UNIVERSITY OF CALIFORNIA,
IRVINE

Sweet Little Nothings; or, Searching for a Pair of Stops, a Pair of Higgs Bosons, and a Pair of New Small Wheels for the Upgrade of the Forward Muon System of the ATLAS Detector at CERN

DISSERTATION

submitted in partial satisfaction of the requirements
for the degree of

DOCTOR OF PHILOSOPHY

in Physics

by

Daniel Joseph Antrim

Dissertation Committee:
Professor Anyes Taffard, Chair
Professor Daniel Whiteson
Professor Arvind Rajaraman

2019

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

That "we don't accomplish anything in this world alone" is never more true than in the case of attempting a doctorate in physics, especially in the sub-field of experimental high energy particle physics that, nowadays, likely has you stretched across the globe in your work on large-scale experiments operated by international collaborations.

I have had the opportunity to be based at CERN for essentially the entirety of my doctorate. This would not have been possible without the many people who have helped me along the way. Their number is many, too far to count or list here. They know who they are, though, and without them I would not have succeeded. Not professionally and most definitely not personally.

As a result, the importance of my friends and family has never felt stronger. Without them none of this means a thing and I am surely indebted to them forever. They have given me the opportunity to climb mountains (figuratively and literally), not feel alone in the world, and to establish a true home in France and Switzerland. The pain and suffering that a young researcher sometimes feels, whether in pushing out a physics analysis or in confronting the rather unscientific bureaucracies inherent in a large collaboration like ATLAS, are far outweighed by their measure. The biggest takeaways from my time at CERN are therefore not technical at all and boil down to a few of the "strongest of all warriors" that are requisite in our dealings and goings-on about the world: patience and compassion. I have been damn lucky to have been surrounded by exemplars of these two during my doctorate. I hope to follow along in their footsteps as we move forward.

CURRICULUM VITAE

Daniel Joseph Antrim

Education

2013–2019 | **University of California, Irvine** — Ph.D. in Physics
Degree emphasis: Experimental High Energy Particle Physics
Advisor: Professor Anyes Taffard

2008–2013 | **University of California, Los Angeles** — B.Sc. in Physics
Degree emphasis: Mathematical Physics
Distinguishments:
Departmental honors (Physics and Astronomy)
Latin honors (Cum laude)

2004–2008 | **Woodbridge High School** — High School Diploma

Employment

2014–2019 | **University of California, Irvine Dept. of Physics and Astronomy**
Graduate Student Researcher
University of California, Irvine (UCI)

2013–2014 | **University of California, Irvine Dept. of Physics and Astronomy**
Teaching Assistant
University of California, Irvine (UCI)

2011–2013 | **University of California, Los Angeles Dept. of Physics and Astronomy**
Assistant to Professor Robert J. Finkelstein
University of California, Los Angeles (UCLA)

2011–2012 | **Rosenfeld Management Library**
Junior manager of library resources
University of California, Los Angeles (UCLA)

2010–2011 | **Institute of Geophysics and Planetary Physics (IGPP)**
Undergraduate researcher
University of California, Los Angeles (UCLA)

Significant Research Experience

As a member of the ATLAS Collaboration, 2014–Present

(Note: Labels to follow in ***bold italics*** denote official ATLAS designations/roles.)

ATLAS Physics Analysis Roles

Analysis Contact & Paper Editor Started and led the full Run-2 ATLAS search for evidence of double Higgs boson production in the final state where one Higgs boson decays via the dileptonic decay of the WW^*, ZZ^*, or $\tau\tau$ channels and the other decays via the $b\bar{b}$ channel

Analysis Contact Led Run-2 ATLAS searches for evidence of the supersymmetric partner to Standard Model top quark in dilepton final states

Led ATLAS search for evidence of supersymmetric partners to Standard Model Electroweak gauge bosons and leptons in dilepton final states for the ATLAS LHC Run-1 Supersymmetry Summary Analysis

Technical Roles within the ATLAS Collaboration

Trigger and data-acquisition software expert responsible for the commissioning and integration of the software-based readout driver ("software-ROD") for the Phase-1 New Small Wheel upgrade for the ATLAS muon spectrometer

Developer of the VERSO data-acquisition and electronics configuration & calibration software used for characterising the front-end electronics for the Phase-1 New Small Wheel Upgrade project and used in MicroMegas detector test-beam, cosmic-stand, test-bench, and chamber integration data-taking scenarios

Publications

247 published papers as part of the ATLAS Collaboration (9300+ citations [Inspire])

(Note: Items that are marked with a star ("★") denote publications in which I held a leading analyzer, analysis contact, and/or paper-editor role)

Peer-Reviewed Publications from the ATLAS Collaboration

Search for the non-resonant Higgs boson pair production in the $bb\ell\nu\ell\nu$ final state with the ATLAS detector in pp collisions at $\sqrt{s} = 13$ TeV, Physics Letters B. (submitted) ★

Search for the electroweak production of supersymmetric particles in final states with two or three leptons at $\sqrt{s} = 13$ TeV with the ATLAS detector, Eur. Phys. J. C 78 (2018) 995

Search for direct top squark pair production in final states with two leptons in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector, Eur. Phys. J. C77 (2017) 898 ★

Search for the electroweak production of supersymmetric particles in $\sqrt{s} = 8$ TeV pp collisions with the ATLAS detector, Phys. Rev. D 93, 052002 (2016) ★

Public Documents from the ATLAS Collaboration

Search for direct top squark pair production in final states with two leptons in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector, ATLAS-CONF-2017-034 (2017) ★

Non-ATLAS Peer-Reviewed Public Documents

System on a board readout for the VMM ASIC, Nuclear Instruments and Methods in Physics (in preparation) ★

Conferences and Seminar Presentations

Searches for sleptons with the ATLAS detector
Barcelona, Spain
26^{th} Int'l Conference on Supersymmetry and Unification of Fundamental Interactions
July 24^{th}, 2018

Searches for SUSY and Higgs pair production in dilepton events at ATLAS
University of California, Irvine
Departmental Particle Physics Seminar
January 17^{th}, 2018

Development of Trigger and Readout Electronics for the ATLAS New Small Wheel Detector Upgrade
Beijing, China
Int'l Conference on Technology and Instrumentation in Particle Physics 2017
May 22^{nd}, 2017

Honors and Awards

University of California, Irvine (Graduate)

US-ATLAS Outstanding Graduate Student Achievement Award, 2019

US-LHC Users Organization Lightning Round Talk Award, 2018

Technical Skills

General and Data-Analysis Oriented

Programming Languages: C/C++ (low- and high-level), Python, bash, Qt (graphical user-interface development)

Data formats: ROOT, HDF5

Machine learning (skilled in the use of the Keras and Tensorflow libraries)

Statistics and Frequentist limit setting procedures for ATLAS analyses

Analysis framework development and maintenance (skilled in and familiar with the ATLAS analysis software and xAOD–based event data model)

Technical/Hardware Oriented

Trigger/DAQ software, as related to High Energy Particle Physics experiments (within the ATLAS environment and standalone)

Network-based computing (experience in communication over both the UDP and TCP protocols)

Asynchronous and multi-threaded programming principles (experience primarily within C++)

Readout electronics and detector instrumentation, as related to particle detectors

Familiar hardware protocols: HDLC, SPI, I2C

Advising

University of California, Irvine

Alexander Armstrong (Graduate)

Syed Faizanul Haque (Undergraduate)

Outreach and Education

Official CERN Tour Guide

Leading tours through the CERN facilities, including underground to the ATLAS experimental cavern, to engage with and educate the public

Teaching Experience

Teaching Assistant (University of California, Irvine)

Physics 3LC (2 academic quarters): Laboratory on optics, radioactivity, & atomic physics

Physics 7D (1 academic quarter): Electromagnetism for Physicists & Engineers

ABSTRACT OF THE DISSERTATION

Sweet Little Nothings; or, Searching for a Pair of Stops, a Pair of Higgs Bosons, and a Pair of New Small Wheels for the Upgrade of the Forward Muon System of the ATLAS Detector at CERN

By

Daniel Joseph Antrim

Doctor of Philosophy in Physics

University of California, Irvine, 2019

Professor Anyes Taffard, Chair

This thesis reports on two searches for physics beyond the Standard Model (SM), performed using data collected from the $\sqrt{s} = 13\,\text{TeV}$ proton-proton (pp) collisions recorded by the ATLAS detector at CERN between the years 2015–2018, the period of LHC Run 2. Both searches are performed in final states requiring exactly two leptons, where a lepton is considered to be either an electron or muon, and enriched in b-tagged jets. The first is a search for a relatively light supersymmetric partner to the SM top-quark, referred to as the 'stop quark', under the assumption of R-parity conserving Supersymmetry. No statistically significant discrepancies are observed in the analysis' regions. As a result, exclusion limits are set on the allowed mass range of the stop quark. The second is a search for evidence of Higgs bosons (h) being produced in pairs, a process predicted in the SM and that is sensitive to the Higgs self-coupling parameter. Given the small predicted cross-section for this process, the analysis searches for enhanced production of Higgs pairs, to a level observable already in the Run 2 data and above that predicted by the SM alone. The analysis is performed in the dilepton final state of the $hh \to bbWW$ channel, and is the first time that this search has been performed by the ATLAS experiment. No evidence for Higgs boson pairs are observed in the Run 2 data and upper limits are therefore set on the cross-section of the SM Higgs boson pair production process. In addition to these two physics analyses, this thesis discusses the upgrade of the forward muon system of the ATLAS detector, referred to as the 'New Small Wheel' (NSW) Upgrade, that aims for an installation in the ATLAS cavern beginning in 2020. Emphasis is given to discussion of the frontend readout electronics, both on their readout and calibration.

Chapter 1

The Standard Model of Particle Physics

If you wish to make an apple pie from scratch, you must first invent the universe.

–Carl Sagan, *Cosmos: A Personal Voyage*

All models are wrong, but some are useful.

–George Box

As it stands, what has become known as the 'Standard Model (SM) of Particle Physics' is nothing less than one of the greatest achievements of mankind, due to both the magnitude by which it has changed our perception of the underlying nature of the universe and to the methods and tinkerings by which this nature was unveiled by many clever physicists whose history has become veritable lore. In terms of imagination and insight, it is second only to the special and general theories of relativity – though the fields are nevertheless intricately intertwined.

Not considering the scientific progress made in the 18^{th} and 19^{th} centuries, and ignoring the ancient Greeks despite their fabled invention of atomic theory, the physical insights and major work that led to the current picture of elementary particle physics described by the SM began with the *annus mirabilis* papers of Albert Einstein in the year 1905 [1, 2, 3]. In these papers, Einstein was able to shed light on the quantization of electromagnetic radiation (building off of the seminal work of Max Planck [4]) and introduce the special theory of relativity. These works laid the conceptual and philosophical groundwork for the major breakthroughs in fundamental physics of 20^{th} century physics: from the 'old quantum theory' of Bohr and Sommerfeld in the early 1900's to the equivalent wavefunction and matrix-mechanics formulations of Schrödinger and Heisenberg that coalesced

into 'modern' quantum mechanics in the mid-1920's. The modern approach, non-relativistic at its heart, provided a sufficient mathematical and interpretable framework in which to work and match predictions to observed phenomena, old and new. It has for the most part remained unchanged and is the quantum mechanics that is taught to students at both the undergraduate and graduate level to this very day. It is the theory that has since revolutionised all aspects of the physical sciences and technologies that dictate our everyday-lives. In the mid-1920's, however, despite large efforts put forth by the forbears of modern quantum mechanics, the quantum-mechanical world had yet to be made consistent with Einstein's theory of relativity — a requirement that must be met for all consistent physical theories of nature. It was the insight of Paul Dirac who was finally able to successfully marry the theory of the quantum with that of relativity when he introduced his relativistic quantum-mechanical treatment of the electron in 1927 and 1928 [5, 6].[1] This work provided the starting point for a decades-long search of a consistent quantum-mechanical and relativistic treatment of electrodynamics, known as *quantum electrodynamics* (QED). The search for QED ended at the end of the 1940's with the groundbreaking work of Dyson, Feynman, Schwinger, and Tomanaga [9, 10, 11, 12, 13, 14, 15, 16] that introduced the covariant and gauge invariant formulation of QED — the first such relativistic quantum field theory (QFT). QED allowed the physicists to make predictions that agreed with observation at unprecedented levels of accuracy and has since led to the adoption of its language and mathematical toolkit as the foundational framework in which to construct models that accurately describe nature.[2] The SM is no less than an ultimate conclusion of these works: a consistent set of relativistic quantum field theories, using the language developed by Feynman et al., that describes essentially all aspects of the known particles and forces that make up the observed universe.

1.1 Particles and Forces

There are four known fundamental forces at work in the universe: electromagnetism, the weak interaction, the strong interaction, and gravity. Our understanding of the existence of each of these forces has essentially been arrived at empirically, with physicists following experimental clues, and their basic behaviors deduced after long trials of effort. The SM encompasses all of these forces except for gravity, which currently is only described by the classical (i.e. not quantum) theory of geometrodynamics, or general relativity. The gravitational interaction is incredibly weak in comparison to the others, however, and is not relevant to the types of particle interactions that we are currently sensitive to in particle physics experiments. Electromagnetism is by far the most familiar, as it is the force most commonly experienced and is what is at work in our everyday life (reaction forces between objects on tables and chairs, friction, wall-plugs, batteries, DNA structure,

[1] A complete history of the people and ideas involved in the development of the modern theory of Quantum Mechanics can be found in references [7, 8], and the references therein.

[2] For a complete discussion of the developments leading up to QED, see the fabulous book by S. Schweber [17].

etc...) and is typically what students are first presented with in their physics studies. The weak force is responsible for things like radioactive decay, which makes possible the process of nuclear β-decay and the nuclear fission process that fuels the sun, for example. The strong force is what binds protons and neutrons together, and thus is responsible for holding together most of the (ordinary) matter in the universe.[3]

The forces mediate the interactions between the matter particles, which we use to deduce their presence. The SM predicts fundamental, point-like particles that appear in two general classes depending on whether they have integral spin ($\mathcal{S} \in [0, 1, 2, ...]$) or half-integral spin ($\mathcal{S} \in [1/2, 3/2, ...]$); the former are referred to as *bosons* and the latter as *fermions*. In the SM, the particles that are responsible for making up matter are all spin-1/2 fermions and are either *leptons* or *quarks*; within each class there are three generations (or families) that are essentially copies of the first. The forces in the SM are interpreted as being mediated by spin-1 bosons, referred to as the *gauge bosons*. The leptons and quarks all experience the weak force, but only the quarks experience the strong interaction. All electrically charged particles interact with the electromagnetic interaction.

The particles of the SM are described as quantum fields whose dynamics are described by the SM Lagrangian from which the equations of motions can be derived. The particles, and by extension the SM Lagrangian that describes them, are found to be invariant under transformations of spacetime (space translations, rotations, Lorentz boosts) and three internal transformations described by unitary transformations: $\mathcal{P}\times\mathcal{SU}(3)_C\times\ \mathcal{SU}(2)_L\times\ \mathcal{U}(1)_Y$. This is illustrated in Figure 1.1. The strong force is described by a local $\mathcal{SU}(3)$ symmetry that acts only on the particles that have *color charge*. The term "color" arises from the fact that the color charge is found to exist in three varieties which have been labelled as red (r), blue (b), or green (g), and due to the fact that "colorless" states are formed when all three are combined (r+g+b), just like with visible light that humans are familiar with, or when states are formed of color-anti-color pairs (r+$\bar{\text{r}}$). For this reason, the QFT describing the strong force is called Quantum *Chromo*dynamics (QCD), and is mediated by eight *gluons* (G). The particles subject to the weak force are invariant under weak-isospin $\mathcal{SU}(2)$ transformations, mediated by the three W bosons (W_1, W_2, W_3). The $\mathcal{U}(1)$ transformations, mediated by the B boson, preserve weak-hypercharge, Y. The $\mathcal{SU}(2)$ symmetry is respected only by the left-handed chiral particles (leptons or quarks), with the right-handed chiral particles not participating. There is additionally a single scalar (i.e. spin-0) field, the Higgs field, that is an $\mathcal{SU}(2)$ doublet, about which more will be described shortly. The particle content thus described is presented in detail in Table 1.1. The $\mathcal{SU}(2)$ left-handed chiral fields appear as doublets and are grouped in and 'up-down' pair (e.g. (u_{L}, d_{L}) or (e_{L}, $\nu_{e,\text{L}}$)) whereas the right-handed chiral fields, living in the singlet representation of $\mathcal{SU}(2)$, do not (e.g. u_{R})). Note that the SM does not allow for right-handed neutrinos (a term like ν_{R} does not appear).

[3] 'Ordinary' to distinguish from dark matter, for example.

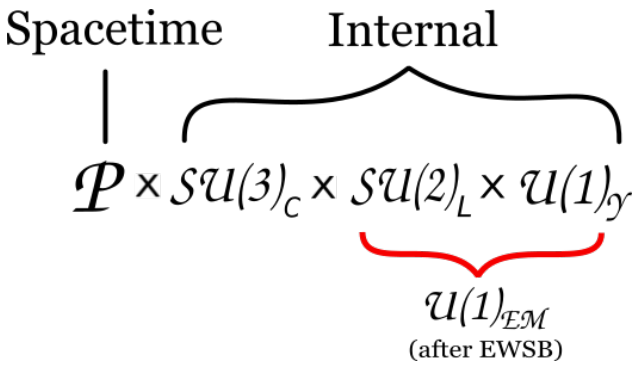

Figure 1.1: The spacetime and internal gauge structure of the SM. $\mathcal{P}$ refers to the Poincaré symmetry group. $\mathcal{SU}(3)_c$ refers to the $\mathcal{SU}(3)$ symmetry of the color sector of QCD and $\mathcal{SU}(2)_L \times \mathcal{U}(1)_Y$ refers to the left-handed chiral symmetry of the electroweak interaction. After spontaneous symmetry breaking due to the Higgs mechanism, the $\mathcal{SU}(2)_L \times \mathcal{U}(1)_Y$ symmetry reduces to the $\mathcal{U}(1)_{EM}$ symmetry of electromagnetism.

The SM Lagrangian is shown in Equation 1.1 and describes the complete content of the SM: encompassing all interactions between the known particles and the symmetries that they obey.

$$\mathcal{L}_{\text{SM}} = -\frac{1}{4} \sum_{\text{gauge}} F^i_{\mu\nu} F^{i\,\mu\nu} - \sum_f \bar{f} \gamma^\mu D_\mu f + (D_\mu \phi)^\dagger (D^\mu \phi) - \mu^2 \phi^\dagger \phi - \lambda (\phi^\dagger \phi)^2 \tag{1.1}$$

The first term of Equation 1.1 is a sum over the three internal gauge groups, and $F^a_{\mu\nu} = \partial_\mu A^a_\nu - \partial_\nu A^a_\mu + g f^{abc} A^b_\mu A^c_\nu$, where A_μ is one of the three gauge fields, g is the associated gauge coupling parameter, and a sum over i is implied. The f^{abc} are the so-called *structure constants* of the gauge group. For Abelian groups like $\mathcal{U}(1)$, $f^{abc} = 0$. For non-Abelian gauge groups like $\mathcal{SU}(2)$ and $\mathcal{SU}(3)$, $f^{abc} \neq 0$. For example, for $\mathcal{SU}(2)$ the structure constants are nothing more than the Levi-Civita totally anti-symmetric tensor, ε_{ijk}, giving for the weak gauge force:

$$\boldsymbol{\mathcal{W}}_{\mu\nu} = \partial_\mu \boldsymbol{\mathcal{W}}_\nu - \partial_\nu \boldsymbol{\mathcal{W}}_\mu - g_2 \boldsymbol{\mathcal{W}}_\mu \times \boldsymbol{\mathcal{W}}_\nu \tag{1.2}$$

where $\boldsymbol{\mathcal{W}}_\mu$ is the vector of the three weak gauge fields (W_1, W_2, and W_3) and g_2 is their associated gauge coupling. The non-zero f^{abc} of non-Abelian gauge groups means that the gauge bosons of the weak and strong interactions can interact with themselves due to terms appearing in Equation 1.1 that contain only the gauge bosons.

The second term of Equation 1.1 describes the lepton and quark kinetic energies and their interactions with the gauge fields. The f refer to the fermion fields (quarks and leptons) and the corresponding sum is over all species of fermion. D_μ is the gauge covariant derivative, and for the

SM is given by:

$$D_\mu = \partial_\mu - ig_1\frac{Y}{2}\mathcal{B}_\mu - ig_2\frac{\tau^i}{2}\mathcal{W}^i_\mu - ig_3\frac{\lambda^a}{2}G^a_\mu \tag{1.3}$$

where g_1, g_2, and g_3 are the gauge coupling constants for $\mathcal{U}(1)_Y$, $\mathcal{SU}(2)_L$, and $\mathcal{SU}(3)_C$, respectively, that give the overall strength of the associated coupling. Summation over repeated indices is implied and the τ^i (λ^a) are the three (eight) generators of the $\mathcal{SU}(2)$ ($\mathcal{SU}(3)$) gauge group, with $i \in [1, 2, 3]$ ($a \in [1, .., 8]$), and are typically represented by the Pauli (Gell-Mann) matrices. Note that the form of Equation 1.3 is strictly mandated by the requirement that the theory be *gauge invariant*, i.e. that transformations of the fields under the internal symmetries of Figure 1.1 leave the action of $\mathcal{L}_{\text{SM}}$ unchanged.

The last three terms in Equation 1.1 are all terms including the Higgs field, ϕ, and will be discussed in detail in Section 1.3.

Inspection of Equation 1.1 will reveal two things. The first thing that one may notice is that it does not appear to describe electromagnetism as it does not have a term representing the photon, the familiar mediator of the electromagnetic interaction. The second, and perhaps more immediately obvious, thing is that no mass terms appear in $\mathcal{L}_{\text{SM}}$: all fields appear to have zero mass! Both of these facts are counter to our everyday experience: we know electromagnetism is real and that matter, at the very least, is massive. In the next few sections we will see how these apparent issues are resolved.

Table 1.1: The particle content of the SM and their transformation properties under the SM gauge groups, prior to electroweak symmetry breaking. The representations of each of the gauge groups are shown in the three-right columns. The $\mathcal{U}(1)$symmetry of weak-hypercharge transformations is one-dimensional and the column gives the weak-hypercharge $\mathcal{Y}$ associated with each field. For $\mathcal{SU}(3)$and $\mathcal{SU}(2)$, **1** refers to the field belonging to the associated singlet representation, **2** to the doublet representation, **3** to the triplet representation, and **8** to the octet representation.

	Field Label	Content	Spin	$\mathcal{U}(1)$ $(=\mathcal{Y})$	$\mathcal{SU}(2)$	$\mathcal{SU}(3)$
Quarks	Q_i	$(u_{\rm L},\, d_{\rm L})$, $(c_{\rm L},\, s_{\rm L})$, $(t_{\rm L},\, b_{\rm L})$	1/2	1/3	**2**	**3**
	$u_{{\rm R},i}$	$u_{\rm R}$	1/2	4/3	**1**	**3**
	$d_{{\rm R},i}$	$d_{\rm R}$	1/2	$-2/3$	**1**	**3**
Leptons	L_i	$(e_{\rm L},\, \nu_{e,{\rm L}})$, $(\mu_{\rm L},\, \nu_{\mu,{\rm L}})$, $(\tau_{\rm L},\, \nu_{\tau,{\rm L}})$	1/2	-1	**2**	**1**
	$e_{{\rm R},i}$	$e_{\rm R}$, $\mu_{\rm R}$, $\tau_{\rm R}$	1/2	-2	**1**	**1**
Gauge Fields	B	B	1	0	**1**	**1**
	W	$(W_1,\, W_2,\, W_3)$	1	0	**3**	**1**
	G	G_a, $a \in [1, .., 8]$	1	0	**1**	**8**
Higgs Field	ϕ	$(\phi^+,\, \phi^0)$	0	1	**2**	**1**

1.2 The Electroweak Theory

It was the work of Glashow, Weinberg, and Salam (GWS) that ultimately put forth a consistent picture of the chiral weak force and its unification with electromagnetism [18, 19, 20]. As a result, the theory of particles and fields that respect the $\mathcal{SU}(2)_L \times \mathcal{U}(1)_Y$ gauge invariance of the SM is sometimes referred to as 'GWS theory', but is more typically known as the electroweak theory. Since all matter particles are subject to the electroweak interaction, but only a subset of the particles that have color charge (the quarks) are subject to the strong interaction described by QCD, the study of the SM can essentially be partitioned into two parts: the part that deals with the dynamics and interactions of colored objects (the 'QCD part', $\mathcal{L}_{\rm QCD}$) and the part that deals with electroweak interactions, including the Higgs (the 'Electroweak part', $\mathcal{L}_{\rm Electroweak}$). Given the broad reach of the electroweak interaction, in the early days GWS theory was considered the heart of the SM and why GWS were awarded the Nobel prize in 1979.[4] In this section we will focus on the $\mathcal{SU}(2)_L \times \mathcal{U}(1)_Y$ portion of $\mathcal{L}_{\rm SM}$.

The first thing to remember is that the electroweak theory is *chiral*, i.e., it distinguishes between

[4] Actually, the acceptance of the GWS theory as the de-facto SM of the time was not widely held until some years after its publication, when t'Hooft proved that it was renormalizable [21, 22]. Such a complete understanding in the QCD sector would not come until later, in the 1970's, with the work of Gross, Wilczek, and Politzer [23, 24].

left- and right-chiral fermion fields. For conceptual clarity, it can be useful to take the massless (relativistic) limit of fermions to get an idea of what chirality represents. For a massless fermion field, the chirality is equivalent to the perhaps more-familiar *helicity*, defined as the projection of its spin onto its momentum (direction of motion). The helicity of left-handed (right-handed) massless fermions is positive (negative), meaning that their spin is parallel (anti-parallel) to its momentum. Fermion fields, then, are commonly defined inclusive of their handedness, with the left- and right-handed components projected out using the P_L and P_R projection operators,

$$f_{\mathrm{L}} = P_L f = \frac{1}{2}(1-\gamma_5)f \quad \text{and} \quad f_{\mathrm{R}} = P_R f = \frac{1}{2}(1+\gamma_5)f. \tag{1.4}$$

Focusing only on the first generation of the leptons (the discussion holds equally well for the second and third generations, as well as for the quarks), we can gather the $\mathcal{U}(1)$ terms of Equation 1.1,

$$\begin{aligned} -\mathcal{L}_{\mathrm{ferm}}(\mathcal{U}(1), \text{leptons}) &= \bar{L} i\gamma^\mu (i g_1 \frac{Y_L}{2} B_\mu) L + \bar{e}_R i\gamma^\mu (i g_1 \frac{Y_R}{2} B_\mu) e_R \\ &= \frac{g_1}{2}[Y_L(\bar{\nu}_L \gamma^\mu \nu_L + \bar{e}_L \gamma^\mu e_L) + Y_R \bar{e}_R \gamma^\mu e_R] B_\mu, \end{aligned} \tag{1.5}$$

where $L = (\nu_L, e_L)$ is used in going from the first to second line. Likewise, gathering the associated $\mathcal{SU}(2)$ terms and noting that $\tau^i W^i$ is a 2×2 matrix since the τ^i are $\mathcal{SU}(2)$ generators (e.g. the Pauli matrices) gives,

$$\begin{aligned} -\mathcal{L}_{\mathrm{ferm}}(\mathcal{SU}(2), \text{leptons}) &= \bar{L}\, i\gamma^\mu [i g_2 \frac{\tau^i}{2} W^i_\mu]\, L \\ &= -\frac{g_2}{2}\left[\bar{\nu}_L \gamma^\mu \nu_L W^0_\mu - \sqrt{2}\bar{\nu}_L \gamma^\mu e_L W^+_\mu - \sqrt{2}\bar{e}_L \gamma^\mu \nu_L W^-_\mu - \bar{e}_L \gamma^\mu e_L W^0_\mu\right], \end{aligned} \tag{1.6}$$

where we have used the following re-definitions of the $\mathcal{SU}(2)$ gauge fields,

$$W^+_\mu = \frac{1}{\sqrt{2}}\left(-W^1_\mu + iW^2_\mu\right) \qquad W^-_\mu = \frac{1}{\sqrt{2}}\left(-W^1_\mu - iW^2_\mu\right) \qquad W^0_\mu = W^3_\mu. \tag{1.7}$$

In principle, Equations 1.5 and 1.6 describe completely all electroweak interactions between matter and the gauge fields of $\mathcal{SU}(2)_L \times \mathcal{U}(1)_Y$. We would like to make the correspondence between these equations and what we know to empirically exist: the electromagnetic interaction and the presence of a massive, charged mediator of the weak nuclear force responsible for nuclear β-decay, for example. From the theory of QED, it is a-priori known what the form of the interaction of the neutral photon and the electron should look like. Inspecting all charge-preserving (i.e. neutral) terms of Equation 1.5 and 1.6, it can be seen that the B_μ and ${W^0}_\mu$ fields have this expected fermion coupling, suggesting a re-definition as follows,

$$\begin{pmatrix} A_\mu \\ Z_\mu \end{pmatrix} = \begin{pmatrix} \cos\theta_W & \sin\theta_W \\ -\sin\theta_W & \cos\theta_W \end{pmatrix} \begin{pmatrix} B_\mu \\ W^0_\mu \end{pmatrix}, \tag{1.8}$$

where we have used $Y_L = -1$ and define the relations between the $\mathcal{SU}(2)$ and $\mathcal{U}(1)$ couplings as,

$$\sin\theta_W = \frac{g_1}{\sqrt{g_1^2+g_2^2}} \qquad \cos\theta_W = \frac{g_2}{\sqrt{g_1^2+g_2^2}} \qquad e = \frac{g_1 g_2}{\sqrt{g_1^2+g_2^2}}. \tag{1.9}$$

The angle θ_W is known as the *Weinberg angle*, or the *weak mixing angle*. It quantifies the amount of *gauge mixing* that occurs between the neutral $\mathcal{SU}(2)_L \times \mathcal{U}(1)_Y$ gauge fields, B_μ and ${W^0}_\mu$.

The above algebra allows to re-write the electroweak Lagrangian, now describing interactions between the fermions and the newly defined A_μ, Z, and ${W^\pm}$, as,

$$\begin{aligned}\mathcal{L}_{\text{ferm, first-gen.}} = & \underbrace{\sum_{f\in\nu_e,e,u,d} eQ_f\left(\bar{f}\gamma^\mu f\right)A_\mu}_{\text{Neutral, }\sim\text{ EM}} \\ & + \underbrace{\frac{g_2}{\cos\theta_W}\sum_{f\in\nu_e,e,u,d}\left[\ \bar{f}_L\gamma^\mu f_L\left(T_f^3 - Q_f\sin^2\theta_W\right) + \bar{f}_R\gamma^\mu f_R\left(-Q_f\sin^2\theta_W\right)\ \right]Z_\mu}_{\text{Neutral weak interaction}} \\ & + \underbrace{\frac{g_2}{\sqrt{2}}\left[\ \left(\bar{u}_L\gamma^\mu d_L + \bar{\nu}_{e,L}\gamma^\mu e_L\right)W_\mu^+ + h.c.\ \right]}_{\text{Charged weak interaction}}\end{aligned} \tag{1.10}$$

The first term of Equation 1.10 has the expected form expected from QED, describing the interaction between a neutral gauge boson and fermions and allows us to interpret the parameter e, introduced in Equation 1.9, as the coupling of electromagnetism (electric charge) with Q_f as the fermion's electric charge quantum number (in units of e). The A_μ arrived at via the gauge mixing of Equation 1.8 then must correspond to the photon of electromagnetism. The second term of Equation 1.10 predicts the existence of an additional neutral gauge boson, the Z boson, with its couplings to the left- and right-handed fermions dictated by the $\mathcal{SU}(2)_L \times \mathcal{U}(1)_Y$ gauge mixing. The quantity T_f^3 is the fermion field's quantum number associated with the third component of weak-isospin ($\mathcal{SU}(2)$). The third term of Equation 1.10 involves charged weak-interactions involving the ${W^\pm}_\mu$ gauge bosons that transform the up- and down-type fields of the left-handed $\mathcal{SU}(2)$ doublet fields into each other.

The terms involving ${W^\pm}_\mu$ in Equation 1.10 are of the form $\bar{\nu}_L\gamma^\mu e_L$ which, using the chiral projection operators (Equation 1.4), can be written as follows,

$$\bar{\nu}_L\gamma\mu e_L = \frac{1}{2}\bar{\nu}\gamma^\mu(1-\gamma_5)e, \tag{1.11}$$

showing that the charged weak interactions involving ${W^\pm}_\mu$ are the coherent sum of vector (γ^μ) and axial-vector ($\gamma^\mu\gamma_5$) bilinear covariants: this is the famous *V-A* charged-current interaction of Fermi's nuclear β-decay. It is this *V-A* form that results in the charged interactions of the weak force not being invariant under chiral transformations ($f_R \leftrightarrow f_L$): they involve only the left-chiral

fermion fields. For this reason, *parity* is said to be maximally violated by the weak interaction.[5] This result, as presented in the above, is due to our having injected it into our assumption on the field content in the first place out of hindsight. There is no first-principles reason why the weak interactions should be this way, however, and historically it was arrived at empirically.

What we have principally shown in this section is that, in order for the $\mathcal{SU}(2)_L \times \mathcal{U}(1)_Y$ content of the SM Lagrangian to correspond to what is found experimentally, it is expected that the gauge fields of the underlying symmetries mix. Specifically, the neutral $\mathcal{SU}(2)$ gauge field (${W^0}_\mu$) mixes with that of the $\mathcal{U}(1)$ gauge symmetry (B_μ) by an amount dictated by the Weinberg angle, θ_W, resulting in descriptions of interactions consistent with the experimentally observed photon and Fermi decay, as well as a prediction of a neutral Z-boson. The fermion electric charge is seen to be dependent on this mixing and can be related to the underlying $\mathcal{SU}(2)$ and $\mathcal{U}(1)$ gauge symmetries by the Gell-Mann-Nishijima relation,

$$Q_f = T_f^3 + \frac{1}{2}Y, \tag{1.12}$$

This relation summarises well the result of the $\mathcal{SU}(2)_L \times \mathcal{U}(1)_Y$ gauge mixing and allows one to infer that the electromagnetism of common experience is related to the weak interaction and is in fact just one aspect of a unified electroweak interaction.

We have thus shown that the SM predicts the existence of the familiar electromagnetic force and potentially provides an additional mediator (the ${W^\pm}_\mu$) for the charged weak interaction that, prior to the formulation of GWS, was lacking a consistent physical description. However, it is still not evident how the SM can support the experimental fact that fermions have mass and that the predicted mediator of the charged weak-nuclear force (the ${W^\pm}_\pm$) *must* be massive given the very short range of the interaction. In order for such mass terms to be allowed in $\mathcal{L}_{\text{SM}}$, we need the Higgs mechanism.

1.3 The Higgs Mechanism and Electroweak Symmetry Breaking

The missing mass-terms for the fields in $\mathcal{L}_{\text{SM}}$ are provided by the Brout-Englert-Higgs (BEH) mechanism [25, 26, 27]. Before describing the specifics of the BEH mechanism, we should first describe the problem of why $\mathcal{L}_{\text{SM}}$ doesn't support general mass terms for any of the fields in the first place. That is, for example, why can't a fermion term like $m\bar{f}f$ exist in $\mathcal{L}_{\text{SM}}$?

Adding mass terms to $\mathcal{L}_{\text{SM}}$ for the fermions explicitly breaks the underlying $\mathcal{SU}(2)$ gauge symmetry. This can be understood if we recognize the experimentally supported fact that the left-handed

[5] A parity transformation refers to inverting a field's space coordinates as $\vec{x} \to -\vec{x}$.

fermions appear as $\mathcal{SU}(2)$ doublets and that the right-handed fermions as singlets,

$$\begin{aligned} m\bar{f}f &= m\bar{f}(P_L + P_R)f \\ &= m\bar{f}P_L P_L f + m\bar{f}P_R P_R f \\ &= m\left(\bar{f}_R f_L + \bar{f}_L f_R\right), \end{aligned} \tag{1.13}$$

where we have used identity relations of the projection operators P_L and P_R and the fact that $\bar{f}P_L = \bar{f}_R$ (and vice-versa). The last line of Equation 1.13 involve terms mixing $\mathcal{SU}(2)$ doublets with $\mathcal{SU}(2)$ singlets. Such a term is therefore not allowed if we wish to keep the $\mathcal{SU}(2)$ gauge symmetry intact.

Mass terms for the gauge bosons, of the form $mB_\mu B^\mu$, also do not work. For the Abelian $\mathcal{U}(1)$ symmetry, for example, gauge invariance implies invariance of $\mathcal{L}_{\text{SM}}$ under transformations of the form $B'_\mu \to B_\mu - \partial_\mu \chi/g$. Such a mass term for the gauge bosons is clearly not invariant under such a transformation. Even forgoing this fact, adding such a term would quickly lead to non-renormalizable divergences appearing in the theory, due to the longitudinal field components that appear in massive field propagators, rendering $\mathcal{L}_{\text{SM}}$ meaningless.

The BEH mechanism provides a way out of this problem. It refers to the introduction of a spin-0 field, the Higgs field (Table 1.1), to the SM along with its corresponding interaction terms to $\mathcal{L}_{\text{SM}}$: the last three terms of Equation 1.1. The final two terms make up what is referred to as the Higgs potential and can be expressed as,

$$V(\phi) = -\mu^2\phi^2 - \lambda\phi^4 \tag{1.14}$$

The Higgs field is an $\mathcal{SU}(2)$ doublet and it can be seen that the interactions described by Equation 1.14 respect $\mathcal{SU}(2)$ gauge symmetry. If $\mu^2 > 0$, nothing all too interesting occurs and Equation 1.14 describes a self-interacting, complex scalar field. If we take $\mu^2 < 0$, however, then the classical potential described by Equation 1.14 has non-zero minima located at $\phi = \pm v$ with $v = \sqrt{-\mu^2/\lambda}$. This is illustrated in Figure 1.2. We see that the stable equilibrium point ϕ_0 of the Higgs potential, the *Higgs vacuum expectatin value* (vev), is not at $\phi = 0$ but at v,

$$\phi_0 = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 \\ v \end{pmatrix} \tag{1.15}$$

The choice of Equation 1.15 to represent the Higgs vacuum is motivated by the requirement that the vacuum not be electrically charged — a fact motivated very much by experiment and everyday experience — so the up-type $\mathcal{SU}(2)$ component of the Higgs field, ϕ^+ (Table 1.1), is chosen to be zero for ϕ_0. The choice of an electrically neutral vacuum sets the rest of the $\mathcal{SU}(2)_L \times \mathcal{U}(1)_Y$ structure of the complex Higgs field since, by the Gell-Mann-Nishijima relation (Equation 1.12) and charge conservation, a neutral $\mathcal{SU}(2)_L \times \mathcal{U}(1)_Y$ field should have down-type $\mathcal{SU}(2)$ quantum numbers and

$\mathcal{U}(1)$ hypercharge $Y = 1$,

$$Q = T_3 + \frac{1}{2}Y \rightarrow Q_{\phi_0} = -\frac{1}{2} + \frac{1}{2} \times 1 = 0. \tag{1.16}$$

Note that Equation 1.15 states that only one component of the Higgs $\mathcal{SU}(2)$ doublet attains a non-zero vev. This clearly means that the $\mathcal{SU}(2)$ gauge symmetry is not respected by the choice of $\mu^2 < 0$ and that the electroweak $\mathcal{SU}(2)_L \times \mathcal{U}(1)_Y$ symmetry is *spontaneously broken*.[6] The Higgs field acquiring a non-zero vev is then referred to as the *electroweak symmetry breaking* (EWSB) of the SM.

To further examine the physical consequences of EWSB, we perturb the Higgs field about its minimum value,

$$\phi(x) \propto \begin{pmatrix} 0 \\ \frac{1}{2}(v + h(x)) \end{pmatrix}, \tag{1.17}$$

where $h(x)$ correspond to excitations of the Higgs field that represent the physically observable Higgs boson. Inserting Equation 1.17 into the $D_\mu\phi$ terms of Equation 1.1, one eventually works through the algebra and obtains,

$$|D_\mu\phi(x)|^2 = \frac{1}{8}v^2 g_2^2 \left[\left(W_\mu^1\right)^2 + \left(W_\mu^2\right)^2 \right] + \frac{1}{8}v^2 \left(g_1 B_\mu - g_2 W_\mu^3\right)^2. \tag{1.18}$$

Using the field re-definitions for the W_μ, A_μ and Z_μ introduced in Section 1.2, we see that this can be re-written as (modulo factors of 2),

$$|D_\mu\phi(x)|^2 \propto \left(\frac{1}{2}v g_2\right)^2 W_\mu^+ W^{-\,\mu} + \left(\frac{1}{2}v\sqrt{g_1^2 + g_2^2}\right)^2 Z_\mu Z^\mu + (0)^2 A_\mu A^\mu, \tag{1.19}$$

which provide, clearly, mass terms for the electroweak gauge bosons:

$$M_W = \frac{1}{2}v g_2, \qquad M_Z = \frac{1}{2}v\sqrt{g_1^2 + g_2^2}, \qquad M_A = 0. \tag{1.20}$$

The expression for the masses acquired by the $W^\pm$ and Z gauge bosons in Equation 1.20 is expected by Goldstone's theorem [28] which states that for every broken continuous symmetry one expects an associated massless scalar field (a 'Goldstone boson') to appear in the theory. The fact that the $W^\pm$ and Z acquire mass after EWSB is then interpreted as these fields having acquired longitudinal field components by 'eating' the Goldstone boson degrees of freedom associated with the breaking of $\mathcal{SU}(2)_L$. The BEH mechanism refers specifically to this means of the gauge bosons acquiring

[6] A symmetry of a Lagrangian is said to be 'spontaneously' broken if the Lagrangian of the underlying theory respects the symmetry but it gets broken through dynamical means or if the lowest-energy state (vacuum) does not respect the symmetry.

mass via 'eating' the Goldstone bosons.

The fact that the Higgs vev respects charge conservation (Equation 1.16) means that $\mathcal{L}_{\text{SM}}$, after EWSB, still respects a local $\mathcal{U}(1)$ gauge symmetry; although now this is the $\mathcal{U}(1)$ gauge symmetry associated with electromagnetism, $\mathcal{U}(1)_{EM}$, as opposed to that of weak-hypercharge, $\mathcal{U}(1)_Y$. This indicated in Figure 1.1.

There are also additional terms involving the now-massive $W^{\pm}$ and Z bosons and $h(x)$ in the expansion of $|D_\mu \phi(x)|^2$ of Equation 1.18 (not shown) that describe the gauge bosons' interactions with the observable Higgs boson, involving terms of the form hVV and $hhVV$ ($V \in (W, Z)$) whose coupling strengths depend on the gauge boson masses (Equation 1.20):

$$\mathcal{L}_{h-VV} \propto \frac{M_V^2}{v} \qquad \mathcal{L}_{hh-VV} \propto \frac{M_V^2}{v^2}. \tag{1.21}$$

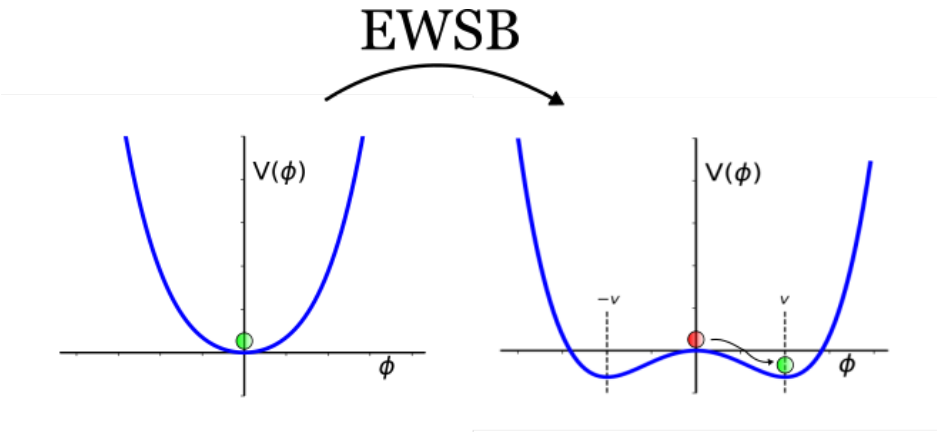

Figure 1.2: Illustration of electroweak symmetry breaking (EWSB). ***Left***: Higgs potential with $\mu^2 > 0$ with stable equilibrium at $\phi = 0$. ***Right***: With $\mu^2 < 0$, $\phi = 0$ is no longer a stable equilibrium and the Higgs attains a non-zero vacuum expectation value at $\pm v$ — breaking the $\mathcal{SU}(2)_L \times \mathcal{U}(1)_Y$ gauge symmetry of the electroweak sector of the SM.

As opposed to 'eating' gauge degrees of freedom as in the case of the $W^{\pm}$ and Z bosons, the fermion masses are obtained by adding additional interaction terms to $\mathcal{L}_{\text{SM}}$ between the fermions and Higgs fields,

$$\mathcal{L}_{f-h} = y_f \left(\bar{L} \phi e_R^- + \phi^\dagger \overline{e^-}_R L \right). \tag{1.22}$$

Since both L and ϕ are $\mathcal{SU}(2)$ doublets, adding the right-handed $\mathcal{SU}(2)$ singlet terms do not spoil the $\mathcal{SU}(2)$ symmetry. When the Higgs field acquires a non-zero vev after EWSB, we can insert

Equation 1.17 into Equation 1.22 to obtain expressions for the fermions masses,

$$m_f = y_f \frac{v}{\sqrt{2}}, \tag{1.23}$$

where the y_f are referred to as the fermion *Yukawa couplings*, and are free parameters of the SM that need to be measured. Additional interactions arise between the fermions and $h(x)$ whose couplings are related to the fermion masses as,

$$\mathcal{L}_{f-h} \propto \frac{m_f}{v} \bar{f} f h. \tag{1.24}$$

The general form Equation 1.22 implies that the y_f are matrices representing the Higgs-fermion Yukawa couplings. They can be diagonalized by performing the proper unitary transformations between the weak- and mass-bases of the fermion fields. In the case of the leptons, this rotation is the identity: the lepton's weak eigenbasis is the same as their mass eigenbasis. This is mainly due to the extraordinarily large mass difference between the charged and neutral leptons within each lepton generation [29]. Within the quark-sector, however, the mass- and weak-basis fermion fields differ. This implies that the diagonalization procedure results in mixing among the weak eigenstates of the quark fields to produce the observed mass eigenstates; i.e. the quark mass-eignstates (d, s, b) are coherent mixtures of the weak eigenstates (d', s', b').[7] This allows for the flavor-changing processes that are present in charged weak interactions, allowing for interaction terms involving the decay of a quark of one family into that of another family. The amount of mixing in the quark sector is dictated by a 3×3 unitary matrix known as the *Cabibbo-Kobayashi-Maskawa* (CKM) matrix [30] $\mathcal{V}_{\text{CKM}}$, the general form of which has four free parameters: three mixing angles and a complex phase, δ. The off-diagonal terms of the CKM matrix and the value of the mixing angles dictate the amount of flavor mixing in the quark sector. The complex phase δ allows for charge-parity (CP) symmetry violating effects to occur. In fact, this term is the *only* term of the SM that allows for CP-violation — an effect important for providing interactions that are asymmetric between matter and anti-matter fields.[8]

The remaining terms of $V(\phi)$ (Equation 1.14) involve only the Higgs field. After EWSB and the Higgs field acquiring vev we obtain,

$$V(\phi) \to V(\phi)_{\text{EWSB}} = -\lambda\nu^2 h^2 - \lambda\nu h^3 - \frac{1}{4}\lambda h^4 + \text{const.} \tag{1.25}$$

where we have ignored the terms already discussed above. The first term of Equation 1.25 is the

[7] The mixing can be parametrised as either occurring between the up-type, down-type, or a mixture of up- and down-type fields of each $\mathcal{SU}(2)$ doublet. Without loss of generality and for simplicity, it is usually given with respect to the down-type quark fields as shown here.

[8] Charge Parity (CP) symmetry refers to the invariance of a theory with respect to swapping particles with their corresponding anti-particles and, additionally, inverting a field's spatial coordinates, $\psi(\vec{x}) \to \psi(-\vec{x})$. The former is the 'C' symmetry and the latter is the 'P' symmetry.

Higgs boson mass term, the second and third are the triple and quartic Higgs self-couplings,

$$\underbrace{m_h = \sqrt{-2\mu^2} = \sqrt{2\lambda v^2}}_{\text{Higgs boson mass}} \qquad \underbrace{\mathcal{L}_{hhh} \propto \frac{m_h^2}{v} \qquad \mathcal{L}_{hhhh} \propto \frac{m_h^2}{v^2}}_{\text{Triple and quartic Higgs self-couplings}} . \tag{1.26}$$

1.4 The Complete Standard Model, Successes and Shortcomings

The physical field content of the SM, after EWSB, is detailed in Table 1.2 Also shown in Table 1.2 are the electric charge associated with each particle field, Q, the relevant couplings, and particle masses.

Table 1.2: The particle content of the SM after the process of electroweak symmetry breaking. Shown for each particle species are the associated electric charge, Q, coupling, and mass (approximate). The y_i are the Yukawa coupling (Equation 1.24), α_{EM} is the QED coupling constant ('fine structure constant'), $\mathcal{V}$ indicates the parameters of the CKM matrix, and α_s is the QCD coupling constant. The quantities λ and μ are the Higgs self-coupling parameter and mass terms, respectively, appearing in Higgs potential terms (Equation 1.14).

	Physical Field	Q	Coupling	Mass [GeV]
Quarks	u, c, t	$2/3$	$(y_i =)$ 1×10^{-5}, 7×10^{-3}, 1	2×10^{-3}, 1.27, 173
	d, s, b	$-1/3$	$(y_i =)$ 3×10^{-5}, 5×10^{-4}, 0.02	4×10^{-4}, 0.10, 4.18
Leptons	e, μ, τ	-1	$(y_i =)$ 3×10^{-7}, 6×10^{-4}, 0.01	5×10^{-4}, 0.106, 1.777
	ν_e, ν_μ, ν_τ	0	–	–
Bosons	γ	0	$\alpha_{\text{EM}} \simeq 1/137$	0
	Z	0	$\sin\theta_W \simeq 0.5$	91.2
	(W^+, W^-)	$(+1, -1)$	$\mathcal{V}_{\text{CKM}}$	80.4
	G	0	$\alpha_s \simeq 0.1$	0
Higgs	h	0	λ, μ	125.09

1.4.1 Successes of the Standard Model

The particle content presented in Table 1.2 represents the current picture of the visible matter content of the Universe. It is a concise picture. With this particle content in hand, and the description of their fundamental particle interactions (Equation 1.1), the predictive power of the SM is immense.

Figure 1.3 shows a summary of LHC measurements of cross-sections for various production processes at 7, 8, and 13 TeV. The agreement of these measurements with the theoretical predictions provided by the SM, spanning over 10 orders of magnitude, is an incredible testament to the power of the SM and the toolkit provided by QFT. The power of the SM, and its internal consistency, is illustrated in Figure 1.4, which shows the results of indirect measurements of the W-boson and top-quark masses based on fits to electroweak precision data. When all measurements other than those on m_h are included, the constraints of the SM only allow for a small region of the (m_{top}, M_W) parameter space. Adding the Higgs mass measurement results only shrinks this allowed area. The fact that these indirect measurements agree so well with the direct measurements paints a picture of the SM as being a fundamentally complete picture of the known physical phenomena. It is likely that if the SM were not accounting for particular types of interactions or fields, the level of agreement between the direct measurements and those obtained indirectly via fits to electroweak precision data would not be to the level seen in Figure 1.4.

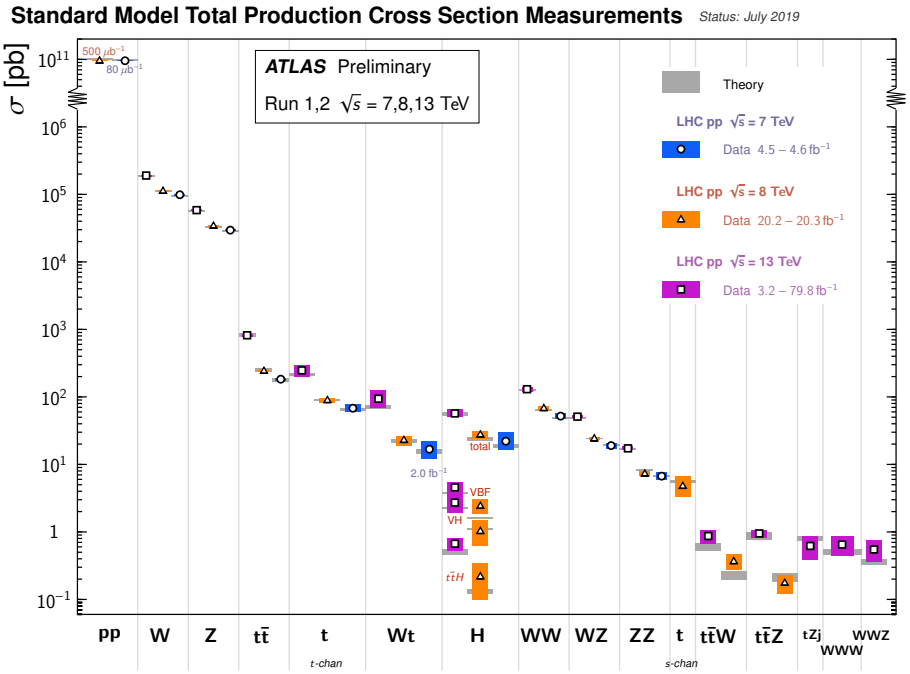

Figure 1.3: Summary of several Standard Model total production cross section measurements, corrected for branching fractions, compared to the corresponding theoretical expectations. Figure taken from Ref. [31].

With the discovery a Higgs boson like particle with a mass at 125 GeV in 2012 [36, 37], the final piece of the SM described by Equation 1.1 is potentially found. The 2012 Higgs discovery meant the start of a very long experimental program, focused on studying this new particle and confirming its role as being the fundamental scalar boson, ϕ, appearing in the SM.

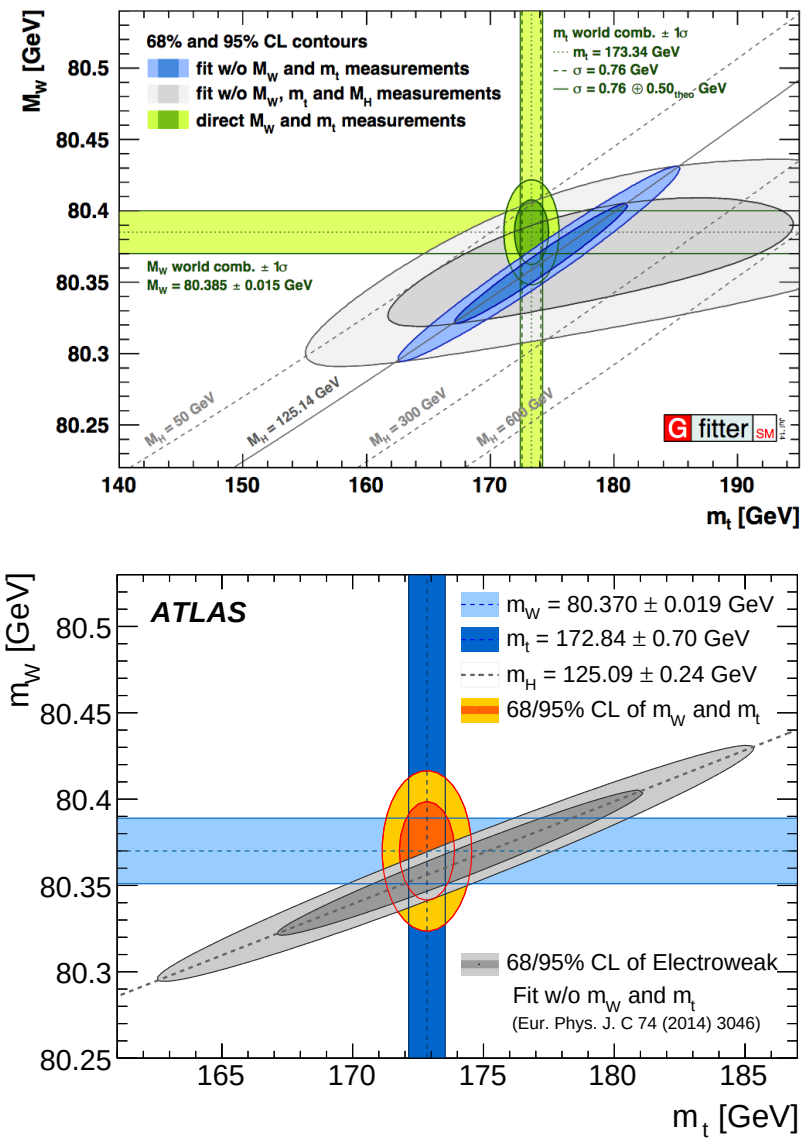

Figure 1.4: Contours at 68% and 95% CL obtained from scans of M_W versus m_{top}, for the global electroweak fits in comparison to the direct measurements. ***Top***: Including Higgs boson mass measurements in the indirect fit [32, 33] (blue) or excluding them (grey). Figure taken from Ref. [34]. ***Bottom***: Including the latest precision W-boson mass measurement from the ATLAS experiment. Figure taken from Ref. [35].

It should be stressed that the terms associated with the Higgs potential appearing in the SM Lagrangian, given by Equation 1.14, are by no means fundamental. They do not have to appear in this way. The terms appearing in Equation 1.14 take the form they do since they could lead to masses for the fermions and gauge bosons. There is no fundamental symmetry motivating their precise form. The BEH mechanism is inspired by that of super-conductivity, in which the formation of composite (i.e. not elementary) scalar particles — the Cooper pairs — occurs. The fact that the

same type of phase transition should describe the generation of masses for the elementary particles of the SM, and that it should presuppose the existence of an *elementary* scalar boson, was simply left as one of the last open questions of the SM. In a sense, the truth of the form underlying Equation 1.14 was not important to BEH. The more important takeaway was that there *could* be a mechanism by which the SM particles acquired mass without disrupting the fundamental gauge structure of the SM that had already held up to experimental scrutiny.

It is then up to the experiments to verify that the 125 GeV scalar boson discovered in 2012 is responsible for the BEH mechanism as described in Section 1.3. The form of the Higgs potential as defined in Equation 1.14 makes very clear predictions on the form and strengths of the couplings to the known fundamental particles: the gauge bosons and fermions, with couplings to the Higgs predicted to take the forms of Equations 1.21 and 1.24, respectively. It is then up to the ATLAS and CMS experiments to verify that the new particle couples to these 'old' particles just as predicted. The coupling strengths also dictate the Higgs decay rates into specific SM particles, as indicated in Figure 1.5. Any deviation with respect to the SM-predicted values in the measurement of the Higgs decay branching ratios, or in the values of the fermion or gauge coupling strengths and their dependence on the particle masses, would indicate that the particle discovered in 2012 is not the Higgs boson as predicted in the SM.

All of the measurements of the properties of the 125 GeV particle made by the ATLAS and CMS experiments are so far in fairly good agreement with the SM prediction of a $m_h = 125$ GeV Higgs boson [38, 39, 40, 41, 42, 43, 44, 45, 46]. The agreement with the SM prediction, over a wide variety of measurements, is illustrated in Figure 1.6 which shows the measurements of the fermion and gauge couplings and of the cross-sections of the leading Higgs production mechanisms and decay branching ratios. Within the precision of these measurements, the SM is fully supported by the experiments and it appears as though the particle discovered in 2012 is in fact the particle as predicted by BEH.

In addition to the Higgs couplings to the SM fermions and gauge bosons, the SM Lagrangian predicts terms describing the Higgs *self*-couplings. These terms are described by the λ parameter in Equation 1.14 and appearing in Table 1.2. The Higgs self-coupling parameter λ has a value predicted by the SM, as seen in Equation 1.26, which is fixed by the Higgs boson mass. The parameter λ is directly responsible for providing the structure of the Higgs potential, indicated by Equation 1.25, and therefore plays a fundamental role in EWSB. Measuring how the 125 GeV Higgs boson couples and decays to the SM fermions and gauge bosons, then, is only a necessary requirement for confirming that the particle is the one predicted by the SM: it is not sufficient. To truly confirm that the 125 GeV boson is indeed that predicted by the SM, and that EWSB is described by the BEH mechanism, the Higgs self-coupling parameter will have to be measured experimentally and confirmed to take its predicted value given by Equation 1.26.

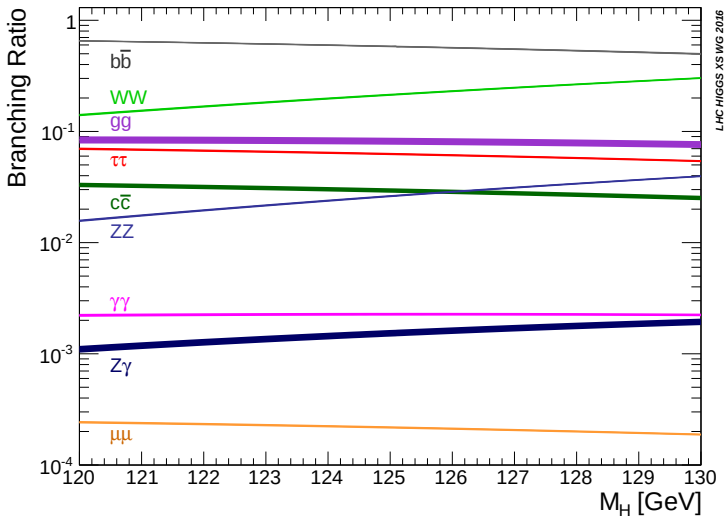

Figure 1.5: Predicted branching ratios for an SM-like Higgs boson with $m_h = 125\,\text{GeV}$. Figure taken from Ref. [47].

Direct measurement of λ proceeds only through the observation of events in which Higgs bosons are produced in pairs. At the LHC, this process occurs predominantly via gluon-gluon fusion through the two diagrams illustrated in Figure 1.7. The triangle and box diagrams shown in Figure 1.7 represent destructively interfering amplitudes. As a result, the cross-section associated with the production of Higgs boson pairs is exceedingly low and statistically significant observation of this process is not expected to occur until near the end of the lifetime of the LHC. The study of Higgs boson pairs is the subject of the analysis to be presented in Chapter 9.

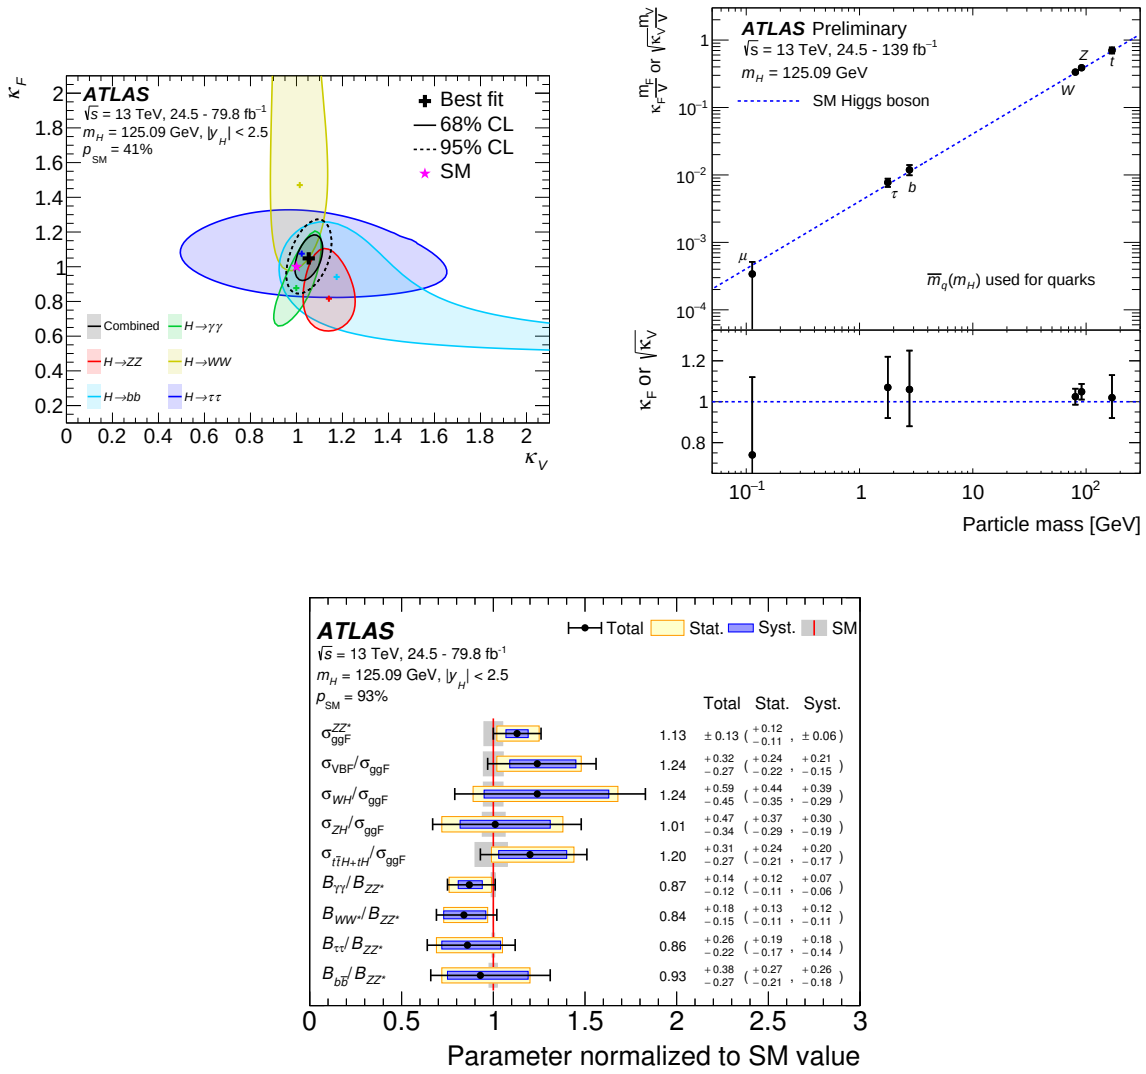

Figure 1.6: Higgs precision measurements of couplings to SM particles. Figures taken from Ref. [48]. ***Left***: Combined measurement of fermion and gauge-boson Higgs coupling modifiers, κ_f and κ_V (assumed to be universal across fermion and gauge-boson species in the result pictured). Values of κ_f or κ_V equal to 1 correspond to the SM prediction for the Higgs' couplings to these particles. ***Right***: Measured values of the Higgs fermion and gauge-boson coupling parameters as a function of the fermion and gauge-boson masses. The blue dashed line shows the SM prediction (Equations 1.21 and 1.24). ***Bottom***: Measurements of Higgs production cross sections and (relative) decay branching ratios.

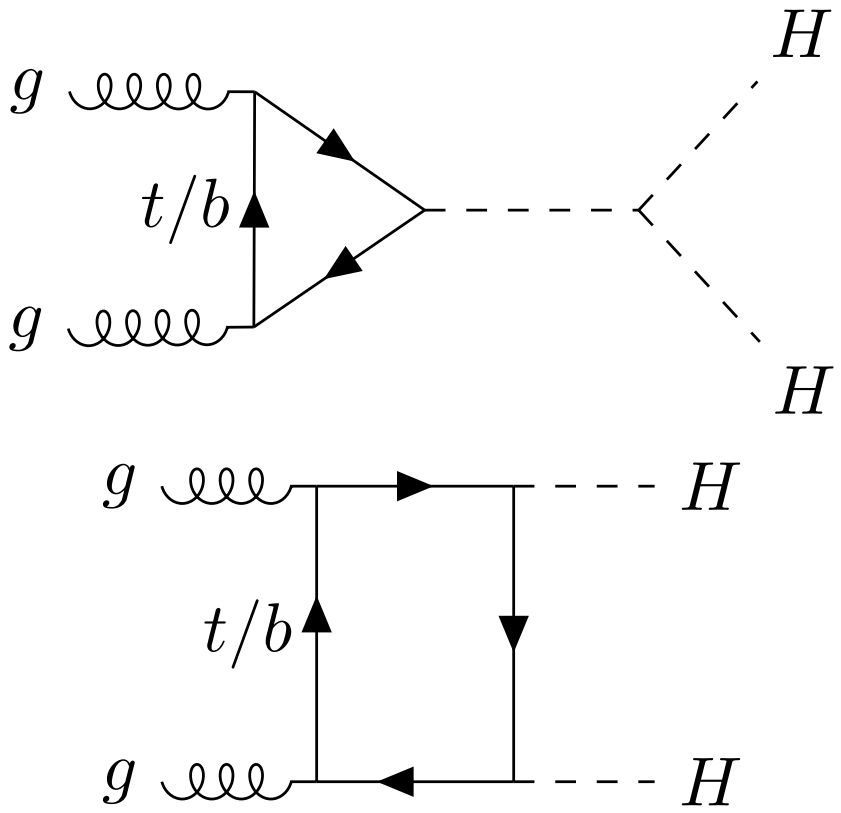

Figure 1.7: Representative Feynman diagrams that contribute at leading order in QCD to the non-resonant production of Higgs boson pairs. ***Top***: The 'triangle' diagram, sensitive to the Higgs self-coupling, λ. ***Bottom***: The 'box' diagram, sensitive to the (squared) Yukawa couplings to the third generation fermions entering the loop.

1.4.2 Perceived Shortcomings of the Standard Model and Open Questions

Despite the impressive results and predictive power of the SM illustrated in the previous section, the SM is unable to provide a complete picture of what we now consider the observed Universe. A few items that are clearly not explained within the framework of the SM are described below. Chapter 2 goes on to introduce an extension to the SM that provides explanations for many of these short comings of the SM.

Existence of Dark Matter (DM) and Dark Energy (DE)

Most of the experimental evidence that we currently have that supports the notion that physics beyond the SM (BSM) exists comes not from collider-based particle physics experiments, but from astrophysical observations. Astrophysical observations suggest that the majority of the matter content of the Universe is composed of a non-luminous, weakly interacting type of matter referred to as 'Dark Matter' (DM) and that the expansion of the Universe is accelerating, perhaps due to presence of 'Dark Energy' (DE) [49, 50]. There has not yet been experimental proof of the particle nature of DM, but many theories suggest that it fall under the class of 'Weakly Interacting Massive Particle' (WIMPs), i.e. that it be 'WIMP-like'. The particle nature of DE is completely unknown and without a fundamental description.

Massive Neutrinos

The observation of neutrino oscillations [51] provides evidence in support of neutrinos having nonzero masses. This directly contradicts the massless neutrino hypothesis of the SM.

Matter-Antimatter Asymmetry

The only parameter in the SM that allows for CP violation is the CP-violating phase in the CKM matrix, which is unable to account for the amount of CP violation needed to account for the observed asymmetry in the amount of matter over antimatter in the Universe. Additional CP violating effects are necessary to account for this observed asymmetry and should be relevant to early-Universe cosmology [52].

The Hierarchy Problem

The Hierarchy Problem refers to the fact that the electroweak sector, through the scalar Higgs boson, is sensitive to high energy cut-off scales nearing the Planck Mass, $M_P \approx 10^{18} - -10^{19}\,\text{GeV}$. This is evident when computing the higher-order corrections to the Higgs mass,

$$m_h^2 = m_{h,0}^2 + \Delta m_h^2,$$

which are found to be quadratically divergent due to the fact that the Higgs is not protected

by any fundamental internal symmetries:

$$\Delta m_h^2 = -\frac{|y_f|^2}{16\pi^2}\left[2\Lambda^2 + \mathcal{O}\left(m_f^2 \ln\left(\frac{\Lambda}{m_f}\right)\right)\right], \tag{1.27}$$

where y_f is the fermion Yukawa coupling, m_f is the associated fermion mass, and Λ is the ultra-violate cut-off scale. Similar terms also appear for the SM gauge bosons, resulting in their own contributions like Equation 1.27, but given the near-unity Yukawa coupling of the top-quark, the divergence is most evident from the top-quark loop contributions to Δm_h^2, such as that shown in Figure 1.8. The apparent contradiction that the electroweak scale at

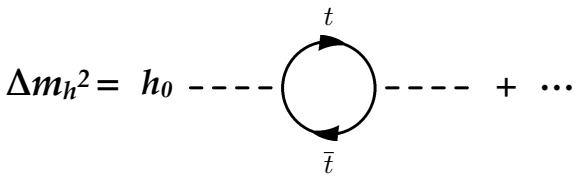

Figure 1.8: Top-quark loop contribution to the higher-order computation of the Higgs mass that leads to unregulated divergent terms such as Equation 1.27.

which the Higgs boson is experimentally known to exist should also be unavoidably sensitive to M_P, attempting to drive the Higgs boson mass to values 16 orders of magnitude larger than the measured one, is at the heart of the Hierarchy Problem. The high energy scales at which SM calculations fail are typically assumed to be the scales at which new physics arise and take over as a more complete description of natural phenomena. That the Higgs boson is apparently driven to these scales motivates the arguments in favor of their being new physics at or near the electroweak scale which act to cancel the quadratically divergent terms appearing in the Higgs mass corrections. That cancellations of terms on the order of 10^{30} (Λ^2) should be possible introduces the philosophical concerns of *naturalness*, which ponder whether or not Nature is such that the additional sources of high energy physics — not accounted for by the SM — can conspire in such a way as to allow for such perfectly fine-tuned cancellation.

Chapter 2

Physics Beyond The Standard Model

The only consolation he drew from the present chaos was that his theory managed to explain it.

–Thomas Pynchon, *V.*

As seen in Section 1.4, the SM provides an incredible toolkit from which the majority of the observed physical phenomena can be described. As we saw, however, there are several open questions whose answers must be pursued if we are to truly understand the underlying nature of the Universe. Given the scope of these open questions, and the apparent inconsistencies exposed by the Hierarchy Problem, their solutions are all but guaranteed to further revolutionize our view of the Universe. As a result, there have been a multitude of proposals of new theories of particles and fields over the preceding decades that attempt to solve some of these remaining mysteries.

With the discovery of the 125 GeV Higgs boson, the physics community is seemingly on the verge of gaining further insight into the nature of these problems. Any new source of physics beyond the SM (BSM) will necessarily have much to say about the electroweak and Higgs sectors of the SM, especially in view of the Hierarchy Problem's hinting at the electroweak scale being so near the portal to new physics. These sectors tightly constrain the forms that BSM physics can take, as well, if they are to remain consistent with the precise confirmation of the SM predictions over the past 50 to 60 years. One such constraint, a beautiful result of the electroweak theory and one that is fundamental to the nature of EWSB, is that of the ρ parameter of the SM, defined in Equation 2.1. The precise equality of Equation 2.1 can be seen to hold by using the relationships between the electroweak gauge mixings and couplings in Equations 1.9 and 1.20. The $\rho = 1$ equality holds to very high precision even when including higher order corrections to the gauge boson masses and interactions, a fact that his been verified by precision measurements [53, 54].

$$\rho = \frac{M_W}{M_Z \cos\theta_W} \underset{\text{SM}}{=} 1 \tag{2.1}$$

The requirement that $\rho = 1$ restricts the form any new BSM physics attempting to alter the electroweak sector can take. It is interesting to note that the last equality in Equation 2.1 can hold even for Higgs sectors that are far more complex than the minimal one predicted by the SM, and that are perhaps composed of *multiple* Higgs doublets. This can be seen by casting the ρ parameter into its more general form, shown in Equation 2.2:

$$\rho = \frac{\sum_{i=1}^{n} \left[T_{3,i}\left(T_{3,i}+1\right) - \frac{1}{4}Y_i^2\right] v_i}{\sum_{i=1}^{n} \frac{1}{2}Y_i^2 v_i}, \tag{2.2}$$

where n runs over the number of scalar Higgs doublets in the theory, T_3 and Y are the $\mathcal{SU}(2)$ and $\mathcal{U}(1)$ gauge quantum numbers, respectively, associated with each doublet, and the v quantities are the vacuum expectation values of each Higgs boson. In the SM, with only the single Higgs doublet having $Y = 1$ and $T_3 = 1/2$ (Table 1.1), the equality $\rho = 1$ clearly holds. Any extended Higgs sectors whose multiple Higgs doublets have $\mathcal{SU}(2)$ and $\mathcal{U}(1)$ gauge quantum numbers satisfying $\rho = 1$ are in principle allowed within the constraints of current precision electroweak measurements.

The possibility that the Higgs sector may be more complex without disrupting strict relationships predicted very cleary by the SM has led to the development of many fields of research. One such potential for BSM physics is the class of theories that minimally extend the Higgs sector by only including a single additional scalar double ($n = 2$ in Equation 2.2). Such theories are referred to as Two Higgs-Doublet Models ('2HDMs'), and their phenomenology is rich [55]. The two Higgs doublets are typically referred to as 'H_u' and 'H_d', where the subscripts indicate to which component of the $\mathcal{SU}(2)$ fermions they couple and give mass, and are complex scalar fields with eight degrees of freedom between the two of them. In these models, the three degrees of freedom needed for the $W^{\pm}$ and Z bosons to acquire mass after EWSB leave open five additional degrees of freedom. As a result, 2HDMs are characterized by an extended Higgs sector in which there are *five* physical massive Higgs bosons: two neutral components (h^0 and H^0), one neutral pseudoscalar particle (A^0), and two that are electrically charged ($H^{\pm}$). The neutral component h^0 is typically considered to be lightest Higgs boson and is identified as the 125 GeV boson discovered in 2012. Other than this assumption on h^0, the mass spectrum of these five Higgs bosons and their phenomenology is dependent on the specific instantiation of the 2HDM in question. As it turns out, one of the most promising all-encompassing theories of new physics, Supersymmetry (SUSY), is a 2HDM. Section 2.1 provides a brief introduction to SUSY that will be relevant as background to the analyses presented in Chapters 8 and 9.

2.1 Supersymmetry

SUSY is a theoretical framework that allows for the fermionic (spin-1/2) and bosonic (spin-1) fields to mix. That is, it allows one to relate the fields typically associated with matter with those typically associated with the forces and interactions by introducing the SUSY generators Q that allow to interchange a fermion with a boson:

$$Q\,|\text{fermion}\rangle = |\text{boson}\rangle\,, \qquad Q\,|\text{boson}\rangle = |\text{fermion}\rangle \tag{2.3}$$

As spin is fundamentally associated with a *spacetime* symmetry, the notion that the fermions and bosons could be related in such a way was famously forbidden by the 'no-go' theorem of Coleman and Mandula [56] which stated that the structure of the SM was maximal. That is, there could be no other form for the mixing of spacetime and internal symmetries other than by the direct product formulation of the SM: $\mathcal{P}\times\mathcal{SU}(3)_C\times\ \mathcal{SU}(2)_L\times\ \mathcal{U}(1)_Y$. The work of Ref. [57] provided a way out, however, and was able to show that spacetime and internal symmetries could mix more generally if the generators, Q, were *fermionic* — with spin no greater than 1/2 — rather than bosonic, as is the case in the SM. That four dimensional SUSY was shown to be possible, and that matter-type fields and particle-type fields are possibly two sides of a more fundamental coin, is an exciting proposition and opened up SUSY as one of the more enticing avenues for the search for new physics.

SUSY proposes that each SM particle be partnered with a new particle, its 'superpartner', whose spin differs by 1/2 unit *less* than the SM partner. These superpartners are then grouped into the 'supermultiplets' on which the SUSY algebra operates. The structure of the supermultiplets is as follows,

$$\underbrace{\begin{pmatrix} f_{1/2} \\ \tilde{\phi}_0 \end{pmatrix}}_{\substack{\text{Chiral} \\ \text{Supermultiplet}}} \qquad \underbrace{\begin{pmatrix} V_1 \\ \tilde{V}_{1/2} \end{pmatrix}}_{\substack{\text{Gauge} \\ \text{Supermultiplet}}} \qquad \begin{cases} f_{1/2}: & \text{Spin-1/2 Weyl fermion} \\ \tilde{\phi}_0: & \text{Complex scalar} \\ V_1: & \text{Spin-1 boson} \\ \tilde{V}_{1/2}: & \text{Spin-1/2 Weyl fermion} \end{cases}, \tag{2.4}$$

with the spin-1/2 fermions of the SM being paired with their scalar SUSY partner in the 'chiral supermultiplets' and the spin-1 gauge bosons of the SM being paired with fermionic SUSY partners in the 'gauge supermultiplets'. The newly added SUSY particles are labelled with a '$\sim$' above their label so as to distinguish them relative to the well-known SM counterparts. The structure of these multiplets implies a *doubling* of the particle content of the SM, which has broad consequences on the ways that SUSY can potentially manifest itself and necessarily implies a rich phenomenology.

The additional particle content implied by a minimal instantiation of SUSY, referred to as the Minimally Supersymmetric SM ('MSSM'), is presented in Table 2.1. The squarks and sleptons are the SUSY partners of the quarks and leptons. The neutralinos and charginos are the electroweak partners to the SM electroweak gauge bosons. It should be noted that the squarks and sleptons are scalar, meaning that they are inherently not chiral: the subscripts L and R appearing in their name in Table 2.1 refer instead to the SM particle to which they are partnered. The f_L and f_R can generally mix to form the physical fields, f_1 and f_2; however, it is generally assumed that only those in the third generation mix appreciably. The physical fields have subscripts $\{1, 2, ...\}$ with the lower numbers referring to the particles with the lower mass. For example, $m_{\tilde{t}_1} < m_{\tilde{t}_2}$. Table 2.1 also shows that the SUSY partners to the two Higgs doublets, $\tilde{H}_d^0$ and $\tilde{H}_u^0$, mix with the SUSY partners of the electroweak gauge bosons to form the mass eigenstates of the charginos and neutralinos. The structure of this electroweak gauge mixing in the MSSM leads to complicated mass hierarchies of the $\tilde{\chi}^0$ and $\tilde{\chi}^\pm$ particles, meaning that direct searches for these particles in particle collider experiments will be challenging.

In most MSSM scenarios, it is assumed that an additional symmetry, referred to as 'R-parity', exists and that the SM particles have R-parity quantum number $+1$ and the SUSY particles -1. For the types of SUSY relevant to the present thesis, it is assumed that R-parity is absolutely conserved. SUSY scenarios having conserved R-parity are referred to as 'R-parity conserving (RPC) SUSY' scenarios.[1] RPC SUSY scenarios imply that the lightest supersymmetric particle ('LSP') must be absolutely stable. The current lack of any directly observed additional particle content in the Universe implies that the LSP, while stable, must also be weakly interacting and electrically neutral. As a result, the LSP is typically taken to be $\tilde{\chi}_1^0$. One major consequence, then, of RPC SUSY is that it provides a natural candidate for WIMP-like[2] DM via the $\tilde{\chi}_1^0$ fermion.

An additional consequence of SUSY is that the Hierarchy Problem, described in Section 1.4, can potentially be resolved entirely. This can be seen by computing the additional loop contributions to the Higgs mass corrections arising as a result of fundamental scalars. The additional loops from such scalar particles are illustrated in Figure 2.1. The scalar contributions take the form of Equation 2.5,

$$\left(\Delta m_h^2\right)_S = \frac{\lambda_S}{16\pi^2}\left[2\Lambda^2 - \mathcal{O}\left(m_S^2 \ln\left(\frac{\Lambda}{m_S}\right)\right)\right], \tag{2.5}$$

where m_S is the scalar mass and λ_S is its Yukawa-type Higgs coupling parameter. These scalar loop corrections are seen to suffer the same type of quadratic divergence as in the case of Equation 1.27 but they enter *with an opposite overall sign.* Combining the fermionic and scalar loops from particles

[1] Alternatively, those scenarios that violate R-parity conservation are referred to as R-parity violating, or RPV, SUSY scenarios.

[2] The abbreviation 'WIMP' stands for 'weakly-interacting massive particle'.

Table 2.1: Additional particle content of the MSSM, including the extended SM (R-parity +1) Higgs sector.

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

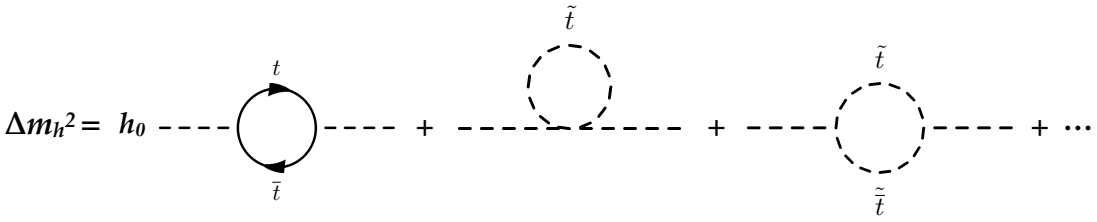

Figure 2.1: Loop contributions to the higher-order computation of the Higgs mass as a result of the SM top-quark, t, and its scalar superpartner, $\tilde{t}$.

in the same chiral supermultiplet leads to an overall correction of the form shown in Equation 2.6:

$$\Delta m_h^2 \underset{\text{w/ SUSY}}{\propto} \left(m_f^2 - m_S^2\right) \ln\left(\frac{\Lambda}{m_S}\right) + 3m_f^2 \ln\left(\frac{m_S}{m_f}\right) + \mathcal{O}\left(\frac{1}{\Lambda^2}\right) \tag{2.6}$$

The remarkable feature that Equation 2.6 shows is that if SUSY is *exact*, that is, if the particles within the same supermultiplet enter with the same mass ($m_f = m_S$) and have the same couplings ($y_f = \lambda_f$), then the problematic divergence underlying the Hierarchy Problem disappears! This is a very suggestive result that we immediately know, however, not to be the case in reality given the fact that additional particles with the exact same masses of the known SM particles would have been produced in the laboratory environment ages ago. If SUSY is to exist as described by the MSSM (Table 2.1), it must then be a *broken symmetry* wherein the masses of the SUSY partners are larger than their SM counterpart — large enough such that they would have escaped experimental detection in high energy particle colliders up until the present day.

As mentioned in Section 1.4.2, however, the experimental observation of a relatively light Higgs boson is suggestive of new physics entering at the electroweak scale. If this assumption is taken to heart, and if SUSY is to explain away the Hierarchy Problem, among other things, then the implied

cancellation in Equation 2.6 should still hold, at least approximately. This implies that the masses of at least the lightest SUSY partners cannot be too large and should be no larger than $\mathcal{O}(\text{TeV})$,

$$|m_f^2 - m_S^2| \leq \mathcal{O}(\text{TeV}^2). \tag{2.7}$$

As the leading contribution to the loop corrections to the Higgs mass arise due to the SM top-quark loops, the relation in Equation 2.7 implies that the scalar SUSY partners to the top-quark should not be much larger than $\mathcal{O}(\text{TeV})$. If this is true — that TeV-scale SUSY partners to the SM top-quark exist — they should in-principle be able to be produced at the center-of-mass collision energies available at the LHC.

The above discussion places a fundamental importance on the SUSY partner to the SM top quark. Apart from simply producing them directly within the *pp* collisions at the LHC, their presence may be observed *indirectly*. In Section 1.4.1 the fundamental importance of the production of Higgs boson pairs was introduced, with the leading production mechanisms in the SM indicated by the diagrams in Figure 1.7. If the SUSY partner to the top-quark exists, and is not too heavy as suggested by Equation 2.7, then it may contribute additional diagrams leading to final state pairs of Higgs bosons. Potential, additional single-loop contributions of this type are illustrated in Figure 2.2. Depending on the exact configuration of the MSSM, then, the presence of the scalar partner of the SM top-quark could potentially lead to non-SM values for the Higgs self-coupling parameter λ and also to *enhanced* rates of Higgs boson pairs being produced at the LHC [58]. This latter fact is illustrated in Figure 2.3, which shows the general trend — relatively independent of the MSSM configuration — that low-mass scalar partners to the top-quark can lead to appreciable increases in the Higgs boson pair-production cross-section. This is an interesting consequence of the scalar partners of the top-quark being relatively light. Given the exceedingly low SM-predicted cross-section for Higgs boson pair-production at the LHC, observation of Higgs boson pairs in the current and next runs of the LHC may imply that SUSY is indeed hiding near the electroweak scale.

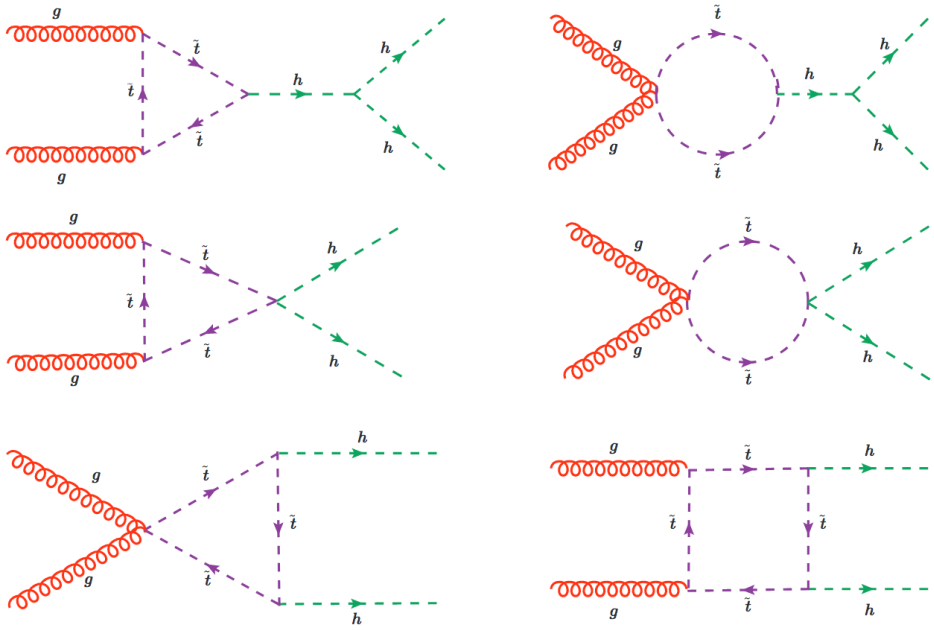

Figure 2.2: Single-loop stop-quark diagrams in the MSSM that contribute to the non-resonant production of Higgs boson pair, in addition to those already predicted in the SM (Figure 1.7). Figure taken from Ref. [58].

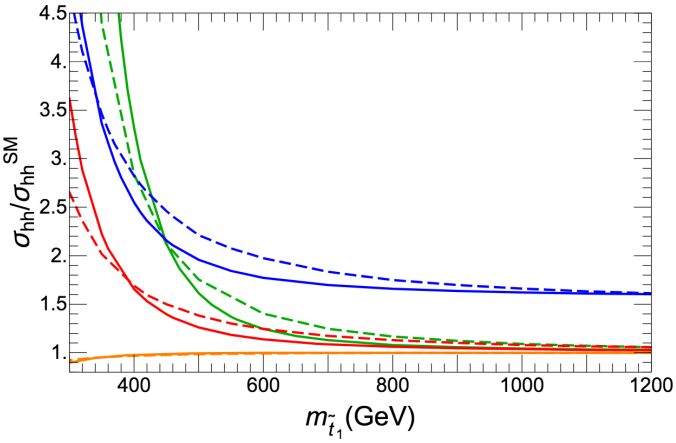

Figure 2.3: Higgs pair production cross-section normalized to the SM prediction as a function of the mass of the lighter stop quark, $\tilde{t}_1$. Further details are in Ref. [58], the source of the figure.

Chapter 3

Experimental Setup

> *I know of no more encouraging fact than the unquestionable ability of man to elevate his life by a conscious endeavour. It is something to be able to paint a particular picture, or to carve a statue, and so to make a few objects beautiful; but it is far more glorious to carve and paint the very atmosphere and medium through which we look, which morally we can do.*
>
> –Henry David Thoreau, *Walden*

The work to be described in the present thesis was done at CERN[1], the particle physics laboratory located along the French-Swiss border just outside of Geneva, Switzerland. CERN is comprised of almost 18,000 personnel, of which over 13,000 are researchers in the field of experimental particle physics. It is a truly international workplace, with the personnel comprised of representatives of over 110 nationalities and who are either working directly for CERN[2] or for their respective home institutions — universities or national labs — located across more than 70 countries worldwide [59]. These researchers will generally work at any of the independent experiments located along the various beamlines that network throughout the CERN campus (see Figure 3.1).

At the time of writing, there are four large experiments[3] taking place currently at CERN, all located along the Large Hadron Collider (LHC): ALICE [60], LHCb [61], CMS [62], and ATLAS [63]. The CMS and ATLAS detectors are general purpose detectors, with broad research programs, whereas

[1] The acronym CERN was historically derived from '*Conseil européen pour la recherche nucléaire'*. Nowadays, 'CERN' has become a standalone name for the lab itself and is currently referred to as the '*Organisation européenne pour la recherche nucléaire*'; or, in English: the '*European Organisation for Nuclear Research.*'

[2] Of the roughly 18,000 researchers in experimental particle physics, only about 5% are employed directly by CERN itself.

[3] For the most part, one can interchange the words 'detector' and 'experiment' when referencing large-scale, long-term particle physics experiments such as those that have taken place over the past few decades: the detectors tend to take on the role of representing the entire collaboration of physicists, engineers and associated personnel, as well as the entire scope of the associated research programs.

the ALICE and LHCb detectors are specialised for the study of heavy-ion collisions and b-hadron physics, respectively.

This chapter will present a brief introduction to the workings of the LHC in Section 3.1. In Section 3.2, given that the present author is a member of the ATLAS collaboration, a detailed description of the various components that make up the ATLAS detector will be presented.

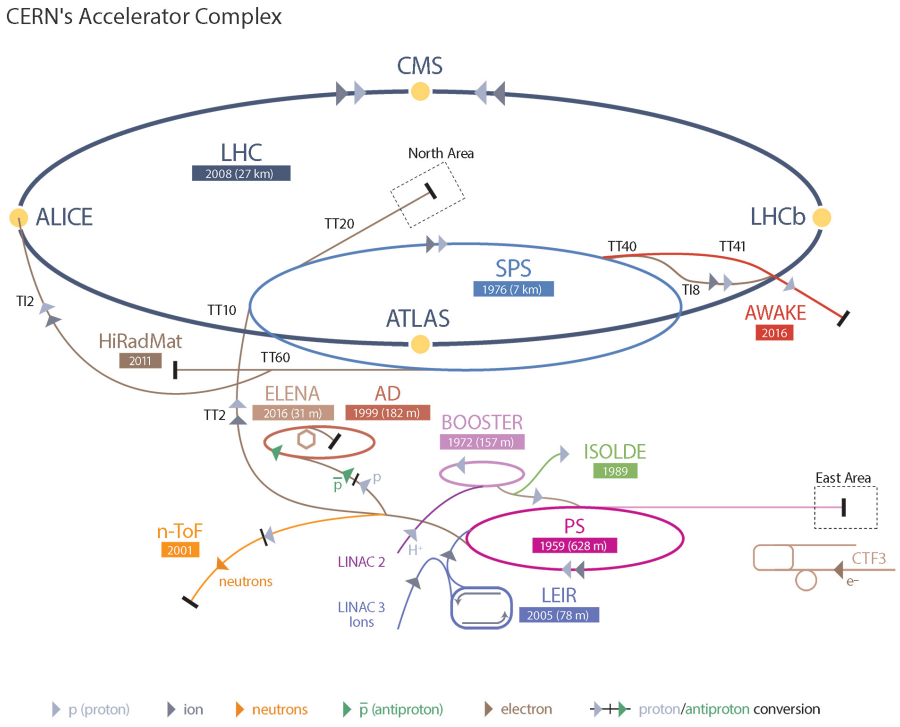

Figure 3.1: Illustration of the various beamlines, accelerator and storage rings, and experimental points that the CERN accelerator complex is home to. The protons that circulate through the LHC, and that are eventually made to collide inside the ATLAS detector, follow the path: Linac 2 → Booster → Proton Synchotron (PS) → Super Proton Synchotron (SPS) → LHC. Figure taken from Ref. [64].

3.1 The Large Hadron Collider

The LHC [65] is a circular particle accelerator, with a 27 kilometer circumference, located at an average distance of 100 meters beneath the surface of the Earth. It is nominally used for proton-proton (pp) collisions, wherein two counter-rotating beams of protons are made to collide head-on at specific interaction points (IP) along the 27 kilometer ring, but can also be run in heavy-ion

configurations wherein proton-lead (p-Pb) or lead-lead (Pb-Pb) collisions take place.[4][5] The pp collisions take priority over those of the heavy-ions, with the collisions each year consisting of only a few weeks in the winter for the heavy-ion configurations and typically six to seven months for the pp configuration. The LHC is designed to accelerate protons to a center-of-mass energy of $\sqrt{s} = 14\,\text{TeV}$.

3.1.1 Machine Design and Layout

Machine Composition

The LHC was planned as the successor to the Large Electron Positron (LEP) collider [66, 67], which was in operation between the years of 1989 to 2000. LEP is still the most powerful lepton collider to date, having maximal electron-positron center-of-mass collision energies of 209 GeV. After LEP, the particle physics community knew that the next collider at CERN needed to have $\mathcal{O}(10)$ TeV collision energies; either to be able to probe from all angles any new physics discovered at LEP and/or the Tevatron [68], or to provide the necessary power to search for still-elusive hints of BSM physics. At the very least, given a non-discovery of the Higgs boson at LEP and the Tevatron, the community would need a discovery machine powerful enough to produce electroweak-scale Higgs bosons and an $\mathcal{O}(10)$ TeV hadron collider — as we now know — is sufficient for this job.

In order to increase center-of-mass collisions energies, collider designs can take two routes: they can either increase in size, that is, have larger circumferences (radii), or they can increase the strength of the magnetic fields used to keep the circulating charged particles in orbit. This can be seen by first considering the expression for the relativistic cyclotron frequency, ω, of a particle moving in a circular orbit,

$$\omega = \frac{qB}{\gamma m}, \tag{3.1}$$

where m is the particle's rest mass, B is the magnitude of the magnetic field experienced by the particle, q is the particle's electric charge, and γ is the relativistic Lorentz factor, $\gamma = \sqrt{1-\beta^2} = \sqrt{1-(v/c)^2}$, with v the particle's velocity and c the speed of light. Using Equation 3.1, it can be seen that a particle of higher energy confined to a fixed circular orbit necessarily has a higher angular velocity by relating the particle's angular velocity to its kinetic energy:

$$E_{\text{kin}} \propto mv^2 = m(\omega R)^2 = \frac{q^2 B^2 R^2}{m\gamma^2}. \tag{3.2}$$

In planning the construction of the LHC, the costs in civil engineering and real-estate works that

[4] The specific lead (Pb) species used in collisions is $^{208}_{82}$Pb.

[5] More rarely, the LHC can also be used to circulate gold (Au) ions. There are even plans to have proton-oxygen (p-O) runs in the future, which will allow for the LHC experiments to provide research that potentially complements dark matter research based on cosmic-ray air showers.

would be required to construct a larger tunnel in which to house the LHC ring (increasing R) far outweighed the costs of research into and development of magnet systems strong enough to bend the multi-TeV particles along the beam orbit prescribed by the already-existing LEP tunnel (increasing B). The desired center-of-mass collision energy of $\mathcal{O}(10)$ TeV, the fact that the LHC would be a hadron (proton) collider, and the fact that the LHC would be using the existing LEP tunnel dictate the required magnetic field strength needed to keep the protons in stable orbits at the LHC. This is seen by using Equation 3.2, solving for B, and comparing the LHC and LEP design parameters,

$$\begin{aligned}\frac{B^2_{\text{LHC}}}{B^2_{\text{LEP}}} &= \frac{(E_{\text{LHC}} m_{\text{LHC}} \gamma^2_{\text{LHC}})/(q^2_{\text{LHC}} R^2_{\text{LHC}})}{(E_{\text{LEP}} m_{\text{LEP}} \gamma^2_{\text{LEP}})/(q^2_{\text{LEP}} R^2_{\text{LEP}})} \\ &= (E_{\text{LHC}}/E_{\text{LEP}}) \times (m_{\text{LHC}}/m_{\text{LEP}}) \times (\gamma^2_{\text{LHC}}/\gamma^2_{\text{LEP}}) \times (q^2_{\text{LEP}}/q^2_{\text{LHC}}) \times (R^2_{\text{LEP}}/R^2_{\text{LHC}}) \\ &\approx (1\,\text{TeV}\ /0.2\,\text{TeV}) \times (m_p/m_e) \times (1) \times (1) \times (1) \\ &\approx 10^4,\end{aligned} \tag{3.3}$$

which shows that the strength of the LHC bending magnets must be on the order of 100$\times$ the strength of those used at LEP. The magnetic fields experienced by the electron and positron beams at LEP were 0.22 Tesla. By Equation 3.3, the LHC bending magnets should achieve magnetic field strengths on the order of 10 Tesla in order to achieve the desired collision energies. The maximum achievable magnetic field of conventional ferromagnets is about 2 Tesla. To meet the $\sqrt{s}$ $\approx$10 TeV design goal, the magnet system used by the LHC to confine the protons to their circular orbits must then be composed of *superconducting* electromagnets. The entire magnet system of LEP was therefore removed and replaced with superconducting niobium-titanium (Nb-Ti) alloy based electromagnets which are superconducting at temperatures below 10 K. To reach temperatures below this 10 K threshold, the LHC magnets are housed in cryostats that allow for the Nb-Ti elements to be fully submerged in a bath of superfluid Helium at a temperature of 1.9 K [69]. In total, the LHC contains more than 120 tonnes of superfluid Helium. It goes without saying that there is a significant amount of resources and person power at CERN devoted to the refrigeration and cryogenics systems that are required for the LHC to run.

Additionally, the fact that LEP was a *particle-antiparticle* collider meant that the counter-rotating beams could be made to occupy a single ring: the same magnetic field could produce counter-rotating beams of electrons and positrons within the same beam pipe.[6] As a result, the LEP beam tunnel was constructed with only a single ring in mind and is relatively narrow: the LEP tunnel, and therefore LHC tunnel, is only $\approx$3.7 m wide on average. As the LHC is a *particle-particle* collider, it necessarily requires *two* magnetic fields of opposing polarity to circulate one of its beams in the clockwise direction and the other in the counter-clockwise direction. Given the limited space

[6]The electrons and positrons at LEP were vertically separated within the beam pipe by electrostatic separators placed throughout the LEP ring. Turning off these separators is, to first approximation, how the LEP operators would get the electrically-attracting electrons and positrons to collide.

in the tunnel, however, it is not possible to house two separate rings of superconducting bending magnets with all of the services that they require *in addition* to the requisite minimal space needed for personnel and maintenance access. This forced the need of the so-called '2-in-1' design of the main bending magnets of the LHC, wherein the two beam pipes are housed in the same cryostat in which the counter-rotating beams are held in their respective orbits by coupled magnetic fields. An illustration of this now-iconic design of the LHC bending (dipole) magnets and surrounding cryostat and containment structure is illustrated in Figure 3.3. Each of the 15 meter long superconducting dipole electromagnets of the LHC responsible for constraining the protons to their circular orbits has currents of 11850 Amperes flowing through it and achieves magnetic field strengths of 8.33 Tesla.

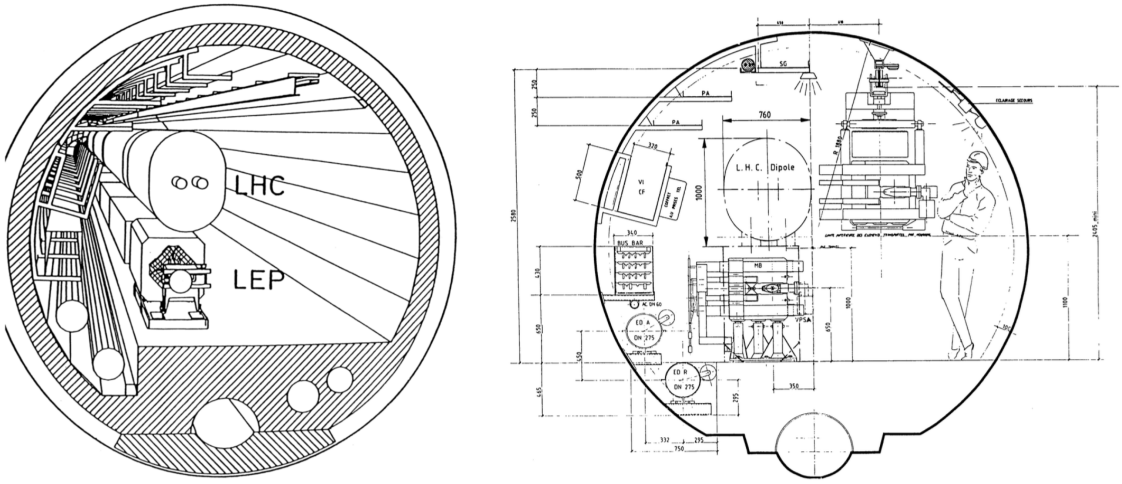

Figure 3.2: Figures are taken from Ref. [70]. ***Left***: Illustration comparing the size of a prototype '2-in-1' LHC dipole configuration to the LEP dipole and how they fit inside of the LEP/LHC tunnel. Note that prior to LHC operation, the LEP magnets will have been removed: the two are shown side-by-side for comparison purposes only. ***Right***: Cross-sectional view of the LEP/LHC tunnel with a comparison of the LHC '2-in-1' dipole on top of the LEP dipole. An illustration of an average size person is shown for scale. Also shown is the service crane in use, to give an idea of the size required for potential maintenance access. Clearly, two single-bore, superconducting rings each similar in size to the LEP dipole would not fit comfortably in the tunnel. The LHC '2-in-1' design fits in nearly the same area as the LEP dipoles while additionally being able to contain both particle beams.

Connecting the Dots

The LHC is essentially a chain of superconducting magnets of the type described in the previous paragraphs, where the bending (dipole) magnets critical to the LHC design were introduced. The chain is laid in a double-octagonal structure, illustrated in Figure 3.4. There are eight octants, at the center of which the LHC ring is straight and does not curve. The LHC curvature occurs at the boundaries of each of the octants and is primarily made up of bending (dipole) magnets. The straight sections are where the interaction regions are located and are referred to as 'Points',

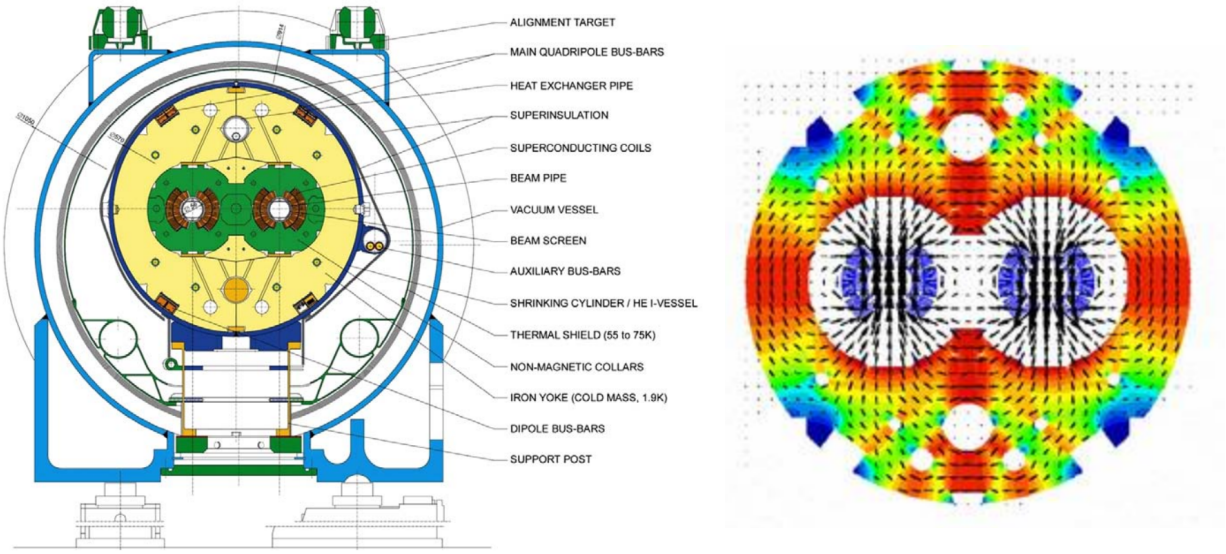

Figure 3.3: Figures taken from Ref. [65, 71]. ***Left***: Cross-sectional view of an LHC dipole bending magnet, with relevant parts indicated. The protons orbit inside of the beam pipes, each of which has a diameter of roughly 3 cm. It is interesting to note that the non-magnetic steel collars (in green) are of critical import to the success of the magnet systems. They are required to prevent the dipole structure from being deformed or torn apart due to the intense magnetic forces tending to push the two beam-pipes apart as a result of their counter-rotating electromagnetic currents. These forces amount to about 400 tonnes per meter of dipole when in full operation — almost equivalent in magnitude to the weight of a Boeing 747. ***Right***: Magnetic field lines of the coupled dipole fields that bend the counter-rotating proton beams and keep them in their circular orbits around the LHC ring.

numbered 1 to 8. Points 1, 2, 5, and 8 are where the four large LHC experiments are located. Points 1 and 5 are home to the services and underground areas of the general purpose experiments, ATLAS and CMS, respectively. The underground experimental caverns associated with Point 1 and 5 were not present for LEP and had to be constructed after LEP was retired in 2000. Figure 3.5 provides an illustration of how the surface and underground areas are situated at Point 1. Points 2 and 8 host the services and underground areas of the ALICE and LHCb experiments, respectively. At these Points, Points 1, 2, 5, and 8, the counter-rotating beams are made to collide. The remaining Points, Points 3, 4, 6, and 7, are host to various beam 'services' necessary for the operation of the LHC. Point 3 and 7 host the beam betatron and momentum cleaning ('beam collimation') systems, respectively. Point 4 hosts the superconducting radio-frequency (RF) systems which accelerate the beams to their nominal collision energies. Point 6 is the location of the beam-abort system — the so-called 'beam dump' — where the LHC beams may be removed very quickly from the LHC ring by using *kicker* magnets [72] that divert the beams out of the LHC ring in a safe manner. The beams may be dumped if the LHC wishes to refill with protons (or heavy-ions) and needs to remove any remnants of the previous fill, in case of beam instabilities observed in the LHC ring, or if one

of the experiments signals the need for a beam dump (in case of beam stability or detector issues observed at the associated IP).

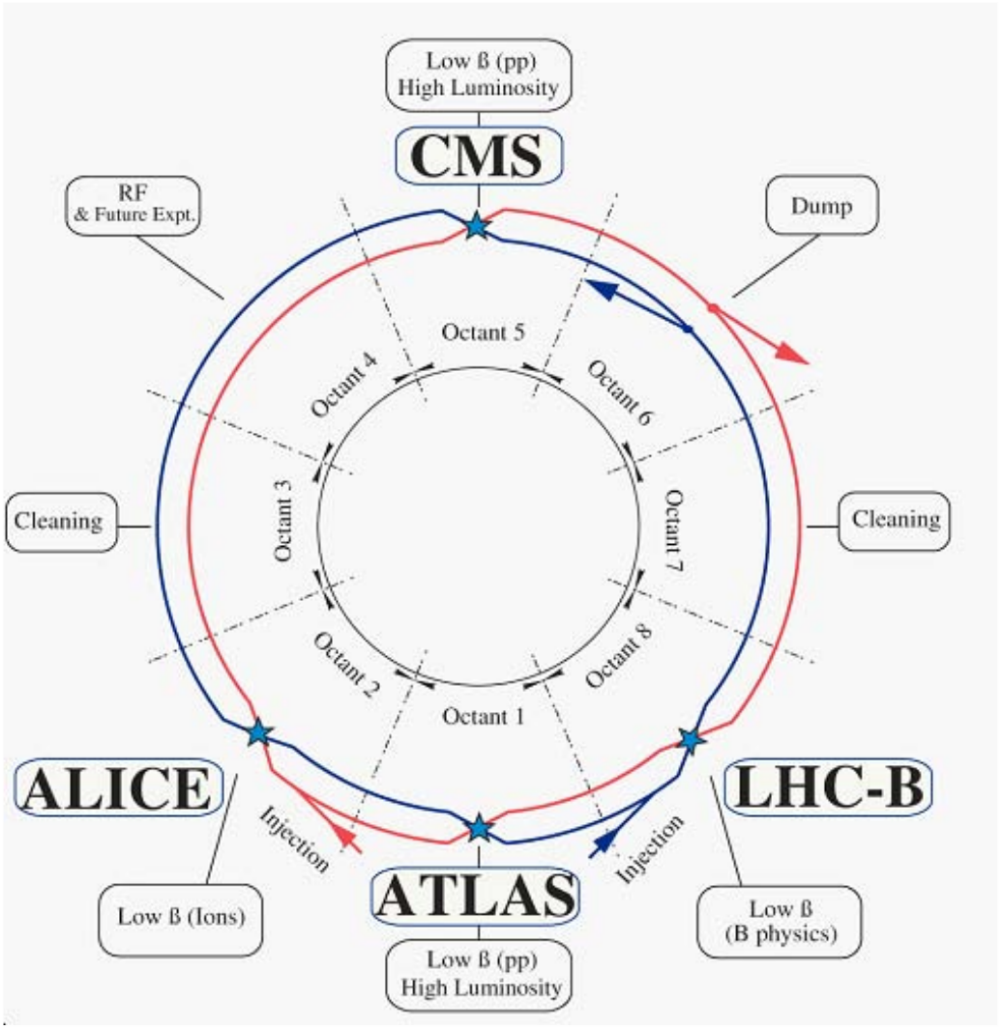

Figure 3.4: Layout of the LHC and its two counter-rotating beams. Beam 1 is in blue and rotates counter-clockwise. Beam 2 is in red and rotates clock-wise. At the center of each octant is a straight section which houses the experimental caverns or LHC beam facilities. At the boundaries of each octant are located the curved sections. Figure taken from Ref. [65].

3.1.2 Injection Chain and Bunch Structure

We now have an idea of how the proton beams relevant to the work in this thesis are made to circulate in the LHC ring. In this section we will briefly describe the initial source of the protons and how they are introduced into the LHC ring. The LHC relies on a series of pre-acceleration steps that bring the initial low-energy protons to energies sufficient enough to begin their journey through the LHC. The sum-total of these steps is referred to as the LHC *injection chain* [74]. The

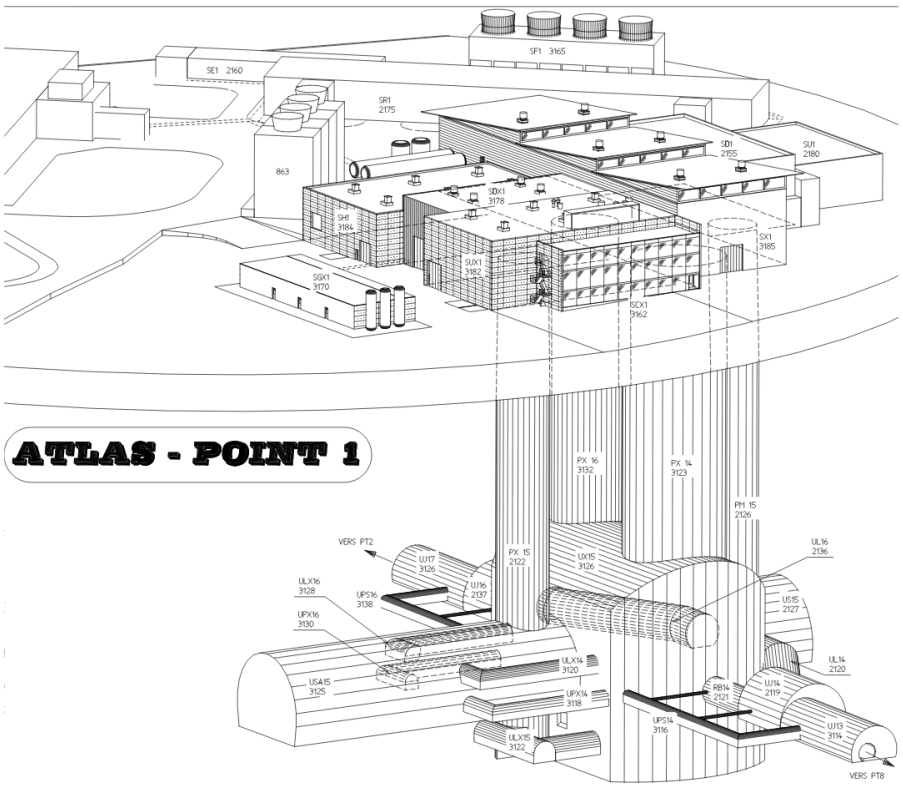

Figure 3.5: Diagram showing the surface buildings and services and underground areas of Point 1, where the ATLAS experiment is located. The LHC ring can be seen at the bottom, with its directions indicated by the 'VERS PT8 (2)' arrows pointing towards Point 8 (2). The experimental cavern in which the ATLAS detector sits is UX15. The control room for the ATLAS experiment, whereat operators can monitor and control the state of the ATLAS detector, is located 100 m above UX15 in the building SCX1. Figure taken from Ref. [73].

components of the LHC injection chain form the heart of the CERN accelerator complex illustrated in Figure 3.1. For *pp* collisions in the LHC, the protons are initially sourced from Hydrogen atoms that are released at the start of Linac 2. The Hydrogen atoms are immediately stripped of their electrons after passing through the *duoplasmatron* ion source [75]. The protons are then passed through Linac 2, a linear accelerator, which accelerates the protons to 50 MeV. They then enter the Proton Synchotron Booster (PSB), a circular storage ring composed of four stacked rings, which accelerates the protons to 1.4 GeV. The PSB injects the protons into the Proton Synchotron (PS) which accelerates them to 25 GeV. The Super Proton Synchotron (SPS) receives the protons from the PS and accelerates them to 450 GeV. At this point the protons have sufficient energy to be injected into the LHC. There are two injection points into the LHC since, up until this stage, the protons are circulating in the same direction: one injection point sends protons into the counter-clockwise beamline of the LHC, and the other into the clockwise beamline. Until all of the protons from a single *fill* make their way into the LHC, they will circulate at the injection energy

of 450 GeV. After the filling completes[7], the superconducting RF cavities located at Point 4 will begin to accelerate the protons to their final collision energies.[8] The acceleration is achieved by increasing the frequency of the RF oscillations; however, given that a 450 GeV proton is already ultra-relativistic, the adjustment of the frequency needed to get to the collision energies is not large but still has consequences for the design of the detector and measurement apparatus along the LHC ring [76].

The proton beams circulating the LHC are not a continuous stream of protons; rather, they are grouped into what are referred to as *bunches*. The protons arrive at the LHC in these bunches which are initially prepared in the smaller machines that make up the LHC injection chain and then are kept in their final *bunch structure* by the RF cavities. The accelerating RF cavities provide an accelerating electromagnetic field that oscillates longitudinally. The bunches, each composed of roughly 10^{11} protons, are then made to oscillate longitudinally in so-called *synchotron oscillations* around the central node of the RF oscillation as they circulate through the LHC ring. The proton bunches are then effectively 'shaped' by the oscillating RF field: protons in a bunch lagging behind or that are ahead of those particles at the center of the bunches will be accelerated or decelerated accordingly so as to be pushed back into the center of the bunch. The LHC RF cavities have an oscillation frequency of 400 MHz which defines the boundaries in which proton bunches can lie. These boundaries are called *RF buckets* and, along with the circumference of the LHC, dictate the number of proton bunches that can potentially fit in the LHC. The relationship between the RF oscillations and the RF bucket and bunch structure is illustrated in Figure 3.6. In total, approximately 35640 RF buckets exist when the LHC is in operation. Not all buckets contain proton bunches, however. In fact, at the time of the writing of the present thesis, RF buckets filled with proton bunches have a minimal separation of 10 RF buckets, meaning that following an RF bucket containing a proton bunch there is at least 9 unfilled RF buckets. This corresponds to a minimal time between proton bunches — the *bunch spacing* — of 25 nanoseconds. At the time of the present thesis, the operating conditions of the LHC maximally allow for 2808 25 ns-spaced bunches.[9] The bunch-spacing and overall bunch structure of an LHC fill is not only decided by the operators of the LHC but also by what the detectors at Points 1,2,5, and 8 can tolerate. This is because shorter bunch spacing means higher intensity and multiplicity of collisions occurring at each of these IP. A 25 nanosecond bunch spacing corresponds to a maximal pp collision rate of 40 MHz. The detectors at each of the IP have been designed with this collision rate in mind and anything higher may push them beyond their design limits.

[7] A standard LHC fill takes on the order of 4 minutes per ring.

[8] If all goes smoothly, this acceleration stage takes roughly 20 minutes.

