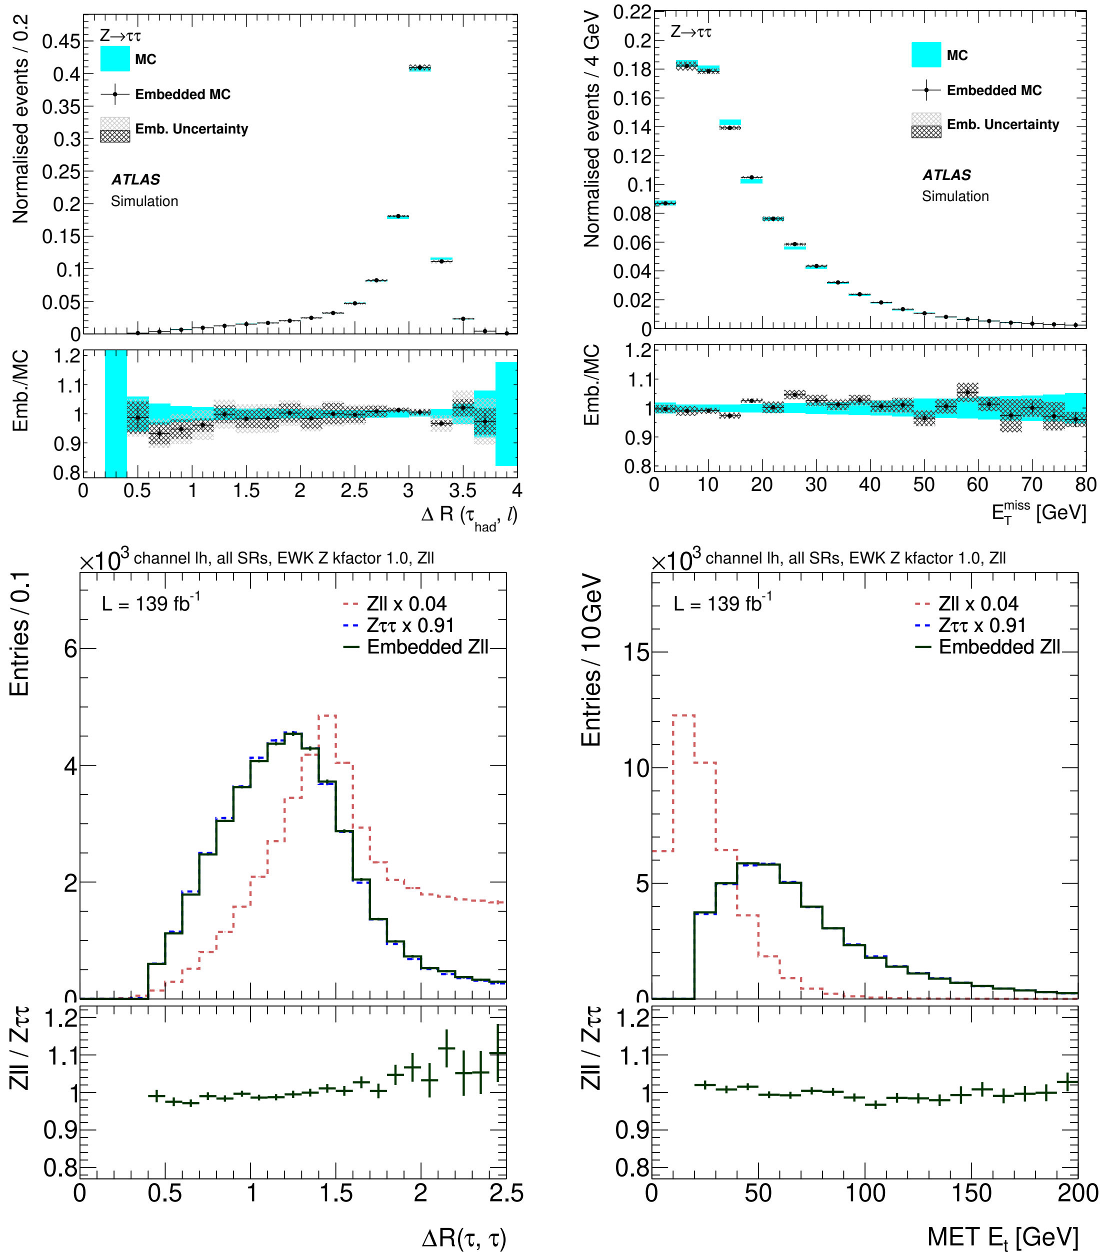

Figure 9.13: Comparison between the Run 1 embedding (top) and the object-level embedding (bottom) in the $\tau_\ell\tau_{\rm had}$ channel for the ΔR between the two leptons (left) and the missing transverse energy of the event (right).

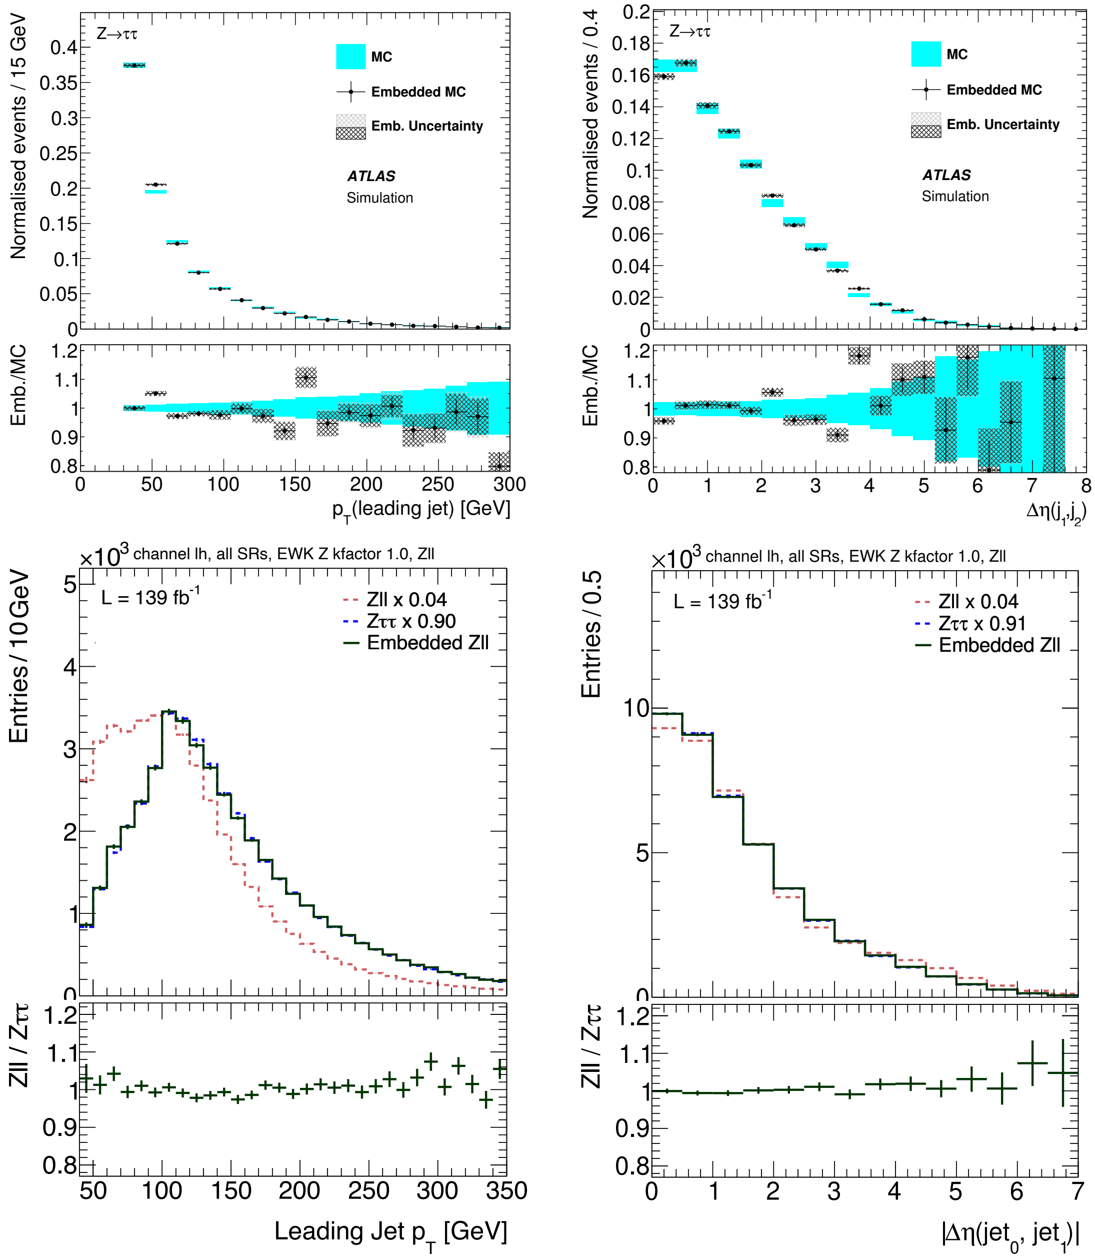

Figure 9.14: Comparison between the Run 1 embedding (top) and the object-level embedding (bottom) in the $\tau_\ell\tau_{\rm had}$ channel for the $p_{\rm T}$ of the leading jet (left) and the $\Delta\eta$ between the two leading jets (right).

9.5.6 Usage of the embedded control regions in the $H \to \tau\tau$ analysis

Events from the embedded $Z \to \ell\ell$ control region can be used to model the acceptance of $Z \to \tau\tau$ in the various signal regions as long as the embedded events provide an accurate description of the $Z \to \tau\tau$ process. The acceptance rates for both processes are assessed by comparing the $Z \to \tau\tau$ yield in each of the various signal regions of the analysis to that of embedded $Z \to \ell\ell$ in the control regions. The results are summarized in Figure 9.15. The regions studied are the eight boost bins, the two VBF bins and the two VH bins. Data-MC comparisons are shown and illustrate the NFs for the Z process that are observed in the various CRs.

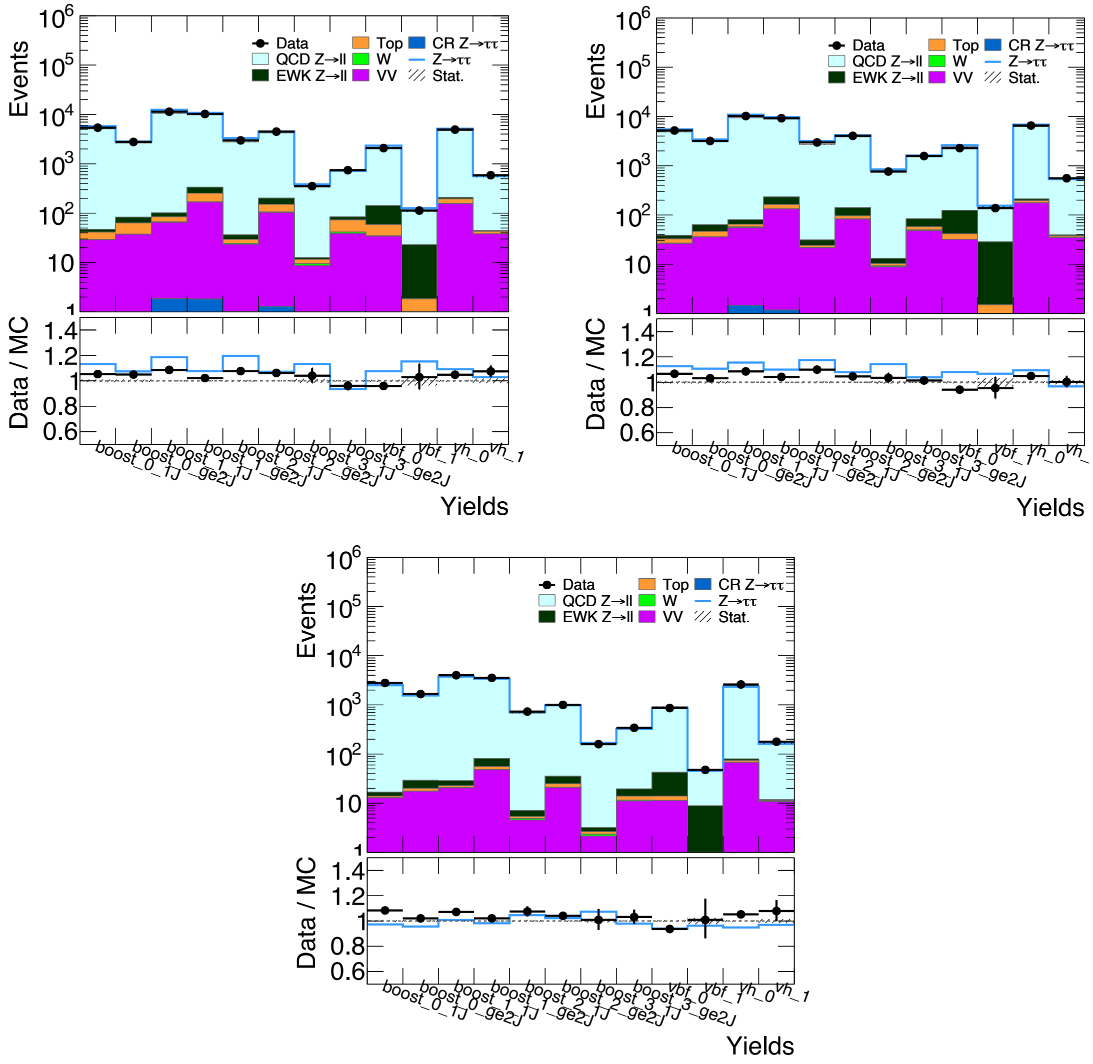

Figure 9.15: $Z \to \tau\tau$ (blue lines) yields in the SRs of the analysis and embedded $Z \to \ell\ell$ (light blue bands) yields in the associated CRs in the $\tau_{\rm had}\tau_{\rm had}$ (top left), $\tau_\ell\tau_{\rm had}$ (top right) and $\tau_\ell\tau_\ell$ (bottom) channels. Other processes contributing to the $Z \to \ell\ell$ control regions (coloured bands) and embedded data yields (black points) in the CRs are also included. The blue lines on the ratio plots indicate differences in acceptance between embedded $Z \to \ell\ell$ and $Z \to \tau\tau$ in the various regions.

While in the $\tau_\ell\tau_\ell$ channel the yields are similar between the two processes, discrepancies of approximately 12% can be seen in the $\tau_{\rm had}\tau_{\rm had}$ and $\tau_\ell\tau_{\rm had}$ channels. When fixing the overall normalisation of the embedded $Z \to \ell\ell$ process to that of $Z \to \tau\tau$ in each channel, the relative variations per region can be extracted as in Figure 9.16. No significant bias is seen for events in the $\tau_\ell\tau_\ell$ channel, while slight discrepancies appear between events with low and high jet multiplicities in the boost regions of the $\tau_{\rm had}\tau_{\rm had}$ and $\tau_\ell\tau_{\rm had}$ channels.

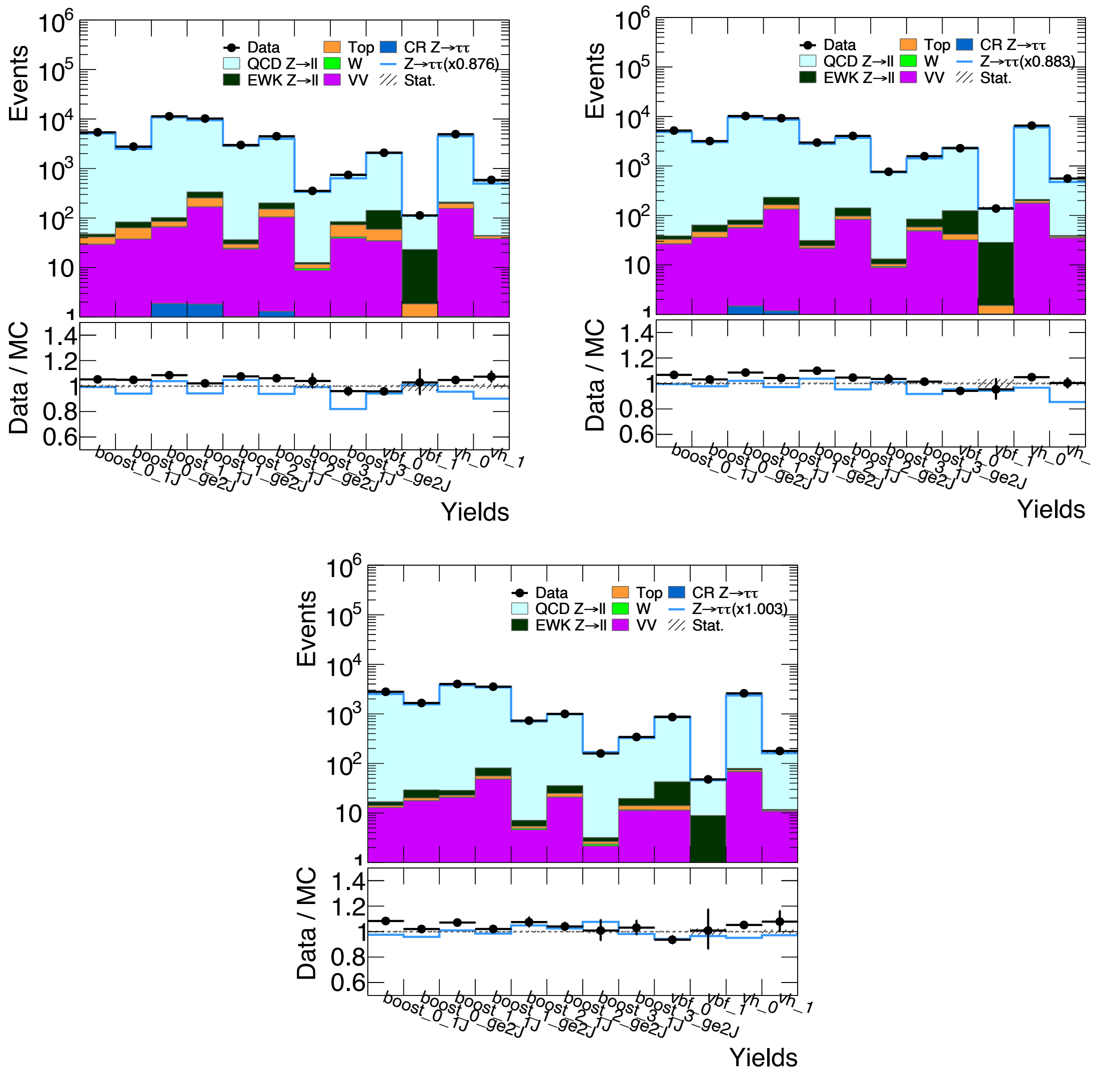

Figure 9.16: $Z \to \tau\tau$ (blue lines) yields in the SRs of the analysis and embedded $Z \to \ell\ell$ (light blue bands) yields in the associated CRs in the $\tau_{\rm had}\tau_{\rm had}$ (top left), $\tau_\ell\tau_{\rm had}$ (top right) and $\tau_\ell\tau_\ell$ (bottom) channels. Other processes contributing to the $Z \to \ell\ell$ control regions (coloured bands) and embedded data yields (black points) in the CRs are also included. In this case the normalisation of $Z \to \tau\tau$ is constrained to be exactly equal to that of the embedded $Z \to \ell\ell$ so that shape differences can be isolated.

Additionally, the migrations between truth and reconstructed bins used in the STXS framework are compared between the $Z \to \tau\tau$ and the embedded $Z \to \ell\ell$ samples. An appropriate modelling of these migrations is important to ensure that no bias is introduced to the STXS measurements. The embedding procedure replicates the migrations in the boost signal regions (bins of boson transverse momentum and jet multiplicity) within 2%.

9.5.6.1 Embedding-specific systematic uncertainties

As shown in Section 9.5.4, disagreements in terms of overall normalisation and in shapes of distributions are observed between $Z \to \tau\tau$ and embedded $Z \to \ell\ell$ events. This is likely due to the simplifications within the embedding method, such as the omission of effects like spin correlations and corrections to the direction of the light lepton, as well as potential differences in the event selection (e.g. through limitations that don't enable proper treatment of the overlap removal within embedding).

Since embedding is used to adjust the normalisation of the $Z \to \tau\tau$ contribution in the signal regions based on the differences between data and MC observed in the embedded control regions (see Figure 9.24), only mismodellings that affect the value of the normalisation factors need to be considered. In particular, non-closure in the variables used to define STXS bins (Higgs boson p_{T} and jet multiplicity in the boost regions, and the VBF and VH tagger scores) are of interest. Embedding-specific systematic uncertainties are derived to account for this non-closure. They are computed as follows:

1. The size of the mismatches in acceptance across adjacent bins is quantified by comparing their relative yields for both $Z \to \tau\tau$ and embedded $Z \to \ell\ell$ Monte Carlo.

2. This mismatch can be addressed by migrating events across those bins in the embedded $Z \to \ell\ell$ control regions until the relative yields are identical to those of $Z \to \tau\tau$. Therefore the next step is to measure the direction and size of this migration (i.e. how many events need to be moved from one category to the other).

3. Since adjacent bins also have different normalisation factors, data events must be migrated at a rate slightly different than that of the Monte Carlo events. As the data-Monte Carlo agreement is seen to vary smoothly as a function of the variable studied, a conservative way to estimate this effect consists of taking the average data-to-Monte Carlo ratio from the bin events are taken from (or being transferred to), as detailed in Equation 9.1. Taking the VBF tagger score as an example, for bins adjacent to the transition region between the the two VBF categories, the data-to-Monte Carlo ratio is seen to vary only slightly. Thus, taking the average data-to-Monte Carlo normalisation factor for the entire bin when performing the calculation will yield a conservative estimate that does not rely on specific assumptions on the nature of the migrations.

4. The equation used to calculate the data-to-Monte Carlo normalisation factors after the migrations described above is

$$\left(\frac{\text{Data}}{\text{MC}}\right)_X^{\text{post-migration}} = \frac{\left(\frac{\text{Data}}{\text{MC}}\right)_X \times N_X^{\text{events}} + \left(\frac{\text{Data}}{\text{MC}}\right)_Y \times N_{\text{migrated}}^{\text{events}}}{N_X^{events} + N_{\text{migrated}}^{\text{events}}} \tag{9.1}$$

   where X is the bin under study, Y is the bin where events are being migrated to (or from), and $N_{\text{migrated}}^{\text{events}}$ is the number of events that need to be transferred (it is negative in cases where events are removed from bin X). A special treatment is used for the Higgs boson p_{T}-based signal regions, as 4 bins are used for this variable instead of 2. Since migrations can happen in either direction, the estimates are calculated in two series: the first one assumes that all migrations have to come from interactions with the bin above, and the second one, from interactions with the bin below. For all bins, the largest of these 2 values provides an upper-bound on the size of the effect. This does not account for very specific cases where the migrations could cascade across the regions; in such cases, it is difficult to provide a precise estimate. However, as those could only affect the two central bins, where the data-to-Monte Carlo normalisation factor differences are extremely small, this possible effect is neglected.

The values obtained through this procedure are shown in Table 9.10 for the VBF tagger and in Table 9.11 VH tagger in the VBF and VH categories. For the boost signal regions, the values obtained for the jet multiplicity migrations are presented in Table 9.12. Those for the Higgs boson p_{T} migration (not shown) are of similar size. All uncertainties are found to be below the 1% level. These uncertainties define the degree of precision to which the embedded $Z \to \ell\ell$ normalisation in the data can be applied to $Z \to \tau\tau$ signal region events. An overall uncertainty of 1% uncorrelated across all control regions was chosen as input to the fit; while this choice is slightly conservative, it ensures that the pruning procedure[15] will not remove this degree of freedom, which could otherwise lead to experimental NPs being artificially pulled to compensate for the observed differences in normalisation.

Channel	Category	Migration	Absolute $\Delta(\frac{\text{Data}}{\text{MC}})$	Migration $\Delta(\frac{\text{Data}}{\text{MC}})$
$\tau_e\tau_\mu$	vbf 0	0.011%	10.257%	0.001%
	vbf 1	-0.307%	9.303%	0.029%
$\tau_\ell\tau_{\text{had}}$	vbf 0	0.084%	3.730%	0.003%
	vbf 1	-2.023%	3.596%	0.074%
$\tau_{\text{had}}\tau_{\text{had}}$	vbf 0	0.138%	8.325%	0.011%
	vbf 1	-3.885%	7.685%	0.311%

Table 9.10: Impact of the VBF tagger score non-closure on the data-to-Monte Carlo normalisation factors. The third column indicates the relative size of the migration needed, the fourth the relative difference in data-to-Monte Carlo normalisation factors between the two bins considered, and the fifth the relative magnitude of the migration effect on the NFs.

