Measurements of the Higgs Boson in the $H \to \tau\tau$ Decay Channel

Jacob Howard
The Queen's College
University of Oxford

Submitted in partial fulfillment of
the requirements for the degree of
Doctor of Philosophy

Hilary 2015

Abstract

The generation of vector boson mass via the Higgs mechanism in the Standard Model has been confirmed by the 2012 discovery of a candidate Higgs boson in the $H \to WW$, $H \to ZZ$, and $H \to \gamma\gamma$ decay channels. In contrast, the Yukawa couplings hypothesized to provide the mass of fermions in the Standard Model have yet to be observed. The $H \to \tau\tau$ decay channel currently provides the best opportunity for observing these couplings.

This thesis describes two separate but related searches for Higgs boson decays in the $H \to \tau\tau$ decay channel using proton-proton collisions recorded by the ATLAS detector. The first analysis is a general search for all Higgs boson production mechanisms leading to a $H \to \tau\tau$ decay using 4.5 fb^{-1} of 7 TeV and 20.3 fb^{-1} of 8 TeV proton-proton collision data. A deviation from the background-only hypothesis is observed with a significance of 4.5σ for a hypothetical Higgs boson mass of $m_H = 125$ GeV — a strong indication of a $H \to \tau\tau$ signal. For the same mass point, the best fit value for the signal strength is found to be $1.43^{+0.43}_{-0.37}$ × the Standard Model expectation. The second analysis is a search for Higgs boson production in association with a vector boson using 20.3 fb^{-1} of 8 TeV proton-proton collision data. Results in the $Z_{ll}H \to \tau_l\tau_h$ channel indicate limits of 9.14 × the Standard Model expectation for VH signal production at $m_H = 125$ GeV. In addition, two studies on enhancement of computing performance in the ATLAS trigger and data analysis pipeline are presented.

To my parents, my grandparents, and my wife.

Acknowledgements

In the course of my D.Phil., I've had the pleasure of working with many talented individuals, without whom I would have been completely lost. I've also been lucky enough to call many of them my friends.

I would first like to thank my advisors, Chris Hays and Sinead Farrington, for their immense help in setting the course of my studies, as well as their continued guidance.

For their help with work on GPUs, I would like to thank Dmitry Emeliyanov, John Baines, Tim Bristow, Ben Wynne, and Andy Washbrook.

My work in ATLAS has also benefited immeasurably from the advice and assistance of many individuals involved with the $H \to \tau\tau$ analyses, but I would particularly like to thank Dimitris Varouchas, Carl Jeske, and Henrik Öhman for their camaraderie and persistent support.

Tim, Dave, Chase, and Johannes — I would never have made it through this process without your constant friendship, empathy, and immense help. I won't forget the awesome times we've had together, and I know there will only be more of them in the future.

Last, but most importantly, I must thank my amazing partner and best friend. Sashka — you were the least expected part of the last four years, but I don't know how I ever thought I could have done this without you. Я кохаю тебе — дуже дуже!

If we fail to anticipate the unforeseen or expect the unexpected in a universe of infinite possibilities, we may find ourselves at the mercy of anyone or anything that cannot be programmed, categorized, or easily referenced.

FOX MULDER

Contents

List of Figures

List of Tables

List of Algorithms

Chapter 1

Introduction

The Standard Model of particle physics is the culmination of over a century's worth of development, owing its success to both profound theoretical insights and daunting experimental research. With the discovery of a candidate Higgs boson by experiments at the Large Hadron Collider (LHC) in 2012 [1, 2], a crucial and long-sought component of this model may finally be within experimental grasp. It is now of paramount importance to understand the properties of this newly discovered particle, in particular whether it adheres to the theoretical predictions of the Standard Model Higgs boson.

Initial measurements of this particle, in particular its mass [3], spin [4], and couplings [5], indicate a strong consistency with the predicted properties of the Higgs boson. In addition to the measurements being made in the original discovery channels, it is also important that measurements of this particle be made in other decay channels, particularly those other than the electroweak bosons. The $H \to \tau\tau$ and $H \to b\bar{b}$ decay channels will provide the best opportunity for measurements of the particle's fermionic decays in the lepton and quark sectors, respectively.

The upgrades currently underway at the LHC will see increases in both the instantaneous luminosity and center-of-mass energy of the proton beams used to generate collisions. These upgrades will provide orders of magnitude more data, as well as access to higher energy frontiers, allowing more precise measurement of fermionic Higgs decay modes in the future.

This thesis presents the results of two separate but related $H \to \tau\tau$ searches at the ATLAS experiment at the LHC. The first is a general search for all Higgs boson production modes leading to a $H \to \tau\tau$ decay. This search is divided into three analyses, each focusing on one of the possible decay modes of the $\tau\tau$ system: $H \to \tau_h\tau_h$, $H \to \tau_l\tau_h$, and $H \to \tau_l\tau_l$.[1] These three analyses are collectively referred to herein as the "main analyses." The second search focuses on Higgs boson production in the VH associated production mode. This search is divided into four analyses, each focusing on a particular combination of Higgs production mode and $\tau\tau$ decay mode: $W_{l\nu}H \to \tau_h\tau_h$, $W_{l\nu}H \to \tau_l\tau_h$, $Z_{ll}H \to \tau_h\tau_h$, and $Z_{ll}H \to \tau_l\tau_h$.[2] These four analyses are collectively referred to herein as the "VH analyses."

This thesis focuses on the $H \to \tau_l\tau_h$ and $Z_{ll}H \to \tau_l\tau_h$ analyses, though some information common to all $H \to \tau\tau$ analyses is also presented. Chapter 2 provides an overview of the Standard Model with a focus on the theoretical model of electroweak unification. Chapter 3 gives an account of the LHC and ATLAS experimental apparatuses, with Chapter 4 describing the techniques used to select and reconstruct events recorded by the ATLAS detector. In Chapter 5, summaries of two studies to enhance the computational performance of the ATLAS trigger and data analysis pipeline are given. The main and VH analyses are detailed in Chapters 6-8, with results given in Chapter 9, and brief conclusions given in Chapter 10.

---

[1] Here τ_l denotes a leptonically decaying τ lepton, while τ_h denotes a hadronically decaying τ lepton.

[2] The fully leptonic VH decay channels, $W_{l\nu}H \to \tau_l\tau_l$ and $Z_{ll}H \to \tau_l\tau_l$, are not considered due to extremely limited statistics.

Chapter 2

Theoretical motivation

Over the past century, the Standard Model of particle physics has emerged as one of the most successful theories in the history of science. It has provided predictions of physical quantities to unprecedented accuracy, and is the most rigorously tested and successful physical theory [6].

The Standard Model is founded on the theoretical framework of quantum field theory. It consists of explicit constructions of three out of the four known fundamental forces within this framework. The Standard Model also catalogs an assortment of particles that constitute known matter and gauge bosons that mediate interactions.

Despite its overwhelming success, the Standard Model remains an incomplete and merely effective theory. It also has yet to be validated in its entirety, most notably missing full confirmation of the Higgs mechanism of electroweak symmetry breaking.

This chapter provides an outline of the Standard Model, in particular its particle and force content, as well as its limitations. Particular focus is given to the historical motivation and theoretical development of electroweak symmetry breaking via the Higgs mechanism. In addition, potential Higgs boson search channels at the LHC are enumerated, and the significance of the $H \to \tau\tau$ decay channel is discussed.

2.1 The Standard Model of particle physics

The Standard Model is formulated as a relativistic, locally gauge-invariant quantum field theory. Quantum field theories describe particles as quantized fluctuations of a field and interactions as couplings between fields. Local gauge invariance mandates that the dynamics of these fields are invariant under gauge transformations that may be functions of x^μ.

The Standard Model is locally gauge-invariant (symmetric) under transformations defined by the group SU(3) × SU(2) × U(1). The components of this internal symmetry correspond to the strong [SU(3)] and electroweak [$\mathrm{SU}(2)_L \times \mathrm{U}(1)$] forces. These symmetries, via Noether's theorem [7], give rise to conserved currents of strong (color) charge, weak isospin, and weak hypercharge — the dynamics of which describe particle interactions. The enforcement of *local* gauge invariance gives rise to massless vector bosons, which act as the quanta of these currents and mediators of the corresponding interactions. Due to the process of electroweak symmetry breaking, ostensibly via the Higgs mechanism, the $W^\pm$ and Z vector bosons acquire mass, while their cousin, the photon γ, remains massless.

In the Standard Model Lagrangian, fermionic matter fields are described by Dirac spinors and bosonic gauge fields are described by vector fields. The Higgs boson has the distinction of being the only fundamental scalar particle. Masses are encoded as coefficients of quadratic self-interaction terms. Each of the Standard Model forces has a different characteristic strength, encoded in the Lagrangian by coefficients known as "coupling constants."

Naïve treatment of higher-order loop diagrams, such as that in Figure 2.1, leads to logarithmically divergent terms in quantum field theory calculations of gauge boson interactions. A procedure known as "renormalization" was developed to recast these infinities as energy-dependent cancellations that occur as corrections to the "bare" masses and coupling constants [8], resulting in the "physical" masses and coupling

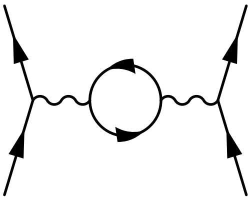

Figure 2.1: A higher-order QED loop diagram.

constants which are actually measured. The physical (measurable) values of coupling constants vary with the energy transfer of a given process. As a result, they are referred to as "running couplings."

The procedure of renormalization was originally developed in the context of quantum electrodynamics, but in 1971 't Hooft showed [9, 10] that non-abelian Yang-Mills gauge theories [11], which form the foundation of electroweak theory and quantum chromodynamics, are also renormalizable.

2.1.1 Particle content

The particle content of the Standard Model is outlined in Table 2.1. The model contains six types of leptons across three generations, six types of quarks across three generations, four types of mediator particles, and one Higgs particle. For each quark or lepton there exists a corresponding anti-particle with the same mass but all internal quantum numbers reversed.

The lepton sector is made up of three charged particles, the e, μ, and τ, and a corresponding neutrino for each. The charged particles of the lepton sector participate in both weak and electromagnetic interactions, while the neutrinos participate only in weak interactions. No leptons carry color charge, and consequently they are excluded from strong interactions. Each lepton (anti-lepton) is assigned a "lepton number" of $+1$ (-1), a quantity which is conserved in weak and electromagnetic interactions.

In the Standard Model, neutrinos are treated as massless particles, and thus their flavor should be conserved. However, experiments have conclusively shown [13] that

Type	Name	Symbol	Generation	Mass	Spin	Charge
Lepton	Electron	e	1	0.511 MeV	$1/2$	-1
	e-neutrino	ν_e	1	<225 eV	$1/2$	0
	Muon	μ	2	105.658 MeV	$1/2$	-1
	μ-neutrino	ν_μ	2	<0.19 MeV	$1/2$	0
	Tau	τ	3	1776.82 MeV	$1/2$	-1
	τ-neutrino	ν_τ	3	<18.2 MeV	$1/2$	0
Quark	Up	u	1	2.3 MeV	$1/2$	$+2/3$
	Down	d	1	4.8 MeV	$1/2$	$-1/3$
	Charm	c	2	1.275 GeV	$1/2$	$+2/3$
	Strange	s	2	95 MeV	$1/2$	$-1/3$
	Top	t	3	173.5 GeV	$1/2$	$+2/3$
	Bottom	b	3	4.18 GeV	$1/2$	$-1/3$
Gauge boson	Photon	γ	-	0	1	0
	Gluon	g	-	0	1	0
	W	$W^\pm$	-	80.385 GeV	1	±1
	Z	Z	-	91.1876 GeV	1	0
Higgs boson	Higgs	H	-	125.7 GeV	0	0

Table 2.1: Particle content of the Standard Model [12].

neutrinos do have mass, albeit small, by observing oscillation between different neutrino flavors. Consequently, lepton flavor violation does occur, and is introduced in the Standard Model in the form of the PMNS matrix. The existence of *global* lepton number conservation depends on the nature of neutrinos, and violation of this conservation law will occur if neutrinos are their own anti-particles (so-called "Majorana fermions").

The quark sector is made up of six fractionally charged quarks: the u, d, c, s, t, and b. In addition to electromagnetic and weak isospin charge, the quarks also carry color charge, and consequently participate in strong, weak, and electromagnetic interactions. Quark bound states (hadrons and mesons) are assigned a "baryon number," determined by their quark/anti-quark content, $B = \frac{1}{3}(n_q - n_{\bar{q}})$. Baryon number is conserved in all interactions.

The gauge bosons act as mediators for the Standard Model interactions — the photon for the electromagnetic force, the $W^\pm$ and Z bosons for the weak force, and the gluon for the strong force. In addition to couplings with the lepton and quark matter fields, the non-abelian gauge symmetries of the Standard Model lead to self-couplings of the gluon field, as well as various coupling combinations of the photon, $W^{\pm}$, and Z bosons.

The Higgs boson is the quantum of the Higgs field, which provides the mechanism for electroweak symmetry breaking and mass generation in the Standard Model. It couples directly to all massive particles in the Standard Model, including itself, with coupling parameters proportional to the mass of the respective particle.

2.1.2 Quantum electrodynamics

Quantum electrodynamics (QED) is the quantum field theory formulation of electromagnetism. It is the oldest and best understood of all Standard Model theories [8], and to date it has provided the most accurate predictions of any physical theory [6].

QED extends Maxwell's classical (yet already Lorentz-invariant) formulation of electromagnetism [14] to the quantum regime, where charge and current are quantized, and the vector potential A^{μ} is reinterpreted as a mediator field. The QED Lagrangian density is given by

$$\mathcal{L}_{\mathrm{QED}} = \overline{\psi}(i\gamma^{\mu}\partial_{\mu} - m_{\psi})\psi - g_e\overline{\psi}\gamma^{\mu}A_{\mu}\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu}, \tag{2.1}$$

where

$$F_{\mu\nu} = \partial_{\mu}A_{\nu} - \partial_{\nu}A_{\mu}. \tag{2.2}$$

Here ψ is the Dirac spinor of a charged particle, m_{ψ} is the corresponding particle mass, A_{μ} is the photon vector field, and $F_{\mu\nu}$ is the field tensor, describing the dynamics of a free massless vector field. The first term is simply the Dirac equation for a massive free particle, the second term describes the QED interaction vertex, and the third term describes a free massless vector field.

QED is invariant under the internal symmetry $U(1)$, which corresponds to a gauge transformation of the vector potential:

$$A_\mu \rightarrow A_\mu + \partial_\mu \lambda, \tag{2.3}$$

with a simultaneous transformation of the Dirac field:

$$\psi \rightarrow e^{-ig_e \lambda} \psi, \tag{2.4}$$

where λ is some function of x^μ.

The quantum of the mediator field is the photon, γ. The photon couples to any particle with electric charge, including the charged leptons, quarks, and $W^\pm$ bosons. The primary interaction vertex is shown in Figure 2.2. In this vertex, a charged particle e enters, emits or absorbs a photon, and exits. The type and flavor of the particle remain unchanged. The coupling constant for this vertex, g_e, is proportional to the electric charge of the particle coupling to the photon.

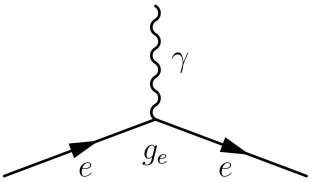

Figure 2.2: Primary QED interaction vertex.

2.1.3 GWS electroweak theory and the Higgs mechanism

In the context of the Standard Model, the electromagnetic force is unified with the weak force into a combined electroweak theory. This unification allows the weak and electromagnetic forces to be seen as arising from the same underlying symmetry.

2.1.3.1 The weak force

The weak force was originally proposed to explain β-decay in radioactive nuclei. In the decay, assumed at that time to be

$$n \rightarrow p^+ + e^-, \tag{2.5}$$

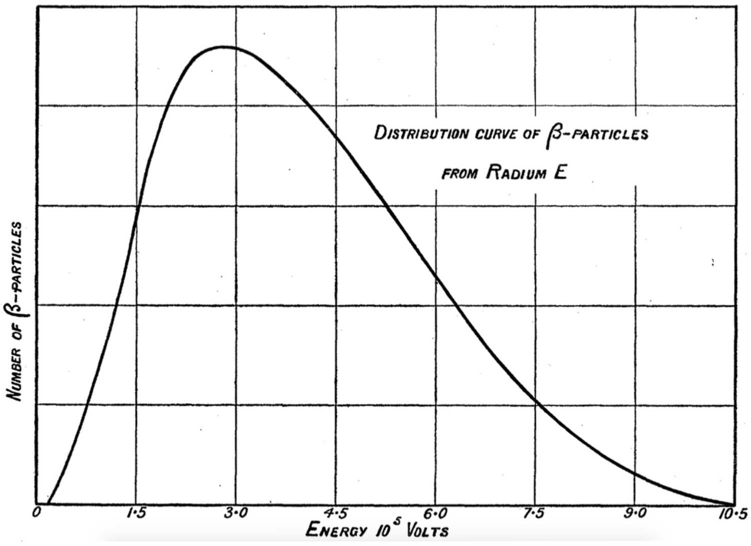

Figure 2.3: The energy spectrum of β-particles (electrons) in Radium E (the historical name for ^{210}Bi) β-decay [15].

it was observed [15] that the energy of the emitted electron was not constant. Instead, a continuous spectrum of energies was observed, as illustrated in Figure 2.3. This was problematic because the energy of the electron in a two-body decay is kinematically determined to be a fixed value:

$$E_e = \frac{m_P^2 - m_D^2 + m_e^2}{2m_P}. \tag{2.6}$$

Here m_P is the mass of the parent nucleus, m_D is the mass of the daughter nucleus, and m_e is the mass of the electron.

Pauli proposed [16] that the spectrum could be explained by the emission of a second, electrically neutral particle (the neutrino) [1]:

$$n \to p^+ + e^- + \overline{\nu_e}. \tag{2.7}$$

Pauli's theory was given a quantitative formulation by Fermi [17], who treated the β-decay interaction as an effective four-fermion vertex with coupling constant G_F, the modern value of which is given by $G_F/(\hbar c)^3 = 1.166 \times 10^{-5}$ GeV^{-2}. Fermi's theory was the first effective formulation of the weak force.

---

[1] At the time, of course, it was not known that multiple neutrino types existed, much less that an anti-neutrino would be necessary for conservation of lepton number.

Fermi originally assumed that the interaction was parity-conserving, but this assumption was challenged by Lee and Yang [18] in their effort to explain the so-called τ-θ puzzle (in which the K^+ demonstrated parity-violating decay modes). An experiment carried out by Wu [19] using spin-aligned ^{60}Co atoms indicated a preferential direction for the emission of electrons, thereby confirming parity violation in the weak sector.

Full elucidation of the parity-violating structure of the weak interaction did not come until 1958, when Feynman and Gell-Mann published a paper [20] exploring the potential forms of the parity-violating weak currents. They began by considering the most general form of the Lagrangian for the four-fermion β-decay interaction:

$$\sum_i C_i(\overline{\psi}_n O_i \psi_p)(\overline{\psi}_\nu O_i \psi_e), \tag{2.8}$$

where the operators O_i (1, γ^5, γ^μ, $\gamma^\mu\gamma^5$, and $\sigma^{\mu\nu} = (i/2)(\gamma^\mu\gamma^\nu - \gamma^\nu\gamma^\mu)$) combine with the wave functions to form the five possible bilinear covariants, each with distinct parity behavior (scalar S, pseudoscalar P, vector V, axial vector A, and antisymmetric tensor T, respectively). The parity violation observed by Wu indicated the existence of some axial component, although the experimental evidence was wildly conflicting at the time. Feynman and Gell-Mann, presuming that the interaction involved in β-decay was the same as that involved in muon decay, took muon decay measurements as strong evidence for a $V - A$ structure of the parity-violating component of the weak force. This conclusion was confirmed by subsequent repetition of the conflicting β-decay experiments.

2.1.3.2 Early electroweak unification

Unification of the weak and electromagnetic interactions was first proposed by Schwinger in 1957 [21], who suggested that the weak and electromagnetic interactions might both be generated by a triplet of vector fields. His work was followed by that of Glashow, who proposed the modern $\mathrm{SU}(2)_L \times \mathrm{U}(1)$ construction of four vector mediator fields:

two massive charged mediators, one massive neutral mediator, and one massless neutral mediator. At the time, no mechanism was known for generating the masses of the three massive vector mediator fields, because introduction of a Proca mass term of the form

$$m^2 A^\mu A_\mu \tag{2.9}$$

destroys gauge invariance. Consequently, Glashow's formulation was limited to the idea of "partial symmetries," where only a portion of the Lagrangian is invariant under gauge transformations. These partial symmetries create corresponding "partial conservation laws," which hold only when masses are neglected. Using these conservation laws, Glashow was able to generate (through mixing of the underlying weak isotriplet and weak isosinglet) the parity-violating structure of the weak interaction, while still maintaining the parity-preserving vector structure of electromagnetic interactions.

2.1.3.3 Goldstone's theorem

Before a complete formulation of electroweak unification could be constructed, a mechanism for generating the masses of the vector bosons was necessary. Initial attempts centered on the idea of spontaneous symmetry breaking, where a non-zero vacuum has a different symmetry than that of the Lagrangian. For example, consider a complex scalar field

$$\phi = \frac{\phi_1 + i\phi_2}{\sqrt{2}} \tag{2.10}$$

and a Lagrangian density

$$\mathcal{L} = (\partial^\mu \phi^*)(\partial_\mu \phi) + \mu^2 \phi^* \phi - \lambda^2 (\phi^* \phi)^2, \tag{2.11}$$

which is invariant under transformations of the form

$$\phi \to e^{i\alpha} \phi. \tag{2.12}$$

This yields a potential of the form

$$\mathcal{U}(\phi) = -\mu^2 \phi^* \phi + \lambda^2 (\phi^* \phi)^2. \tag{2.13}$$

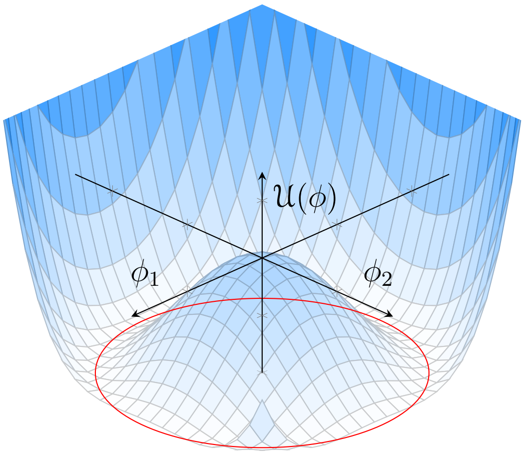

Figure 2.4: A hypothetical Higgs potential, with continuous minima indicated by the red circle.

Such a potential, often referred to as the "Mexican hat" potential, is shown in Figure 2.4. In this potential, there is a continuum of minima along the circle defined by

$$\phi_1^2 + \phi_2^2 = \frac{\mu^2}{\lambda^2}. \tag{2.14}$$

Because the Feynman calculus is a perturbative procedure around a ground state, a ground state must be chosen around which to expand. This choice leads to no loss of generality in the procedure, but *any* choice of ground state that does not share the symmetry of the Lagrangian will "spontaneously break" the true symmetry of the theory.

As an example, one might choose to expand about

$$\phi_1 = \frac{\mu}{\lambda} \qquad\qquad \phi_2 = 0. \tag{2.15}$$

New fields may be chosen representing fluctuations about this ground state:

$$\eta = \phi_1 - \frac{\mu}{\lambda} \qquad\qquad \xi = \phi_2. \tag{2.16}$$

Using these fields, the complex field may be rewritten as

$$\phi = \frac{1}{\sqrt{2}}\left[\frac{\mu}{\lambda} + \eta + i\xi\right], \tag{2.17}$$

which can be substituted into the Lagrangian, yielding

$$\mathcal{L} = \left[\frac{1}{2}(\partial^\mu\eta)(\partial_\mu\eta) - \mu^2\eta^2\right] + \left[\frac{1}{2}(\partial^\mu\xi)(\partial_\mu\xi)\right] + \text{couplings} + \text{constant}. \tag{2.18}$$

The spontaneous breaking of the original symmetry results in a massive real scalar field, η, with mass

$$m_\eta = \sqrt{2}\mu. \tag{2.19}$$

Unfortunately, this procedure also results in the generation of a massless scalar field, ξ. This is problematic in the context of electroweak unification because no such massless scalars have been observed. The appearance of a massless scalar in this procedure is, in fact, guaranteed by Goldstone's theorem [22], which states that one or more massless scalar bosons (termed "Goldstone bosons") will appear whenever a continuous global symmetry is broken.

2.1.3.4 The Higgs mechanism

In the early 1960s, it occurred to Higgs [23, 24] and a number of others [25, 26] that the procedure of spontaneous symmetry breaking might work under the enforcement of *local* gauge invariance. For example, the Lagrangian in Equation 2.11 might be made invariant under transformations of the form

$$\phi \to e^{i\alpha(x)}\phi. \tag{2.20}$$

This can be accomplished by the introduction of a massless vector gauge field, A_μ, which transforms as

$$A_\mu \to A_\mu + \frac{1}{g}\partial_\mu\alpha(x), \tag{2.21}$$

and the substitution of the covariant derivative,

$$\mathcal{D}_\mu = \partial_\mu - igA_\mu, \tag{2.22}$$

into the Lagrangian ($\partial_\mu \to \mathcal{D}_\mu$) to yield

$$\mathcal{L} = [(\partial^\mu + igA^\mu)\phi^*]\,[(\partial_\mu - igA_\mu)\phi] + \mu^2\phi^*\phi - \lambda^2(\phi^*\phi)^2 - \frac{1}{4}F^{\mu\nu}F_{\mu\nu}. \tag{2.23}$$

Figure 2.5: An incorrect "conversion" vertex resulting from a term which is bilinear in different fields [8].

Here g is a coupling constant, and the last term of the Lagrangian is the field tensor necessary to describe the dynamics of a free vector field.

As this Lagrangian shares the same ground states as that in Equation 2.11, it can again be rewritten in terms of the fields η and ξ, and becomes

$$\begin{aligned}\mathcal{L} = &\left[\frac{1}{2}(\partial^\mu \eta)(\partial_\mu \eta) - \mu^2 \eta^2\right] + \left[\frac{1}{2}(\partial^\mu \xi)(\partial_\mu \xi)\right] \\ &+ \left[-\frac{1}{4}F^{\mu\nu}F_{\mu\nu} + \frac{1}{2}\left(g\frac{\mu}{\lambda}\right)^2 A^\mu A_\mu\right] \\ &- i\left(g\frac{\mu}{\lambda}\right)(\partial^\mu \xi)A_\mu \\ &+ \text{couplings} + \text{constant}.\end{aligned} \tag{2.24}$$

Importantly, the gauge field A_μ has acquired a mass:

$$m_A = g\frac{\mu}{\lambda}. \tag{2.25}$$

Once again, however, the Goldstone boson ξ remains. Moreover, it appears to couple to A_μ via the term

$$-i\left(g\frac{\mu}{\lambda}\right)(\partial^\mu \xi)A_\mu. \tag{2.26}$$

This corresponds to a "conversion" vertex of the form shown in Figure 2.5. This is an indication that the fields in Equation 2.24 do not correspond to true physical particles [8].

The solution to this problem exploits the local gauge invariance of the Lagrangian. Equation 2.17 can, to first order, be equivalently rewritten in terms of new fields h and θ as

$$\phi = \frac{1}{\sqrt{2}}\left(\frac{\mu}{\lambda} + h(x)\right)e^{i\frac{\lambda}{\mu}\theta(x)}. \tag{2.27}$$

By choice of a particular gauge,

$$\alpha(x) = -\frac{\lambda}{\mu}\theta(x), \tag{2.28}$$

the azimuthal mode θ can be eliminated:

$$\phi \to \frac{1}{\sqrt{2}}\left(\frac{\mu}{\lambda} + h(x)\right). \tag{2.29}$$

This component is absorbed by the gauge field:

$$A_\mu \to A_\mu - g\frac{\lambda}{\mu}\partial_\mu\theta(x). \tag{2.30}$$

Thus, in this gauge, the Lagrangian becomes

$$\begin{aligned} \mathcal{L} = &\left[\frac{1}{2}(\partial^\mu h)(\partial_\mu h) - \mu^2 h^2\right] \\ &+ \left[-\frac{1}{4}F^{\mu\nu}F_{\mu\nu} + \frac{1}{2}\left(g\frac{\mu}{\lambda}\right)^2 A^\mu A_\mu\right] \\ &+ \text{couplings} + \text{constant}. \end{aligned} \tag{2.31}$$

The massless Goldstone boson has disappeared from the Lagrangian, absorbed into the definition of A_μ. This creates a third degree of freedom for A_μ, longitudinal polarization, as would be expected for a *massive* gauge boson. A single physical massive scalar remains, the only remnant of ϕ, and is known as the “Higgs field,” the quantum of which is known as the “Higgs boson.” Generation of gauge field mass via the reformulation of a Lagrangian about a non-zero ground state, combined with the removal of Goldstone bosons via a gauge transformation, is a technique known as the “Higgs mechanism.”

2.1.3.5 GWS electroweak theory

Using the Higgs mechanism, Weinberg [27] and Salam [28] expanded on the work of Glashow to form the modern Glashow-Weinberg-Salam (GWS) construction of electroweak unification.

Unlike the example of Higgs and Goldstone given in the previous sections, which broke a U(1) symmetry, GWS theory starts with a system that is invariant under both SU(2) transformations in weak isospin space as well as U(1) transformations in weak hypercharge space.

The formulation begins by considering a set of three weak isospin currents, $\boldsymbol{j}_\mu$, coupling to a weak isotriplet of massless vector bosons, $\boldsymbol{W}^\mu$, along with a weak hypercharge current, j_μ^Y, coupling to a massless isosinglet vector boson, B^μ:

$$-i(g\boldsymbol{j}_\mu \cdot \boldsymbol{W}^\mu + \frac{g'}{2} j_\mu^Y B^\mu). \tag{2.32}$$

The fields, as well as the $V - A$ structure of the weak force, are encoded in the electroweak Lagrangian by

$$\begin{aligned} \mathcal{L}_{\text{EW}} = & -\frac{1}{4}\boldsymbol{W}_{\mu\nu}\boldsymbol{W}^{\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu} \\ & + \overline{f_L}\gamma^\mu(i\partial_\mu - g\frac{1}{2}\boldsymbol{\sigma} \cdot \boldsymbol{W}_\mu - g'\frac{Y}{2}B_\mu)f_L \\ & + \overline{f_R}\gamma^\mu(i\partial_\mu - g'\frac{Y}{2}B_\mu)f_R, \end{aligned} \tag{2.33}$$

where $\boldsymbol{W}_{\mu\nu}$ and $B_{\mu\nu}$ are the field tensors for the free vector fields, f_L represents a left-handed lepton or quark fermionic doublet, f_R represent a right-handed fermionic singlet, $\boldsymbol{\sigma}$ are the Pauli matrices (the three-dimensional representation of SU(2)), and g and g' are electroweak coupling parameters.

In order to generate masses, as well as to cause these underlying fields to mix, four real scalar fields are introduced in a complex isospin doublet:

$$\phi = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix}. \tag{2.34}$$

This doublet is given a non-zero vacuum expectation value and coupled to the $\boldsymbol{W}^\mu$ and B^μ fields through the addition of a term in the Lagrangian:

$$\mathcal{L}_{\text{Higgs}} = \left|(i\partial_\mu - g\frac{1}{2}\boldsymbol{\sigma} \cdot \boldsymbol{W}_\mu - g'\frac{Y}{2}B_\mu)\phi\right|^2 - \mathcal{U}(\phi), \tag{2.35}$$

where $\mathcal{U}(\phi)$ is again the Mexican hat potential given in Equation 2.13. As before, the Lagrangian must be rewritten in terms of perturbations about some ground state of this potential, and a gauge is chosen that again removes all but one degree of freedom from ϕ:

$$\phi \to \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix}. \tag{2.36}$$

Here the vacuum expectation value, $v = \frac{\mu}{\lambda}$, is defined by the parameters of the potential. Although any choice of ground state will generate masses for the vector bosons, a choice of ground state that leaves the Lagrangian invariant under some subset of symmetries will leave the gauge bosons associated with those symmetries massless [29]. Accordingly, the choice can be engineered to generate mass in the weak sector while leaving the photon massless.

Expansion of Equation 2.35 about this minimum yields

$$\begin{aligned} \mathcal{L}_{\text{Higgs}} =& \frac{1}{2}(\partial^\mu h)(\partial_\mu h) - \mu^2 h^2 \\ &+ \frac{1}{4} v^2 g^2 W^{\mu+} W^-_\mu \\ &+ \frac{1}{8} v^2 (g^2 + g'^2) Z^\mu Z_\mu \\ &+ 0 A^\mu A_\mu \\ &+ \text{couplings} + \text{constant}, \end{aligned} \tag{2.37}$$

where

$$W^\pm_\mu = \frac{W^1_\mu \mp i W^2_\mu}{\sqrt{2}} \qquad Z_\mu = \frac{g W^3_\mu - g' B_\mu}{\sqrt{g^2 + g'^2}} \qquad A_\mu = \frac{g' W^3_\mu + g B_\mu}{\sqrt{g^2 + g'^2}} \tag{2.38}$$

represent the physical $W^\pm$, Z, and photon fields, which are a mixture of the underlying fields. The $W^\pm$ and Z bosons have both acquired a mass while the photon remains conspicuously massless. Of the four scalar fields originally introduced to create the Higgs potential, three have been consumed to generate the longitudinal polarization modes of the new massive particles, and only a single massive scalar Higgs field remains.

One major observable consequence of this theory is the ratio of the $W^\pm$ to Z boson masses. This ratio, derived from Equation 2.37, is given by

$$\frac{m_W}{m_Z} = \frac{g}{\sqrt{g^2 + g'^2}} = \cos\theta_W. \tag{2.39}$$

The parameter θ_W is known as the "Weinberg angle," and its measurement via neutral current processes and Z pole observables gives precise agreement with the independently measured ratio of the $W^{\pm}$ and Z boson masses [12].

The same Higgs mechanism can be also used to generate the masses of the fermions. This is achieved by adding a Yukawa coupling between the fermion fields and Higgs field to the electroweak Lagrangian:

$$\mathcal{L}_{\text{Higgs-Yukawa}} = -G_f \left[\overline{f_L} \phi f_R + \overline{f_R} \phi^{\dagger} f_L \right], \tag{2.40}$$

where G_f is a coupling parameter that depends on the fermion type. When rewritten about the ground state in Equation 2.36, this becomes

$$\mathcal{L}_{\text{Higgs-Yukawa}} = -\frac{G_f}{\sqrt{2}} v(\overline{f_L} f_R + \overline{f_R} f_L) - \frac{G_f}{\sqrt{2}} (\overline{f_L} f_R + \overline{f_R} f_L) h. \tag{2.41}$$

The first term of Equation 2.41 is a Dirac mass term, with

$$m_f = \frac{G_f v}{\sqrt{2}}, \tag{2.42}$$

while the second term represents a coupling with the Higgs field, which, due to the presence of G_f in Equation 2.42, should be proportional to the mass of the fermion in question.

While the GWS theory has proved tremendously successful in its predictions of the $W^{\pm}$ and Z boson masses, as well as their ratio, it provides no mechanism for predicting the Higgs boson mass. This problem stems from the unknown parameters of the Higgs potential. At present, experimental searches for the Higgs boson offer the only avenue for determination of its mass. Nevertheless, other electroweak observables can be used to constrain the most likely value for the Higgs boson mass [30]. These fits, along with pre-discovery experimental exclusion regions, are shown in Figure 2.6.

2.1.4 Quantum chromodynamics

Quantum chromodynamics (QCD) is the quantum field theory formulation of the strong nuclear force. QCD is invariant under the internal symmetry SU(3), which

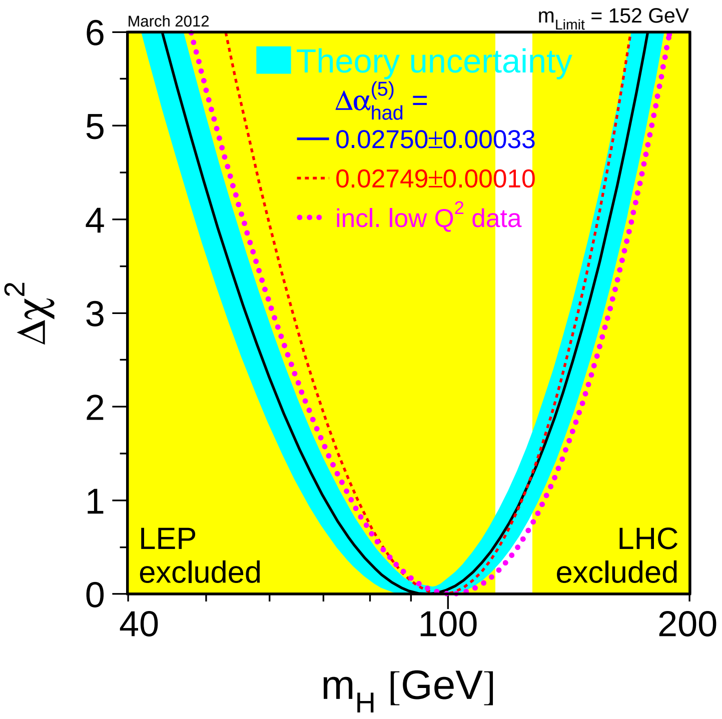

Figure 2.6: Best-fit values for the Higgs boson mass with experimental exclusions indicated [30].

corresponds to rotations in color space. The mediator of the strong force is the gluon, g. The gluon couples to any particle with color charge, including quarks, anti-quarks, and other gluons.

The portion of the Standard Model Lagrangian governing strongly interacting particles is given by

$$\mathcal{L}_{QCD} = \bar{q}(i\gamma^\mu\partial_\mu - m_q)q - g_s\bar{q}\gamma^\mu\lambda_a G^a_\mu q - \frac{1}{4}G^a_{\mu\nu}G^{\mu\nu}_a. \tag{2.43}$$

Here q represents a quark matter field (one for each quark flavor and color), λ_a are the Gell-Mann matrices (the eight-dimensional representation of SU(3)), and G^a_μ are the gluon vector fields. The primary quark/gluon interaction vertex is shown in Figure 2.7.

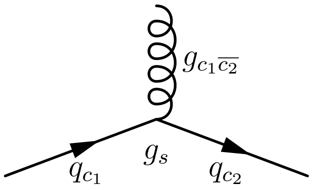

Figure 2.7: Primary QCD interaction vertex.

In this vertex, a quark carrying color c_1 enters, emits a gluon with one unit of color c_1 and one unit of anti-color $\overline{c_2}$, and exits with the same flavor but color c_2. The actual color composition of gluons can be more complex than is represented here, because adjoint representations of SU(3) must be formulated using at least some linear combinations of color-anticolor states to avoid creation of color singlet states. The color charge carried by gluons also gives rise to the three and four gluon vertices shown in Figure 2.8.

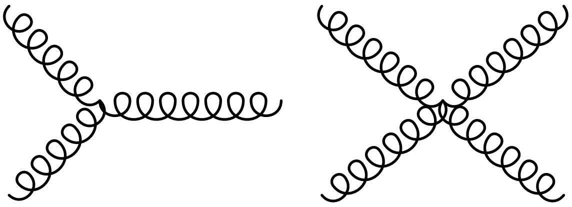

Figure 2.8: Three and four gluon QCD vertices.

Due to the algebraic structure of its underlying symmetry group and the quark content of the model, QCD exhibits a behavior known as "asymptotic freedom," where the interaction strength between particles *grows* as the distance between them increases. This gives rise to the phenomenon of "confinement," in which the existence of quarks and gluons is restricted to the bound states of hadrons. As quarks and gluons are pulled apart, they fragment into other hadrons, which in the context of a particle detector create particle showers known as "jets." Figure 2.9 shows a three jet event in the OPAL detector, where two quarks fragmented into jets, and one quark radiated a gluon that fragmented into a third jet.

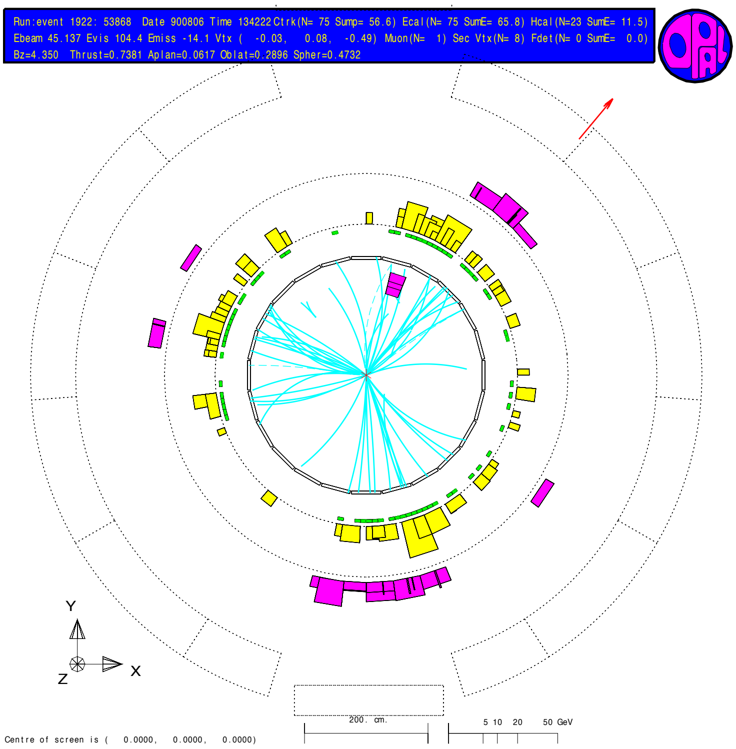

Figure 2.9: A three jet event in the OPAL detector [31].

2.2 Shortcomings of the Standard Model

While the Standard Model's agreement with experiment is nothing short of spectacular, it is very clearly *not* a complete and final theory.

2.2.1 Gravity

The most obvious shortcoming of the Standard Model is the glaring omission of gravity. To date, no suitable quantum theory of gravity has been formulated, primarily due to the fact that quantum gravity appears to be non-renormalizable. However, this absence is not entirely catastrophic to the theory, as gravity is many orders of magnitude weaker than the strong, weak, and electromagnetic forces at the energies and timescales associated with typical particle interactions.

2.2.2 The Hierarchy Problem

Another subtle but perhaps more important issue is the so-called "hierarchy problem." While not technically a deficiency of the Standard Model, the hierarchy problem arises when trying to formulate the Standard Model as part of a larger "Grand Unified Theory," where concern arises over the vast disparity in strength between the weak force and gravity. Theories that postulate unification of the fundamental forces at energies around the Planck scale necessarily invite large quantum corrections to the bare mass of the Higgs boson. In order to generate a Higgs boson mass on the order of ~ 100 GeV, these corrections, on the order of 10^{16} GeV or higher [6], would need to cancel to incredible precision. Such cancellations appear highly unnatural in the absence of any mechanism to explain them.

One proposed solution to the hierarchy problem is a class of theories known as "supersymmetry." Supersymmetry, or SUSY for short, postulates that each particle of the Standard Model has a corresponding "superpartner" particle with the same mass but different spin. In the case of fermions, the superpartners are bosons, and in the case of bosons, the superpartners are fermions. In such a scheme, the loop corrections to the Higgs boson mass cancel between particles and their superpartners, and no unnatural fine-tuning is required for the Higgs boson mass to have its observed value. Complete cancellation would require superpartners to have the same mass as their corresponding partners, a property which experimental searches have ruled out. Consequently, if SUSY exists, it must be a broken symmetry, and thus the loop cancellations induced by superpartners are not exact, merely mitigating the hierarchy problem into a so-called "little hierarchy problem."

2.2.3 Non-perturbative QCD

A more practical and pressing problem in Standard Model physics arises in the context of QCD interactions. In particular, the perturbative calculus that underlies most calculations in quantum field theory breaks down for low-energy QCD calculations, where the coupling constant, α_s, becomes large. The typical approach to this problem involves a process known as "factorization," where the scattering cross section for a QCD process is expressed as the convolution of two terms: a "hard" term, which can be calculated using standard perturbative methods, and a non-perturbative parton density function, values of which can be measured by experiment [32]. An important consequence of this procedure is the introduction of another calculational scale, known as the "factorization scale," which acts as a sort of cut-off between hard and soft QCD processes [32].

2.2.4 Parameters depending on experiment

Perhaps the most unsettling aspect of the Standard Model is its large number of parameters that must be taken solely from experiment. In particular, the Standard Model offers no *ab initio* prediction of particle masses nor coupling strengths. It also offers no explanation for the apparent generational structure of the quarks and leptons, nor for the disparities in masses between generations. It will be the job of future theories to further elucidate these aspects, though the Standard Model will likely remain the low-energy limit of a larger unified theory.

2.3 Potential Higgs boson search channels at the LHC

The 7 and 8 TeV proton-proton collisions of the LHC provide several Higgs boson production mechanisms, and the ATLAS and CMS detectors are both capable of identifying this production by searching a wide variety of decay channels [33, 34].