[9] The number of allowed bunches is significantly lower than the 35640 RF buckets with 25 ns bunch-spacing potentially allow for. This is due, in part, to the non-trivial bunch-structure typically employed but also in large part to the fact that there is a $\approx 30\,\mu$s *abort gap* in the LHC ring where no filled RF buckets exist. The abort gap is a number of continuous unfilled RF buckets that allows the ramp up of the kicker magnets used for the beam dump to occur in the absence of filled buckets. In this way, the kicker magnet ramp up does not disturb the structure of the circulating proton beams. Only after this ramp up is finished should the kicker magnets disturb the beams.

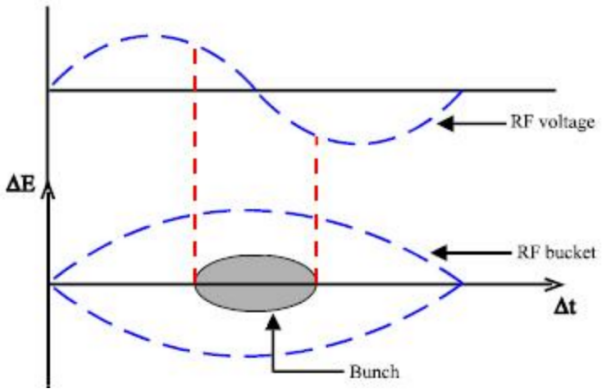

Figure 3.6: Illustration of the particle bunch structure in a particle collider such as the LHC. The particles are accelerated by radio-frequency (RF) oscillations whose amplitude is illustrated in the upper plot. The RF bucket's boundary, illustrated in the lower plot, is defined by a full period of the RF oscillation and the particle bunch formation, depicted in grey, occurs at the central node of the oscillation. The area occupied by the particle bunch is related to the beam's longitudinal *emittance*. Figure taken from Ref. [77].

3.1.3 The Concept of Luminosity

In designing a particle collider, the collision energy is not the only important parameter. Equally important is the value of the instantaneous *luminosity* that can be achieved by the collider. An expression for the instantaneous luminosity, $\mathcal{L}$, is given by,

$$\mathcal{L} = \frac{N^2 n_b f}{4\pi\sigma_x\sigma_y} \cdot S, \tag{3.4}$$

where N is the number of particles per bunch, n_b is the number of colliding bunches, f is the bunch revolution frequency, $\sigma_{x,y}$ are the transverse beam widths in the Gaussian approximation, and S is a reduction factor that accounts for geometric factors such as the non-zero crossing-angle of the colliding beams [74, 78]. The instantaneous luminosity, $\mathcal{L}$, can be seen by Equation 3.4 to have units of $\mathrm{cm}^{-2}\mathrm{s}^{-1}$ and can be conceptually thought of as the outgoing flux of particles per unit area and time after a bunch crossing in which successful pp collisions occur. The LHC is designed to deliver collisions to the high luminosity IP (Figure 3.4) at $\mathcal{L} = 10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$. Accurate knowledge of $\mathcal{L}$ is of the utmost importance for collider design and operation. Not only does it parametrise the potential collision rate once the collider beam and bunch structure are decided, but it allows for the accurate prediction of the number of collision events, N_{proc}, associated with a particular physics process with cross-section σ_{proc},

$$N_{\mathrm{proc}} = \sigma_{\mathrm{proc}} \int \mathcal{L}\,\mathrm{d}t \equiv \sigma_{\mathrm{proc}} \cdot L, \tag{3.5}$$

where L is referred to as the *integrated luminosity* and has units of cm^{-2}. A common unit for integrated luminosity is the *barn*, with symbol 'b': one barn is defined as $10^{-24}\,\text{cm}^{-2}$. The datasets collected so far by the LHC experiments are such that the *femtobarn* (fb), $10^{-39}\,\text{cm}^{-2}$, is relevant.

3.1.4 Operation of the Large Hadron Collider

The LHC has been in stable operation since 2009. It operates in so-called *runs*: multi-year periods of roughly continuous data-taking. As CERN shuts down during the winter months, each run is segmented each year with a several month long shutdown in the winter with a ramp-up period in the spring. During these shorter shutdowns, maintenance and upgrades may take place. In between a given run there is a multi-year break, a *long shutdown*, in which large(er)-scale maintenance and upgrades of both the LHC and the experiments can take place. At the time of writing, there has so far been two runs of the LHC, Run 1 and Run 2. Run 1 took place during the years 2009–2012 and Run 2 during 2015–2018. The integrated luminosities for each of the data taking years between Run 1 and Run 2 is shown in Figure 3.7. The scheduled LHC timeline is illustrated in Figure 3.8, inclusive of the High-Luminosity LHC period after which the LHC magnet system will have been fully upgraded.

The data relevant to the work presented in this thesis were collected in both Run 1 and Run 2 of the LHC, specifically that data collected in the years 2012–2018. The luminosities, instantaneous and integrated, as well as the center-of-mass collision energies, $\sqrt{s}$, for these data-taking periods are shown in Table 3.1. Also shown in Table 3.1 are the average values of the mean number of interactions per bunch crossing, $\langle\mu\rangle$, observed during each data-taking year. The quantity $\langle\mu\rangle$ is related to the amount of *pileup* observed during data-taking. Pileup is caused by overlapping pp interactions within the same (*in-time* pileup) or neighboring (*out-of-time* pileup) bunch-crossing(s) at the interaction point. The pileup scales with the instantaneous luminosity. Distributions of $\langle\mu\rangle$ are shown in Figure 3.7 for the Run 2 data-taking period.

	Run 1	**Run 2**			
Year	2012	2015	2016	2017	2018
Collision energy, $\sqrt{s}$ [TeV]	8	13			
Peak Luminosity, $\mathcal{L}$ [$\text{cm}^{-2}\text{s}^{-1}$] ($\times 10^{34}$)	0.77	0.5	1.4	2.1	2.1
Integrated Luminosity, L [fb^{-1}]	20.2	3.2	33.0	44.3	59.9
Mean number of of interactions per bunch crossing, $\langle\mu\rangle$	20.7	13.4	25.1	37.8	36.1

Table 3.1: Summary parameters for the data-taking periods relevant to the work presented in this thesis. The integrated luminosity is that relevant to performing physics analysis and potentially differs with respect to the total integrated luminosity delivered to ATLAS by the LHC (Figure 3.7) due to the application of strict quality criteria on the data prior to its use in physics analyses.

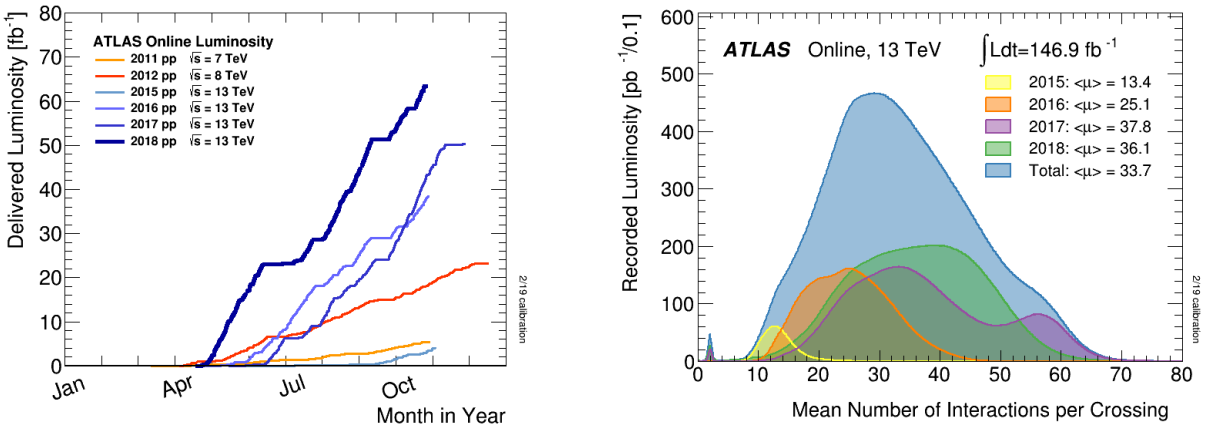

Figure 3.7: ***Left***: The ATLAS integrated luminosity during the data-taking years 2011–2018. ***Right***: The observed average number of *pp* interactions per bunch-crossing, $\langle\mu\rangle$, observed by ATLAS during the Run 2 data-taking years, 2015–2018.

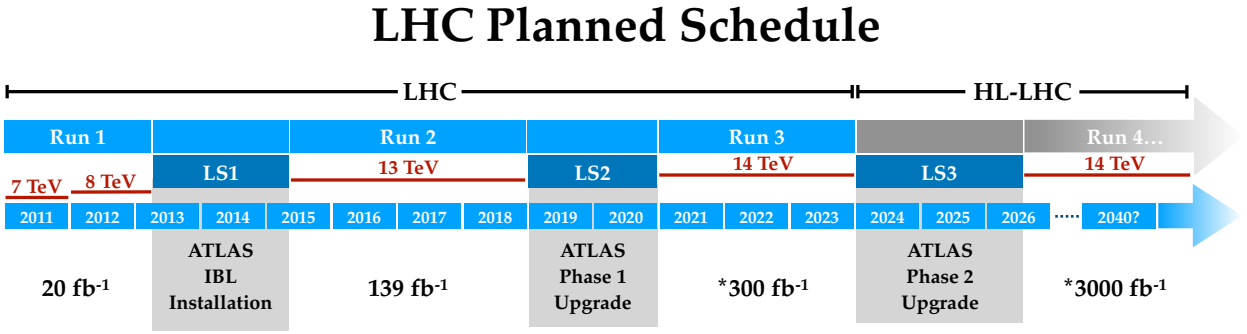

Figure 3.8: Schedule of the LHC operation. At the time of writing, the LHC is in the middle of 2019 during its second Long Shutdown (LS) period, LS2, after having delivered an integrated luminosity of 20 fb^{-1} during Run 1 and 139 fb^{-1} during Run 2 to the ATLAS detector. The predicted integrated luminosity estimates delivered to Point 1 and 5 for Run-3 and for the High-Luminosity LHC (HL-LHC) era are 300 fb^{-1} and 3000 fb^{-1} (after $\approx$ 10 years of running), respectively. The center-of-mass collision energies for each period are given in the red text.

3.2 The ATLAS Detector

In this section we will extend our focus to the ATLAS detector, the general purpose particle detector located at Point 1 of the LHC ring (see Figure 3.5). Roughly cylindrical in shape, coaxial with the beam-pipe, the ATLAS detector is 44 m long and 25 m tall. It is by far the largest such detector ever built and, generally, is the largest and most complex device ever constructed. Being general purpose in scope, the ATLAS detector is hermetic and has nearly 4π radians of solid angle coverage around the *pp* collision point. Such detectors are commonly designed to have various subsystems — *subdetectors* — which are designed for the identification of specific types of

particles and interactions. They tend to be layered about the interaction point and cylindrically symmetric since the *pp* interactions taking place within the detector have no preferred direction in the plane transverse to the direction in which the proton beams are travelling. A view of the ATLAS detector and its subdetectors is provided by Figure 3.10. In the following we will briefly describe each subsystem in turn, describing first the detectors located nearer to the *pp* collision and proceeding outwards.

3.2.1 The ATLAS Coordinate System

The ATLAS detector uses a right-handed coordinate system with the origin located at the geometric center of the detector. The x-axis points to the center of the LHC ring, the y-axis points upwards and away from the center of the Earth, and the z-axis is along the beam-pipe. The side associated with positive (negative) z is referred to as the 'A' ('C') side of the detector.[10] Due to its cylindrical symmetry, ATLAS also uses the cylindrical coordinates, (r, ϕ, z), with ϕ the azimuthal angle about the z-axis and having $\phi = 0$ along the positive x-axis. The spherical polar angle, θ, is defined with respect to the z-axis, having $\theta = 0$ parallel to the beam-pipe and $\theta = \pi/2$ in the xy-plane transverse to the beam-pipe. The pseudorapidity, η, is commonly used when describing systems of particles or locations within the detector and is defined as $\eta = -\ln[\tan(\theta/2)]$. The relationship between pseudorapidity and polar angle is illustrated in Figure 3.9. Large (small) values of η correspond to the *forward* (*central*) region of the detector. The rapidity, y, is related to η and is defined as $y = \frac{1}{2}\ln[(E + p_z)/(E - p_z)]$. The pseudorapidity of a particle traversing the detector is equal to its rapidity if the particle is massless or ultra-relativistic; otherwise, they are different. The comparison between a particle's pseudorapidity and rapidity is illustrated in Figure 3.9. The coordinates used to describe systems of particles are typically described by their four-momenta: (p_x, p_y, p_z) or, equivalently, (p_T, η, ϕ). A distance metric commonly used to describe the distance between two systems of particles in the detector is $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}$. The ΔR quantity using y instead of η is also sometimes used and will be indicated by ΔR_y.

3.2.2 The Inner Detector

The innermost subdetector of ATLAS is the Inner Detector (ID) [80]. The ID covers the region $|\eta| < 2.5$ and is composed, in order of increasing radial distance from the beam-pipe, of the silicon pixel detector, the silicon-strip semiconductor tracker (SCT), and the transition radiation tracker (TRT). These detectors enable the reconstruction of the tracks associated with the $\mathcal{O}(1000)$ charged particles emerging from each *pp* bunch collision occurring every 25 ns. An illustration of the ID

[10] 'A' for 'airport', since this is the side pointing towards Geneva International Airport, and 'C' for either 'Crozet' or 'Charly's', depending on who you ask, since this is the side pointing towards the town of Crozet and/or Charly's Pub in the town of Saint-Genis-Pouilly.

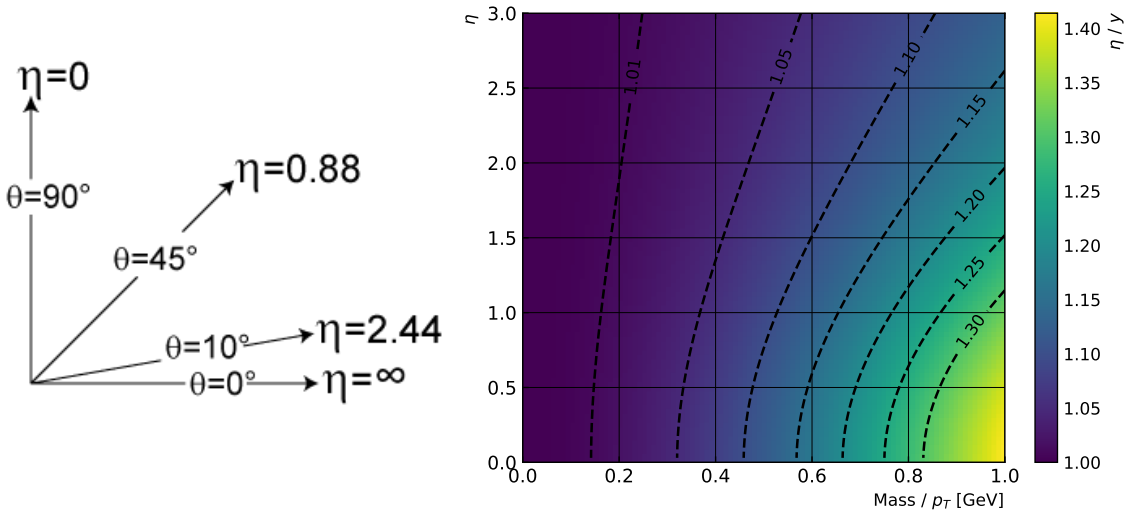

Figure 3.9: ***Left***: Illustration of the relationship between the pseudorapidity, η, and polar angle, θ, defined as the angle with respect to the beam-axis (z-axis). ***Right***: Distribution of the ratio of a particle's pseudorapidity to its rapidity, η/y, as a function of its pseudorapidity (y-axis) and the ratio of its mass to its transverse momentum, p_T (x-axis).

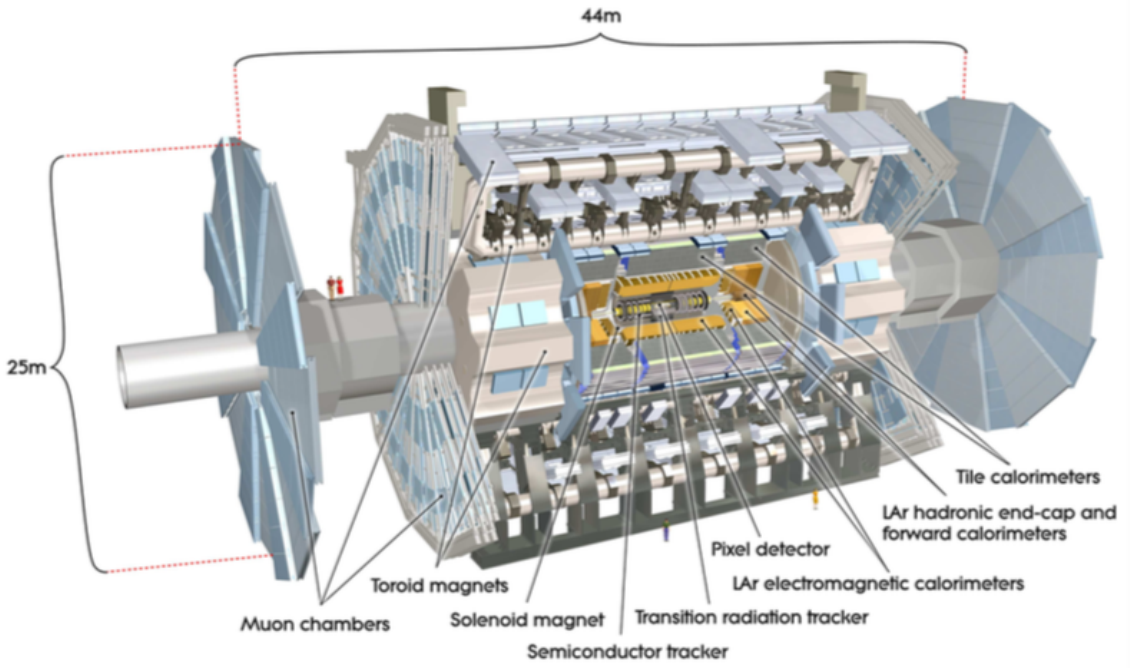

Figure 3.10: Cut-away view of the ATLAS detector with sub-systems indicated. Shown for comparison are figures of average-height humans standing at the feet of the detector and standing on the forward shielding between the big wheels of the forward muon system. Figure taken from Ref. [63].

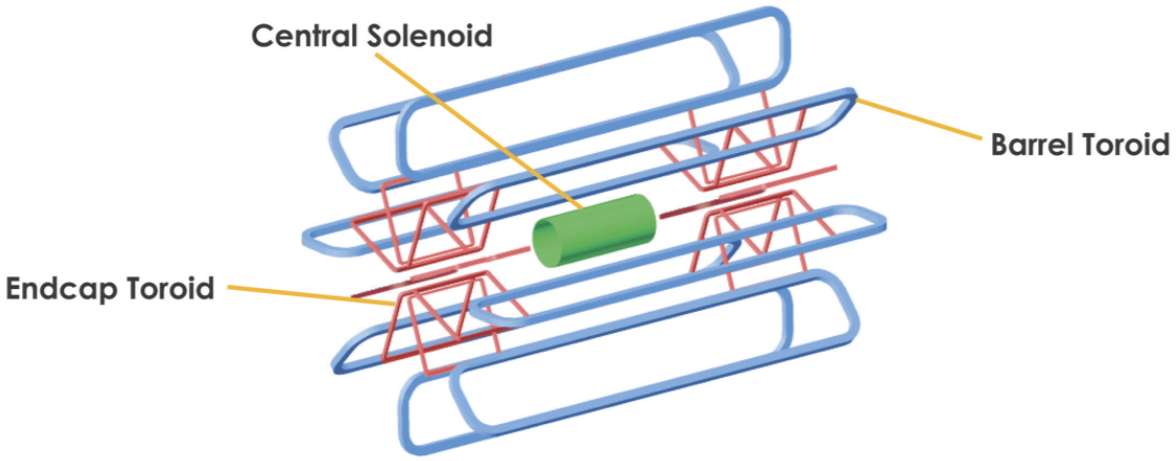

Figure 3.11: A view of the ATLAS magnet system. Shown are the 2 T solenoid magnet in green, the barrel toroid system in blue, and endcap toroid magnets in red. Figure taken from Ref. [79].

and its subdetectors is shown in Figure 3.12. Additional, more detailed views of the barrel and endcap sections of the ID are shown in Figure 3.13. Figure 3.14 shows further information about the physical locations in r, z, and η of each of the components that make up the ID. The ID is situated inside of the central solenoid, indicated in Figure 3.11, which provides an axial 2 T magnetic field and extends over a length of 5.3 m with a diameter of 2.5 m. The bending of charged particles in the xy-plane due to the presence of the solenoidal field allows for their momenta to be measured using the curvature of their reconstructed tracks.

The Pixel Detector and IBL

The pixel detector is the innermost subdetector of the ID, situated very near to and surrounding the beam-pipe. It is composed of three separate sections: a barrel section and two end-cap sections. The barrel section of the pixel detector has a cylindrical geometry and the end-cap sections are disks centered on the beam-pipe. The barrel section has four layers, each with increasing radius, and there are three disks in each of the end-caps. This ID geometry, shown in Figure 3.13, covers the region $|\eta| < 2.5$.

The pixel detector, being so near the pp collisions, is subject to the highest particle fluxes of any other subsystem. As a result, it is built to have very fine granularity: its sensing elements consist of 250 μm thick detectors housing pixels of reverse-biased n-type silicon semiconductor material, each having a nominal size of $50 \times 400\,\mu m^2$. In total, there are roughly 80 million channels read out from the pixel detector alone. This allows for the pixel detector's fine spatial hit resolution of 10 μm in $(r - \phi)$ and 115 μm along z.

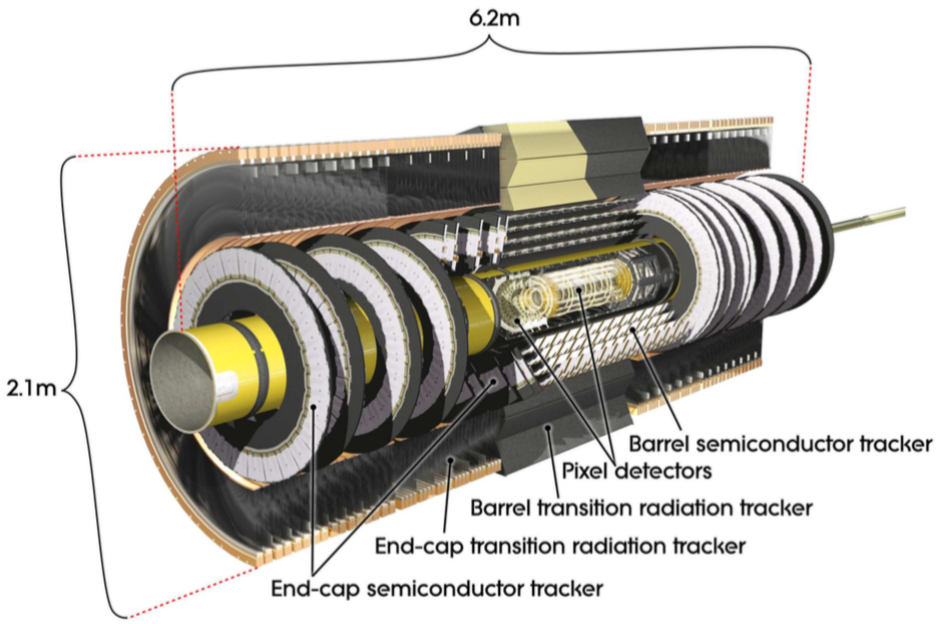

Figure 3.12: Cross-sectional view of the ATLAS inner detector. Shown are the barrel and end-cap portions of the pixel, SCT, and TRT detectors. Figure taken from Ref. [63].

The innermost layer of the pixel detector's barrel section is referred to as the *Insertable B-Layer* (IBL), and was installed at the beginning of the Run 2 data-taking period [81]. It corresponds, essentially, to the instrumentation of the ATLAS beam-pipe, as seen in Figure 3.15, and is located at a radial distance of 3.3 cm. It alone accounts for 8 million readout channels of the pixel detector — resulting in an ultra precise spatial hit resolution of 8 μm in $(r-\phi)$ and 40 μm along z. Beyond improving the overall measurements and reconstruction of charged particle tracks, the IBL was installed in order to improve the performance of secondary vertex reconstruction — an essential ingredient to the algorithms associated with the reconstruction and identification of jets originating from the decays of b-hadrons whose decays occur at radial distances frequently beyond that of the IBL, as illustrated in Figure 5.19..

The Semiconductor Tracker

The semiconductor tracker (SCT), like the pixel detector, uses silicon semiconductor-based sensing elements. It surrounds the pixel detector, as illustrated in Figure 3.13, and has similar barrel and end-cap geometries. The barrel section of the SCT is composed of 4 cylindrical layers and the end-caps consist of 9 disks. The silicon sensing elements are in a strip-like geometry with 80 μm strip pitch. The strips in the barrel section run parallel to the beam-pipe and those in the end-caps

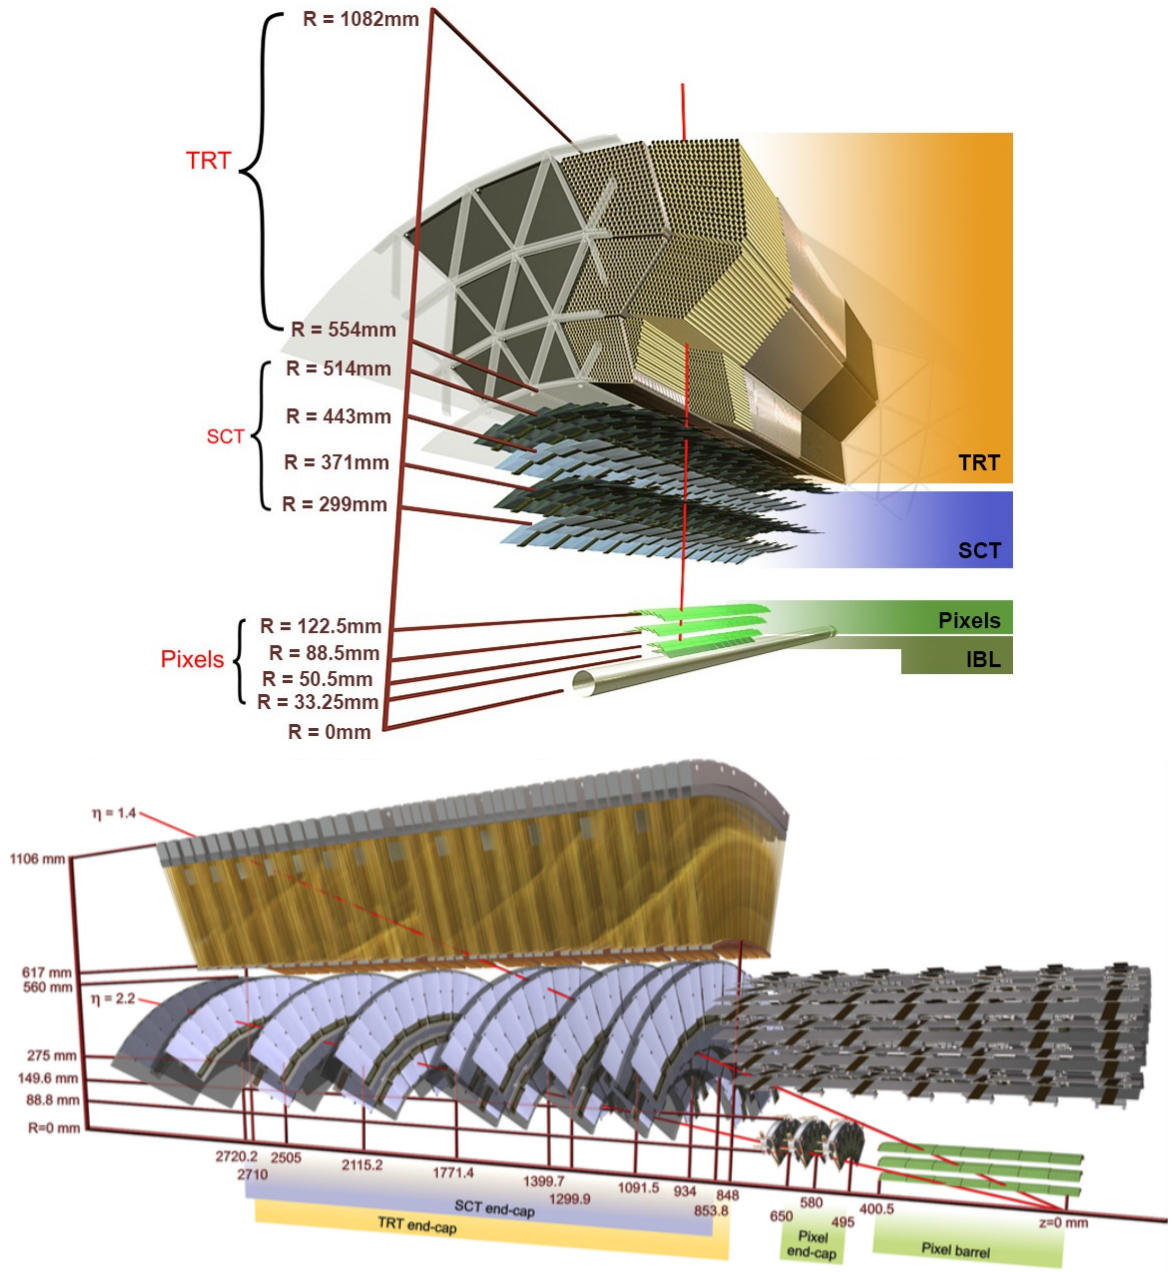

Figure 3.13: Cut-away views of the barrel (***top***) and end-cap (***bottom***) portions of the ATLAS inner detector, with each of the three subdetectors indicated along with their envelopes in r and in z. Figure taken from Ref. [63].

are perpendicular, extending along the radial direction.[11] The spatial hit resolution of the SCT is

[11] The SCT layers in both the barrel and end-cap sections additionally contain small-angle (40 mrad) stereo strips (c.f. Figure 6.8) to allow for measurement of both $(r - \phi)$ and z information.

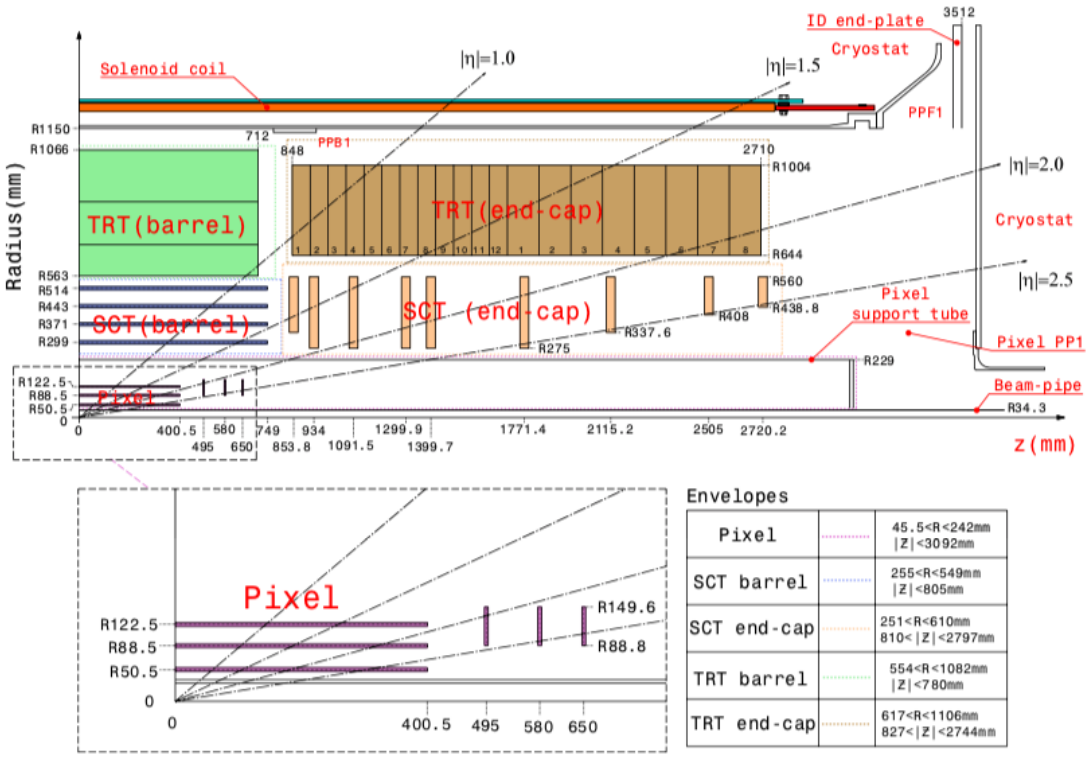

Figure 3.14: Detailed drawing of the locations of the subsystems that make up the ATLAS ID. Positions in r, z, and η are shown. Figure taken from Ref. [63].

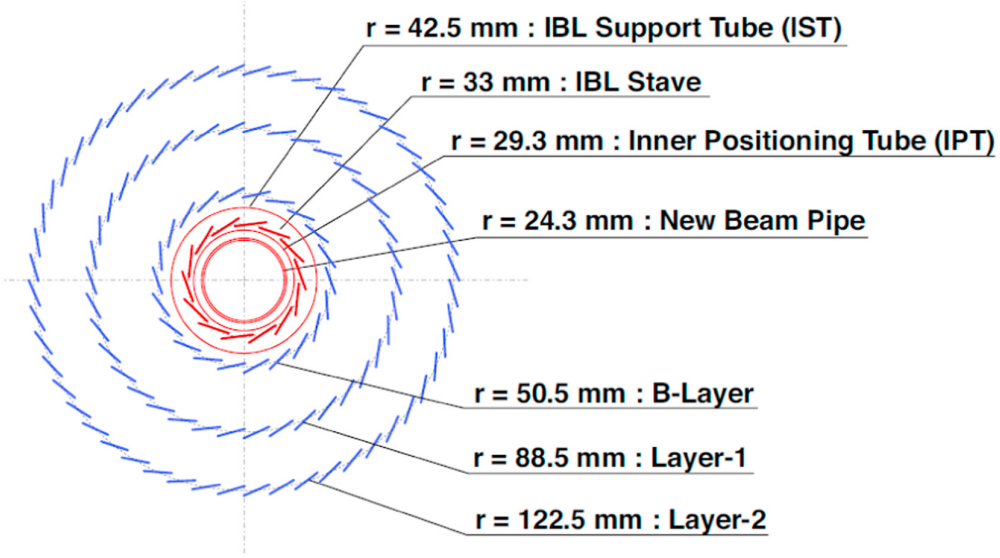

Figure 3.15: Transverse view of the barrel section of the pixel detector, showing the innermost layer, the Insertable B-Layer (IBL), and its support structure (red) as well as the three surrounding layers (blue). Figure taken from Ref. [82].

$17\,\mu m$ in $(r-\phi)$ and $580\,\mu m$ along z.

The Transition Radiation Tracker

The outermost layer of the ATLAS ID, surrounding the SCT, is the transition radiation tracker (TRT). The TRT is a tracking volume designed around the proportional drift tube concept and is composed of polyimide drift tubes (straws) that are 4 mm in diameter. The barrel section of the TRT contains up to 73 layers of 144 cm-long straws aligned parallel to the beam-pipe while the end-cap section has 160 straw planes composed of 37 cm long straws arranged radially into wheels (see Figure 3.13). Each straw of the TRT has a $31\,\mu m$ -width gold-coated wire at its center which acts as anode and is grounded, while the inner walls of each straw are kept at a potential of approximately -1.5 kV. A track hit in a given straw is the result of the 70% Xe – 27% CO_2 – 3% O_2 gas mixture contained in the straw volume being ionised and the resulting electrons (ions) drifting to the center wire (inner wall) of the straw. The induced current from the drifting charge is converted to an electrical signal and read out.

On average, a single charged-particle track leaves 36 hits in the TRT. The TRT only provides $(r-\phi)$ information (no z information), for which it has a per-straw hit resolution of $130\,\mu m$. The relatively poor hit resolution, when compared to the silicon based tracking detectors, is compensated by the large number of hits per track which lead to very long measured track lengths as compared to the pixel and SCT detectors. Additionally, the straws are embedded in and individually separated by a polypropylene fiber which induces transition-radiation photons to be produced. The amount and pattern of transition radiation depends on the mass of the passing particle: the passage of an electron will produce significantly more transition radiation than heavier charged particles, such as the copiously-produced pion. Information provided by the TRT therefore provides additional discrimination power between electrons and pions and enhances the performance of ATLAS' electron identification algorithms that primarily depend on information coming from the calorimeter systems (Section 3.2.3).

3.2.3 Calorimeter Systems

The ATLAS calorimeter systems are situated outside of the ID and central solenoid and are tasked with the measurement and containment of showers from electrically charged and neutral particles. A view of the calorimeter systems is provided by Figure 3.16. Broadly speaking, there are two types of calorimeters based on their purpose: electromagnetic and hadronic calorimeters. The electromagnetic calorimeter system has η coverage that matches the inner-detector and is optimized for precision measurements of electrons and photons. The hadronic calorimeter system has readout cells that are generally of coarser granularity as compared to the electromagnetic calorimeter and is

designed to meet the requirements for jet and missing transverse momentum measurements. Besides classification by physics purpose, the calorimeter system can also be broken into two classes based on detector technology: either based on gaps of cooled liquid-argon [83] or on scintillating tiles as the active media [84].

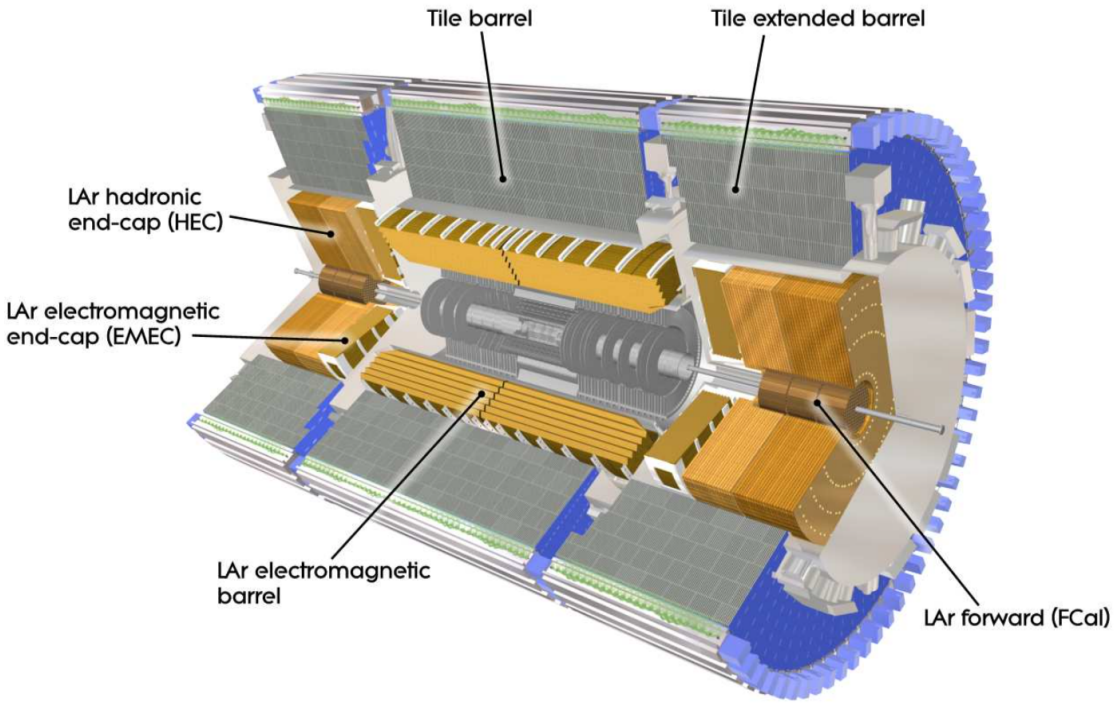

Figure 3.16: Cut-away view of the ATLAS calorimeter system, with liquid-argon and scintillating-tile subsystems indicated. Figure taken from Ref. [63].

Electromagnetic Calorimeter

The electromagnetic (EM) calorimeter is a high-granularity lead/liquid-argon (LAr) sampling calorimeter situated outside of the ID and sharing the same cryostat as the central solenoid. It consists of barrel and end-cap sections that cover the entire range within $|\eta| < 3.2$ and is illustrated in Figure 3.16. The structures of the electromagnetic barrel and end-cap calorimeters are shown in Figure 3.17. The EM calorimeter is designed in an accordion type structure to provide full coverage in ϕ. The cooled LAr fills the gaps between layers of the accordion structure. Passing particles from the interaction point undergo scattering and bremsstrahlung processes as they pass through the lead absorbers. The resulting electrons and photons ionise the LAr, producing drift electrons and ions whose signals are read out by the interleaved readout electrodes. The 2.1 mm drift gap has an operating voltage of $\approx 2\,\text{kV}$. The electromagnetic calorimeter is > 22 radiation lengths (X_0), ensuring that the majority of electrons and photons are completely contained within the

EM calorimeter. The majority of the EM energy, amounting to approximately 16 X_0, is contained within the second sampling layer (see Figure 3.17). The fine granularity of the readout, indicated in Figure 3.17, was designed with the ability to distinguish individual photons arising from $\pi^0 \to \gamma\gamma$ decays. The ability to distinguish photons pairs so precisely is also important for the main Higgs boson decay channel used for its discovery, $h \to \gamma\gamma$.

In the region $|\eta| < 1.8$, a so-called *presampler* detector is used to correct for the energy lost by electrons and photons due to material interactions occurring upstream of the EM calorimeter. It is a single LAr layer, with width 1.1 cm (0.5 cm) in the barrel (end-cap).

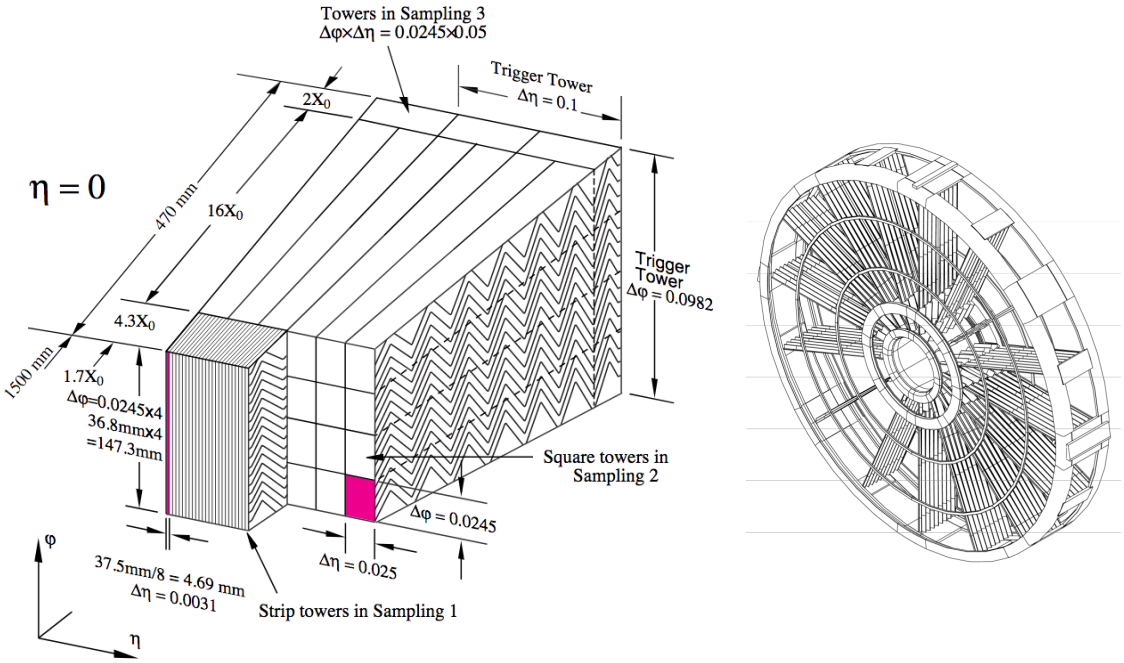

Figure 3.17: Figures taken from Ref. [83]. ***Left***: Cut-away view of the barrel electromagnetic calorimeter and its accordion structure. Indicated are the geometry and absorption properties of the three sampling layers. Also indicated is the granularity of the electrode readout in $\Delta\phi \times \Delta\eta$ in each layer. ***Right***: Diagram of the electromagnetic end-cap calorimeter accordion wheel structure (only a sub-set of the accordion structure is shown).

Hadronic Calorimeter

The barrel section of the hadronic calorimeter is composed of a lead/scintillating-tile type detector whereas the end-cap hadronic calorimeter is based on copper/LAr-based technology.

The lead/scintillating-tile calorimeter (the 'tile calorimeter') is located just beyond the EM calorimeter. It is composed of a barrel section, covering $|\eta| < 1.0$, and two extended barrels that cover

$0.8 < |\eta| < 1.7$ (see Figure 3.16). It is a sampling calorimeter using steel as the passive absorber and scintillating plastic tiles as the active media. The tile calorimeter is composed of modules in which the scintillating tiles are situated in $(r - \phi)$ within the steel absorbers, as shown in Figure 3.18. The detector is segmented radially into three layers and the readout of the scintillation light, using wavelength-shifting fibers that are fed into photomultiplier tubes (PMT) situated along the outer radii, is organized in a projective geometry, also illustrated in Figure 3.18. In the barrel (extended barrel) section, most of the hadronic energy is captured by the first (last) two layers which account for ≈ 5.5 (6) hadronic interaction lengths (λ) of the ≈ 7 in total.

The hadronic end-cap (HEC) calorimeter consists of two wheels per end-cap, situated behind the electromagnetic end-cap calorimeter, and provides calorimetric coverage in the range $1.5 < |\eta| < 3.2$. A view of the HEC can be seen in Figures 3.16 and 3.19. The HEC calorimeter is built from layers of copper plates interleaved with 8.5 mm LAr gaps which provide the active medium for this sampling calorimeter. The readout structure is obtained by dividing the gaps into separate drift zones for which there are dedicated readout electrodes. This readout structure is arranged in a projective geometry.

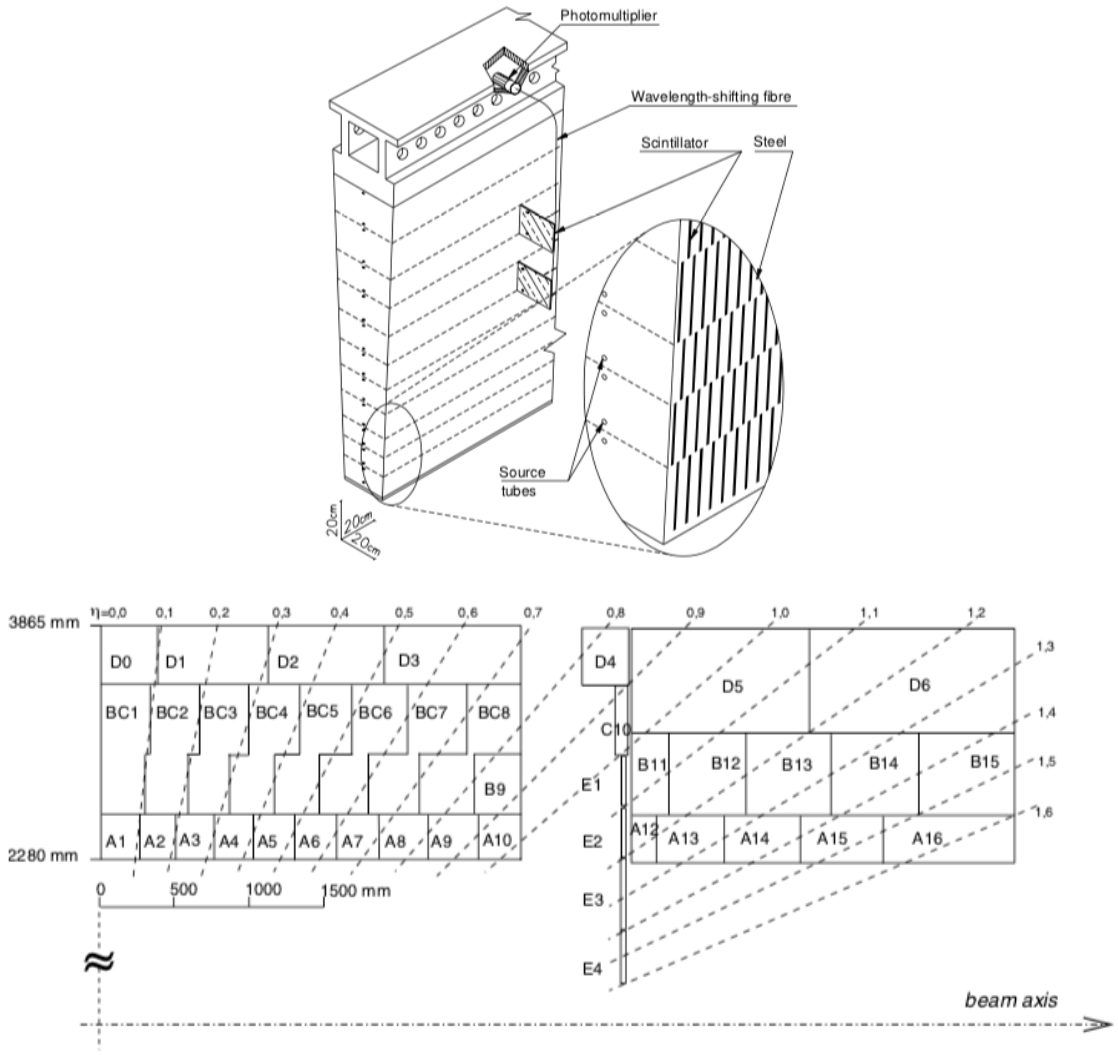

Figure 3.18: Figures taken from Ref. [84]. ***Top***: A view of a tile calorimeter module with its interleaved steel absorbers and scintillating tiles and PMT readout. Also indicated are the source tubes through which radioactive Cesium (Cs) sources are passed for calibration purposes [85]. ***Bottom***: Illustration of the segmentation of the projective readout of both the barrel and extended barrel tile calorimeter.

Forward Calorimeter

The forward calorimeter (FCal) system [86] provides calorimetric coverage to high $|\eta|$, between $3.1 < |\eta| < 4.9$, furthering the hermeticity of the detector. As indicated in Figure 3.19, FCal consists of three layers in the z direction: an electromagnetic layer (FCal 1) and two hadronic layers (FCal 2 and FCal 3). All three layers use LAr as the active medium but differ in their passive media. FCal 1 uses copper for its absorber, chosen for its heat removal properties, while FCal 2 and FCal 3 use tungsten, chosen to provide high containment and minimisation of the

lateral spread of hadronic showers. The FCal modules consist of matrices of the passive material with regularly spaced readout tubes oriented parallel to the beam-pipe that are filled with the cooled LAr.

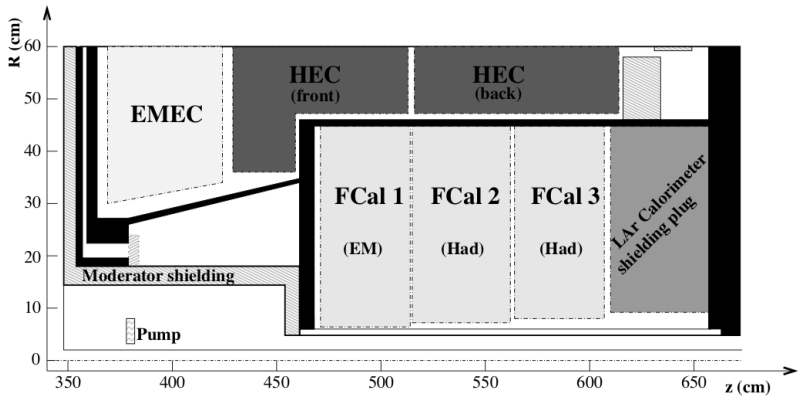

Figure 3.19: View of the forward calorimeter (FCal) system. Portions of the electromagnetic and hadronic end-cap systems are also shown. Figure taken from Ref. [86].

3.2.4 The Muon Spectrometer

Surrounding the calorimeters is the muon spectrometer (MS) [87], responsible for the detection of high-momentum, minimum-ionizing muons originating from the *pp* interaction. The MS is based on the magnetic deflection of muon tracks, allowing for their momentum determination. The bending of the muon trajectories is provided by the large superconducting air-core toroid magnet system, illustrated in Figure 3.11, consisting of a large barrel toroid over the range $|\eta| < 1.4$ and end-cap toroid systems in the range $1.6 < |\eta| < 2.7$. The superconducting toroid magnet provides an average field of 4 T. The magnetic field bending strength is roughly constant in η, except in the region in which the transition between the barrel and end-cap toroids takes place ($1.4 < |\eta| < 1.6$). A view of the ATLAS detector is shown in Figure 3.20, where it can be seen that the volume enclosed by the MS takes up most of the available volume outside of the calorimeter systems in the underground experimental cavern at Point 1. It should be noted that the overall design of the superconducting toroid structure, dictated by the performance requirements of the MS, is what gives ATLAS its large size and essentially drove the original design of all subdetectors discussed in the previous sections.

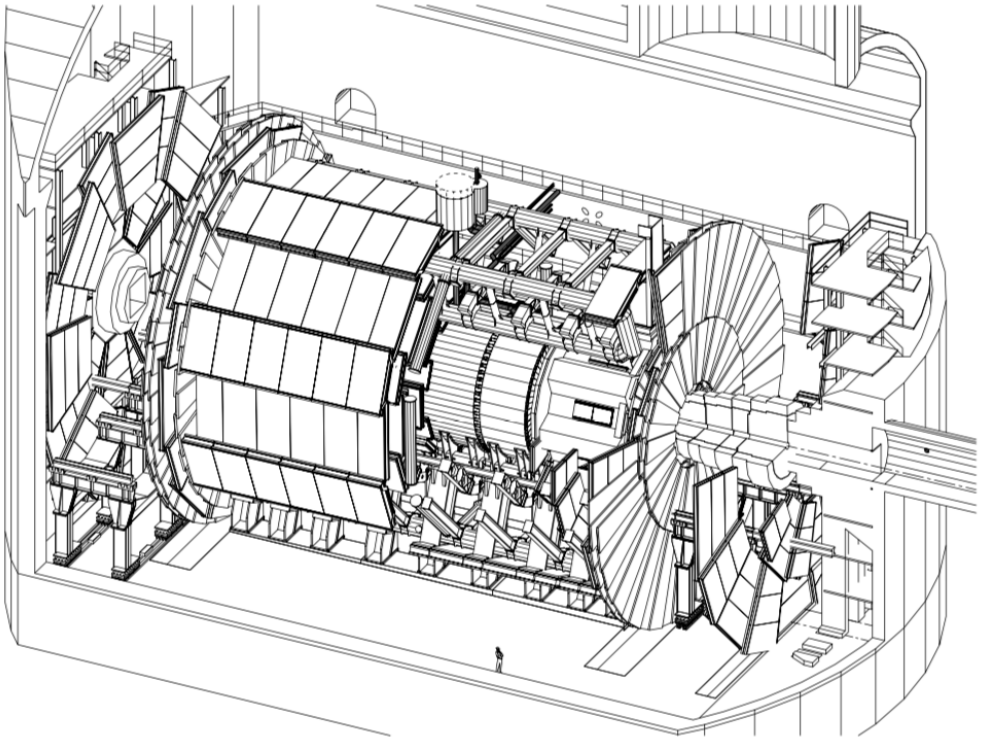

Figure 3.20: A view of the ATLAS detector inside the underground experimental area UX15. The cut-away view exposes the toroid structure and calorimeter system. Notice that the outermost muon stations in the forward regions are located at the extreme ends of the cavern. Figure taken from Ref. [87].

There are four types of gaseous radiation detector used in the MS, and their chamber layout is based on the concept of projective towers. The chambers follow the structure of the toroid magnet structure and have a 16-fold segmentation in azimuth, shown in Figure 3.21. They are arranged in large and small sectors, with the large sectors covering the regions between the coils of the toroid and the small sectors the azimuthal range in which the coils sit. The detector types can be broken into two classes and are either *precision* or *trigger* chambers. The precision chambers are composed of Monitored Drift Tube (MDT) [88] and Cathode Strip Chamber (CSC) [89] detectors and allow for the precise measurement the muon tracks as they traverse the MS, specifically the precise measurement in the bending plane of these tracks so as to allow for accurate determination of the muon momenta through their curvature. The trigger chambers are composed of Resistive Plate Chamber (RPC) [90] and Thin Gap Chamber (TGC) [91] detectors and have fast signal formation and readout times, allowing for accurate assignment of a passing muon to a specific pp bunch crossing. Both types of detectors exist in the barrel and end-cap sections of the MS and there are typically at least three layers of precision-type chambers over the entire $|\eta|$ range of the MS in order to allow for the sagitta measurement of the muon tracks necessary for momentum determination. The number of precision chamber hits over the entire range in $\eta - \phi$ of the MS is

shown on the left side of Figure 3.22. In the regions $|\eta| \sim 0$ and $|\eta| \sim 1.2$ there are noticeable drops in chamber coverage in order to allow for ID and calorimeter services and in the transition region between the barrel and end-cap, respectively, as seen on the right side of Figure 3.22.

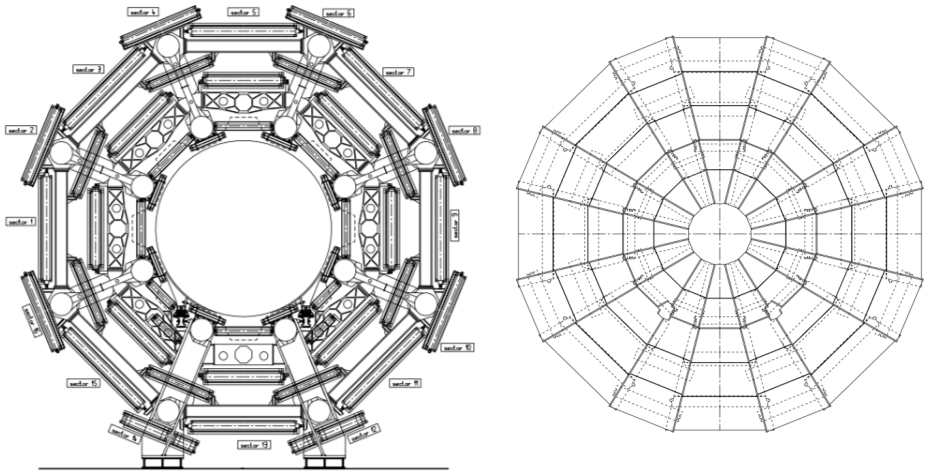

Figure 3.21: View of the 16-fold segmentation of the muon spectrometer in the barrel (***left***) and end-cap (***right***). Clearly seen in both is the arrangement of the detector chambers into large and small sectors, allowing for complete coverage in azimuth. The view of the end-cap is that only of the MDT chambers located at $z \approx 13\,\text{m}$. Figures taken from Ref. [87].

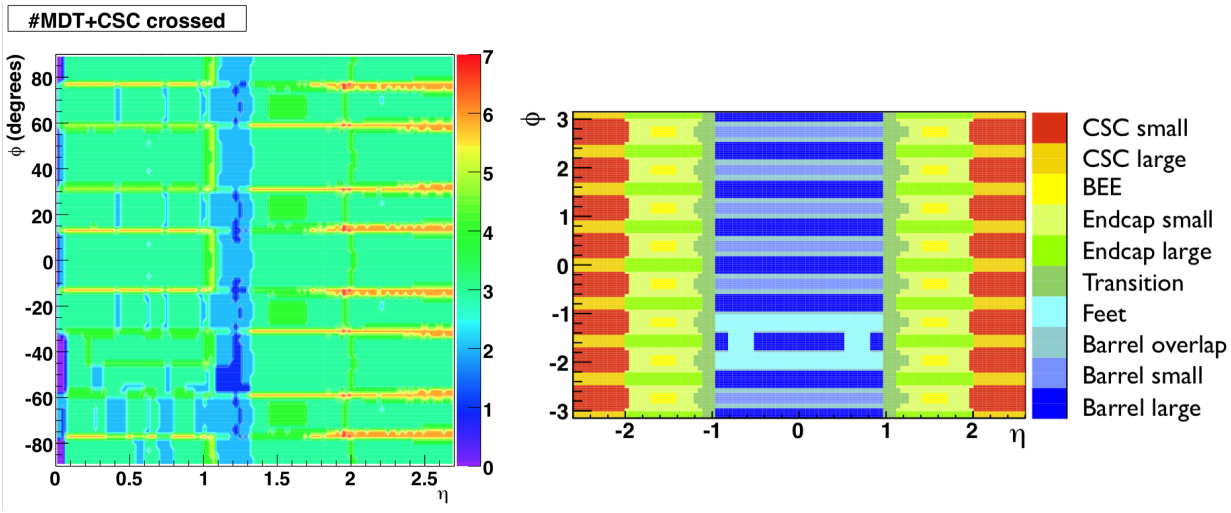

Figure 3.22: ***Left***: Number of precision muon chambers (MDT and CSC) traversed by a muon passing through the muon spectrometer as a function of η and ϕ. The regions of high numbers of crossings (> 4) correspond to the regions of overlap between the large and small sectors. Figure taken from Ref. [92]. ***Right***: Location in $\eta - \phi$ of several regions of the MS. Figure taken from Ref. [93].

The layout of the muon chambers and the corresponding detector technologies in the barrel and end-cap sections is shown in Figure 3.23. Here we will briefly describe each, starting with those in

the barrel section.

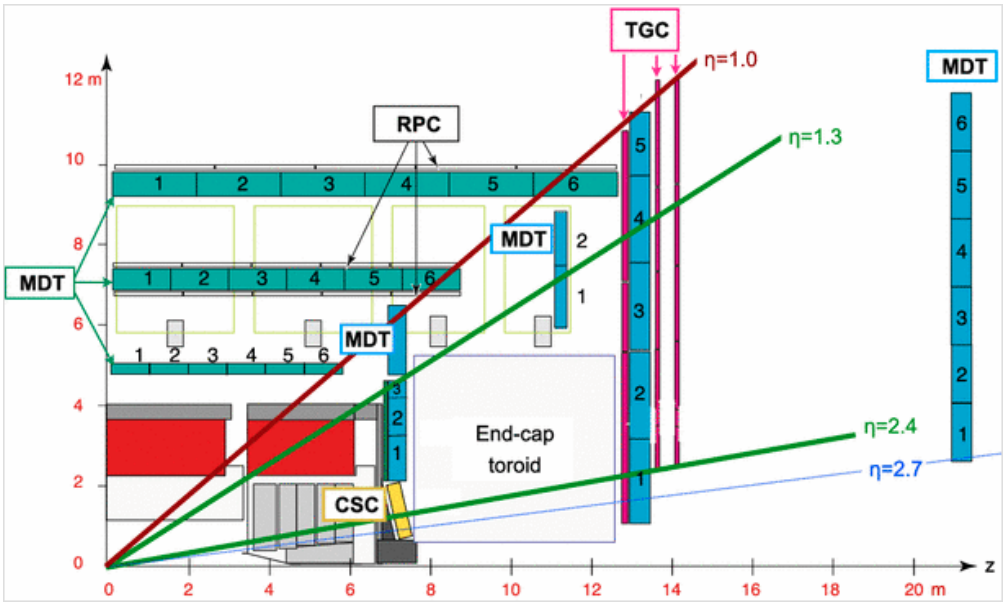

Figure 3.23: A view in the $r - z$ plane of a quadrant of the muon spectrometer (MS). Indicated by color are the four detector technologies used in the MS: MDT (blue), RPC (grey), TGC (red), and CSC (yellow). The light grey boxes at $6 < r < 9\,\text{m}$ indicate the location of the barrel toroid structures. Also shown are the envelopes in $|\eta|$ of the barrel, small wheel, and big wheel sections of the MS. Figure taken from Ref. [94].

Muon Spectrometer: Barrel

The muon chambers in the barrel section of the MS are rectangular in shape and arranged in 3 cylindrical shells, concentric about and parallel to the beam-axis at radial distances of 5, 7.5, and 10.5 m (see Figure 3.21). The precision chambers in the barrel section are composed of MDT chambers with tubes perpendicular to the beam-axis and parallel to the toroidal magnetic field, allowing for precision measurement along η. The MDT tubes are 3 cm in diameter and contain a 93% Ar – 7% CO_2 gas mixture with a single tungsten-rhenium wire operated at 3 kV. Traversal of a minimum ionising particle (MIP) ionises the gas within the tube, and the signal of the resulting ionisation charge is read out. The typical spatial resolution of a single MDT tube is below 100 μm. The MDT chambers are built as multi-layers of many MDT tubes which allows for the improvement of the spatial resolution down to 50 μm when the information from the individual layers is combined. An MDT double multi-layer chamber is shown in Figure 3.24. Also illustrated in this figure is the principle by which the tube hits in a given MDT multi-layer are used to form tracklets which aid in the process of muon track-building.

The chambers responsible for constructing muon trigger primitives in the MS barrel are the RPC

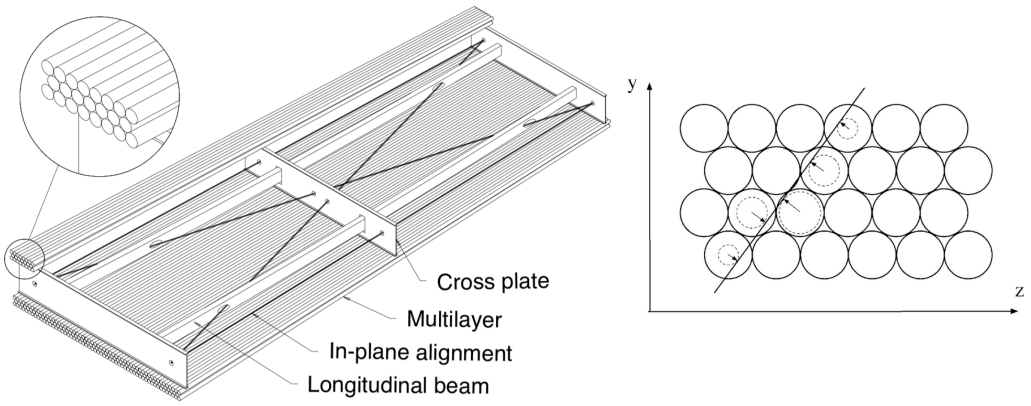

Figure 3.24: ***Left***: Illustration of a double-multilayer MDT chamber with its internal alignment and support structure exposed. A zoom-in on the multilayer of MDT tubes is shown. Figure taken from Ref. [87]. ***Right***: Illustration of the multilayer MDT tracklet-fitting algorithm [95].

chambers, whose principle of operation is shown on the left of Figure 3.26. The RPC gap is 2 mm, filled with tetrafluorethane ($C_2H_2F_4$), and is lined with parallel plate electrodes operated at a potential difference of 9.8 kV. This high operating potential and gas mixture allows for a timing resolution of 2 ns. Readout strips in x and y collect the induced charge from the ionisation events within the gap and provide additional spatial information for track and trigger-primitive building.

Muon Spectrometer: End-cap

The end-cap muon chambers, located in $1 < |\eta| < 2.7$, are arranged in 4 rings — *wheels* — extending radially and concentric with the beam axis at $z \approx 7.5, 10, 14, 22$ m from the pp interaction point. The wheel at $z \approx 7.5$ m, located on the IP-side of the end-cap toroid, is referred to as the 'Small Wheel' and those at $z > 10$ m are referred to as the 'Big Wheels'. That at $z \approx 10$ m, situated above the end-cap toroid, is an intermediate muon station composed of MDT chambers and has generally lower coverage than the Small and Big Wheels.

As in the barrel section, the primary precision measurement in the end-caps is provided by MDT chambers which are located in all four wheels of the end-cap. The MDT tubes are oriented azimuthally in order to obtain precision measurement in η. At the region $2 < |\eta| < 2.7$, in the innermost muon station in the end-cap that experiences the highest background rates, the precision muon measurement is provided by the CSC chambers at low radii. The CSC detectors are multi-wire proportional chambers, illustrated in Figure 3.25, with cathode strips perpendicular to anode wires and operated with $Ar/CO_2/CF_4$ gas mixtures. Passing MIPs result in ionisation events whose signals along the strips and wires are subsequently readout. As compared to the MDT chambers, the CSC detectors can resolve spatial information in both η and ϕ and, due to their relatively

high granularity readout structure, can sustain the higher background rates experienced in this very forward region of the detector. The CSC sectors are multi-layered (4-layers) and can achieve spatial hit resolutions on the order of 60 μm.

The trigger chambers in the end-cap are composed of the TGC detectors. Like the CSC, the TGC is a multi-wire proportional chamber with a gas mixture of CO_2 and n-pentane (n-C_5H_{12}). An illustration of the operating principles of a TGC detector is shown in Figure 3.26. The graphite cathodes and wires, with 1.4 mm separation, are held at a potential difference of 2.9 kV. This high potential difference and anode/cathode geometry allows for signals to be readout with a timing resolution of 4 ns. The signals from the drift electrons, collected along the wires, and the induced charge on the strips located behind the G-10 layer are read out and provide two-dimensional spatial information that can be used both in track and trigger-primitive building.

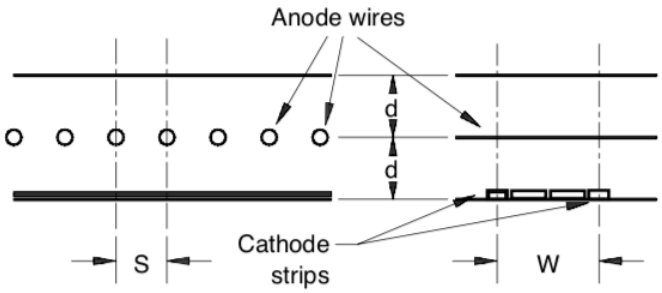

Figure 3.25: Diagram showing the main components of a cathode-strip chamber (CSC). On the ***left*** (***right***) is a view parallel (perpendicular) to the anode wires and perpendicular (parallel) to the cathode strips. Figure taken from Ref. [87].

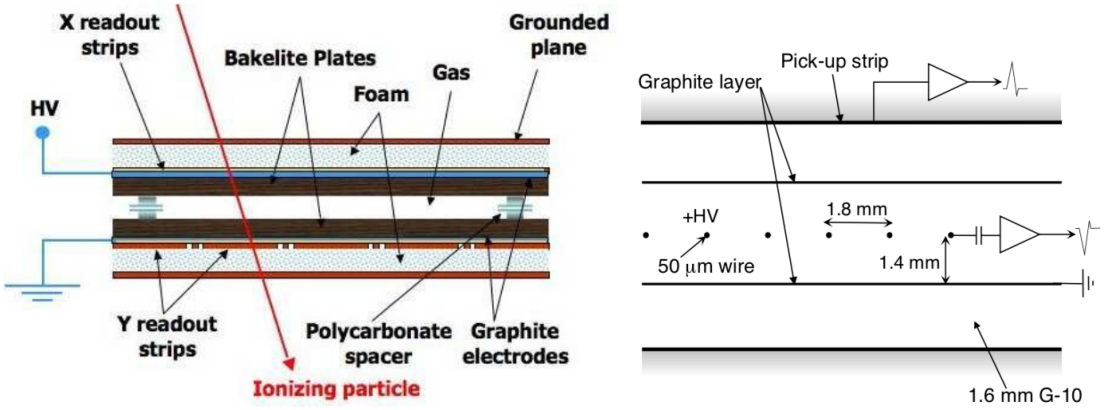

Figure 3.26: Muon trigger chambers. Figures taken from Ref. [87]. ***Left***: Illustration of a resistive plate chamber (RPC) and its principle of operation. ***Right***: Diagram showing the main components of a thin-gap chamber (TGC).

3.2.5 Trigger and Data Acquisition

During Run 2 operation between 2015–2018, the LHC delivered pp collisions to ATLAS at instantaneous luminosities of $10^{34}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, at a bunch spacing of 25 ns, giving 33.7 pp interactions per bunch crossing on average (see Figure 3.7). These values correspond to roughly 10^9 pp interactions per second. It is not possible for the ATLAS detector and data storage facilities to both respond to and record every one of these interactions. In fact, from a physics perspective it is not necessarily desirable to record every single interaction. The vast majority of such interactions arise from uninteresting, soft collision processes which are not likely to contain, for example, decays of Higgs bosons or of new particles not accounted for in the SM. For this reason, the ATLAS detector employs an *online*[12] selection strategy to select potentially interesting candidate events to be further processed and considered for permanent storage. This online selection strategy is referred to as the *trigger* system [96].

The ATLAS Run 2 trigger system consists of two levels: a hardware-based low-level trigger, referred to as the *Level-1* (L1) trigger, and a second level software-based high-level trigger (HLT) [97]. The L1 trigger uses relatively coarse-grained measurements from the calorimeters and MS. It performs the first level of selection, reducing the initial input 40 MHz rate of events by accepting events at a maximal rate of 100 kHz. The L1 trigger performs searches for coarse proxies of interesting physics objects: leptons, photons, and jets. It triggers on electrons and photons based on energy deposits in the EM calorimeter, limited to $|\eta| < 2.5$. The hadronic calorimeter provides jet candidates to the L1 trigger system via calorimeter 'towers' made up of trigger elements constructed by a sliding window algorithm. Each trigger element is constructed by calculating energy sums of calorimeter cells in $\eta - \phi$. Muon-based L1 triggers are based on coincidences of hits along the layers of the MS that form projective towers, or *roads*, consistent with high-p_T muons.

The candidate events selected by the L1 trigger system are forwarded to the HLT. The HLT system is composed of a Level 2 (L2) trigger and the event filter (EF). The L2 system is similar to the L1 trigger, but performs more refined measurements on the objects and regions of the detector that resulted in the initial L1 trigger's decision to accept the event. The EF is purely software based, using the ATLAS Athena reconstruction framework [98] to perform high level object reconstruction and identification using algorithms similar to those used in the offline environment and described in Chapter 5. The HLT accept rate is roughly 1 kHz. The accepted events are sent to CERN's permanent storage facilities and are made ready for the offline analysis. An overview of the trigger system is shown in Figure 3.27.

[12]The 'online' environment refers to that of the ATLAS detector during runtime. The 'offline' environment refers to anytime in which the data being inspected or analysed is not *at that time* being recorded by ATLAS but instead has already been stored to permanent storage and is readily accessible at any time.

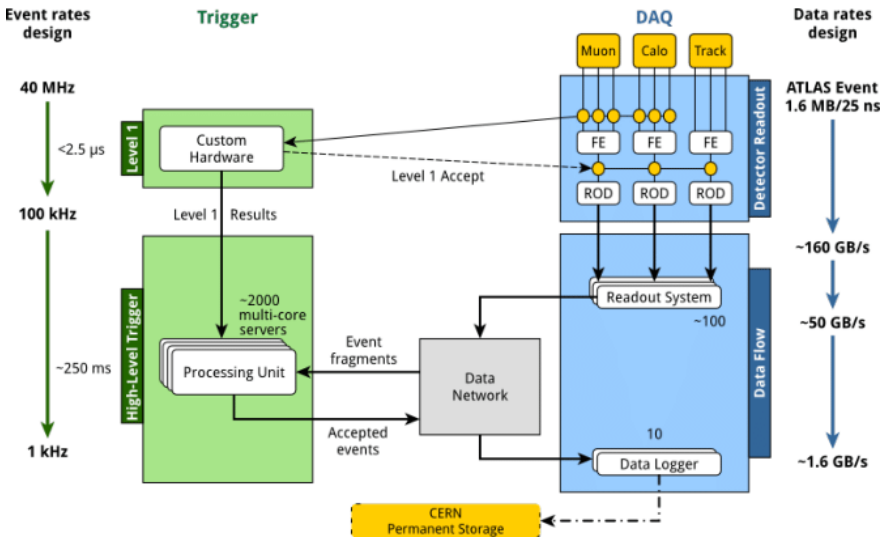

Figure 3.27: Overview of the ATLAS Run 2 trigger and data-acquisition architecture. Data from the muon and calorimeter systems are used for the Level 1 (L1) trigger, reducing the input event rate from 40 MHz to 100 kHz. The data accepted by the L1 trigger are forwarded to the readout drivers (RODs) [96] which, among other things, re-shuffle the raw data into the standardized ATLAS event format [99]. The events selected by the HLT at a rate of 1 kHz are pulled from the RODs and then forwarded to the permanent storage. Figure taken from Ref. [97].

Chapter 4

Simulation of pp Collision Events

> *[The Analytical Engine] might act upon other things besides number, were objects found whose fundamental relations could be expressed by those of the abstract science of operations, and which should be also susceptible of adaptations to the action of the operating notation and mechanism of the engine. Supposing, for instance, that the fundamental relations of pitched sounds in the science of harmony and of musical composition were susceptible of such expression and adaptations, the engine might compose elaborate and scientific pieces of music of any degree of complexity or extent.*
>
> –Ada Augusta, Countess of Lovelace

In order to be able to make predictions with which comparisons to the observed data recorded by the ATLAS data can be made, the need for an accurate simulation infrastructure arises. Simulation of the pp collision process must be made in order to estimate predicted rates of specific SM processes so that, in searches for new physics, one can construct reliable models of both the *background-only* and *background-plus-signal* hypotheses; the former including only those processes described by the SM and the latter including additional BSM physics processes on top of the SM ones. With both models in hand, tests of compatibility between the data recorded by the ATLAS detector can be made and quantitative statements about the likelihood of either of the two models can be made.

Simulation of the pp collisions in a general purpose high-energy particle-physics detector like ATLAS implies a precise knowledge of the underlying physics driving the pp interactions; namely, the theory of QCD. Generally speaking, physics processes driven primarily by electroweak physics (the decays of W and Z bosons, for example) are more well understood and the choices made in their simulation typically have a smaller impact on the analyses to be presented in the current thesis and for this reason the next sections give emphasis to the simulation methods particular to QCD physics. In Section 4.1 an introduction to the concepts underlying the simulation of QCD processes will be given. The simulation of the physics processes rely on detailed Monte-Carlo (MC) programs for

their implementation, a few of which that are relevant to the work presented in this thesis will be described in Section 4.2. Section 4.3 then describes how these MC-based simulations of the underling *pp* collision processes are put on the same footing as the recorded data via the accurate simulation of the ATLAS detector and its response to the *pp* collision products. Section 4.4 briefly describes how the simulation of pileup events is performed.

4.1 QCD Factorization and Fragmentation

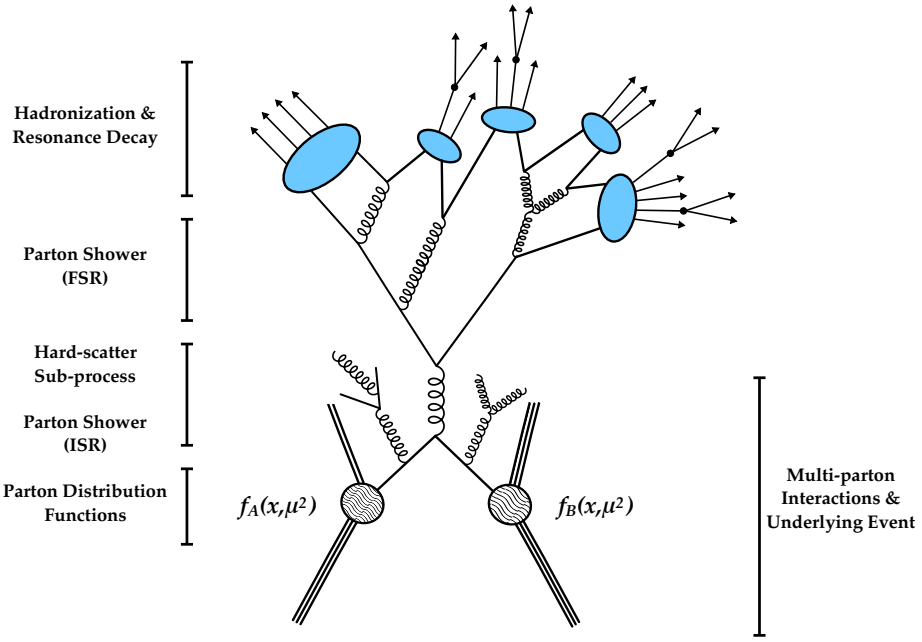

Figure 4.1: Cartoon illustration of the Monte-Carlo *pp* event simulation process. The observable color-neutral hadrons at the end result of hardonization are shown at the top, indicated by arrows. Example resonance decays could be $K_S \to \pi^0\pi^0$, $\Lambda^+ \to \pi^+ n$, or a B-hadron decay chain. The gluons emitted from the incident partons constitute the starting point of the initial state radiation (ISR). In order to avoid clutter, the hadronization of the ISR-emitted gluons and quarks is omitted. The blue ovals indicate the location of hadronization, simulated by phenomenological models as discussed in the text.

Before simulation of the ATLAS detector response can take place, accurate calculation and simulation of the production of final state particles as a result of a *pp* collision must occur. The steps of such simulation are illustrated in Figure 4.1. The breakdown of steps illustrated in Figure 4.1 represent the main characterising features of QCD, namely:

- Asymptotic freedom (confinement)

- Renormalization group equations (running of scales / scale dependence)
- Factorization (separation of high- and low-energy (perturbative and non-perturbative) regimes)

The asymptotic freedom and confinement of QCD state that the strength of the strong coupling, α_s, becomes weak at small distances (high energy scales, Q^2) and strong at a large distance. It is asymptotic freedom that enables the use of perturbative QCD (pQCD) to calculate the hard-scatter (large Q^2) reactions in QCD. Additionally, asymptotic freedom makes the color force dependent on the choice of scale, Q^2, making it necessary to perform computations with respect to a reference scale. This forces the introduction of arbitrary scales, μ, that are the reference scales enabling the pQCD computations to proceed but on which physical observables should not depend. This last requirement leads to the renormalization group equations that relate pQCD expressions at a given scale, μ, to physical observables.

The factorization implied by the above allows us to separate hard-scatter processes from the soft (low energy) parts of QCD that cannot be calculated perturbatively. These nonperturbative regimes occur at the characteristic scales of confinement onset, at $Q \sim \Lambda$ ($\sim 200\,\text{MeV}$). At scales above $m_{\text{hadron}} \sim 1\,\text{GeV}$, the perturbative treatment becomes viable.

$$\sigma_{pp\to X} = \sum_{A,B} \int_0^1 \mathrm{d}x_A \int_0^1 \mathrm{d}x_B \; f_A(x_A, \mu_F^2)\, f_B(x_B, \mu_F^2) \times \hat{\sigma}_{AB\to X}(x_A p_A, x_B p_B, \mu_F^2, \mu_R^2) \tag{4.1}$$

This factorization of QCD implies that calculations of pp collision processes take the form of Equation 4.1, where the sum runs over the partons of types A and B that exist in the incident protons and that contribute to the hard-scatter process. The PDFs, $f_i(x_i, \mu_F^2)$, absorb the non-perturbative aspects of the computation and parameterise the probability to find a given parton in the proton carrying a fraction of the proton's momentum, x_i, evaluated at a given energy scale $Q^2 = \mu_F^2$. As the PDFs describe the non-perturbative initial conditions of the partons initiating the hard-scatter processes, they are incalculable and must be obtained from subsidiary measurements or from altogether different experiments that are dedicated to their measurement [100, 101]. The PDFs are measured at specific scales Q^2 and are then extrapolated to the energy regime relevant to the physics process being calculated. This extrapolation is governed by the Dokshitser-Gribov-Lipatov-Altarelli-Parisi (DGLAP) evolution equations which evolve the PDFs from one scale to another [102, 103, 104]. The parton-level cross-section, $\hat{\sigma}_{AB\to X}$, describes the hard-scatter between the initiating partons and can be calculated using the machinery of perturbation theory. Both the PDFs and parton-level cross-section are dependent on the *factorization scale*, μ_F, which defines the cut-off scale below which phenomena are absorbed in the PDFs and above which phenomena contribute to the hard-scatter. As with all calculations involving the perturbation theory of QFT, the parton-level cross-section also depends on the renormalization scale, μ_R, defining the ultra-violet cut-off scale. Examples of proton PDFs, measured using data inclusive of 14 TeV data from the

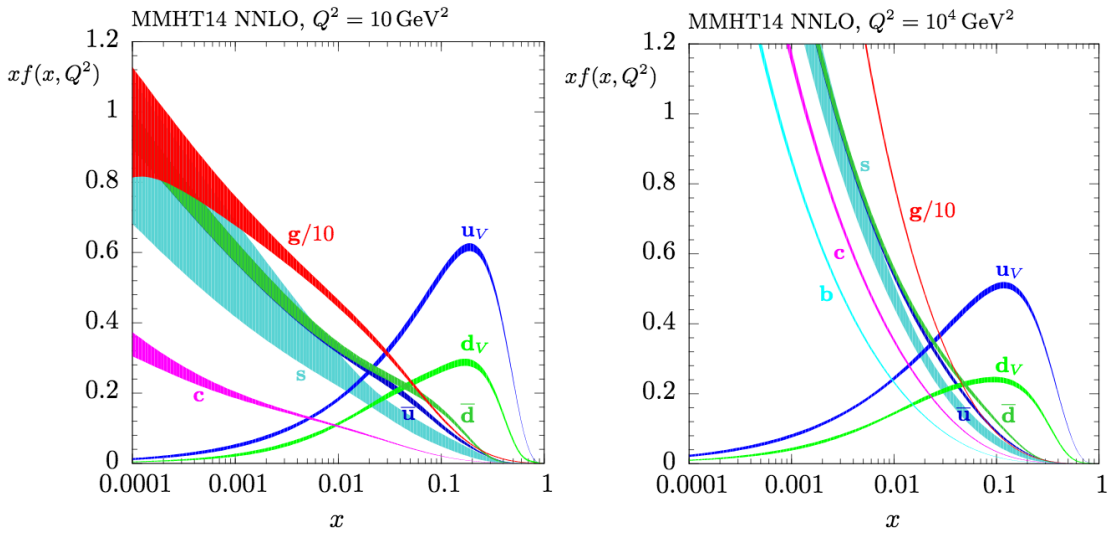

Figure 4.2: Proton PDFs evaluated at energy scales $Q^2 = 10\,\mathrm{GeV}^2$ (***left***) and $Q^2 = 10^4\,\mathrm{GeV}^2$ (***right***). The bands indicate the 68% confidence-level uncertainty bands. The valence-quark components are indicated with a subscript 'v' and the sea-quark components without. Figures taken from Ref. [105].

LHC, are provided in Figure 4.2.

The remaining elements of the pp event simulation described by Figure 4.1 are the parton shower (PS) and hadronization steps, the latter being non-perturbative and treated with phenomenological models. The PS simulates the successive emission of quarks and gluons from the partons in the final (initial) state of the simulated process. In the collinear factorization limit, the $n+1$-parton (post-emission) cross-section is related to the n-parton (pre-emission) cross-section as:

$$\mathrm{d}\sigma_{n+1} \approx \mathrm{d}\sigma_n \mathrm{d}P_a(z,q^2) \approx \mathrm{d}\sigma_n \frac{\alpha_s}{2\pi}\frac{\mathrm{d}q^2}{q^2}\mathrm{d}z\,\mathcal{P}_{ab}(z), \tag{4.2}$$

where $\mathrm{d}P_a(z,q^2)$ is the probability that parton a will split into two partons at scale q^2, with the parton b carrying the fraction z of the momentum of the initial parton a. The $\mathcal{P}_{ab}(z)$ are Altarelli-Parisi splitting functions [102] describing the possible parton branchings: $q \to gq$, $g \to gg$, or $g \to q\bar{q}$. The PS evolution follows by repeated implmentation of Equation 4.2, leading to arbitrarily many parton splittings and therefore potentially arbitrarily many particles in the final state. The basic idea, then, is to evolve the partons as above until they are below a certain energy scale, $q^2 = Q_0^2$, which is the scale below which the hadronization process takes place. In typical PS MC

programs, the PS evolution is performed via the construction of the Sudakov form factor,

$$\Delta_a(q_1^2, q_2^2) = \exp\left(- \sum_{b \in \{q,g\}} \int_{q_2^2}^{q_1^2} \int_{z_{\min}}^{z_{\max}} \frac{\alpha_s}{2\pi} \frac{\mathrm{d}q^2}{q^2} \mathrm{d}z \, \mathcal{P}_{ab}(z) \right), \tag{4.3}$$

which represents the probability that a parton evolves from energy scale q_1 to the lower scale q_2 *without* splitting.

The PS simulation of the final-state radiation (FSR) operates by following a *forward evolution* whereby partons initially at scale Q^2 emit radiation at the scale q_2 determined by sampling Equation 4.3. This process is repeated and if, for any of the partons in the final state, the q_2^2 value is below $Q_0^2 \approx 1\,\mathrm{GeV}^2$, the shower development is terminated and hadronization takes place. For the simulation of initial-state radiation (ISR), a *backward evolution* occurs wherein the radiation is emitted by the initiating partons. In this case the final low-energy scale is that of the ancestor partons fragmenting from the PDFs, gaining energy to take part in the hard-scatter, and the initial energy scale is the high-energy scale of the hard-scatter. As the ISR shower evolution is sensitive to the PDFs, it is dependent on the factorization scale μ_F since the value of the PDF is scale-dependent.

As the partons reach the hadronization scale, the confining nature of QCD takes over and the non-perturbative regime of QCD is reached once again. At this stage, the color-neutral hadrons that are observable within the detector take form. This formation process is described by phenomenological models that capture the general features of QCD. The main hadronization models used in the MC simulation relevant to the current analysis are the Lund string model [106] and cluster hadronization model [107]. Unstable particles produced as a result of the hadronization process are also decayed at this point; in some cases relying on the use of dedicated programs such as EvtGen in the case of b-hadron decays [108].

The so-called *spectator partons*, indicated by the outgoing lines from the PDFs in Figure 4.1, are those partons not directly involved in the pp hard-scatter sub-process. They may undergo soft interactions, resulting in the underlying event (UE). The physics processes underlying the UE are driven by low-energy phenomena and must therefore be simulated using phenomenological modes, in much the same way as the hadronization simulation. These models are characterised by tuneable parameters which are optimised by dedicated measurements and comparisons to observed data in the experiments [109]. There is also processes of multiple-parton interactions (MPI) in which there are additional, relatively soft, parton-parton scattering processes in addition to the hard-scatter subprocess. The processes of UE and MPI result in long-distance QCD color connection effects that lead to additional radiation of quarks and gluons, also simulated using tuneable phenomenological models provided by specific MC generators [110, 111].

4.2 Monte-Carlo Event Generators

There are a handful of MC-based event generators dedicated for LHC physics and the implementation of the processes described in the previous section, illustrated in Figure 4.1 [112]. In the following we reference those relevant to the work to be presented in the current thesis.

4.2.1 Matrix Element Generators

Matrix element (ME) MC generators provide calculations of the matrix element corresponding to the underlying *pp* hard-scatter sub-processes. They are generally interfaced to separate PS MC generators that perform the subsequent fragmentation.

MadGraph5_AMC@NLO [113] is an MC event generator providing automated computation of matrix elements at LO, with NLO computations aided by the MC@NLO method [114].

Powheg-Box [115] is a library for implementing NLO calculations in MC programs using the Powheg formalism [116] to match the matrix element computation to that of the PS.

4.2.2 General Purpose Event Generators

General purpose MC event generators are capable of performing matrix element calculations as well as handling the fragmentation and UE modelling.

Pythia [117, 118] is a standalone event generator capable of simulating over 200 processes, specifically $2 \to 1$, $2 \to 2$, and (in a few cases) $2 \to 3$ processes. In the present thesis Pythia is used primarily as a PS/UE generator, with separate MC generators performing the matrix element computation.

Herwig [119, 120] is a general purpose event generator emphasising the detailed simulation of QCD parton showers. In the work to be presented in this thesis, Herwig is used only as a PS generator.

Sherpa [121, 122] is a general purpose MC generator capable of both LO and NLO computations and contains a detailed PS generator based on the Catani-Seymour dipole formalism [123]. In the present analysis, when Sherpa is used for the matrix element computation it is also used for the PS.

4.2.3 Afterburners and Specialists

There exist several so-called MC-generator *afterburners* which are run after one of the steps detailed in Figure 4.1, typically after the matrix-element computation prior to running the PS algorithms or at the end, supplementing the hadronization step in order to improve the specific description of physics processes therein.

EvtGen [124] is an MC generator that is interfaced to and runs after the PS and hadronization step. It implements a detailed description of the physics of b-hadron decays, including models of their semi-leptonic decays, CP-violating effects, and accurate modelling of the angular distributions and correlations relevant to sequential decays of b-hadrons.

MadSpin [125] is a tool based on the formalism of Ref. [126] that handles the decays of heavy resonances with NLO matrix element computations provided by MadGraph5. It ensures that off-shell effects and spin-correlations amongst child particles (from the heavy resonance decay) are accounted for properly.

4.3 Simulation of the Detector Response

The end result of the steps discussed above and outlined in Figure 4.1 is a collection of four-vectors of all stable particles after hadronization. By itself, this collection of particle four-vectors is useful for studying physics processes at the so-called *truth*- or *particle*-level, allowing one to understand specific processes without the effects of the ATLAS detector's geometric acceptance and response folded in.

Ultimately, however, one wishes to use the MC simulation to make predictions on the collection of specific physics processes being simulated. To use the MC simulated events in this way, they need to be compared to the recorded data and therefore must be analyzed after the full ATLAS reconstruction process, described in Section 5, so as to ensure meaningful comparisons are made. To accomplish this, a detailed GEANT4 [127] model of the ATLAS detector is used which includes not only the sensing elements of each of the subdetectors but also includes simulation of inactive material such as support structures, cabling, and services [128]. Custom algorithms are developed for each of the ATLAS subdetectors which convert the simulated GEANT4 energy depositions into hits on the corresponding detectors, modelling the detailed nature of each subdetector's response to incident particles. For example, the simulation of the MDT drift tubes (Section 3.2.4) includes realistic simulation of signal formation as a result of the ion drift and avalanche evolution as well as a realistic simulation of the response of the readout electronics that produce the digitized output signals. The result of the simulation, then, is a set of simulated digitized output signals which may

then be treated by the same reconstruction algorithms used during actual data taking, including the trigger. The paths of the particle-level and fully reconstructed simulated *pp* collision events, as well as the paths of the data events, through the simulation and reconstruction infrastructure are illustrated in Figure 4.3. The reconstruction stage refers to the process of reconstructing the high-level physics objects and is described in detail in Chapter 5.

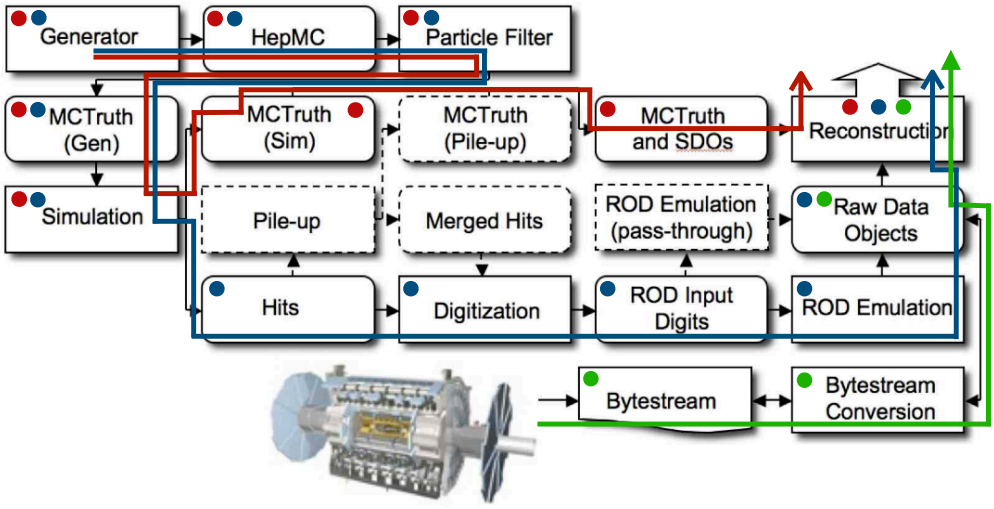

Figure 4.3: Paths of the ATLAS *pp* event simulation, starting from the MC event generation all the way to the full ATLAS reconstructed event (Chapter 5). The red path and markers indicate that of the particle-level (truth-level) events, while the blue path and markers indicate that of the fully reconstructed simulated events. The green path and markers show the path that real data takes when reconstructed and recorded by ATLAS. The acronym SDO (ROD) stands for 'Simulated Data Object' ('ReadOut Driver'). Original figure taken from Ref. [128].

4.4 Simulation of Pileup

As mentioned previously, there are two main types of pileup to consider: in-time pileup and out-of-time pileup. The former mainly consists of soft QCD interactions and is modelled similarly to the UE/MPI processes. The out-of-time pileup is modelled by separately generating a large sample of minimum-bias *pp* collision events using the Pythia MC event generator. For each generated hard-scatter process of interest, a given number of these minimum-bias events' detector hits are overlaid onto those from the simulated hard-scatter event prior to digitization and reconstruction. The number of overlaid minimum-bias events is based on a sampling of the pileup distribution for a given data-taking period, corresponding to the $\langle\mu\rangle$ profile shown in Figure 3.7. Overlaying the simulated events in this way, at the hit level and prior to digitization, allows for a precise simulation of the effects of pileup on the detector response. This process is illustrated in Figure 4.4.

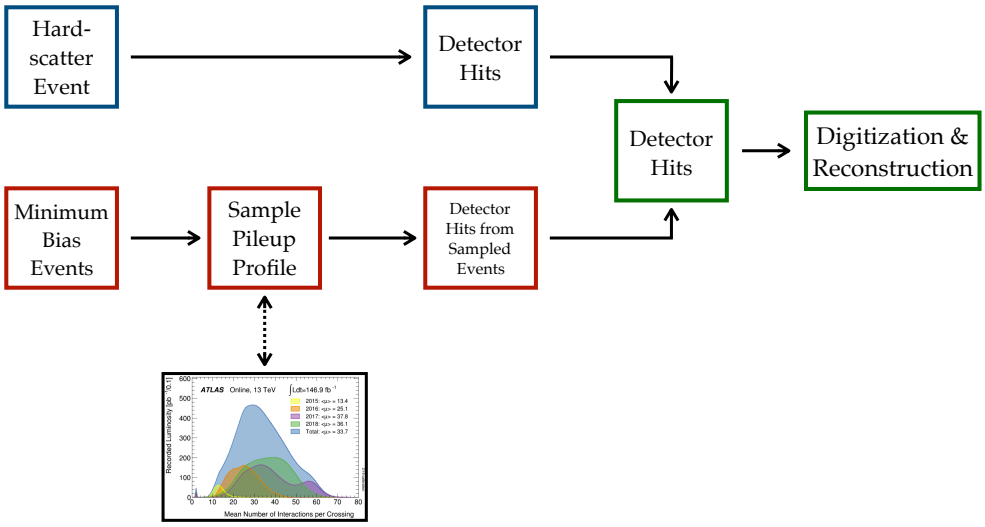

Figure 4.4: Cartoon illustration of the steps for performing the MC-MC pileup overlay used for simulating the out-of-time pileup in ATLAS. Events are sampled from a collection of simulated minimum-bias events generated using the Pythia MC event generator. For each generated hard-scatter process, the number of overlaid simulated minimum-bias events is determined by sampling the expected pileup profile for the given data-taking period (Figure 3.7). The simulated detector hits from both the main hard-scatter process and from the overlaid pileup events are combined prior to the remaining digitization and reconstruction steps of the ATLAS detector simulation chain (Figure 4.3).

Chapter 5

Physics Building Blocks and their Reconstruction

If you can put your five fingers through it it is a gate, if not a door.

–Stephen Dedalus, in James Joyce's *Ulysses*

In order to convert the multitude of electrical signals read out by the subdetectors of ATLAS as a result of a sucsessful trigger (c.f. Section 3.2.5) into well-defined and meaningful representations of the underlying physics process that initiated them, at the level required for performing high-quality physics analysis, several steps of reconstruction and identification must take place. The physics analyses presented in the current work involve the use of leptons, jets, and the so-called missing transverse momentum, $\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}$. The methods used to deduce the presence of these objects within the ATLAS detector will be discussed in this chapter. Section 5.1 introduces the reconstruction of charged-particle tracks and pp interaction vertices within the ID, both of which are used as low-level seeds or inputs to the reconstruction of the high-level physics objects to be discussed in the subsequent sections. Section 5.2 goes on to discuss the reconstruction of the charged leptons relevant to the current work: electrons and muons. Sections 5.3 and 5.3.4 describe the reconstruction of jet objects and the identification of jets arising from the decay of heavy-flavor hadrons, respectively. Section 5.4 then goes on to describe the reconstruction of $\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}$, which relies on an accurate description of leptons and jets. The methods used for reconstructing the leptons and jets are not one hundred percent accurate: detector information arising due to an electron may leave signatures similar to those of a jet, for example, and thus spoil their unambiguous description. Where relevant, in the following we will discuss the methods by which the reconstruction and identification of the physics objects is made more precise and how high levels of confidence about their actual presence within the detector are achieved.

5.1 Charged-Particle Tracks and Primary Vertices

The reconstruction of charged-particle tracks ('tracking') and primary interaction vertices ('vertexing') is based on information provided by the ID, primarily by the pixel and SCT subdetectors [129, 130, 131, 132, 133]. Charged-particles produced in pp collisions will leave signals — *hits* — on the different layers of the ID. The aim of tracking is to translate these layer hits into *spacepoints* which are then combined to form a track following the particle's traversal trough the ID. Given its highly granular readout, the pixel detector provides three dimensional spacepoints from each layer hit while the back-to-back readout strips on each layer of the SCT must be combined, using the stereo-angle information from the second set of strips, to give three dimensional spacepoint information. The hit information provided by the TRT straws is two-dimensional in nature, providing only $r-\phi$ information in the barrel section and $\phi - z$ information in the end-caps.

Within the solenoidal magnetic field of the ID, charged-particle tracks follow helical trajectories in the plane transverse to the beam-pipe (xy-plane) and can be fully characterised by five *track (perigee) parameters*:

$$(d_0, z_0, \phi, \theta, q/p)\,, \tag{5.1}$$

where d_0 (z_0) is the transverse (longitudinal) impact parameter, ϕ and θ are the azimuthal and polar coordinate, respectively, of the track at the point at which d_0 and z_0 are defined, q/p is the ratio of the particle charge to the magnitude of its momentum. The charge of a track is determined by its curvature within the magnetic field. The track parameters are defined with respect to their associated primary vertex, whose reconstruction will be described shortly. An illustration describing the track parameters is provided by Figure 5.1.

The primary track reconstruction algorithm used in ATLAS follows an *inside-out* pattern recognition procedure and first starts with information provided by track *seeds*, composed of a few spacepoints, in the silicon detectors which then are extended outwards into the TRT [129]. The inside-out approach accounts for the majority of tracks reconstructed in ATLAS but it is complemented by an *outside-in* approach that starts with the TRT hits and moves inwards [129]. This latter approach is useful in recovering those tracks with ambiguous or missing inner-layer pixel hits; for example, in the case of photon conversions or long-lived neutral particle decays.

The collection of reconstructed tracks is used as input to the primary vertex reconstruction. Primary vertex reconstruction follows a so-called *adaptive vertex fitting* (AVF) [131, 133] procedure and occurs in two steps: primary vertex finding, in which tracks are associated to a particular vertex candidate, and vertex fitting, which involves the reconstruction of the actual vertex position and its errors. After the vertex fitting stage, the tracks associated with a given vertex are refit with the constraint of the vertex position and its errors. The track refitting can update the track parameters

(Equation 5.1) associated with the tracks. Only vertices with at least two charged particle tracks with $p_T > 400\,\text{MeV}$ are considered.

In the high luminosity collisions at the LHC there will generally be multiple primary vertices associated with each pp bunch crossing. A physics *event* in ATLAS, then, is chosen as the set of processes originating from the pp interaction associated with the *hardest* primary vertex — the *primary hard-scatter vertex* — taken as that primary vertex with the highest sum of squared p_T of tracks originating from that vertex. The subsequent event reconstruction takes place around the primary hard-scatter vertex and only those objects originating from it are taken as relevant when reconstructing the physics objects in the event. Any additional primary vertices are considered as *pileup vertices*.

The presence of so-called *secondary*, *tertiary*, and so on..., vertices are also important and will be described in Section 5.3.4.

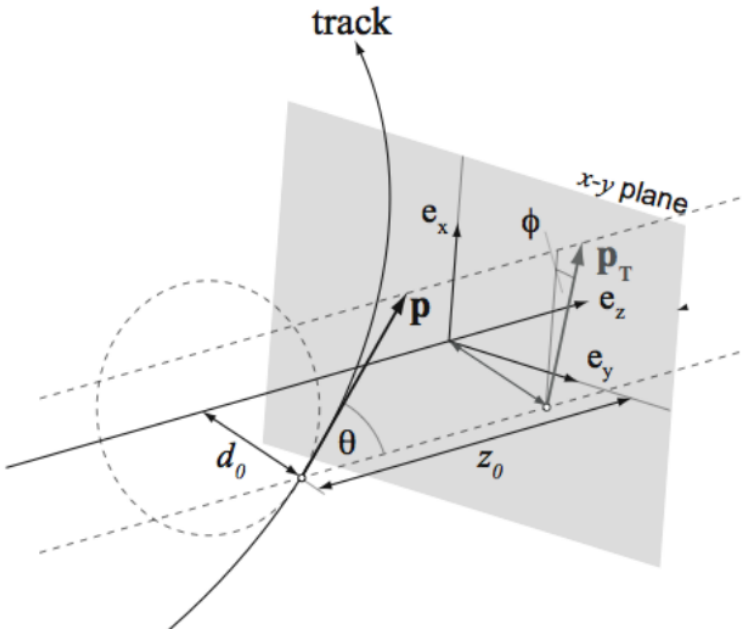

Figure 5.1: Illustration of the relationship between the track parameters and associated track. In this scenario, the hard scatter primary vertex is located at $(e_x, e_y, e_z) = (0, 0, 0)$, though this is not generally the case.

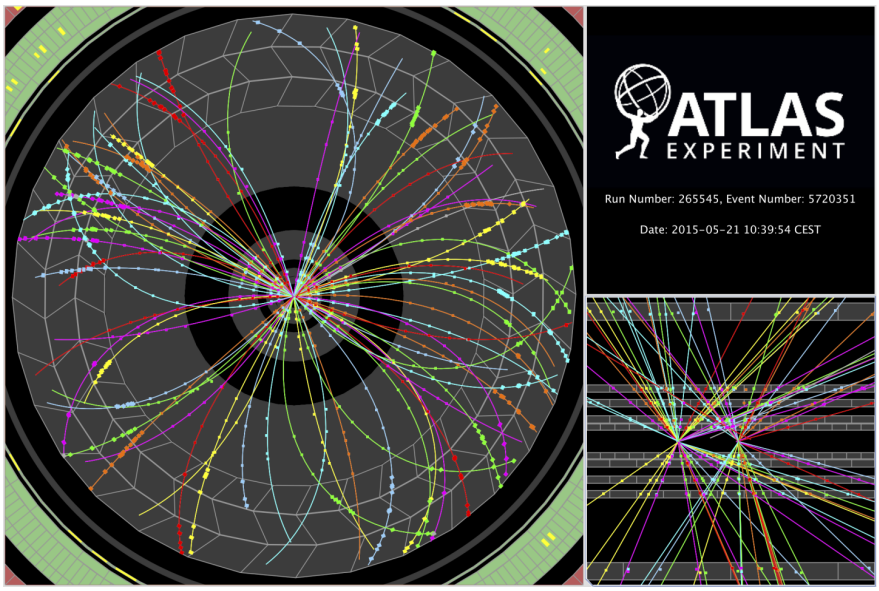

Figure 5.2: Event display of a low-pileup event recorded at the start of Run 2, in early 2015. ***Left***: Transverse view of the ID. Seen in color are the reconstructed tracks traversing the inner layers of the pixel detector, SCT, and TRT. The colored dots are all reconstructed spacepoints used as input to the track fitting procedure. ***Right, lower***: View in $r-z$ of the same pp bunch-crossing event as on the left. Two reconstructed primary vertices are clearly observed. On average, in Run 2 there were roughly 30 primary vertices reconstructed per event, with up to ≈ 65 occurring at maximum.

5.2 Electrons and Muons

Electrons and muons, being charged particles, leave identifiable tracks within the ID. As a result, their reconstruction involves the use of the tracks and vertices described in the previous section, using them essentially as initial seeds for their complete reconstruction. Electron reconstruction, described in Section 5.2.1, complements the track information provided by the ID with calorimetric information provided by the EM calorimeter (Section 3.2.3) and with knowledge about the pattern of transition radiation expected to occur in the TRT as a result of passing electrons. Muon reconstruction, described in Section 5.2.2, revolves around stitching together the tracks reconstructed in the ID with those tracks independently reconstructed in the MS layers at large radii.

5.2.1 Electrons

Electron Reconstruction

The reconstruction of electron candidates is based on three components which characterise the signature of electrons: localised clusters of energy deposits found within the EM calorimeter, charged-

particle tracks identified in the ID, close-matching (in (η, ϕ)) of the tracks to the clusters that form the final electron candidates [134]. It is generally possible to match multiple tracks to the same electromagnetic cluster, all originating from the same primary electron produced in the hard-scatter. This is due to the fact that electrons lose significant amounts of energy to bremsstrahlung photons as they interact with and traverse the ID. These radiated photons can then undergo conversion to electron-positron pairs, which, too, can undergo further bremsstrahlung. The positrons, electrons, and photons are usually emitted in a very collimated fashion and thus deposit most of their energy in a localised fashion within the calorimeter.

The search for localised energy deposits in the EM calorimeter is performed by following a sliding window algorithm over the individual cells whose dimensions are defined by the second sampling layer of the EM calorimeter (Figure 3.17). Electron candidates are seeded by localised energy deposits whose summed transverse energy, across all layers of the EM calorimeter, is greater than 2.5 GeV [134]. These clusters act as seeds for the matching of reconstructed ID tracks. The reconstructed tracks are refit using a Gaussian Sum Filter (GSF) method [135] that accurately accounts for the bremsstrahlung energy losses characteristic of electron traversal and are then matched to the localised clusters using the cluster barycenter as the point of reference to match in $\eta - \phi$. If there is no GSF-track candidate matching to the EM calorimeter cluster seed, then the cluster is marked as an unconverted photon. The cluster is marked as a converted photon if a matched GSF-track candidate exists but is not associated with the primary hard-scatter vertex.

Starting after 2016, the ATLAS electron reconstruction algorithm was updated to reconstruct electron energy deposits using a dynamic EM topo-clustering algorithm similar to that used in jet reconstruction (Section 5.3) [136], as opposed to the sliding window algorithm described above.

Electron Identification

Once electron candidates are reconstructed, they are selected based on various levels of identification. A further set of identification criteria is required on top of the reconstruction so as to improve the selection of true electrons originating from the primary hard-scatter vertex — so-called *prompt* electrons — over *non-prompt* sources of reconstructed electrons such as those originating from photon conversions or the misidentification of charged pions that leave electron-like tracks in the ID. This identification criteria is based on the construction of a multivariate likelihood (LH) and is referred to as the *electron likelihood identification*. The inputs to the LH are listed in Table 5.1 and include measurements from the tracking system in the ID, calorimetric information, and quantities that combine the tracking and calorimetric information [136].

The electron LH is based on the products for the signal and background probability density functions (PDFs) associated with the set of inputs in Table 5.1:

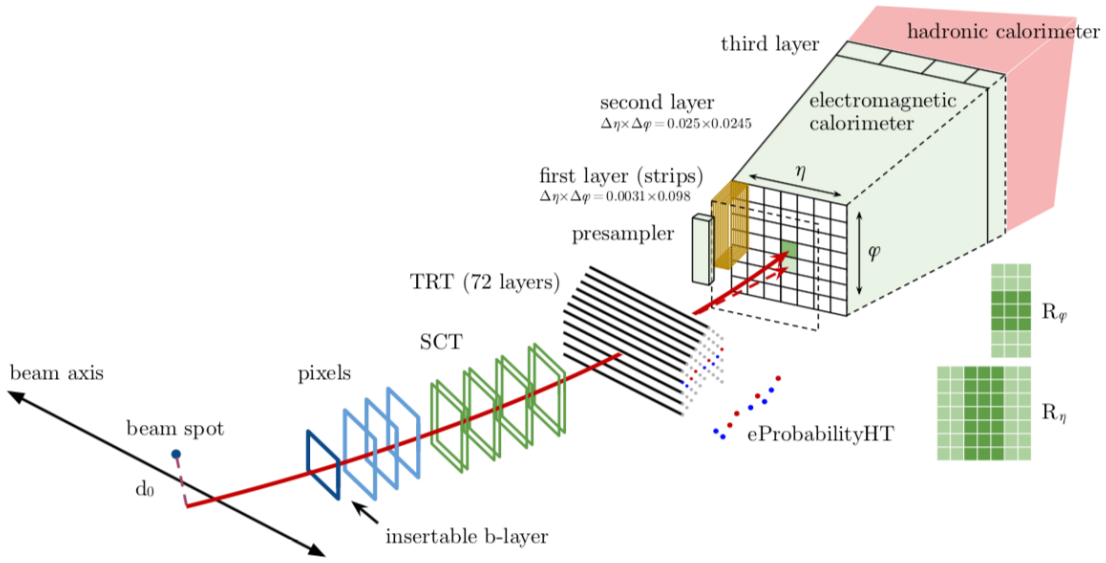

Figure 5.3: Cartoon illustrating the quantities going into the electron identification algorithm, listed in Table 5.1. Figure taken from Ref. [134].

$$L_{S\,(B)}(\mathbf{x}) = \prod_{i=1}^{n} P_{S\,(B),i}(x_i), \tag{5.2}$$

where $\mathbf{x}$ is the vector of quantities listed in Table 5.1 and the $P_{S\,(B),i}(x_i)$ are the values of the PDF for quantity i at value x_i for the signal (S) and background (B). The likelihoods are built using simulation and the signal is composed of samples of prompt electrons and the background is built from a combination of jets that mimic the signature of prompt electrons, electrons from photon conversions, and non-prompt electrons from the decay of hadrons containing heavy-flavours [136]. The final electron LH discriminant, shown in Figure 5.4, is based on a transformed version of the ratio,

$$d_L = \frac{L_S}{L_S + L_B}, \tag{5.3}$$

where the transformation acts to spread d_L to values not bounded by 0 and 1, motivated by the need to have well-defined working points based on selections on d_L.

There are four such fixed values of the final LH discriminant that are used to define four working points corresponding to increasing thresholds on the final LH discriminant: VeryLoose, Loose, Medium, and Tight. The efficiencies to identify prompt electron candidates are measured using samples of $Z \to ee$ and $J/\psi \to ee$ following a tag-and-probe approach. They are found, for electron candidates with $E_T > 40\,\text{GeV}$, to be 93%, 88%, and 80% for the Loose, Medium, and Tight

working points, respectively [136].

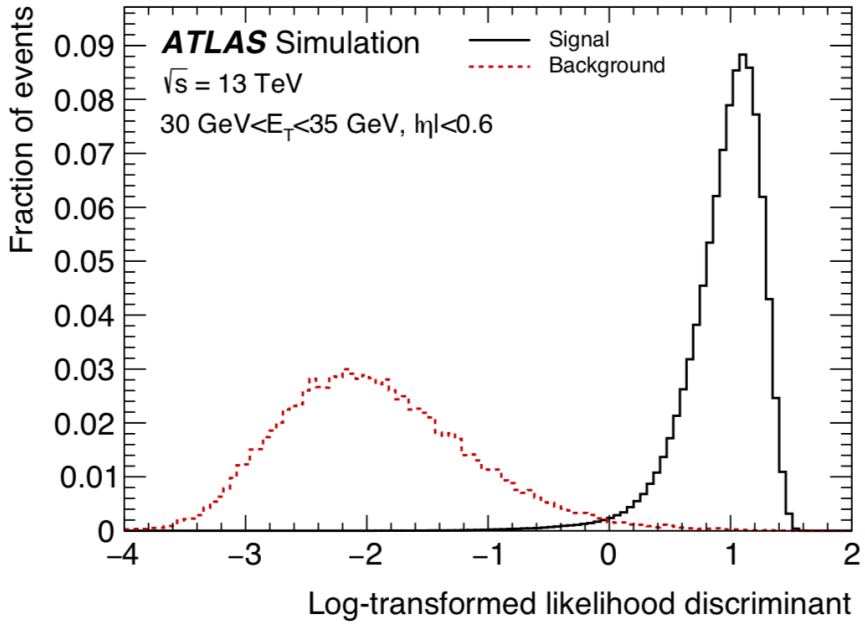

Figure 5.4: Transformed LH-based electron identification discriminant for electron candidates with $30\,\text{GeV} < E_T < 35\,\text{GeV}$ and $|\eta| < 0.6$. Figure taken from Ref. [134].

Table 5.1: Variables used as input to construct the electron identification likelihood. From Ref. [136].

Input Type	**Name**	**Description**
Hadronic Leakage	R_{had1}	Ratio of E_T in the first layer of the hadronic calorimeter to E_T of the EM cluster
	R_{had}	Ratio of E_T in the hadronic calorimeter to E_T of the EM cluster
Third layer of EM calorimeter	f_3	Ratio of the energy in the third layer to the total energy in the EM calorimeter. Only used for $E_T < 30\,\text{GeV}$ and $\lvert\eta\rvert \leq 2.37$.
Second layer of EM calorimter	$w_{\eta 2}$	Lateral shower width, $\sqrt{(\sum E_i \eta_i^2)/(\sum E_i) - ((\sum E_i \eta_i)/(\sum E_i))^2}$, where E_i is the energy and η_i is the pseudorapidity of cell i and the sum is calculated within a window of 3×5 cells centered at the electron cluster position.
	R_ϕ	Ratio of the energy in 3×3 cells over the energy in 3×7 cells centered at the electron cluster position.
	R_η	Ratio of the energy in 3×7 cells over the energy in 7×7 cells centered at the electron cluster position.
First layer of EM calorimeter	w_{stot}	Shower width, $\sqrt{(\sum E_i (i - i_{\text{max}})^2)/(\sum E_i)}$, where i runs over all strips in a window of $\Delta\eta \times \Delta\phi \approx 0.0625 \times 0.2$, corresponding typically to 20 strips in η, and i_{max} is the index of the highest-energy strip. Used only for $E_T > 150\,\text{GeV}$.
	E_{ratio}	Ratio of the energy difference between the maximum energy deposit and the energy deposit in a secondary maximum in the cluster to the sum of these energies.
	f_1	Ratio of the energy in the first layer to the total energy in the EM calorimeter.
Track conditions	n_{Blayer}	Number of hits in the innermost pixel layer.
	n_{Pixel}	Number of hits in the pixel detector.
	n_{Si}	Total number of hits in the pixel and SCT detectors.
	d_0	Transverse impact parameter relative to the beam-spot.
	$\lvert d_0/\sigma(d_0)\rvert$	Significance of transverse impact parameter defined as the ratio of d_0 to its uncertainty.
	$\Delta p/p$	Momentum lost by the track between the perigee and the last measurement point divided by the momentum at perigee.
TRT	eProbabilityHT	Likelihood probability based on transition radiation in the TRT.
Track-cluster matching	$\Delta\eta_1$	$\Delta\eta$ between the cluster position in the first layer and the extrapolated track.
	$\Delta\phi_{\text{res}}$	$\Delta\phi$ between the cluster position in the second layer of the EM calorimeter and the momentum-rescaled track, extrapolated from the perigee, times the charge q.
	E/p	Ratio of the cluster energy to the track momentum. Used for $E_T > 150\,\text{GeV}$.

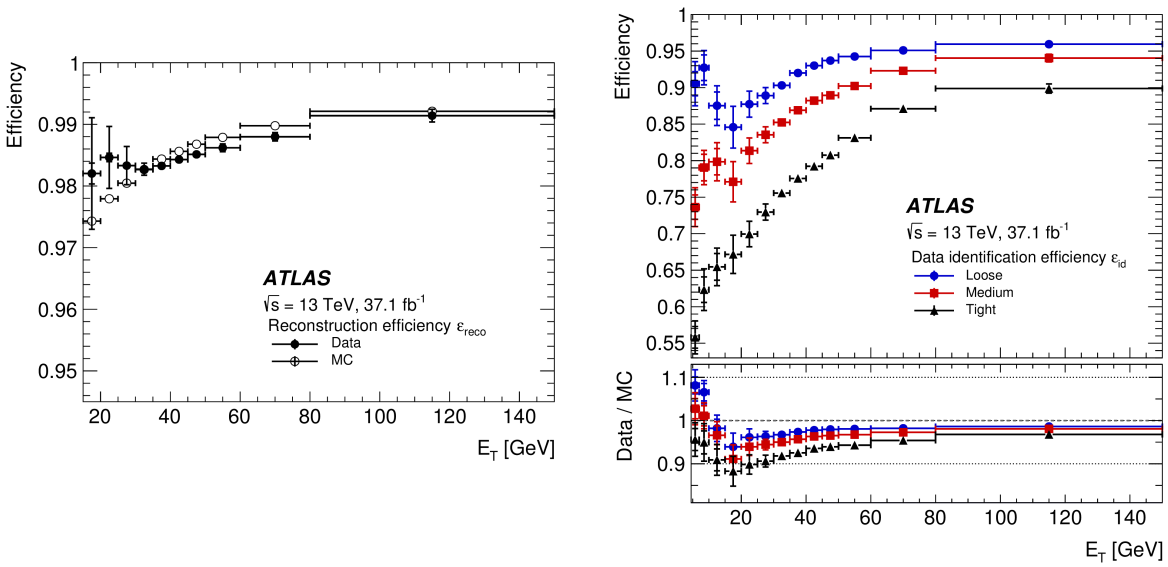

Figure 5.5: Electron reconstruction and identification efficiencies as measured in 2015–2016. Figures taken from Ref. [134]. ***Left***: Electron candidate reconstruction efficiency, measured in simulation and in data, as a function of the candidate E_T. ***Right***: Electron identification efficiency measured in data, as a function of electron E_T, for the three standard LH identification working points LOOSE (blue), MEDIUM (red), and TIGHT (black). The lower panel shows the ratio of the efficiency measured in data over that measured in simulation.