Channel	Category	Migration	Absolute $\Delta(\frac{\text{Data}}{\text{MC}})$	Migration $\Delta(\frac{\text{Data}}{\text{MC}})$
$\tau_e\tau_\mu$	vh 0	0.137%	2.659%	0.004%
	vh 1	-4.079%	2.591%	0.110%
$\tau_\ell\tau_{\text{had}}$	vh 0	-0.436%	3.595%	0.016%
	vh 1	9.282%	3.729%	0.317%
$\tau_{\text{had}}\tau_{\text{had}}$	vh 0	-0.130%	2.083%	0.003%
	vh 1	2.120%	2.040%	0.042%

Table 9.11: Impact of the VH tagger score non-closure on the data-to-Monte Carlo normalisation factors. The third column indicates the relative size of the migration needed, the fourth the relative difference in data-to-Monte Carlo normalisation factors between the two bins considered, and the fifth the relative magnitude of the migration effect on the NFs.

[15] To speed up the fitting procedure, very small uncertainties that will not noticeably impact the final result are removed from the fit model through a procedure called pruning. In the $H \to \tau\tau$ analysis, the pruning threshold is set at an impact of 1%.

Channel	Category		Migration	Absolute $\Delta(\frac{\text{Data}}{\text{MC}})$	Migration $\Delta(\frac{\text{Data}}{\text{MC}})$
$\tau_e\tau_\mu$	boost 0	1J	-0.397%	5.821%	0.023%
	boost 0	ge2J	0.634%	6.181%	0.039%
	boost 1	1J	-0.731%	4.733%	0.035%
	boost 1	ge2J	0.796%	4.968%	0.039%
	boost 2	1J	-0.274%	3.138%	0.009%
	boost 2	ge2J	0.198%	3.239%	0.006%
	boost 3	1J	-4.489%	2.396%	0.113%
	boost 3	ge2J	2.190%	2.340%	0.050%
$\tau_\ell\tau_{\text{had}}$	boost 0	1J	-0.323%	3.338%	0.011%
	boost 0	ge2J	0.514%	3.453%	0.018%
	boost 1	1J	-1.803%	3.941%	0.072%
	boost 1	ge2J	1.940%	4.103%	0.078%
	boost 2	1J	-4.012%	4.845%	0.203%
	boost 2	ge2J	2.849%	5.092%	0.141%
	boost 3	1J	-4.885%	2.031%	0.104%
	boost 3	ge2J	2.376%	2.073%	0.048%
$\tau_{\text{had}}\tau_{\text{had}}$	boost 0	1J	-1.245%	0.255%	0.003%
	boost 0	ge2J	2.451%	0.256%	0.006%
	boost 1	1J	-3.854%	5.923%	0.237%
	boost 1	ge2J	4.098%	6.296%	0.248%
	boost 2	1J	-5.105%	1.207%	0.065%
	boost 2	ge2J	3.436%	1.222%	0.041%
	boost 3	1J	-8.825%	8.203%	0.794%
	boost 3	ge2J	4.133%	8.937%	0.355%

Table 9.12: Impact of the jet multiplicity non-closure on the data-to-Monte Carlo normalisation factors. The fourth column indicates the relative size of the migration needed, the fifth the relative difference in data-to-Monte Carlo normalisation factors between the two bins considered, and the sixth the relative magnitude of the migration effect on the NFs.

9.5.6.2 Other systematic uncertainties

Systematic uncertainties related to the various objects in the event (such as those pertaining to $\tau_{\text{had-vis}}$ reconstruction identification or to jet calibration) are propagated through the embedding procedure and their impact is assessed, as variations that affect embedded $Z \to \ell\ell$ and $Z \to \tau\tau$ events differently would challenge the ability of the former to adequately model the latter.

Uncertainties can impact the closure of the embedding procedure in two ways:

- Changes to the embedded $Z \to \ell\ell$ Monte Carlo and data distributions (via changes to the embedded $Z \to \ell\ell$ events and/or changes to the folding and unfolding efficiencies as well as embedding parametrisations).
- Changes to the $Z \to \tau\tau$ Monte Carlo distributions.

Uncertainties that affect the Monte Carlo distributions typically have similar impacts on $Z \to \tau\tau$ and embedded $Z \to \ell\ell$ since their phase spaces are so similar. Therefore, they tend to have little impact on the ratios shown in Figure 9.15 and thus have little overall impact on the procedure. For example, the impact of the CKKW matching scale systematic variation[16] on the $Z \to \tau\tau$ and embedded $Z \to \ell\ell$ yields is almost identical, as shown in Figure 9.17. The small residuals, computed as the difference between the impact of the uncertainty on $Z \to \tau\tau$ and on embedded $Z \to \ell\ell$, are computed and propagated to the $H \to \tau\tau$ measurement.

Other uncertainties can have an impact on both data and Monte Carlo, in which case the impact on the embedded data and embedded Monte Carlo is compared and the residual is computed. Since the final expected $Z \to \tau\tau$ yield corresponds to the $Z \to \tau\tau$ Monte Carlo expectation multiplied by the normalisation factor obtained from the embedded $Z \to \ell\ell$ and data distributions, the impact of the uncertainty on the final yield is assessed via its impact on the $Z \to \tau\tau$ Monte Carlo yield and via the residual of the changes to the embedded $Z \to \ell\ell$ Monte Carlo and data yields (i.e. the impact on the NFs measured). This type of uncertainty is illustrated in Figure 9.18 by variations of the $\tau_{\text{had-vis}}$ identification efficiency and in Figure 9.19 for the muon reconstruction efficiency.

As was mentioned in Section 8.6, a good control region is one whose associated extrapolation and statistical uncertainties are lower than the original modelling uncertainties. In the case of the embedded control regions, the extrapolation uncertainties are seen to be small since most residuals are extremely small. The statistical uncertainties are also quite small due to large $Z \to ee$ and $Z \to \mu\mu$ yields. In contrast, typical modelling uncertainties such as CKKW tend to be quite large, as can be seen in Figure 9.17. The embedded CRs are therefore an excellent addition to the $H \to \tau\tau$ analysis.

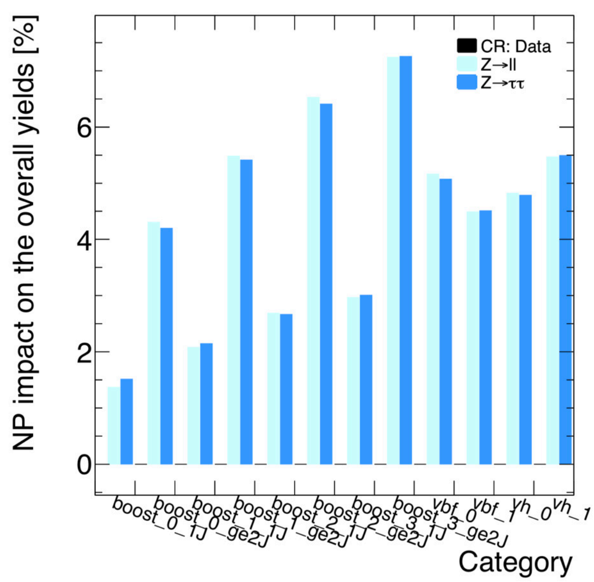

Figure 9.17: Impact of the variation of the CKKW systematic uncertainty on the $Z \to \tau\tau$ (blue) and embedded $Z \to \ell\ell$ yields (light blue). Since the effect is almost identical in both cases, this source of uncertainty is not expected to significantly affect the ratios shown in Figure 9.15.

[16]CKKW is one of the theoretical uncertainties considered for the Z process in the $H \to \tau\tau$ analysis.

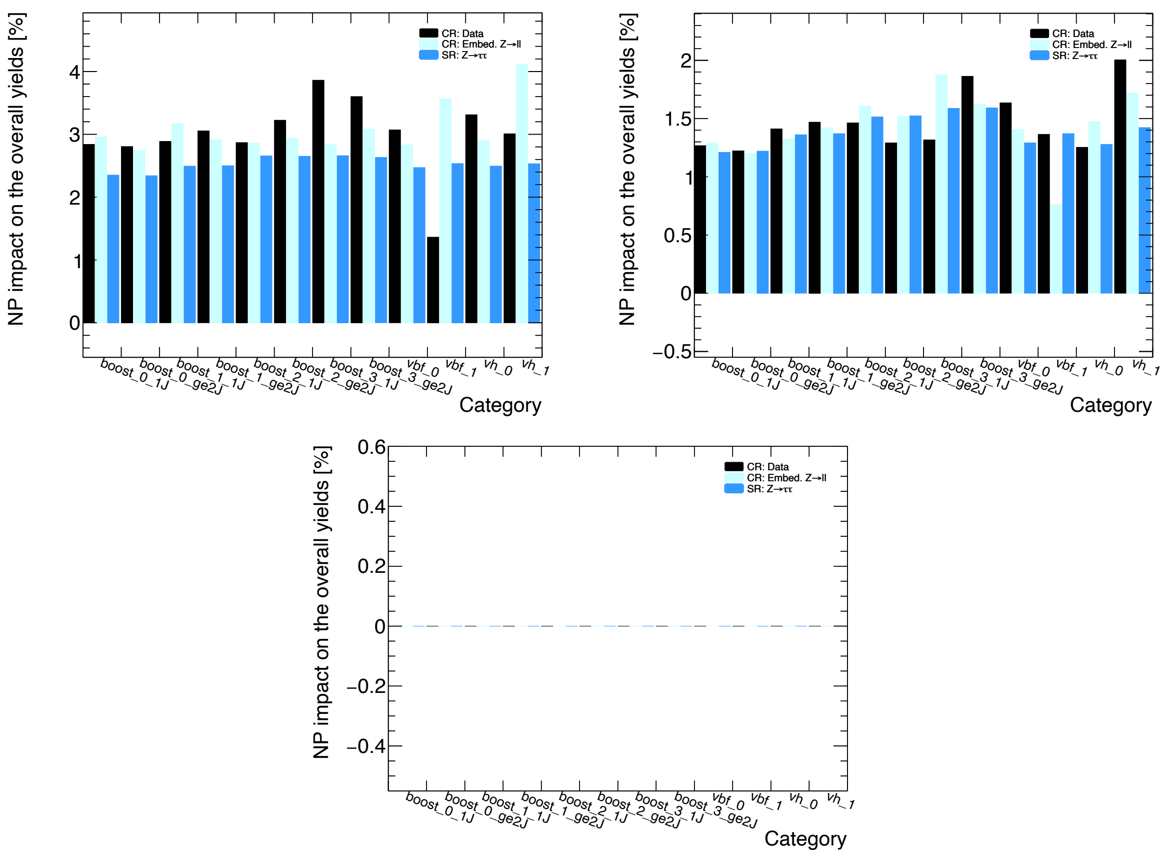

Figure 9.18: Impact of the variation of the $\tau_{\text{had-vis}}$ identification efficiency on the $Z \rightarrow \tau\tau$ Monte Carlo (blue) and embedded $Z \rightarrow \ell\ell$ data (black) and Monte Carlo (light blue) yields. The $\tau_{\text{had}}\tau_{\text{had}}$ (top left), $\tau_{\ell}\tau_{\text{had}}$ (top right) and $\tau_e\tau_\mu$ (bottom) channels are shown separately. The difference between the impact on the embedded $Z \rightarrow \ell\ell$ MC and embedded data is used to assign a corresponding uncertainty to the NF obtained in the embedded control regions.

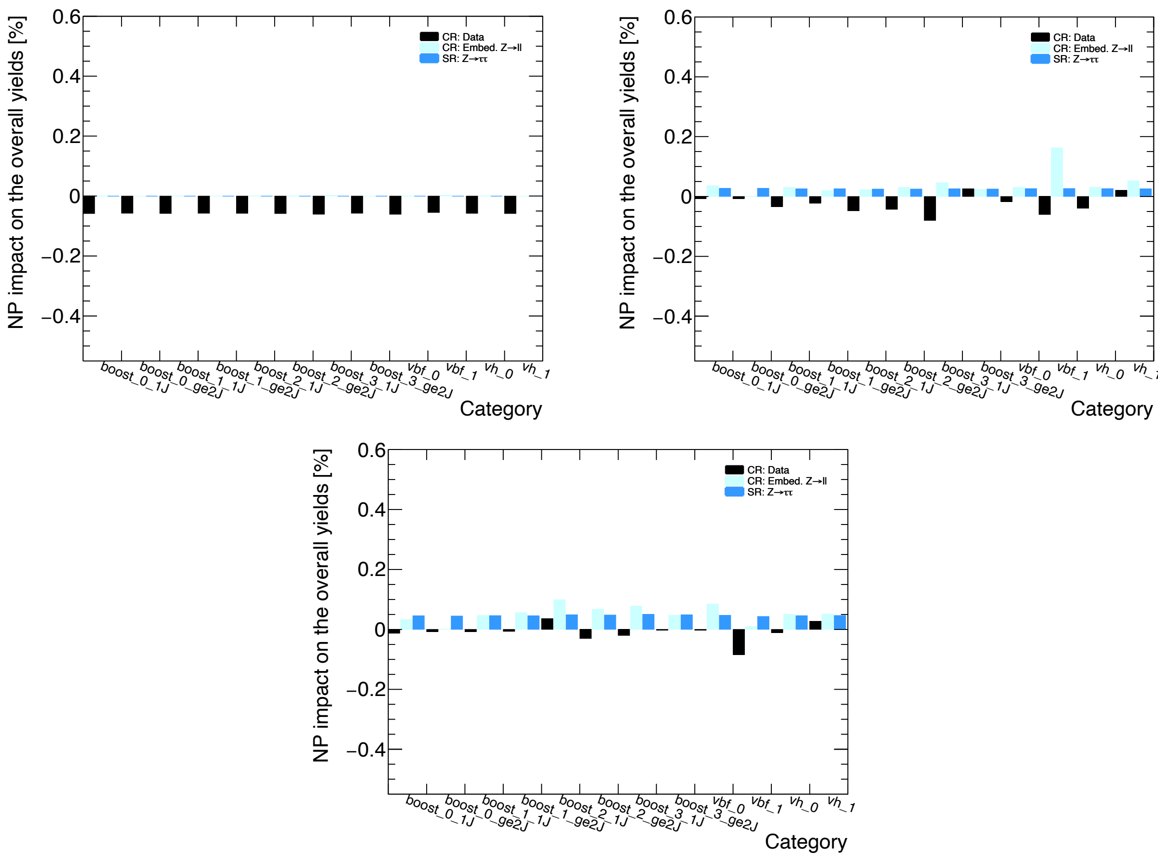

Figure 9.19: Impact of the variation of the muon reconstruction efficiency on the $Z \to \tau\tau$ Monte Carlo (blue) and embedded $Z \to \ell\ell$ data (black) and Monte Carlo (light blue) yields. This variation affects both the unfolding and the embedding parametrisations. The $\tau_{\rm had}\tau_{\rm had}$ (top left), $\tau_\ell\tau_{\rm had}$ (top right) and $\tau_e\tau_\mu$ (bottom) channels are shown separately. The difference between the impact on the embedded $Z \to \ell\ell$ MC and embedded data is used to assign a corresponding uncertainty to the NF obtained in the embedded control regions.

9.5.7 Results in the embedded control regions

As discussed at the beginning of this chapter, the embedded control regions are used in the analysis for two purposes:

- Assess the modelling of the main distributions that the analysis is sensitive to. This is a qualitative study.
- Constrain the normalisation of the Z process. This requires the use of the CRs in the fit to determine appropriate Z NFs.

9.5.7.1 Modelling of key variables in the CRs

A few key distributions are shown in the figures below for the $\tau_\ell\tau_{\text{had}}$ channel. These plots now include all other background contributions to the control region, for completeness. A data-to-Monte Carlo comparison is included and can be used to assess the quality of the modelling for each variable of interest.

For the boost signal regions, the transverse momentum of the reconstructed boson is observed (in the left part of Figure 9.20) to be well modelled with an overall normalisation factor of approximately 1.1, while the jet multiplicity distribution of the right side of Figure 9.20 shows a large trend in the data-to-MC comparison, with an over-estimation of the high jet multiplicity contributions in MC. In the VBF signal regions, both the dijet invariant mass (Figure 9.21, left) and the transverse momentum of the Higgs-plus-dijet system (Figure 9.21, right) distributions are seen to be poorly modelled, again with the MC over-estimating the data contributions at high values. This was observed in previous rounds of analysis and was of concern due to the use of these variables in the VBF tagger. The distribution of the VBF tagger score at the left of Figure 9.22 shows that the modelling is quite good despite these issues. The VH tagger score in the VH signal regions is shown on the right of Figure 9.22, with good agreement for the shape of the distributions and no significant phase-space dependence of the overall Z normalisation. Finally, the two $t\bar{t}H$ taggers are shown in the $t\bar{t}H$ signal regions in Figure 9.23. Good agreement can be observed for both the $Z \rightarrow \tau\tau$ and $t\bar{t}$ processes (no embedding is used for the latter).

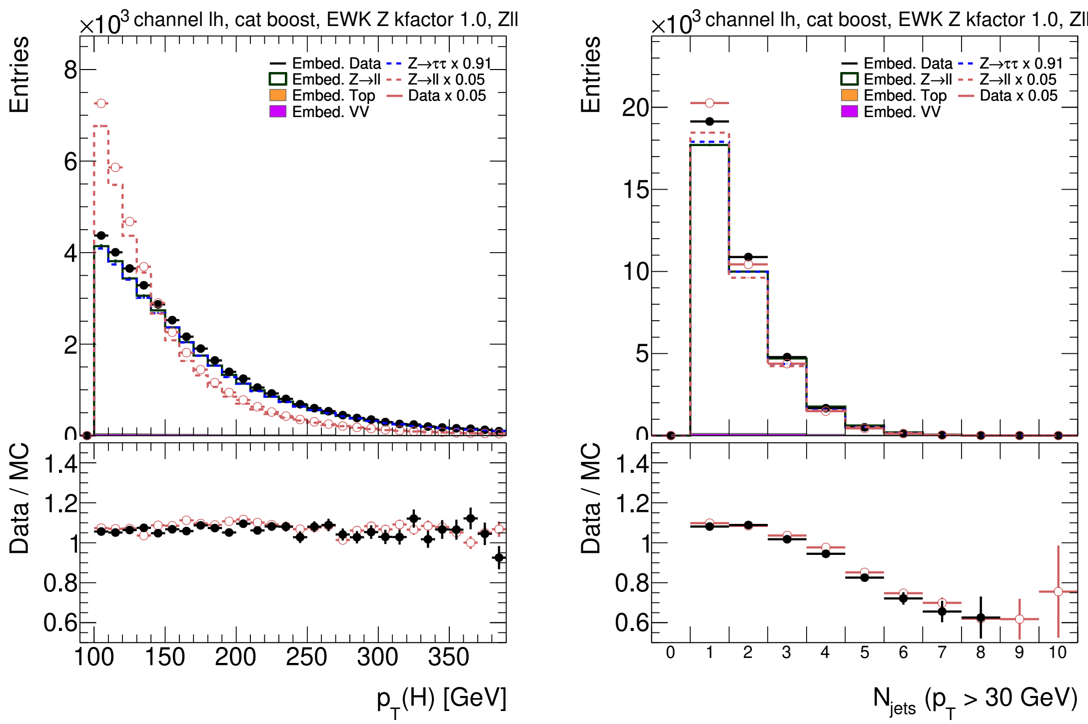

Figure 9.20: Transverse momentum of the reconstructed boson (left) and jet multiplicity (right) in the boost signal regions for the $\tau_\ell\tau_{\text{had}}$ channel. As before, the contributions from $Z \to \tau\tau$ in the SRs (dashed blue lines), unembedded MC and data in the source regions (red lines and dots), and embedded $Z \to \ell\ell$ MC and data in the control regions (coloured lines and black dots) are required to have the same overall normalisations and the scaling is shown in the legend.

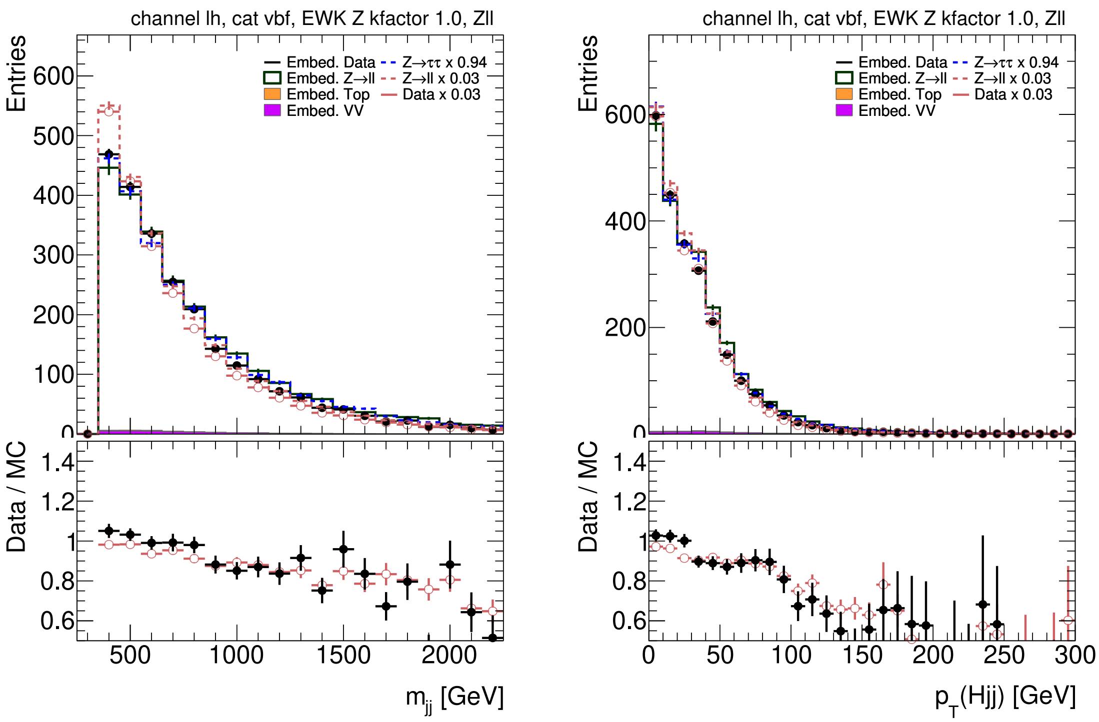

Figure 9.21: Invariant mass of the dijet system (left), transverse momentum of the Higgs and dijet system (right) in the VBF signal, source and control regions for the $\tau_\ell\tau_{\rm had}$ channel.

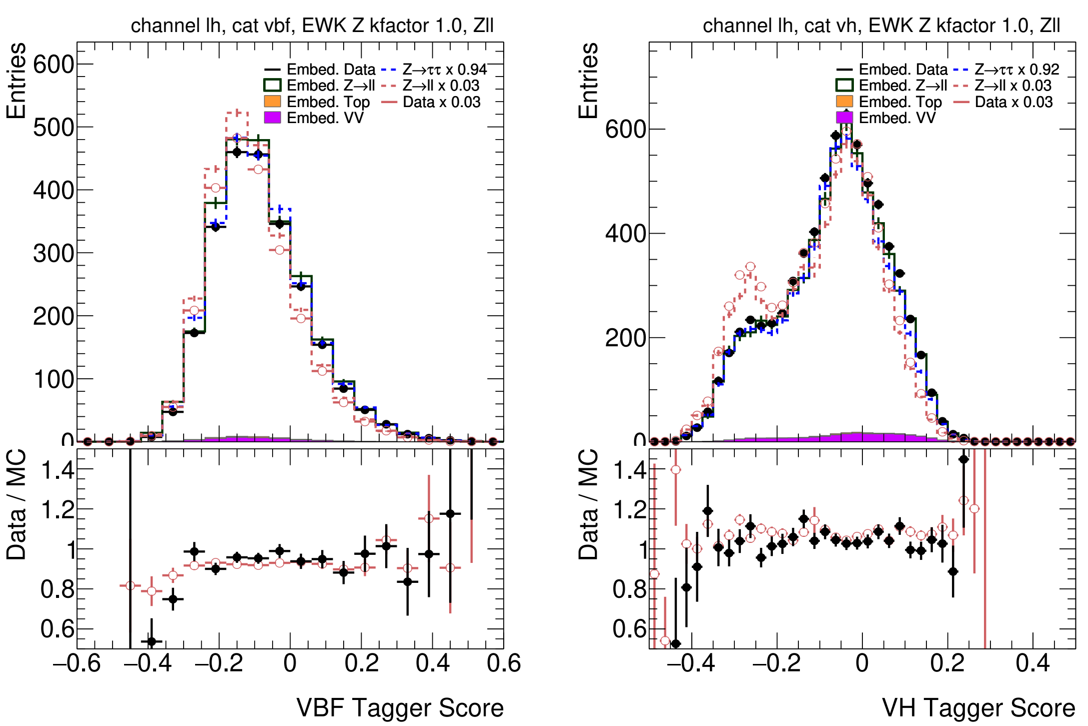

Figure 9.22: Tagger scores in the VBF (left) and VH (right) signal, source and control regions for the $\tau_\ell\tau_{\rm had}$ channel.