There are four primary Higgs boson production mechanisms at the LHC: gluon-gluon fusion (ggF), vector boson fusion (VBF), and two associated production modes. The associated production modes include one in which the Higgs boson is produced in association with a vector boson (VH, where V is a $W^{\pm}$ or Z boson) and another

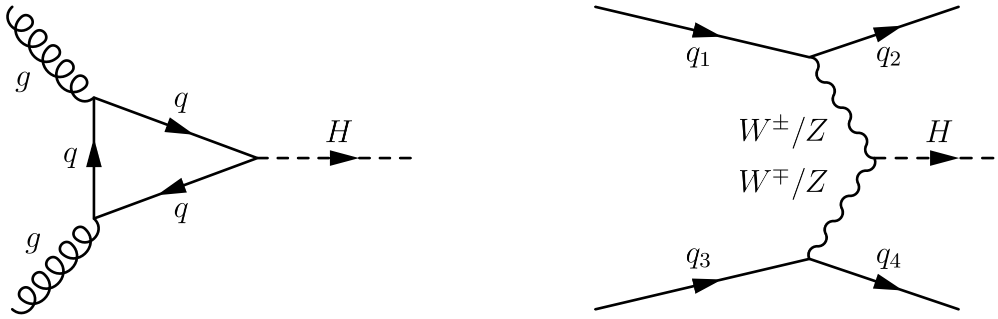

(a) Gluon-gluon fusion

(b) Vector boson fusion

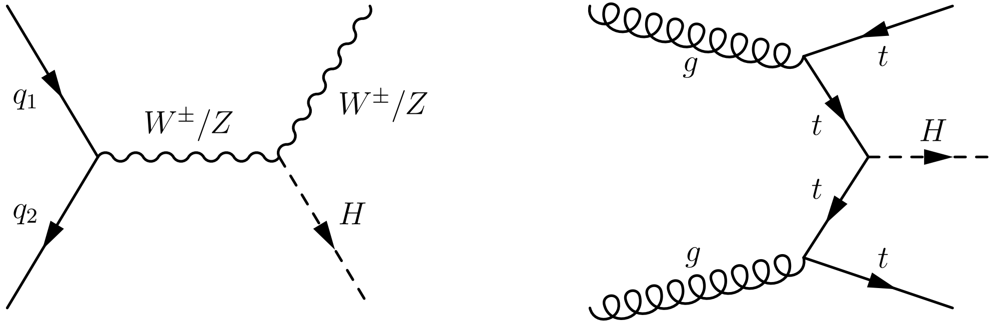

(c) VH associated production

(d) ttH associated production

Figure 2.10: Various Higgs boson production modes at the LHC, including (a) gluon-gluon fusion, (b) vector boson fusion, (c) VH associated production, (d) ttH associated production.

where the Higgs boson is produced in association with a $t\bar{t}$ pair (ttH). These processes are illustrated in Figure 2.10.

The cross sections for these production modes as a function of hypothetical Higgs boson mass are shown in Figure 2.11. The ggF and VBF processes are the dominant production modes at the LHC for $m_H > 100$ GeV.

In addition to multiple production modes, there is also a large number of decay modes that can be studied at the LHC. The Higgs boson branching ratios as a function of hypothetical Higgs boson mass are shown in Figure 2.12. Low-branching-ratio channels, such as $H \to \gamma\gamma$ and $H \to ZZ$, provide the cleanest probes of Higgs decays, while high-branching-ratio channels, such as $H \to b\bar{b}$, $H \to WW$, and $H \to \tau\tau$, provide more signal events at the cost of more complicated event topologies and larger backgrounds.

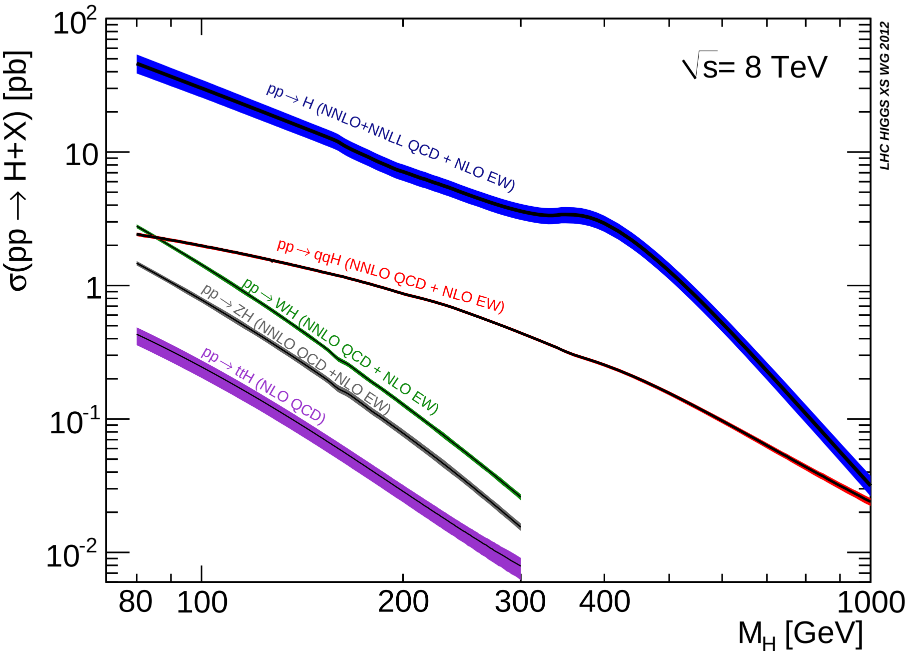

Figure 2.11: Higgs boson production cross sections for proton-proton collisions at $\sqrt{s}$ =8 TeV as a function of hypothetical Higgs boson mass [35].

2.3.1 Importance of the $H \to \tau\tau$ decay channel

The $H \to \tau\tau$ decay channel provides an important probe of Higgs physics at the LHC. In particular, it is currently the only accessible measurement of the Higgs boson coupling to leptons, providing direct evidence of mass generation in the lepton sector via the Higgs mechanism. Analysis of the VH production modes may also provide a measurement of the coupling ratio g_{HWW}/g_{HZZ} [36].

Quite fortuitously, the $H \to \tau\tau$ branching ratio in the vicinity of $m_H = 125$ GeV is relatively high compared to other decay modes. The τ leptons themselves decay hadronically with a branching ratio of $\sim 65\%$ and leptonically with a branching ratio of $\sim 35\%$ [12]. This leads to a rich collection of fully hadronic, fully leptonic, and mixed leptonic-hadronic final states, all of which can be feasibly exploited.

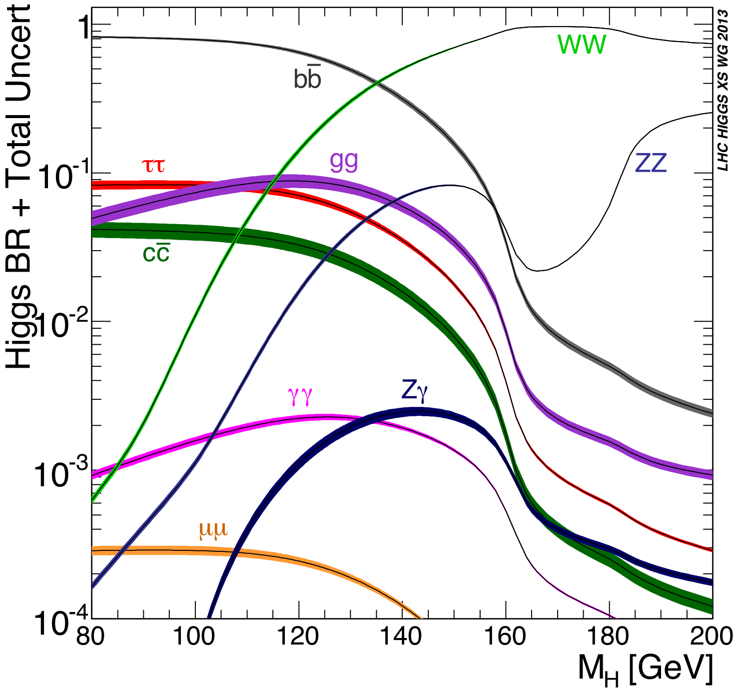

Figure 2.12: Higgs boson branching ratios as a function of hypothetical Higgs boson mass [35].

Chapter 3

Experimental setup

The $H \to \tau\tau$ analyses were performed on proton-proton collisions at the LHC using data collected by the ATLAS detector. This chapter provides an overview of the design specifications for both the collider and detector. The ATLAS luminosity measurement procedure is also outlined. The collision and data taking periods in 2011 and 2012, collectively referred to as "Run I," are summarized in Section 6.2.

3.1 The Large Hadron Collider

The LHC is a proton-proton and heavy-ion collider. The motivation for and design of the LHC are described in [37]. It is designed to facilitate the discovery of TeV-scale physics beyond the Standard Model through high center-of-mass energy collisions. The high-intensity beams of the LHC also allow for more precise probes of known Standard Model processes.

The LHC is part of a larger accelerator complex at CERN, which also includes multiple linear accelerators, the Proton Synchrotron Booster, the Low Energy Ion Ring (LEIR), the Proton-Synchrotron (PS), and the Super Proton Synchrotron (SPS). A diagrammatic representation of the CERN accelerator complex is given in Figure 3.1. These accelerators, some dating back to the 1950s, are used sequentially to gradually increase the energy of particles to a level suitable for injection into the LHC. The LHC itself is 26.7 km in circumference and was constructed in the tunnel originally built for

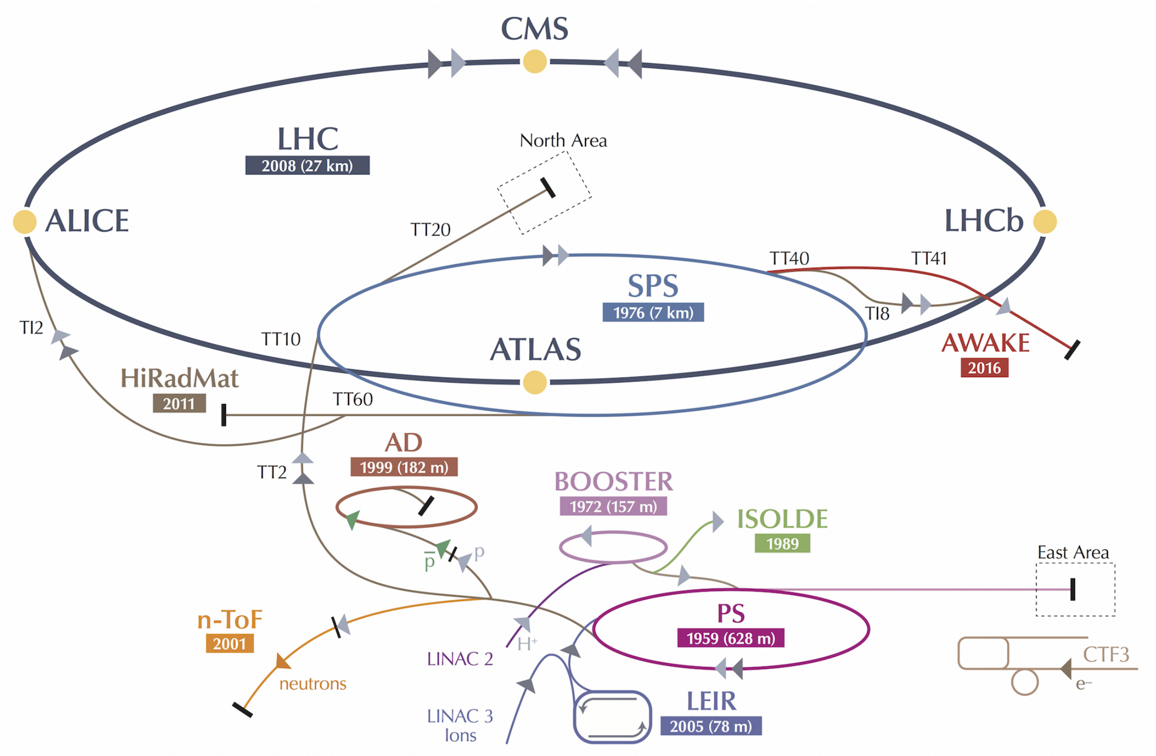

Figure 3.1: The CERN accelerator complex [38].

the Large Electron-Positron (LEP) collider. Having originally been constructed for a leptonic matter-antimatter collider, the tunnel imposed certain design constraints uncommon for a hadronic matter-matter collider.

The LHC consists of two counter-rotating proton-proton beams, each with a design energy of 7 TeV, for a combined center-of-mass collision energy of 14 TeV. Each beam is designed to contain 2,808 proton bunches, with each bunch containing 1.1×10^{11} protons. Colliding at a rate of 40 MHz yields a design instantaneous luminosity of 10^{34} $\text{cm}^{-2}\text{s}^{-1}$. This high luminosity necessitates the use of two matter beams, which must be spatially separated. This is in contrast to matter-antimatter colliders, such as the Tevatron, where the matter and antimatter beams may share the same beam pipe. The LHC is also designed to collide lead ion beams at a center-of-mass energy of 5.5 TeV and a peak instantaneous luminosity of 10^{27} $\text{cm}^{-2}\text{s}^{-1}$. The maximum energy of the LHC beams is limited by the strength of the magnetic dipole field in the storage ring, with a field of 8.33 T corresponding to an energy of 7 TeV. The maximum beam intensity is limited by a number of factors, including beam-beam interactions affecting each particle at collision points, the aperture of the beam screen, and heat load due to synchrotron radiation. Energy and luminosity are also both limited by the ability to safely discharge the beams.

The LHC lattice is primarily composed of 1,232 dipole magnets, although a variety of other magnets affecting beam parameters are also used. A twin-bore superconducting dipole magnet design was chosen to allow side-by-side placement of the proton beams, a necessity due to the size of the LHC tunnel. The LHC also includes 16 400 MHz superconducting radio-frequency (RF) cavities used to accelerate particle beams.

The LHC hosts a variety of experiments using the same proton and lead-ion beams, each with a different range of capabilities and physics goals. There are four primary experiments situated at the four beam-crossing points of the LHC. Two of these experiments are general-purpose detectors: ATLAS, short for A Toroidal LHC ApparatuS, and CMS, short for Compact Muon Solenoid. Both are designed to detect and study a range of Standard Model and beyond the Standard Model physics. There are also two special purpose detectors: LHCb, which focuses on the study of b-quark physics, and ALICE, short for A Large Ion Collider Experiment, which is dedicated to the study of the quark-gluon plasma that forms in the wake of heavy ion collisions.

3.2 The ATLAS detector

The ATLAS detector is a general-purpose detector situated at Point 1 of the LHC. The motivation for and design of the ATLAS detector is described in [39]. The goals of the ATLAS experiment include both discovery of new physics as well as precision study of known physics processes. One primary design benchmark for the ATLAS

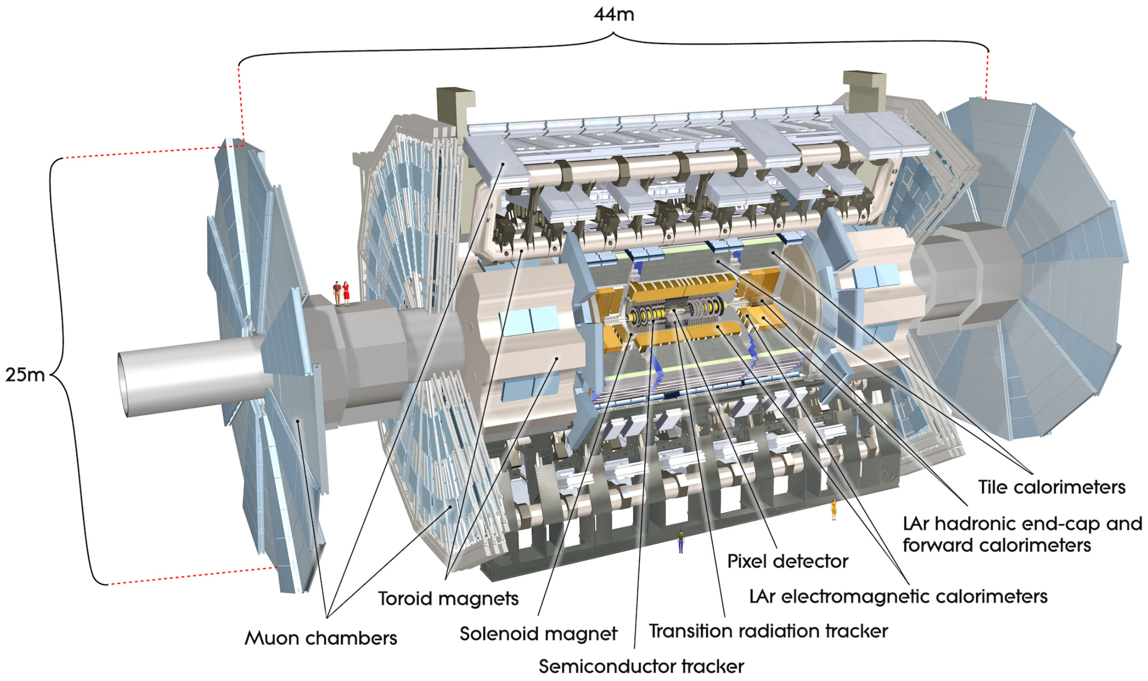

Figure 3.2: The ATLAS detector with subsystems labeled [39].

detector was its performance in detection and identification of the Standard Model Higgs boson.

The detector is designed as a cylindrically nested structure composed of multiple independent subsystems. The cylindrical “barrel” layers of the detector are complemented by circular “end-cap” layers, and together these layers work to surround the collision point as hermetically as possible. Each of the subsystems of the ATLAS detector serves a unique purpose. A diagram of the primary components of the ATLAS detector is given in Figure 3.2. In addition to the physical subsystems of the detector, hardware and software triggers are used to filter events recorded by the detector.

The coordinate system used when describing the detector originates from the nominal beam collision point in the center of the detector. The z axis runs along the beamline of the detector through the barrel section, while the x axis points toward the center of the LHC and the y axis points toward the top of the detector. A cylindrical coordinate system is most often used, with the $R - \phi$ plane perpendicular to the

z axis. The R coordinate is measured from the beam axis and the azimuthal angle ϕ is measured from the x axis in the direction of the y axis. The polar angle θ is measured relative to the z axis. For high-momentum particles, where $p \gg m$ $(E \approx p)$, pseudorapidity η is used to describe a particle's trajectory relative to the beamline:

$$\eta = -\ln\left[\tan\left(\frac{\theta}{2}\right)\right]. \tag{3.1}$$

For massive particles, rapidity y is used:

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

The metric for $\eta - \phi$ space, ΔR, is given by

$$\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}. \tag{3.3}$$

Two additional quantities, the transverse and longitudinal impact parameters, are also defined for tracks. The transverse impact parameter, denoted d_0, is the distance between the track and the beamline at perigee (the point along the track which is closest to the beamline). The longitudinal impact parameter, denoted z_0, is the z coordinate of the track at perigee.

One important design consideration for the ATLAS detector is its robustness to the effects of so-called "pile-up." There are two types of pile-up that affect the detector: "in-time" pile-up and "out-of-time" pile-up. In-time pile-up is due to additional inelastic proton-proton collisions in the same bunch crossing, which produce tracks and energy deposits in the detector. Out-of-time pile-up is due to proton-proton interactions in previous or future bunch crossings, which are problematic because the readout time of some detector systems, specifically calorimeters, is longer than the bunch crossing period. In addition to design considerations aimed at mitigating these effects, analysis-level corrections to account for pile-up are applied as described in Chapter 4.

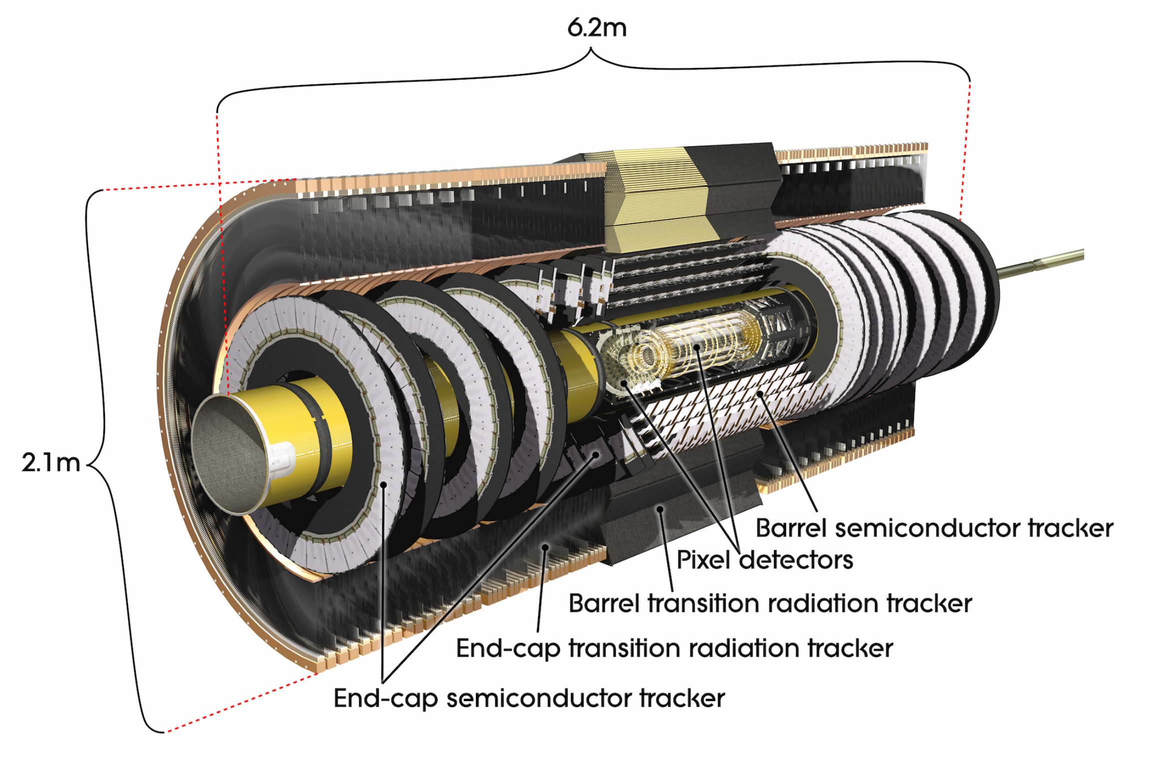

Figure 3.3: The ATLAS Inner Detector with subsystems labeled [39].

3.2.1 Inner Detector

The function of the Inner Detector is to provide precise, high-resolution tracking information. The Inner Detector is composed of three subsystems: the Pixel Detector, the Semi-Conductor Tracker (SCT), and the Transition Radiation Tracker (TRT). Tracking coverage is provided by the Pixel Detector and SCT up to $|\eta| = 2.5$ and by the TRT up to $|\eta| = 2.0$. A diagram of the Inner Detector components is shown in Figure 3.3.

The Inner Detector is permeated by a magnetic field generated by the ATLAS solenoid magnet system. This allows the Inner Detector to determine charged particle momentum, which is inversely proportional to track curvature, as well as particle charge, which is indicated by the direction of track curvature.

The position resolution of the Inner Detector decreases with distance from the beamline. The innermost layers of the Inner Detector are pixel layers, providing precise reconstruction of the collision vertex, surrounded by the lower-resolution SCT layers and the TRT. This approach allows for precise tracking without excess cost, material thickness, or readout bandwidth.

3.2.1.1 Pixel Detector

The Pixel Detector is a silicon-based detector that detects the transit of charged particles through silicon pixels. Particles passing through the pixels ionize the silicon, leaving a detectable charge distribution. The Pixel Detector is a modular design composed of three barrel layers complemented by three end-cap layers at each end of the detector. Each layer is a circular arrangement of staves mounted with multiple individual pixel modules. All pixel modules within the ATLAS detector are identical rectangular arrays of pixels. There is a total of 1,744 pixel modules within the ATLAS detector. Each pixel module contains 47,232 pixels, with $\sim 90\%$ of pixels 50 μm × 400 μm in size and the remainder 50 μm × 600 μm in size. Every pixel module contains 46,080 unique readout channels, resulting in approximately 80.4 million readout channels for the Pixel Detector. The Pixel Detector provides an intrinsic accuracy of 10 μm in the $R - \phi$ plane, with the barrel and end-cap layers providing an intrinsic accuracy of 115 μm along the z and R axes, respectively.

3.2.1.2 Semi-Conductor Tracker

The SCT is similar in purpose and function to the Pixel Detector, but is designed to be located farther from the beam axis with lower resolution requirements. The SCT is a modular design composed of four barrel layers complemented by nine end-cap layers at each end of the detector. Like the Pixel Detector, each SCT layer is composed of staves mounted with multiple individual modules. Unlike the Pixel Detector, SCT modules are composed of two bonded layers of silicon strips, with the two strip layers oriented at a small relative angle. This allows reconstruction of particle transits through the module in two dimensions, albeit with some ambiguity. It is also worth noting that, unlike the Pixel Detector, SCT modules are not completely homogeneous. In particular, modules in the barrel layers, which are made of rectangular arrays of silicon strips, differ from those in the end-cap layers, where the silicon strips are arranged along annular segments. There is a total of 8,176 SCT modules within the ATLAS detector, resulting in approximately 6.3 million readout channels for the SCT. The SCT provides an intrinsic accuracy of 17 μm in the $R-\phi$ plane, with the barrel and end-cap layers providing an intrinsic accuracy of 580 μm along the z and R axes, respectively.

3.2.1.3 Transition Radiation Tracker

The TRT is a straw tube tracker surrounding the SCT. The TRT only provides $R-\phi$ position information, with an intrinsic accuracy of 130 μm per straw. On average, the TRT yields 36 hits per particle track, and has a total of approximately 351,000 readout channels. The large number of hits and longer track length coverage enhance the precision of particle momenta measurements. Transition radiation photons, generated by electrons passing through the TRT, provide additional criteria for electron reconstruction and identification. In the barrel region, the TRT straws are arranged parallel to the beamline, and in the end-cap region, the TRT straws are arranged radially.

3.2.2 Magnet system

ATLAS contains two magnet systems: a solenoid and a toroid. The fields generated by these magnets cause charged particle tracks to curve, allowing the Inner Detector and Muon Spectrometer to determine their charge and momentum.

3.2.2.1 Solenoid magnet

The ATLAS solenoid magnet surrounds the Inner Detector, providing charged particle track curvature primarily for the Inner Detector. It generates an axial 2 T magnetic field which permeates the Inner Detector systems. It is 5.3 m long and 2.5 m in diameter, providing a near-uniform field in the central region, $|\eta| < 1.6$. At $|\eta| > 1.6$, the momentum measurement degrades by less than 10% due to non-uniformity of the magnetic field.

3.2.2.2 Toroid magnet

The ATLAS toroid magnet is arranged with eight-fold azimuthal symmetry around the calorimeter system. Its primary purpose is to provide bending of muons for momentum measurements by the Muon Spectrometer. It is composed of three parts: one barrel portion, with a field primarily permeating $|\eta| < 1.4$, and two end-cap portions, with fields primarily permeating $1.6 < |\eta| < 2.7$. A photograph of the barrel portion of the ATLAS toroid is shown in Figure 3.4. The bending power of the toroid is calculated in terms of the integrated field strength between the innermost and outermost muon-chamber planes, ranging from 1.5 to 5.5 Tm in $|\eta| < 1.4$ and 1 to 7.5 Tm in $1.6 < |\eta| < 2.7$. A smaller but non-zero magnetic field exists in the $1.4 < |\eta| < 1.6$ region due to overlap of the barrel and end-cap magnetic fields.

3.2.3 Calorimetry

There are two primary calorimeters in the ATLAS detector: the electromagnetic calorimeter and the hadronic calorimeter. A forward calorimeter also provides coverage of high-pseudorapidity regions. All use different techniques to measure energy deposits in the detector. Precise energy measurement is essential for the reconstruction and calibration of particle and jet energy, as well as E_T^{miss}. The layout of the ATLAS calorimeter components is shown in Figure 3.5.

Hermeticity is essential for ensuring proper jet energy and E_T^{miss} calibration, as well as minimizing the interaction of “punch-through” particles (jet debris leaking through the calorimeter) with the muon system. Anywhere from 9.7 to 11 interaction

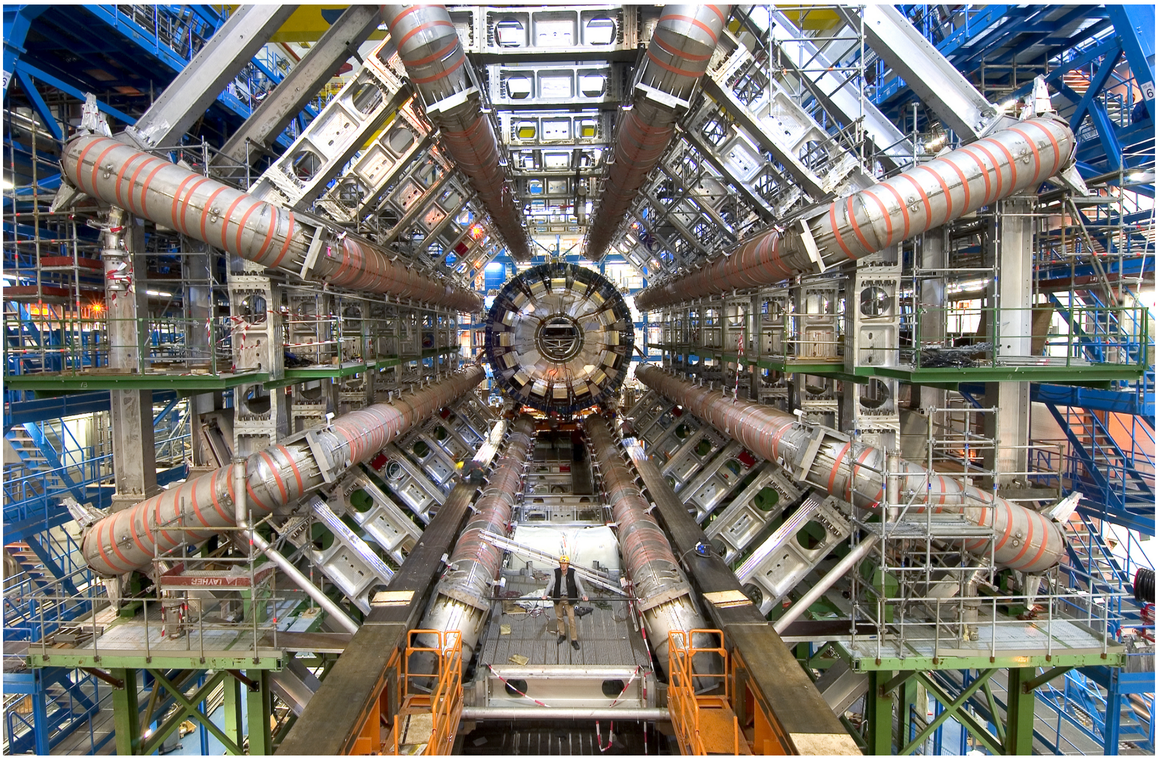

Figure 3.4: The barrel segment of the ATLAS toroid magnet [39].

lengths of active material across the η range of the detector ensure jet punch-through is reduced to a level below that of other muon backgrounds.

3.2.3.1 Electromagnetic calorimeter

The electromagnetic calorimeter is a liquid argon based sampling calorimeter. It uses alternating layers of lead absorber and liquid argon active material to measure energy deposits of electrons and photons. The calorimeter is composed of accordion-shaped electrodes and absorber plates to provide isotropic coverage in ϕ.

The barrel region segment of the calorimeter is separated into two identical cylindrical components separated by 4 mm at $z = 0$, and covers the range $|\eta| < 1.475$. The end-cap segments of the calorimeter are separated into two coaxial wheels covering $1.375 < |\eta| < 2.5$ and $2.5 < |\eta| < 3.2$.

The barrel segment and outer-wheel of the end-cap segment are divided into three layers, each with decreasing granularity. The first layer provides extremely fine seg-

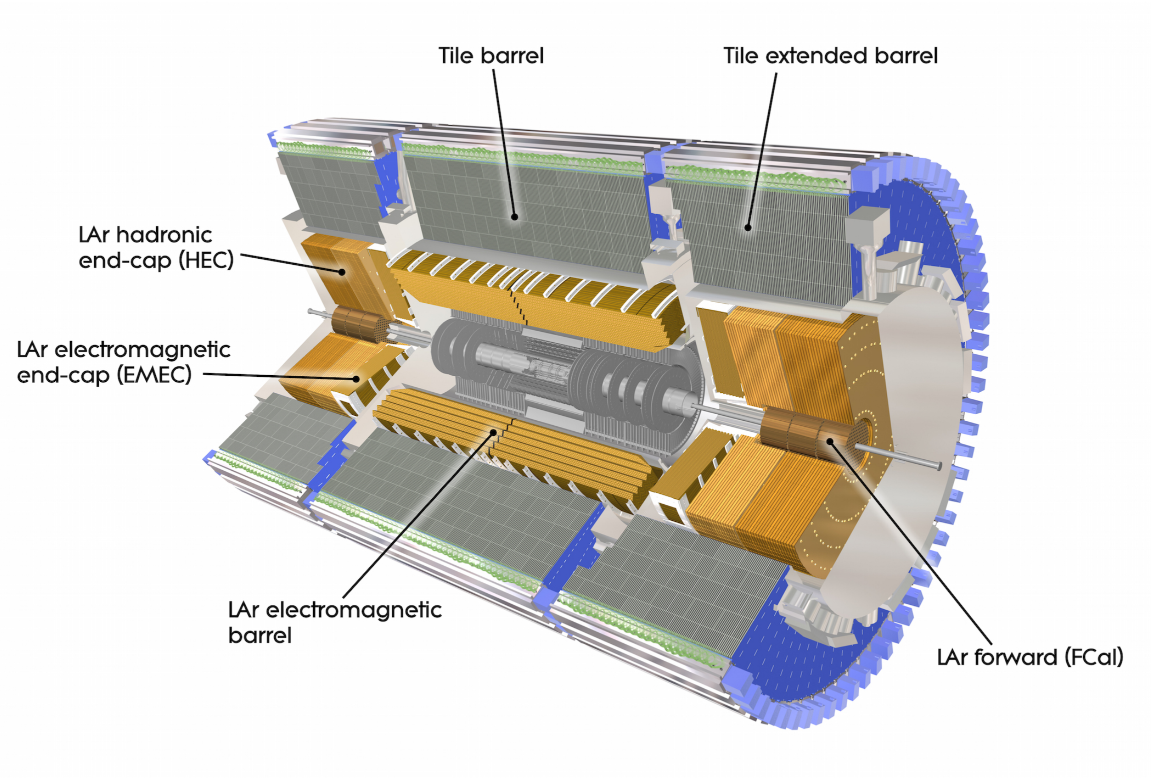

Figure 3.5: The ATLAS calorimeter system [39].

mentation in η, with strip cells covering regions of 0.0031×0.098 in $\eta \times \phi$ space. The second layer, where the majority of energy is deposited, uses square cells of size 0.025×0.0245 in $\eta \times \phi$ space. The third layer, which absorbs the tails of electromagnetic showers, is segmented into cells of 0.05×0.0245 in $\eta \times \phi$ space. The electromagnetic calorimeter has approximately 180,000 readout channels.

The electromagnetic calorimeter is also complemented by two argon presampler layers in the region $|\eta| < 1.8$. These presamplers provide an energy measurement for particles that begin to shower before reaching the electromagnetic calorimeter.

3.2.3.2 Hadronic calorimeter

The hadronic calorimeter is a sampling calorimeter used to measure the energy of hadronic decay products in the ATLAS detector. It consists of a large barrel segment, two smaller extended barrel segments, and two end-cap components. It is positioned directly outside of the electromagnetic calorimeter. The barrel segments of the calorimeter are composed of alternating layers of steel and scintillating tiles, which are read out by wavelength-shifting fibers and photomultiplier tubes.

The central barrel segment covers $|\eta| < 1.0$ and the extended barrel segments cover $0.8 < |\eta| < 1.7$. The end-cap segments of the calorimeter are liquid-argon-based, using copper as an absorber. They are separated into two coaxial wheels covering $1.5 < |\eta| < 3.2$.

The fibers that read out light from the scintillating tiles are grouped in a manner that divides the hadronic calorimeter into three layers of cells, with the first two layers having dimensions of 0.1×0.1 in $\eta \times \phi$ space and the third layer having dimensions of 0.2×0.1 in $\eta \times \phi$ space. The hadronic calorimeter has approximately 15,000 readout channels.

3.2.3.3 Forward calorimeter

The forward calorimeter provides energy measurements in the high-pseudorapidity region, $3.1 < |\eta| < 4.9$. It uses liquid argon as its active material, sharing a cryostat with the electromagnetic and hadronic calorimeter end-caps. There are three modules in each end-cap of the forward detector. The first module uses copper as absorption material and is primarily designed for electromagnetic measurements. The latter two modules use tungsten as absorption material and are optimized for hadronic measurements.

3.2.4 Muon system

The ATLAS muon system is divided into two components: a tracking system (the Muon Spectrometer), which provides precise measurements of muon momenta, and a triggering mechanism, which allows for high-speed identification of muons in an event. The ATLAS muon system components are shown in Figure 3.6.

As with the Inner Detector, the muon momentum measurement relies on the magnetic deflection of charged particles, in this case provided by the ATLAS toroid

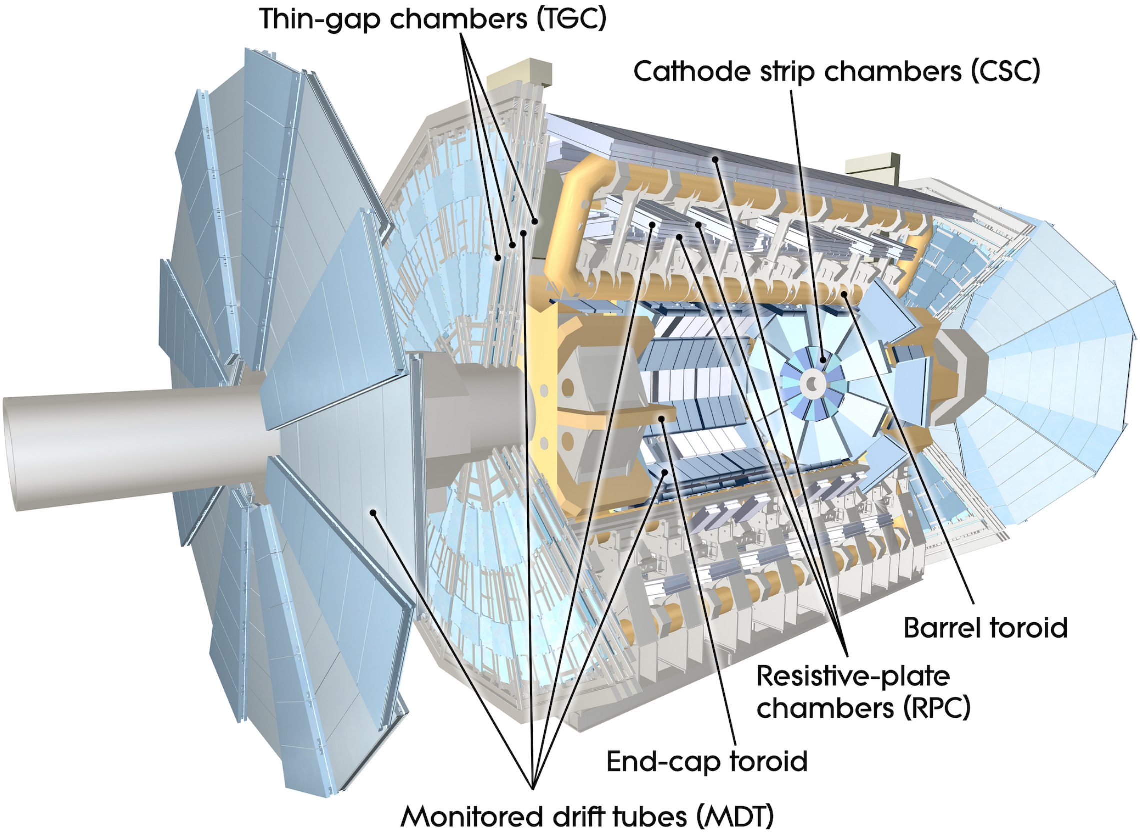

Figure 3.6: The ATLAS Muon Spectrometer [39].

magnets. Tracking in the barrel region is provided by three layers of monitored drift tubes (MDTs) covering $|\eta| < 2.7$, supplemented in the first layer of the end-caps by cathode strip chambers (CSCs) covering $2.0 < |\eta| < 2.7$. The CSCs provide better performance at higher rates, as well as better time resolution. The MDTs have a spatial resolution of 35 μm per chamber, while the CSCs provide a spatial resolution of 40 μm in the y-z plane and 5 mm in the transverse plane. The Muon Spectrometer is designed to measure the momenta of 1 TeV tracks to 10% accuracy and to provide a precise standalone measurement of low-p_T muons.

The muon trigger system, which covers $|\eta| < 2.4$, is capable of triggering on muons with a time resolution of 15–25 ns. It uses resistive plate chambers (RPCs) in the barrel and thin gap chambers (TGCs) in the end-caps to detect the presence of muons in an event for use in the Level 1 trigger.

3.2.5 Trigger

The tremendous event rate generated by the LHC necessitates rapid yet stringent selection criteria on proton-proton collisions in order to meet the target event recording rate of ~ 200 Hz. ATLAS utilizes a three-tier trigger system to select events based on criteria of interest, such as large E_T^{miss}, high-p_T muons, jets, electrons, photons, or hadronically decaying τ leptons.

The level 1 (L1) trigger is hardware-based, using a combination of calorimeter and muon trigger information to select events. As part of its selection process, it identifies solid angle Regions of Interest (RoIs), where interesting features have been flagged for analysis by subsequent trigger levels. It does not use any tracking information. It is designed to reduce the event rate to 75 kHz in a processing window of 2.5 μs.

The level 2 (L2) trigger is software-based, and has access to all detector data within RoIs. It uses approximations of offline algorithms to reduce the event rate to approximately 3.5 kHz in a processing window of 40 ms.

The third component of the ATLAS trigger system is the Event Filter (EF). It is a software-based trigger using near-offline analysis procedures to reduce the event rate to 200 Hz in a processing window of ~ 4 s.

The level 2 trigger and Event Filter are collectively referred to as the "High-Level Trigger" (HLT).

3.2.6 Luminosity measurement

The ATLAS luminosity measurement procedure is summarized in [40]. Accurate integrated luminosity measurement is essential for event rate modeling. Uncertainties associated with the integrated luminosity measurement are some of the main systematic uncertainties for many analyses, including $H \to \tau\tau$, and are detailed in Chapter 8.

In ATLAS, integrated luminosity is measured in time segments referred to as "lumi blocks." Each lumi block is 60 s in duration and assumed to have the same average luminosity. Various quality and consistency criteria are applied to lumi blocks, and the total recorded integrated luminosity is determined by the sum of good lumi blocks.

The instantaneous luminosity for a storage ring collider, such as the LHC, is given by

$$\mathscr{L} = \frac{\mu n_b f_r}{\sigma_{\text{inel}}}, \tag{3.4}$$

where μ is the average number of inelastic collisions per bunch crossing, n_b is the number of bunch pairs colliding per revolution, f_r is the revolution frequency, and σ_{inel} is the proton-proton inelastic cross section. Because the ATLAS detector is not perfectly efficient, the calculation must be recast in terms of visible quantities:

$$\mathscr{L} = \frac{\mu_{\text{vis}} n_b f_r}{\sigma_{\text{vis}}}, \tag{3.5}$$

where $\mu_{\text{vis}} = \varepsilon\mu$ and $\sigma_{\text{vis}} = \varepsilon\sigma_{\text{inel}}$ represent quantities adjusted for detector and reconstruction algorithm efficiency ε.

μ_{vis} is an experimentally observable quantity that can be determined on a bunch-by-bunch basis or by integrating over recorded data. On average, $\mu_{\text{vis}} > 1$, precluding the use of simple event-counting algorithms and necessitating the use of hit-counting algorithms to measure a signal proportional to μ_{vis}. In ATLAS, hit-counting is performed by the Beam Conditions Monitor (BCM) and Luminosity Measurement Using Cherenkov Integrating Detector (LUCID). The BCM uses four diamond sensors, arranged around the beam pipe at $z = \pm 184$ cm, to measure incident particles. LUCID uses sixteen aluminum tubes filled with C_4F_{10}, arranged around the beam pipe at $z = \pm 17$ m and covering $5.6 < |\eta| < 6.0$, to measure generation of Cherenkov radiation by particles using photomultipliers.

Calibration and determination of σ_{vis} is performed by dedicated van der Meer scans only a few times over the course of collider operation. These scans use stepped separation of the proton beams to measure beam shape parameters Σ_x and Σ_y, from which the luminosity can be determined (and Equation 3.5 calibrated) using the relation

$$\mathscr{L} = \frac{n_b f_r n_1 n_2}{2\pi\Sigma_x\Sigma_y}, \tag{3.6}$$

where n_i is the number of protons per bunch in beam i.

Chapter 4

Triggering and object reconstruction

A variety of physics objects are used in the $H \to \tau\tau$ analyses, including electrons, muons, jets, hadronically decaying τ leptons, and E_T^{miss}. The electrons and muons generated by leptonically decaying τ leptons are effectively indistinguishable from so-called "prompt" electrons and muons, and thus the identification of τ_l decays is equivalent to identifying their associated leptonic decay products.

Although each analysis uses slightly different object selection criteria, the triggers and reconstruction algorithms used to identify these objects are the same in all three cases. Because triggers and reconstruction algorithms have different performance for data and Monte Carlo, scale factors must be derived to correct for differing efficiencies.

This chapter summarizes the triggers and reconstruction algorithms used to identify events and physics objects of interest. The measurements of corresponding efficiencies and scale factors are also outlined. Finally, the computation of E_T^{miss} is detailed, along with a description of the object overlap removal procedure. A discussion of the data preparation and track reconstruction algorithms underlying object reconstruction can be found in Chapter 5. Event selections based on triggers and reconstructed physics objects are discussed in Chapters 6 and 7, while uncertainties associated with trigger and object reconstruction efficiencies are discussed in Chapter 8.