5.2.2 Muons

Muon Reconstruction

The reconstruction of muon candidates is performed by combining the tracking capabilities of the ID and the MS [137]. Muon reconstruction first starts with the independent reconstruction of charged-particle tracks within the ID and the MS. The independently formed tracks are subsequently combined to form a complete track representing the traversal of a muon through the full detector. The muon track in the ID is reconstructed like any other charged-particle track (Section 5.1).

Muon reconstruction within the MS starts with a pattern finding phase, looking for hit patterns each of the muon chambers to then form track segments. The track segments between different MS layers are then fit together to form muon track candidates. At least two matching track segments are required in order to build a muon track candidate, except in the transition region between the barrel and end-cap where single track-segment candidates can be used. Once a muon track candidate is formed from the combined segments, a global χ^2 fit is performed to improve the association of hits to each muon candidate. The χ^2 is repeated several times, removing outlying

hits as necessary, until a threshold is met for all associated hits.

There are several algorithms used to combine the muon track candidates in the ID and MS, each using different sets of information related to the ID, MS, and calorimeters. At the time of the current work, there are four standard combination algorithms used each based on the subdetectors used in their construction:

- **Combined Muon (CB)** This type of muon is formed with a global refit using all muon track candidate hits in the ID and the MS. Hits may be added or removed from the MS track candidate during the refit. Muons are reconstructed following an outside-in pattern recognition algorithm, in which the muon is first reconstructed in the MS and extrapolated inwards to the ID hits. A complementary, albeit non-standard, inside-out algorithm also exists.

- **Segment-tagged Muon (ST)** An ID track is classified as a muon if, once extrapolated to the MS, it is associated with at least one local track segment in the MDT or CSC chambers. Segment-tagged muons are used when a muon candidate crosses only one layer of the MS chambers, either because of their low p_T or because they fall into un-instrumented regions of the MS.

- **Calorimeter-tagged Muon (CT)** An ID track is classified as a muon if it is matched to an energy deposit in the calorimeter that is compatible with a minimum ionising particle. Calorimeter-tagged muons have the lowest purity, but recover acceptance in regions of the MS that are only partially instrumented to allow for cabling and services to the calorimeter and ID systems, particularly in the region $|\eta| < 0.1$.

- **Extrapolated Muon (ME)** This type of muon is based only on the track candidates formed in the MS and a loose requirement that the track candidate be pointing back towards the IP. Extrapolated muons are mainly used to extend the acceptance of muon reconstruction into the region $2.5 < |\eta| < 2.7$ that is not covered by the ID acceptance.

Muon Identification

Muon identification refers to the act of applying additional quality criteria on the reconstructed muon candidates in order to mainly suppress contamination from background sources that mimic muon signatures, such as pion and kaon decays, while ensuring high rates for the acceptance of prompt muons. There are three standard muon identification working points in ATLAS, each a subset of the previous one, and are referred to as the Loose, Medium, and Tight muon identification working points. Medium muons are the default in ATLAS analyses and can only be

composed of CB and ME muons. As all muons used in the present thesis are MEDIUM muons, only this identification working point will be described in detail.

Reconstructed muon candidates originating from non-prompt sources such as in-flight decays of charged hadrons in the ID, are often characterised by the presence of a 'kink' in the reconstructed muon's track. It is therefore expected that the independent momentum measurements made in the ID and MS may be incompatible for non-prompt sources of muon candidates. The muon identification criteria, then, make use of quantities that relate the ID and MS muon track candidates. These quantities are described in Table 5.2. MEDIUM muons have a rather loose selection on the compatibility between the ID and MS momentum measurements and, with respect to those quantities in Table 5.2, are only required to have a q/p significance less than 7. On top of requirements on those quantities described in Table 5.2, the muon identification working points place additional requirements on the number and type of hits in the ID and MS. All identification working points require, in the ID, that there be at least 1 hit in the pixel subdetector, at least 5 hits in the SCT subdetector, less than 3 silicon holes,[1] and at least 10% of the TRT hits originally assigned to the muon track candidate exist after the combined reconstruction. MEDIUM muons further require that the CB muons have at least 3 hits in at least two MDT layers, except in $|\eta| < 0.1$ where tracks with at least one MDT layer but no more than one MDT hole are allowed. The ME Medium muons are required to have at least 3 MDT or CSC layer hits, and are employed only in $2.5 < |\eta| < 2.7$.

Table 5.2: Quantities used in the muon identification.

Quantity Name	**Measurement**	**Description**
q/p significance	$\lvert(q/p)^{\mathrm{ID}} - (q/p)^{\mathrm{MS}}\rvert / \sqrt{\sigma^{\mathrm{MS}}_{p_T} + \sigma^{\mathrm{ID}}_{p_T}}$	Absolute value of the difference between the ratio of the charge and momentum of the muon candidates measured in the ID and MS, divided by the quadrature sum of the corresponding uncertainties.
ρ'	$\lvert p_T^{\mathrm{MS}} - p_T^{\mathrm{ID}}\rvert / p_T^{\mathrm{Combined}}$	Absolute value of the difference between the transverse momentum measurements in the ID and the MS, divided by that of the combined muon candidate.
χ^2_{norm}	–	Normalized χ^2 of the combined muon track fit

[1] A missing hit is considered a 'hole' only if it falls between hits successfully assigned to a given track.

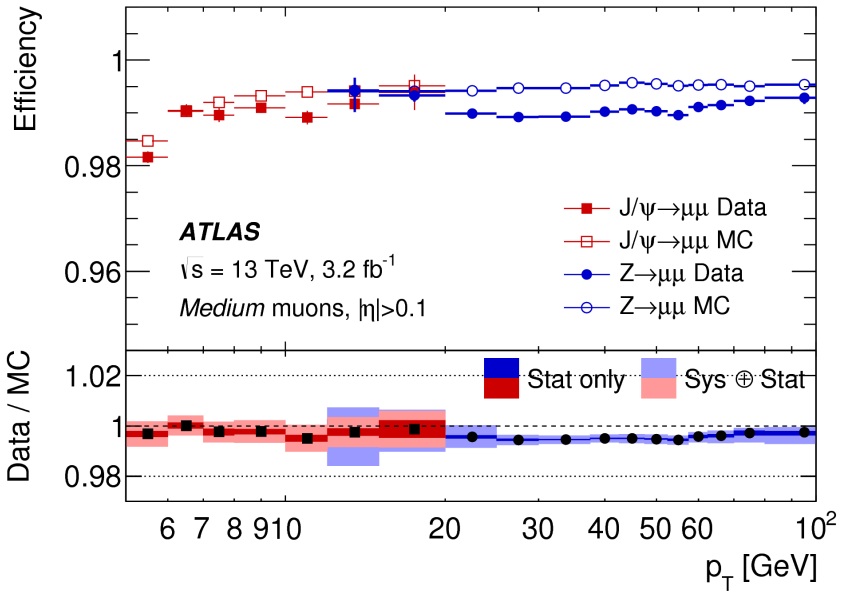

Figure 5.6: Reconstruction efficiency for Medium muons, based on data collected in 2015. Figure taken from Ref. [137].

5.2.3 Lepton Isolation

Prompt electrons and muons originating from the decays of heavy particles, such as W, Z, or Higgs bosons, are often produced well isolated from other particles. Leptons originating from non-prompt sources, such as from the semileptonic decay of b-hadrons within a b-jet (Section 5.3.4) are embedded within a jet (Section 5.3) and are typically surrounded by many additional tracks or calorimeter energy deposits. In order to ensure that the leptons selected in the analyses are leptons arising from prompt sources from the primary pp hard-scatter vertex, leptons must satisfy specific *isolation criteria*. The isolation criteria relevant to the present thesis are based on track- and calorimeter-based isolation variables. The track-based isolation variables, $p_T^{\text{varconeXX}}$, are defined as the sum of the p_T of selected tracks within a cone centered around the lepton, excluding that lepton's tracks, where the surrounding cone has a variable radius, is defined as follows,

$$\Delta R = \min\left(\frac{10}{p_T\,[\text{GeV}]}, \Delta R_{\text{max}}\right) \tag{5.4}$$

where ΔR_{max} is the maximum cone size. The variable radius compensates for the fact that leptons produced in the decay of high-momentum heavy particles tend to have other decay produces very close to the lepton's direction of travel. The isolation variable $p_T^{\text{varcone20}}$ ($p_T^{\text{varcone30}}$) has $\Delta R_{\text{max}} = 0.2$ (0.3). The calorimeter-based isolation variables are similarly defined as the track-based ones, but are based on the transverse calorimeter energy deposits within a surrounding cone (and not associated with the lepton). The calorimeter-based isolation variable considered for the present thesis is

$E_T^{\text{topocone20}}$, and has a surrounding cone with a fixed radius of $\Delta R = 0.2$. The isolation criteria that all leptons relevant to the present thesis must satisfy is based on a two-dimensional, varying selection on track- and calorimeter-based isolation variables and is called the 'GradientLoose' isolation working point. The GradientLoose isolation working point, and its associated isolation variables used for the two-dimensional selection, are defined in Table 5.3.

Table 5.3: Isolation variables (second and third columns) and the isolation criteria (right column) for the electrons and muons relevant to the present thesis. When the isolation variables are divided by the associated lepton's p_T, they are referred to as 'relative isolation variables'.

	Isolation Variables		**Isolation Selection**
	Electrons	**Muons**	**GradientLoose**
Track-based	$p_T^{\text{varcone20}}/p_T^e$	$p_T^{\text{varcone30}}/p_T^\mu$	2D selection such that:
Calorimeter-based	$E_T^{\text{topocone20}}/p_T^e$	$E_T^{\text{topocone20}}/p_T^\mu$	$\geq 95\%$ ($\geq 99\%$) efficient at $p_T \geq 25$ (≥ 60) GeV

The efficiency for leptons to satisfy the isolation criteria outlined in Table 5.3 are measured both in MC and in data, as shown in Figure 5.7 (in the specific case of muons), using $Z/\gamma^* \to \ell\ell$ (high-p_T) and $J/\psi \to \ell\ell$ events (low-p_T). Data-to-MC SFs are derived as the ratio of the data and MC efficiencies (bottom panel of Figure 5.7) and are applied during the analyses, on a *per lepton* basis, to correct the MC-based efficiency to satisfy the isolation selection criteria.

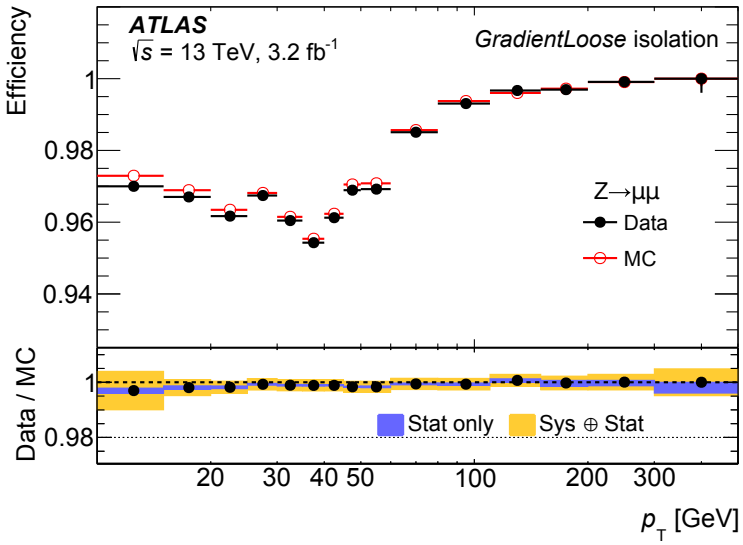

Figure 5.7: Isolation selection efficiency for the GradientLoose isolation working point for muons satisfying the Medium identification, measured in data and in MC, as a function of the muon p_T. The bottom panel shows the ratio of efficiencies measured in data and in MC, giving the efficiency scale-factor. Figure taken from Ref. [137].

5.3 Jets

Due to the confining nature of QCD, color-charged quarks and gluons produced in the initial pp interactions do not exist as free states for observably meaningful timescales and therefore do not leave unambiguous signatures in the detector. Instead, their production is characterised by the radiation of additional quarks and gluons roughly collinear with the initiating colored particles. The radiation pattern of these colored objects is dictated by the color field that binds them and eventually results in the production of color-neutral hadrons. The collimated spray of hadrons as a result of this *hadronisation* process leads to the phenomenology of *jets*, which are the macroscopically observable signature of produced quarks and gluons. The reconstruction of jets refers to any suitable, i.e. physically meaningful and stable, method for grouping together, or *clustering*, the end-products of the hadronisation process in such a way that the properites of the initiating quarks or gluons, such as their quantum numbers and/or kinematics, can be inferred from the resulting clustered object. The standard method for reconstructing jets in ATLAS will be introduced in Section 5.3.1. In Section 5.3.2, the steps taken to turn these reconstructed jets into accurate representations of the initiating quarks and/or gluons, so that they can be used meaningfully in physics analyses, will be discussed.

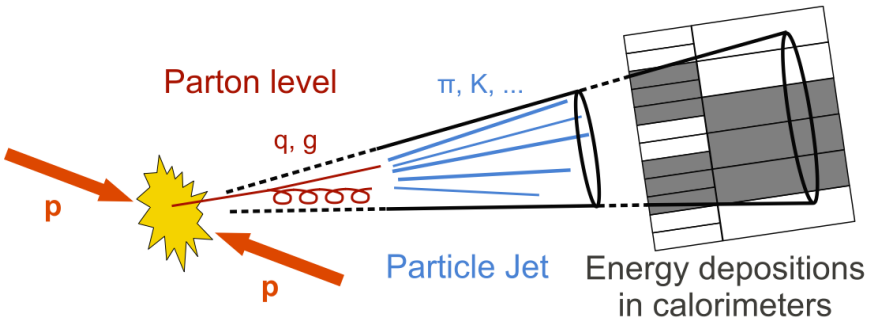

Figure 5.8: Illustration of the jet formation process, beginning with the initiating quark and/or gluons (partons) which hadronise to form particle jets discernible by the tracking detectors in the ID and calorimeter jets defined by energy depositions in the calorimeter systems.

5.3.1 Jet Reconstruction

Topological Cell Clustering

The process of jet reconstruction begins first with the clustering of the lowest level calorimeter elements, *calorimeter cells*, corresponding to the readout channels in the LAr and tile calorimeters. Figure 5.9 gives an idea of the calorimeter cell granularity across the calorimeter system. The clustering algorithm used by ATLAS is a three-dimensional *topological clustering* algorithm [138,

139]. The highly granular calorimeter system used in ATLAS, with its finely segmented lateral readout and longitudinal sampling layers, allows for the subsequent topological clusters ('topo-clusters') to capture in detail the energy-flow details of jets. Topo-cluster formation begins by first identifying so-called *seed cells* which have a rather high signal-to-noise ratio (S/N), $S/N > 4$. Here, the signal is defined as the absolute value of the calorimter-cell energy measurement, $|E|$, and the noise is defined as the sum in quadrature of the RMS of the electronics and expected pileup noise contributions. Cells neighboring the seed cells satisfying $S/N > 2$ are then collected into the topo-cluster. A neighboring cell is defined in three-dimensions as either the calorimeter cells directly adjacent within the same calorimeter layer as the seed cell, or, if in adjacent layers or in different calorimeter sub-systems, cells having at least partial overlap in the (η, ϕ) plane with the seed. The final set of cells, the perimeter cells, satisfying $S/N \geq 0$ are then collected. This last threshold essentially collects all those cells surrounding already-collected cells within each layer. Figure 5.10 illustrates the concept of topological cell clustering as described here.

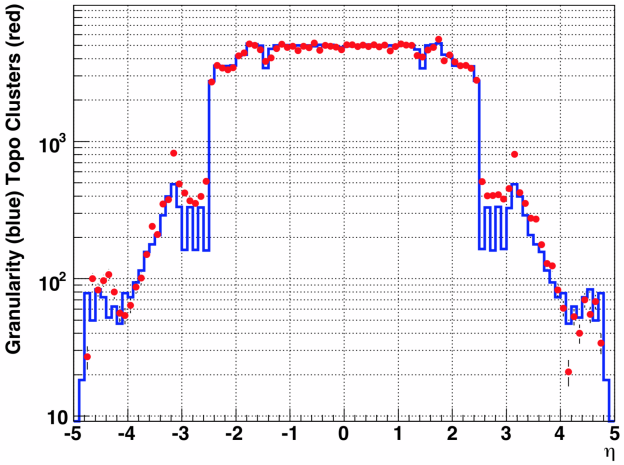

Figure 5.9: The blue histogram shows the average calorimeter cell granularity, i.e. number of calorimeter cells per $\Delta\eta = 0.1$, as a function of detector η. The red points show an approximation of the blue histogram based on calculations of the expected noise per calorimeter cell. Figure taken from Ref. [138].

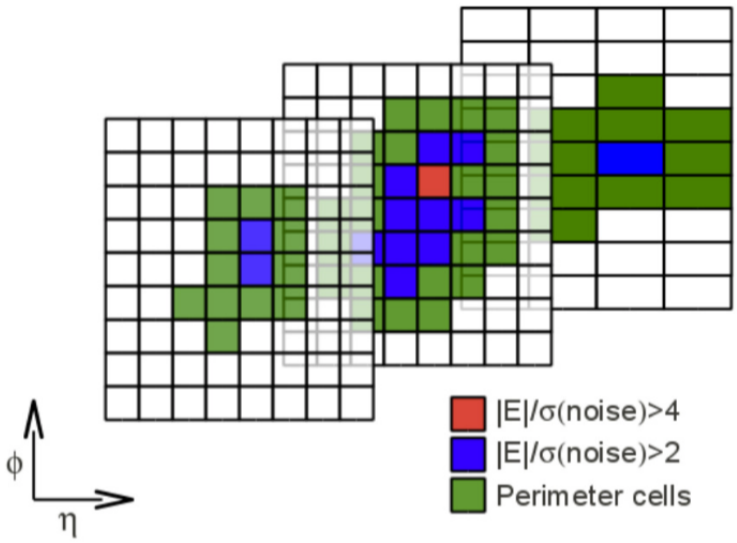

Figure 5.10: Illustration of calorimeter-cell topological clustering across the three layers of the hadronic calorimeter. Indicated are the cells satisfying the signal-to-noise requirements for the seed (red), neighbor (blue), and perimeter (green) cells that make up the final three-dimensional topo-cluster. Figure taken from Ref. [140].

Jet Finding

Once the set of topo-clusters is formed, the process of jet finding begins. As there is no single unique way to define a jet, there is a wide variety of jet finding algorithms whose purpose is to associate jet constituents — here, the calorimeter-cell topo-clusters — to form the final object representing the jet. The default jet finding algorithm used by ATLAS is the *anti-k_t* jet clustering algorithm [141]. The anti-k_t algorithm belongs to the more general class of sequential recombination algorithms and is favored for its infrared and collinear (IRC) safe properties as well as the fact that it tends to produce rather simple jets, geometrically, that are circular in the $y-\phi$ plane, as seen in Figure 5.11. IRC safety in jet finding algorithms refers to the property that neither additional collinear splitting of jet constituents (e.g. the initiating or radiated partons) nor soft emissions should change the clustered jet. IRC-safe jets are therefore robust against these divergent regimes of QCD, sensitive to arbitrary calculational choices made in perturbation theory, and makes them physically meaningful observable objects with which one can make predictions.

The anti-k_t algorithm takes as input constituents the topo-clusters described in Section 5.3.1 and computes the quantities,

$$d_{ij} = \min\left(\frac{1}{k_{T,i}^2}, \frac{1}{k_{T,j}^2}\right)\frac{\Delta R_{ij}^2}{R^2}, \tag{5.5}$$

$$d_{iB} = \frac{1}{k_{T,i}^2}, \tag{5.6}$$

with $\Delta R_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$, R is a parameter whose value regulates the radial extent of the jet, and $k_{T,i}$ is the transverse momentum of the i^{th} constituent. The d_{ij} and d_{iB} quantities are 'distance' metrics used in the clustering of input topo-cluster constituents. The former represents the 'distance' between the i^{th} and j^{th} constituent while latter represents the 'distance' between the i^{th} consituent and the beam-line, introduced to distinguish between constituents originating from the primary hard-scatter vertex and those originating from soft proton interaction remnants. The work to be discussed in the present thesis sets $R = 0.4$, which is the standard used in ATLAS.

The anti-k_t algorithm proceeds by clustering those constituents whose inter-distance is smallest, thereby tending to cluster higher-$p_T$constituents together, which can be seen by inspection of Equation 5.5 and 5.6. If, of the set of input constituents, the smallest distance is a d_{ij}, the associated constituents indicated by i and j are recombined to form a single constituent in the list that replaces them both. If the smallest distance is a d_{iB}, then the constituent indicated by index i is removed from the set of constituents and is considered as a complete jet. This process repeats, starting with the now smaller (due to successful constituent recombination or removal) set of constituents, until no constituents are left. The result of this process is a set of recombined constituents that represent jets, as illustrated in Figure 5.11.

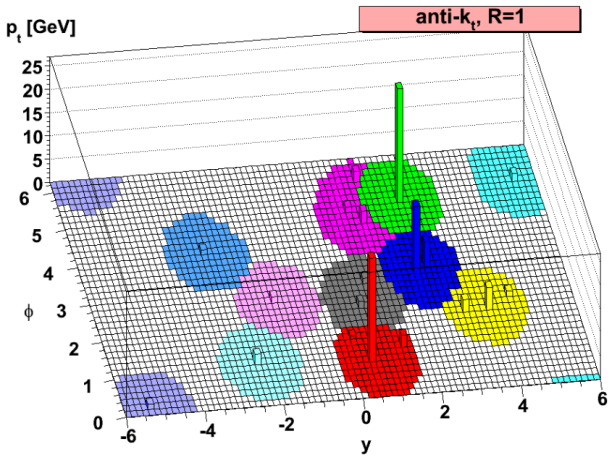

Figure 5.11: An illustration of jet constituents clustered by the anti-k_t algorithm. Seen are the energetic constituents. The filled and colored circles represent areas populated by soft jet consituents, and represent the jet *catchment area* [142] whose size is dictated by the R parameter in the anti-k_t algorithm (Equation 5.5). Figure taken from Ref. [141].

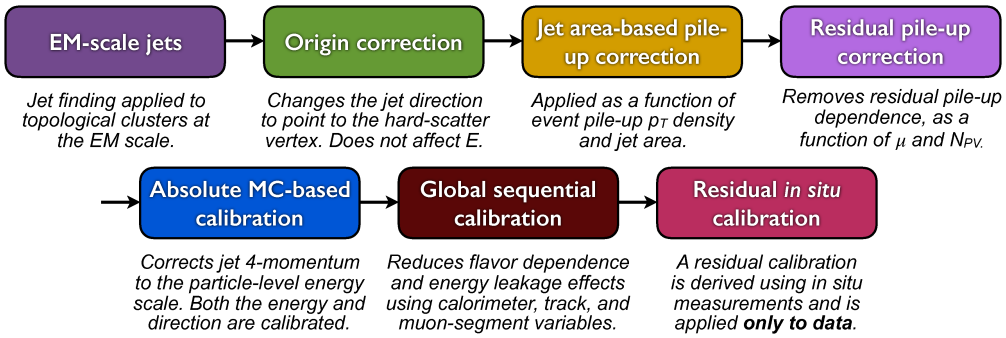

Figure 5.12: Flowchart representing the sequence of steps taken in the jet calibration. Figure taken from Ref. [143].

5.3.2 Jet Calibration

The jets reconstructed following the steps described in Section 5.3.1 are objects clustered at the electromagnetic (EM) scale, which correctly measures the energy of electromagnetic showers but does not accurately account for energy depositions characteristic of hadronic particle decays and interactions. These jets are therefore referred to as 'EM-scale' jets. To correctly assign meaningful energy and momentum measurements to the reconstructed jets that correspond to the energies and momenta of the initiating, underlying particle-level jets, several *jet energy scale* (JES) calibration steps are taken [143]. The steps are detailed in the flowchart in Figure 5.12 and will be briefly described in the following text. The measurements made at each of these steps are subtle and account for many effects not present in the case of the reconstruction of electrons and muons, for example, due to the fact that jets are rather complicated collective phenomena whose measurements rely primarily on a single subsystem (the calorimeters). Not only are electrons and muons generally simpler objects, seeded by comparatively unambiguous tracks within the ID, their energy and momentum measurements are the result of a combination of well-defined measurements made by two independent subsystems (ID measurements combined with EM calorimeter or the MS) that provide independent cross-checks on the validity of the measurements made by each. For the most part, this is not the case of the reconstruction of jets within ATLAS and as a result many of the choices made in the calibration of EM-scale jets have non-negligible impact in the analyses that will be discussed in this thesis, whereas the analogous choices for electrons and muons have minimal impact. For this reason the jet calibration procedure outlined in Figure 5.12 will be outlined briefly in the text that follows.

Jet Origin Correction

The reconstructed EM-scale jets are built with the assumption that they originate from the geometric center of the detector, as opposed to the primary hard-scatter vertex from which the

initiating partons arise. The so-called jet origin correction, therefore, refers to recalculating the jet four-momentum vector by adjusting it in such a way that it points to the primary hard-scatter vertex. This correction acts primarily to improve the η resolution of jets. This procedure is only one hundred percent accurate, of course, under the assumption that all of the jet constituents going into the EM-scale jet reconstruction originated from the hard-scatter vertex, as opposed to some fraction having come from a pileup vertex, for example.

Pileup Corrections

Jets are extended object with relatively large *catchment areas* [142] that make them susceptible to pileup effects. Several corrections, therefore, to the jet energy are taken in order to account for contributions to the EM-scale jet reconstruction arising from both in-time and out-of-time pileup interactions.

The first pileup correction is an area-based correction which subtracts the per-event pileup contribution to the p_T of each jet based on the jet's area, where the jet area is defined as in Ref. [142]. This pileup contribution is taken as the median p_T density, ρ, of jets in the $\eta - \phi$ plane and can be thought of as a baseline 'noise' term distributed evenly throught the calorimeter that contributes to a jet's reconstructed energy.

After the area-based pileup correction is made, there still remains residual dependence of the reconstructed jet p_T on the number of reconstructed primary vertices and on the number of interactions, μ. These dependences are measured by performing linear fits of the jet p_T as a function of each quantity, binned as a function of the detector η, η_{det}.

The final, pileup-corrected jet p_T is given by Equation 5.7:

$$p_T^{\text{Corr}} = p_T^{\text{Reco}} - \underbrace{\rho \times A}_{\text{Area-based}} - \underbrace{\alpha \times (N_{\text{PV}} - 1) - \beta \times \mu}_{\text{Residual}}, \tag{5.7}$$

where A is the jet area, and the α and β terms are derived from the linear fits mentioned above and are $\alpha = \partial p_T / \partial N_{\text{PV}}$ and $\beta = \partial p_T / \partial \mu$, respectively. The former accounts for effects arising as a result of in-time pileup and the latter for those due to out-of-time pileup. The effect of the pileup corrections is shown in Figure 5.13, where it can be seen that the area-based correction is an overall offset, as expected, after which residual pileup dependencies of the jet p_T as a function of $|\eta_{\text{det}}|$ are still observed. This residual pileup dependence is greater in the forward regions of the detector where pileup and background activity is largest.

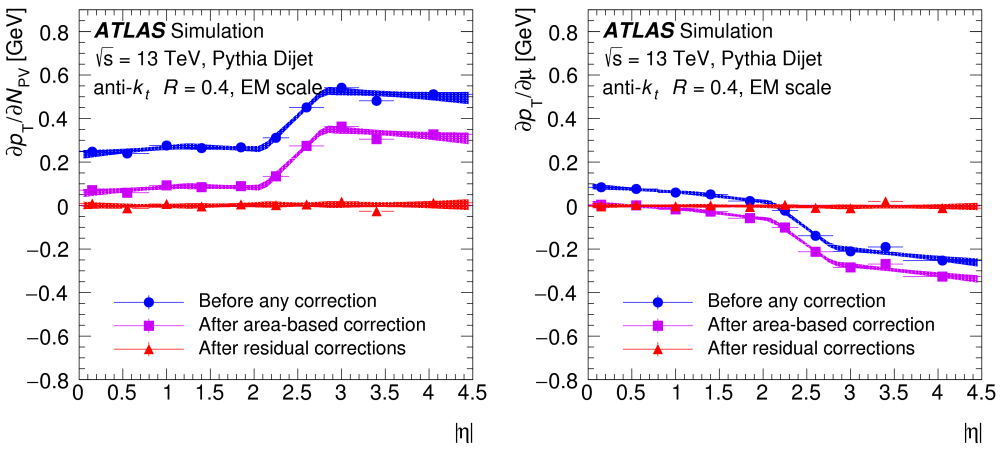

Figure 5.13: Dependence of the p_T of EM-scale reconstructed jets on N_{PV} (in-time pileup) (***left***) and on μ (out-of-time pileup) (***right***). The blue curves show the dependence prior to any pileup corrections, the purple curves are after the area-based correction, and the red curves are the final dependence after the full pileup correction described in Equation 5.7 is taken into account. Figures taken from Ref. [143].

Absolute Jet Energy Scale and η Correction

This correction corrects the EM-scale reconstructed jet to the true energy scale based on particle-level jets and is therefore purely MC-based. Particle-level jets are jets clustered using the anti-k_t algorithm using the generator-level particles at the end of the hadronisation step as the input constituents, and therefore represent the reconstructed jet prior to its interaction with the calorimeter (see Figure 5.8). The correction accounts for mismodelling of the inactive material within the detector, radiation not accounted for in the reconstructed calorimeter-based jet due to the clustering algorithm not accepting it ('out-of-cone radiation'), non-compensation of the hadronic calorimeters, and for effects related to detector geometry or transitions between calorimeter technologies.

The correction is derived by matching the EM-scale reconstructed jets, in simulation, to the particle-level jets and deriving the average energy response, $E^{\mathrm{Reco}}/E^{\mathrm{Truth}}$, where E^{Truth} is the energy of the particle-level jet. The inverse of this energy response is taken as a correction to the EM-scale reconstructed calorimeter jets in simulation. An additional correction accounts for biases observed in the EM-scale reconstructed jet η, which is largest in regions wherein a jet is likely to encompass two calorimeter regions or technologies which result in different energy responses. The η correction is derived as the difference between the reconstructed and particle-level jet η values and applied to the EM-scale reconstructed calorimeter jets as in the case of the energy-response correction. The average energy response, as a function of EM-scale jet p_T, and the η correction is shown in Figure 5.14. After this correction stage, jets are referred to as being at the EM+JES scale.

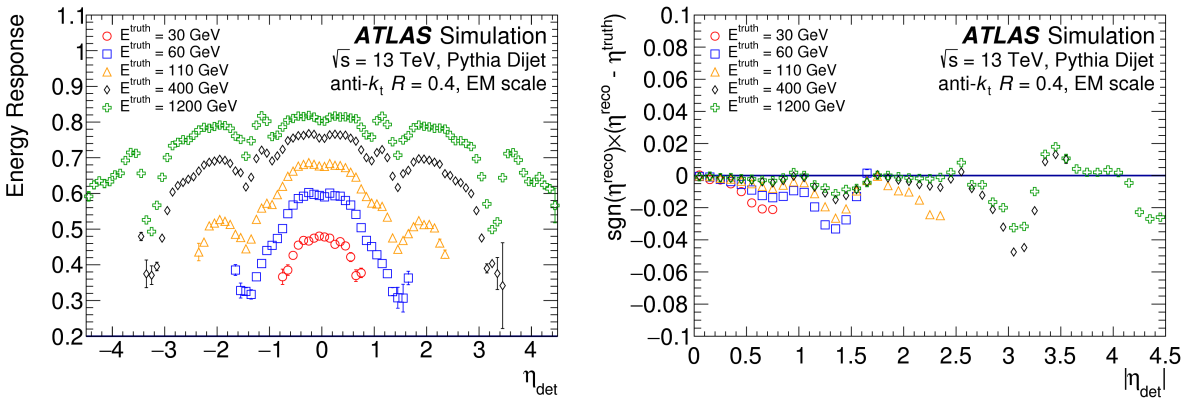

Figure 5.14: Figures taken from Ref. [143]. ***Left***: Average energy response as a function of jet detector η, η_{det}. The colors correspond to different energy regimes for the particle-level jet to which the EM-scale reconstructed jet is matched. The inverse of the response is the final correction and can be seen to be largest for lower-p_T jets. ***Right***: Difference in η for the EM-scale reconstructed jet and the particle-level jet to which it is matched. The bias is clearly seen, with values typically negative, and it being largest for $|\eta_{\text{det}}| \sim 1.4$ (~ 3.1), corresponding to the barrel-endcap (endcap-forward) transition regions.

Global Sequential Calibration

The so-called Global Sequential Calibration (GSC) is a catch-all correction to account for remaining dependencies of the EM+JES-scale jets on the jet shower shapes as well as fluctuations in the jet flavor composition and inter-jet energy distribution. This correction improves the handling of fluctuations in the composition of the particles that initiate the jet; for example, correcting for the differences expected between quark- and gluon-initiated jets. The former (quark-initiated) jets are typically more collimated with fewer, but higher-p_T, hadronic constituents. The latter (gluon-initiated) jets typically contain many more, softer-p_T, particles and have wider transverse profiles (and therefore do not traverse as far into the calorimeter). The GSC has five stages, each following a numerical inversion of a corresponding jet response as in the case(s) described in Section 5.3.2, but are based on observables sensitive to the jet shower profile and growth within the calorimeter as well as on the number and type of tracks associated with the reconstructed jet. The use of tracking information from muons in this correction additionally helps with correcting the energy response of jets that are not fully contained in the calorimeter but leak into the MS, so-called *punch-through jets*. Figure 5.15 illustrates the concept of jet punch-through.

In-situ Corrections

The last stage of the jet calibration accounts for remaining differences in the EM+JES jet response between data and simulation and is primarily derived using events from data, as opposed to MC

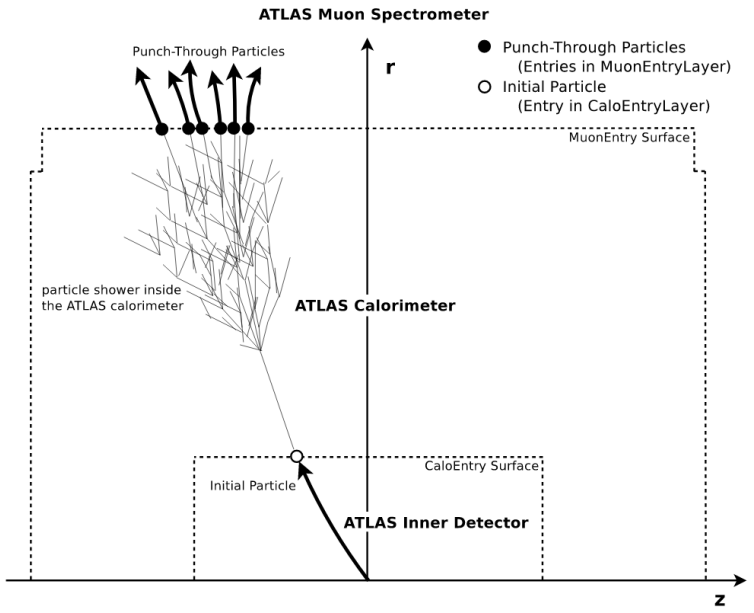

Figure 5.15: Illustration of jet punch-through. High momentum particles produced in the shower of an energetic jet within the hadronic calorimeter escape into the muon spectrometer, leaving detectable signatures in the muon chambers. It can be seen that energy and/or momentum may be left unaccounted for in the final event reconstruction or be assigned to separate muon objects as opposed to the initiating jet from which the punch-through particles arise. Such effects disrupt the proper assignment of energy-momentum to the jet and can spoil the overall momentum balance in the event. Figure taken from Ref. [144].

simulation.[2] There are two classes of correction, the first being the so-called η-intercalibration step which corrects the average response of forward jets ($0.8 < |\eta| < 4.5$) to that of well-measured central jets using dijet events in which the two jets (the one in the forward region and the other in the central region) are back-to-back in ϕ. The second class of corrections is based on the method of balancing an EM+JES scale jet against a well-measured reference object. The balance methods use only central jets ($|\eta| < 0.8$) and the choice of the reference object provides the p_T scale to which the response correction applies. Balance methods using well-measured photons and Z boson decays to leptons ($Z \to ee$ and $Z \to \mu\mu$) as the reference objects are used to derive the response corrections for jets with p_T up to 950 GeV. Response corrections for jets covering p_T ranges up to 2 TeV are derived using a multijet balance method in which a high-p_T central jet is balanced against a reference system of lower-p_T jets. The final response correction as a function of jet p_T derived using these in-situ methods is shown in Figure 5.16.

[2]In this case, the term 'in-situ' means that the corresponding corrections are derived using data, as opposed to using events from Monte-Carlo simulation.

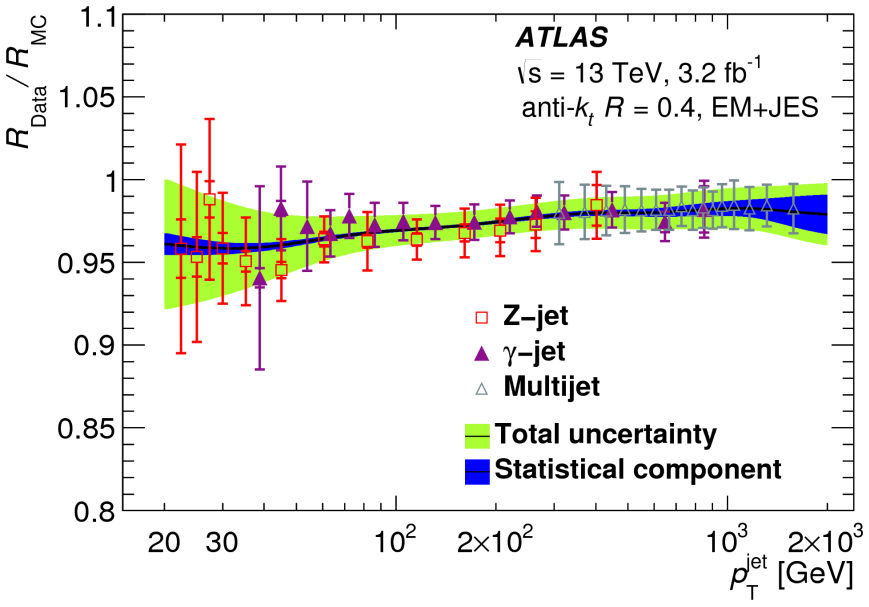

Figure 5.16: Ratio of the EM+JES jet response in data to that in MC simulation as a function of jet p_T. The different markers indicate the measurement contributions from the different reference-object used to derive the in-situ response corrections. The final response correction is given by the black line. The total uncertainty on the response is given by the green band and the blue band indicates only the statistical component of the uncertainty. Figure taken from Ref. [143].

5.3.3 Pileup Jet Suppression

Pileup interactions often result in reconstructed jets in the final state that are not associated with the primary hard-scatter vertex. Such jets act to contaminate the selection of jets for analyses targetting physics processes that naturally lead to jets. A multivariate algorithm has been defined that allows analyses to discriminate between these pileup jets and hard-scatter jets such that the former may be discarded from the collection of reconstructed jets and the purity of the latter may be increased. This algorithm is called the 'Jet Vertex Tagger' (JVT) and it relies on using the tracks associated with reconstructed jets in order to give an indication of how likely a jet is to be due to a pileup interaction or not [145].

JVT is based on two quantities, 'corrJVF' and 'R_{p_T}', to build a k-nearest neighbors (k-NN) multivariate discriminant.[3] These two quantities are based on the tracks that are associated with the reconstructed jets. As tracks naturally provide clear indication as to their originating interaction point, the idea is that jets with tracks predominantly arising from the primary hard-scatter vertex are more likely to be true hard-scatter jets, and not pileup jets.

[3]The quantity 'JVF', appearing in corrJVF, stands for 'Jet Vertex Fraction'.

The quantity 'corrJVF' is defined as follows,

$$\text{corrJVF} = \frac{\sum_k p_T^{\text{trk}_k}(\text{PV}_0)}{\sum_\ell p_T^{\text{trk}_\ell}(\text{PV}_0) + \langle p_T^{\text{PU}} \rangle}, \tag{5.8}$$

with the numerator and first term in the denominator being the scalar p_T sum of the tracks that are associated with the jet and that originate from the primary hard-scatter vertex (PV_0). The term $\langle p_T^{\text{PU}} \rangle$ is a term proportional to the average scalar p_T sum of the tracks that are associated with the jet but that originate from pileup vertices. The quantity 'corrJVF' is expected to have values near one for jets originating from the hard-scatter and values near zero for pileup jets.

The second quantity, R_{p_T}, is defined as follows,

$$R_{p_T} = \frac{\sum_k p_T^{\text{trk}_k}(\text{PV}_0)}{p_T^{\text{jet}}}, \tag{5.9}$$

and represents the p_T sum of the tracks associated with the jet and that originate from PV_0 divided by the jet's calibrated transverse momentum that includes the pileup subtraction (Section 5.3.2). R_{p_T} is peaked at zero for pileup jets, where little to no p_T from tracks from PV_0 are expected. For jets arising from the primary hard-scatter, R_{p_T} is the charged p_T fraction and generally has larger values than for pileup jets. Distributions of the corrJVF and R_{p_T} quantities are shown in Figure 5.17.

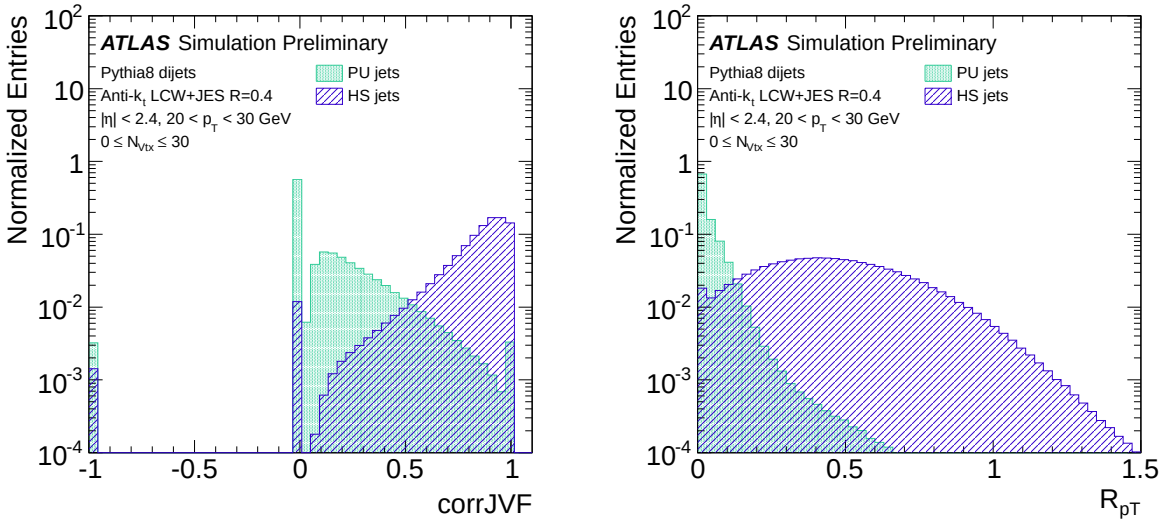

Figure 5.17: Distributions of the variables used in the definition of JVT for pileup and hard-scatter jets. Figures taken from Ref. [145]. ***Left***: corrJVF. Values of -1 are assigned to jets with no associated tracks. ***Right***: R_{p_T}.

The JVT discriminant is constructed using corrJVF and R_{p_T} to construct a two-dimensional likelihood, based on a k-NN algorithm trained on samples of pileup and hard-scatter jets. Each point in the two-dimensional plane gives the relative probabilty for the jet to be a hard-scatter jet. Distributions of the final two-dimensional likelihood and JVT for a sample of pileup and hard-scatter jets are shown in Figure 5.18. JVT values nearer one indicate a higher probability for the given jet to be a hard-scatter jet, whereas lower values indicate that the jet is likely to be a pileup jet.

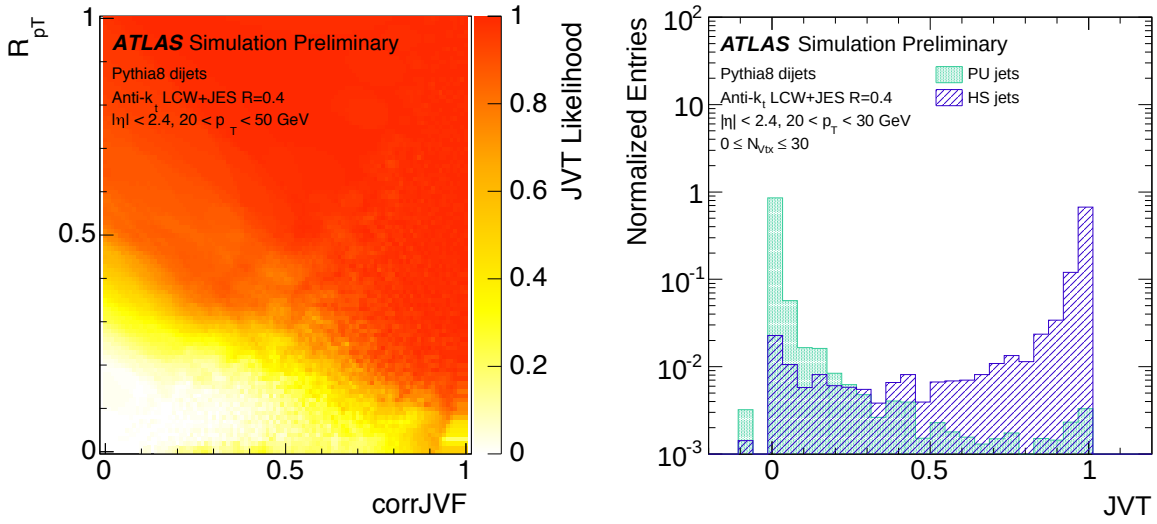

Figure 5.18: Figures taken from Ref. [145]. ***Left***: The two-dimensional JVT likelihood based on the corrJVF and R_{p_T} variables. ***Right***: A distribution of JVT for a sample of pileup and hard-scatter jets. Negative JVT values correspond to jets having no associated tracks.

In order to suppress pileup jets in the physics analyses to be described in Chapters 8 and 9, each jet is required to satisfy JVT > 0.59 which gives a 92% selection efficiency for hard-scatter jets. This requirement is applied only to jets with $p_T < 60\,\text{GeV}$ and within the tracking volume of the ID ($|\eta| < 2.4$) since the contribution of pileup jets at higher p_T values is considered to be negligible. The efficiencies and corresponding data-to-MC efficiency correction scale-factors (SFs), correcting the MC efficiency to that of the observed data, of this cut are derived using $Z \to \mu\mu$ events.

5.3.4 Flavor Tagging of Jets

The ability to identify jets containing heavy-flavored hadrons, i.e. jets containing b- and c-flavored hadrons, is an important aspect to many of the most critical measurements and analyses being done at the large LHC experiments. The SM top-quark is the heaviest known elementary particle, the second-most recent elementary particle to be discovered, and, given its importance to electroweak and Higgs physics, is an object subjected to high levels of precise study at the LHC. The top-quark

decays before hadronization timescales and therefore its decay products, which are a W-boson and a b-quark nearly 100% of the time, carry away information characterising its properties. Being able to characterise the jets initiated by the hadronization of these b-quarks, then, is of critical import if the physics of the top-quark wish to be understood. The recent discovery of an SM-like Higgs boson, with a mass of $m_h = 125\,\text{GeV}$, decays to a pair of b-quarks nearly 60% of the time. Without question, then, the thorough study of the Higgs boson necessitates the ability to precisely identify the presence of the pair of b-initiated jets from its decays. Additionally, we will see in subsequent chapters that the presence of b-quark initiated jets is a characteristic signature of many BSM physics scenarios. The identification of these types of jets, a process referred to as 'flavor tagging', is of the utmost importance to analyses performed with the ATLAS detector as well as in the work to be presented in this thesis. Jets tagged as having likely arisen as a result of the hadronization of an initiating b-(c-)quark are referred to as 'b-tagged jets' ('c-tagged jets'), or simply as 'b-jets' ('c-jets'). All other jets then are assumed to have arisen from the decay of light-flavor quarks and are referred to as 'light-flavor jets' ('light-jets').

Heavy-flavor tagging of jets relies on the relatively long lifetimes of the b- and c-hadrons which initiate them. The typical b-hadron lifetime is $\tau \approx 1.6\,\text{ps}$ ($c\tau \approx 450\,\mu m$), which leads to b-hadrons traversing typically macroscopic distances away from the primary hard-scatter vertex before they decay. As seen in Figure 5.19, b-hadrons with transverse momenta on the order of 50 GeV will travel nearly half a centimeter before decaying. Also seen in Figure 5.19, those b-hadrons with p_T values nearing 250 GeV will actually decay outside of the beam-pipe within the region of the IBL. Given the high spatial resolution of the ID pixel and SCT detectors discussed in Section 3.2.2, the presence of these long-lived particles should be detectable via the presence of at least one secondary decay vertex corresponding to the point at which the heavy-flavored hadron decays. Figure 5.20 illustrates the standard topology of a heavy-flavor initiated jet with a secondary decay vertex that is displaced with respect to the primary hard-scatter vertex and is the source of displaced tracks.

The algorithms used to identify heavy-flavor initiated jets, then, rely on information characterising the long lifetimes of the initiating particles and on the presence of secondary decay vertices within the jet. These high-level algorithms used to identify these jets take as input information provided by taggers that rely on low-level information based on the presence of displaced tracks and secondary vertices. These low-level taggers will be introduced in Section 5.3.4 and the construction of the high-level tagger, used in the analyses to be presented in this thesis, will be presented in Section 5.3.4.

Low Level Taggers and Inputs

The final algorithm used to identify b-tagged jets take as input the outputs from several low-level b-tagging algorithms. There are two classes of low-level algorithms: those that rely on the impact parameter information of the tracks associated with the jets and those that rely on the explicit

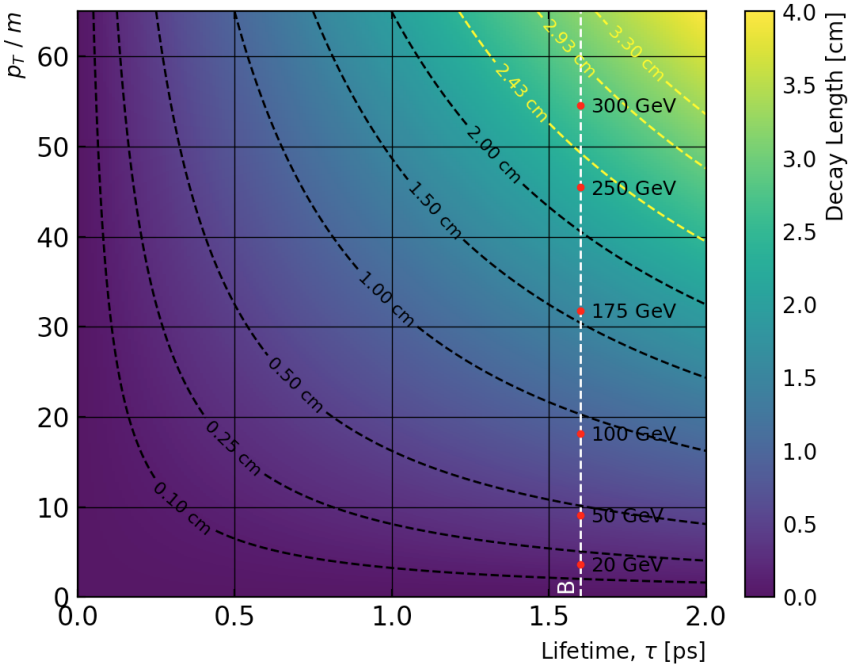

Figure 5.19: Particle decay length as a function of its lifetime and transverse momentum normalised to its rest mass. The white-dashed line indicates the average lifetime of B-hadron species, taken as 1.6 ps, with a mass taken to be 5.5 GeV [54]. The red-dots along the B-hadron line indicate locations for specific transverse momenta for the decaying B-hadron. The yellow contours indicate the locations of the IBL (Figure 3.15), with 2.43 cm corresponding to the beam-pipe radius.

reconstruction of secondary decay vertices within the jet [146]:

- **IP2D and IP3D** The IP2D and IP3D algorithms make use of the signed transverse impact parameter significance of tracks to construct discriminating variables. The IP3D algorithm additionally makes use of the longitudinal impact parameter significance. The algorithms rely on constructing log-likelihood ratios (LLR) taking as inputs probability density functions (PDFs) for b-, $c-$, and light-flavor jet probabilities on a per-track basis. See Figure 5.20 for an illustration of the signed impact parameter. Tracks associated with secondary decay vertices in heavy-flavor jets tend to have positively signed impact parameters since they are produced *within* the jet cone. Signed impact parameter quantities tend to be symmetric for light-flavor jets. See Figure 5.21.

- **Secondary Vertex Finding Algorithm (SV1)** The SV1 algorithm [147] reconstructs a single displaced secondary vertex within a jet, starting from the set of all possible two-track vertices while rejecting tracks likely to be associated with non-heavy-flavored long-lived particles (K_s or Λ), photon conversions, or vertices due to detector material interactions. The inclusive secondary vertex and associated tracks are then used to construct discriminating observables sensitive to the differences between b-, c-, and light-flavor jets.

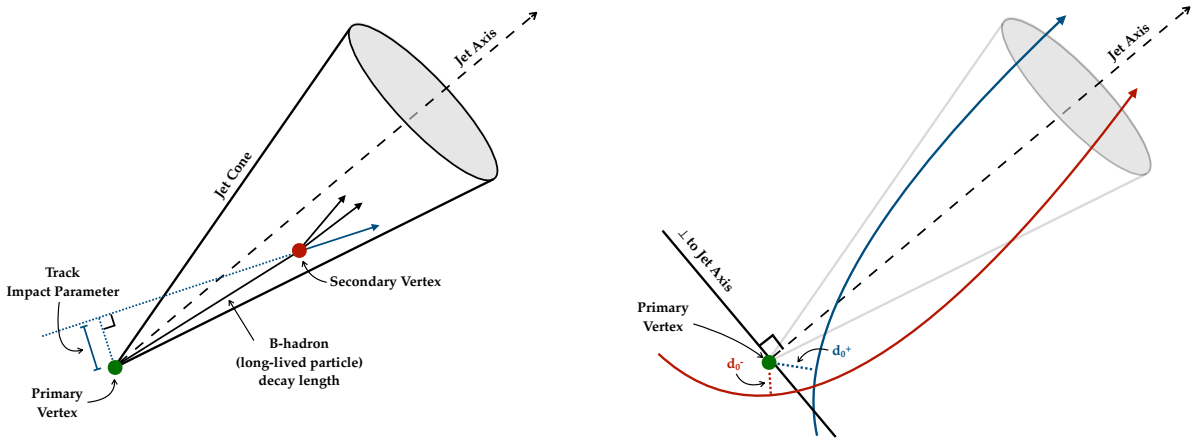

Figure 5.20: ***Left***: Topology of a b-jet. The b-hadron produced near the primary hard-scatter vertex (green dot), initiating the b-jet, has a long lifetime and decays a macroscopic distance away from the primary hard-scatter vertex to produce a secondary vertex (red dot) from which additional tracks are produced and subsequently reconstructed. The tracks originating from the secondary vertex will have larger impact parameters relative to the primary hard-scatter vertex as compared to tracks originating from the primary hard-scatter vertex. ***Right***: Illustration of signed impact parameter, specifically the signed transverse impact parameter d_0. The blue track has a transverse impact parameter whose projection onto the jet axis is in the direction in which the jet momentum points and is given a positive d_0. The red track's transverse impact parameter's projection onto the jet axis points opposite to the jet momentum and is given a negative d_0. The curvature of the tracks is exaggerated for illustrative purposes.

- **Multi-vertex Finding Algorithm (JetFitter, JF)** The JetFitter algorithm [148] attempts to reconstruct the full b-hadron decay chain using a modified Kalman filter [149] which assumes a common line on which the primary, b-hadron, and c-hadron decay vertices lie. As with SV1, the construction of many discriminating observables related to the reconstructed set of vertices are used to build templates for b-, c-, and light-flavor jets.

Figure 5.21 provides an example of a few observables provided by the low-level tagging algorithms.

High Level Tagger: MV2

The low-level taggers discussed in the previous section provide a set of many useful and complementary observables. In an attempt to make the most efficient use of all the information provided by this set of observables, a boosted decision tree (BDT) algorithm is used to combine the outputs of these low-level algorithms. This algorithm, referred to as the MV2 b-tagging algorithm, is trained using the ROOT Toolkit for Multivariate Data Analysis (TMVA) [151]. In the analyses based on data recorded by ATLAS between 2015–2016, the MV2 algorithm was trained using jets from a simulated sample of top-quark pair production events. For the analyses based on the full Run 2

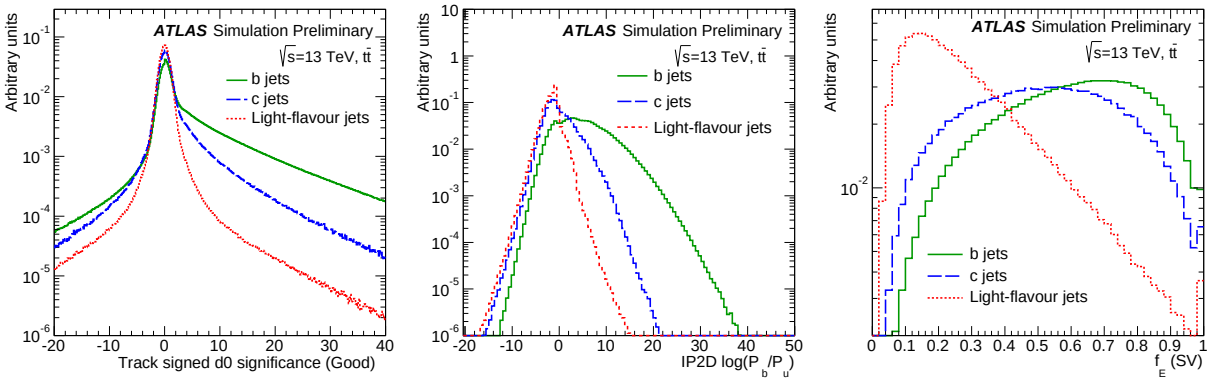

Figure 5.21: Examples of a few low-level quantities used in the ATLAS flavor tagging algorithms. The blue histograms are distributions associated with b-jets, green are those of c-jets, and red are those of light-flavor jets. Figures taken from Ref. [150]. ***Left***: Two-dimensional (signed) d_0 significance for tracks matched to jets. *Middle*: IP2D b-jet log-likelihood ratio. ***Right***: Energy fraction, defined as the energy of the tracks in the displaced vertex reconstructed by the SV1 algorithm relative to the energy of all tracks in the jet.

data recorded by ATLAS, up to and including the year 2018, the MV2 algorithm was retrained using a sample composed of jets both from top-quark pair production events as well as from simulated events of a BSM physics scenario of a heavy Z' decaying to $b\bar{b}$. The latter was included in the retraining so as to allow the algorithm to have in its training sample high-p_T jets that are not present in the SM top-quark pair production events, as illustrated in Figure 5.22.

Figure 5.22 shows a distribution of the MV2 algorithm's output. The MV2 score is computed on a per-jet basis, using the set of low-level inputs listed in Table 5.4. Working points, defined with different target-efficiencies for accepting b-tagged jets, are defined by selection thresholds on the MV2 discriminant. The standard ATLAS b-tagging working points are defined for accepting b-jets with $p_T > 20\,\text{GeV}$ with average efficiencies of 60%, 70%, 77%, and 85%. The working points are based on selections made on the MV2 output score and are defined in Table 5.5, along with the rejection factors[4] for c-jets, τ-jets[5], and light-flavor jets.

b-jet Identification Calibration

As the efficiencies for the MV2 algorithm detailed in Table 5.5 are based entirely on MC simulation, a calibration procedure is performed to correct the MC-based efficiencies to those observed in data. This is necessary to get an accurate prediction in simulation of the rate of b-tagged jets occurring

[4]The rejection factor is defined as the inverse of the efficiency. A rejection factor of 100 means, therefore, that the associated object is accepted — on average — 1 out of every 100 times that it appears.

[5]Hadronically decaying τ leptons are accepted by the b-tagging algorithms at rates higher than light-flavor jets due to the non-negligible decay length of the τ lepton which give them b-like characteristics.

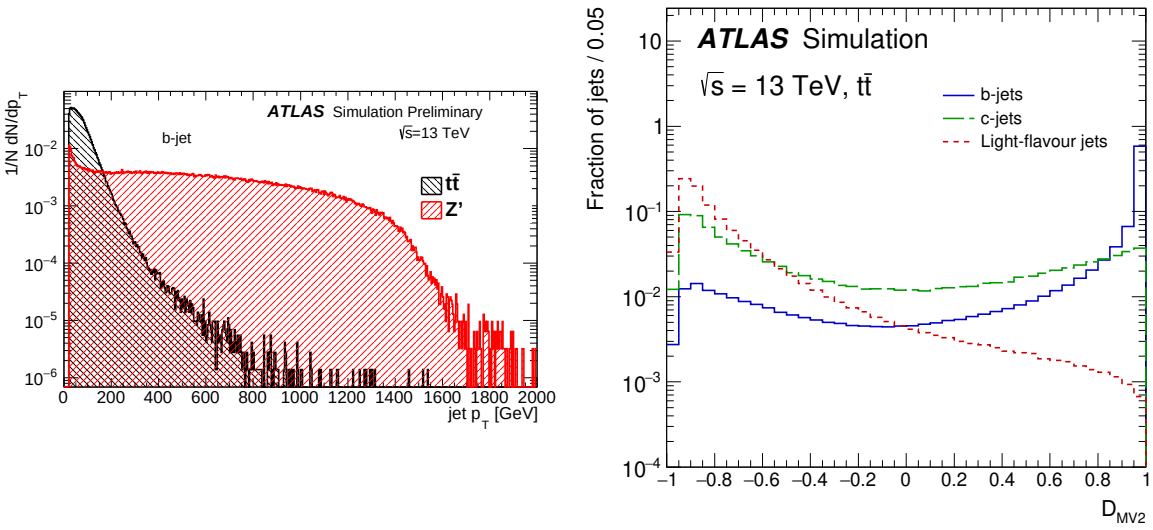

Figure 5.22: Figures taken from Ref. [150]. ***Left***: Distribution of jet p_T for jets used in the training of the MV2 BDT algorithm. The Z' sample of jets is only included in the MV2 training relevant to the analyses based on the full Run 2 dataset collected by ATLAS. ***Right***: Distribution of the BDT-based MV2 b-tagging algorithm output score, shown for b-jets (blue), c-jets (green), and light-flavor jets (red).

in data. The efficiencies (and rejection factors) in Table 5.5 are therefore measured in data. The result of the calibration is a correction scale-factor, applied on a per-jet basis, that is defined as $\text{SF} = \varepsilon_{\text{Data}}/\varepsilon_{\text{MC}}$, where $\varepsilon_{\text{Data(MC)}}$ are the measured efficiencies in data (MC) for a jet to be tagged as a b-jet, measured as a function of the jet p_T and η. The b-tagging SF are derived using a sample enriched in b-tagged jets arising from SM top-quark pair production [146]. The b-tagging efficiencies in MC and data as well as the data-to-MC efficiency SF are shown in Figure 5.23. The rate of c-, τ-, and light-flavor jets to be identified as b-tagged jets corresponding to the rejection factors listed in Table 5.5, also have corresponding correction scale-factors using events in data enriched in c-jets and mis-tagged light-flavor jets [146, 152, 153]. These latter scale-factors correct the rate of mis-tagging (i.e. rate of identifying jets as b-jets when they are not initiated by b-hadrons).

Table 5.4: Variables used as input to the high-level tagger MV2c10.

Input Source	**Input Name**	**Description**
Kinematics	p_T(jet)	Jet transverse momentum
	η(jet)	Jet pseudorapidity
IP2D, IP3D	$\log(p_b/p_{\text{light}})$	Likelihood ratio between the b- and light-jet hypotheses
	$\log(p_b/p_c)$	Likelihood ratio between the b- and c-jet hypotheses
	$\log(p_c/p_{\text{light}})$	Likelihood ratio between the c- and light-jet hypotheses
Secondary Vertex (SV1)	m_{SV}	Invariant mass of tracks at the secondary vertex assuming pion masses
	f_E	Fraction of the charged jet energy in the secondary vertex
	N_{TrkAtVtx}	Number of tracks used in the secondary vertex
	N_{2TrkVtx}	Number of two-track vertex candidates
	L_{xy}	Transverse distance between the primary and secondary vertices
	L_{xyz}	Distance between the primary and secondary vertices
	S_{xyz}	Distance between the primary and secondary vertices divided by its uncertainty
	ΔR(jet, SV)	ΔR between the jet axis and the direction of the secondary vertex relative to the primary vertex
JetFitter (JF)	N_{2TrkVtx}	Number of two-track vertex candidates (prior to JetFitter decay-chain fit)
	m_{JF}	Invariant mass of tracks from displaced vertices assuming pion masses
	S_{xyz}	Significance of the average distance between the primary and displaced vertices
	f_E	Fraction of the charged jet energy in the secondary vertices
	$N_{\text{1-trk vertices}}$	Number of displaced vertices with one track
	$N_{\geq\text{2-trk vertices}}$	Number of displaced vertices with more than one track
	N_{TrkAtVtx}	Number of tracks from displaced vertices with at least two tracks
	$\Delta R(\vec{p}_{\text{jet}}, \vec{p}_{\text{vtx}})$	ΔR between the jet axis and the vectorial sum of the momentum of all tracks attached to displaced vertices

Table 5.5: Working points defined for the MV2 b-jet identification algorithm. The cut thresholds on the MV2 discriminant associated with a given b-jet efficiency (working point) are given in the second column. The rejection factors for c-, τ-, and light-flavor jets are shown in the three rightmost columns. The precise MV2 discriminant thresholds are dependent on the calibration, and differ between the analyses based only on the data collected in the years 2015-2016 and the full Run 2 dataset including the years 2017 and 2018. The MV2 threshold values shown here are those corresponding to the full Run 2 dataset.

b-jet efficiency	MV2 selection	Rejection Factor		
		c-jet	τ-jet	Light-flavor jet
60%	> 0.94	23	140	1200
70%	> 0.83	8.9	36	300
77%	> 0.64	4.9	15	110
85%	> 0.11	2.7	6.1	25

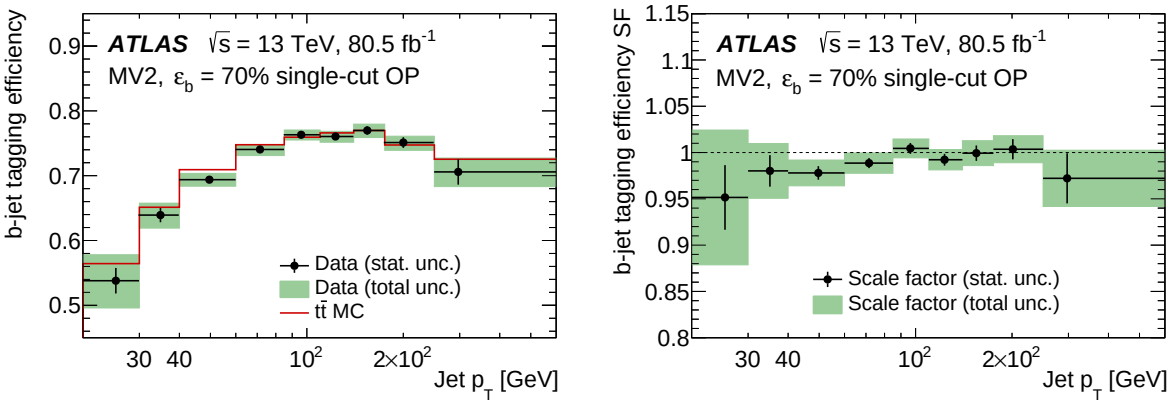

Figure 5.23: Figures taken from Ref. [146]. ***Left***: b-jet tagging efficiency as a function of jet p_T for the 70% WP of the MV2 b-jet tagging algorithm in MC (top-quark pair production, in red) and data (black points). Efficiency values below and above 70% occur, but the efficiency averaged over the full range shown is roughly 70%. ***Right***: b-jet tagging efficiency correction scale-factors for the 70% WP of the MV2 b-jet tagging algorithm as a function of jet p_T for the same samples as on the *left*.

5.4 The Missing Transverse Momentum

The colliding protons within each successful *pp* interaction are assumed to be colliding *head on*, with momentum only parallel to the beam-axis. Momentum conservation then implies that the vectorial sum of the transverse momenta of all objects originating from the primary hard-scatter vertex should be exactly zero. The missing transverse momentum, $\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}$, whose magnitude is denoted by $E_{\mathrm{T}}^{\mathrm{miss}}$, is calculated each event and its components are defined as the negative vectorial sum of the reconstructed objects associated with the event [154],

$$E_{x\,(y)}^{\mathrm{miss}} = E_{x\,(y)}^{\mathrm{miss,\,electron}} + E_{x\,(y)}^{\mathrm{miss,\,photon}} + E_{x\,(y)}^{\mathrm{miss},\,\tau} + E_{x\,(y)}^{\mathrm{miss,\,jets}} + E_{x\,(y)}^{\mathrm{miss,\,muon}} + E_{x\,(y)}^{\mathrm{miss,\,soft}}, \tag{5.10}$$

where each of the $E_{x\,(y)}^{\mathrm{miss},\,i}$ terms ($i \in \{\mathrm{electron}, \mathrm{photon}, \tau, \mathrm{jets}, \mathrm{muon}\}$) are the negative sum of the momenta in the x- (y-) direction for the respective (calibrated) objects discussed in previous sections. The $E_{x\,(y)}^{\mathrm{miss,\,soft}}$ term is track-based and is reconstructed from the transverse momentum of reconstructed tracks originating from the primary hard-scatter vertex but not associated with any of the hard objects indicated in the other terms appearing in Equation 5.10. The total $E_{\mathrm{T}}^{\mathrm{miss}}$ and its direction in azimuth is then given by,

$$E_{\mathrm{T}}^{\mathrm{miss}} = \sqrt{\left(E_{x}^{\mathrm{miss}}\right)^{2} + \left(E_{y}^{\mathrm{miss}}\right)^{2}}, \tag{5.11}$$

$$\phi^{\mathrm{miss}} = \arctan\left(E_{y}^{\mathrm{miss}} / E_{x}^{\mathrm{miss}}\right). \tag{5.12}$$

Any imbalance in the visible momenta in a given event will result in nonzero $E_{\mathrm{T}}^{\mathrm{miss}}$. Nonzero values of $E_{\mathrm{T}}^{\mathrm{miss}}$ may indicate the presence of weakly interacting particles that do not leave detectable signatures in ATLAS, such as neutrinos from leptonically decaying W bosons in the case of the SM, or may indicate the presence of additionally weakly interacting particles such as the weakly interacting dark matter candidates that many BSM physics scenarios predict. Precise measurement of the $E_{\mathrm{T}}^{\mathrm{miss}}$, then, is of the utmost importance to many searches for BSM physics. Searches for R-parity conserving Supersymmetry (Section 2.1) are commonly characterised by large values of $E_{\mathrm{T}}^{\mathrm{miss}}$ due to a multiplicity of stable weakly interacting SUSY particle in the final state. Additionally, detector mismodelling, detector noise, limited detector coverage, or miscalibration of the reconstructed objects used in the reconstruction of $\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}$ can also contribute nonzero values to $E_{\mathrm{T}}^{\mathrm{miss}}$. The measurement resolution of the $E_{\mathrm{T}}^{\mathrm{miss}}$ is also susceptible to pileup effects, and generally degrades as the levels of pileup increase.

Chapter 6

The Phase 1 New Small Wheel Upgrade Project

Every enterprise requires commitment to common goals and shared values. Without such commitment there is no enterprise; there is only a mob.

–Peter F. Drucker, *Essential Drucker*

So it goes...

–Kurt Vonnegut, *Slaughterhouse Five*

During its first two periods of operation — Runs 1 and 2 — the LHC operated at or exceeded its design goals. In Run 2, in 2018, ATLAS observed a peak of more than 60 pp interactions per bunch crossing with nominal running conditions for physics data taking and instantaneous luminosities exceeding 2×10^{34} cm^{-2} s^{-1}. These instantaneous luminosity levels were those originally envisaged as the end goal for data taking *after* LS2, during which the Phase 1 Upgrades are planned to take place both for the LHC and the large experiments, as shown in Figure 3.8. During LS2, which at the time of writing is currently taking place, the LHC plans to upgrade its injection chain to allow for delivering ever more increased peak instantaneous luminosities to the experiments. During LS3, the LHC will have its entire magnet system overhauled and will be upgraded to become the High Luminosity LHC (HL-LHC) allowing for subsequent LHC runs to deliver increased bunch densities to the experiments, thereby achieving peak luminosities far exceeding those observed thus far [155].

In the following sections, the large-scale Phase 1 Upgrade to the ATLAS detector's forward muon system will be introduced. The upgrade is a replacement of the current MDT- and CSC-based Small Wheel section of the ATLAS MS endcap, located at $z \approx 7\,\text{m}$ (Figure 3.23), and is the largest

such upgrade targeting HL-LHC data-taking conditions to take place so far amongst the LHC experiments. The detector upgrade is referred to as the New Small Wheel (NSW) and is planned for completion by the end of LS2 [156]. The NSW is planned to last for the remainder of the ATLAS detector's lifetime, and is designed with the HL-LHC data-taking conditions in mind. In Section 6.1 the reasons necessitating an upgrade to the forward muon system of ATLAS will be presented, with emphasis on how the current system will not be able to maintain high levels of performance in the Run 3 and HL-LHC era. Section 6.2 will then proceed to describe the detector technologies that make up the NSW, and how they will meet the foreseen challenges of Run 3 and beyond, as well as their layout within the ATLAS detector. Sections 6.3-6.4.3 will introduces the aspects of the NSW that the current author predominantly worked on; namely, the development of the software interface to the NSW front-end readout electronics and detectors necessary for both validation and integration of the NSW readout infrastructure.

6.1 The Need for an Upgraded Forward Muon System

The luminosity increases planned to take place during Run 3 and the HL-LHC era means that ATLAS will be subjected to exceedingly high particle rates. As mentioned previously, the peak number of interactions per bunch-crossing observed during Run 2 operations was $\langle\mu\rangle \approx 60$. Given the recent experience of the LHC delivering above expectation, the increases in collision intensity planned for Run 3 imply a peak number approaching $\langle\mu\rangle \approx 80-100$. That of the HL-LHC, nearing luminosities of $10^{35}\,\mathrm{cm}^{-2}\mathrm{s}^{-1}$, implies peaks approaching $\langle\mu\rangle \approx 200$ or more, depending on the beam configuration. While most of the detectors in the ATLAS muon system were designed with large enough safety margins to handle the foreseen rates, the inner-most forward station — the Small Wheel — will exceed its design capabilities. The current Small Wheel covers $1.0 < |\eta| < 2.7$ (Figure 3.23). At $|\eta| = \pm 2.7$, rates of $\approx 20\,\mathrm{kHz/cm}^2$ are expected, far higher than what the currently-installed MDT and CSC detectors and readout electronics can handle. Moreover, the end-cap region of ATLAS covers nearly 63% of the ATLAS muon system rapidity. High detector performance in this region, therefore, is of the utmost importance in order to achieve the long-term physics goals of ATLAS and the upgraded LHC.

The end-cap Level-1 muon trigger logic is based on track segments formed by hits in the TGC chambers in the middle muon station at $z \approx 13\,\mathrm{m}$, the first 'Big Wheel' (Figure 3.23). The p_T of muon candidates, used to define the thresholds at which the triggers accept an event, is determined by the angle of the track segment with respect to the direction pointing back to the IP. In the forward region of ATLAS there is a significant background rate composed of low energy particles — mainly protons, neutrons, and photons — generated as a result of material interactions between the Small and Big Wheels. Figure 6.1 illustrates the levels of material interactions within the ATLAS cavern that lead to this so-called 'cavern background'. These low energy particles can

produce 'fake' triggers by traversing the end-cap chambers at angles similar to that of real high-p_T muons originating from the IP. This situation is illustrated in Figure 6.2, wherein it is seen that a significant fraction of the reconstructed muons in the forward muon system cannot be matched to a trigger candidate from the ID. These 'fake' triggers amount to roughly 90% of the total muon triggers in ATLAS. With the foreseen luminosity upgrades of the LHC, the number of such triggers in the end-cap region will exceed the *total* ATLAS Level-1 trigger rate budget of 100 kHz (only 20 kHz are budgeted for Level-1 muon triggers).

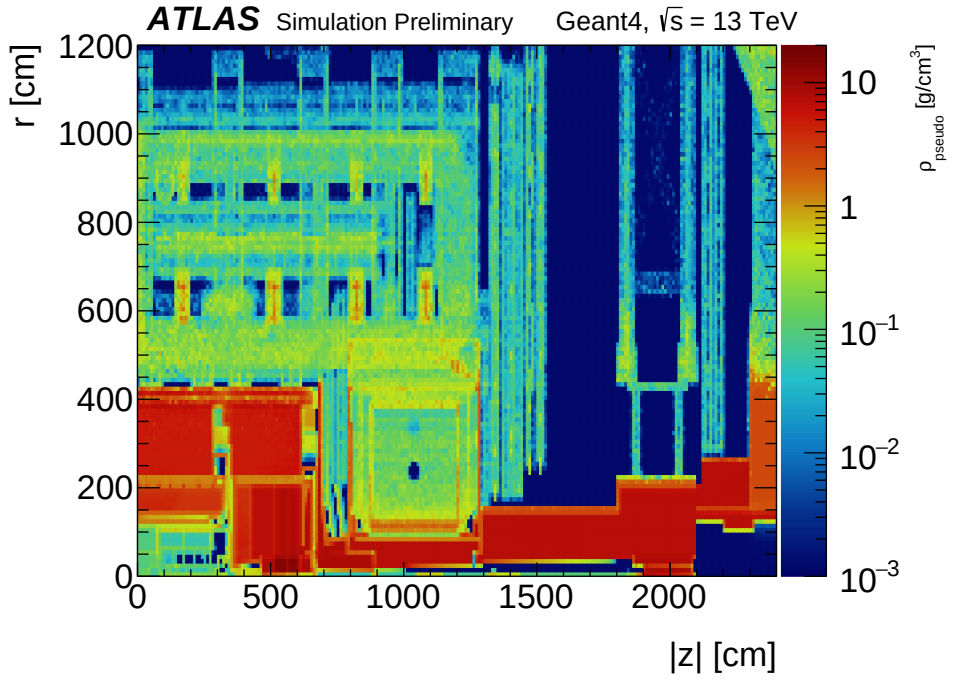

Figure 6.1: Simulation of the average material density in materials within a quadrant of the ATLAS cavern. The forward-calorimeter cryostat, beam-pipe, and shielding ($z \approx 6.8$ m) in front of the Small Wheel experience high energy deposition. The largest rates observed by the muon system are those experienced by the CSC detectors in the current Small Wheel, which can be seen in the above to lie directly behind and above regions of high material interaction. Figure taken from Ref. [157].

The NSW will take part in the Level-1 muon trigger by providing tracking information to aid the TGC-based Level-1 muon trigger logic currently in use for providing Level-1 muon trigger candidates in the forward region. The inclusion of such trigger information from the NSW will help to discriminate against 'fake' triggers due to background particles, as illustrated in Figure 6.2, and will also improve the overall muon p_T measurement in the forward muon trigger which will enable the Level-1 muon triggers to maintain low p_T-thresholds. If there were no NSW upgrade, in order to cope with the high rate of Level-1 muon triggers, the p_T thresholds of the single-muon triggers would have to be increased from a minimum of ≈ 20 GeV to thresholds well above 40 GeVin order to drop the total rate of muon triggers (inclusive of those in the barrel and end-cap) below the

20 kHz budget. Alternatively, the forward muon trigger system could be removed altogether. In this latter case, all muon triggers would be seeded by Level-1 muon candidates in the barrel only. Both of these scenarios lead to drastic reduction in physics performance. Major Higgs-physics goals of ATLAS, for example, would take large hits if either of these two non-NSW scenarios were enacted. Increasing the Level-1 muon trigger p_T-thresholds results in a loss in efficiency of more than 30% for the $Wh, h \to bb$ process — a key channel necessary for the observation and study of the Higgs couplings to b-quarks. A drop in efficiency on the order of 50% is expected for the same process if the trigger thresholds are kept low but the forward muon trigger system is removed. This is illustrated in Figure 6.3, both for the $Wh, h \to bb$ scenario discussed but also for the $h \to WW^*$ production channel, an important source of Higgs bosons in ATLAS.

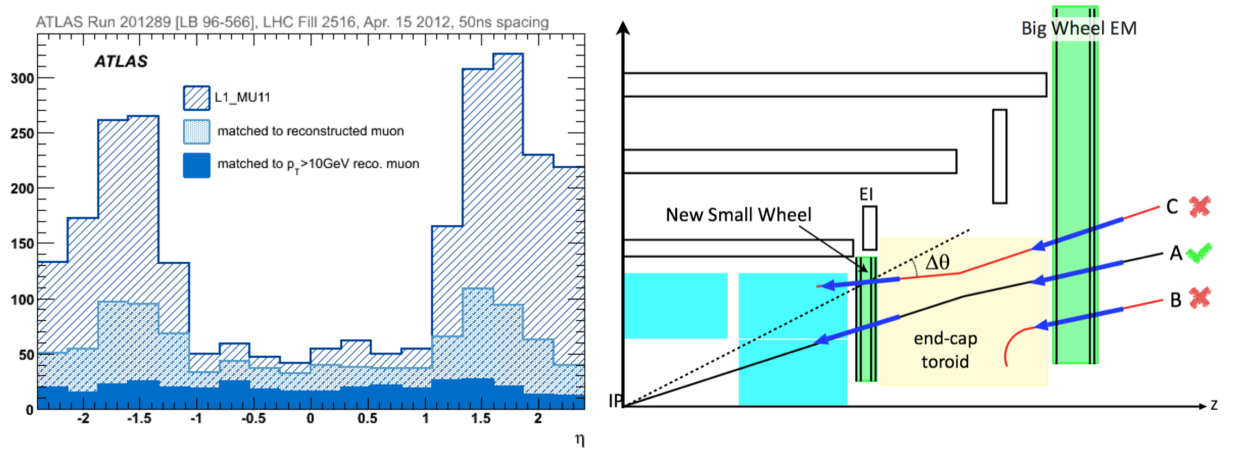

Figure 6.2: Figures taken from Ref. [156]. ***Left***: Level-1 muon trigger rates as a function of η. L1_MU11 refers to the rate of the Level-1 accept rate. After requiring that the muon trigger candidates responsible for the Level-1 muon trigger firing are matched to a real muon, the rate drops significantly, illustrating the rate of 'fake' muon triggers in the forward muon system. ***Right***: Illustration of the principle of operation of the NSW trigger system. The aim of the NSW trigger is to provide accurate pointing information (via the illustrated $\Delta\theta$ computation) within the innermost muon system to aid the trigger logic of the Big Wheel and thereby reduce the acceptance of 'fake' triggers that do not have coincident pointing track segments in both wheels. With the NSW trigger logic, only the muon trigger candidate labelled 'A' will result in a Level-1 muon trigger being fired. The track candidate 'C' ('B') fails to fire a trigger due to non-IP-pointing (missing) track-segments in the NSW, characteristic of cavern background particles.

In addition to the trigger performance of the current forward muon system being insufficient for the foreseen data-taking challenges, the tracking performance of both the MDT and CSC detectors of the current Small Wheel will also degrade with the increased luminosities. For example, the MDT detectors are characterised by long signal drift-times that are hundreds of nanoseconds long. With both these long characteristic drift times and the increased particle rates predicted to occur with the LHC upgrades, the occupancies that will be experienced by the MDT chambers in the current Small Wheel will result in significant degradation of their tracking capabilities. At the high rates

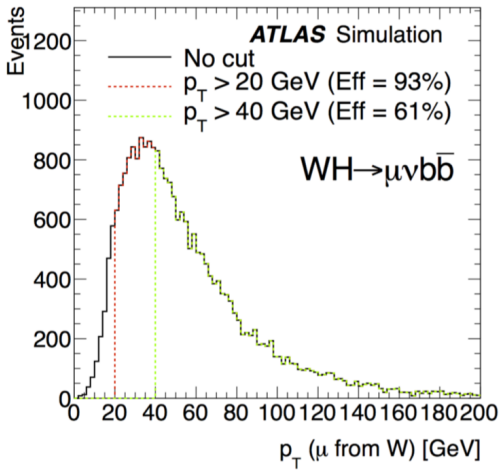

L1MU threshold (GeV)	$H \to b\bar{b}$ (%)	$H \to W^+W^-$ (%)
$p_T > 20$	93	94
$p_T > 40$	61	75
$p_T > 20$ barrel only	43	72
$p_T > 20$ with NSW	90	92

Figure 6.3: Figures taken from studies contained in Ref. [156]. ***Left***: Truth-level p_T distribution of the trigger-muon in the Wh analysis, showing the trigger thresholds and corresponding truth-level selection efficiency of those thresholds. ***Right***: Selection efficiencies for several trigger-threshold scenarios with and without the NSW for the $Wh, h \to b\bar{b}$ and $h \to WW^*$ processes.

expected in the future runs of the (HL-)LHC, the tracking efficiencies of the MDT chambers will approach 50% or worse (lower), as illustrated in Figure 6.4. This leads to significant reductions in the capabilities of the inner-most muon stations in the forward region to provide a third high-resolution muon measurement necessary for the sagitta measurement used for reconstructing muon momenta.

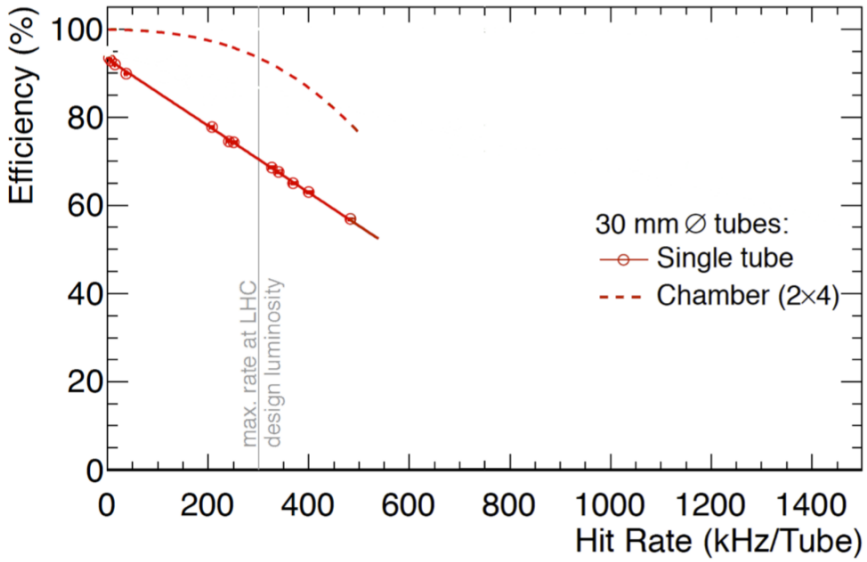

Figure 6.4: Measurements of MDT tube hit (solid line) and track-segment efficiency (dashed-line) as a function of tube hit rate. The vertical line indicating the maximum LHC rate corresponds to an instantaneous luminosity of $1 \times 10^{34}\text{cm}^{-2}\text{s}^{-1}$. Figure taken from Ref. [156].

6.2 The New Small Wheel Detector

In this section, a description of the NSW will be presented. The layout of the NSW will be described in Section 6.2.1 and then a description of the MicroMegas (MM) [158] and Small-strip Thin Gap Chamber (sTGC) detector technologies that will compose the NSW will be given in Sections 6.2.2 and 6.2.3, respectively. The MM are thought of as providing primarily high precision muon tracking while the sTGC are thought of as primarily providing trigger primitives for the forward Level-1 muon trigger system. As will be seen, though, both the MM and sTGC technologies will provide both high-quality tracking and trigger primitives. The NSW must last for the remainder of the ATLAS detector's lifetime, throughout both Run 3 and the entirety of the HL-LHC era. In the following sections, the design decisions of the NSW — related to both its layout and detector choices — will be presented in light of their ability to address the particular challenges described in the previous section.

6.2.1 Geometry and Layout

The NSW consists of 16 detector planes, separated into two 'multilayers'. Each multilayer will be composed of four detector planes of each of the two detector technologies that make up the NSW: the Small-strip Thin Gap Chamber (sTGC) and MicroMegas[1] (MM) detectors. The sTGC are based on multiwire chamber technology, operated using CO_2:n-pentane gas mixture, and the MM are a type of micropattern gaseous detector (MPGD) [159], operating using an Ar:7%CO_2 gas mixture. The detectors are arranged into sectors in azimuth, following a similar layout as the current Small Wheel with alternating large and small sectors, as illustrated in Figure 6.5. The organisation of the detector technologies in each sector is illustrated in Figure 6.6, showing the sTGC–MM–MM–sTGC layout of the detector quadruplets, with a 50 mm spacer frame separating the two halves. Each sector of the NSW will have a length nearing 5 m, giving the NSW a diameter of nearly 10 m.

The current Small Wheel muon system covers the range $1.0 < |\eta| < 2.7$, shown in Figure 3.23. The NSW will not cover the entirety of this range, however, and will only cover the range $1.3 < |\eta| < 2.7$ while the already-existing MDT chambers in the Small Wheel at $1.0 < |\eta| < 1.3$ will remain in operation after the NSW installation. The NSW will additionally extend the Level-1 muon trigger system acceptance to include the region $2.4 < |\eta| < 2.7$, which is currently not the case.

The large number of detection planes provided by the NSW adds a large degree of redundancy within each of the quadruplets of a given technology: if part of a layer, or even a complete layer, is non-functioning the remaining layers within the quadruplet will compensate and a loss in performance

[1] The name 'MicroMegas' is a loose acryonym for 'MICROMEsh GAseous Structure'.

will be minimal. There is additional redundancy provided by the fact that both MM and sTGC detector technologies will be used for both tracking and trigger functionalities: if an entire or part of a quadruplet becomes non-functioning, the tracking and/or trigger responsibilities of the lost detector will be partially covered by the presence of the still-functioning quadruplet(s)/layer(s) of the other detector technology covering the same or overlapping η range. The degree of redundancy built into the NSW detector is such that the performance goals of ATLAS with respect to the forward muon system can be met for the remainder of the experiment's lifetime of 15 years or more. With generally few opportunities for meaningful repair and maintenance access, this redundancy is a necessity for these timescales.

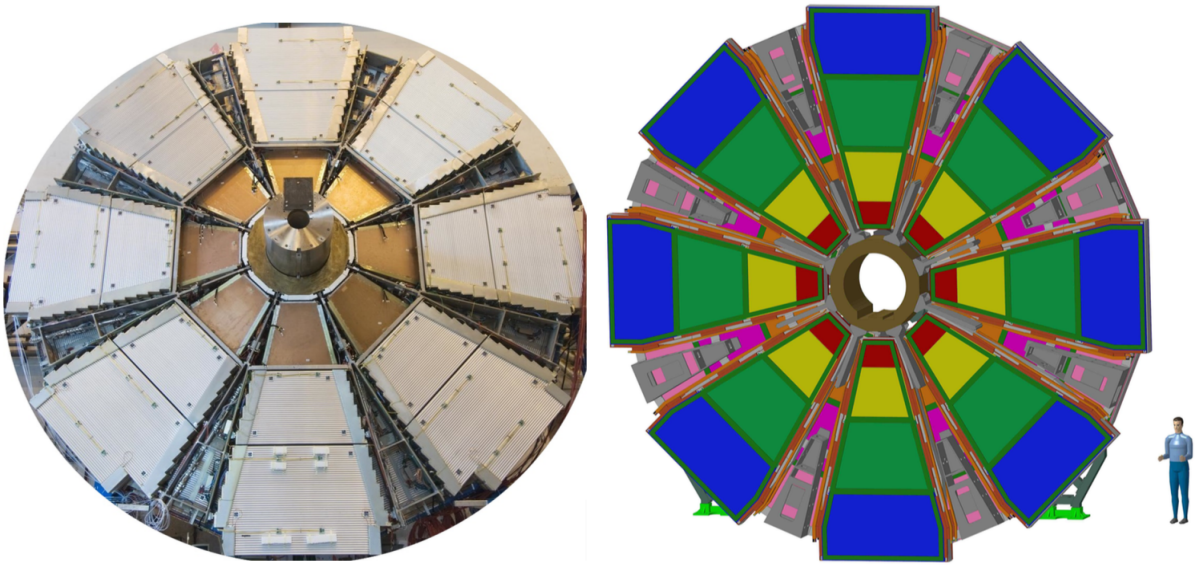

Figure 6.5: ***Left***: Current Small Wheel detector (prior to installation in ATLAS), with CSC detectors (copper color) at low radii and MDT chambers (silver color) at higher radii. ***Right***: Geometry of the NSW, with a view of the large-sector side. The gaps in azimuth between the large sectors are instrumented with detectors on the side facing into the page.

6.2.2 The Micromegas Detectors

As mentioned above, the MM detectors are primarily designed with the high-precision tracking requirements of the HL-LHC in mind, requiring a per-layer spatial hit resolution better than 100 μm. The MM detectors are characterised by readout strips with very fine segmentation and good time resolution. Due to this fine time resolution, they will be able to complement the trigger scheme based on the sTGC detector technology.

MM detectors are characterised by two asymmetric regions. The standard MM detector consists of a planar *drift* electrode, a gas gap of a few millimeters in thickness acting as a conversion and drift region, and a thin metallic mesh at $\mathcal{O}(100)\,\mu m$ above the readout electrodes. The region between

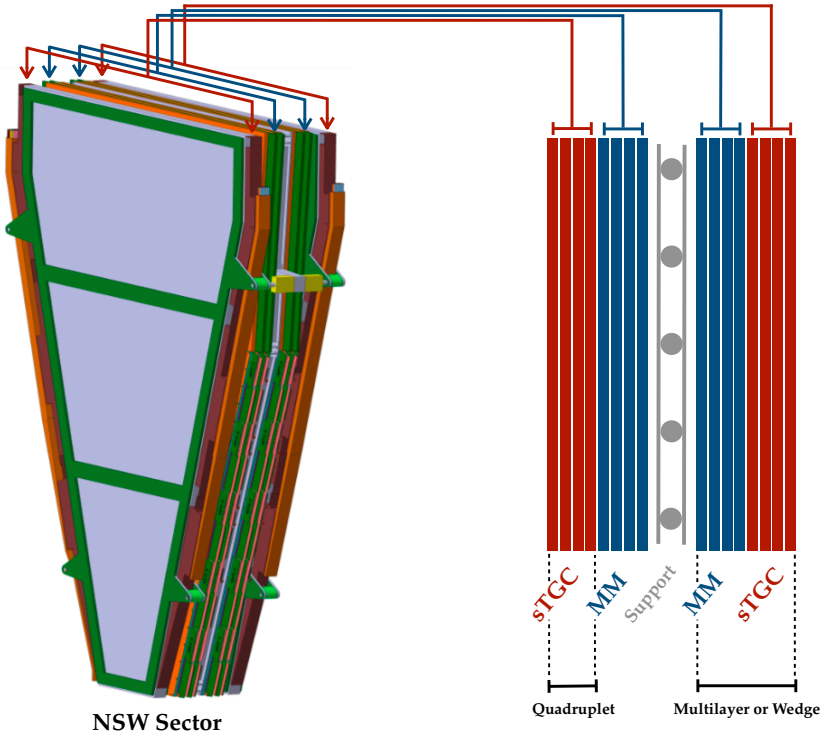

Figure 6.6: On the left is a mechanical drawing of an NSW sector, with the specific detector components illustrated on the right, composed of 16 detector planes: 8 MM layers sandwiched between 4 sTGC layers on either side. The base component of an NSW detector technology is a single detector plane, or layer, four of which are comprised in a single unit referred to as a 'quadruplet'. A single side of the NSW, composed of an MM and sTGC quadruplet, is referred to as a 'wedge' and a sector is referred to as a 'double wedge'.

the mesh and readout electrodes is the amplification region wherein gain factors on the order of 10^4 are achievable.

The electric potentials within the drift and avalanche regions are maintained at a few hundred V/cm and $\approx 50\,\mathrm{kV/cm}$, respectively. Charged particles traversing the drift space ionise the gas and the electrons, liberated in the ionisation process, drift towards the mesh at timescales on the order of tens of nanoseconds. The electron avalanche takes place in the amplification region in about a nanosecond, resulting in a fast current pulse on the readout electrode.

The MM detectors in the NSW are *resistive-strip* MM detectors, characterised by an insulating layer over the readout electrodes. The insulating layer acts to protect the sensitive readout electrodes from sparking events that reduce the detector performance over time and that are expected to happen frequently given the very large gas amplification. Resistive strips on top of the insulating layer collect the avalanche electrons and induce signals on copper readout strips embedded beneath the insulating layer. The geometry (strip pitch and width) of the copper readout strips need not

necessarily be the same as that of the resistive strips. Resistive-strip MM detectors are able to sustain higher amplification and particle rates, a necessary characteristic for targeting the high-luminosities foreseen at the HL-LHC. The design and principle of operation of the resistive-strip MM technology is shown in Figure 6.7, in which incident MIPs traversing perpendicular and at an angle relative to the detector plane are illustrated.

Each layer of an MM detector in the NSW is composed of 1024×8 parallel readout strips, giving a highly granular spatial readout allowing for high resolution position measurements to be made. Although all readout strips of a given MM layer are parallel, the MM quadruplets will be capable of two-dimensional readout and will provide both r and ϕ coordinate information of traversing particles. The two dimensional readout is achieved through the use of a small-angle stereo readout, in which MM layers are tilted by fixed relative angles of ± 1.5 degrees. This means of a two-dimensional readout is unlike that achieved by the CSC detectors in the current Small Wheel, in which the readout plane consists of perpendicular wires and strips providing the two dimensional readout information. This perpendicular readout is susceptible to high rates of so-called *ghost hits*, illustrated in the left side of Figure 6.8, and leads to high levels of track-building ambiguities. In the high particle rates expected at the HL-LHC, the multiplicities of these ghost hits would lead to an unacceptable degradation in tracking performance. The combination of the small-angle stereo readout and high strip-multiplicities on each MM layer will enable the MM detectors to sustain these high particle rates. In the right side of Figure 6.8, the layout of the MM layers in the NSW as regards their readout strip orientation (i.e. tilted or not) is illustrated.

Typical definitions of hit locations within strip-based detectors are based on the centroid method, using the charge-weighted strip position to define the spatial location of the hit. This method works optimally for particles incident at angles perpendicular (zero inclination) to the detector readout plane. The planar NSW detectors, however, will be subjected to particles originating from the IP that follow inclined tracks. In the MM detectors, the centroid method provides worsening spatial resolution as the incident angles of these inclined tracks increase due to the charges from the multiple ionisation events of a single incident particle being spread across multiple readout strips (Figure 6.7). To overcome this, the MM hit reconstruction will use hit timing information to use the 5 mm conversion gap as a time-projection chamber (TPC). In the NSW, this two-dimensional TPC hit reconstruction is referred to as the micro-TPC ('μ-TPC') reconstruction method. The use of the timing information of each of the hits recorded by the readout strips allows for the position above the readout plane to be reconstructed, thereby allowing for a mini-track to be reconstructed that follows the ionisation history of a single incident particle. From this mini-track, a well-defined metric for defining the hit location for inclined tracks on each MM layer can be defined, as illustrated on the left side of Figure 6.9. In the NSW, the MM hit location will be determined through the combined use of both the charge-centroid and μ-TPC methods, allowing for sufficient spatial resolution spanning the relevant incident angles, as illustrated on the right side of Figure 6.9.

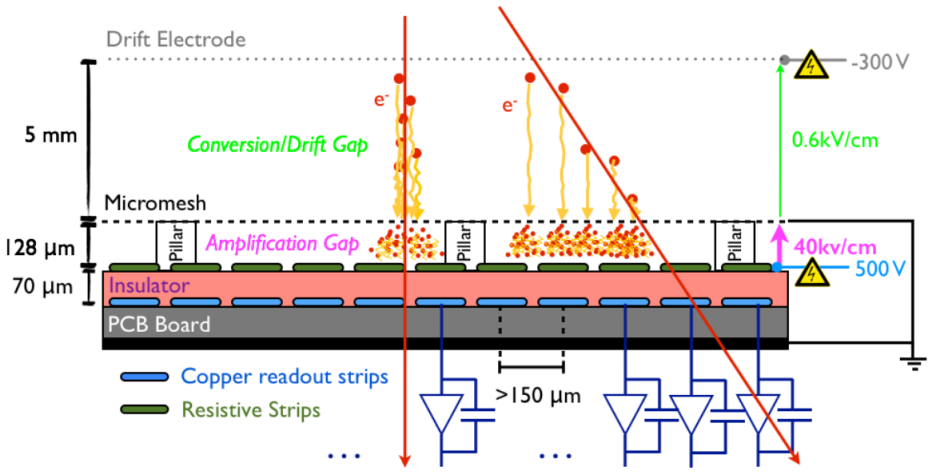

Figure 6.7: Illustration of the operating principle of a resistive-strip MM detector chamber. Figure taken from Ref. [156].

6.2.3 The Small-strip Thin Gap Chamber Detectors

The sTGC are those detectors stated as being primarily used for producing trigger primitives to be used by the Level-1 muon trigger system. As triggering detectors, they have fast signal formation and readout times owing to their characteristic high operating electric fields. This enables them to provide accurate bunch crossing identification (i.e. assigning hits to a specific bunch crossing). The average drift time of ionisation electrons within the sTGC chambers is less than 25 ns, with the earliest cluster arrival times well below 10 ns. Figure 6.10 shows the distribution of the measured drift times in an sTGC, with 95% of the measured and simulated data being inside a 25 ns (the LHC bunch crossing frequency) time window. The sTCG detectors also provide good online angular resolution — better than 1 mrad — for the track segments reconstructed across the layers of an sTGC quadruplet, meaning fairly good spatial resolution. The high quality angular resolution allows for good p_T determination based on the trigger primitives provided to the Level-1 trigger. Offline spatial reconstruction for track segments provided by the sTGC detectors will also be good, matching the sub-100 μm threshold for the HL-LHC precision tracking requirements in the forward region of the muon system.

The basic sTGC detector is shown in Figure 6.11. It consists of a grid of 50 μm gold-plated tungsten wires with 1.8 mm pitch sandwiched between two cathode planes at a distance of 1.4 mm. On one plane there are the readout pads and on the other the strips. The strips have a 3.2 mm pitch, which is smaller than that of the current TGC detectors used in the current forward muon system (Figure 3.23), which is the origin of the name of the sTGC detectors. The wires and strips are perpendicular, with the wires stretching along r and the strips along ϕ.

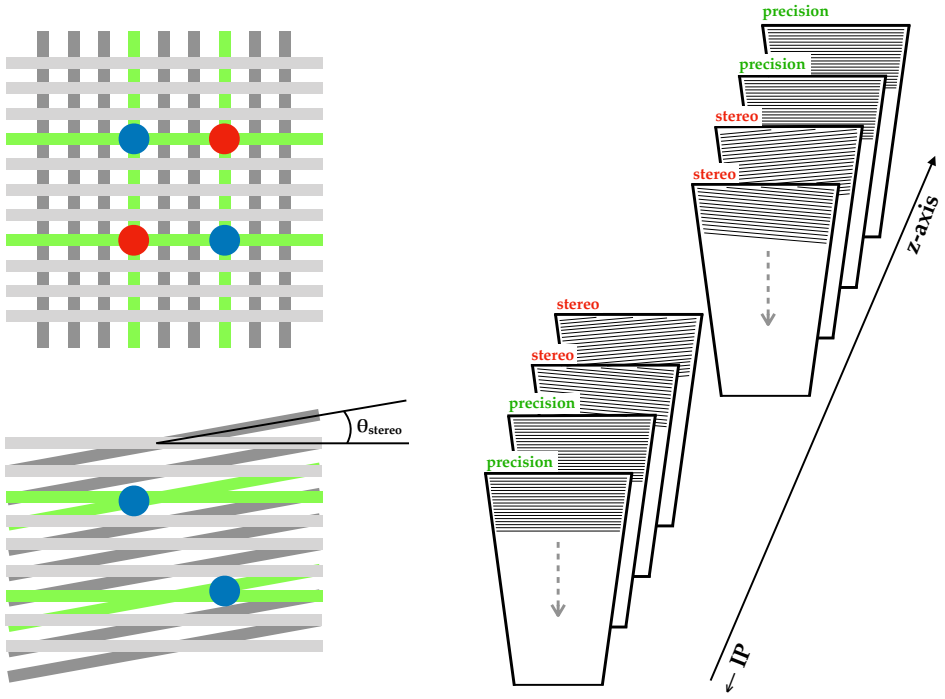

Figure 6.8: ***Left***: Illustration of the small-angle stereo two-dimensional readout principle. The upper panel shows two layers with readout strips perpendicular to each other. The lower panel shows two layers with readout strips tilted at a small angle, θ_{stereo}, relative to one another. In both cases, true detector hits are indicated by the blue dots. The readout strips registering the hits are indicated in green. In the case with perpendicular readout, there are two ghost hits indicated by the red dots. For perpendicular readout strips, there will generally be $N_{\text{ghost}} = (N^2 - N)$ ghost hits for N real hits when the readout strips on the two layers have equal pitches as in the case of the MM detectors in the NSW. On the lower panel, the green readout strips registering the two hits do not cross as a result of the small stereo angle and there are no ghost hits. ***Right***: Illustration of the layout of the 'precision' and 'stereo' readout layers of the two MM quadruplets in a given NSW sector. The precision layers, so-called since their readout strips are along ϕ and measure directly the coordinate in the bending plane of passing muons relevant for p_T determination, are the outer two layers of each quadruplet relative to the spacer frame. The stereo layers are those with the readout strips tilted at $\pm 1.5^\circ$ relative to the precision layers and are the two layers in each quadruplet nearest the central spacer frame. One stereo layer in each quadruplet has stereo angle $+1.5^\circ$ and the other -1.5°, making for $\Delta\theta = 3^\circ$ between the two stereo layers in each quadruplet.

The pads are used to form coincidences between the layers of the sTGC quadruplets in order to form projective trigger primitives corresponding to muon tracks pointing roughly back to the IP. There are ≈ 1300 pads per sTGC layer. The collected charge from all components of the sTGC — the pads, strips, and wires — associated with a traversing particle are read out and used for the subsequent offline precision muon-track reconstruction. Two-dimensional readout is achieved by the r and ϕ information provided by the strips and wires, respectively. To help cope with the foreseen

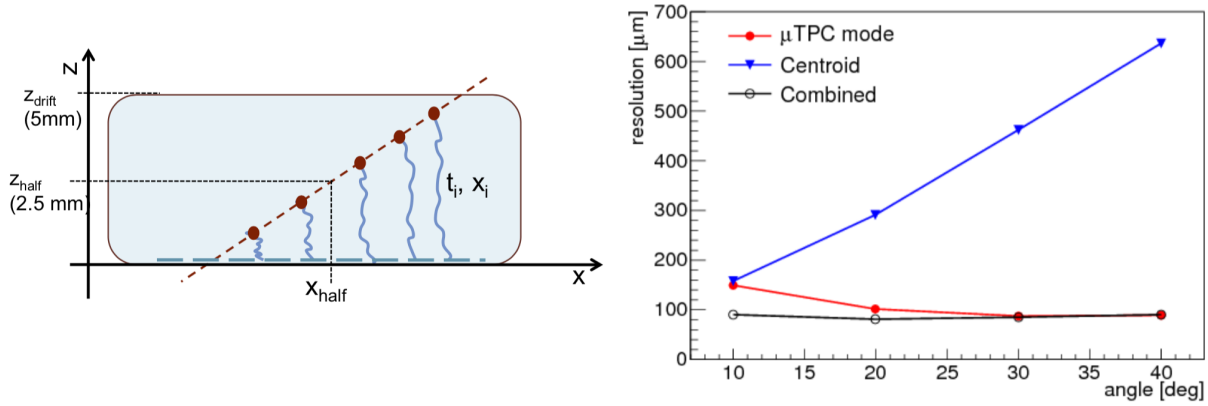

Figure 6.9: Figures taken from Ref. [156]. ***Left***: Illustration of the 'x_{half}' method of defining a hit position for an inclined track: the hit position is defined as the position on the readout plane corresponding to the half-height location along the reconstructed tracklet in the MM conversion gap. ***Right***: Expected MM spatial resolution with charge centroid method (blue triangles), μ-TPC method (filled red circles), and the combination of the two (black open circles) as a function of particle incident angle. The combination of the two methods achieves the greater-than 100 μm single-layer spatial resolution required for HL-LHC operation for the expected range of incident angles.

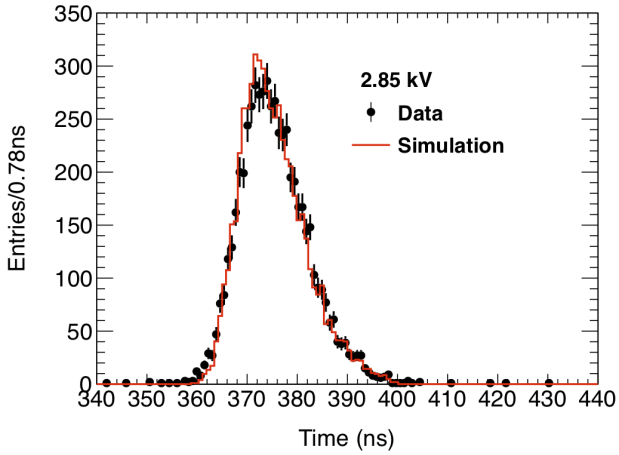

Figure 6.10: Due to the very thin gap and the intense electric field applied, $\approx 95\,\%$ of the measured drift times of passing MIPs are confined within a window of 25 ns, showing that the sTGC are capable of reliable bunch crossing assignment. Figure taken from Ref. [156].

high particle rates, and to reduce the types of tracking ambiguities described in Section 6.2.2 due to combinatorial hits, only those wires and strips covering the projective area of pads registering the passage of the particle are read out. Additionally, wires are grouped together (i.e. multiple wires correspond to a single readout channel) as only a rough measurement of the azimuthal coordinate — not relevant for p_T determination in the trigger— is needed. With the charge information collected from the three sub-components, the sTGC is able to achieve very good spatial resolutions

approaching $\approx 80\,\mu m$ ($\approx 100\,\mu m$) per layer for perpendicular (inclined) tracks.

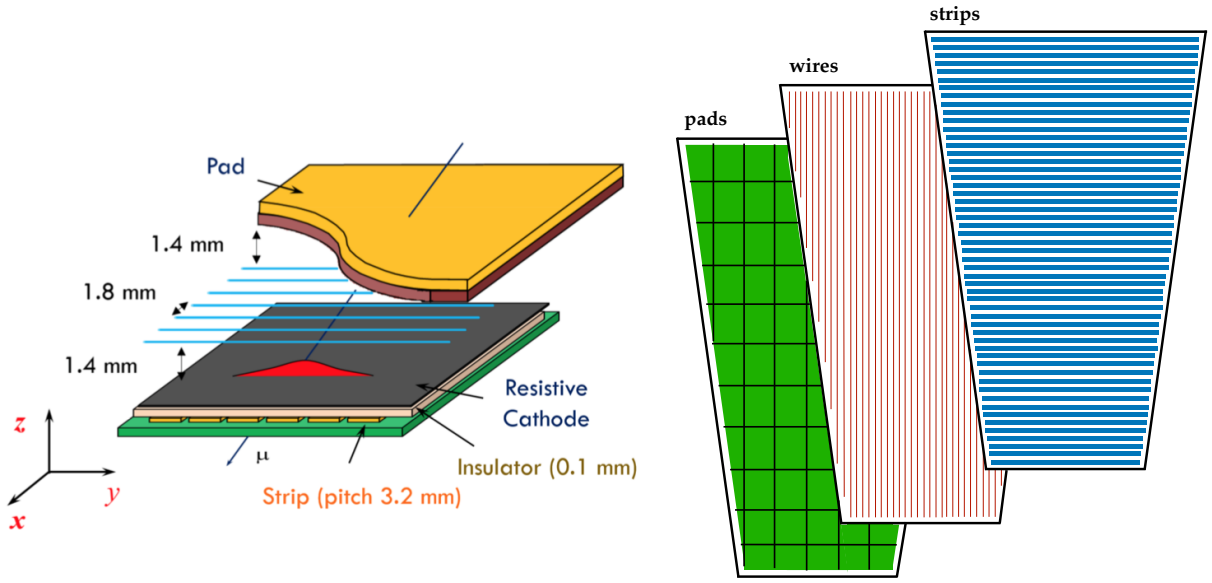

Figure 6.11: ***Left***: Graphical representation of the internal structure of an sTGC detector. The passage of a traversing MIP (black line) induces signals on the anode wires, readout pads, and strips behind the resistive cathode planes as a result of the drifting of ionisation charges and their subsequent multiplication (in red). Figure taken from Ref. [160]. ***Right***: Illustration of the geometry of the sTGC pads, wires, and strips and their relative layout in a single sTGC layer. The strips (blue) are along ϕ and provide the measurement of the precision momentum coordinate in the particle bending plane. The wires are along r and provide measurements sensitive to ϕ. The pads provide information necessary for building trigger coincidences between the layers of an sTGC quadruplet and determine which subset of wires and strips will subsequently be readout upon a Level-1 trigger accept. The drawing is not to scale and does not illustrate the actual segmentation of the readout components in a given layer, only their relative orientation.

6.3 NSW Readout Electronics and Detector Instrumentation

To a very large extent, high-performance detectors are realised and made possible only by their interfacing to equally high-performance readout electronics; that is, high quality electronics enable high quality detectors. In this section, therefore, an introduction to the frontend electronics[2] relevant to the NSW will be given. There is a large set of both frontend and backend electronics specifically designed and constructed for the NSW [161]. Here, though, we will focus primarily on

[2]'Frontend' electronics are those housed directly on the detectors themselves and are responsible for reading out the detector signals and potentially many other responsibilities, such as calibration, configuration, and detector controls (temperature monitoring, etc...). 'Backend' electronics refer to those not necessarily located in the experimental cavern housing the detectors but are in the nearby service areas or on the surface (c.f. Figure 3.5) and are dedicated, for example, to the implementation of detector trigger logic and data decoding and event building (i.e. gathering all data from detector hits associated with a given event).

the frontend electronics. The NSW frontend electronics revolve around the operation of a family of rather complex application specific integrated circuits (ASICs), each targetting a specific (set of) purpose(s) related to data-acquisition or constructing trigger primitives:

VMM [162, 163, 164] Frontend readout ASIC used for precision and fast trigger data signal collection in both MM and sTGC detectors[3]

ROC [156, 161] 'ReadOut Controller' ASIC, responsible for buffering and aggregating precision data from the VMM

TDS [165] 'Trigger Data Serializer' ASIC, processes trigger signals from VMMs on the sTGC detectors and prepares trigger primitives for the NSW Level-1 muon trigger logic

ART [156, 166] 'Address in Real Time' ASIC, processes trigger signals from VMMs on the MM detectors and prepares trigger primitives for the NSW Level-1 muon trigger logic

The VMM ASIC is the primary ASIC of the NSW: all other ASICs listed above take as input the outputs of the VMM and, in this respect, are secondary. The work of the current author on the NSW frontend electronics is based almost exclusively on the characterisation and validation of the VMM ASIC. For these reasons, focus will be given on describing the VMM in Section 6.3.1. For the interested reader, further information on the ROC, TDS, and ART ASICs is provided in the references above. The interface of the NSW frontend electronics to the associated detectors is dictated for the most part by the geometry, type, and number of the detectors' readout elements (MM: strips, sTGC: strips, wires, and pads) and foreseen space limitations in the ATLAS detector. The work in the present thesis concerns the frontend boards related to the MM detectors. This is mainly due to the fact that the VMM was initially conceptualised as a frontend ASIC specialised for MPGD detectors and its initial prototypes were studied under this context by the RD51 collaboration at CERN which is devoted to the development of MPGD technologies, such as the MM detectors.[4] In Section 6.3.2, then, a description of the relevant front-end boards used in the study, validation, and integration of the VMM ASIC, as well as those that will be used in the MM detectors of the NSW, will be given. The analogous VMM-based frontend boards for the sTGC detectors are objectively more complicated than those of the MM detectors for purely technical and uninteresting reasons. The discussion below, however, is detector-technology agnostic and the sTGC frontend description can be found elsewhere [156].

[3]The acronym for which 'VMM' stands is not informative as in the case of the other NSW ASICs. The 'V' in VMM refers to the first initial of Venetios Polychronakos who is an ATLAS physicist involved in the VMM design, among many other things. The 'MM' in VMM refers to the MM detectors. The reason for this acronym is the tradition of the main designer of the VMM, Gianluigi De Geronimo, to name his projects based on the initial person requesting the project to be designed and that project's purpose.

[4]For more information on the RD51 collaboration, visit their homepage: `http://rd51-public.web.cern.ch/rd51-public/`.

6.3.1 The VMM ASIC

The VMM is a custom ASIC that can be used in a variety of charge-interpolated tracking detectors. There have been several iterations of the VMM over the years, starting with the VMM1 [162] and moving up to the VMM3 [163, 164]. An illustration of the evolution of complexity of the VMM ASIC is shown in Figure 6.12. The VMM1 was a purely analog readout ASIC considered to be a prototype of the initial VMM functionalities. The VMM2 was an extensive upgrade of the VMM1, adding the necessary digital logic and functional blocks required for the NSW detectors. The VMM3 is the final version of the ASIC[5] providing enhanced functionality required for data taking in ATLAS as well as addressing bugs found in the operation of the VMM2. VMM3 is the version that will be used in the NSW once installed in ATLAS. The VMM implements all logic using triple modular redundancy (TMR) to protect itself from single event upsets (SEUs) that are expected to be quite common in the harsh radiation environments in which the NSW will be situated.

The VMM is composed of 64 discrete frontend channels, each to be connected to a detector readout element. A block diagram illustrating the functional blocks of the VMM3 and its channels is shown in Figure 6.13. Each channel has a dedicated charge amplifier ('CA' in Figure 6.13) and signal shaping functionality, threshold discriminator, test-pulse injection capacitor with adjustable amplitude provided by a 10-bit digital-to-analog converter (DAC), and precise signal amplitude and timing measurements digitally readout via internal per-channel analog-to-digital converters (ADCs). The amplitude measurement of the input pulse is digitised by a 10-bit ADC ('PDO', for 'peak detector output', in Figure 6.13) and the timing estimation of the signal pulse's peak is digitised by an 8-bit ADC ('TDO', for 'time detector output', in Figure 6.13). The digitised output of these amplitude and timing measurements are calculated by these internal ADCs within 200 ns and stored in internal buffers of the VMM. The buffered data is stored until selected for readout by another ASIC housed on the same frontend board as the VMM, the ROC, based on the buffered data's associated bunch crossing. The PDO and TDO data are those used in the precision hit reconstruction in the detectors. The hit timing information provided by the TDO is critical only for the MM detectors which rely primarily on the μ-TPC hit reconstruction method, as illustrated in Figure 6.9. For the sTGC, the accurate timing information is not required for precision tracking and, to increase readout bandwidth, the timing information will likely not be serialised in the data output from the sTGC frontend electronics.