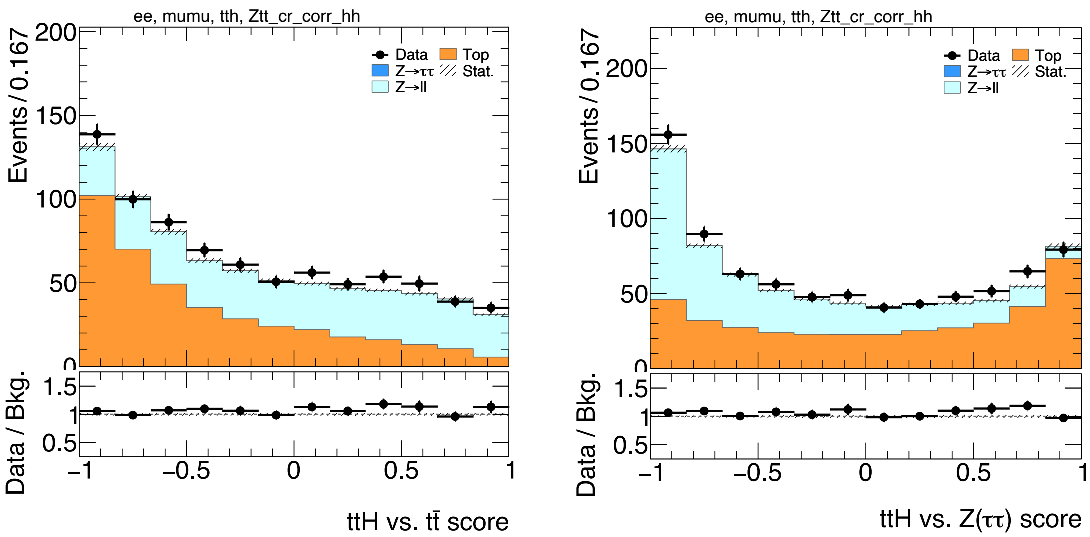

Figure 9.23: Data to Monte Carlo comparison for the tagger scores of $t\bar{t}H$ vs $t\bar{t}$ (left) and $t\bar{t}H$ vs Z (right) in the $t\bar{t}H$ control regions.

9.5.7.2 Normalisation of the Z process

Figure 9.24 illustrates the $Z \to \tau\tau$ normalisation as estimated by the embedded samples in each signal region of the $H \to \tau\tau$ analysis for each decay channel. In this case, the 36 signal regions described in Section 9.4.2.2 are considered as single bins. The $t\bar{t}H$ signal regions are not considered since large contributions from the top process make it hard to assess the $Z \to \tau\tau$ normalisation in the embedding control regions. These normalisation factors are used as inputs to the fit discussed in Section 9.7.

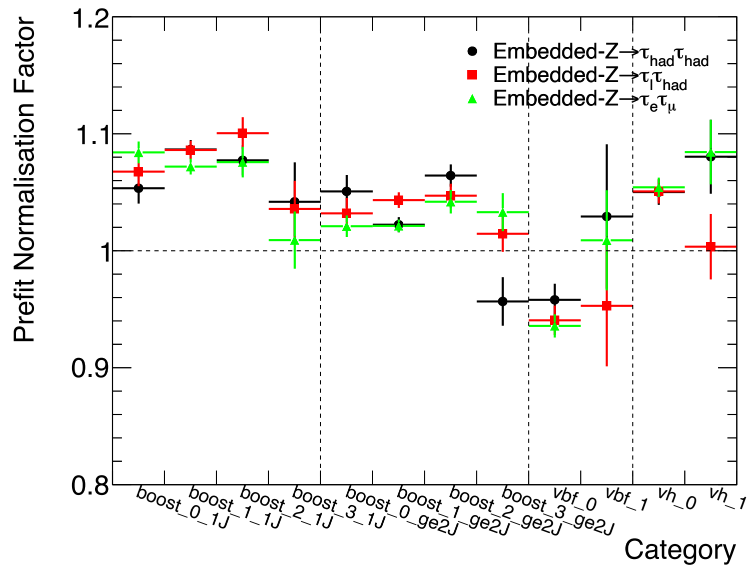

Figure 9.24: $Z \to \tau\tau$ normalisation factor estimated using the embedded control regions matching each signal region of the $H \to \tau\tau$ analysis for each decay channel and category. The normalisation factor is measured in each bin as the difference between data and the small backgrounds (diboson, top, ...) divided by the embedded $Z \to \ell\ell$ simulation sample yield.

9.5.8 Z electroweak contributions

Another known source of mismodelling for the Z process concerns the Z electroweak (EWK) contributions. In this analysis, they are evaluated using Sherpa v2.2.1 (at leading order). In previous measurements, such as the one carried out in the context of the electroweak production of dijets in association with a Z boson [73], this generator was seen to under-predict the data by a large factor. In [73] the predictions were compared to NLO predictions from Powheg Pythia8 and Herwig, with the latter reproducing the data with better accuracy as shown in Figure 9.25.

In the $H \to \tau\tau$ analysis, Sherpa and Powheg Pythia8 samples are available for the $Z \to \mu\mu$ and $Z \to ee$ processes while Herwig samples are only available for $Z \to ee$. The $Z \to \tau\tau$ sample is only available in Sherpa v2.2.1, so this generator is the only one that can be used to estimate the Z background contributions to the $H \to \tau\tau$ measurement. Sherpa v2.2.1 predicts very small electroweak contributions in most of the signal regions of the analysis (below 5%) except in the tight VBF phase space where it accounts for approximately 20% of the total Z contributions. To ensure that the EWK contribution is appropriately assessed, a comparison of all three generators using the object-level embedding method with $Z \to ee$ samples is shown in Figure 9.26. It is shown for the $\tau_{\rm had}\tau_{\rm had}$,

$\tau_\ell \tau_{\text{had}}$ and $\tau_e \tau_\mu$ channels combined.

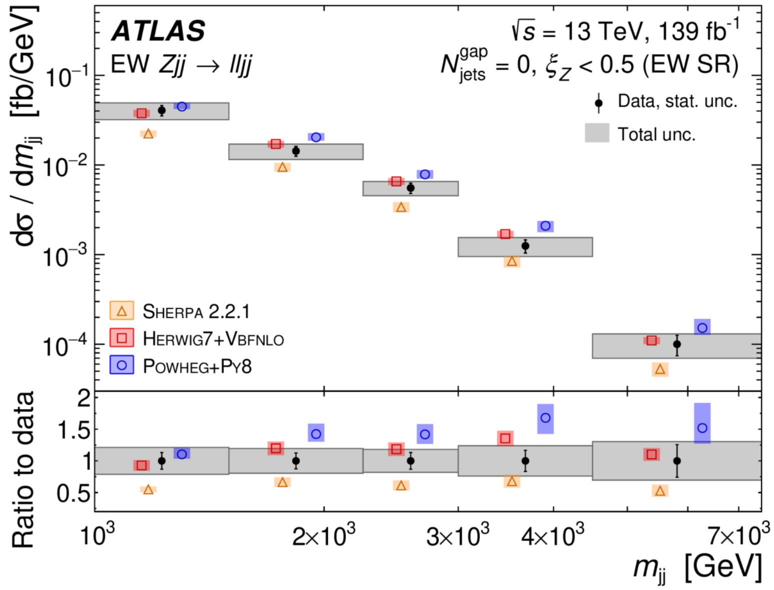

Figure 9.25: Differential cross sections for EW Zjj production as a function of m_{jj}. The unfolded data (black points) are compared with the predictions produced by HERWIG (red), POWHEG PYTHIA8 (blue) and SHERPA v2.2.1 (orange). Image from [73].

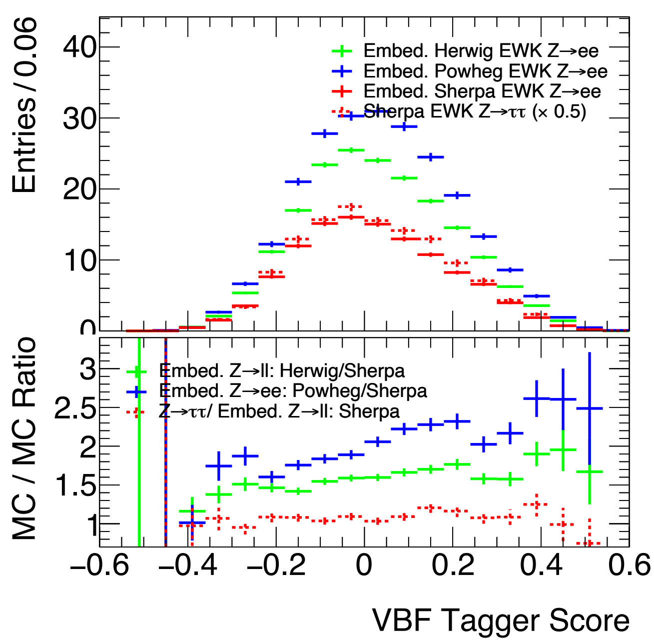

Figure 9.26: VBF tagger distribution in the $\tau_{\text{had}}\tau_{\text{had}}$, $\tau_\ell \tau_{\text{had}}$ and $\tau_e \tau_\mu$ channels combined using $Z \to ee$ events from SHERPA (red), POWHEG PYTHIA8 (blue) and HERWIG (green), with $Z \to \tau\tau$ from SHERPA (dashed red) shown for comparison. As the embedding is only performed on $Z \to ee$ events for this figure (instead of $Z \to ee + Z \to \mu\mu$), the $Z \to \tau\tau$ prediction needs to be scaled down by a factor 2 to be compared to the embedded events.

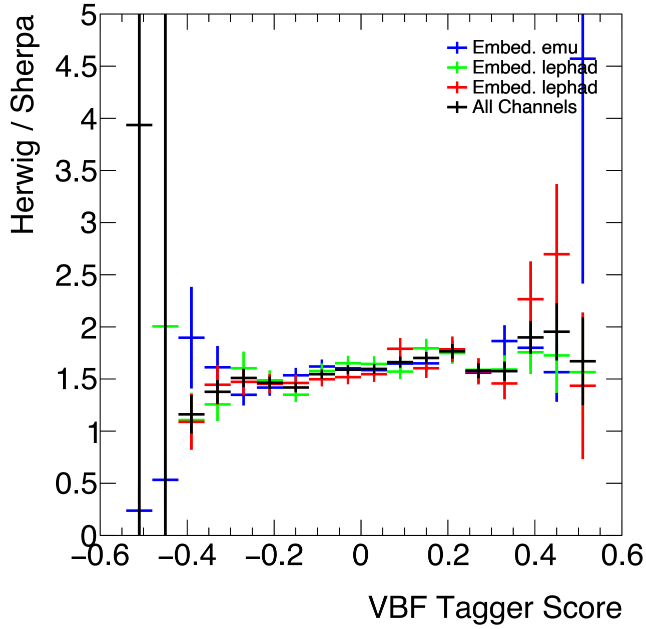

Figure 9.27: Comparison of the Sherpa and Herwig contributions as a function of the VBF tagger score in the $\tau_e\tau_\mu$ (blue), $\tau_\ell\tau_{\text{had}}$ (green) and $\tau_{\text{had}}\tau_{\text{had}}$ (red) channels individually, compared to the all-channel combination (black).

As discussed, in each signal region of the $H \to \tau\tau$ analysis the $Z \to \tau\tau$ prediction is scaled using a dedicated normalisation factor obtained from a $Z \to \ell\ell$ control region designed to mimic the kinematic phase space of the specific signal region considered. Consequently, the exact fraction between the EWK and QCD production of Z events is determined by the data in the control region. However an additional uncertainty in the Sherpa EWK $Z \to \tau\tau$ prediction needs to be considered in the signal regions to account for the impact of the missing orders in the calculation of the $m_{\tau\tau}^{\text{MMC}}$ lineshape and propagated to the $H \to \tau\tau$ analysis. Using the Herwig NLO prediction as reference, the LO Sherpa prediction can be rescaled through a NLO/LO correction factor of approximately 1.7 (see Figure 9.27). The $m_{\tau\tau}^{\text{MMC}}$ shape of the Z processes (QCD+EWK) is compared with and without this scaling factor on the EWK component. This comparison leads to a negligible difference in the low-score VBF categories where the EWK fraction is small. However it leads to a sizeable uncertainty in the high-score VBF categories, where the Z electroweak background has sizeable contribution. Figure 9.28 shows the impact of this uncertainty in the $\tau_{\text{had}}\tau_{\text{had}}$ channel.

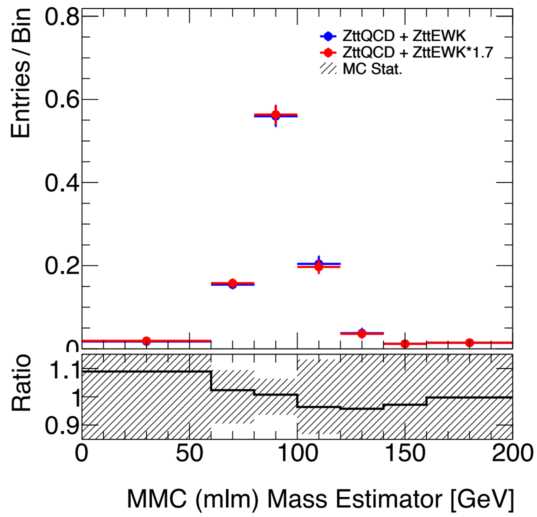

Figure 9.28: $m_{\tau\tau}^{\text{MMC}}$ lineshape systematic uncertainty for $Z \to \tau\tau$ electroweak in the high-score VBF category of the $\tau_{\text{had}}\tau_{\text{had}}$ channel.

9.6 Overview of background estimations

Unlike the $t\bar{t}H \to$ multilepton signal regions, the $H \to \tau\tau$ regions are dominated by contributions from events with real and prompt light leptons and $\tau_{\text{had-vis}}$. The bulk of background estimates for the $H \to \tau\tau$ analysis are therefore based on Monte Carlo simulation predictions, except for the fake estimates which must still be data-driven due to the insufficient precision and accuracy of the estimations from MC as previously discussed in the context of the $t\bar{t}H \to$ multilepton analysis. The various background contributions are shown in Figure 9.29 for all signal regions combined. In order of importance, the main backgrounds are:

- $Z \to \tau\tau$: 70-90% of the background contributions, depending on the signal region. These contributions are irreducible, although some sensitivity can be achieved through the use of the $m_{\tau\tau}^{\text{MMC}}$. As discussed, the modelling of the Z process is thoroughly validated using $Z \to \ell\ell$ control regions, where its overall normalisation is also determined. The object-level embedding technique is used to provide control regions with a phase space that is as similar as possible as that of $Z \to \tau\tau$ in the signal regions.

- Fake leptons (e, μ, τ_{had}): 5-20% of background contributions. Fake and isolation factors similar to the ones described for the $t\bar{t}H \to$ multilepton analysis are used in channels with $\tau_{\text{had-vis}}$, while the matrix method is used to estimate fake rates in the $\tau_e\tau_\mu$ channel. The matrix method works as follows. The number of events that pass the analysis' selection N_{TT} has contributions from events with two real (N_{rr}), one real and one fake (N_{rf} and N_{fr}), and two fake light leptons (N_{ff}). The matrix method utilizes objects in a control region defined using looser (L) lepton selections as well as L-to-T selection efficiencies (ϵ) for real and fake light leptons measured in MC and data, respectively, to estimate those contributions by inverting the matrix

$$\begin{bmatrix} N_{TT} \\ N_{TL} \\ N_{LT} \\ N_{LL} \end{bmatrix} = \begin{bmatrix} \epsilon_r\epsilon_r & \epsilon_r\epsilon_f & \epsilon_f\epsilon_r & \epsilon_f\epsilon_f \\ \epsilon_r\bar{\epsilon}_r & \epsilon_r\bar{\epsilon}_f & \epsilon_f\bar{\epsilon}_r & \epsilon_f\bar{\epsilon}_f \\ \bar{\epsilon}_r\epsilon_r & \bar{\epsilon}_r\epsilon_f & \bar{\epsilon}_f\epsilon_r & \bar{\epsilon}_f\epsilon_f \\ \bar{\epsilon}_r\bar{\epsilon}_r & \bar{\epsilon}_r\bar{\epsilon}_f & \bar{\epsilon}_f\bar{\epsilon}_r & \bar{\epsilon}_f\bar{\epsilon}_f \end{bmatrix} \begin{bmatrix} N_{rr} \\ N_{rf} \\ N_{fr} \\ N_{ff} \end{bmatrix}$$

where $\bar{\epsilon} = 1 - \epsilon$.

- Top: $< 5\%$ of background contributions (except in the $t\bar{t}H$ signal regions, where it reaches 30-50%). Dedicated control regions are used to validate the modelling and determine the normalisation of top processes.

Other backgrounds include V+jets, diboson and $t\bar{t}V$. Their contributions are generally very small and simply estimated using MC simulation.

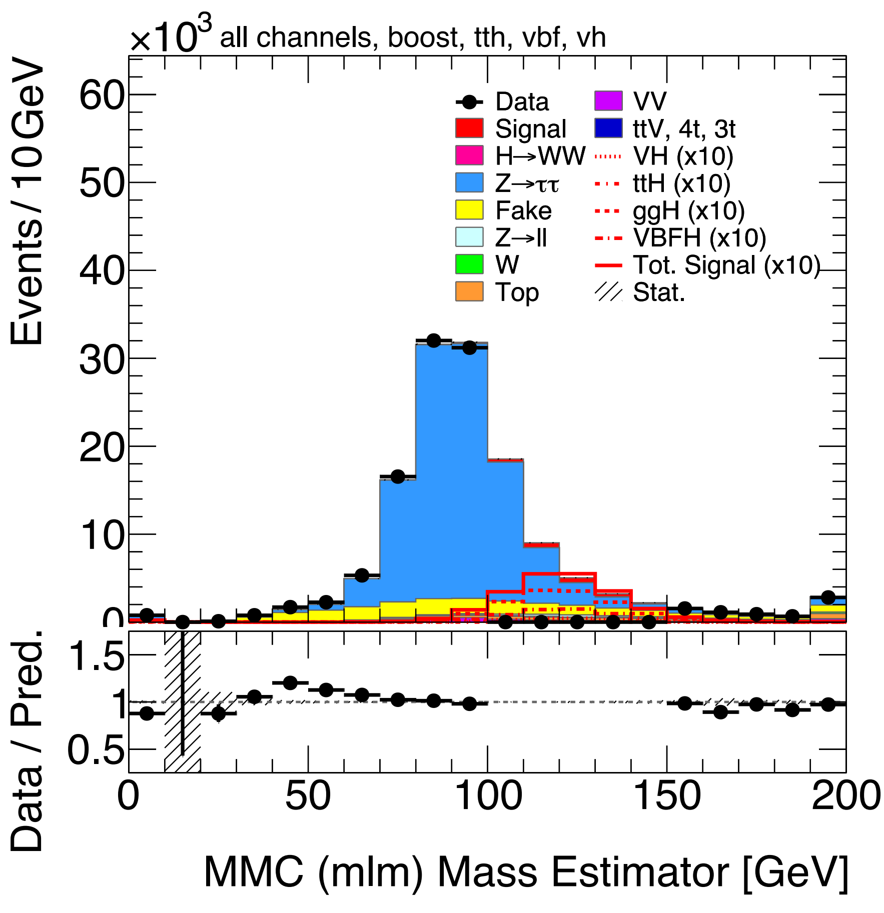

Figure 9.29: Background contributions to the $H \to \tau\tau$ analysis, shown for all signal regions combined with the data blinded in the region between 100 and 150 GeV. The main contributions correspond to the $Z \to \tau\tau$ process (blue), processes with fake leptons (yellow) and top processes (orange). Only statistical uncertainties are considered here.

9.7 Fit strategy

With all background contributions properly assessed, the $H \to \tau\tau$ cross section measurement can be performed. Section 9.2 outlined the three types of measurements that are performed in the $H \to \tau\tau$ analysis. They correspond to the following fit setups, with different numbers of parameters of interest:

- An inclusive $pp \to H \to \tau\tau$ cross section.
- Inclusive production mode cross sections (ggH, VBF, VH and $t\bar{t}H$).
- Cross sections in STXS bins.

All fits are based on the STXS-binned result, but differ in how the POI are defined in the likelihood by relying on the merging of parameters when moving towards the more inclusive measurements. The fitting procedure is for the most part identical to that described in Chapter 8. Note that these results aim to measure cross sections (σ), unlike Chapter 8 where the signal strength parameter μ was the parameter of interest. This change reduces the impact of theory uncertainties on the measurement since the cross section prediction σ_{SM} and its related uncertainties are not directly part of the POI. Results are however shown with respect to the SM prediction central value, to facilitate comparisons.

To discriminate against the $Z \to \tau\tau$ backgrounds, the fits are performed using the $m_{\tau\tau}^{\mathrm{MMC}}$ as a binned discriminating variable in the 38 signal regions of the analysis. The binning of the $m_{\tau\tau}^{\mathrm{MMC}}$ distribution is optimised individually in each SR. The various control regions, on the other hand, are used as single-bin histograms. There are 36 embedded $Z \to \tau\tau$ control regions (one per signal region, except for the two $t\bar{t}H$ regions) and 6 top control regions (one for ggH, one for VBF and one for VH for $\tau_e\tau_\mu$ and $\tau_\ell\tau_{\mathrm{had}}$). Additionally, the $t\bar{t}H$ signal regions themselves are used to constrain the normalisation of the Z and top processes. In total, this corresponds to 37 normalisation factors for the Z process and 7 for the top process: one for each process in each control region and one for each process from the joint $t\bar{t}H$ signal regions. This fit model is illustrated in Figure 9.30.

Three types of fit are presented below. The first is a fit performed using an Asimov dataset to assess the sensitivity of the analysis and to evaluate the contributions from the various sources of uncertainty considered. These results are used to optimize the binning of the measurement. The second is a fit using data from the control regions and from the high and low $m_{\tau\tau}^{\mathrm{MMC}}$ areas of the signal regions, where the signal contributions are negligible. This type of fit allows a determination of the top and Z NFs and also of various NPs within the phase space of the analysis. It can be used to identify tensions with the subsidiary measurements that are used as input to the fit. Of particular interest here are the Z process normalisation and the $\tau_{\text{had-vis}}$ reconstruction and identification efficiency NPs. Finally, with the fit model set, the main sources of uncertainty understood and the sensitivity of the analysis determined, a final fit with unblinded data is performed with all NFs and NPs re-fitted using the additional data. The various cross section measurements outlined above are extracted from this final fit.

9.7.1 Signal purity

Figure 9.31 shows the production mode purity of the 38 signal regions, with the $\tau_e\tau_\mu$, $\tau_\ell\tau_{\mathrm{had}}$ and $\tau_{\mathrm{had}}\tau_{\mathrm{had}}$ channels merged. All regions are dominated by the contributions they aim to select, except for the VH regions with low tagger score where a large ggH contamination is seen. The relative background contributions for each of the signal regions are shown in Figure 9.32.