4.1 Electrons

Electrons play a key role in $H \to \tau\tau$ analyses, particularly in those analysis channels with leptons in their final state. Electrons are used for triggering and selection, and as inputs to $\tau\tau$ invariant mass reconstruction. The triggering, reconstruction process, identification criteria, and absolute energy calibration for electrons are documented in [41] and [42]. Reconstruction efficiency measurements for 2011 operation are given in [43] and for 2012 operation in [44].

4.1.1 Triggering

The L1 trigger computes the transverse energy of deposits in the electromagnetic calorimeter in towers of cells, each with size 0.1×0.1 in $\eta \times \phi$ space. Adjacent pairs of towers that pass a particular E_T threshold are used to seed the L2 trigger, where tracking information is added for further discrimination. Finally, the EF uses approximations of the offline reconstruction and identification algorithms described in Section 4.1.2 with slightly looser selections to identify events for offline reconstruction.

The efficiency of electron triggers is measured using a "tag-and-probe" method in $Z \to ee$ data. Tag-and-probe methods work by isolating an unbiased high-purity sample of a particular type of event, typically by identifying a resonance such as a Z boson or J/ψ meson. Events in this sample contain a well-identified "tag" object, such as an electron or muon, and a loosely identified "probe" object, which serves as a test object. The efficiency for a given reconstruction process or selection can then be determined by measuring the number of probe objects which are reconstructed or pass the selection. Efficiencies with respect to offline reconstruction for the various electron trigger levels as a function of electron E_T and η are shown in Figure 4.1 for the combined electron trigger used in the 8 TeV $H \to \tau_l\tau_h$ and $Z_{ll}H \to \tau_l\tau_h$ analyses.

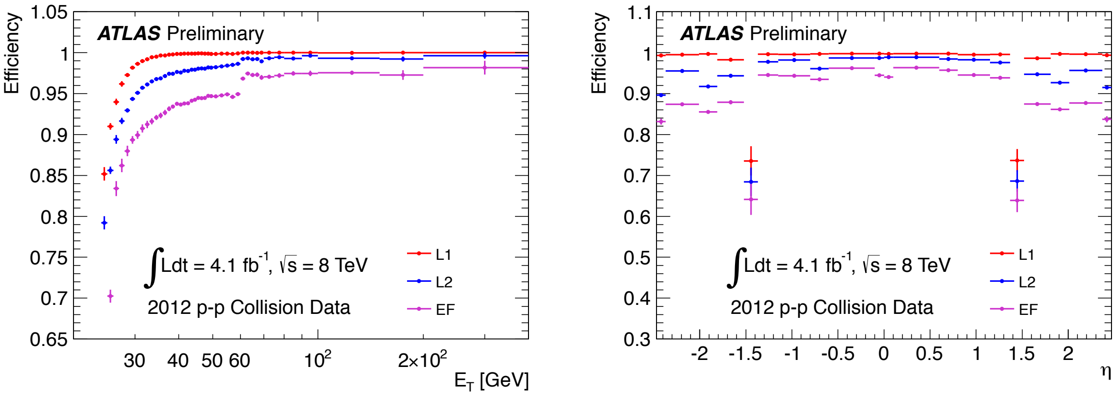

(a) (b)

Figure 4.1: L1, L2, and EF trigger efficiencies with respect to offline reconstruction as a function of (a) electron E_T and (b) electron η measured using $Z \rightarrow ee$ tag-and-probe methods with 4.1 fb^{-1} of 8 TeV data for the combined `EF_e24vhi_medium1 || EF_e60_medium1` trigger [45].

4.1.2 Reconstruction

Electron reconstruction is based on deposits in the electromagnetic calorimeter paired with tracks from the Inner Detector. In the forward regions of the detector ($|\eta| > 2.5$), where tracking information is unavailable, the reconstruction algorithms are effectively unable to distinguish between electrons and photons due to their similar deposits in the electromagnetic calorimeter.

Reconstruction itself is a three-step process. First, a cluster recognition algorithm identifies clusters of energy deposits within the electromagnetic calorimeter that are possibly due to an electron. Second, a track association algorithm attempts to pair clusters with tracks from the Inner Detector. Finally, after an electromagnetic cluster has been paired with a corresponding track, the cluster size and shape are reoptimized taking into account the trajectory of the track.

4.1.2.1 Cluster recognition

The cluster recognition algorithm uses a sliding-window scan with a window size of 3×5 cells (3 cells across in the η direction, 5 across in ϕ), each with size 0.025×0.0245

in $\eta \times \phi$ space. To qualify as a seed for reconstruction, the transverse energy of the cluster must be greater than 2.5 GeV.

4.1.2.2 Track association

The track association algorithm works by extrapolating tracks with $p_T > 0.5$ GeV into the center layer of the electromagnetic calorimeter. Tracks are considered a match if $|\Delta\eta| < 0.05$ and $|\Delta\phi| < 0.05$ between the track and the cluster barycenter, except in the direction of curvature of the track, where the requirement is loosened to $|\Delta\phi| < 0.1$. In the case of multiple matching tracks, the one with the smallest ΔR from the cluster barycenter is chosen. If no matching track is found, the cluster is considered a photon candidate. Discrimination between prompt electrons and converted photons is performed by checking for nearby tracks originating from a displaced vertex and by examining the innermost hits of the associated track in the Inner Detector.

4.1.2.3 Cluster size optimization

After clusters have been associated with a track, their sizes are optimized to take into account differences in the expected energy distribution in different regions of the detector. In the barrel region the cluster size is increased to 3×7 cells, while in the end-cap region the cluster size is increased to 5×5 cells, each cell having size 0.025×0.0245 in $\eta \times \phi$ space.

4.1.3 Energy calibration

The electron candidate energy is computed using its deposits in the electromagnetic calorimeter. It is corrected by estimating energy loss in passive material before the electromagnetic calorimeter (using the presampler) and to the sides of and behind the electromagnetic calorimeter.

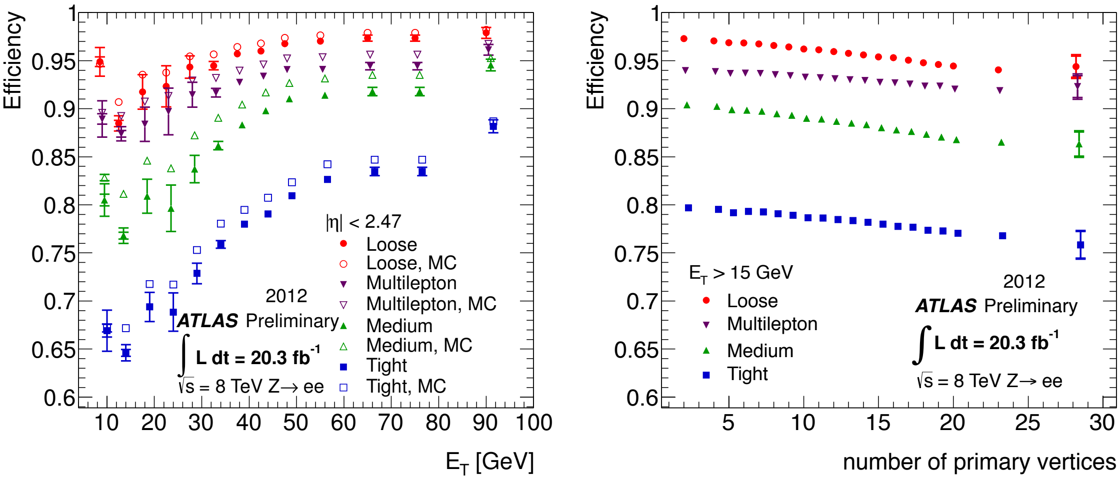

(a) (b)

Figure 4.2: Electron identification efficiencies as a function of (a) electron E_T and (b) the number of primary vertices measured using $Z \rightarrow ee$ tag-and-probe methods with 20.3 fb^{-1} of 8 TeV data and corresponding Monte Carlo samples [46].

The response of the electromagnetic calorimeter to electrons and photons is derived in ATLAS using test-beam measurements. The final electron energy scale calibration uses *in situ* techniques based on $Z \rightarrow ee$ and $J/\psi \rightarrow ee$ decays.

4.1.4 Identification and selection

Electron identification criteria are based on electromagnetic calorimeter and track variables. Cuts on these variables are optimized based on cluster η and E_T. Three working points are defined in terms of these cut values: loose, medium, and tight. Each of these working points adds new cuts and tightens previous ones, trading identification efficiency for background rejection. The efficiency of these working points is measured using $Z \rightarrow ee$ tag-and-probe methods. Identification efficiencies as a function of electron E_T and the number of primary vertices are shown for each working point in Figure 4.2.

In the VH analyses, electrons are required to pass the medium identification criteria. This includes cuts on shower-shape variables, hadronic calorimeter leakage, track and track-cluster matching, transverse impact parameter, and transition radiation identified by the TRT. In the $H \to \tau_l \tau_h$ analysis, electrons are required to pass tight identification, which tightens the cuts used in the medium identification, adds additional criteria based on track quality and the ratio of cluster energy to track momentum, and places a veto on photon conversions. After identification, additional criteria are applied using calorimeter information to ensure purity of the reconstructed electrons.

In the $H \to \tau\tau$ analyses, electron selection is restricted to $|\eta| < 2.47$ to ensure coverage by the Inner Detector. Electrons in the calorimeter crack region, $1.37 < |\eta| < 1.52$, are excluded from analysis due to poor identification performance, but are still used in the overlap removal procedures described in Section 4.6. In the main analyses, electrons are required to have $p_T > 15$ GeV, a requirement that is relaxed in the VH analyses to $p_T > 10$ GeV. The VH analyses also place a threshold on the longitudinal impact parameter, $|z_0| < 10$ mm. In the $H \to \tau_l \tau_h$ and VH analyses, two additional isolation criteria are placed on electrons. The first is a track-based isolation:

$$I^e_{p_T}(R_I) = \frac{\sum_{\Delta R < R_I} p_T^{\text{track}}}{p_T^e}. \tag{4.1}$$

Here the sum is over all tracks within an η-ϕ cone of radius R_I surrounding the electron (excluding the electron track). It is required that $I^e_{p_T}(0.4) < 0.06$ for the $H \to \tau_l \tau_h$ analysis and $I^e_{p_T}(0.4) < 0.2$ for the VH analyses. The second isolation requirement is calorimeter-based:

$$I^e_{E_T}(R_I) = \frac{\sum_{\Delta R < R_I} E_T^{\text{cell}}}{p_T^e}. \tag{4.2}$$

Here the sum is over clusters of energy deposits in the electromagnetic and hadronic calorimeters within a cone of radius R_I surrounding the electron. It is required that $I^e_{E_T}(0.2) < 0.06$ for the $H \to \tau_l \tau_h$ analysis and $I^e_{E_T}(0.2) < 0.2$ for the VH analyses.

4.2 Muons

Like electrons, muons also play an important role in the $H \to \tau\tau$ analyses with leptons in the final state, where they are used for triggering, selection, and mass reconstruction. The complete reconstruction process, identification criteria, and efficiency measurements for muons in the ATLAS detector are described in detail in [39] and [47]. Muons have a dedicated hardware-based triggering system and their momenta are measured to high-precision by the Muon Spectrometer. Both systems are described in Section 3.2.4. Different algorithms can subsequently be used to reconstruct a variety of muon types.

4.2.1 Reconstruction

Reconstruction of muons is performed by a family of algorithms collectively referred to as the "STACO" algorithms. These algorithms use information from the Muon Spectrometer, Inner Detector, and calorimeters. There are four types of reconstructed muons in ATLAS:

- **Stand-alone**: Stand-alone muons are those identified solely by the Muon Spectrometer. These muons are required to have at least two hits in the muon tracking chambers to allow extrapolation to the beamline. They are primarily useful in regions where tracking is not provided by the Inner Detector.
- **Combined**: Combined (CB) muons are those identified by both the Inner Detector and the Muon Spectrometer. Tracks are reconstructed separately for each subsystem and then matched and merged. This is the primary type of reconstructed muon used within ATLAS.
- **Segment-tagged**: Segment-tagged (ST) muons are identified by extrapolating a track from the Inner Detector to a hit in the Muon Spectrometer. This type

of reconstruction is useful in cases where a muon crosses only one layer of the Muon Spectrometer, either due to a low p_T or a trajectory with little coverage.

- **Calorimeter-tagged**: Calorimeter-tagged (CaloTag) muons are those identified by extrapolating a track from the Inner Detector into a calorimeter energy deposit consistent with a muon. These muons have low purity and are used primarily in regions where coverage is not provided by the Muon Spectrometer.

4.2.2 Energy calibration

Although the Muon Spectrometer provides a highly accurate measurement of muon momenta, additional corrections are needed to account for interaction with passive material, multiple scattering, limited resolution or misalignment of the Muon Spectrometer, and uncertainties related to the magnetic field. Final calibration of muon momenta is performed using *in situ* techniques based on $Z \to \mu\mu$ and $J/\psi \to \mu\mu$ decays.

4.2.3 Identification and selection

The efficiency of muon reconstruction is measured using $Z \to \mu\mu$ and $J/\psi \to \mu\mu$ tag-and-probe methods and shown in Figure 4.3.

As with electron identification, loose, medium, and tight identification thresholds are defined, with classification determined by the STACO algorithms based on the number of hits and tagging segments in the Muon Spectrometer. In the $H \to \tau_l\tau_h$ and $Z_{ll}H \to \tau_l\tau_h$ analyses, loose combined muons are used. After identification, additional criteria are applied on the Inner Detector track associated with the muon.

In the $H \to \tau\tau$ analyses, muon selection is limited to $|\eta| < 2.5$ due to the use of combined muons and their reliance on the Inner Detector. In the main analyses, muons are required to have $p_T > 10$ GeV, whereas for the VH analyses they are only

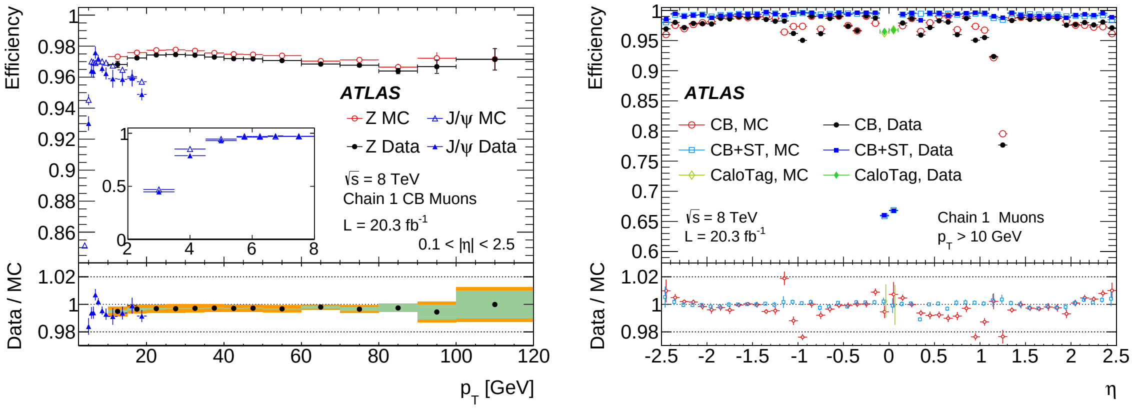

(a) (b)

Figure 4.3: Muon reconstruction efficiencies as a function of (a) muon p_T and (b) muon η measured using $Z \to \mu\mu$ and $J/\psi \to \mu\mu$ tag-and-probe methods with 20.3 fb^{-1} of 8 TeV data and corresponding Monte Carlo samples [47]. "Chain 1" indicates usage of the STACO algorithms.

required to have $p_T > 6$ GeV. The VH analyses also place a requirement on the longitudinal impact parameter, $|z_0| < 10$ mm. In the $H \to \tau_l\tau_h$ and VH analyses, the same p_T and E_T isolation criteria that are applied to electrons are also applied to muons: $I^{\mu}_{p_T}(0.4) < 0.06$ and $I^{\mu}_{E_T}(0.2) < 0.06$ for the $H \to \tau_l\tau_h$ analysis, and $I^{\mu}_{p_T}(0.4) < 0.06$ and $I^{\mu}_{E_T}(0.2) < 0.2$ for the VH analyses.

4.3 Jets

Jets play an important role in the $H \to \tau\tau$ analyses, both in identifying signatures of production modes (such as VBF) and in seeding the reconstruction and identification of hadronically decaying τ leptons. The reconstruction and identification of jets are documented in [48]. Jet reconstruction is a multi-step process and involves careful calibration of calorimeter energy deposits in order to faithfully reconstruct jet energy.

The first step of jet reconstruction is the formation of three-dimensional topological clusters of calorimeter cells. The energy deposits in these clusters are then calibrated using a local hadronic calibration scheme, which accounts for differences in calorimeter response due to shower type and location in the detector. Finally, jets are reconstructed using calibrated clusters as input, and the final jet energy scale (JES) calibration is performed. After jet reconstruction is complete, a b-tagging procedure is applied; this is used in the $H \to \tau\tau$ analyses to help identify those jets originating from top quark processes.

4.3.1 Clustering and local hadronic calibration

The input objects to the jet reconstruction algorithms are three-dimensional topological clusters of calorimeter cells referred to as "topo-clusters" [49]. The clustering algorithm used to generate these clusters is designed to capture the low-energy tails of particle showers in the calorimeter while suppressing noise from electronics and pile-up. The formation of clusters begins with the identification of seed cells, which are cells that have a signal-to-noise ratio, t, above a large threshold, $t_{\text{seed}} = 4$. Each seed cell is used as the start of a cluster, with immediately adjacent cells added to the cluster if they have a signal-to-noise ratio above a low threshold, $t_{\text{cell}} = 0$. If any of these adjacent cells have a signal-to-noise ratio above a medium threshold, $t_{\text{neighbor}} = 2$, they are used as secondary seeds to continue expansion of the cluster. Once clusters are identified, a second algorithm searches for local energy deposit maxima in cells of the cluster. If more than one such maximum is found, the cluster is split by reforming separate clusters using the maxima as seeds.

After topo-clusters have been identified, their energies are recalibrated using a local hadronic calibration scheme [50]. This recalibration is necessary because the calorimeter response for hadronic energy deposits, mostly due to charged pions and neutrons, is different than the response for electromagnetic deposits, mostly due to electrons and photons. Additional corrections must be made for hadronic energy deposits not captured in topo-clusters and for the varying amounts of passive material in front of the calorimeter in different regions of the detector.

The appropriate corrections for a given topo-cluster are determined by classifying

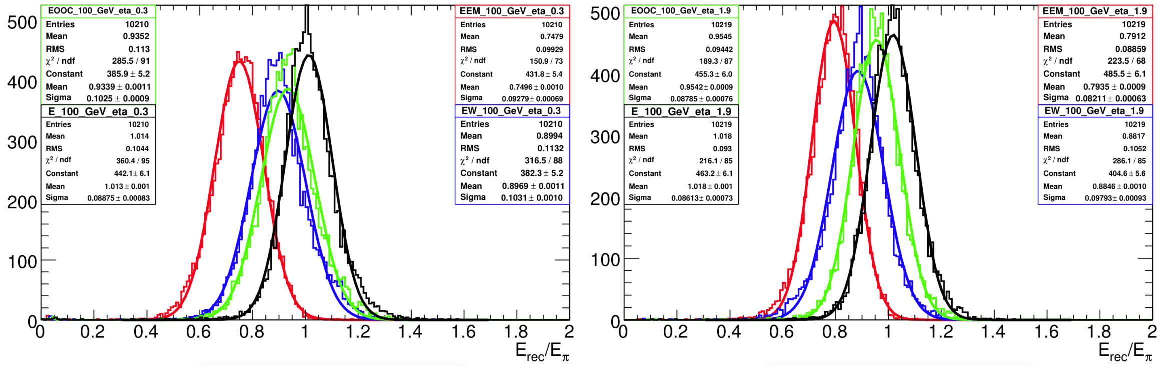

(a) $0.2 < |\eta| < 0.4$ (b) $1.8 < |\eta| < 2.0$

Figure 4.4: Ratio of reconstructed to simulated pion energy at subsequent stages of the local hadronic calibration procedure for 100 ± 30 GeV pions in different η regions. Stages shown include without any corrections (red), with hadronic scale calibrations added (blue), with out-of-cluster corrections added (green), and finally with passive material corrections added (black) [50].

the cluster as either hadronic or electromagnetic by means of cluster shape variables. These variables include various energy-weighted cluster moments,

$$\langle x \rangle = \frac{1}{E_{\text{cluster}}} \cdot \sum_{\text{cells}} E_{\text{cell}} \cdot x_{\text{cell}}, \tag{4.3}$$

where x is an observable such as η, ϕ, or distance from shower axis and center, as well as variables related to cluster isolation and the fraction of energy deposited in the electromagnetic calorimeter and the most active cell(s) of the cluster. The corresponding correction factors are determined from simulations of charged pions using Geant4 [51]. Performance of the calibration scheme at various stages is shown in Figure 4.4.

4.3.2 Reconstruction

After topo-cluster identification and calibration, jets in ATLAS are reconstructed using the anti-k_t algorithm [52]. The anti-k_t algorithm is a type of sequential recombination algorithm. A sequential recombination algorithm iterates over objects in a detector (e.g. particle tracks or calorimeter energy deposits) and either merges these objects or halts once some criteria have been met.

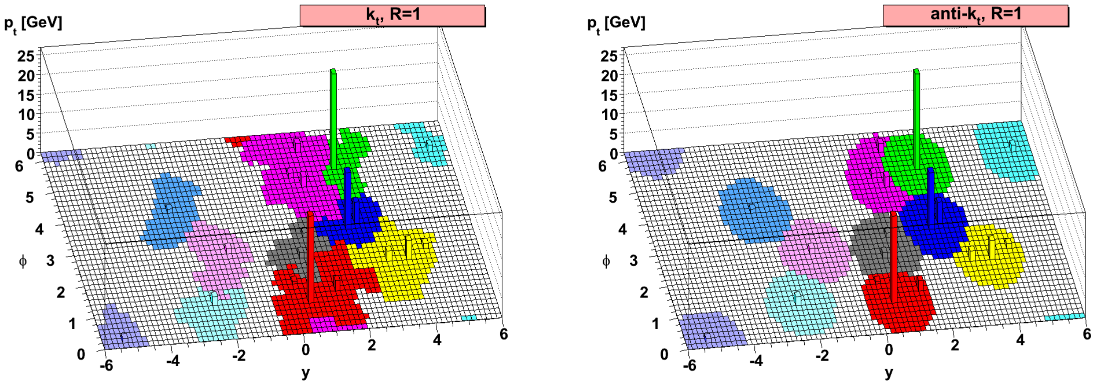

(a) k_t (b) anti-k_t

Figure 4.5: Comparison of k_t and anti-k_t clustering algorithms with size parameter $R = 1$ [52].

One important aspect of sequential recombination algorithms is whether or not they produce stable results (i.e. the same resulting jet objects) in the face of collinear splittings (collinear safety) or soft gluon emissions (infrared safety). Unlike other sequential recombination algorithms, such as k_T [53] or Cambridge-Aachen [54], the anti-k_t algorithm offers both collinear and infrared safety. It also produces more conical jets than other algorithms, as shown in Figure 4.5.

The anti-k_t algorithm takes a single parameter, R, which determines the size of reconstructed jets. For the $H \to \tau\tau$ analyses, $R = 0.4$ is used. The algorithm defines an inter-object metric:

$$d_{ij} = \min\left(\frac{1}{{p_T^i}^2}, \frac{1}{{p_T^j}^2}\right)\frac{\Delta R_{ij}^2}{R^2}, \tag{4.4}$$

as well as a single-object metric relative to the beam:

$$d_{iB} = \frac{1}{{p_T^i}^2}. \tag{4.5}$$

At each iteration of the algorithm, d_{ij} is computed for all pairs of objects, i and j, and d_{iB} is computed for each object. If the smallest of these values belongs to two objects, they are merged into a single object, while if these smallest of these values belongs to a single object, it is classified as a jet and removed from the list of objects considered by the algorithm. The algorithm stops when all objects have been classified as jets.

4.3.3 Energy calibration

After reconstruction, jets are subjected to further energy corrections. There are four primary steps: pile-up subtraction, an origin correction, simulation-derived energy corrections, and, in the case of calorimeter jets in data, an *in situ* calibration is applied based on momentum balance with some reference object.

4.3.3.1 Pile-up subtraction

Pile-up subtraction removes additional energy deposits due to additional proton-proton collisions in the same or nearby bunch crossings. This correction is computed using Monte Carlo simulations, and is parameterized in terms of

- The number of primary vertices in the event (a measure of in-time pile-up);
- The expected number of interactions per bunch crossing (a measure of out-of-time pile-up);
- η_{jet}; and
- p_T^{jet}.

4.3.3.2 Origin correction

The origin correction modifies the jet so that it points to the primary vertex from which it originates rather than the nominal interaction point. However, this only affects the direction of the jet, not the energy.

4.3.3.3 Simulation-derived energy correction

The first part of the correction to the jet energy is based on Monte Carlo simulation of calorimeter response in events that include the effects of pile-up. It is derived as a function of η_{jet} and jet energy at the local hadronic calibration energy scale. The simulated calorimeter response used to derive these corrections is shown in Figure 4.6.

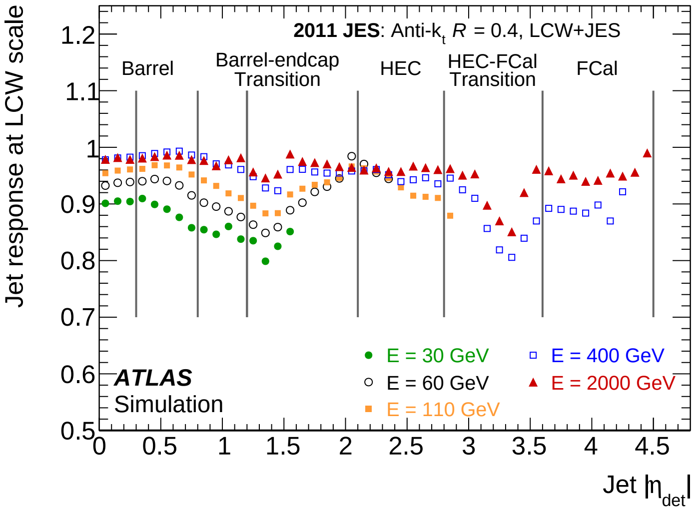

Figure 4.6: The simulated average calorimeter response for jets at various energies calibrated using the local hadronic calibration scheme for 7 TeV samples [48].

4.3.3.4 *In situ* calibration

The second part of the correction to the jet energy is performed using *in situ* techniques based on the p_T balance between a jet and a well-measured reference object (a photon or Z boson). For forward jets, the correction is derived using the p_T balance between two jets, one of which is a well-measured jet in the central part of the detector. The ratio of the average response in data to the average response in MC derived using these methods is shown in Figure 4.7.

4.3.4 Identification and selection

Fake jets may originate from a variety of sources, including protons colliding with gas in the beam pipe, scattering events from farther down the beam pipe, cosmic ray muons, and calorimeter noise. To reject fake jets from these backgrounds, identification criteria based on energy reconstruction quality, energy deposits in the direction of the shower, and reconstructed tracks are applied. From these criteria, four working points are defined: looser, loose, medium, and tight. The looser jet identification working point is used in the $H \to \tau\tau$ analyses. The efficiencies of these working

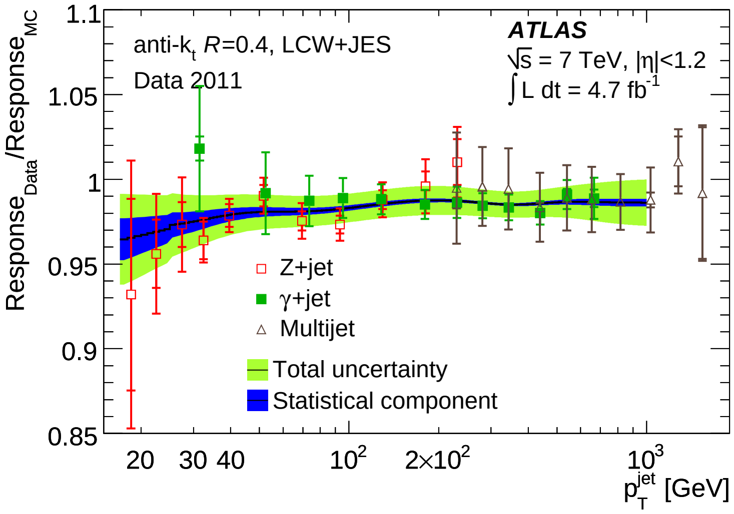

Figure 4.7: The ratio of average jet response in data to average jet response in Monte Carlo, both using local hadronic calibration, as a function of p_T^{jet} measured using various *in situ* techniques with 4.7 fb^{-1} of 7 TeV data [48].

points are measured using tag-and-probe studies and are shown for two jet η regions in Figure 4.8.

In the $H \to \tau\tau$ analyses, only jets with $|\eta| < 4.5$ are used. The VH analyses additionally require $p_T^{\text{jet}} > 30$ GeV. All $H \to \tau\tau$ analyses place a cut on the jet vertex fraction (JVF), defined by

$$\text{JVF} = \frac{\sum_{\text{TPV}} p_T^{\text{track}}}{\sum_{\text{TAV}} p_T^{\text{track}}}, \tag{4.6}$$

where TPV is the set of tracks associated with the jet originating from the primary vertex and TAV is the set of tracks associated with the jet originating from any vertex. Placing a lower bound on the JVF helps to reduce the effects of pile-up. For 7 TeV samples, jets with $|\eta| < 2.4$ are required to have $|\text{JVF}| > 0.75$. For 8 TeV samples, jets with $|\eta| < 2.4$ and $p_T < 50$ GeV are required to have $|\text{JVF}| > 0.5$.

4.3.5 *b*-tagging

In order to reduce background processes due to top quarks (which decay almost exclusively to Wb), a b-tagging technique is used to identify jets originating from b-quarks. The complete ATLAS b-tagging procedure is described in [55].

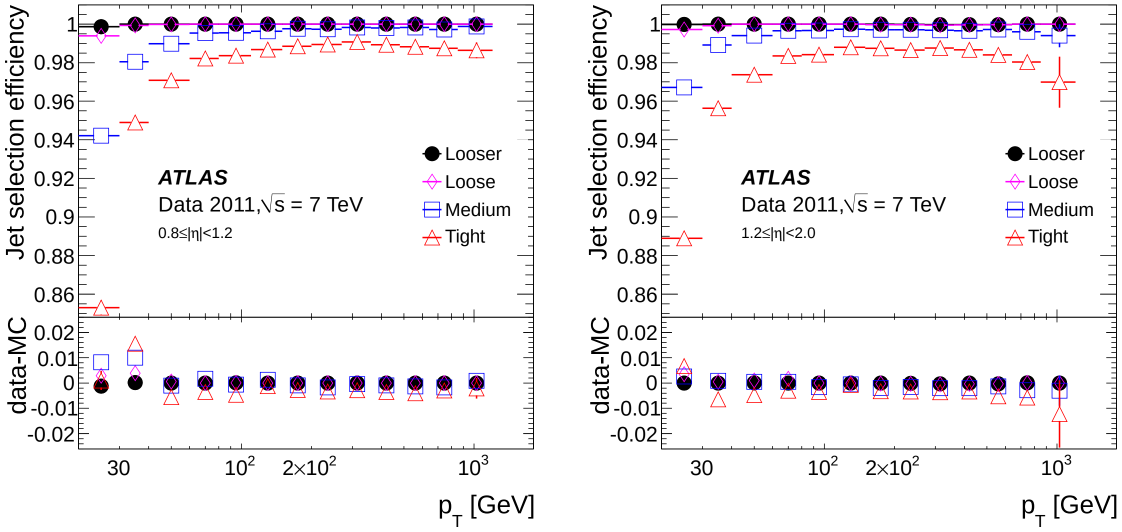

(a) $0.3 < |\eta| < 0.8$ (b) $1.2 < |\eta| < 2.0$

Figure 4.8: Jet identification efficiencies as a function of jet p_T for various η regions [48].

There are three algorithms, each utilizing different techniques, used for b-tagging in ATLAS:

- **IP3D**: Uses information on the impact parameter of the tracks in the jet;
- **SV1**: Attempts to reconstruct displaced secondary decay vertices characteristic of b-quark decays; and
- **JetFitter**: Uses information on the topology of weak b-quark decays.

These three algorithms are combined, using a neural network, to form a single discriminant known as the “MV1 score.” Because these algorithms require tracking information, b-tagging is only possible for jets in $|\eta| < 2.5$, where coverage is provided by the Inner Detector.

In the main analyses, a working point is chosen for which the b-tagging algorithms are 70% efficient on average for jets with $p_T > 15$ GeV in simulated $t\bar{t}$ events. The b-tagging efficiency as a function of p_T^{jet} at this working point is shown in Figure 4.9.

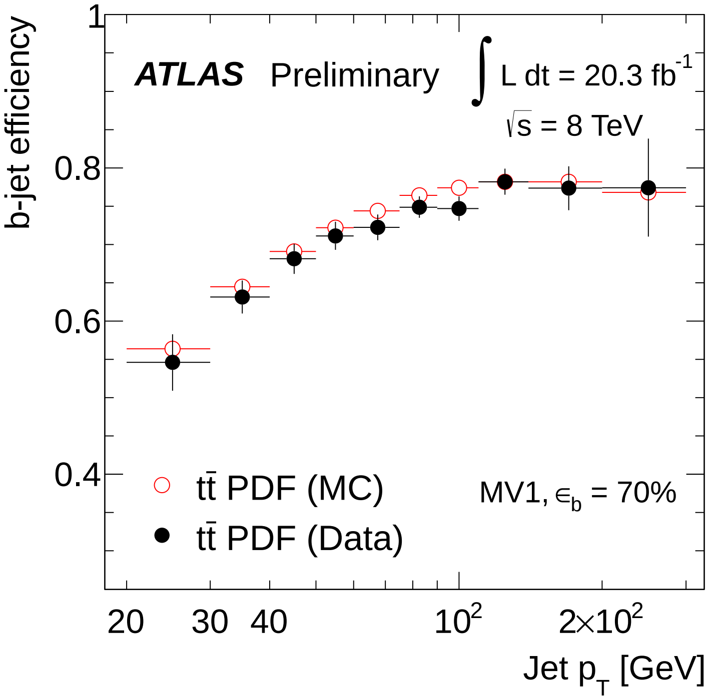

Figure 4.9: b-tagging efficiency in data compared with that in Monte Carlo for 20.3 fb^{-1} of 8 TeV data at the 70% efficiency working point [56].

The probability of misidentifying a jet originating from a light quark as originating from a b-quark has been shown to be 0.1–0.5%, depending on the p_T and η of the jet [57], as shown in Figure 4.10.

4.4 Hadronically decaying τ leptons

Hadronically decaying τ leptons are perhaps the most important physics objects under consideration in the $H \to \tau\tau$ analyses. Hadronic τ decays encompass $\sim 65\%$ of all τ lepton decays. These hadronic decays are characterized by the presence of one charged pion $\sim 72\%$ of the time, three charged pions $\sim 22\%$ of the time, and charged kaons in the remainder of cases. Neutral pions may also be present. Because such a large fraction of τ_h decays are characterized by one or three charged tracks, the $H \to \tau\tau$ analyses focus exclusively on "one-prong" and "three-prong" τ_h decays.

For most of the $H \to \tau\tau$ analyses, hadronic τ objects are used for triggering, event

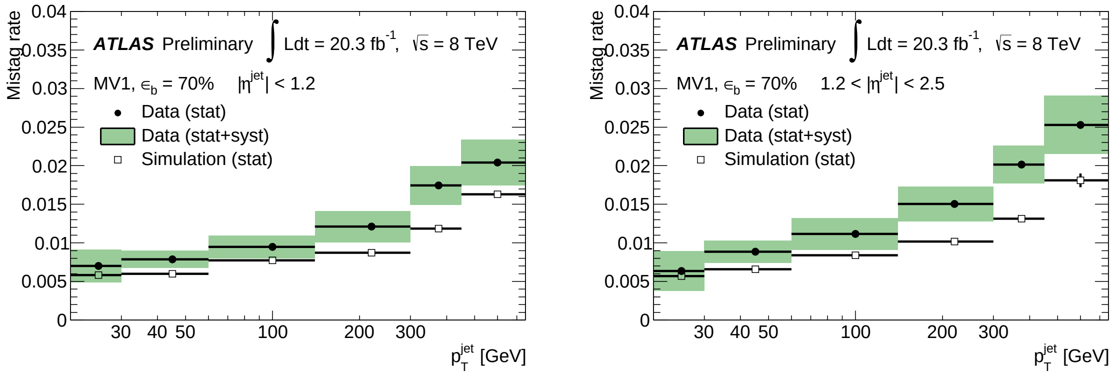

(a) $|\eta| < 1.2$ (b) $1.2 < |\eta| < 2.5$

Figure 4.10: Misidentification rates for light quarks incorrectly identified as b-quark jets as a function of jet p_T in different η regions for 20.3 fb^{-1} of 8 TeV data and corresponding Monte Carlo samples at the 70% efficiency working point [57].

selection, or both. This includes $H \to \tau_h\tau_h$, $H \to \tau_l\tau_h$, $Z_{ll}H \to \tau_h\tau_h$, $Z_{ll}H \to \tau_l\tau_h$, $W_{l\nu}H \to \tau_h\tau_h$, and $W_{l\nu}H \to \tau_l\tau_h$. Hadronic τ triggering, reconstruction, and energy calibration for the ATLAS detector in Run I are described fully in [58].

4.4.1 Triggering

The hadronic τ trigger plays an important role in all analyses with a τ_h signature in the final state. A variety of hadronic τ triggers were developed for Run I data taking, including $\tau_h + E_T^{\mathrm{miss}}$ triggers, double τ_h triggers, and combined τ_h + light lepton triggers.

The L1 trigger identifies τ_h candidates as 2×2 clusters of trigger towers, where each tower has size 0.1×0.1 in $\eta \times \phi$ space and extends across both the electromagnetic and hadronic calorimeters. An additional isolation region is defined as the towers between this core region and a larger 4×4 window. The E_T of the core region must be above a given threshold and the E_T of the isolation region must be less than 4 GeV for an event to continue to the L2 trigger.

The L2 trigger is seeded with τ_h candidates by the L1 trigger. In addition to calorimeter information, tracking information from the pixel and SCT detectors is used to further refine selections. To suppress the effects of pile-up in the L2 trigger, where vertex reconstruction is not available, only tracks with $|\Delta z_0| < 2$ mm and $\Delta R < 0.1$ with respect to the highest-p_T track inside the τ_h RoI are considered. Further criteria were chosen to reject multijet backgrounds based on combined tracking and calorimeter variables. These variables were carefully chosen to perform consistently across the varying pile-up conditions experienced in 8 TeV data taking. Variables that relied heavily on the electromagnetic calorimeter, such as those used in 7 TeV data taking, showed poor stability with respect to the number of primary vertices, as illustrated in Figure 4.11. A study was performed to identify variables that provided strong discrimination in addition to robustness against pile-up. These variables include:

- $f_{\text{iso}}^{\text{L2}}$: The ratio of the scalar sum of the p_T of tracks within the isolation region around the τ_h candidate, $0.1 < \Delta R < 0.3$, to that of tracks in the core region, $\Delta R < 0.1$;
- $f_{\text{cent}}^{\text{L2}}$: The ratio of E_T in a cone of radius $R = 0.1$ around the τ_h candidate to that in a cone of radius $R = 0.4$;
- $f_{\text{track}}^{\text{L2}}$: The ratio of the p_T of the leading track in a cone of radius $R = 0.3$ around the τ_h candidate to the total E_T in a cone of the same size around the τ_h candidate; and
- $R_{\text{track}}^{\text{L2}}$: The p_T-weighted ΔR of τ_h candidate tracks in a cone of radius $R = 0.3$ around the τ_h candidate.

The EF implementation for the τ_h trigger is largely identical to the offline reconstruction and identification procedures described in Sections 4.4.2 and 4.4.4. The EF does not have full access to vertex reconstruction, and thus a technique similar to the L2 $|\Delta z_0|$ cut is used, along with a cut on $|\Delta d_0|$. In addition, a different set of variables is used for the boosted decision tree discriminant.

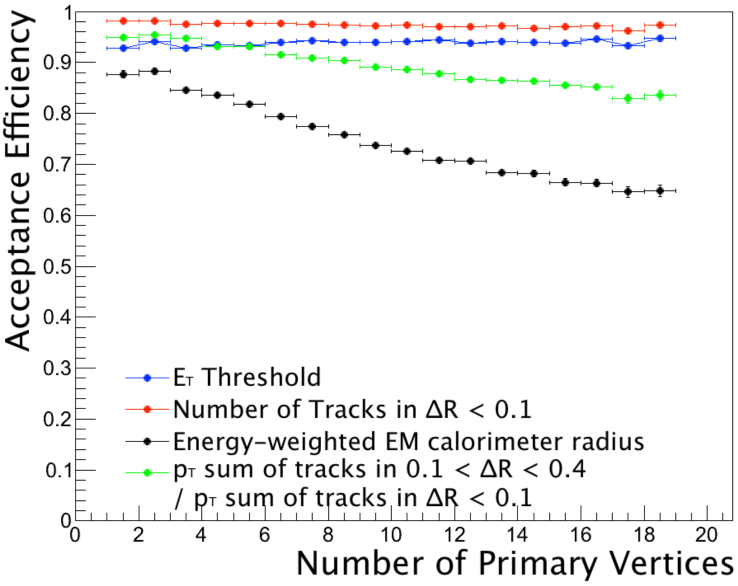

Figure 4.11: The efficiencies of 2011 L2 trigger selections for simulated 8 TeV $Z \to \tau\tau$ events as a function of the number of primary vertices. Significant performance degradation was observed for variables based on the electromagnetic calorimeter due to increased pile-up.

The efficiency of the τ_h trigger is measured using tag-and-probe methods in $Z \to \tau\tau$ samples, with additional background modeling and subtraction necessary due to the intrinsic impurity of the $Z \to \tau\tau$ sample. Efficiencies for all three levels with respect to offline identification as a function of $p_T^{\tau_h}$ and the number of primary vertices in the event are shown in Figure 4.12.

4.4.2 Reconstruction

Offline τ_h reconstruction is seeded by all jets reconstructed with the anti-k_t algorithm with a distance parameter $R = 0.4$. Two additional requirements are applied: $p_T^{\text{jet}} > 10$ GeV and $|\eta_{\text{jet}}| < 2.5$, the latter of which is necessary to ensure tracking coverage. Once τ_h seeds have been identified, reconstruction proceeds in three steps: τ vertex association, determination of τ_h four-momentum, and track association.

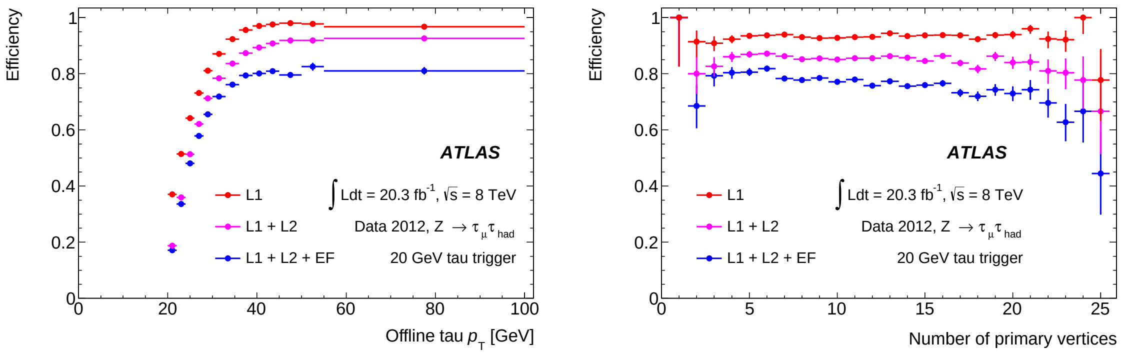

(a) (b)

Figure 4.12: The efficiency for each level of the τ_h trigger with respect to the offline identification efficiency as a function of (a) $p_T^{\tau_h}$ and (b) the number of primary vertices in the event measured using $Z \to \tau\tau$ tag-and-probe methods with 20.3 fb^{-1} of 8 TeV data [58].

4.4.2.1 τ vertex association

The τ decay vertex, which may be displaced from the primary vertex in the event, must be identified in order to reduce the effects of pile-up and correctly compute reconstruction quantities. All reconstructed vertices in the event are considered candidates for the τ vertex. For each vertex, a cone of radius $R = 0.2$ is constructed from the vertex in the direction of the τ_h candidate. For each vertex cone, the ratio of the scalar p_T sum of tracks within the cone associated with the corresponding vertex to that of tracks within the cone associated to any vertex is computed. The vertex with the largest such ratio is chosen as the τ vertex. Only tracks with $p_T > 1$ GeV are considered.

4.4.2.2 Determination of τ_h four-momentum

After the τ vertex has been identified, the four-momentum of the τ_h candidate is then determined by recalculating the η and ϕ directions of the jet barycenter in the τ vertex coordinate system. The four-momentum of clusters within $\Delta R < 0.2$ of the recomputed barycenter are then summed to compute the four-momentum of the τ_h candidate. At this stage, the topo-clusters constituting the jet and the resultant τ_h

candidate four-momentum are calibrated using only local hadronic calibration.