In addition to the precision data useful for reconstructing high-level muon tracks provided by the PDO and TDO measurements, the VMM can also operate at a faster mode used for providing data needed for the building of trigger primitives. In the trigger mode referred to as the 'Address-in-Real-Time' (ART) mode, the VMM outputs the *first* channel address (e.g. MM strip number) on which

[5] In actuality, there have been two versions of the VMM3: the VMM3 and the 'VMM3a'. The VMM3a is the true final version that will be used in the NSW and addresses issues found during the testing and validation of the first round of the VMM3 production

an above-threshold peak was detected. An alternative trigger mode relies on the VMM's ability to perform a fast digitisation of the signal pulse amplitude using a 6-bit ADC, providing a coarse but rapid measurement of the signal amplitude. Additionally, the VMM can output fast timing signals, such as a flag indicating a signal over threshold ('Time-Over-Threshold', or TOT as in Figure 6.13). The ART trigger scheme is used for the MM trigger primitives and the combination of the coarse 6-bit amplitude and fast-timing signals is used for building the sTGC trigger primitives.

One of the powers of the VMM is its highly configurable nature. This, of course, makes the ASIC quite complex in terms of its design but is advantageous in terms of its range usefulness. Some of the main configurable items are the peaking (integration) time of the signal shaper (25, 50, 100, and 200 ns), adjustable thresholds per-VMM configured by an on-VMM 10-bit (DAC), adjustable channel gains (0.5, 1, 3, 4.5, 6, 9, 12, and 16 mV/fC), time-to-amplitude conversion (TAC) ramp time (60, 100, 350, and 650 ns) relevant for the timing measurements, and channel threshold trimmers provided by 5-bit DACs that adjust each channel's individual threshold around the globally-configured (per-VMM) threshold. The VMM input can also be configured to handle either positive or negative input signals. This latter fact is necessary for the NSW, in which the sTGC and MM detectors will induce signals of opposite polarity on their readout elements.[6]

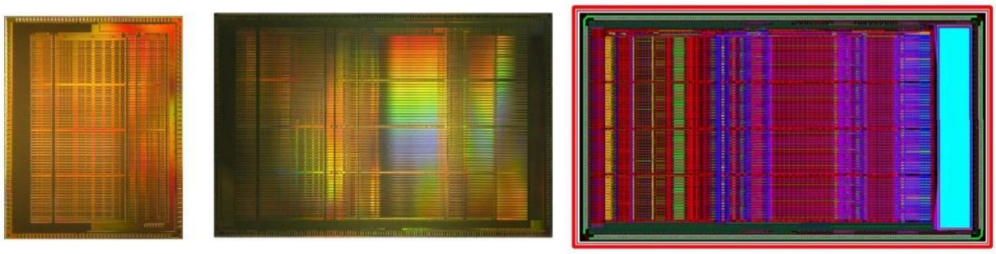

Figure 6.12: Evolution of the VMM ASIC. Shown are the silicon dies or routing of the the VMM1 (***left***), to the VMM2 (***middle***), and the VMM3 (***right***), the final version that will be used in the NSW. The VMM1 has an area of 50 mm^2 with $\approx$ 500 k MOSFETs, the VMM2 area is 115 mm^2 with $>$ 5 M MOSFETs, and the VMM3 area is 130 mm^2 with $>$ 6 M MOSFETs.[7]

6.3.2 Frontend-electronics Boards for the NSW

There have been two classes of frontend boards encountered throughout the work on the NSW electronics represented in this thesis. The first class relates to a type of frontend board that is not specific to the NSW but which houses one (or several) VMM ASICs. The second class relates to prototype frontend boards following the design to be used in the NSW to readout the MM detectors, the so-called MMFE8 frontend board. Figure 6.14 provides pictures of these two classes

[6] The MM strips and sTGC wires produce signals of opposite polarity as the sTGC strips and pads.

[7] 'MOSFET' stands for metal-oxide-semiconductor field-effect transistor', the most widely used transistor in digital and analog electronics.

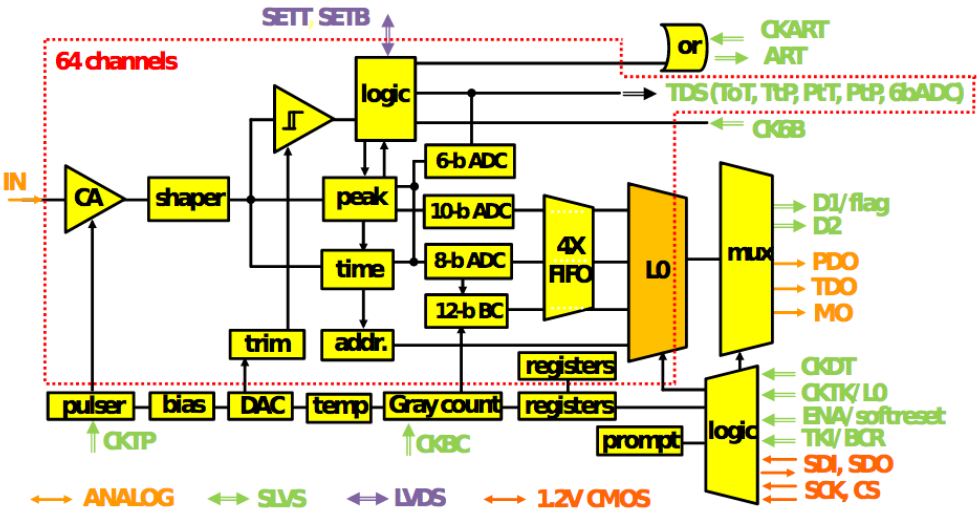

Figure 6.13: Architecture of the VMM3. The items contained within the dotted red line are repeated for each of the 64 channels of the VMM. Figure taken from Ref. [164].

of frontend boards with relevant parts indicated.

The general purpose frontend boards, such as the GPVMM in Figure 6.14, are those used primarily for the direct study and characterisation of the VMM ASIC. The prototype MMFE8 frontend boards allow for a more realistic interface to the MM detectors to be studied as well as for the overall design and layout of the hardware components on the boards themselves to be studied and optimised. One such optimisation was the placement of the DC-to-DC converters used for the on-board power distribution.[8] DC-to-DC converters are potential sources of electronic noise and stray electromagnetic fields as a result of their high switching frequencies. When placed directly on the frontend boards, their proximity to the sensitive readout elements means that the matter of their placement and shielding needs to be handled carefully.

The types of boards discussed here have allowed for the study of the VMM ASIC both on- and off-detector, where in the latter case prototype MM detectors were used. Over the timespan of the present thesis, many iterations of each type of frontend board were encountered, with subsequent iterations adapting to the evolution of the VMM, for example (Figure 6.12). In addition to housing the VMM ASIC, the frontend boards house a Xilinx field-programmable gate array (FPGA). One of the main purposes of the FPGA is to aggregate the data being output from the VMM and prepare it for being sent off board via standard network protocols over Ethernet to the backend data-acquisition system. Additionally, the FPGA responds to commands and requests submitted by users via the software interface. Such commands may have as endpoint the FPGA itself or they

[8] In the final MMFE8 prototype frontend boards, and in those to be used in the NSW, the DC-to-DC converters are packaged in the so-called FEAST ASIC developed by CERN. Further information at `https://project-dcdc.web.cern.ch/project-dcdc/Default.html`

may require the FPGA to forward them to the VMM ASIC. These concepts are described further in Section 6.4.1.

The GPVMM board has a simple connector designed in such a way that lends itself to be easily adapted to many types of detectors in labs and teststands. The MMFE8 board, both the prototypes and those to be used in the NSW, use a ZEBRA[9] elastomeric connector that allows for a seamless interconnect between the readout strips on the MM detector layers and the VMM channel inputs. The principle of the ZEBRA connector is illustrated in Figure 6.15.

The NSW frontend boards are interfaced directly to the detectors. Both the GPVMM-type and MMFE8 frontend boards have been used on small prototype MM detectors, with active area of $10 \times 10\,\mathrm{cm}^2$. A picture of an MMFE8 connected to such a detector is shown in Figure 6.16. On the full-scale MM detectors to be used in the NSW, the frontend boards are situated along the edge of each detector layer in a given multilayer. A drawing of a complete MM quadruplet with frontend electronics boards is shown in Figure 6.17.

[9]More information on ZEBRA-type connectors can be found online at `https://en.wikipedia.org/wiki/Elastomeric_connector`

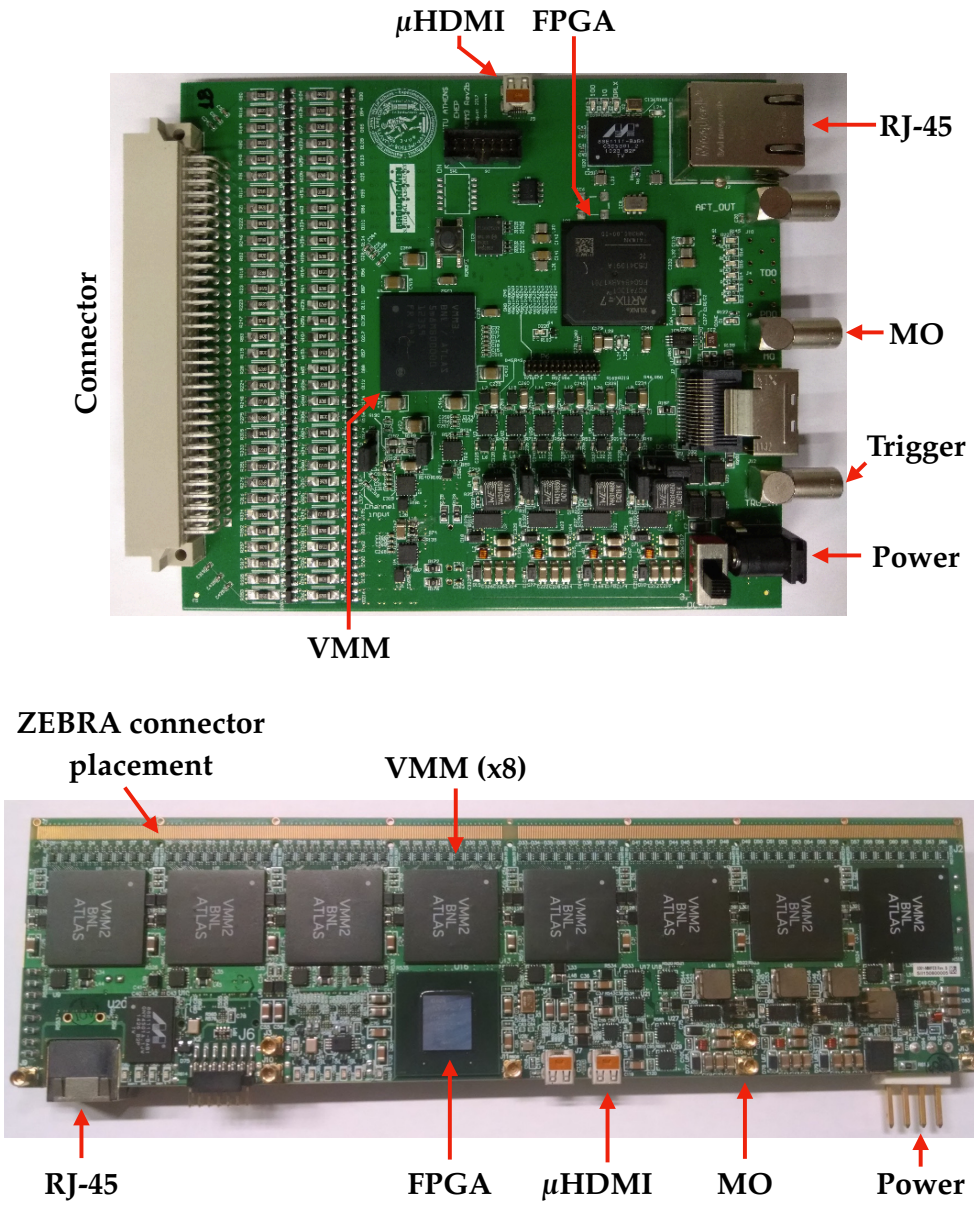

Figure 6.14: Pictures of two VMM-based frontend boards. The RJ-45 connectors provide network I/O via standard Ethernet cables. The μHDMI connectors provide inputs for additional signalling purposes, such as external trigger signals. 'MO' refers to the multiplexed monitoring output of the VMM (see Figure 6.13), which samples the analog signals internal to the VMM, such as the shaped input signals, prior to their digitisation. ***Top***: General purpose VMM (GPVMM) board housing a single VMM ASIC. The MO and trigger I/O are provided by LEMO connectors. ***Bottom***: Prototype MMFE8 board with its 8 VMM ASICs. The ZEBRA connector used for the MM detectors is described in the text. On this board the MO output is accessible primarily by the use of an oscilloscope probe.

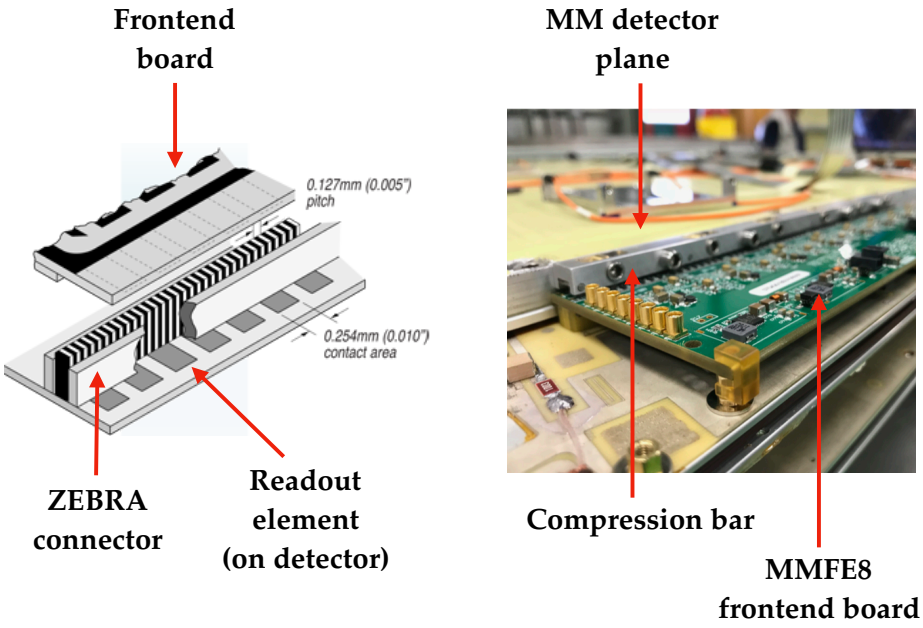

Figure 6.15: ***Left***: Illustration of the ZEBRA connector concept. The ZEBRA connector is composed of a high density of alternating conductive (white) and non-conductive (black) layers that make contact with the detector readout elements on bottom and frontend board sensing elements on top. ***Right***: Picture of the ZEBRA connector being used with an MMFE8 frontend board on a full-scale MM detector. When interfaced to the MM detector, the MMFE8 frontend board is situated such that the VMMs are downward facing and are not therefore visible as pictured. The connection of the ZEBRA connector is realised via the tightening of a compression bar mounted on the MM detector chamber. The resulting downward pressure forces the ZEBRA connector to be securely sandwiched between the readout elements of the MM detector and VMM channel inputs located on the frontend board, bringing them in direct electrical contact.

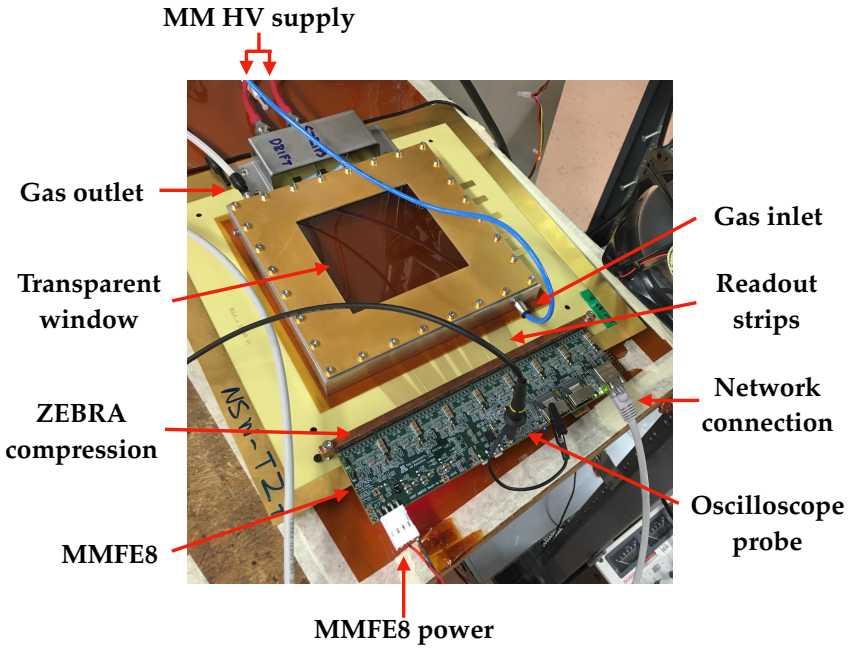

Figure 6.16: An MMFE8 interfaced to a small $10 \times 10\,\mathrm{cm}^2$ prototype MM detector chamber in the RD51 lab at CERN. The transparent window allows for visual inspection of the drift mesh and reduction in material through which incident particles pass. The red cables at the top provide the high voltage (HV) power supply to the MM drift mesh and the resistive strips. The routing of the detector strips to the MMFE8 can be seen on the chamber backplane.

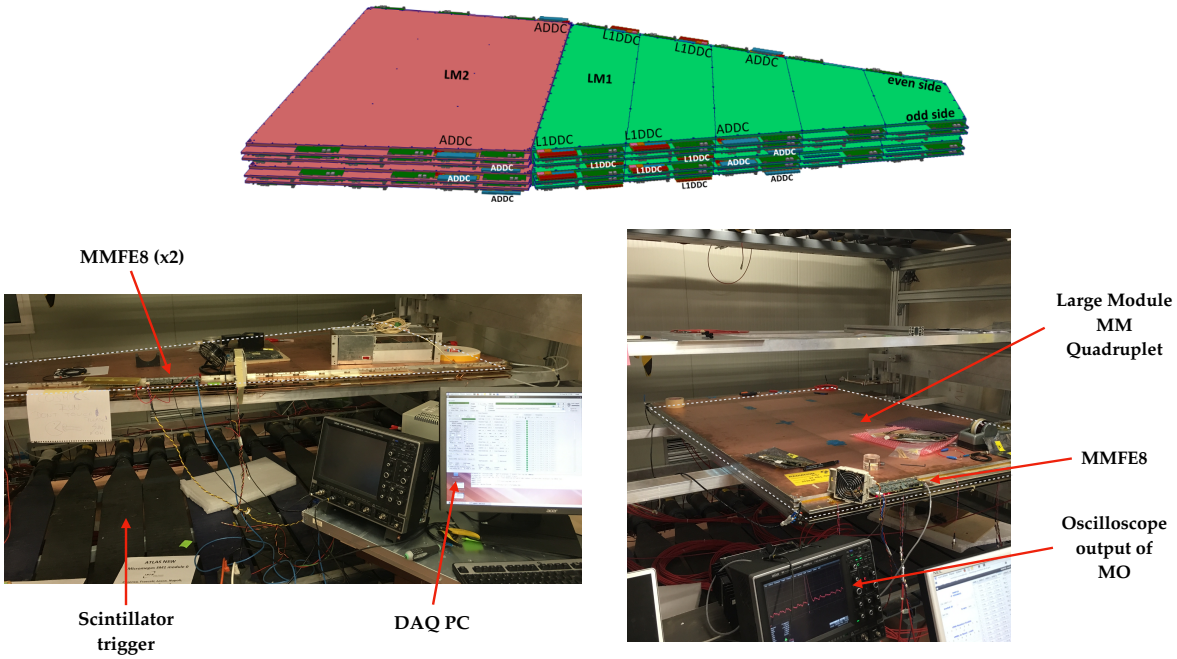

Figure 6.17: ***Top***: Mechanical drawing of a MM large sector quadruplet. MM quadruplets are segmented into two modules, here indicated as LM1 and LM2. Indicated along the edge of the quadruplet are the locations of the frontend boards interfaced directly with the detector. The green frontend boards are the MMFE8s housing the VMM ASICs. Information about the L1DDC and ADDC boards can be found in Ref. [161]. ***Bottom, left***: The first small-sector MM quadruplet module (only the lower half, analogous to the LM1 in the Top drawing), 'SM0', on the cosmic-ray test-stand located in the RD51 lab at CERN. The SM0 module is partially instrumented, with only a few MMFE8 boards interfaced as indicated. The outline of the SM0 module is given by the grey dashed line. The lower half of the scintillator-trigger system is seen beneath the SM0 module. ***Bottom, right***: The first LM2-type MM quadruplet module on the cosmic-ray test-stand located in the RD51 lab at CERN. The outline of the LM2 module is given by the grey dashed line. A shaped VMM signal pulse can be seen on the oscilloscope.

6.4 Development of Configuration, Data-acquisition, and Calibration Software for the Validation of the NSW Frontend

To facilitate the testing and readout of the VMM ASIC, in view of its evolution and the approaching installation of the NSW, a complete configuration and data-acquisition (DAQ) system has been built. The system, referred to as the 'VMM Readout System' (VRS), consists of a flexible firmware and software infrastructure that allows for interfacing to frontend boards housing the VMM ASIC, such as those shown in Figure 6.14.

A minimal instantiation of the VRS system is illustrated in Figure 6.18 in which the DAQ PC, hosting the VRS software, communicates directly to a set of frontend boards via the network.[10] The frontend boards need not be interfaced to a detector, as signals can be injected via the test-pulse injection capacitors on the VMM channels. The firmware, loaded onto the FPGA on each frontend board, implements the necessary logic for handling the network communication between the software and the VMM ASIC. The types of communication between the software will be described in subsequent sections, but primarily consist of configuration commands being forwarded to either the FPGA or the VMM ASIC. The firmware handles the logic necessary for packaging the readout data being output by the VMMs on the associated frontend boards. Once packaged, it then serialises the data over the network back to the DAQ PC where it can be monitored and promptly stored on disk for later analysis. In this minimal setup, the clock and trigger signals driving the VMM readout are generated independently on each frontend board by the on-board FPGA. Communication is made between the DAQ PC and a frontend board via a unique network (IP) address associated with each frontend board that is exposed to the network by the firmware.

The VRS system is extensible and performant enough that it may handle an arbitrary number of frontend boards and detectors. This is illustrated in Figure 6.19 in which several detectors, with their own set of frontend boards, are being controlled, configured and readout via the VRS system. In this case, an external clock and trigger system send their signals to a 'VRS Supervisory Board' (VSB) which handles the synchronous transmission of these signals to the independent groups of frontend boards located on the separate detectors. The VSB also acts to multiplex the communication between the DAQ PC and the many grouped frontend boards: a single network connection is made between the DAQ PC and VSB, which handles the forwarding of commands to specific frontend boards. This and other VSB functionalities are implemented by dedicated electronics as well as by firmware loaded on an on-board FPGA.

The work relevant to the present thesis primarily concerns the development of the software infrastructure of the VRS system. In the next sections aspects of the VRS software will be introduced

[10] Network communication in the VRS system is implemented following the User Datagram Protocol (UDP) network communication protocol. Physical network connections are made using standard Ethernet cables and on-board RJ-45 connectors.

and described.

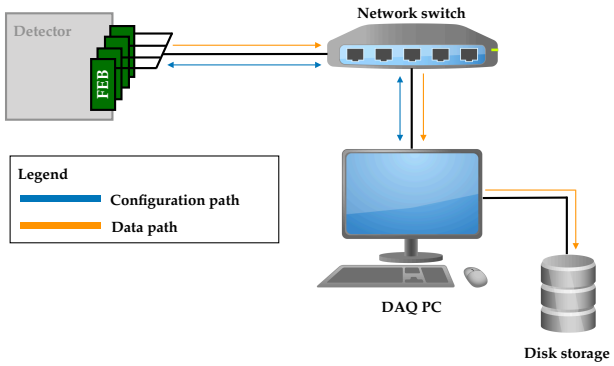

Figure 6.18: A minimal VRS setup, in which a set of frontend boards (FEBs) are connected directly to the hosting DAQ PC. This setup is that typically used in labs and test benches, in which direct study of the VMM ASIC can be performed.

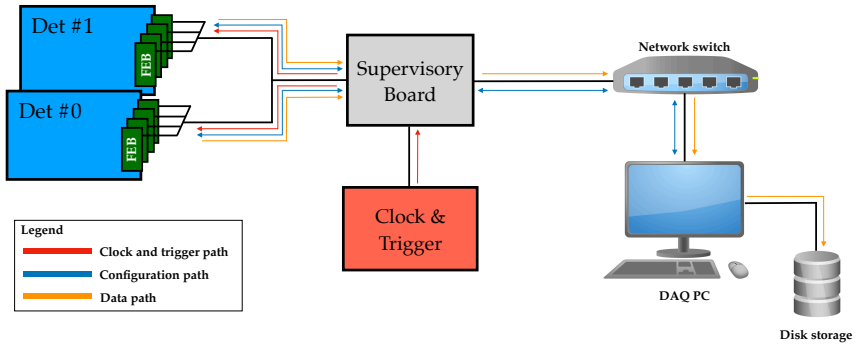

Figure 6.19: A standard VRS setup, in which detector-grouped frontend boards (FEBs) receive external clock and trigger signals that are transmitted synchronously via the intermediate VRS Supervisory Board (VSB). This setup is used mainly for data taking scenarios in testbeams or in cosmic-ray stands.

6.4.1 VERSO

The software residing on the DAQ PC side (server side) of the VRS system, illustrated in Figure 6.18 and 6.19, provides the main interface by which users can control the overall state of the frontend electronics, either on- or off-detector, as well as control the data acquisition and monitoring. The high-level software suite containing these functionalities is referred to as 'VERSO', an acronym for 'VMM Embedded ReadOut Software'. In this section an overview of VERSO's role within the VRS system will be provided.

The user interface (UI) and software backend to VERSO are written entirely in the C++ pro-

gramming language, with the graphical user interface (GUI) relying primarily on the use of the Qt framework [167]. The VERSO GUI is shown in Figure 6.20. VERSO has two main responsibilities: orchestrating configuration processes and data-acquisition (DAQ). Both functionalities are performed using a custom-made communication protocol implemented in the UDP/IP communication protocol. As illustrated in Figures 6.18 and 6.19, all communication from VERSO is sent over the network to either FPGAs located directly on the frontend boards or on the VRS supervisory board, depending on the data-taking situation. The logical blocks indicated in Figure 6.20 are described as follows:

Run Control This block sets up the underlying configuration and DAQ that VERSO implements. From here the user can initiate and terminate the data taking sessions ('runs'). VERSO supports several frontend types, either differing in the version of the VMM or in the implementation of the firmware (selected by the 'VMM2', 'VMM3', or 'L0 R/O' buttons). The 'Setup' and 'Config' fields allow a user to load configuration files describing the detector-readout-element-to-VMM-channel mapping (detector geometry) and VMM ASIC configuration, respectively.[11]

Network Communication This block describes the network addresses of the frontend electronics and handles the sustained connection to them.

FPGA and Operational Parameters This block sets the configuration parameters for the FPGA on the connected frontend boards. Such parameters, for example, are the frequency and width of the VMM channel test pulse signal and the bunch crossing clock driving the readout of the VMM ASIC.

Message Reporting Reports messages visually so that the user may acknowledge the current state of the system. Sends messages also to files stored on the DAQ PC.

VMM Configuration Panel The 'Global Registers' and 'Channel Registers' panels allow the user to set the configuration parameters of the VMM ASICs on the frontend boards. An instance of the VMM Channel Registers panel is shown in Figure 6.21.

VMM Calibration Panel This panel allows the user to schedule various VMM calibration routines. The user may also provide files containing calibration constants/parameters, derived from previous calibration runs, that get loaded into the associated VMM configuration bitstreams during subsequent VMM configuration processes.

[11] The detector-readout-element-to-VMM-channel mapping refers to both the geometric description of any detectors to which VERSO is communicating and the correspondence, for example, between an MM detector's strip location to the frontend board and VMM channel responsible for reading out that MM strip's signals. This correspondence is required if one wishes to construct high level objects, such as particle tracks, traversing through several detector layers and with each layer read out by different frontend boards.

The configuration of the FPGA is sent over the network from VERSO and follows an address-value mechanism whereby each configurable parameter of the FPGA has its value stored within the FPGA at a specific memory register. VERSO sets an FPGA configurable parameter by forwarding a register address followed by its corresponding value as specified by the user. The FPGA sends acknowledgement packets over the network to VERSO upon each message received, indicating whether it was able to successfully perform the requested configuration action or not. If not, VERSO can log this information and re-attempt transmission of the lost or mis-handled packet containing the configuration specification. The VMM ASIC receives its own configuration following a Serial Peripheral Interface (SPI) protocol.[12] An individual VMM ASIC configuration specification is rather large and is nearly 2 kilobits long. This large configuration bitstream is due to the highly configurable nature of the VMM ASIC, especially so given that each of the 64 channels of a single VMM has 24 bits worth of configuration information. The specification of the VMM configuration bitstreams is set by the user in the 'Global' and 'Channel' configuration panels on the GUI, shown in Figure 6.20. The 'Global' VMM parameters are those that are not specific to an individual channel; for example, the VMM channel threshold, channel gain, or the specification of the VMM signal shaper's integration time. The VERSO software handles the construction of the configuration bitstreams for each VMM ASIC and forwards them to the corresponding FPGA that is directly connected to the VMM to be configured. Upon receipt of a VMM configuration bitstream, the FPGA forwards it to the SPI input of the specified VMM. The FPGA knows to which of the (potentially several) on-board VMMs to forward the configuration based on a custom addressing dataframe parsed by the FPGA that VERSO prepends to each configuration bitstream.

The VERSO DAQ backend is responsible for handing the readout of the frontend boards. VMM hit data is sent over the network and received by VERSO. The hit data may be that of a detector being readout by the VMM-based frontend electronics or simulated data created by the individual VMM channel test-charge capacitors that can inject signal onto each of the VMM channel inputs. In the latter case, the frequency and characteristics (width, amplitude, etc...) of the signals are specified via the FPGA configuration described above, since the FPGA orchestrates the triggering of the VMM channel test-charge injection signals.

VERSO's DAQ functionalities are implemented following a single-producer/single-consumer architecture. In such an architecture, at the start of each run, the software constructs two threads with independent tasks: a 'listener' thread and an 'event builder' thread. The listener thread is responsible for receiving the raw digitised data packets over the network from the frontend electronics and buffering them in a queue. The listener thread indexes each of the data fragments in the queue based on their corresponding event or trigger identifier. VMM hit data corresponding to the same trigger will all be given the same such identifier. For each received trigger identifier, the event builder thread collects the associated data fragments from the queue so that all fragments from a given event may be handled at a single time. The event builder thread is responsible for decoding

[12] `https://en.wikipedia.org/wiki/Serial_Peripheral_Interface`

Figure 6.20: ***Top***: Main VERSO user interface, with different logical blocks indicated. ***Bottom***: VERSO dataflow monitor, showing the rate of VMM channel hits for all connected frontend boards and VMMs. Also displayed are the rate at which individual hits and events are recorded, where an 'event' is a collection of hits associated with the same trigger. The data being shown in this graph correspond to hits generated with the VMM channel test pulse injection with the output data recording disabled. The DAQ efficiency, defined in Equation 6.1, can be observed on this graph by dividing the blue line by the red line.

and inspecting the raw data contained in these fragments and building an output data structure that represents high-level constructs related to the event as a whole: associating the event's VMM hits with specific frontend boards and/or detector readout elements, for example. It is this latter aspect — the gathering together of all data fragments corresponding to a specific trigger — that is referred to as 'event building'. Once such an event is built, VERSO stores it on disk in the ROOT `TTree` format so that it may be used in later offline analysis.

Given the asynchronous nature of the network communication, as well as the fact that the UDP/IP network protocol does not ensure a fixed latency, events based on the same trigger are not guar-

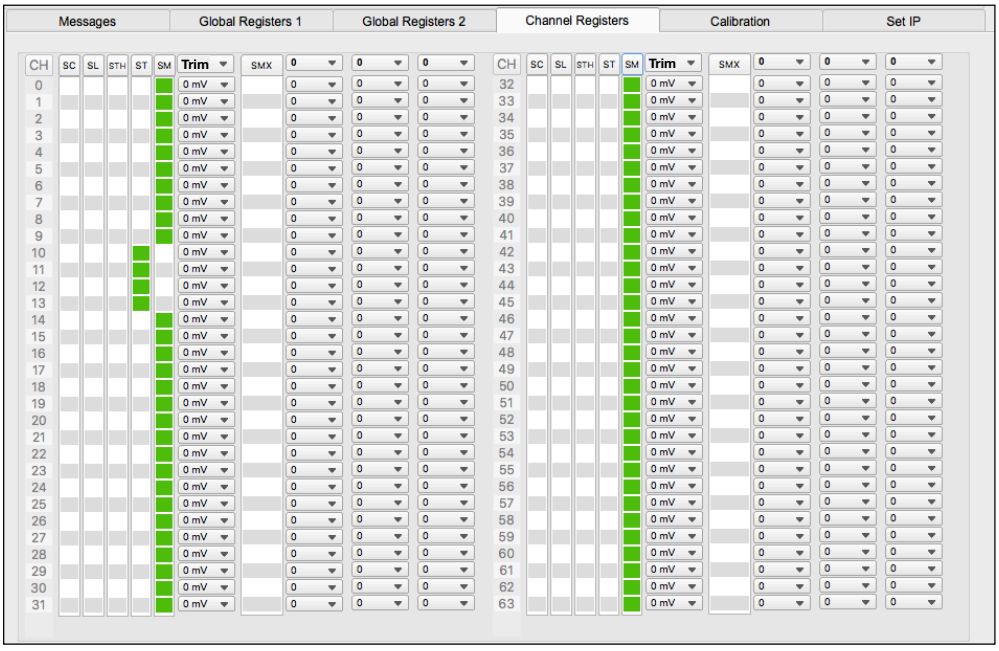

Figure 6.21: The VERSO 'Channel Registers' panel, showing the configuration specification of each of the 64 channels of an individual VMM. The 'ST' ('SM') flags activate the corresponding channel's internal test-charge capacitor (masking). In the example shown, VMM channels 10-13 (inclusive) will have their test-charge capacitor activated. All other channels are masked and are effectively disabled. The 'Trim' configuration refers to the 5-bit channel threshold trimming, specifying by how much the individual channel's threshold should be moved relative to the global VMM threshold value.

anteed to be captured by VERSO within the same UDP/IP network packet if there are multiple frontend boards being read out. The use of the single-producer/single-consumer model is adapted to such a case, with the producer (listener) thread indexing and filling the queue and the consumer thread performing the event building. In VERSO, the queue in which indexed data fragments are stored is implemented as a lock-free[13] concurrent queue. As a result, the queue may be accessed and manipulated by each thread without affecting or blocking the other, allowing for smooth DAQ operation. The event building thread can then periodically inspect the queue and wait for a configurable amount of time for the expected number of data fragments from each of the frontend elements before finishing the event building process. Once the event building process is completed, the event building thread starts the process over and begins gathering the data fragments from the next-received trigger number identifier. The use of the multi-threaded architecture, pivoting around a fast concurrent queue implementation,[14] has allowed VERSO to achieve 100% DAQ efficiency in

[13] `https://en.wikipedia.org/wiki/Non-blocking_algorithm`
[14] `http://moodycamel.com/blog/2013/a-fast-lock-free-queue-for-c++`

all use-cases encountered (Section 6.4.3). Here the DAQ efficiency is defined as,

$$\varepsilon_{\text{DAQ}} = \frac{N_{\text{recorded}}}{N_{\text{expected}}}, \tag{6.1}$$

where N_{recorded} is the number of events stored on disk by the event building thread and N_{expected} is the total number of event data fragments sent over the network to VERSO from the frontend electronics. The latter number, N_{expected}, can be determined via counters located on the FPGAs on the frontend.

6.4.2 Calibration of the VMM ASIC

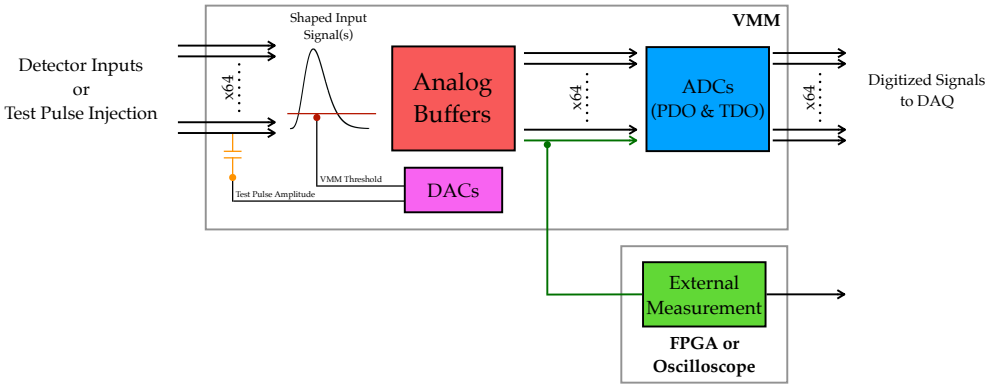

Figure 6.22: Illustration of the concept of using an external measurement for calibration of the VMM internals. For the automated calibration routines described in the text, the ADC provided by the FPGA is used as the external measurement device but an oscilloscope probe can be used as well. For calibration purposes, the external measurements are of the VMM channel analog outputs prior to digitisation via the internal VMM ADCs (green), the VMM threshold DAC analog output (red circle), and the output of the DAC setting the test pulse amplitude (yellow) provided by the test charge capacitor on each VMM channel. Not indicated, external measurements are also made of each of the VMM channel's discriminator thresholds.

As mentioned above, the VERSO software performs automated calibrations of the frontend electronics. The types of signals being collected on the NSW detectors are characterised by very low signal amplitudes and are therefore typically noise dominated, with much of the noise component derived mainly from the frontend electronics themselves. It is therefore necessary to optimise the readout functionality of the VMM ASIC by calibrating it and ensuring that its signal-to-noise ratio, S/N, is maximised during data taking. Additionally, to achieve the best performance in the reconstruction of particle tracks in the NSW detectors, several key aspects such as the charge amplitude and timing measurement of the VMM need to be understood and well calibrated. In this section a few of the calibration routines supported by VERSO will be discussed.

Many of the VMM calibrations rely on an external measurement of certain VMM characteristics to be made. These measurements are made using absolute measurements independent of any of the internal VMM measurement mechanisms. The VMM has built-in multiplexing capabilities that allow for it to route analog signals to the MO output described in previous sections. Specific for calibration, the VMM can route to MO any of the following:

Channel Analog Output: Analog levels of each channel's analog level prior to digitisation via the internal ADCs

Threshold DAC: Analog output of the DAC controlling the global VMM threshold

Test Pulse DAC: Analog output of the DAC controlling the test pulse amplitude

Channel Threshold Trimmers: Analog measurement of each channel's discriminator's threshold voltage

On the frontend boards housing an FPGA, as in Figure 6.14, the external measurement is made by routing the MO output of a VMM to the built-in 12-bit ADC of the Xilinx FPGA, referred to as the xADC. The xADC has a full scale range of 1 V, providing a digital measurement resolution of roughly 0.24 mV. The digitised samples produced by the analog measurements of the xADC are forwarded to the VERSO DAQ software so that they can be stored and subsequently processed by the calibration algorithms. As the xADC is housed in the on-board FPGA, the VERSO software can control its operation and access its measurements using similar configuration and DAQ functionalities as described in the previous section. Alternatively, the external measurement of the MO can be made via an oscilloscope probe. The use of the oscilloscope probe does not lend itself to automated calibration but provides for direct inspection of the analog levels as well as a reference on the measurements of the xADC. An illustration of these external measurement mechanisms is provided in Figure 6.22.

DAC Calibration

Each VMM has two important DACs: one for setting the global VMM threshold and one for setting the amplitude of the injected test-pulse charge via the test-pulse capacitor on each channel. Each is 10-bit, meaning that it can be set to integer values between 0-1023 (inclusive). In order for the user to know in absolute terms what the VMM threshold or test-pulse amplitude is set to, i.e. in terms of analog voltage values and not digital counts, the analog level of each DAC is measured using the xADC. At each configured DAC value, between 0 and 1023, analog measurements are made of the DAC output and the relationship between the output analog value and configured digital value can be obtained. This is illustrated in Figure 6.23 for a single VMM's threshold DAC. A linear relationship is assumed and from a linear fit the $DAC\text{-}mV$ relation, giving the mapping from

the configured digital DAC value to its analog value observed by the VMM channels, is obtained from the slope of the line. The DAC-mV relation between VMMs is not necessarily the same, and therefore must be measured for every single VMM in the system in order to be able to configure them all with the same *absolute* threshold or to inject test pulses with common amplitudes across VMMs. Given the known capacitance of the test-pulse injection capacitor of the VMM channels, the value of the injected charge for a configured test-pulse amplitude can be obtained via the $Q = C \cdot V$ relation, where C is the injection capacitor's capacitance and V is the test-pulse DAC output amplitude.

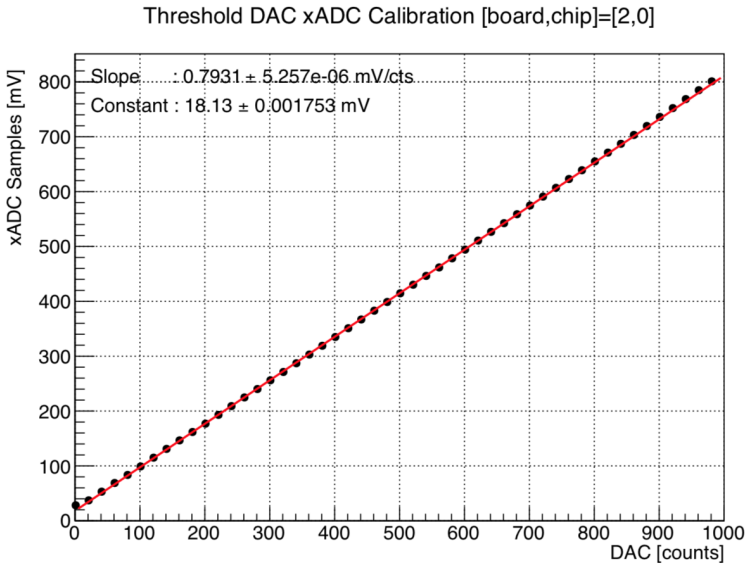

Figure 6.23: VMM global threshold measurement, using measurements made by the xADC, as a function of the set digital DAC value. The slope and y-intercept ('Constant') of the linear fit are provided. The slope gives the DAC-mV relation that allows for converting the digtal DAC configured value to an analog level at the VMM inputs.

Measurement of Channel Baselines and Noise

External measurements via the xADC can be made of each channel's analog output (i.e. prior to digitisation) when no signals are present at the channel input. Doing so allows one to obtain a measurement of the channel *baseline*, i.e. the ambient level when no signal is present. Fluctuations about the baseline, then, give a measurement of the inherent channel *noise*. Measurements of both of these quantities are important, both when the frontend electronics are and are not interfaced to a detector. Measurements of the noise when the frontend electronics are not interfaced to a detector give a measurement of the inherent electronics noise characteristic of the frontend board. Interfacing the frontend electronics to a detector introduces additional capacitive effects that can introduce new sources of noise that must also be studied and perhaps reduced. Knowledge of both

the absolute baseline and noise levels of each VMM allows for the optimal threshold to be set for the VMMs. In order to be sensitive to small-amplitude signals, one would set the threshold just above the baseline since the baseline defines the effective zero-level of an observable signal. In practice, one sets the VMM threshold a few times the magnitude of the measured noise above the baseline: typically 2 to 3 times the level of the measured noise above the baseline. Figure 6.24 shows an example of the measurements of a VMM channel's baseline and noise, taken by the xADC. The xADC samples each channel's baseline and from the width of the distribution of these measurements the noise level can be obtained. Figure 6.25 shows the evolution of the measured noise as a function of the VMM's configured gain, both when the frontend board housing the VMM ASIC is and is not interfaced to a small prototype MM detector having ∼30 pF capacitance.

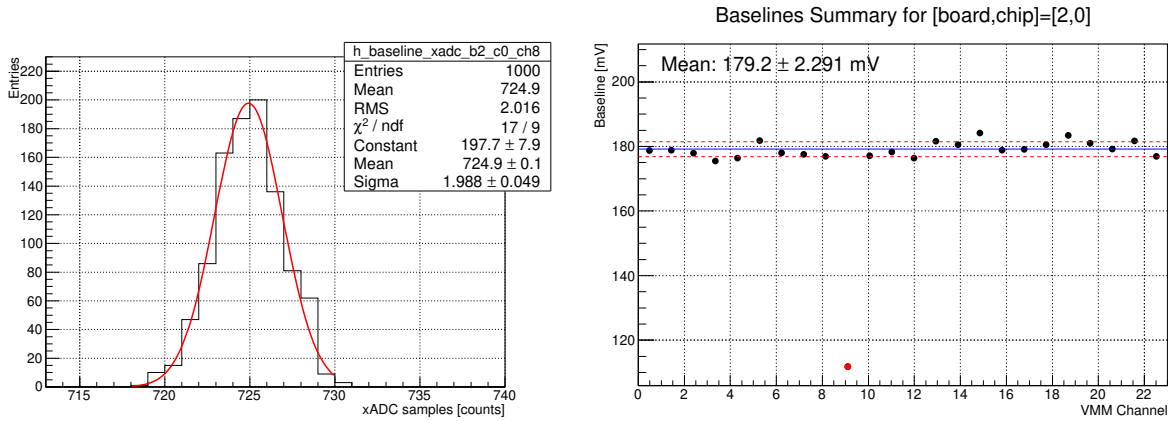

Figure 6.24: ***Left***: 1000 xADC measurements of a single VMM channel's input baseline. The width of the Gaussian fit (red) gives the channel's noise level. ***Right***: Summary of the baseline measurements for 24 channels of a VMM. The error on the reported mean baseline, which is indicated by the blue line, is given by the RMS spread of the baseline measurements across all channels and is indicated by the red dashed lines. The channel with the red point (channel 9) is a faulty (dead) channel.

Channel Threshold Trimming

The global VMM threshold inherently has channel-to-channel variation. In order to ensure that the discriminators of all VMM channels in the system fire at the same value, the VMM has a per-channel threshold trimming functionality included. This functionality allows for each channel's threshold to be adjusted ('trimmed') in 32 steps that cover a range of approximately 32 mV. To equalize the channel discriminators, external measurements made via the xADC are performed. For each channel, measurements of the channel discriminator's threshold are taken at each of the 32 steps in the allowed range of trimming. Once done for every VMM channel in the system, a metric for equalizing their thresholds is needed. In the example of Figure 6.26, the metric is to bring all channels as close as possible to the globally set VMM threshold. Once this is done,

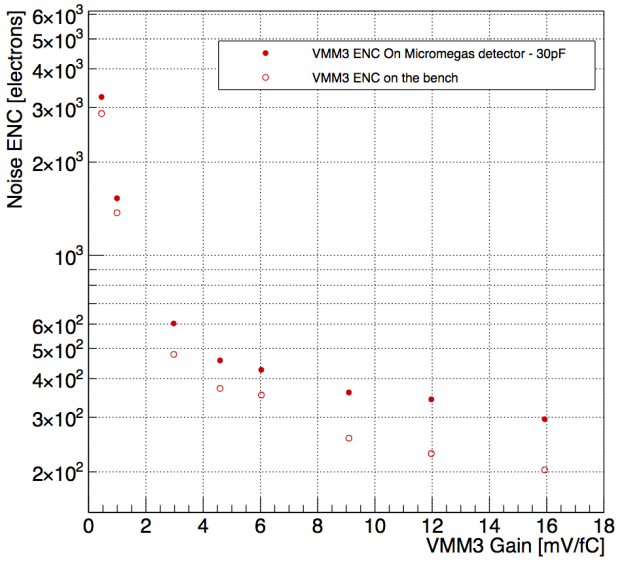

Figure 6.25: Equivalent noise charge (ENC) of the VMM3 as a function of its configured gain value. Noise measurements are performed with a GPVMM-type frontend board both on- (filled circles) and off-detector (unfilled circles). The detector used is a small $10 \times 10\,\mathrm{cm}^2$ prototype MM chamber.

all VMM channels in the system ideally have equal thresholds and will have equivalent responses to input signals. As seen in Figure 6.26, prior to this calibration there is quite some variation, which shows the importance of this fine-grained calibration procedure. The amount of threshold trimming necessary to bring a channel to its equalized state is independent of the globally set VMM threshold; therefore, once the calibration procedure is performed and the subsequent channel threshold trimming is applied, the global threshold can be adjusted with confidence that all channels in the system are likewise adjusted coherently.

Signal Amplitude Gain and Pedestal

The output of the PDO for each channel must also be calibrated. The PDO calibration aims to correct the per-channel response to a given input signal. Its calibration relies on the injection of test pulses on each channel. Measurements of the output PDO for each channel can be taken as a function of the amplitude of the test pulse, which is configured via the 10-bit test pulse DAC. This is illustrated on the left side of Figure 6.27. Linear fits to the resulting measurements provide a measure of the channel gain, which will generally differ with respect to the globally set VMM channel gain. The right side of Figure 6.27 shows the channel-by-channel slope (gain) variation measured for a single VMM. From the DAC-mV relation derived from the DAC calibration described above, the observed slope can be obtained in absolute terms (i.e. in terms of mV/fC) to give the channel gain. This can then be used to correct the signal amplitude measurements in later analysis; for example, when constructing particle tracks. From the measurements above, relying on the linear

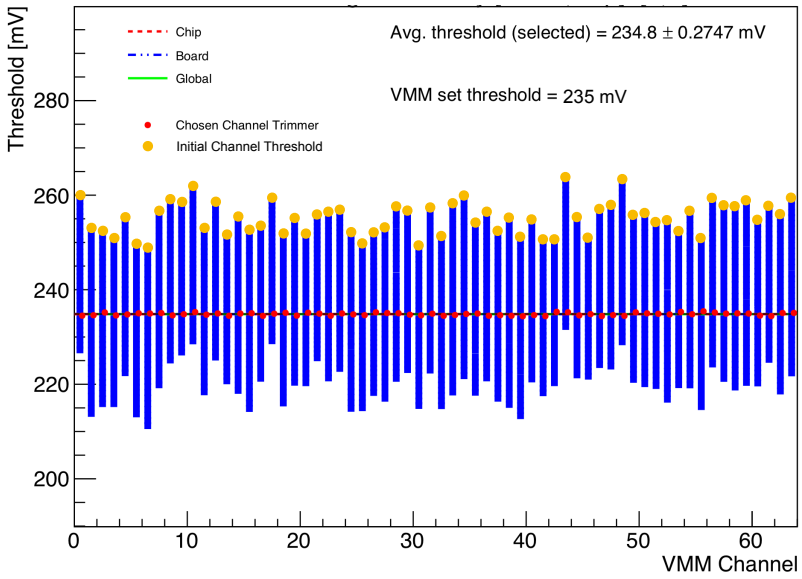

Figure 6.26: Summary plot of a VMM channel threshold trimming calibration. For this set up, a single VMM is calibrated and has a global threshold set to 235 mV using the threshold DAC. The channel-by-channel variation of the threshold is seen by the yellow dots, which indicate the per-channel threshold measurement made by the xADC prior to any threshold trimming. The blue columns indicate the entire threshold range accessible by each channel by scanning the full range of channel trimmers and taking xADC measurements at each. It can be seen that most channels have a full range of approximately 32 mV. The calibration algorithm shown here picks for all VMM channels the amount of trimming that allows for them to have thresholds as near as possible to the globally set threshold. In this case, the post-calibration threshold of all channels, indicated by the red dots, is very nearly 235 mV, and is 234.8 ± 0.27 mV on average.

fits, the PDO pedestal can be measured as the extrapolated y-intercept. Due to the VMM's internal ADC having a built-in baseline subtraction, the PDO pedestal will differ with respect to the baseline as described Section 6.4.2. The PDO pedestal is an offset that must be subtracted from the PDO amplitude measurements prior to any analysis or track fitting. The VMM can be configured to output PDO measurements for those channels that are neighboring a channel that has an above-threshold signal but that do not themselves have an above-threshold signal. With this option enabled, individual channels may have injected test-pulses and have the neighboring channels' PDO measurements recorded. The neighboring channels readout in this manner thus provide an alternative method for measuring the PDO pedestal. The neighbor-enabled method for obtaining VMM channel PDO pedestals is actually more efficient as a significantly smaller amount of data samples need to be recorded since one does not need to scan over values of the test pulse amplitude: pedestals obtained in this manner can be obtained for any single value of the test pulse amplitude since they are obtained from those channels *not* being pulsed. Figure 6.28 shows the PDO pedestals measured for the channels of a single VMM using the linear-fit extrapolation and neighbor-enabled methods. The two methods agree quite well with one another.

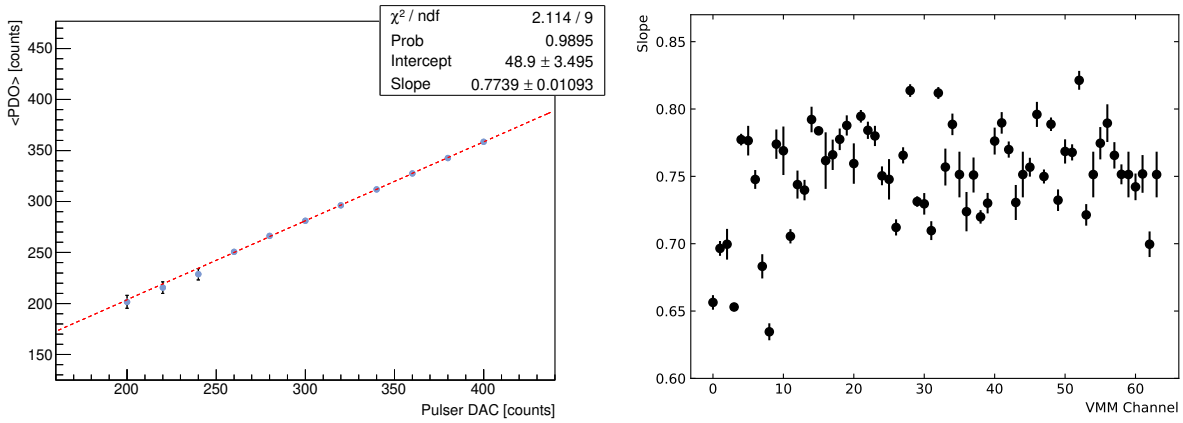

Figure 6.27: ***Left***: Average PDO response as a function of test pulse amplitude, for a single VMM channel. The slope gives a measure of the channel gain. ***Right***: Measured slopes of the PDO response to injected test pulses for 64 channels of a VMM.

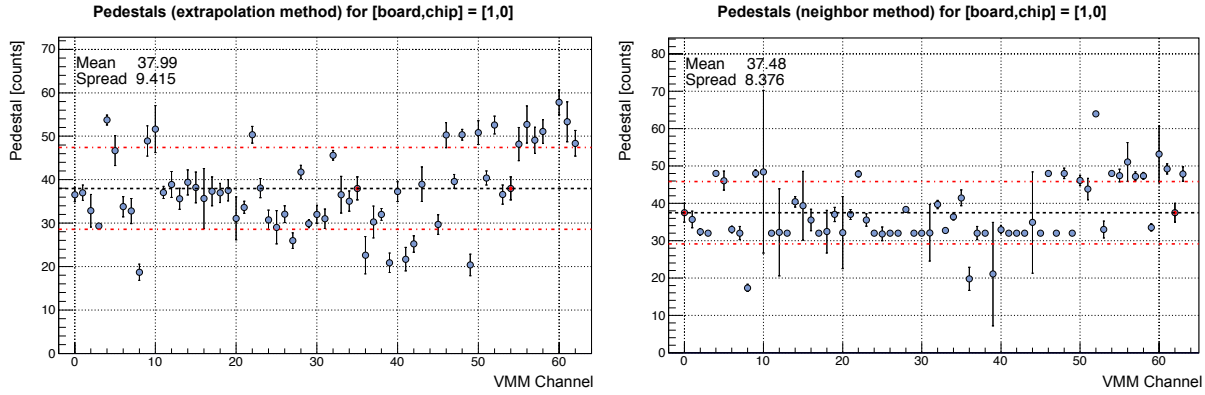

Figure 6.28: Measurements of PDO pedestals of each channel of a VMM. ***Left***: Using the linear-fit extrapolation method. The mean pedestal measured across all channels is 37.99, in units of PDO counts. ***Right***: Using the neighbor method. The mean pedestal measured across all channels is 37.48, in units of PDO counts.

Timing Measurement Calibration

The time measurement of the MM readout is of the utmost importance for providing high-resolution tracks since the track-primitive building in the MM detectors relies primarily on the μ-TPC reconstruction method, as illustrated on the right side of Figure 6.9. The TDO output must then be calibrated if high resolution track-primitives are to be constructed.

The TDO provides signal timing information and is the amplitude of a latched time-to-analog converter (TAC) voltage ramp that is initiated either when the signal peak amplitude is found

(coinciding with the time at which the PDO value is latched) or when the signal amplitude first crosses above threshold. The former is referred to as 'timing-at-peak' and the latter as 'timing-at-threshold'. The operation of the TAC ramp and subsequent TDO measurement is illustrated in Figure 6.29. The maximum TAC ramp time is configurable, and if the maximum is reached the output TDO value will be saturated at its maximum value. In nominal VMM operation, the TAC ramp is halted at the following falling-edge of the bunch-crossing clock driving the VMM readout and read out as the TDO. In an alternative configuration of the frontend electronics, the FPGA which provides the bunch-crossing clock to the VMM can be configured such that it halts the bunch-crossing clock for a fixed amount of time once the TAC ramp is initiated. After this fixed amount of time, the bunch-crossing clock is re-initiated and the first falling edge then halts the TAC ramp. This *fixed latency* method of acquiring the TDO allows for a fine-grained calibration of the timing measurement of the VMM. The timing calibration is performed by scanning the bunch-crossing clock's latency (i.e. the amount of time for which it is halted) and measuring the TDO output, per channel, as a result of injected test pulses. This is illustrated in Figure 6.30 for two values of the TAC ramp time: 100 ns and 350 ns. Linear fits to the resulting TDO measurements as a function of this fixed latency then give the per-channel conversion from TDO output (counts) to absolute time (e.g. nanoseconds). The timing calibration using the nominal TDO measurement method is performed similarly. However, instead of scanning the bunch-crossing clock latency, the time difference between the test pulse clock driving the injection of the test pulse charge and the bunch crossing clock is scanned. From Figure 6.29 it can be seen that by skewing the test pulse injection in this manner, for larger values of the time difference the TAC ramp time will be nearer to the next subsequent bunch crossing clock and therefore will be halted sooner. The resulting TDO measurements for several values of the amount of time skewing allows for the timing calibration to be performed.

Once the TDO conversions are known, the inherent VMM timing resolutions can be studied and further measurement improvements can be made. Figure 6.31 shows the TDO timing resolution in absolute time as a function of the injected charge. It can be seen that for input signals with charges $Q \approx \mathcal{O}(10)$ fC that sub-nanosecond timing resolutions are approached. Given that the MM detector avalanche gains are on the order of 10^4, this leads to the VMM having sub-nanosecond timing resolutions for avalanches produced from 1 to 2 drift electrons.

The timing measurement is highly sensitive to the signal shaping, as can be seen in Figure 6.31. This is due to the fact that the shaped signal's peak and threshold-crossing becomes less pronounced for larger shaper integration times, leading to poorer timing resolution. The effect of the signal shaping also introduces non-negligible timewalk, as seen in Figure 6.32, which shows the effect of timewalk for a VMM configured for timing-at-threshold. The timewalk effects are exaggerated for timing-at-threshold since the point at which a given signal rises above threshold is more susceptible to noise, in addition to the integration time effects described previously. For a given timing configuration, the measurement of the timewalk for a given signal amplitude is used to correct the timing information

provided by the VMM TDO measurement and can thus improve the subsequent track-primitive reconstruction.

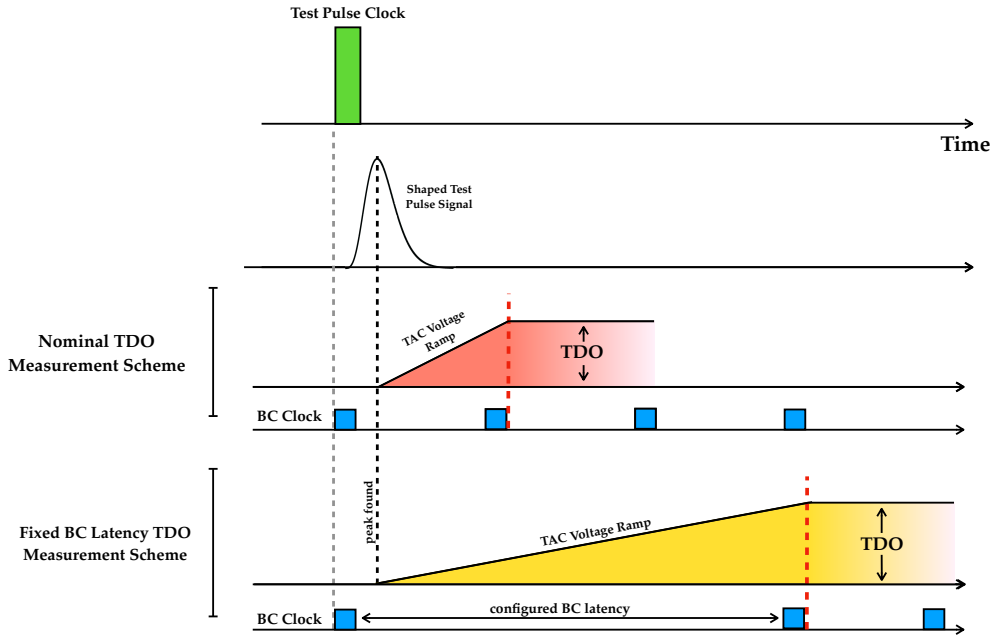

Figure 6.29: Illustration of the TDO measurement for an injected test pulse. The test pulse is injected following the test pulse clock (green), after which the subsequent shaped signal is formed. In the timing-at-peak scenario, once the peak-detector circuitry in the VMM fires, the TAC voltage ramp is initiated. In the nominal TDO measurement (red), the bunch crossing clock frequency remains unchanged and will be kept at its 40 MHz frequency in LHC conditions. In this configuration, the next falling edge of the bunch crossing clock after the TAC voltage ramp begins halts the ramp, and the TAC voltage is digitised as the TDO value. While in this mode, the TAC ramp is confined to times smaller than that given by the bunch crossing clock frequency (25 ns in the case of the LHC 40 MHz clock). In the fixed latency mode of acquiring the TDO measurement (yellow), the bunch crossing clock is halted for a configurable amount of time after which it is resumed and its next falling edge halts the TAC voltage ramp. In the fixed latency mode, the TAC can ramp to its maximal configured value (60, 100, 350, or 650 ns in the VMM3) if the configured BC latency is at least as large as this value.

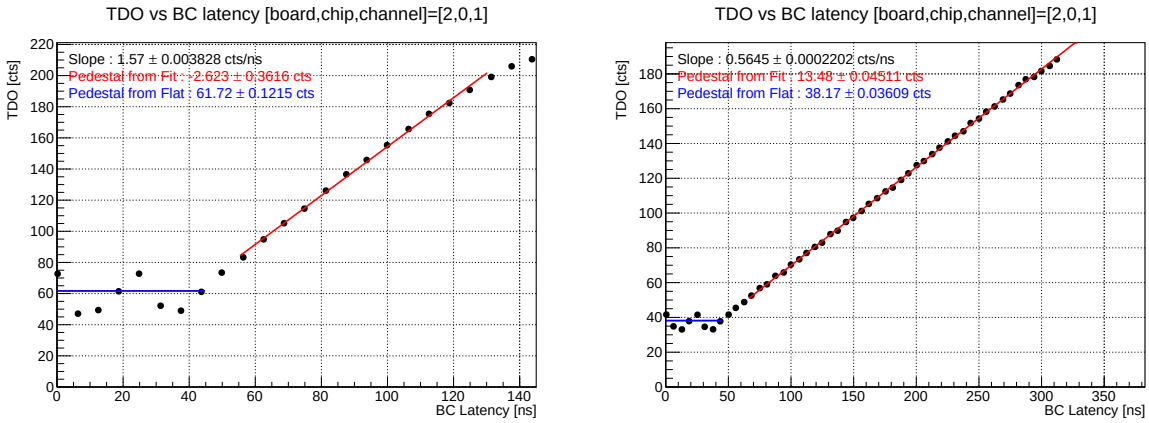

Figure 6.30: TDO calibration using the fixed-latency method for a TAC ramp time of 100 ns (***left***) and 350 ns (***right***). The flat portion, fit with a blue line, are 50 ns long and correspond to the configured VMM shaper's integration (peaking) time. As can be inferred by Figure 6.29, proper timing measurements for times smaller than the integration time cannot be obtained. The linear fit for latencies beyond the peaking time give the timing calibration for the channel via the slope which gives the conversion between the TDO digitised count values to absolute time in nanoseconds. It can be seen that the TDO conversion constant is larger for smaller TAC ramp times, which follows from the fact that the TDO 8-bit range is the same irrespective of the configured TAC ramp.

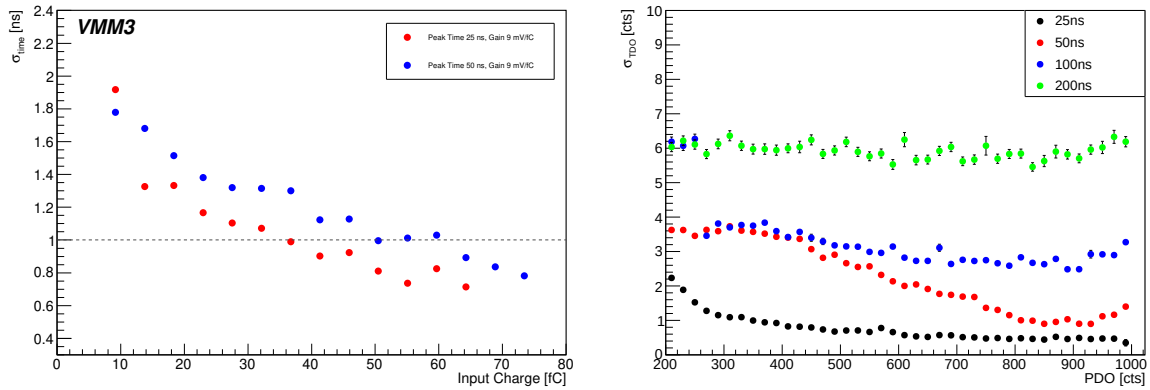

Figure 6.31: VMM timing resolution as a function of the injected test pulse charge for various integration times. ***Left***: Absolute timing resolution as a function of absolute charge for shaper integration times of 25 ns (red) and 50 ns (blue). ***Right***: TDO resolution as function of injected test-pulse amplitudes, over the full PDO range, for various shaper integration times: 25 ns (black), 50 ns (red), 100 ns (blue), and 200 ns (green).

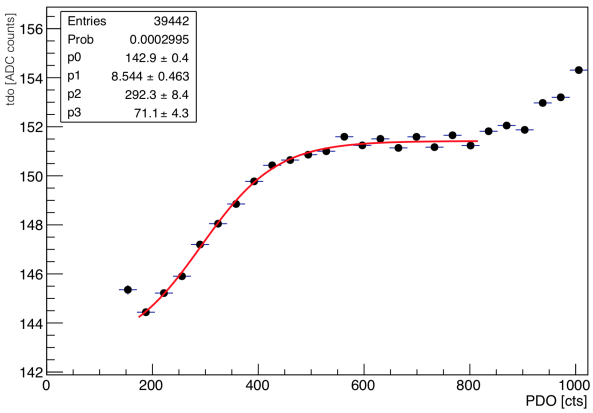

Figure 6.32: Measured timewalk for a given VMM channel configured for timing-at-threshold. The x-axis is the measured PDO value of the channel and the y-axis is the TDO value, giving the time information of the signal. The red line is a fit to a sigmoid-like function: TDO $= p_0/(p_1 + p_2 * e^{p_3 \times \mathrm{PDO}})$. The injected signals all occur at the same time, but due to the timewalk effects the measured time provided by the TDO measurement are biased to larger values as the measured amplitude increases.

6.4.3 Use Cases

Throughout the timespan of the present thesis, the VRS and, specifically, the VERSO software infrastructure has been used in a wide variety of settings and for many purposes. In this section a few such use-cases will be highlighted.

Test Benches

For the majority of studies of the frontend electronics, namely the frontend boards and the VMM ASIC, detectors are not needed. For this reason, one of the primary uses of the VRS and VERSO DAQ and calibration software is in testbenches in electronics labs. A typical setup is shown in Figure 6.33. The use of VRS in the testbench allows for quick and direct inspection of the frontend electronics via an oscilloscope. In such a setup, the response of the frontend electronics to communications sent from VERSO to the FGPA (or the FPGA to the VMM), and vice versa, can immediately be seen by visual means. Calibration and off-detector performance of the frontend electronics and VMM can also be studied in the 'clean' scenarios provided by the testbench setup. Use in labs and testbenches has allowed for the VRS software and firmware to be fully validated and for the validation of the NSW frontend boards as they have evolved.

Cosmic Stand

The VRS system has also been used to collect cosmic muon data from instrumented detectors, on cosmic stands located in NSW labs at CERN and elsewhere. One such example is shown in Figure 6.17, showing the large-scale MM chambers in the cosmic stand located at the RD51 lab at CERN. Another, more recent (at the time of writing) example is that of the final MM detector integration site whereat the final NSW MM chambers are being assembled. As the final DAQ system that will be used once the NSW is installed in ATLAS is not yet fully ready, at the time of writing, the VRS system was used to configure and readout the frontend electronics interfaced to the final full-scale MM double-wedges. For the final integration, noise tests of the electronics on the final MM detectors that will be installed in ATLAS are needed in order to verify the base performance of the large-scale detectors. The double-wedge MM detectors are additionally placed onto a cosmic-ray stand and initial data was able to be collected with the VRS and VERSO DAQ software. The VRS system, in this case, helped allow for the final NSW MM integration procedures to be put into place while the full-scale DAQ infrastructure to be used once the NSW is installed in ATLAS was made available.

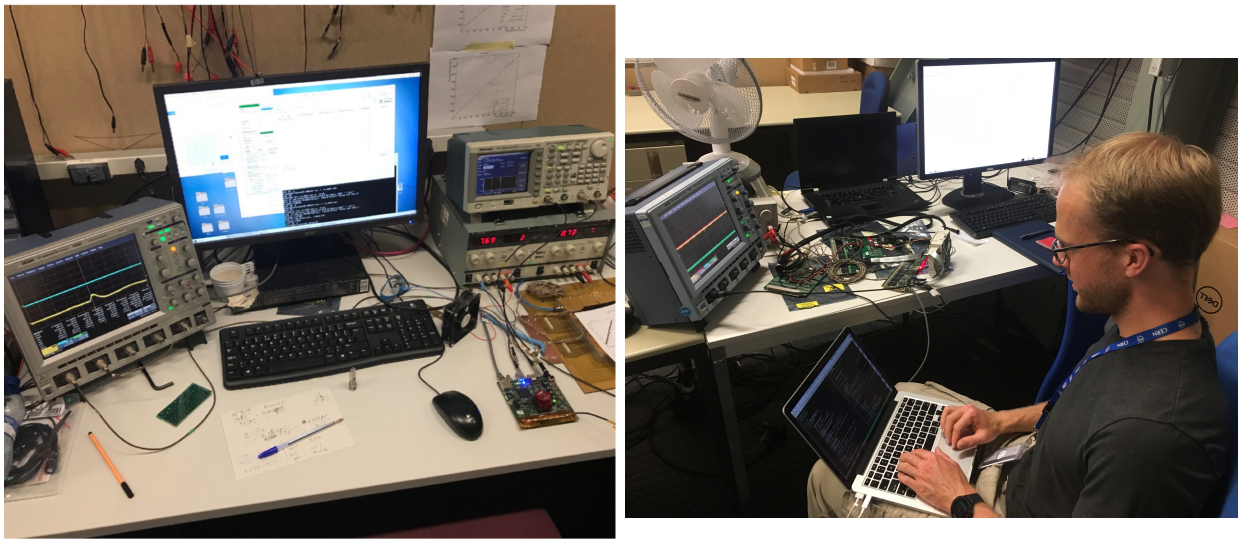

Figure 6.33: ***Left***: A typical VRS testbench setup. The frontend board housing the VMM is seen in the picture on the lower right (GPVMM-type frontend board). An oscilloscope is available for probing the frontend board and VMM outputs visually. The frontend board is connected to the shown PC via a direct Ethernet connection. A function generator producing input signals is at the right, being used to test the injection of signals independently from the VMM channel test charge capacitors. ***Right***: Another example of a typical testbench setup, with the present author hard at work. In this case, VERSO is hosted on a single laptop with a direct Ethernet connection to the frontend board. The frontend board in this case is an MMFE8-type frontend board. In the case shown, the VRS firmware on the frontend board is being debugged. The PC at the upper right, connected to the frontend board via a mini-SAS connection, is responsible for writing and uploading the firmware to the on-board FPGA.

Frontend Board Validation

The design of the frontend board to be used in the NSW was being finalised throughout the timespan of this thesis. In order to validate that the frontend board met the standards required for operation in the NSW, specifically that it did not introduce too much electronics noise as a result of its design, the VRS system and VERSO software were required. The prototype boards, then, for the NSW are all outfitted with an FPGA and Ethernet connector as in Figure 6.14 so that they can be interfaced with the VRS readout system. The main goals of the frontend board validation are to ensure proper routing between the VMM ASIC and the rest of the frontend board elements which can be tested by verifying the VMM's response to specific configuration commands. Noise tests are also performed by measuring the frontend board noise as described in Section 6.4.2. Such an example of a noise tests using prototype MMFE8s while interfaced to full-scale MM detector chambers is shown in Figure 6.17. Performing such tests allow, also, to verify that the physical connections between the frontend electronics and detectors, such as the ZEBRA connectors in the case of the MM detectors, are without fault. Any such fault would appear as a faulty set of channels

Figure 6.34: ***Left***: Assembled NSW MM sector (double-wedge) at the MM integration site at CERN. ***Right***: MM sector on the cosmic-ray test-stand at the MM integration site at CERN. The detector is instrumented with MMFE8 frontend boards with Ethernet readout achieved via the VRS and VERSO DAQ software, hosted on the PC located just behind the cosmic stand.

only when the electronics are interfaced to the detector, for example.

Testbeams

From 2016 until the present time of writing, the VRS system and VERSO DAQ software have been used as the main DAQ infrastructure for test beam campaigns at CERN. The testbeam area at CERN is located at the CERN Prevessin site, in the North Area. Proton beams from the SPS are made to collide with a fixed target upstream of the detectors and the collision products are sent through several beamlines in the North Area along which detectors can be placed. The beams from the SPS are extracted roughly every 40 seconds and particles are sent down the beamlines towards the detectors for roughly 5 second 'spills'. In such cases, the particle rates observed at the detectors are on the order of several 100 kHz. The main purpose of these testbeam campaigns was to verify the performance of both the VMM ASIC and prototype NSW frontend boards in high-rate environments. The VRS system and VERSO DAQ software have succeeded in recording

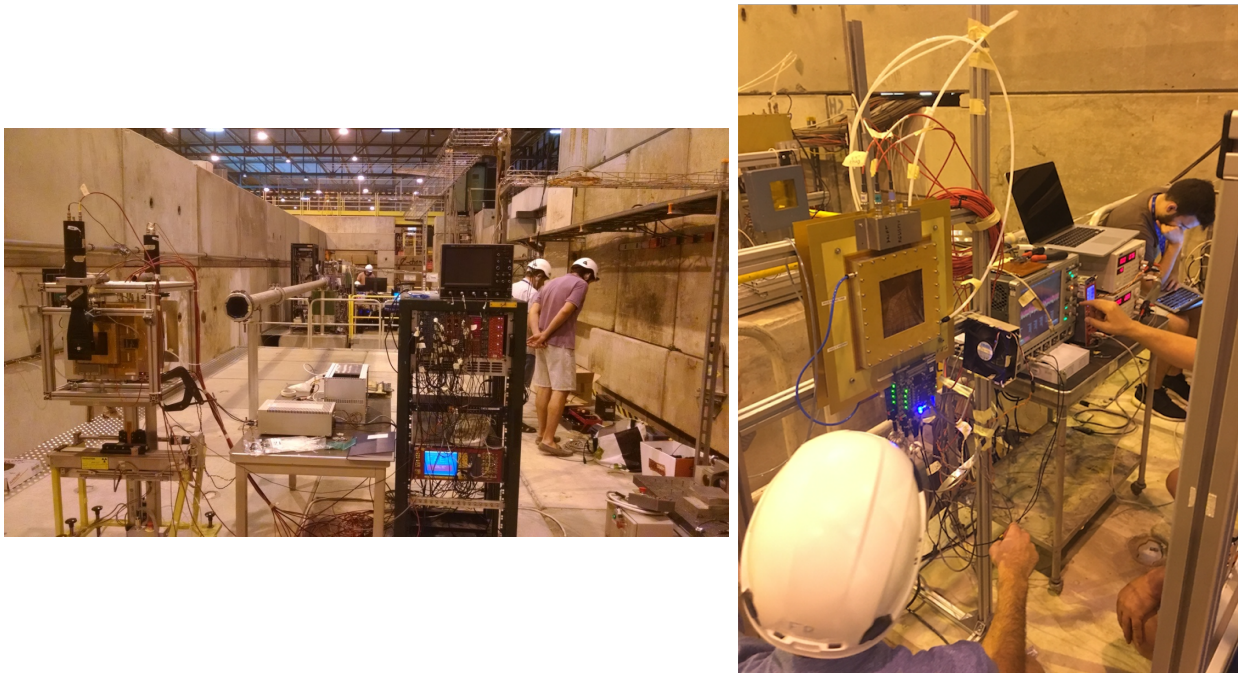

Figure 6.35: ***Left***: Typical testbeam setup at the CERN North Area, located at the Prevessin Site at CERN. The small prototype MM chambers are housed in the aluminum frame on the left with scintillator trigger paddles on either side. The beam tunnel can be seen extending into the distance. The trigger and high voltage systems are seen on the rack in the middle of the picture. The DAQ PC is housed in a control room (not shown) some 10-15 m outside of the beam area. ***Right***: A close up of a small prototype MM chamber at a test beam. The GPVMM-type frontend board is interfaced at the bottom. Here it can be seen that there are two MM chambers back-to-back, each with a frontend board attached at their bottom. The cables providing high voltage (in red) and gas (white) to the MM chambers are seen, as well.

data at 100% efficiency at these rates, allowing for high volumes of data to be collected and used in NSW detector performance and electronics studies. For such data, used to validate the frontend electronics and tracking performance of NSW-type detectors in realistic data-taking environments, the calibration routines described in Section 6.4.2, and implemented within the VERSO software infrastructure, are of the utmost importance.

VMM Production Validation

Throughout the work presented in this thesis, the VMM ASIC has gone through 3 versions. In order to validate that a given VMM production is adequate — meaning that there is no systematic fault in the ASIC production, for example — large scale tests of all ASICs produced in manufacturing test runs are required. In many cases, the automated calibration routines supported by VERSO were required to perform these production tests. In one such case, the manufacturing process of

the silicon wafers used to produce the VMM was changed and in order to verify that the resulting changes observed on the VMM performance were due to this change, a direct study of the VMM characteristics prior to the changed process being applied to the silicon was needed. This is shown in Figure 6.36, in which a silicon wafer containing pre-cut and pre-thinned VMM ASICs are being interfaced directly with the VRS and VERSO DAQ software. The interfacing of the VERSO DAQ and calibration software is achieved by the use of a custom-made adapter board provided to the silicon inspection machine, shown on the right side of Figure 6.36, that makes all the necessary connections to the VMM ASIC pins by making direct contact with the silicon. The silicon wafer is moved beneath the adapter card via the inspection machine so that it may inspect different locations on the silicon wafer in which the VMM ASICs are located.

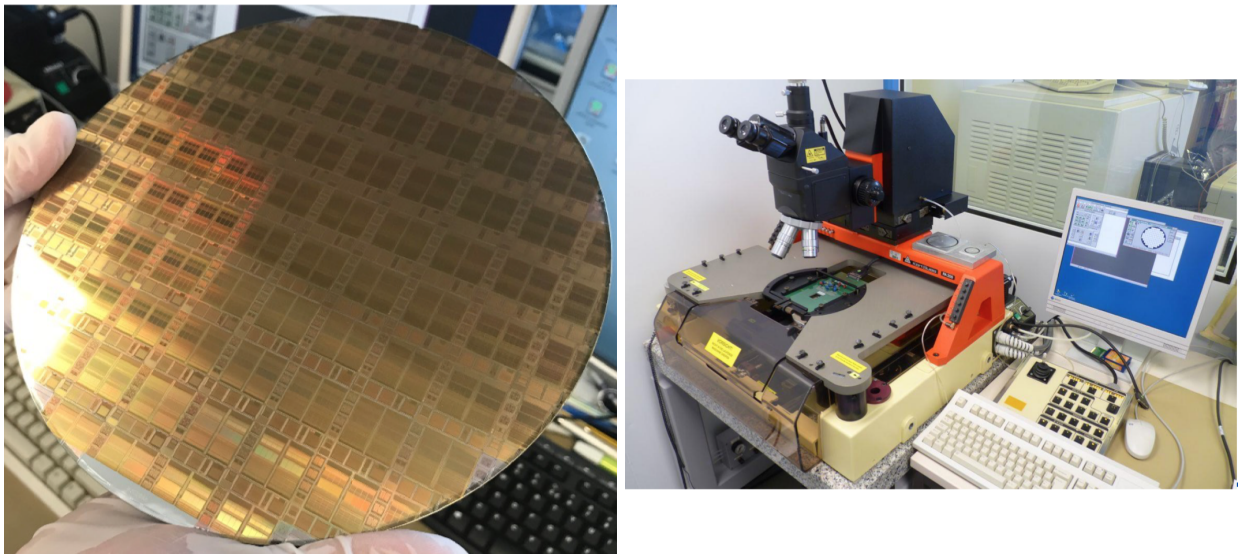

Figure 6.36: Silicon wafer testing lab at CERN. ***Left***: Silicon wafer die containing VMM ASIC chips prior to being cut. Additional ASICs are shared on the same silicon wafer, as well, so as to make efficient use of space on the silicon and of CHFs. ***Right***: The silicon wafer die is inserted into the machine on the left. Visual inspection of the silicon can be made via the microscope. The machine allows for adapter boards to be used to interface to the ASICs directly, via the software on the PC that controls the machine.

Chapter 7

Common Elements in the Search for New Physics

Our 'Age of Anxiety' is, in great part, the result of trying to do today's jobs with yesterday's tools!

–Marshall McLuhan

There are many ways in which physicists study and analyse the data collected by the ATLAS detector. Once the gathered data has been aggregated and the physics objects therein have been reconstructed (Chapter 5), an analyst has at their fingertips access to the stuff of high energy particle physics from which they can test their hypotheses or perform measurements of fundamental quantities. In the present work, several analyses representing the search for BSM physics will be discussed in detail. The details of each of these analyses differ quite a bit, but each follow a general *analysis strategy* that will be discussed further in Section 7.1. Physics analyses start with a stated goal in mind; for example, a measurement of some quantity to a desired level of precision or, as in the analyses to be presented, a statement about the existence of a new particle or theory of new physics. The methods by which these statements are to be made are intertwined with, and generally dictate the initial design of, the overall analysis strategy and are statistical in nature. In Section 7.4 a discussion of the statistical methods used in the analyses to be presented will be given, highlighting how statements about new physics are made and how broad-strokes physics conclusions can be drawn from them. Section 7.2 follows up on the topic of how one estimates particular sources of SM backgrounds for which the MC simulation cannot be relied upon to give reasonable predictions, particularly on those processes that lead to leptons being produced in the detector that do not originate from the pp hard-scatter processes of interest.

7.1 General Analysis Strategy

In this section the analysis strategy used in the searches for new physics to be presented in Chapters 8 and 9 will be given. The general analysis workflow for designing an analysis is outlined in the following sub-sections.

7.1.1 Target the Signal

The search for a particular source of new physics, such as a particular model of SUSY (Chapter 2), begins first with the thorough understanding of the signatures that the new physics model will leave in the ATLAS detector. This generally requires a strict definition of the *final state* of the new physics model that one wishes to look for; for example, deciding to search for evidence of SUSY via the production of the SUSY partners to the SM top-quark in final states having exactly two leptons (electrons or muons) instead of exactly zero or exactly one lepton, as in Chapter 8.[1] Once a new physics model has been chosen, along with its final state, there is a well-defined *signal* to be looked for in the data recorded by the ATLAS detector. The production and decay of the sought-for signal is then simulated via MC methods in the exact same manner as for the SM processes, as described in Chapter 4. In physics analyses, the physics processes not inclusive of the sought-for signal processes are referred to as the *background* processes.

The simulation of the signal process allows one to study the kinematics of the signal in detail, in order to get an overall feel for what phase space the signal inhabits. Knowledge of both the signal final state and its kinematics therein informs the analyst about the specific SM background processes that are likely to be relevant to the analysis. For example, if the sought-for signal decays to two leptons with opposite electric charge that are of the same flavor (both leptons are electrons or both are muons, for example) it is very likely that the SM processes inclusive of Z-boson production will be relevant, since this is one of the main Z-boson decay final states, as opposed to the production of a single W-boson whose decay does not lead directly to final states with two leptons. Knowledge of the dominant SM background processes, then, allows one to determine how the phenomenology and kinematics of the signal differ with respect to those of the relevant backgrounds by comparing the simulated events of each. The aim of this is to be able to define a basis of kinematic observables that allows for the discrimination between the signal and background. From such a basis of observables, one can define regions of phase space in which the signal-to-background ratio is large, such that the likelihood of observing the presence of the signal is (ideally) maximal. Such regions of increased

[1] Performing searches for new physics by the partitioning of specific new physics models by their resulting final states allows for separate, independent dedicated analyses to be carried out for each possible final state with the idea that each one will be more sensitive to the presence of the new physics in their respective final state than would be a single analysis attempting to target all possible final states of the new physics production. The results of the independent analyses' searches can be statistically combined once they are finished, leading to enhanced sensitivities to the new physics scenario in question that is more or less independent of the final state.

signal purity[2] are referred to as *signal regions* (SR). As an example, take the case where there is a single discriminating variable in our basis of useful kinematic observables. One would apply a selection on this observable in such a way that pp collision events satisfying this selection are likely to be enhanced in signal events. This is one-dimensional SR case is illustrated in Figure 7.1.

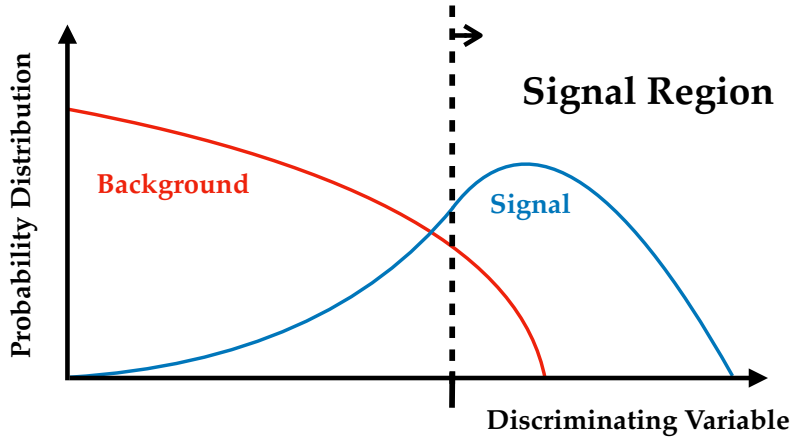

Figure 7.1: Signal region concept illustrated in the case of a one-dimensional selection made on a discriminating kinematic observable. The dominant SM background (red) is characterised by typically small values of the discriminating variable whereas the signal (blue) has values that extend beyond that of the background. The signal region in this case is defined by requiring pp collision events to have values of the discriminating variable that are larger than the value indicated by the dashed vertical line, where the signal purity is enhanced. The y-axis represents the probability distribution of the background and signal processes, not their absolute yield for a given range of the quantity on the x-axis.

7.1.2 Define the Trigger Strategy

Once an analysis understands the final state and its kinematics that it will be searching for, the strategy for gathering the pp collision data consistent with that final state must be defined. This requires that the analysis make a choice about which triggers to use for recording the collision events in ATLAS, a process referred to as defining the analysis' *trigger strategy*. The analyses to be presented in Chapters 8 and 9 are based on searches for signals with leptons (electrons or muons) in the final state. As a result, the triggers these analyses use are all based on signatures of high-p_T leptons.

The lepton triggers used in the analyses to be presented trigger are based on the p_T of the leptons. They are configured to have a given p_T threshold and if a lepton is identified in the online trigger system during the pp collisions to have a p_T value at or above this threshold, the trigger 'fires' and

[2]The 'purity' of a process is defined as the fraction of a given process in a region of phase space, relative to the sum of all processes (inclusive of the process in question).

the event is recorded (see Section 3.2.5). Lepton triggers are characterised by a sharp 'turn-on', meaning that they become fully efficient very near the online threshold at which they operate. Here, 'efficiency' is defined as the ability for the trigger to make a decision to record the event when there is actually an object satisfying its requirements. A trigger is 100% efficiency if it has a threshold of 10 GeV and it fires for every single lepton with a p_T at or above 10 GeV. The maximum attainable trigger efficiency is never exactly 100% due to non-100% coverage of the trigger system, however. Examples of the trigger efficiencies, measured in data and in MC, are shown in Figure 7.2 which shows the trigger efficiency turn-on curves for representative electron and muon triggers. It can be seen, for example, that the electron trigger efficiency reaches a plateau very near its configured threshold of 28 GeV. As described in Section 3.2.5, the trigger system has two-levels: Level 1 (L1) and HLT. The efficiencies for the trigger at L1 and HLT differ, primarily due to the lower lepton momentum resolution achieved at L1. This can be seen in the right side of Figure 7.2, which shows the trigger efficiency turn-on also for the L1 trigger that seeds the HLT trigger. The rise in efficiency, with respect to the offline reconstructed object's p_T, is shallower at L1 than at the HLT as a result of the poorer muon momentum resolution at L1.

As can be seen in the left side of Figure 7.2, the MC description of the trigger response is fairly good. However, analyses typically require that their offline reconstructed objects on which they base their trigger strategy have p_T values that lie on the trigger efficiency plateau. This is because the trigger efficiency plateau represents a region in which the trigger response is stable (i.e. unchanging) in both data and MC and can therefore be easily calibrated in the offline analysis with scale-factors that account for differences in the measured trigger efficiencies in data and in MC simulation. The description in the MC of the region at the trigger threshold, which has non-vertical rise, is potentially difficult to model accurately and therefore the calibrations are not typically derived for these regions in which the efficiency is not at its plateau. This is illustrated in Figure 7.3, showing the typical case of a stable and flat response in both the observed data and MC in the region of the trigger efficiency plateau but a potentially difficult to characterise response during the turn-on phase of the trigger efficiency. Given, however, the relatively good momentum resolution for leptons at L1, the offline p_T requirements on leptons can be kept very near the online p_T thresholds, typically within 1 GeV or so. Triggers based on jets or on the missing transverse momentum, however, typically require offline p_T requirements that are significantly higher than the online thresholds due to the fact that the online reconstruction of these objects in the trigger is not at the level of precision attainable in the offline reconstruction. This is illustrated by the trigger efficiency curves shown in Figure 7.4, showing the efficiency curves for triggers based on the reconstruction of the missing transverse momentum. These $E_{\mathrm{T}}^{\mathrm{miss}}$-based triggers all have online thresholds near 80 GeV, but they do not reach their efficiency plateau until offline $E_{\mathrm{T}}^{\mathrm{miss}}$ values nearing 250 GeV.

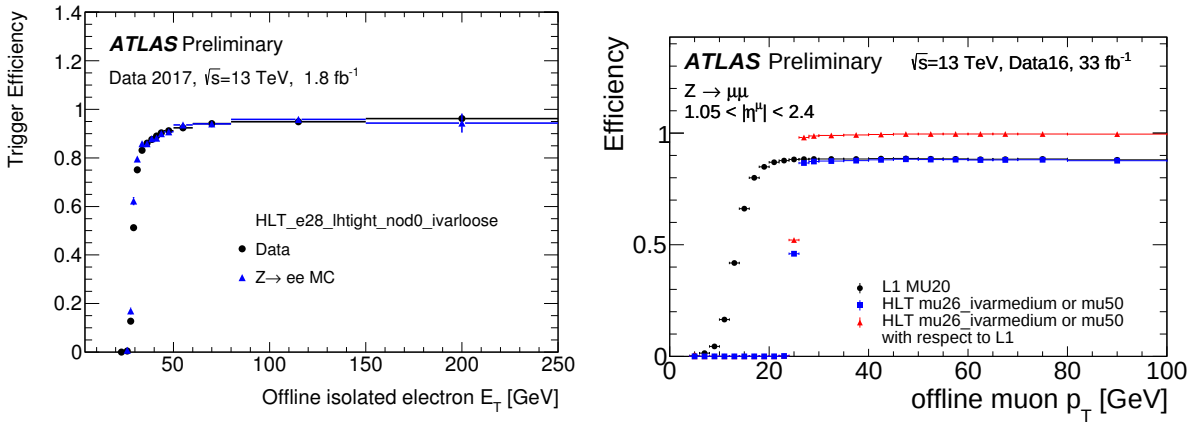

Figure 7.2: Figures showing the measured trigger efficiency as a function of the associated offline object for a representative electron trigger (***left***, from Ref. [168]) and muon trigger (***right***, from Ref. [169]).

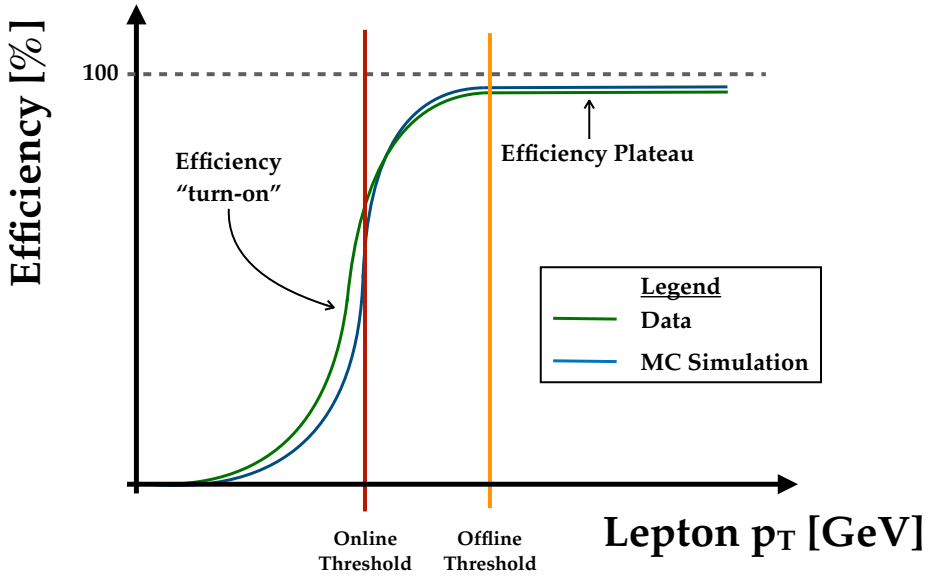

Figure 7.3: Cartoon illustrating the principle of a trigger efficiency 'turn on' curve. The efficiency for the lepton to fire the trigger, as a function of the p_T of the offline object, is plotted as a function of the offline object's p_T. At the online level, given the typically poorer lepton momentum resolution, there is generally not a perfectly sharp (i.e. vertical) turn on at the p_T threshold of the trigger. Instead there is an 'S'-curve, with the efficiency increasing with a steep slope until it reaches a point where it flattens out. This latter point is referred to as the trigger efficiency 'plateau'. Offline analyses typically apply offline p_T requirements on their objects such that they are always on the plateau of the associated trigger used for event selection.

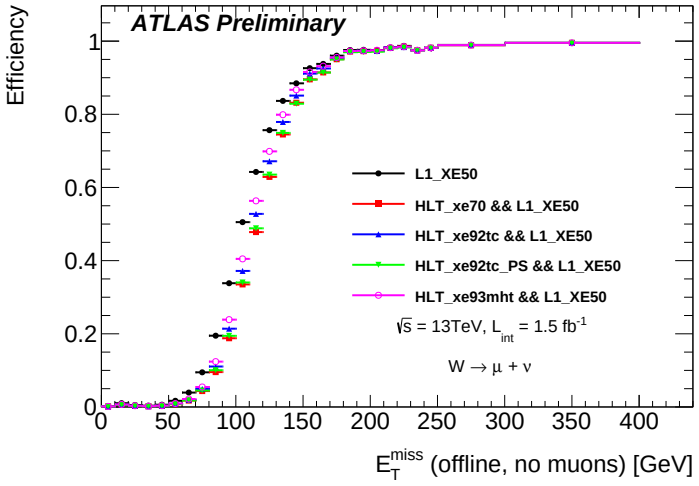

Figure 7.4: Trigger efficiency turn-on curve for typical triggers based on the missing transverse momentum. Figure taken from AtlasPublic/MissingEtTriggerPublicResults.

7.1.3 Background Estimation and the Control Region Method

The general principle behind searches for new physics is to define a SR, or a set of SRs, and then make predictions about how the signal and background behave therein. Such predictions can then be compared to the data actually recorded by the ATLAS detector and the statistical procedures described in Section 7.4 can be used to make statements about whether or not — or to what degree — the data is likely to contain the specified signal. The emphasis, then, in physics analyses is on the understanding and precise estimation of the backgrounds. Without being able to properly estimate the contribution of the background processes to the events observed in the SRs, well-defined predictions cannot therein be made, resulting in ineffective analyses.

The process of estimating the backgrounds in an analysis' SRs is aptly referred to as *background estimation*. There are many background estimation methods that are used. There exist general background estimation techniques, applicable to a wide range of SM processes, as well as more dedicated estimation techniques that are specific to a smaller subset of SM processes. Most rely on the MC simulation of the SM processes, either as the primary source of providing the prediction of a given SM process in an analysis' SR(s) or secondarily, as a means of providing a cross-check on or input to a prediction obtained using the observed data obtained in auxiliary measurements as the primary source. The high levels of accuracy imposed upon the ATLAS MC simulation infrastructure is derived from the large and dominant role that the MC simulation plays in the background estimation procedures in almost all analyses performed by the ATLAS experiment.

The same background estimation strategy is used in each of the analyses to be presented in the present thesis and is as follows. Once the SRs designed to capture the sought-for signal are defined, MC simulation can be used to determine the overall contribution of all SM background processes. In such a way, the MC simulation can be used to understand which SM processes in the SR(s) are dominant and which are sub-dominant. The former are those whose relative contribution to the total background prediction are large and the latter are those for which this quantity is small or negligible.

As illustrated in Figure 7.1 for the one-dimensional example, SR(s) are typically defined by events populating the *tails* of discriminating observables or probe extreme regions of phase space. Such regions typically exhibit low cross-section (background rates) and are regions of phase space for which the underlying theoretical inputs to the MC may be less well-understood as compared to the bulk of the phase space. The MC simulation by itself, therefore, may not be able to adequately describe the background processes in the SRs defined in these regions, failing to describe either the overall cross-section of specific processes or the actual shape of the discriminating observables' distributions therein. The former results in a failure in the overall predicted *rate*, or normalisation, of a specific backgrounds' contribution to the event yields in the regions and the latter results in a failure in predicting the background process' *acceptance*. In order to increase the confidence in the

background estimates in such SR(s), the analyses in the present thesis make use of the so-called *control region method*. This method is characterised by defining a (set of) region(s) in which there is (are) high purities of the dominant background process(es). These regions are referred to as *control regions* (CRs) and are ideally defined using the same basis of observables used to define the analysis' SRs. The CRs are defined to be orthogonal to the SRs, meaning that no events that satisfy the requirements of the SRs populate the CRs. The observed data in the CRs, which are enriched in a specific background process, are used to derive factors that correct the cross-section predictions of the MC estimates of the dominant background processes for which the CRs are defined. The per-process normalisation corrections, μ_p, can essentially be thought of as those factors that adjust the process' normalisation in such a way as to cover any discrepancy between the observed data yield and MC prediction for the process in question:

$$\mu_p = \frac{N^{\mathrm{CR}}_{\mathrm{data}} - \sum\limits_{\substack{i \\ i \neq p}} N^{\mathrm{CR}}_{\mathrm{MC},\,i}}{N^{\mathrm{CR}}_{\mathrm{MC},\,p}}, \tag{7.1}$$

where 'p' indicates the process for which the CR is defined, $N^{\mathrm{CR}}_{\mathrm{data}}$ is the observed data yield in the CR, and $N^{\mathrm{CR}}_{\mathrm{MC},\,j}$ is the predicted yield in MC for the background process j. If there is more than one process for which a normalisation correction factor is being derived, and therefore more than one CR, the normalisation factors are constrained by the process' contribution across all CRs in which it is present and the expression in Equation 7.1 is expanded into a system of equations,

$$\begin{aligned} N^{\mathrm{CR1}}_{\mathrm{data,sub}} &= \mu_i N^{\mathrm{CR1}}_i + \mu_j N^{\mathrm{CR1}}_j + ... \\ N^{\mathrm{CR2}}_{\mathrm{data,sub}} &= \mu_i N^{\mathrm{CR2}}_i + \mu_j N^{\mathrm{CR2}}_j + ... \\ &\vdots \end{aligned} \tag{7.2}$$

where $N^{a}_{\mathrm{data,sub}}$ is the observed data yield in the region a with the MC predictions for those processes not having a dedicated CR subtracted (analogous to the numerator appearing in Equation 7.1), μ_p are the normalisation factors for each process being solved for, and N^{a}_{p} are the MC predictions for process p in region a. Each process' dedicated CR ideally exhibits both a high purity of the given process and a relatively large number of events[3] and therefore is the only CR that has any real constraining power on the process' normalisation correction factor. If this is true, the expression in Equation 7.1 generally holds true for each process with a CR, even in the case of multiple CRs and normalisation factors.

As mentioned above, the CRs are ideally defined using the same basis of kinematic observables as used in the definition of the SR. When this is the case, it is more likely that the constructed CRs probe a similar kinematic phase space as that of the SRs. It is important that the CRs are

[3] A too large difference in the numbers of events observed in the CRs, as compared to the SRs, may indicate that the CRs are kinematically very different from the SRs, however.

kinematically similar to the SRs so that the correction factors derived in them are representative of the SRs; that is, that the underlying root cause of the need for the correction is the same in both the CRs, where the corrections are derived, and the SRs, where the corrections are applied. If an SR requires high numbers of jets (high event activity), for example, but the CR is defined to have zero jets (low event activity) then any normalisation correction derived in the CR may be correcting for physics effects that are not relevant to the phase space probed by the SR. In such a case, extrapolation uncertainties will generally be incurred in the final background estimate in the SR.

In addition to the CRs, so-called *validation regions* (VRs) are typically defined. The VRs are typically kinematically more similar to the SRs than the CRs, while still maintaining orthogonality between the CRs and the SRs. VRs are defined for each CR and allow for one to validate the extrapolation of the CR-derived normalisation correction for each process in a region more similar to the SR. As they are kinematically closer to the SRs, VRs are generally less pure in the specific process for which they are defined, and will also have generally fewer events, as compared to the associated CRs. The validation is done by comparing the post-corrected MC prediction of the backgrounds to the observed data in the VRs, ensuring that both the overall normalisation of the backgrounds agrees with the observed data as well as the overall shape of the relevant observables used in the definition of the SRs. The relationship between the CR, VR, and SR is illustrated in the one-dimensional case in Figure 7.5.

When constructing a set of CRs and VRs, it is important to do so using the right set of observables out of the total basis of observables from which the SRs are defined. It is important that the shape as predicted by the simulated background process for which the normalisation correction is being derived reproduces that of the observed data. If this is not the case, then the extrapolation from the CR to the SR suffers. This is illustrated in Figure 7.6 for two scenarios in which the MC-based background prediction of the shape of the observable used for defining the various regions both agrees and does not agree with the observed data. If the MC simulation for the specific background for which a normalisation correction is being derived has monotonic shape mis-modellings, as in the case of Scenario B in Figure 7.6, the normalisation correction will generally not be applicable in the SR and may lead to SR background estimates with false over- or under-predictions of the data. In the latter case, a false excess in data may be observed and lead to mis-statements about the likelihood of the existence of new physics in the SRs. In the former case, a false over-prediction will lead to too-prematurely excluding the possibility for new physics to arise when it may in fact exist. For this reason, dedicated studies on the dependence of the derived normalisation corrections on the set of observables used to define the CRs and VRs should generally be made so that the analysis avoids these susceptibilities to shape mis-modellings in the MC. Of course, this can become challenging in the general case where the SRs are defined using a large basis of potentially correlated observables and/or when one wishes to define several CRs to correct multiple processes' normalisations.

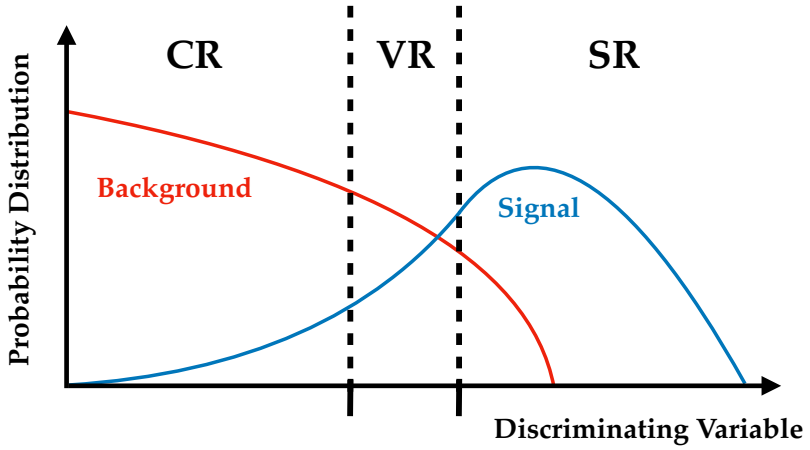

Figure 7.5: Illustration of the control region method, in the one-dimensional case analogous to that presented in Figure 7.1. The control region (CR) is pure in the background process but is defined kinematically alongside the signal region (SR). A validation region (VR), ideally still with high background purity, is defined between the CR and SR and is used to validate the extrapolation of the background estimate from the CR to the SR. The y-axis represents the probability distribution of the background and signal processes, not their absolute yield for a given range of the quantity on the x-axis.

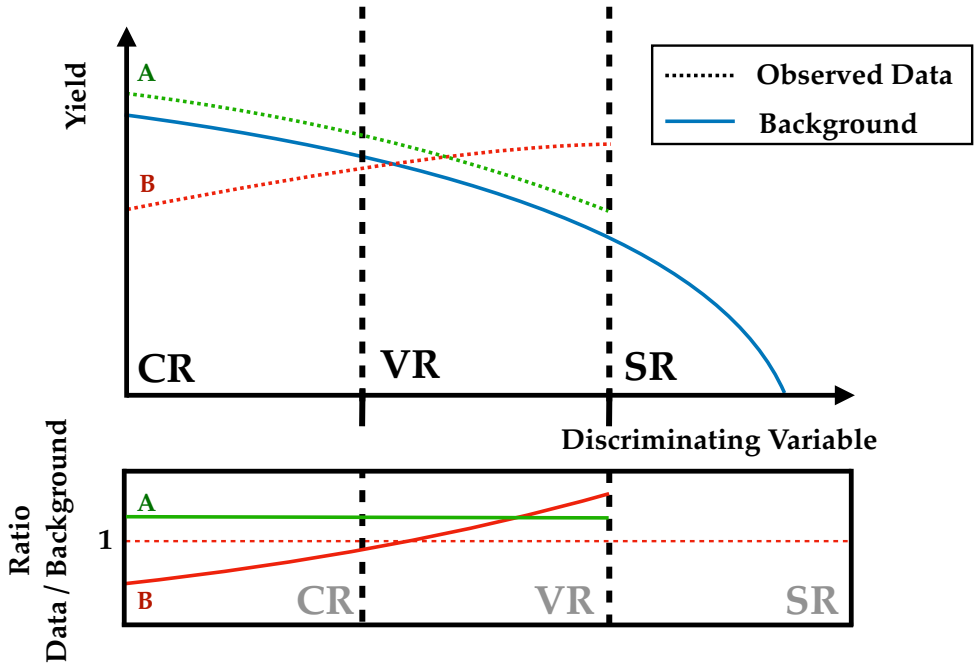

Figure 7.6: Illustration of CR extrapolation scenarios in the control region method. In Scenario A (green data) the predicted shape of the discriminating variable used to define the CR, VR, and SR agrees well with the observed data in both the CR and VR, as seen by the flat data-to-background ratio in the bottom. In Scenario B (red data) the predicted shape of the discriminating variable differs with respect to that of the observed data, leading to an observed slope in the data-to-background ratio. In Scenario A, the normalisation correction derived for the background process in the CR will be well extrapolated into the VR and gives confidence in its applicability in the SR. In Scenario B, the CR-derived normalisation factor for the background process will pull the background prediction in the wrong direction when extrapolated to the VR, making the data-to-background agreement even worse and reducing its applicability in the SR.

7.2 Estimation of Sources of Fake and Non-prompt Leptons

Despite both the high levels of accuracy achieved by the ATLAS simulation infrastructure and the lepton reconstruction and identification algorithms described in Chapter 5, sources of misidentified reconstructed leptons still exist and lead to an additional source of backgrounds to the analyses discussed in Chapters 8 and 9. These background sources of leptons are broken down into two categories:

- **Fake leptons**: Cases in which signals in the ATLAS detector are selected as being leptons when in fact there is no real lepton present
- **Non-prompt leptons**: When real, genuine leptons are identified but they are not leptons originating from the primary pp hard-scatter interaction process of interest

In the subsequent discussion, the term 'fake' will be used in reference to the two categories listed above, unless specified otherwise.

The contribution of backgrounds leading to sources of fake leptons are generally predicted using methods based on the observed data — referred to as 'data-driven' methods — and arise from various sources and mechanisms. The sources of fake leptons will be described in Section 7.2.1, separately for electrons and muons. Section 7.2.2 provides some reasoning for why a data-driven approach is generally taken for estimating these backgrounds. Sections 7.2.3 and 7.2.4 go on to describe the two data-driven approaches taken in the analyses to be presented in this thesis for estimating fake lepton contributions: the so-called 'Matrix Method' and the 'Same-sign Extrapolation Method', respectively.

7.2.1 Sources of Fake Leptons

The types and sources of fake leptons generally have different experimental signatures than those leptons that genuinely originate from the pp hard-scatter. However, due to the non-perfect lepton identification and isolation algorithms, such sources are able to contaminate the various regions of an analysis. The rates of contamination are generally quite low for the analyses to be presented, but their inclusion in the background estimates of the analyses has measurable consequences nevertheless.

The analyses to be presented in the current thesis make use of b-tagging algorithms to identify jets originating from b-hadrons. Fake leptons, both electrons and muons, can originate from the semi-leptonic decays of b- and c-quarks within these b-tagged jets, following $b \to \ell$ or cascade-type $b \to c \to \ell$ decays of the B hadrons within the jets. The leptons resulting from such decays are

typically embedded within or very close to the originating reconstructed jet object and the lepton isolation requirements are intended to reduce this type of background. The subsequent paragraphs will describe additional sources of fake electrons and muons, which generally differ between the two lepton species.

Sources of Fake Electrons

As described in Section 5.2.1, electrons are reconstructed based on the presence of well-reconstructed tracks in the ID matched to deposited energy clusters in the EM calorimeter. Light-flavor jets, originating from the production of light quarks (u, d, s), or gluon jets, which are associated with a large number of tracks due to their increased radiation pattern, are able to fake electrons as they leave tracks in the ID as well as subsequent energy depositions in both the EM and hadronic calorimeters. This background, due to mis-identified jets, is typically suppressed by the use of lepton isolation and by jet shower-shape information used in the electron identification: the hadronic shower shapes and radial extent differ with respect to the electromagnetic shower produced by a genuine electron.

An additional large source of fake electrons is due to photon conversion processes, $\gamma \to e^+e^-$, and other electromagnetic scattering processes that happen as a result of detector material interactions. These processes leave both tracks in the ID and electromagnetic energy depositions in the EM calorimeter which are difficult to distinguish from genuine electrons. Neutral hadron decays, such as the $\pi^0 \to e^+e^-\gamma$ Dalitz decay, also lead to electron-like signatures. This decay of the π^0 only has a branching fraction of just over 1% [54], but given the large production of π^0 states in the pp collision this has the potential to be a relevant source of fake electrons. These electromagnetic sources of fake electrons are distinguished by their generally larger impact parameters relative to genuine prompt electrons.

Sources of Fake Muons

As described in Section 5.2.2, muons are primarily reconstructed via the combination of tracking information provided by the ID and MS, and, generally speaking, they should be the only particle species to reach the MS. In addition to the semi-leptonic decays of heavy-flavored jets described above, however, there are several sources of fake muons. Highly energetic jets can have elongated shower profiles that reach the outer radii of the hadronic calorimeter, with a non-zero chance of exiting the calorimeter and resulting in particle leakage into the MS. Such cases are referred to as calorimeter punch-through, and have been illustrated in Figure 5.15. Punch-through particles can leave signatures similar to charged muons whose subsequent MS tracks are associated with a track in the ID, leading to a reconstructed combined muon faking a genuine muon. An additional

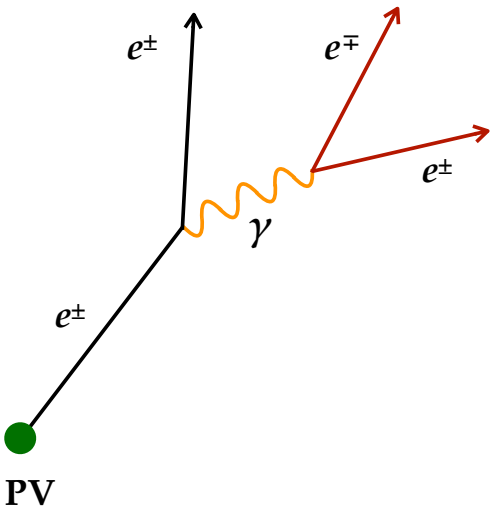

Figure 7.7: Illustration of electron bremsstrahlung radiation with photon conversion exhibiting 'trident' topology.

source of fake muons come from the in-flight decays of charged hadrons, such as the $K^{\pm}$ and $\pi^{\pm}$, that can decay to $\mu^{\pm}\nu$. These sources of non-prompt muons typically result in low-p_T muons and are characterised by combined tracks with exhibiting a kink topology, briefly mentioned when discussing the muon combined reconstruction in Section 5.2.2. The p_T dependence of muon cross-sections for various sources of muons (prompt, fake, and non-prompt) is shown in Figure 7.8, along with an illustration of the kinked-track topology characterising the in-flight decays of Kaons and charged pions.

7.2.2 The Need for a Data-driven Approach

In the analyses to be presented in Chapters 8 and 9, relatively tight identification working points are used for electrons and muons. As a result, the contamination of fake leptons in these analyses is relatively minor. Although small, their contamination does have measurable effects and so their contribution must be accounted for in order to achieve accurate estimates of the backgrounds in each analysis.

Several methods exist to estimate the background rates arising from the sources of fake leptons, those relying on data-driven methods or those based entirely on the MC simulation. Relying on the MC simulation of these sources of fake leptons, described in previous sections, means to rely entirely on the GEANT4 simulation of the ATLAS detector and on the MC generation and showering processes to accurately predict the rates of these processes. There are several problems with this approach and they are (non-exhaustively) as follows. Given the very small region of phase space

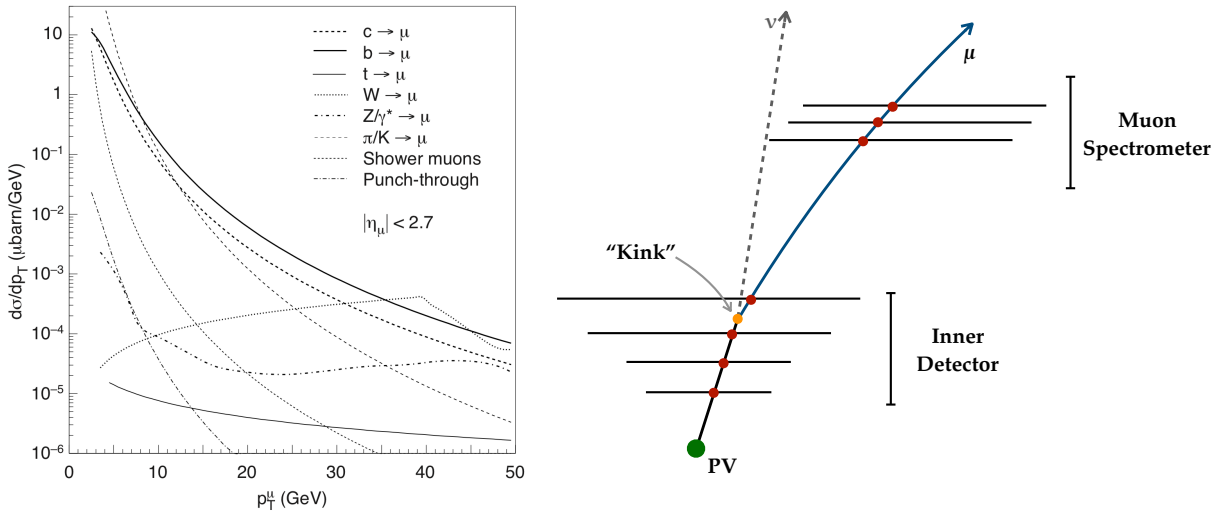

Figure 7.8: ***Left***: Transverse momentum dependence of muon cross-sections for muons originating from various prompt and non-prompt sources. Figure taken from Ref. [87]. ***Right***: Illustration of a reconstructed non-prompt muon resulting from a kinked-track topology. A produced $K^{\pm}$ or $\pi^{\pm}$ is produced and decays in-flight to a muon and muon-neutrino. The point at which the hadron decays is indicated by the yellow dot. The red circles indicate detector hits in the ID and MS layers indicated by the horizontal black lines.

being probed by the analysis, the number of MC events needed to appropriately sample the sources of production of fake leptons as described above would be prohibitively large if a statistically relevant sample is desired. An accurate prediction of the production rates of several of these fake lepton sources would require an accurate underlying theoretical model of many processes, such as heavy-flavor jet fragmentation, which is challenging. Additionally, many sources of fake leptons arise as a result of detector material interactions or as a result of subtle and difficult-to-model failure modes of the detector response. The accurate prediction of the rate of electron bremsstrahlung and photon conversions, for example, requires high levels of precision in the simulation and measurement of the active and passive material in the ATLAS detector and cavern, which is not necessarily possible. The rates of jets being mis-identified as electrons and jet punch-through, for example, require that the MC simulation of the calorimeter response and shower evolution are accurately modelled. The MC simulation is not expected to perform to the degree at which these subtle, and comparatively rare, effects are accurately predicted. For this reason, data-driven approaches are typically taken for estimating the background rates of these fake lepton sources. In the analyses to be presented, two data-driven approaches are taken. In the search described in Chapter 8, the Matrix Method is used. In the search described in Chapter 9, the Same-sign Extrapolation Method is used. These methods are introduced in Sections 7.2.3 and 7.2.4, respectively.

7.2.3 The Matrix Method

The Matrix Method, discussed thoroughly in Ref. [170], is one of the most common methods used in ATLAS analyses for estimating backgrounds due to processes containing fake leptons. It is characterised by the definition of two levels of lepton selection:

> **Tight Leptons**: Those leptons passing all reconstruction, identification, and isolation criteria as the leptons used in the final analysis' results
>
> **Loose Leptons**: Leptons requiring similar selections as the Tight leptons but typically with either, or both, identification and isolation criteria relaxed

The Tight leptons are a subset of the Loose, by definition. In the analysis described in Chapter 8, the Loose leptons are defined by loosening only the lepton identification working points. Generally speaking, both samples of Loose and Tight leptons will contain both fake and real leptons.[4] The Matrix Method consists of measuring a set of efficiencies: the *real* (*fake*) *efficiencies*, ε_r (ε_f), defined as the efficiency for a real (fake) electron or muon that satisfies the Loose selection criteria to also satisfy the Tight selection criteria. This is illustrated in Figure 7.9. As illustrated, both the Loose and Tight lepton samples will contain contributions of both fake and real leptons. The Matrix Method can be generalised to final states with any number of leptons. In the discussion to follow, we will discuss that of final states with two leptons: the dilepton Matrix Method.