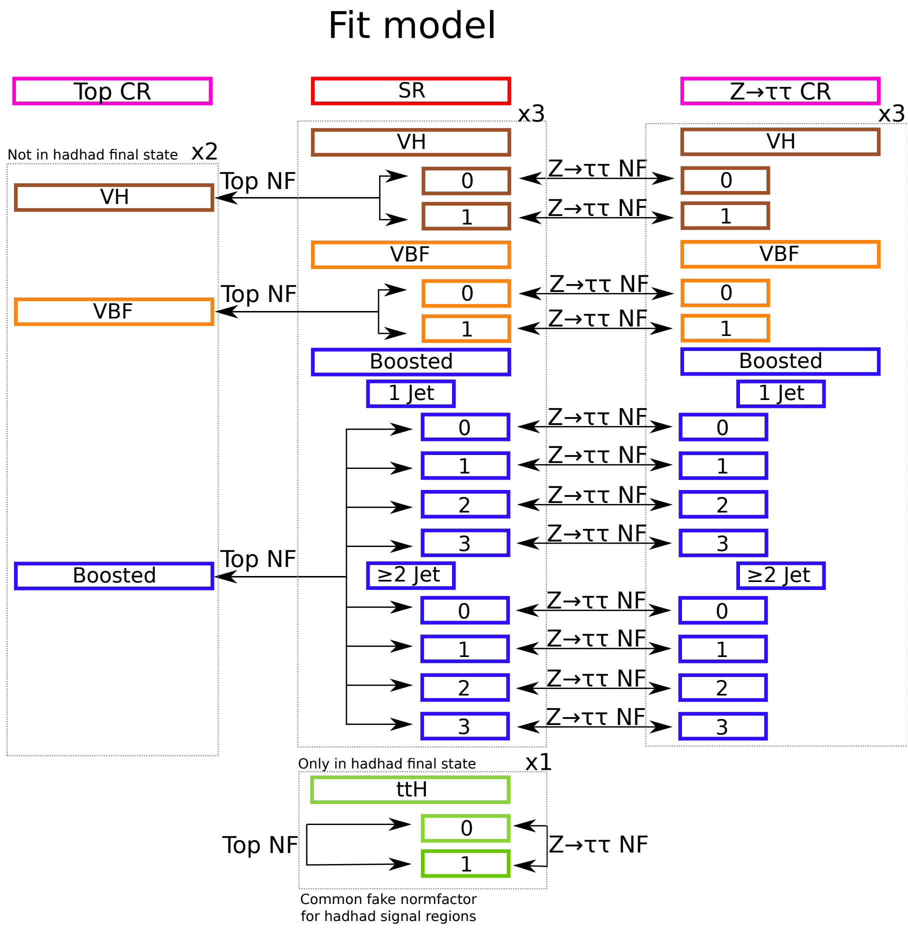

Figure 9.30: Fit model for the $H \to \tau\tau$ cross section measurement.

9.7.2 Systematic uncertainties

The systematic uncertainties considered for this measurement are similar to the experimental and theoretical uncertainties described for the $t\bar{t}H \to$ multilepton analysis, but include improved subsidiary measurements for the various NPs, performed using the full Run 2 dataset. The impact of the NPs is considered both on the overall normalisation of the contributions they are related to as well as on the shape of the $m_{\tau\tau}^{\text{MMC}}$. All embedding systematic uncertainties described in Section 9.5.6, and the uncertainties related to the data-driven estimates are also considered. The most relevant experimental and theoretical uncertainties for the $H \to \tau\tau$ measurements are discussed below.

9.7.2.1 Experimental uncertainties

Uncertainties related to the reconstruction, identification and calibration of $\tau_{\text{had-vis}}$ are important for the $H \to \tau\tau$ cross section measurements. Uncertainties on the $\tau_{\text{had-vis}}$ energy scale include uncertainties derived from a χ^2 fit of the visible mass in $Z \to \mu\tau_{\text{had-vis}}$ events in data, as well as uncertainties related to the difference between the energy and momentum measurements. Uncertainties related to the comparison of truth and reconstructed $\tau_{\text{had-vis}}$ energy distributions in POWHEG PYTHIA8 $Z \to \tau\tau$

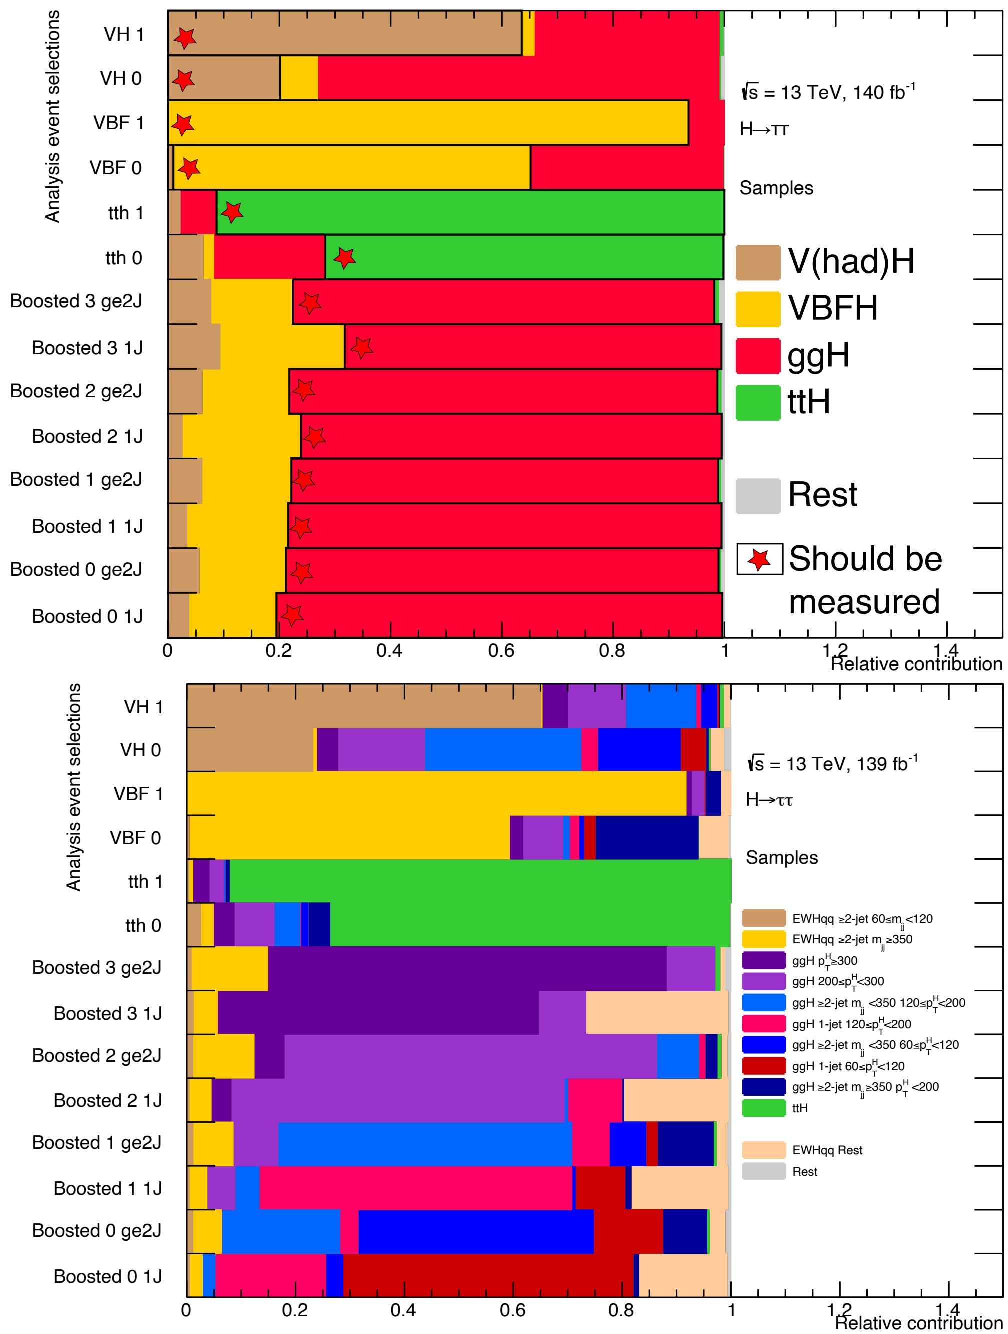

Figure 9.31: Production mode purity of the various signal regions of the analysis, shown inclusively for the tau decay modes for the 4 POI setup (top) and 9 POI setup (bottom).

samples and in samples with different geometries and physics lists[17] are also included. The $\tau_{\text{had-vis}}$ reconstruction uncertainty is derived using MC simulation by comparing the efficiencies obtained in the Powheg Pythia8 $Z \to \tau\tau$ to those obtained in the samples with different geometries and physics lists. The $\tau_{\text{had-vis}}$ identification efficiency and trigger efficiency are measured in data using a tag-and-probe method, as described in Chapter 6 for the identification efficiency, and the uncertainties on these

[17]A physics list collects all the particles, physics processes and production thresholds used in Geant4.

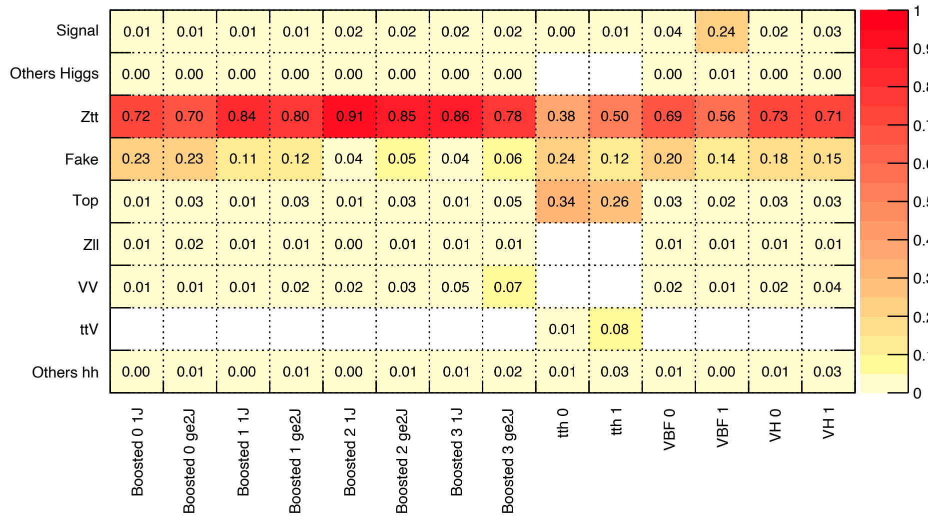

Figure 9.32: Relative background contributions in each of the signal regions of the analysis for the $\tau_e\tau_\mu$, $\tau_\ell\tau_{\text{had}}$ and $\tau_{\text{had}}\tau_{\text{had}}$ channels combined.

measurements are included as NPs. Finally, the electron veto efficiency is obtained in data for both true electrons and true $\tau_{\text{had-vis}}$ using e-$\tau_{\text{had-vis}}$ tag-and-probe methods and in-situ in the $\tau_{\text{1-prong}}$ channel of the $H \to \tau\tau$ analysis (within an $m_{\tau\tau}^{\text{MMC}}$ window of 70-100 GeV) respectively.

Uncertainties related to the energy scale and resolution of jets and of the soft term of $E_{\text{T}}^{\text{miss}}$ are also included. These uncertainties can be quite large for events where a significant contribution from pile-up jets degrades the $m_{\tau\tau}^{\text{MMC}}$ estimation. As mentioned at the beginning of this chapter, the inclusion of a pile-up jet in the $E_{\text{T}}^{\text{miss}}$ calculation significantly biases the $m_{\tau\tau}^{\text{MMC}}$. The energy requirements for jets outlined in the event selections described in Section 9.4 are carefully tuned to reduce this effect. However, variations to the jet energy scale and resolution change the overall effect of these requirements and can result in pile-up jets migrating into the signal regions through a change to their energy that brings them above the requirement threshold. Such migrations therefore significantly affect the $m_{\tau\tau}^{\text{MMC}}$ lineshape and therefore significantly impact the sensitivity of the analysis. Additionally, a changes to the jet energy resolution will typically result in changes to the $E_{\text{T}}^{\text{miss}}$ resolution and thereby affect the width of the $m_{\tau\tau}^{\text{MMC}}$ distribution.

9.7.2.2 Theory uncertainties

The theory uncertainties considered for signal processes are: QCD scale uncertainties, non-perturbative calculation uncertainties (mainly related to the underlying event and to the hadronisation model), and uncertainties on input parameters such as α_S and the PDFs. These uncertainties can impact the total predicted cross section, but also the acceptance into the various signal regions (due for example to jet bin migrations originating from missing higher orders in the calculations) and the shape of the $m_{\tau\tau}^{\text{MMC}}$. They are evaluated in the same way as outlined for the $t\bar{t}H \to$ multilepton analysis, except for a specific treatment of the QCD scale uncertainties which is needed for the ggH process due to accidental cancellations that occur when the scales are varied. The process for evaluating the ggH scale uncertainties follows the recommendations of the LHCXSWG [97].

For the $Z \to \tau\tau$ and $Z \to \ell\ell$ samples, the uncertainties consist of those on the PDFs, α_S, QCD and resummation scales and parton matching. The first is evaluated by producing 100 replicas of event weights of the nominal PDF set used, with the mean and standard deviation used as central value and uncertainty[18]. The central value is also compared to the one obtained when alternative PDF sets are used. The QCD scales are varied as outlined for the $t\bar{t}H \to$ multilepton analysis. The effects of the resummation scale (QSF) and jet-to-parton matching (CKKW) uncertainties are evaluated for this analysis using a truth-level parameterisation derived in Z events as a function of jet multiplicity and transverse momentum of the Z boson.

9.8 Asimov fit results

9.8.1 Inclusive $pp \to H \to \tau\tau$ and production mode cross sections

The sources of uncertainty for the the inclusive fit are listed in Table 9.13. The largest contributions correspond to signal theory uncertainties, MC statistics as well as uncertainties related to jets and $E_{\mathrm{T}}^{\mathrm{miss}}$, due to the sensitivity of the $m_{\tau\tau}^{\mathrm{MMC}}$ to the presence of pile-up jets.

The $\tau_e\tau_\mu$ channel is the least sensitive one, as can be expected from the smaller branching ratio, from the presence of additional neutrinos that complicate the computation of the $m_{\tau\tau}^{\mathrm{MMC}}$ and from the larger background rates. This result highlights the importance of $\tau_{\text{had-vis}}$ reconstruction, identification and calibration for $H \to \tau\tau$ measurements.

Parameter of interest	$\sigma_H \times \mathrm{BR}_{\tau\tau}$
Source of uncertainty	$\Delta\sigma/\sigma_{\mathrm{SM}}$
Signal theory	±0.086
Jet + $E_{\mathrm{T}}^{\mathrm{miss}}$	±0.038
MC stat.	±0.035
Tau	±0.021
Luminosity	±0.018
Fake lepton ($e, \mu, \tau_{\mathrm{had}}$) estimates	±0.018
Z theory	±0.017
b-jets	±0.009
Electrons + muons	±0.005
Total syst.	±0.111
Stat. only	±0.063
Total	±0.127

Table 9.13: Expected results and contributions from systematic uncertainties for the inclusive $pp \to H \to \tau\tau$ cross section measurement, ranked by their impact on the measurement and using an Asimov dataset with $\mu = 1$. The categories "Jet + $E_{\mathrm{T}}^{\mathrm{miss}}$", "Tau", "$b$-jets" and "Electrons + muons" contain uncertainties related to the reconstruction, calibration, identification and triggering of these objects (as applicable).

Because of the large increase in the size of the dataset and of the MC samples available, the improved subsidiary measurements of the NPs provided by the combined performance groups and many analysis

[18]Due to a problem in Sherpa sample productions, the α_S and 100 PDF replicas of the events weights are broken in $Z \to \mu\mu$ samples. The uncertainties derived using $Z \to ee$ samples have therefore been transferred to the $Z \to \mu\mu$ samples.

improvements, this result is more precise than the previous measurement using 36 fb^{-1} of data, which had an expected sensitivity of $1.00^{+0.28}_{-0.25}$.

Similar results are observed for the production mode cross section setup, as outlined in Table 9.14. The most sensitive channel is VBF due to high signal purity, followed by ggH, which benefits from large statistics.

Production mode	ggH	VBF	VH	$t\bar{t}H$
Source of uncertainty	$\Delta\sigma/\sigma_{\mathrm{SM}}$			
Signal theory	±0.229	±0.107	±0.082	±0.208
Jet + $E_{\mathrm{T}}^{\mathrm{miss}}$	±0.121	±0.044	±0.140	±0.203
MC stat.	±0.090	±0.043	±0.229	±0.308
Tau	±0.051	±0.017	±0.043	±0.058
Luminosity	±0.021	±0.016	±0.012	±0.015
Fake lepton $(e, \mu, \tau_{\mathrm{had}})$ estimates	±0.050	±0.024	±0.093	±0.031
Z theory	±0.037	±0.019	±0.042	±0.043
b-jets	±0.009	±0.008	±0.007	±0.077
Electrons + muons	±0.013	±0.002	–	±0.006
Total syst.	±0.285	±0.129	±0.325	±0.562
Stat. only	±0.137	±0.115	±0.493	±0.879
Total	±0.315	±0.172	±0.583	±1.012

Table 9.14: Expected results and contributions from systematic uncertainties for the production mode cross section measurements. The categories "Jet + $E_{\mathrm{T}}^{\mathrm{miss}}$", "Tau", "$b$-jets" and "Electrons + muons" contain uncertainties related to the reconstruction, calibration, identification and triggering of these objects (as applicable).

9.8.2 Cross sections in STXS bins

Inclusive and production mode cross section measurements can be expanded upon using the STXS framework [99]. In the context of the full Run 2 $H \to \tau\tau$ measurement, the binning from Stage 1.2 is used as reference with a few adjustments made to reflect the sensitivity of the analysis. The binning for the analysis is compared to the reference STXS binning in Figures 9.33-9.35. Note that the VBF and VH production modes are jointly referred to as the electroweak $qq \to H + qq$ production since the two processes become ambiguous for $V \to qq$.

The STXS measurement consists of 6 ggH bins, 2 electroweak $qq \to H + qq$ bins and 1 $t\bar{t}H$ bin. The results of the STXS fit are consistent with those of the other two cross section measurements, with the same three sources of uncertainty playing a dominant role as highlighted in Table 9.15. In some bins, uncertainties related to the fake lepton estimations and Z theory uncertainties also become important.

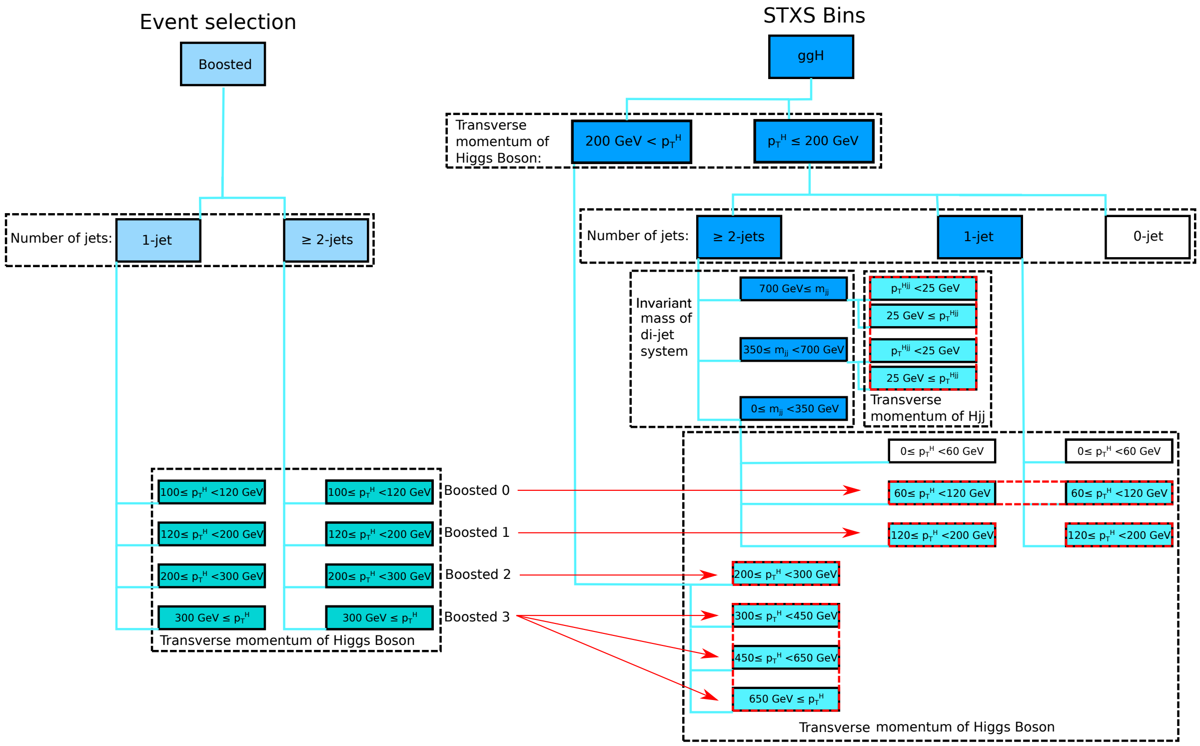

Figure 9.33: Bins of the $H \to \tau\tau$ measurement (left) compared to STXS Stage 1.2 bins (right) for the ggH production mode. Red arrows highlight how the bins of the $H \to \tau\tau$ measurement relate to those of Stage 1.2.

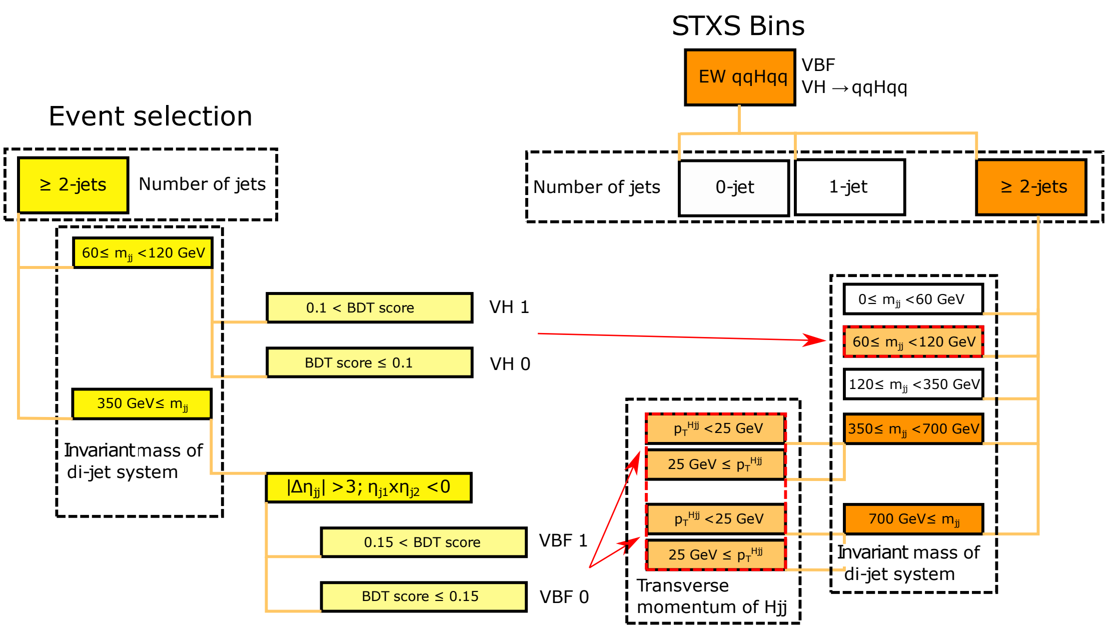

Figure 9.34: Bins of the $H \to \tau\tau$ measurement (left) compared to STXS Stage 1.2 bins (right) for the electroweak $qq \to H + qq$ production mode. Red arrows highlight how the bins of the measurement relate to those of Stage 1.2.