4.4.2.3 Track association

Tracks are associated with the τ_h candidate if they are within $\Delta R < 0.2$ of the τ_h candidate direction and satisfy the following criteria:

- $p_T > 1$ GeV;
- Two hits in the pixel layer;
- Seven hits in the pixel and SCT layers combined;
- $|\Delta d_0| < 1.0$ mm with respect to the τ vertex; and
- $|\Delta z_0 \sin(\theta)| < 1.5$ mm with respect to the τ vertex.

4.4.3 Energy calibration

Additional offline corrections are applied to adjust the reconstructed τ_h candidate energy from local hadronic calibration scale to the final τ_h energy scale. These corrections are derived from $W \to \tau\nu$ and $Z \to \tau\tau$ Monte Carlo simulation. The effects of these corrections are shown in Figure 4.13. A small additional correction is also made for τ_h candidates in the transition region between the barrel and end-cap electromagnetic calorimeters, along with an η-dependent pile-up subtraction for all τ_h candidates.

4.4.4 Identification and selection

The τ_h reconstruction algorithms do little to reject multijet backgrounds. Consequently, a boosted decision tree (BDT) is used to discriminate between τ_h objects and normal jets. The operation of BDTs is described in full in Section 6.5.4. There are separate BDTs trained for one-prong and three-prong τ_h candidates. The list of

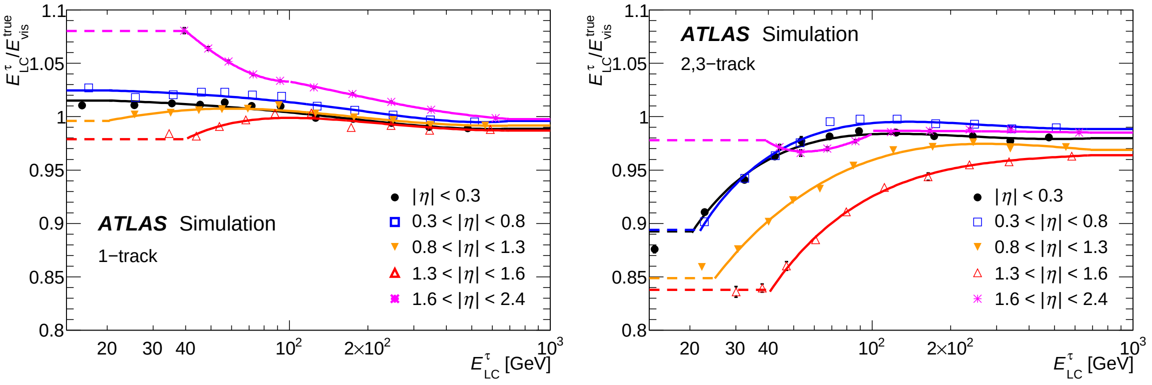

(a) One-prong τ_h candidates (b) Two- and three-prong τ_h candidates

Figure 4.13: The ratio of τ_h energy for candidates using local hadronic calibration to their true visible energy as a function of $E_T^{\tau_h}$ at the local hadronic calibration scale, calculated using simulated samples [58].

variables used in these BDTs, as well as those at EF, is given in Table 4.1. They include:

- f_{cent}: The ratio of E_T in a cone of radius $R = 0.1$ around the τ_h candidate to that in a cone of radius $R = 0.2$;
- f_{track}: The ratio of the p_T of the leading track in a cone of radius $R = 0.2$ around the τ_h candidate to the total E_T in a cone of the same size around the τ_h candidate;
- $R^{\mathrm{L2}}_{\mathrm{track}}$: The p_T-weighted ΔR of τ_h candidate tracks in a cone of radius $R = 0.4$ around the τ_h candidate;
- $S_{\mathrm{lead\ track}}$: The transverse impact parameter of the leading track in a cone of radius $R = 0.2$ around the τ_h direction divided by its corresponding uncertainty;
- $N^{\mathrm{iso}}_{\mathrm{track}}$: The number of tracks in the isolation region around the τ_h candidate, $0.2 < \Delta R < 0.4$;
- ΔR_{max}: The maximum distance between a track in the core region around the τ_h candidate, $\Delta R < 0.2$, and the τ_h direction;

Variable	Event Filter		Offline	
	1-prong	3-prong	1-prong	3-prong
f_{cent}	●	●	●	●
f_{track}	●	●	●	●
$R_{\text{track}}^{\text{L2}}$	●	●	●	●
$S_{\text{lead track}}$	●		●	
$N_{\text{track}}^{\text{iso}}$	●		●	
ΔR_{max}		●		●
$S_{\text{T}}^{\text{flight}}$		●		●
m_{track}		●		●
$m_{\pi^0+\text{track}}$			●	●
N_{π^0}			●	●
$p_T^{\pi^0+\text{track}}/E_T$			●	●

Table 4.1: Boosted decision tree input variables for the Event Filter and the offline hadronic τ identification algorithms for 1-prong and 3-prong τ_h candidates [58]. Definitions are given in Section 4.4.4.

- $S_{\text{T}}^{\text{flight}}$: The displacement in the transverse plane of the secondary decay vertex, divided by its uncertainty;
- m_{track}: The invariant mass of the four-momenta of all tracks within a cone of radius $R = 0.4$ around the τ_h direction;
- $m_{\pi^0+\text{track}}$: The invariant mass of tracks and neutral pions within a cone of radius $R = 0.2$ around the τ_h direction;
- N_{π^0}: The number of neutral pions reconstructed within a cone of radius $R = 0.2$ around the τ_h direction; and
- $p_T^{\pi^0+\text{track}}/E_T$: The ratio of p_T estimated using track and π^0 information to the calorimeter E_T measurement.

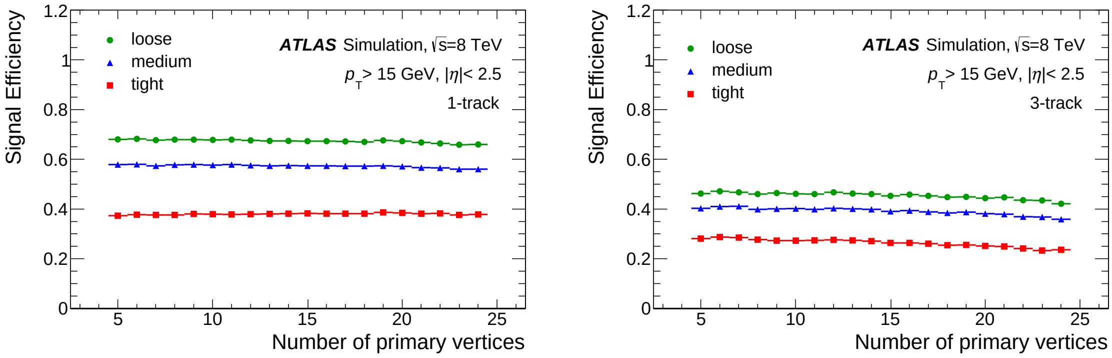

(a) One-prong τ_h

(b) Three-prong τ_h

Figure 4.14: The efficiency for each offline τ_h reconstruction working point as a function of the number of primary vertices in events measured using $Z \to \tau\tau$ tag-and-probe methods with 20.3 fb^{-1} of 8 TeV data [58].

Three working points are defined based on the resultant BDT score: loose, medium, and tight. The efficiency for each working point, measured in $Z \to \tau\tau$ tag-and-probe studies, can be seen in Figure 4.14. Additional BDTs are also implemented to reject misidentified electrons and muons. In the $H \to \tau\tau$ analyses, medium τ_h objects are used. It is required that they have $p_T > 20$ GeV and $|\eta| < 2.47$. In the VH analyses, the requirement is relaxed to $|\eta| < 2.5$ for three-prong taus.

4.5 E_T^{miss} reconstruction

The presence of E_T^{miss} in $H \to \tau\tau$ final states is expected due to the presence of one or more neutrinos per τ decay. The calculation of E_T^{miss} in ATLAS is described in full in [59].

E_T^{miss} reconstruction is performed by summing the E_T of calorimeter cells, calibrated according to the reconstructed objects with which they are associated. The precedence for cell association is:

1. Electrons
2. Photons

3. Hadronically decaying τ leptons

4. Jets

5. Muons

Cells which are not associated to any object are also considered, and their contribution is referred to as "CellOut." The formula for E_T^{miss} is then given by

$$E_{T_{x,y}}^{\text{miss}} = \sum_{o \in \text{objects}} E_{T_{x,y}}^{\text{miss},o}, \tag{4.7}$$

where objects include electrons, photons, hadronically decaying τ leptons, jets, muon calorimeter deposits, and CellOut cells. This summation is performed over the range $|\eta| < 4.5$. The vector sum of the muon transverse momenta is also added to this sum.

4.6 Overlap removal

In some cases, multiple objects may be reconstructed and identified using the aforementioned techniques while corresponding to the same physical object in the detector. If this happens, an "overlap removal" scheme is used to select objects based on the relative likelihood of their corresponding reconstruction and identification being correct. Objects are treated as overlapping if they lie within $\Delta R < 0.2$ of one another. Once one of the overlapping objects is selected for use in the analysis, the remaining overlapping objects are discarded. In the $H \to \tau\tau$ analyses, objects are selected with the following precedence:

1. Muons

2. Electrons

3. Hadronically decaying τ leptons

4. Jets

Several looser object requirements are used to achieve a more stringent overlap removal. In the $H \rightarrow \tau_l \tau_h$ analysis, electrons are only required to pass the loose identification threshold to count toward overlap removal (the requirement for electrons in the calorimeter crack region remains at the medium threshold due to poor electron identification in this region). Additionally, muons are only required to have $p_T > 2$ GeV in order to veto τ_h objects. In the VH analyses, muons are required to have $p_T > 4$ GeV to count toward overlap removal of any object.

Chapter 5

Computing techniques

Several factors motivated the study and implementation of advanced computing techniques to facilitate a more optimized and efficient data recording and analysis pipeline.

Firstly, the demanding operating conditions of the ATLAS trigger system place considerable constraints on the execution time of the HLT software. To mitigate this issue, a feasibility study was performed to assess the potential of Graphics Processing Units (GPUs) to accelerate data preparation algorithms in the ATLAS HLT. This research was done in the context of a larger study aimed at accelerating the entire HLT data preparation and tracking pipeline using GPUs shared between multiple HLT processes.

Secondly, the large data volume and complex analysis requirements of the $H \to \tau\tau$ analyses motivated studies into the optimization of analysis quantity computation. To increase the performance of the analysis code, a framework was developed to evaluate expressions on data and Monte Carlo samples using a just-in-time (JIT) compiler.

This chapter provides a summary of these studies. In particular, the theoretical motivation, implementation, and results are discussed for each.

Level	Type	Input rate	Output rate	Execution time
1	Hardware	20 MHz	70 kHz	2.5 μs
2	Software	70 kHz	6.5 kHz	90 ms
Event Filter	Software	6.5 kHz	600 Hz	1 s

Table 5.1: Peak trigger processing rates and execution times for ATLAS Run I.

5.1 Optimization of HLT data preparation with graphics processing units

The ATLAS trigger system is one of the most computationally demanding steps in the entire analysis pipeline. The peak trigger processing rates and corresponding execution times for ATLAS Run I are summarized in Table 5.1.

The execution time for trigger algorithms is highly constrained by these conditions. While parallelism and round-robin load-balancing in the trigger server farm do allow for an execution time longer than the event period, limitations on cooling in the ATLAS cavern combined with budgetary constraints make unlimited horizontal scaling neither technologically nor economically viable.

The fractions of time spent in various trigger processing steps are given in Table 5.2. Data preparation and tracking are, at present, two of the most time consuming steps in the entire trigger processing chain. Consequently, these algorithms are ideal candidates for optimization.

Although the current trigger implementation is sufficient, future operating conditions will require significantly enhanced performance and throughput. A comparison

Processing Step	Level 2 time fraction	Event Filter time fraction
Data preparation and tracking	73%	52%
Calorimeter	17%	24%
Muon	7%	16%
Other	3%	8%

Table 5.2: Processing time distributions for various HLT steps in ATLAS Run I.

Year or Machine	Peak instantaneous luminosity
2010	2.1×10^{32} cm^{-2}s^{-1}
2011	3.65×10^{33} cm^{-2}s^{-1}
2012	7.73×10^{33} cm^{-2}s^{-1}
2015	10^{34} cm^{-2}s^{-1}
High Luminosity LHC	10^{35} cm^{-2}s^{-1}

Table 5.3: Past and projected LHC peak instantaneous luminosities.

of past and projected LHC instantaneous beam luminosities is given in Table 5.3. As instantaneous luminosity increases, the collision rate and occupancy of hits within Inner Detector modules will increase proportionally. This increase will lead to unacceptable performance of the existing trigger algorithms, many of which exhibit $\mathcal{O}(n^2)$ behavior or worse.

In this section, a feasibility study is presented on the usage of GPUs to accelerate trigger data preparation and tracking algorithms. An overview of the existing algorithms is given, along with a brief discussion on general purpose computing on GPUs. The resultant GPU-ported algorithms and performance are presented, along with discussion of integration complexities and considerations.

5.1.1 Data preparation

The Level 2 and Event Filter triggers utilize tracking data from the Pixel Detector and SCT in their reconstruction algorithms. This data must be acquired, decoded, and processed into a form suitable for tracking algorithms to use. This process is referred to as "data preparation," and is the first step performed in HLT processing.

Data preparation technically begins in the Level 1 trigger, where solid-angle Regions of Interest (RoIs) are identified by the Level 1 trigger using calorimeter and muon detector data. After the RoIs are verified by the Level 2 trigger, the data for an event's RoIs are copied over a network connection from a set of transient storage locations known as Readout Buffers (ROBs). The ROBs transmit data in fragments of a highly compressed bytestream format, which must first be decoded into "hits"

Module 1 Header
Module 1 Hit 1
...
Module 1 Hit M_1
Module 1 Trailer
...
Module N Header
Module N Hit 1
...
Module N Hit M_N
Module N Trailer

Figure 5.1: Bytestream format for a single ROB fragment [60].

representing particle transits through the pixel and SCT modules. These particle transits in pixel and SCT modules typically activate multiple adjacent silicon cells. Consequently, adjacent cell activations must be clustered to identify actual particle transits.

5.1.1.1 Bytestream decoding

The first step in data preparation, bytestream decoding, involves the transformation of the compact transport representation of pixel and SCT data into data structures suitable for clustering. An outline of the bytestream format is given in Figure 5.1. The bytestream consists of a series of words, which are 16-bit in the case of the SCT detector and 32-bit in the case of the Pixel Detector. The words fall into three categories: headers, hits, and trailers. Headers encode module identifiers, hits encode activated silicon cells within the modules, and trailers serve as delimiters to separate module data.

The existing decoding algorithm is given in Algorithm 1. The algorithm is entirely serial, although it does present immediate potential for parallelization at the RoI and ROB fragment levels. Further parallelization at the word level is complicated by a parsing state that is mutated during iteration. The algorithm exhibits $\mathcal{O}(n)$ behavior, where n is the number of input bytestream words.

```
Input: Bytestream data by RoI
Output: Collections of hits organized by module
foreach RoI r in input do
    foreach ROBFragment f in r do
        InitializeROBFragmentParsingState();
        foreach Word w in f do
            type = GetWordType(w);
            if type == header then
                // Mutates ROB fragment parsing state to set hit container
                MarkModuleBegin(w);
            end
            else if type == hit then
                // Mutates hit container
                RecordModuleHit(w);
            end
            else if type == trailer then
                // Mutates ROB fragment parsing state to unset hit container
                MarkModuleEnd(w);
            end
        end
    end
end
```

Algorithm 1: Pixel and SCT bytestream decoding algorithm.

5.1.1.2 Clustering

Once module hits have been identified, they must be combined into clusters representing particle transits in the modules. This is trivial in the case of the SCT modules, where silicon cells are arranged one-dimensionally and adjacency can be determined by a simple loop. The pixel modules present a more complex scenario, requiring a nested loop over hits and existing clusters.

The existing clustering algorithm for pixel modules is given in Algorithm 2. The algorithm is entirely serial, and does not lend itself to trivial parallelization. The clustering algorithms exhibit $\mathcal{O}(n^2)$ behavior for the pixel detector and $\mathcal{O}(n)$ behavior for the SCT detector, where n is the number of activated silicon cells in the detector's modules.

After clustering is complete, the resultant clusters are translated and rotated

```
Input: Collections of hits organized by module
Output: Collections of clusters
foreach Module m in input do
    clusters = [];
    foreach Hit h in m do
        cluster = none;
        foreach Cluster c in clusters do
            foreach Hit e in c do
                if Adjacent(h, e) then
                    if cluster != none then
                        Merge(cluster, c);
                    end
                    else
                        c.add(h);
                        cluster = c;
                    end
                end
            end
        end
        if cluster == none then
            new_cluster = [h];
            clusters.add(new_cluster);
        end
    end
    foreach Cluster c in clusters do
        ComputeCenter(c);
    end
end
```

Algorithm 2: Pixel module hit clustering algorithm.

to match the module's position and orientation within the detector. This process generates "spacepoints" that can be used for track reconstruction.

5.1.2 Tracking

The spacepoints identified by the data preparation algorithms are used by the trigger tracking algorithms to reconstruct particle tracks in the Inner Detector. There are two algorithms used by the ATLAS trigger for tracking: IDScan and SiTrack [61]. For the purposes of this study, the SiTrack algorithm was chosen due to its combinatoric

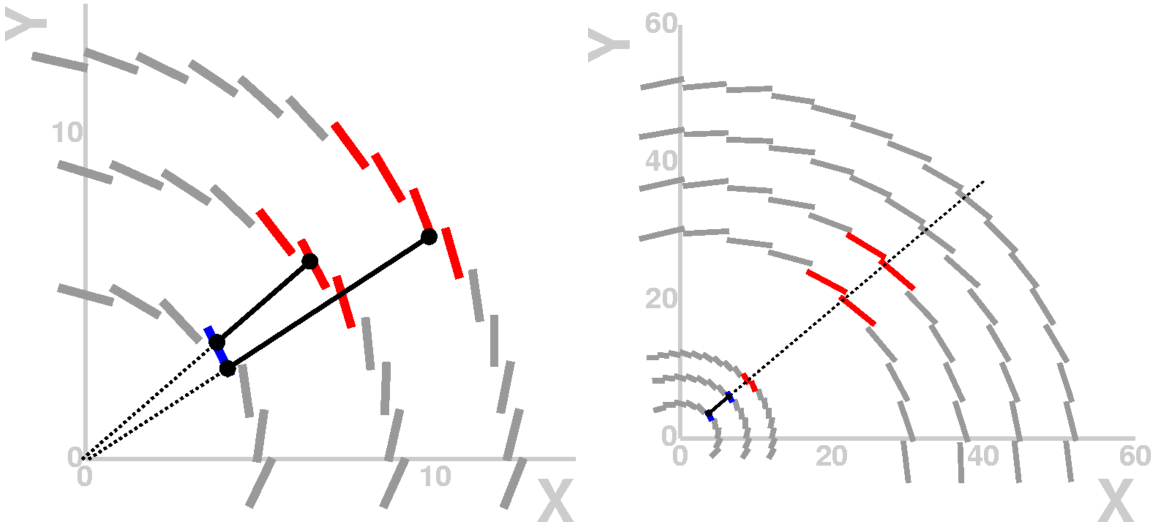

(a) Track seed formation

(b) Track seed extension

Figure 5.2: SiTrack track seed formation and extension procedures [33].

approach to track formation and more immediate potential for parallelization.

5.1.2.1 Track seed formation and extension

The first step in the tracking process consists of seed formation and extension. Track seeds are formed by combinatorially generating pairs of spacepoints from the inner layers of the Inner Detector. A cut is placed on the transverse impact parameter of the straight line joining a seed's two spacepoints; this reduces the number of seeds to be processed, but places a limit on the lowest reconstructable track p_T [33]. This process is illustrated in Figure 5.2a.

After track seeds have been generated, they are extended into the outer layers of the Inner Detector to form triplets of spacepoints. This process is illustrated in Figure 5.2b.

5.1.2.2 Track merging and clone removal

In order to form the final particle tracks used in the trigger, spacepoint triplets with the same trajectories are merged to form track candidates. It is often the case that the resultant track candidates will share the same outer layer spacepoints but different seed layer spacepoints. Such groups of tracks are referred to as “clones.” In order to remove clones, a de-duplication algorithm iterates over track candidates and merges clone tracks into a single final track, removing the additional clone tracks.

5.1.3 General purpose computing on graphics processing units

General purpose computing on graphics processing units (GPGPU computing) is a relatively modern technique, allowing programs to leverage the hundreds or thousands of cores present on modern GPUs to perform Single Instruction Multiple Thread (SIMT) computations in a massively parallel fashion.

Modern GPGPU computing is facilitated by software development kits (SDKs) such as CUDA and OpenCL. While the precise implementation details differ, the architecture of GPGPU computing remains largely isomorphic across SDKs. Computation functions, called “kernels,” are written in a restricted C or C++ dialect and compiled for the architecture of the target computing device. Data to be processed are then copied from host memory (RAM associated with the CPU) to device memory (memory associated with the GPU). Kernels are then run in parallel on hundreds or thousands of threads, all efficiently managed by dedicated thread hardware on the GPU. These threads are organized into multidimensional abstractions called “blocks,” where each thread (and its associated kernel) can access information about its position within the block structure. Thread position information can be used to map an instance of the kernel onto corresponding data from the input array(s). Multiple kernels may be executed to sequentially transform data, without the data ever leaving GPU memory. When kernel execution is complete, the resultant data is copied back to host memory.

Despite the simplicity of the programming model, a large number of complications arise due to the specialized nature of GPU architectures.

Firstly, GPU memory is organized into a complex hierarchy of constant, global, block-local, and thread-local memory, each with different access performance, paradigms, and penalties. For memory access to be efficient, complex and often device-dependent restrictions are placed on read and write sizes, address boundaries, and the order of the requesting threads. While newer architectures reduce the requirements for so-called "coalesced" memory access, the desire to support a heterogeneous array of device types and architectures makes this perhaps the most difficult problem with GPGPU computing.

Secondly, unlike a CPU, where separate threads of execution can execute different instructions in parallel on different cores, GPU thread hardware requires all threads to execute the same instruction. If branching occurs, execution is serialized, with threads paused and executed sequentially until their instructions realign. This presents a problem for code with significant branching logic (e.g. `if/else` statements, `for/while` loops, etc). As such, one must be careful to minimize branching in code ported from the CPU.

Thirdly, the restrictions placed on the dialect of C or C++ used to program kernels often include a lack of scalar types, a standard library, or support for dynamic memory allocation, adding further complication to the porting of CPU-oriented code.

Finally, performance optimizations are typically device or architecture specific, and require intimate knowledge of the underlying computing hardware. While compilers for GPGPU programming are becoming increasingly intelligent, the "close to the metal" approach required for efficient GPGPU performance is often a limiting factor for code maintainability.

5.1.4 GPU-accelerated algorithms

For the purposes of this study, the NVIDIA CUDA platform was chosen due to its performance, features, support, hardware availability, as well as its previous use in ATLAS [62, 63]. A subsequent study was performed using OpenCL to test a more heterogeneous array of hardware, the results of which are also presented. Both implementations employ the same underlying parallel algorithms, with small modifications made for the particular language dialect.

5.1.4.1 Bytestream decoding

The GPU-ported version of the bytestream decoding algorithm is largely based on the original CPU version. On the GPU, however, different bytestream fragments are mapped to distinct thread blocks, and the processing of each word within the bytestream fragment is mapped to a separate thread, allowing for simultaneous decoding of all words within the fragment.

The positional context dependence of words within a bytestream fragment is mitigated through the use of fast, block-local shared memory, in which threads can rapidly scan neighboring words to determine their associated module.

A persistent global output buffer is created to store the decoded hits so that they may be used by the subsequent clustering routine. Atomic incrementation routines are used to ensure coordinated access to this array by threads.

In practice, limitations on thread block size require threads to be re-used across multiple words. This places the complexity of the parallelized algorithm at $\mathcal{O}(\frac{n}{m})$, where n is the number of activated silicon cells in the detector's modules and m is the number of threads available on the GPU. With advancements in hardware allowing for larger numbers of threads, this should place the algorithm closer to constant-time performance.

5.1.4.2 Clustering

The GPU version of the clustering algorithm is not a port of the CPU version. Instead, it is based on a cellular automaton originally designed to cluster energy deposits in the LHCb calorimeter [64].

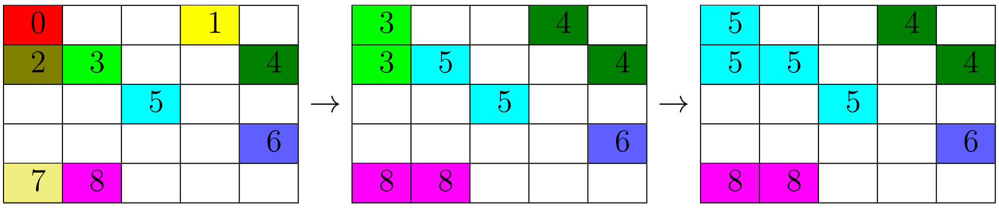

Figure 5.3: Evolution of cellular automaton clustering for a two-dimensional pixel array [65].

The algorithm operates on the output of the bytestream decoding algorithm, copying hits from the global array to block-local memory. The cellular automaton then tags all initial hits with a unique identifier (in this case an integer corresponding to the thread index). At each evolutionary step, the tag of a hit in the automaton is set to the largest tag of any adjacent hit (which is rapidly computable due to all hits of a given module being in block-local memory). Clustering is complete when the automaton stops evolving, at which point clusters can be easily extracted by a trivial iteration over the tagged hits. An illustration of the process is given in Figure 5.3.

5.1.4.3 Track seed formation and extension

The GPU track seed formation and extension algorithms operate in a fashion similar to their CPU counterparts. The primary difference is that 16 × 16 or 32 × 32 thread blocks are used to efficiently evaluate the Cartesian product of spacepoints from different layers of the detector. Intermediate global arrays are used to store the intermediate track seeds and the track triplets. A one-dimensional thread block is then used to merge triplets to create track candidates. The algorithms are illustrated in Figure 5.4.

5.1.4.4 Track merging and clone removal

The GPU version of track merging and clone removal is made difficult by the large number of track pairs, $\frac{N\times(N-1)}{2}$, which typically do not fit into the faster memory locations on the GPU. Consequently, the algorithm divides the process into two steps.

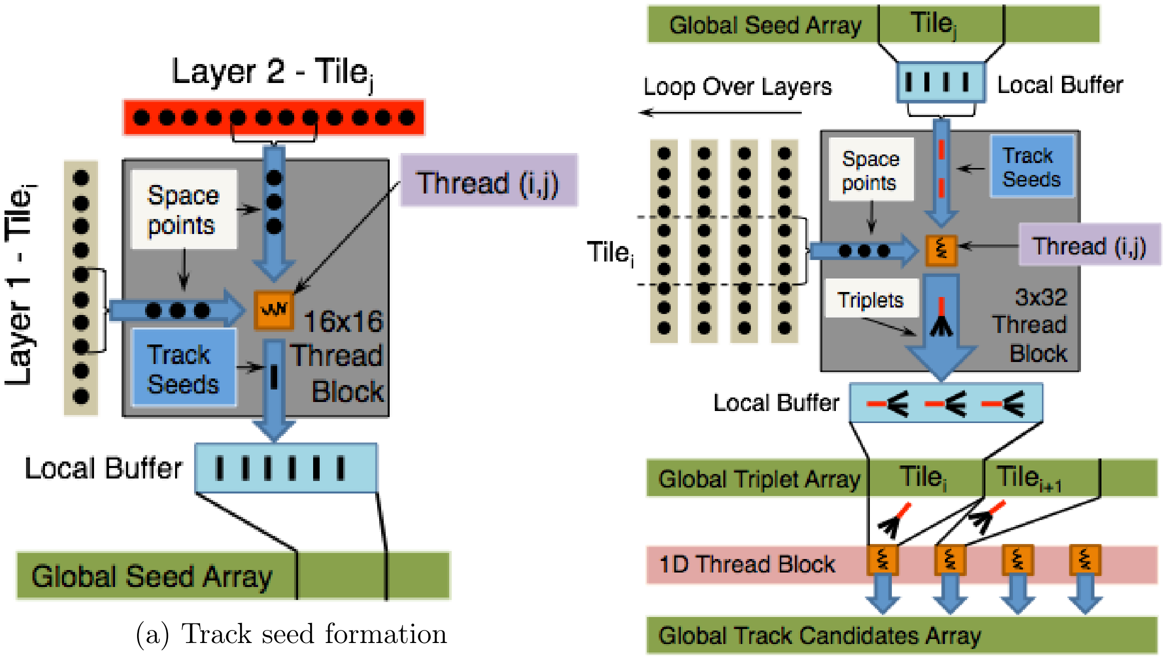

(a) Track seed formation

(b) Track seed extension

Figure 5.4: Graphical representation of GPU track candidate formation [65].

First, tracks with the same extension spacepoints but different seeds are identified and merged. Second, duplicate tracks are removed. The approach is one of brute force, with each thread iterating through a range of input tracks. This iteration is performed on global memory, which is very slow. Nevertheless, the algorithm shows performance gains over the CPU due to the large number of candidate tracks.

5.1.5 Porting complexities and considerations

A number of complexities arose in the process of porting trigger algorithms to the GPU, due to both the intrinsic nature of GPGPU programming, as well as factors in the original CPU code development.

The primary roadblock to GPU porting was the state of the existing CPU implementation. The existing code was inherently serial, with a complex “spaghetti-code” structure and minimal documentation and/or comments. In addition, a complex C++ class hierarchy prevented a more direct port of the existing code. The solution to these problems came in the form of dedicated manpower and reverse engineering. Approximately one year of full-time development work, divided amongst multiple GPU-experienced developers, was required to port these four algorithms.

Another complication arose due to the limited on-device memory available to the GPU. While the process of hit word decoding and cluster-to-spacepoint conversions are computationally trivial, they do require access to a hardware map of the detector. This map encodes information about the silicon cell layout within modules, as well as the position and orientation of modules within the detector. Within the ATLAS trigger software framework, this hardware map occupies several gigabytes of memory — far too much to store directly on the GPU. The map also uses several custom C++ data structures to encode module information and facilitate module lookups. To solve this problem, a slimmed version of the hardware map was developed, using only those C++ features available to CUDA. To facilitate efficient lookup of modules, a perfect hash function was generated to efficiently map module identifiers to their corresponding location in memory.

5.1.6 Integration with trigger software framework

Beyond the complexities involved with porting the CPU code, several complications arose with the integration of the GPU-ported trigger algorithms and the existing trigger software framework. First, the trigger software framework uses a custom build system, making it difficult to integrate the CUDA compiler and runtime libraries. Second, multiple trigger instances are typically run on each node in the trigger server farm, which necessitated a GPU resource sharing paradigm.

The solution to both of these problems was the development of a specialized client-server architecture allowing for Remote Procedure Call (RPC) invocation of GPU kernels. An illustration of this architecture is given in Figure 5.5. In this architecture, RPC communication is performed using shared memory as a transport.

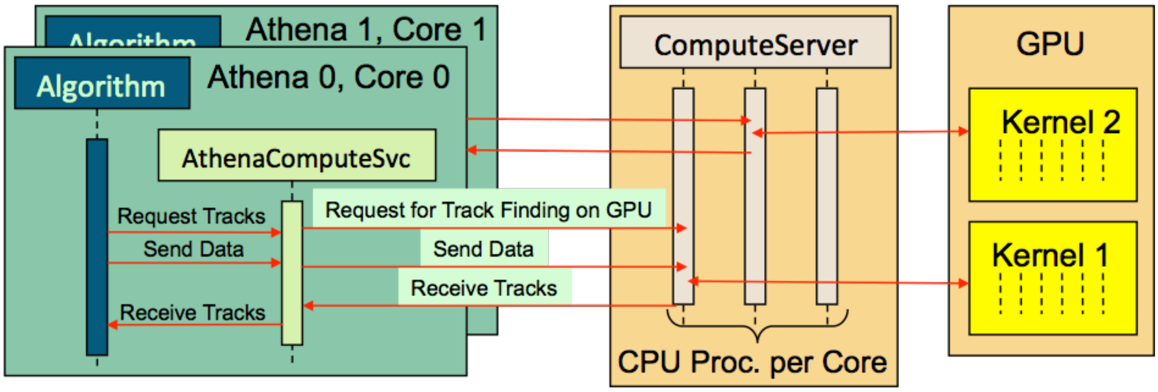

Figure 5.5: Client server architecture for GPU resource sharing [65].

Fortuitously, this shared memory buffer can also be used as the host memory buffer for transferring data to the GPU, negating the need for host-to-host copies. The client software is based exclusively on Application Programming Interfaces (APIs) available in the Portable Operating System Interface (POSIX) standard, allowing GPU integration into the trigger software without the need to link additional custom libraries. The server component can also maintain a single copy of the requisite hardware map to share amongst multiple kernels, rather than overwhelming the GPU memory with individual copies.

5.1.7 Performance results

Performance of the GPU-ported algorithms was measured on a variety of hardware and compared with a baseline CPU implementation extracted from the trigger software framework. For this evaluation, simulated high-luminosity (2×10^{34} cm^{-2}s^{-1}) $t\bar{t}$ data were used.

A comparison of data preparation performance is shown in Figure 5.6. For full-detector data preparation processing, the processing time is reduced by a factor of 26x when using the GPU algorithm. In accordance with Amdahl's law [66], the GPU-parallelized algorithm exhibits almost constant computational complexity, a result of the highly parallelizable nature of the algorithm and the large ratio of threads to

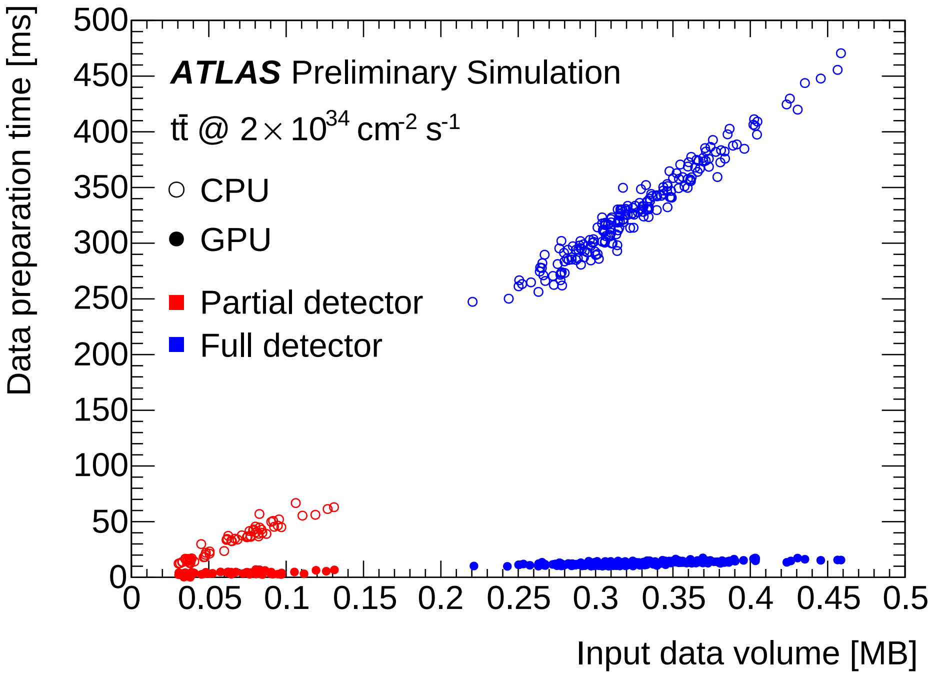

Figure 5.6: Data preparation performance comparison between an NVIDIA Tesla C2050 GPU and a single-threaded Intel E5620 CPU.

input data.

A comparison of tracking performance is shown in Figure 5.7. For the largest number of tracking input spacepoints, the processing time is reduced by a factor of 12x when using the GPU algorithm. The improvement in scaling behavior is not as drastic as in the data preparation case, primarily due to the fact that a smaller fraction of the tracking algorithm code is inherently parallelizable.

Further comparison of data preparation performance on GPUs is given in Table 5.4. The results show a promising pattern of performance evolution across device specifications. This suggests that although CPU frequency increases may be stalled due to thermal limitations, once parallelized, the algorithms again enter a regime in which performance improvements come “for free” with hardware advances.

The performance of the GPU client-server architecture was also benchmarked, and results are shown in Figure 5.8. The GPU performance does show some saturation at high trigger job multiplicity, although total throughput is still significantly higher

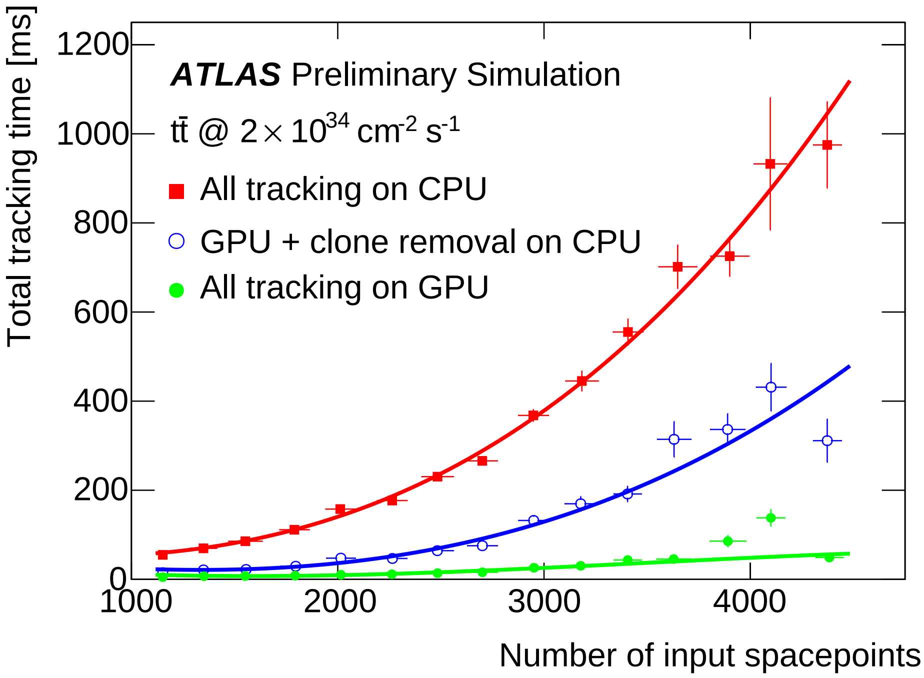

Figure 5.7: Tracking performance comparison between an NVIDIA Tesla C2050 GPU and a single-threaded Intel E5620 CPU. An additional hybrid test was performed where all processing steps were performed on the GPU, except for track clone removal, which was performed on the CPU.

than that of the CPU. This saturation is likely due to the limited number of streaming multiprocessors on the GPU. In practice, however, this saturation is not problematic, because trigger job resources typically impose more pressing constraints on trigger nodes due to their high memory usage. This problem may also be addressed without code modification by future hardware advances.

An OpenCL port of the GPU bytestream decoding and clustering algorithms was benchmarked on a variety of NVIDIA and AMD GPU and CPU hardware. The results of these benchmarks are given in Table 5.5. The results indicate reasonable performance on alternative hardware, though there is some performance disparity between the CUDA and OpenCL implementations, even on identical hardware. This disparity is most likely explained by the device-agnosticism of OpenCL's computing paradigms and the consequential inability to optimize certain memory access operations.

Device	Architecture	Cores	Core speed	Event processing time	Event processing rate
C1060	Tesla	240	1300 MHz	26.8 ms	37.3 Hz
GT 630M	Fermi	96	800 MHz	18.3 ms	54.5 Hz
GT 650M	Kepler	384	835 MHz	17.2 ms	58.2 Hz
C2050	Fermi	448	1150 MHz	9.87 ms	101 Hz
K20	Kepler	2496	706 MHz	7.83 ms	128 Hz
K40	Kepler	2880	745 MHz	6.39 ms	156 Hz

Table 5.4: Data preparation performance and specification comparison for a variety of NVIDIA GPU hardware.

Device	Platform	Cores	Core speed	Pixel processing time	SCT processing time
NVIDIA C2050	GPU (CUDA)	448	1150 MHz	3.2 ms	3.6 ms
NVIDIA C2050	GPU (OpenCL)	448	1150 MHz	3.9 ms	4.0 ms
AMD v7800	GPU (OpenCL)	1440	700 MHz	8.3 ms	7.7 ms
AMD 6276	CPU (OpenCL)	32	2300 MHz	19.0 ms	12.1 ms

Table 5.5: Data preparation performance and specification comparison for a variety of NVIDIA and AMD GPU and CPU hardware using CUDA and OpenCL implementations of the ATLAS data preparation algorithms.

5.2 Enhancement of analysis performance using JIT compilation techniques

A large fraction of high energy physics analysis constructs, such as counting or histogramming, require the evaluation of an expression in the context of an associative data structure representing an event. For example, such an expression may be used to select events passing some criteria, or to compute a quantity derived from an event. These expressions may contain symbols representing properties of the data or, more generally, functions of these properties used to compute derived quantities.

In this section, a study is presented on the usage of JIT compilation techniques to enhance the performance of expression evaluation in physics analyses. An overview

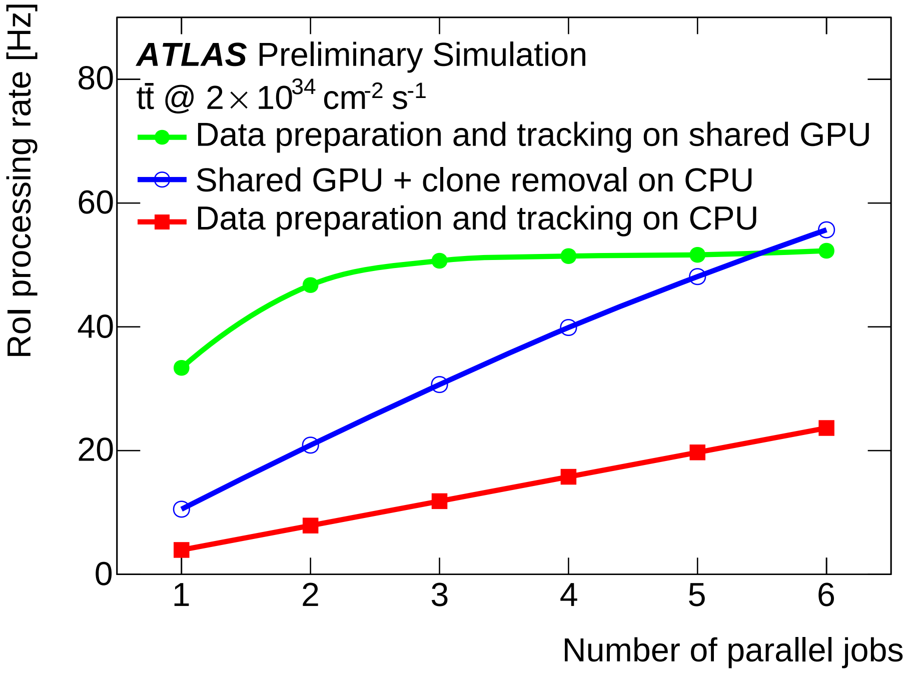

Figure 5.8: RoI throughput performance comparison between an NVIDIA Tesla C2050 GPU accessed through the client-server architecture and a single-threaded Intel E5620 CPU. An additional hybrid test was performed where all processing steps were performed on the GPU, except for track clone removal, which was performed on the CPU.

of existing techniques is presented, along with a brief discussion on JIT computing and an outline of the test implementation. Results of using JIT-compiled expression evaluation are presented, along with discussion of considerations for practical usage.

5.2.1 JIT compilation

Modern advances in computing, in particular those oriented toward dynamic languages such as JavaScript and Python, have greatly advanced compiler technology and techniques for dynamic code evaluation [67, 68, 69, 70]. These advances have tremendous potential to alter the way in which many analysis quantities are evaluated.

In particular, there has been a shift from execution relying on interpreters and virtual machines to execution relying on dynamic code generation. This has, in large part, been driven by the so-called “browser wars,” where the execution speed of JavaScript has been paramount for the user-facing performance of complex and dynamic web applications. Python’s own recent position at the forefront of scientific computing has also motivated studies into enhancement of its performance.

Traditionally, source code translation and/or execution has been performed by an interpreter, a virtual machine, or a compiler that translates directly to machine code. Each of these techniques has both benefits and drawbacks.

An interpreter evaluates code by iterating over its instructions sequentially, interpreting them at run time and performing the corresponding operation on the CPU. While this is the most dynamic approach from a development and deployment perspective, it is also the least performant.

A virtual machine is a hybrid approach between run time code evaluation and full machine code compilation. A virtual machine provides an architecture-independent instruction set, usually in an effort to provide a cross-platform abstraction layer. At compile time, source code is translated into instructions from this set. This allows for some optimizations to be performed at compile time. At run time, the virtual machine instruction set is translated into operations on the actual CPU. This approach provides much of the same flexibility as source code interpretation, but with a smaller overhead at run time, imposed only by the translation of the virtual machine instructions.

A compiler that translates directly to machine code is by far the most performant of these three techniques. Such compilers are designed to generate code for a specific CPU architecture, typically x86 or x86_64, and always require ahead-of-time compilation. This incurs a considerable penalty in terms of development and deployment flexibility.

The desire to bring dynamic languages such as JavaScript and Python to the forefront of performance has motivated the development of JIT compilation techniques for the dynamic generation of machine code. Such techniques allow for programs to evaluate their own performance and generate native machine code at run time to facilitate execution that might otherwise be handled by an interpreter or virtual machine. This approach brings with it the flexibility of dynamic languages, both in terms of development cycle and deployment, as well as the performance characteristic of native machine code.