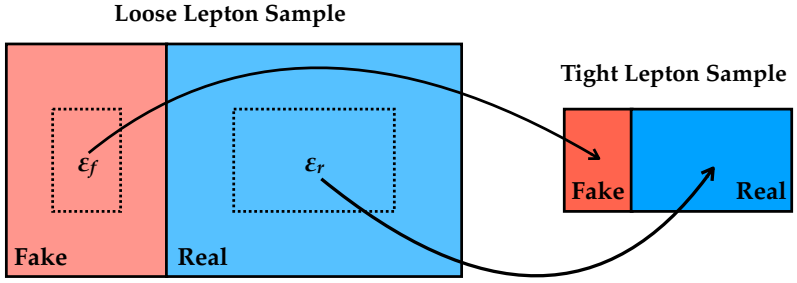

Figure 7.9: Illustration of the loose and tight lepton samples used in the Matrix Method.

The real efficiencies, ε_r, are measured in data using Z-boson tag-and-probe methods, requiring the probe lepton to satisfy the Tight lepton selection and to be matched to the trigger. The probe lepton must satisfy the Loose lepton selection. The fraction of Loose probe leptons to pass the Tight lepton selection then gives a measure of ε_r.

The fake efficiencies, ε_f, are measured in data using events with different-flavor leptons, where one lepton is an electron and the other is a muon, that have the same electric charge. A tag-and-

[4] Genuine, prompt leptons originating from the pp hard-scatter interaction point are typically referred to as 'real' leptons in order to distinguish them, semantically, from fake and non-prompt leptons.

probe method similar to that used in the measurement of ε_r is used and relies on the fact that a comparatively small amount of SM processes can result in same-sign and different-flavor events. Therefore, when a probe satisfies the Tight selection it is very likely to be the case that the probe lepton is a fake lepton. An additional component of the ε_f is measured by additionally requesting that there be at least one b-tagged jet in the same-sign, different-flavor selection. This allows for ε_f to be measured in a region enriched in fake leptons originating from semi-leptonic decays of heavy-flavor jets. The various measurements of ε_f are combined following an averaging scheme, weighted according to the composition of fake lepton sources expected to contaminate the SRs.

The real and fake efficiencies are not global quantities. Instead, they are measured as a function of both the lepton p_T and η such that they may track the effects of changes in detector response and material interaction over the p_T and η ranges relevant to the leptons used in the analysis.

Once ε_r and ε_f are obtained, the fake lepton background can be obtained by inverting the equation relating the measured quantities (Tight versus Loose), taken from the observed data, in terms of those that we want to know (Fake versus Real):

$$\begin{pmatrix} N_{TT} \\ N_{TL} \\ N_{LT} \\ N_{LL} \end{pmatrix} = M \begin{pmatrix} N_{LL}^{RR} \\ N_{LL}^{RF} \\ N_{LL}^{FR} \\ N_{LL}^{FF} \end{pmatrix} \tag{7.3}$$

where in the sub- and super-scripts, the first (second) index refers to the leading (sub-leading) lepton. The sub-script 'T' ('L') refers to the lepton passing the Tight (Loose) lepton selection. The super-script 'R' ('F') indicates whether or not the lepton is a real (fake) lepton. For example, the quantity N_{LL}^{FR} is the number of events in which the leading lepton in the sample of Loose leptons is fake and the sub-leading is real. The matrix M is given by,

$$M = \begin{pmatrix} \varepsilon_{r,1}\,\varepsilon_{r,2} & \varepsilon_{r,1}\,\varepsilon_{f,2} & \varepsilon_{f,1}\,\varepsilon_{r,2} & \varepsilon_{f,1}\,\varepsilon_{f,2} \\ \varepsilon_{r,1}\,\overline{\varepsilon_{r,2}} & \varepsilon_{r,1}\,\overline{\varepsilon_{f,2}} & \varepsilon_{f,1}\,\overline{\varepsilon_{r.2}} & \varepsilon_{f,1}\,\overline{\varepsilon_{f,2}} \\ \overline{\varepsilon_{r,1}}\varepsilon_{r,2} & \overline{\varepsilon_{r,1}}\,\varepsilon_{f,2} & \overline{\varepsilon_{f,1}}\,\varepsilon_{r,2} & \overline{\varepsilon_{f,1}}\,\varepsilon_{f,2} \\ \overline{\varepsilon_{r,1}}\,\overline{\varepsilon_{r,2}} & \overline{\varepsilon_{r,1}}\,\overline{\varepsilon_{f,2}} & \overline{\varepsilon_{f,1}}\,\overline{\varepsilon_{r,2}} & \overline{\varepsilon_{f,1}}\,\overline{\varepsilon_{f,2}} \end{pmatrix} \tag{7.4}$$

where the notation $\overline{\varepsilon}$ indicates $(1-\varepsilon)$.

The number of events with double-fake and single-fake leptons satisfying the analysis' Tight selection (N_{TT}^{FF} and $N_{TT}^{RF} + N_{TT}^{FR}$, respectively) can then be obtained from the number of events with double-fake and single-fake leptons satisfying the Loose selection (N_{LL}^{FF} and $N_{LL}^{RF} + N_{LL}^{FR}$,

respectively) through inversion of Equation 7.3 and noting the following relations:

$$N_{TT}^{RR} = \varepsilon_{r,1}\,\varepsilon_{r,2} \times N_{LL}^{RR} \tag{7.5}$$

$$N_{TT}^{RF} = \varepsilon_{r,1}\,\varepsilon_{f,2} \times N_{LL}^{RF} \tag{7.6}$$

$$N_{TT}^{FR} = \varepsilon_{f,1}\,\varepsilon_{r,2} \times N_{LL}^{FR} \tag{7.7}$$

$$N_{TT}^{FF} = \varepsilon_{f,1}\,\varepsilon_{f,2} \times N_{LL}^{FF}. \tag{7.8}$$

The quantities appearing on the right-hand-side of Equations 7.5-7.8 depend entirely on the observed data. The number of events in the analysis' Tight selection that have at least one fake lepton is then given by the sum $N_{TT}^{RF} + N_{TT}^{FR} + N_{TT}^{FF}$. This gives a total integrated yield for the fake background contribution in the analysis' Tight selection; however, from Equations 7.6-7.8 a set of per-event weighting factors ('fake weights'), depending only on the $\varepsilon_r(p_T, \eta)$ and $\varepsilon_f(p_T, \eta)$ efficiency factors for the two leptons, can be defined. These fake weights allow for kinematic distributions of the fake lepton background sources to be populated by appropriately applying them to events in the data sample consisting of leptons satisfying the analysis' Loose selection.

7.2.4 Same-sign Extrapolation Method

The Same-sign Extrapolation Method is another method for estimating the numbers of events with fake and non-prompt leptons. This method is well described in Refs. [171, 172]. The method is tailored for dilepton analyses that require the two leptons to have opposite electric charge and takes advantage of the fact that the production mechanisms for the fake leptons described above are generally uncorrelated with the charges of leptons in the event. This means that the rates of the various sources of fake leptons will generally be the same in the sample of oppositely-charged dilepton events and in the sample of dilepton events in which the leptons have the same charge. The former are referred to as opposite-sign (OS) events and the latter as same-sign (SS) events. This assumption is not entirely correct, however, since the underlying sources of the fake and/or non-prompt leptons are not all charge symmetric.[5] For example, a dominant source of mis-identified leptons in dileptonic events arises from the production of semi-leptonically decaying top-quark pairs, in which the sub-leading reconstructed electrons arise from a mis-identified jet. This process is charge-symmetric since converted photons produced in jets are equally likely to give rise to a reconstructed e^+ or e^- and are uncorrelated to the sign of the real lepton. In trident decays whereby bremsstrahlung photons (radiated from the lepton children of the top-quarks) undergo conversion processes, however, there is a partial charge correlation of the reconstructed (non-prompt) electron with the charge of the original lepton, and hence conversion processes contribute more to the OS sample of events. These two cases are illustrated in Figure 7.10. The lepton definitions within the

[5] 'Charge-symmetric' means that the dilepton final state of a given process may be either OS or SS, with equal probability of each case occurring.

OS and SS event samples are the same. That is, they use the same reconstruction, identification, and isolation working points. This is compared to the Matrix Method (Section 7.2.3), in which the two samples of leptons used in the technique's implementation differ by their lepton definitions.

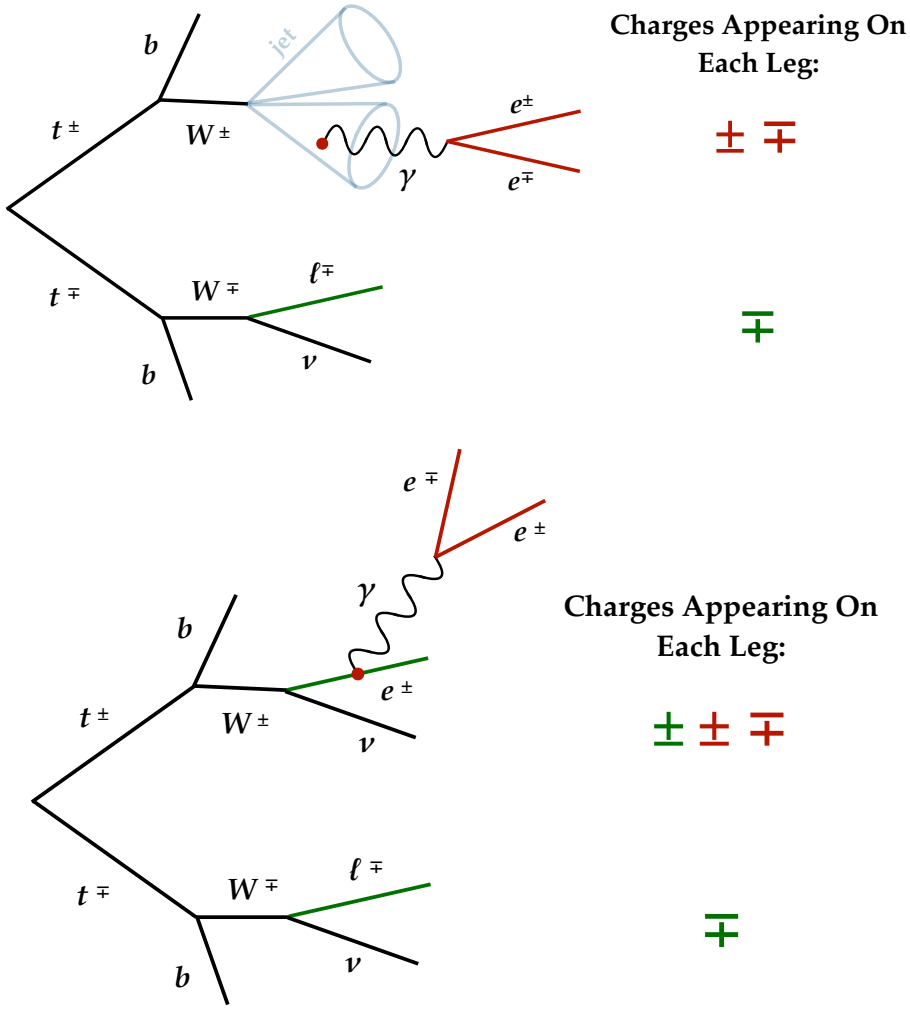

Figure 7.10: Illustration of charge-asymmetry in fake lepton production arising in top-quark pair production events, leading to the rate of photon conversion sources of fake electrons being larger in OS dilepton events. ***Top***: Shower photons arising from decays within one of the jets in semi-leptonic decays of top-quark pairs may convert to an electron-positron pair, leading to one of the jets being reconstructed as an electron. Looking at the charge possibilities of each side of the top-quark pair decay, the event is equally likely to be classified as either an OS or SS event. ***Bottom***: Trident events arising in dileptonic top-quark pair production events can lead to a non-prompt electron from the photon conversion being selected as one of the event's candidate leptons. The overall charge possibilities of the lepton charges on the side of the trident decay are correlated with the charge of the initial lepton. Looking at the charge combinations possible between both sides of the top-quark pair decay, the event is more likely to be classified as an OS event.

The method also assumes that the rate of dilepton events in which *both* leptons are fake is negligible, and that the composition of the fake background contributing to the dilepton final state is composed of events in which one of the leptons is real. In the majority of cases, the sub-leading lepton is the fake lepton. For this reason, when using the MC simulation to gain predictions of the composition of the fake backgrounds, the MC events in which only one of the leptons is fake are considered. This will become clear in the discussion to follow and in the specific implementation described in Chapter 9.

The general method works as follows. Since, as described above, the sources of the various sources of fake backgrounds in the OS and SS samples of events are generally the same, the method relies on using the sample of SS events to provide a template of the fake backgrounds to be used in the OS selections. The contribution of fakes to each of the regions (CR, VR, or SR) in the analysis is estimated by subtracting the prediction of the prompt (real) SM backgrounds from the observed data in the associated SS selections, defined similarly to the OS selections used in the analysis but with the charge requirements inverted. The ratio of the number of OS to SS events with fake leptons, $f^{SS\to OS}$, is taken entirely from MC and is applied to the SS data that has had the prompt-MC contribution subtracted in order to extrapolate this number to the OS regions. This SS extrapolation is described by Equation 7.9:

$$\begin{aligned} N_{\text{OS}}^{\text{fake}} &= f^{SS\to OS} \times N_{\text{SS}}^{\text{fake}} \\ &= \frac{N_{\text{MC,OS}}^{\text{fake}}}{N_{\text{MC,SS}}^{\text{fake}}} \times (N_{\text{data,SS}} - N_{\text{MC,SS}}^{\text{real}}) \end{aligned} \tag{7.9}$$

In addition to providing the overall yields of the fake backgrounds in a given region, the method described by Equation 7.9 provides the means to inspect the kinematics of the predicted fake backgrounds by simply computing Equation 7.9 on a bin-by-bin basis when populating histograms of kinematic observables.

Further details on the implementation of the Same-sign Extrapolation Method, described by Equation 7.9, will be given in Chapter 9. In particular, the implicit sensitivities to the MC simulation will be described, as well as additional extrapolations specific to the analysis.

7.3 Systematic Uncertainties

There are many sources of systematic uncertainty ('systematics') that affect the results of the analyses to be presented in Chapters 8 and 9. The considered sources of systematic uncertainty are listed in Table 7.1. There are uncertainties related to the overall event (e.g. the uncertainty on the luminosity measurement), on the reconstruction and identification of the physics objects described in Section 5, and in the MC simulation of both the SM background and signal processes.

Brief descriptions of the sources of systematic uncertainty appearing in Table 7.1 are given in Sections 7.3.1-7.3.4.

Table 7.1: Summary of the sources of systematic uncertainties affecting the measurements in the analyses discussed in Chapters 8 and 9. Sources of uncertainty in red (blue) pertain only to the search presented in Chapter 8 (9). Those in black are considered in both analyses. For the uncertainties related to the SM background modelling, it is indicated whether or not they are computed using the Transfer Factor Method (Section 7.3.2).

Event-level			
Luminosity Measurement			
Pile-up Modelling			
Object Reconstruction			
Jets	**Flavor Tagging**	**Leptons**	$E_{\mathrm{T}}^{\mathrm{miss}}$
Jet Energy Scale (JES)	b-tag Eff.	Reconstruction Eff.	Soft-term Resolution
Jet Energy Resolution (JER)	Mis-tag Eff.	Identification Eff.	Soft-term Scale
Pile-up Suppression (JVT)		Isolation Eff.	
		Trigger Eff.	

Background Modelling		
Souce of Uncertainty	**Affected Background Processes**	**Transfer Factor Approach?**
Hard-Scatter Generation	$t\bar{t}$, Wt, Z+HF	Yes
Fragmentation	$t\bar{t}$, Wt	Yes
Additional Radiation (ISR and FSR variation)	$t\bar{t}$, Wt	Yes
PDF Choice and Uncertainty	$t\bar{t}$, Wt, Z+HF, VV	Yes
Scale (μ_R, μ_F) Variations	$t\bar{t}$, Wt, VV, Z+HF	Yes
Cross-section Uncertainty	$t\bar{t}$, Wt	No
$t\bar{t}$ Interference Uncertainty	Wt	Yes
Prompt Subtraction & Fake Composition	Fake Lepton Background	No
SS-OS Extrapolation	Fake Lepton Background	No

Signal Modelling
Fragmentation
Scale (μ_R, μ_F) Variations
PDF Choice and Uncertainty
Cross-section Uncertainty

7.3.1 Experimental Uncertainties

Event-wide Uncertainties

Event-wide (process-independent) uncertainties affecting the overall normalisation of the processes relate to both the luminosity and pileup measurements. The uncertainty on the integrated luminosity used to normalise all MC simulated processes is derived following the methodology described in Ref. [173]. For the 2015+2016 (full Run 2) dataset, relevant to the analysis described in Chapter 8 (9), this uncertainty was found to be 2.1% (1.7%). The uncertainty on the luminosity measurement does not affect processes whose SR normalisation is constrained by data in dedicated CRs.

An uncertainty is considered on re-weighting the pileup distributions in the MC simulation. The re-weighting is applied in order to correct for the differences in the actual pile-up distributions observed in data and those assumed at the time of producing the MC simulation (Section 4.4).

Jet Uncertainties

The systematic uncertainties on reconstructed jet objects are related to the jet energy resolution (JER), jet energy scale (JES), and JVT. There are many sources of uncertainties related to the JES and JER, each related to a specific part of the JES and JER calibration measurements, as described in Section 5.3.2. They arise from the techniques and corrections derived in MC, including statistical, detector, modelling effects, jet flavor compositions, pileup corrections, and η-dependence effects. The effects of the JES and JER uncertainties are among the more dominant sources of uncertainty on the final analysis results in both analyses to be presented. Given the complexity of the JES and JER calibrations, there are nearly 100 components associated with their uncertainties that must be incorporated into an analysis' measurement uncertainty. In general, the leading uncertainty in searches such as those to be presented in the current thesis is statistical in nature — related to the limited statistics in data and in the MC simulation used for the final SR predictions — and not related to the systematic uncertainties. For this reason, the searches described in Chapters 8 and 9 used a reduced set of JES and JER uncertainties that are derived following a Principal Component Analysis (PCA) designed to capture only the dominant components — and their correlations — of the total set of JES and JER uncertainty components that are relevant to the phase space being probed by the analyses. In the analysis presented in Chapter 8 (9), this uncertainty reduction process leads to only 4 (34) separate components of the combined JES and JER systematic uncertainty.

Flavor Tagging Uncertainties

There are uncertainties in the jet flavor tagging efficiencies, as well as in the measured mis-tagging efficiencies associated with c- and light jets. They are a mixture of statistical, experimental, and modelling uncertainties incurred during the flavor tagging calibration procedures. The uncertainties enter into the analyses through their impact on the scale-factors, described in Section 5.3.4, that are applied in the analysis. Given the importance of b-tagged jets in the final states of the signal processes in the analyses described in Chapters 8 and 9, these uncertainties have non-negligible impact on the analyses' results.

Lepton Uncertainties

Uncertainties on the measurement of leptons correspond to the electron and muon reconstruction, identification, trigger, and isolation efficiencies in a manner similar to the flavor tagging in that systematic variations incurred in the associated scale-factor measurements are applied in the analysis. Additional uncertainties related to the lepton kinematics due to the resolution and scale of the electron (muon) energy (momentum) measurement are considered. The muon momentum measurement uncertainties are derived for both the ID and MS measurement of the combined muons used in the analyses.

Missing Transverse Momentum, $E_{\mathbf{T}}^{\mathbf{miss}}$

Systematic variations of the $E_{\mathrm{T}}^{\mathrm{miss}}$ are coherently incurred as a result of the systematic variations, described above, being applied to the objects provided as input to the $E_{\mathrm{T}}^{\mathrm{miss}}$ calculation: the leptons and jets. Additional uncertainties related to the scale and resolution of the soft-term of the $E_{\mathrm{T}}^{\mathrm{miss}}$ calculation are also considered. Given that the analyses considered in Chapters 8 and 9 are characterised by real sources of $E_{\mathrm{T}}^{\mathrm{miss}}$, the soft-term component plays a small role in the magnitude of the $E_{\mathrm{T}}^{\mathrm{miss}}$ and therefore its uncertainties have negligible impact on the analyses. Generally, the $E_{\mathrm{T}}^{\mathrm{miss}}$ uncertainties are sub-dominant.

Fake Lepton Estimate

As will be seen in Chapters 8 and 9, the overall contribution of the fake background processes to the analyses to be presented in this thesis is very small. As a result, the systematic uncertainties related to this background are almost always negligible in impact on the analyses' final results. However, given the subtle nature of these data-driven estimates, we describe the methods by which systematic uncertainties are derived for them. They are as follows:

Matrix Method (Section 7.2.3): There are three primary sources of systematic uncertainty ascribed to the fake estimate derived from the Matrix Method in the analysis described in Chapter 8. The first is related to the limited statistics in the region(s) used for the determination of the fake efficiencies, ε_r and ε_f. The second is related to the 'prompt subtraction': the evaluation of the portion of real events contaminating the regions in which the fake efficiencies are derived, which is evaluated using the MC simulation. The component of real lepton sources is varied by $\pm 30\%$ and the impact on the resulting fake efficiencies is propagated to to the final analysis. The third component is related to the compositional differences of the background sources leading to fake leptons in the region in which the fake efficiencies are derived and

in the regions (SRs) in which the fake estimate is applied. To assess the systematic related to compositional differences, alternative regions in which the leading components of the fake backgrounds (e.g. the heavy flavor component) are varied are defined and the impact on the measured fake efficiencies is used to define an uncertainty.

Same-sign Extrapolation Method (Section 7.2.4): As with the Matrix Method, there are three primary sources of uncertainty prescribed to the fake background estimate derived using the Same-sign Extrapolation Method. The first, as with the Matrix Method, is related to the prompt subtraction in which the rate of contamination of prompt processes is varied. In the analysis described in Chapter 9, the real contamination in the same-sign regions is varied by $\pm 50\%$, and the impact on the final fake background estimate is used to quantify an uncertainty. Following the studies in Refs. [171, 172], an additional uncertainty (on top of the statistical component) on the extrapolation from the SS to the OS regions is included by varying the $f^{SS \to OS}$ factors by $\pm 20\%$ and assessing the impact on the final fake background prediction. The statistical uncertainty related to the extrapolation over the d_{hh} discriminant, described in Chapter 9, is applied as an additional uncertainty on this fake background estimate. Uncertainties related to the composition of fakes in the SS and OS regions are found to be small, and are covered by the extrapolation factor uncertainties already described.

7.3.2 The Transfer Factor Method

For estimating the impact of modelling uncertainties on the SM backgrounds, described in Section 7.3.3, that have dedicated CRs to constrain their overall normalisation in the SRs, the so-called Transfer Factor (TF) Method is used. Since the purpose of the CRs is to constrain the process' normalisation using the observed data, we do not want the systematic variations to directly impact the processes' normalisations within the SRs. Instead, the impact of the systematic variations is assessed via their effect on the SM processes' acceptance in the CRs and SRs. If a given systematic variation for a given SM process affects the process in a coherent manner across both the CR and SR, then the resulting affect of the systematic variation on the analysis should be reduced. This will nearly be guaranteed if the phase spaces being probed by the CR and SR are similar. The more dissimilar the phase space being probed by the CR and SR, the larger the expected impact of a given systematic variation due to the larger kinematic extrapolation required. This can

conceptually be seen by considering Equation 7.10:

$$\begin{aligned} N_p^{\mathrm{SR}} &= \mu_p \times N_{p,\,\mathrm{MC}}^{\mathrm{SR}} \\ &= \left(\frac{N_{p,\mathrm{data}}^{\mathrm{CR}}}{N_{p,\,\mathrm{MC}}^{\mathrm{CR}}} \right) \times N_{p,\,\mathrm{MC}}^{\mathrm{SR}} \\ &= N_{p,\mathrm{data}}^{\mathrm{CR}} \times \left(\frac{N_{p,\,\mathrm{MC}}^{\mathrm{SR}}}{N_{p,\,\mathrm{MC}}^{\mathrm{CR}}} \right) \\ &= N_{p,\mathrm{data}}^{\mathrm{CR}} \times \underbrace{\tau_p}_{\substack{\mathrm{Transfer} \\ \mathrm{Factor}}} , \end{aligned} \tag{7.10}$$

where 'p' is the process for which the CR and normalisation factor are defined, μ_p is the CR-derived normalisation factor (c.f. Equation 7.1), $N_{p,\mathrm{data}}^{\mathrm{CR}}$ is the observed data in the CR with the MC simulation for all processes that are not process p subtracted, and $N_{p,\,\mathrm{MC}}^{\mathrm{SR}}$ is the MC-based SR prediction of process p. The quantity τ_p is the process' TF that extrapolates the observed data in the CR to the SR. It can be seen that if the MC simulation response for a given process for a given systematic variation is the same across both the CR and SR, that the TF will be unchanged as a result of the systematic variation and therefore the predicted contribution of this process in the SR will be unaffected by the systematic uncertainty. If kinematics differ across the CR and SR, the acceptance for a given process may vary in going from the CR and SR (or vice versa) and therefore such a cancellation is not likely to occur due to the larger extrapolation required.

It can be seen, then, that for processes whose SR normalisation is derived in dedicated CRs, that the TF appearing in Equation 7.10 quantifies the impact of acceptance variations between the CR and SR. Systematic variations of the SM backgrounds with dedicated CRs, then, are quantified by their impact on the resulting TF values. This is contrary to assessing their impact by measuring the change in a process' SR prediction by simply comparing the SR predictions before and after a given systematic variation is applied. The uncertainties ascribed to SM processes via the TF Method are computed as follows,

$$\Delta\tau = \frac{|\tau_{\mathrm{nominal}} - \tau_{\mathrm{variation}}|}{\tau_{\mathrm{nominal}}}, \tag{7.11}$$

where τ_{nominal} ($\tau_{\mathrm{variation}}$) is the TF computed using the nominal (systematically varied) prediction of the process. The quantity $\Delta\tau$ is then taken as a fractional uncertainty on the corresponding process' SR predicted yield.

Sources of uncertainty stated as following the TF approach in Table 7.1 are quantified following Equation 7.11. The others are assessed simply by taking the impact of the variation on the process' predicted yields in each of the regions appearing in the analyses.

7.3.3 Background Modelling Uncertainties

Uncertainties in the modelling of specific processes, SM or otherwise, are typically assessed by comparing the nominal MC simulation for the processes in question to that of an MC simulation with certain theoretical or phenomenological parameters varied. In this way, one can assess the impact of the underlying assumptions made in the MC simulation on the analyses' final results.

Top-quark Pair and Single-top Wt Production

The production of SM top-quark pairs, $t\bar{t}$, is by far the most dominant SM background in both of the analyses described in Chapters 8 and 9. Some of the largest uncertainties in both analyses described therein are related to the modelling of the $t\bar{t}$ process. The process in which a single top-quark is produced in association with a W-boson plays a large role in the analysis presented in Chapter 9. The MC simulation of these two processes is done using the same MC generation steps and, as a result, the methods by which their systematic evaluation is performed are the same. Here we describe the systematic variations used in the analyses for both of these processes.

Variation of the hard-scatter generation is performed by comparing the $t\bar{t}$ and Wt samples produced using Powheg for the matrix element generation to that using aMC@NLO, but keeping the showering and fragmentation model the same in both (Pythia8). The effects of the choice of fragmentation and hadronization model is assessed by comparing the use of Pythia8, used in the nominal $t\bar{t}$ and Wt predictions, to that of Herwig, all the while keeping Powheg for the hard-scatter generation in both cases. The description of the ISR and FSR provided by the hard-scatter generation is ascribed an uncertainty by varying the underlying parameters that describe the characteristic energy scales at which additional radiation (beyond that described by the matrix-element hard-scatter) is produced. In Powheg, this is controlled primarily by the `hdamp` parameter, and to assess the impact of this parameter's value on the analyses it is varied up and down by a factor of 2 relative to the nominal scenario. The impact of the choice of PDF used in the hard-scatter generation is assessed by varying the choice of PDF from the nominal, which is that of the NNPDF collaboration [174], to that of the MMHT [105], CT14 [175], and PDF4LHC [176] PDF sets and taking the envelope of the variations. The PDF error set, comprised of 100 separate components and provided by the NNPDF collaboration, is used to assign an additional uncertainty on the nominal PDF used. To assess the impact of the finite order in QCD at which the MC simulation is taken, and sensitivity to missing higher-order terms, the factorization and renormalization scales, μ_F and μ_R, are varied in all pairings possible in which either is varied by a factor of two up or down. The resulting *scale uncertainties* are derived by taking the envelope of the resulting variations.

Only in the analysis presented in Chapter 9 is the uncertainty related to the theoretical prediction of the $t\bar{t}$ and Wt processes taken into account. The former is taken to be $\pm 5.82\%$ and the latter

$\pm 5.32\%$ [177, 178, 179]. The analysis described in Chapter 9 considers the sum of the $t\bar{t}$ and Wt processes as a single background and constrains the normalisation of their sum using a dedicated CR. The uncertainties on these processes' cross-section, therefore, are taken into account to allow for the *composition* of the combined estimate ($t\bar{t} + Wt$), as opposed to its normalisation, to vary within the theoretical uncertainties.

An additional uncertainty, considered in both analyses presented in the current thesis, is related to the non-trivial quantum interference between the NLO predictions of the $t\bar{t}$ and $Wt + b$ processes relevant to both analyses. This uncertainty is assessed by comparing the estimates of the Wt process simulated under the so-called Diagram Removal (DR) and Diagram Subtraction (DS) schemes that are used in the NLO calculation of the single-top Wt process. The DR and DS schemes, and their comparison as a means of assessing a systematic uncertainty as just described, are fully described in Refs. [180, 181]. This *interference uncertainty* is negligible in the analysis presented in Chapter 8, given the relatively small contamination of the Wt process. However, it is an important uncertainty in the analysis presented in Chapter 9. In this latter analysis, the Wt background contributes at a rate equal to that of the $t\bar{t}$ process, making the assessment of their interference a subtle one that will be described in Chapter 9.

Z-boson Production in Association with Heavy-Flavor Jets

Processes involving the production of a SM Z-boson in association with two or more heavy-flavor jets are important for the analysis presented in Chapter 9. These processes are referred to simply as 'Z+HF'.

The PDF and scale uncertainties on the Z+HF process are computed in exactly the same fashion as for the $t\bar{t}$ and Wt processes described in the previous section. Additional uncertainties on the parton shower merging scales are assessed by varying the CKKW merging scales [182] used by the Sherpa MC generator used for the Z+HF simulation. To assess the impact of different MC hard-scatter generation and fragmentation models, the nominal sample produced using Sherpa is compared to a Z+HF simulation produced using Madgraph+aMC@NLO for the hard-scatter generation. This last source of uncertainty is important for the analysis described in Chapter 9 and amounts to an uncertainty of $\pm 15\%$ on the Z+HF background estimate in the analysis' SRs. The other sources of uncertainty on the Z+HF background are minor in comparison.

Diboson Production

The SM production of boson pairs — diboson production ('VV') — plays an important role only in the analysis described in Chapter 8. The nominal VV background estimate is simulated using the

SHERPA MC generator. Uncertainties arising as a result of variations in the factorization and renormalization scales, PDF variations, and uncertainties in the CKKW merging scales are considered. The methods by which these uncertainties are quantified are described above. These systematic uncertainties amount to about a $\pm 9\%$ uncertainty in the SR prediction in the analysis described in Chapter 8, but are generally smaller than the modelling uncertainties on the $t\bar{t}$ background process.

7.3.4 Signal Modelling Uncertainties

Uncertainties on the signal modelling are assessed in both analyses presented in Chapters 8 and 9. The methods by which the uncertainties arising due to variations in the μ_F and μ_R scales, as well as due to variations in the PDF choice, are assessed in exactly the same manner as for the SM backgrounds described in Section 7.3.3. In both searches, these uncertainties on the signal are negligible. For the analysis described in Chapter 8, the uncertainties arising as a result of varying the fragmentation model are found to be negligible for the phase space probed by the SUSY parameter space, described by the masses of the supersymmetric top-quark and LSP, in which that search is performed. In the analysis described in Chapter 9, uncertainties arising as a result of the choice in fragmentation model are assessed by varying the choice of parton shower generator from the nominal HERWIG to an alternative description provided by PYTHIA and have little impact on the analysis' final result. The uncertainties on the theoretical cross-sections on the signal models appearing in the two analyses are generally small and are found to not have a large impact on either of the analyses' final results.

7.4 Hypothesis Testing and Statistics

This section describes the statistical procedures used in the analyses to be presented in Chapters 8 and 9 that allow for conclusions to be drawn about the compatibility of the observed data with theories of BSM physics. The statistical inference tools described are inherently Frequentist and are, for the most part, the *de facto* standard for physics analyses searching for evidence of BSM physics at the large experiments at the LHC. Their widespread adoption by the experiments at the LHC does not indicate the philosophical merit of Frequentist inference methodology, but rather highlights the technically simple implementation of Frequentist hypothesis testing that allows for physics analyses to not get bogged down in some of the details associated with Bayesian analyses, computational or otherwise. Indeed, most people by default think and interact with the world around them in a Bayesian manner. Taking the path of least resistance, physicists have tended to opt for the simpler implementation of reporting their results, which, at the end of the day, tend to not lose out much in terms of the picture of the objective truth that they draw [183]. Section 7.4.1 will describe, in somewhat general terms, what hypotheses tests are and the way in which they

are performed in ATLAS. Sections 7.4.2-7.4.4 describe the details by which the measurements and systematic uncertainties of an analysis are transcribed into the language of the hypothesis test described in Section 7.4.1 using a likelihood-based test statistic.

7.4.1 Hypothesis Testing and the CL_s Construction

Hypothesis testing starts with the unambiguous formulation of the hypothesis being tested. In the search for evidence of BSM physics, there are two hypothesis pitted against one another. The first is the *null hypothesis*, denoted H_0, which is the hypothesis subject to the test and corresponds to the SM hypothesis. The null hypothesis is commonly referred to simply as the background-only (B) hypothesis. The second hypothesis is the *alternate hypothesis*, denoted H_1, and corresponds to the SM with the addition of the BSM physics process being sought out. The hypothesis H_1 is commonly referred to as the signal-plus-background (S+B) hypothesis. In both searches presented in Chapters 8 and 9, H_0 is taken to be the SM. In the search presented in Chapter 8, H_1 is taken to be a specific instantiation of the MSSM (Section 2.1), with specific masses of the stop quark and LSP. In the search presented in Chapter 9, H_1 is taken to be the non-resonant production of Higgs boson pairs. In this latter case, the H_1 hypothesis is indeed a process predicted by the SM but it is one that is not included in the H_0 hypothesis.

The Test Statistic and p-Values

In order to perform a hypothesis test in the Frequentist arena, a *test statistic*, $q(x)$, is defined. A test statistic is defined using the analysis' measurements x alone and is used in order to define metrics by which the observed data is said to agree with one of the two hypothesis, either H_0 or H_1. In Section 7.4.2, the exact form of the likelihood used in modern LHC experiments, and that used in the analyses discussed in Chapters 8 and 9, will be introduced. Here we will discuss general features of Frequentist test statistics and introduce some of the language that will be used later on when discussing the results of the analyses.

The conclusions eventually drawn about a given hypothesis are based on the observed value of $q(x)$ and where this value lies in relation to the pre-defined *critical region*. The critical region is defined by a cut value, q_c, on the distribution of $q(x)$ under a specified hypothesis. In the one-sided tests to be considered in the present thesis, H_1 will tend to have larger values of $q(x)$ as compared to H_0. The critical region defines two important parameters associated with the hypothesis test. The

first is the quantity α, which is referred to as the *significance level*, and is defined as follows,

$$\int_{q_c}^{+\infty} f(q|H_0)\,\mathrm{d}q = \alpha, \tag{7.12}$$

where $f(q|H_0)$ is the probability distribution for the test statistic under the background-only hypothesis. The quantity α reports the probability for the background-only hypothesis (the SM) to be rejected when it is actually true. This is commonly referred to as the Type I error rate. The second quantity is β and is defined as,

$$\int_{-\infty}^{q_c} f(q|H_1)\,\mathrm{d}q = \beta, \tag{7.13}$$

where $f(q|H_1)$ is the probability distribution for the test statistic under the signal-plus-background hypothesis. The quantity β gives the probability to reject the signal-plus-background hypothesis when it is actually true. This is commonly referred to as the Type II error rate. The quantity $(1-\beta)$ is referred to as the *power of the test*. The better a given physics analysis is at being able to discriminate between the signal and background, i.e. to have clear separation between the H_0 and H_1 hypotheses, the smaller (larger) is β (the power of the test).

For simplicity, the two hypotheses H_0 and H_1 can be generalised by introducing a so-called 'signal strength' parameter, μ, which acts as a multiplicative factor on the signal cross-section appearing in H_1. The hypothesis H_0, then, corresponds to the case $\mu = 0$ and that of H_1 corresponds to $\mu = 1$. With this general notation, then, the test statistic under either hypothesis is labelled as q_μ.

Once a test statistic is specified, and its expected distribution under a given hypothesis is obtained, p-values can be defined in order to compute the probability that the observed data originates from the considered hypothesis (value of μ). They are computed as follows,

$$p_\mu = \int_{q_{\mu,\mathrm{obs}}}^{+\infty} f(q_\mu|\mu)\,\mathrm{d}q_\mu, \tag{7.14}$$

where $q_{\mu,\mathrm{obs}}$ is the observed value of the test statistic in data and $f(q_\mu|\mu)$ is the probability density function of q_μ assuming hypothesis μ. A particular case of Equation 7.14 is that of p_0, which quantifies the agreement of the data with the background-only hypothesis ($\mu = 0$). The p_μ-value associated with a given hypothesis (μ-value) is typically converted into the equivalent corresponding Gaussian significance, Z, defined as the number of standard deviations that correspond to an upper-tail probability of p_μ. This is illustrated in Figure 7.11.

As the value of p_μ gets smaller, the confidence that the assumed hypothesis (value of μ) is true

decreases. At a certain point, it becomes acceptable to say that the assumed hypothesis is incompatible with reality and the hypothesis described by the particular value of μ is said to be *excluded*. In the particle physics community, the conventional threshold to take for the value of p_μ at which point a hypothesis is said to be excluded is $p_\mu = 0.05$, corresponding to $Z = 1.64$ as illustrated in Figure 7.11. This choice of the p_μ-value at which point exclusion is said to occur defines the critical region, described above, of the test. The value of 0.05 corresponds to the significance level of the test (c.f. Equation 7.12), and is referred to a hypothesis test being performed at the 95% confidence level (CL) (i.e. CL $\equiv (1 - \alpha)$)

In order to claim that new physics has been seen, the null hypothesis ($\mu = 0$) must be rejected. The thresholds at which new physics can be said to have been observed and discovered are much more stringent than that used for the exclusion of a specified hypothesis. Incompatibilities with the null-hypothesis at the level of $p_0 = 1.3 \times 10^{-3}$ and $p_0 = 2.9 \times 10^{-7}$ are required in order to state that observation and discovery, respectively, of new phenomena has occurred. These thresholds, illustrated in Figure 7.11, for claiming observation and discovery are the fabled '3σ' and '5σ' p_0-value criterion adopted by the particle physics community.

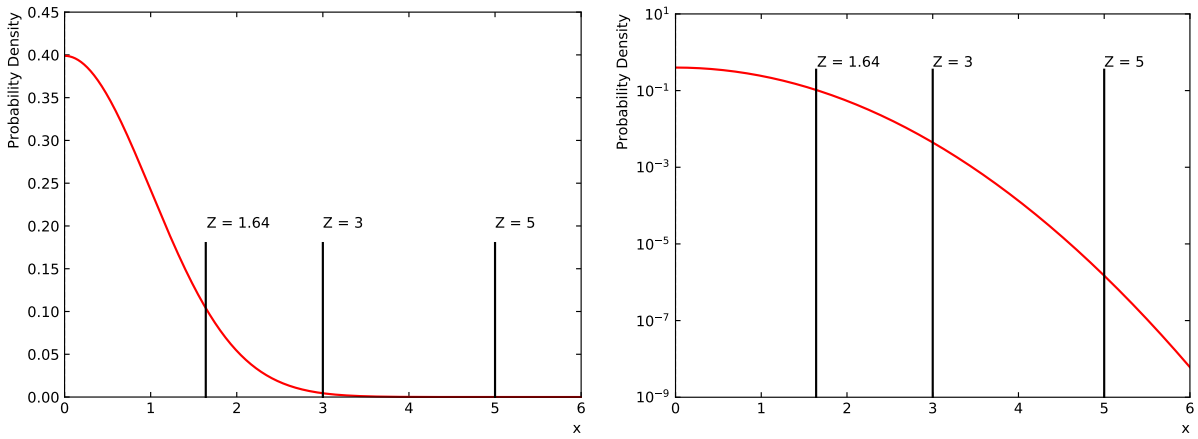

Figure 7.11: Gaussian tail significance levels corresponding to specific p-values. The significance level of $Z = 1.64\sigma$ corresponds to $p = 0.05$, that of $Z = 3\sigma$ to $p = 1.3 \times 10^{-3}$, and that of $Z = 5\sigma$ to $p = 2.9 \times 10^{-7}$. The area under the tail to the right of each indicated significance level corresponds to the associated p-value. ***Left***: Linear y-scale. ***Right***: Logarithmic y-scale.

The CL_s Construction

In searches for new physics, the statement that a given signal hypothesis has been excluded is an important one. Once made by the LHC experiments, the specific signal model is essentially considered no longer important to be searched for. Therefore, the metrics by which the experiments make claims of exclusion have surrounding them a wide-ranging literature discussing the merits and

drawbacks of the many such metrics that have been proposed over the years. The bare p_μ-value, for example, extracted from the observed data is subject to statistical fluctuations and it can lead to unphysical exclusions when a downward fluctuation in the observed number of events occurs. This could lead to a premature exclusion of a broad region of new physics that would perhaps no longer be looked into by future analyses or experiments.

The standard metric used by the LHC experiments today is known as '$\mathrm{CL_s}$' [184, 185], and is constructed in such a way as to reduce the likelihood of excluding signal hypotheses that a search is not a-priori sensitive to. The $\mathrm{CL_s}$ metric is given by,

$$\mathrm{CL}_s = \frac{p_\mu}{1 - p_0}, \tag{7.15}$$

where the quantities p_μ and p_0 quantify the compatibilities between the data and the signal-plus-background and background-only hypotheses, respectively. Downward fluctuations in data, as those described above, will lead to larger values of p_0; thus leading to larger values of $\mathrm{CL_s}$ that avoid premature exclusion.

At the LHC, the $\mathrm{CL_s}$ metric is used primarily for performing hypothesis tests aimed at claiming exclusion. The standard null-hypothesis p_0-value is still used for claiming observation and discovery, as described above. A given signal hypothesis with $\mu = 1$ is considered excluded at 95% CL when $\mathrm{CL_s} \leq 0.05$. Note that this prescription for exclusion, $\mathrm{CL_s} \leq \alpha$, is generally a stronger requirement than the standard prescription, $p_\mu \leq \alpha$. The $\mathrm{CL_s}$ metric is also used to compute *upper limits*. A 95% CL upper limit on a given signal hypothesis specified by μ is the largest value of μ satisfying $\mathrm{CL_s} \geq 0.05$. The interpretation being that this corresponds to the largest possible signal cross-section that is unable to be excluded and therefore smaller values of μ, corresponding to smaller signal cross-sections, are still consistent with the observed data and cannot therefore be excluded. The process of scanning μ hypotheses and computing the $\mathrm{CL_s}$ in order to find an upper limit on μ is illustrated in Figure 7.12.

7.4.2 The Profile Likelihood Ratio Test Statistic

The test statistic associated with many of the LHC experiments, including ATLAS, and the one used in the analyses to be presented in Chapters 8 and 9 is based on a likelihood ratio. The construction of the test statistic is described in this section.

In the analyses to be presented, so-called 'counting experiments' are performed wherein only the numbers of events, from data and the predicted background and signal, are used as input. These numbers are taken from all relevant regions in the analysis: the CRs and the SRs. The expected number of events populating each region is given by the following:

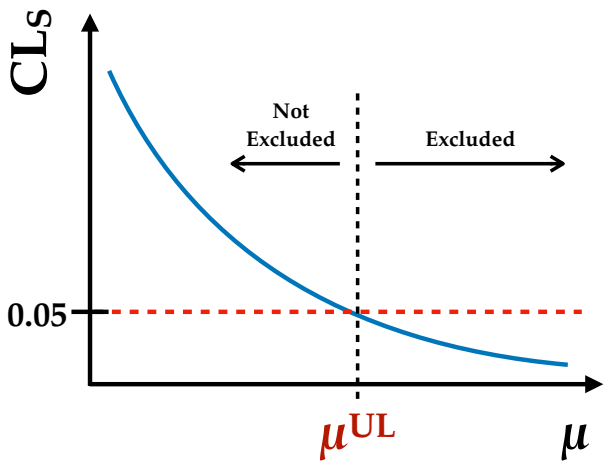

Figure 7.12: An upper limit scan on the signal strength parameter μ associated with a signal hypothesis. The CL_s, given by Equation 7.15, is recomputed for a range of μ values describing a given signal hypothesis. This is shown by the blue line. The μ value at which the CL_s curve crosses the line $CL_s = 0.05$, μ^{UL}, is the upper limit on μ for the signal hypothesis. Values of μ smaller than μ^{UL} remain compatible with the observed data, while those values greater than μ^{UL} are excluded at 95% CL.

$$N_r^{\text{exp}}(\mu_{\text{sig}}, \boldsymbol{\mu}_{\text{bkg}}, \boldsymbol{\theta}) = \mu_{\text{sig}} \cdot N_{r,\,\text{sig}}^{\text{exp}}(\boldsymbol{\theta}) + \sum_{b \in \text{bkg}} \mu_b \cdot N_{r,\,b}^{\text{exp}}(\boldsymbol{\theta}), \tag{7.16}$$

where $\boldsymbol{\theta}$ are a set of fit nuisance parameters (NP) associated with the systematic uncertainties described in Section 7.3, $\boldsymbol{\mu}_{\text{bkg}}$ are normalisation factors associated with the background processes (indexed by 'b'), μ_{sig} is the signal-strength modifier associated with the signal hypothesis, $N_{r,\,\text{sig}}^{\text{exp}}$ is the predicted signal yield in region r, and $N_{r,\,b}^{\text{exp}}$ is the predicted background yield in region r for process b. The predicted number of events for each process, signal or background, depend on the $\boldsymbol{\theta}$ parameter vector since systematic variations can adjust the overall normalisation of a given process or adjust the acceptance of a given process in the phase space probed by the regions indexed by r. Both of these effects result in a change in the predicted rate of a process in a given region.

The observed data yield in each region of the analysis is expected to obey Poisson statistics. Therefore, the likelihood function $L(\mu_{\text{sig}}, \boldsymbol{\mu}_{\text{bkg}}, \boldsymbol{\theta})$ is constructed as a product of Poisson probability terms:

$$L_0(\mu_{\text{sig}}, \boldsymbol{\mu}_{\text{bkg}}, \boldsymbol{\theta}) = \prod_{r \in \text{regions}} \frac{[N_r^{\text{exp}}(\mu_{\text{sig}}, \boldsymbol{\mu}_{\text{bkg}}, \boldsymbol{\theta})]^{N_r^{\text{obs}}}}{N_r^{\text{obs}}!} \cdot \exp\left[N_r^{\text{exp}}(\mu_{\text{sig}}, \boldsymbol{\mu}_{\text{bkg}}, \boldsymbol{\theta})\right], \tag{7.17}$$

where N_r^{obs} is the observed data yield in region r.

It is standard practice to parametrize the systematic uncertainties associated with the measurements of the N_r^{exp} in such a way that $\boldsymbol{\theta} = 0$ (all of the θ_i equal to zero) corresponds to the central (nominal) value of the set of parameters associated with the uncertainty (e.g. the nominal value of JES). The values $\theta_i = \pm 1$ represent shifts in the parameter values by their $\pm 1\sigma$ variation, as defined by the systematic uncertainty (e.g. systematic shifts in the value of the JES). The systematic uncertainties described in Section 7.3 are typically computed such that, at the level of performing a physics analysis, only the $\pm 1\sigma$ shifts about the central value in the associated parameter are known. Means of interpolation and extrapolation, needed to obtain a smoothly varying response in the numbers of predicted events as the θ parameters vary continuously between their $\pm 1\sigma$ values, are provided by the HistFactory toolkit [186]. The effects of the NP associated with each of the analysis' systematic uncertainties are implemented as *constraint terms* in the likelihood appearing in Equation 7.17. These terms are typically implemented as Gaussians centered on zero and with a width of one:

$$L(\mu_{\text{sig}}, \boldsymbol{\mu}_{\text{bkg}}, \boldsymbol{\theta}) = L_0(\mu_{\text{sig}}, \boldsymbol{\mu}_{\text{bkg}}, \boldsymbol{\theta}) \cdot \underbrace{\prod_{i=1}^{n} \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{\theta_i^2}{2}\right)}_{\text{Constraints}}. \tag{7.18}$$

The best estimates (maximum likelihood estimates, MLE) for the parameters (μ_{sig}, $\boldsymbol{\mu}_{\text{bkg}}$, $\boldsymbol{\theta}$) are determined in the fit to the observed data, $\boldsymbol{N}^{\text{obs}}$, via the maximization of the likelihood function given by Equation 7.18. Typically, and equivalently, the negative log likelihood, $-\ln L$, is minimized. Technically, the minimization of Equation 7.18 in the analyses to be presented in Chapters 8 and 9 is done using Miniuit via an interface provided by RooFit [187, 188]. The values of all parameters prior to (after) the minimization procedure are referred to as the 'pre-fit' ('post-fit') values. The pre-fit values for the μ_{bkg} parameters are set to 1 and the θ_i are set to 0, corresponding to the central values of the Gaussian-shifted parameters associated with the systematic uncertainties.

The test statistic defined for the analyses to be presented is based on the following *profile likelihood ratio*,

$$\lambda(\mu_{\text{sig}}) \equiv \frac{L(\mu_{\text{sig}}, \hat{\hat{\boldsymbol{\mu}}}_{\text{bkg}}, \hat{\hat{\boldsymbol{\theta}}})}{L(\hat{\mu}_{\text{sig}}, \hat{\boldsymbol{\mu}}_{\text{bkg}}, \hat{\boldsymbol{\theta}})}, \tag{7.19}$$

where in the numerator the parameters ($\boldsymbol{\mu}_{\text{bkg}}, \boldsymbol{\theta}$) are fit to their MLE values for a given value of the assumed value of the signal-strength parameter μ_{sig}. In the denominator, the full set of parameters (μ_{sig}, $\boldsymbol{\mu}_{\text{bkg}}$, $\boldsymbol{\theta}$) are equal to their MLE values. The likelihood ratio defined in Equation 7.19 is used to define the final test statistic relevant to the analyses,

$$q_{\mu_{\text{sig}}} = -2 \ln \lambda(\mu_{\text{sig}}), \tag{7.20}$$

which is used to perform the hypothesis tests described in Section 7.4.1.

Profile Likelihood Fit Determination of Background Constraints

We note here that the quantities described by $\boldsymbol{\mu}_{\text{bkg}}$ in Equations 7.16–7.20 correspond to normalisation correction factors for a given SM background process. In the analyses to be presented in Chapters 8 and 9, for the background processes for which a dedicated CR is not defined, the associated μ_b are set to 1 and are not allowed to vary during the fit (i.e. their pre-fit and post-fit values are equal). For those processes for which a dedicated CR is defined, the associated μ_b parameters are analogous to the normalisation correction factors described in Section 7.1.3 and are unconstrained in the fit procedure. They are therefore said to 'freely float' during the fit and their post-fit value does not necessarily correspond to their pre-fit one. Given the low numbers of events typically expected in the SRs, and the relatively large numbers and purities expected in the CRs, the post-fit values of the μ_b typically correspond to those values expected from the simple computations provided by Equation 7.1 and/or Equation 7.2. An additional note is that in the signal-plus-background hypotheses with $\mu_{\text{sig}} \neq 0$, if the CRs have a non-negligible contamination of signal events, however, this correspondence will generally not be true since the post-fit values of the μ_b will depend on a given value of the μ_{sig} parameter. The desire to construct a robust background-only model is one of the motivating factors, then, in designing analyses in such a way as to minimize the signal contamination in the CRs.

Profiling

The Gaussian constraint terms in Equation 7.18, centered on 0 and with widths of 1, are equivalent to transforming the $\boldsymbol{\theta}$ parameters associated with the systematic uncertainties as follows,

$$\boldsymbol{\theta}' = \frac{\boldsymbol{\theta} - \hat{\boldsymbol{\theta}}}{\hat{\boldsymbol{\sigma}}}, \tag{7.21}$$

where the quantity $\boldsymbol{\sigma}$ corresponds to the initial widths of the $\pm 1\sigma$ variations associated with the systematic variation represented by $\boldsymbol{\theta}$. Equation 7.21 normalises all NP to facilitate the easy comparison of their pre- and post-fit values and uncertainties. For a given NP, then, a post-fit value and uncertainty near 0 and 1, respectively, indicate that the data was not able to adjust the NP. Changes in the post-fit NP values (uncertainties) are referred to as NP 'pulling' ('profiling'). NP values may be pulled far from zero in order to maximize the overall agreement of the background prediction with data during the fit procedure. NP may be profiled in such a way that the post-fit uncertainties on the NP are smaller than the initial estimate of the uncertainty provided by auxiliary measurements and data. Such a case indicates that the initial prior on the impact of the systematic uncertainty was too large and that the uncertainties may be reduced so as to be more compatible with the range allowed by the data statistics relevant for the analysis at hand. This potential for the profiling mechanism to reduce the overall impact of systematic uncertainties on

an analysis' results is seen as a general benefit of the profile likelihood prescription described by Equations 7.19 and 7.20. The interpretation of this is that the profiling mechanism allows for the impact of the systematic uncertainties to be more accurately characterised in the phase space in which they are being applied, as opposed to that of the auxiliary measurements in which they are initially derived.

7.4.3 Toy Examples of the Profile Likelihood Fit

The best way to get a feel for the profiling mechanisms just described is via a few simple examples, to which we turn now.

Profile Likelihood Fit Example 1a

This example illustrates the mechanisms of both NP profiling and pulling. We set up a dummy analysis, with two regions and two backgrounds only. The regions are referred to as 'Region 1' and 'Region 2' and the backgrounds are 'Bkg 1' and 'Bkg 2'. A summary of the predicted and observed yields, as well as the uncertainties on the predicted yields, is as follows:

- **Region 1**
  - Background composition: 80 events from Bkg 1
  - Observed data: 100 events
  - Systematic Uncertainties:
    - 'Norm. Bkg. 1': A 50% systematic uncertainty on the predicted yield of Bkg 1
- **Region 2**
  - Background composition: 100 events from Bkg 2
  - Observed data: 100 events
  - Systematic Uncertainties:
    - 'Norm. Bkg. 2_1': A 10% systematic uncertainty on the predicted yield of Bkg 2
    - 'Norm. Bkg. 2_2': An additional 10% systematic uncertainty on the predicted yield of Bkg 2

From this information, a likelihood as described by Equation 7.18 is constructed. In this setup, all of the μ parameters are set to 1 for the backgrounds and NP constraint terms parametrized by θ terms for each of the three uncertainties are included. The pre-fit situation is illustrated by the

left side of Figure 7.13. A profile-likelihood fit to the observed data is performed and the results for the fitted parameters are shown in Figure 7.14 and by the right side of Figure 7.13.

Focusing first on Region 1, we see that after the fit the pre-fit 20 event excess is removed by the increase in Bkg. 1's predicted yield. This is explained by the NP on Bkg. 1 being pulled by nearly 0.5, as seen in Figure 7.14. Using Equation 7.21, a pull by 0.5 in this NP is precisely 20 events ($\theta^{\text{post-fit}} \times (\text{pre-fit uncertainty} \times N^{\text{pre-fit}}) \rightarrow 0.5 \times (0.5 \times 80) = 20$). The post-fit uncertainty has also been reduced to that allowed by the data statistics, which is driven by the NP on Bkg. 1 being profiled from its initial uncertainty of ± 1 to ± 0.25 ($\Delta\theta^{\text{post-fit}} \times (\text{pre-fit uncertainty} \times N^{\text{pre-fit}}) \rightarrow 0.25 \times (0.5 \times 80) = 10$), also seen in Figure 7.14.

Looking to the post-fit results of Region 2, we see that, as with Region 1, the overall background uncertainty has been profiled so as to be compatible with that allowed by the data statistics. The main difference with respect to Region 1, however, is the fact that there are two NP constraints on Bkg. 2. Each of the NP constraint terms is already at the level of 10% that corresponds to the data statistics. There is therefore a redundancy in the NP constraints on Bkg. 2 and the fit develops a correlation between them such that their combined post-fit impact on Bkg. 2 is as seen on the right side of Figure 7.13. The correlation between the two constraints is seen in Figure 7.14 to be $\rho = -0.5$, which, in this simple scenario, can be expected to be the case by the following:

$$\begin{aligned} \sigma_1 &= \sigma_2 = \sigma, \\ \sigma_{1\oplus 2} &= \sqrt{\sigma_1^2 + \sigma_2^2 + 2\sigma_1\sigma_2\rho} = \sigma, \end{aligned}$$

where, in the second line, we expect σ to be at the level of the data statistics. This requires $\rho = -0.5$ and is indeed what the profile-likelihood fit ends up with.

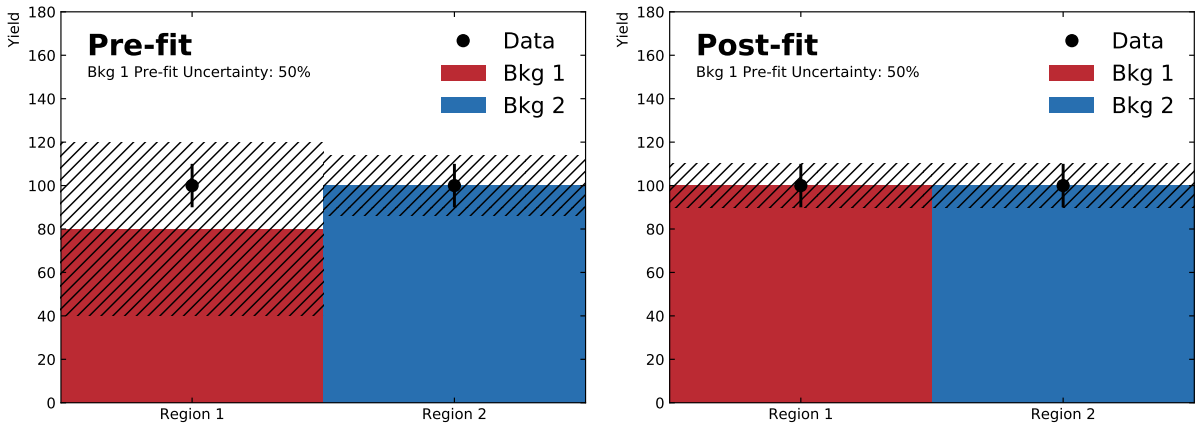

Figure 7.13: ***Left***: Pre-fit scenario for Example 1a, described in the text. The hatched areas indicate the uncertainty in the predicted yields. ***Right***: Post-fit scenario for Example 1a, described in the text. The hatched areas indicate the uncertainty in the predicted yields.

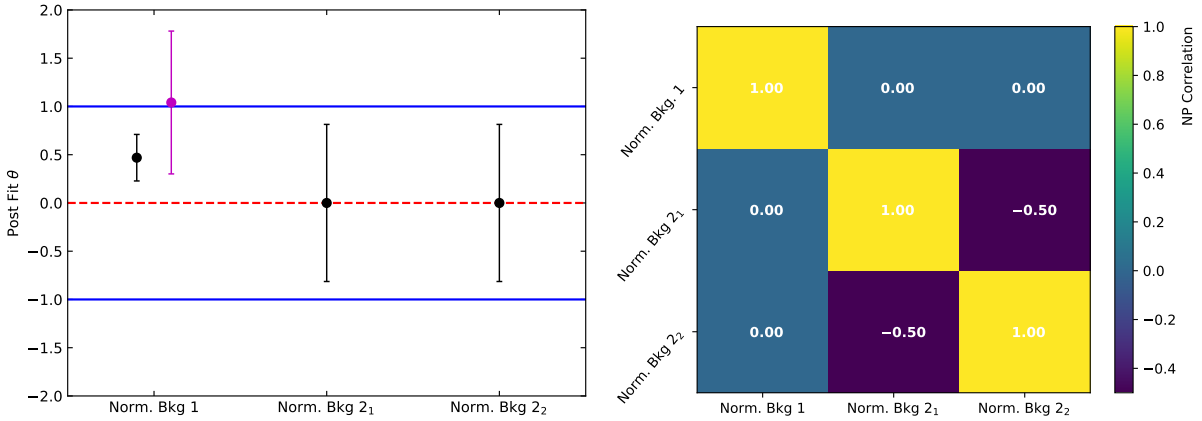

Figure 7.14: ***Left***: Post-fit values for the parameters entering the profile-likelihood fit described in Example 1a. The post-fit result for the NP 'Norm. Bkg. 1' in pink corresponds to that of the fit configuration described in Example 1b, below. All other parameters (in black) correspond to those in the fit configuration described in Example 1a. ***Right***: Post-fit correlation matrix for the parameters entering the profile-likelihood fit described in Example 1a.

Profile Likelihood Fit Example 1b

In the previous example we saw that, through the pulling of a constraint term, the complete removal of a discrepancy between the observed data and the predicted background was made possible. This reinforces the idea that the predicted numbers of events are dependent upon the NP entering the fit ($N^{\text{exp}} \rightarrow N^{\text{exp}}(\boldsymbol{\theta})$) and that their post-fit values can change even if their is no explicit normalisation-correcting μ terms in the fit. Of course, the large uncertainty on Bkg. 1's prediction in Example 1a equates to a loose constraint term and allows for a large degree of flexibility in the fit for it to be pulled in such a way as to have the post-fit prediction perfectly line up with the data being fit to. In this current example we reproduce the fit configuration introduced in Example 1a, but instead the uncertainty on Bkg. 1 in Region 1 is reduced to 10%. The pre- and post-fit distributions of the observed and predicted events, with their uncertainties, are shown in Figure 7.15. The results of the fit for Region 2 are exactly the same as in Example 1a since we have not altered the parameters describing this region.

Given the tighter constraint on the predicted yield of Bkg. 1 in Region 1, as compared to Example 1a, the NP describing Bkg. 1's normalisation uncertainty does not have enough freedom to be pulled to fully cover the 20 event excess. The post-fit value of the NP is shown in pink in Figure 7.14. It is pulled to a value of 1.04, which corresponds to just above 8 events as seen in Figure 7.15. The smaller uncertainty on Bkg. 1's prediction in this example, as compared to Example 1a, corresponds to a higher level of confidence in its pre-fit value. The associated NP constraint term, therefore, should not be allowed the freedom to contradict this high degree of confidence by shifting the prediction to the same extent as that in Example 1a.

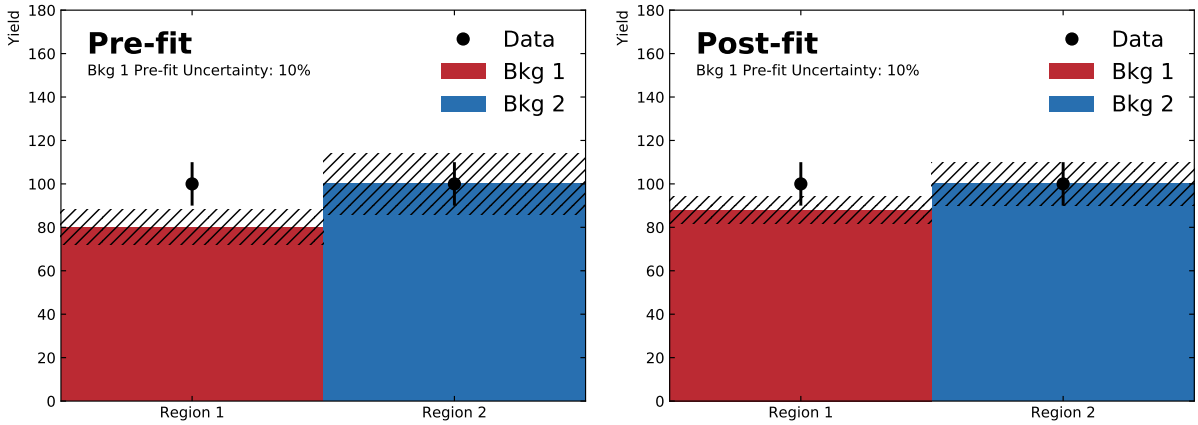

Figure 7.15: ***Left***: Pre-fit scenario for Example 1b, described in the text. The hatched areas indicate the uncertainty in the predicted yields. ***Right***: Post-fit scenario for Example 1b, described in the text. The hatched areas indicate the uncertainty in the predicted yields.

Profile Likelihood Fit Example 2

In this example we again set up the same fit configuration as in Example 1a. However, a signal process 'Signal' with 20 events predicted in Region 1 is added. The signal-strength parameter μ_{sig}, acting as a freely floating normalisation term for the signal process, is now included in the fit as well. A profile-likelihood fit to the observed data is once again performed. The pre- and post-fit distributions of the observed data and predicted backgrounds, with their uncertainties, are shown in Figure 7.16.

From the post-fit distributions shown in Figure 7.16, we see the expected result that the post-fit uncertainties on the predicted backgrounds are such that they match the precision allowed by the data statistics. As we have not changed anything in Region 2, relative to Example 1a and Example 1b, we do not expect the post-fit values of the NP associated with Bkg. 2 to have changed relative to those examples. The post-fit value of the NP associated with Bkg. 1, however, differs quite substantially relative to the earlier examples, as is seen in Figure 7.17. We see that the NP term 'Norm. Bkg. 1' is neither pulled nor profiled. Instead, the unconstrained signal-strength parameter μ_{sig} is able to 'eat up' the remaining degrees of freedom needed for the reduction in the background prediction's uncertainty. Since the pre-fit predicted yield of the signal process is such that the complete background prediction in Region 1 matches the data, the post-fit value for μ_{sig} is equal to its pre-fit value of 1. It's uncertainty, however, is ± 1.90 and corresponds to a post-fit uncertainty on the signal process' prediction of nearly 40 events (1.90×20). This precision on the signal matches the pre-fit uncertainty of of Bkg. 1, with its predicted 80 events and 50% uncertainty. This illustrates the fact that the μ_{sig} term is completely degenerate with the NP associated with Bkg. 1: any variation in the NP term can be fully compensated by the unconstrained μ_{sig} term.

This compensation allows for the background precision in Region 1 to be reduced nearly to that of the data statistics, as seen in Figure 7.16, and is characterised by the nearly perfect anti-correlation between the μ_{sig} and Bkg. 1 NP constraint terms, shown in the right side of Figure 7.17. This ability for μ_{sig} to compensate the NP by preventing it from being pulled is a general feature of the addition of unconstrained normalisation-scaling μ terms in these types of fits, regardless of whether they are associated with signal or background processes.

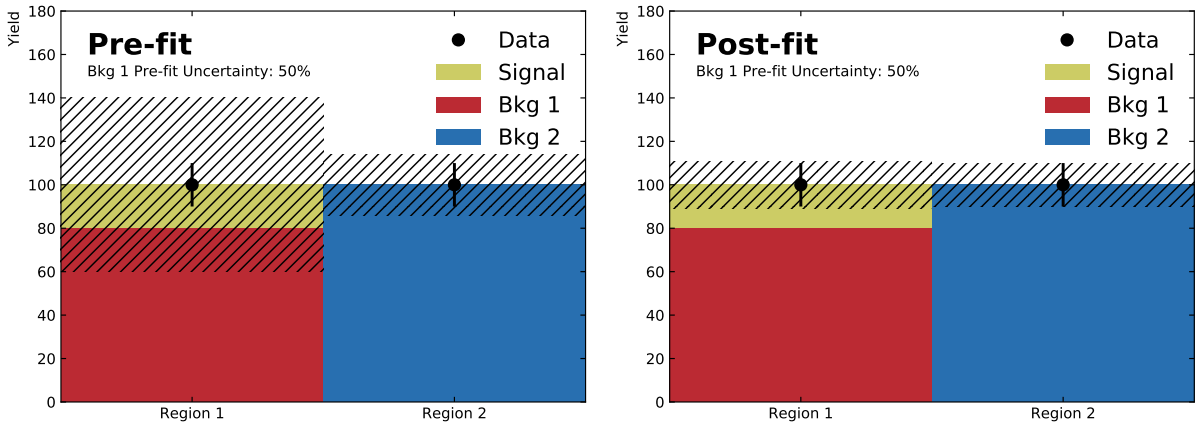

Figure 7.16: ***Left***: Pre-fit scenario for Example 2, described in the text. The hatched areas indicate the uncertainty in the predicted yields only on the background processes. ***Right***: Post-fit scenario for Example 2, described in the text. The hatched areas indicate the uncertainty in the predicted yields only on the background processes.

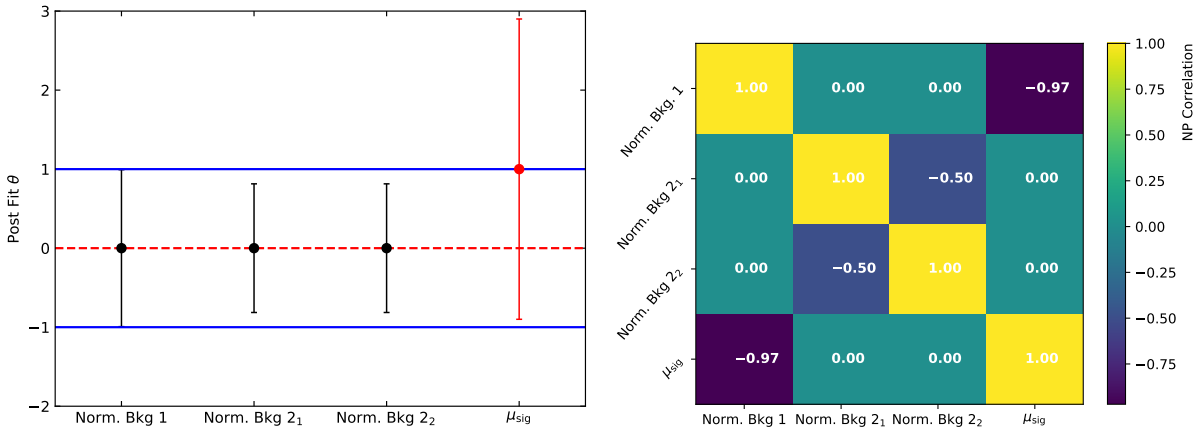

Figure 7.17: ***Left***: Post-fit values for the parameters entering the profile-likelihood fit of Example 2, described in the text. ***Right***: Post-fit correlation matrix for the parameters entering the profile-likelihood fit of Example 2, described in the text.

It is interesting to make note of the fact, exemplified by the above discussion, that the precision on a given process' μ term, if allowed to vary in the fit, is highly dependent on the precision of the prediction of the other backgrounds. This, of course, is intuitive and has meaningful consequences

for searches for new physics, in which the signal-plus-background hypotheses are characterised by μ_{sig} terms associated with the sought-for signal processes. Large uncertainties in background predictions result in less precise statements about the presence of signal in the analyses' regions. This, then, reduces the statistical power (c.f. Equation 7.13) of the hypothesis tests being performed and, if the uncertainties are large enough, prevent the analyses from making meaningful statements about the compatibility of the signal-plus-background hypotheses with the observed data.

As an illustration, Figure 7.18 shows the CL_s computed as a function of varying signal-to-background ratio, for several background uncertainty hypotheses, for a fixed number of predicted background and data events in a single-region analysis. The number of signal events needed to cross into the critical $\text{CL}_s < 0.05$ region, at which point a signal hypothesis may be considered excluded at 95% CL, is quite sensitive to the uncertainty on the background prediction. Again, this is intuitive mainly from the perspective that reduced levels of precision in the background prediction correspond directly to increased levels of ambiguity as to the point at which a signal becomes apparent.

Figure 7.18 also illustrates the increase in the upper limit on a signal process' cross-section, as a result of increased background uncertainty, in cases where no significant excess in data over the predicted background level is observed. For example, imagine that the sought-for signal process has a predicted cross-section leading to an S/B value of 0.4 in Figure 7.18. At a level of background uncertainty of 30%, with a $\text{CL}_s \approx 0.35$, this process can neither be said to be excluded nor to actually exist. At this background uncertainty, the analysis would find an upper-limit on the signal process' μ_{sig} term nearing a value of 2, since the CL_s crossing point is near $S/B = 0.8$. Using this terminology, then, when no meaningful discrepancies (characterised by very low p_0-values) between the predicted background and observed event counts are seen, the upper-limit value under the hypothesis for a given signal process gives a measure of 'how far' the analysis is from being sensitive to that process. If the upper-limit value is 200 times the predicted cross-section, the analysis is not sensitive to the process at all. However, if the upper-limit value is $\mathcal{O}(1)$ times the process' predicted cross-section, the analysis is entering the regime wherein it will be able to resolve the presence of the signal and can begin to make meaningful statements about its likelihood of existence. If an upper-limit value corresponds to a cross-section value that is *less* than the predicted one, the process — as predicted — is excluded.

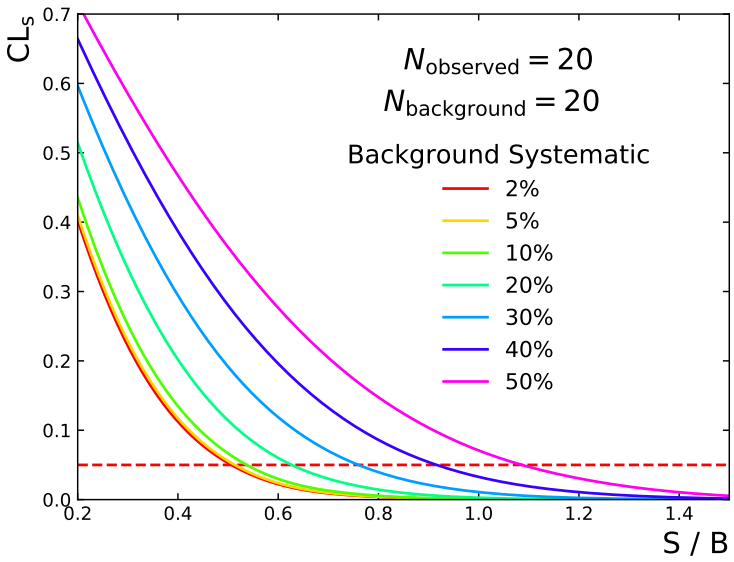

Figure 7.18: Dependence of computed CL_s on the signal-to-background ratio for a single-region analysis in which 20 events are predicted and 20 events are observed, shown for varying levels of uncertainty on the predicted background.

7.4.4 Asymptotic Formulation of Likelihood Ratio Tests

In Section 7.4.3 we have made use of the profile likelihood ratio fits and even shown some examples of hypothesis testing. We have not shown, however, how the probability density functions of the test statistics described by Equation 7.20 are obtained. The probability density functions are needed in order to compute the p_μ-values as indicated by Equation 7.14. There are two general approaches taken to obtain such distributions. The first is by generating so-called 'pseudo-experiments', in which the predictions of the hypotheses are sampled within their uncertainties many times, with a value of the corresponding test statistic computed for each sample, such that a distribution of the q_μ may be obtained. This method can be computationally intensive if high-levels of precision are required for the analysis. The method used in the searches presented in Chapters 8 and 9 relies on asymptotic approximations of the sampling distributions of the profile-likelihood ratio test statistic of Equation 7.20. The complete discussion and description of the asymptotic formulae used to represent Equation 7.20 is given in Ref. [189].

The key takeaway from Ref. [189], is that in the large N limit, with N corresponding to the data sample size, the sampling distributions of the profile likelihood test statistics take the form of non-central chi-square distributions with single degrees of freedom,

$$f(q_\mu|\Lambda) = \frac{1}{2\sqrt{q_\mu}}\frac{1}{\sqrt{2\pi}}\left[\exp\left(-\frac{1}{2}(\sqrt{q_\mu}+\sqrt{\Lambda})^2\right)+\exp\left(-\frac{1}{2}(\sqrt{q_\mu}-\sqrt{\Lambda})^2\right)\right], \tag{7.22}$$

where the non-centrality term, Λ, is given by the following:

$$\Lambda = \frac{(\mu - \hat{\mu})^2}{\sigma^2}. \tag{7.23}$$

Example distributions of the test statistic probability density function are shown in Figure 7.19, using Equation 7.22, illustrating the differences between the background-only and signal-plus-background hypotheses for a case in which there is a large predicted rate of signal and for the case in which there is not. As described by Equation 7.14, the p_μ and p_0 values needed, among other things, for the computation of $\mathrm{CL_s}$ can be obtained from these distributions for a given value of q_μ by integrating the areas to the right of the specified q_μ under the distributions associated with the $S + B$ and B hypotheses, respectively.

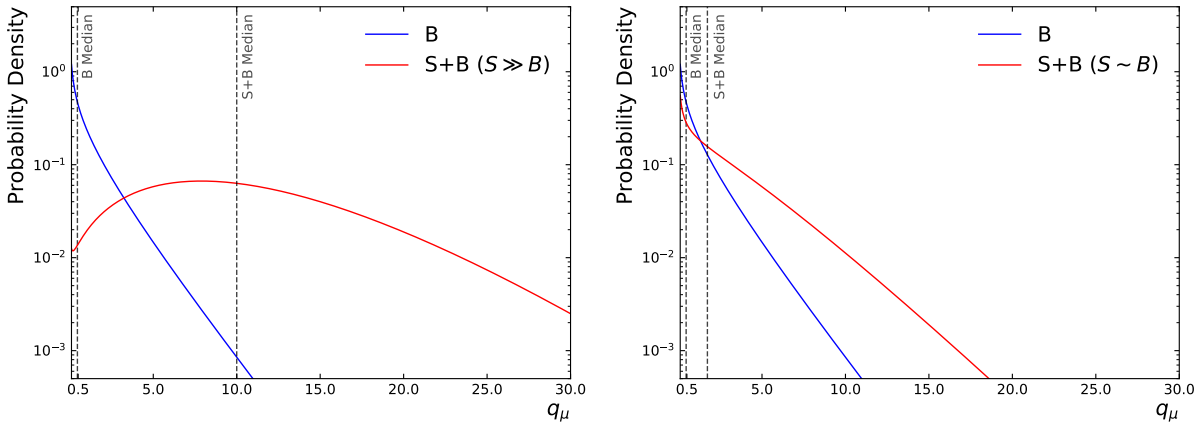

Figure 7.19: Probability densities of the profile likelihood test statistic of Equation 7.20 obtained under the asymptotic formulae described in Ref. [189] and shown in Equation 7.22. The background-only hypothesis ($\mu = 0$) is shown in blue, and the signal-plus-background hypothesis ($\mu \neq 0$) is shown in red. ***Left***: The case in which the predicted amount of signal is large compared to the background rate. ***Right***: The case in which the predicted signal rate is similar to that of the background.

Chapter 8

The Search for the Supersymmetric Top-quark

> *It will be remembered that the 18^{th} century was, on the whole, addicted to an ascending series of living forms shading by insensible degrees into each other and leading onto man. There was no consideration of the fact that this might be reading into Nature a greater unity than she actually possessed. It led inevitably to some highly questionable taxonomy produced in the effort to compress all life into positions upon a single stairway.*
>
> –Loren Eiseley, *Darwin's Century*

> *Cease, cows, life is short.*
>
> –Gabriel García Márquez, *One Hundred Years of Solitude*

As described in Chapter 2, the lighter mass eigenstate of the superpartner of the top quark, $\tilde{t}_1$, plays an important role in helping solve the Hierarchy problem. Most SUSY scenarios prefer $\tilde{t}_1$ masses not much larger than 1 TeV, so as to keep them at the scale EWSB. If the $\tilde{t}_1$ exists, and is indeed lighter than the TeV scale, it may easily be produced in the 13 TeV pp collisions occurring at the LHC. Searches for stops therefore play a prominent role in the searches for SUSY in ATLAS.

This chapter presents a search for the production of stop quarks. The SUSY model considered here satisfies R-parity conservation, as do the majority of SUSY searches at the LHC, which forces the $\tilde{t}_1$ to be produced in pairs and to follow a decay chain ending in the stable LSP. The LSP in these models is assumed to be the lightest neutralino, $\tilde{\chi}_1^0$.

SUSY searches performed in ATLAS are done using so-called 'simplified models', wherein all pa-

rameters of the MSSM are selected except for the masses of the sparticles relevant to the decays of the sparticle being searched for. For the case presented in this chapter, a search for $\tilde{t}_1$ is described and the simplified models used to describe the SUSY signal process are parametrised by two free parameters: the mass of the $\tilde{t}_1$ and the mass of the $\tilde{\chi}_1^0$. With these two parameters defined, the kinematics and production cross-section are fixed. In the models considered, the $\tilde{t}_1$ is classified into three regions of phase space according to the mass of the $\tilde{\chi}_1^0$, into which it decays. The potential decays of the $\tilde{t}_1$ are illustrated by the decay diagrams in Figure 8.1. showing the cases of the two-, three-, and four-body decays of the $\tilde{t}_1$. These decays are associated with three different regions of parameter space in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane, illustrated in Figure 8.2. The analysis to be discussed in this chapter reports the results of a search for the production of $\tilde{t}_1$ particles following the three-body decay, in which the mass-difference between the $\tilde{t}_1$ and $\tilde{\chi}_1^0$, $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0)$, is less than (greater than) the mass of the SM top-quark ($W + b$ system). This is the middle region of phase space indicated in Figure 8.2, $\tilde{t}_1 \to bW\tilde{\chi}_1^0$. In this region, the two-body $\tilde{t}_1$ decays are kinematically suppressed and the $\tilde{t}_1$ decays via an off-shell top-quark or $\tilde{\chi}_1^\pm$ to the three-body $bW\tilde{\chi}_1^0$ final state.

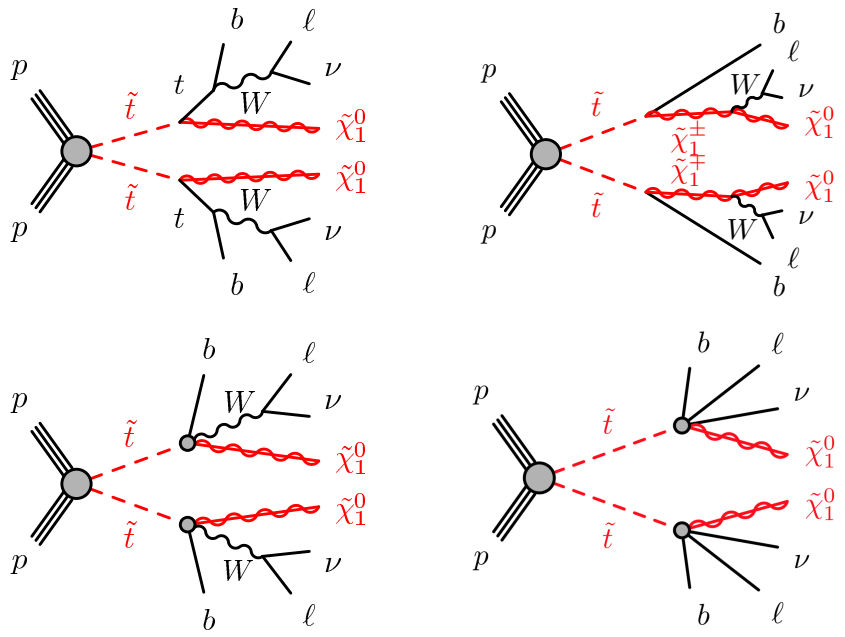

Figure 8.1: Decay diagrams for the $\tilde{t}_1$ relevant to the two-lepton final state of $\tilde{t}_1$ pair-production. ***Top***: Diagrams leading to the two-body decay of the $\tilde{t}_1$, either into $\tilde{t}_1 \to t\tilde{\chi}_1^0$ (*top left*) or $\tilde{t}_1 \to b\tilde{\chi}_1^\pm$ (*top right*). ***Bottom left***: Three-body decay of the $\tilde{t}_1$, $\tilde{t}_1 \to bW\tilde{\chi}_1^0$. ***Bottom right***: Four-body decay of the $\tilde{t}_1$, $\tilde{t}_1 \to bff'\tilde{\chi}_1^0$.

Figure 8.3 shows the results of the searches for $\tilde{t}_1$ particles performed in Run 1 of the LHC by the ATLAS experiment. This figure indicates the regions of SUSY parameter space, under the hypothesis of the simplified model parametrised by the masses of the $\tilde{t}_1$ and $\tilde{\chi}_1^0$ particles, that are excluded at 95% CL by the zero, one, and two lepton analyses. It can be seen that the Run 1 analyses targeting the two-lepton final states did not effectively cover the three-body regime, at

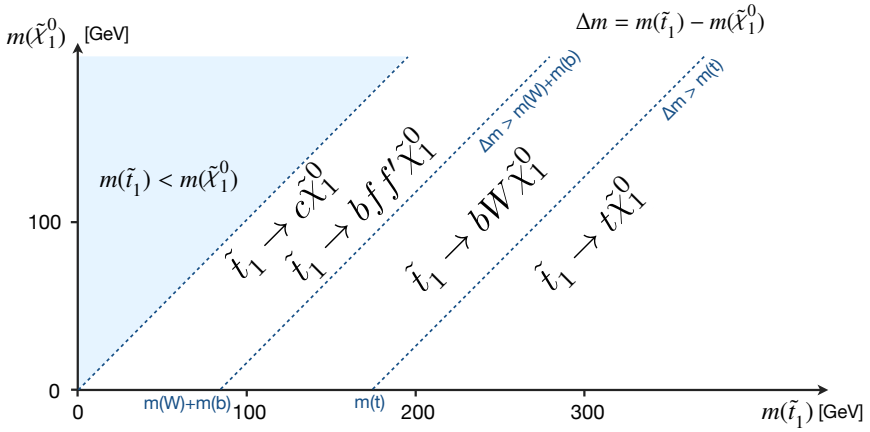

Figure 8.2: Kinematic boundaries in the $(\tilde{t}_1,\tilde{\chi}_1^0)$ plane, indicating the three kinematic regimes associated with the decay of the $\tilde{t}_1$: the two-body $\tilde{t}_1 \to t\tilde{\chi}_1^0$ decay ($\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) > m_t$), the three-body $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ decay ($\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) < m_t$), and the four-body $\tilde{t}_1 \to bff'\tilde{\chi}_1^0$ decay ($\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) < m_b + m_W$). Here $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) = m_{\tilde{t}_1} - m_{\tilde{\chi}_1^0}$.

least in comparison to all analyses in the two-body and four-body regions of the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$-plane.

Section 8.1 provides a brief introduction to the type of plots shown in Figure 8.3, which are the general SUSY exclusion plots that are presented to summarise the results of SUSY searches based on simplified models by both the ATLAS and CMS experiments. Section 8.2 then begins by describing the kinematics and phenomenology of the three-body $\tilde{t}_1$ decays that are being targeted in the SUSY analysis presented in this chapter. With the signal kinematics in mind, Section 8.3 describes the overall event selection and object definition used in the analysis, so that the analysis' event sample can be understood. Section 8.4 describes how the phenomenology described in Section 8.2 motivates a particular choice of discriminating observables that are sensitive to the three-body decay of the $\tilde{t}_1$ particle, as well as the definition of the analysis' SRs. Section 8.5 describes the analysis' background estimation strategy, particular the definition of CRs and VRs used to constrain the dominant SM backgrounds. Section 8.6 closes the chapter with the results of the search for the three-body $\tilde{t}_1$ in the two-lepton final state.

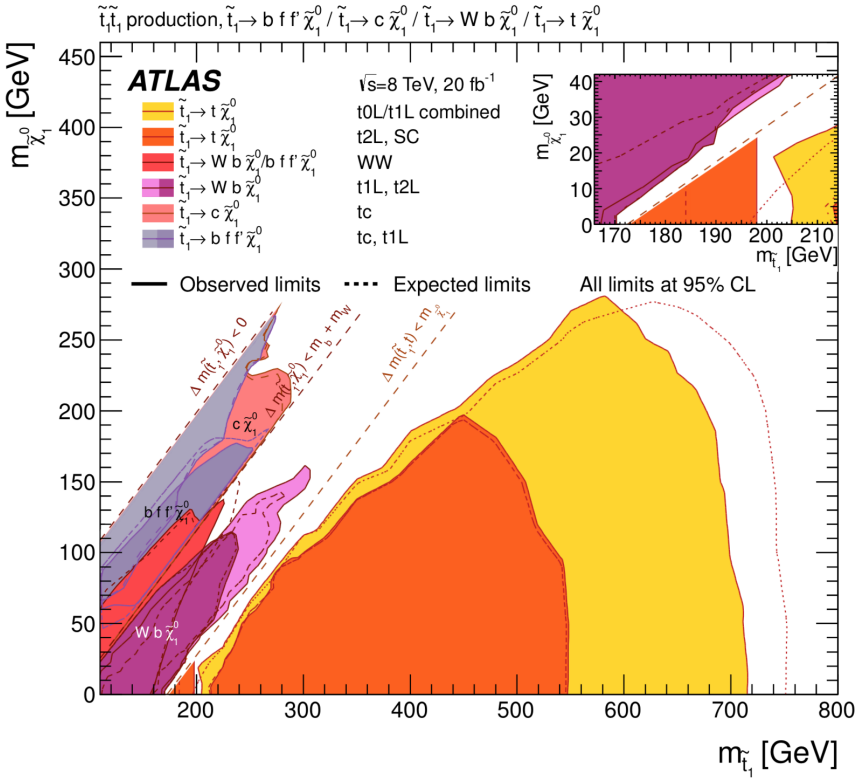

Figure 8.3: Summary of SUSY $\tilde{t}_1$ searches in the $(\tilde{t}_1, \tilde{\chi}_1^0)$ plane at the end of Run 1. The searches relevant to the work in the present thesis are indicated by the coverage of purple and dark orange: 't1L,t2L' and 'WW', respectively. Figure taken from Ref. [190].

8.1 Aside: SUSY Exclusion Plots

In this section we describe the typical SUSY exclusion plot, illustrated in Figure 8.4. These plots are used when presenting the results of SUSY searches based on the assumptions made in simplified models of the MSSM, parametrised only by the masses of the pair-produced sparticle and the stable sparticle appearing in the final state (the LSP), assuming R-parity conservation. The main things to observe in this plot are as follows:

1. The region contained *within* the contours is the region of SUSY parameter space, under the assumption of the simplified model, excluded at 95% CL
2. The region in which the mass of the produced sparticle is less than the mass of the LSP is kinematically forbidden
3. The analysis sensitivity (i.e. power to exclude) decreases as the mass of the produced sparticle increases

4. The analysis sensitivity (i.e. power to exclude) decreases as $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \to 0$ ($\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \to m_X$, with m_X a mass scale related to an intermediate state appearing in the sparticle decay)

5. The expected 95% CL exclusion limits are reported with their $\pm 1\sigma$ uncertainty band, related to the precision of the analyses based on their systematic uncertainty evaluation

6. The observed 95% CL exclusion limit is not reported with an error

Item (3) in the above is related to the fact that the cross-section of the SUSY sparticle production cross-section is a function of only the mass of the pair-produced sparticle in the simplified models assumed here. The production cross-section decreases in proportion to the mass of the sparticle, indicated in Figure 8.5, which leads to analyses becoming less sensitive to the scenarios with larger masses of the produced sparticles.

Item (4) is also relatively straightforward. As the masses of the produced sparticle and final-state LSP become nearer to one another, and even degenerate, kinematic phase space suppression results in there being less energy-momentum available to be transferred to the final state particles. This results in final state particles having typically lower energies and momenta, making them difficult to trigger on and reconstruct efficiently. This inefficiency reduces an analysis' sensitivity to these kinematic regimes, since the final states become experimentally inaccessible. This scenario also appears when the produced sparticle follows a more complex decay chain, as in the cases illustrated in Figure 8.1, where intermediate mass scales such as the mass of the SM top-quark or W-boson are introduced. Near these regions, the kinematics of the particles decaying from the objects defining these intermediate mass scales change rapidly. As a result, analyses targeting particle kinematics on one side of such a boundary cannot simultaneously be sensitive to the particle kinematics on the opposing side of the boundary. For this reason, it is common practice to define separate analyses targeting each side of a given boundary, as indicated in Figure 8.3.

The expected limits in Item (5) in the above are related to running the analysis' hypothesis tests on the SUSY scenarios indicated by the position on the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$-plane, and using for $\boldsymbol{N}^{\text{obs}}$ (c.f. Equations 7.17 and 7.18) the value in the SRs for the predicted background; that is, the observed data counts in the analysis' signal regions are substituted with the value of the total predicted background in the signal regions. The observed data is still used in the CRs, which provide the normalisation correction factors as described in Section 7.4.2. The observed limits discussed in Item (6), on the other hand, take as $\boldsymbol{N}^{\text{obs}}$ the actual observed data counts in the analysis' SRs. Due to statistical fluctuations in the observed data, the expected and observed 95%CL exclusion contours do not typically overlap perfectly.

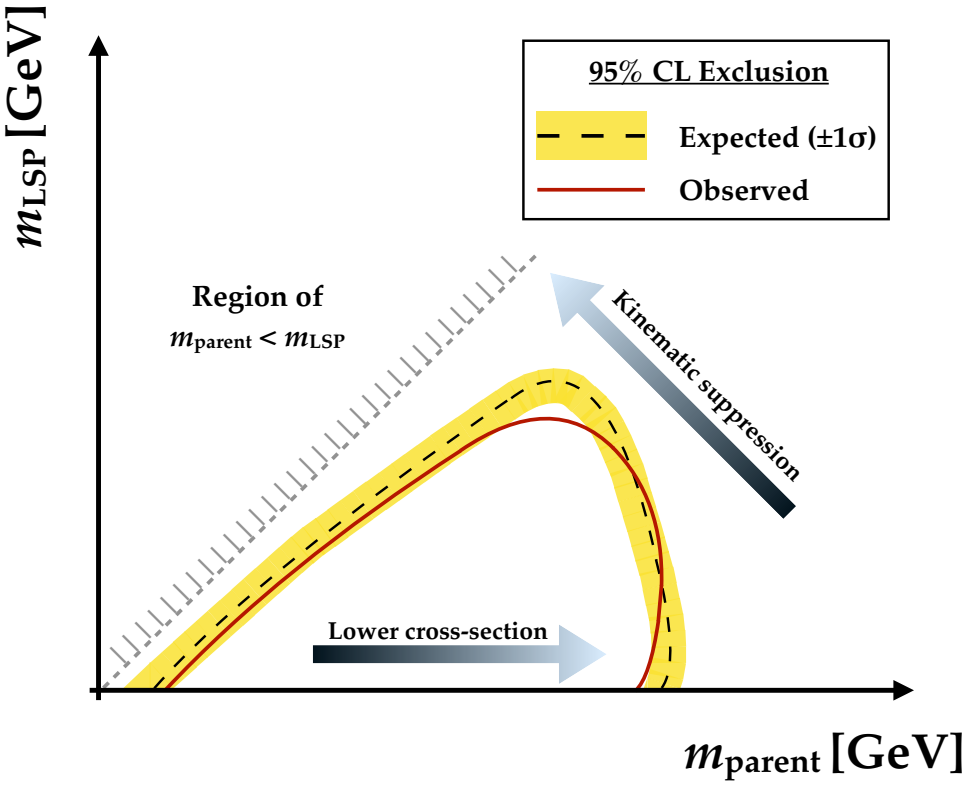

Figure 8.4: Cartoon illustrating an example of a typical two-dimensional exclusion plot for a SUSY simplified model. The mass of the pair-produced SUSY particle is on the x-axis and the mass of the LSP to which it decays is on the y-axis. The region bounded by (i.e. inside of) the dashed-black (red) curves indicate the expected (observed) region of SUSY parameter space excluded at the 95% CL. The $\pm 1\sigma$ error band is shown on the expected limits, and is based on the systematic uncertainties included in the analysis. SUSY exclusion curves such as these usually exhibit a triangular shape as a result of two competing effects, indicated by the arrows: 1) the reduction in production cross-section as the mass of the produced SUSY particle increases, and 2) kinematic suppression occurring when the phase space available to the final state decreases as a result of $\Delta m(x, y) = m_x - m_y \to 0$. Kinematic suppression results in the final state objects, such as leptons and b-tagged jets in the case presented in the current thesis, having very little momenta, thereby making it difficult to efficiently identify the event as being consistent with the given SUSY hypothesis being searched for.

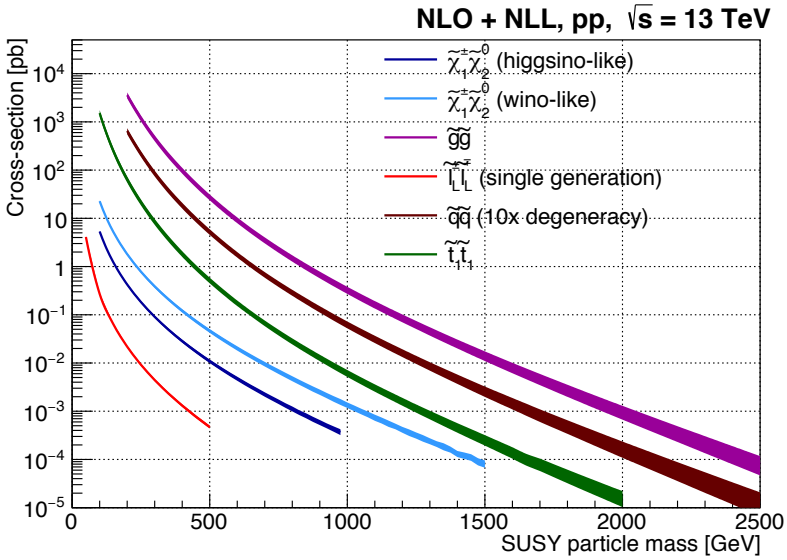

Figure 8.5: Production cross-section for the SUSY sparticles indicated in the legend, as a function of their mass, computed as in Ref. [191]. The bands indicate the uncertainty due to the theoretical prediction of the production process. The green band indicates the $\tilde{t}_1$ production cross-section, relevant to the analysis described in the current chapter. For comparison, the production cross-section for SM top-quark pair-production is $\approx 830\,\text{pb}$ at $13\,\text{TeV}$.

8.2 Description and Phenomenology of the Signal

8.2.1 Aside: SUSY Signal Grids

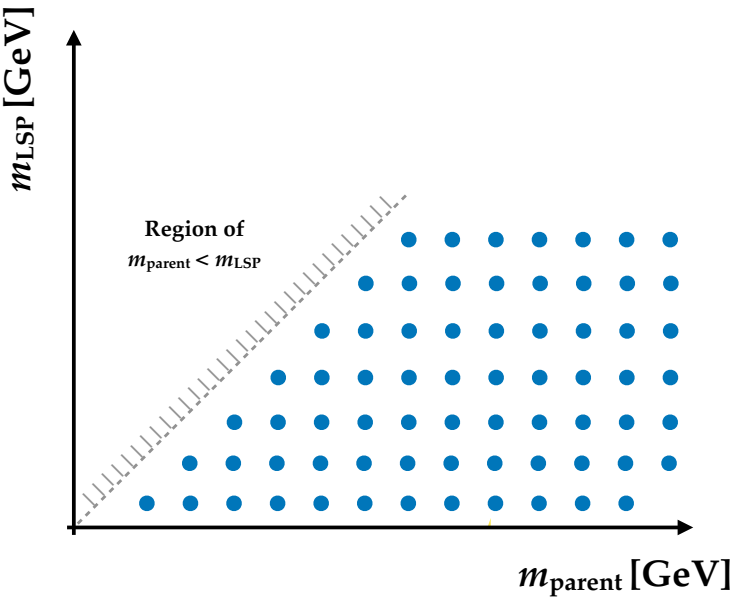

Figure 8.6: Illustration of a SUSY signal grid. The blue points indicate points in the simplified model parameter space for which an individual MC-simulated sample is produced and available for use in analyses searching for the given type of SUSY.

Before the search for the $\tilde{t}_1$ can begin, MC simulated samples representing this process must be obtained. Once a given simplified model of the MSSM is chosen, specific points in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$-plane are chosen and for each a separate MC simulated sample is produced. Each point represents a different hypothesis of SUSY, characterised by the specific masses of the $\tilde{t}_1$ and $\tilde{\chi}_1^0$ sparticles. In ATLAS, this process of choosing points in the SUSY parameter space to produce MC simulated samples is one requiring high levels of deliberation and discussion. This is due to the fact that the simulation of complicated physics processes requires large amounts of ATLAS' CPU resources, of which any given analysis group only has so much allocated, and so, among other things, judicious choices about the number of events to be simulated at each point in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$-plane must be made, as well as careful consideration of previous analyses' exclusion coverage in the same parameter space. For this reason, the resulting SUSY 'signal grids' are *sparse*, as illustrated in Figure 8.6. Depending on how sparse the produced signal grid is, how the points are grouped, or how large each MC-sample at each point is (in terms of number of MC events simulated), the resulting analyses' conclusions can result in unphysical or even discontinous exclusion regions, which is clearly the case in the Run 1 searches for the three-body decay of the $\tilde{t}_1$, seen in Figure 8.3. When interpreting the results of SUSY searches, then, it is important to keep these facts in mind so as to

not over-interpret the results' statements about any fine-grained details of SUSY parameter space to which the analyses are generally not sensitive. Each point in a SUSY signal grid allows, rather, an analysis to make general statements about an *extended*, but still local, region of parameter space under the assumed simplified model.

8.2.2 Description of the Three-body Decay Final State

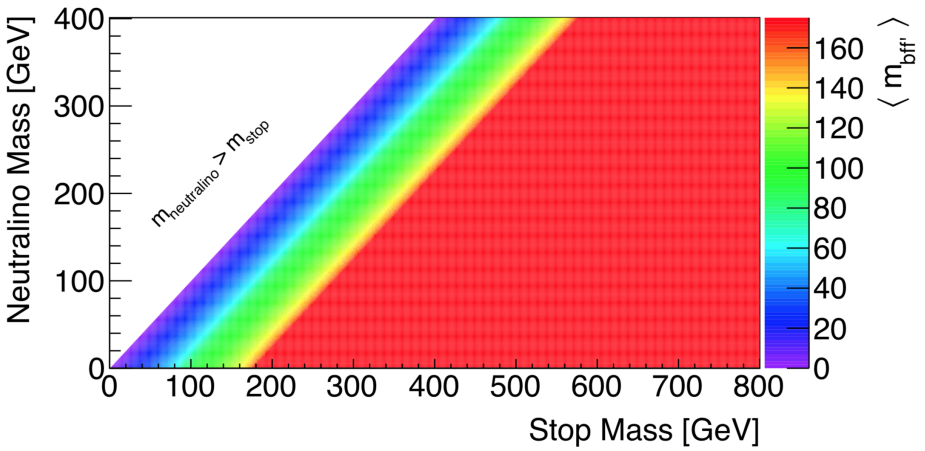

Figure 8.7: Average invariant mass of the b-quark and SM fermions (ff') in the final state of the $\tilde{t}_1 \to bff'\tilde{\chi}_1^0$ decay, as a function of the mass of the $\tilde{t}_1$ and $\tilde{\chi}_1^0$ particles. Figure taken from Ref. [192].

In the three-body decays, the $\tilde{t}_1$ is assumed to decay via an off-shell SM top quark. This is a kinematically suppressed region of phase space, as illustrated in Figure 8.7, which shows the average invariant mass of the $\tilde{t}_1$ decay products across the two-, three-, and four-body decay regions of the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$-plane (c.f. Figure 8.2). In Figure 8.7 the onset of the restricted phase space is clearly visible as one drops beneath the $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) = m_{\text{top}}$ line, at which point the final state particle kinematics change rapidly.

The restricted phase space of the three-body $\tilde{t}_1$ decays has a large impact on the analyses searching for the $\tilde{t}_1$. The b-quark children of the $\tilde{t}_1$ decays have decreasing momenta as one moves from the $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) = m_{\text{top}}$ kinematic boundary to that of $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) = m_W$. As a result, the b-jets that arise as a result of the b-quark hadronization process are difficult to identify efficiently given the minimum 20 GeV p_T thresholds typical of the ATLAS jet reconstruction and flavor-tagging identification algorithms. This is illustrated in Figure 8.8, showing the b-tagged jet multiplicities for various $\tilde{t}_1$ decay scenarios and how they compare to the dominant SM backgrounds. As $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \to m_W$, the flavor-tagging algorithms are unable to identify the b-jets, and the final state is kinematically similar to that of SM WW production.

The $\tilde{\chi}_1^0$ children of the $\tilde{t}_1$ decays, as with the b-quarks, also exhibit diminishing momenta as

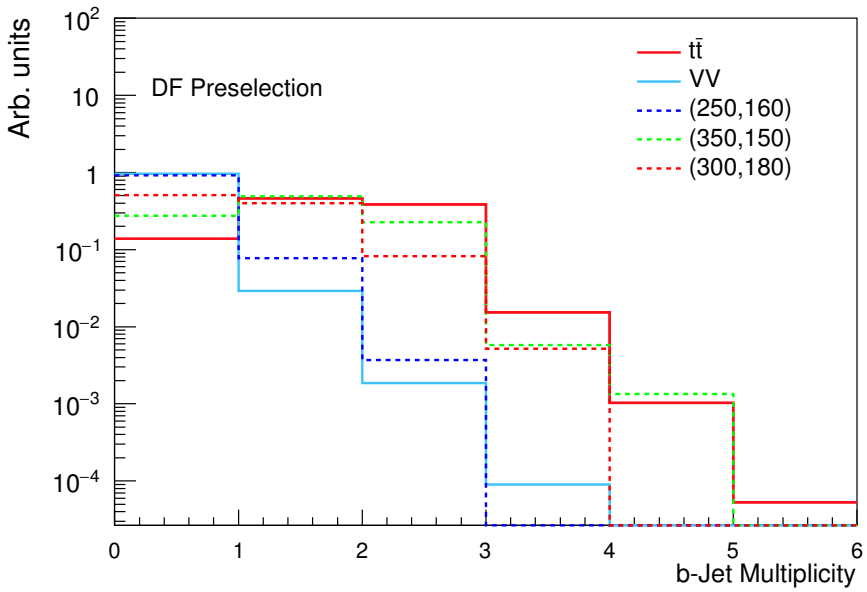

Figure 8.8: Normalised b-tagged jet multiplicity distributions for three $\tilde{t}_1$ signals (dashed lines), indicated by their position in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$-plane in the legend, and SM $t\bar{t}$ (solid red red) and diboson (VV) (solid blue line) production. It can be seen that the $\tilde{t}_1$ decays with $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \sim m_{\text{top}}$ have b-tagged jet multiplicities similar to that of SM $t\bar{t}$production. The $\tilde{t}_1$ decays with $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \rightarrow m_W$ tend to have zero reconstructed b-tagged jets. The selection, 'DF Preselection', in the plot is that described in Table 8.3.

one traverses the three-body region from the $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) = m_{\text{top}}$ kinematic boundary to that of $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) = m_W$. As a result, the three-body $\tilde{t}_1$ decays do not have the typical signature of large amounts of missing transverse momentum, typical of R-parity conserving SUSY decays. The reduced impact of the $\tilde{\chi}_1^0$ particles on the magnitude of the missing transverse momentum furthers the three-body $\tilde{t}_1$ decays' resemblance to those of SM WW production, in that the primary carriers of missing transverse momentum are the neutrinos from the leptonic decays of the W bosons.

The three-body region allows for the on-shell production of the W bosons. For this reason, the leptons and neutrinos from their decay are not kinematically suppressed in the same way as the b-quarks and $\tilde{\chi}_1^0$. Their kinematics are therefore typical, in scale, to those of SM $t\bar{t}$ and WW production.

8.3 Event Selection and Object Definitions

In this section, the definitions of the event sample and of the reconstructed objects used in the analysis are given.

8.3.1 Trigger Selection

As described in Section 8.2.2, the W-bosons in the three-body $\tilde{t}_1$ decays are on-shell and therefore lead to leptons whose kinematics are not kinematically suppressed. Each of the leptons is expected to have similar kinematics, given the relatively uncorrelated nature of their decay, with transverse momenta on average at the scale of $m_W/2 \sim 40\,\text{GeV}$. For this reason, the analysis can rely entirely on triggers based on the leptons. The set of triggers used in the search for $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ is composed entirely of dilepton triggers, having online lepton p_T thresholds of of $10 - 18\,\text{GeV}$.

8.3.2 Object Definitions

Here we describe the definitions of the leptons and jets used in the search for the production of $\tilde{t}_1$ quarks. The lepton (electron and muon) definitions are given in Table 8.1 and the jet definitions are given in Table 8.2. Discussion of the working points for lepton identification and isolation is given in Section 5.2. That of jets is found in Section 5.3 and 5.3.4, for jets, generically, and b-tagged jets, respectively.

The p_T requirements of the leptons are such that they be on the plateau of the trigger efficiency, as described in Section 7.1.2. The lepton isolation requirements are rather loose, given that a large contamination of fake and non-prompt leptons is not expected as a result of the requirement of two relatively high-p_T leptons. The requirements on the impact parameter quantities $|d_0/\sigma_{d_0}|$ and $|z_0 \times \sin\theta|$ further ensure that the leptons are likely to be prompt and have originated from the primary hard-scatter vertex.

Table 8.1: Lepton definitions for the 2015+2016 analysis searching for the $\tilde{t}_1$ quark.

	Leptons			
	Electrons		**Muons**	
	Baseline	**Signal**	**Baseline**	**Signal**
p_T requirement [GeV]	$(>10, >10)$	$(>25, >20)$	$(>10, >10)$	$(>25, >20)$
$\lvert\eta\rvert$ requirement	< 2.47		< 2.4	
Identification WP	`Loose`	`Medium`	`Medium`	
Isolation	`GradientLoose`			
$\lvert d_0/\sigma_{d_0}\rvert$	$--$	< 5	$--$	< 3
$\lvert z_0 \times \sin\theta\rvert$ [mm]	$--$	< 0.5	$--$	< 0.5

Table 8.2: Jet definitions for the 2015+2016 analysis searching for the $\tilde{t}_1$ quark.

	Jets	***b*-tagged Jets**
p_T requirement [GeV]	> 20	
$\|\eta\|$ requirement	< 2.8	< 2.4
Pileup suppression	JVT > 0.59 if $p_T < 60\,\text{GeV}$ and $\|\eta\| < 2.4$	
Flavor-tagging WP	$--$	77%

8.3.3 Standard Event Pre-selection

The starting sample of events, referred to as the analysis' *preselection* level, is defined simply and requires that events satisfy the following basic 'data quality' (DQ) requirements:

- No issues in any of the ATLAS subsystems were detected when the event was recorded
- A primary vertex with at least 2 tracks must be present in the event
- No cosmic muons[1] can be in the event
- If a poorly reconstructed jet is found in the event, likely to have arisen from stochastic noise bursts or issues in the calorimeter system, the event is rejected

Additionally, all events entering the analysis must also have two baseline level leptons (c.f. Table 8.1) which have opposite electric charges, and must have a dilepton invariant mass $m_{\ell\ell} > 20\,\text{GeV}$. This latter requirement on $m_{\ell\ell}$ removes backgrounds due to low-mass Z-boson resonances, such as the J/ψ ($m \sim 5\,\text{GeV}$) and Υ ($m \sim 11\,\text{GeV}$) mesons. The complete preselection is summarized in Table 8.3.

Table 8.3: Preselection summary for the 2015+2016 analysis searching for the $\tilde{t}_1$ quark. Signal level object requirements (Table 8.1 and 8.2), as well as trigger requirements, are applied to the sample of events satisfying these requirements.

Selection	**SF Preselection**	**DF Preselection**
Lepton multiplicity	$==$ 2 signal leptons	
Dilepton flavor	ee or $\mu\mu$	$e\mu$ or μe
DQ requirements	satisfied	
Dilepton charge requirement	opposite charge (OS)	
Dilepton invariant mass, $m_{\ell\ell}$ [GeV]	> 20	

[1] A cosmic muon is a reconstructed muon that is not consistent with having originated from the primary vertex, and therefore likely to have originated from a cosmic ray air shower. If a reconstructed muon does not satisfy both $|z_0| < 1\,\text{mm}$ and $|d_0| < 0.2\,\text{mm}$, it is considered a cosmic muon.

8.4 Signal Selection

8.4.1 The Recursive Jigsaw Reconstruction

The $\tilde{t}_1$ search presented here makes use of of a high-level event reconstruction technique referred to as the 'Recursive Jigsaw Reconstruction' (RJR) technique [193]. This technique is a generalisation of the methods introduced in Ref. [194], referred to as the 'Super Razor'.

The RJR technique takes the measured information in an event — the observable objects and their four momenta, as well as the missing transverse momentum — and, through a series of assumptions and algorithms is able to reorganize those objects in such a way as to be interpreted as 'fitting' into a user-specified, or rather user-*imposed*, decay tree. Indeed, the technique's name derives from the methods it employs to perform this reorganization. The 'recursion' implied in the technique's name refers to the way in which its rules and algorithms are implemented in order to *navigate* through the decay tree: stepping through the decay tree one level at at time, using information only from the current reference frame associated with the object in the decay tree in that level, to determine the properties of the subsequent ones. This effectively means that conclusions drawn at the 'upper levels' of the decay about the kinematic properties of the particles present are not lost or altered as one traverses further down the decay chain. The 'jigsaw' in the technique's name refers to the sets of rules and/or algorithms that are applied — so-called 'Jigsaw Rules' — in order to either piece together or break apart the reconstructed event provided the kinematic information (lab-frame four-vectors) that the user has provided as input to the decay tree assumption.

As an illustrative example of this decay tree 'fitting', consider the $t\bar{t}$ and $\tilde{t}_1 \to b\tilde{\chi}_1^{\pm}$ decays illustrated in Figure 8.9, which provides an example of an RJR decay tree imposition. As will be outlined later, the RJR technique determines the relative velocities (boosts) between the different states or rest-frames present in a given RJR decay tree. These boosts between frames are illustrated diagrammatically as the black lines in the lower portion of Figure 8.9 connecting the various states, both intermediate ('Decay States') and final ('Visible States' or 'Invisible States'), which are represented as the circles. Although the two processes differ, they both fit topologically the RJR decay tree in Figure 8.9. In both cases, the final state is composed of two b-tagged jets, two leptons, and missing transverse momentum and one can intuitively see how to allocate these objects' four-momenta to the states in the RJR decay tree: the $V_{1\,(a,b)}$ states are the b-tagged jets, the $V_{2\,(a,b)}$ states are the leptons, and the $I_{a,b}$ represent the neutrinos (plus the LSPs in the case of the $\tilde{t}_1 \to b\tilde{\chi}_1^{\pm}$ signal). Although the two processes may fit the same RJR tree, they are of course quite different kinematically (e.g. the $I_{a,b}$ states for $\tilde{t}_1 \to b\tilde{\chi}_1^{\pm}$ are composed of additional *massive* particles in addition to the neutrinos) and any variable derived using the four-vectors derived from this decay tree imposition has the potential to tease out these differences between the signal and background processes.

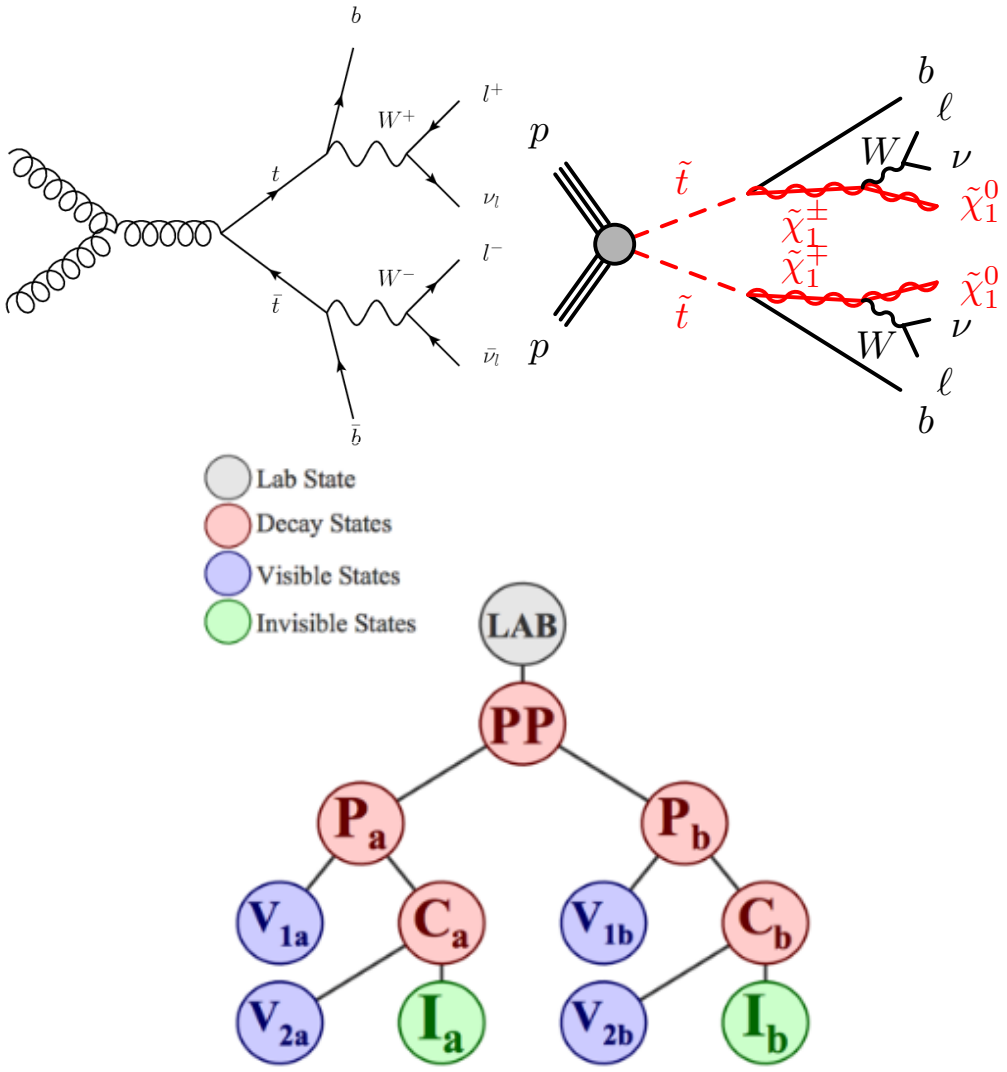

Figure 8.9: An example of an RJR decay tree interpretation of physics processes. The RJR decay tree (***bottom***) can be fit to both the $t\bar{t}$ decay (***upper left***) or the $\tilde{t}_1 \to b\tilde{\chi}_1^\pm$ process (***upper right***). Each of the upper processes *topologically* matches that of the RJR decay tree, but the underlying differences in their kinematics means that kinematic observables derived from this RJR decay tree may provide means of discrimination between the two. Each circle in the RJR decay tree represents a reconstructed reference frame, characterised by its own four-vector information.

The RJR decay tree imposed in the present analysis searching for the three-body decay of the $\tilde{t}_1$ quark is presented in Figure 8.10. The visible system $(V_a + V_b)$ is provided only the two leptons in the event. The invisible system $(I_a + I_b)$ is provided the missing transverse momentum.