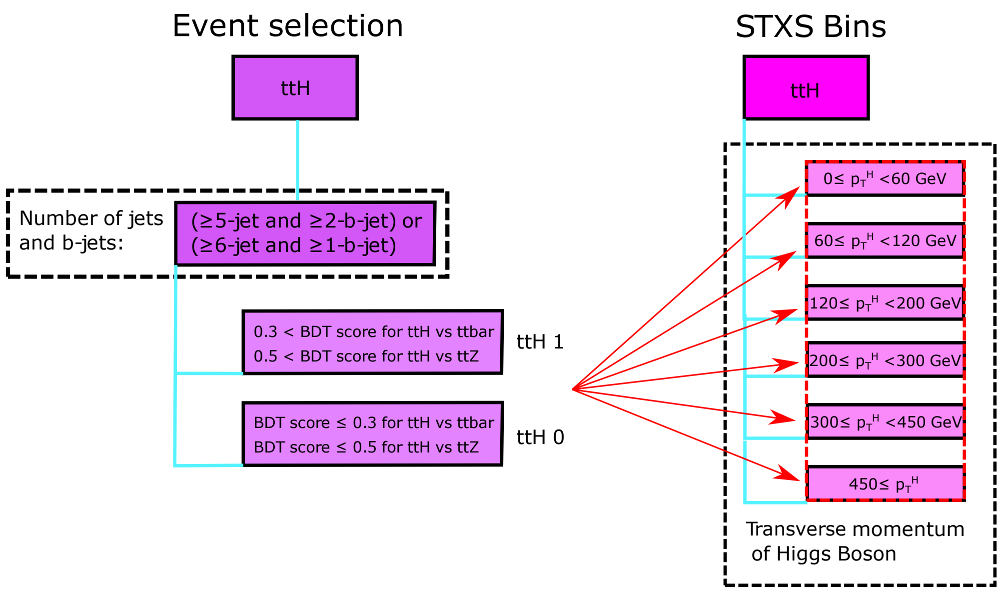

Figure 9.35: Bins of the $H \to \tau\tau$ measurement (left) compared to STXS Stage 1.2 bins (right) for the $t\bar{t}H$ production mode. Red arrows highlight how the bins of the measurement relate to those of Stage 1.2.

Production mode	ggH	ggH	ggH	ggH	ggH	ggH	$qq \to H + qq$	$qq \to H + qq$	$t\bar{t}H$
# of jets	≥ 2	≥ 1	1	≥ 2	≥ 1	≥ 1	≥ 2	≥ 2	
m_{jj} (GeV)	≥ 350	(< 350)		< 350			$\geq 60, < 120$	≥ 350	
p_T^H (GeV)	< 200	$\geq 60, < 120$	$\geq 60, < 120$	$\geq 120, < 200$	$\geq 200, < 300$	≥ 300			
Source of uncertainty	$\Delta\sigma/\sigma_{\mathrm{SM}}$								
Signal theory	±0.721	±0.137	±0.225	±0.321	±0.261	±0.347	±0.056	±0.123	±0.209
Jet + $E_{\mathrm{T}}^{\mathrm{miss}}$	±0.732	±0.335	±0.270	±0.301	±0.167	±0.108	±0.135	±0.036	±0.203
MC stat.	±0.665	±0.374	±0.240	±0.259	±0.109	±0.130	±0.225	±0.049	±0.307
Tau	±0.375	±0.181	±0.117	±0.089	±0.043	±0.069	±0.038	±0.010	±0.057
Luminosity	±0.037	±0.038	±0.029	±0.014	±0.019	±0.020	±0.012	±0.016	±0.016
Fake lepton estimates	±0.404	±0.214	±0.064	±0.086	±0.036	±0.048	±0.093	±0.014	±0.029
Z theory	±0.290	±0.334	±0.191	±0.113	±0.030	±0.051	±0.041	±0.019	±0.044
b-jets	±0.036	±0.066	±0.044	±0.028	±0.022	±0.022	±0.021	±0.009	±0.077
Electrons + muons	±0.037	±0.017	±0.039	±0.047	±0.013	±0.020	±0.015	±0.002	-
Total syst.	±1.330	±0.665	±0.496	±0.589	±0.350	±0.441	±0.309	±0.145	±0.607
Stat. only	±1.077	±0.476	±0.361	±0.460	±0.268	±0.411	±0.486	±0.141	±0.879
Total	±1.695	±0.812	±0.609	±0.741	±0.440	±0.602	±0.568	±0.201	±1.012

Table 9.15: Expected results and contributions from systematic uncertainties for the cross section measurements in STXS bins. The categories "Jet + $E_{\mathrm{T}}^{\mathrm{miss}}$", "Tau", "$b$-jets" and "Electrons + muons" contain uncertainties related to the reconstruction, calibration, identification and triggering of these objects (as applicable).

9.9 Sideband fit results

Additional fits are performed including the control regions and studies are made by also including the regions with low and high $m_{\tau\tau}^{\mathrm{MMC}}$ values (sidebands), where the signal contributions are low, to get a more accurate estimate of the background contributions. An improvement to the subsidiary measurements is achieved for many NPs, in particular those that the analysis is quite sensitive to such as the jet energy scale and resolution.

In the sideband fit studies, large deviations (pulls) are specifically observed for the $\tau_{\text{had-vis}}$ reconstruction and identification efficiency NPs. This is due to the fact that the inclusion of the embedded CRs in the fit allows for an independent measurement of the Z process normalisation, which complements the simultaneous measurement of the Z normalisation and the $\tau_{\text{had-vis}}$ efficiencies in the sideband regions. Indeed, because of the different kinematic regimes considered in the various SRs, without the embedding CRs it would be difficult to resolve whether differences between the $\tau_e\tau_\mu$, $\tau_\ell\tau_{\text{had}}$ and $\tau_{\text{had}}\tau_{\text{had}}$ channels are due to a mismeasurement of $\tau_{\text{had-vis}}$ efficiencies in the subsidiary measurements provided by the combined performance group or because of a kinematic-dependent mismodelling of the Z process. Because the reconstruction and identification $\tau_{\text{had-vis}}$ efficiency measurements provide similar degrees of freedom to the likelihood model of the fit, both are pulled in the sideband fit. The $\tau_{\text{had-vis}}$ efficiency NPs are shown in Table 9.16. The Z NFs used in this fit are those of Figure 9.24.

Nuisance parameter	Pull ± constraint
RNN $\tau_{\text{had-vis}}$ identification efficiency	-1.50 ± 0.79
$\tau_{\text{had-vis}}$ reconstruction efficiency	-1.27 ± 0.84
$\tau_{\text{had-vis}}$ electron veto efficiency	-0.00 ± 0.99

Table 9.16: The $\tau_{\text{had-vis}}$ nuisance parameters from the fit including the top and Z control regions as well as the sidebands. Pulls are observed for the reconstruction and RNN identification efficiency NPs.

Since the combined performance group measures the $\tau_{\text{had-vis}}$ reconstruction efficiency and its uncertainty using only MC samples[19], some doubts remain as to the reliability of the measurement. At the time of writing, due to the observed pulls an in-situ measurement is being performed for the $\tau_{\text{had-vis}}$ reconstruction efficiency in the $H \to \tau\tau$ sidebands by inflating the uncertainty on the subsidiary measurements by a factor of 4 to give the fit additional freedom. This results in efficiency measurements that differ from those provided by approximately 4% per $\tau_{\text{had-vis}}$. A similar in-situ measurement will also be performed for the final unblinded fit.

With the NPs allowed to take on the values determined by the sideband and control regions, the modelling near the Higgs peak is excellent, as can be observed in Figures 9.36-9.38 for the boost region with $120 \leq p_{\mathrm{T}}(H) < 200$ (boost 1) and 1 jet, increasing the confidence that the likelihood model appropriately describes the processes within the phase space of the $H \to \tau\tau$ analysis.

[19] Unlike the other $\tau_{\text{had-vis}}$ efficiency measurements, which are performed with data.

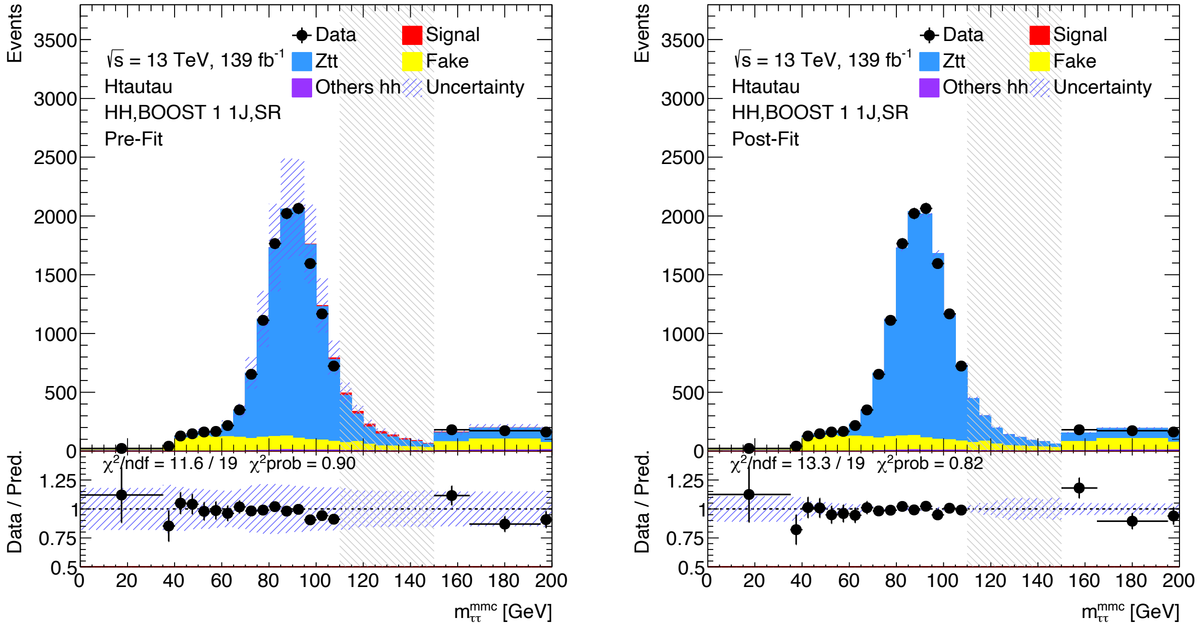

Figure 9.36: Distribution of the $m_{\tau\tau}^{\mathrm{MMC}}$ before (left) and after (right) the sideband fit in the boost 1 region with 1 jet of the $\tau_{\mathrm{had}}\tau_{\mathrm{had}}$ channel, with all systematic uncertainties taken into account. Data bins with contributions from $H \to \tau\tau$ events above 5% or within the range 110-150 GeV are blinded.

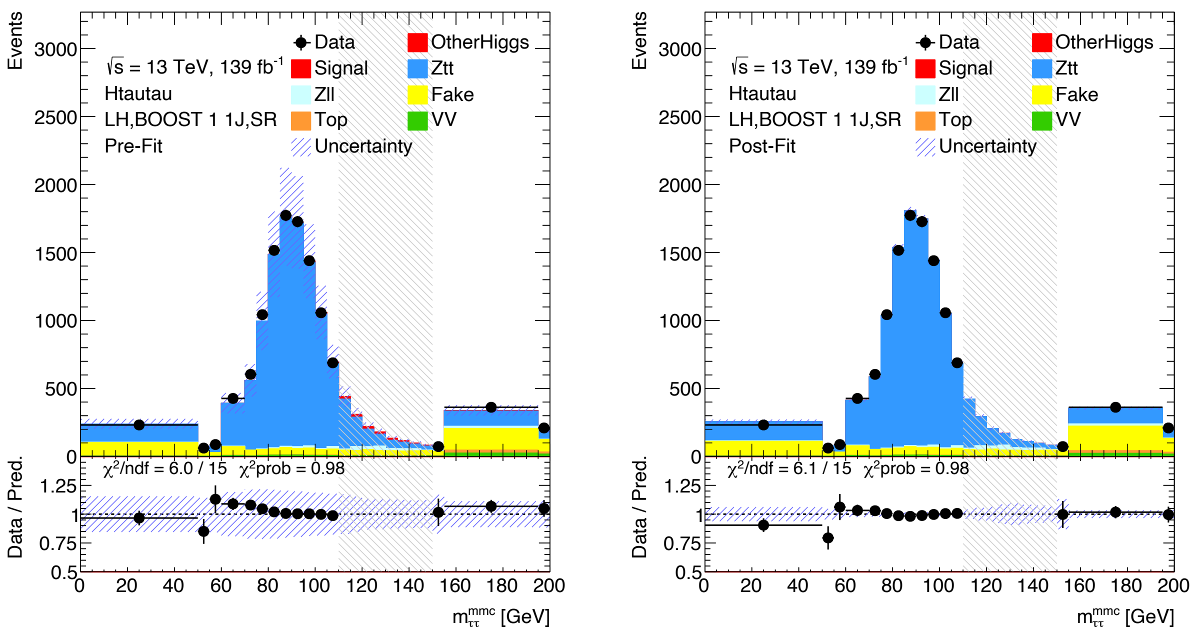

Figure 9.37: Distribution of the $m_{\tau\tau}^{\mathrm{MMC}}$ before (left) and after (right) the sideband fit in the boost 1 region with 1 jet of the $\tau_{\ell}\tau_{\mathrm{had}}$ channel, with all systematic uncertainties taken into account. Data bins with contributions from $H \to \tau\tau$ events above 5% or within the range 110-150 GeV are blinded.

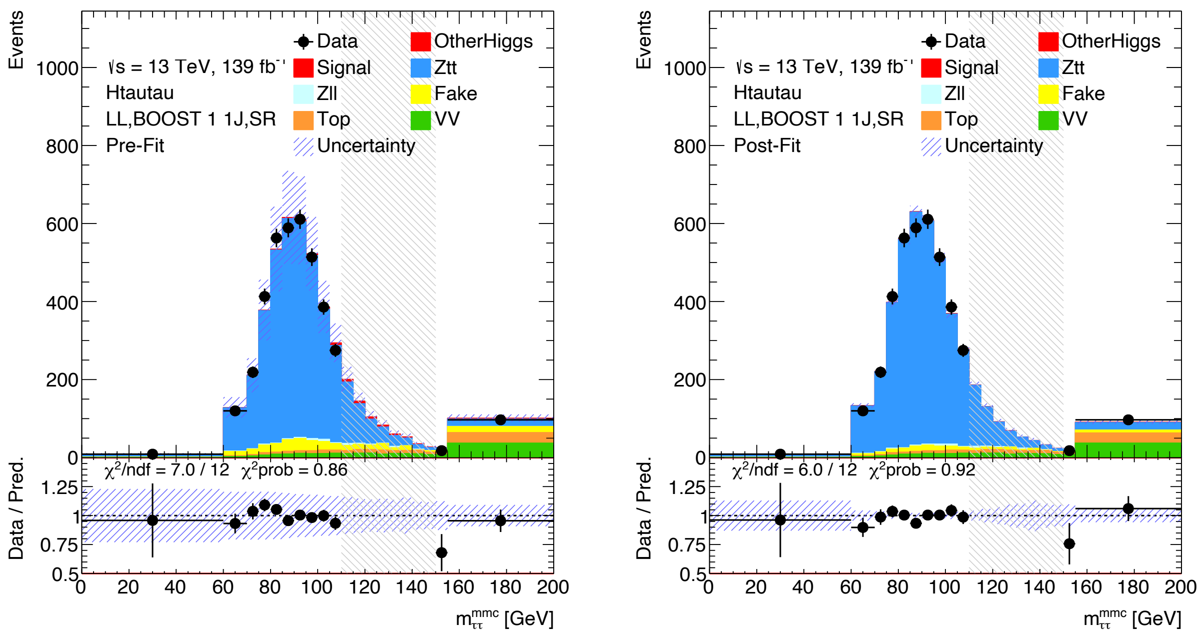

Figure 9.38: Distribution of the $m_{\tau\tau}^{\mathrm{MMC}}$ before (left) and after (right) the sideband fit in the boost 1 region with 1 jet of the $\tau_e\tau_\mu$ channel, with all systematic uncertainties taken into account. Data bins with contributions from $H \to \tau\tau$ events above 5% or within the range 110-150 GeV are blinded.

9.10 Unblinded fit results

With the data unblinded, the inclusive $pp \to H \to \tau\tau$ cross section is measured to be $\sigma_{pp\to H\to\tau\tau} = 3.15 \pm 0.20$ (stat.) $^{+0.40}_{-0.36}$ (syst.) pb, in agreement with the SM prediction of 3.46 ± 0.13 pb. The measurement is shown in Figure 9.39 for the three channels combined, normalised to the SM expectation. For comparison, the measurement performed using the CMS detector with the full Run 2 dataset found a value of $\sigma_{pp\to H\to\tau\tau} = 2.960^{+0.394}_{-0.370}$ pb [80].

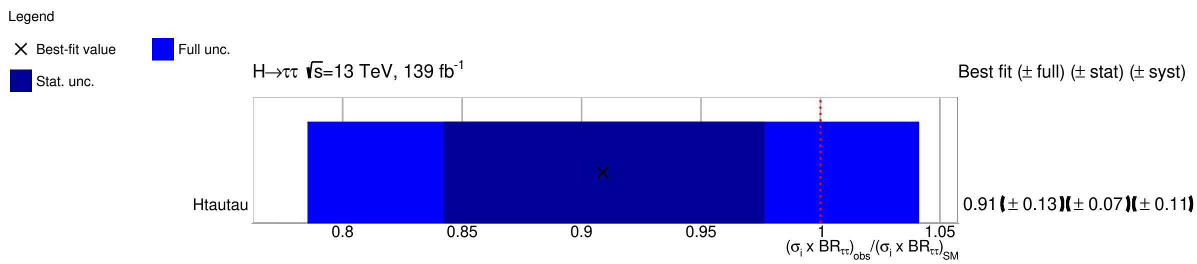

Figure 9.39: Inclusive $pp \to H \to \tau\tau$ cross section measurement, normalised to the SM prediction (red line).

Finally, the results for the production mode measurements are presented in Figure 9.40 and the ones for the STXS measurements are shown in Figure 9.41. These results are compatible with the SM expectation as well as with each other within 2 standard deviations in all cases.

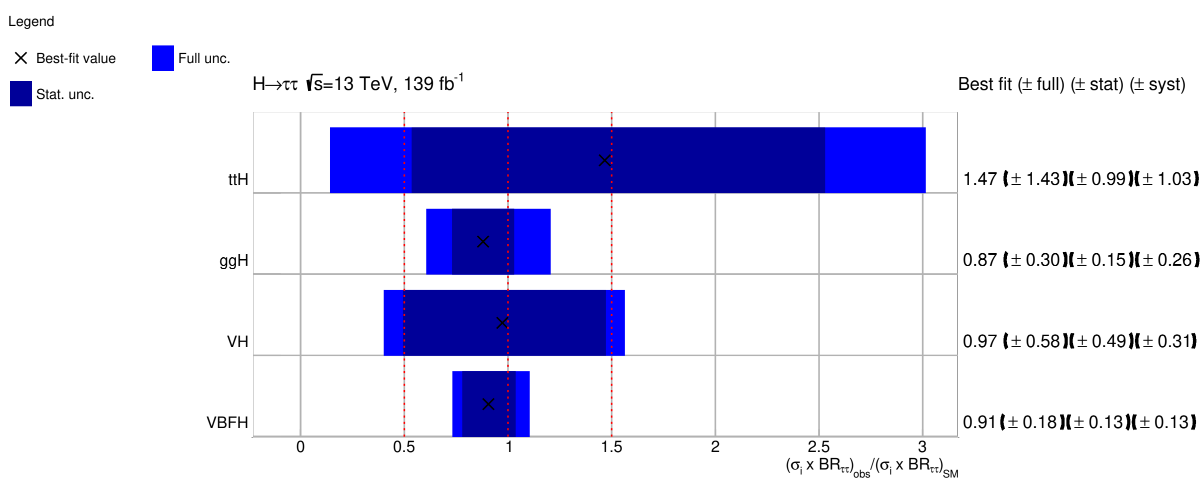

Figure 9.40: Inclusive production mode cross sections: the measurements for $t\bar{t}H$, ggH, VH and VBF are shown from top to bottom and are normalised to the SM predictions (red line at 1).

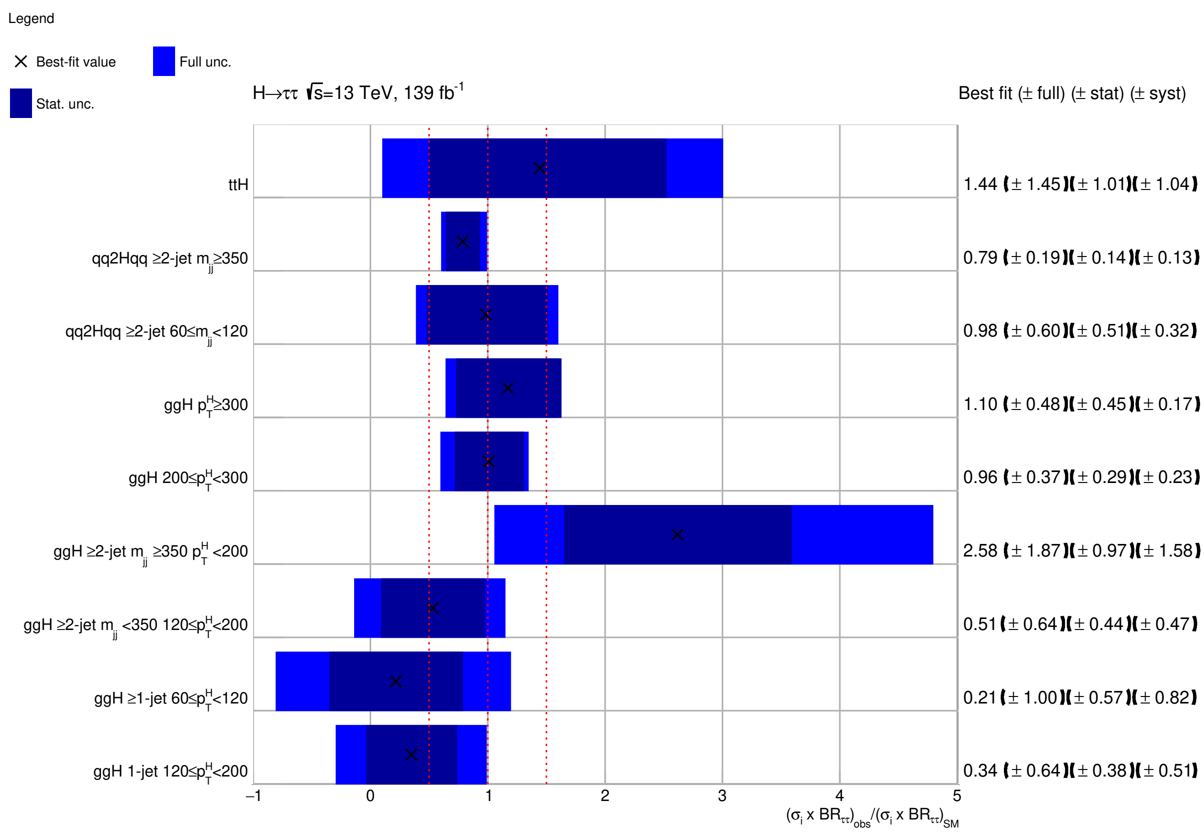

Figure 9.41: Cross section measurements in the 9 STXS bins, normalised to the SM prediction (red line at 1).

Chapter 10

High-Luminosity LHC

The LHC collaborations have so far enjoyed several successful years of collisions leading to the discovery of the Higgs boson and providing constantly-improving precision measurements of the Standard Model. The LHC machine operates at a centre of mass energy of 13 TeV, a design luminosity of 1×10^{34} cm^{-2}s^{-1} (with a peak instantaneous luminosity that actually reached 2.14×10^{34} cm^{-2}s^{-1} during Run 2) and an average number of interactions per bunch crossing of around 40 [28]. A total of approximately 20 fb^{-1} were collected in the first run, with a further 140 fb^{-1} collected during the second run. Building on the success of Run 2, with the discovery of processes such as $t\bar{t}H$ and in-depth studies of previously-measured processes such as $Z \to \tau\tau$ and $H \to \tau\tau$, more runs are planned for the next decade. Run 3 is scheduled to start in 2022 and end in 2024 and future runs are planned for 2027 and beyond. However, further progress will require an upgrade to the LHC accelerator complex. As it stands, it will take another 10 years of running before the statistical uncertainty on measurements can be halved, and many regions of phase space will remain outside of the LHC's reach. While changes to the LHC machine to reach higher centre-of-mass energies are out of the question in the short term, an increase in the accelerator luminosity would allow for an increase in the statistical sensitivity to new physics and for improvements upon the current Standard Model measurements.

The High-Luminosity LHC (HL-LHC) will be installed during the Long Shutdown between Run 3 and Run 4. It will allow for instantaneous luminosities of up to 7.5×10^{34} cm^{-2}s^{-1}, and an integrated luminosity of up to 3000 fb^{-1} over a ten-year period [24]. The upgrades occurring during the shutdown between Runs 3 and 4 are called "Phase-II" upgrades. To cope with the increased particle flux and radiation rates at the HL-LHC due to the increased instantaneous luminosity, the ATLAS detector will also need to be upgraded since it was not built to function in HL-LHC conditions. For the ATLAS detector, the Phase-II upgrades [19, 24] will, among other things, consist in the installation of a new all-silicon tracker (ITk). The ITk upgrade is expected to provide efficient track and vertex reconstruction and precise momentum determination with the HL-LHC conditions, at an increased coverage out to $|\eta| = 4.0$, and will be an essential component of pile-up mitigation strategies. While the ITk will have pixel and strip sections matching the current performance despite the increased pile-up rates, it will no longer contain a TRT.