5.2.2 Existing evaluation techniques

Many high energy physics analyses are performed with the ROOT data analysis framework. In ROOT, expression evaluation on events is typically performed (indirectly) by the `TTreeFormula` expression evaluator [71]. `TTreeFormula` uses a custom interpreter that supports a subset of the C++ language grammar. ROOT also facilitates the use of bespoke C++ code for the evaluation of expressions, although this approach is far less flexible than the evaluation of expressions at run time.

Additionally, `TTreeFormula`'s implementation is such that expression evaluation is strongly coupled to data loading. This problem is exacerbated by ROOT's primary APIs for accessing `TTreeFormula`, namely `TTree::Draw` and `TTree::Project`, which enforce single-use access of data loaded by `TTreeFormula`.

5.2.3 JIT computing on high energy physics data

This study explored the use of the Numexpr JIT compiler [72] to accelerate the evaluation of expressions on data for high energy physics analyses. Numexpr dynamically generates machine code for evaluating expressions on NumPy [73] arrays. NumPy arrays are highly efficient multidimensional homogeneously typed arrays written as a C extension for Python. Numexpr is also capable of utilizing vectorized instructions to operate on multiple chunks of data at once.

Because the use of Numexpr decouples the evaluation of expressions from data loading, Numexpr also allows for multithreaded evaluation of its generated code.

Concordantly, results for multithreaded expression evaluation are also presented.

Expressions used in the $Z_{ll}H \rightarrow \tau_l\tau_h$ were benchmarked for a simulated 8 TeV $t\bar{t}$ sample. Expressions of varying size and complexity were evaluated on different event sample sizes. Comparisons were made using `TTreeFormula`, bespoke C++, and Numexpr-enhanced Python. To load data into a format suitable for dynamic evaluation, the root_numpy [74] package was used.

5.2.4 Performance results

Execution times for expressions of varying complexity are shown in Figure 5.9. Data loading time is excluded from these times, though the overhead of run time compilation is included. For each configuration, five trials were performed and their times averaged. The standard error of each configuration is shown. The numerical results of different techniques were compared and no significant differences in precision were observed.

The JIT-compiled expression evaluation exhibits better performance than ROOT in all cases, and in some cases even surpasses the performance of the pre-compiled C++ code due to Numexpr's use of vectorized instructions. For large expressions, the JIT compiler exhibits poorer performance, though this can be mitigated through multithreading and, potentially, use of a different compiler infrastructure.

5.2.5 Considerations for practical usage

Although Numexpr provides considerable performance enhancement of expression evaluation, it does have several drawbacks that might cause complications in naïve usage.

Firstly, Numexpr does not support out-of-core computation, so the entire dataset under evaluation must be loaded into memory. However, packages do exist that allow for out-of-core evaluation of larger-than-memory datasets using Numexpr [75].

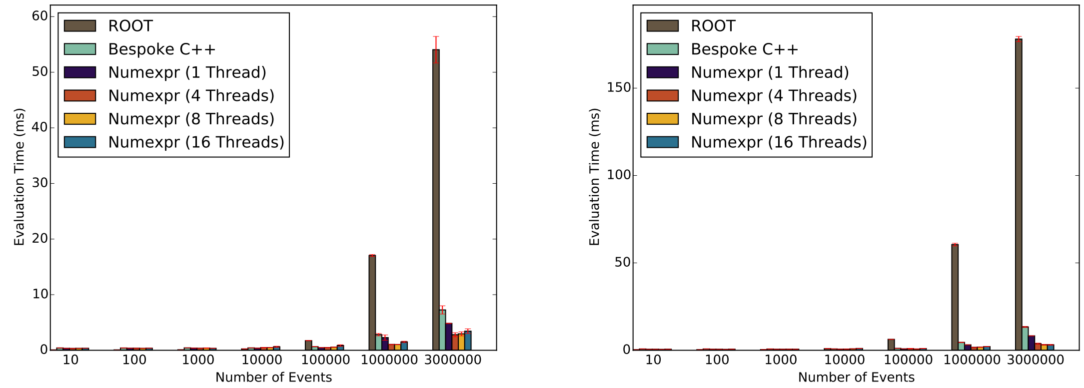

(a) 1 expression tree leaf

(b) 2 expression tree leaves

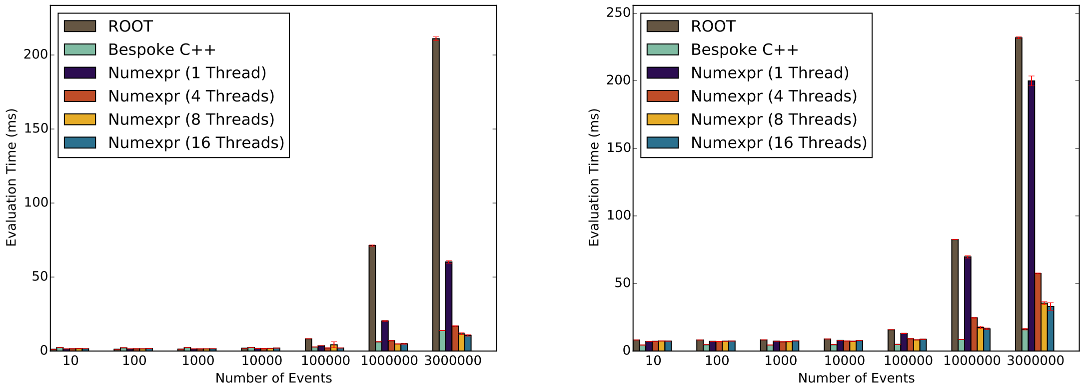

(c) 7 expression tree leaves

(d) 23 expression tree leaves

Figure 5.9: Expression evaluation times for different evaluation techniques with a variety of expressions on simulated 8 TeV $t\bar{t}$ events. Times are averaged across five trials, with standard error bars indicated.

Secondly, Numexpr does not allow for the evaluation of arbitrary functions on event properties. New intrinsics can be implemented in the Numexpr compiler, but no programmatic facility is provided for this purpose. Nonetheless, functions can be evaluated externally (potentially even using Numexpr), and the results joined with the original data to be incorporated into Numexpr's evaluation infrastructure.

Finally, the current infrastructure for loading ROOT files into Numpy arrays, root_numpy, is not as performant as ROOT's native interface. This is primarily due to the implementation of the Python/ROOT bindings, for which a more efficient implementation could be developed. Nonetheless, the loading time is within an order of magnitude of ROOT's native speed. Given that multiple expressions are typically evaluated against a single dataset (e.g. in the case of multiple analysis categories or histogram distributions), the reusability of pre-loaded data still provides extensive performance benefits over ROOT's single-use data loading and expression evaluation infrastructure.

Chapter 6

Overview of the $H \to \tau\tau$ analyses

Each of the $H \to \tau\tau$ analyses uses slightly different techniques due to differing event topologies and backgrounds. This chapter provides a brief overview of the analysis types, datasets, blinding strategy, and signal extraction techniques used in these analyses. This overview is followed by a detailed description of the discriminating variables and event selections used in the $H \to \tau_l\tau_h$ and $Z_{ll}H \to \tau_l\tau_h$ analyses.

6.1 Analysis types

The main analyses, documented in their entirety in [76], are each a multivariate analysis (MVA) utilizing a boosted decision tree (BDT) discriminant to separate signal and background events, as detailed in Section 6.5.4. An additional cut-based analysis is also performed in each of the main analysis channels, although the cut-based results are used only as a cross-check due to their significantly lower sensitivity. Both analysis types provide distinct advantages. The BDT provides stronger discrimination power at the cost of opaque selection logic, while the cut-based analysis provides a more traditional and physics-driven signal extraction procedure.

The VH analyses will be documented in a future ATLAS publication. They are implemented solely as cut-based analyses due to the large statistics required to train a BDT and limitations on manpower.

In the MVA, a fit is performed on the output score distribution of the BDT. While selections on event topology are used to divide events into categories based roughly on their production mode, selections on event kinematics are left intentionally loose so that the responsibility for signal and background discrimination is delegated primarily to the BDT. Thus, unlike traditional cut-based analyses, the majority of background discrimination and rejection in the MVA is performed *after* signal region categorization, encoded in the final BDT output distribution.

In the cut-based analyses, a fit is performed on the reconstructed $\tau\tau$ invariant mass distribution, $m_{\tau\tau}^{\mathrm{MMC}}$. The mass reconstruction is performed using a maximum likelihood scan detailed in Section 6.5.3.5. Due to the presence of an extra neutrino in the WH search channel, the late-projected transverse mass, M_{2T} [77], is used as an alternative to the mass reconstruction scan. In contrast to the MVA, the majority of background discrimination and rejection in the cut-based analyses is performed *during* signal region categorization by the selection that defines the region.

6.2 Analysis datasets

The main analyses used all proton-proton collision data recorded by the ATLAS detector during Run I. Quality criteria were applied to the data, including requirements on beam quality and detector subsystem status, resulting in the 4.5 fb^{-1} of 7 TeV (2011) data and 20.3 fb^{-1} of 8 TeV (2012) data used in the analyses.[1] The cumulative integrated luminosity over the course of Run I is shown in Figure 6.1a. In-time pile-up conditions, as indicated by the average number of interactions per bunch crossing, are shown for 7 TeV and 8 TeV operation in Figure 6.1b.

The VH analyses used only the 8 TeV dataset, primarily due to the relatively low integrated luminosity of the 7 TeV dataset.

---

[1] In the ATLAS experiments, “good” lumi blocks are organized into “Good Runs Lists” (GRLs). The GRLs used in these analyses are:

`data11_7TeV.periodAllYear_DetStatus-v36-pro10-02_CoolRunQuery-00-04-08_All_Good.xml`

`data12_8TeV.periodAllYear_DetStatus-v61-pro14-02_DQDefects-00-01-00_PHYS_StandardGRL_All_Good.xml`

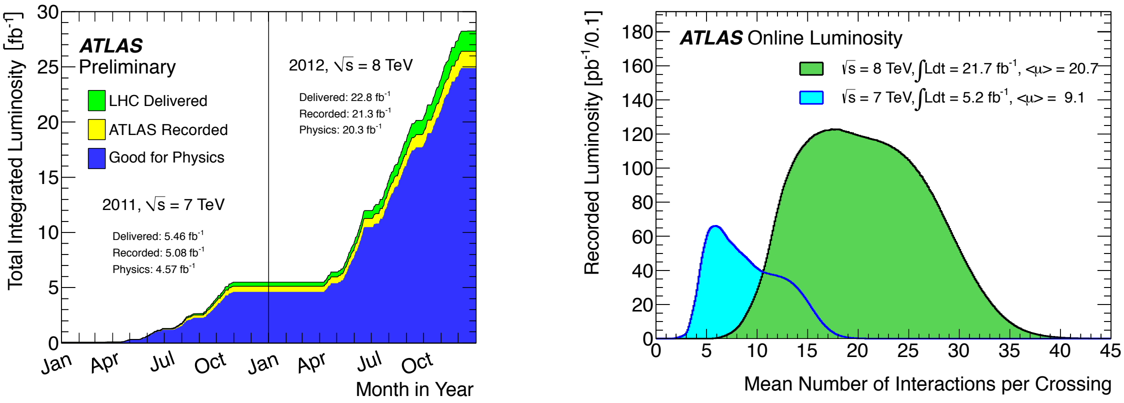

(a) Integrated luminosity

(b) In-time pile-up

Figure 6.1: Properties of Run I data taking including (a) the cumulative integrated luminosity delivered, recorded, and passing data quality criteria and (b) the average number of proton-proton interactions per bunch crossing.

6.3 Blinding strategy

In order to avoid human-induced bias in the analysis results, distributions of discriminating variables (BDT and $m_{\tau\tau}^{\mathrm{MMC}}$) entering the fit were partially or completely blinded for data in the signal regions. For the BDT score distribution, where signal-like events are expected to occupy upper bins, data in lower bins were unblinded until the integrated content of Monte Carlo signal events in the unblinded bins represented 30% of the expected signal contribution to the region. For the $m_{\tau\tau}^{\mathrm{MMC}}$ distribution, data events were blinded in the range 90–160 GeV in the main analyses and over the entire $m_{\tau\tau}^{\mathrm{MMC}}$ range in the VH analyses.

6.4 Hypothesis testing and statistical procedures

There are three primary statistical procedures in the search for and measurement of a $H \to \tau\tau$ signal:

- **Limit setting**: Putting an upper bound on the maximum signal strength that would not be ruled out by a given measurement;

- **Detection**: Measuring an event rate greater than that of known background processes; and
- **Measurement**: Determining a confidence interval around the best-fit value of a model parameter.

All three of these procedures are performed via hypothesis testing. In a statistical sense, hypothesis testing consists of quantifying measurement significance via a p-value (or equivalent Gaussian significance Z) for a combination of hypothesis and measurement. A p-value represents the probability that a different measurement would display equal or worse agreement with the hypothesis.

In order to calculate a p-value, the sampling distribution (probability density function of potential measurement values) is required. For a simple counting experiment, this might be the probability distribution of potential counts. More generally, a test statistic (some function of the measurement value) may be used, and the calculation of the p-value based on the sampling distribution of the test statistic. The p-value is calculated by integrating the sampling distribution from the measured value to the upper bound of the sampling distribution domain in the direction characterized as displaying “worse” agreement.

Before a measurement is made, thresholds are set on detection significance (the significance assuming a “null” hypothesis of background processes only) and exclusion significance (the significance assuming a null hypothesis of background *and* signal processes). In particle physics, a common detection significance threshold is $p = 2.87 \times 10^{-7}$, which corresponds to a Gaussian significance of $Z = 5\sigma$, and a common exclusion threshold is $p = 0.05$ (exclusion at the 95% confidence level). For measurements, two-sided confidence intervals are set at 68% and 95% confidence levels, corresponding to p-values of 0.16 and 0.025, respectively, on each side of the sampling distribution.

In the $H \to \tau\tau$ analyses, test statistics built around the profile likelihood ratio are used for limit setting, detection, and measurement. The description of the likelihood function, test statistics, and asymptotic formulae for test statistic sampling distributions in the large sample limit are given in [78].

6.4.1 Profile likelihood ratio

In order to generate a test statistic, a hypothesis must be reified in terms of a quantitative model. This model must encode both parameters of interest as well as nuisance parameters whose values must be taken from a fit. In the $H \to \tau\tau$ analyses, the parameter of interest is the signal strength, $\mu = \sigma/\sigma_{\mathrm{SM}}$, which represents the ratio of actual Higgs boson production rate to the rate predicted by the Standard Model. In addition, various nuisance parameters, $\boldsymbol{\theta}$, enter the fit model, encoding background normalizations (described in Chapter 7) as well as experimental, theoretical, and modeling uncertainties (described in Chapter 8).

In the $H \to \tau\tau$ analyses, measurements are extracted from histograms of discriminating distributions (BDT score and $m_{\tau\tau}^{\mathrm{MMC}}$). The likelihood function for a given measurement with N histogram bins, n_i, is given by the product of Poisson distributions:

$$L_{S+B}(\mu, \boldsymbol{\theta}|n_i) = \prod_{i=1}^{N} \frac{(\mu s_i + b_i)^{n_i}}{n_i!} e^{-(\mu s_i + b_i)}, \tag{6.1}$$

where s_i and b_i are the Standard Model expectation values for signal and background bin content, respectively, and may be functions of $\boldsymbol{\theta}$.

Additional measurements may be performed to constrain nuisance parameters in a maximum likelihood fit. The likelihood functions for these measurements depend on the nuisance parameter under consideration, but are typically Gaussian distributions. In general,

$$L_{NP}(\theta|X) = f_X(\theta, \dot{\theta}, \Delta\theta), \tag{6.2}$$

where f is some function of the nuisance parameter, θ, the nominal value of the nuisance parameter, $\dot{\theta}$, the uncertainty associated with θ, $\Delta\theta$, and depends on the measurement value X.

The combined likelihood function may then be expressed as

$$L(\mu, \boldsymbol{\theta}) = \prod_{n \in SR} L_{S+B}(\mu, \boldsymbol{\theta}|n) \prod_{n \in CR} L_{S+B}(\mu, \boldsymbol{\theta}|n) \prod_{\theta \in \boldsymbol{\theta}} L_{NP}(\theta|X), \tag{6.3}$$

where the first product is over signal region histograms, the second product is over control region histograms, and the third product is over nuisance parameters.

When performing the likelihood fit, one is generally not interested in the fit values of nuisance parameters. If it is possible to express the nuisance parameters in terms of the signal strength, the likelihood function may be recast to depend solely on μ, allowing extraction of the parameter of interest without calculating the best fit values of nuisance parameters. Such a modified likelihood is known as the "profile likelihood."

This technique may be used to construct the profile likelihood ratio:

$$\lambda(\mu) = \frac{L(\mu, \hat{\hat{\boldsymbol{\theta}}})}{L(\hat{\mu}, \hat{\boldsymbol{\theta}})}. \tag{6.4}$$

Here $\hat{\hat{\boldsymbol{\theta}}}$ is the value of $\boldsymbol{\theta}$ that maximizes L for a specified value of μ. $\hat{\mu}$ and $\hat{\boldsymbol{\theta}}$ are the values of μ and $\boldsymbol{\theta}$ that maximize L unconditionally. The value of the profile likelihood ratio is bounded in the range $[0, 1]$, with 0 indicating a value of μ to be least likely and 1 indicating the best agreement between the data and the hypothetical μ.

6.4.2 Test statistics and asymptotic formulae

The test statistics used in the $H \to \tau\tau$ analyses vary depending on the hypothesis being tested. In limit setting and detection, a one-sided test statistic is used that bounds the potential values of μ, while for measurement, a two-sided test statistic is used.

The sampling distribution for each test statistic can be computed with asymptotic formulae that are valid in the large sample size limit, where the negative log-likelihood is shown by Wald [79] to follow a Gaussian distribution:

$$-2\ln\lambda(\mu) = \frac{(\mu - \hat{\mu})^2}{\sigma^2} + \mathcal{O}(1/\sqrt{N}). \tag{6.5}$$

The standard deviation, σ, can be determined by using a special representative data set, referred to as the "Asimov data set," for which properties of the distribution $-2\ln\lambda(\mu)$ are known. Traditional alternative techniques involve the determination of the sampling distribution through large quantities of Monte Carlo simulations.

6.4.2.1 Limit setting

In limit setting, the p-value of the signal-plus-background ($\mu = 1$) hypothesis may be used to set limits at a specified confidence level (though, as will be subsequently discussed, the alternative CL$_s$ method is used in practice). Expected limits are computed under the assumption of the alternative background-only ($\mu = 0$) hypothesis.

A test statistic, q_μ, is used that only considers values of $\hat{\mu} \leq \mu$ as displaying disagreement with the hypothesis:

$$q_\mu = \begin{cases} -2\ln\lambda(\mu) & \hat{\mu} \leq \mu \\ 0 & \hat{\mu} > \mu. \end{cases} \tag{6.6}$$

The p-value of the signal-plus-background hypothesis can then be computed by

$$p_\mu = \int_{q_\mu^{obs}}^{\infty} f_{q_\mu}(q_\mu|\mu)dq_\mu, \tag{6.7}$$

where f_{q_μ} is the sampling distribution for q_μ.

A typical procedure for setting a 95% confidence level upper limit on μ would entail finding the smallest value of μ for which $p_\mu \leq 0.05$. This procedure has the unfortunate property that hypotheses with small expected signal yields may be excluded by a downward statistical fluctuation of the data [80]. The CL$_s$ method [81] was developed to counter this problem by reweighting the exclusion p-value according to whether or not the background-only hypothesis also disagrees with the data:

$$p_{\mathrm{CL}_s} = \frac{p_\mu}{1 - p_b}, \tag{6.8}$$

where p_b is defined by

$$p_b = \int_{-\infty}^{q_\mu^{obs}} f_{q_\mu}(q_\mu|0)dq_\mu. \tag{6.9}$$

The modified p-value, p_{CL_s}, does not suffer from the same "over-exclusion" behavior, and is used as an alternative to p_μ to compute exclusion limits at the desired confidence level.

6.4.2.2 Detection

In signal detection, the p-value of the background-only hypothesis is used to quantify detection significance. Expected sensitivity to the presence of a signal is computed under the assumption of the alternative signal-plus-background hypothesis.

It is assumed that $\mu \geq 0$ for any true value of μ. Concordantly, a test statistic, q_0, is defined such that its value is 0 for cases where $\hat{\mu} < 0$:

$$q_0 = \begin{cases} -2\ln\lambda(0) & \hat{\mu} \geq 0 \\ 0 & \hat{\mu} < 0. \end{cases} \tag{6.10}$$

Thus, if the best fit value $\hat{\mu}$ is found to be below 0, the data are not considered to disagree with the background-only hypothesis. The p-value of the background-only hypothesis can then be computed by

$$p_0 = \int_{q_0^{obs}}^{\infty} f_{q_0}(q_0|0)dq_0, \tag{6.11}$$

where f_{q_0} is the sampling distribution for q_0.

6.4.2.3 Measurement

The unconditional maximum likelihood estimator, $\hat{\mu}$, is used as the best-fit value for the model parameter μ. Because both upward and downward fluctuations of the data will result in smaller p-values, confidence intervals around this value are computed using a two-sided test statistic:

$$t_\mu = -2 \ln \lambda(\mu). \tag{6.12}$$

The p-value representing compatibility of the observed data with $\hat{\mu}$ can then be computed by

$$p_\mu = \int_{t_\mu^{obs}}^{\infty} f_{t_\mu}(t_\mu|\hat{\mu}) dt_\mu, \tag{6.13}$$

where f_{t_μ} is the sampling distribution for t_μ. This allows for the construction of two-sided confidence intervals around $\hat{\mu}$.

6.4.3 Implementation and nuisance parameter treatment

The fit model for the $H \to \tau\tau$ analyses was implemented using the HistFactory software package [82]. Separate fit models were constructed for the main and VH analyses. In both cases, likelihood fits were performed across all analysis channels. Optimization of histogram bin sizes and edges was performed to maximize signal sensitivity while ensuring reliable and robust event statistics. A variety of checks were also performed to identify unwanted nuisance parameter behavior, such as strong correlations, nuisance parameters pulled far from their nominal values in the fit, and nuisance parameters which exhibited double minima in negative log-likelihood distributions.

In the main analyses, a variety of pruning, smoothing, and symmetrization techniques were applied to shape systematics to avoid fit instabilities arising from low event statistics in systematic samples:

- **Shape pruning**: A χ^2 test was performed between up/down shape variations and the nominal sample, and the variation ignored if the χ^2 value was greater than a channel-dependent threshold.

- **Smoothing**: The `TH1::Smooth` method of ROOT was used to smooth the ratio of variation to nominal shape, which was then propagated back to the shape variation.
- **Partial symmetrization**: Bins of systematics where both variations were found to lie on one side of the nominal bin value were adjusted so that the smaller variation took the nominal value.
- **Significance pruning**: Shape systematics were discarded when, for all bins, the magnitude of the variation ($|\text{up} - \text{down}|$) was less than 10% of the statistical uncertainty for that bin.

Additionally, normalization systematics in the main analyses were required to have an effect larger than 0.05% in either the up or down variation. They were also adjusted such that if both were found to lie on one side of the nominal value, the smaller variation was replaced with the nominal value.

In the $Z_{ll}H \to \tau_l\tau_h$ analysis, where backgrounds were largely estimated using data-driven methods, such adjustments did not prove necessary for fit stability.

6.5 Discriminating variables

A number of discriminating variables were utilized in the $H \to \tau_l\tau_h$ and $Z_{ll}H \to \tau_l\tau_h$ analyses. These variables were used for event selection, categorization into signal regions, and as input to the discriminating distributions used in the fit. Each of these variables may be roughly characterized as pertaining to one of three event aspects: topology, kinematics, or reconstructed mass.

6.5.1 Topological variables

Topological variables played an important role in the selection and categorization of events. In the main analyses, N_{jets} provided an important classifier by which to categorize events based on production mode. For example, the VBF production mode is expected to result in two opposite-direction high-p_T forward jets (the remnants of the quarks that radiated the vector bosons), while boosted Higgs events would be expected to have at least one jet off which the Higgs boson recoils. N_{leptons}, while important for defining event selections in the $H \to \tau\tau$ analyses, was also useful for vetoing certain backgrounds (such as diboson contributions to the $H \to \tau_l\tau_h$ analysis) where extra leptons would be expected.

In addition, two variables were constructed in the $H \to \tau_l\tau_h$ analysis as quantitative metrics of the event topology for input to the BDT:

- E_T^{miss}-ϕ-centrality: This variable is defined as the scalar sum of the components of the unit vector in the direction of E_T^{miss} in a transverse plane that has been transformed so that the two τ leptons are orthogonal. It takes a maximum value of $\sqrt{2}$ if the E_T^{miss} is directly between the two τ leptons in ϕ, falling off to a value of 1 if the E_T^{miss} is aligned with one of the τ leptons in ϕ, and tailing off to < 1 when the E_T^{miss} is outside the ϕ range delimited by the two τ leptons.

- X-η-centrality: This class of variables is used in VBF events to quantify the pseudorapidity of an object, X, relative to the two opposite-direction high-p_T forward jets in the event. It is defined as

  $$C_{\eta_1,\eta_2}(\eta_X) = \exp\left[\frac{-1}{(\eta_1 - \eta_2)^2}\left(\eta_X - \frac{\eta_1 + \eta_2}{2}\right)^2\right], \tag{6.14}$$

  where η_1 and η_2 are the pseudorapidities of the jets and η_X is the pseudorapidity of the object in question. It takes a maximum value of 1 when the object in question is between the two jets in pseudorapidity, falling off to a value of $1/e$ when the object is at the same pseudorapidity as one of the jets, and tailing off to $< 1/e$ when the pseudorapidity of the object is outside of the pseudorapidity range delimited by the two jets.

6.5.2 Kinematic variables

In addition to topological variables, several kinematic variables were used to select and categorize events. Many of these were also used as inputs to the BDT. For $H \to \tau_l\tau_h$, these variables included:

- p_T^H;
- $p_T^{j_1}$ (the p_T of the leading jet);
- $p_T^{j_2}$ (the p_T of the subleading jet);
- $\Delta R_{\tau_h,l}$;
- $p_T^{\tau_h} + p_T^l + p_T^{\text{jets}}$;
- $\frac{p_T^{\tau_h}}{p_T^l}$;
- $\eta_{j_1} \times \eta_{j_2}$;
- $|\eta_{j_1} - \eta_{j_2}|$; and
- $p_T^{\text{total}} = p_T^{\tau_h} + p_T^l + p_T^{\text{jets}} + E_T^{\text{miss}}$.

Due to the limited statistics in the $Z_{ll}H \to \tau_l\tau_h$ analysis, the only kinematic selection performed was on the p_T of the τ_h decay products. This provided an important discriminant against fake τ_h backgrounds.

6.5.3 Mass variables

A variety of different reconstructed mass variables were used in the $H \to \tau_l\tau_h$ and $Z_{ll}H \to \tau_l\tau_h$ analyses. Some of these variables represent full reconstructed masses, while others represent partially reconstructed masses that are bounded by the true mass of some parent particle.

6.5.3.1 Invariant mass of the $Z \to ll$ system

In the $Z_{ll}H \to \tau_l\tau_h$ analysis, the invariant mass of the leptons associated with the Z boson, $m_{l_{Z,1}l_{Z,2}}$, provided an important tag of the ZH signature. The invariant masses for all pairs of same-flavor and opposite-sign leptons in a candidate event were calculated, and the leptons that exhibited an invariant mass closest to the Z boson mass were identified as originating from the Z boson. This identification of the Z mass provided a clean and simple criteria by which to reject fake leptons.

6.5.3.2 Dijet mass

In the $H \to \tau_l\tau_h$ analysis, the invariant mass of the two leading jets, m_{j_1,j_2}, was used as an input to the BDT. This variable is useful for identifying VBF events, where the two high-p_T forward jets in the event are expected to have a large invariant mass.

6.5.3.3 Transverse mass

In the $H \to \tau_l\tau_h$ analysis, the transverse mass, m_T, between the lepton and the E_T^{miss} was used as a discriminant against $W+\text{jets}$ backgrounds. For two objects in an event, the transverse mass is defined as

$$m_T = \sqrt{(E_T^1 + E_T^2)^2 - (\vec{p}_T^{\,1} + \vec{p}_T^{\,2})^2}. \tag{6.15}$$

The transverse mass is a partial mass, meaning that it is bounded by the true invariant mass of the parent particle. It is particularly useful in the case where one or more components of the daughter momenta is unknown, such as the non-transverse components of the invisible decay products. Because m_T is bounded by the invariant mass of the parent particle, it provides a useful discriminant between $\tau \to l\nu\bar{\nu}$ and $W \to l\nu$ decays.

6.5.3.4 Visible mass

The invariant mass of the visible τ lepton decay products was also useful for constructing a discriminant against fake backgrounds. Like m_T, it is only a partial mass, not accounting for the invisible τ lepton decay products. While it could in principle also be used as an approximation of the full invariant mass of the $\tau\tau$ system, a more accurate mass reconstruction technique was used in the $H \to \tau\tau$ analyses.

6.5.3.5 Missing Mass Calculator

The motivation for and implementation of the full $\tau\tau$ invariant mass reconstruction technique used in the $H \to \tau\tau$ analyses is detailed in [83]. Reconstruction of the full $\tau\tau$ invariant mass is complicated by the presence of one or two neutrinos per τ decay. Traditionally, analyses have relied on either partial reconstruction of the $\tau\tau$ invariant mass, such as visible mass or transverse mass, or approximations based on assumptions about the invisible decay products, such as the collinear mass approximation [84]. Each of these techniques has various drawbacks, including poor mass resolution, strong assumptions about the $\tau\tau$ system, or even lack of calculability in some cases.

In an effort to overcome the limitations imposed by such approximations, the $H \to \tau\tau$ analyses used a technique referred to as the "Missing Mass Calculator" (MMC). It has the benefit of making no assumptions about the $\tau\tau$ decay kinematics, while still constructing a complete, accurate, unbiased invariant mass.

Operation of the MMC relies on four known constraints:

$$\begin{aligned}
E_{T_x}^{\text{miss}} &= p_{\text{miss}_1}\sin(\theta_{\text{miss}_1})\cos(\phi_{\text{miss}_1}) + p_{\text{miss}_2}\sin(\theta_{\text{miss}_2})\cos(\phi_{\text{miss}_2}) \\
E_{T_y}^{\text{miss}} &= p_{\text{miss}_1}\sin(\theta_{\text{miss}_1})\sin(\phi_{\text{miss}_1}) + p_{\text{miss}_2}\sin(\theta_{\text{miss}_2})\sin(\phi_{\text{miss}_2}) \\
m_{\tau_1}^2 &= m_{\text{miss}_1}^2 + m_{\text{vis}_1}^2 + 2E_{\text{vis}_1}E_{\text{miss}_1} - 2p_{\text{vis}_1}p_{\text{miss}_1}\cos(\Delta\theta_{\text{vis-miss}_1}) \\
m_{\tau_2}^2 &= m_{\text{miss}_2}^2 + m_{\text{vis}_2}^2 + 2E_{\text{vis}_2}E_{\text{miss}_2} - 2p_{\text{vis}_2}p_{\text{miss}_2}\cos(\Delta\theta_{\text{vis-miss}_2}).
\end{aligned} \tag{6.16}$$

Here $E_{T_x}^{\text{miss}}$ and $E_{T_y}^{\text{miss}}$ are the x and y components of E_T^{miss}, p_{vis_i} and p_{miss_i} are the momenta of the visible and invisible decay products of the ith τ, θ_{miss_i} and ϕ_{miss_i} are the polar and azimuthal angles of the invisible decay products of the ith τ, $\Delta\theta_{\text{vis-miss}_i}$ is the angle between p_{vis_i} and p_{miss_i}, and m_{τ_i}, m_{vis_i}, m_{miss_i} are the masses of ith τ, its visible decay products, and its invisible decay products, respectively.

These constraints are combined with a 2–4 dimensional likelihood function to solve for 6–8 kinematic parameters (the three components of neutrino momentum from each τ decay, plus an invariant mass of the $\nu\overline{\nu}$ system in the case of a leptonic

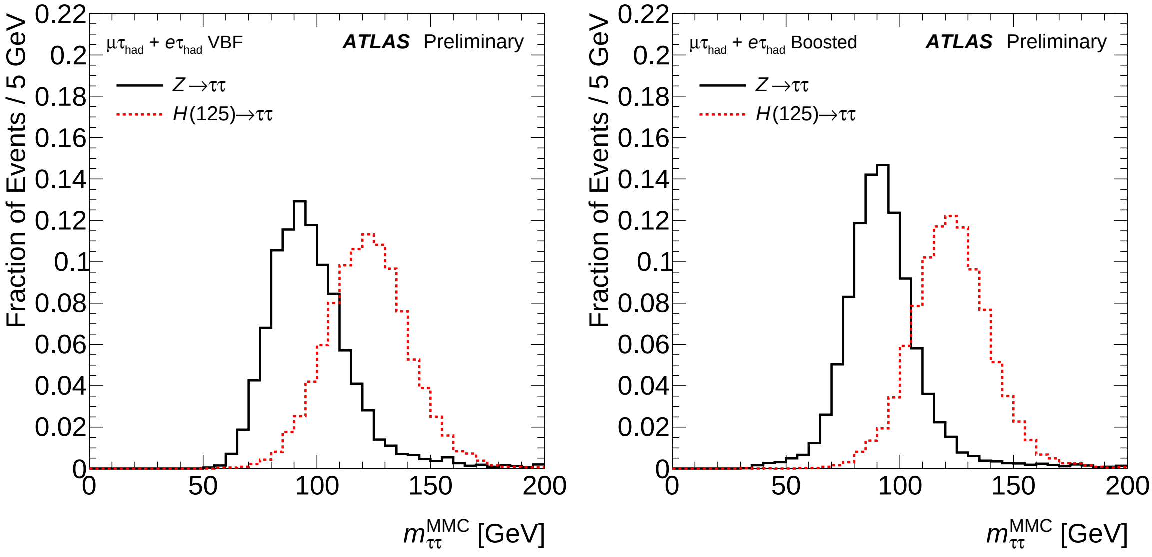

(a) VBF (b) Boosted

Figure 6.2: The reconstructed $m_{\tau\tau}^{\text{MMC}}$ for $H \to \tau\tau$ ($m_H = 125$ GeV) Monte Carlo and $Z \to \tau\tau$ embedding events (see Section 7.3.2) in the two signal regions of the $H \to \tau_l\tau_h$ analysis [76].

τ decay). In order to weight hypothetical decay parameters, the likelihood function uses information about τ decay kinematics, such as E_T^{miss} resolution and the ΔR distribution between the visible and invisible τ decay products. The mass likelihood scan is performed in a space of $(\phi_{\text{miss}_1}, \phi_{\text{miss}_2})$, with additional dimensions added for $\nu\bar{\nu}$ invariant masses as necessary. The most likely mass is chosen as the reconstructed mass.

The MMC converges for ~99% of $H \to \tau\tau$ and $Z \to \tau\tau$ events. It also provides strong mass resolution and discrimination, as illustrated in Figure 6.2.

6.5.4 Boosted decision trees

The discriminating variables described in this section were also used as input to a BDT classifier for more stringent discrimination in the main analyses. A decision tree [85] is a classifier that uses recursive partitioning of samples (based on a set of sample properties) to perform binary or multiclass classification. Decision trees are trained using labeled samples, with each selection node constructed to optimize some separation criteria for the training samples it evaluates.

Decision trees are inherently unstable, and thus a small change in the training samples can lead to a large fluctuation in the performance of the classifier. A technique known as "boosting" is used to increase the stability as well as performance of decision trees. Boosting uses an ensemble of iteratively trained decision trees, each trained to correctly classify samples that were misclassified by its predecessor. The result is referred to as a "boosted decision tree." For each sample, the output of the BDT is a value in the range $[-1, +1]$, with -1 being most background-like and $+1$ being most signal-like. Two boosting algorithms were used in the MVA: AdaBoost [86] for the $H \to \tau_h\tau_h$ channel and gradient boosting [87] for the $H \to \tau_l\tau_h$ and $H \to \tau_l\tau_l$ channels. Gradient boosting is effectively a generalization of AdaBoost to arbitrary loss functions, though both algorithms provide similar performance in the context of the $H \to \tau\tau$ analyses. BDTs within the $H \to \tau\tau$ analyses were implemented using the Toolkit for Multivariate Data Analysis with ROOT (TMVA) [88]. A separate BDT was used for each analysis category.

Event properties used for BDT training and classification in the $H \to \tau_l\tau_h$ channel are given in Table 6.1. The list was generated by considering all potential event properties and pruning those that were not well-modeled or did not have a significant positive effect on BDT discrimination power. Distributions of BDT input variables in the VBF and boosted signal regions can be found in Appendix A, except for $m_{\tau\tau}^{\text{MMC}}$, distributions of which are shown in Section 6.7.

6.6 Triggers

To identify events for consideration, single electron and muon triggers were used to seed the $H \to \tau_l\tau_h$ and $Z_{ll}H \to \tau_l\tau_h$ analyses. Trigger p_T thresholds were chosen to be

Variable	VBF	Boosted
$\Delta R_{\tau_h,l}$	●	●
m_T	●	●
E_T^{miss}-ϕ-centrality	●	●
$m_{\tau\tau}^{\text{MMC}}$	●	●
$p_T^{\tau_h} + p_T^l + p_T^{\text{jets}}$		●
$\frac{p_T^{\tau_h}}{p_T^l}$		●
m_{j_1,j_2}	●	
$\eta_{j_1} \times \eta_{j_2}$	●	
$\lvert\eta_{j_1} - \eta_{j_2}\rvert$	●	
l-η-centrality	●	
p_T^{total}	●	

Table 6.1: Boosted decision tree training and classification variables and their usage within the $H \to \tau_l\tau_h$ signal regions [76].

as low as possible while avoiding pre-scaled triggers.[2] A summary of the triggers used for 7 TeV and 8 TeV data taking periods is given in Table 6.2. The first component of the trigger name indicates its level in the trigger chain, with EF corresponding to the Event Filter. The second component of the trigger name indicates the object type (electron or muon) and corresponding p_T threshold, as well as any isolation requirements and additional trigger chain information. The third component indicates the trigger identification criteria and revision.

For the 7 TeV $H \to \tau_l\tau_h$ analysis, the $\tau_e\tau_h$ channel is seeded by single electron triggers with a p_T threshold varying between 20–22 GeV, while the $\tau_\mu\tau_h$ channel is seeded by single muon triggers with a p_T threshold of 18 GeV. For the 8 TeV $H \to \tau_l\tau_h$ and $Z_{ll}H \to \tau_l\tau_h$ analyses, more complex triggers are used that allow for either isolated lower-p_T objects with or higher-p_T objects without isolation requirements. The 8 TeV

---

[2] Pre-scaled triggers are triggers that only accept every nth event passing the trigger criteria, where n is a configurable (potentially even dynamic) value. This results in a rate reduction that is typically useful for low-p_T threshold triggers, where high rates are expected. This behavior is in contrast to non-pre-scaled triggers, where every event passing the trigger criteria is recorded.

Collision energy	Type	Period	Trigger
7 TeV	Single Electron	A–J	`EF_e20_medium`
		K	`EF_e22_medium`
		L	`EF_e22_medium1`
	Single Muon	A–I	`EF_mu18_MG`
		J–L	`EF_mu18_MG_medium`
8 TeV	Single Electron	All	`EF_e24vhi_medium1 || EF_e60_medium1`
	Single Muon	All	`EF_mu24i_tight || EF_mu36_tight`

Table 6.2: A summary of the triggers used across 7 TeV and 8 TeV data taking periods in the $H \to \tau_l \tau_h$ and $Z_{ll}H \to \tau_l \tau_h$ analyses. A summary of ATLAS trigger nomenclature is given in Section 6.6.

electron trigger p_T thresholds are 24 GeV with isolation and 60 GeV without isolation. The 8 TeV muon trigger p_T thresholds are 24 GeV with isolation and 36 GeV without.

6.7 Event selection and categorization

After candidate events are identified by the single lepton triggers, they are subjected to a series of selection criteria designed to categorize them into different signal regions. The event selections for the MVA and VH analyses are designed to impose the minimum number of requirements necessary for identification and categorization of events. In the MVA, this minimal selection is designed to allow more efficient discrimination by the BDT. In the VH analyses, a minimal selection is necessitated by the limited number of events available, though it is still quite stringent due to requirements on the number of leptons in the event.

The background estimation methods used to construct the plots in this section, as well as additional selections designed to isolated background-rich control regions, are detailed in Chapter 7.

6.7.1 $H \to \tau_l \tau_h$

In the $H \to \tau_l \tau_h$ analysis, events are first required to pass a basic topological preselection, and then subsequently divided into two signal regions with additional selection criteria: VBF and boosted. The categorization into signal regions is done in a cascading fashion, so that events are only considered for inclusion in the boosted category if they have already failed the VBF categorization selection.

6.7.1.1 Preselection

The $H \to \tau_l \tau_h$ preselection imposes only basic requirements on event topology:

- $N_l = 1$;
- $N_{\tau_h} = 1$;
- $Q_{\tau_h} \cdot Q_l < 0$ (the lepton and hadronically decaying τ lepton should be oppositely charged); and
- $m_{\tau\tau}^{\mathrm{MMC}} > 0$ ($m_{\tau\tau}^{\mathrm{MMC}}$ is set to a negative value if the likelihood scan fails to converge).

Plots of various distributions at $H \to \tau_l \tau_h$ preselection can be seen for the electron channel in Figure 6.3 and the muon channel in Figure 6.4.

6.7.1.2 Post-preselection

All signal regions include two additional post-preselection criteria: a veto on $W +$ jets events and a veto on b-jet events. The $W +$ jets veto is implemented by a cut on transverse mass: $m_T < 70$ GeV. The b-jet veto is implemented by a cut on the b-tagging MV1 score at a working point which is 70% efficient for rejection of b-jet events.

6.7.1.3 VBF

The VBF signal region is designed to isolate events generated by the VBF production mode, characterized by two opposite-direction high-p_T forward jets originating from the quarks which radiated the vector bosons. VBF selection is defined relative to preselection by additionally requiring:

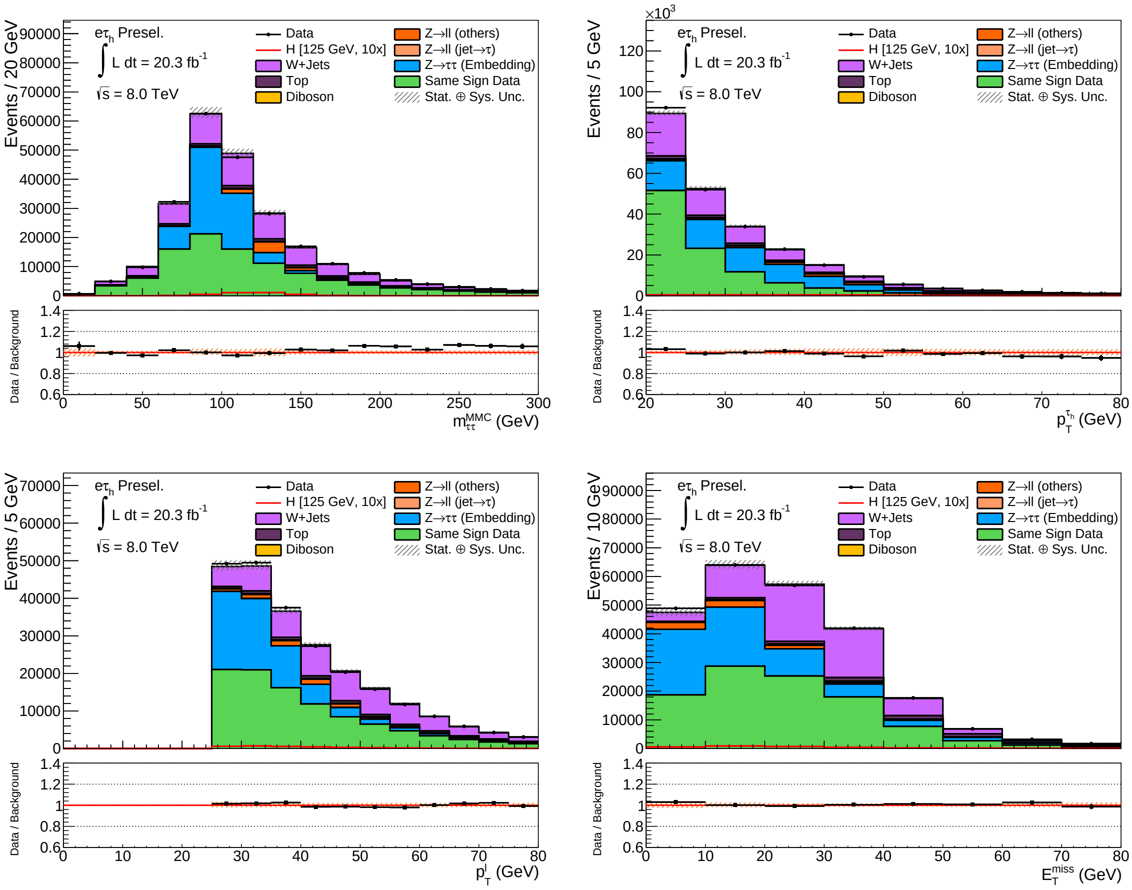

Figure 6.3: Various distributions at preselection in the electron channel of the $H \rightarrow \tau_l\tau_h$ analysis with 20.3 fb^{-1} of 8 TeV data.

- $N_{\text{jets}} \geq 2$;
- $p_T^{j_1} > 50$ GeV;
- $p_T^{j_2} > 30$ GeV; and
- $|\eta_{j_1} - \eta_{j_2}| > 3.0$.