The RJR decay tree in Figure 8.10 is the most general RJR decay tree one can make for targeting R-parity conserving SUSY models. This decay tree makes the least number of assumptions on the decay of the pair-produced sparticles and their children (if any). This is beneficial, as developing a search (whose a-priori intent is for discovery of new physics) that is largely dependent on the model

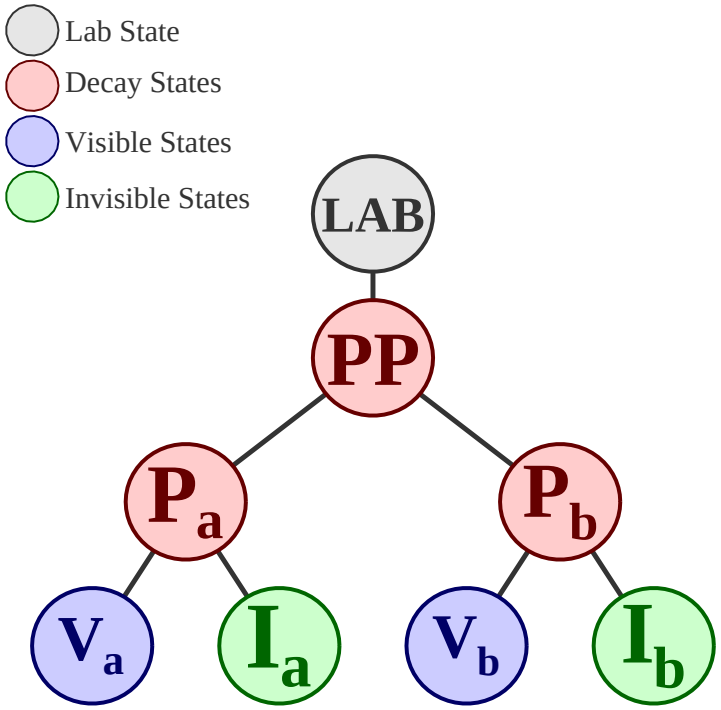

Figure 8.10: RJR decay tree assumption used in the 2015+2016 analysis searching for the three-body decay of the $\tilde{t}_1$ quark, $\tilde{t}_1 \to bW\tilde{\chi}_1^0$. It is the most general decay tree for R-parity conserving SUSY scenarios in which pair-produced sparticles (P_a and P_b) each decay to visible states (V_a and V_b) and invisible states (I_a and I_b). Each of the final states (V_i and/or I_i) may, in principle, be composed of multiple particles. In the present analysis, only the two leptons in the event are provided as inputs to the V_i and both are required to be non-empty (i.e. V_a is one of the leptons, V_b the other). The missing transverse momentum in the event is decomposed via a set of Jigsaw Rules into the states I_a and I_b.

assumed for designing the analysis' selection criteria greatly limits the applicability and scope of that search. For example, one can imagine using the decay tree of Figure 8.9 also for the three-body decays considered in the present search, $\tilde{t}_1 \to bW\tilde{\chi}_1^0$. However, the decay tree in Figure 8.9 assumes that there is enough information to kinematically separate the b-tagged jets and leptons into the $V_{1(a,b)}$ and $V_{2(a,b)}$ states, respectively, therefore requiring at least four visible objects to be reconstructed. As discussed in Section 8.2.2, this is not the case for the three-body decays of the $\tilde{t}_1$ relevant for the present analysis and so Figure 8.9 is not well-suited for this search. What we want, in the end, is a means to separate from the background the presence of an R-parity conserving SUSY signal while making as few assumptions as possible.

As mentioned, the goal of the RJR technique is to provide a full accounting of all of the states specified by the user in the imposed RJR decay tree; or, at least a close and/or optimal approximation of them. Unfortunately, it is in general impossible to do this as a result of the presence of multiple weakly interacting particles in the final state and, less generally, because of combinatoric

ambiguities that make it difficult to assign objects to specific locations in the decay (e.g. associating b-jets to the correct side of the decay). The RJR technique and its so-called Jigsaw Rules provide a means to systematically resolve these unknowns on an event-by-event basis.

In order to resolve combinatoric ambiguities amongst the reconstructed visible objects so that they may be grouped 'correctly', one can employ a Jigsaw Rule for performing a recursive partitioning of the visible objects until that grouping which minimizes the visible mass of each side of a decay vertex is found. This is not necessarily the true grouping, but it is a method that is found to reproduce well, on average, the correct grouping of objects.

In the RJR decay tree used in the search for $\tilde{t}_1 \to bW\tilde{\chi}_1^0$, illustrated in Figure 8.10, there are 6 under-constrained degrees-of-freedom (DOF) due to the weakly interacting particles/missing transverse momentum. These are (2 DOF for each) as follows:

- The longitudinal momenta of the I_i states: $p_{I_i,z}$
- The splitting of the missing transverse momentum between the I_i states: $\boldsymbol{p}_{I_a,T}+\boldsymbol{p}_{I_b,T} = \boldsymbol{E}_T^{\text{miss}}$
- The mass of the di-invisible system composed of I_a and I_b: m_I (m_{I_a} and m_{I_b})

The RJR technique sets out to provide values for these unknowns through assumptions (constraints) that it makes for determining the relative velocities (boosts) between rest-frames indicated by the circles in Figure 8.10. The recursive element of the technique is most obvious in this respect: the algorithm moves from the first known reference frame (the lab frame) and traverses down the decay tree using information only from the current frame to determine the boosts for moving into the next reference frame. That is, with respect to Figure 8.10, information only from the lab frame is used to move to the PP (center-of-mass, COM) frame and information from this COM frame is used to move to either of the respective P_i frames. There are specific Jigsaw Rules at each step of this navigating of the RJR decay tree that determine *how* the information in each of these frames is used simultaneously to move to the next and to provide estimates of the under-constrained DOF mentioned above.

The first step that the RJR technique takes in constructing the boosts relating the reference frames is to determine the mass of the invisible system, m_I, composed of the states I_a and I_b. This information is required to be able to consistently perform boosts between the reference frames while keeping each side, or *hemisphere*, of the decay balanced against the other. This means that, generally, the di-invisible system will have non-trivial opening angles between the I_i states in order to balance the visible system frame-by-frame. This means that it will attain a mass[2] To satisfy this balancing requirement, the RJR technique takes for m_I the smallest Lorentz-invariant mass

[2]From special relativistic mechanics, the mass of a system comprised of two subsystems, 1 and 2, is generally given by $m_{12}^2 = m_1^2 + m_2^2 + 2(E_1E_2 - |\vec{p}_1||\vec{p}_2|\cos\theta_{12})$, where θ_{12} is the opening angle between $\vec{p}_1$ and $\vec{p}_2$.

consistent with the input lab-frame observables that will accommodate the subsequent boosts and also prevent the states I_i from becoming tachyonic. For the decay tree in Figure 8.10 this turns out to be,

$$m_I^2 = m_V^2 - 4m_{V_a}m_{V_b}, \tag{8.1}$$

where V refers to the total visible system comprised of V_a and V_b.

Next, the RJR technique determines the boost from the lab to the PP (COM) frame. To determine this boost, we must determine the longitudinal momentum of the invisible system, $p_{I,z}^{\text{Lab}}$, which is one of the under-constrained DOF. The RJR technique chooses this value such that the rapidity of the visible and invisible systems are equal. This choice results in our choice of the PP (COM) reference frame being a longitudinally boost-invariant choice, meaning that all subsequent observables derived in the RJR decay tree will also be longitudinally boost-invariant. Additionally, one can see that this also forces the mass of the PP system, m_{V+I}, to take its minimum value.[3]

With the longitudinal momentum of the invisible system determined, along with its mass, we have the expression for the boost relating the lab frame to the PP (COM) frame:

$$\vec{\beta}_{PP}^{\text{LAB}} = \frac{\vec{p}_{PP}^{\text{LAB}}}{E_{PP}^{\text{LAB}}} = \frac{\vec{p}_V^{\text{LAB}} + \vec{p}_I^{\text{LAB}}}{E_V^{\text{LAB}} + \sqrt{|\vec{p}_I^{\text{LAB}}|^2 + m_I^2}}. \tag{8.2}$$

The boost in Equation 8.2 affords provides us the observables in the PP (COM) reference frame. We must use information in this frame to determine how to move to the individual sparticle frames, P_i. This will provide us with the last set of information regarding how the missing transverse momentum is to be shared between the states I_i. Here, the RJR technique makes the assumption that $m_{V_a} = m_{V_b}$, which, in the context of the processes considered, is reasonable. As we are in the COM frame, this choice for the masses of the V_i states dictates that the boosts of each of their individual reference frames be equal in magnitude and anti-parallel. For our decay tree, a possible solution for this is chosen to be:

$$\vec{\beta}_{P_i}^{PP} = \frac{\vec{p}_{V_a}^{PP} - \vec{p}_{V_b}^{PP}}{E_{V_a}^{PP} + E_{V_b}^{PP}}. \tag{8.3}$$

Under this boost, all observables subsequently defined are *contra-boost invariant*, meaning that, as in the case for the longitudinal boost invariance described previously, all observables are therefore insensitive (on average) to the fact that this boost from the PP (COM) frame to the P_i frames is not necessarily the true boost. Indeed, this final contra-boost invariance is the main motivation behind the structure of Equation 8.3.

[3]This is seen using the previous footnote and knowledge that $\vec{E}_T^{\text{miss}}$ must balance the visible system's transverse momentum.

Now that we have the boost $\vec{\beta}_{P_i}^{PP}$ to the P_i frames, one can determine the splitting of invisible momentum by requiring $\vec{p}_{V_i}^{P_i} + \vec{p}_{I_i}^{P_i} = 0$, which is the final under-constrained DOF of the RJR decay tree relevant to the analysis being presented.

With the application of the RJR technique, then, all of the unknowns in the event are thus determined (approximately), and one has access to kinematic information (four vectors) for every level of the RJR decay tree that has been imposed (Figure 8.10).

8.4.2 Discriminating Observables

In this section we describe the basis of kinematic observables, defined using the states defined by the RJR, that are used to define the SRs, CRs, and VRs in the present analysis.

With the RJR technique providing us with an approximation of the center-of-mass frame of the $\tilde{t}$ system (the state PP in Figure 8.10), a natural variable to define is the total energy, or invariant mass, available to those objects in our decay tree. With the entire system defined, this can be obtained easily and we refer to this approximate COM energy as m_{PP}, shown in Figure 8.11.

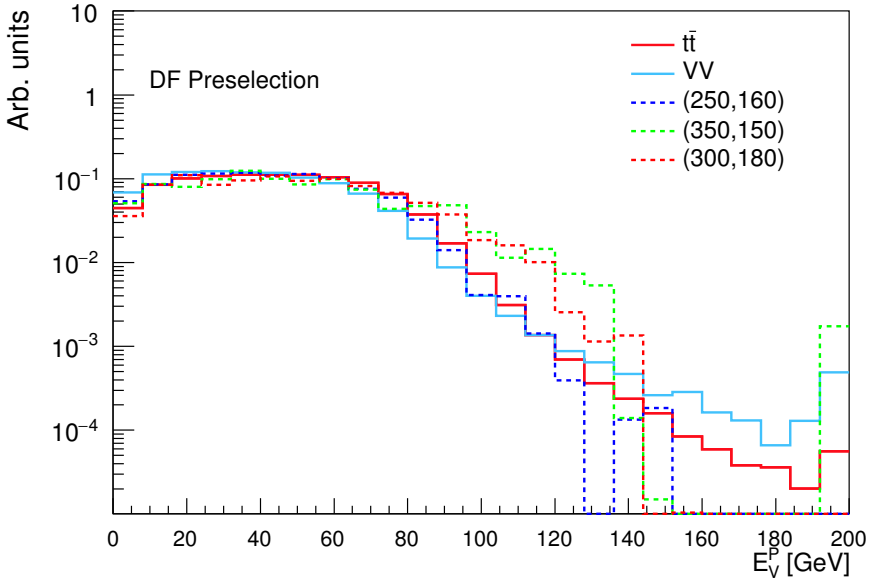

Figure 8.11: Normalised distribution of the m_{PP} variable for several $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ signal hypotheses, labeled with respect to their assumed $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ values in the legend. The SM $t\bar{t}$ and diboson (VV) background processes are also shown for comparison.

In addition to m_{PP}, we can define the associated transverse momentum of the COM frame, $|\vec{p}_T^{PP}|$, shown in Figure 8.12. Used in conjunction, the two variables m_{PP} and $|\vec{p}_T^{PP}|$ can be used to build

the ratio, R_{p_T}, defined as:

$$R_{p_T} = \frac{|\vec{p}_T^{PP}|}{|\vec{p}_T^{PP}| + m_{PP}/4}. \tag{8.4}$$

Distributions of R_{p_T} are provided in Figure 8.13. The more central nature of our signal hypothesis wherein the pair-produced sparticles decay to heavy final state objects and distribute more of their energy to the transverse plane can be seen in Figure 8.13.

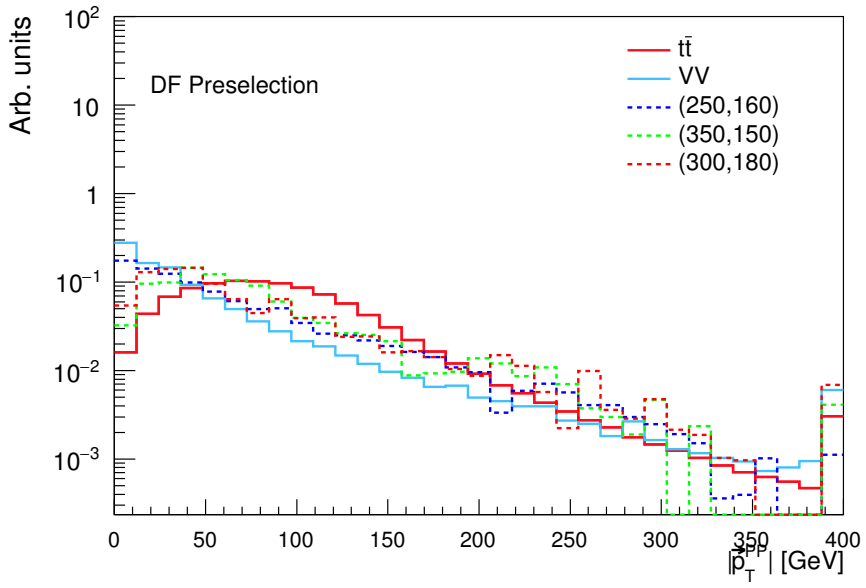

Figure 8.12: Normalised distribution of the $|\vec{p}_T^{PP}|$ variable for several $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ signal hypotheses, labeled with respect to their assumed $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ values in the legend. The SM $t\bar{t}$ and diboson (VV) background processes are also shown for comparison.

The next variable that can be defined uses the relative velocity of the PP frame as seen in the LAB frame, as well as the total visible system $V = V_a + V_b$ as seen in the PP frame. This is the azimuthal angle between the velocity, or boost, from the LAB to PP frame and the visible system V (which, in our case, is simply the dilepton system) as seen in the PP (COM) frame. This variable, $\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP})$, is shown in Figure 8.14.

The key features of the angular quantity $\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP})$ can be seen in Figure 8.14. The $\tilde{t}_1$ signals clearly peak more strongly near values of $\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP}) \sim \pi$ and are roughly overlapping for all values of $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0)$, meaning that the behavior of $\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP})$ is loosely independent on $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0)$. Rather, $\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP})$ tends to be more dependent on the assumptions gone into constructing the underlying RJR decay tree in Figure 8.10, from which the quantity $\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP})$ is ultimately derived. As discussed, with the dilepton system $V = V_a + V_b$, Equation 8.1 results in the assumption $m_I = m_V$. For the massive $\tilde{\chi}_1^0$ particles in the invisible final state of our $\tilde{t}_1$ signal decays, this assumption is clearly incorrect and means that the boost from the LAB frame to the

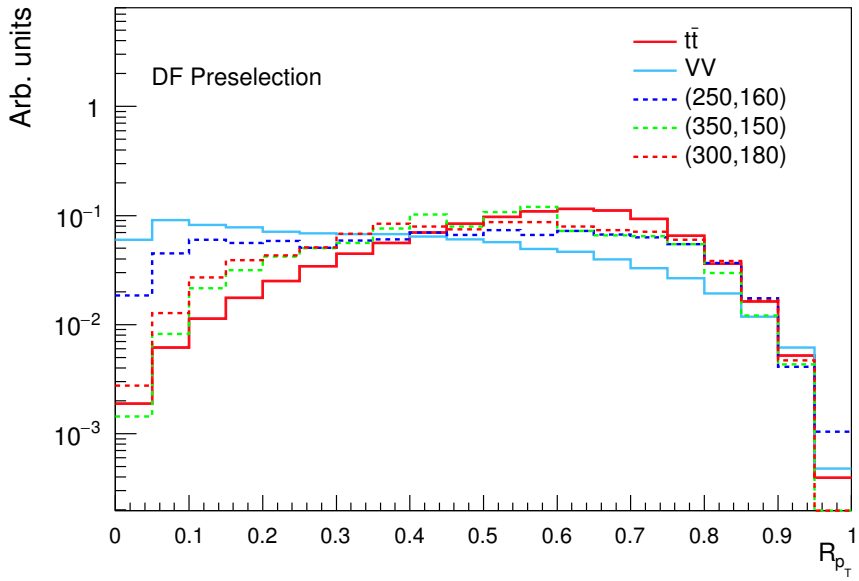

Figure 8.13: Normalised distribution of the R_{p_T} variable for several $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ signal hypotheses, labeled with respect to their assumed $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ values in the legend. The SM $t\bar{t}$ and diboson (VV) background processes are also shown for comparison.

approximate PP (COM) frame tends to be an *over*-boost since the assumption of equal masses assumes too much momentum for the invisible system. In reality, on the other hand, much of the $\tilde{t}_1$ energy-momentum goes into the *mass* of the invisible system. This over-boosting can be seen in Equation 8.2 with the under-predicted m_I in the denominator. This results in the system V to point, on average, in the opposite (anti-parallel) direction with respect to the boost in Equation 8.2 when viewed in the approximate COM frame, PP.

We can also define a variable that is sensitive to the (squared) mass differences between the pair-produced $\tilde{t}_1$ and final state $\tilde{\chi}_1^0$. This observable, corresponding to the energy of one of the V_i (one of the leptons) with respect to the sparticle P_i frames of reference. This variable, E_V^P, can be interpreted as the approximate energy of the state V_i in the frame corresponding to P_i and has an endpoint following,

$$E_V^P \propto \frac{m_{P_i}^2 - m_{I_i}^2}{m_{P_i}}, \tag{8.5}$$

which follows essentially from special relativistic kinematics. Distributions of E_V^P are given in Figure 8.15.

The next Jigsaw variable to be used is the so-called 'visible shape' of the PP (COM) frame. In the limiting case of $m_{V_a} = m_{V_b} = 0$, in which the current analysis finds itself, this quantity corresponds to the (inverse) boost factor (γ) associated with the boost from the PP frame to the P_i frames and

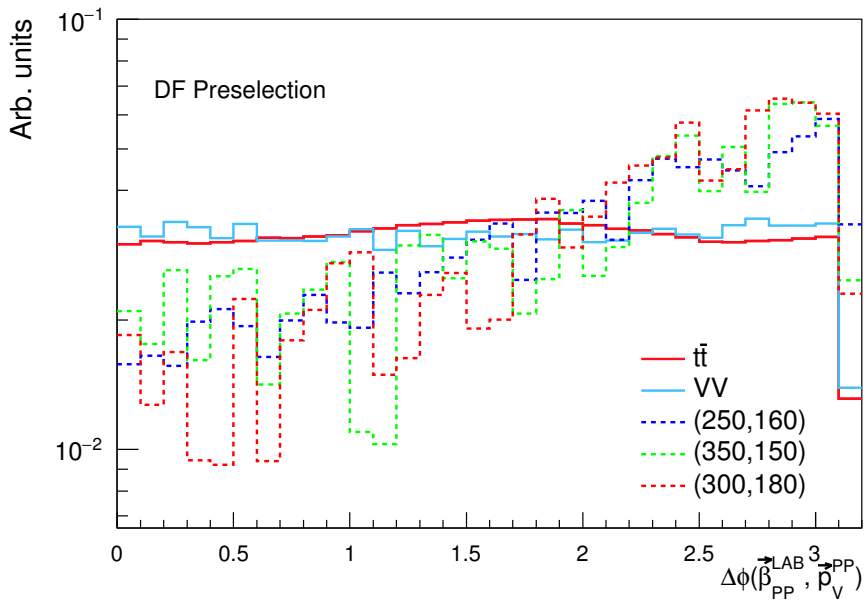

Figure 8.14: Normalised distribution of the $\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP})$ variable for several $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ signal hypotheses, labeled with respect to their assumed $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ values in the legend. The SM $t\bar{t}$ and diboson (VV) background processes are also shown for comparison.

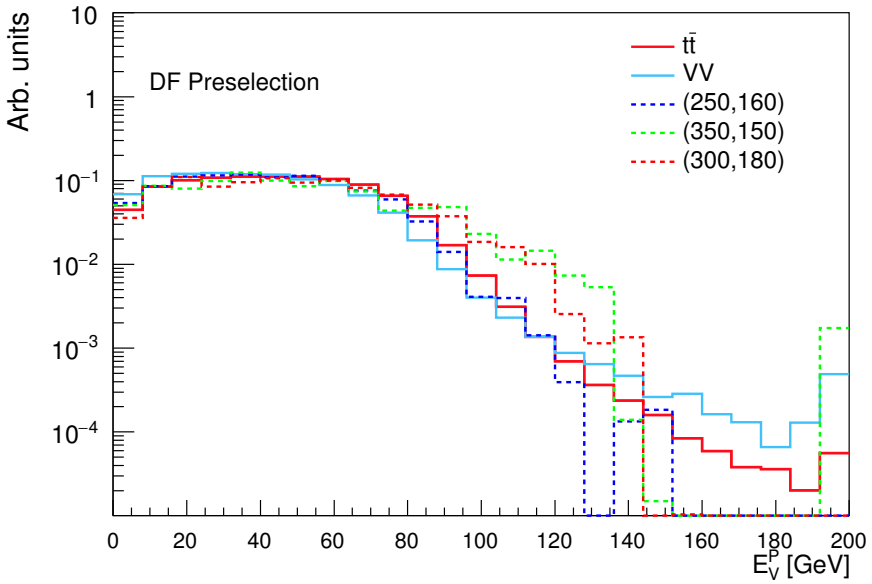

Figure 8.15: Normalised distribution of the E_V^P variable for several $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ signal hypotheses, labeled with respect to their assumed $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ values in the legend. The SM $t\bar{t}$ and diboson (VV) background processes are also shown for comparison.

is defined as follows,

$$\text{PP Visible Shape} \to 1/\gamma_P^{PP} \equiv \frac{\sqrt{2(|\vec{p}_{V_a}^{PP}||\vec{p}_{V_b}^{PP}| + \vec{p}_{V_a}^{PP} \cdot \vec{p}_{V_b}^{PP})}}{|\vec{p}_{V_a}^{PP}| + |\vec{p}_{V_b}^{PP}|.} \qquad (8.6)$$

The quantity $1/\gamma_P^{PP}$ is a measure of how the two visible systems, V_i, are distributed. It tends towards unity when V_a and V_b (the two leptons) are equal in momenta and collinear while tending towards zero when they are back-to-back or having different momenta. In this sense, the variable earns its name as a measure of the 'visible shape' of the decay. In decays with massive particles contributing to the missing transverse momentum, the visible final state legs V_a and V_b will, on average, tend to not be back-to-back since, in addition to balancing each other, they much balance the invisible particles. This effect is exaggerated in two cases: (1) when the decay is highly 'active' (as in $t\bar{t}$), and (2) when the invisible system is composed of heavy particles. The effect (1) is essentially due to our including in the states V_a and V_b only the final state leptons: the quantity $1/\gamma_P^{PP}$ is blind to any jets in the event and the 'shape' it corresponds to is in response to balancing these jets as well. The effect (2) results in the visible legs recoiling harder off of the heavier invisible particles, thus tending to be collinear. Distributions of the $1/\gamma_P^{PP}$ quantity are shown in Figure 8.16.

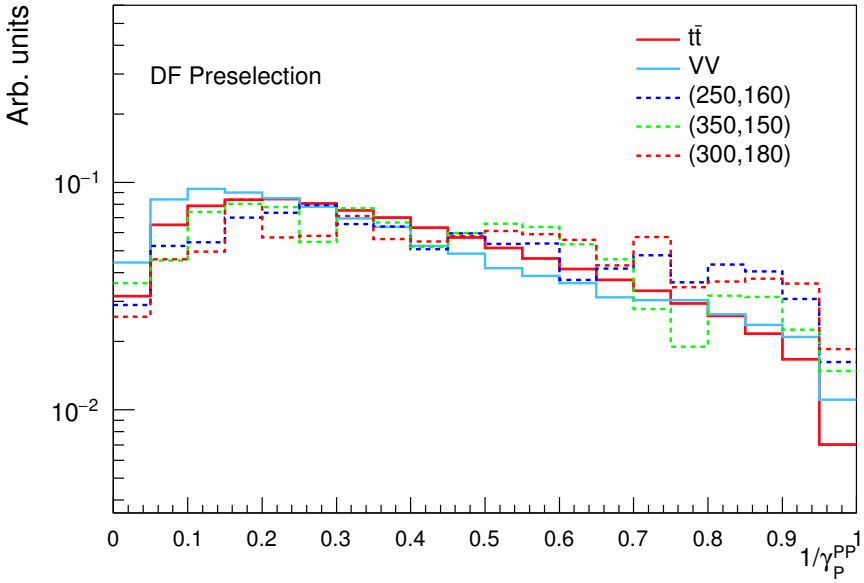

Figure 8.16: Normalised distribution of the $1/\gamma_P^{PP}$ variable for several $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ signal hypotheses, labeled with respect to their assumed $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ values in the legend. The SM $t\bar{t}$ and diboson (VV) background processes are also shown for comparison.

The variables discussed so far are relatively independent of the underlying signal model that we are attempting to search for. Their utility relies simply on their being massive particles contributing to the missing transverse momentum.

We now discuss a variable that, in order for it to be useful in the search for $\tilde{t}_1$ particles, makes an assumption on the nature of the pair-produced sparticle that we are searching for. This variable, $\cos\theta_b$, is angular in nature that is sensitive to the decay angle of a given lepton with respect to the beam axis in it's production frame.

We define $\cos\theta_b$ using only the two final state leptons. Each of the leptons is boosted into the frame of the dilepton system, after which we take the difference in pseudorapidities of the leptons in this frame, $\Delta\eta_{\ell\ell} = \eta^+ - \eta^-$, where η^+ (η^-) refers to the pseudorapidity of the positively (negatively) charged lepton. The quantity $\cos\theta_b$ is then defined as,

$$\cos\theta_b = \tanh(\Delta\eta_{\ell\ell}/2), \tag{8.7}$$

and is discussed further in Ref. [195]. As $\cos\theta_b$ is sensitive to the projection angle of the lepton from its parent $\tilde{t}_1$, it is sensitive to the spin of the $\tilde{t}_1$. This can be seen in Figure 8.17, where we see that for the decay of the spin-zero $\tilde{t}_1$, the angular distribution is flatter and peaks at 0. For the SM $t\bar{t}$ and WW backgrounds, however, $\cos\theta_b$ tends to peak at values ± 1.

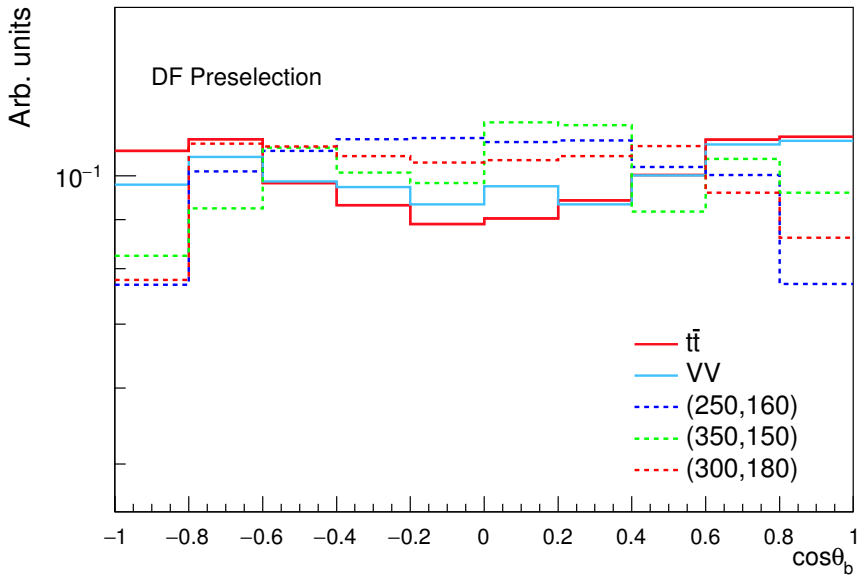

Figure 8.17: Normalised distribution of the $\cos\theta_b$ variable for several $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ signal hypotheses, labeled with respect to their assumed $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ values in the legend. The SM $t\bar{t}$ and diboson (VV) background processes are also shown for comparison.

In addition to the potential separation power between the decays of scalar $\tilde{t}_1$ particles and the dominant SM backgrounds that $\cos\theta_b$ exhibits in Figure 8.17, it should also be noted that it is a longitudinally boost-invariant quantity. It depends only on the *difference* in pseudorapidities of the leptons, $\Delta\eta_{\ell\ell}$, and thus is insensitive to our ignorance of the overall z-momenta in the event. In this sense, the quantity $\cos\theta_b$ follows very nicely the variables obtained from the RJR decay tree, which are also motivated by longitudinal (and contra-) boost invariance.

8.4.3 Signal Region Definitions

The definition of the SRs in the search for the three-body decay of the $\tilde{t}_1$ revolve around two observations:

1. The $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ kinematics change rapidly as one traverses from the region near $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) = m_{\text{top}}$ to that of $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) = m_W$, as illustrated in Figure 8.7
2. The angular quantities $\cos\theta_b$ and $\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP})$, sensitive to the presence of the scalar $\tilde{t}_1$ and massive $\tilde{\chi}_1^0$, respectively, allow for tight separation between the $\tilde{t}_1$ signal and dominant SM background processes

The first item will ultimately result in the definition of two orthogonal SRs, one generally more sensitive to the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ decays nearer to the $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) = m_{\text{top}}$ boundary and the other being sensitive to those nearer to the $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) = m_W$ boundary. The primary method for defining SRs targeting these differing regimes is motivated by Figure 8.8, which shows that the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ decays with $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0)$ nearer to m_W, the reconstructed b-tagged jet multiplicity tends to 0. For this reason, the SRs targeting the $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \to m_W$ will require that the candidate events have zero b-tagged jets, i.e. they apply a b-tagged jet *veto* ('b-veto'). The SRs targeting the $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \to m_{\text{top}}$, on the other hand, will require that there be at least one b-tagged jet in the candidate events.

The second item listed above is best illustrated by the distribution in Figure 8.18. These two-dimensional distributions illustrate that the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ signal tends to populate a different region in the two-dimensional space defined by $(\cos\theta_b, \Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP}))$ than the dominant dilepton SM backgrounds composed of $t\bar{t}$ and diboson production. For this reason, the SRs will be further characterised by applying a selection within the two-dimensional basis defined by the $\cos\theta_b$ and $\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP})$ variables. Figure 8.19 shows similar distributions as those presented in Figure 8.18, but illustrates that the relationship observed in the latter persist after requiring that there be b-tagged jets in the events (the former distributions require a b-veto). The fact that the $(\cos\theta_b, \Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP}))$ relationship holds roughly independent of the b-tagged jet multiplicity means that the same selection in this two-dimensional plane can be applied for the SRs targeting both $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ regimes, $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \to m_W$ and $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \to m_{\text{top}}$.

With the above discussion in mind, the complete basis of observables used to define the SRs are given by the following,

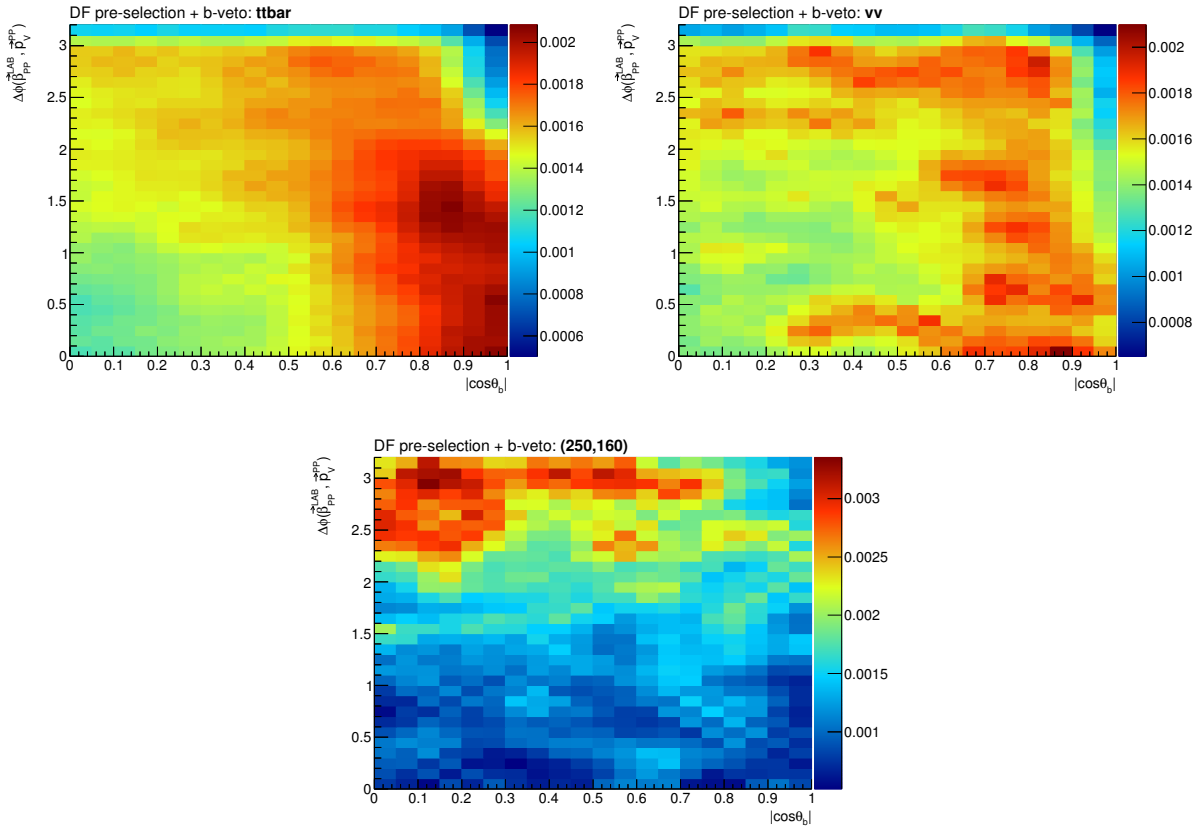

Figure 8.18: Two-dimensional relationship between $\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP})$ and $\cos\theta_b$ for SM $t\bar{t}$ (***upper left***), SM diboson processes (***upper right***), and the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ decay with $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (250, 160)$ GeV (***bottom***). The selection applied to the events populating these distributions follows the DF Preselection outlined in Table 8.3 with the additional requirement that there be no b-tagged jets in the event. The distributions are normalized, so only the shapes are relevant for comparison.

Table 8.4: Basis of kinematic observables for the definition of analysis regions in the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ search.

R_{p_T}
$1/\gamma_P^{PP}$
$\cos\theta_b$
$\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP})$
E_V^P
b-tagged jet multiplicity

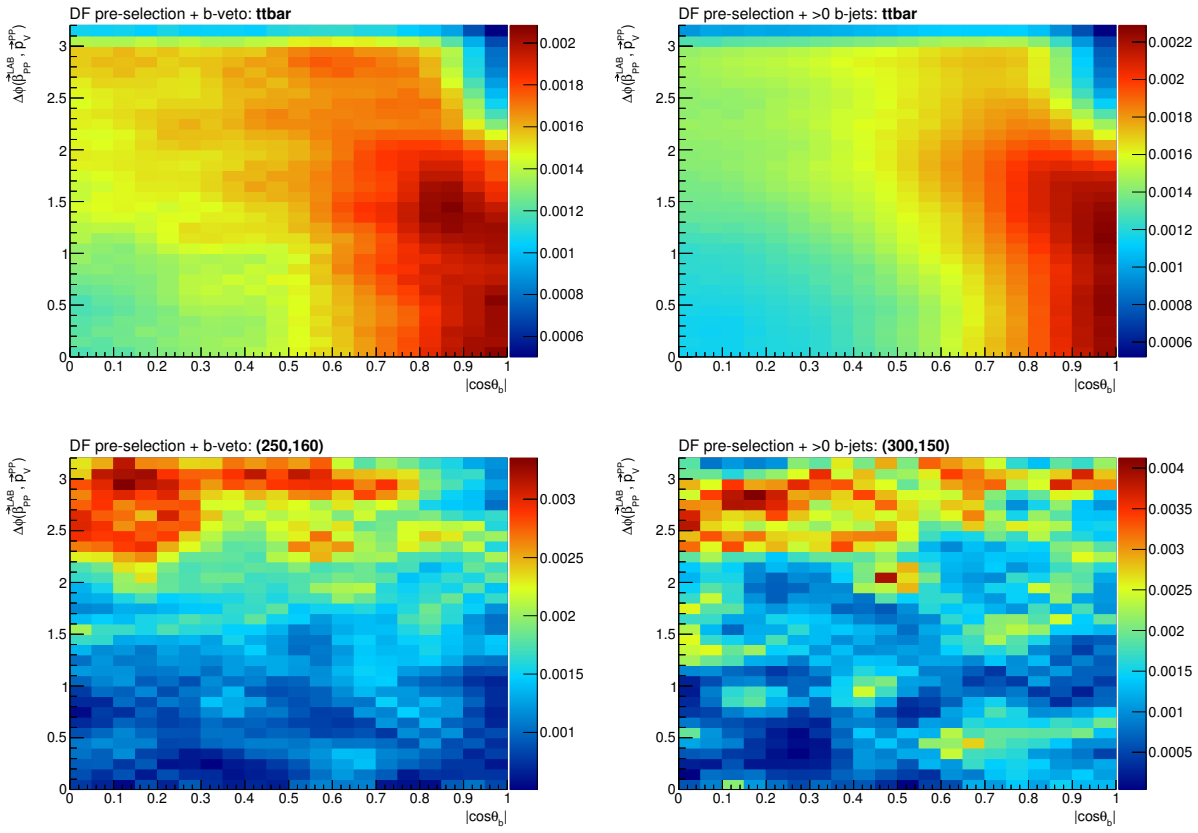

Figure 8.19: Two-dimensional relationship between $\Delta\phi(\vec{\beta}_{PP}^{\mathrm{LAB}}, \vec{p}_V^{PP})$ and $\cos\theta_b$ for SM $t\bar{t}$ (***upper***), the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ decay with $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (250, 160)\,\mathrm{GeV}$ (***lower left***), and the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ decay with $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (350, 150)\,\mathrm{GeV}$ (***lower right***). The selection applied to the events populating these distributions are required to satisfy the DF Preselection outlined in Table 8.3, with those on the left having an additional requirement that there be no b-tagged jets in the events and those on the right requiring that there be at least one b-tagged jet in the event. The distributions are normalized, so only the shapes are relevant for comparison.

Optimization Procedure

In order to choose the precise selection thresholds on the quantities listed in Table 8.4, an optimization procedure is followed. The procedure used in the search for the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ decay is a brute force method, and is as follows. A multi-dimensional scan is performed over the subset of those quantities in listed in Table 8.4, which are listed in Table 8.5. The scan is done using objects (leptons and jets) that satisfy the signal level criteria, detailed in Tables 8.1 and 8.2.

Table 8.5: Kinematic observables used in the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ grid search optimization procedure.

R_{p_T}
$1/\gamma_P^{PP}$
m
b

In Table 8.5, the quantities 'm' and 'b' are related to the slope and y-intercept of the line in the two-dimensional space $(\cos\theta_b, \Delta\phi(\vec{\beta}_{PP}^{\,\mathrm{LAB}}, \vec{p}_V^{\,PP}))$, as illustrated in Figure 8.20. For each of the quantities

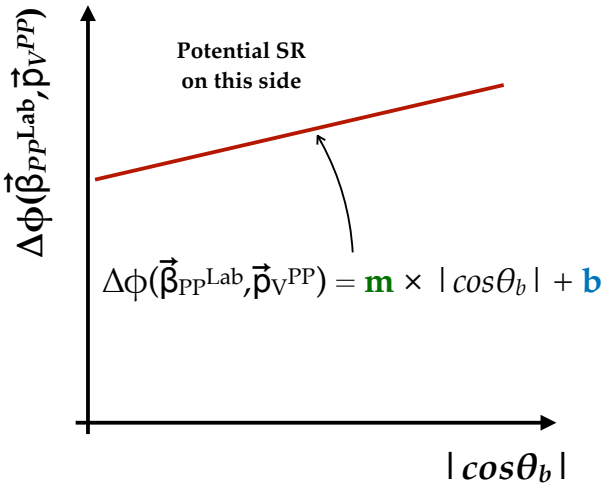

Figure 8.20: Illustration of the two-dimensional selection in the $(\cos\theta_b, \Delta\phi(\vec{\beta}_{PP}^{\,\mathrm{LAB}}, \vec{p}_V^{\,PP}))$-plane.

in Table 8.5, a fine-grained set of threshold steps is defined and a multi-dimensional scan over all allowed step values for all four quantities is performed. At each step, a region is defined with the

requirements as follows,

$$\begin{aligned} R_{p_T} &> \text{step}_{\{R_{p_T}\}} \\ 1/\gamma_P^{PP} &> \text{step}_{\{1/\gamma_P^{PP}\}} \\ \text{m} &= \text{step}_{\{\text{m}\}} \\ \text{b} &= \text{step}_{\{\text{b}\}}, \end{aligned}$$

and a metric for evaluating the sensitivity to the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ process of the given region is calculated. The metric used in the search for the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ process is an approximation to the Gaussian significance associated with the discovery p_0-value. The larger this quantity is, the larger the expected sensitivity to the given $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ decay. The significance is computed using the RooStats library and takes as input the predicted number of signal events, the predicted number of background events, and a level of uncertainty on the background event yield.[4] Since the actual background uncertainty is unknown at the point of designing an analysis' search regions, it is usually taken to be a conservative value. In the search described here, the relative background uncertainty, Δb_{opt}, was taken to be,

$$\Delta b_{\text{opt}} = \sqrt{(0.30)^2 + (\delta s)^2},$$

where δs is the relative statistical uncertainty on the MC-simulated signal sample. The quantity δs is added in order to act as a penalty term for cases in which the MC statistics in the signal sample become too small to give reliable estimates of the significance.

The optimisation scan is performed using only the dominant expected backgrounds, SM $t\bar{t}$ and diboson production, as input. For the signal samples used to provide the estimated signal yield input, the scan is performed three times for each of the three *benchmark* signal samples with $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (250, 160)$, $(300, 150)$, and $(300, 120)$ GeV, in order to assess how well each of the regions in the step procedure perform across the three-body phase space with differing values for $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0)$.

At the end of the scan, there are $\mathcal{O}(10^4)$ different regions checked. The best region is chosen based on its maximization of the significance metric as well as maintaining that there be at most 30% statistical uncertainty on the signal and background MC samples once the selection for the associated region is applied. This choice is made taking into account the three benchmark $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ signal scenarios, with the hope of obtaining a continuous sensitivity throughout the three-body regime for scenarios with different values of $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0)$.

[4] The precise significance metric used is RooStats::NumberCountingUtils::BinomialObsZ

Table 8.6: Final values obtained after performing the brute-force scan over the parameters listed in Table 8.5.

Quantity	Final Value Obtained from Scan
R_{p_T}	0.7
$1/\gamma_P^{PP}$	0.7
m	0.9
b	1.6

Final Signal Region Definitions

On top of the selections defined as a result of the brute-force scan, listed in Table 8.6, an E_V^P requirement is added. As the kinematic endpoint of the E_V^P variable is sensitive to the (squared) mass differences between the $\tilde{t}_1$ and $\tilde{\chi}_1^0$ (c.f. Equation 8.5), it is expected that the specific cut value on this quantity will be lower for the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ decay scenarios with lower values of $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0)$. The specific values of E_V^P are determined by running a similar optimisation scan as detailed in the previous section, but with only the E_V^P quantity being scanned over and with all other cuts fixed to those detailed in Table 8.6. This is done twice, once for the benchmark signal point with $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (250, 160)\,\text{GeV}$ and a second time with $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (350, 120)\,\text{GeV}$. The former (latter) defines the E_V^P cut used for the SR targeting the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ decay scenarios with $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0)$ nearer to m_W (m_{top}).

The final SR definitions used in for the search for the three-body decay of the $\tilde{t}_1$ are detailed in Table 8.7. There are four SRs defined in total, two of which target the $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) \to m_W$ region and two target the $\Delta m(\tilde{t}_1, \tilde{\chi}_1^0) = m_{\text{top}}$ region. The former are defined by a b-tagged jet veto, and the latter are inclusive of events with b-tagged jets. Within each of these two categories, the SRs are further broken down by the dilepton flavor: whether the two leptons in the events have the same-flavor (SF) or different-flavor (DF). In order to reduce backgrounds in which the two leptons arise from the decays of Z-bosons, a Z-veto is applied to the SF SRs by requiring that the dilepton invariant mass be outside a 20 GeVwindow centered on the Z-boson mass.

The total expected background yield and composition in terms of specific SM processes is given in Table 8.8 for each of the SRs defined in Table 8.7. It can be seen, as assumed in the text so far, that SM $t\bar{t}$ and diboson production processes are expected to be the main SM processes contaminating the SRs.

Table 8.7: Final signal region definitions for the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ search. Selections are made after the pre-selection requirements defined in Table 8.3. The quantities 'm' and 'b' are illustrated in Figure 8.20 and refer to the two-dimensional selection: $\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP}) > \text{m} \times |\cos\theta_b| + \text{b}$. The lepton p_T requirements are chosen so as to be on the dilepton trigger efficiency plateau.

Signal Region Definitions for the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ Search				
	Common Selection			
Lead lepton p_T[GeV]	> 25			
Sub-lead lepton p_T[GeV]	> 20			
R_{p_T}	> 0.7			
$1/\gamma_P^{PP}$	> 0.7			
m	0.9			
b	1.6			
	SRw-SF	**SRw-DF**	**SRt-SF**	**SRt-DF**
Dilepton invariant mass, $m_{\ell\ell}$ [GeV]	$\lvert m_{\ell\ell} - 91.2\rvert > 20$	no req.	$\lvert m_{\ell\ell} - 91.2\rvert > 20$	no req.
b-tagged jet multiplicity	Exactly 0		> 0	
E_V^P	> 95		> 110	

Table 8.8: Expected SM background processes in the SRs for the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ search, defined in Table 8.7. The bottom two rows show the predicted yields for two of the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ benchmark scenarios. The errors on the quoted yields are due to statistical and systematic uncertainties.

Process	**SRw-SF**	**SRw-DF**	**SRt-SF**	**SRt-DF**
Total Expected SM	10.94 ± 3.60	8.74 ± 3.45	3.00 ± 1.59	4.62 ± 2.38
$t\bar{t}$	4.32 ± 2.42	4.73 ± 2.82	2.41 ± 1.54	3.38 ± 2.31
Single-top	0.31 ± 0.19	0.23 ± 0.07	0.12 ± 0.015	0.15 ± 0.09
Diboson	3.76 ± 1.71	3.03 ± 1.33	0.16 ± 0.05	0.00 ± 0.00
Z+jets	1.52 ± 0.79	0.05 ± 0.01	0.10 ± 0.03	0.00 ± 0.00
$t\bar{t} + V$	0.04 ± 0.01	0.11 ± 0.03	0.21 ± 0.06	0.24 ± 0.09
Fakes	0.99 ± 0.52	0.59 ± 0.29	0.00 ± 0.00	0.82 ± 0.23
$\tilde{t}_1 \to bW\tilde{\chi}_1^0$, $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (250, 160)$ GeV	15.06 ± 1.78	18.93 ± 1.86	3.34 ± 0.84	4.06 ± 0.97
$\tilde{t}_1 \to bW\tilde{\chi}_1^0$, $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (300, 150)$ GeV	4.47 ± 1.09	7.03 ± 1.45	11.43 ± 2.00	12.86 ± 1.97

8.5 Estimation of the Standard Model Backgrounds

In this section, we describe the methods used for the estimation of the SM background contamination in the SRs. By Table 8.8, it can be seen that the expected SM background contamination to the SRs in the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ search is primarily composed of events from the $t\bar{t}$ and diboson processes. Given that these are the dominant backgrounds for the analysis, their estimation is performed using the control region method, described in Section 7.1.3. That is, dedicated CRs and VRs are defined for each of the two processes in order to provide a normalisation correction for their MC prediction in the SRs. All other background processes, being subdominant, have their predicted contribution to the SR background taken directly from the MC simulation. The estimation of the contribution of sources leading to fake and non-prompt leptons is performed using the Matrix Method, described in Section 7.2.3.

Sections 8.5.1 and 8.5.2 describe the background estimate for the $t\bar{t}$ and diboson processes, respectively.

The CR and VR definitions for both the $t\bar{t}$ and diboson background are based on the SR definitions given in the previous section. The CRs are defined primarily by inverting the two-dimensional selection made in the $(\cos\theta_b, \Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP}))$-plane, and maintaining similar selections as in the SRs for the other variables. The VRs, on the other hand, are defined by inverting the selections on the non-angular variables relative to those made in the SRs.

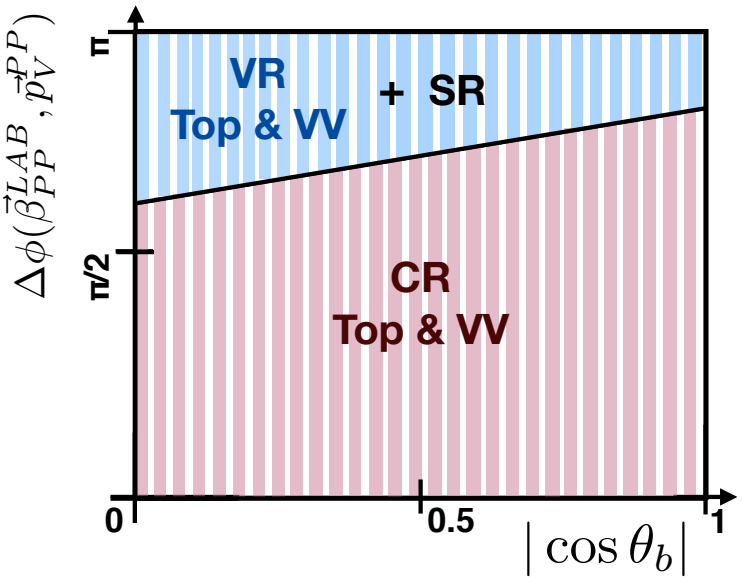

Figure 8.21: Illustration of the CR and VR strategy used in the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ search. The defining characteristic for the definition of these regions is based on the region in the $(\cos\theta_b, \Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP}))$-plane that they select. The CR inverts the requirements on these quantities relative to the SRs, while the VR has the same requirements as in the SRs but inverts selections made on the other observables.

8.5.1 Top-quark pair production

The CRs and VRs designed to derive and validate the semi-data-driven normalisation correction factor for the $t\bar{t}$ background process are called CR-Top and VR-Top, respectively, and are defined in Table 8.9. The strategy for the CR and VR selections in the $(\cos\theta_b, \Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP}))$ plane are described in the previous section. Several of the selections on the kinematic quantities relative to those in the SRs (c.f. Table 8.7) are relaxed. In both CR-Top and VR-Top, the E_V^P requirement is relaxed to $E_V^P > 80\,\text{GeV}$ and the requirement on the $1/\gamma_P^{PP}$ quantity is removed. In VR-Top, the R_{p_T} requirement is inverted relative to that used in the SRs. Given that the $t\bar{t}$ background is flavor symmetric, only different-flavor events are allowed to populate CR-Top and VR-Top, in order to avoid contamination from Z-boson processes. As a result, no additional requirement on $m_{\ell\ell}$ is

made in these regions. For increased purity, CR-Top requires that there be at least one b-tagged jet, while VR-Top applies a veto in order to be orthogonal to CR-Top.

VR-Top is defined to have zero b-tagged jets, while CR-Top requires at least one. In dedicated studies, it has been verified that the $t\bar{t}$ normalisation correction derived in the b-jet rich region CR-Top is well extrapolated to separate validation regions, and is rather independent of the b-tagged jet multiplicity. This gives confidence that VR-Top can be used as an appropriate check on the $t\bar{t}$ normalisation correct factor and that it's extrapolation to the SRs, which have differing requirements on the b-tagged jet multiplicity, is reasonable.

Distributions of several key observables in CR-Top are shown in Figures 8.22-8.23.

Table 8.9: Definitions of the CR and VR for the $t\bar{t}$ background process for the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ search.

	Regions	
Variable	**CR-Top**	**VR-Top**
Dilepton Flavor	DF	DF
$m_{\ell\ell}$ [GeV]	no req.	no req.
Lead lepton p_T [GeV]	> 25	> 25
Sub-lead lepton p_T [GeV]	> 20	> 20
b-tagged jet multiplicity	> 0	Exactly 0
E_V^P [GeV]	> 80	> 80
R_{p_T}	> 0.7	< 0.7
$1/\gamma_P^{PP}$	no req.	no req.
$(\cos\theta_b, \Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP}))$	$\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP}) < 0.9 \times \lvert\cos\theta_b\rvert + 1.6$	$\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP}) > 0.9 \times \lvert\cos\theta_b\rvert + 1.6$

Kinematic Distributions in CR-Top

8.5.2 Diboson Production

The SM diboson processes are composed of WW, $ZW + WZ$, and ZZ production. The dominant process for the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ SRs is WW, as discussed in the text. In order to constrain the WW component, in addition to those components with a Z boson, the diboson CRs are split into two, one targeting the different-flavor enriched component of the diboson background (predominantly WW) and one in which the same-flavor component ($ZW + WZ$ and ZZ) is enriched.

The same-flavor and different-flavor diboson CRs (VRs), CR-VV-SF and CR-VV-DF (VR-VV-SF and VR-VV-DF), respectively, are defined in Table 8.10. All of the regions apply a b-tagged jet veto. Although there are SRs that are inclusive of b-tagged jets (SRt-SF and SRt-DF), the diboson background is negligible in them and so the normalisation correction factors derived in the diboson CRs can be extrapolated with confidence into the SRs with a b-tagged jet veto (SRw-SF and SRw-DF). There is no difference between the same-flavor and different-flavor CRs, apart from the dilepton flavor requirements and Z-veto. The diboson VRs have the same, or inclusive, R_{p_T}

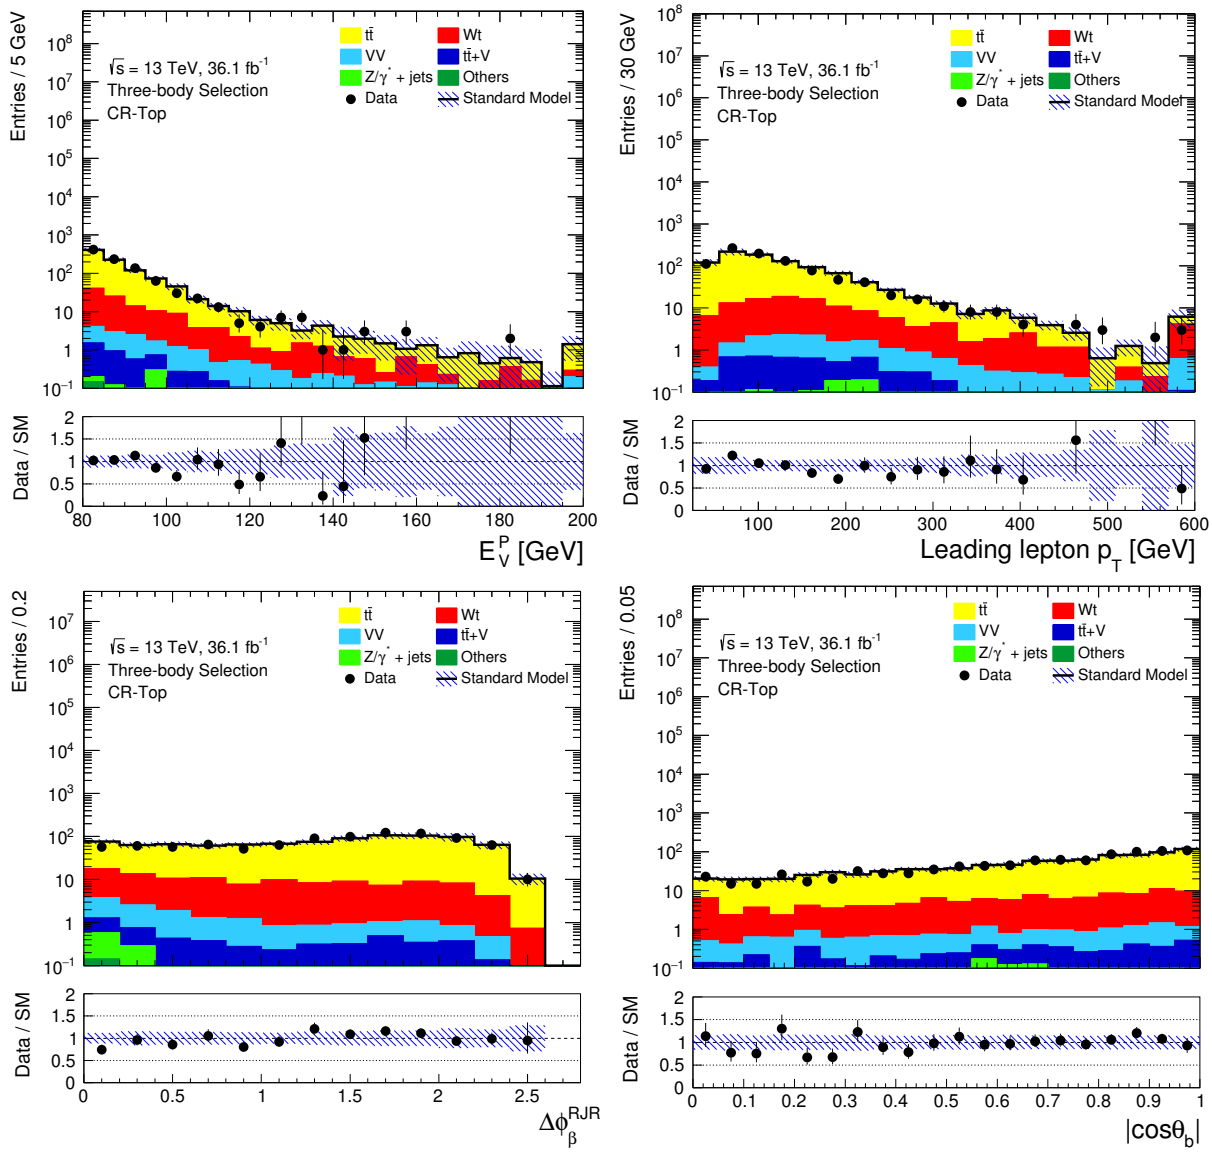

Figure 8.22: Distributions of E_V^P (***upper left***), leading lepton p_T (***upper right***), $\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP})$ (***lower left***), and $|\cos\theta_b|$ (***lower right***) in the $t\bar{t}$CR, CR-Top. The error on the SM processes includes statistical and systematic uncertainties. The post-fit normalization correction factors for the $t\bar{t}$and diboson processes have been applied.

and $1/\gamma_P^{PP}$ selections as the CRs but have orthogonal E_V^P requirements that move them closer the SR selections.

Distributions of several key observables in CR-VV-SF (CR-VV-DF) are shown in Figures 8.24-8.25 (Figures 8.26-8.27).

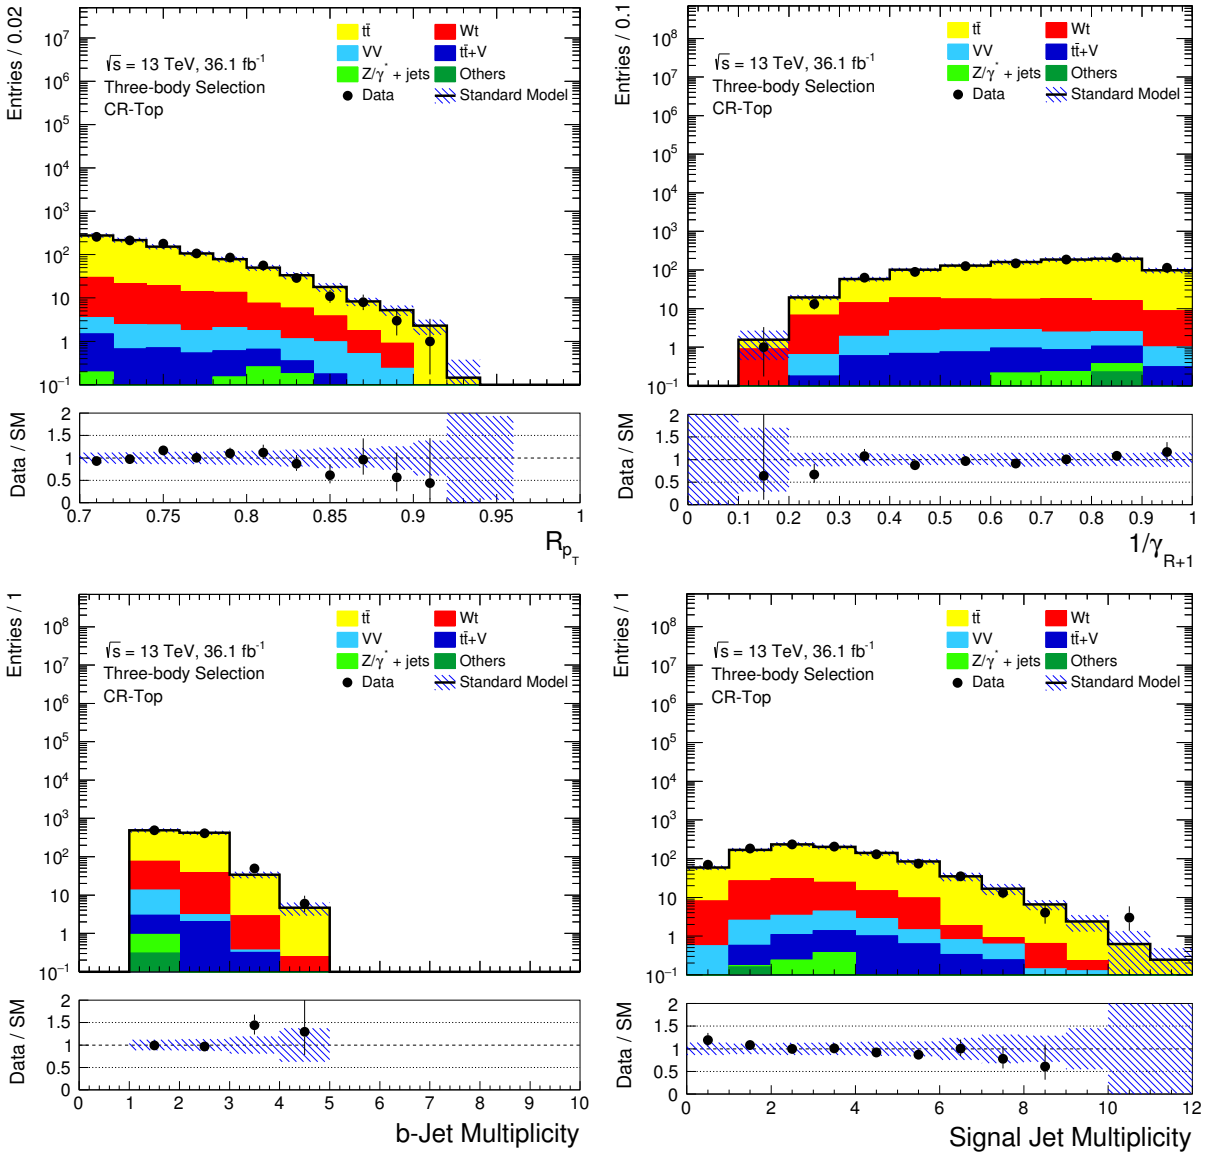

Figure 8.23: Distributions of R_{p_T} (***upper left***), leading lepton $1/\gamma_P^{PP}$ (***upper right***), b-tagged jet multiplicity (***lower left***), and non-b-tagged jet multiplicity(***lower right***) in the $t\bar{t}$CR, CR-Top. The error on the SM processes includes statistical and systematic uncertainties. The post-fit normalization correction factors for the $t\bar{t}$and diboson processes have been applied.

Table 8.10: Definitions of the CR and VR for the diboson background processes for the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ search.

	Regions			
Variable	**CR-VV-DF**	**CR-VV-SF**	**VR-VV-DF**	**VR-VV-SF**
Dilepton Flavor	DF	SF	DF	SF
$m_{\ell\ell}$ [GeV]	no req.	$\lvert m_{\ell\ell} - 91.2\rvert > 10$	no req.	$\lvert m_{\ell\ell} - 91.2\rvert > 10$
Lead lepton p_T [GeV]	> 25	> 25	> 25	> 25
Sub-lead lepton p_T [GeV]	> 20	> 20	> 20	> 20
b-tagged jet multiplicity	Exactly 0	Exactly 0	Exactly 0	Exactly 0
E_V^P [GeV]	> 50	> 70	$\in (50, 95)$	$\in (60, 95)$
R_{p_T}	< 0.5	< 0.5	< 0.7	< 0.4
$1/\gamma_P^{PP}$	> 0.7	> 0.7	> 0.7	> 0.7
$(\cos\theta_b, \Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP}))$	$\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP}) < 0.9 \times \lvert\cos\theta_b\rvert + 1.6$		$\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP}) > 0.9 \times \lvert\cos\theta_b\rvert + 1.6$	

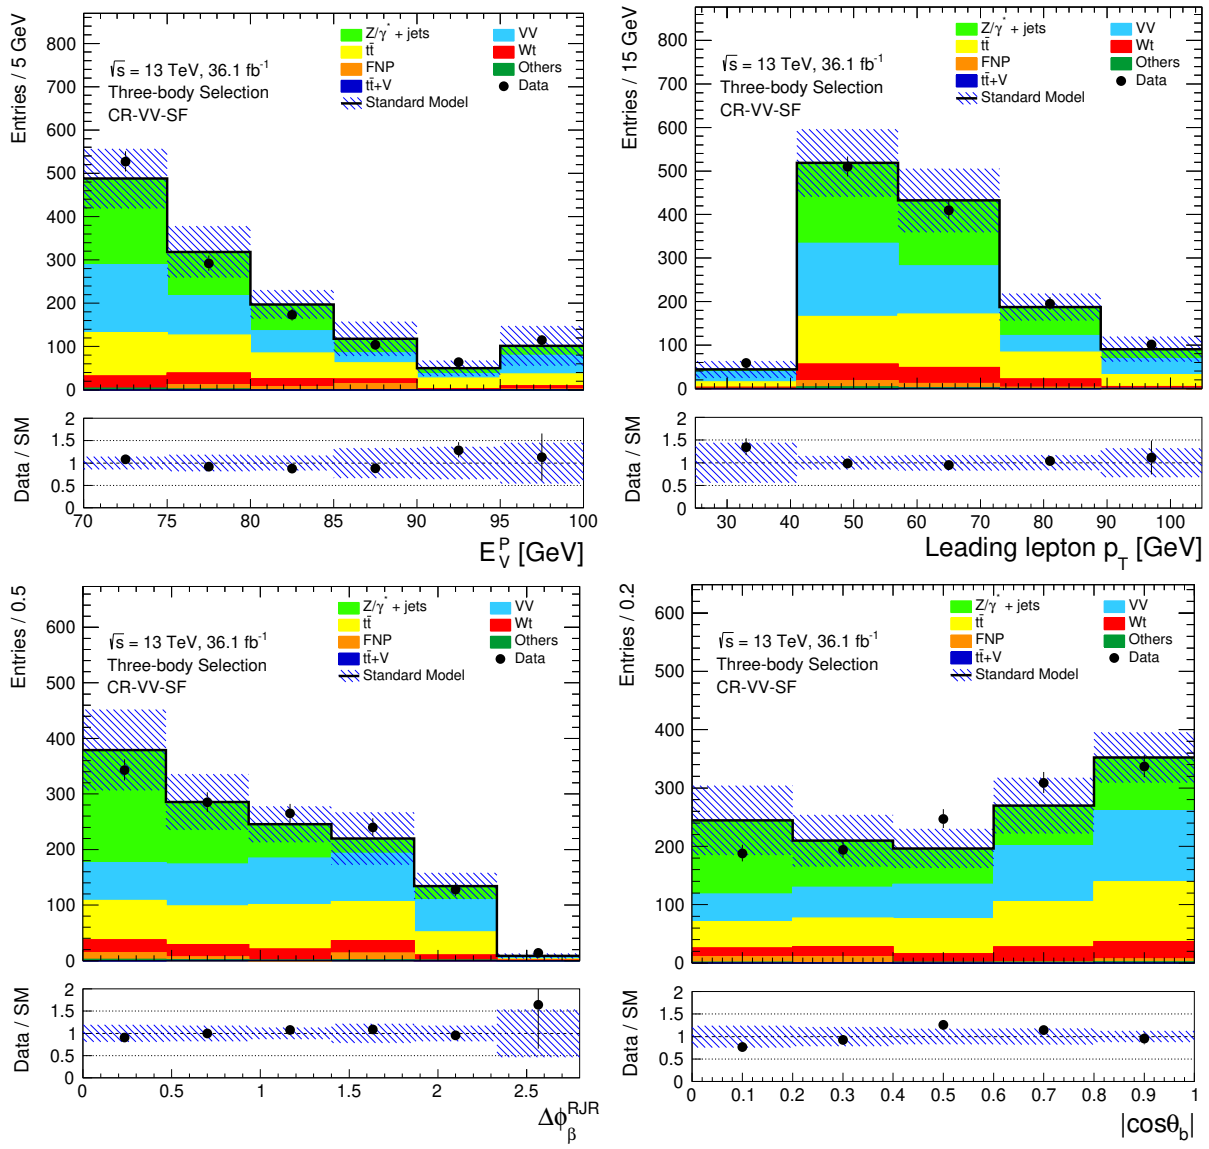

Figure 8.24: Distributions of E_V^P (***upper left***), leading lepton p_T (***upper right***), $\Delta\phi(\vec{\beta}_{PP}^{\text{LAB}}, \vec{p}_V^{PP})$ (***lower left***), and $|\cos\theta_b|$ (***lower right***) in the same-flavor diboson CR, CR-VV-SF. The error on the SM processes includes statistical and systematic uncertainties. The post-fit normalization correction factors for the $t\bar{t}$and diboson processes have been applied.

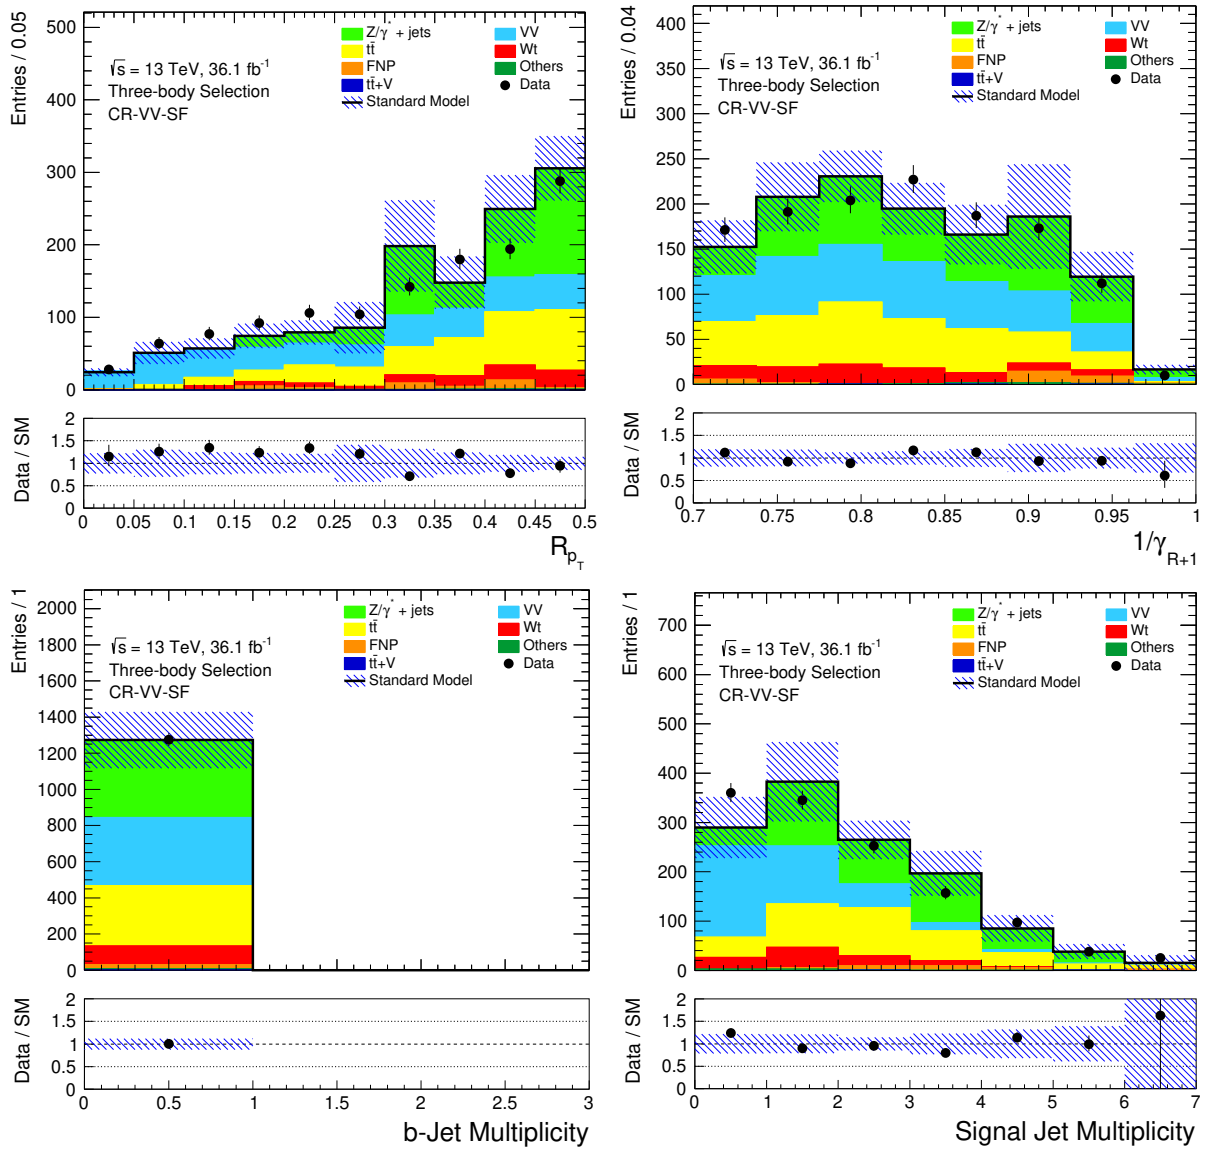

Figure 8.25: Distributions of R_{p_T} (***upper left***), leading lepton $1/\gamma_P^{PP}$ (***upper right***), b-tagged jet multiplicity (***lower left***), and non-b-tagged jet multiplicity(***lower right***) in the same-flavor diboson CR, CR-VV-SF. The error on the SM processes includes statistical and systematic uncertainties. The post-fit normalization correction factors for the $t\bar{t}$and diboson processes have been applied.

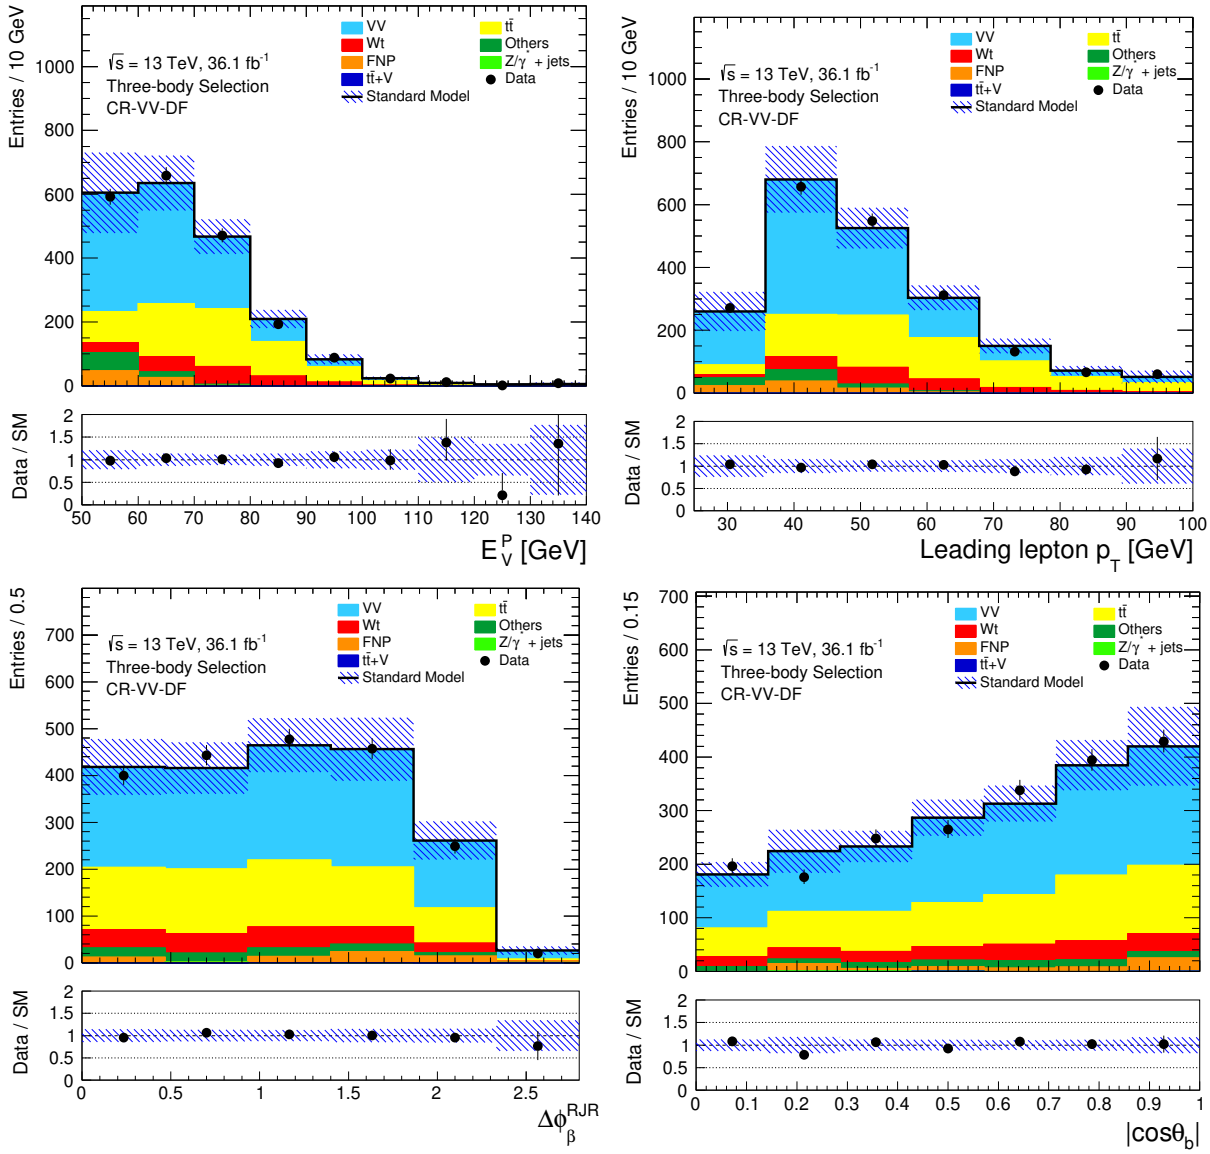

Figure 8.26: Distributions of E_V^P (***upper left***), leading lepton p_T (***upper right***), $\Delta\phi(\vec{\beta}_{PP}^{\mathrm{LAB}}, \vec{p}_V^{PP})$ (***lower left***), and $|\cos\theta_b|$ (***lower right***) in the different-flavor diboson CR, CR-VV-DF. The error on the SM processes includes statistical and systematic uncertainties. The post-fit normalization correction factors for the $t\bar{t}$and diboson processes have been applied.

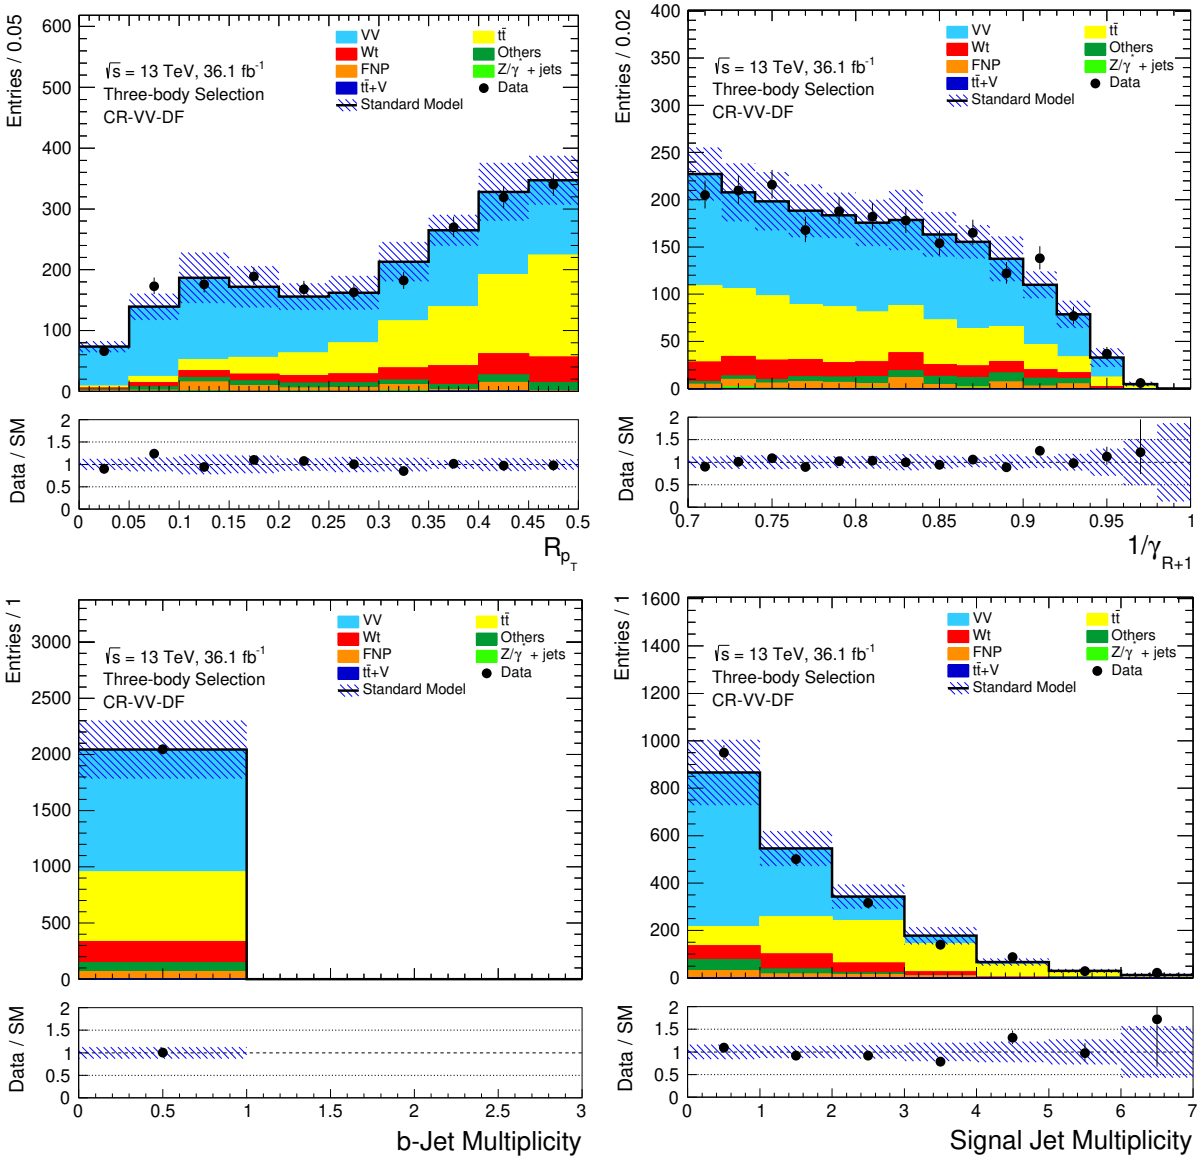

Figure 8.27: Distributions of R_{p_T} (***upper left***), leading lepton $1/\gamma_P^{PP}$ (***upper right***), b-tagged jet multiplicity (***lower left***), and non-b-tagged jet multiplicity(***lower right***) in the different-flavor diboson CR, CR-VV-DF. The error on the SM processes includes statistical and systematic uncertainties. The post-fit normalization correction factors for the $t\bar{t}$and diboson processes have been applied.

8.5.3 Background-only Fit

In order to assess the impact of the CRs on the background estimation in the SRs, a so-called 'background-only' fit is performed. A background-only fit is profile-likelihood fit, as described Section 7.4.2, in which the only regions contributing to the likelihood (c.f Equation 7.18) are the analysis' CRs. The result of running a background-only fit to data in the CRs is shown in Table 8.11, which shows the MC predicted yields for the background processes both before and after the background-only fit is performed, as well as the observed data counts, in each of the CRs and VRs in the analysis. The post-fit yields in the CRs are expected to agree with the observed data counts, since the latter are used as constraints in the fit model and there are as many freely-floating parameters in the fit (3 μ factors) as there are CRs; therefore, the fit has enough freedom to cover any discrepancy between the observed data and pre-fit MC prediction of the background processes. The agreement observed between the post-fit MC and the observed data in the VRs shows that the extrapolation, at least in terms of the corrected MC's normalisation, is performing well.

The normalisation correction factors for the $t\bar{t}$ and diboson processes obtained in the background-only fit are listed in Table 8.12. They are generally consistent with one.

Table 8.11: Yields in the $t\bar{t}$ and diboson CRs and VRs for the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ search for the main background processes contributing to the analysis. The lower-portion of the table are the yields before the background-only fit to data in the CRs, without the normalisation corrections applied. The upper-portion of the table are those taken after the background-only fit to data. The errors on the quoted numbers are due to the statistical and experimental systematic uncertainties.

Process	**CR-Top**	**CR-VV-DF**	**CR-VV-SF**	**VR-Top**	**VR-VV-DF**	**VR-VV-SF**
Observed Data	951	2046	1275	1197	1896	783
Post-fit Total SM	951.00 ± 30.84	2046.05 ± 45.23	1275.17 ± 35.68	1231.78 ± 86.59	2013.57 ± 116.49	780.44 ± 117.32
Post-fit $t\bar{t}$	833.03 ± 32.85	619.74 ± 111.40	333.62 ± 60.91	733.87 ± 64.91	754.94 ± 78.17	127.14 ± 22.09
Post-fit Diboson (DF)	11.51 ± 2.43	1093.28 ± 125.83	0.00 ± 0.00	331.13 ± 82.57	886.53 ± 168.16	0.00 ± 0.00
Post-fit Diboson (SF)	0.00 ± 0.00	0.00 ± 0.00	378.94 ± 124.32	0.00 ± 0.00	0.00 ± 0.00	380.00 ± 141.58
Post-fit Single-top	101.10 ± 9.73	186.47 ± 27.99	103.47 ± 17.43	111.52 ± 14.49	151.88 ± 14.37	36.40 ± 5.62
Post-fit $t\bar{t}+V$	4.35 ± 0.42	0.39 ± 0.07	0.36 ± 0.07	1.27 ± 0.22	0.42 ± 0.13	0.05 ± 0.02
Post-fit Z+jets	0.70 ± 0.22	$1.83^{+2.55}_{-1.83}$	428.58 ± 92.55	$0.47^{+0.85}_{-0.47}$	$0.39^{+0.71}_{-0.39}$	191.37 ± 78.38
Post-fit Single-higgs	0.31 ± 0.13	78.95 ± 9.17	6.23 ± 1.06	$0.44^{+0.52}_{-0.44}$	54.98 ± 4.34	9.40 ± 1.11
Post-fit Fakes	0.00 ± 0.00	65.37 ± 2.22	23.96 ± 1.25	53.09 ± 1.92	164.42 ± 5.68	36.09 ± 3.04
Total SM	905.14 ± 16.54	1988.38 ± 110.43	1248.21 ± 123.42	1184.39 ± 71.23	1952.92 ± 61.72	764.65 ± 99.06
$t\bar{t}$	787.43 ± 11.29	585.87 ± 102.00	315.39 ± 55.87	693.71 ± 53.51	713.64 ± 66.57	120.19 ± 20.17
Diboson (DF)	11.25 ± 1.62	1069.46 ± 12.45	0.00 ± 0.00	323.89 ± 50.27	867.18 ± 77.75	0.00 ± 0.00
Diboson (SF)	0.00 ± 0.00	0.00 ± 0.00	370.13 ± 15.48	0.00 ± 0.00	0.00 ± 0.00	371.16 ± 34.38
Single-top	101.10 ± 9.80	186.49 ± 28.25	103.49 ± 17.59	111.52 ± 14.60	151.88 ± 14.46	36.41 ± 5.66
$t\bar{t}+V$	4.35 ± 0.42	0.39 ± 0.07	0.36 ± 0.07	1.27 ± 0.23	0.42 ± 0.13	0.05 ± 0.02
Z+jets	0.70 ± 0.23	$1.83^{+2.57}_{-1.83}$	428.65 ± 93.15	$0.48^{+0.86}_{-0.48}$	$0.40^{+0.72}_{-0.40}$	191.36 ± 78.78
Single-higgs	0.31 ± 0.13	78.96 ± 9.26	6.23 ± 1.07	$0.44^{+0.53}_{-0.44}$	54.98 ± 4.37	9.40 ± 1.12
Fakes	0.00 ± 0.00	65.37 ± 2.23	23.96 ± 1.26	53.09 ± 1.92	164.42 ± 5.70	36.09 ± 3.04

Table 8.12: Normalisation correction factors for the $t\bar{t}$ ($\mu_{t\bar{t}}$), same-flavor diboson ($\mu_{\text{VV-SF}}$), and different-flavor diboson ($\mu_{\text{VV-DF}}$) processes derived from the background-only fit to the CRs. The errors on the quoted numbers are due to the statistical and experimental systematic uncertainties entering the fit.

$\mu_{t\bar{t}}$	1.06 ± 0.05
$\mu_{\text{VV-SF}}$	1.02 ± 0.35
$\mu_{\text{VV-DF}}$	1.02 ± 0.12

8.6 Results

A profile likelihood fit is performed, including all of the CRs and SRs of the analysis. In order to assess the compatibility of the observed data in the SRs with the predicted SM background, the observed data in the SRs is also provided in the fit as a constraint. The result of running this fit, with the added SRs and the observed data therein, is shown in Table 8.13. Even though we have added the additional data constraint from the SRs, the post-fit normalisation correction factors derived for the $t\bar{t}$ and diboson processes are consistent with those derived in the background-only fit configuration described in Section 8.5.3, due to the larger statistical power of the data counts observed in the CRs, as opposed to the low numbers observed in the SRs.

It can be seen in Table 8.13 that there is no significant deviation between the observed data counts in the SRs as compared to the post-fit prediction of the total SM background. The largest null-hypothesis p_0-value in these regions, giving the incompatibility of the observed data with the predicted SM background, is associated with the region SRt-SF and is $p_0 = 0.15$, still quite far from the $p_0 = 0.05$ critical value. Figure 8.28 shows a few kinematic distributions in the SRs with the post-fit MC prediction compared to the observed data. In these plots, the SR selection on the variable being plotted has been removed.

As there is no significant deviation between the prediction and observed data, hypothesis tests are performed in order to assess which regions of the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ parameter space can be excluded. As described in Section 7.4, we use the CL_s construction to determine whether a given point in the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane is excluded at 95% CL. Hypothesis tests are run over the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ signal grid, with all four SRs and three CRs included in the profile likelihood along with the observed data in all regions included. The results of the hypothesis tests are illustrated in Figure 8.29, showing the exclusion regions in the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ region of the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane.

Table 8.13: Observed and predicted yields in the SRs for the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ search.

	Regions			
Process	**SRw-SF**	**SRw-DF**	**SRt-SF**	**SRt-DF**
Observed Data	4	6	6	6
Post-fit Total SM	9.83 ± 3.57	7.84 ± 3.09	3.10 ± 1.42	4.40 ± 1.87
Post-fit $t\bar{t}$	4.18 ± 1.65	4.57 ± 2.18	2.54 ± 1.34	3.58 ± 1.77
Post-fit Diboson	3.43 ± 2.27	2.88 ± 1.44	0.16 ± 0.09	0.04 ± 0.03
SM	9.52 ± 3.41	7.53 ± 2.84	2.95 ± 1.34	4.20 ± 1.77
$t\bar{t}$	3.95 ± 1.55	4.31 ± 2.06	2.40 ± 1.28	3.38 ± 1.68
Diboson	3.35 ± 1.58	2.82 ± 1.22	0.16 ± 0.06	0.04 ± 0.03
Single-top	0.31 ± 0.29	0.23 ± 0.12	0.12 ± 0.05	0.14 ± 0.08
Z+jets	1.47 ± 1.47	0.05 ± 0.01	0.10 ± 0.03	0.00 ± 0.00
$t\bar{t} + V$	0.03 ± 0.01	0.06 ± 0.02	0.18 ± 0.05	0.24 ± 0.07
Fakes	0.42 ± 0.28	0.06 ± 0.06	0.00 ± 0.00	0.41 ± 0.09

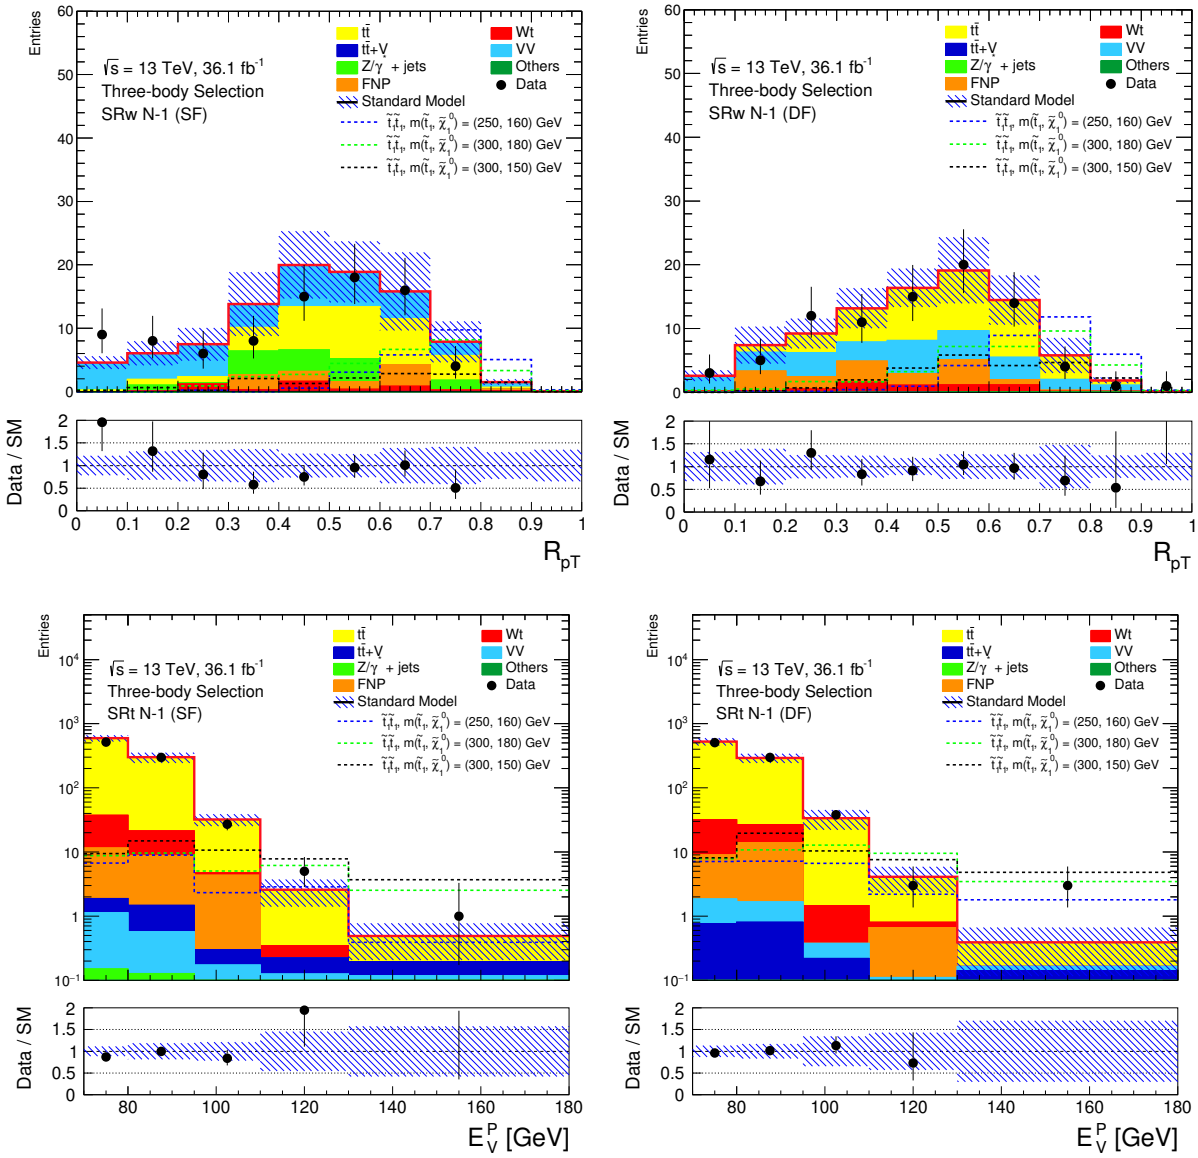

Figure 8.28: ***Top***: Distributions of R_{p_T} in SRw-SF (*left*) and SRw-DF (*right*), with the observed data. ***Bottom***: Distributions of E_V^P in SRt-SF (*left*) and SRt-DF (*right*), with the observed data. The selection on the variable being plotted, in the corresponding SR, has been removed from the selection applied to the events populating the histogram bins. The hatched bands indicate the statistical and systematic uncertainty on the background estimate. The dashed lines show the MC simulated signal processes for the three benchmark $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ signal samples at $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0}) = (250, 160)$, $(350, 180)$, and $(350, 150)$ GeV.

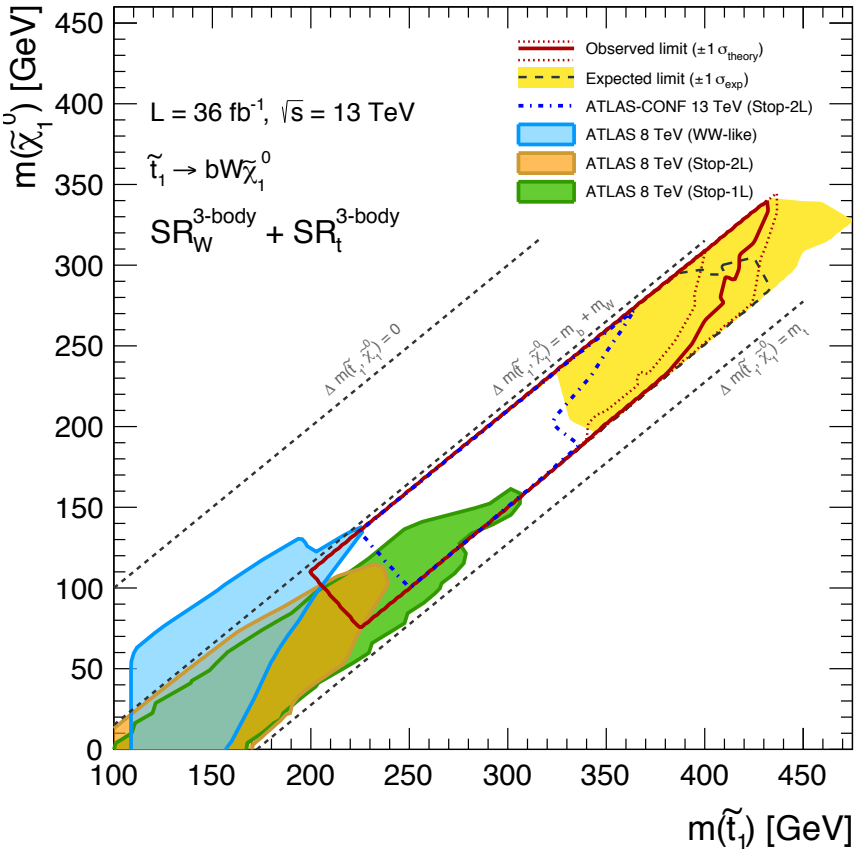

Figure 8.29: 95% CL exclusion contours for the 2015+2016 search for the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ process. Shown for comparison are the Run 1 8 TeV results for the same $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ scenario, both in the one ('Stop 1L') and two ('Stop 2L' and 'WW-like') lepton channels. The $\pm 1\sigma_{\text{theory}}$ lines on the observed limits correspond to varying the $\tilde{t}_1 \to bW\tilde{\chi}_1^0$ signal predicted cross-section values up and down within their theoretical uncertainty and re-running the hypothesis tests. The dashed blue exclusion contour is a previous iteration of the current analysis, but based on a reduced dataset ($3.2\,\text{fb}^{-1}$, based only on data collected in 2015, as opposed to $36\,\text{fb}^{-1}$) and is shown also for comparison purposes.

Chapter 9

The Search for the Non-resonant Production of Higgs Boson Pairs

It was miraculous... Anybody could do it; it required no brains at all. It merely required no character.

–Joseph Heller, *Catch-22*

...the ass will carry his load, but not a double load; ride not a free horse to death.

–Miguel de Cervantes, *Don Quixote*

In this chapter, the search for the non-resonant production of Higgs boson pairs (hh) will be discussed. As described in Section 1.4, the SM predicts non-resonant hh production but with an exceedingly small overall cross-section of 31.05 fb that is $\mathcal{O}(10^3)$ times smaller than that of the dominant production modes of single Higgs bosons [47, 196, 197, 198, 199, 200, 201, 202]:

$$\frac{\sigma_h}{\sigma_{hh}} = \frac{4.852 \times 10^3 \,\text{fb}}{31.05 \,\text{fb}} = 1.56 \times 10^3 \tag{9.1}$$

The need to complete our understanding of the 125 GeV Higgs boson, discovered in 2012, requires that meaningful samples of events containing Higgs pairs be obtained, as the study of these hh events are the only way to have direct measurement of and ability to constrain the Higgs self-coupling parameter λ responsible for the structure of the SM vacuum and the onset of EWSB. At the time of writing, at the end of LHC Run 2, the current pp collision datasets available to the general purpose LHC experiments are not large enough to observe the hh process as predicted in the SM, as made obvious by the relation in Equation 9.1. With 139 fb^{-1} of data collected during Run 2 of the LHC by the ATLAS detector, only $\mathcal{O}(1000)$ hh events will have occurred (using the

SM-predicted cross-section). Taking into account the non-100% acceptance of the detector, the dominant branching fractions of the Higgs decays, and the experimentally viable channels for hh decays to be observed, there is no possibility to observe a clearly distinguished hh signal. With this said, however, there are clear motivations for developing a robust hh search program with the current pp collision data. The first being that, given such a small predicted cross-section, *any* observed signal consistent with hh production in the present data will be an unambiguous sign of new physics that modifies the Higgs sector in such a way as to lead to enhanced hh production. The second motivation lies in the fact that, considering the LHC timeline (c.f. Figure 3.8), the forecasted amount of data required to have a 5σ evidence of hh production, as predicted by the SM, may yet be out of reach even at the end of the HL-LHC era. The goal, then, is to maximize the sensitivity to hh production as soon as possible so that the analyses of the future pp collision data are as effective as possible at accumulating hh candidate events so that the nature of EWSB may have a chance of being directly scrutinized at the LHC.

With this understanding, the ATLAS searches for hh production in both Run 1 and for most of Run 2 have focused on the channels sensitive to the hh decays that have the clearest experimental signatures. Given both the high branching fraction for the SM $h \to bb$ decay, and the distinctive signature of a resonance structure in the bb-system invariant mass at $m_h = 125\,\text{GeV}$, the most relevant channels to search for hh production will generally have one of the Higgs bosons decaying via the $h \to bb$ channel, with the second Higgs typically decaying via the dominant decay modes typical for the study of single Higgs production: $h \to bb$, $h \to \tau\tau$, $h \to WW$, and $h \to \gamma\gamma$. The relative branching fractions of the hh decays are indicated in Figure 9.1.

At the time of writing, the searches for SM-like non-resonant hh production have observed no discrepancies between the SM backgrounds and the observed data. That is, there has been no statistically significant observation of events consistent with hh production. Given this, analyses searching for SM-like non-resonant hh production typically quote 95% CL upper-limits (UL) on the hh production cross-section. Such UL, as discussed in Section 7.4, give a measure of how sensitive these analyses are to observing the process as predicted by the SM. The UL are therefore typically quoted as a ratio of the cross-section UL over that of the cross-section predicted by the SM: $\sigma^{\text{UL}}/\sigma^{\text{SM}}$. The most sensitive analyses in ATLAS, considering only those based on the analysis of the partial Run 2 data collected in 2015–2016, are the $hh \to bbbb$ ('$4b$') [203], $hh \to bb\tau\tau$ ('$bb\tau\tau$') [204], and $hh \to bb\gamma\gamma$ ('$bb\gamma\gamma$') [205] channels. Additional searches, also based on the 2015–2016 data collected by ATLAS, have also been performed in the $hh \to WW^*WW^*$ ('$4W$') [206], $hh \to WW^*\gamma\gamma$ ('$WW^*\gamma\gamma$') [207], and in the single-lepton channel of the $hh \to bbWW^*$ ('$bbWW^*$') [208] channels. The results of these analyses, and their combination, are presented in terms of $\sigma^{\text{UL}}/\sigma^{\text{SM}}$ in Figure 9.2. It can be seen that the $4b$, $bb\tau\tau$, and $bb\gamma\gamma$ analyses are driving the combined sensitivity to the hh production process, with the others having little impact (all having $\sigma^{\text{UL}}/\sigma^{\text{SM}} > 100$). The observed (expected) $\sigma^{\text{UL}}/\sigma^{\text{SM}}$ value based on the combination of all six analyses is 6.9 (10), using $\sigma_{hh}^{\text{SM}} = 33.5\,\text{fb}$ as the SM assumption for the hh production cross-section.

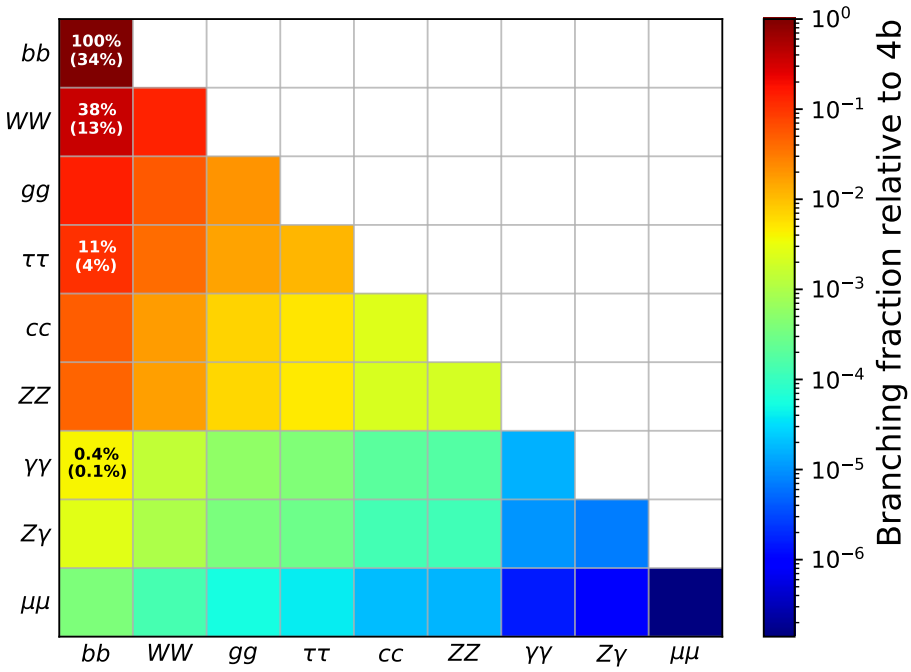

Figure 9.1: Branching factions for hh decays, with one Higgs decaying via the channel indicated on the x-axis and the other decaying via the channel indicated on the y-axis. The branching fractions indicated by the color are normalized to that of the $hh \to bbbb$ decay. Where there is text on the figure, the upper number indicates the branching fraction relative to the $hh \to bbbb$ decay and the lower number in parenthesis indicates the absolute branching fraction value (i.e. not normalised to $hh \to bbbb$).

From Figure 9.1, one sees that the $bbWW^*$ channel is the second-leading decay mode available for hh production. This is a rich hh channel as there are three decay modes available to the WW^* system: the fully hadronic, semi-leptonic, or dileptonic decay. The only experimentally viable signatures with hope of sensitivity to the hh process in the $bbWW^*$ decay channel will be those with at least a single lepton. The fully hadronic mode will be dominated by QCD backgrounds, making this a challenging search channel. The clean signature of a final state lepton will help remove contamination from QCD backgrounds. Given that the dilepton $h \to WW^*$ channel is one of the main channels in which the study of single Higgs production proceeds [209, 210, 211], one also expects, perhaps naively, the dilepton $bbWW^*$ channel to have been fully explored in the 2015–2016 dataset of LHC Run 2. However, only the single lepton channel had been studied in this dataset and, as can be seen in Figure 9.2, it is one of the least sensitive channels. In fact, the single lepton $bbWW^*$ channel is less sensitive than the $4W$ and $WW^*\gamma\gamma$ channels. This latter fact is counter-intuitive, especially when considering the $4W$ scenario, since one expects that the resonant bb system available to the single lepton $bbWW^*$ scenario to at least give this channel an upper-hand at being able to select hh candidate events. This lack of sensitivity in the single lepton $bbWW^*$ search channel, based on analysis of the 2015–2016 dataset, is likely due to the strategy

that this analysis used, in the sense that it has not yet been properly optimized to target the SM-like non-resonant hh production process. Additional complications arise in this channel that are experimental in nature, such as ambiguities in identifying the lepton when it overlaps with one of the reconstructed jets from the hadronically decaying W boson.

As a result of this discussion, an analysis searching for SM-like non-resonant hh production in the dilepton channel of the $hh \to bbWW^*$ decay is presented in this chapter. This analysis is based on the full Run 2 pp collision data collected by ATLAS between the years 2015–2018, totalling $139\,\text{fb}^{-1}$, and is a new addition to the searches for hh in ATLAS. It acts to improve the sensitivity to SM-like non-resonant hh production in the $bbWW^*$ channel that was lacking in the 2015–2016 analyses.

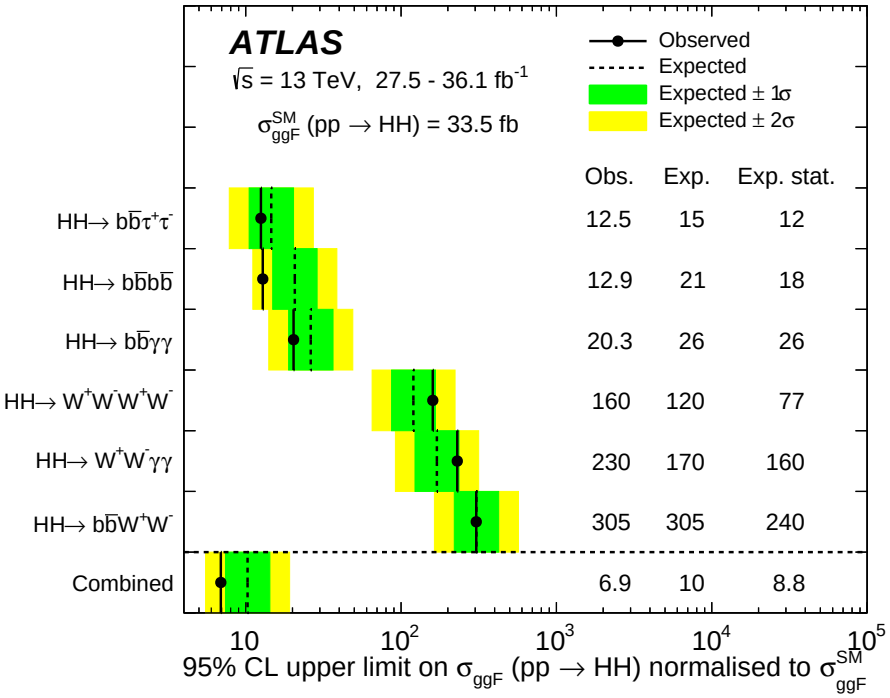

Figure 9.2: Summary of 95% CL cross-section upper-limits for non-resonant, SM-like hh production in ATLAS, based on the searches performed using $36\,\text{fb}^{-1}$ of data collected between 2015–2016. The $bbWW^*$ result is based on the single lepton channel of the $bbWW^*$ decay, where only one of the W-bosons decays leptonically and the other decays hadronically. Figure taken from Ref. [212]. The analyses in Ref. [212] use $\sigma_{hh}^{\text{SM}} = 33.5\,\text{fb}$.

9.1 The Dilepton $hh \to bbWW^*$ Signal Process

In the dilepton decay of the $hh \to bbWW^*$ channel, the WW^* system cannot be fully reconstructed in order to produce, for example, a WW^* invariant mass observable that has a resonance structure

at $m_h = 125\,\text{GeV}$. This is due to the kinematically under-constrained final state in which the $E_{\text{T}}^{\text{miss}}$ is a result of the double neutrino system arising from the leptonically decaying W bosons. Additionally, the final state is the same as that of SM $t\bar{t}$ production. As a result, the dominant SM background to the analysis is expected to be SM $t\bar{t}$ and top-quark processes.

There are crucial *topological* differences between the observable final states of dilepton SM $t\bar{t}$ production and that of $bbWW^*$, however. The dileptonic $bbWW^*$ final state is characterised by 'Higgs hemispheres'. One hemisphere contains the two b-tagged jets from the $h \to bb$ decay and is generally opposite in the transverse (xy) plane to the second hemisphere that contains the two leptons and $E_{\text{T}}^{\text{miss}}$ from the $h \to WW^*$ decay. In the case of $t\bar{t}$, the b-tagged jets, leptons, and neutrinos tend to be distributed more uniformly within the event, leading to events that do not typically exhibit the same back-to-back hemispheres topology as the $bbWW^*$ signal process. This is illustrated in Figure 9.3.

An additional feature of the $hh \to bbWW^*$ decay, available only in the dilepton final state, is related to the spin-structure of the $h \to WW^*$ decay. The Higgs boson is a fundamental scalar particle with spin-0 while the W bosons are spin-1 particles. In the $h \to WW^*$ decay, then, the spins of the W bosons must be oppositely aligned in order to conserve angular momentum. The spin-information encoded in the anti-aligned W-boson spins is preserved in their leptonic decays, governed by the parity-violating electroweak interaction. This is illustrated in Figure 9.4, showing the $h \to WW^* \to \ell\nu\ell\nu$ decay chain for two possible orientations of the W boson spins. The parity-violating weak decays require that the ν ($\bar{\nu}$) be left-(right-)handed which ultimately forces the charged leptons to emerge from the W-boson decays in the same direction. As a result, the charged leptons from the WW^* decays will be relatively collinear, as will the neutrinos, which further enhances the Higgs hemisphere signature of the dilepton $hh \to bbWW^*$ final state. These effects are most clearly observed in the dilepton invariant mass, $m_{\ell\ell}$, and in the opening angle between the two leptons, $\Delta\phi_{\ell\ell}$, shown in Figure 9.5. The analyses studying single Higgs production in the $h \to WW^*$ channel rely primarily on these effects, particularly on their impact on the $m_{\ell\ell}$ distribution, in order to separate out the Higgs events from the SM backgrounds [209, 210, 211].

The presence of these Higgs hemispheres in the signal, in which the momentum flow of the WW^* system is correlated with that of the bb system's, will be relied upon to inform the analysis' choice of kinematic observables that are sensitive to the dilepton $hh \to bbWW^*$ signal process.

In the following we present additional kinematic distributions illustrating the unique topologies of the dilepton $hh \to bbWW^*$ signal process as compared to the dominant top-quark background processes.

Figure 9.6 shows distributions related to the bb system. It can be seen from the $|\Delta\phi_{bb}|$ and ΔR_{bb} distributions that the two b-tagged jets are collinear in the signal process as compared to the background processes. For the $t\bar{t}$ background, these tendencies in $|\Delta\phi_{bb}|$ and ΔR_{bb} are reversed,

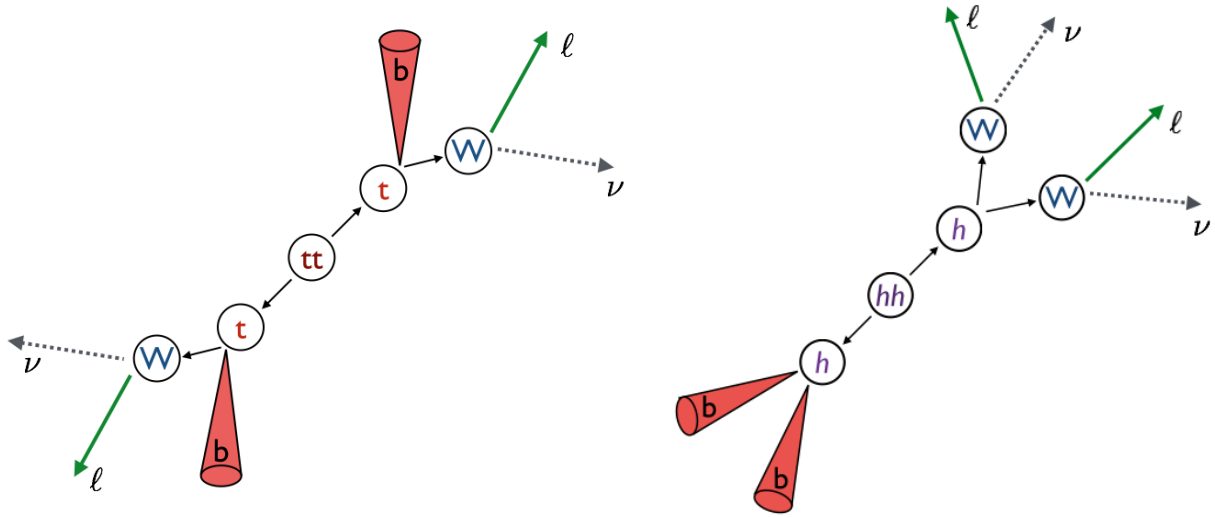

Figure 9.3: Dilepton $bbWW$ shapes. ***Left***: SM top-quark pair production. ***Right***: $hh \to bbWW^*$.

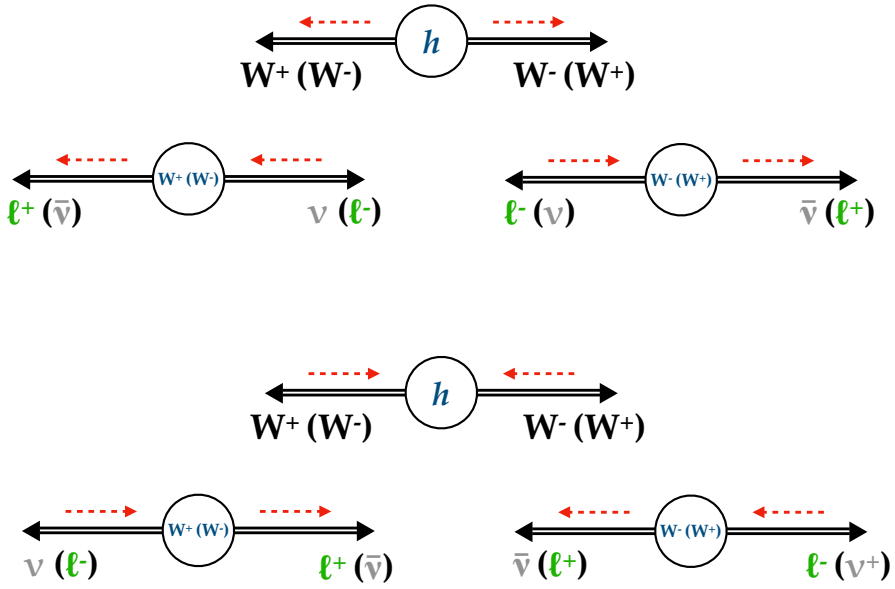

Figure 9.4: Illustration of the correlation in lepton momentum direction as a result of the decay of the spin-zero Higgs boson to W-boson pairs, which decay via the parity-violating electroweak interaction to the final state leptons. Double-lined arrows indicate the direction in which the particle travels and the dashed red arrows indicate the orientation of the particle spin-angular momentum relative to their direction of travel. ***Top***: Scenario with W-boson spins outward. ***Bottom***: Scenario with W-boson spins inward.

as expected from their different topology with the b-tagged jets on opposing sides of the event. For single-top Wt, similar trends as observed in $t\bar{t}$are observe but, since the second b-tagged jet in

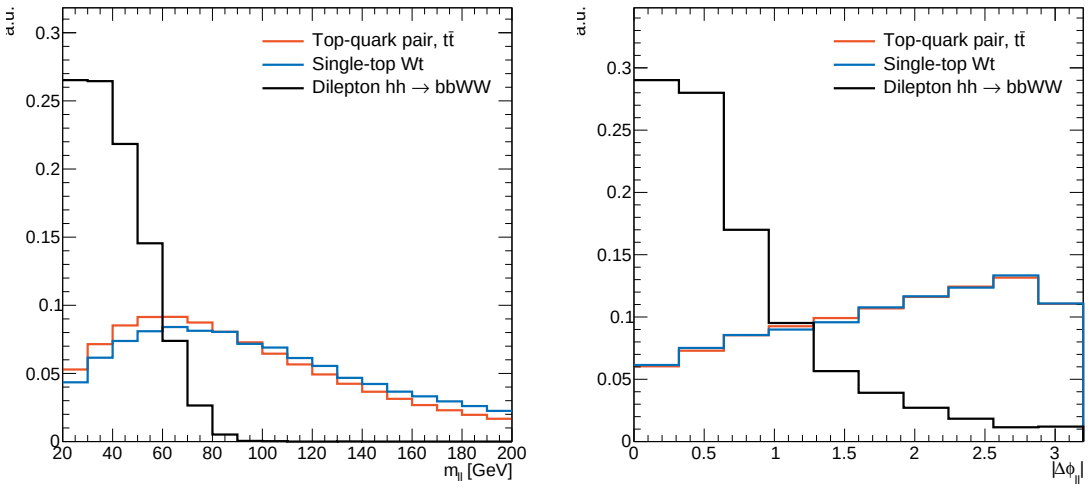

Figure 9.5: Normalized distributions showing the shapes of kinematic distributions for the SM top-quark processes ($t\bar{t}$ and single-top Wt) and the dilepton $hh \to bbWW^*$ signal process. ***Left***: Dilepton system invariant mass, $m_{\ell\ell}$. ***Right***: $|\Delta\phi|$ between the two leptons, $|\Delta\phi_{\ell\ell}|$.

the Wt process is due to higher-order contributions (or a light-flavor jet mistakenly identified as a b-jet), the b-tagged jets are often more collinear. The invariant mass of the leading two b-tagged jets, m_{bb}, is nicely peaking at the mass of the Higgs boson for the signal, as expected.