Due to these changes, the various algorithms presented in Chapter 5 must be adapted before they can be used at the HL-LHC. This chapter will be dedicated to a discussion of the changes that will be made to the $\tau_{\text{had-vis}}$ algorithms introduced in Chapter 6. In particular, changes to the acceptance are required, extending from $|\eta| = 2.5$ in Run 2 to $|\eta| = 4.0$ for the HL-LHC. Reliance on TRT variables for track reconstruction and selection, as well as for electron-τ discrimination must be addressed. The high pile-up rates expected require that the track selection and $\tau_{\text{had-vis}}$ identification algorithms be restructured. The PanTau neutral pion reconstruction suffers from the HL-LHC conditions and must also be re-tuned. Finally, the resolution of the calorimeter-based energy calibration suffers significantly due to the high pile-up rates in the forward regions, introducing limitations to the use of forward $\tau_{\text{had-vis}}$. The updates to the all of these algorithms are described in Section 10.1.

An estimation of the performance of object reconstruction, identification and calibration at the HL-LHC is useful for many reasons. First, it informs the design decisions for several sub-detector components. For example, the pixel pitch was chosen to be 25×100 μm in the barrel of the first layer of the ITk instead of the proposed 50×50 μm due to the observed impact on track reconstruction efficiencies. It also informs HL-LHC machine design decisions. In particular, Section 10.2 highlights the impact of various pile-up rate and density scenarios on the performance of $\tau_{\text{had-vis}}$ algorithms and provides performance benchmarks associated to various LHC environments. Second, object performance estimates allow for more accurate sensitivity estimates to be performed by physics analysis groups, allowing for a better understanding of the possible physics reach of the HL-LHC. Performance functions are used as input for some of the sensitivity projections and represent a parametrised estimate of reconstruction and identification efficiencies, fake rates, and a smearing of calibrated energies. These functions are derived based on the performance observed in the studies presented in the following sections. Finally, the study of $\tau_{\text{had-vis}}$ reconstruction, identification and calibration at the HL-LHC provides insight into the limitations of the current algorithms and helps guide future developments and improvements.

10.1 Tau lepton reconstruction, identification and calibration

10.1.1 Samples

Samples with Drell-Yan $\to \tau\tau$, Drell-Yan $\to ee$ and dijet events overlaid with additional events to replicate the high pile-up conditions of $\langle\mu\rangle = 200$ are used for HL-LHC studies. Changes to the ATLAS detector geometry, such as the introduction of the ITk, are included in the simulation. The samples can be used in a first step to study the performance of the reconstruction, identification and calibration algorithms from Run 2 in HL-LHC conditions. In a second step, they may be used to optimize the parameters of these algorithms to recover any performance losses, or to improve upon the existing algorithms for further performance gains.

The HL-LHC samples are compared to their Run 2 counterparts in Table 10.1. The comparison illustrates the higher demands on the $\tau_{\text{had-vis}}$ reconstruction and identification algorithms due to increased rates of pile-up, in particular resulting in increased amounts of tracks to sort through at reconstruction level and of $\tau_{\text{had-vis}}$ candidates to sort through at the identification stage. It is important to note that a robust track selection is expected to help with identification of pile-up jets through an accurate classification of tracks into the pile-up category[1].

[1] Recall from Chapter 6 that the track selection step for $\tau_{\text{had-vis}}$ reconstruction relies on a 2-step classification into four categories: τ tracks, conversion tracks, tracks from the underlying event and pile-up tracks.

	Run 2	HL-LHC
Number of $\tau_{\text{had-vis}}$ candidates per event	4 ± 2	46 ± 6
Number of tracks per $\tau_{\text{had-vis}}$ candidate in the event	23	599

Table 10.1: Increase in the number of tau candidates in Drell-Yan $\rightarrow \tau\tau$ events within the Run 2 and HL-LHC samples (of which two are real $\tau_{\text{had-vis}}$), and number of tracks in the event divided by the number of $\tau_{\text{had-vis}}$ candidates.

10.1.2 Reconstruction

A track selection MVA similar to the one described in Section 6.1 is trained using the same variable set as in Run 2 but removing TRT-related variables. The Drell-Yan $\rightarrow \tau\tau$ sample mentioned in Section 10.1.1 is used for this step. This training results in a poor reconstruction efficiency for $\tau_{\text{had-vis}}$ in HL-LHC conditions (also referred to as upgrade conditions in this chapter). This is due to the fact that pile-up tracks tend to be softer than those of $\tau_{\text{had-vis}}$ coming from the hard-scatter. The algorithm "learns" to tighten the selection on momentum-related variables to compensate for the increased rate of pile-up tracks in the HL-LHC, leading to a decrease in reconstruction efficiency particularly for $\tau_{\text{had-vis}}$ candidates with 3 prongs, given that the tracks of 3-prong $\tau_{\text{had-vis}}$ are on average softer than those of 1-prong $\tau_{\text{had-vis}}$.

In order to improve the performance, a cut-based track selection is implemented before the training reduces the amount of pile-up tracks. Tracks that remain are more similar to those of $\tau_{\text{had-vis}}$ candidates, allowing the algorithm to focus on the remaining differences. This pre-training selection is based on the Run 1 track selection [38], although with looser requirements, and aims to select τ-like tracks. It includes requirements on the track momentum ($p_{\text{T}} > 1$ GeV), the number of hits within the ITk (at least 6 of which at least one should be in the pixel sub-system) and the impact parameters ($d_0 < 1$ mm and $z_0 \sin\theta < 1.5$ mm).

The classification from the MVA results in reconstruction efficiencies of 86% and 61% for 1- and 3-prong $\tau_{\text{had-vis}}$ respectively, with the 3-prong $\tau_{\text{had-vis}}$ efficiency suffering with respect to the 1-prong case due to all three tracks needing to be properly reconstructed, and because these tracks tend to be softer and less isolated than 1-prong $\tau_{\text{had-vis}}$ tracks, leading on average to poorer scores of the track selection algorithm. In comparison, the Run 2 performance is of 95% (83%) reconstruction efficiency for 1-(3-) prong candidates, indicating that a significant degradation in the performance occurs when attempting to use the algorithm in HL-LHC conditions despite a complete re-training. The loss of performance is due to a three-fold increase in misclassification of τ tracks as pile-up tracks resulting in a 3-4% efficiency loss for τ tracks despite the implementation of the cut-based selection. Misclassification of pile-up tracks as τ tracks, while also increased in the HL-LHC conditions, has a lower impact on the performance of this algorithm. Indeed, in general the track selection algorithm is more likely to miss a τ track than to incorrectly select a pile-up track as a τ track, as can be observed in Figure 10.1. When restricting the upgrade training to a $|\eta| < 2.5$ region, the efficiencies are of 91% and 73% indicating that a large part of the loss in performance is due to the extended kinematic reach of upgrade $\tau_{\text{had-vis}}$ reconstruction, but that the reconstruction performance from Run 2 is not recovered even within the more central η range. Track selection efficiencies are one of the main limiting factors to the use of $\tau_{\text{had-vis}}$ in the forward region at the HL-LHC. This is not expected to have a large impact on the physics potential of the HL-LHC, since many physics signatures (such as those of $H \rightarrow \tau\tau$) result in $\tau_{\text{had-vis}}$ being predominantly produced within the central region of the detector, but it is still interesting to note this effect.

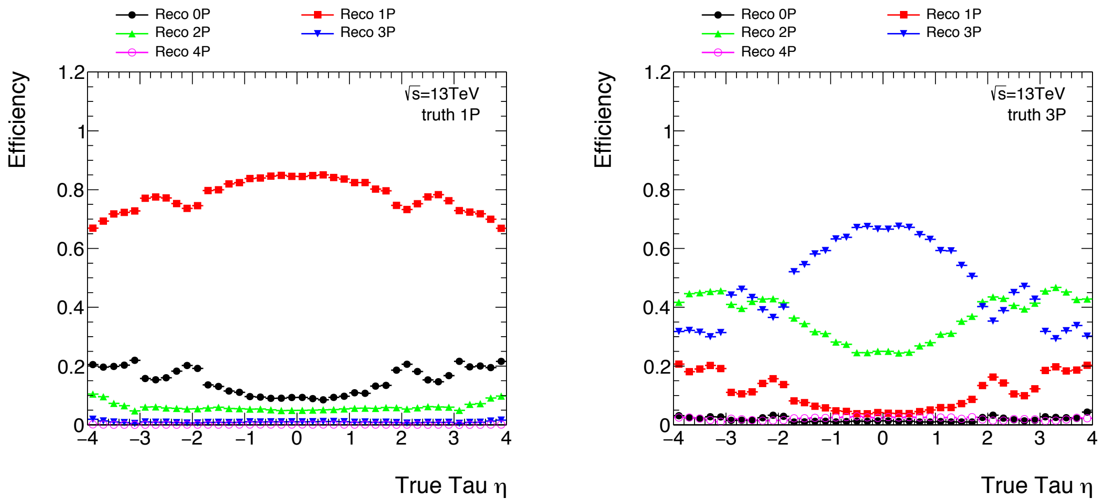

Figure 10.1: Reconstruction efficiency for true 1- (left) and 3-prong (right) $\tau_{\text{had-vis}}$ with a track selection MVA re-tuned for the HL-LHC conditions and the upgraded ATLAS detector. When incorrectly reconstructed, 1-prong $\tau_{\text{had-vis}}$ usually fall into the 0-prong category (black) while 3-prong $\tau_{\text{had-vis}}$ fall into the 2-prong category (green).

A lot of effort has gone into improving track reconstruction in parallel to these studies. Improved track reconstruction will result in lower rates of pile-up tracks and improve the performance of the track selection for τ_{had}. The results presented here are therefore a pessimistic estimate of what will be possible at the HL-LHC.

10.1.3 Energy calibration

For the energy calibration of $\tau_{\text{had-vis}}$ objects, a calorimeter-based calibration similar to the one presented in Section 6.6 yields poor results. In this calibration a the linear regression that aims to account for pile-up contributions to the $\tau_{\text{had-vis}}$ energy is performed using

$$E_{\text{pile-up}}(N_{\text{PV}}, |\eta|, n_P) = A(|\eta|, n_P) \times (N_{\text{PV}} - \langle N_{\text{PV}} \rangle). \tag{10.1}$$

The linearity assumption is seen to be appropriate in most of the $|\eta|$ range except for $2.5 < |\eta| < 4.0$ where the energy contribution from pile-up appears to be correlated with the transverse momentum of the $\tau_{\text{had-vis}}$ candidate (see Figure 10.2).

The breakdown of the linearity assumption in the forward region translates into a degradation of the energy resolution, shown in Figure 10.3. While the $\tau_{\text{had-vis}}$ energy resolution in Run 2 allows for the use of taus with transverse momenta down to 20 GeV, in upgrade conditions the resolution quickly degrades below 30 GeV especially at high η.

These results indicate that the calorimeter-based energy calibration cannot yield results similar to those achieved in Run 2. In particular, $\tau_{\text{had-vis}}$ candidates with p_{T} below 30 GeV can only be used for HL-LHC physics analysis if a different method is used for energy calibration. First attempts at re-training the MVA calibration show similar losses in performance for $|\eta| > 2.5$ in particular due to the

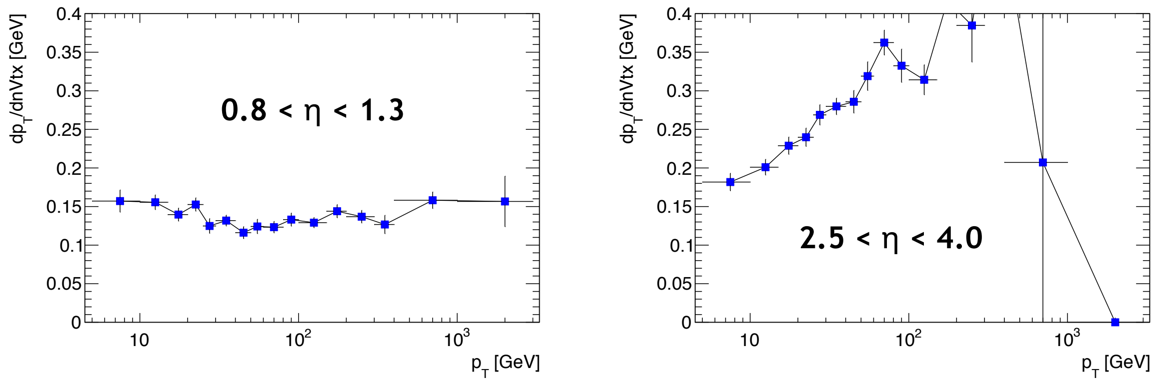

Figure 10.2: Energy deposits from pile-up divided by the number of primary vertices in the event, as a function of the p_{T} of the $\tau_{\text{had-vis}}$ candidate. In the more central regions (left), the energy contribution from pile-up is linear and only depends on the number of vertices in the event. However, in the forward region (right) it is also correlated with the the transverse momentum of the $\tau_{\text{had-vis}}$ candidate.

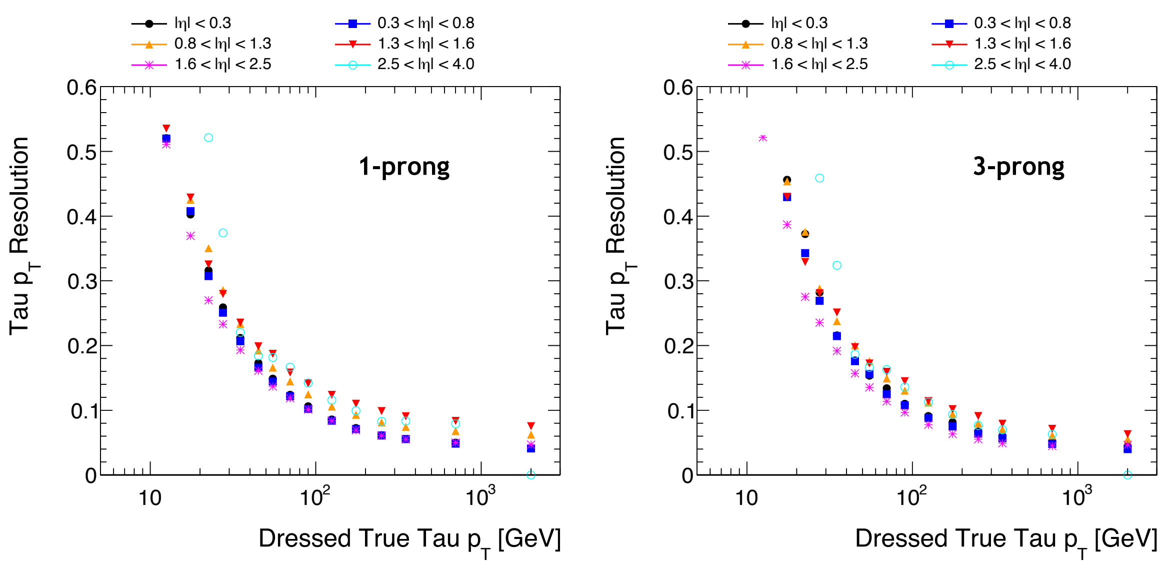

Figure 10.3: Energy resolution when using the baseline correction for 1-prong (left) and 3-prong (right) $\tau_{\text{had-vis}}$ candidates. The dressed p_{T} refers to the momentum of the $\tau_{\text{had-vis}}$ including all photons and charged leptons within the $\tau_{\text{had-vis}}$ cone. The resolution is approximately 25% for 30 GeV $\tau_{\text{had-vis}}$. In comparison, the resolution of the calorimeter-based calibration for 30 GeV $\tau_{\text{had-vis}}$ in Run 2 is of 14% (see Figure 6.4).

inability of Pan Tau to identify neutral pions in the forward region[2]. Even within the central region this calibration provides resolutions approximately 1.5 times worse than those obtained in Run 2, suggesting that alternate methods to calibrate $\tau_{\text{had-vis}}$ need to be considered for the HL-LHC.

[2]The ATLAS calorimeters do not have fine enough segmentation in the forward region to enable neutral pion identification, especially given the HL-LHC pile-up conditions.

10.1.4 Identification

10.1.4.1 Tau identification using Boosted Decision Trees

The BDT identification algorithm inherits information only from those tracks classified as τ-like by the track selection algorithm. To independently assess the ability of the $\tau_{\text{had-vis}}$ identification algorithm to cope with the increased pile-up rates, the track selection algorithm before re-tuning is used in a first step. The algorithm presented in Section 6.2 is re-tuned using the same upgrade-specific Drell-Yan $\rightarrow \tau\tau$ samples that were discussed before. The $\tau_{\text{had-vis}}$ candidates are selected via a truth-matching procedure. The background sample consists of jets from a similar upgrade-specific dijet sample. Both pile-up and hard-scatter jets are considered. The overall performance is poor, with much higher background rates than those observed in Run 2 for all working points.

The cause is once again found to be the increased pile-up levels expected at the HL-LHC. Two distinct populations are clearly visible in the background sample, with pile-up and hard-scatter jets being assigned different average output scores (see Figure 10.4). Since hard-scatter jets are kinematically closer to $\tau_{\text{had-vis}}$, the $\tau_{\text{had-vis}}$ identification algorithm is only able to provide strong discrimination against pile-up jets, with hard-scatter jets ultimately falling into the $\tau_{\text{had-vis}}$ category and contributing to a high fake rate. The plot of the centrality of the energy of the jet[3] in Figure 10.5, an input variable of the identification algorithm, illustrates this kinematic difference between hard-scatter and pile-up jets.

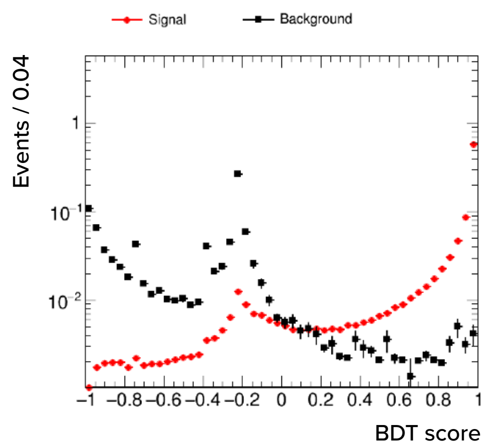

Figure 10.4: Output score distribution, with positive scores indicating a more τ-like object and negative scores indicating jet-like objects. A clear two-peak structure can be observed, with pile-up jets clustering around a score of -1 and hard-scatter jets clustering around -0.2.

Since the re-tuning of the track selection algorithm described above significantly reduces the amount of pile-up jets surviving the reconstruction step, it is added back in. Additionally, it is noted that many physics analyses study final states where the hard-scatter interaction vertex can be reconstructed using the final state topology. The $\tau_{\text{had-vis}}$ identification algorithm can therefore be re-trained with this information: a selection requiring that the $\tau_{\text{had-vis}}$ vertex be no more than 2 mm away from the primary vertex of the event is added to the reconstruction to suppress the pile-up jet contribution. The identification algorithm is re-trained on the new set of reconstructed $\tau_{\text{had-vis}}$ candidates. These

[3] The energy centrality is defined as the fraction of transverse energy (at the EM scale) deposited in the region within $\Delta R < 0.1$ of the barycentre's axis to that deposited within $\Delta R < 0.2$.

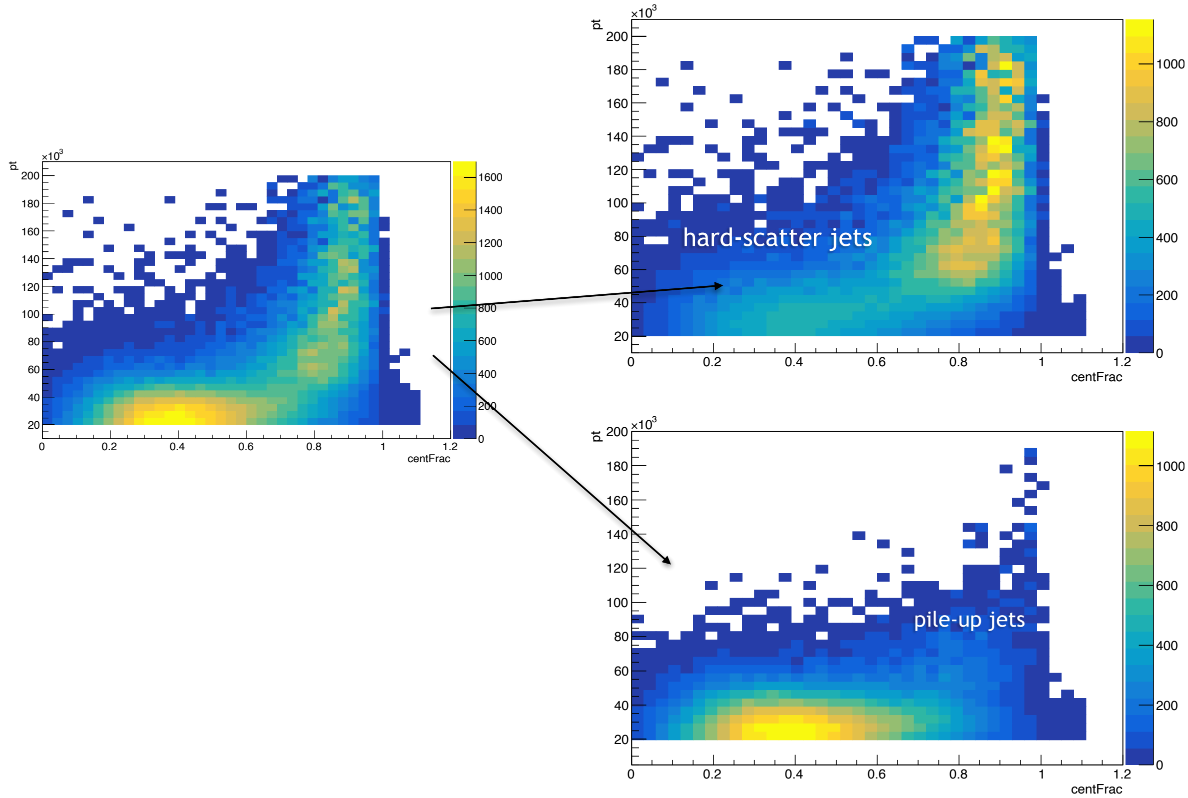

Figure 10.5: Energy centrality (centFrac, left) as a function of the transverse momentum of the jet. Two distinct populations can be seen, as evidenced by the splitting into hard-scatter (top right) and pile-up (bottom right) jets. The hard-scatter jet distribution is similar to that of $\tau_{\text{had-vis}}$ (not pictured here).

two changes significantly improve the performance of the $\tau_{\text{had-vis}}$ identification algorithm, leading it to out-perform the Run-2 algorithm for all working points as shown in Figure 10.6. The performance in the forward region is significantly worse than that at $|\eta| < 2.5$. The $\tau_{\text{had-vis}}$ identification efficiencies are shown in Figure 10.7 for all three working points, and corresponding background rates appear in Figure 10.8.

Further performance improvements would require significant changes to the BDT, such as a re-optimisation of its training parameters, its architecture, or the variables and objects considered. Since the RNN $\tau_{\text{had-vis}}$ identification algorithm resulted in significant performance improvements when it was implemented at the end of Run 2, further upgrades studies focus on this version of the identification algorithm.

10.1.4.2 Tau identification using Recurrent Neural Networks

The algorithm is trained using $\tau_{\text{had-vis}}$ candidates within $|\eta| < 4.0$. As before, a selection ensuring a τ vertex that matches that of the hard-scatter is added before the training, to facilitate the comparison between the performance of the optimal HL-LHC BDT to the HL-LHC RNN. While the RNN algorithm outperforms the BDT one for 1-prong $\tau_{\text{had-vis}}$, achieving twice the background rejection rates for same tau identification efficiency, no improvement can be seen in the case of 3-prong $\tau_{\text{had-vis}}$.