If events fail these selection cuts, they are considered for inclusion in the boosted signal region. Events which do pass VBF selection are subject to an additional requirement on the visible mass of the τ decay products: $m_{\tau\tau}^{\text{vis}} > 40$ GeV. Events that pass VBF selection but fail this cut are not used.

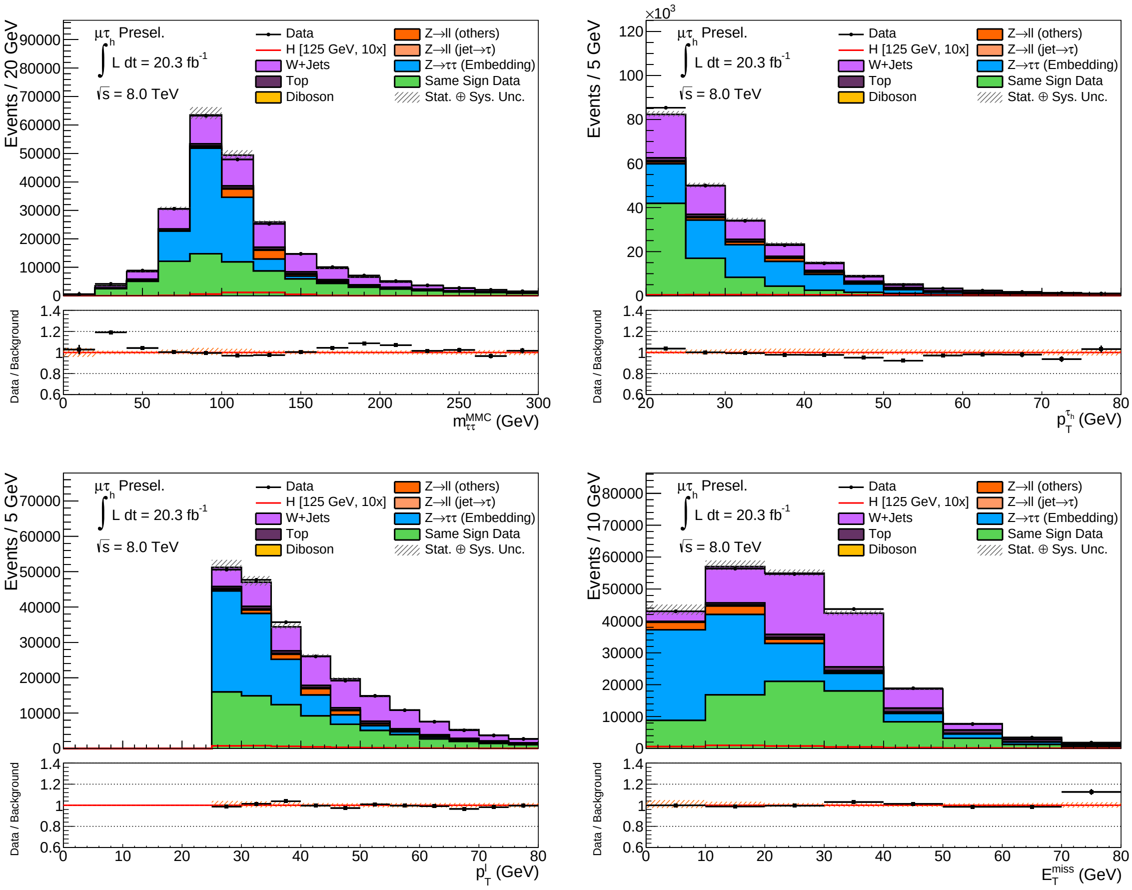

Figure 6.4: Various distributions at preselection in the muon channel of the $H \to \tau_l\tau_h$ analysis with 20.3 fb^{-1} of 8 TeV data.

Plots of various distributions in the $H \to \tau_l\tau_h$ VBF signal region can be seen in Figure 6.5.

6.7.1.4 Boosted

Only events that fail VBF selection cuts are considered for inclusion in the boosted signal region. They are required to have $p_T^H > 100$ GeV. Boosted Higgs events are primarily dominated by ggF events, but also include some contribution from VBF and VH. The boosted Higgs topology is useful due to the small angle between the invisible τ decay products, which leads to better reconstruction of E_T^{miss} and mass variables which depend on E_T^{miss}. In non-boosted Higgs events, the contributions of back-to-back τ decay products can cancel, leading to an incorrect reconstruction of

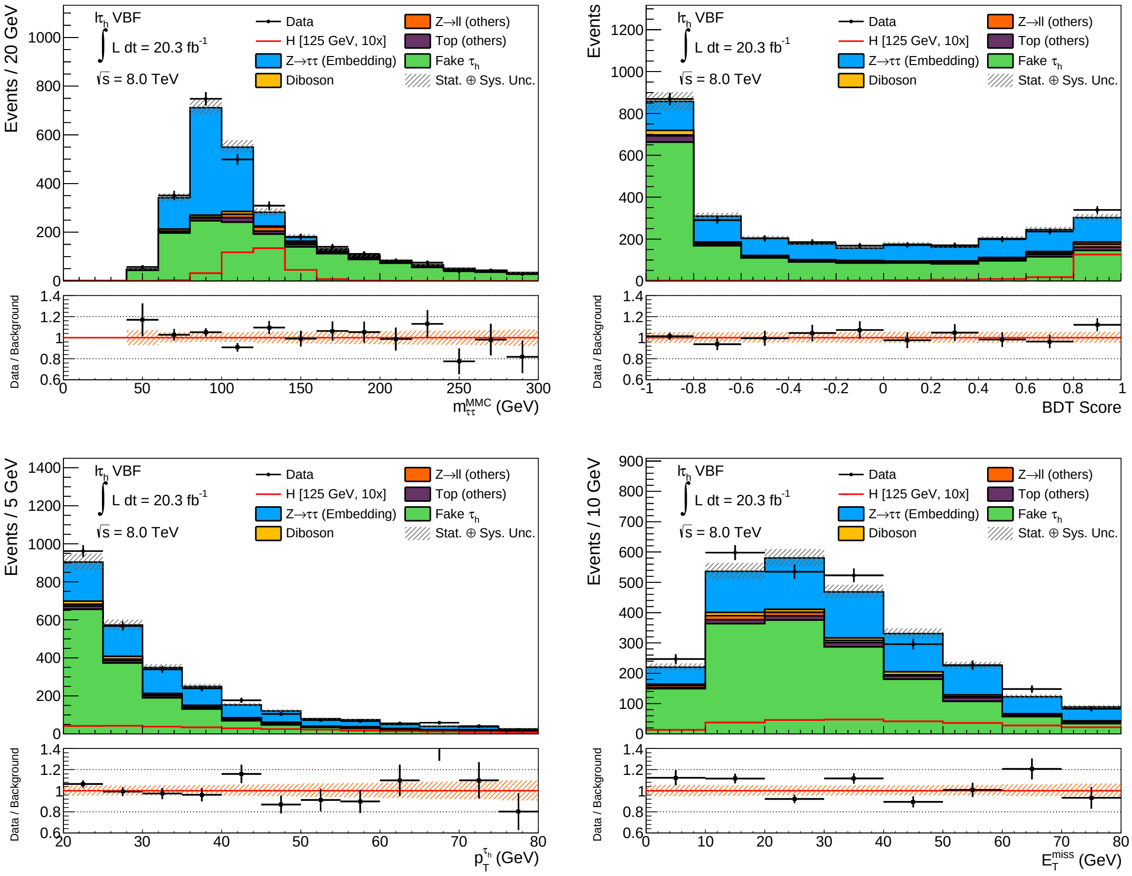

Figure 6.5: Various distributions in the unblinded $H \to \tau_l\tau_h$ VBF signal region with 20.3 fb^{-1} of 8 TeV data. These plots are constructed using pre-fit background estimation.

E_T^{miss} and mass variables.

Plots of various distributions in the $H \to \tau_l\tau_h$ boosted signal region can be seen in Figure 6.6.

6.7.2 $Z_{ll}H \to \tau_l\tau_h$

Like the $H \to \tau_l\tau_h$ analysis, the $Z_{ll}H \to \tau_l\tau_h$ analysis imposes basic preselection cuts to categorize events into analysis channels. However, due to the singular production mode and limited statistics, only a single signal region is used.

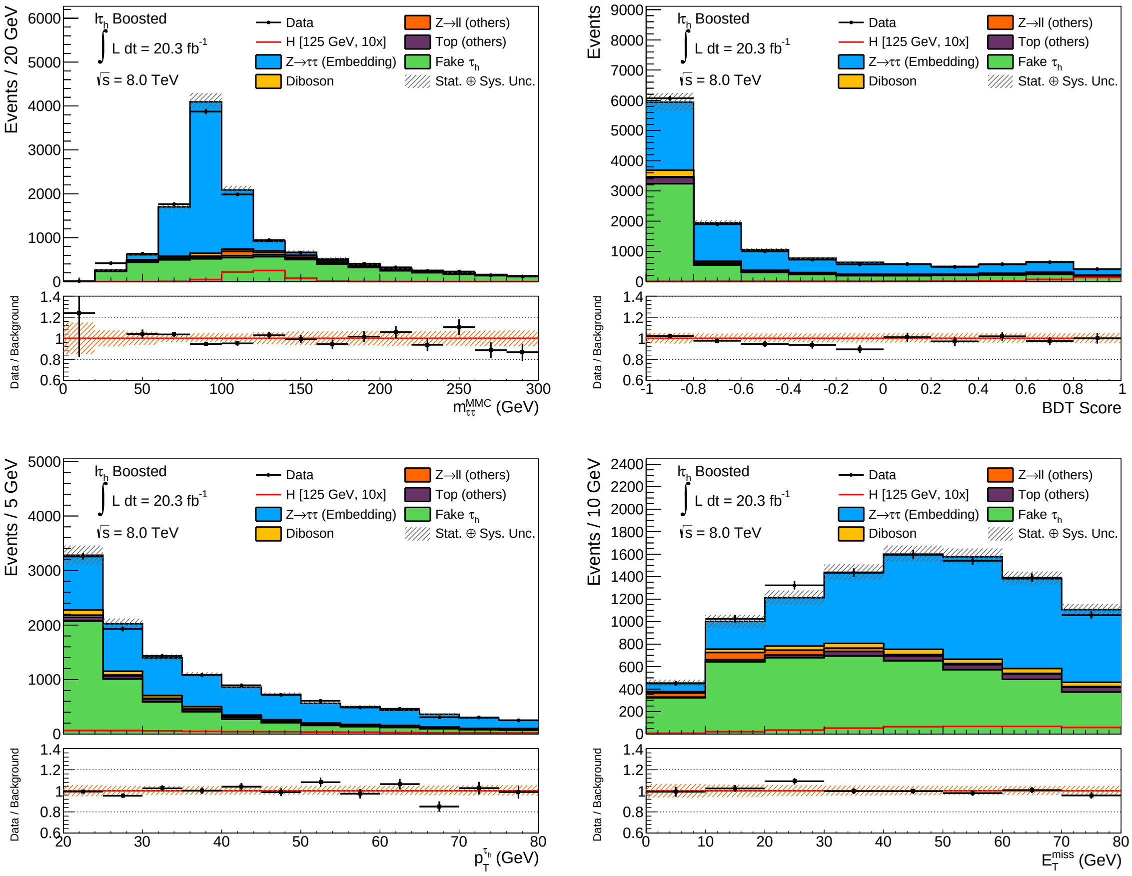

Figure 6.6: Various distributions in the unblinded $H \rightarrow \tau_l\tau_h$ boosted signal region with 20.3 fb^{-1} of 8 TeV data. These plots are constructed using pre-fit background estimation.

6.7.2.1 Preselection

The $Z_{ll}H \rightarrow \tau_l\tau_h$ preselection places requirements on both the Z boson decay and the $\tau\tau$ system:

- $N_l = 3$;
- $N_{\tau_h} = 1$;
- Flavor$(l_{Z,1})$ = Flavor$(l_{Z,2})$ (here $l_{Z,1}$ and $l_{Z,2}$ represent the leading and subleading leptons associated with the Z boson decay);
- $Q_{l_{Z,1}} \cdot Q_{l_{Z,2}} < 0$;

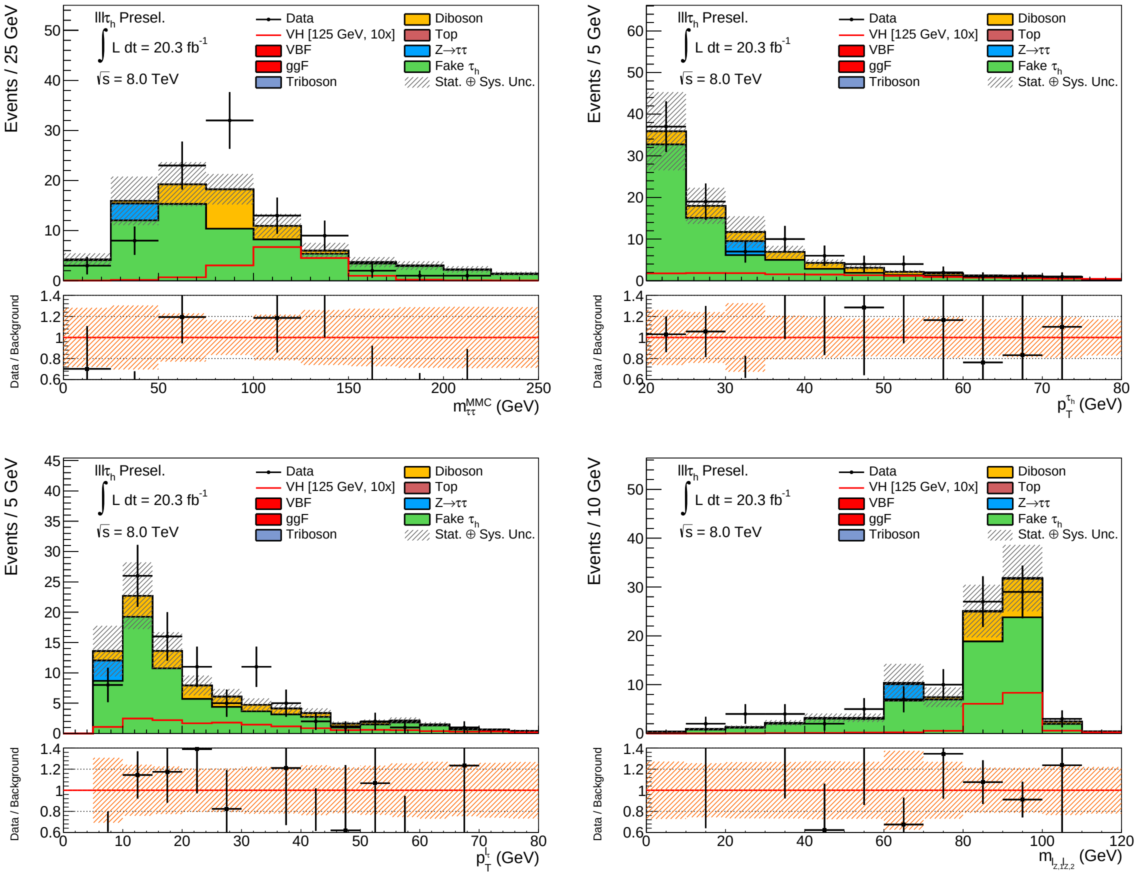

Figure 6.7: Various distributions at preselection in the $Z_{ll}H \to \tau_l\tau_h$ analysis with 20.3 fb^{-1} of 8 TeV data.

- $Q_{l_\tau} \cdot Q_{\tau_h} < 0$; and
- $m_{\tau\tau}^{\mathrm{MMC}} > 0$.

In order to determine which two of the three required leptons should be associated with the Z boson decay, the invariant mass of all pairs of same-flavor opposite-sign leptons in the event is computed and the pair with an invariant mass closest to the Z boson mass chosen. The remaining lepton is then identified as the τ_l decay.

Plots of various distributions at $Z_{ll}H \to \tau_l\tau_h$ preselection can be seen in Figure 6.7.

6.7.2.2 Signal region

The $Z_{ll}H \rightarrow \tau_l\tau_h$ signal region imposes three additional requirements to reduce backgrounds due to fake leptons and τ_h objects:

- 81 GeV $< m_{l_{Z,1}l_{Z,2}} <$ 101 GeV;
- Tight identification criteria on the electron associated with the τ_l decay; and
- $p_T^{\tau_h} >$ 30 GeV.

The latter two cuts were chosen to maximize the ratio $S/\sqrt{B}$ in the signal region, where S is the expected signal content and B is the expected background content.

Plots of various distributions in the $Z_{ll}H \rightarrow \tau_l\tau_h$ signal region can be seen in Figure 6.8.

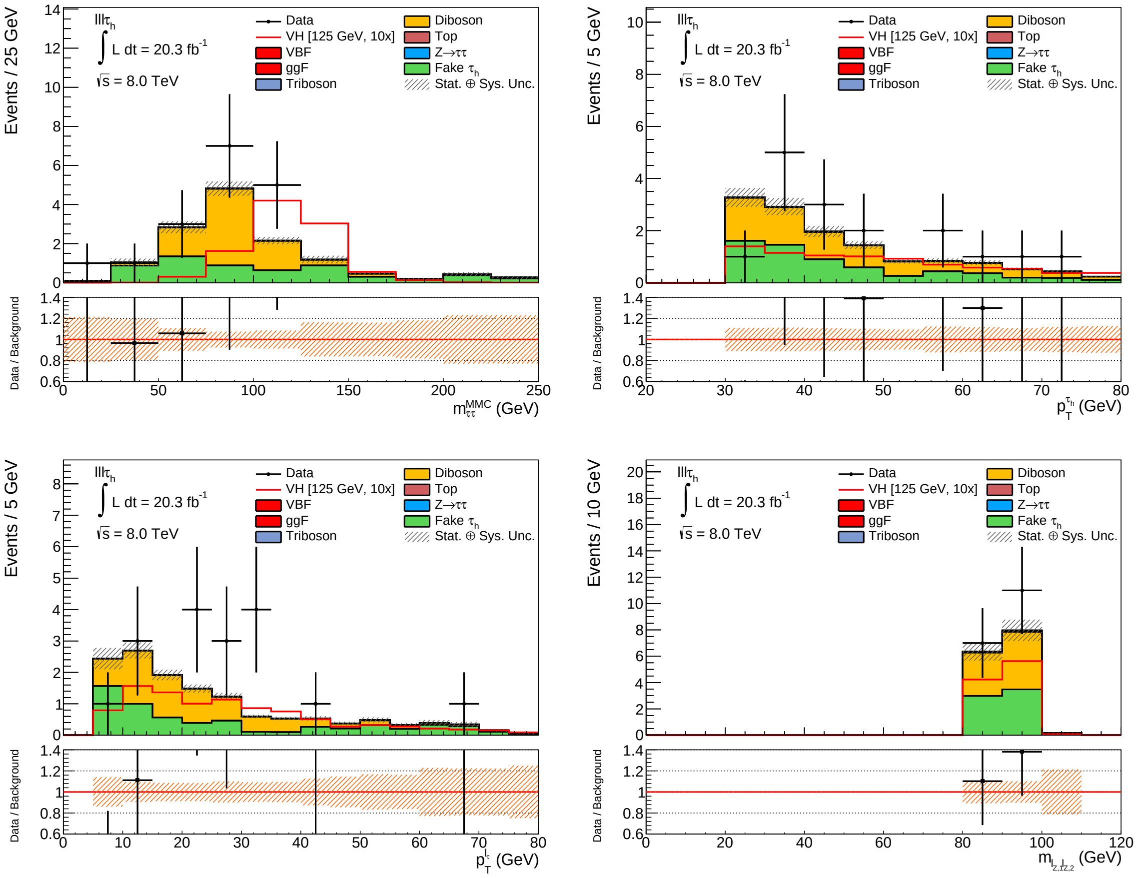

Figure 6.8: Various distributions in the signal region of the $Z_{ll}H \rightarrow \tau_l\tau_h$ analysis with 20.3 fb^{-1} of 8 TeV data. These plots are constructed using pre-fit background estimation.

Chapter 7

Signal and background modeling

Accurate modeling of both signal and background processes is essential for the detection, identification, or exclusion of a Higgs boson. These processes, in particular their event yields and shape contributions in discriminating distributions, must be well-understood in order for the statistical methods of Chapter 6 to be meaningful.

This chapter summarizes the simulation and data-driven techniques used to model signal and background processes in the $H \to \tau_l \tau_h$ and $Z_{ll}H \to \tau_l \tau_h$ analyses. It begins with a description of the signal processes considered in the $H \to \tau\tau$ analyses. This is followed by an enumeration of background processes and their mechanisms of inclusion in event selections. Finally, the Monte Carlo and hybrid "embedding" samples used to model various processes in the analyses are summarized, and an account is given of the two separate techniques used to model the large QCD multijet background.

7.1 Signal processes

There are four signal processes considered in the main analyses: ggF, VBF, WH, and ZH. In the VH analyses, only the WH and ZH processes are considered signal processes, while the ggF and VBF processes are considered "background" signal processes. ttH is not considered in any $H \to \tau\tau$ analysis due to its negligible contribution.

Because the mass of the Higgs boson coupling to τ leptons is not currently known, a variety of signal process samples are produced at different hypothetical Higgs boson mass points. The mass search range is guided by previous experimental exclusions, and is chosen to be 100 GeV to 150 GeV, with mass points spaced at 5 GeV intervals.

7.2 Background processes

A variety of processes can mimic the event topology of the $H \to \tau\tau$ and $VH \to \tau\tau$ signal processes. Background processes in the $H \to \tau\tau$ analyses are split into two general categories:

- **Irreducible backgrounds**: Those processes that generate real l, τ_l, and τ_h objects satisfying the event selections; and
- **Fake backgrounds**: Those processes that are included in event selections due to the presence of fake l, τ_l, or τ_h objects.

Irreducible backgrounds can only be differentiated from signal processes to a limited extent using cuts on event kinematics and discriminating variables (such as reconstructed invariant mass). However, because they contain real objects, often generated through processes nearly identical to the signal processes, these backgrounds cannot be completely attenuated.

In contrast, fake backgrounds can be reduced by better object reconstruction and identification criteria. The primary fake objects considered in the $H \to \tau_l\tau_h$ and $Z_{ll}H \to \tau_l\tau_h$ analyses are jets misidentified as hadronically decaying τ leptons. τ_h objects may also be faked by light leptons. In the case of the $H \to \tau_l\tau_h$ and $Z_{ll}H \to \tau_l\tau_h$ analyses, where a light lepton is required, these fake τ_h objects can be produced in association with either a real light lepton (e.g. one coming from W or Z boson processes) or a jet misidentified as a light lepton.

Both irreducible and fake backgrounds have a significant contribution to the $H \to \tau\tau$ analyses. Additionally, each background may contribute in multiple ways.

The large assortment of mechanisms by which backgrounds can enter the event selections necessitates a variety of different modeling techniques, which are described in Section 7.3.

7.2.1 $Z \to \tau\tau$

The $Z \to \tau\tau$ background is the primary irreducible background in the $H \to \tau_l\tau_h$ analysis. Capable of producing real τ_l and τ_h signatures, the $Z \to \tau\tau$ background is particularly important due to its similarity to the $H \to \tau\tau$ signal processes. Discrimination from signal using variables built on invariant mass reconstruction of the $\tau\tau$ system is only viable if the mass of the Higgs boson is significantly different from that of the Z boson.

In the $H \to \tau_l\tau_h$ analysis, modeling of the $Z \to \tau\tau$ background is validated in a dedicated control region. This control region is defined relative to a given signal region by imposing the following cuts:

- $m_T < 40$ GeV; and
- $m_{\tau\tau}^{\text{MMC}} < 100$ GeV.

Distributions of the BDT output score in the VBF and boosted $Z \to \tau\tau$ control regions are shown for 8 TeV samples in Figure 7.1.

7.2.2 QCD multijet

QCD processes provide a significant source of fake τ_h objects that enter into the $H \to \tau_l\tau_h$ and $Z_{ll}H \to \tau_l\tau_h$ analyses. Although the rates for jets faking electrons and muons are very low, the very large gg cross section at the LHC leads to a non-negligible contribution of fake τ_h and fake light lepton objects being produced in tandem. The QCD multijet background is the largest source of fake τ_h objects in these analyses and is notoriously difficult to model accurately with Monte Carlo. Consequently, two data

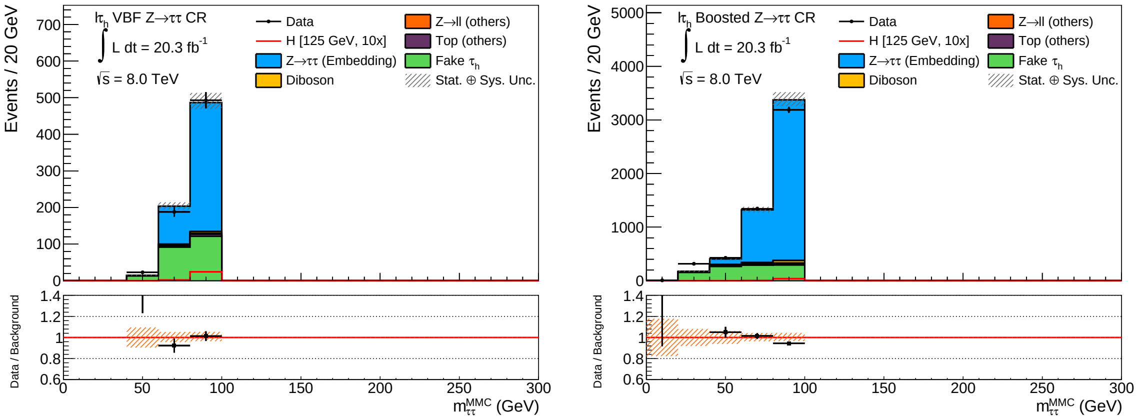

(a) VBF (b) Boosted

Figure 7.1: The $m_{\tau\tau}^{\mathrm{MMC}}$ distributions in the (a) VBF and (b) boosted $Z \to \tau\tau$ control regions with 20.3 fb^{-1} of 8 TeV data.

driven methods are used to account for the QCD multijet contribution, as described in Section 7.3.

7.2.3 $W+$jets

Aside from pure QCD contributions, backgrounds where jet-faked τ_h objects are produced in combination with real leptons also generate a significant number of background events. The largest such background is the $W+$jets background. The primary mode of contribution is due to a leptonically decaying W boson combined with a jet faking a τ_h signature. The W boson can also decay to a τ lepton, but the majority of $W+$jets events that pass event selection are due to the W-generated τ decaying leptonically combined with a jet-faked τ_h.

In the $H \to \tau_l\tau_h$ analysis, modeling of the $W+$jets background is validated in dedicated control regions. These control regions are defined relative to signal regions by inverting the $W+$jets veto in the corresponding selection, thereby requiring $m_T > 70$ GeV. Distributions of the BDT output score in the VBF and boosted $W+$jets control regions are shown for 8 TeV samples in Figure 7.2.

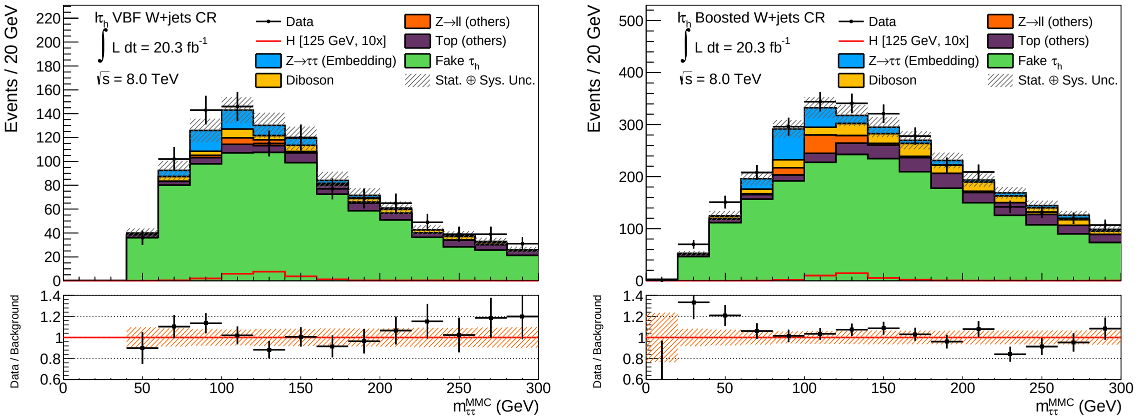

(a) VBF (b) Boosted

Figure 7.2: The $m_{\tau\tau}^{\mathrm{MMC}}$ distributions in the (a) VBF and (b) boosted $W+$jets control regions with 20.3 fb^{-1} of 8 TeV data.

7.2.4 $Z \to ll$

In the $Z_{ll}H \to \tau_l\tau_h$ analysis, $Z \to ll$ processes combined with QCD multijet processes form the largest source of fake τ_h backgrounds. In the $H \to \tau_l\tau_h$ analysis, the $Z \to ll$ background can contribute fake background events in two ways. First, although the analysis vetoes the presence of a secondary light lepton, one of the leptons may be missed by object reconstruction or fail identification, and a QCD jet may fake a τ_h signature. Alternatively, one of the leptons may fake a τ_h signature.

$Z \to ll$ modeling is validated in the $H \to \tau_l\tau_h$ analysis by replacing the dilepton veto with a requirement of two opposite-sign same-flavor leptons. Distributions of the BDT output score in the VBF and boosted $Z \to ll$ control regions are shown for 8 TeV samples in Figure 7.3.

7.2.5 Top

Backgrounds due to top quark events contribute both irreducible and fake-driven backgrounds. The top quark, which decays to a W boson and down-type quark (primarily b), is capable of generating true τ_h and τ_l signatures. Additionally, jets originating from top quarks can be misidentified as τ_h objects. This leads to a large

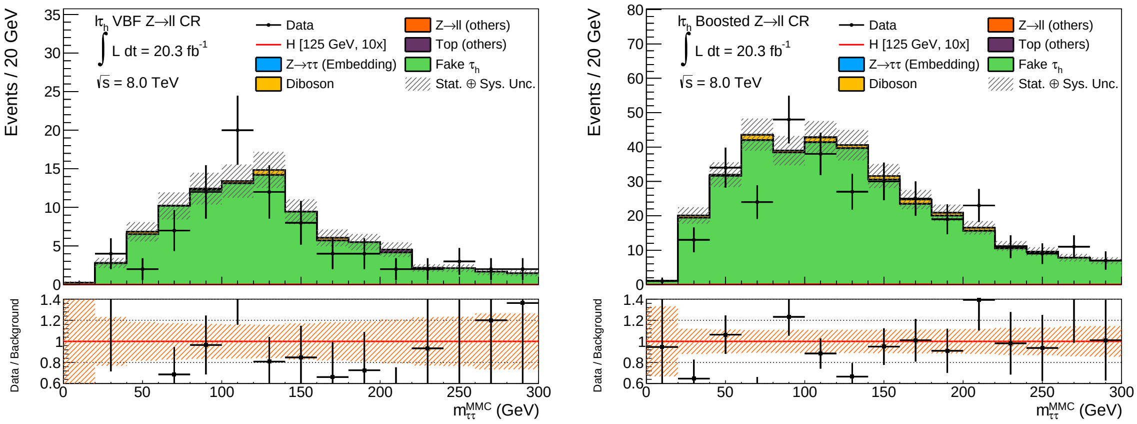

(a) VBF

(b) Boosted

Figure 7.3: The $m_{\tau\tau}^{\mathrm{MMC}}$ distributions in the (a) VBF and (b) boosted $Z \to ll$ control regions with 20.3 fb^{-1} of 8 TeV data.

variety of modes in which top quark processes can fake $H \to \tau_l \tau_h$ signatures, due to both $t\bar{t}$ and single top production.

In the $H \to \tau_l \tau_h$ analysis, modeling of the top background is validated in dedicated control regions. These control regions are defined relative to signal regions by inverting the MV1-based b-jet veto in the corresponding selection. Distributions of the BDT output score in the VBF and boosted top control regions are shown for 8 TeV samples in Figure 7.4.

7.2.6 Diboson and triboson

Diboson production (WW, WZ, and ZZ) provides only a small background contribution to the $H \to \tau_l \tau_h$ analysis, but it is the major irreducible background for all VH analyses. In the case of $H \to \tau_l \tau_h$, the secondary light lepton veto provides protection from these backgrounds, although hadronic decays of one of the bosons may allow events to bypass this criteria. The triboson background is not considered in the $H \to \tau_l \tau_h$ analysis, and provides only a minuscule contribution to the $Z_{ll}H \to \tau_l \tau_h$ analysis.

In the $Z_{ll}H \to \tau_l \tau_h$ analysis, modeling of the diboson background is validated in

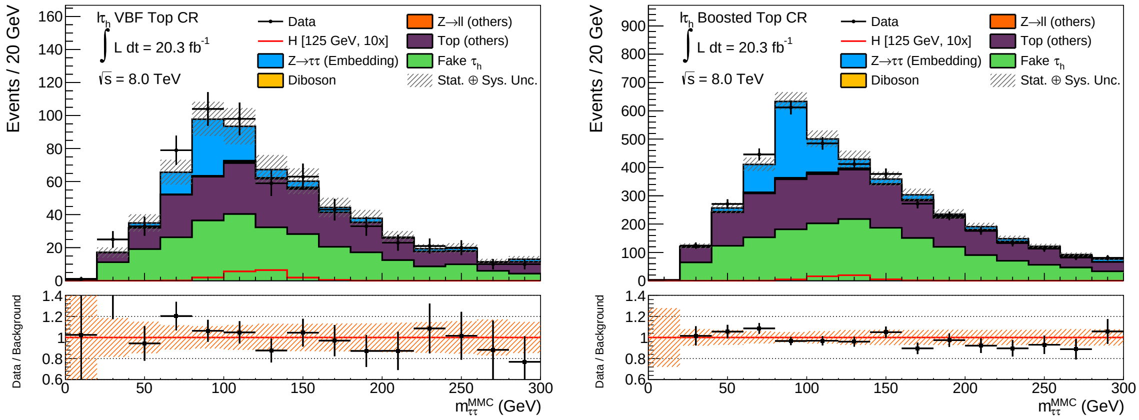

(a) VBF

(b) Boosted

Figure 7.4: The $m_{\tau\tau}^{\mathrm{MMC}}$ distributions in the (a) VBF and (b) boosted top control regions with 20.3 fb^{-1} of 8 TeV data.

a dedicated control region. This control region is defined relative to preselection by requiring 81 GeV $< m_{l_{Z,1}l_{Z,2}}, m_{\tau\tau}^{\mathrm{MMC}} <$ 101 GeV. The $m_{\tau\tau}^{\mathrm{MMC}}$ and $m_{l_{Z,1}l_{Z,2}}$ distributions in the $Z_{ll}H \to \tau_l\tau_h$ diboson control regions are shown in Figure 7.5.

7.3 Modeling techniques

Several different techniques were employed to model the various signal and background contributions to the analyses. Monte Carlo samples were used to model all signal contributions, as well as the contributions of backgrounds not containing fake τ_h objects. The large irreducible $Z \to \tau\tau$ background was modeled using a special hybrid technique involving simulated τ decays embedded in data events. Finally, the large QCD multijet fake τ_h background, which is poorly described by Monte Carlo, was modeled using two different data-driven methods.

7.3.1 Monte Carlo

ggF and VBF samples were generated using POWHEG [89], with parton showering and underlying event modeling handled by PYTHIA (version 6 [90] for 7 TeV and version 8 [91] for 8 TeV). WH and ZH samples were generated using PYTHIA (version 6

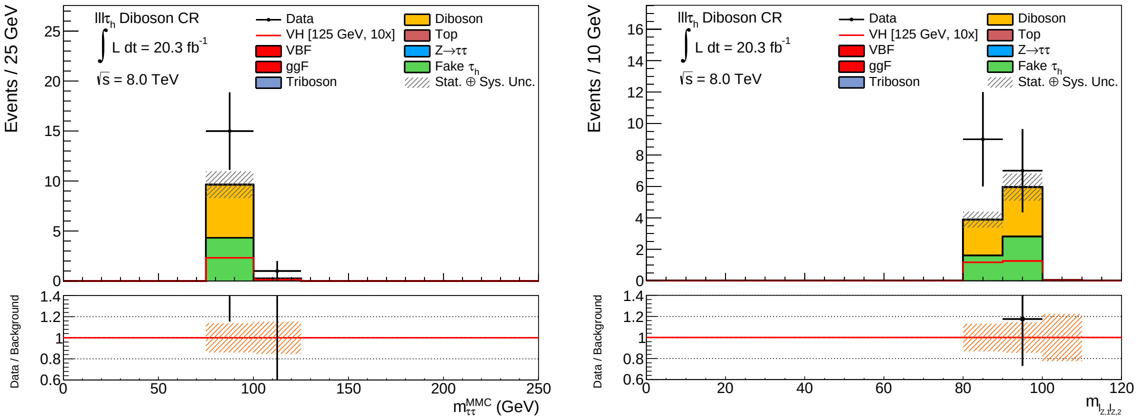

(a) $m_{\tau\tau}^{\text{MMC}}$ (b) $m_{l_{Z,1}l_{Z,2}}$

Figure 7.5: The (a) $m_{\tau\tau}^{\text{MMC}}$ and (b) $m_{l_{Z,1}l_{Z,2}}$ distributions in the $Z_{ll}H \to \tau_l\tau_h$ diboson control region.

for 7 TeV and version 8 for 8 TeV).

$W+$jets and $Z \to ll$ samples were generating using ALPGEN [92] with PYTHIA 6 handling parton showering and underlying event modeling. $t\bar{t}$ samples were generated using POWHEG, while single top samples were generated using MC@NLO [93] and AcerMC [94]. In the main analyses, WW samples were generated using ALPGEN and gg2WW [95], while WZ and ZZ samples were generated using HERWIG [96]. In the VH analyses, higher-statistic WW, $W\gamma$, WZ, ZZ samples were taken from the ATLAS $H \to WW$ analysis [97], where they were generated using combinations of gg2WW, HERWIG, POWHEG, PYTHIA, JIMMY [98], and MadGraph [99]. For all top and main analyses' diboson samples, HERWIG was used for parton showering while the underlying event was modeled with JIMMY. Triboson samples for the VH analyses were generated using MadGraph using PYTHIA 6 for parton showering and underlying event modeling.

The CT10 [100] parton distribution functions (PDFs) were used for the POWHEG, gg2WW, and MC@NLO generators, while the CTEQ6L1 [101] PDFs were used with the PYTHIA, ALPGEN, MadGraph, AcerMC, and HERWIG. In all cases, TAUOLA [102] was used to simulate τ lepton decays, while PHOTOS [103] was used to simulate bremsstrahlung from charged leptons. Detector simulation was performed using Geant4.

A variety of correction procedures were also applied to samples as recommended by ATLAS physics combined performance groups. Finally, samples were generated with a variety of different pile-up conditions to reproduce the differing conditions of each data taking period and weighted to the fraction of integrated luminosity represented by each data taking period.

7.3.1.1 Monte Carlo normalization

Monte Carlo samples were normalized to match expected yields in data using their effective cross section and the amount of integrated luminosity being modeled. The normalization factor is given by

$$f_p = \frac{n_p^{\mathrm{data}}}{n_p^{\mathrm{MC}}} = \frac{\mathscr{L} \cdot \sigma_p^{\mathrm{effective}}}{n_p^{\mathrm{MC}}}, \tag{7.1}$$

where the index p refers to the process being modeled by the Monte Carlo sample in question, n_p^{data} is the number of events for the process expected in data, n_p^{MC} is the number of Monte Carlo events generated for the process, $\mathscr{L}$ is the integrated luminosity of the data, and $\sigma_p^{\mathrm{effective}}$ is the effective cross section for the process in question.

In the fake factor method, described in Section 7.3.3, the normalization of the top component not coming from jet → τ_h fakes was allowed to float freely in the final fit as a nuisance parameter.

7.3.2 $Z \to \tau\tau$ embedding

In the $H \to \tau_l \tau_h$ analysis, the irreducible $Z \to \tau\tau$ background is estimated using a hybrid data and Monte Carlo technique known as “embedding.” This technique offers a data-driven approach to estimating this background, in spite of the fact that a pure and high-statistics $Z \to \tau\tau$ sample cannot be readily obtained from data (both because of the difficulties associated with τ identification and the potential overlap with the $H \to \tau\tau$ signal).

The technique begins by selecting a high-purity $Z \to \mu\mu$ sample from data. $Z \to \mu\mu$ and $Z \to \tau\tau$ events have identical event topologies and kinematics, except for the different masses and decays of the muons and τ leptons. The small coupling of the Higgs boson to muons also means there will be little signal contamination in the sample. The selection criteria for the $Z \to \mu\mu$ sample include:

- Various criteria on the number of hits within the Inner Detector;
- $I^{\mu}_{p_T}(0.4) < 0.2$ for each muon;
- $|\eta_\mu| < 2.5$ for each muon;
- $p_T^{\mu_1} > 15$ GeV for 7 TeV data, 20 GeV for 8 TeV data;
- $p_T^{\mu_2} > 15$ GeV for 7 TeV and 8 TeV data;
- Muons have a common primary vertex;
- Muons are oppositely charged; and
- $m_{\mu\mu} > 40$ GeV.

In the event that more than one muon pair is identified per event, the pair with an invariant mass closest to the Z boson mass is used.

After the $Z \to \mu\mu$ events are identified, the muons are “subtracted” by removing muon tracks and any corresponding energy deposits in the calorimeters. A TAUOLA-simulated τ decay is then used to replace the muon, with the kinematics of the τ set to be identical to the muon.

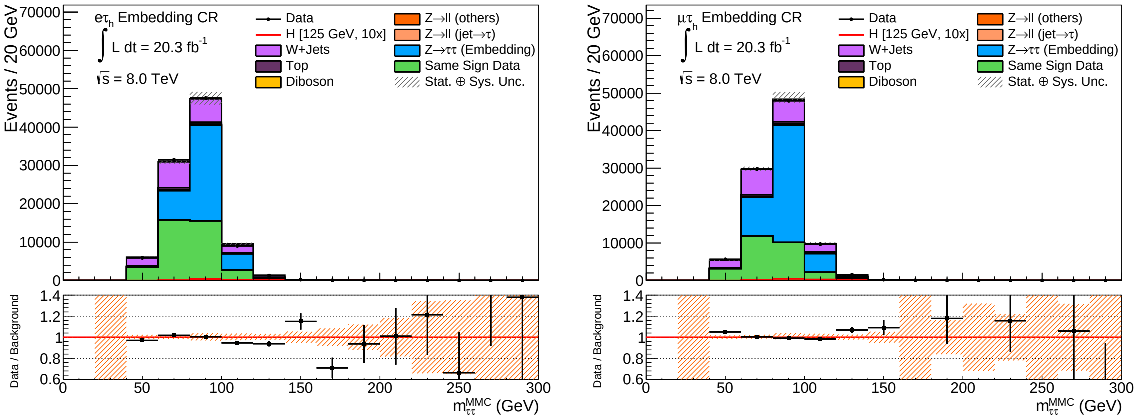

(a) $e\tau_h$ (b) $\mu\tau_h$

Figure 7.6: The $m_{\tau\tau}^{\text{MMC}}$ distribution in the $Z \to \tau\tau$ embedding normalization control region for the (a) electron and (b) muon channels after application of the embedding normalization factors with 20.3 fb^{-1} of 8 TeV data.

7 TeV		8 TeV	
$\tau \to e\nu_\tau\overline{\nu_e}$	$\tau \to \mu\nu_\tau\overline{\nu_\mu}$	$\tau \to e\nu_\tau\overline{\nu_e}$	$\tau \to \mu\nu_\tau\overline{\nu_\mu}$
1.02	0.70	0.62	0.53

Table 7.1: Embedding normalization factors in the $H \to \tau_l\tau_h$ analysis.

7.3.2.1 $Z \to \tau\tau$ embedding normalization

The resulting $Z \to \tau\tau$ embedding sample must be normalized to data. For modeling checks, this normalization is computed in a special control region defined relative to preselection by requiring 40 GeV $< m_{\tau\tau}^{\text{vis}} <$ 70 GeV. The $m_{\tau\tau}^{\text{MMC}}$ distributions in this control region for electron and muon channels with normalization factors applied for 8 TeV samples are shown in Figure 7.6.

The normalization factors derived for the $Z \to \tau\tau$ embedding background depend on the collision energy and type of lepton from the τ_l decay and are given in Table 7.1. For the final fit, the normalization of the $Z \to \tau\tau$ background was allowed to float freely as a nuisance parameter.

7.3.3 Fake factor method

In the $H \to \tau_l\tau_h$ and $Z_{ll}H \to \tau_l\tau_h$ analyses, the "fake factor" method was used as the primary estimate of the fake jet $\to \tau_h$ background resulting from QCD multijet, $W+\text{jets}$, $Z \to ll$ (jet $\to \tau_h$), and top (jet $\to \tau_h$) processes. The fake factor method is useful in cases where extrapolation from a control region is not possible due to unreliable statistics. This is particularly relevant for the low-statistics signal regions of the main analyses, as well as the low-statistics VH analyses.

The fake factor method uses a selection of τ_h candidates that pass all selection criteria for a given region but fail τ_h identification. These objects are referred to as "anti-τ_h" objects. The exact criteria for identifying these anti-τ_h objects are somewhat arbitrary, but there are two goals:

1. Maintain sufficient statistics for reliable extrapolation; and
2. Select objects with sufficiently similar kinematics to the target signal region.

The τ_h identification BDT score provides a continuous selection criterion by which anti-τ_h objects can be identified. Because the flavor composition (and consequently the kinematics) of jets faking τ_h objects is highly dependent on BDT score (see Figure 7.7), careful consideration must be made to meet both of these goals. In all three of these analyses, the anti-τ_h definition is taken to be those τ_h candidates with a BDT score above 0.7 × the loose BDT score working point and below the medium BDT score threshold.