Figure 9.7 shows a few distributions related to the dilepton, $E_{\mathrm{T}}^{\mathrm{miss}}$, and $E_{\mathrm{T}}^{\mathrm{miss}}$ + dilepton systems. We can see from the $|\Delta\phi(\boldsymbol{p}_{\mathrm{T}}^{\mathrm{miss}}, \boldsymbol{p}_{\mathrm{T}}^{\ell\ell})|$ and $\Delta R_{\ell\ell}$ distributions that for the signal hh process, the leptons are collinear as expected, with the leptons contained almost entirely within a cone of radius $\Delta R_{\ell\ell} = 1$. For the top-quark backgrounds, there is a tendency for the leptons to be back-to-back, with $|\Delta\phi_{\ell\ell}|(\Delta R_{\ell\ell}) \to \pi$.

The $|\Delta\phi|$ between the $E_{\mathrm{T}}^{\mathrm{miss}}$ and dilepton system, $|\Delta\phi(\boldsymbol{p}_{\mathrm{T}}^{\mathrm{miss}}, \boldsymbol{p}_{\mathrm{T}}^{\ell\ell})|$, shows that the WW^* system is collinear, as well. For the SM top-quark backgrounds, $|\Delta\phi(\boldsymbol{p}_{\mathrm{T}}^{\mathrm{miss}}, \boldsymbol{p}_{\mathrm{T}}^{\ell\ell})|$ tends slowly towards π. In the case of single-top Wt, however, there is a tendency for $|\Delta\phi(\boldsymbol{p}_{\mathrm{T}}^{\mathrm{miss}}, \boldsymbol{p}_{\mathrm{T}}^{\ell\ell})|$ to gather towards lower values due to similar effects as described previously: the more collinear b-tagged jets force the WW^* system to also become collinear so that it balances the bb system.

To take further advantage of the topological differences mentioned in the preceding, we can construct an observable that is sensitive to the overall distribution of the momentum flow in the event. This is to capture the more 'planar' nature of the dilepton $hh \to bbWW^*$ decay as compared to that of the SM top-quark backgrounds. This observable, H_{T2}^{R}, is defined as follows,

$$H_{\mathrm{T2}}^{R} = \frac{|\boldsymbol{p}_{\mathrm{T}}^{\mathrm{miss}} + \boldsymbol{p}_{\mathrm{T}}^{\ell,0} + \boldsymbol{p}_{\mathrm{T}}^{\ell,1}| + |\boldsymbol{p}_{\mathrm{T}}^{b,0} + \boldsymbol{p}_{\mathrm{T}}^{b,1}|}{|\boldsymbol{p}_{\mathrm{T}}^{\mathrm{miss}}| + |\boldsymbol{p}_{\mathrm{T}}^{\ell,0}| + |\boldsymbol{p}_{\mathrm{T}}^{\ell,1}| + |\boldsymbol{p}_{\mathrm{T}}^{b,0}| + |\boldsymbol{p}_{\mathrm{T}}^{b,1}|}, \tag{9.2}$$

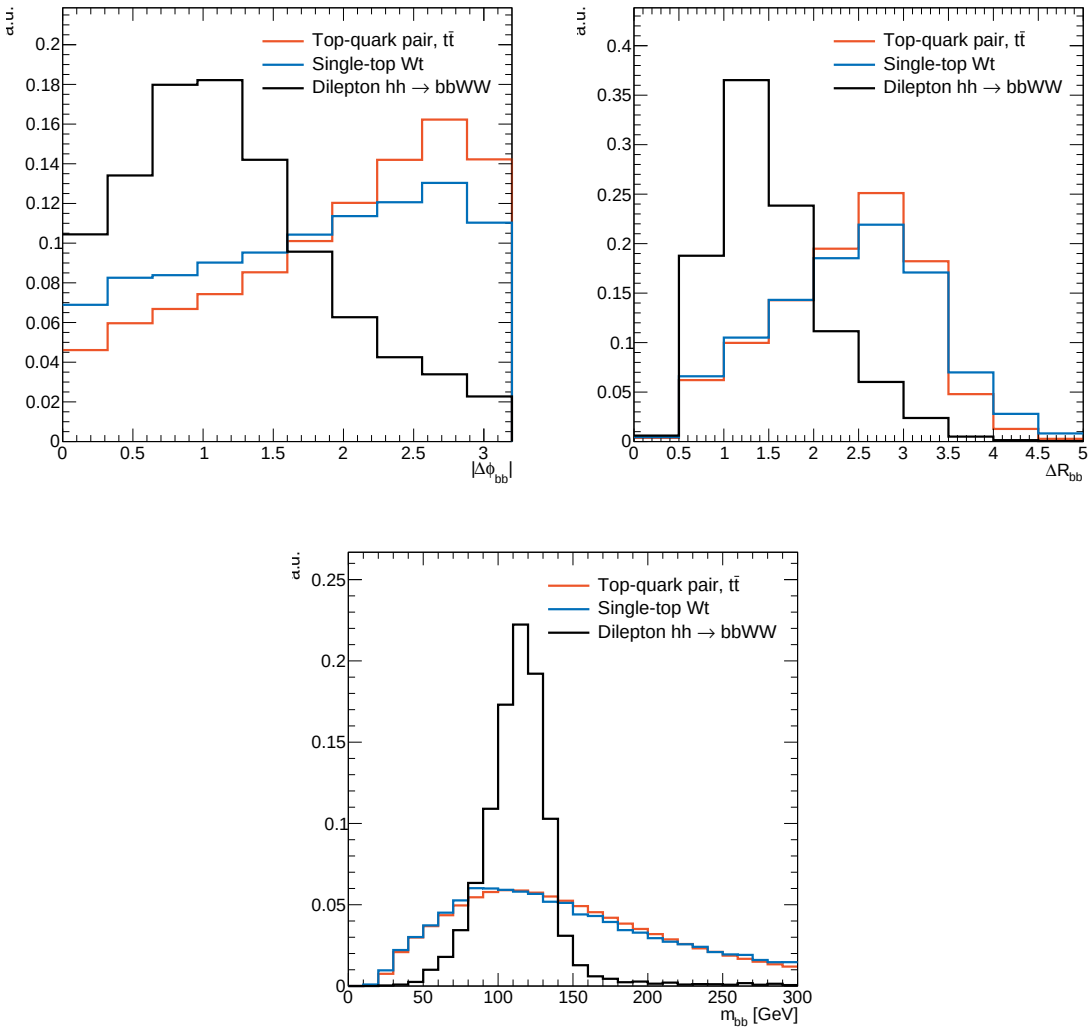

Figure 9.6: Normalized distributions showing the shapes of kinematic distributions for the SM top-quark processes ($t\bar{t}$ and single-top Wt) and the dilepton $hh \to bbWW^*$ signal process. ***Left***: $|\Delta\phi|$ between the two b-tagged jets, $|\Delta\phi_{bb}|$. ***Right***: ΔR between the two b-tagged jets, ΔR_{bb}. ***Bottom***: Invariant mass of the leading two b-tagged jets, m_{bb}.

where the $\boldsymbol{p}_{\mathrm{T}}^{\ell(b),0\{1\}}$ are the transverse momenta of the leading {subleading} lepton (b-tagged jet).

The numerator of Equation 9.2 contains two parts, each the magnitude of a vectorial sum of transverse momenta: the first part containing the objects associated with the WW^* decay and the second with those associated with the bb decay. For the dilepton $hh \to bbWW^*$ signal, then, the numerator is the sum of the magnitudes of each of the Higgs boson's momenta. The denominator appearing in Equation 9.2 is the scalar sum of the transverse momenta of each of the final state objects considered in teh analysis. By taking this ratio, a measure of the overall spread of the momentum in the event is obtained.

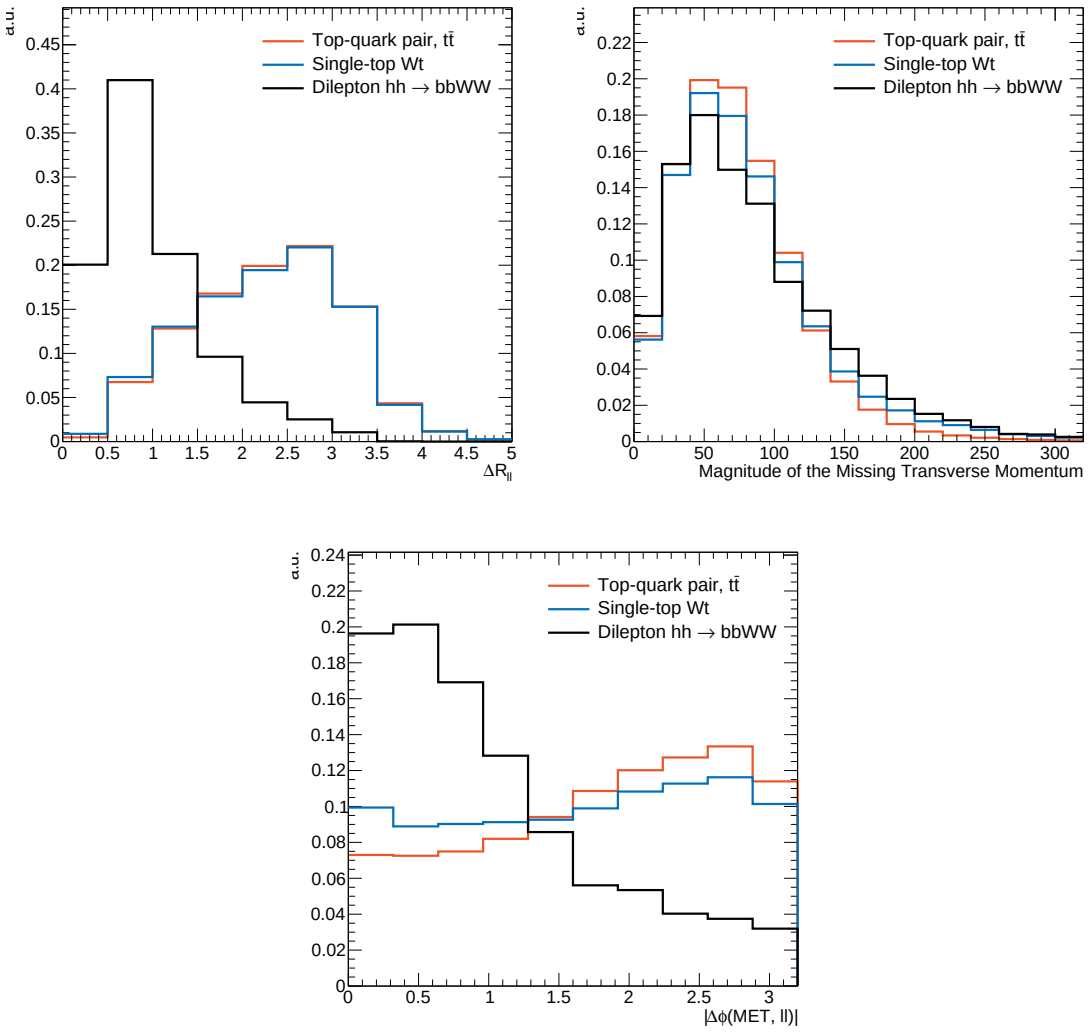

Figure 9.7: Normalized distributions showing the shapes of kinematic distributions for the SM top-quark processes ($t\bar{t}$ and single-top Wt) and the dilepton $hh \to bbWW^*$ signal process. ***Left, upper***: $|\Delta\phi|$ between the two leptons, $|\Delta\phi_{\ell\ell}|$. ***Right, upper***: ΔR between the two leptons, $\Delta R_{\ell\ell}$. ***Left, lower***: Magnitude of the missing transverse momentum, $E_{\mathrm{T}}^{\mathrm{miss}}$. ***Right, lower***: $\Delta\phi$ between $E_{\mathrm{T}}^{\mathrm{miss}}$ and dilepton system, $|\Delta\phi(\boldsymbol{p}_{\mathrm{T}}^{\mathrm{miss}}, \boldsymbol{p}_{\mathrm{T}}^{\ell\ell})|$.

For the relevant SM top-quark processes, with the visible final state objects distributed more evenly about the event, the quantities appearing in the numerator of Equation 9.2 will tend to be smaller, due to vectorial cancellation, than in the case of the hh signal. Due to the Higgs hemispheres in the signal, with separately collinear WW^* and bb systems, the numerator in Equation 9.2 will be large. It follows, then, that the quantity H_{T2}^{R} should tend towards unity for the dilepton $hh \to bbWW^*$ signal process and towards lower values for the SM top-quark processes. This behavior is illustrated in Figure 9.8.

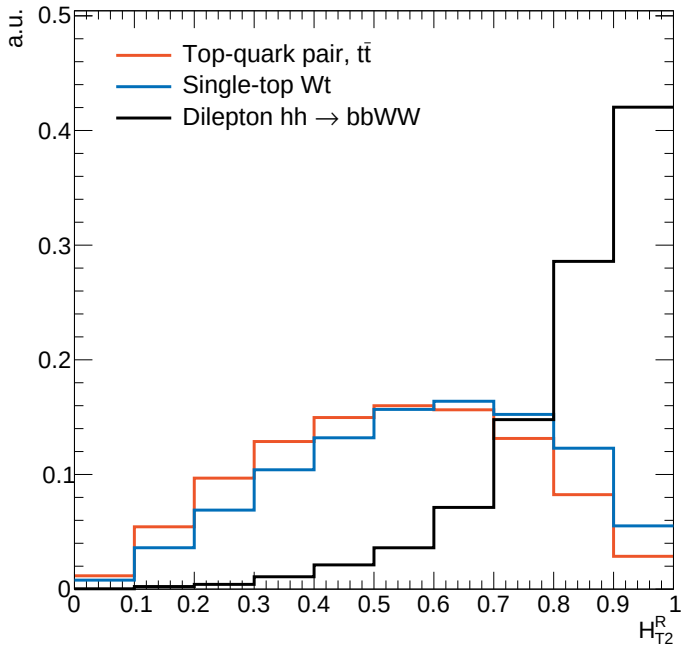

Figure 9.8: Normalized distributions showing the shape of the H_{T2}^{R} observable for the SM top-quark processes ($t\bar{t}$ and single-top Wt) and the dilepton $hh \to bbWW^{*}$ signal process.

We also consider a variant of the m_{T2} variable [213, 214, 215, 216], a transverse-mass variable that can be used in events where there are two or more invisible particles, as is our case with the two neutrinos. Generally, m_{T2} assumes that the decay is symmetric, as in $t\bar{t}$, where there are two decay legs, each with a visible and invisible child particle decay. The m_{T2} observable is generically defined as follows,

$$m_{\mathrm{T2}}^{2} \equiv \min_{\boldsymbol{q}_{\mathrm{T}}^{\mathrm{miss},\,a}+\boldsymbol{q}_{\mathrm{T}}^{\mathrm{miss},\,b}=\boldsymbol{p}_{\mathrm{T}}^{\mathrm{miss}}} \left[\max\left\{m_{\mathrm{T}}^{2}(\boldsymbol{p}_{\mathrm{T}}^{\mathrm{vis},\,a}, \boldsymbol{q}_{\mathrm{T}}^{\mathrm{miss},\,a}), m_{\mathrm{T}}^{2}(\boldsymbol{p}_{\mathrm{T}}^{\mathrm{vis},\,b}, \boldsymbol{q}_{\mathrm{T}}^{\mathrm{miss},\,b})\right\}\right], \tag{9.3}$$

where 'a' ('b') indicate one of the two sides of the decay and the $\boldsymbol{p}_{\mathrm{T}}^{i}$ are the transverse momenta of the visible objects on side $i \in (a, b)$, and $\boldsymbol{q}_{\mathrm{T}}^{i}$ is a partition of the total $\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}$ given to side $i \in (a, b)$. The minimization is carried out over all partitionings of the $\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}$ into the $\boldsymbol{q}_{\mathrm{T}}^{i}$ and the m_{T} are the typical transverse masses,

$$m_{\mathrm{T}}^{2}(\boldsymbol{p}_{\mathrm{T}}^{\mathrm{vis}}, \boldsymbol{q}_{\mathrm{T}}^{\mathrm{miss}}) = (E_{\mathrm{T}}^{\mathrm{vis}} + |\boldsymbol{q}_{\mathrm{T}}^{\mathrm{miss}}|)^{2} - |\boldsymbol{p}_{\mathrm{T}}^{\mathrm{vis}} + \boldsymbol{q}_{\mathrm{T}}^{\mathrm{miss}}|^{2} \tag{9.4}$$

The m_{T2}-based observable defined for the present analysis is referred to as 'm_{T2}^{bb}', and is an m_{T2} construction in which the visible objects are the two b-tagged jets in the event. This means that side 'a' ('b'), using the notation of Equation 9.3, contains one of the b-tagged jets as $\boldsymbol{p}_{\mathrm{T}}^{\mathrm{vis},\,\mathrm{a}}$ ($\boldsymbol{p}_{\mathrm{T}}^{\mathrm{vis},b}$).

Given the fact that in the SM top-quark backgrounds, the neutrinos (approximated by the $\boldsymbol{q}_{\mathrm{T}}^{\mathrm{miss},i}$) and the b-tagged jets originate from the same origin (a top-quark), the m_{T2}^{bb} observable tends to have a kinematic bound at roughly the mass of the top-quark, m_{top}. In the SM $t\bar{t}$ process, where both b-tagged jets originate from the decay of a top-quark, this kinematic endpoint is strict. In the case of single-top Wt, however, where only one b-tagged jet arises from the decay of a top-quark, the kinematic endpoint at m_{top} is not so strict and the m_{T2}^{bb} observable can extend beyond m_{top}. Figure 9.9 illustrates the m_{T2}^{bb} observable.

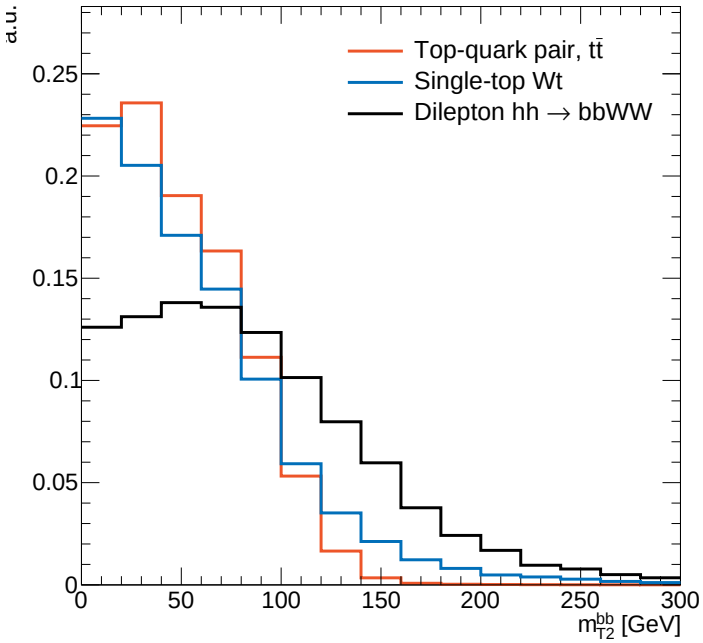

Figure 9.9: Normalized distributions showing the shape of the m_{T2}^{bb} observable for the SM top-quark processes ($t\bar{t}$ and single-top Wt) and the dilepton $hh \to bbWW^*$ signal process.

9.2 Event Selection and Object Definitions

9.2.1 Trigger Strategy

The $h \to WW^*$ decay, from which the leptons in the final state arise, is characterized by a sub-leading lepton which has typically quite low p_T as a result of one of the W-bosons being off-shell. If the analysis is to record the dilepton $hh \to bbWW^*$ events, then, it is important to minimize the p_T thresholds on at least one of the leptons in the event so that the signal events do not get discarded at the level of the trigger. As a result, the trigger strategy used in the dilepton $hh \to bbWW^*$ analysis is based on a combination of single and dilepton triggers. The single lepton

triggers typically enforce higher p_T thresholds on the single leading lepton than those encountered in the dilepton triggers, but they do not make any requirement on the sub-leading lepton candidate. Events with sub-leading leptons below the dilepton triggers' thresholds, then, have a chance of being recorded and kept for use in the analysis.

The typical p_T thresholds for the single-lepton triggers used in the present analysis are $26-28\,\text{GeV}$. There are two types of dilepton triggers used in the analysis: 'symmetric' and 'asymmetric'. Symmetric triggers enforce trigger p_T thresholds that are roughly the same for both the leading and subleading online lepton, and are typically around 20 GeV. The asymmetric triggers require relatively low p_T thresholds (lower than single lepton p_T thresholds) on the leading lepton, but very low thresholds (typically around 7 GeV) on the sub-leading lepton. The asymmetric triggers are only defined for dilepton events where the leptons have different-flavor ($e\mu$ or μe events).

The trigger strategy is illustrated in Figure 9.10. The single lepton triggers are given priority as they make no kinematic requirements on the sub-leading leptons in the event. Both leptons in the event are checked for having fired corresponding single lepton trigger before checking the dilepton triggers. If the event has a different-flavor lepton pair, the asymmetric triggers are given priority over the symmetric ones due to their having lower p_T thresholds on the sub-leading lepton. The trigger strategy outlined in Figure 9.10 is $\sim 95\%$ efficient for events with leading (subleading) lepton p_T values of 28 GeV (20 GeV) and $\sim 90\%$ efficient for events with subleading leptons having p_T values below 20 GeV.

9.2.2 Object Definitions

This section describes the definitions and working points of the leptons and jets used in the search for dilepton $hh \to bbWW^*$. The lepton definitions are given in Table 9.1 and the jet definitions are given in Table 9.2. Discussion of the working points for lepton identification and isolation is given in Section 5.2. That of jets is found in Section 5.3 and 5.3.4, for jets, generically, and b-tagged jets, respectively.

The p_T requirements of the leptons depends on the trigger used to select the event. That is, the point in the flowchart in Figure 9.10 at which a trigger fires defines the offline p_T threshold applied to the lepton: it is taken to be 2 GeV above the threshold of the fired trigger. Given the use of a wide set of single lepton and dilepton triggers, this threshold can vary between 19-28 GeV for the leading leptons and 10-17 GeV for the sub-leading leptons. As in the analysis discussed in Chapter 8, the lepton isolation requirements are rather loose, given that a large contamination of fake and non-prompt leptons is not expected as a result of the requirement of two relatively high-p_T leptons. The requirements on the impact parameter quantities $|d_0/\sigma_{d_0}|$ and $|z_0 \times \sin\theta|$ further ensure that the leptons are likely to be prompt and have originated from the primary hard-scatter

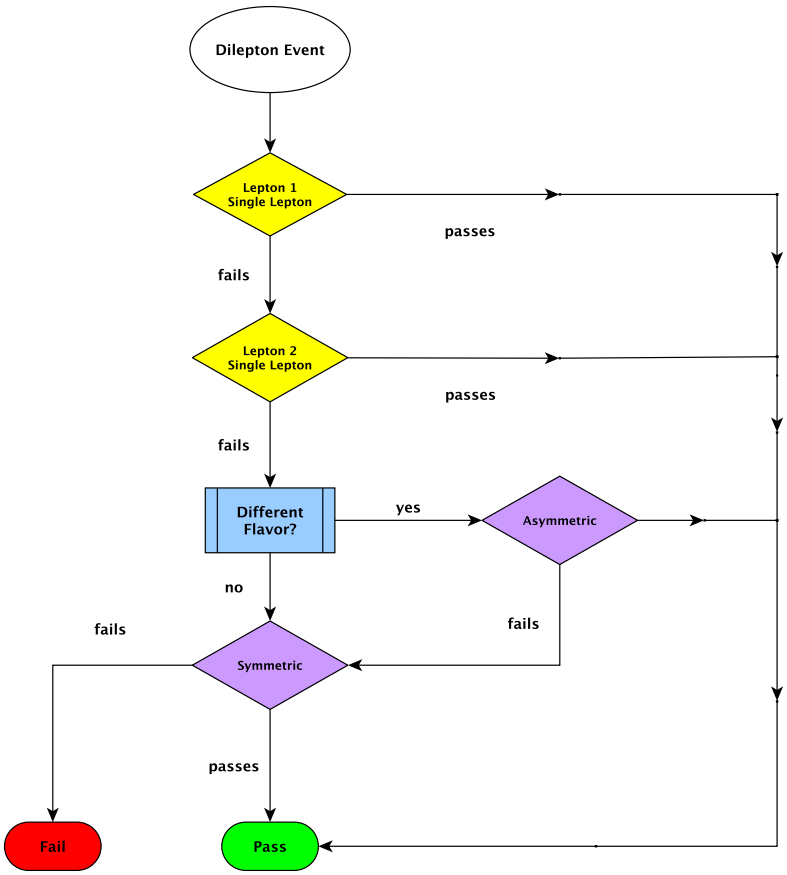

Figure 9.10: Flowchart illustrating the implementation of the trigger strategy used in the search for the dileptonic $hh \to bbWW^*$. The diamonds indicate points at which a specific trigger is checked. The yellow diamonds refer to the single lepton triggers associated with one of the leptons in the event. The purple diamonds refer to testing of dilepton triggers. In the figure, 'Asymmetric' and 'symmetric' refers to the asymmetric and symmetric different-flavor dilepton triggers, respectively. The red 'Fail' and green 'Pass' terminations refer to the event failing and succeeding, respectively, the trigger logic. Only events that pass are kept and considered in the analysis.

vertex.

9.2.3 Corrections to the b-tagged Jet Momenta

In the search for dileptonic $hh \to bbWW^*$, we apply the so-called 'muon-in-jet' energy correction to all b-tagged jets. This correction is used to correct the four-vector of all b-tagged jets on a *per-jet* basis by attempting to recover unaccounted for energy and momentum due to semileptonic b-decays within the b-tagged jets. As illustrated in Figure 9.11, in b-quark decays occuring in the b-tagged jets where you have at least a single decaying W-boson, you will have the possibility of having at least one lepton contained within the b-tagged jet. The W-boson decays at essentially equal rates

Table 9.1: Lepton definitions for the analysis searching for dilepton $hh \to bbWW^*$.

	Leptons			
	Electrons		**Muons**	
	Baseline	**Signal**	**Baseline**	**Signal**
p_T requirement [GeV]	$(> 10, > 10)$	Trig. threshold +2 GeV	$(> 10, > 10)$	Trig. threshold +2 GeV
$\|\eta\|$ requirement	< 2.47		< 2.4	
Identification WP	`Loose`	`Tight`	`Medium`	
Isolation	`GradientLoose`			
$\|d_0/\sigma_{d_0}\|$	$--$	< 5	$--$	< 3
$\|z_0 \times \sin\theta\|$ [mm]	$--$	< 0.5	$--$	< 0.5

Table 9.2: Jet definitions for the analysis searching for dilepton $hh \to bbWW^*$.

	Jets	***b*-tagged Jets**
p_T requirement [GeV]	> 20	
$\|\eta\|$ requirement	< 2.8	< 2.4
Pileup suppression	`JVT` > 0.59 if $p_T < 60\,\text{GeV}$ and $\|\eta\| < 2.4$	
Flavor-tagging WP	$--$	70%

to all flavors of leptons, however we cannot efficiently the electrons and hadronic τ decays occuring inside of jets. We can, however, identify muons originating from in-flight decays within jets, as they do not deposit a significant amount of energy within the jet and escape the calorimeter. The electron and τ energy deposits, on the other hand, will for the most part be contained within the jet and so their energy is, at least partially, accounted for in the jet reconstruction.

As the escaping muon's energy and momenta are not fully taken into account when reconstructing the b-tagged jets, it is important to see if we can recover this lost energy. The aim being to improve the b-tagged jet energy response and therefore improve the resolution of the invariant mass of the reconstructed bb-system, m_{bb}, which peaks at $m_h = 125\,\text{GeV}$ in the hh signal being searched for. As will be seen in Section 9.3, the current analysis requires that signal events have m_{bb} within a small window about m_h. If we can improve the m_{bb} resolution, then we can increase the signal purity of this selection.

The muon-in-jet correction proceeds as follows. Any muons satisfying the baseline requirements described in Table 9.1 that are also with $\Delta R = 0.4$ of a b-tagged jet are first identified. The *nearby* muon with the highest p_T then has its four-momentum 'added back' to the reconstructed b-tagged jet. Muons deposit, on average, $3-4\,\text{GeV}$ of energy in the calorimeter. Once a muon's four momentum has been added back to the b-tagged jet in question, the calorimeter energy deposit of the muon is subtracted away from the b-tagged jet's four momenta so as to avoid double counting

of the calorimeter energy: once from the muon, now considered part of the b-tagged jet, and a second time from the energy deposit already having been used in the jet reconstruction.

Figure 9.12 shows the m_{bb} distribution for the dilepton $hh \to bbWW^*$ signal process before and after performing the muon-in-jet correction as described above. With the m_{bb} selections described in Section 9.3, the muon-in-jet correction increases the hh signal acceptance by roughly $5-7\%$ (depending on any additional selections being applied on top of the m_{bb} ones).

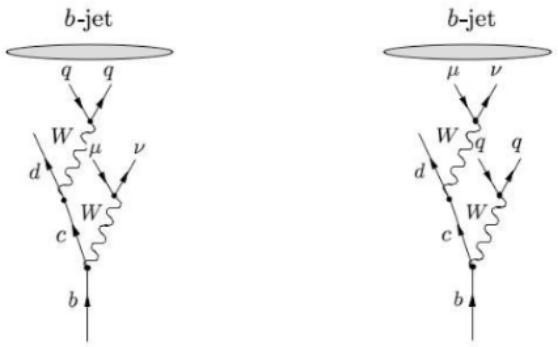

Figure 9.11: Diagrams illustrating the semileptonic weak decays of the b- and c-quarks initiating the reconstruction of a b-tagged jet. The muons and neutrinos escape the calorimeter and their momenta and energy are not accounted for in the b-tagged jet reconstruction.

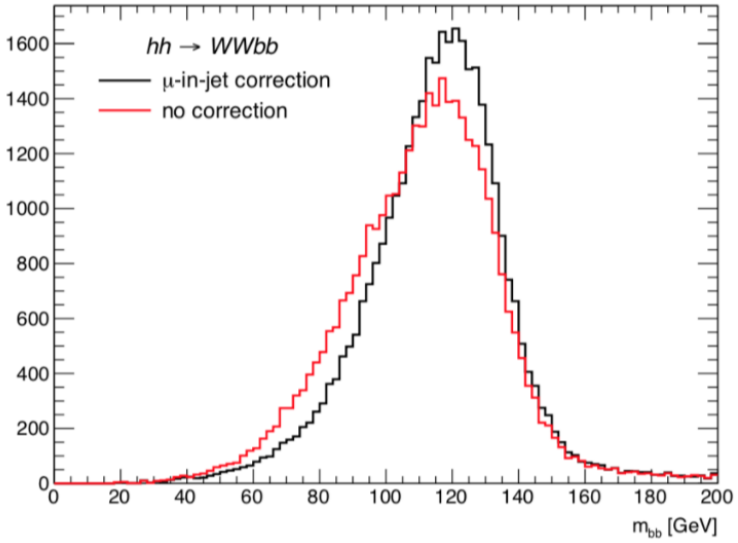

Figure 9.12: The invariant mass of the leading two b-tagged jets in the dilepton $hh \to bbWW^*$ signal process before (red) and after (black) the muon-in-jet momentum correction. Only the analysis' preselection requirements are applied to the events being shown in the distribution.

9.2.4 Standard Event Pre-selection

The preselection requirements applied to all events in the search for dilepton $hh \to bbWW^*$ is the same as that described in Chapter 8, and fully detailed in Table 8.3.

9.3 Neural Network Based Signal Selection

The current analysis makes use of a multi-output classifier, one that does not simply classify a single process against a single background label, but rather a classifier that provides multiple output labels with each pertaining to a distinct class or process. One of the easiest ways to build such a classifier is to take a multi-variate approach that is by default suitable for multi-output classification: neural networks.

9.3.1 Classifier Architecture

The analysis makes use of a deep-learning, neural network based approach. The classifiers that we build are trained to classify pp collision events according to four potential class labels, inspired by the dominant expected background processes:

1. Dilepton non-resonant $hh \to bbWW^*$
2. SM top-quark processes ($t\bar{t} + Wt$), 'Top'
3. SM Z+jets processes, $Z \to \{ee, \mu\mu\}$
4. SM Z+jets processes, $Z \to \tau\tau$

The classifier is trained with separate labels for the $Z \to \{ee, \mu\mu\}$ and $Z \to \tau\tau$ processes as these lead to clearly different final state kinematics. The dilepton final state that we eventually select in the analysis is composed only of electrons and/or muons. The $Z \to \tau\tau$ process contributes only in the cases where both τ leptons decay leptonically. The electrons and muons from these τ decays have very different kinematic signatures as compared to those from the direct decays of the Z-bosons. Allowing the classifier to learn how to distinguish between these Z decays improves its overall performance to separate the hh signal process from the backgrounds.

We construct the neural network architecture using the Keras [217] library, using Tensorflow [218] as a backend. An illustration of the neural network architecture is given in Figure 9.13. The network inputs are passed through a dense (fully-connected) layer, which is trained with a dropout layer, and then a second dense layer. The final activation is a softmax activation [219].

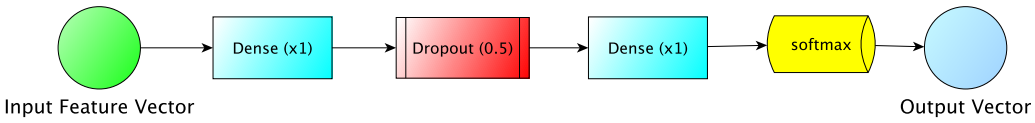

Figure 9.13: Illustration of the neural network graph employed in the analysis. The input feature vector has a length of 35 and the output vector is length 4, one for each of the targeted processes.

Each of the dense layers are 250 nodes wide and have their weights randomly initialized by sampling from a truncated normal distribution centered on zero with a width given by $\sqrt{1/N_{\text{inputs}}}$, where N_{inputs} is the number of input features (the length of the input vector). The activation functions for each of the dense layers are rectified linear units ('ReLu') [220].

Using an output layer with a softmax activation function allows one to interpret the outputs as each representing a probability[1] for the output's associated class (hh, Top, $Z \to \{ee, \mu\mu\}$, or $Z \to \tau\tau$) given the inputs and for this reason it is commonly used for multi-class neural network classifiers. The association of the softmax activation with a class probability can be seen by its definition,

$$a_j = \frac{e^{z_j}}{\sum_k e^{z_k}}, \tag{9.5}$$

where a_j is the activation of the j^{th} output neuron, the z_i are the inputs to the output layer, and k runs over all output neurons. It can be seen that if one sums over all outputs of a layer whose activation is given by Equation 9.5 that the sum is equal to one. Thus, the outputs of the softmax layer can be seen as a probability distribution. For this reason, in the discussion to follow, we refer to the outputs of our neural network as 'p_i', where i has four possibilities for each of the four outputs: $i \in \{hh, \text{Top}, Z \to ee/\mu\mu, Z \to \tau\tau\}$.

The use of dropout layers during the training process of is a form of statistical learning regularization that is reminiscent of ensemble methods in the non-deep-learning arena, such as random forests [221]. They act to randomly disable a tunable fraction of inputs during various points in the training stage [222]. This tunable fraction is referred to as the *dropout rate*. The use of dropout regularization prevents nodes within the network from co-adapting too much, thus reducing the effects of overtraining. This is illustrated in Figure 9.14. During each batch of events forwarded to the network during the training phase, the dropout layer disables portions of the network and

[1] The use of the term 'probability' here is not entirely correct, as the outputs are not *strictly* probabilities given the fact that the network does not know about the full set of possible processes that will be presented to it when it is used outside of the training stage and in the analysis. This is clear from the fact that the sum appearing in the denominator of Equation 9.5 is not over the full set of processes that will be provided as input to the network when it is used in the analysis. The sum is only over those processes that we have decided to train the classifier to provide labels for. Additionally, the labelled processes are presented in artificial proportions during the training process, described in Section 9.3.2, which may further prevent the network output's from being viewed as proper probabilities.

thereby presents a modified network to the inputs. Conceptually, then, using dropout during training is similar to training a set of very many, different *weak* neural networks. During test time, at the time when the neural network is actually being used in the analysis, the network's weights, which have been determined after training over the set of thinned networks, are scaled down by the dropout rate. This is illustrated in Figure 9.15.

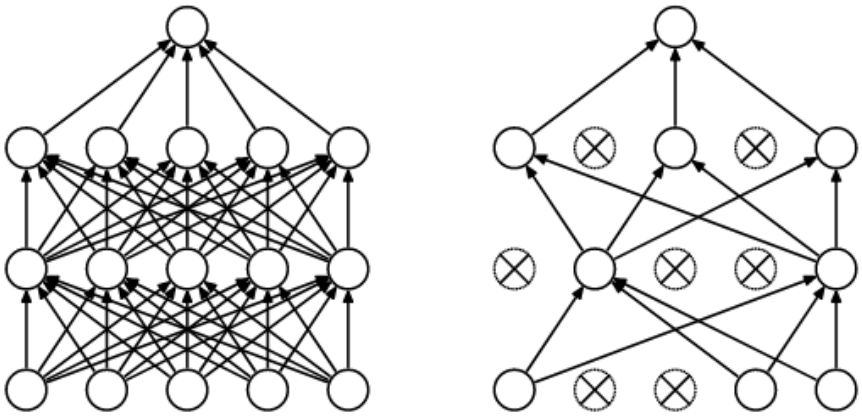

Figure 9.14: Illustration of dropout regularization. Figures taken from Ref. [222]. ***Left***: A standard neural network with two fully-connected layers. ***Right***: An example of a thinned network produced by applying dropout to the network on the left. The units with 'X' have been dropped.

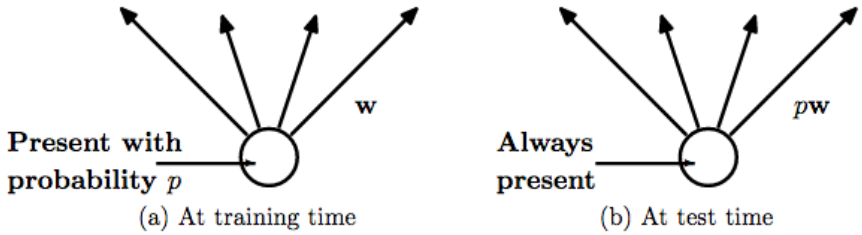

Figure 9.15: Illustration of the dropout rate effect on the network weights. Figures taken from Ref. [222]. ***Left***: A node in a fully-connected layer at training time is present in the network with a probability equal to the dropout rate and is connected to the next layer with weights represented by $\boldsymbol{w}$. ***Right***: At test time, the node is present with 100% probability but its weights are scaled down by the dropout rate, $p\boldsymbol{w}$.

As mentioned above, the use of dropout regularization prevents nodes within the network from co-adapting too much and forces the network to learn more robust features that are useful in conjunction with many different random subsets of the other nodes. That is, dropout regularization ensures that the model is robust against the loss of any individual "piece of evidence" and is found to reduce the effects of overtraining, which improves the generalizability of the trained classifier.

The neural network classifier used in the present analysis, illustrated in Figure 9.13, uses a single

dropout layer acting on the first fully-connected node and is given a dropout rate of 50%.

During training, the loss metric is the categorical crossentropy and the Adam optimization algorithm [223] is used.[2]

9.3.2 Definition of the Training Sample

Prior to training a multi-variate classifier, such as the neural network described in Section 9.3.1, a well-defined set of simulated events for all labelled processed must be set aside for use in the training. Ideally, these events are not to be used in the actual analysis, when the trained classifier is used only for inference. These two sets are referred to as the *training set* and *testing set*, respectively.

For defining our split between training and testing sets, the present analysis uses the 'hold-out method' wherein a fraction of the training events are set aside for use in evaluating classifier performance metrics during the classifier training. This is illustrated in Figure 9.16, where we use, out of our entire set of MC simulated events, a percentage ($a\%$) for defining the training sample, with subsets of this set aside for evaluation of the training performance.

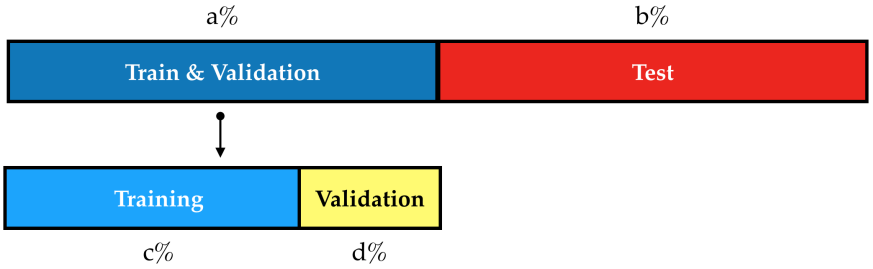

Figure 9.16: Illustration of the sample split used for training and testing the classifier. 'Test' refers to the sample of events used at the analysis level (classifier inference only). 'Training' and 'Validation' events are used only during the training stage: Those events marked as 'Training' are used for training of the classifier and those marked as 'Validation' are used to evaluate the network performance metrics throughout the training process. a is the percentage of the sample used for training and validation, b the percentage used in the analysis, c is the percentage of the Training data used for network training, and d is the percentage of the Training data used for evaluating performance metrics. This gives $a\% + b\% = 100\%$ and $c\% + d\% = a\%$.

During the training, the network is presented with inputs from MC-simulated events from each of the four processed described above, each of which having a corresponding network output label. The method used ensures an equal balance of processes in the training sample and revolves around the total number of events existing in the MC sample of the signal process. The number of events corresponding to 50% of the total number of events in the signal MC sample is set aside from each

[2] More on categorical cross-entropy: https://ml-cheatsheet.readthedocs.io/en/latest/loss_functions.html#cross-entropy

of the four samples to be used for training: hh, Top, $Z \to \{ee, \mu\mu\}$, and $Z \to \tau\tau$. The sum of all of these events corresponds to the 'Training' sample in Figure 9.16. As the neural network is trained, we want to make sure that at any given moment in the training process that the network has a roughly equal chance of being presented with an event from any one of the four possible processes. This ensures that no biases are learned by the network simply as a result of having an imbalanced training sample. To avoid this, the training sample is randomly shuffled such that events in the Training sample are evenly distributed. This breakdown and shuffling of the training events is illustrated in Figure 9.17. For the neural network trained in the present analysis, 20% of the Training sample is set aside to make up the Validation sample.

In addition to equally balancing the number of representative examples of each of the classes that the neural network will be trained to identify, the MC events are not weighted during the training. That is, the MC events are not weighted to their typical MC event weights associated with their cross-section, reconstruction efficiencies, etc... nor are they given specific global class weights in the training. All processes therefore have equal weight and equal sampling from the perspective of the training of the classifier. In the end, these parameters are additional *hyperparameters* entering the training that must individually be tuned, if it is desired to do so. The choices made here, with respect to the process sampling and weighting, were made with the aim of designing the simplest training procedure possible.

All events used in the training procedure are required to satisfy the analysis' preselection requirements defined in Section 9.2 as well as having at least one b-tagged jet, as detailed in Table 9.3. This is a rather loose selection — looser than the one that will eventually be applied in the definition of the analysis' SRs, CRs, and VRs — but allows for larger numbers of events to be used in the training process.

Table 9.3: Selection applied to events used in the neural network classifier training. In addition to this selection, the events must pass the standard preselection defined in Section 9.2.

Observable	**Selection**
Analysis preselection	Applied
b-tagged jet multiplicity	≥ 1

9.3.3 Input Features and their Preprocessing

The thirty five inputs to the neural network are detailed in Table 9.4. They are composed of low level features of the visible final state objects, such as the lepton and jet momenta, as well as observables sensitive to the Higgs hemisphere topology characterising the dilepton $hh \to bbWW^*$ signal and described in Section 9.1.

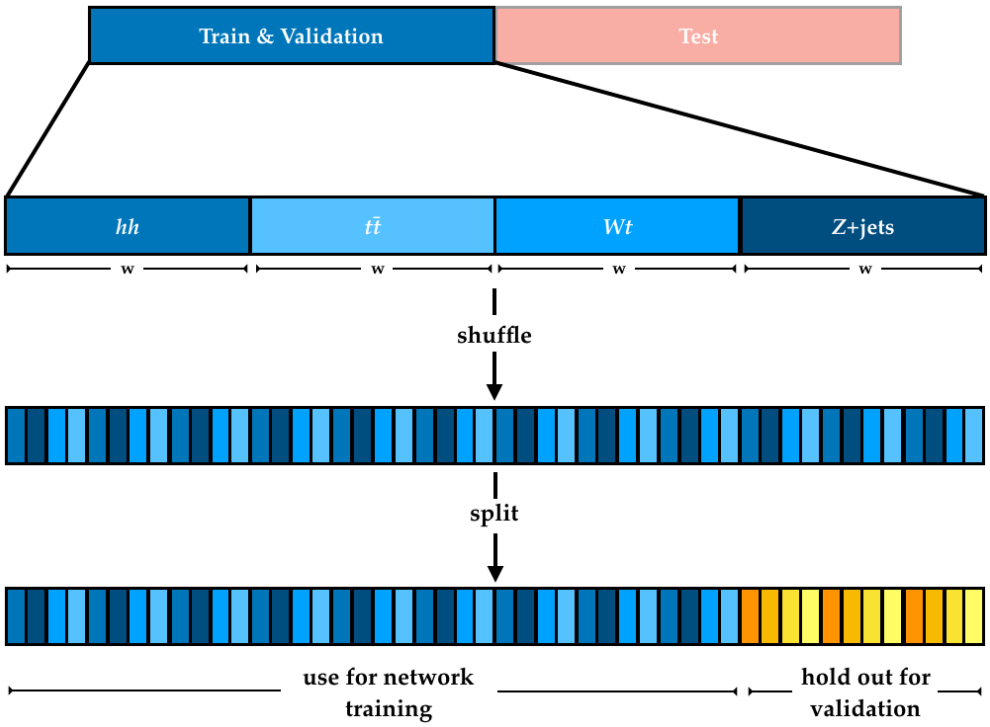

Figure 9.17: Detailed illustration of how the allocated training data is split into separate training and validation data samples. The sample size, or width $\boldsymbol{w}$, is defined by the signal MC sample: $\boldsymbol{w}$ is defined to be half of the events available in the signal MC. This width defines the sample sizes of the remaining three processes for which the neural network classifier provides an output label. Before the training process begins, these samples are randomly shuffled such that each process is distributed evenly throughout the entire set of events set aside for training. The validation sample is defined by holding out 20% of the shuffled data. Labels and scales in the figure are illustrative only.

As discussed in Section 9.3.2 and indicated in Table 9.3, the neural network classifier is trained on events with the minimal requirement that they pass the analysis preselection requirements and that there be at least one b-tagged jet in the event. One can see by comparing the training event selection in Table 9.3 and the network inputs in Table 9.4, that there will be cases where several of the input features will be ill-defined. Specifically, there will be cases wherein a training event has only a single b-tagged jet in which case the observables listed in Table 9.4 that require at least two b-tagged jets are not defined. What is done in these cases is to fix the two b-tagged jet observables in Table 9.4 for these one b-tagged jet events to the *mean* of the corresponding observable as seen in those events that have at least two b-tagged jets. The mean for each observable is defined using as sample of events all those events that make up the training sample, inclusive of all four of the trained-against processes. That is, distributions of all observables are constructed for each of four processes (hh, Top, $Z \to \{ee, \mu\mu\}$, $Z \to \tau\tau$) for those events with at least two b-tagged jets. For each observable, the four process-specific distributions are summed together to produce the final

distribution from which the mean value is extracted. During the network training, these mean values are used for the observables requiring at least two b-tagged jets to be defined in those events only having a single b-tagged jet. This procedure, illustrated in Figure 9.18, is done only during the training phase of the classifier.

A second stage of input feature pre-processing is applied to all features provided to the classifier, done at both the training and in the testing stage of the analysis. This stage is a standardization step and involves, for each event, the shifting of each input feature by its mean value and scaling its value by its pre-shifted variance,

$$x' = \frac{x - \langle x \rangle}{\sigma_x}.$$

This type of standardization is fairly standard practice, and is done so that all inputs provided to the neural network are on a similar range — centered at zero and with common spread — thereby making it easier for the classifier to learn the relationships (i.e. correlations) between the inputs, as opposed to their absolute scales. These standardization schemes additionally aid in numerically stabilizing the neural network's gradient-descent-based optimization procedures employed during the training phase. The standardization described above and used in the present analysis is illustrated in Figure 9.19.

Table 9.4: Description of the variables used as inputs to the DNN classifier.

(p_T, η, ϕ)	p_T, η, and ϕ of the leptons, leading two signal jets, and leading two b-tagged jets
Dilepton flavour	Whether the event is composed of two electrons, two muons, or one of each
$\Delta R_{\ell\ell}$, $\lvert\Delta\phi_{\ell\ell}\rvert$	ΔR and magnitude of the $\Delta\phi$ between the two leptons
$m_{\ell\ell}$, $p_T^{\ell\ell}$	Invariant mass and the transverse momentum of the dilepton system
$E_{\mathrm{T}}^{\mathrm{miss}}$, $E_{\mathrm{T}}^{\mathrm{miss}}$-$\phi$	Magnitude of the missing transverse momentum vector and its ϕ component
$\lvert\Delta\phi(\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}, \mathbf{p}_{\mathrm{T}}^{\ell\ell})\rvert$	Magnitude of the $\Delta\phi$ between the $\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}$ and the transverse momentum of the dilepton system
$\lvert\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}} + \mathbf{p}_{\mathrm{T}}^{\ell\ell}\rvert$	Magnitude of the vector sum of the $\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}$ and the transverse momentum of the dilepton system
Jet multiplicities	Numbers of b-tagged and non-b-tagged jets
$\lvert\Delta\phi_{bb}\rvert$	Magnitude of the $\Delta\phi$ between the leading two b-tagged jets
m_{T2}^{bb}	m_{T2} using the leading two b-tagged jets as the visible inputs and $\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}}$ as invisible input
H_{T2}	Numerator appearing in Equation 9.2, $\lvert\boldsymbol{p}_{\mathrm{T}}^{\mathrm{miss}} + \boldsymbol{p}_{\mathrm{T}}^{\ell,0} + \boldsymbol{p}_{\mathrm{T}}^{\ell,1}\rvert + \lvert\boldsymbol{p}_{\mathrm{T}}^{b,0} + \boldsymbol{p}_{\mathrm{T}}^{b,1}\rvert$
H_{T2}^{R}	As defined in Equation 9.2

Note that it is this standardization procedure which motivates the method that is used for augmenting the two b-tagged jet observables in the one b-tagged jet events. By setting the ill-defined observables in the on b-tagged jet events to the mean of the observables as defined in the two b-tagged jet events, we theoretically prevent the neural network from learning from that specific observable on such events since after standardization they will be shifted towards zero and any spurious correlation will be minimized. In practice, however, it has been found that the specific choice of the scheme used for augmenting the two b-tagged jet observables in one b-tagged jet events has little meaningful impact on the performance of the classifier. The one described above has been

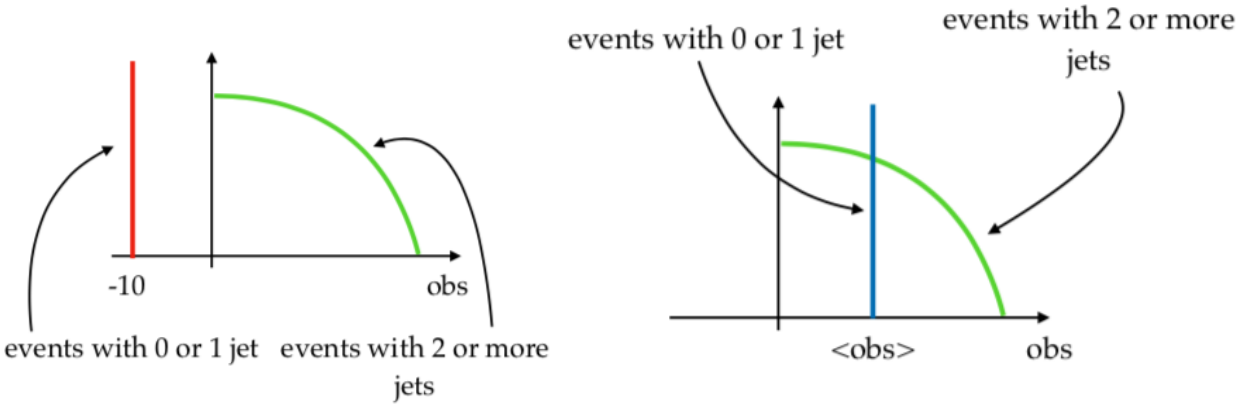

Figure 9.18: Illustration of setting events with ill-defined observables to the mean of those observables as measured in events in which the observables are properly defined. ***Left***: A counter-example in which ill-defined events' observables are set to a fixed default value outside of the physical range (a common practice). This would bias the mean of the properly defined events to the left. ***Right***: Setting the ill-defined events' observables to the mean of the observable as seen in events where the observable is properly defined. Doing this does not affect the overall mean of the observable in events where the observable is well-defined, nor does it lead to a discontinuity in the distributions of observables being provided to the classifier during its training.

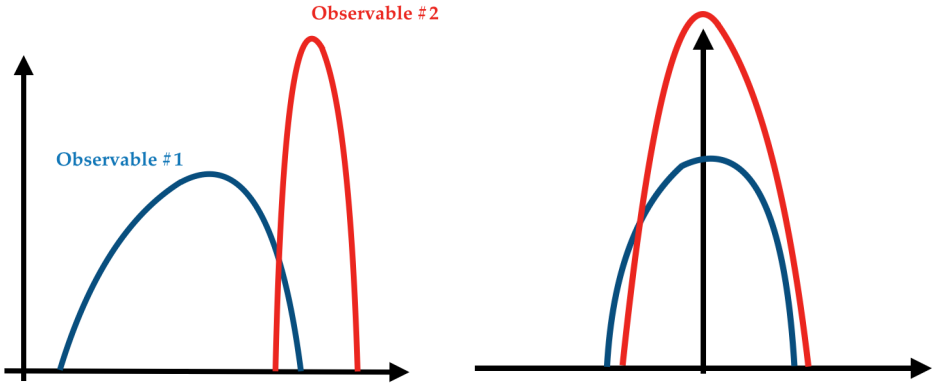

Figure 9.19: Illustration of the input feature standardization that occurs for all input features, for all events, prior to the features being fed into the neural network classifier (for both training and testing). ***Left***: Two different observables (overlaid for comparison) that have quite different shapes. Observable #1 has a larger variance and is shifted to the left with respect to Observable #2, which is a quite localized observable. ***Right***: After standardization, both observables have roughly the same mean (zero) and their widths have been scaled such that they are similar.

used as it is the simplest, practically speaking.

The pre-processing steps, happening at each event and in order, are as follows:

1. Augment two b-tagged jet observables for one b-tagged jet events (if applicable)

2. Perform input feature standardization

3. Feed inputs to network (either for training or for use in the analysis)

9.3.4 Training Procedure

For illustration purposes, Figure 9.20 illustrates the process of training a neural network with many batches of events, the sum total of which comprise the entirety of the events set aside for training. An epoch refers to the point at which every single event in the training sample has been fed into the network. The neural network classifier described in Section 9.3.1 is trained using the sample of events described in Section 9.3.2 using a batch size of 2,000 events, with a maximum number of training epochs set to a large number ($\mathcal{O}(100)$).

Neural network classifiers typically train for many epochs, meaning that every event in the training sample is used multiple times in the training procedure. The training procedure used in the present analysis relies on an 'early-stopping' metric that is imposed during the training phase such that, if the network's learning begins to plateau, it will stop the training whether or not the maximum number of epochs has been reached. The metric used to determine whether the early-stopping criteria has been met is the neural network loss value as evaluated on the held-out validation sample of events. The aim of the network, during the training phase, is to minimize the loss. Therefore, once the metric used for defining the early-stopping criteria hits a (minimum) plateau it will stop after a configurable number of epochs in which no new minimum is found. This configurable number, referred to as the 'patience', is set to 20 for the classifier used in the present analysis. Figure 9.21 shows several training performance metrics evaluated on both the training and held-out validation data, from which it can be seen that setting the patience parameter to 20 allows the network to maximize its learning before effects of overtraining begin to set in.

9.3.5 Classifier Discriminants

The multi-output classifier described in Section 9.3.1 provides four outputs for each event once provided an input feature vector. In our case, the four outputs are taken to loosely represent the probabilities, or strengths, that the set of inputs correspond to one of the four labels for which the classifier has been trained to distinguish. We refer to these outputs, then, as p_{hh}, p_{Top}, $p_{Z-\text{SF}}$, and $p_{Z-\tau\tau}$ for the hh, Top, $Z \to \{ee, \mu\mu\}$, and $Z \to \tau\tau$ process labels, respectively. Representative distributions for each of the p_i, for the corresponding signal and background processes, are shown in Figure 9.22.

From the set of p_i, we build a set of four composite discriminants, each of which combines information contain within them all. These discriminants are log-ratio discriminants of the output scores,

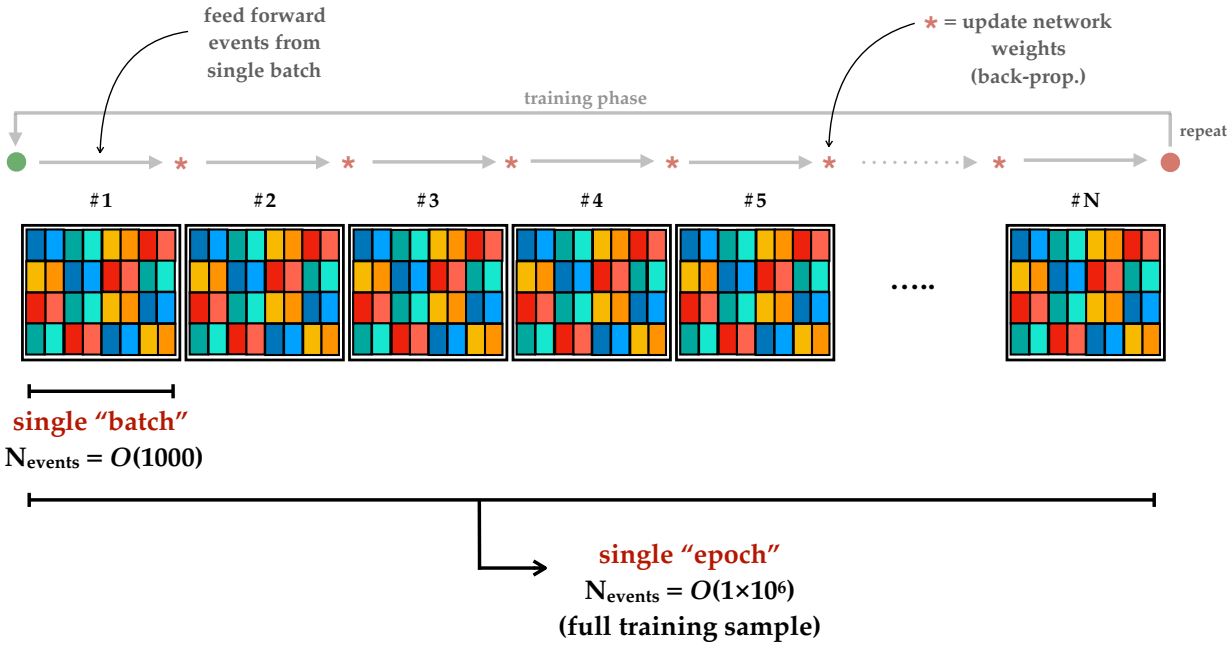

Figure 9.20: A neural network training sample is broken into many sub-samples referred to as 'batches'. In the training phase of the neural network, the network parameters (weights and biases) are only updated after each batch is fed into the network. A training 'epoch' refers to the point in the training at which all of the events in the entire training sample have been fed forward through the network for training purposes. The size of each batch (i.e. the number of events that are in a batch), as well as the number epochs to train for, are configurable parameters that must be optimized. See Ref. [219] for more information.

defined in Equation 9.6. Representative distributions for each of the composite discriminants are shown in Figure 9.23. The discriminant d_{hh} is the primary observable around which the SRs, CRs, and VRs are defined in the present analysis.

$$d_Q = \ln\left(\frac{p_Q}{p_X + p_Y + p_Z}\right) \tag{9.6}$$

9.3.6 Signal Region Definition

The definitions of the SRs, SR-SF and SR-DF, targeting the dilepton $hh \to bbWW^*$ signal process are tabulated in Table 9.5 and are as follows. The dominant expected backgrounds in the analysis are those for which separate output labels have been defined in the construction of the neural network classifier described in the previous sections: SM top-quark processes ($t\bar{t}$ and single-top

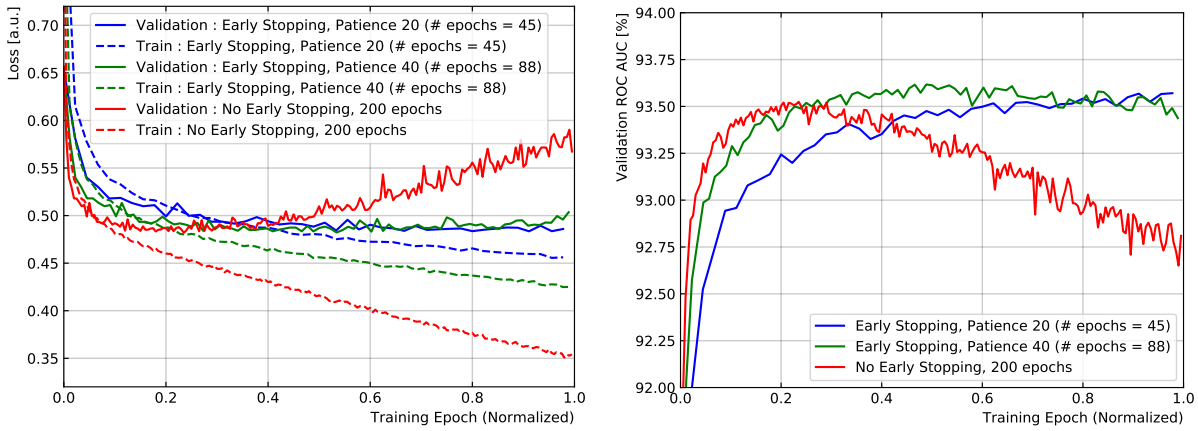

Figure 9.21: ***Left***: The neural network loss as a function of the number of training epochs. The solid lines report the loss evaluated on the held-out validation sample of events, while the dashed lines report that evaluated using the sample of events used for training. The larger discrepancy observed between the validation and training loss observed in the cases where the training is allowed to proceed for a larger number of epochs indicate larger amounts of overtraining. ***Right***: Receiver operating characteristic area-under-the-curve (ROC AUC), evaluated using the held-out validation sample of events, as a function of the number of training epochs. Larger values of ROC AUC indicate better network classification performance. The drop in ROC AUC observed in the green and red curves are a result of increased levels of overtraining, resulting in the classifier being unable to generalize to data other than that used for training.

Wt) and Z+jets processes. The top-quark processes are flavor symmetric, however the Z+jets processes tend primarily to populate only the same-flavor dilepton final states, with a smaller component in the different-flavor dilepton final state due to the leptonic τ decays in the $Z \to \tau\tau$ process. For this reason, we define two classes of SR: one requiring SF dilepton events and the other requiring DF dilepton events. The SR events are additionally required to have at least two b-tagged jets, with the invariant mass of the two leading b-tagged jets having an invariant mass consistent with the mass of the Higgs boson, $m_{bb} \in [110, 140]$ GeV. Taking advantage of the lepton correlations, illustrated by Figure 9.5, we further require that the dilepton invariant mass of the SR events be relatively low: $m_{\ell\ell} < 60$ GeV. This requirement on $m_{\ell\ell}$ effectively acts to remove the majority of background events arising from Z-boson processes.

The SR definitions further rely on placing a cut on the d_{hh} discriminant. The specific choice of cut, in both SR-SF and SR-DF, are determined by computing the expected 95% CL cross-section upper-limit with the selections described in the previous paragraph applied while scanning over varying selections made on the d_{hh} discriminant. The cut thresholds on the d_{hh} discriminant in both SR-SF and SR-DF that minimize the cross-section upper-limit value are chosen as the final values to define SR-SF and SR-DF, and are indicated in Table 9.5.

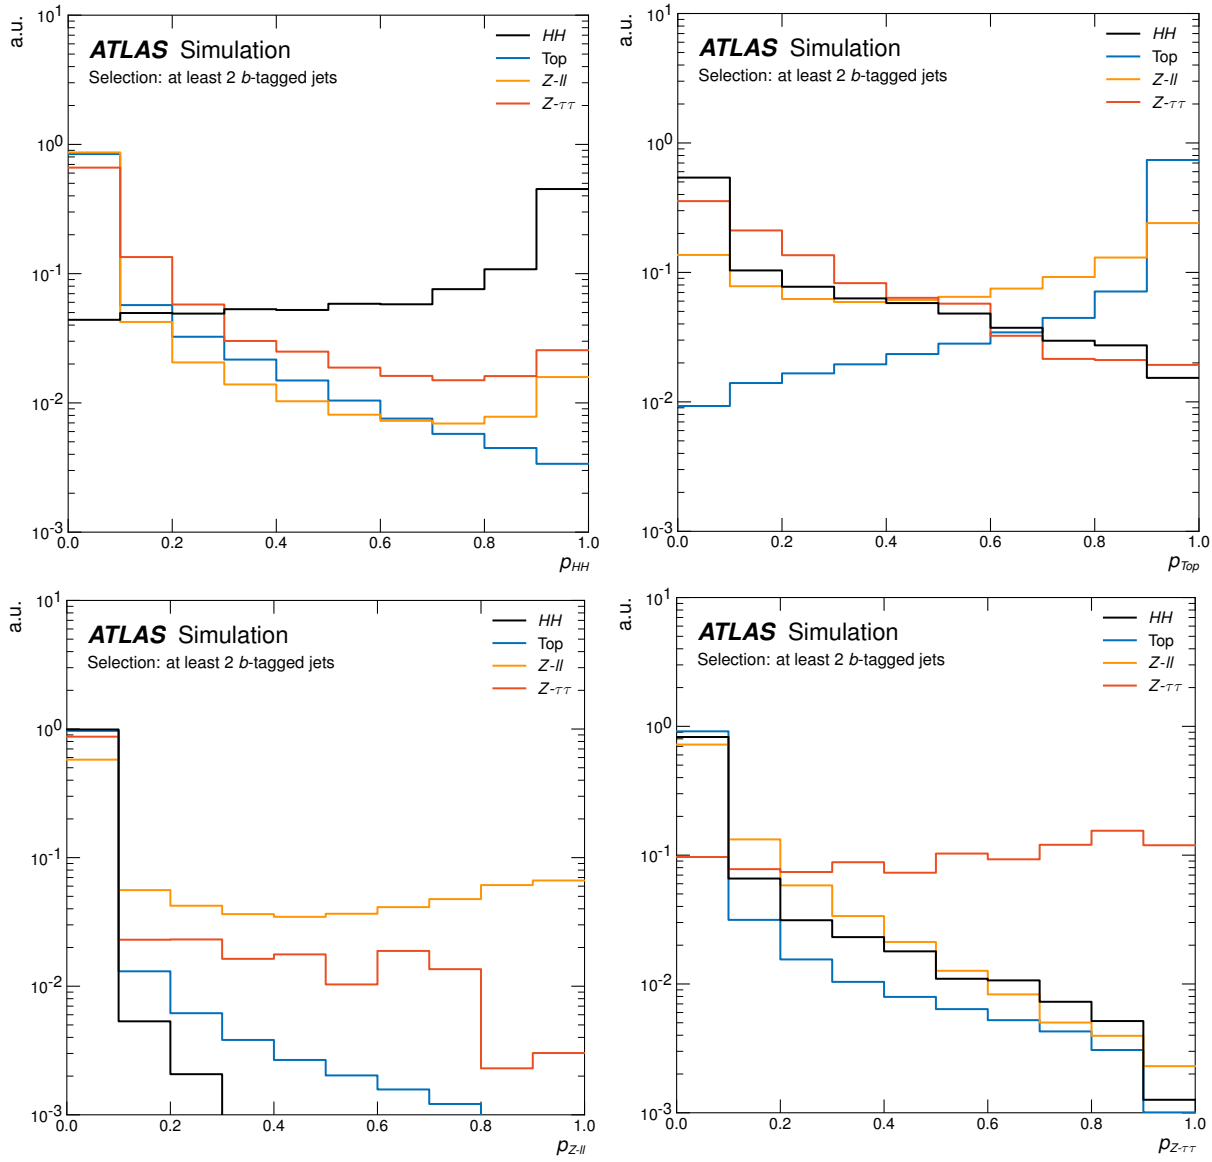

Figure 9.22: Normalized distributions of the four network outputs, shown for the dilepton $hh \to bbWW^*$ signal (black), Top (blue), $Z \to \{ee, \mu\mu\}$ (yellow), and $Z \to \tau\tau$ (orange) processes. From the top left and moving clock-wise: p_{hh}, p_{Top}, $p_{Z-\tau\tau}$, and $p_{Z-\text{SF}}$.

Table 9.5: SRs for the search for the dilepton $hh \to bbWW^*$ process.

	Region	
Observable	**SR-SF**	**SR-DF**
Dilepton Flavor	ee or $\mu\mu$	$e\mu$ or μe
b-tagged jet multiplicity	≥ 2	≥ 2
bb-system invariant mass, m_{bb} [GeV]	$\in [110, 140]$	$\in [110, 140]$
Dilepton invariant mass, $m_{\ell\ell}$ [GeV]	< 60	< 60
d_{hh}	> 5.45	> 5.55

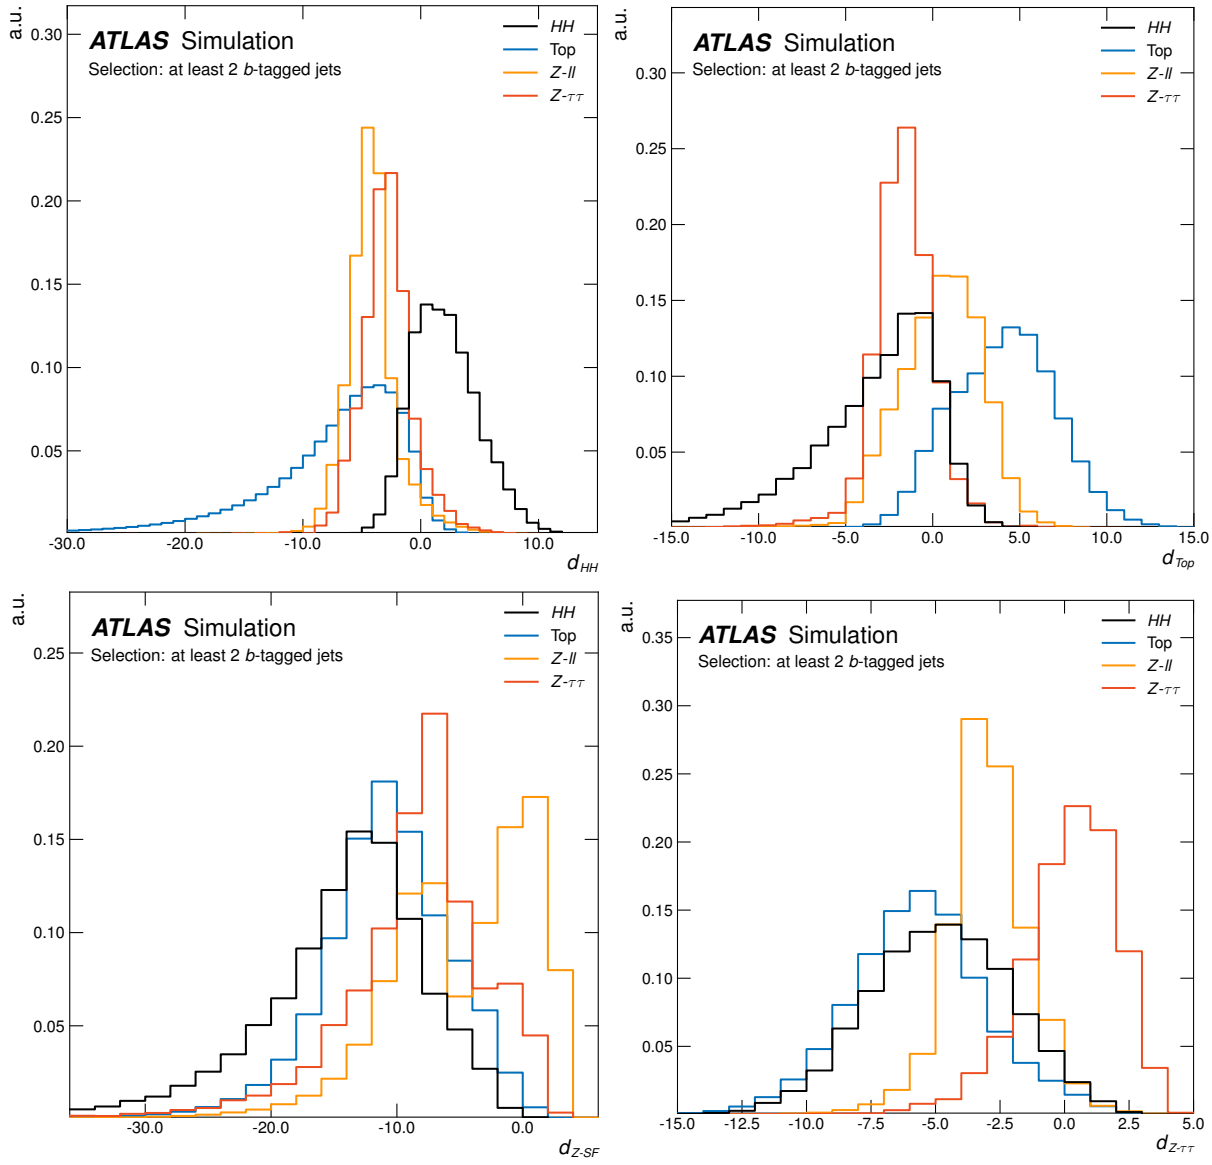

Figure 9.23: Normalized distributions of the four composite discriminants, shown for the dilepton $hh \to bbWW^*$ signal (black), Top (blue), $Z \to \{ee, \mu\mu\}$ (yellow), and $Z \to \tau\tau$ (orange) processes. From the top left and moving clock-wise: d_{hh}, d_{Top}, $d_{Z-\tau\tau}$, and $d_{Z-\text{SF}}$.

9.4 Estimation of the Standard Model Backgrounds

9.4.1 Top Quark Backgrounds

We define a CR for the SM background that composed of the sum of the $t\bar{t}$ and single-top Wt process, referred to simply as the Top background. In reality, already at NLO, these two processes become difficult to distinguish between one another due to their leading both to the $WbWb$ final state with interfering diagrams. These interference effects are not properly accounted for in the MC simulation. This ambiguity in their separate definition is the motivating reason for combining them into a single process 'Top' and treating their sum as a single process, since this is how they behave in data. Towards this aim, we define CR-Top to provide a data-driven normalization correction on the combined Top background. The aim of CR-Top is to provide a correction to the Top background that is representative of that in the SRs defined in Table 9.5. For this to be the case, CR-Top is defined so that the Top background composition is very similar to that found in the SRs. That is, the relative proportion of the simulated $t\bar{t}$ and single-top Wt in the Top background appearing in SR-SF and SR-DF is the same as that appearing in the Top background in CR-Top. In this way, the normalisation correction will correct for any interference effects not properly accounted for in the MC simulation.

The definition of the associated CR and VR, CR-Top and VR-Top, are given in Table 9.6. As with CR-Top, VR-Top is defined to have similar Top background composition as the SRs. CR-Top and VR-Top select DF and SF events, respectively. Given that the Top background is flavor symmetric, this is an appropriate extrapolation to make. Both CR-Top and VR-Top are defined to be orthogonal to the SR selections by inverting the Higgs mass requirement defined by the m_{bb} observable. For CR-Top, the m_{bb} requirements are inverted entirely, while for VR-Top only the upper portion of the m_{bb} distribution is selected since, in the SF dilepton events, the lower tail of m_{bb} is contaminated by Z-boson processes. The d_{hh} requirements, relative to those in the SRs, are relaxed so as to increase the available statistics.

Figures 9.24-9.26 show distributions of the relevant observables in CR-Top. Figure 9.27 shows the distribution of the composite discriminant d_{hh} in CR-Top without the requirement on d_{hh} applied, so that the full shape of the discriminant may be studied.

Table 9.6: Definitions of the CR and VR for the Top quark backgrounds ($t\bar{t}$ and single-top Wt) for the search targeting the dilepton $hh \to bbWW^*$ process.

	Region	
Observable	**CR-Top**	**VR-Top**
Dilepton Flavor	$e\mu$ or μe	ee or $\mu\mu$
b-tagged jet multiplicity	≥ 2	≥ 2
m_{bb} [GeV]	< 100 or > 140	> 140
$m_{\ell\ell}$ [GeV]	< 60	< 60
d_{hh}	> 4.5	> 4.5

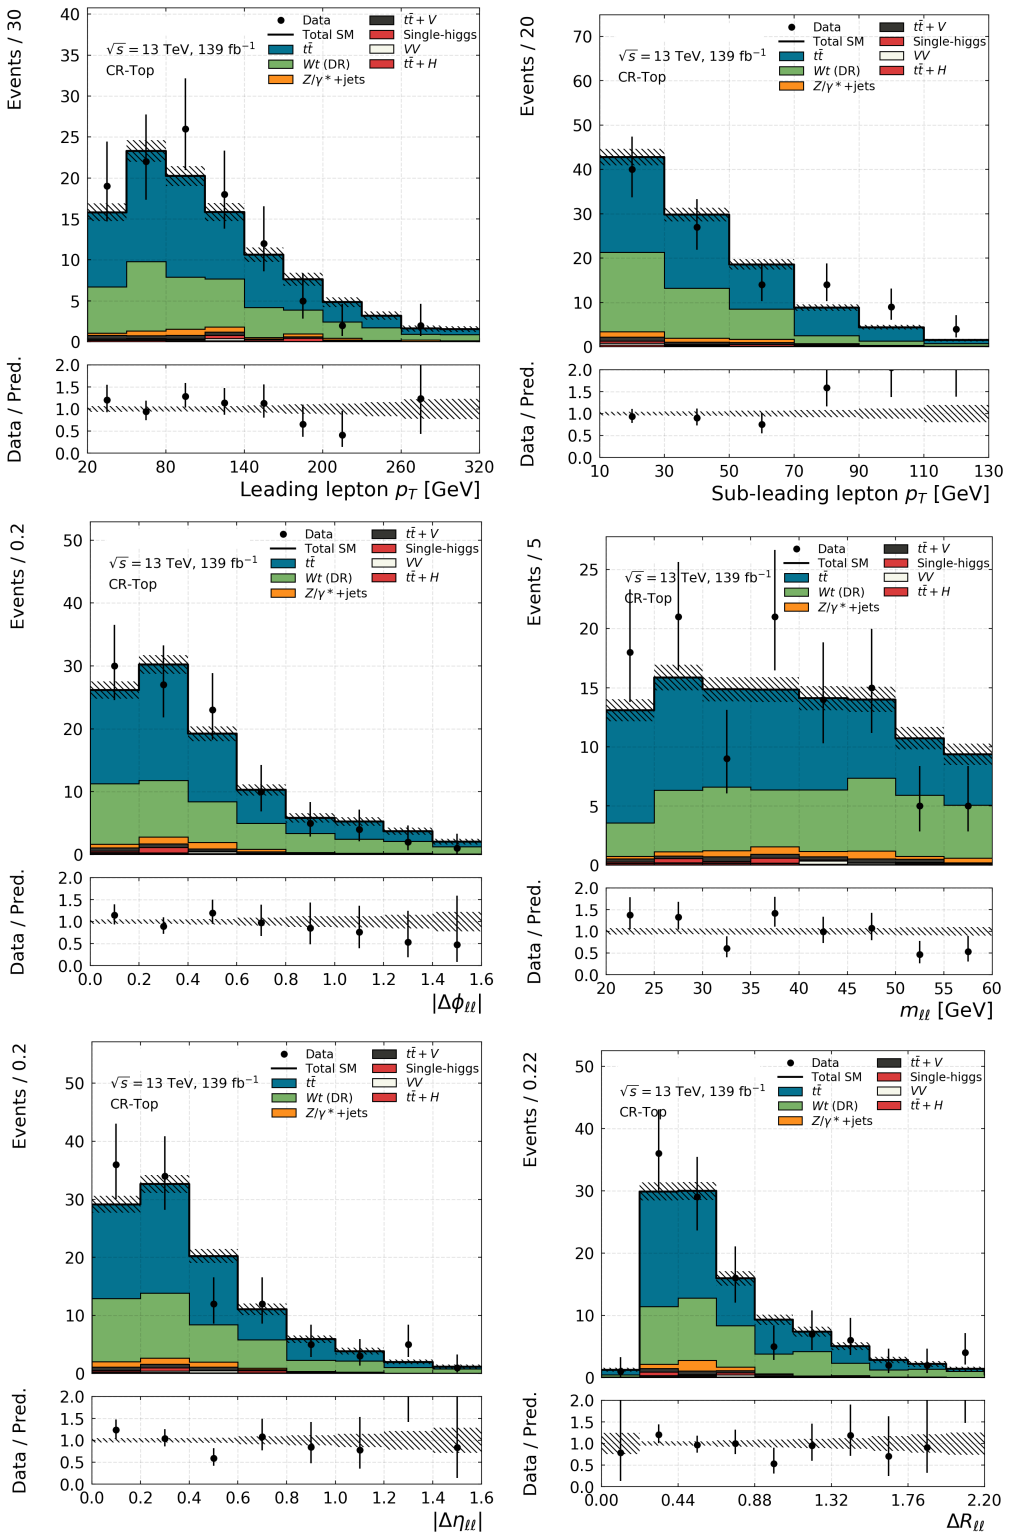

Figure 9.24: Kinematic distributions in the Top control region, CR-Top. The error bands include only the statistical uncertainty. The normalization factors obtained from the background-only fit (Table 9.13) are applied to the Top ($t\bar{t}$ and Wt) and Z+jets MC processes.

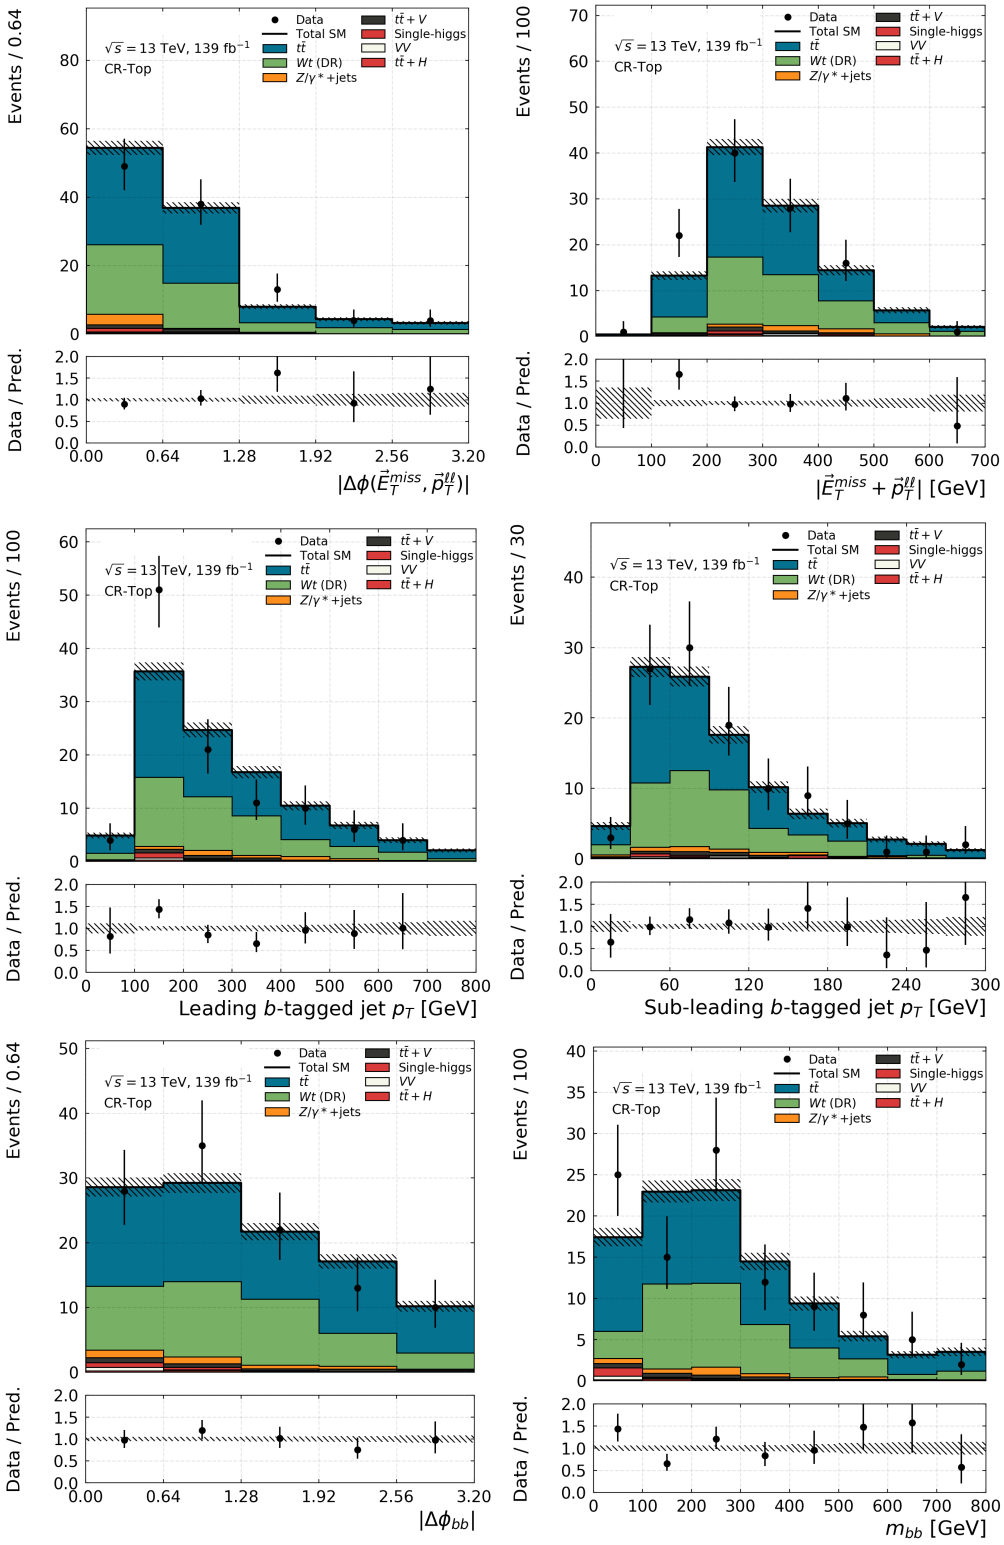

Figure 9.25: Kinematic distributions in the Top control region, CR-Top. The error bands include only the statistical uncertainty. The normalization factors obtained from the background-only fit (Table 9.13) are applied to the Top ($t\bar{t}$ and Wt) and Z+jets MC processes.

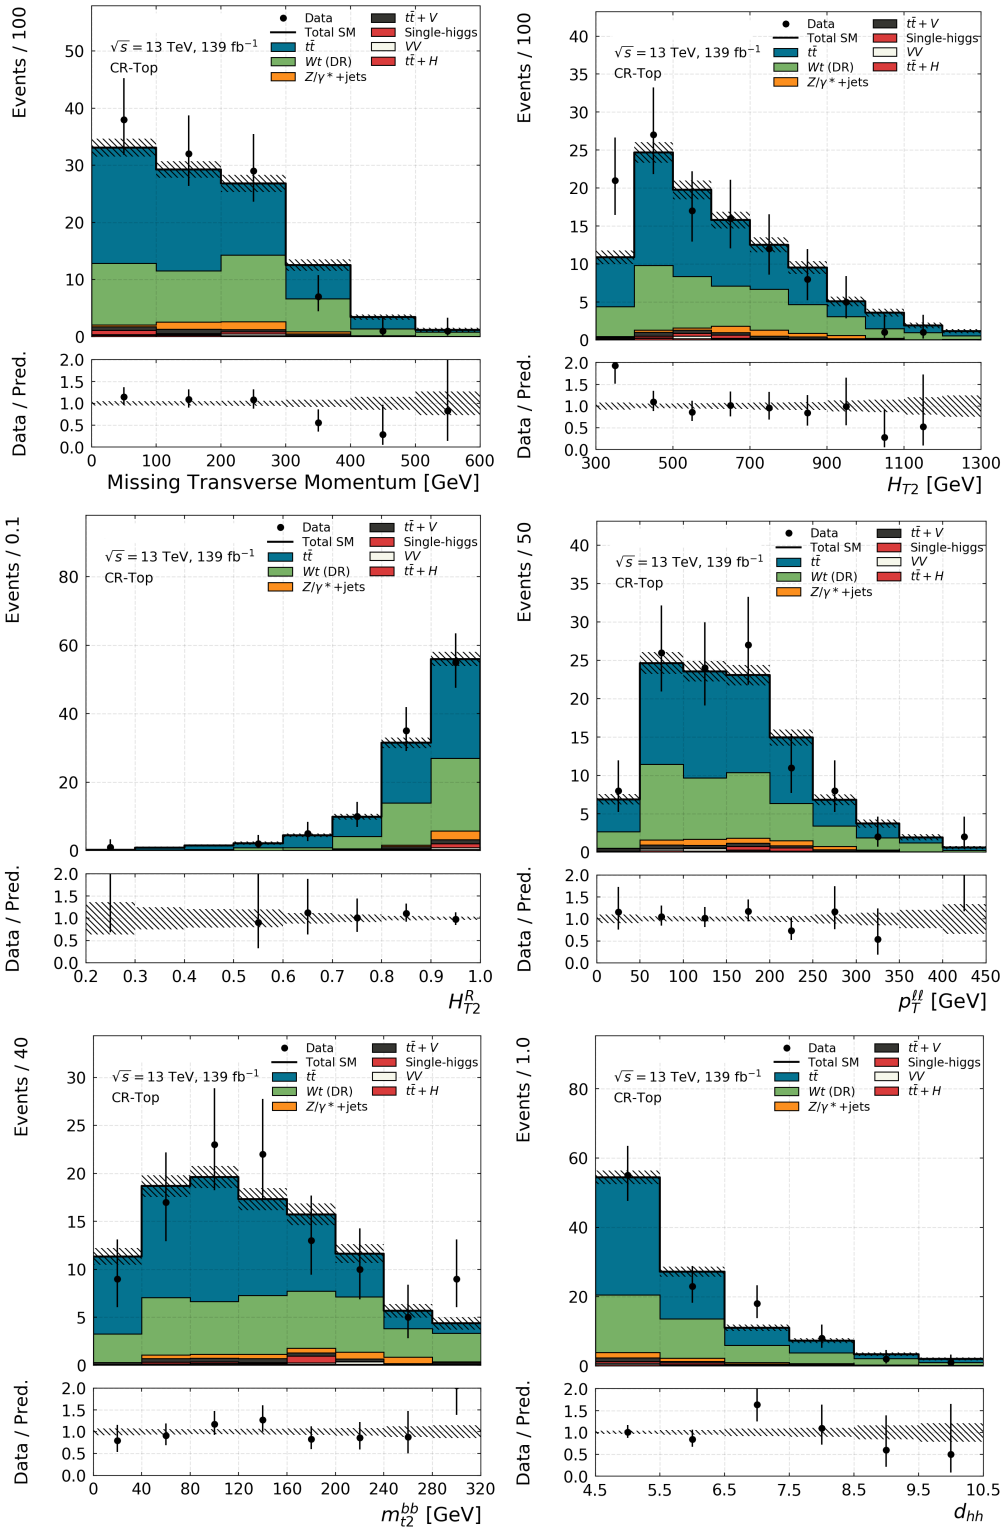

Figure 9.26: Kinematic distributions in the Top control region, CR-Top. The error bands include only the statistical uncertainty. The normalization factors obtained from the background-only fit (Table 9.13) are applied to the Top ($t\bar{t}$ and Wt) and Z+jets MC processes.

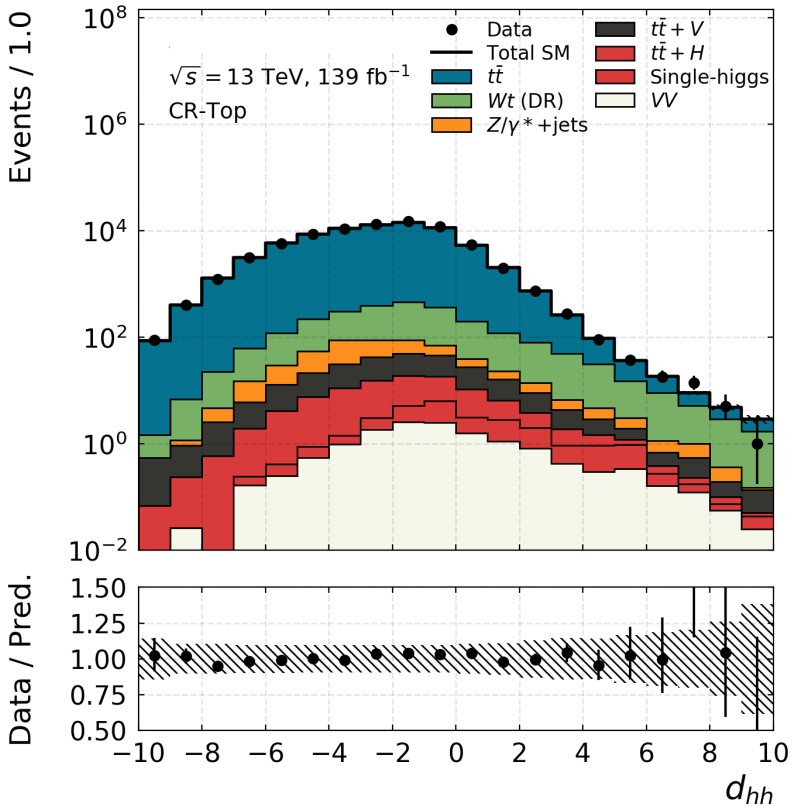

Figure 9.27: d_{hh} distribution in the Top control region, CR-Top, without the d_{hh} selection applied. The error band in the ratio includes both statistical and systematic uncertainty while the error band in the histogram includes only statistically uncertainty. The normalization factors obtained from the background-only fit (Table 9.13) are applied to the Top ($t\bar{t}$ and Wt) and Z+jets MC processes.

9.4.2 Z Boson Production in Association with Heavy Flavor Jets

In the ≥ 2 b-tagged jet regions of phase space that are being probed by SR-SF and SR-DF (c.f. Table 9.5), there remains non-negligible contamination from Z+jets processes. As the cross-section of Z-boson production in association with heavy flavor jets is known to be mismodelled by MC predictions [224, 225], especially in the relatively collinear b-tagged jet regimes sensitive to the theoretically challenging $g \to bb$ processes that $h \to bb$ analyses are probing, we define a CR and VR enriched in the Z+heavy-flavor processes in order to constrain the MC prediction with data.

We define the Z+heavy-flavor process as those Z-boson processes whose leading two b-tagged jets are identified as arising either from b- or c-hadrons, using the MC truth-level information to identify the jets as such. The corresponding definitions of the Z+heavy-flavor and Z+light-flavor processes are given in Table 9.7.

Table 9.7: Definition of the Z+heavy-flavor and Z+light-flavor backgrounds. Shown are the truth-level association of the leading two reconstruction level b-tagged jets (in MC) to b- or c-hadrons or light-flavor processes (mis-tags), indicated by ℓ.

Process	**Truth-level b-tagged Jet Pair Origin**
Z+heavy-flavor	bb, bc, or cc
Z+light-flavor	$b\ell$, $c\ell$ or $\ell\ell$

The Z+heavy-flavor regions are inclusive in dilepton flavor and are defined by inverting the $m_{\ell\ell} < 60$ GeV requirement made in the SRs. The Z+heavy-flavor CR, CR-Z+HF, selects events on the Z-boson mass pole with a relatively tight window in $m_{\ell\ell}$ centered at $m_Z = 91.2\,\text{GeV}$. The Z+heavy-flavor VR, VR-Z+HF, selects events in side-bands of $m_{\ell\ell}$ around the window selected by CR-Z+HF. Both CR-Z+HF and VR-Z+HF retain the requirements on the b-tagged jets that are applied in the SRs, so as to ensure that the b-tagged jet kinematics between these regions are similar. As with CR-Top and VR-Top, the requirement on d_{hh} is relaxed in order to increase the statistics available for these regions. Independent studies have been performed to study the dependence of the Z+heavy-flavor normalisation correction factor derived in CR-Z+HF on the d_{hh} selection, showing that there is no statistically significant dependence on the cut. The loose d_{hh} selection is therefore adequate and allows for improved statistical precision in the definition of the Z+heavy-flavor process' normalization correction.

Figures 9.28-9.30 show distributions of the relevant observables in CR-Top. Figure 9.31 shows the distribution of the composite discriminant d_{hh} in CR-Top without the requirement on d_{hh} applied, so that the full shape of the discriminant may be studied.

Table 9.8: Definitions of the CR and VR for the Z+heavy flavor process for the search targeting the dilepton $hh \to bbWW^*$ process.

	Region	
Observable	**CR-Z+HF**	**VR-Z+HF**
Dilepton Flavor	$ee, \mu\mu, e\mu$, or μe	ee or $\mu\mu$
b-tagged jet multiplicity	≥ 2	≥ 2
m_{bb} [GeV]	$\in [110, 140]$	$\in [110, 140]$
$m_{\ell\ell}$ [GeV]	$\in [81.2, 101.2]$	$\in [71.2, 81.2]$ or $\in [101.2, 115]$
d_{hh}	> 0	> 0

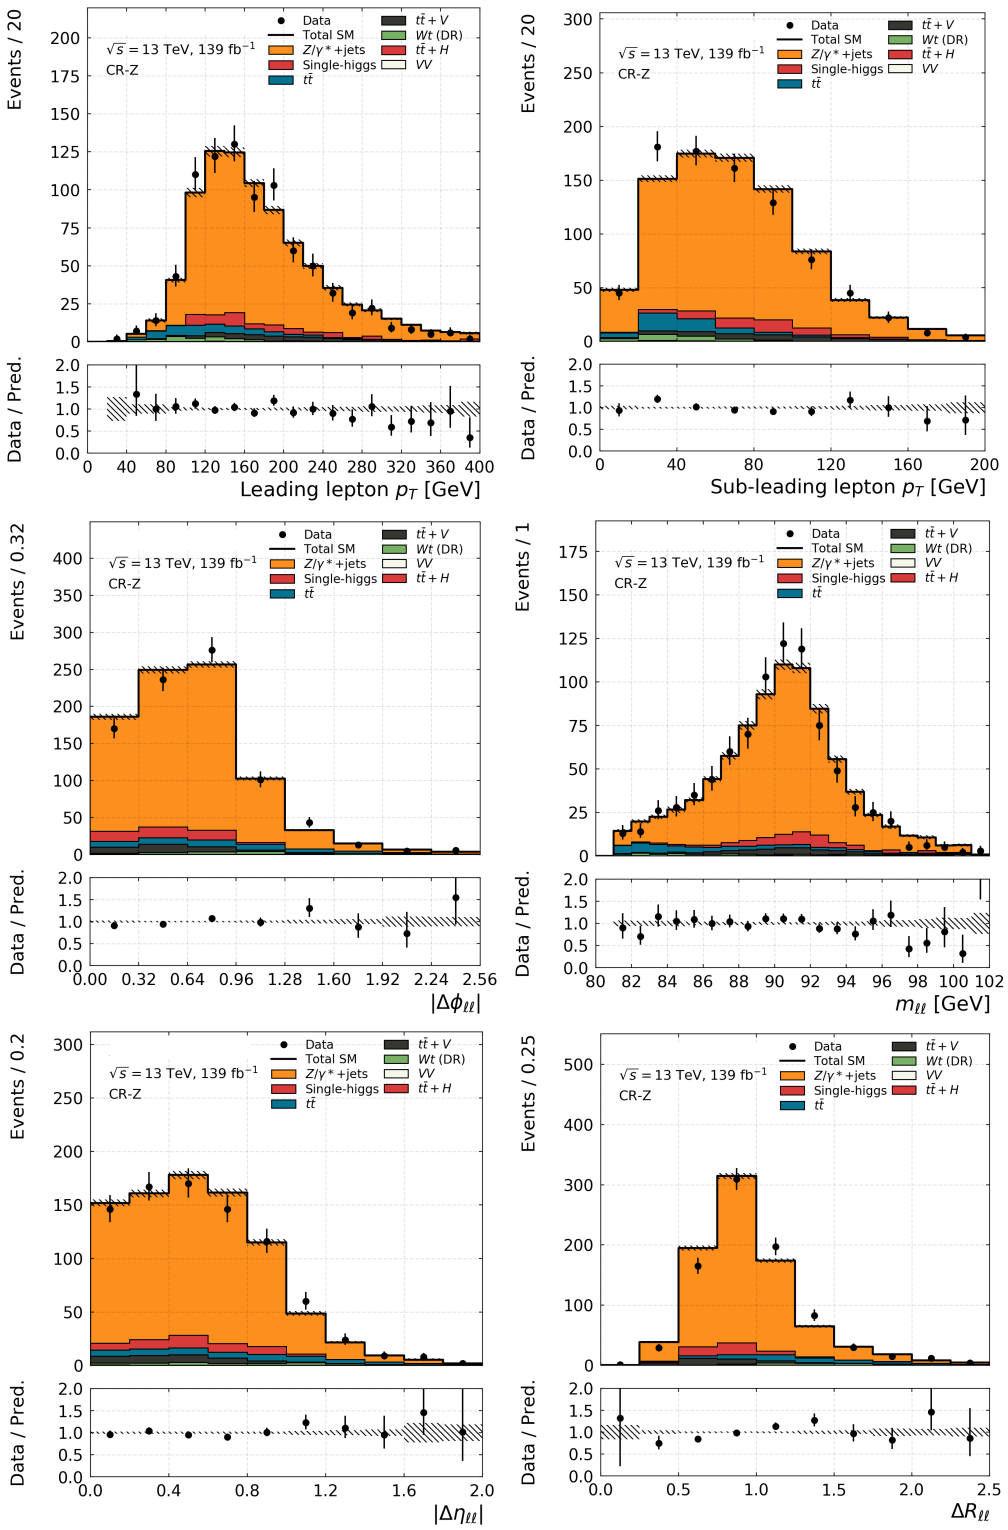

Figure 9.28: Kinematic distributions in the Z+heavy flavor control region, CR-Z+HF. The error bands include only the statistical uncertainty. The normalization factors obtained from the background-only fit (Table 9.13) are applied to the Top ($t\bar{t}$ and Wt) and Z+jets MC processes.

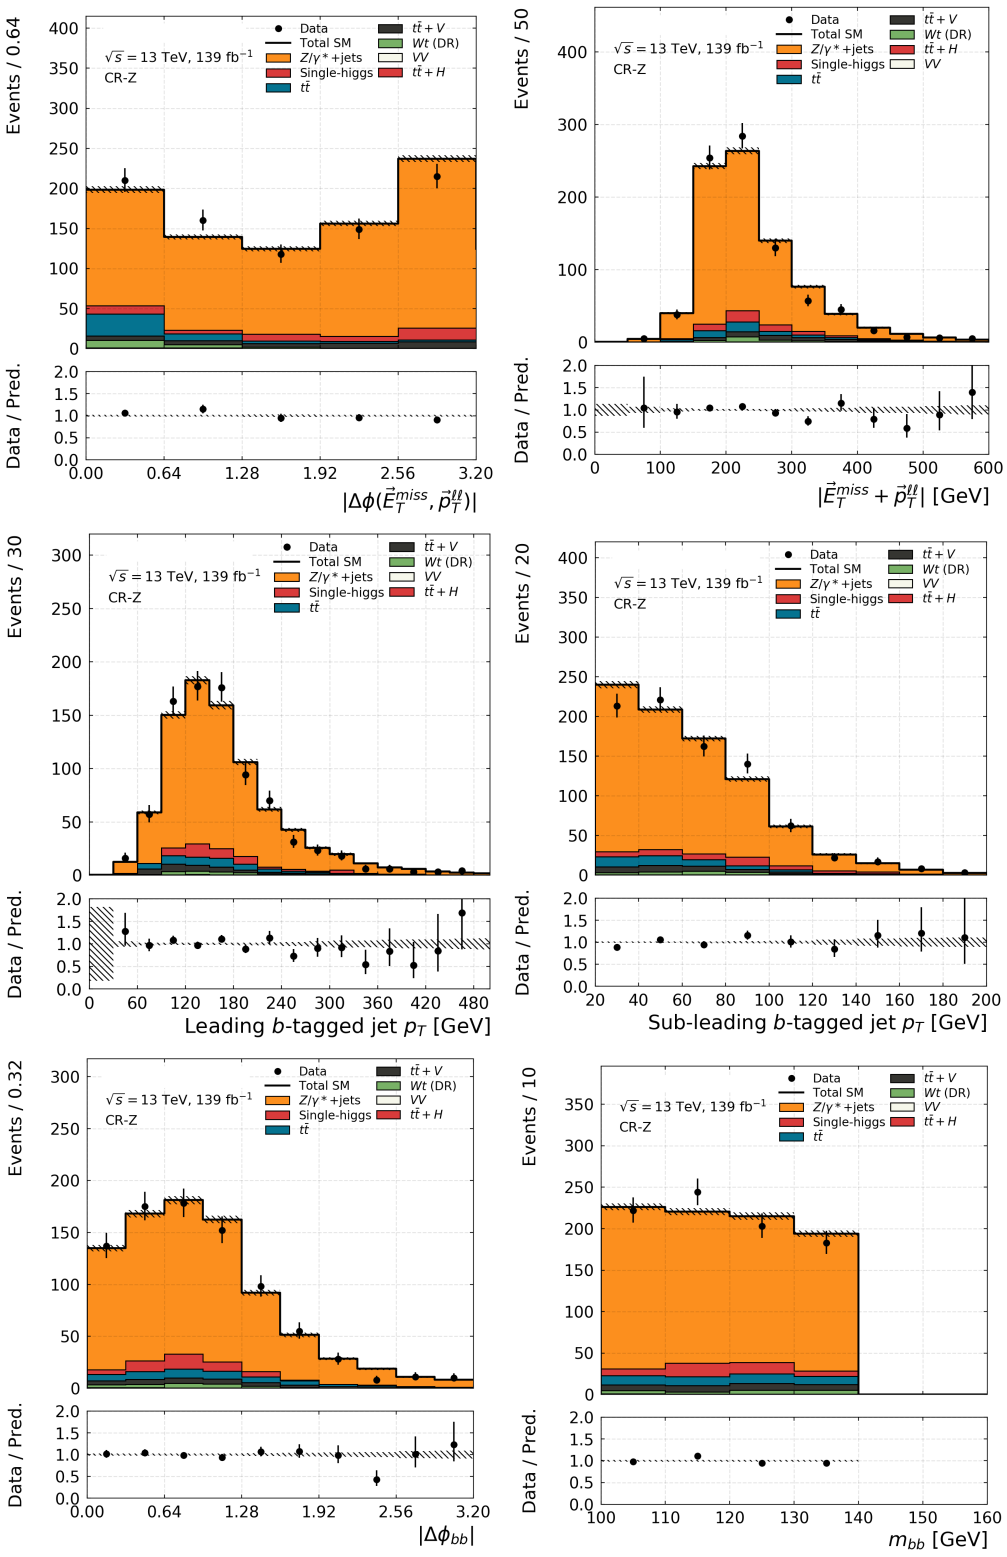

Figure 9.29: Kinematic distributions in the Z+heavy flavor control region, CR-Z+HF. The error bands include only the statistical uncertainty. The normalization factors obtained from the background-only fit (Table 9.13) are applied to the Top ($t\bar{t}$ and Wt) and Z+jets MC processes.

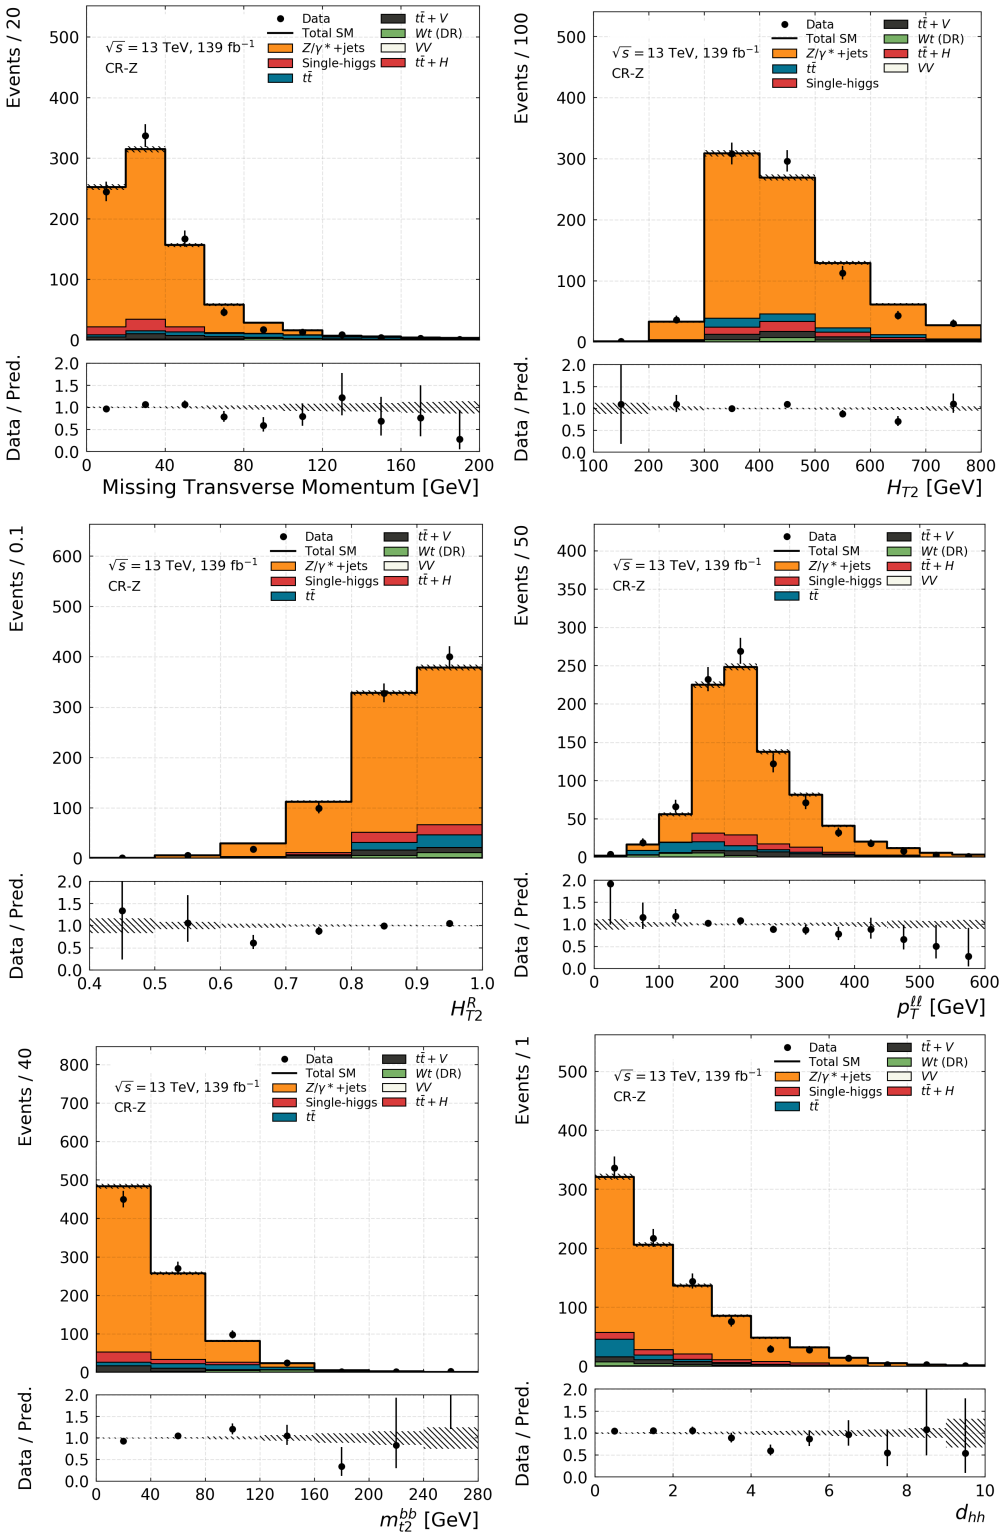

Figure 9.30: Kinematic distributions in the Z+heavy flavor control region, CR-Z+HF. The error bands include only the statistical uncertainty. The normalization factors obtained from the background-only fit (Table 9.13) are applied to the Top ($t\bar{t}$ and Wt) and Z+jets MC processes.

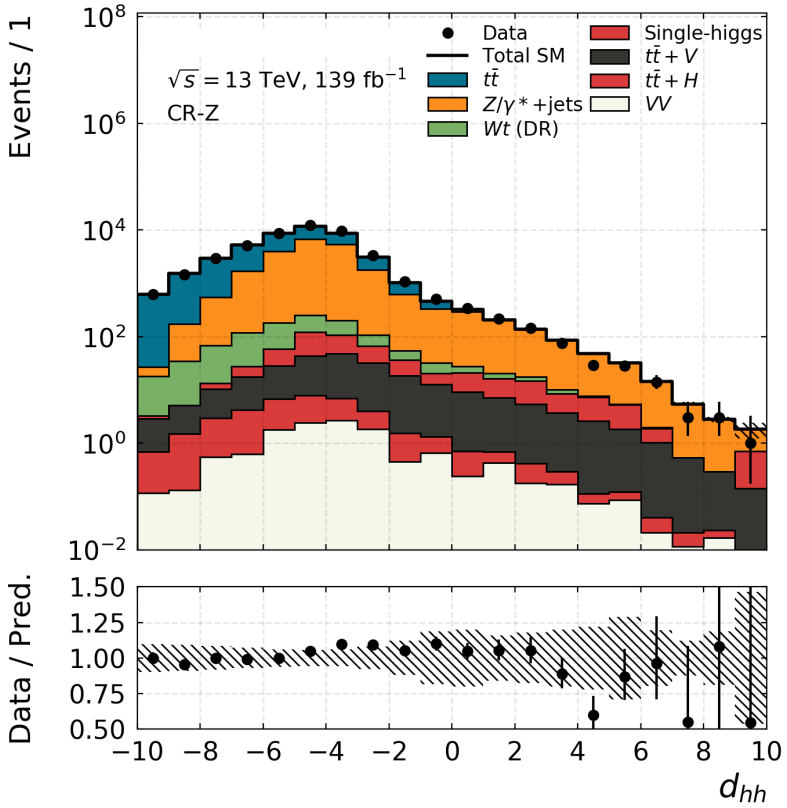

Figure 9.31: d_{hh} distribution in the Z+HF CR without the d_{hh} selection of the Z+HF CR applied. The error band in the ratio include both statistical and systematic uncertainties while the error band in the histogram includes statistical uncertainty only. The normalization factors obtained from the background-only fit (Table 9.13) are applied to the Top ($t\bar{t}$ and Wt) and Z+jets MC processes.

9.4.3 Sources of Fake and Non-Prompt Leptons

In this section we describe the estimation of sources of fake and non-prompt leptons. The analysis uses the Same-sign Extrapolation Method, described in detail in Section 7.2. In order to perform the estimate of fake leptons in this manner, we define a set of fake lepton control regions which are defined similarly as the analysis CRs and SRs, but with the opposite-sign (OS) lepton requirement inverted, instead requiring same-sign (SS) lepton pairs. This is as described in Section 7.2. On top of this, for the present analysis, we further invert the analysis' d_{hh} requirements in the SS regions. This is done so that we may increase the SS sample statistics further than if we were to rely solely on using the SS sample with all other selections remaining the same as their OS counterparts. In this way, we extrapolate over the d_{hh} quantity in addition to the lepton charge requirement and Equation 7.9 is modified to include the factor $\varepsilon^{d_{hh}}$, as indicated in Equation 9.7.

$$\begin{aligned} N_{\mathrm{OS}}^{\mathrm{fake}} &= \varepsilon^{d_{hh}} \times f^{SS\to OS} \times N_{\mathrm{SS}}^{\mathrm{fake}} \\ &= \frac{N_{\mathrm{MC,dhh+}}^{\text{non-prompt}}}{N_{\mathrm{MC,dhh-}}^{\text{non-prompt}}} \times \frac{N_{\mathrm{MC,OS}}^{\mathrm{fake}}}{N_{\mathrm{MC,SS}}^{\mathrm{fake}}} \times (N_{\mathrm{data,SS}} - N_{\mathrm{MC,SS}}^{\mathrm{real}}), \end{aligned} \tag{9.7}$$

where $N_{\mathrm{MC,dhh+}}$ ($N_{\mathrm{MC,dhh-}}$) gives the numbers of events in the non-inverted (inverted) d_{hh} selections. Additional systematic uncertainties related to the measurement of the $\varepsilon^{d_{hh}}$ quantity are assessed on top of those already described in Section 7.3.

Using truth level information available in the MC, the fake and non-prompt sources of leptons are categorized either as being due to semi-leptonic decays of heavy-flavor hadrons (resulting in non-prompt electrons or muons), non-prompt electrons due to photon conversions, or grouped into a catch-all category not falling under any of the others already defined ('other').

Figures 9.32 and 9.33 show the data and MC distributions in the same-sign selections used for deriving the same-sign extrapolation for the lepton transverse momenta, dilepton invariant mass, and d_{hh} observables. The fake estimate provided entirely by MC is generally able to recover the observation in each region, providing confidence that the extrapolation factor $f^{SS\to OS}$, derived using the MC, is correct.

Figure 9.34 shows the shapes of the data-driven fakes based on the SS data sample compared to the MC non-prompt estimate in the SS and OS versions of the SR selections. It is seen that the data-driven fake sample, with the prompt MC events subtracted, has a compatible shape with the purely MC-based fake estimate for the observable of interest, d_{hh}. This provides further support in the extrapolation from the SS to OS events as well as in the extrapolation over the d_{hh} observable.