Since the RNN algorithm considers more than only the tracks classified as τ tracks, a new training

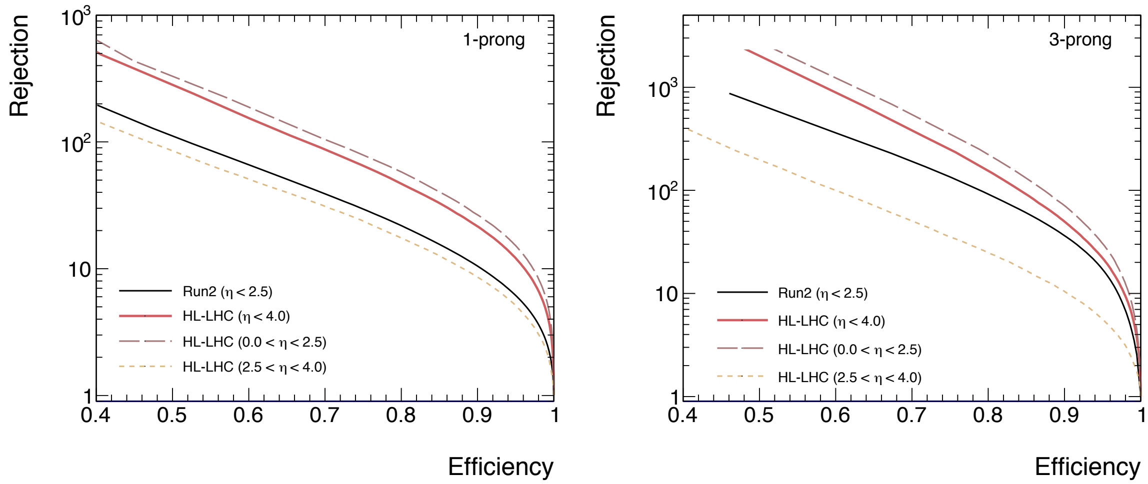

Figure 10.6: Efficiency of tau identification as a function of jet background rejection, for the Run 2 (black) and upgrade (red) tunings. The upgrade algorithm performance is split, with curves for $|\eta| < 2.5$ (dashed burgundy, equivalent to the $|\eta|$ range of Run 2) and $|\eta| > 2.5$ (dashed orange, only available for the upgrade re-tuning) shown separately as well. Here, the Run 2 performance refers to the performance of the tau identification algorithm during Run 2 and not the performance of the Run-2-like algorithm in HL-LHC conditions.

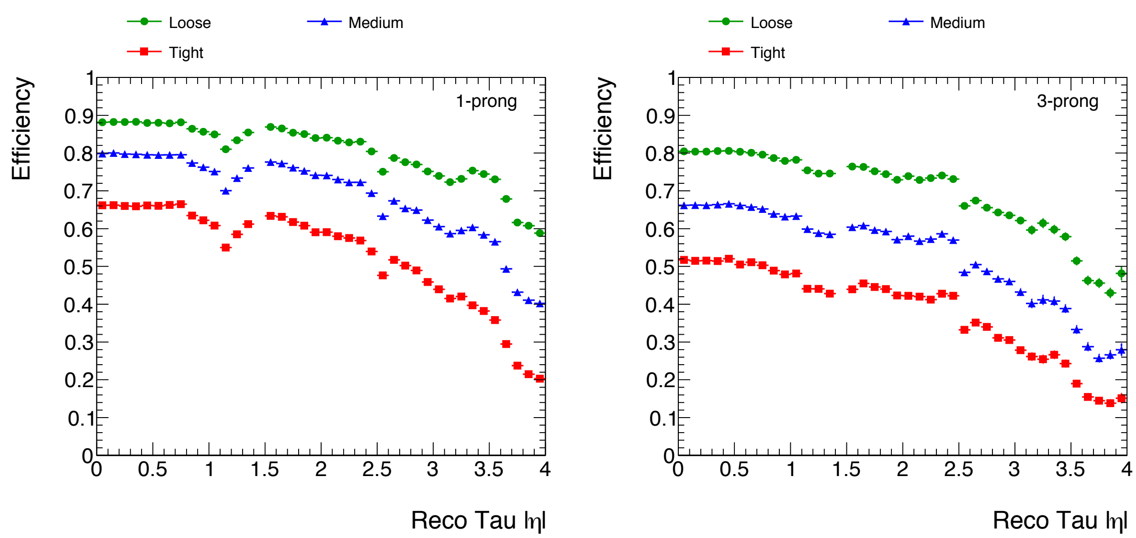

Figure 10.7: Tau identification efficiency for true $\tau_{\text{had-vis}}$ candidates for the loose (green), medium (blue) and tight (red) working points as a function of $|\eta|$ for 1- (left) and 3-prong (right) $\tau_{\text{had-vis}}$ candidates with $p_{\text{T}} > 20$ GeV.

that includes the three track classification scores is performed. The addition of the classification scores provides the RNN with information about the pile-up activity within the $\tau_{\text{had-vis}}$ candidate as well as information about the $\tau_{\text{had-vis}}$ itself. This update makes the vertex requirement obsolete and improves the performance of the RNN $\tau_{\text{had-vis}}$ identification algorithm with respect to the BDT, as illustrated in

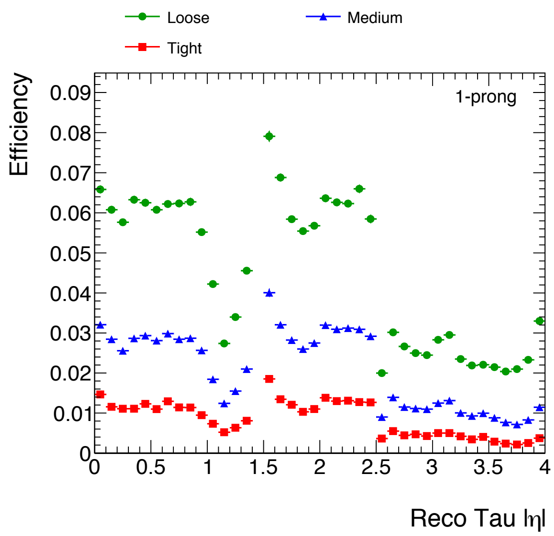

Figure 10.8: Tau identification efficiency for the jet background for the loose (green), medium (blue) and tight (red) working points as a function of $|\eta|$ for 1- (left) and 3-prong (right) $\tau_{\text{had-vis}}$ candidates with $p_{\text{T}} > 20$ GeV.

Figure 10.9. The 3-prong $\tau_{\text{had-vis}}$ identification in particular also benefits from this change.

Despite improvements in performance for both 1- and 3-prong $\tau_{\text{had-vis}}$ candidates when using the RNN, certain issues remain. Specifically, in the forward region ($2.0 < |\eta| < 4.0$), the identification efficiency drops significantly, as was the case for the BDT. Additionally, the somewhat large dependence (of approximately 20%) on the number of vertices in the event is still present when using the RNN, as shown in Figure 10.10.

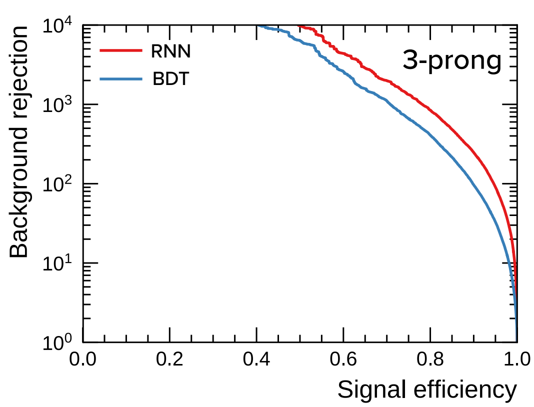

Figure 10.9: $\tau_{\text{had-vis}}$ reconstruction efficiency as a function of jet rejection for the upgrade tunings of the BDT $\tau_{\text{had-vis}}$ identification algorithm (blue) and its RNN counterpart (red) for 1-prong (left) and 3-prong (right) candidates.

While improved track classification in the forward region will have a positive impact on the performance of the $\tau_{\text{had-vis}}$ identification algorithm, the performance in the forward region is likely to remain affected by the high pile-up rates in that region.

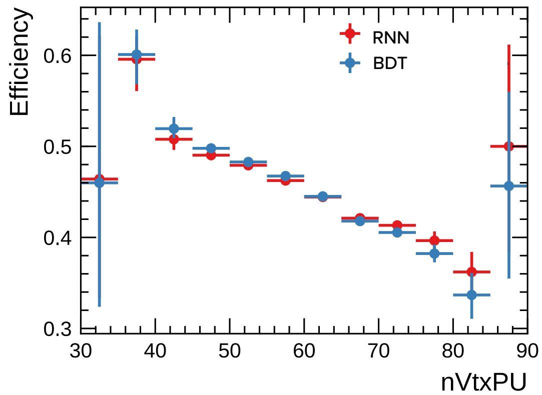

Figure 10.10: Identification efficiency as a function of the number of vertices in the event (right) for the tight working point of the RNN $\tau_{\text{had-vis}}$ identification algorithm re-tuned for upgrade (red), including the track classification score in the training and looking at the full η range. The BDT tight working point is shown in blue for reference.

10.2 Pile-up density studies

There are two scenarios considered for the peak instantaneous luminosities expected at the HL-LHC: 5×10^{34} cm^{-2}s^{-1} and 7.5×10^{34} cm^{-2}s^{-1} corresponding to an average of up to 140 and 200 collisions respectively. The studies presented so far are based on the higher $\langle\mu\rangle = 200$ estimate and provide a lower bound on the expected physics capabilities of the HL-LHC.

Pile-up mitigation strategies are key to the performance of particle reconstruction, identification and calibration algorithms as highlighted in this chapter. The $\tau_{\text{had-vis}}$ algorithms depend on pile-up through two effects. The first is that the calorimeter is sensitive to the overall pile-up levels. For example, as highlighted in the previous section, pile-up energy subtraction is very difficult to perform reliably in the forward region of the ATLAS detector at the HL-LHC. The second involves the pile-up identification strategy used. In the case of $\tau_{\text{had-vis}}$, this strategy is based on the associations of the clusters within the $\tau_{\text{had-vis}}$ to different collisions using the information about the tracks within the object, in particular their impact location along the beam axis. This strategy can in general provide reliable results as the ITk's high granularity enables excellent track-to-vertex association. However, because this strategy takes advantage of the spread in the position of vertices within a bunch crossing due to the finite length of each proton bunch, the density of vertices in a given event has a direct impact on the pile-up mitigation. Since this density is determined by the properties of the colliding beams, variations to the accelerator design can result in significant changes to the performance of the algorithms discussed so far.

This is important because the final running conditions of the HL-LHC have not yet been determined. While a baseline configuration exists, several alternatives are under consideration. These alternatives could provide advantages in terms of running conditions, risk mitigation and costs reductions. These alternatives would also affect the pile-up rates and the pile-up density at the HL-LHC. The studies presented in the previous section are expanded to assess the dependence of the performance on the pile-up rates and density, and provide information to the HL-LHC machine team so that the alternative configurations can be evaluated based on their impact on the physics potential of the HL-LHC.

The track selection and $\tau_{\text{had-vis}}$ identification algorithms optimised for $\langle\mu\rangle = 200$ conditions at the nominal pile-up density are used, but applied to $Z \to \tau\tau$ and dijet samples prepared specifically for this study. A plot of jet rejection[4] versus $\tau_{\text{had-vis}}$ efficiency is shown in Figure 10.11. The performance is compared in $Z \to \tau\tau$ and dijet events for scenarios with nominal and high pile-up densities ρ. A decrease of the order of 2-3% in identification efficiency at same rejection for all working points is seen for the scenario with higher pile-up density. This degradation impacts one- and three-prong $\tau_{\text{had-vis}}$ candidates equally.

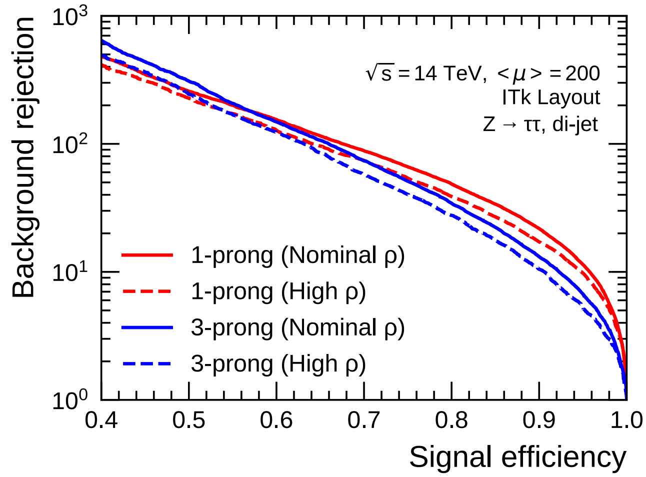

Figure 10.11: Jet rejection as a function of $\tau_{\text{had-vis}}$ identification efficiency for the algorithm optimised for a scenario with lower pile-up density ("nominal ρ", solid lines), compared to a scenario with increased pile-up density ("high ρ", dashed lines). One-prong (red) and three-prong (blue) performance is shown separately.

Figure 10.12 shows the degradation of the jet rejection for the loose identification working point as a function of pile-up density, which affects 3-prong $\tau_{\text{had-vis}}$ candidates more significantly as expected. Since higher overall pile-up rates also result in increased pile-up densities, this figure also provides information about the performance for scenarios with $\langle\mu\rangle = 140$ and $\langle\mu\rangle = 200$. Four regions are identified for reference. The region with $\rho = [1, 1.5]$ corresponds approximately to $\langle\mu\rangle = 140$ and low pile-up densities (achievable for example if crab cavities were used at the HL-LHC). The region with $\rho = [1.5, 2]$ corresponds instead to $\langle\mu\rangle = 200$ with similar densities. Higher density scenarios are associated to the regions $\rho = [2, 3]$ and $\rho = [3, 4]$ approximately for $\langle\mu\rangle = 140$ and $\langle\mu\rangle = 200$ respectively.

Similar studies were performed by the CMS collaboration and the impact of these projections on physics results was also assessed. These studies, in conjunction with running condition, risk mitigation and cost considerations, will help inform the final decision regarding the HL-LHC conditions.

[4]The jet rejection is defined as the inverse of the fake rate, where the fake rate represents the ratio of reconstructed jets with one (three) tracks that pass the RNN selection working point to the total amount of reconstructed jets with one (three) tracks.

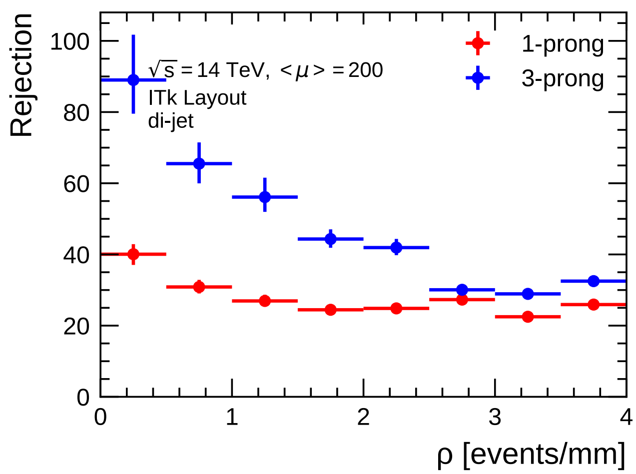

Figure 10.12: Jet rejection as a function of average pile-up density for the loose τ identification RNN working point, with corresponding efficiencies of 85% (75%) for 1(3)-prong candidates. The one-prong (red) and three-prong (blue) performances are shown separately.

10.3 Physics potential of the HL-LHC

The larger datasets obtained with the HL-LHC will allow for measurements of Higgs couplings to second generation fermions[5] and improvements to all couplings measurements. Projections of the precision of many measurements are provided in Yellow Reports [6], and rely on the assumption that the Run 2 performance of the various reconstruction, identification and calibration algorithms can be maintained at the HL-LHC. The assumption of similar performance is based, among other things, on the results presented in this chapter. These results highlight the ability of the current algorithms to provide excellent $\tau_{\text{had-vis}}$ reconstruction and identification; the slight degradation observed is likely to be offset with the use of re-tuned and improved tracking algorithms. While larger performance losses are expected for the $\tau_{\text{had-vis}}$ energy calibration due to the increased pile-up rates, significant progress has been made on the topic in the context of jet calibrations. Methods such as constituent subtraction [8, 9], SoftKiller [17] and PUPPI [10] have been shown to provide excellent performance with regards to pile-up mitigation. The use of these methods in $\tau_{\text{had-vis}}$ reconstruction and energy calibration could allow for the recovery of (or even an improvement upon) the Run 2 performance for $\tau_{\text{had-vis}}$.

The Yellow Report presents many results, of which a few are particularly interesting in the context of this thesis. Projections of the production mode cross section measurements using $H \to \tau\tau$ decays are shown in Table 10.2 and represent improvements of a factor of two on the results presented in Chapter 9. Projections for the ratios of cross sections, discussed in Chapter 8 the context of the $t\bar{t}H \to$ multilepton analysis, are illustrated in Figure 10.13. The uncertainty on the ratio is of approximately 7% at the HL-LHC, in comparison to the 25% uncertainty on the current result. Finally, Figure 10.14 shows a projection of the precision that can be achieved at the HL-LHC for Higgs couplings to fermions and bosons, which is four times smaller than that of current results. All of these results present significant improvements to the current measurements and highlight the extensive physics potential of the HL-LHC.

Aside from these improvements, the HL-LHC will also provide sensitivity to new processes. In particular, evidence for the Higgs self-coupling will likely be reached through measurements of di-Higgs production cross sections. Since channels such as $HH \to bb\tau\tau$ will contribute significantly to these measurements, excellent τ reconstruction, identification and calibration will be essential at the HL-LHC.

Production mode	Source of uncertainty	Fractional error	
		Full Run 2	HL-LHC
ggH	Syst.	±0.28	±0.11
	Stat.	±0.14	±0.03
	Total	±0.32	±0.12
VBF	Syst.	±0.13	±0.07
	Stat.	±0.11	±0.03
	Total	±0.17	±0.08

Table 10.2: Uncertainty on the inclusive production mode cross section measurements using the full Run 2 dataset and using 3000 fb^{-1} of data at the HL-LHC. Only the results for VBF and ggH are included since the projections were made using an earlier $H \to \tau\tau$ result where no VH or $t\bar{t}H$ measurements were available. HL-LHC values from [6].

[5] For example, the $H \to \mu\mu$ process is expected to be observed with a significance of over 9σ [62].

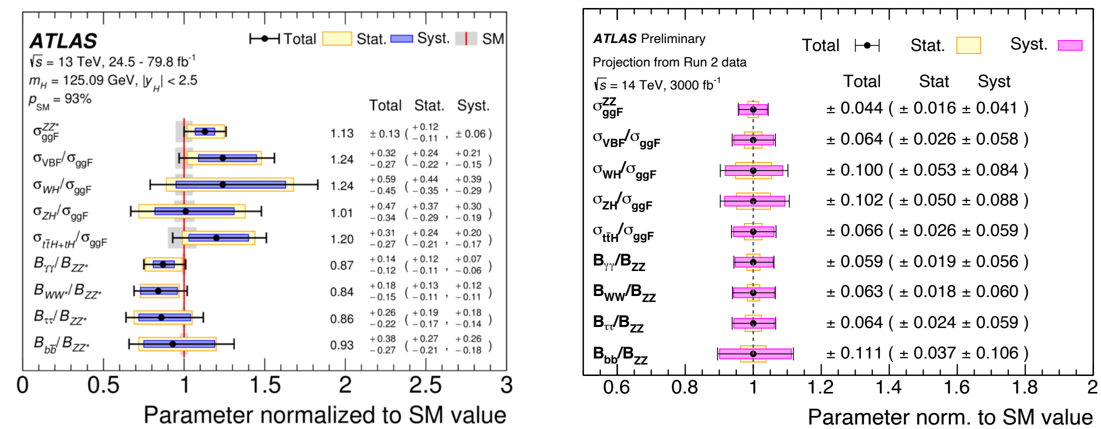

Figure 10.13: Ratios of cross sections and branching fractions with a partial Run 2 dataset (left) and using 3000 fb^{-1} of data at the HL-LHC (right). Note the change of scale on the x-axis. Images from [72] and [6].

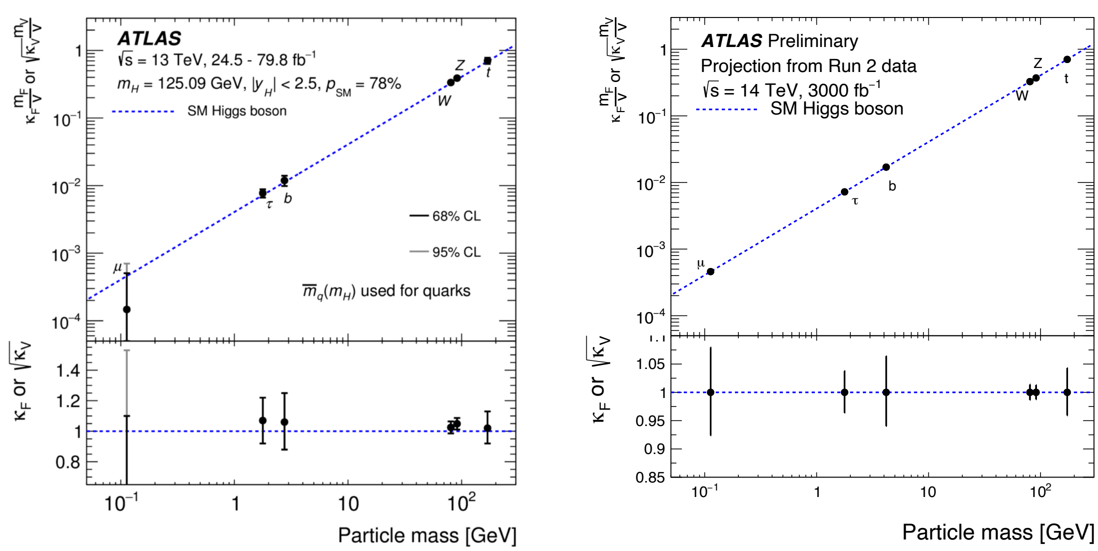

Figure 10.14: Reduced coupling strength modifiers[6] $\kappa_F \frac{m_F}{v}$ for fermions ($F = t,\ b,\ \tau,\ \mu$) and $\sqrt{\kappa_V}\frac{m_V}{v}$ for weak gauge bosons ($V = W,\ Z$) as a function of their masses m_F and m_V and of the VEV (v = 246 GeV) for a partial Run 2 dataset (left) and using 3000 fb^{-1} of data at the HL-LHC (right). Note the change of scale on the ratio plots. Images from [72] and [6].

Chapter 11

Conclusion

Measurements of the properties of the Higgs boson have been conducted by the ATLAS and CMS collaborations since the discovery of this boson during the first run of the LHC in 2012. During Run 1, the Higgs boson's mass was determined with great precision, its spin and parity were found to be 0^+, and the strength of its couplings to other particles were constrained to between 10 and 50%. The second run of the LHC built upon these results, with the addition of new measurements such as the cross section of Higgs production in association to vector bosons and the confirmation of its coupling to down-type quarks through the observation of $H \to b\bar{b}$. In this thesis, an emphasis was put on the measurements of Higgs couplings to third generation fermions.

The interaction of Higgs bosons with top quarks is large, and its strength can be probed at tree-level through the measurement of the $t\bar{t}H$ production cross section. Because of its small contribution to the total Higgs production cross section, the $t\bar{t}H$ process could not be observed within the Run 1 dataset. The Run 2 $t\bar{t}H$ search benefited from a combination of more advanced analysis methods and a larger dataset. The inclusion of new channels, such as the $3\ell + 1\tau_{\text{had}}$ channel described in this thesis, further enhanced the sensitivity to this process, with the search for $t\bar{t}H$ culminating in the observation of the process using a partial Run 2 dataset. The search for the $t\bar{t}H$ process was discussed in particular in the context of channels with $\tau_{\text{had-vis}}$ in the final state, highlighting difficulties associated to $\tau_{\text{had-vis}}$ reconstruction and identification in low-statistics channels, specifically with respect to the estimation of fake backgrounds.

With the existence of the $t\bar{t}H$ process established, its cross section was measured to be $\sigma_{t\bar{t}H} = 670 \pm 90$ (stat.) $^{+110}_{-100}$ (syst.) fb, in agreement with the SM prediction of 508 ± 45 (PDF+α_S) $^{+29}_{-47}$ (QCD scale) fb. This cross section measurement complemented the ggH cross section measurement, which is also sensitive to the Yukawa coupling to the top quark. The ratio $\sigma_{t\bar{t}H+tH}/\sigma_{ggH}$ was found to be $1.20^{+0.31}_{-0.27}$ when normalised to its SM value, indicating compatibility with the SM predictions for the ggH loop contents.