The selected anti-τ_h objects are then extrapolated into the signal region by means of the eponymous fake factor:

$$n^{\text{SR}}_{\text{fake-}\tau_h} = (n^{\text{SR,data}}_{\text{anti-}\tau_h} - n^{\text{SR,others}}_{\text{anti-}\tau_h}) \times \text{FF}, \tag{7.2}$$

where

$$\text{FF} = \frac{n^{\text{CR}}_{\text{identified-}\tau_h}}{n^{\text{CR}}_{\text{anti-}\tau_h}}. \tag{7.3}$$

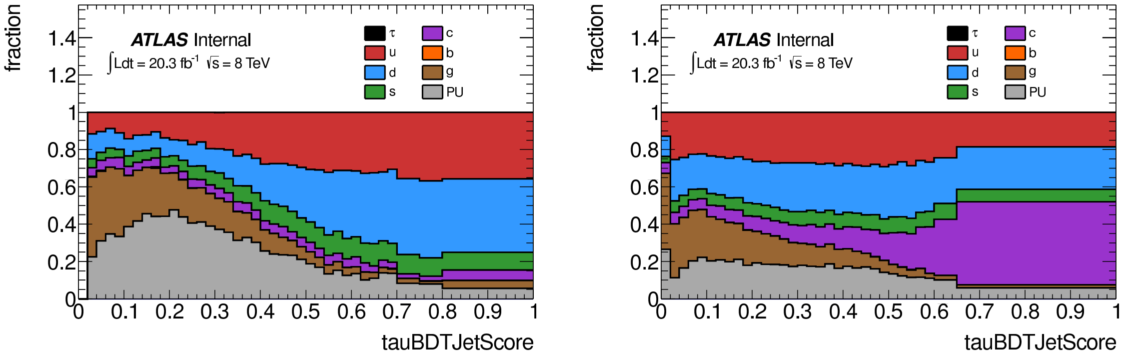

(a) 1-prong anti-τ_h objects (b) 3-prong anti-τ_h objects

Figure 7.7: Fraction of anti-τ_h candidates originating from each parton type in $W+$jets simulated events in the $H \to \tau_l \tau_h$ 0-jet control region [104].

The $n^{\mathrm{SR,others}}_{\mathrm{anti}\text{-}\tau_h}$ component accounts for real τ_h objects accidentally that fail identification. It also includes $l \to \tau_h$ fakes and other backgrounds, which are estimated purely from Monte Carlo or embedding. These backgrounds include $Z \to \tau\tau$, $Z \to ll$ ($l \to \tau_h$), top (non-jet $\to \tau_h$), diboson, and triboson. Events may also contribute more than one anti-τ_h candidate, in which case each anti-τ_h candidate is considered in a separate anti-τ_h event. The control regions used for fake factor derivation are chosen to disallow any contribution of signal processes to $n^{\mathrm{CR}}_{\mathrm{identified}\text{-}\tau_h}$.

The value of the fake factor is dependent on the p_T of the anti-τ_h object, the number of prongs in the anti-τ_h, and the region in which the method is being applied. In the VH analyses, the fake factor is also a function of η_{τ_h}. Furthermore, the value of the fake factor depends on the process generating the anti-τ_h object, and thus the fake factor must be computed as a weighted average of process-specific fake factors in order to maintain the proper ratio of quark and gluon contributions and event kinematics in the anti-τ_h sample. In the VH analyses, anti-τ_h events are largely dominated by $Z \to ll$ (jet $\to \tau_h$) processes, so it is not necessary to calculate a “combined” fake factor. However, in the $H \to \tau_l \tau_h$ analysis, several processes contribute non-negligible numbers of anti-τ_h events, necessitating a combined fake factor.

The combined fake factor is given by

$$\mathrm{FF}(p_T, n_{\mathrm{prong}}, \mathrm{region}) = \sum_{b \in \mathrm{backgrounds}} R_b \cdot \mathrm{FF}_b(p_T, n_{\mathrm{prong}}, \mathrm{region}), \tag{7.4}$$

where R_b is the fraction of anti-τ_h events contributed by the background b and FF_b is the fake factor specific to that background. R_b is calculated using Monte Carlo estimates for the anti-τ_h contribution of $Z \to ll$ (jet $\to \tau_h$) and top (jet $\to \tau_h$) processes. For $W+$jets, this fraction is estimated using a data-driven technique involving the anti-τ_h $W+$jets control region and a Monte Carlo-based extrapolation from the anti-τ_h $W+$jets control region to the anti-τ_h signal region:

$$R_{W+\mathrm{jets}} = \frac{n^{\mathrm{SR},W+\mathrm{jets}}_{\text{anti-}\tau_h}}{n^{\mathrm{SR},W+\mathrm{jets}}_{\text{anti-}\tau_h} + n^{\mathrm{SR,QCD}}_{\text{anti-}\tau_h}}, \tag{7.5}$$

where $n^{\mathrm{SR},W+\mathrm{jets}}_{\text{anti-}\tau_h}$ is estimated by

$$n^{\mathrm{SR},W+\mathrm{jets}}_{\text{anti-}\tau_h} = n^{\mathrm{WCR,data}}_{\text{anti-}\tau_h} \times \frac{n^{\mathrm{SR},W+\mathrm{jets\ MC}}_{\text{anti-}\tau_h}}{n^{\mathrm{WCR},W+\mathrm{jets\ MC}}_{\text{anti-}\tau_h}}, \tag{7.6}$$

and $n^{\mathrm{SR,QCD}}_{\text{anti-}\tau_h}$ is estimated by

$$n^{\mathrm{SR,QCD}}_{\text{anti-}\tau_h} = n^{\mathrm{SR,data}}_{\text{anti-}\tau_h} - (n^{\mathrm{SR},W+\mathrm{jets}}_{\text{anti-}\tau_h} + n^{\mathrm{SR,others}}_{\text{anti-}\tau_h}). \tag{7.7}$$

Here “others” refers to the contribution of $Z \to ll$ (jet $\to \tau_h$) and top (jet $\to \tau_h$). The value of R_{QCD} is taken to make the ratios sum to unity:

$$R_{QCD} = 1 - R_{W+\mathrm{jets}} - R_{\mathrm{top\ (jet}\to\tau_h)} - R_{Z\to ll\ (\mathrm{jet}\to\tau_h)}. \tag{7.8}$$

R_b factors for the $H \to \tau_l \tau_h$ analysis are given in Table 7.2.

The background-specific fake factors, FF_b, are computed using the process-specific control regions defined relative to the signal region being modeled. Fake factors for each of the three analyses are shown in Figures 7.8 and 7.9.

One additional correction was necessary for accurate computation of E_T^{miss} using the fake factor method. Because anti-τ_h objects fail τ_h identification, they are

	7 TeV		8 TeV	
	Boosted	VBF	Boosted	VBF
$R_{W+\text{jets}}$	0.75	0.60	0.62	0.46
R_{QCD}	0.13	0.24	0.26	0.40
$R_{Z\to ll\ (\text{jet}\to\tau_h)}$	0.06	0.13	0.05	0.11
$R_{\text{top}\ (\text{jet}\to\tau_h)}$	0.06	0.03	0.07	0.03

Table 7.2: Fraction of anti-τ_h candidates from processes contributing to the fake τ_h background in the $H \to \tau_l\tau_h$ signal regions [104].

treated as jets when calculating E_T^{miss}. This leads to disagreement with data due to inconsistent calibration and treatment of fake τ_h objects in the E_T^{miss} computation. Consequently, the E_T^{miss} for anti-τ_h objects in the 8 TeV analyses was recomputed using the τ_h treatment. Due to technical constraints, an alternative approach was used in the 7 TeV $H \to \tau_l\tau_h$ analysis, where events were instead reweighted as a function of the projection of E_T^{miss} along the direction of $p_T^{\tau_h}$ normalized by $p_T^{\tau_h}$:

$$\frac{E_T^{\text{miss}}}{p_T^{\tau_h}}\cos(\phi_{\tau_h} - \phi_{E_T^{\text{miss}}}). \tag{7.9}$$

7.3.4 OS – SS method

An alternative data-driven technique for estimating the QCD multijet background contribution was used in some regions of the $H \to \tau_l\tau_h$ analysis, in particular the preselection and the embedding normalization regions. This provided not only an important cross-check of the fake factor method, but also a more stringent test of Monte Carlo modeling. It also allowed bootstrapping of the normalization factor calculations necessary for the fake factor method.

The technique, known as the "OS – SS" method, uses an extrapolation from the same-sign (SS) variation of the signal region to the standard opposite-sign (OS) signal region. No signal events will enter the same-sign region, but it will be rich in QCD multijet events that exhibit little or no charge correlation between the τ decay products. It will also contain some contribution from processes where a real lepton is produced in association with a jet → τ_h fake. In general, the less charge correlation

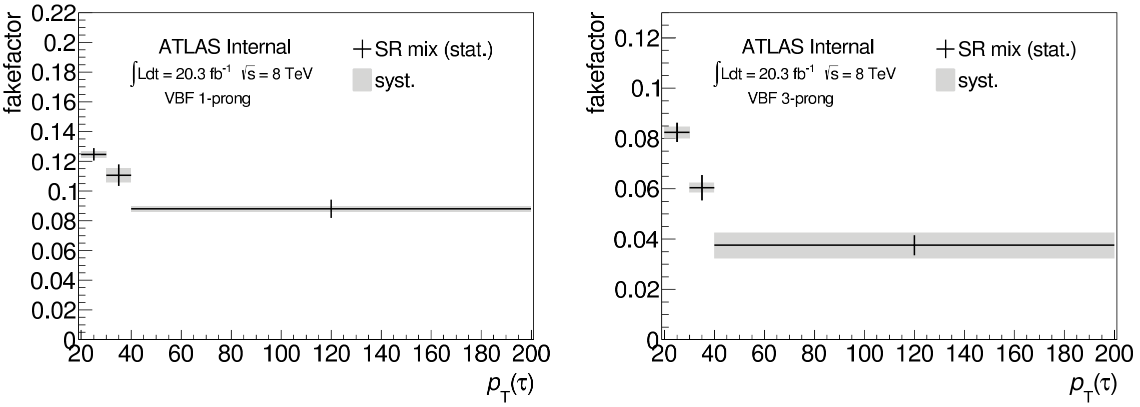

(a) One-prong anti-τ_h, VBF

(b) Three-prong anti-τ_h, VBF

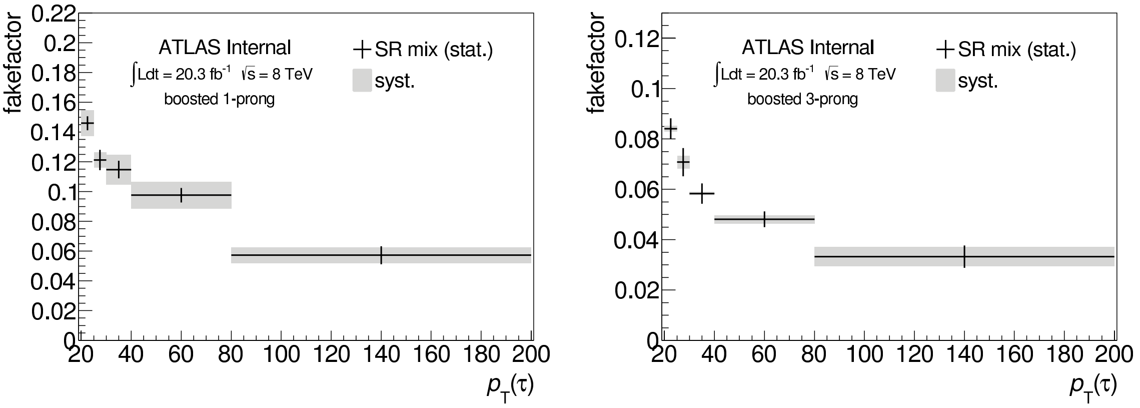

(c) One-prong anti-τ_h, Boosted

(d) Three-prong anti-τ_h, Boosted

Figure 7.8: Fake factors, as a function of anti-τ_h p_T, for (a) one-prong and (b) three-prong anti-τ_h objects in the VBF signal region and (c) one-prong and (d) three-prong anti-τ_h objects in the boosted signal region [104].

between τ decay products for a given background, the larger the component of that background estimated by the same-sign component will be. Any same-sign contribution from a process other than the QCD multijet background is subtracted from the corresponding opposite-sign Monte Carlo or embedding estimate — hence the name "opposite-sign minus same-sign."

The extrapolation factor between the same-sign region and the opposite-sign region, r_{QCD}, is necessary to account for differences in flavor composition leading to differing charge correlation and jet $\rightarrow \tau_h$ fake rates. For example, no charge correlation would be expected for gg events, but significant charge correlation would be

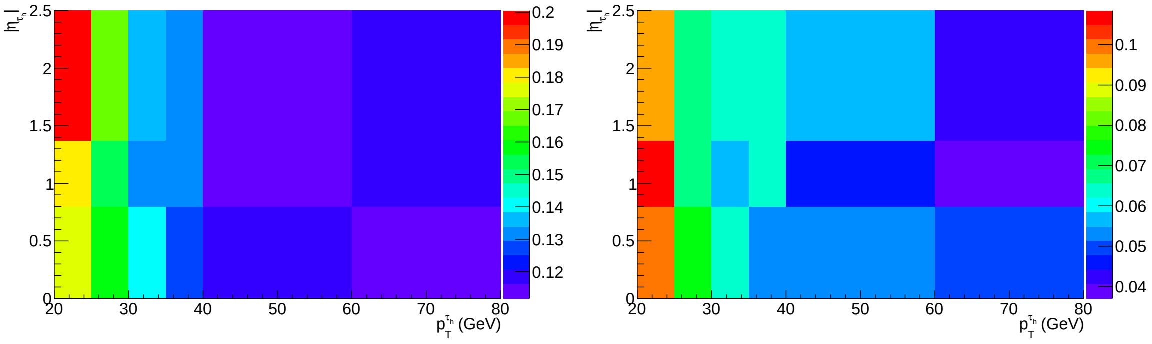

(a) One-prong anti-τ_h

(b) Three-prong anti-τ_h

Figure 7.9: Fake factors, as a function of anti-τ_h p_T and η, for (a) one-prong and (b) three-prong anti-τ_h objects in the ZH analyses.

7 TeV		8 TeV	
$\tau \rightarrow e\nu_\tau\overline{\nu_e}$	$\tau \rightarrow \mu\nu_\tau\overline{\nu_\mu}$	$\tau \rightarrow e\nu_\tau\overline{\nu_e}$	$\tau \rightarrow \mu\nu_\tau\overline{\nu_\mu}$
1.0	1.2	1.0	1.1

Table 7.3: Values of r_{QCD} in the $H \rightarrow \tau_l\tau_h$ analysis.

expected for $q\bar{q}$ events. The differing τ_h identification probability of gluon-initiated jets and quark-initiated jets also leads to some imbalance. The value of r_{QCD} is measured in a dedicated QCD control region and depends on the collision energy and the type of lepton from the τ_l decay:

$$r_{\mathrm{QCD}} = \frac{n_{\mathrm{OS}}^{\mathrm{QCDCR}}}{n_{\mathrm{SS}}^{\mathrm{QCDCR}}}. \tag{7.10}$$

Values of r_{QCD} are given in Table 7.3.

The key assumption of the OS − SS method is that the shape of the QCD multijet background distributions will be the same between the same-sign region and the opposite-sign region. Additionally, it is assumed that the extrapolation factor between the same-sign and opposite-sign regions is the same between the QCD control region in which it is measured and the signal region in which it is applied.

Once r_{QCD} is measured, the complete background model can be estimated by

$$n_{\mathrm{OS}} = r_{\mathrm{QCD}} \cdot n_{\mathrm{SS}}^{\mathrm{data}} + \sum_{b \in \mathrm{backgrounds}} n_{\mathrm{OS-SS}}^{b}, \tag{7.11}$$

where

$$n^b_{\mathrm{OS-SS}} = k^b_{\mathrm{OS}} \cdot n^b_{\mathrm{OS}} - r_{\mathrm{QCD}} \cdot k^b_{\mathrm{SS}} \cdot n^b_{\mathrm{SS}}. \tag{7.12}$$

The k-factors in Equation 7.12 are applied to some Monte Carlo and embedded samples to adjust for incorrectly modeled event rates and τ_h fake rates. These k-factors are sign-dependent due to differences in the flavor composition of the opposite-sign and same-sign regions. Their values also depend on the collision energy and type of lepton resulting from the τ_l decay. In the case of $W + \mathrm{jets}$, they are also region-dependent. The k-factors are derived in their respective process-specific control regions by normalizing Monte Carlo yields to data. For $W + \mathrm{jets}$, k-factors for a given selection must be derived in the $W + \mathrm{jets}$ control region defined relative to that selection. All other k-factors are derived in control regions defined relative to preselection. k-factors are derived first for $W + \mathrm{jets}$, then top, and finally $Z \to ll$ (jet $\to \tau_h$). This order allows $W + \mathrm{jets}$ to be modeled accurately in the top control region. For $Z \to ll$ (jet $\to \tau_h$), no difference in event count is expected between the opposite-sign and same-sign regions, and hence $k_{\mathrm{OS}}^{Z\to ll\ (\mathrm{jet}\to\tau_h)}$ and $k_{\mathrm{SS}}^{Z\to ll\ (\mathrm{jet}\to\tau_h)}$ are the same. For all other backgrounds, the Monte Carlo or embedding modeling are assumed to be accurate, and the k-factors are set to unity (in the case of $Z \to \tau\tau$, the normalization is already determined by factors applied to the embedding samples, as described in Section 7.3.2.1). The values of the k-factors used in the $\mathrm{OS} - \mathrm{SS}$ method for the 7 TeV and 8 TeV BDT-based analyses are presented in Table 7.4.

		7 TeV				8 TeV			
Background	Region	$\tau \rightarrow e\nu_\tau\overline{\nu_e}$		$\tau \rightarrow \mu\nu_\tau\overline{\nu_\mu}$		$\tau \rightarrow e\nu_\tau\overline{\nu_e}$		$\tau \rightarrow \mu\nu_\tau\overline{\nu_\mu}$	
		k_{OS}	k_{SS}	k_{OS}	k_{SS}	k_{OS}	k_{SS}	k_{OS}	k_{SS}
$Z \rightarrow \tau\tau$	All	1.0							
$W + \mathrm{jets}$	Preselection	0.61	0.89	0.59	0.74	0.85	1.02	0.80	0.94
	$Z \rightarrow \tau\tau$ CR	0.61	0.89	0.59	0.74	0.85	1.02	0.80	0.94
	Top CR	0.69	0.89	0.66	0.79	0.95	1.13	0.91	0.99
$Z \rightarrow ll$ (jet $\rightarrow \tau_h$)	All	0.63				0.83			
$Z \rightarrow ll$	All	1.0							
Top	All	2.05	3.70	1.87	0.73	0.90	1.18	0.84	0.99
Diboson	All	1.0							

Table 7.4: k-factors for OS − SS background estimation in the $H \rightarrow \tau_l\tau_h$ analysis.

Chapter 8

Systematic uncertainties

A variety of experimental, theoretical, and modeling uncertainties are considered in the $H \to \tau\tau$ analyses. These uncertainties enter into the final result as nuisance parameters in the fit model. This chapter provides an overview of the most significant systematic uncertainties affecting the $H \to \tau_l\tau_h$ and $Z_{ll}H \to \tau_l\tau_h$ analyses.

8.1 Experimental uncertainties

Most experimental uncertainties that affect the shapes and normalizations of distributions used in the $H \to \tau\tau$ analyses are due to the calibration and finite resolution of energy scales in the ATLAS detector. Uncertainties arising in luminosity and tag-and-probe efficiency measurements also create normalization uncertainties for some Monte Carlo backgrounds.

8.1.1 Luminosity measurement

The uncertainty on the integrated luminosity is due to a variety of uncertainties in the van der Meer scans described in Section 3.2.6. These include uncertainties on beam alignment, beam vibration, beam stability, bunch consistency and population, modeling, reference luminosity and length calibrations, beam-beam interactions, and dependence on the average number interactions per event [40]. The luminosity uncertainty affects only the normalization of those Monte Carlo backgrounds that are not normalized to data. Its value is determined to be ±1.8% for 7 TeV data and ±2.8% for 8 TeV data.

8.1.2 Lepton energy and momentum resolution

Uncertainties on electron energy and muon momentum and their corresponding resolutions result from the *in situ* calibration techniques using $Z \rightarrow ll$ and $J/\psi \rightarrow ll$ events. These corrections are generally on the order of 1% [41, 47]. They are propagated to the analyses by varying and smearing the electron energy and muon momentum.

8.1.3 Lepton efficiency measurements

The uncertainties that arise in tag-and-probe studies used to measure the efficiencies of lepton triggers, reconstruction, identification, and isolation are propagated to these analyses by varying the corresponding scale factors accordingly. The uncertainties in these tag-and-probe studies are described in their corresponding documentation listed in Sections 4.1 and 4.2. They are due to, among other things, the choice of selection used to define the tag object as well as uncertainties on background estimation in the tag-and-probe sample used for the measurement. These uncertainties lead to normalization effects of $< 2\%$.

8.1.4 Jet energy scale and resolution

Multiple uncertainties arise in the energy calibration of jets, including uncertainties on the *in situ* energy calibration, intercalibration across η bins, pile-up, as well as uncertainties on jet flavor composition and the varying response of the detector to different jet flavors [48]. Determination of the jet energy resolution is performed via two different *in situ* methods, with uncertainties arising from experimental selections and Monte Carlo modeling [105]. Uncertainties were accounted for by varying and smearing jet energy. The combined effects of these uncertainties lead to normalization effects of $< 10\%$, as well as shape changing effects in BDT score and $m_{\tau\tau}^{\mathrm{MMC}}$ distributions.

8.1.5 b-tagging efficiency

Scale factors due to the differing performance of b-tagging algorithms for data and Monte Carlo are derived using dileptonic top pair events [56]. A variety of uncertainties in the b-tagging efficiency measurements due to modeling, pile-up, jets, leptons, and E_T^{miss} lead to normalization variations of $< 1\%$ in the $H \to \tau_l \tau_h$ analysis.

8.1.6 Hadronic τ energy scale and momentum resolution

The uncertainty on the energy scale and resolution of τ_h decays was estimated using two separate methods, both described in [58]. The first method, referred to as the "deconvolution method," works by propagating the uncertainties associated with each of the τ_h decay products to the final τ_h energy scale. The second method is an *in situ* technique based on reconstruction of the visible mass peak in $Z \to \tau_\mu \tau_h$ decays. Both of these methods give consistent results, with uncertainties on the τ_h energy scale ranging from 2–4% depending on the number of tracks associated with the τ_h decay. The momentum resolution for reconstructed τ_h candidates is calculated as the standard deviation of a Gaussian fit to the distribution of $\frac{p^{\mathrm{rec}} - p_{vis}^{true}}{p_{vis}^{true}}$, with its corresponding uncertainty determined by smearing the energy of p_{vis}^{true} and altering the τ_h showering model. These uncertainties were accounted for by varying and smearing τ_h energy. The combined effects of these uncertainties lead to normalization effects of $< 3\%$ as well as shape changing effects in BDT score and $m_{\tau\tau}^{\mathrm{MMC}}$ distributions.

8.1.7 Hadronic τ identification

The dominant uncertainties on the $Z \to \tau\tau$ tag-and-probe methods used to measure τ_h trigger, reconstruction, and identification scale factors are described in [58]. These include uncertainties on the background estimation used to isolate the $Z \to \tau\tau$ sample, the τ_h energy scale and shower model, the underlying event, and the statistics of the samples used. The uncertainties on the scale factors for τ_h identification are 2.5% for one-prong τ_h objects and 4% for three-prong τ_h objects.

8.1.8 E_T^{miss}

The uncertainty on E_T^{miss} is due primarily to the individual uncertainties on the energy scales of the E_T^{miss} components, specifically electrons, hadronically decaying τ leptons, and jets. These uncertainties are propagated through the computation of E_T^{miss} described in Section 4.5.

8.2 Signal theoretical uncertainties

A number of theoretical uncertainties affect the shapes and normalizations of signal distributions in the $H \to \tau\tau$ analyses. These are due to the limited order at which Higgs production cross sections can be calculated, as well as assumptions made when generating the Higgs signal samples.

8.2.1 Higher-order QCD corrections

Uncertainties on the cross sections of ggF, VBF, and VH Higgs boson production were estimated to account for higher-order QCD contributions. In all cases, the uncertainty was estimated by varying the renormalization and factorization scales of Monte Carlo generators by a factor of 2 around the nominal value of $m_H = 125$ GeV. The effect on the cross section was estimated within each signal region by applying approximations of the signal region cuts to generator-level quantities. For the VBF and VH processes, this was accomplished using the VBF@NLO [106] and HAWK [107] Monte Carlo generators, respectively. The effects for VBF and VH production are small, on the order of 2–4%. Corrections to the ggF cross section are more substantial, with next-to-leading-order (NLO) effects increasing the cross section by $\sim 100\%$ and next-to-next-to-leading-order (NNLO) effects further increasing the cross section by $\sim 25\%$ [108]. Due to their size, a more complex technique was used to assess the uncertainties of ggF QCD corrections and correlations across different signal regions in the main analyses.

8.2.1.1 Jet-binned ggF QCD scale uncertainties

The VBF and boosted signal regions of the main analyses impose mutually exclusive selections on the number of jets present in an event, thereby partitioning events into bins of exclusive jet multiplicity. These selections are imposed by cuts on kinematic variables, typically the p_T of jets identified in an event.[1] These cuts impose anticorrelations between the uncertainties of neighboring jet bins due to migration of events between bins. More importantly, however, these cuts also introduce cancellations in the perturbative expansions of exclusive jet bin cross sections. Such cancellations can mask the effects of scale variation on exclusive jet bin cross sections, leading to an underestimation of QCD scale uncertainty.

An example of these cancellations, given in [109], is that of the exclusive 0-jet cross section. This cross section can be expressed as the difference between the total inclusive cross section and the inclusive 1-jet cross section:

$$\sigma_0(p^{\text{cut}}) = \sigma_{\text{total}} - \sigma_{\geq 1}(p^{\text{cut}}). \tag{8.1}$$

Here p represents the kinematic variable used to discriminate jet bins, with separation made at the value p^{cut}. The application of this cut introduces logarithmic terms, $L = \ln(p^{\text{cut}}/m_H)$, in the perturbative expansion of $\sigma_{\geq 1}(p^{\text{cut}})$:[2]

$$\sigma_{\geq 1}(p^{\text{cut}}) \simeq \sigma_B \left[\alpha_s(L^2 + L + 1) + \alpha_s^2(L^4 + L^3 + L^2 + L + 1) + \mathcal{O}(\alpha_s^3 L^6)\right], \tag{8.2}$$

---

[1] The jet requirement in the boosted signal region is actually implicit, a result of the $p_T^H > 100$ GeV requirement, which effectively necessitates a jet to enforce p_T balance.

[2] For values of $p^{\text{cut}} \ll m_H$, these logarithmic terms can be large enough to overcome α_s suppression and dominate the expansion, although this is not problematic for present experimental values.

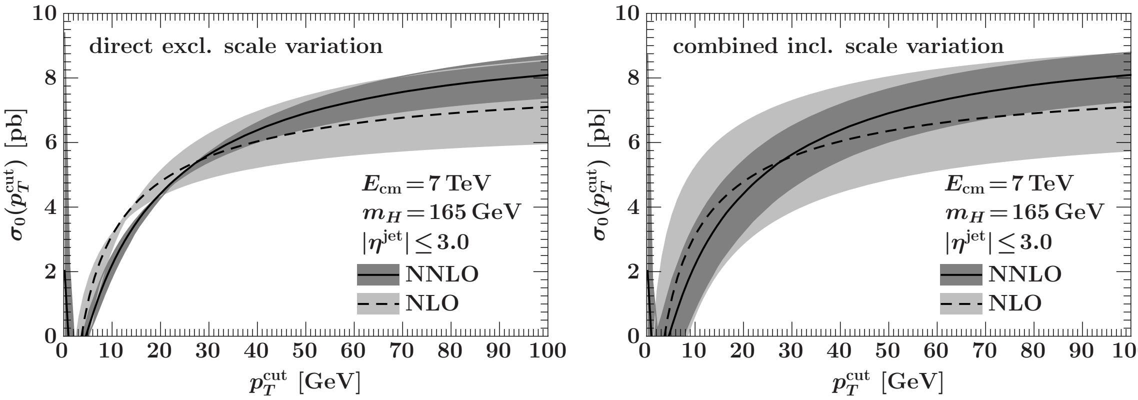

(a) Direct exclusive scale variation

(b) Combined inclusive scale variation

Figure 8.1: The cross section and associated uncertainties for $gg \to H + 0$ jets at NLO and NNLO, with uncertainties calculated via (a) direct exclusive jet bin scale variation and (b) combined inclusive scale variation [110].

where σ_B is the tree-level cross section. When evaluating σ_0 via Equation 8.1, these logarithmic terms can provide partial or complete cancellation with the terms in $\sigma_{\rm total}$:

$$\begin{aligned}\sigma_0(p^{\rm cut}) \simeq \sigma_B\{ &\left[1 + \alpha_s + \alpha_s^2 + \mathcal{O}(\alpha_s^3)\right] \\ &- \left[\alpha_s(L^2 + L + 1) + \alpha_s^2(L^4 + L^3 + L^2 + L + 1) + \mathcal{O}(\alpha_s^3 L^6)\right]\}.\end{aligned} \tag{8.3}$$

These cancellations mask the true uncertainties on the cross section σ_0 when computed by direct exclusive jet bin scale variation. This can be seen in Figure 8.1a, where uncertainties on σ_0 as a function of $p^{\rm cut}$ exhibit a "pinching" effect, becoming completely attenuated at points between 10 GeV and 30 GeV.

Special consideration must be given to correctly assess uncertainties and their anticorrelations for exclusive jet bin cross sections in order to avoid underestimation due to cancellations. The technique used to compute these uncertainties in the main analyses, known as the "Stewart-Tackmann method," is described in [109]. Instead of using direct exclusive scale variation, it computes uncertainties on the exclusive jet bin cross section, σ_n, using the uncertainties on inclusive jet bin cross sections, $\sigma_{\geq n}$ and $\sigma_{\geq n+1}$. The method begins by using the relation

$$\sigma_n = \sigma_{\geq n} - \sigma_{\geq n+1}. \tag{8.4}$$

The key approximation of the method is to assume that uncertainties on $\sigma_{\geq n}$ are uncorrelated for different n:[3]

$$\Delta\sigma_n^2 = \Delta\sigma_{\geq n}^2 + \Delta\sigma_{\geq n+1}^2. \tag{8.5}$$

The effects of this technique can be seen in Figure 8.1b, where the size of uncertainties on σ_0 do not depend on p^{cut}.

In the main analyses, calculation begins by defining the cross section of events in the boosted category as

$$\sigma_{\text{boosted}} = \sigma_{p_T^H > 100 \text{ GeV}} - \sigma_{p_T^H > 100 \text{ GeV\&\&VBF}}, \tag{8.6}$$

where $\sigma_{p_T^H > 100 \text{ GeV}}$ corresponds to those events that pass the boosted signal region selection regardless of the VBF veto and $\sigma_{p_T^H > 100 \text{ GeV\&\&VBF}}$ corresponds to those events that pass both VBF and boosted selections. Using the Stewart-Tackmann method, the uncertainty for σ_{boosted} can be expressed as

$$\Delta\sigma_{\text{boosted}}^2 = \Delta\sigma_{p_T^H > 100 \text{ GeV}}^2 + \Delta\sigma_{p_T^H > 100 \text{ GeV\&\&VBF}}^2. \tag{8.7}$$

Two nuisance parameters are then defined to encode the perturbative ggF uncertainties (and their correlations) for the VBF and boosted signal regions. The first, `QCDScale_ggH1in`, applies to ggF events in the boosted signal region. It encodes the uncertainty associated with the cross section on the inclusive 1-jet cross section, $\sigma_{p_T^H > 100 \text{ GeV}}$. The second, `QCDScale_ggH2in`, applies to ggF events in both the VBF and boosted signal regions. It encodes the uncertainty associated with the inclusive 2-jet cross section, while enforcing the anti-correlation of uncertainties between the VBF and boosted signal regions. In the case of the main cut-based analyses, where a veto is applied on tertiary jets, an additional nuisance parameter, `QCDScale_ggH3in`, is defined by the same method to take into account migrations between the 2-jet and $\geq$3-jet regions.

---

[3] This assumption can lead to overestimation of the uncertainty on σ_n, but it will not lead to an underestimation due to cancellations.

Energy	Region	Cross section × branching ratio × acceptance (pb)	Fractional uncertainty
7 TeV	$\sigma_{p_T^H > 100\ \mathrm{GeV}}$	38.99	27.0%
	$\sigma_{p_T^H > 100\ \mathrm{GeV\&\&VBF}}$	2.14	68.7%
	VBF	9.52	23.9%
8 TeV	$\sigma_{p_T^H > 100\ \mathrm{GeV}}$	54.59	26.3%
	$\sigma_{p_T^H > 100\ \mathrm{GeV\&\&VBF}}$	3.30	67.6%
	VBF	14.02	24.0%

Table 8.1: Cross sections and fractional uncertainties for ggF production of 125 GeV Higgs bosons in 7 TeV and 8 TeV proton-proton collisions calculated via renormalization and factorization scale variation using MCFM 6.8 with dynamic scale calculation for various N_{jets} regions of the $H \to \tau_l \tau_h$ analysis.

Computation of the relevant cross sections were performed using MCFM 6.8 [111] with dynamic scale calculations enabled. To maintain the same order in α_s, the leading-order $H + 2$ jets process was used in combination with the NLO $H + 1$ jet process in the calculation of σ_{boosted} (Equation 8.6). ggF cross sections and uncertainties, which enter into the calculation of these nuisance parameters for the $H \to \tau_l \tau_h$ analysis, are shown in Table 8.1.

8.2.1.2 VBF signal region p_T^H dependence

In the main analyses, where separate BDTs are trained to classify events in the boosted and VBF categories, migration of ggF events from one region to another may affect the shape of the BDT output score distribution. This is an important consideration in the VBF region, where several jet-related variables are used as input to the BDT. The effects of this shape uncertainty are accounted for by reweighting events up and down by the scale uncertainties on ggF production as a function of $\Delta\phi_{Hjj}$, the difference in ϕ between the Higgs and dijet systems, which acts as a metric to constrain N_{jets}.

8.2.2 Higher-order electroweak corrections

To account for contributions from NLO electroweak processes, the VBF and VH samples were both reweighted using corrections factors calculated as a function of p_T^H. These factors were calculated as the difference between the truth-level p_T^H spectrum of POWHEG+PYTHIA and PYTHIA samples and samples generated by HAWK.

8.2.3 Underlying event

Uncertainties related to the modeling of the underlying event and subsequent parton shower in the main analyses were estimated by comparing acceptance of POWHEG+PYTHIA and POWHEG+HERWIG samples for ggF and VBF Higgs production. The resulting normalization uncertainties were 11.2% for the $H \to \tau_l \tau_h$ analysis.

8.2.4 Parton distribution functions

To assess the effects of different parton distribution functions on acceptance of signal events, samples were reweighted at reconstruction level to compare the differences between the nominal values of CT10, MSTW2008NLO [112], and NNPDF [113], as well as the systematic variations of the CT10 distributions. In addition to shape effects on the BDT output score, normalization effects on the order of 3–5% were observed.

8.2.5 $H \to \tau\tau$ branching ratio

The uncertainty on the branching ratio of Higgs bosons to pairs of τ leptons were taken from the LHC Higgs cross section working group [35]. They are mass-dependent, ranging from 3–7%.

8.3 Background modeling uncertainties

Aside from the normalization of the $Z \to \tau\tau$ and top backgrounds, which were left to float in the final fit, several other background modeling nuisance parameters entered the fit model. These include shape systematics associated with the $Z \to \tau\tau$ embedding method and normalization systematics associated with the fake factor method.

8.3.1 Theoretical uncertainties

Additional QCD scale and PDF uncertainties were also accounted for in those background Monte Carlo samples that were not normalized to data, specifically diboson and $Z \to ll$. These uncertainties range from 1–5%.

8.3.2 $Z \to \tau\tau$ embedding

Two separate shape uncertainties were considered for the $Z \to \tau\tau$ embedding procedure: the effect of muon track isolation in the $Z \to \mu\mu$ sample and the subtraction of energy deposits in calorimeter cells.

To estimate the effect of the muon track isolation requirement used in selection of the $Z \to \mu\mu$ sample, two additional $Z \to \tau\tau$ embedding variations were generated. One of these variations removed the muon track isolation requirement while the other used tighter isolation requirements: $I^{\mu}_{p_T}(0.4) < 0.06$ and $I^{\mu}_{E_T}(0.2) < 0.04$.

To estimate the uncertainty of calorimeter cell subtraction, two other $Z \to \tau\tau$ embedding variations were generated with subtraction scaled up and down. The scaling factor used was 0.3 for the 7 TeV analysis and 0.2 for the 8 TeV analysis.

8.3.3 Fake factor method

In the $H \to \tau_l\tau_h$ and VH analyses, there are generally two components that enter into the uncertainty for the fake factor method:

- The statistical uncertainty of the samples used to compute the fake factors; and

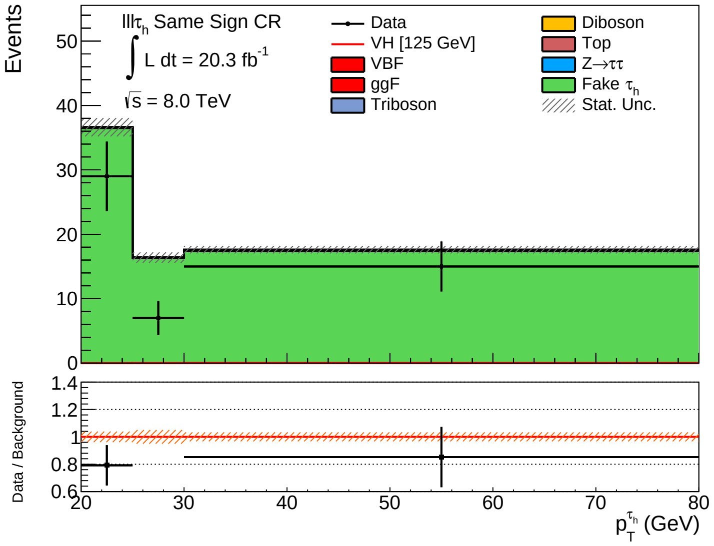

Figure 8.2: $p_T^{\tau_h}$ distributions in the same-sign control region of the $Z_{ll}H \rightarrow \tau_l\tau_h$ analysis with 20.3 fb^{-1} of 8 TeV data.

- The systematic uncertainty on the method itself.

The former is evaluated in all cases by varying the fake factor values within their statistical uncertainty. The latter generally requires consideration of the specific implementation.

In the $H \rightarrow \tau_l\tau_h$ analysis, where the fake factor is computed as the combination of multiple process-specific fake factors, the uncertainty is evaluated by varying the individual R_b up and down by a factor of 2.

In the ZH analyses, where fakes are dominated almost uniquely by contributions from $Z \rightarrow ll$ (jet $\rightarrow \tau_h$) events, an alternative approach is used. In this case, the disagreement with data seen in the fake-dominated same-sign control region, shown in Figure 8.2, is taken as the uncertainty. The disagreement is calculated as a function of leading τ_h p_T in bins of $p_T^{\tau_h}$ corresponding to the binning of the fake factors. High-$p_T^{\tau_h}$ bins above 30 GeV were merged to mitigate large statistical uncertainties due to low bin content.

Additional closure checks using Monte Carlo were made and showed no significant disagreement, except in the 7 TeV $H \to \tau_l \tau_h$ analysis, where a 10% uncertainty on fake backgrounds is assigned in the boosted signal region.

Chapter 9

Results

This chapter presents the results of the statistical tests described in Chapter 6 for the $H \to \tau_l\tau_h$ and $Z_{ll}H \to \tau_l\tau_h$ analyses. Full combination results of the $H \to \tau_l\tau_h$ analysis with the $H \to \tau_h\tau_h$ and $H \to \tau_l\tau_l$ analyses are also shown.

9.1 $H \to \tau_l\tau_h$

The unblinded BDT score distributions for the signal regions of the $H \to \tau_l\tau_h$ analysis are shown in Figure 9.1. These distributions, shown using the best-fit values of background and signal normalizations, display an excess of events with respect to the background-only hypothesis in high BDT score bins, where the majority of the signal contribution is expected.

The 95% confidence limit for exclusion of a $H \to \tau_l\tau_h$ signal, as a function of hypothetical Higgs boson mass, is shown in Figure 9.2a. The observed limits indicated a $\sim 2\sigma$ upward deviation from the expected limit across the mass range under consideration, with an observed (expected) limit of 1.85 (0.89) × the expected Standard Model cross section for $m_H = 125$ GeV.

The expected and observed significance of deviations from the background-only hypothesis as a function of m_H in the $H \to \tau_l\tau_h$ analysis are shown in Figure 9.2b. For $m_H = 125$ GeV, a 2.3σ (2.4σ) deviation from the background-only hypothesis is observed (expected).

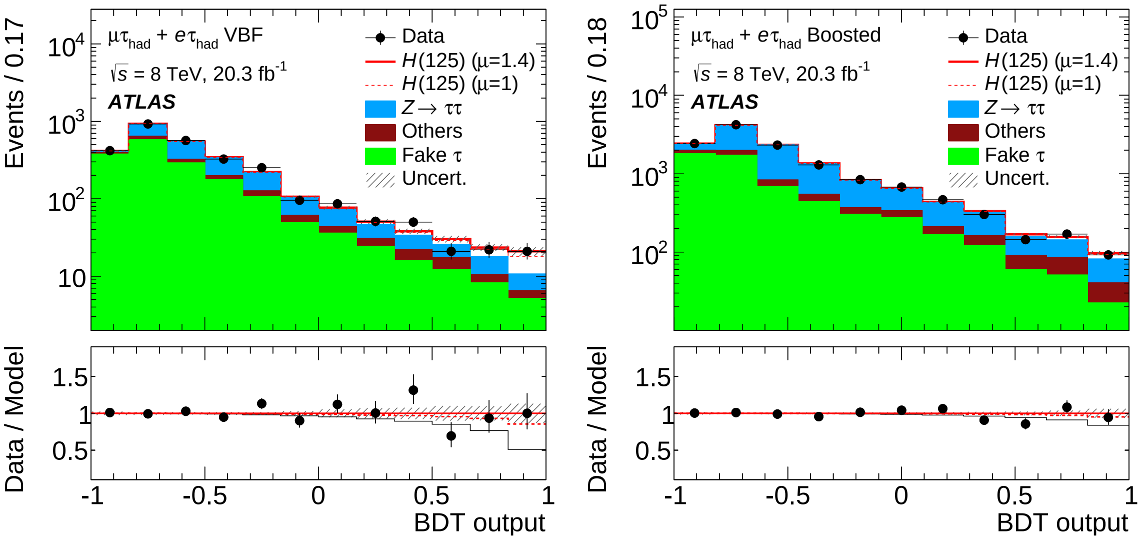

(a) VBF

(b) Boosted

Figure 9.1: The unblinded BDT score distribution in the (a) VBF and (b) boosted signal regions of the $H \to \tau_l \tau_h$ analysis with 20.3 fb^{-1} of 8 TeV data [76]. Background distributions are shown using their best-fit values, while signal distributions are shown using both nominal and best-fit values. BDT score bins in the combined fit are non-uniform, but are resized in these distributions so that their widths are evenly distributed along the BDT score axis.

The most likely values of the signal strength parameter, μ, as a function of m_H, are shown in Figure 9.2c, with a best-fit value of $\mu = 0.98^{+0.50}_{-0.44}$ for $m_H = 125$ GeV.

9.2 Main analyses combination

The combination of the $H \to \tau_h \tau_h$, $H \to \tau_l \tau_h$, and $H \to \tau_l \tau_l$ analyses formed the final results of the $H \to \tau\tau$ analysis effort. The main analyses observed (expected) a deviation from the background-only hypothesis at a significance of 4.5σ (3.4σ). A breakdown of the expected and observed signal significance for each signal region and analysis channel is given in Table 9.1. The most likely value of μ was found to be $1.43^{+0.43}_{-0.37}$ for a hypothetical Higgs boson mass of $m_H = 125$ GeV. The two-dimensional likelihood contours for signal strength in the $(\mu^{\tau\tau}_{\text{ggF}}, \mu^{\tau\tau}_{\text{VBF+VH}})$ plane are shown in Figure 9.3.

A visualization of the combined analysis results is given in Figure 9.4. This vi-

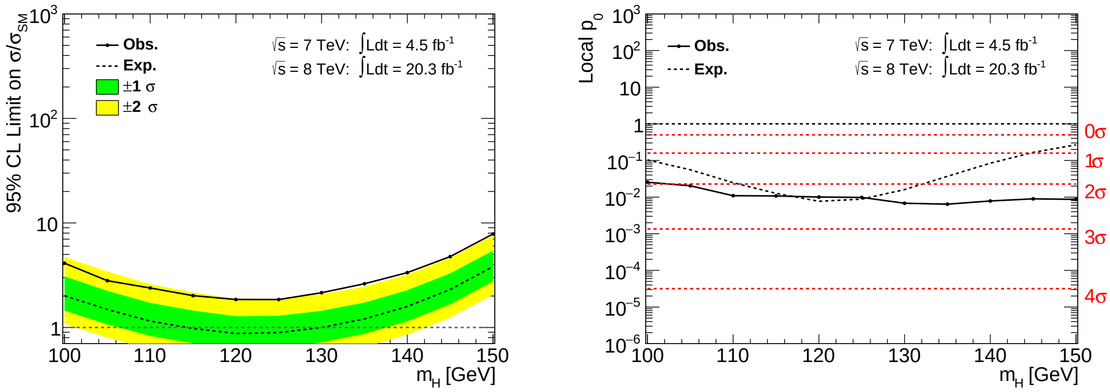

(a) Exclusion limits

(b) Signal significance

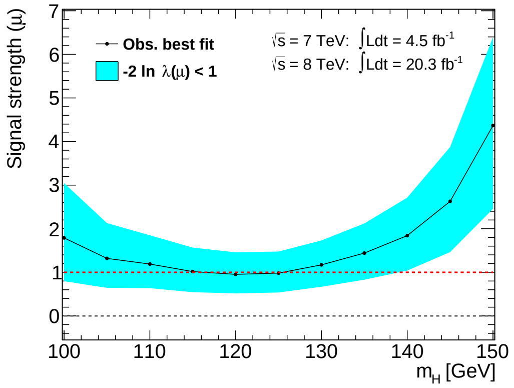

(c) Best-fit μ

Figure 9.2: The results of the $H \to \tau_l \tau_h$ analysis with 4.5 fb^{-1} of 7 TeV data and 20.3 fb^{-1} of 8 TeV data as a function of hypothetical Higgs boson mass. Figure (a) shows expected and observed limits on $H \to \tau_l \tau_h$ signal strength at the 95% confidence level using the CL$_s$ method. Figure (b) shows expected and observed significance of deviations from the background-only hypothesis. Figure (c) shows the best-fit value of the signal strength parameter μ given the observed data.

sualization shows events plotted as a function of the signal-to-background ratio in their BDT score bin. Data are shown against signal and background distributions calculated by fixing the signal strength parameter to various values.

The cut-based variants of the $H \to \tau_h \tau_h$, $H \to \tau_l \tau_h$, and $H \to \tau_l \tau_l$ analyses were also combined as a cross check of the MVA combination. Although less sensitive than the MVA, the cut-based analyses still observed (expected) a 3.2σ (2.5σ) deviation from the background-only hypothesis for $m_H = 125$ GeV. The best-fit value of the

Channel	Category	Expected Significance (σ)	Observed Significance (σ)
$H \to \tau_h\tau_h$	VBF	1.70	2.23
	Boosted	0.82	2.56
	Total	1.99	3.25
$H \to \tau_l\tau_h$	VBF	2.11	2.23
	Boosted	1.11	1.01
	Total	2.33	2.33
$H \to \tau_l\tau_l$	VBF	1.15	1.88
	Boosted	0.57	1.72
	Total	1.25	2.40
Combined		3.43	4.54

Table 9.1: The expected and observed significances of a $H \to \tau\tau$ signal in each of the main analysis channels and signal regions in the combined fit with 4.5 fb^{-1} of 7 TeV data and 20.3 fb^{-1} of 8 TeV data, adapted from [76].

signal strength parameter in the cut-based analyses was found to be $\mu = 1.43^{+0.55}_{-0.49}$ for $m_H = 125$ GeV, in agreement with the MVA combination. A visualization of the $m_{\tau\tau}^{\mathrm{MMC}}$ distribution in the combined cut-based analyses is shown in Figure 9.5.

9.3 $Z_{ll}H \to \tau_l\tau_h$

The 95% confidence limit for exclusion of a $Z_{ll}H \to \tau_l\tau_h$ signal, as a function of hypothetical Higgs boson mass, is shown in Figure 9.6a. The observed limits do not indicate any significant deviation from the expected limits across the mass range under consideration. For $m_H = 125$ GeV, a limit on signal strength of 9.14 (8.08) $\times$ the expected Standard Model cross section is observed (expected).

The expected and observed significance of deviations from the background-only hypothesis as a function of m_H in the $Z_{ll}H \to \tau_l\tau_h$ analysis are shown in Figure 9.6b. For $m_H = 125$ GeV, a 0.3σ (0.3σ) deviation from the background-only hypothesis is observed (expected).

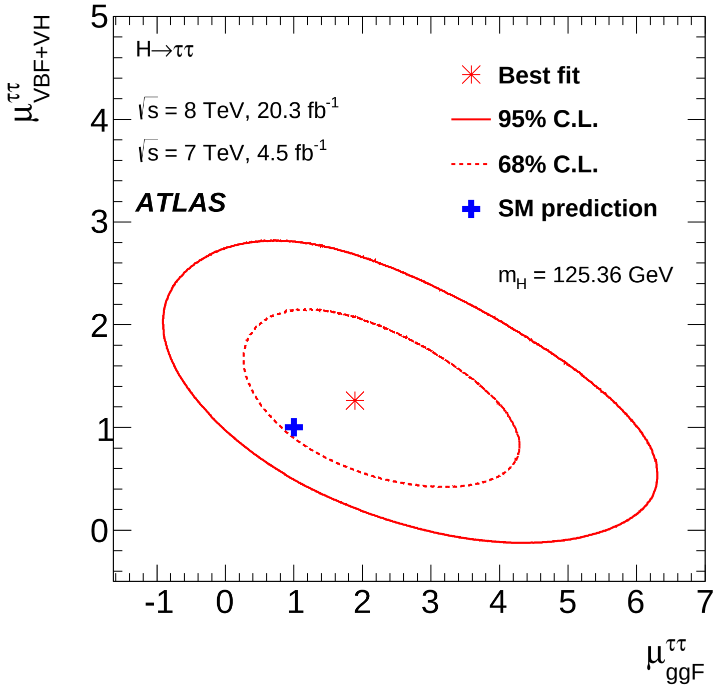

Figure 9.3: The best-fit value of the signal strenth parameter in the $(\mu^{\tau\tau}_{\mathrm{ggF}}, \mu^{\tau\tau}_{\mathrm{VBF+VH}})$ plane, with likelihood contours shown at the 68% (dashed) and 95% (solid) confidence levels and the Standard Model expectation indicated by the blue cross at $(1, 1)$ [76].

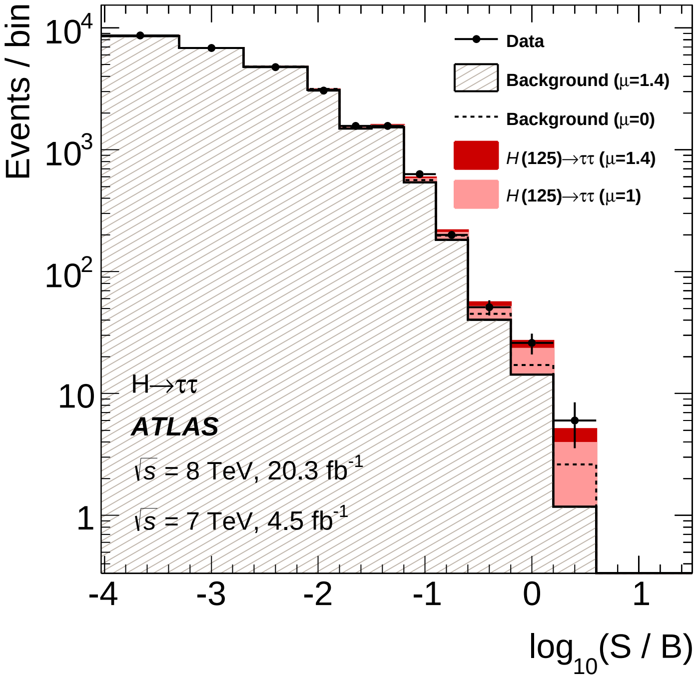

Figure 9.4: Event yields as a function of $\log_{10}(S/B)$, where S and B are the signal and background content, respectively, in the associated BDT score bin [76]. Background and signal contributions are the results of running the fit with different fixed signal strength parameters, including $\mu = 0$, $\mu = 1$, and $\mu = 1.4$.

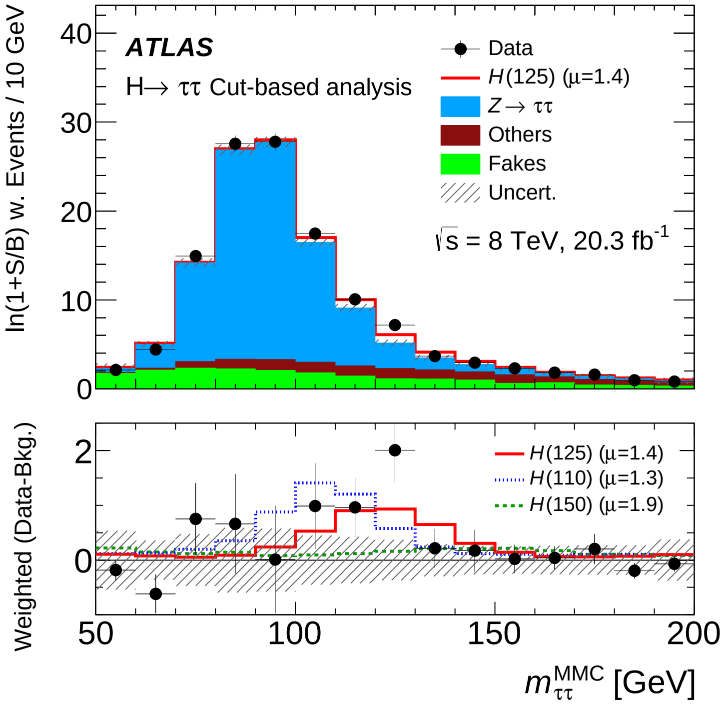

Figure 9.5: The combined $m_{\tau\tau}^{\mathrm{MMC}}$ distribution for $H \to \tau_h\tau_h$, $H \to \tau_l\tau_h$, and $H \to \tau_l\tau_l$ cut-based analyses, with entries weighted by $\ln(1 + S/B)$, where S and B are the signal and background content, respectively, in the signal regions from which entries originate [76]. Background and signal predictions are shown using their best-fit values. The bottom panel shows the difference between weighted data and weighted background entries, with various signal hypotheses shown for different values of m_H, each weighted by their best-fit signal strength.

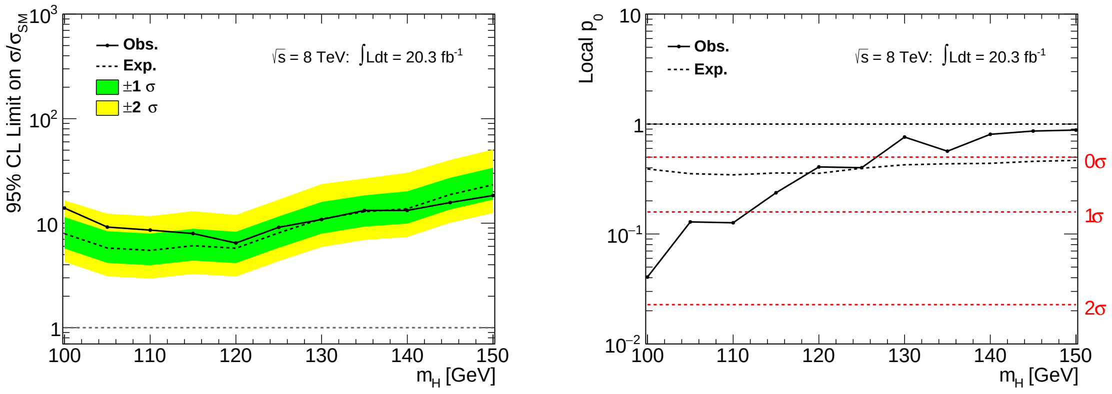

(a) Exclusion limits

(b) Signal significance

Figure 9.6: The results of the $Z_{ll}H \to \tau_l \tau_h$ analysis with 20.3 fb^{-1} of 8 TeV data as a function of hypothetical Higgs boson mass. Figure (a) shows expected and observed limits on $Z_{ll}H \to \tau_l \tau_h$ signal strength at the 95% confidence level using the CL$_s$ method. Figure (b) shows expected and observed significance of deviations from the background-only hypothesis.

Chapter 10

Conclusion

The results of the $H \to \tau\tau$ analyses presented herein have provided strong evidence for the coupling of the newly discovered Higgs boson to τ leptons. While still shy of the 5σ threshold generally regarded as proof of a signal, circumstantial evidence would suggest that the significance of this observation will only grow as more data becomes available. The strength of this signal is also in agreement with the value predicted by the Standard Model. Run II of the LHC will provided an important opportunity to expand these measurements to a larger dataset and improve upon sensitivity in the $H \to \tau\tau$ channel.

In addition, the $Z_{ll}H \to \tau_l\tau_h$ analysis provides some hope that the larger statistics of Run II will provide indication of a VH signal. The fake-dominated nature of this analysis will be an important consideration for Run II, where higher instantaneous luminosity will make τ_h identification even more challenging.

Moving forward, these analyses will ideally expand their range of measurements to include the mass and spin of the Higgs boson, hopefully indicating agreement with results seen in other Higgs decay channels. With a significantly larger dataset, the VH analyses may also provide a secondary probe of HWW and HZZ couplings.

Perhaps the largest challenge awaiting future analyses at the LHC is the computational challenge associated with recording and processing the unprecedented amount of data that will be generated. Serious consideration must be given to software technologies developed outside the sphere of high energy physics. Adaptation, painful as it may be for current analysts, will be essential for the scalability of future analyses.

Nevertheless, Run I of the LHC and the ATLAS experiment provided clear confirmation of an essential and fundamental physical process. With access to higher collision energies in Run II, one might hope not for further confirmation, but rather the indication of new physics through disagreement with Standard Model expectations.

Appendix A

BDT input variables

This appendix contains plots of BDT input variables for the signal regions of the $H \rightarrow \tau_l \tau_h$ analysis. VBF signal region plots are shown in Figures A.1 and A.2. Boosted signal region plots are shown in Figure A.3.

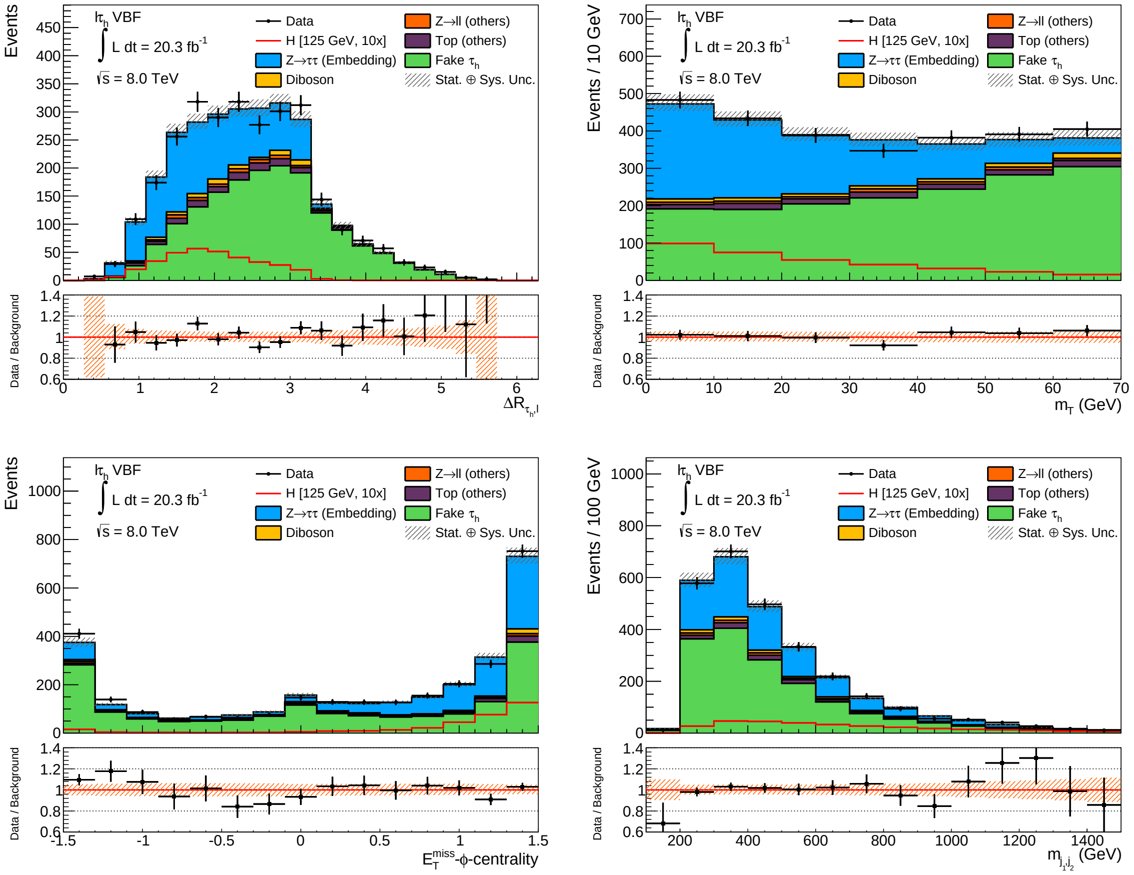

Figure A.1: Various BDT input distributions in the $H \to \tau_l\tau_h$ VBF signal region with 20.3 fb^{-1} of 8 TeV data. These plots are constructed using pre-fit background estimation.

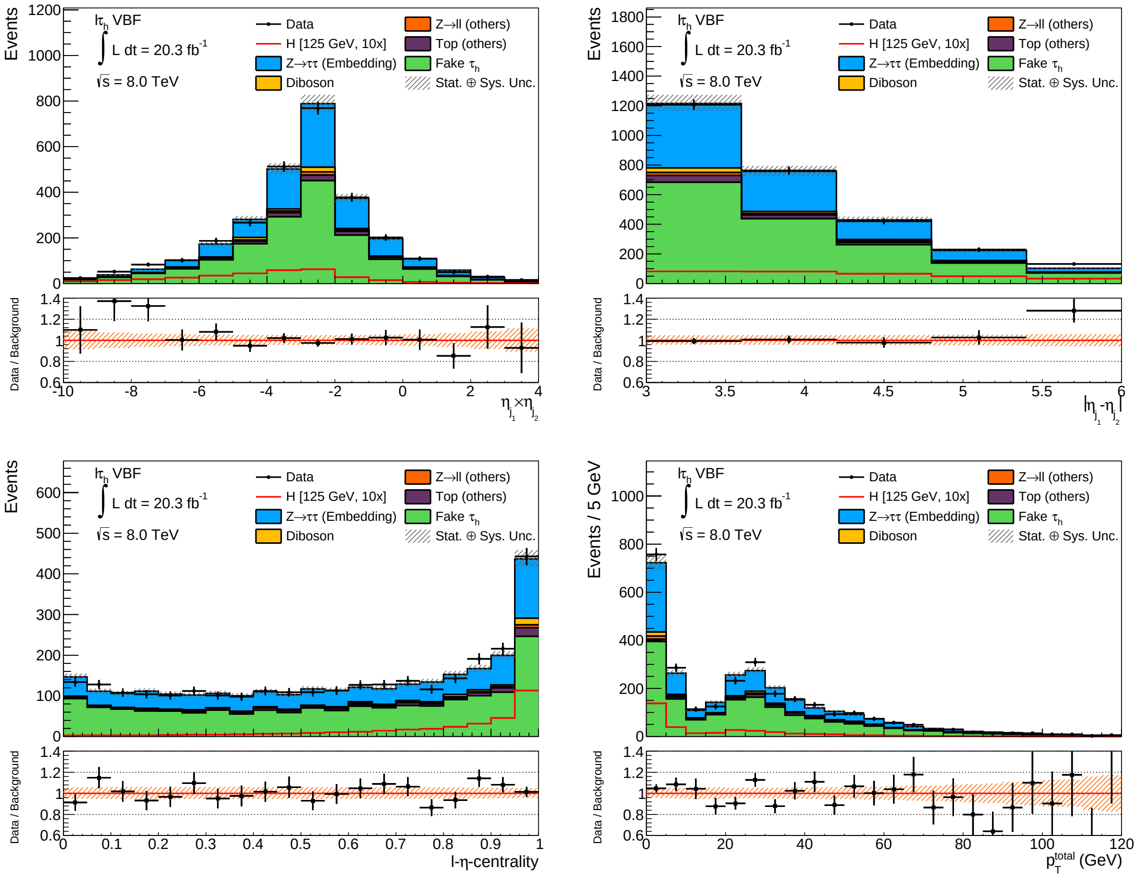

Figure A.2: Various BDT input distributions in the $H \rightarrow \tau_l \tau_h$ VBF signal region with 20.3 fb^{-1} of 8 TeV data. These plots are constructed using pre-fit background estimation.

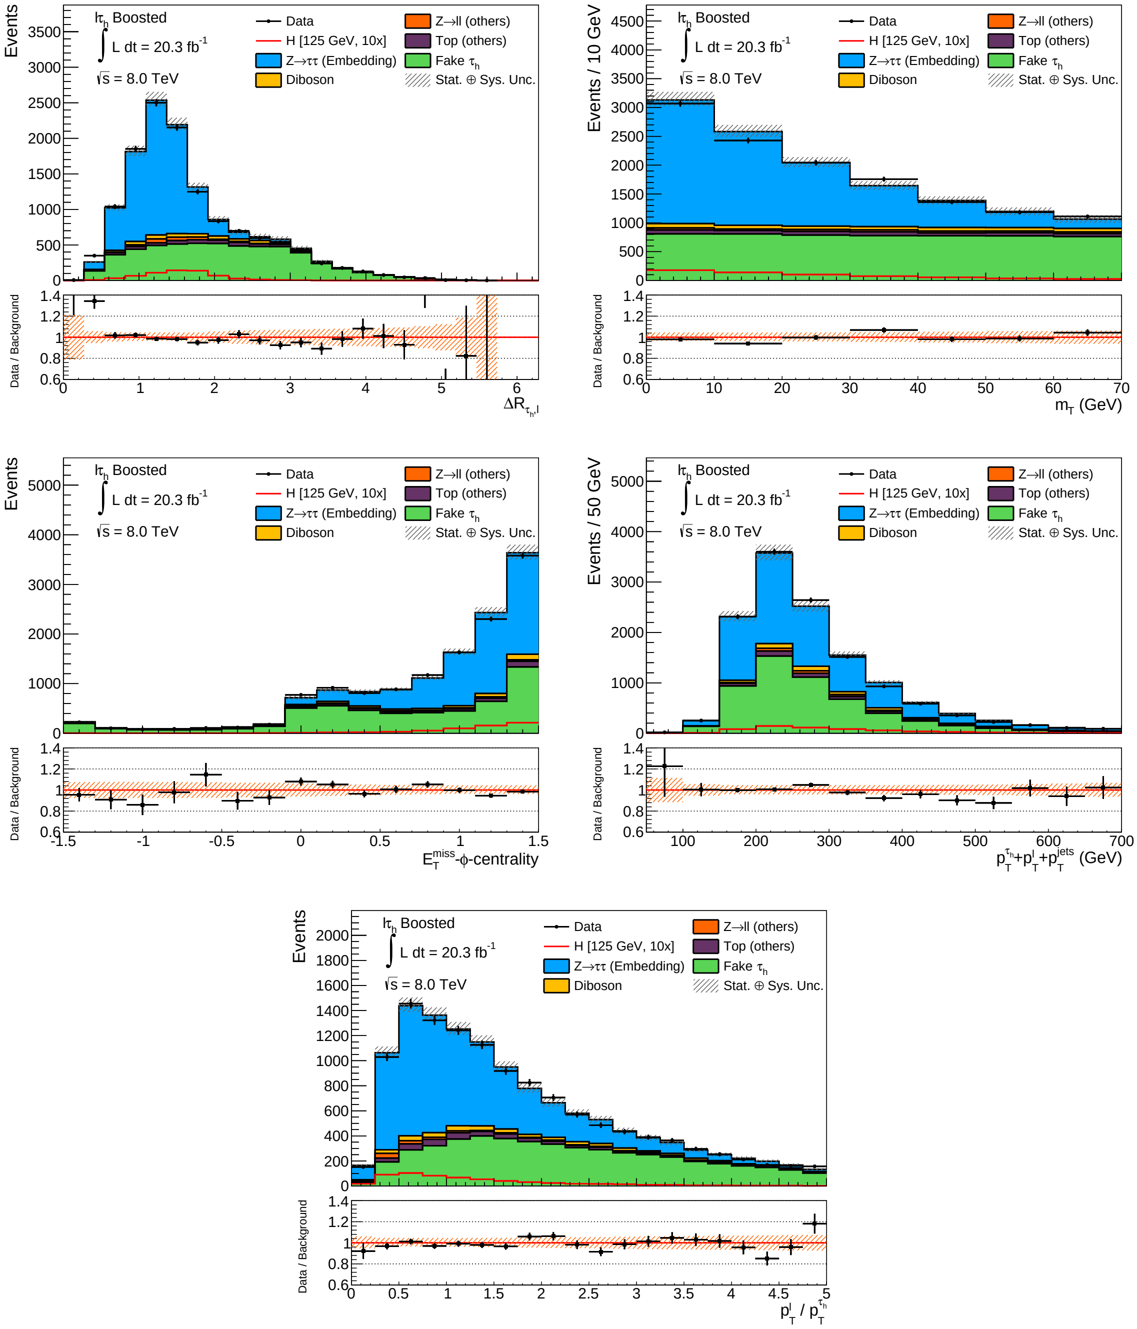

Figure A.3: Various BDT input distributions in the $H \rightarrow \tau_l \tau_h$ boosted signal region with 20.3 fb^{-1} of 8 TeV data. These plots are constructed using pre-fit background estimation.

References

[1] 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. B*, 716(1):1–29, 2012.

[2] 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(1):30–61, 2012.

[3] The ATLAS Collaboration. Measurement of the Higgs boson mass from the $H \to \gamma\gamma$ and $H \to ZZ^* \to 4\ell$ channels with the ATLAS detector using 25 fb^{-1} of pp collision data. 2014.

[4] The ATLAS Collaboration. Evidence for the spin-0 nature of the Higgs boson using ATLAS data. *Phys. Lett. B*, 726(1-3):120–144, 2013.

[5] The ATLAS Collaboration. Updated coupling measurements of the Higgs boson with the ATLAS detector using up to 25 fb^{-1} of proton-proton collision data. In *49th Rencontres de Moriond on QCD and High Energy Interactions*, 2014.

[6] M. E. Peskin and D. V. Schroeder. *An Introduction to Quantum Field Theory.* Addison-Wesley, 1995.

[7] E. Noether. Invariante variationsprobleme. *Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse*, 1918:235–257, 1918.

[8] D. Griffiths. *Introducton to Elementary Particles.* Wiley-VCH, 2004.

[9] G. 't Hooft. Renormalization of massless Yang-Mills fields. *Nucl. Phys. B*, 33(1):173–199, 1971.

[10] G. 't Hooft. Renormalizable Lagrangians for massive Yang-Mills fields. *Nucl. Phys. B*, 35(1):167–188, 1971.

[11] C. Yang and R. Mills. Conservation of isotopic spin and isotopic gauge invariance. *Phys. Rev.*, 96:191–195, 1954.

[12] K. A. Olive et al. Review of particle physics. *Chin. Phys. C*, 38(9):090001, 2014.

[13] The T2K Collaboration. Observation of electron neutrino appearance in a muon neutrino beam. *Phys. Rev. Lett.*, 112:061802, 2014.

[14] J. C. Maxwell. A dynamical theory of the electromagnetic field. *Phil. Trans. R. Soc. Lond.*, 155:459–512, 1865.

[15] C. D. Ellis and W. A. Wooster. The average energy of disintegration of radium e. *Proc. R. Soc. A*, 117(776):109–123, 1927.

[16] W. Pauli. *Institut international de physique Solvay. Structure et propriétés des noyaux atomiques. Rapports et discussions du septième conseil de physique tenu à Bruxelles du 22 au 29 octobre 1933, sous les auspices de l'Institut international de physique Solvay, publiés par la commission administrative de l'Institut*, page 324. impr.-édit. Gauthier-Villars, 55, quai des Grands-Augustins, 1934.

[17] E. Fermi. Versuch einer theorie der β-strahlen. i. *Zeitschrift für Physik*, 88(3-4):161–177, 1934.

[18] T. Lee and C. Yang. Parity nonconservation and a two-component theory of the neutrino. *Phys. Rev.*, 105:1671–1675, 1957.

[19] C. S. Wu, E. Ambler, R. Hayward, D. Hoppes, and R. Hudson. Experimental test of parity conservation in beta decay. *Phys. Rev.*, 105:1413–1415, 1957.

[20] R. Feynman and M. Gell-Mann. Theory of the Fermi interaction. *Phys. Rev.*, 109:193–198, 1958.

[21] J. Schwinger. A theory of the fundamental interactions. *Ann. Phys.*, 2(5):407–434, 1957.

[22] J. Goldstone. Field theories with superconductor solutions. *Il Nuovo Cimento*, 19(1):154–164, 1961.

[23] P. W. Higgs. Broken symmetries, massless particles and gauge fields. *Phys. Lett.*, 12:132–133, 1964.

[24] P. W. Higgs. Broken symmetries and the masses of gauge bosons. *Phys. Rev. Lett.*, 13:508–509, 1964.

[25] F. Englert and R. Brout. Broken symmetry and the mass of gauge vector mesons. *Phys. Rev. Lett.*, 13:321–323, 1964.

[26] G. Guralnik, C. Hagen, and T. Kibble. Global conservation laws and massless particles. *Phys. Rev. Lett.*, 13:585–587, 1964.

[27] S. Weinberg. A model of leptons. *Phys. Rev. Lett.*, 19:1264–1266, 1967.

[28] A. Salam. *Elementary Particle Theory.* Almqvist and Wiksell, 1968.

[29] F. Halzen and A. D. Martin. *Quarks & Leptons: An Introductory Course in Modern Particle Physics.* John Wiley & Sons, Inc., 1984.

[30] The LEP Electroweak Working Group. The LEP electroweak working group.

[31] The OPAL Collaboration. OPAL event displays.

[32] R. Devenish and A. Cooper-Sarkar. *Deep Inelastic Scattering*. Oxford University Press, 2004.

[33] The ATLAS Collaboration. *Expected performance of the ATLAS experiment: detector, trigger and physics*. CERN, Geneva, 2009.

[34] The CMS Collaboration. CMS physics: Technical design report volume 2: Physics performance. *J. Phys. G*, 34(CERN-LHCC-2006-021. CMS-TDR-8-2):995–1579. 669 p, 2007.

[35] LHC Higgs Cross Section Working Group, S. Heinemeyer, C. Mariotti, G. Passarino, and R. Tanaka (Eds.). Handbook of LHC Higgs cross sections: 3. Higgs properties. *CERN-2013-004*, CERN, Geneva, 2013.

[36] C. Boddy, S. Farrington, and C. Hays. Higgs boson coupling sensitivity at the LHC using $H \to \tau\tau$ decays. *Phys. Rev. D*, 86:073009, 2012.

[37] L. Evans and P. Bryant. LHC machine. *J. Instrum.*, 3(08):S08001, 2008.

[38] J. Haffner. The CERN accelerator complex. 2013.

[39] The ATLAS Collaboration. The ATLAS experiment at the CERN Large Hadron Collider. *J. Instrum.*, 3(08):S08003, 2008.

[40] The ATLAS Collaboration. Improved luminosity determination in pp collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector at the LHC. *Eur. Phys. J. C*, C73(8):2518, 2013.

[41] The ATLAS Collaboration. Electron performance measurements with the ATLAS detector using the 2010 LHC proton-proton collision data. *Eur. Phys. J. C*, C72:1909, 2012.

[42] The ATLAS Collaboration. Electron and photon energy calibration with the ATLAS detector using LHC run 1 data. *Eur. Phys. J. C*, C74(10):3071, 2014.

[43] The ATLAS Collaboration. Electron reconstruction and identification efficiency measurements with the ATLAS detector using the 2011 LHC proton-proton collision data. *Eur. Phys. J. C*, C74(7):2941, 2014.

[44] The ATLAS Collaboration. Electron efficiency measurements with the ATLAS detector using the 2012 LHC proton-proton collision data. Technical Report ATLAS-CONF-2014-032, CERN, Geneva, 2014.

[45] M. Wielers, R. Mantifel, A. Tricoli, and P. Bell. Single electron trigger performance plots. Technical Report ATL-COM-DAQ-2012-146, CERN, Geneva, 2012.

[46] S. Heim, K. Lohwasser, L. Iconomidou-Fayard, G. Pasztor, J. Kretzschmar, and M. Delmastro. Update of electron efficiency plots (September 2013): For approval. Technical Report ATL-COM-PHYS-2013-1287, CERN, Geneva, 2013.

[47] The ATLAS Collaboration. Measurement of the muon reconstruction performance of the ATLAS detector using 2011 and 2012 LHC proton-proton collision data. *Eur. Phys. J. C*, 74(11), 2014.

[48] The ATLAS Collaboration. Jet energy measurement and its systematic uncertainty in proton-proton collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. 2014.

[49] W. Lampl, S. Laplace, D. Lelas, P. Loch, H. Ma, S. Menke, S. Rajagopalan, D. Rousseau, S. Snyder, and G. Unal. Calorimeter clustering algorithms: Description and performance. Technical Report ATL-LARG-PUB-2008-002. ATL-COM-LARG-2008-003, CERN, Geneva, 2008.

[50] T. Barillari, E. Bergeaas Kuutmann, T. Carli, J. Erdmann, P. Giovannini, K. J. Grahn, C. Issever, A. Jantsch, A. Kiryunin, K. Lohwasser, A. Maslennikov, S. Menke, H. Oberlack, G. Pospelov, E. Rauter, P. Schacht, F. Spanó, P. Speckmayer, P. Stavina, and P. Strízenec. Local hadronic calibration. Technical Report ATL-LARG-PUB-2009-001-2. ATL-COM-LARG-2008-006. ATL-LARG-PUB-2009-001, CERN, Geneva, 2008.

[51] S. Agostinelli et al. Geant4 — a simulation toolkit. *Nucl. Instrum. Methods Phys. Res. A*, 506(3):250–303, 2003.

[52] M. Cacciari, G. P. Salam, and G. Soyez. The anti-k_t jet clustering algorithm. *J. High Energy Phys.*, 2008(04):063, 2008.

[53] S. Catani, Y. Dokshitzer, M. Seymour, and B. Webber. Longitudinally invariant k_t-clustering algorithms for hadron-hadron collisions. *Nucl. Phys. B*, 406(1–2):187–224, 1993.

[54] Y. Dokshitzer, G. Leder, S. Moretti, and B. Webber. Better jet clustering algorithms. *J. High Energy Phys.*, 1997(08):001, 1997.

[55] The ATLAS Collaboration. Commissioning of the ATLAS high-performance b-tagging algorithms in the 7 TeV collision data. Technical Report ATLAS-CONF-2011-102, CERN, Geneva, 2011.

[56] The ATLAS Collaboration. Calibration of b-tagging using dileptonic top pair events in a combinatorial likelihood approach with the ATLAS experiment. Technical Report ATLAS-CONF-2014-004, CERN, Geneva, 2014.

[57] The ATLAS Collaboration. Calibration of the performance of b-tagging for c and light-flavour jets in the 2012 ATLAS data. Technical Report ATLAS-CONF-2014-046, CERN, Geneva, 2014.

[58] The ATLAS Collaboration. Identification and energy calibration of hadronically decaying tau leptons with the ATLAS experiment in pp collisions at $\sqrt{s} = 8$ TeV. 2014.

[59] The ATLAS Collaboration. Performance of missing transverse momentum reconstruction in ATLAS studied in proton-proton collisions recorded in 2012 at 8 TeV. Technical Report ATLAS-CONF-2013-082, CERN, Geneva, 2013.

[60] C. P. Bee, D. Francis, L. Mapelli, R. McLaren, G. Mornacchi, J. Petersen, and F. J. Wickens. The raw event format in the ATLAS trigger & DAQ. Technical Report ATL-DAQ-98-129, CERN, Geneva, 1998.

[61] S. Armstrong, J. T. Baines, et al. Algorithms for the ATLAS high-level trigger. *IEEE Trans. Nucl. Sci.*, 51(3):367–374, 2004.

[62] M. Rovatsou. Acceleration analysis on GPUs for the ATLAS experiment at CERN - SIMT design of the high level trigger Kalman fitter. Master's thesis, University of Edinburgh, 2010.

[63] C. Jones. Porting the z-finder algorithm to GPU. Master's thesis, University of Edinburgh, 2010.

[64] V. Breton, N. Brun, and P. Perret. A clustering algorithm for the LHCb electromagnetic calorimeter using a cellular automaton. 2001.

[65] D. Emeliyanov and J. Howard. GPU-based tracking algorithms for the ATLAS high-level trigger. *J. Phys.: Conf. Ser.*, 396(1):012018, 2012.

[66] G. M. Amdahl. Validity of the single processor approach to achieving large scale computing capabilities. In *Proceedings of the April 18-20, 1967, Spring Joint Computer Conference*, AFIPS '67 (Spring), pages 483–485, New York, NY, USA, 1967. ACM.

[67] Mozilla Developer Network. *SpiderMonkey.* Mozilla, https://developer.mozilla.org/en-US/docs/Mozilla/Projects/SpiderMonkey/Internals, 2014.

[68] Google. *V8 JavaScript Engine.* Google, https://developers.google.com/v8/design, 2014.

[69] The WebKit Open Source Project. *JavaScriptCore.* http://trac.webkit.org/wiki/JavaScriptCore, 2014.

[70] The PyPy Team. *PyPy.* http://pypy.org/performance.html, 2014.

[71] The ROOT Team. *ROOT: An Object-Oriented Data Analysis Framework.* CERN, http://root.cern.ch/drupal/content/users-guide, 5.34.00 edition, 2013.

[72] The Numexpr Team. *Numexpr.* https://github.com/pydata/numexpr, 2014.

[73] S. van der Walt, S. Colbert, and G. Varoquaux. The NumPy array: A structure for efficient numerical computation. *Comput. Sci. Eng.*, 13(2):22–30, 2011.

[74] rootpy developers and contributors. *root_numpy.* http://rootpy.github.io/root_numpy/index.html, 2014.

[75] F. Alted, I. Vilata, et al. PyTables: Hierarchical datasets in Python, 2002–.

[76] The ATLAS Collaboration. Evidence for the Higgs-boson Yukawa coupling to tau leptons with the ATLAS detector. 2015.

[77] A. Barr, T. Khoo, P. Konar, K. Kong, C. Lester, K. Matchev, and M. Park. Guide to transverse projections and mass-constraining variables. *Phys. Rev. D*, 84:095031, 2011.

[78] G. Cowan, K. Cranmer, E. Gross, and O. Vitells. Asymptotic formulae for likelihood-based tests of new physics. *Eur. Phys. J. C*, 71(2), 2011.

[79] A. Wald. Tests of statistical hypotheses concerning several parameters when the number of observations is large. *Trans. Amer. Math. Soc.*, 54(3):426–482, 1943.

[80] G. Zech. Upper limits in experiments with background or measurement errors. *Nucl. Instrum. Methods Phys. Res. A*, 277(CERN-EP-88-164):608. 6 p, 1988.

[81] A. L. Read. Presentation of search results: the CL_s technique. *J. Phys. G*, 28(10):2693, 2002.

[82] K. Cranmer, G. Lewis, L. Moneta, A. Shibata, and W. Verkerke. HistFactory: A tool for creating statistical models for use with RooFit and RooStats. 2012.

[83] A. Elagin, P. Murat, A. Pranko, and A. Safonov. A new mass reconstruction technique for resonances decaying to $\tau\tau$. *Nucl. Instrum. Methods Phys. Res. A*, 654(1):481–489, 2011.

[84] R. Ellis, I. Hinchliffe, M. Soldate, and J. V. D. Bij. Higgs decay to $\tau^+\tau^-$: A possible signature of intermediate mass Higgs bosons at high energy hadron colliders. *Nucl. Phys. B*, 297(2):221–243, 1988.

[85] L. Breiman, J. Friedman, R. Olshen, and C. Stone. *Classification and Regression Trees.* Wadsworth and Brooks, Monterey, CA, 1984.

[86] Y. Freund and R. E. Schapire. A decision-theoretic generalization of on-line learning and an application to boosting. *J. Comput. Syst. Sci.*, 55(1):119–139, 1997.

[87] J. H. Friedman. Stochastic gradient boosting. *Comput. Stat. Data Anal.*, 38(4):367–378, 2002.

[88] A. Hoecker, P. Speckmayer, J. Stelzer, J. Therhaag, E. von Toerne, and H. Voss. TMVA: Toolkit for multivariate data analysis. *PoS*, ACAT:040, 2007.

[89] S. Alioli, P. Nason, C. Oleari, and E. Re. NLO Higgs boson production via gluon fusion matched with shower in POWHEG. *J. High Energy Phys.*, 2009(04):002, 2009.

[90] T. Sjöstrand, S. Mrenna, and P. Skands. PYTHIA 6.4 physics and manual. *J. High Energy Phys.*, 2006(05):026, 2006.

[91] T. Sjöstrand, S. Mrenna, and P. Skands. A brief introduction to PYTHIA 8.1. *Comput. Phys. Commun.*, 178(11):852–867, 2008.

[92] M. L. Mangano, F. Piccinini, A. D. Polosa, M. Moretti, and R. Pittau. ALPGEN, a generator for hard multiparton processes in hadronic collisions. *J. High Energy Phys.*, 2003(07):001, 2003.

[93] S. Frixione and B. R. Webber. Matching NLO QCD computations and parton shower simulations. *J. High Energy Phys.*, 2002(06):029, 2002.

[94] B. P. Kersevan and E. Richter-Wąs. The Monte Carlo event generator AcerMC versions 2.0 to 3.8 with interfaces to PYTHIA 6.4, HERWIG 6.5 and ARIADNE 4.1. *Comput. Phys. Commun.*, 184(3):919–985, 2013.

[95] T. Binoth, M. Ciccolini, N. Kauer, and M. Krämer. Gluon-induced W boson pair production at the LHC. *J. High Energy Phys.*, 2006(12):046, 2006.

[96] G. Corcella, I. G. Knowles, G. Marchesini, S. Moretti, K. Odagiri, P. Richardson, M. H. Seymour, and B. R. Webber. HERWIG 6: an event generator for hadron emission reactions with interfering gluons (including supersymmetric processes). *J. High Energy Phys.*, 2001(01):010, 2001.

[97] The ATLAS Collaboration. Observation and measurement of Higgs boson decays to WW^* with the ATLAS detector. 2014.

[98] J. Butterworth, J. Forshaw, and M. Seymour. Multiparton interactions in photoproduction at HERA. *Zeitschrift für Physik C: Particles and Fields*, 72(4):637–646, 1996.

[99] J. Alwall, M. Herquet, F. Maltoni, O. Mattelaer, and T. Stelzer. MadGraph 5: going beyond. *J. High Energy Phys.*, 2011(6), 2011.

[100] H.-L. Lai, M. Guzzi, J. Huston, Z. Li, P. M. Nadolsky, J. Pumplin, and C.-P. Yuan. New parton distributions for collider physics. *Phys. Rev. D*, 82:074024, 2010.

[101] J. Pumplin, D. R. Stump, J. Huston, H.-L. Lai, P. Nadolsky, and W.-K. Tung. New generation of parton distributions with uncertainties from global QCD analysis. *J. High Energy Phys.*, 2002(07):012, 2002.

[102] Z. Wąs. TAUOLA the library for τ lepton decay, and KKMC/KORALB/KORALZ/… status report. *Nucl. Phys. B (Proc. Suppl.)*, 98(1–3):96–102, 2001.

[103] N. Davidson, T. Przedzinski, and Z. Wąs. PHOTOS interface in C++: Technical and physics documentation. 2010.

[104] D. Alvarez Piqueras et al. Measurement of the Higgs boson couplings in the $\tau\tau$ final state with the ATLAS detector - supporting note. Technical Report ATL-COM-PHYS-2014-170, CERN, Geneva, 2014.

[105] The ATLAS Collaboration. Jet energy resolution in proton-proton collisions at $\sqrt{s} = 7$ TeV recorded in 2010 with the ATLAS detector. *Eur. Phys. J. C*, C73(3):2306, 2013.

[106] K. Arnold, J. Bellm, G. Bozzi, M. Brieg, F. Campanario, et al. VBFNLO: A parton level Monte Carlo for processes with electroweak bosons – manual for version 2.5.0. 2011.

[107] A. Denner, S. Dittmaier, S. Kallweit, and A. Mück. HAWK 2.0: A Monte Carlo program for Higgs production in vector-boson fusion and Higgs strahlung at hadron colliders. 2014.

[108] S. Catani, D. de Florian, M. Grazzini, and P. Nason. Soft-gluon resummation for Higgs boson production at hadron colliders. *J. High Energy Phys.*, 2003(07):028, 2003.

[109] I. W. Stewart and F. J. Tackmann. Theory uncertainties for Higgs and other searches using jet bins. *Phys. Rev. D*, D85:034011, 2012.

[110] LHC Higgs Cross Section Working Group, S. Dittmaier, C. Mariotti, G. Passarino, and R. Tanaka (Eds.). Handbook of LHC Higgs cross sections: 2. differential distributions. *CERN-2012-002*, CERN, Geneva, 2012.

[111] J. M. Campbell, R. K. Ellis, and G. Zanderighi. Next-to-leading order Higgs + 2 jet production via gluon fusion. *J. High Energy Phys.*, 0610:028, 2006.

[112] A. Martin, W. Stirling, R. Thorne, and G. Watt. Parton distributions for the LHC. *Eur. Phys. J. C*, C63:189–285, 2009.

[113] S. Forte, L. Garrido, J. I. Latorre, and A. Piccione. Neural network parametrization of deep inelastic structure functions. *J. High Energy Phys.*, 0205:062, 2002.