The interaction of Higgs bosons with leptons was also studied in Run 2. Higgs decays to τ leptons are associated to large branching ratios, making the study of the $H \to \tau\tau$ process particularly interesting in this context. The $H \to \tau\tau$ cross section measurement depends on a reliable estimate of its largest background, $Z \to \tau\tau$. Due to the similarity of the two processes, the modelling of this background is difficult to validate. An object-level embedding method was therefore introduced in this thesis. It allows the use of $Z \to ee$ and $Z \to \mu\mu$ events to validate the modelling of $Z \to \tau\tau$ in a phase space similar to the one of the $H \to \tau\tau$ analysis through the implementation of embedding parametrisations that re-scale the momentum of the electrons and muons of $Z \to ee$ and $Z \to \mu\mu$ to better match the distributions expected from the visible τ decay products of $Z \to \tau\tau$. The use of this method, alongside other improvements to the analysis and the availability of a larger dataset, resulted in a measurement of the inclusive $pp \to H \to \tau\tau$ production cross section of $\sigma_{pp\to H\to\tau\tau} = 3.15 \pm 0.20$ (stat.) $^{+0.40}_{-0.36}$ (syst.) pb, in agreement with the SM prediction of 3.46 ± 0.13 pb. Higgs production mode measurements were also included for ggH, VBF, VH and $t\bar{t}H$, and expanded upon the previous $H \to \tau\tau$ measurements of the ggH and VBF cross sections. Finally, binned measurements following the STXS framework were also provided, with no significant disagreements from the SM predictions observed.

Future runs of the LHC will aim to expand upon the results presented here. At the HL-LHC, Higgs couplings to second generation fermions will be possible. With the increased statistics available at the HL-LHC, the precision on known couplings of the Higgs boson will be reduced to below 5%. In combination with measurements performed in regions of phase space probing high energy processes such as vector boson scattering or production of Higgs bosons with high transverse momentum, these results will allow for some of the most stringent tests of the SM yet. Finally, new measurements will be within reach. Extraction of the Higgs self-coupling from di-Higgs processes will be possible for the first time and will allow a measurement of λ. In combination with measurements of m_H and v, this measurement will provide a further probe the self-consistency of the SM. Excellent performance of reconstruction, identification and calibration algorithms will be needed for these results. As was the case for the $t\bar{t}H$ search, the di-Higgs search will rely on the inclusion of many final states, including $HH \to bb\tau\tau$. Maintained performance of $\tau_{\text{had-vis}}$ reconstruction, identification and calibration will therefore continue to be key at the HL-LHC.

Bibliography

[1] F.D. Aaron et al. Combined Measurement and QCD Analysis of the Inclusive $e^+ - p$ Scattering Cross Sections at HERA. *JHEP*, 1001:109, 2010.

[2] S. Agostinelli et al. Geant4: a simulation toolkit. *Nucl. Instrum. Meth. A*, 506:250, 2003.

[3] Reinhard Alkofer. QCD Green functions and their application to hadron physics. *Brazilian Journal of Physics*, 37(1b), 2007.

[4] Guido Altarelli and G. Parisi. Asymptotic Freedom in Parton Language. *Nucl. Phys. B*, 126:298–318, 1977.

[5] The ATLAS and CMS Collaborations. Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV. *Journal of High Energy Physics*, 2016(8), Aug 2016.

[6] The ATLAS and CMS Collaborations. Report on the Physics at the HL-LHC and Perspectives for the HE-LHC. *CERN Yellow Rep.*, 2019.

[7] J. Bernstein. Spontaneous symmetry breaking, gauge theories, the higgs mechanism and all that. *Rev. Mod. Phys.*, 46:7–48, 1974.

[8] P. Berta, L. Masetti, D.W. Miller, and M. Spousta. Pileup and underlying event mitigation with iterative constituent subtraction. *Journal of High Energy Physics*, 2019(8), Aug 2019.

[9] Peter Berta, Martin Spousta, David W. Miller, and Rupert Leitner. Particle-level pileup subtraction for jets and jet shapes. *Journal of High Energy Physics*, 2014(6), Jun 2014.

[10] Daniele Bertolini, Philip Harris, Matthew Low, and Nhan Tran. Pileup per particle identification. *Journal of High Energy Physics*, 2014(10), Oct 2014.

[11] Manfred Böhm, Ansgar Denner, and Hans Joos. *Gauge Theories of the Strong and Electroweak Interaction*. B. G. Teubner, Stuttgart/Leipzig/Wiesbaden, 3 edition, April 2001.

[12] L. Breiman et al. *Classification and Regression Trees*. Chapman Hall, 1984.

[13] Jon Butterworth et al. PDF4LHC recommendations for LHC Run II. *J. Phys. G*, 43:023001, 2016.

[14] Matteo Cacciari and Gavin P. Salam. Pileup subtraction using jet areas. *Physics Letters B*, 659(1-2):119–126, Jan 2008.

[15] Matteo Cacciari, Gavin P Salam, and Gregory Soyez. The anti-kt jet clustering algorithm. *Journal of High Energy Physics*, 2008(04):063–063, Apr 2008.

[16] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. FastJet user manual. *The European Physical Journal C*, 72(3), Mar 2012.

[17] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. Softkiller, a particle-level pileup removal method. *The European Physical Journal C*, 75(2), Feb 2015.

[18] CERN. The Standard Model and beyond. `https://cds.cern.ch/record/1473657/files/SMinfographic_image.png?subformat=`.

[19] The ATLAS Collaboration. ATLAS Phase-II Upgrade Scoping Document. *CERN-LHCC-2015-020*. `http://cds.cern.ch/record/2055248`.

[20] The ATLAS Collaboration. Beyond the Standard Model Physics at the HL-LHC and HE-LHC. *CERN-LPCC-2018-05*. arXiv:1812.07831.

[21] The ATLAS Collaboration. Future physics opportunities for high-density QCD at the LHC with heavy-ion and proton beams. *CERN-LPCC-2018-07*. arXiv:1812.06772.

[22] The ATLAS Collaboration. HL/HE-LHC Physics Workshop Report WG1 : Standard Model physics opportunities. *ATL-COM-PHYS-2018-1673*. `https://cds.cern.ch/record/2649885`.

[23] The ATLAS Collaboration. HL/HE Physics Workshop Report: Higgs physics opportunities. *ATL-COM-PHYS-2018-1658*. `https://cds.cern.ch/record/2649550`.

[24] The ATLAS Collaboration. Letter of Intent for the Phase-II Upgrade of the ATLAS Experiment. *CERN-LHCC-2012-022*. `https://cds.cern.ch/record/1502664`.

[25] The 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. B716 (2012) 1-29*. arXiv:1207.7214.

[26] The ATLAS Collaboration. Observation of Higgs boson production in association with a top quark pair at the LHC with the ATLAS detector. *Phys. Lett. B 784 (2018) 173*. doi:10.1016/j.physletb.2018.07.035.

[27] The ATLAS Collaboration. Opportunities in Flavour Physics at the HL-LHC and HE-LHC. *CERN-LPCC-2018-05*. arXiv:1812.07638.

[28] The ATLAS Collaboration. Peak Luminosity per Fill in 2018. *Public plot*. `https://atlas.web.cern.ch/Atlas/GROUPS/DATAPREPARATION/PublicPlots/2018/DataSummary/figs/peakLumiByFill.pdf`.

[29] The ATLAS Collaboration. Search for the Associated Production of a Higgs Boson and a Top Quark Pair in multilepton final states in pp Collisions at $\sqrt{s} = 13$ TeV with the ATLAS Detector. *CERN-EP-2017-281*. arXiv:1712.08891.

[30] The ATLAS Collaboration. Technical Design Report for the ATLAS Inner Tracker Pixel Detector. *CERN-LHCC-2017-021*. `https://cds.cern.ch/record/2285585`.

[31] The ATLAS Collaboration. Technical Design Report for the Phase-II Upgrade of the ATLAS Tile Calorimeter. *CERN-LHCC-2017-019*. `https://cds.cern.ch/record/2285583`.

[32] The ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *JINST*, 3:S08003. 437 p, 2008. Also published by CERN Geneva in 2010.

[33] The ATLAS Collaboration. *Expected performance of the ATLAS experiment: detector, trigger and physics*. CERN, Geneva, 2009.

[34] The ATLAS Collaboration. Performance of primary vertex reconstruction in proton-proton collisions at $\sqrt{s} = 7$ TeV in the ATLAS experiment. *ATLAS-CONF-2010-069*, 2010.

[35] The ATLAS Collaboration. The ATLAS Simulation Infrastructure. *The European Physical Journal C*, 70(3):823–874, Sep 2010.

[36] The ATLAS Collaboration. The simulation principle and performance of the ATLAS fast calorimeter simulation FastCaloSim. Technical Report ATL-PHYS-PUB-2010-013, CERN, Geneva, Oct 2010.

[37] The ATLAS Collaboration. Particle Identification Performance of the ATLAS Transition Radiation Tracker. Technical Report ATLAS-CONF-2011-128, CERN, Geneva, Sep 2011.

[38] The ATLAS Collaboration. Performance of the Reconstruction and Identification of Hadronic Tau Decays with ATLAS. Technical Report ATLAS-CONF-2011-152, CERN, Geneva, Nov 2011.

[39] The ATLAS Collaboration. Jet energy measurement with the ATLAS detector in proton-proton collisions at $\sqrt{\mathrm{s}} = 7$ TeV. *The European Physical Journal C*, 73(3), Mar 2013.

[40] The ATLAS Collaboration. A neural network clustering algorithm for the ATLAS silicon pixel detector. *Journal of Instrumentation*, 9(09):P09009–P09009, Sep 2014.

[41] The ATLAS Collaboration. Electron and photon energy calibration with the ATLAS detector using LHC Run 1 data. *The European Physical Journal C*, 74(10), Oct 2014.

[42] The ATLAS Collaboration. Modelling $Z \to \tau\tau$ processes in ATLAS with τ-embedded $Z \to \mu\mu$ data. *Journal of Instrumentation*, 10(09):P09018–P09018, Sep 2015.

[43] The ATLAS Collaboration. Study of the spin and parity of the Higgs boson in diboson decays with the ATLAS detector. *Eur. Phys. J. C*, 75(CERN-PH-EP-2015-114):476. 36 p, Jun 2015.

[44] The ATLAS Collaboration. Electron and photon energy calibration with the ATLAS detector using data collected in 2015 at $\sqrt{s} = 13$ TeV. Technical Report ATL-PHYS-PUB-2016-015, CERN, Geneva, Aug 2016.

[45] The ATLAS Collaboration. Measurement of the tau lepton reconstruction and identification performance in the ATLAS experiment using pp collisions at $\sqrt{s} = 13$ TeV. Technical Report ATL-COM-PHYS-2016-929, CERN, Geneva, Jul 2016.

[46] The ATLAS Collaboration. Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV. *The European Physical Journal C*, 76(5), May 2016.

[47] The ATLAS Collaboration. Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector. *The European Physical Journal C*, 76(11), Oct 2016.

[48] The ATLAS Collaboration. Reconstruction of hadronic decay products of tau leptons with the ATLAS experiment. *The European Physical Journal C*, 76(5), May 2016.

[49] The ATLAS Collaboration. Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Physical Review D*, 96(7), Oct 2017.

[50] The ATLAS Collaboration. Jet reconstruction and performance using particle flow with the ATLAS Detector. *The European Physical Journal C*, 77(7), Jul 2017.

[51] The ATLAS Collaboration. Measurement of the tau lepton reconstruction and identification performance in the ATLAS experiment using pp collisions at $\sqrt{s} = 13$ TeV. Technical Report ATLAS-CONF-2017-029, CERN, Geneva, May 2017.

[52] The ATLAS Collaboration. Optimisation and performance studies of the ATLAS b-tagging algorithms for the 2017-18 LHC run. Technical Report ATL-PHYS-PUB-2017-013, CERN, Geneva, Jul 2017.

[53] The ATLAS Collaboration. Performance of the ATLAS track reconstruction algorithms in dense environments in LHC Run 2. *The European Physical Journal C*, 77(10), Oct 2017.

[54] The ATLAS Collaboration. Search for the Standard Model Higgs boson produced in association with top quarks and decaying into a $b\bar{b}$ pair in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Rev. D*, 97(CERN-EP-2017-291. 7):072016. 44 p, Dec 2017.

[55] The ATLAS Collaboration. Technical Design Report for the ATLAS Inner Tracker Strip Detector. Technical Report CERN-LHCC-2017-005. ATLAS-TDR-025, CERN, Geneva, Apr 2017.

[56] The ATLAS Collaboration. Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1. *The European Physical Journal C*, 77(7), Jul 2017.

[57] The ATLAS Collaboration. Calibration of light-flavour b-jet mistagging rates using ATLAS proton-proton collision data at $\sqrt{s} = 13$ TeV. Technical report, CERN, Geneva, Apr 2018.

[58] The ATLAS Collaboration. Evidence for the associated production of the Higgs boson and a top quark pair with the ATLAS detector. *Physical Review D*, 97(7), Apr 2018.

[59] The ATLAS Collaboration. Measurement of the Higgs boson coupling properties in the $H \to ZZ \to 4\ell$ decay channel at $\sqrt{s} = 13$ TeVwith the ATLAS detector. *Journal of High Energy Physics*, 2018(3), Mar 2018.

[60] The ATLAS Collaboration. Measurements of Higgs boson properties in the diphoton decay channel with 36 fb^{-1} of pp collision data at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Rev. D*, 98(arXiv:1802.04146. 5):052005. 136 p, Jan 2018.

[61] The ATLAS Collaboration. Performance of missing transverse momentum reconstruction with the ATLAS detector using proton–proton collisions at $\sqrt{s} = 13$ TeV. *The European Physical Journal C*, 78(11), Nov 2018.

[62] The ATLAS Collaboration. Prospects for the measurement of the rare Higgs boson decay $H \to \mu\mu$ with 3000 fb^{-1} of pp collisions collected at $\sqrt{s} = 14$ TeV by the ATLAS experiment. Technical report, CERN, Geneva, May 2018.

[63] The ATLAS Collaboration. Analysis of $t\bar{t}H$ and $t\bar{t}W$ production in multilepton final states with the ATLAS detector. Technical Report ATLAS-CONF-2019-045, CERN, Geneva, Oct 2019.

[64] The ATLAS Collaboration. ATLAS data quality operations and performance for 2015-2018 data-taking. *JINST*, 15(arXiv:1911.04632. 04):P04003. 43 p, Nov 2019.

[65] The ATLAS Collaboration. Cross-section measurements of the Higgs boson decaying into a pair of τ-leptons in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Physical Review D*, 99(7), Apr 2019.

[66] The ATLAS Collaboration. Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data. *Journal of Instrumentation*, 14(12):P12006–P12006, Dec 2019.

[67] The 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. *The European Physical Journal C*, 79(8), Aug 2019.

[68] The ATLAS Collaboration. Expected performance of the ATLAS detector at the High-Luminosity LHC. Technical Report ATL-PHYS-PUB-2019-005, CERN, Geneva, Jan 2019.

[69] The ATLAS Collaboration. Identification of hadronic tau lepton decays using neural networks in the ATLAS experiment. Technical Report ATL-PHYS-PUB-2019-033, CERN, Geneva, Aug 2019.

[70] The ATLAS Collaboration. Jet energy scale and resolution measured in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. Technical Report ATL-COM-PHYS-2019-281, CERN, Geneva, Apr 2019.

[71] The ATLAS Collaboration. Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run 2 data collected in 2015 and 2016. *The European Physical Journal C*, 79(3), Mar 2019.

[72] The ATLAS Collaboration. Combined measurements of Higgs boson production and decay using up to 80 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment. *Physical Review D*, 101(1), Jan 2020.

[73] The ATLAS Collaboration. Differential cross-section measurements for the electroweak production of dijets in association with a Z boson in proton-proton collisions at ATLAS. 6 2020.

[74] The ATLAS Collaboration. Operation of the ATLAS trigger system in Run 2. *JINST*, 15(arXiv:2007.12539. 10):P10004. 60 p, Jul 2020.

[75] The 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.

[76] The CMS Collaboration. Observation of $t\bar{t}H$ production. *Phys. Rev. Lett. 120, 231801*. doi:10.1103/PhysRevLett.120.231801.

[77] The CMS Collaboration. The CMS experiment at the CERN LHC. The Compact Muon Solenoid experiment. *JINST*, 3:S08004. 361 p, 2008. Also published by CERN Geneva in 2010.

[78] The CMS Collaboration. Precise determination of the mass of the Higgs boson and tests of compatibility of its couplings with the standard model predictions using proton collisions at 7 and 8 TeV. *The European Physical Journal C*, 75(212), 2015.

[79] The CMS Collaboration. Measurement of Higgs boson production and decay to the $\tau\tau$ final state. Technical Report CMS-PAS-HIG-18-032, CERN, Geneva, 2019.

[80] The CMS Collaboration. Measurement of Higgs boson production in the decay channel with a pair of τ leptons. Technical Report CMS-PAS-HIG-19-010, CERN, Geneva, 2020.

[81] T. Cornelissen, M. Elsing, S. Fleischmann, W. Liebig, E. Moyse, and E. Salzburger. Concepts, design and implementation of the ATLAS New Tracking NEWT. *ATL-SOFT-PUB-2007-007*, 2007.

[82] Glen Cowan, Kyle Cranmer, Eilam Gross, and Ofer Vitells. Asymptotic formulae for likelihood-based tests of new physics. *The European Physical Journal C*, 71(2), Feb 2011.

[83] D. de Florian, C. Grojean, F. Maltoni, et al. LHC Higgs Cross Section WG Picture Gallery.

[84] Deng9578. Null and alternative distribution, Mar 2017. Wikimedia image.

[85] Abdelhak Djouadi. The Anatomy of electro-weak symmetry breaking. I: The Higgs boson in the standard model. *Phys.Rept.*, 457:1–216, 2008.

[86] John F. Donoghue, Eugene Golowich, and Barry R. Holstein. *Dynamics of the Standard Model*. Cambridge University Press, 2 edition, 1994.

[87] A. Elagin, P. Murat, A. Pranko, and A. Safonov. A new mass reconstruction technique for resonances decaying to $\tau\tau$. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 654(1):481–489, Oct 2011.

[88] John Ellis. Higgs Physics. (KCL-PH-TH-2013-49. KCL-PH-TH-2013-49. LCTS-2013-36. CERN-PH-TH-2013-315):117–168. 52 p, Dec 2013. 52 pages, 45 figures, Lectures presented at the ESHEP 2013 School of High-Energy Physics, to appear as part of the proceedings in a CERN Yellow Report.

[89] Markus Elsing. Tracking at the LHC. *Lectures given at the University of Freiburg*, 2016.

[90] F. Englert and R. Brout. Broken Symmetry and the Mass of Gauge Vector Mesons. *Phys. Rev. Lett.*, 13:321–322, 1964.

[91] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. Syst. Sci.*, 55, 1997.

[92] R. Fruhwirth. Application of kalman filtering to track and vertex finding. *Nucl. Instrum. Meth*, A262:444–450, 1987.

[93] R. Fruhwirth, W. Waltenberger, and P. Vanlaer. Adaptive vertex fitting. *J. Phys*, G34(N343), 2007.

[94] T Gleisberg, S Höche, F Krauss, M Schönherr, S Schumann, F Siegert, and J Winter. Event generation with SHERPA 1.1. *Journal of High Energy Physics*, 2009(02):007–007, Feb 2009.

[95] A. Graves. *Supervised Sequence Labelling with Recurrent Neural Networks.* Studies in Computational Intelligence 385, Springer, 2012.

[96] V.N. Gribov and L.N. Lipatov. Deep inelastic e p scattering in perturbation theory. *Sov. J. Nucl. Phys.*, 15:438–450, 1972.

[97] LHC Higgs Cross Section Working Group. *Handbook of LHC Higgs Cross Sections: 1. Inclusive Observables.* CERN Yellow Reports: Monographs. CERN, Geneva, 2011. The authors dedicate this Report to the memory of Nicola Cabibbo and Georges Charpak.

[98] LHC Higgs Cross Section Working Group. *Handbook of LHC Higgs Cross Sections: 3. Higgs Properties: Report of the LHC Higgs Cross Section Working Group.* CERN Yellow Reports: Monographs. Jul 2013.

[99] LHC Higgs Cross Section Working Group. *Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector.* 2016.

[100] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble. Global Conservation Laws and Massless Particles. *Phys. Rev. Lett.*, 13:585–587, 1964.

[101] Julie Haffner. The CERN accelerator complex. Complexe des accélérateurs du CERN. Oct 2013. General Photo.

[102] Peter W. Higgs. Broken Symmetries and the Masses of Gauge Bosons. *Phys. Rev. Lett.*, 13:508–509, 1964.

[103] S. Hochreiter and J. Schmidhuber. Long Short-Term Memory. *Neural Computation*, 9:1735, 1997.

[104] David Hohn. *Search for top-quark-pair Higgs prouction with tau leptons using the ATLAS detector at the LHC.* PhD thesis, Bonn, 2018.

[105] R. E. Kalman. A new approach to linear filtering and prediction problems. *Transactions of the ASME-Journal of Basic Engineering*, 82(Series D):35–45, 1960.

[106] J. Neyman and E. S. Pearson. On the Problem of the Most Efficient Tests of Statistical Hypotheses. *Philosophical Transactions of the Royal Society of London. Series A, Containing Papers of a Mathematical or Physical Character*, 231:289–337, 1933.

[107] (Particle Data Group) C. Patrignani et al. Review of Particle Physics. *Chin. Phys. C*, 40(10), 2016.

[108] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to Quantum Field Theory.* Westview Press, 1995.

[109] A. Salzburger, 2011.

[110] AC Team. The four main LHC experiments. Jun 1999.

[111] The ATLAS and CMS Collaborations. Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV. 2015. ATLAS-CONF-2015-044, CMS-PAS-HIG-15-002.

[112] The ATLAS and CMS Collaborations. Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV. *Journal of High Energy Physics*, 2016(8), Aug 2016.

[113] The ATLAS Collaboration. Alignment of the ATLAS Inner Detector and its Performance in 2012. *ATLAS-CONF-2014-047*, 2014.

[114] The ATLAS Collaboration. Constraints on the off-shell Higgs boson signal strength in the high-mass ZZ and WW final states with the ATLAS detector. *The European Physical Journal C*, 75(335), 2015.

[115] The ATLAS Collaboration. Search for the associated production of the Higgs boson with a top quark pair in multilepton final states with the ATLAS detector. *Phys. Lett. B*, 749:519, 2015.

[116] The ATLAS Collaboration. Measurements of the Higgs boson production and decay rates and coupling strengths using pp collision data at $\sqrt{s} = 7$ and 8 TeV in the ATLAS experiment. *Eur. Phys. J. C*, 76(1):6, 2016.

[117] The ATLAS Collaboration. Measurement of the cross-section for electroweak production of dijets in association with a Z boson in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Lett. B*, 775(CERN-EP-2017-115):206–228. 23 p, Sep 2017.

[118] The ATLAS Collaboration. Cross-section measurements of the Higgs boson decaying into a pair of τ-leptons in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Rev. D*, 99(arXiv:1811.08856):072001. 58 p, Nov 2018. Submitted to Phys.Rev.

[119] The CMS Collaboration. Observation of the Higgs boson decay to a pair of tau leptons. CMS-HIG-16-043, CERN-EP-2017-181.

[120] The CMS Collaboration. Search for the associated production of the Higgs boson with a top-quark pair. *JHEP*, 09:087, 2014.

[121] The CMS Collaboration. Search for Higgs boson production in association with top quarks in multilepton final states at $\sqrt{s} = 13$ TeV. CMS-PAS-HIG-17-004, 2017.

[122] Abraham Wald. On Cumulative Sums of Random Variables. *The Annals of Mathematical Statistics*, 15(3):283–296, Sep 1944.

[123] Richard Wigmans. *Calorimetry: energy measurement in particle physics; 2nd ed.* Oxford University Press, Oxford, 2017.

[124] S. S. Wilks. The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses. *The Annals of Mathematical Statistics*, 9(1):60–62, Mar 1938.

[125] P. F. Åkesson, T. Atkinson, M. J. Costa, M. Elsing, S. Fleischmann, A. N. Gaponenko, W. Liebig, E. Moyse, E. Salzburger, and M. Siebel. ATLAS Tracking event data model. *ATL-SOFT-PUB-2006-004*, 2006.