For each region appearing in the analysis, a dedicated measurement of the quantities appearing in

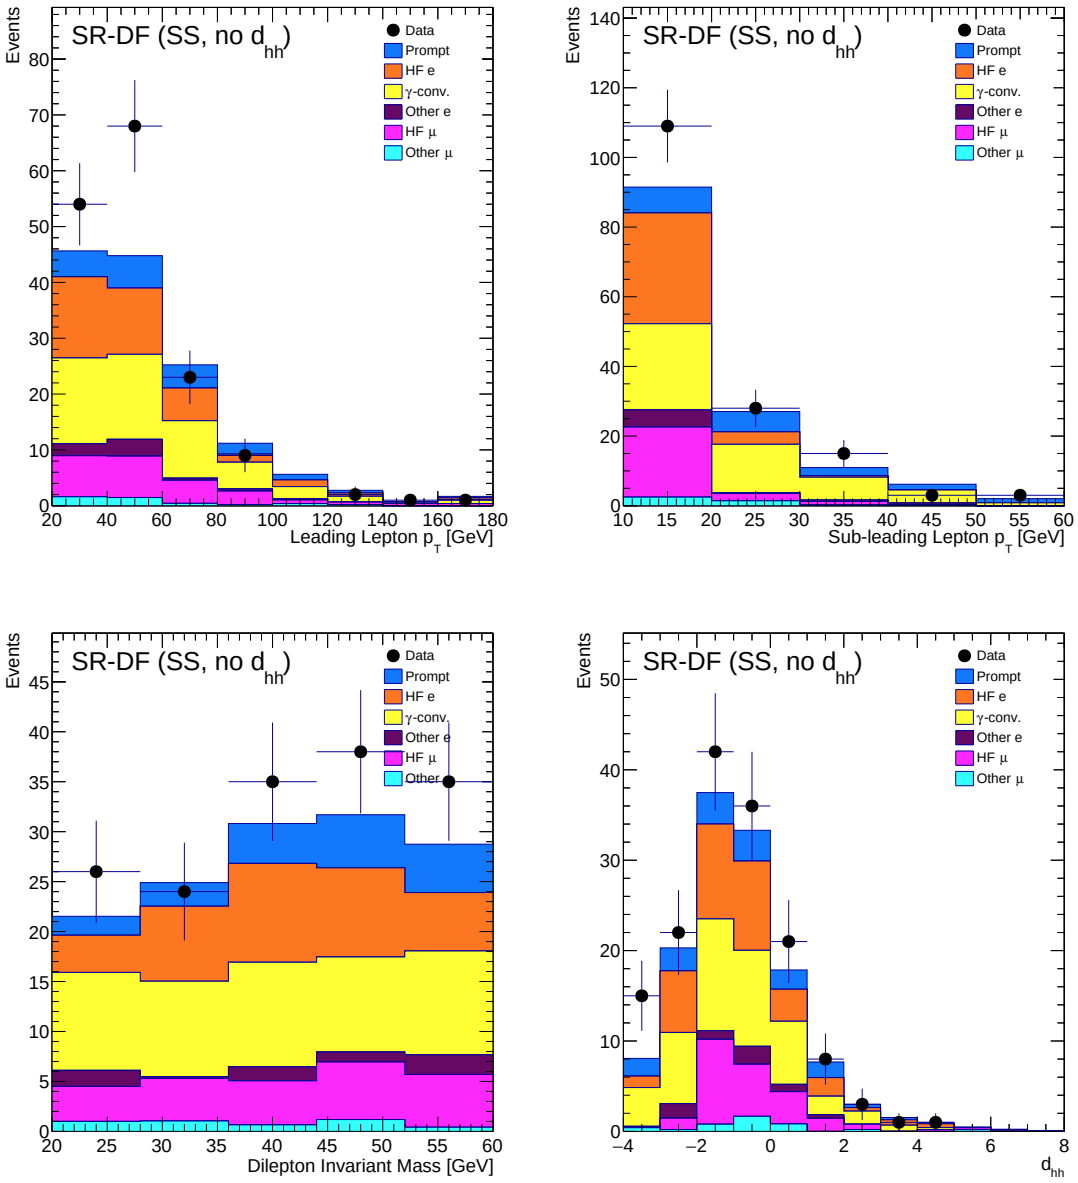

Figure 9.32: Data and MC distributions in the same-sign SR-DF selection with the d_{hh} requirement removed. The non-prompt MC is broken into the categories described in the text. ***top-left***: Leading lepton p_T, ***top-right***: Sub-leading lepton p_T, ***bottom-left***: dilepton invariant mass, and ***bottom-right***: d_{hh}.

Equation 9.7 is made, so that the fake estimate in every region can be obtained. The measured values of all quantities needed for Equation 9.7 are reported in Table 9.11. For each region, the $f^{SS\to OS}$ extrapolation factors are assessed for their dependence on d_{hh}, to check whether or not any additional higher order correction to account for the change in the fake composition as one extrapolates across the d_{hh} selection. This is shown in Figure 9.35 for SR-DF and SR-SF where

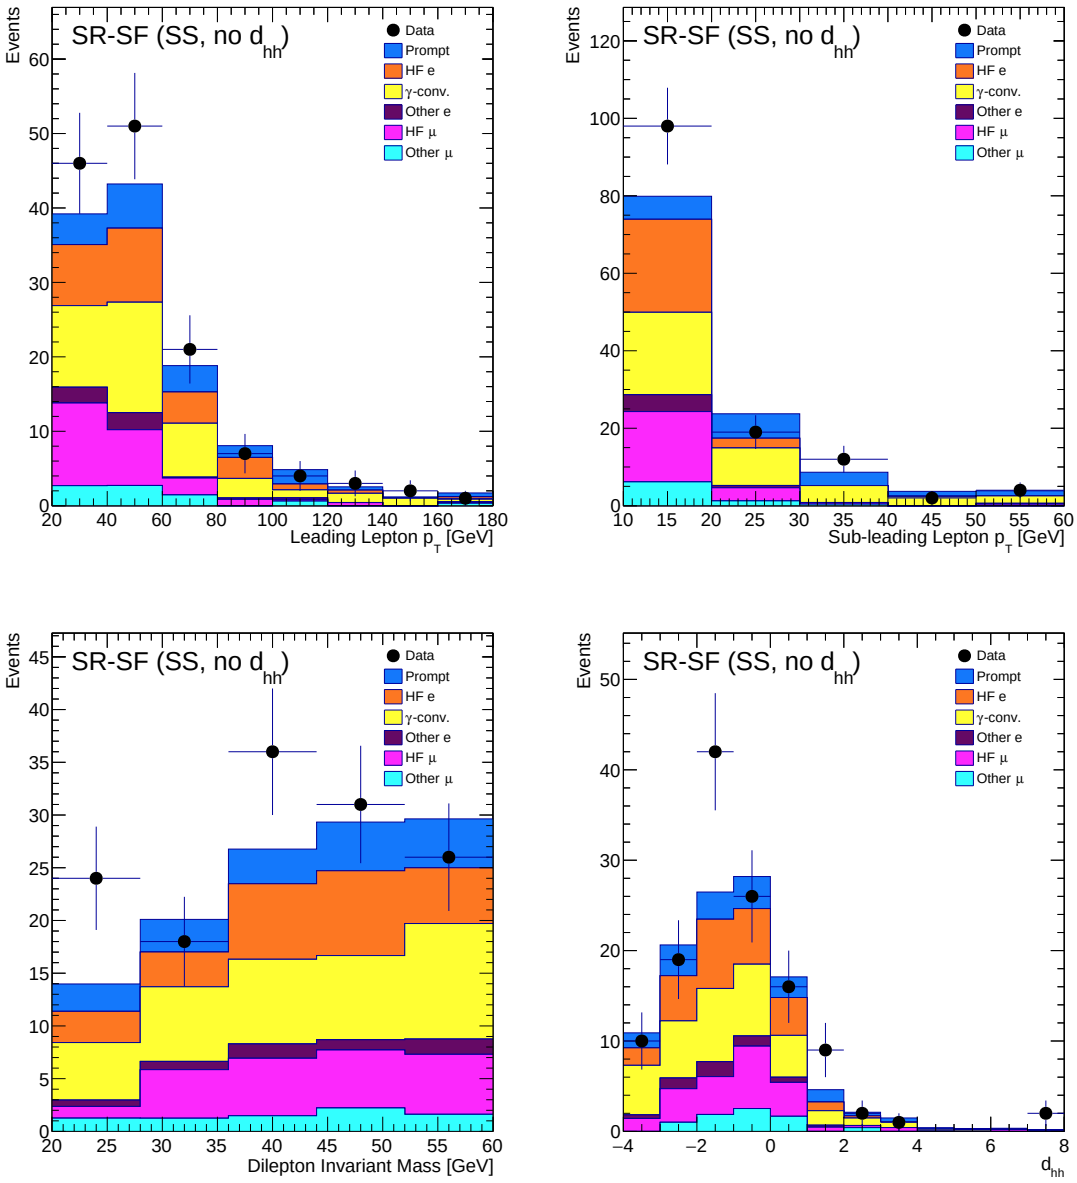

Figure 9.33: Data and MC distributions in the same-sign SR-SF selection with the d_{hh} requirement removed. The non-prompt MC is broken into the categories described in the text. ***top-left***: Leading lepton p_T, ***top-right***: Sub-leading lepton p_T, ***bottom-left***: dilepton invariant mass, and ***bottom-right***: d_{hh}.

it can be seen that there is no statistically significant dependence of the $f^{SS\to OS}$ quantity on d_{hh}. The same is found to hold for the CRs and VRs. Figure 9.35 also indicate the $\pm 20\%$ uncertainty taken on the $f^{SS\to OS}$ factors, defined as the typical spread of the $f^{SS\to OS}$ factors within each of the regions in the analysis.

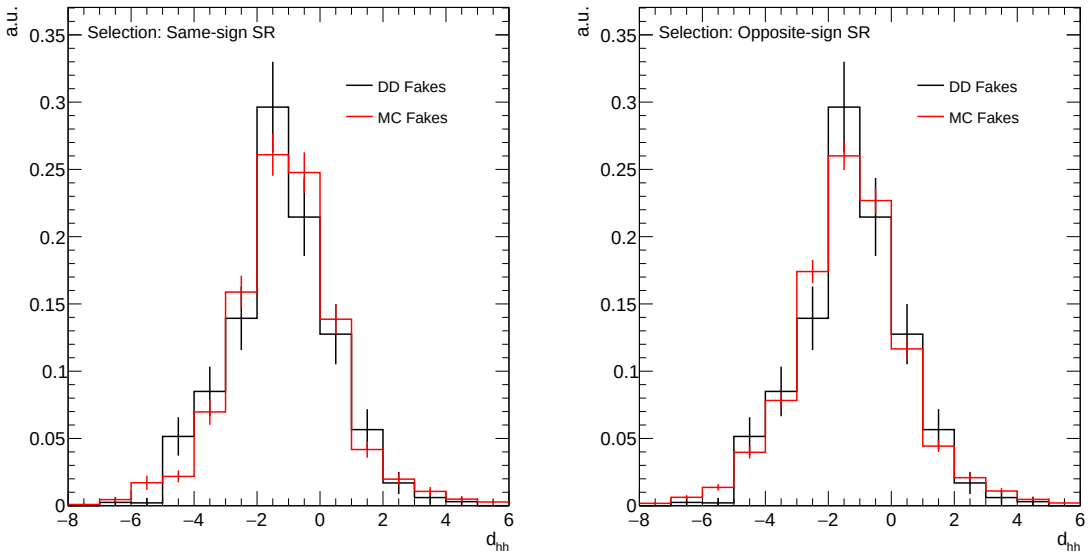

Figure 9.34: Normalized d_{hh} distributions of the same-sign data with the real SM processes subtracted using MC ('DD Fakes') compared to the fake prediction using MC ('MC Fakes') in the SR selections inclusive of dilepton flavor (SF+DF). ***Left***: Same-sign lepton charge requirement. ***Right***: Opposite-sign lepton charge requirement.

In Tables 9.9 and 9.10, the breakdown of the sources of fake leptons in the SS and OS versions of the analysis' signal regions, SR-DF and SR-SF, are presented along with their associated $f^{SS\to OS}$ factors. The single largest component of the fake background, in the OS regions, is due to conversion electrons. However, considering both electrons and muons together, fake leptons are predominantly arising from heavy flavor sources, such as in-flight semileptonic decays of b- and c-flavored hadrons.

Table 9.9: Fake lepton sources based on truth level information in MC for the SS and OS different-flavor signal region selections, with inverted d_{hh} selection. Uncertainties are statistical only. Shown also is the total MC fake, total MC prompt (non-fake), and Total MC (fake + non-fake) yields. The last row shows the observed number of events in each region. The right most column shows the ratio of the opposite-sign to same-sign yields for each of the fake sources.

Source of Fake	**SR-DF (SS, inverted d_{hh})**	**SR-DF (OS, inverted d_{hh})**	**OS/SS**
Heavy flavor e	35.70 ± 2.67	59.04 ± 3.45	1.65 ± 0.16
Conversion e	49.72 ± 3.16	98.65 ± 4.43	1.98 ± 0.15
Other e	6.16 ± 1.11	25.90 ± 2.31	4.20 ± 0.85
Heavy flavor μ	23.00 ± 2.46	52.84 ± 3.27	2.29 ± 0.28
Other μ	4.10 ± 0.91	13.77 ± 1.67	3.36 ± 0.85
Total MC Fake	118.67 ± 5.02	250.19 ± 7.10	2.10 ± 0.11
Total MC Prompt	18.28 ± 1.37	16784.73 ± 58.34	-
Total MC	136.95 ± 5.20	17034.92 ± 58.77	-
Data	158	16720	-

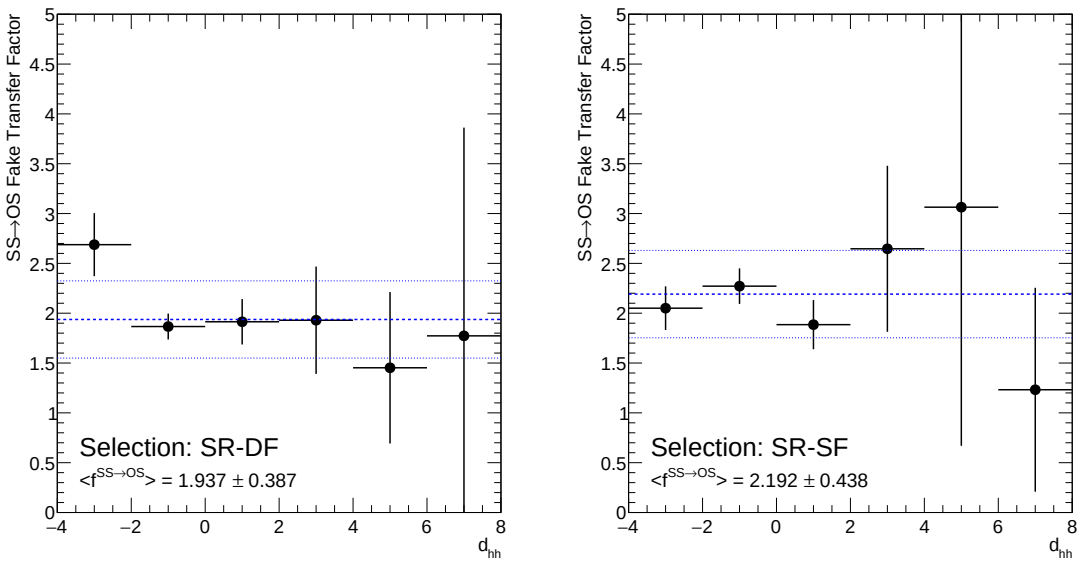

Figure 9.35: The $f^{SS\to OS}$ factors for SR-DF (***left***) and SR-SF (***right***) as a function of d_{hh}. The middle blue dashed lines are the average $f^{SS\to OS}$ across the d_{hh} range and the outer blue lines are the $\pm 20\%$ shift about this average value.

Table 9.10: Fake lepton sources based on truth level information in MC for the SS and OS same-flavor signal region selections, with inverted d_{hh} selection. Uncertainties are statistical only. Shown also is the total MC fake, total MC prompt (non-fake), and Total MC (fake + non-fake) yields. The last row shows the observed number of events in each region. The right most column shows the ratio of the opposite-sign to same-sign yields for each of the fake sources.

Source of Fake	**SR-SF (SS, inverted d_{hh})**	**SR-SF (OS, inverted d_{hh})**	**OS/SS**
Heavy flavor e	26.75 ± 2.30	44.37 ± 2.97	1.66 ± 0.18
Conversion e	39.11 ± 2.88	70.57 ± 3.81	1.80 ± 0.16
Other e	5.16 ± 1.00	16.04 ± 1.75	3.11 ± 0.70
Heavy flavor μ	22.15 ± 2.26	66.12 ± 3.70	2.99 ± 0.35
Other μ	7.82 ± 1.31	22.75 ± 2.17	2.91 ± 0.56
Total MC Fake	101.00 ± 4.62	219.87 ± 6.70	2.18 ± 0.12
Total MC Prompt	18.00 ± 1.28	17903.42 ± 64.27	-
Total MC	119.00 ± 4.80	18123.29 ± 64.62	-
Data	133	17773	-

Table 9.11: A summary of the quantities used in the determination of the fake contribution to the analysis regions and used as input to Equation 9.7. The right-most column gives the actual fake estimate in each of the analysis regions. The uncertainty on $N^{\text{DD fake}}$ includes the effects of the $\pm 20\%$ uncertainty on $f^{SS \to OS}$, $\pm 50\%$ uncertainty taken on the prompt MC subtraction, and the uncertainty on ε^{dhh}.

Region	$N_{\textbf{SS,data}}$	$N_{\textbf{SS,prompt MC}}$	ε^{dhh}	$f^{SS \to OS}$	$N^{\textbf{DD fake}}$
CR-Top	799	104.50 ± 3.13	0.0058 ± 0.0010	1.93 ± 0.39	7.79 ± 3.34
VR-Top	461	54.49 ± 1.89	0.0066 ± 0.0015	1.83 ± 0.37	4.91 ± 2.14
CR-Z+HF	137	40.66 ± 3.73	0.0347 ± 0.0044	2.73 ± 0.55	9.12 ± 5.47
VR-Z+HF	110	29.31 ± 2.55	0.0324 ± 0.0054	2.28 ± 0.46	5.96 ± 3.09
SR-SF	133	18.00 ± 1.28	0.0022 ± 0.0013	2.10 ± 0.42	0.55 ± 0.39
SR-DF	158	18.28 ± 1.37	0.0014 ± 0.0010	2.17 ± 0.43	0.42 ± 0.36

9.4.4 Background-only Fit

As in Section 8.5.3, we perform a background-only fit to the observed data in the CRs defined in the present analysis. The results are shown in terms of their impact on the post-fit predicted yields in the CRs and VRs in Table 9.12. The agreement observed between the post-fit MC and the observed data in the VRs shows that the extrapolation is performing well. The normalization correction factors derived in the fit for the Top and Z+heavy-flavor processes, μ_{Top} and $\mu_{Z+\text{heavy-flavor}}$, respectively, are reported in Table 9.13. The value for μ_{Top} is below one and indicates that the data prefers the combined $t\bar{t}$ and single-top Wt background estimate to be reduced, likely due to destructive interference effects that are not taken into account in the MC prediction. The rather large value of $\mu_{Z+\text{heavy-flavor}}$ is consistent with the measurements reported by both ATLAS and CMS in Refs. [224, 225], indicating that the MC predictions of the Z+heavy-flavor process underpredict the cross-section in kinematic regions similar to those being probed in the present analysis. The value for $\mu_{Z+\text{heavy-flavor}}$ also happens to consistent with that found by the ATLAS search for the $hh \to bb\tau\tau$ [204], which also has a dedicated measurement of a Z+heavy-flavor normalization correction.

Table 9.12: Table showing the pre- and post-fit yields for the contributing MC background processes in the Top and Z+heavy flavor control and validation regions. The Z+heavy-flavor and Z+light-flavor backgrounds are separated into the high-$m_{\ell\ell}$ ($m_{\ell\ell} > 60$ GeV) and low-$m_{\ell\ell}$ ($m_{\ell\ell} < 60$ GeV) components, with the latter labelled as 'Drell-Yan'. The uncertainties on the yield estimates include all uncertainties discussed in Section 7.3.

Region	CR-Top	CR-Z+HF	VR-Top	VR-Z+HF
Observed Data	108	852	171	157
Post-fit Total SM	107.99 ± 10.39	851.94 ± 29.19	161.75 ± 10.30	147.07 ± 10.54
Post-fit $t\bar{t}$	55.32 ± 6.86	40.57 ± 5.50	45.96 ± 6.00	59.32 ± 8.94
Post-fit Single-top Wt	36.34 ± 4.69	14.63 ± 2.13	31.24 ± 4.30	11.58 ± 1.80
Post-fit $t\bar{t} + V$	2.22 ± 0.23	27.60 ± 2.41	4.49 ± 0.42	3.20 ± 0.28
Post-fit Z+heavy-flavor	3.22 ± 0.47	686.48 ± 32.95	26.76 ± 1.59	60.39 ± 3.56
Post-fit Z+light-flavor	0.23 ± 0.08	27.10 ± 8.91	1.37 ± 0.51	2.49 ± 0.77
Post-fit Diboson	0.72 ± 0.12	0.75 ± 0.45	0.27 ± 0.07	0.20 ± 0.05
Post-fit Single-higgs	2.23 ± 0.79	45.76 ± 3.60	0.81 ± 0.09	4.00 ± 0.80
Post-fit Drell-Yan heavy-flavor	0.00 ± 0.00	0.00 ± 0.00	43.57 ± 2.43	0.00 ± 0.00
Post-fit Drell-Yan heavy-flavor	0.00 ± 0.00	0.00 ± 0.00	2.42 ± 1.06	0.00 ± 0.00
Post-fit Fakes	7.71 ± 3.29	9.03 ± 5.38	4.86 ± 2.09	5.89 ± 3.06
Total SM	131.38 ± 6.68	685.45 ± 15.41	163.63 ± 6.25	149.89 ± 8.75
$t\bar{t}$	69.95 ± 4.45	51.30 ± 3.88	58.11 ± 3.79	75.00 ± 7.65
Single-top Wt	45.97 ± 2.71	18.51 ± 1.50	39.51 ± 2.60	14.65 ± 1.45
$t\bar{t} + V$	2.22 ± 0.23	27.59 ± 2.42	4.49 ± 0.42	3.20 ± 0.28
Z+heavy-flavor	2.37 ± 0.33	505.44 ± 5.00	19.70 ± 0.65	44.46 ± 1.51
Z+light-flavor	0.23 ± 0.08	27.10 ± 8.97	1.37 ± 0.52	2.49 ± 0.77
Diboson	0.72 ± 0.12	0.75 ± 0.45	0.27 ± 0.07	0.20 ± 0.05
Single-higgs	2.23 ± 0.79	45.74 ± 3.63	0.81 ± 0.09	4.00 ± 0.80
Drell-Yan heavy-flavor	0.00 ± 0.00	0.00 ± 0.00	32.08 ± 0.83	0.00 ± 0.00
Drell-Yan light-flavor	0.00 ± 0.00	0.00 ± 0.00	2.42 ± 1.07	0.00 ± 0.00
Fakes	7.71 ± 3.29	9.03 ± 5.38	4.86 ± 2.09	5.89 ± 3.06

Table 9.13: Normalization correction factors derived for the Top and Z+heavy-flavor background processes in the search targeting the dilepton $hh \to bbWW^*$ signal process.

μ_{Top}	0.79 ± 0.10
$\mu_{Z\text{+heavy-flavor}}$	1.35 ± 0.07

9.5 Results

As with the analysis presented in Chapter 8, a profile likelihood fit is performed, including all CRs and SRs of the analysis and with the observed data in all regions used as a constraint. The result of running this fit is shown in Table 9.14. The observed data and post-fit MC prediction agrees quite well in SR-SF. In SR-DF there is a slight excess observed in the data relative to the MC prediction. This excess in SR-DF is not statistically significant however, and amounts to only a 1.05σ excess (null p_0-value of 0.15). Figures 9.36-9.39 present kinematic distributions comparing the observed data and post-fit MC prediction for the SM backgrounds in the SR selections. Figures 9.37 and 9.38 show the d_{hh} distributions in detail in SR-DF and SR-SF, respectively.

As there is no significant deviation between the prediction and observed data, indicating that there is no evidence for enhanced production of Higgs boson pairs, 95% CL cross-section upper-limits are derived for the SM-like non-resonant Higgs boson pair production. These results are reported in Table 9.15 and are with respect to the $pp \to hh$ production process, having taken into account the branching ratio for the dilepton $hh \to bbWW^*$ decay. That is, the 95% CL UL reported in Table 9.15 are on σ_{hh}, and not on $\sigma_{hh} \times \text{BR}(hh \to bbWW^*) \times \text{BR}(WW^* \to \ell\nu\ell\nu)$. The effect of the 1.05σ excess in SR-DF is seen in the observed 95% CL UL, which in Table 9.15 is seen to be larger than the expected value by roughly 1σ. Figure 9.40 illustrates the results listed in Table 9.15, while also comparing them to the leading hh searches (c.f. Figure 9.2): $hh \to 4b$, $hh \to bb\tau\tau$, and $hh \to bb\gamma\gamma$. The comparisons with the other analyses are made only with respect to the expected 95% CL cross-section upper-limits. It should be remembered that the other analyses in Figure 9.2 are based only on the partial Run 2 dataset, collected in 2015–2016 and comprised of $36.2\,\text{fb}^{-1}$ of pp collision data. However, the analysis presented in this chapter is the first time that this search has been performed in ATLAS and it is already improving the prospects of searches for hh in the $bbWW^*$ channel by roughly a factor of 10 (c.f. Figure 9.2). One such improvement, based on performing a multi-binned shape analysis wherein the full *shape* of the d_{hh} distribution is used as input to the hypothesis tests — as opposed to simply *cutting* on the d_{hh} distribution and defining simple signal regions as in Table 9.5 — is expected to increase the sensitivity to the $hh \to bbWW^*$ signal process substantially.

One question that the searches for Higgs boson pair production should being thinking about, as they start preparing for the time when they begin to make observation of Higgs boson pairs, is

whether or not they are designed in such a way as to make meaningful statements about the Higgs self-coupling parameter, λ. As shown in Figure 1.7, the leading contributions to Higgs boson pair production at the LHC proceeds via two diagrams, only one of which is sensitive to λ. The final state kinematics resulting from Higgs bosons originating from the decay of the box-diagram process tend to be much harder due to the massive top-quark loop being directly coupled to the outgoing Higgs bosons. As a result, analyses are typically more sensitive to this process as the final state objects are more directly observable and more efficiently reconstructed. This is the case for the analysis presented in this chapter, a fact that is illustrated in Figure 9.41. The hh system invariant mass in the box-diagram decays tends to be larger than that of the triangle diagram. An indication, then, of whether an analysis is tuned to be sensitive to decays relevant for the measurement of λ is whether or not they are sensitive to the lower portion of the m_{hh} distribution. From Figure 9.41, it is clear that the present analysis is more sensitive to the decays from the box-diagram. Given the early stages of the hh search program in ATLAS, however, this is not currently an urgent issue. The preliminary aim for the hh search program has been to define exactly which channels will be effective at observing hh signal events, in general, and to design analyses that can optimally search for the many potential sources of *enhanced* Higgs boson pair production. As the searches become sensitive to the hh production process, as predicted in the SM, the next steps will be to understand how to design the analyses such that they maximally discriminate between the box-type and triangle-type kinematics so as to improve the overall prospects for the measurement of the Higgs self-coupling parameter.

Table 9.14: Observed and predicted yields in the SRs for the dilepton $hh \to bbWW^*$ search.

	Regions	
Process	**SR-SF**	**SR-DF**
Observed Data	16	9
Post-fit Total SM	14.88 ± 2.12	4.88 ± 1.24
Post-fit $t\bar{t}$	2.57 ± 0.97	1.74 ± 0.66
Post-fit Single-top Wt	2.27 ± 0.54	2.04 ± 0.64
Post-fit Z+heavy-flavor	7.76 ± 1.04	0.21 ± 0.05
SM	14.12 ± 2.05	5.83 ± 1.42
$t\bar{t}$	3.26 ± 1.17	2.20 ± 0.80
Single-top Wt	2.88 ± 0.59	2.58 ± 0.75
$t\bar{t} + V$	0.59 ± 0.07	0.07 ± 0.06
Z+heavy-flavor	5.71 ± 0.74	0.15 ± 0.04
Z+light-flavor	0.32 ± 0.14	0.00 ± 0.00
Single-higgs	0.72 ± 0.09	0.32 ± 0.06
Fakes	0.54 ± 0.38	0.42 ± 0.36

In addition to the 95% CL cross-section UL reported in Table 9.15, we also perform a scan over the Higgs self-coupling parameter, λ, employed in the dilepton $hh \to bbWW^*$ signal MC sample so that we may re-run the analysis for various non-SM hypotheses for the value of λ. The variations

Table 9.15: Expected and observed 95% CL upper-limits on the Standard Model, non-resonant hh production cross-section ($\sigma(pp \to hh)$) and on the ratio of the upper-limit on this value to the value predicted in the SM ($\sigma^{\mathrm{SM}} = 31.05$ fb). The observed limits are in the right most column. The $\pm 1\sigma$ and $\pm 2\sigma$ excursions from the median expected value incorporate the effects of the analysis' systematic uncertainties.

	-2σ	-1σ	**Median Expected**	$+1\sigma$	$+2\sigma$	**Observed**
95% CL UL $\sigma(pp \to hh)$ [pb]	0.465	0.641	0.920	1.363	1.990	1.290
$\sigma^{\mathrm{UL}}/\sigma^{\mathrm{SM}}$	14.97	20.66	29.63	43.89	64.02	41.54

in the Higgs self-coupling are parametrized by $\kappa_\lambda = \lambda/\lambda_{SM}$, and this value is scanned across the values $\kappa_\lambda \in [-20, 20]$. Changing the value of κ_λ changes the relative strength of the competing triangle and box diagrams that lead to hh production at LO at the LHC, with only the former sensitive to the Higgs self-coupling parameter. Variations in κ_λ can therefore lead to either the box diagram or the triangle diagram being the dominant source of Higgs boson pairs. This has visible consequences for the final state. The Higgs bosons from the box diagram tend to have harder transverse momenta as compared to those from the triangle diagram, for example. The Higgs bosons from predominantly box diagram production scenarios, then, lead to harder visible final state objects and a different acceptance to the hh signal as compared to the case in which the Higgs bosons arise predominantly from the triangle diagram. Further discussion about the effects of varying κ_λ on the Higgs bosons' kinematics can be found in Ref. [212]. It should be stressed, however, that the present analysis is in no way optimized for non-SM values of κ_λ. For each new value of this parameter we simply perform the hypothesis tests using the same sets of SRs and CRs described in the sections above.

The results of the κ_λ scan are shown in Figure 9.42. It can be seen that the anaysis has little constraining power on the Higgs self-coupling parameter, with the NLO theory prediction for alternative κ_λ hypotheses described in the figure caption. The lack of ability to more powerfully exclude non-SM values of κ_λ lies in the fact that the analysis is optimized for the SM case, in particular. This is illustrated in Figure 9.43, showing the change in d_{hh} as a function of κ_λ. It can be seen that large values of d_{hh} are observed only for those κ_λ values near the SM hypothesis with $\kappa_\lambda \approx 1$. As a result of the harsh selections (cuts) made on d_{hh} in SR-SF and SR-DF (c.f. Table 9.5), sensitivity is lost for essentially all signal hypotheses with non-SM values of κ_λ. Again, this is a result of the analysis having been optimized with only the SM hypothesis in mind. Future versions of this analysis will hopefully improve the sensitivity to non-SM values of κ_λ such that this parameter may be constrained and, at some point, measured precisely.

One such avenue for improving the analysis' sensitivity to alternative κ_λ hypotheses would be to inject knowledge of κ_λ into the neural network training, perhaps using the technique of parametrized machine learning [226]. The brute force method, of course, to which the methods described in

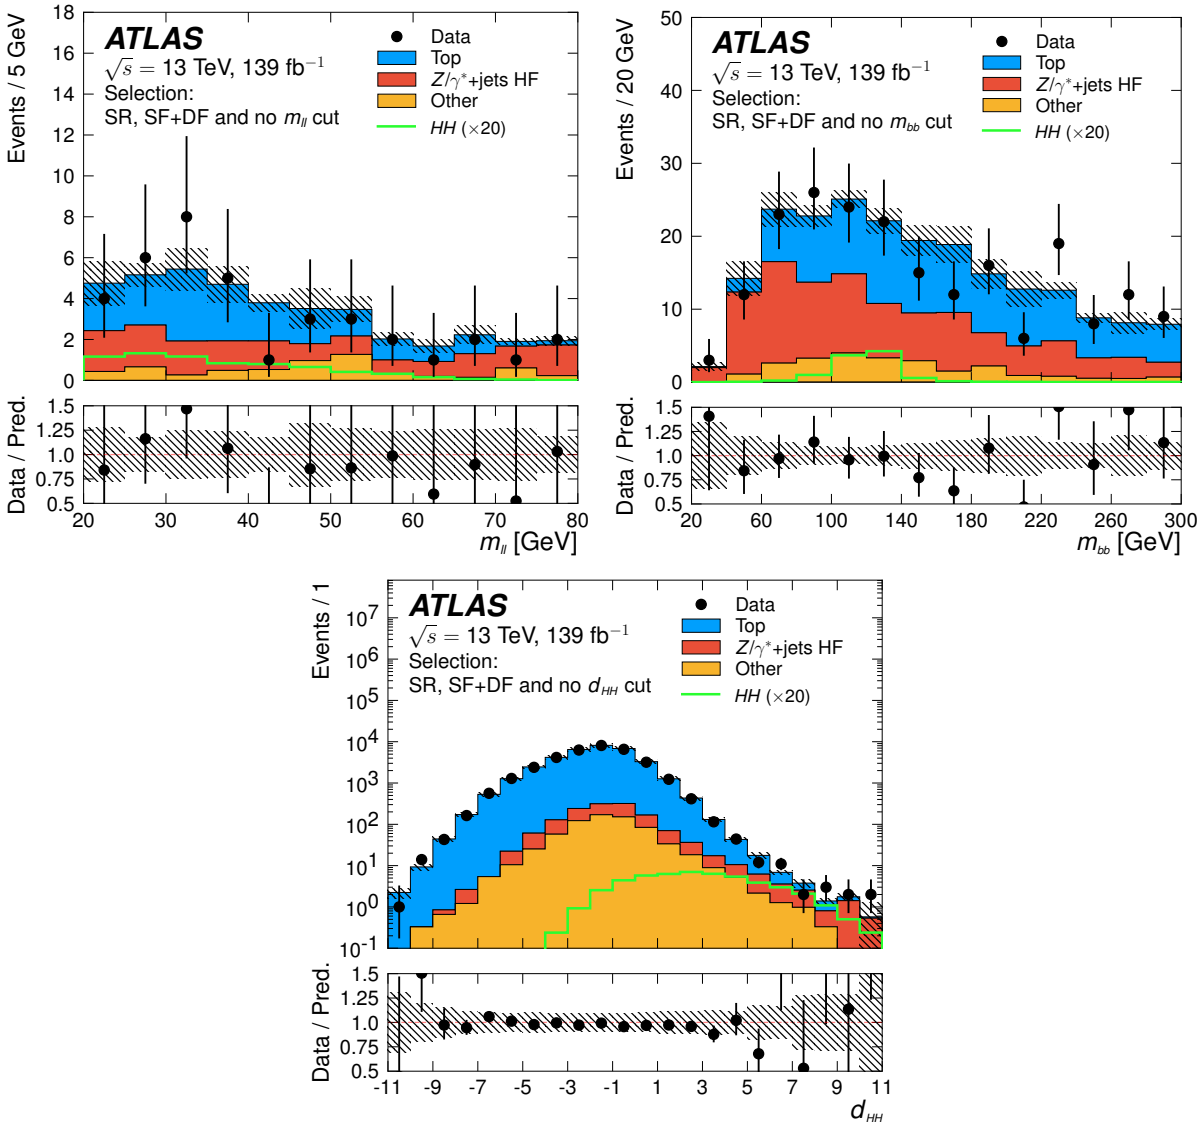

Figure 9.36: Distributions of $m_{\ell\ell}$ (***left***), m_{bb} (***right***), and d_{hh} (***bottom***). The distributions are shown after the fit to data in the control regions under the background-only hypothesis. Each distribution includes both the SF and DF events and imposes the analysis SR requirements on all quantities except for the one being plotted. The SR requirement on d_{hh} in the $m_{\ell\ell}$ and m_{bb} distributions is relaxed to $d_{hh} > 5$. The dilepton $hh \to bbWW^*$ signal, labeled as 'HH', is shown with its cross-section scaled by a factor of 20 relative to the SM prediction for visualization purposes. The ratio of the data to the sum of the backgrounds is shown in the lower panel of each figure. The hatched bands indicate the combined statistical and systematic uncertainty.

Ref. [226] provide an alternative, would be to construct a set of classifiers (potentially many) — each trained using $hh \to bbWW^*$ signal hypotheses with differing values of κ_λ — and then assess the sensitivity to specific values of κ_λ hypotheses using the classifier that has been trained using the hh signal with that value for its Higgs self-coupling parameter.

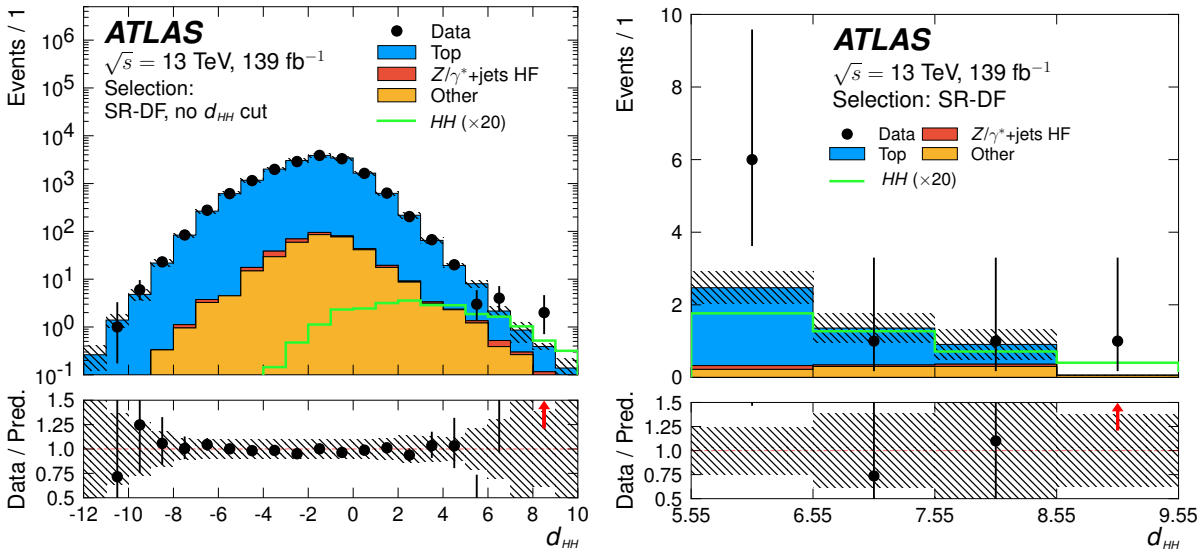

Figure 9.37: Distributions of the d_{hh} observable in SR-DF without the d_{hh} selection applied (***left***) and with the d_{hh} selection ($d_{hh} > 5.55$) applied (***right***). The dilepton $hh \to bbWW^*$ signal, labeled as 'HH', is shown with its cross-section scaled by a factor of 20 relative to the SM prediction for visualization purposes. The ratio of the data to the sum of the backgrounds is shown in the lower panel of each figure. The hatched bands indicate the combined statistical and systematic uncertainty.

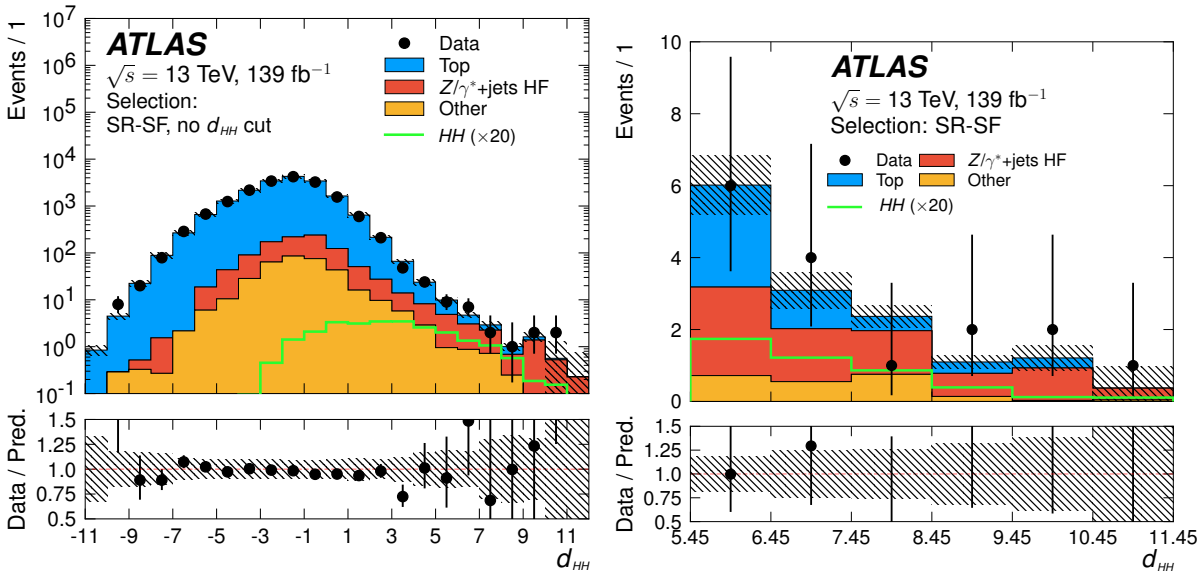

Figure 9.38: Distributions of the d_{hh} observable in SR-SF without the d_{hh} selection applied (***left***) and with the d_{hh} selection ($d_{hh} > 5.45$) applied (***right***). The dilepton $hh \to bbWW^*$ signal, labeled as 'HH', is shown with its cross-section scaled by a factor of 20 relative to the SM prediction for visualization purposes. The ratio of the data to the sum of the backgrounds is shown in the lower panel of each figure. The hatched bands indicate the combined statistical and systematic uncertainty.

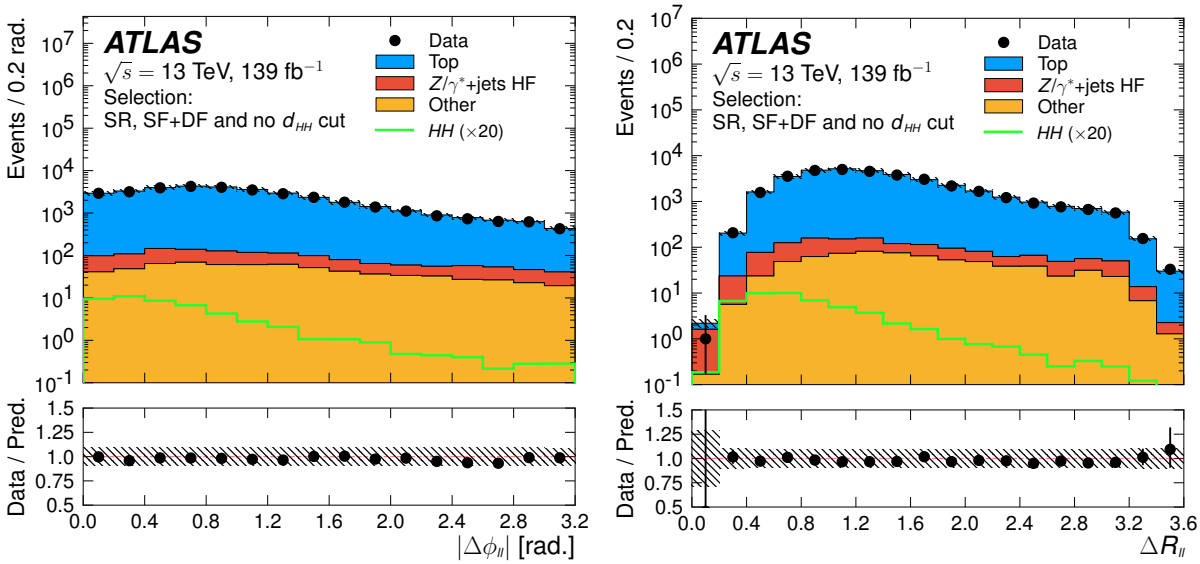

Figure 9.39: Distributions of the $|\Delta\phi_{\ell\ell}|$ and $\Delta R_{\ell\ell}$ quantities in the analysis' SR selections, inclusive of SF and DF dilepton flavors and without the d_{hh} requirements. The dilepton $hh \to bbWW^*$ signal, labeled as 'HH', is shown with its cross-section scaled by a factor of 20 relative to the SM prediction for visualization purposes. The ratio of the data to the sum of the backgrounds is shown in the lower panel of each figure. The hatched bands indicate the combined statistical and systematic uncertainty.

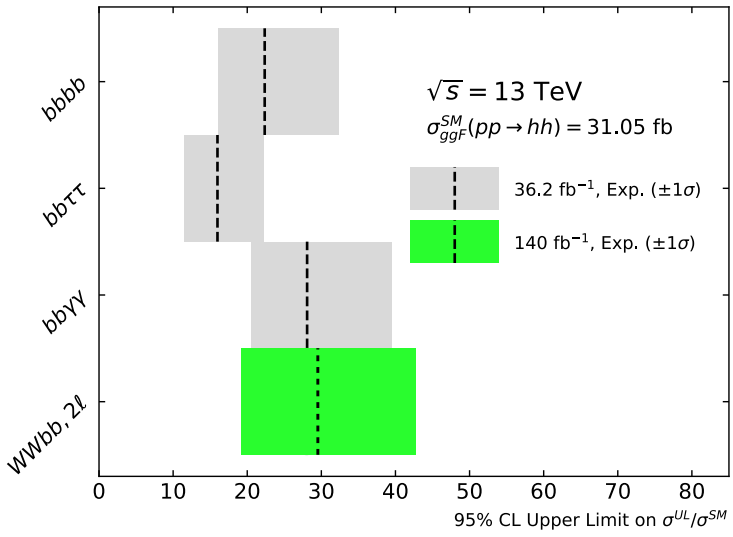

Figure 9.40: Summary plot showing the hh production cross-section upper limit results, normalized to the SM prediction of 31.05 fb, for the analysis presented in this thesis in green, based on the full Run 2 dataset of 139 fb^{-1} of pp collision data. In grey, the partial Run 2 results based on 36 fb^{-1} of data are shown for comparison. Only the expected upper limits are shown.

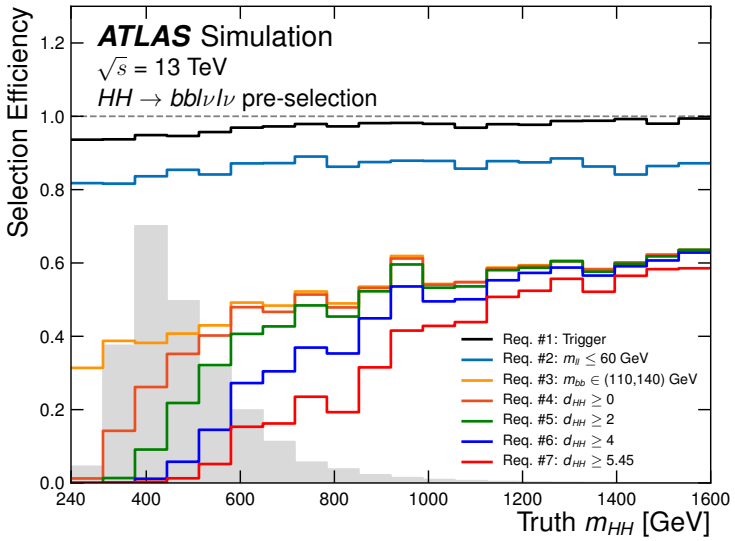

Figure 9.41: Analysis selection efficiency for the dilepton $hh \to bbWW^*$ signal process as a function of the truth-level hh system invariant mass, m_{HH}. Each color indicates an additional selection applied sequentially, and in the order indicated in the legend, with respect to the starting sample of events satisfying the analysis' preselection requirements and having at least two b-tagged jets. For reference, overlaid in grey color, and with arbitrary normalisation, is the truth-level m_{HH} distribution.

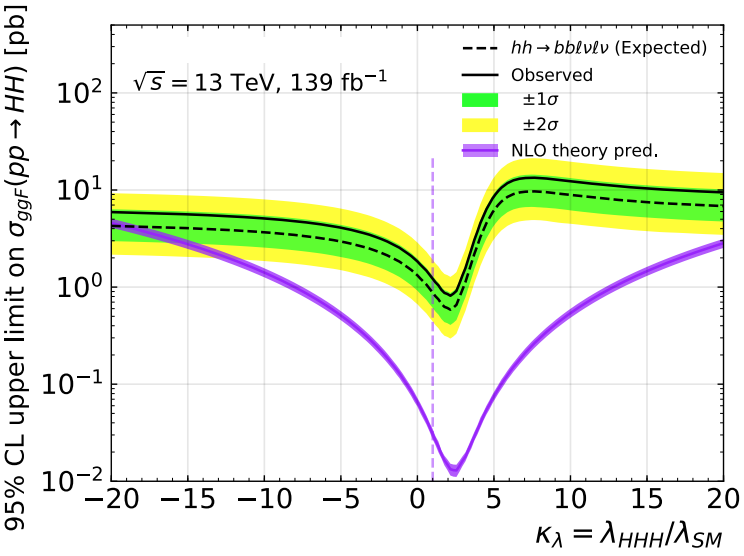

Figure 9.42: Expected and observed cross-section upper-limit for the dilepton $hh \to bbWW^*$ search as a function of the Higgs self-coupling parameter, $\kappa_\lambda = \lambda_{hhh}/\lambda_{hhh}^{\mathrm{SM}}$. The vertical dashed line indicates the SM scenario with $\kappa_\lambda = 1$. The $\pm 1\sigma$ and $\pm 2\sigma$ uncertainty band on the expected upper-limit includes the effects of all experimental and modelling systematic uncertainties. The NLO theory prediction is taken from Ref. [47] and described fully in Ref. [212].

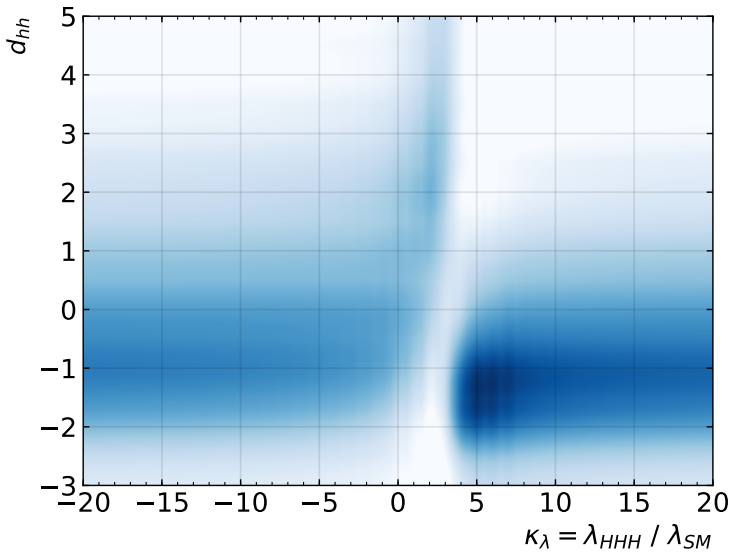

Figure 9.43: The d_{hh} distribution as a function of the Higgs self-coupling parameter, $\kappa_\lambda = \lambda_{hhh}/\lambda_{hhh}^{\mathrm{SM}}$. The trend of the d_{hh} distribution as a function of κ_λ is directly related to the analysis' acceptance to the dilepton $hh \to bbWW^*$ signal under non-SM values of κ_λ, illustrated in Figure 9.42. The darker the blue color means a higher population of the dilepton $hh \to bbWW^*$ signal process. The color white indicates that a given region is not populated at all.

Chapter 10

Concluding Remarks

The point is there ain't no point.

–Cormac McCarthy, *No Country for Old Men*

The effort to understand the universe is one of the very few things that lifts human life a little above the level of farce, and gives it some of the grace of tragedy.

–Steven Weinberg, *The First Three Minutes*

This dissertation has presented work related to the on-going upgrade of the forward muon system of the ATLAS detector — the New Small Wheel — as well as two searches for BSM physics in dilepton final states. The work thus described has taken place within the years 2015–2019, containing the entirety of the LHC Run 2, and during a time of change in the physics program of the ATLAS experiment.

At the time of writing in October 2019, the NSW project is on a critical path, at least with respect to its initial timeline of being fully constructed and installed in ATLAS within the Phase 1 upgrade (c.f. Figure 3.8). In the past few years, the NSW project has seen significant delays. The construction of large-scale MM and sTGC detectors has never been done prior to the NSW, and understanding the construction and behavior of these detectors has been a challenging process that has, at times, required halting chamber production in order to understand emergent phenomena that had not been foreseen. The entire suite of frontend electronics has seen issues, as well, with one of the predominant issues being related to the design of the ASICs themselves. The important readout ASIC, the VMM, has had to go through at least one unforeseen design fix to address subtle issues discovered in the lab, for example. Even at the time of writing, the current understanding

of the yield[1] of the latest version of the VMM ASIC production is unknown. Understanding this is currently of the utmost priority, since the frontend board development and assembly of detector chambers rely first on their being VMMs on-hand before they themselves can move forward and be fully commissioned. If this yield is too low, additional ASICs will have to be manufactured at the silicon foundries, which is a lengthy (and expensive) process. Additionally, the NSW will be the first major detector subsystem of ATLAS to use the upgraded TDAQ system based on the Front-end Link Exchange (FELIX) [227] backend. FELIX, itself, has seen significant delays in its Phase 1 upgrade progress that have led to further slowing of the progress of the NSW. With this in mind, the current best foreseen Phase 1 installation plan for the NSW project will be to install only a single wheel into ATLAS before the start of LHC Run 3, with the second wheel being installed at some point during Run 3, perhaps during one of the winter shutdown periods. It may be likely that neither of the wheels is installed during Phase 1, at which point both will have to be scheduled for installation in Run 3. The NSW is a very ambitious project, and such issues as described here are not indicative of those at the heart of moving it forward but rather a sign of the short window of time that such a large project was aiming to be completed in, as well as its having started at a time of great busyness and distraction at CERN. The current author looks forward to its complete commissioning and smooth operation within the ATLAS experimental cavern, which will be a testament to many, many talented people from around the world having worked together and designed, built, and butted-heads against a massive project of both new ideas and new technologies.

Given the naturalness arguments vying for electroweak-scale SUSY (Section 2.1), the most promising LHC signals for evidence of SUSY have been thought to be obtained in the searches for the direct production of the stop quark, $\tilde{t}_1$. The search for the direct production of pairs of relatively light stops presented in this thesis provides a dilepton strategy for thoroughly investigating this challenging phase space of SUSY, but has so far shown no signs of anything else but consistency with the SM. With the end of LHC Run 2, the SUSY searches for $\tilde{t}_1$ have generally designed analyses that can effectively probe the $(m_{\tilde{t}_1}, m_{\tilde{\chi}_1^0})$ plane. The analyses' strategies are therefore converging and are tending to be bound only by the available statistics in the recorded data and center-of-mass collision energies. Figure 10.1 shows the present status of all searches for the $\tilde{t}_1$ by ATLAS, where it can be seen that present limits extend up to $\tilde{t}_1$ masses of 1 TeV. Given the initial hopes for SUSY being 'just around the corner' with the discovery of a relatively light Higgs boson, the present limits extending to $m_{\tilde{t}_1} \approx 1\,\text{TeV}$ may seem disheartening to those who have been 'waiting in the wings', so to speak. But one must remember the rather biased view of SUSY that such limits present, having been performed under the hypotheses of simplified models of SUSY and with very specific instantiations of the MSSM. In this light, it could be seen as being almost naive to consider searches for $\tilde{t}_1$ disheartening: to have thought it promising in 2009, just as the LHC was turning on, and to think it less so now is a rather impatient view of Nature [228], especially given the many forms that the MSSM can take while remaining consistent with a Higgs boson having a mass of

[1] The term 'yield' here means the fraction of produced VMMs that meet the performance criteria for being sufficient for use in ATLAS, and are not failing specific criteria that are necessary in order to meet the physics goals of ATLAS.

125 GeV [229]. Still, these last points bring skepticism to bear on the effectiveness of the direct searches for SUSY. At the very least, they motivate the continued push for the understanding of the most relevant operating parameters of the MSSM such that searches with well-defined hypotheses can be initiated. Otherwise, the results of the direct searches for the $\tilde{t}_1$ must be interpreted not as making general statements about SUSY, but rather as broad statements about the consistency between observation and prediction in extreme regions of phase space: the SUSY simplified models only defining for the analyses topological handles with which metrics quantifying their sensitivity to these kinematics and final states can be made.

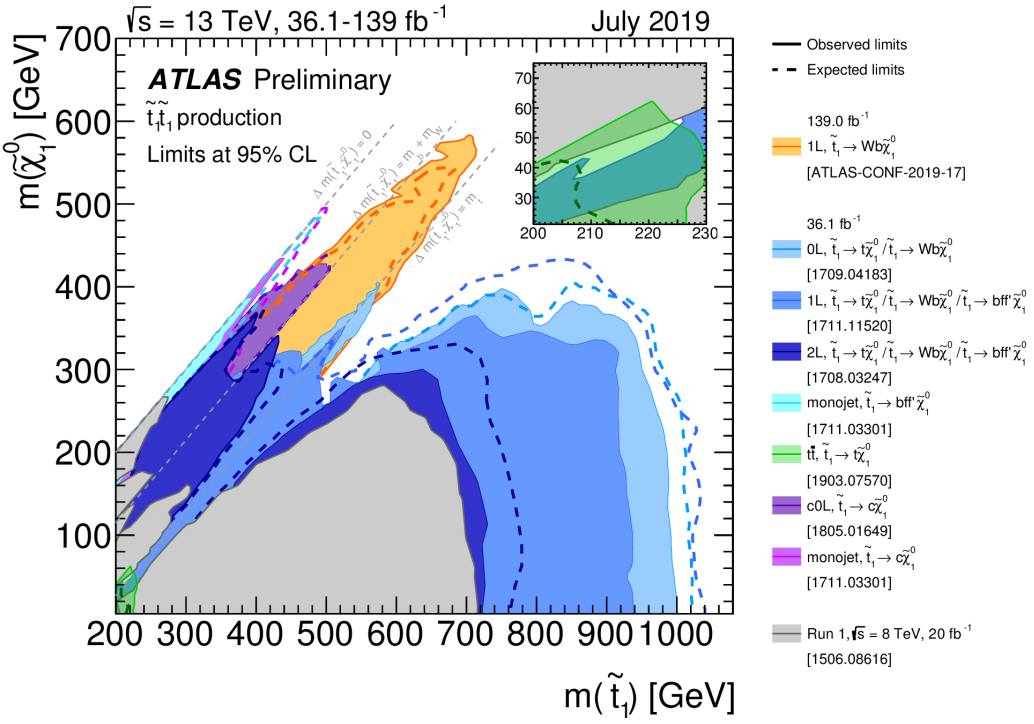

Figure 10.1: Summary of ATLAS' LHC Run 2 exclusions in the $(\tilde{t}_1, \tilde{\chi}_1^0)$ plane, as of July 2019. Figure taken from Ref. [230].

With so many theories of BSM physics having something to say about the Higgs sector of the SM, mainly in their having to address the Hierarchy Problem, the thorough study of this newly discovered particle is of the utmost importance. This is especially true in light of the failure of the direct searches for new physics to find any significant discrepancies in the agreement of the observed data and the SM. Towards this end, the ATLAS program for the search for the production of Higgs boson pairs is approaching a level of maturity that will set itself up for the long-term LHC goal of observing this process. It is primarily through the study of this process, and measurement of the Higgs boson self-coupling parameter, that a complete understanding of whether or not the Higgs sector behaves as the SM predicts can be obtained. The search presented in this thesis for SM-like Higgs boson pair production in the dilepton final state of $hh \to bbWW^*$ is a brand new addition

to this search program, having filled a gap in the hh searches' coverage that had been previously disregarded for one reason or another. It is the first iteration of this search; however, it appears to already be competitive with those search channels previously thought to be the only means of being sensitive to the hh production process. That said, there is much to be improved with the analysis, as touched upon in Chapter 9. In the lead up to — and during — Run 3, as the dilepton $bbWW^*$ channel becomes more well understood, it will likely become more relevant to the overall search for the hh process and in its ability to constrain BSM effects in the Higgs sector that would act to modify the hh production cross-section or to alter the Higgs boson self-coupling parameter. If meaningful statements about the underlying nature of EWSB are to be made by the end of the HL-LHC era, these searches must be active areas of study and their optimality must be ensured as soon as possible. If anything, the addition of a new search channel at the end of LHC Run 2 that was previously not thought to be relevant has sparked a little extra momentum on this front.

Bibliography

[1] A. Einstein. "Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt (On a Heuristic Point of View about the Creation and Conversion of Light)". In: *Annalen der Physik* 322 (1905), pp. 132–148. DOI: 10.1002/andp.19053220607 (cit. on p. 0).

[2] A. Einstein. "Zur Elektrodynamik bewegter Körper (On the Electrodynamics of Moving Bodies)". In: *Annalen der Physik* 322 (1905), pp. 891–921. DOI: 10.1002/andp.19053221004 (cit. on p. 0).

[3] A. Einstein. "Ist die Trägheit eines Körpers von seinem Energieinhalt abhängig? (Does the Inertia of a Body Depend Upon Its Energy Content?)" In: *Annalen der Physik* 323 (1905), pp. 639–641. DOI: 10.1002/andp.19053231314 (cit. on p. 0).

[4] M. Planck. "Ueber das Gesetz der Energieverteilung im Normalspectrum (On the law of the distribution of energy in the normal spectrum)". In: *Annalen der Physik* 309 (1901), pp. 553–563. DOI: 10.1002/andp.19013090310 (cit. on p. 0).

[5] P. A. M. Dirac. "The Quantum Theory of the Electron". In: *Proceedings of the Royal Society of London Series A* 117 (Feb. 1928), pp. 610–624. DOI: 10.1098/rspa.1928.0023 (cit. on p. 1).

[6] Paul A. M. Dirac. "Quantum theory of emission and absorption of radiation". In: *Proc. Roy. Soc. Lond.* A114 (1927), p. 243. DOI: 10.1098/rspa.1927.0039 (cit. on p. 1).

[7] Sigfrido Boffi. "The Rise of Quantum Mechanics". In: *arXiv e-prints* (2008). arXiv: 0806.4515 [physics.hist-ph] (cit. on p. 1).

[8] Jagdish Mehra and Helmut Rechenberg. "The Historical Development of Quantum Theory". In: (2001) (cit. on p. 1).

[9] S. Tomonaga. "On a Relativistically Invariant Formulation of the Quantum Theory of Wave Fields". In: *Progress of Theoretical Physics* 1 (Aug. 1946), pp. 27–42. DOI: 10.1143/PTP.1.27 (cit. on p. 1).

[10] R. P. Feynman. "Space-Time Approach to Quantum Electrodynamics". In: *Physical Review* 76 (Sept. 1949), pp. 769–789. DOI: 10.1103/PhysRev.76.769 (cit. on p. 1).

[11] R. P. Feynman. "The Theory of Positrons". In: *Physical Review* 76 (Sept. 1949), pp. 749–759. DOI: 10.1103/PhysRev.76.749 (cit. on p. 1).

[12] R. P. Feynman. "Mathematical Formulation of the Quantum Theory of Electromagnetic Interaction". In: *Physical Review* 80 (Nov. 1950), pp. 440–457. DOI: 10.1103/PhysRev.80.440 (cit. on p. 1).

[13] J. Schwinger. "On Quantum-Electrodynamics and the Magnetic Moment of the Electron". In: *Physical Review* 73 (Feb. 1948), pp. 416–417. DOI: 10.1103/PhysRev.73.416 (cit. on p. 1).

[14] J. Schwinger. "Quantum Electrodynamics. I. A Covariant Formulation". In: *Physical Review* 74 (Nov. 1948), pp. 1439–1461. DOI: 10.1103/PhysRev.74.1439 (cit. on p. 1).

[15] F. J. Dyson. "The Radiation Theories of Tomonaga, Schwinger, and Feynman". In: *Physical Review* 75 (Feb. 1949), pp. 486–502. DOI: 10.1103/PhysRev.75.486 (cit. on p. 1).

[16] F. J. Dyson. "The S Matrix in Quantum Electrodynamics". In: *Physical Review* 75 (June 1949), pp. 1736–1755. DOI: 10.1103/PhysRev.75.1736 (cit. on p. 1).

[17] S. S. Schweber. *QED and the men who made it: Dyson, Feynman, Schwinger, and Tomonaga*. 1994 (cit. on p. 1).

[18] S. L. Glashow. "Partial Symmetries of Weak Interactions". In: *Nucl. Phys.* 22 (1961), pp. 579–588. DOI: 10.1016/0029-5582(61)90469-2 (cit. on p. 5).

[19] Steven Weinberg. "A Model of Leptons". In: *Phys. Rev. Lett.* 19 (1967), pp. 1264–1266. DOI: 10.1103/PhysRevLett.19.1264 (cit. on p. 5).

[20] Abdus Salam. "Weak and Electromagnetic Interactions". In: *Conf. Proc.* C680519 (1968), pp. 367–377 (cit. on p. 5).

[21] Gerard 't Hooft. "Renormalization of Massless Yang-Mills Fields". In: *Nucl. Phys.* B33 (1971), pp. 173–199. DOI: 10.1016/0550-3213(71)90395-6 (cit. on p. 5).

[22] Gerard 't Hooft. "Renormalizable Lagrangians for Massive Yang-Mills Fields". In: *Nucl. Phys.* B35 (1971). [,201(1971)], pp. 167–188. DOI: 10.1016/0550-3213(71)90139-8 (cit. on p. 5).

[23] David J. Gross and Frank Wilczek. "Ultraviolet Behavior of Non-Abelian Gauge Theories". In: *Phys. Rev. Lett.* 30 (26 June 1973), pp. 1343–1346. DOI: 10.1103/PhysRevLett.30.1343. URL: https://link.aps.org/doi/10.1103/PhysRevLett.30.1343 (cit. on p. 5).

[24] H. David Politzer. "Reliable Perturbative Results for Strong Interactions?" In: *Phys. Rev. Lett.* 30 (26 June 1973), pp. 1346–1349. DOI: 10.1103/PhysRevLett.30.1346. URL: https://link.aps.org/doi/10.1103/PhysRevLett.30.1346 (cit. on p. 5).

[25] F. Englert and R. Brout. "Broken Symmetry and the Mass of Gauge Vector Mesons". In: *Phys. Rev. Lett.* 13 (1964). [,157(1964)], pp. 321–323. DOI: 10.1103/PhysRevLett.13.321 (cit. on p. 8).

[26] Peter W. Higgs. "Broken symmetries, massless particles and gauge fields". In: *Phys. Lett.* 12 (1964), pp. 132–133. DOI: 10.1016/0031-9163(64)91136-9 (cit. on p. 8).

[27] Peter W. Higgs. "Broken Symmetries and the Masses of Gauge Bosons". In: *Phys. Rev. Lett.* 13 (1964). [,160(1964)], pp. 508–509. DOI: 10.1103/PhysRevLett.13.508 (cit. on p. 8).

[28] Jeffrey Goldstone, Abdus Salam, and Steven Weinberg. "Broken Symmetries". In: *Phys. Rev.* 127 (1962), pp. 965–970. DOI: 10.1103/PhysRev.127.965 (cit. on p. 10).

[29] Evgeny Kh. Akhmedov. "Do charged leptons oscillate?" In: *JHEP* 09 (2007), p. 116. DOI: 10.1088/1126-6708/2007/09/116. arXiv: 0706.1216 [hep-ph] (cit. on p. 12).

[30] Makoto Kobayashi and Toshihide Maskawa. "CP Violation in the Renormalizable Theory of Weak Interaction". In: *Prog. Theor. Phys.* 49 (1973), pp. 652–657. DOI: 10.1143/PTP.49.652 (cit. on p. 12).

[31] *Standard Model Summary Plots Summer 2019*. Tech. rep. ATL-PHYS-PUB-2019-024. Geneva: CERN, July 2019. URL: http://cds.cern.ch/record/2682186 (cit. on p. 14).

[32] 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". In: *Phys. Rev.* D90.5 (2014), p. 052004. DOI: 10.1103/PhysRevD.90.052004. arXiv: 1406.3827 [hep-ex] (cit. on p. 15).

[33] CMS Collaboration. "Precise determination of the mass of the Higgs boson and studies of the compatibility of its couplings with the standard model". In: (2014) (cit. on p. 15).

[34] M. Baak et al. "The global electroweak fit at NNLO and prospects for the LHC and ILC". In: *Eur. Phys. J.* C74 (2014), p. 3046. DOI: 10.1140/epjc/s10052-014-3046-5. arXiv: 1407.3792 [hep-ph] (cit. on p. 15).

[35] ATLAS Collaboration. "Measurement of the W-boson mass in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector". In: *Eur. Phys. J.* C78.2 (2018). [Erratum: Eur. Phys. J.C78,no.11,898(2018)], p. 110. DOI: 10.1140/epjc/s10052-018-6354-3,10.1140/epjc/s10052-017-5475-4. arXiv: 1701.07240 [hep-ex] (cit. on p. 15).

[36] ATLAS Collaboration. "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC". In: *Phys. Lett.* B716 (2012), pp. 1–29. DOI: 10.1016/j.physletb.2012.08.020. arXiv: 1207.7214 [hep-ex] (cit. on p. 14).

[37] CMS Collaboration. "Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC". In: *Phys. Lett.* B716 (2012), pp. 30–61. DOI: 10.1016/j.physletb.2012.08.021. arXiv: 1207.7235 [hep-ex] (cit. on p. 14).

[38] ATLAS Collaboration. "Study of the spin and parity of the Higgs boson in diboson decays with the ATLAS detector". In: *Eur. Phys. J.* C75.10 (2015). [Erratum: Eur. Phys. J.C76,no.3,152(2016)], p. 476. DOI: 10.1140/epjc/s10052-015-3685-1,10.1140/epjc/s10052-016-3934-y. arXiv: 1506.05669 [hep-ex] (cit. on p. 16).

[39] ATLAS Collaboration. "Test of CP Invariance in vector-boson fusion production of the Higgs boson using the Optimal Observable method in the ditau decay channel with the ATLAS detector". In: *Eur. Phys. J.* C76.12 (2016), p. 658. DOI: 10.1140/epjc/s10052-016-4499-5. arXiv: 1602.04516 [hep-ex] (cit. on p. 16).

[40] ATLAS Collaboration. "Observation of Higgs boson production in association with a top quark pair at the LHC with the ATLAS detector". In: *Phys. Lett.* B784 (2018), pp. 173–191. DOI: 10.1016/j.physletb.2018.07.035. arXiv: 1806.00425 [hep-ex] (cit. on p. 16).

[41] ATLAS Collaboration. "Observation of $H \to b\bar{b}$ decays and VH production with the ATLAS detector". In: *Phys. Lett.* B786 (2018), pp. 59–86. DOI: 10.1016/j.physletb.2018.09.013. arXiv: 1808.08238 [hep-ex] (cit. on p. 16).

[42] ATLAS and CMS Collaborations. "Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV". In: *JHEP* 08 (2016), p. 045. DOI: 10.1007/JHEP08(2016)045. arXiv: 1606.02266 [hep-ex] (cit. on p. 16).

[43] CMS Collaboration. "Observation of $t\bar{t}H$ production". In: *Phys. Rev. Lett.* 120.23 (2018), p. 231801. DOI: 10.1103/PhysRevLett.120.231801. arXiv: 1804.02610 [hep-ex] (cit. on p. 16).

[44] CMS Collaboration. "Observation of Higgs boson decay to bottom quarks". In: *Phys. Rev. Lett.* 121.12 (2018), p. 121801. DOI: 10.1103/PhysRevLett.121.121801. arXiv: 1808.08242 [hep-ex] (cit. on p. 16).

[45] CMS Collaboration. "Measurements of the Higgs boson width and anomalous HVV couplings from on-shell and off-shell production in the four-lepton final state". In: *Phys. Rev.* D99.11 (2019), p. 112003. DOI: 10.1103/PhysRevD.99.112003. arXiv: 1901.00174 [hep-ex] (cit. on p. 16).

[46] CMS Collaboration. "Constraints on anomalous HVV couplings from the production of Higgs bosons decaying to τ lepton pairs". In: *Submitted to: Phys. Rev.* (2019). arXiv: 1903.06973 [hep-ex] (cit. on p. 16).

[47] D. de Florian et al. "Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector". In: (2016). DOI: 10.2172/1345634,10.23731/CYRM-2017-002. arXiv: 1610.07922 [hep-ph] (cit. on pp. 17, 233, 287).

[48] ATLAS Collaboration. "Combined measurements of Higgs boson production and decay using up to 80 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment". In: (2019). arXiv: 1909.02845 [hep-ex] (cit. on p. 18).

[49] Tamara M. Davis. "Cosmological constraints on dark energy". In: *Gen. Rel. Grav.* 46 (2014), p. 1731. DOI: 10.1007/s10714-014-1731-1. arXiv: 1404.7266 [astro-ph.CO] (cit. on p. 20).

[50] Planck Collaboration. "Planck 2018 results. VI. Cosmological parameters". In: (2018). arXiv: 1807.06209 [astro-ph.CO] (cit. on p. 20).

[51] Y. Fukuda et al. "Evidence for oscillation of atmospheric neutrinos". In: *Phys. Rev. Lett.* 81 (1998), pp. 1562–1567. DOI: 10.1103/PhysRevLett.81.1562. arXiv: hep-ex/9807003 [hep-ex] (cit. on p. 20).

[52] Andrei D Sakharov. "Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe". In: *Soviet Physics Uspekhi* 34.5 (May 1991), pp. 392–393. DOI: 10.1070/pu1991v034n05abeh002497. URL: https://doi.org/10.1070%2Fpu1991v034n05abeh002497 (cit. on p. 20).

[53] J.L. Díaz-Cruz and D.A. López-Falcón. "Probing the mechanism of EWSB with a rho-parameter defined in terms of Higgs couplings". In: *Physics Letters B* 568.3 (2003), pp. 245–253. ISSN: 0370-2693. DOI: https://doi.org/10.1016/j.physletb.2003.06.047. URL: http://www.sciencedirect.com/science/article/pii/S037026930300947X (cit. on p. 22).

[54] K.A. Olive. "Review of Particle Physics". In: *Chinese Physics C* 40.10 (Oct. 2016), p. 100001. DOI: 10.1088/1674-1137/40/10/100001. URL: https://doi.org/10.1088%2F1674-1137%2F40%2F10%2F100001 (cit. on pp. 22, 95, 158).

[55] G. C. Branco et al. "Theory and phenomenology of two-Higgs-doublet models". In: *Phys. Rept.* 516 (2012), pp. 1–102. DOI: 10.1016/j.physrep.2012.02.002. arXiv: 1106.0034 [hep-ph] (cit. on p. 23).

[56] Sidney Coleman and Jeffrey Mandula. "All Possible Symmetries of the S Matrix". In: *Phys. Rev.* 159 (5 July 1967), pp. 1251–1256. DOI: 10.1103/PhysRev.159.1251. URL: https://link.aps.org/doi/10.1103/PhysRev.159.1251 (cit. on p. 24).

[57] Rudolf Haag, Jan T. Łopuszański, and Martin Sohnius. "All possible generators of supersymmetries of the S-matrix". In: *Nuclear Physics B* 88.2 (1975), pp. 257–274. ISSN: 0550-3213. DOI: https://doi.org/10.1016/0550-3213(75)90279-5. URL: http://www.sciencedirect.com/science/article/pii/0550321375902795 (cit. on p. 24).

[58] Peisi Huang et al. "Corrections to di-Higgs boson production with light stops and modified Higgs couplings". In: *Phys. Rev.* D97.7 (2018), p. 075001. DOI: 10.1103/PhysRevD.97.075001. arXiv: 1711.05743 [hep-ph] (cit. on pp. 27, 28).

[59] "CERN Annual Personnel Statistics 2018". In: (2018). URL: https://cds.cern.ch/record/2677223 (cit. on p. 29).

[60] The ALICE Collaboration. "The ALICE experiment at the CERN LHC". In: *Journal of Instrumentation* 3.08 (Aug. 2008), S08002–S08002. DOI: 10.1088/1748-0221/3/08/s08002. URL: https://doi.org/10.1088%2F1748-0221%2F3%2F08%2Fs08002 (cit. on p. 29).

[61] The LHCb Collaboration. "The LHCb Detector at the LHC". In: *Journal of Instrumentation* 3.08 (Aug. 2008), S08005–S08005. DOI: 10.1088/1748-0221/3/08/s08005. URL: https://doi.org/10.1088%2F1748-0221%2F3%2F08%2Fs08005 (cit. on p. 29).

[62] The CMS Collaboration. "The CMS experiment at the CERN LHC". In: *Journal of Instrumentation* 3.08 (Aug. 2008), S08004–S08004. DOI: 10.1088/1748-0221/3/08/s08004. URL: https://doi.org/10.1088%2F1748-0221%2F3%2F08%2Fs08004 (cit. on p. 29).

[63] The ATLAS Collaboration. "The ATLAS Experiment at the CERN Large Hadron Collider". In: *Journal of Instrumentation* 3.08 (Aug. 2008), S08003–S08003. DOI: 10.1088/1748-0221/3/08/s08003. URL: https://doi.org/10.1088%2F1748-0221%2F3%2F08%2Fs08003 (cit. on pp. 29, 42, 44, 45, 46, 48).

[64] Christiane Lefèvre. "The CERN accelerator complex. Complexe des accélérateurs du CERN". Dec. 2008. URL: https://cds.cern.ch/record/1260465 (cit. on p. 30).

[65] Lyndon Evans and Philip Bryant. "LHC Machine". In: *Journal of Instrumentation* 3.08 (Aug. 2008), S08001–S08001. DOI: 10.1088/1748-0221/3/08/s08001. URL: https://doi.org/10.1088%2F1748-0221%2F3%2F08%2Fs08001 (cit. on pp. 30, 34, 35).

[66] *LEP Design Report: Vol. 1: The LEP Injector Chain*. CERN, 1983. URL: https://cds.cern.ch/record/98881 (cit. on p. 31).

[67] *LEP Design Report: Vol. 2: The LEP Main Ring*. CERN, 1984. URL: https://cds.cern.ch/record/102083 (cit. on p. 31).

[68] Stephen Holmes, Ronald S Moore, and Vladimir Shiltsev. "Overview of the Tevatron collider complex: goals, operations and performance". In: *Journal of Instrumentation* 6.08 (Aug. 2011), T08001–T08001. DOI: 10.1088/1748-0221/6/08/t08001. URL: https://doi.org/10.1088%2F1748-0221%2F6%2F08%2Ft08001 (cit. on p. 31).

[69] Juan Casas et al. "Design concept and first experimental validation of the superfluid helium system for the large hadron collider (LHC) project at CERN". In: *Cryogenics* 32 (1992), pp. 118–121. DOI: 10.1016/0011-2275(92)90122-Q (cit. on p. 32).

[70] *ECFA-CERN Workshop on Large Hadron Collider in the LEP tunnel*. CERN. Geneva: CERN, 1984. DOI: 10.5170/CERN-1984-010-V-1. URL: https://cds.cern.ch/record/154938 (cit. on p. 33).

[71] Lyndon Evans. "The Large Hadron Collider". In: *New Journal of Physics* 9.9 (Sept. 2007), pp. 335–335. DOI: 10.1088/1367-2630/9/9/335. URL: https://doi.org/10.1088%2F1367-2630%2F9%2F9%2F335 (cit. on p. 34).

[72] Oliver Sim Brüning et al. *LHC Design Report: Vol. 1: The LHC Main Ring*. CERN Yellow Reports: Monographs. Geneva: CERN, 2004. DOI: 10.5170/CERN-2004-003-V-1. URL: https://cds.cern.ch/record/782076 (cit. on p. 34).

[73] Oliver Sim Brüning et al. *LHC Design Report: Vol. 2: The LHC Infrastructure and General Services*. CERN Yellow Reports: Monographs. Geneva: CERN, 2004. DOI: 10.5170/CERN-2004-003-V-2. URL: https://cds.cern.ch/record/815187 (cit. on p. 36).

[74] Michael Benedikt et al. *LHC Design Report: Vol. 3: The LHC Injector Chain*. CERN Yellow Reports: Monographs. Geneva: CERN, 2004. DOI: 10.5170/CERN-2004-003-V-3. URL: https://cds.cern.ch/record/823808 (cit. on pp. 35, 38).

[75] R Scrivens et al. "Overview of the status and developments on primary ion sources at CERN*". In: CERN-ATS-2011-172 (Sept. 2011), 4 p. URL: https://cds.cern.ch/record/1382102 (cit. on p. 36).

[76] Daniel Boussard and Trevor Paul R Linnecar. *The LHC Superconducting RF System*. Tech. rep. LHC-Project-Report-316. CERN-LHC-Project-Report-316. Geneva: CERN, Dec. 1999. URL: https://cds.cern.ch/record/410377 (cit. on p. 37).

[77] S Baird. *Accelerators for pedestrians*. Tech. rep. AB-Note-2007-014. CERN-AB-Note-2007-014. PS-OP-Note-95-17-Rev-2. CERN-PS-OP-Note-95-17-Rev-2. Geneva: CERN, Feb. 2007. URL: http://cds.cern.ch/record/1017689 (cit. on p. 38).

[78] Werner Herr and B Muratori. "Concept of luminosity". In: (2006). DOI: 10.5170/CERN-2006-002.361. URL: https://cds.cern.ch/record/941318 (cit. on p. 38).

[79] ATLAS Collaboration. *ATLAS magnet system: Technical Design Report, 1*. Technical Design Report ATLAS. Geneva: CERN, 1997. URL: https://cds.cern.ch/record/338080 (cit. on p. 43).

[80] S Haywood et al. "ATLAS inner detector: Technical Design Report, 2". In: Technical Design Report ATLAS (1997). URL: http://cds.cern.ch/record/331064 (cit. on p. 41).

[81] M Capeans et al. *ATLAS Insertable B-Layer Technical Design Report*. Tech. rep. CERN-LHCC-2010-013. ATLAS-TDR-19. Sept. 2010. URL: https://cds.cern.ch/record/1291633 (cit. on p. 44).

[82] M. Backhaus. "The upgraded Pixel Detector of the ATLAS Experiment for Run 2 at the Large Hadron Collider". In: *Nucl. Instrum. Meth.* A831 (2016), pp. 65–70. DOI: 10.1016/j.nima.2016.05.018 (cit. on p. 46).

[83] *ATLAS liquid-argon calorimeter: Technical Design Report*. Technical Design Report ATLAS. Geneva: CERN, 1996. URL: https://cds.cern.ch/record/331061 (cit. on pp. 48, 49).

[84] *ATLAS tile calorimeter: Technical Design Report*. Technical Design Report ATLAS. Geneva: CERN, 1996. URL: https://cds.cern.ch/record/331062 (cit. on pp. 48, 51).

[85] Marija Marjanovic. "ATLAS Tile calorimeter calibration and monitoring systems". In: *21st IEEE Real Time Conference (RT2018) Williamsburg, Virginia, June 11-15, 2018*. 2018. arXiv: 1806.09156 [physics.ins-det] (cit. on p. 51).

[86] ATLAS Collaboration. "The ATLAS Forward Calorimeter". In: *Journal of Instrumentation* 3.02 (Feb. 2008), P02010–P02010. DOI: 10.1088/1748-0221/3/02/p02010. URL: https://doi.org/10.1088%2F1748-0221%2F3%2F02%2Fp02010 (cit. on pp. 51, 52).

[87] ATLAS Collaboration. *ATLAS muon spectrometer: Technical Design Report*. Technical Design Report ATLAS. Geneva: CERN, 1997. URL: https://cds.cern.ch/record/331068 (cit. on pp. 52, 53, 54, 56, 57, 160).

[88] F. Bauer et al. "Construction and Test of MDT Chambers for the ATLAS Muon Spectrometer". In: *Nucl. Instrum. Meth.* A461 (2001), pp. 17–20. DOI: 10.1016/S0168-9002(00)01156-6. arXiv: 1604.02000 [physics.ins-det] (cit. on p. 53).

[89] Theodoros Argyropoulos et al. "Cathode strip chambers in ATLAS: Installation, commissioning and in situ performance". In: *IEEE Trans. Nucl. Sci.* 56 (2009), pp. 1568–1574. DOI: 10.1109/TNS.2009.2020861 (cit. on p. 53).

[90] G. Aielli et al. "The RPC first level muon trigger in the barrel of the ATLAS experiment". In: *Nucl. Phys. Proc. Suppl.* 158 (2006). [,11(2006)], pp. 11–15. DOI: 10.1016/j.nuclphysbps.2006.07.031 (cit. on p. 53).

[91] S. Majewski et al. "A thin multiwire chamber operating in the high multiplication mode". In: *Nucl. Instrum. Meth.* 217 (1983), pp. 265–271. DOI: 10.1016/0167-5087(83)90146-1 (cit. on p. 53).

[92] ATLAS Collaboration. *Muon Reconstruction and Identification: Studies with Simulated Monte Carlo Samples*. Tech. rep. ATL-PHYS-PUB-2009-008. ATL-COM-PHYS-2009-151. Geneva: CERN, Mar. 2009. URL: https://cds.cern.ch/record/1169053 (cit. on p. 54).

[93] A S Mete. "Search for a heavy charged gauge boson decaying to a muon and a neutrino in 1 fb-1 of proton-proton collisions at $\sqrt{s} = 7$ TeV using the ATLAS Detector". Presented 15 Nov 2011. Feb. 2012. URL: https://cds.cern.ch/record/1454661 (cit. on p. 54).

[94] "Performance of the ATLAS muon trigger in pp collisions at $\sqrt{s} = 8$ TeV". In: *The European Physical Journal C* 75.3 (Mar. 2015), p. 120. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-015-3325-9. URL: https://doi.org/10.1140/epjc/s10052-015-3325-9 (cit. on p. 55).

[95] P Branchini et al. "Global Time Fit for Tracking in an Array of Drift Cells: The Drift Tubes of the ATLAS Experiment". In: *Nuclear Science, IEEE Transactions on* 55 (Mar. 2008), pp. 620–627. DOI: 10.1109/TNS.2007.914020 (cit. on p. 56).

[96] Peter Jenni et al. *ATLAS high-level trigger, data-acquisition and controls: Technical Design Report*. Technical Design Report ATLAS. Geneva: CERN, 2003. URL: https://cds.cern.ch/record/616089 (cit. on pp. 58, 59).

[97] William Panduro Vazquez and ATLAS Collaboration. *The ATLAS Data Acquisition system in LHC Run 2*. Tech. rep. ATL-DAQ-PROC-2017-007. 3. Geneva: CERN, Feb. 2017. DOI: 10.1088/1742-6596/898/3/032017. URL: https://cds.cern.ch/record/2244345 (cit. on pp. 58, 59).

[98] *ATLAS Computing: technical design report*. Technical Design Report ATLAS. Geneva: CERN, 2005. URL: https://cds.cern.ch/record/837738 (cit. on p. 58).

[99] C P Bee et al. *The raw event format in the ATLAS Trigger and DAQ*. Tech. rep. ATL-DAQ-98-129. Revised version number 5 submitted on 2016-11-03 11:47. Geneva: CERN, Feb. 2016. URL: https://cds.cern.ch/record/683741 (cit. on p. 59).

[100] Ringaile Placakyte. "Parton Distribution Functions". In: *Proceedings, 31st International Conference on Physics in collisions (PIC 2011): Vancouver, Canada, August 28-September 1, 2011*. 2011. arXiv: 1111.5452 [hep-ph] (cit. on p. 62).

[101] Richard D. Ball et al. "Parton distributions with LHC data". In: *Nucl. Phys.* B867 (2013), pp. 244–289. DOI: 10.1016/j.nuclphysb.2012.10.003. arXiv: 1207.1303 [hep-ph] (cit. on p. 62).

[102] Guido Altarelli and G. Parisi. "Asymptotic Freedom in Parton Language". In: *Nucl. Phys.* B126 (1977), pp. 298–318. DOI: 10.1016/0550-3213(77)90384-4 (cit. on pp. 62, 63).

[103] V. N. Gribov and L. N. Lipatov. "Deep inelastic e p scattering in perturbation theory". In: *Sov. J. Nucl. Phys.* 15 (1972). [Yad. Fiz.15,781(1972)], pp. 438–450 (cit. on p. 62).

[104] Yuri L. Dokshitzer. "Calculation of the Structure Functions for Deep Inelastic Scattering and e+ e- Annihilation by Perturbation Theory in Quantum Chromodynamics." In: *Sov. Phys. JETP* 46 (1977). [Zh. Eksp. Teor. Fiz.73,1216(1977)], pp. 641–653 (cit. on p. 62).

[105] L. A. Harland-Lang et al. "Parton distributions in the LHC era: MMHT 2014 PDFs". In: *Eur. Phys. J.* C75.5 (2015), p. 204. DOI: 10.1140/epjc/s10052-015-3397-6. arXiv: 1412.3989 [hep-ph] (cit. on pp. 63, 171).

[106] Bo Andersson et al. "Parton Fragmentation and String Dynamics". In: *Phys. Rept.* 97 (1983), pp. 31–145. DOI: 10.1016/0370-1573(83)90080-7 (cit. on p. 64).

[107] B. R. Webber. "A QCD Model for Jet Fragmentation Including Soft Gluon Interference". In: *Nucl. Phys.* B238 (1984), pp. 492–528. DOI: 10.1016/0550-3213(84)90333-X (cit. on p. 64).

[108] D. J. Lange. "The EvtGen particle decay simulation package". In: *Nucl. Instrum. Meth.* A462 (2001), pp. 152–155. DOI: 10.1016/S0168-9002(01)00089-4 (cit. on p. 64).

[109] Georges Aad et al. "Measurement of the underlying event in jet events from 7 TeV proton-proton collisions with the ATLAS detector". In: *Eur. Phys. J.* C74.8 (2014), p. 2965. DOI: 10.1140/epjc/s10052-014-2965-5. arXiv: 1406.0392 [hep-ex] (cit. on p. 64).

[110] Torbjorn Sjostrand, Stephen Mrenna, and Peter Z. Skands. "PYTHIA 6.4 Physics and Manual". In: *JHEP* 05 (2006), p. 026. DOI: 10.1088/1126-6708/2006/05/026. arXiv: hep-ph/0603175 [hep-ph] (cit. on p. 64).

[111] J. M. Butterworth, Jeffrey R. Forshaw, and M. H. Seymour. "Multiparton interactions in photoproduction at HERA". In: *Z. Phys.* C72 (1996), pp. 637–646. DOI: 10.1007/BF02909195,10.1007/s002880050286. arXiv: hep-ph/9601371 [hep-ph] (cit. on p. 64).

[112] Andy Buckley et al. "General-purpose event generators for LHC physics". In: *Phys. Rept.* 504 (2011), pp. 145–233. DOI: 10.1016/j.physrep.2011.03.005. arXiv: 1101.2599 [hep-ph] (cit. on p. 65).

[113] J. Alwall et al. "The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations". In: *JHEP* 07 (2014), p. 079. DOI: 10.1007/JHEP07(2014)079. arXiv: 1405.0301 [hep-ph] (cit. on p. 65).

[114] Stefano Frixione and Bryan R. Webber. "Matching NLO QCD computations and parton shower simulations". In: *JHEP* 06 (2002), p. 029. DOI: 10.1088/1126-6708/2002/06/029. arXiv: hep-ph/0204244 [hep-ph] (cit. on p. 65).

[115] Simone Alioli et al. "A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX". In: *JHEP* 06 (2010), p. 043. DOI: 10.1007/JHEP06(2010)043. arXiv: 1002.2581 [hep-ph] (cit. on p. 65).

[116] Stefano Frixione, Paolo Nason, and Carlo Oleari. "Matching NLO QCD computations with Parton Shower simulations: the POWHEG method". In: *JHEP* 11 (2007), p. 070. DOI: 10.1088/1126-6708/2007/11/070. arXiv: 0709.2092 [hep-ph] (cit. on p. 65).

[117] Torbjorn Sjostrand, Stephen Mrenna, and Peter Z. Skands. "A Brief Introduction to PYTHIA 8.1". In: *Comput. Phys. Commun.* 178 (2008), pp. 852–867. DOI: 10.1016/j.cpc.2008.01.036. arXiv: 0710.3820 [hep-ph] (cit. on p. 65).

[118] Torbjörn Sjöstrand et al. "An Introduction to PYTHIA 8.2". In: *Comput. Phys. Commun.* 191 (2015), pp. 159–177. DOI: 10.1016/j.cpc.2015.01.024. arXiv: 1410.3012 [hep-ph] (cit. on p. 65).

[119] G. Corcella et al. "HERWIG 6: An Event generator for hadron emission reactions with interfering gluons (including supersymmetric processes)". In: *JHEP* 01 (2001), p. 010. DOI: 10.1088/1126-6708/2001/01/010. arXiv: hep-ph/0011363 [hep-ph] (cit. on p. 65).

[120] Johannes Bellm et al. "Herwig 7.0/Herwig++ 3.0 release note". In: *Eur. Phys. J.* C76.4 (2016), p. 196. DOI: 10.1140/epjc/s10052-016-4018-8. arXiv: 1512.01178 [hep-ph] (cit. on p. 65).

[121] T. Gleisberg et al. "Event generation with SHERPA 1.1". In: *JHEP* 02 (2009), p. 007. DOI: 10.1088/1126-6708/2009/02/007. arXiv: 0811.4622 [hep-ph] (cit. on p. 65).

[122] Enrico Bothmann et al. "Event Generation with SHERPA 2.2". In: (2019). arXiv: 1905.09127 [hep-ph] (cit. on p. 65).

[123] Steffen Schumann and Frank Krauss. "A Parton shower algorithm based on Catani-Seymour dipole factorisation". In: *JHEP* 03 (2008), p. 038. DOI: 10.1088/1126-6708/2008/03/038. arXiv: 0709.1027 [hep-ph] (cit. on p. 65).

[124] D. J. Lange. "The EvtGen particle decay simulation package". In: *Nucl. Instrum. Meth. A* 462 (2001), p. 152. DOI: 10.1016/S0168-9002(01)00089-4 (cit. on p. 66).

[125] Pierre Artoisenet et al. "Automatic spin-entangled decays of heavy resonances in Monte Carlo simulations". In: *JHEP* 03 (2013), p. 015. DOI: 10.1007/JHEP03(2013)015. arXiv: 1212.3460 [hep-ph] (cit. on p. 66).

[126] Stefano Frixione et al. "Angular correlations of lepton pairs from vector boson and top quark decays in Monte Carlo simulations". In: *JHEP* 04 (2007), p. 081. DOI: 10.1088/1126-6708/2007/04/081. arXiv: hep-ph/0702198 [HEP-PH] (cit. on p. 66).

[127] S. Agostinelli et al. "GEANT4: A Simulation toolkit". In: *Nucl. Instrum. Meth.* A506 (2003), pp. 250–303. DOI: 10.1016/S0168-9002(03)01368-8 (cit. on p. 66).

[128] G. Aad et al. "The ATLAS Simulation Infrastructure". In: *Eur. Phys. J.* C70 (2010), pp. 823–874. DOI: 10.1140/epjc/s10052-010-1429-9. arXiv: 1005.4568 [physics.ins-det] (cit. on pp. 66, 67).

[129] T Cornelissen et al. "Concepts, Design and Implementation of the ATLAS New Tracking (NEWT)". In: ATL-SOFT-PUB-2007-007. ATL-COM-SOFT-2007-002 (Mar. 2007). URL: https://cds.cern.ch/record/1020106 (cit. on p. 70).

[130] "The Optimization of ATLAS Track Reconstruction in Dense Environments". In: ATL-PHYS-PUB-2015-006 (Mar. 2015). URL: https://cds.cern.ch/record/2002609 (cit. on p. 70).

[131] ATLAS Collaboration. "Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC". In: *Eur. Phys. J.* C77.5 (2017), p. 332. DOI: 10.1140/epjc/s10052-017-4887-5. arXiv: 1611.10235 [physics.ins-det] (cit. on p. 70).

[132] "Performance of primary vertex reconstruction in proton-proton collisions at $\sqrt{s}$ =7 TeV in the ATLAS experiment". In: ATLAS-CONF-2010-069 (July 2010). URL: https://cds.cern.ch/record/1281344 (cit. on p. 70).

[133] G Piacquadio, K Prokofiev, and A Wildauer. "Primary vertex reconstruction in the ATLAS experiment at LHC". In: *Journal of Physics: Conference Series* 119.3 (July 2008), p. 032033. DOI: 10.1088/1742-6596/119/3/032033. URL: https://doi.org/10.1088%2F1742-6596%2F119%2F3%2F032033 (cit. on p. 70).

[134] ATLAS Collaboration. "Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton-proton collision data at $\sqrt{s} = 13$ TeV". In: *Eur. Phys. J. C* 79.arXiv:1902.04655. 8 (Feb. 2019), 639. 40 p. DOI: 10.1140/epjc/s10052-019-7140-6. URL: https://cds.cern.ch/record/2657964 (cit. on pp. 73, 74, 75, 77).

[135] "Improved electron reconstruction in ATLAS using the Gaussian Sum Filter-based model for bremsstrahlung". In: ATLAS-CONF-2012-047 (May 2012). URL: https://cds.cern.ch/record/1449796 (cit. on p. 73).

[136] ATLAS Collaboration. "Electron and photon performance measurements with the ATLAS detector using the 2015-2017 LHC proton-proton collision data". In: (2019). arXiv: 1908.00005 [hep-ex] (cit. on pp. 73, 74, 75, 76).

[137] Georges Aad et al. "Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s}$ =13 TeV". In: *Eur. Phys. J.* C76.5 (2016), p. 292. DOI: 10.1140/epjc/s10052-016-4120-y. arXiv: 1603.05598 [hep-ex] (cit. on pp. 77, 80, 81).

[138] W. Lampl et al. "Calorimeter clustering algorithms: Description and performance". In: (2008) (cit. on pp. 82, 83).

[139] Georges Aad et al. "Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1". In: *Eur. Phys. J.* C77 (2017), p. 490. DOI: 10.1140/epjc/s10052-017-5004-5. arXiv: 1603.02934 [hep-ex] (cit. on p. 83).

[140] Javier Montejo Berlingen. "Search for new physics in $t\bar{t}$ final states with additional heavy-flavor jets with the ATLAS detector". Apr. 2016. DOI: 10.1007/978-3-319-41051-7. URL: https://cds.cern.ch/record/2053769 (cit. on p. 84).

[141] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. "The anti-k_t jet clustering algorithm". In: *JHEP* 04 (2008), p. 063. DOI: 10.1088/1126-6708/2008/04/063. arXiv: 0802.1189 [hep-ph] (cit. on pp. 84, 85).

[142] Matteo Cacciari, Gavin P. Salam, and Gregory Soyez. "The Catchment Area of Jets". In: *JHEP* 04 (2008), p. 005. DOI: 10.1088/1126-6708/2008/04/005. arXiv: 0802.1188 [hep-ph] (cit. on pp. 85, 87).

[143] M. Aaboud et al. "Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s}$ = 13 TeV with the ATLAS detector". In: *Phys. Rev.* D96.7 (2017), p. 072002. DOI: 10.1103/PhysRevD.96.072002. arXiv: 1703.09665 [hep-ex] (cit. on pp. 86, 88, 89, 91).

[144] Elmar Ritsch. "Fast Calorimeter Punch-Through Simulation for the ATLAS Experiment". Presented 28 Sep 2011. 2011. URL: https://cds.cern.ch/record/1388275 (cit. on p. 90).

[145] *Tagging and suppression of pileup jets with the ATLAS detector*. Tech. rep. ATLAS-CONF-2014-018. Geneva: CERN, May 2014. URL: https://cds.cern.ch/record/1700870 (cit. on pp. 91, 92, 93).

[146] ATLAS Collaboration. "ATLAS b-jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s}$ = 13 TeV". In: (2019). arXiv: 1907.05120 [hep-ex] (cit. on pp. 95, 98, 100).

[147] "Secondary vertex finding for jet flavour identification with the ATLAS detector". In: ATL-PHYS-PUB-2017-011 (June 2017). URL: https://cds.cern.ch/record/2270366 (cit. on p. 95).

[148] "Topological b-hadron decay reconstruction and identification of b-jets with the JetFitter package in the ATLAS experiment at the LHC". In: ATL-PHYS-PUB-2018-025 (Oct. 2018). URL: https://cds.cern.ch/record/2645405 (cit. on p. 96).

[149] R. Fruhwirth. "Application of Kalman filtering to track and vertex fitting". In: *Nucl. Instrum. Meth.* A262 (1987), pp. 444–450. DOI: 10.1016/0168-9002(87)90887-4 (cit. on p. 96).

[150] "Optimisation of the ATLAS b-tagging performance for the 2016 LHC Run". In: ATL-PHYS-PUB-2016-012 (June 2016). URL: http://cds.cern.ch/record/2160731 (cit. on pp. 97, 98).

[151] "TMVA – Toolkit for Multivariate Data Analysis". In: (). arXiv: physics/0703039 [physics.data-an] (cit. on p. 96).

[152] "Calibration of the ATLAS b-tagging algorithm in $t\bar{t}$ semi-leptonic events". In: ATLAS-CONF-2018-045 (Sept. 2018). URL: https://cds.cern.ch/record/2638455 (cit. on p. 98).

[153] "Calibration of light-flavour jet b-tagging rates on ATLAS proton-proton collision data at $\sqrt{s} = 13$ TeV". In: ATLAS-CONF-2018-006 (Apr. 2018). URL: https://cds.cern.ch/record/2314418 (cit. on p. 98).

[154] "Expected performance of missing transverse momentum reconstruction for the ATLAS detector at $\sqrt{s} = 13$ TeV". In: ATL-PHYS-PUB-2015-023 (July 2015). URL: https://cds.cern.ch/record/2037700 (cit. on p. 101).

[155] G Apollinari et al. *High-Luminosity Large Hadron Collider (HL-LHC): Preliminary Design Report*. CERN Yellow Reports: Monographs. Geneva: CERN, 2015. DOI: 10.5170/CERN-2015-005. URL: https://cds.cern.ch/record/2116337 (cit. on p. 102).

[156] T Kawamoto et al. *New Small Wheel Technical Design Report*. Tech. rep. CERN-LHCC-2013-006. ATLAS-TDR-020. ATLAS New Small Wheel Technical Design Report. June 2013. URL: https://cds.cern.ch/record/1552862 (cit. on pp. 103, 105, 106, 111, 113, 115).

[157] ATLAS Collaboration. *A study of the Pythia 8 description of ATLAS minimum bias measurements with the Donnachie-Landshoff diffractive model*. Tech. rep. ATL-PHYS-PUB-2016-017. Geneva: CERN, Aug. 2016. URL: https://cds.cern.ch/record/2206965 (cit. on p. 104).

[158] Y. Giomataris et al. "MICROMEGAS: A High granularity position sensitive gaseous detector for high particle flux environments". In: *Nucl. Instrum. Meth.* A376 (1996), pp. 29–35. DOI: 10.1016/0168-9002(96)00175-1 (cit. on p. 107).

[159] Tom Francke and Vladimir Peskov. *Micropattern Gaseous Detectors*. Tech. rep. physics/0404035. KTH-2004-04-07. Hamburg: Hamburg TU, Apr. 2004. URL: https://cds.cern.ch/record/729619 (cit. on p. 107).

[160] Konstantinos Ntekas. "Performance characterization of the Micromegas detector for the New Small Wheel upgrade and Development and improvement of the Muon Spectrometer Detector Control System in the ATLAS experiment". Presented 2016. 2016. URL: https://cds.cern.ch/record/2143887 (cit. on p. 114).

[161] Christos Bakalis. "Front-End and Back-End Electronics for the ATLAS New Small Wheel Upgrade". In: *PoS* ICHEP2018 (2019), p. 134. DOI: 10.22323/1.340.0134 (cit. on pp. 114, 115, 123).

[162] G. De Geronimo et al. "VMM1—An ASIC for Micropattern Detectors". In: *IEEE Transactions on Nuclear Science* 60.3 (June 2013), pp. 2314–2321. ISSN: 0018-9499. DOI: 10.1109/TNS.2013.2258683 (cit. on pp. 115, 116).

[163] George Iakovidis. "VMM - An ASIC for Micropattern Detectors". In: *EPJ Web of Conferences* 174 (Jan. 2018), p. 07001. DOI: 10.1051/epjconf/201817407001 (cit. on pp. 115, 116).

[164] Georgios Iakovidis. *VMM3, an ASIC for Micropattern Detectors*. Tech. rep. ATL-MUON-PROC-2018-003. Geneva: CERN, Mar. 2018. URL: https://cds.cern.ch/record/2309951 (cit. on pp. 115, 116, 118).

[165] Jinhong Wang et al. "Design of a Trigger Data Serializer ASIC for the Upgrade of the ATLAS Forward Muon Spectrometer". In: *IEEE Trans. Nucl. Sci.* 64.12 (2017), pp. 2958–2965. DOI: 10.1109/TNS.2017.2771266. arXiv: 1707.05451 [physics.ins-det] (cit. on p. 115).

[166] Martoiu, Sorin. In: (2014). URL: https://indico.cern.ch/event/327350/contributions/760806/attachments/635916/875419/ART_ASIC_review.pdf (cit. on p. 115).

[167] *Qt*. URL: https://www.qt.io/ (cit. on p. 126).

[168] Aranzazu Ruiz-Martinez, Fernando Monticelli, and Samuel David Jones. *Electron and photon trigger efficiency plots with early 2017 data*. Tech. rep. ATL-COM-DAQ-2017-066. Geneva: CERN, June 2017. URL: https://cds.cern.ch/record/2272031 (cit. on p. 150).

[169] Lidia Dell'Asta and Savanna Marie Shaw. *Muon trigger public plots with full 2016 dataset*. Tech. rep. ATL-COM-DAQ-2017-005. Geneva: CERN, Jan. 2017. URL: https://cds.cern.ch/record/2242824 (cit. on p. 150).

[170] "Estimation of non-prompt and fake lepton backgrounds in final states with top quarks produced in proton-proton collisions at $\sqrt{s}$=8 TeV with the ATLAS detector". In: ATLAS-CONF-2014-058 (Oct. 2014). URL: https://cds.cern.ch/record/1951336 (cit. on p. 161).

[171] ATLAS Collaboration. "Measurement of the $t\bar{t}$ production cross-section using $e\mu$ events with b-tagged jets in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: *Phys. Lett. B* 761 (2016), p. 136. DOI: 10.1016/j.physletb.2016.08.019. arXiv: 1606.02699 [hep-ex] (cit. on pp. 163, 169).

[172] ATLAS Collaboration. "Probing the quantum interference between singly and doubly resonant top-quark production in *pp* collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: (2018). arXiv: 1806.04667 [hep-ex] (cit. on pp. 163, 169).

[173] ATLAS Collaboration. "Luminosity determination in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector at the LHC". In: *Eur. Phys. J.* C76.12 (2016), p. 653. DOI: 10.1140/epjc/s10052-016-4466-1. arXiv: 1608.03953 [hep-ex] (cit. on p. 166).

[174] Richard D. Ball et al. "Parton distributions for the LHC Run II". In: *JHEP* 04 (2015), p. 040. DOI: 10.1007/JHEP04(2015)040. arXiv: 1410.8849 [hep-ph] (cit. on p. 171).

[175] Hung-Liang Lai et al. "New parton distributions for collider physics". In: *Phys. Rev. D* 82 (7 Oct. 2010), p. 074024. DOI: 10.1103/PhysRevD.82.074024. arXiv: 1007.2241 [hep-ph] (cit. on p. 171).

[176] Jon Butterworth et al. "PDF4LHC recommendations for LHC Run II". In: *J. Phys. G* 43 (2016), p. 023001. DOI: 10.1088/0954-3899/43/2/023001. arXiv: 1510.03865 [hep-ph] (cit. on p. 171).

[177] Michal Czakon, Paul Fiedler, and Alexander Mitov. "Total Top-Quark Pair-Production Cross Section at Hadron Colliders Through $O(\alpha_S^4)$". In: *Phys. Rev. Lett.* 110 (2013), p. 252004. DOI: 10.1103/PhysRevLett.110.252004. arXiv: 1303.6254 [hep-ph] (cit. on p. 172).

[178] ATLAS and CMS Collaborations. *Combination of ATLAS and CMS results on the mass of the top-quark using up to* 4.9 fb^{-1} *of* $\sqrt{s} = 7$ *TeV LHC data*. ATLAS-CONF-2013-102. 2013. URL: https://cds.cern.ch/record/1601811 (cit. on p. 172).

[179] Nikolaos Kidonakis. "Two-loop soft anomalous dimensions for single top quark associated production with a W^- or H^-". In: *Phys. Rev. D* 82 (2010), p. 054018. DOI: 10.1103/PhysRevD.82.054018. arXiv: 1005.4451 [hep-ph] (cit. on p. 172).

[180] Stefano Frixione et al. "Single-top hadroproduction in association with a W boson". In: *JHEP* 07 (2008), p. 029. DOI: 10.1088/1126-6708/2008/07/029. arXiv: 0805.3067 [hep-ph] (cit. on p. 172).

[181] ATLAS Collaboration. *Simulation of top-quark production for the ATLAS experiment at* $\sqrt{s} = 13$ *TeV*. ATL-PHYS-PUB-2016-004. 2016. URL: https://cds.cern.ch/record/2120417 (cit. on p. 172).

[182] Leif Lönnblad and Stefan Prestel. "Merging multi-leg NLO matrix elements with parton showers". In: *JHEP* 03 (2013), p. 166. DOI: 10.1007/JHEP03(2013)166. arXiv: 1211.7278 [hep-ph] (cit. on p. 172).

[183] Cousins, Robert. "Why isn't every physicist a Bayesian?" In: *American Journal of Physics* 63.398 (1995). DOI: 10.1119/1.17901 (cit. on p. 173).

[184] Alexander L. Read. "Presentation of search results: The CL(s) technique". In: *J. Phys.* G28 (2002). [,11(2002)], pp. 2693–2704. DOI: 10.1088/0954-3899/28/10/313 (cit. on p. 177).

[185] A L Read. "Modified frequentist analysis of search results (the CL_s method)". In: CERN-OPEN-2000-205 (2000). DOI: 10.5170/CERN-2000-005.81. URL: https://cds.cern.ch/record/451614 (cit. on p. 177).

[186] Kyle Cranmer et al. *HistFactory: A tool for creating statistical models for use with RooFit and RooStats*. Tech. rep. CERN-OPEN-2012-016. New York: New York U., Jan. 2012. URL: https://cds.cern.ch/record/1456844 (cit. on p. 179).

[187] ROOT Math Library Team. *Minuit2 Minimization Package*. URL: https://seal.web.cern.ch/seal/snapshot/work-packages/mathlibs/minuit/ (cit. on p. 179).

[188] Wouter Verkerke and David Kirkby. *The RooFit toolkit for data modeling*. 2003. arXiv: physics/0306116 [physics.data-an] (cit. on p. 179).

[189] Glen Cowan et al. "Asymptotic formulae for likelihood-based tests of new physics". In: *Eur. Phys. J.* C71 (2011). [Erratum: Eur. Phys. J.C73,2501(2013)], p. 1554. DOI: 10.1140/epjc/s10052-011-1554-0,10.1140/epjc/s10052-013-2501-z. arXiv: 1007.1727 [physics.data-an] (cit. on pp. 187, 188).

[190] ATLAS Collaboration. "ATLAS Run 1 searches for direct pair production of third-generation squarks at the Large Hadron Collider". In: *The European Physical Journal C* 75.10 (Oct. 2015), p. 510. ISSN: 1434-6052. DOI: 10.1140/epjc/s10052-015-3726-9. URL: https://doi.org/10.1140/epjc/s10052-015-3726-9 (cit. on p. 192).

[191] C. Borschensky et al. "Squark and gluino production cross sections in pp collisions at $\sqrt{s} = 13, 14, 33$ and 100 TeV". In: *The European Physics Journal C* 12 (2014). arXiv: 1407.5066 (cit. on p. 195).

[192] Benjamin Nachman. "Investigating the Quantum Properties of Jets and the Search for a Supersymmetric Top Quark Partner with the ATLAS Detector". PhD thesis. Stanford U., Phys. Dept., 2016. arXiv: 1609.03242 [hep-ex] (cit. on p. 197).

[193] Paul Jackson and Christopher Rogan. "Recursive Jigsaw Reconstruction: HEP event analysis in the presence of kinematic and combinatoric ambiguities". In: *Phys. Rev.* D96.11 (2017), p. 112007. DOI: 10.1103/PhysRevD.96.112007. arXiv: 1705.10733 [hep-ph] (cit. on p. 201).

[194] Matthew R. Buckley et al. "Super-Razor and Searches for Sleptons and Charginos at the LHC". In: *Phys. Rev.* D89.5 (2014), p. 055020. DOI: 10.1103/PhysRevD.89.055020. arXiv: 1310.4827 [hep-ph] (cit. on p. 201).

[195] Tom Melia. "Spin before mass at the LHC". In: *JHEP* 01 (2012), p. 143. DOI: 10.1007/JHEP01(2012)143. arXiv: 1110.6185 [hep-ph] (cit. on p. 211).

[196] S. Dawson, S. Dittmaier, and M. Spira. "Neutral Higgs boson pair production at hadron colliders: QCD corrections". In: *Phys. Rev.* D58 (1998), p. 115012. DOI: 10.1103/PhysRevD.58.115012. arXiv: hep-ph/9805244 [hep-ph] (cit. on p. 233).

[197] S. Borowka et al. "Higgs Boson Pair Production in Gluon Fusion at Next-to-Leading Order with Full Top-Quark Mass Dependence". In: *Phys. Rev. Lett.* 117.1 (2016). [Erratum: Phys. Rev. Lett.117,no.7,079901(2016)], p. 012001. DOI: 10.1103/PhysRevLett.117.079901,10.1103/PhysRevLett.117.012001. arXiv: 1604.06447 [hep-ph] (cit. on p. 233).

[198] Julien Baglio et al. "Gluon fusion into Higgs pairs at NLO QCD and the top mass scheme". In: *Eur. Phys. J.* C79.6 (2019), p. 459. DOI: 10.1140/epjc/s10052-019-6973-3. arXiv: 1811.05692 [hep-ph] (cit. on p. 233).

[199] Daniel de Florian and Javier Mazzitelli. "Higgs Boson Pair Production at Next-to-Next-to-Leading Order in QCD". In: *Phys. Rev. Lett.* 111 (2013), p. 201801. DOI: 10.1103/PhysRevLett.111.201801. arXiv: 1309.6594 [hep-ph] (cit. on p. 233).

[200] Ding Yu Shao et al. "Threshold resummation effects in Higgs boson pair production at the LHC". In: *JHEP* 07 (2013), p. 169. DOI: 10.1007/JHEP07(2013)169. arXiv: 1301.1245 [hep-ph] (cit. on p. 233).

[201] Daniel de Florian and Javier Mazzitelli. "Higgs pair production at next-to-next-to-leading logarithmic accuracy at the LHC". In: *JHEP* 09 (2015), p. 053. DOI: 10.1007/JHEP09(2015)053. arXiv: 1505.07122 [hep-ph] (cit. on p. 233).

[202] Massimiliano Grazzini et al. "Higgs boson pair production at NNLO with top quark mass effects". In: *JHEP* 05 (2018), p. 059. DOI: 10.1007/JHEP05(2018)059. arXiv: 1803.02463 [hep-ph] (cit. on p. 233).

[203] ATLAS Collaboration. "Search for pair production of Higgs bosons in the $b\bar{b}b\bar{b}$ final state using proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: *JHEP* 01 (2019), p. 030. DOI: 10.1007/JHEP01(2019)030. arXiv: 1804.06174 [hep-ex] (cit. on p. 234).

[204] ATLAS Collaboration. "Search for resonant and non-resonant Higgs boson pair production in the $b\bar{b}\tau^+\tau^-$ decay channel in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: *Phys. Rev. Lett.* 121.19 (2018). [Erratum: Phys. Rev. Lett.122,no.8,089901(2019)], p. 191801. DOI: 10.1103/PhysRevLett.122.089901,10.1103/PhysRevLett.121.191801. arXiv: 1808.00336 [hep-ex] (cit. on pp. 234, 279).

[205] ATLAS Collaboration. "Search for Higgs boson pair production in the $\gamma\gamma b\bar{b}$ final state with 13 TeV pp collision data collected by the ATLAS experiment". In: *JHEP* 11 (2018), p. 040. DOI: 10.1007/JHEP11(2018)040. arXiv: 1807.04873 [hep-ex] (cit. on p. 234).

[206] ATLAS Collaboration. "Search for Higgs boson pair production in the $WW^{(*)}WW^{(*)}$ decay channel using ATLAS data recorded at $\sqrt{s} = 13$ TeV". In: *JHEP* 05 (2019), p. 124. DOI: 10.1007/JHEP05(2019)124. arXiv: 1811.11028 [hep-ex] (cit. on p. 234).

[207] ATLAS Collaboration. "Search for Higgs boson pair production in the $\gamma\gamma WW^*$ channel using pp collision data recorded at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: *Eur. Phys. J.* C78.12 (2018), p. 1007. DOI: 10.1140/epjc/s10052-018-6457-x. arXiv: 1807.08567 [hep-ex] (cit. on p. 234).

[208] ATLAS Collaboration. "Search for Higgs boson pair production in the $b\bar{b}WW^*$ decay mode at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: *JHEP* 04 (2019), p. 092. DOI: 10.1007/JHEP04(2019)092. arXiv: 1811.04671 [hep-ex] (cit. on p. 234).

[209] ATLAS Collaboration. "Measurements of gluon-gluon fusion and vector-boson fusion Higgs boson production cross-sections in the $H \to WW^* \to e\nu\mu\nu$ decay channel in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: *Phys. Lett.* B789 (2019), pp. 508–529. DOI: 10.1016/j.physletb.2018.11.064. arXiv: 1808.09054 [hep-ex] (cit. on pp. 235, 237).

[210] ATLAS Collaboration. "Measurement of fiducial differential cross sections of gluon-fusion production of Higgs bosons decaying to WW$^* \to e\nu\mu\nu$ with the ATLAS detector at $\sqrt{s} = 8$ TeV". In: *JHEP* 08 (2016), p. 104. DOI: 10.1007/JHEP08(2016)104. arXiv: 1604.02997 [hep-ex] (cit. on pp. 235, 237).

[211] ATLAS Collaboration. "Observation and measurement of Higgs boson decays to WW* with the ATLAS detector". In: *Phys. Rev.* D92.1 (2015), p. 012006. DOI: 10.1103/PhysRevD.92.012006. arXiv: 1412.2641 [hep-ex] (cit. on pp. 235, 237).

[212] ATLAS Collaboration. "Combination of searches for Higgs boson pairs in pp collisions at $\sqrt{s}$ =13 TeV with the ATLAS detector". In: (2019). arXiv: 1906.02025 [hep-ex] (cit. on pp. 236, 283, 287).

[213] Alan Barr, Christopher Lester, and Phil Stephens. "A variable for measuring masses at hadron colliders when missing energy is expected; mT2 : the truth behind the glamour". In: *Journal of Physics G: Nuclear and Particle Physics* 29.10 (2003), p. 2343. URL: http://stacks.iop.org/0954-3899/29/i=10/a=304 (cit. on p. 242).

[214] Christopher G. Lester. "The stransverse mass, MT2, in special cases". In: *Journal of High Energy Physics* 2011.5 (May 2011), p. 76. ISSN: 1029-8479. DOI: 10.1007/JHEP05(2011)076. URL: https://doi.org/10.1007/JHEP05(2011)076 (cit. on p. 242).

[215] Daniel R. Tovey. "On measuring the masses of pair-produced semi-invisibly decaying particles at hadron colliders". In: *Journal of High Energy Physics* 2008.04 (2008), p. 034. URL: http://stacks.iop.org/1126-6708/2008/i=04/a=034 (cit. on p. 242).

[216] Chrisopher G. Lester and Benjamin Nachman. "Bisection-based asymmetric M_{T2} computation: a higher precision calculator than existing symmetric methods". In: *JHEP* 03 (2015), p. 100. DOI: 10.1007/JHEP03(2015)100. arXiv: 1411.4312 [hep-ph] (cit. on p. 242).

[217] Chollet et al. "Keras". In: (2015). URL: https://keras.io (cit. on p. 248).

[218] "TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems". In: (2015). Software available from tensorflow.org. URL: https://www.tensorflow.org/ (cit. on p. 248).

[219] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. http://www.deeplearningbook.org. MIT Press, 2016 (cit. on pp. 248, 257).

[220] Vinod Nair and Geoffrey E. Hinton. "Rectified Linear Units Improve Restricted Boltzmann Machines". In: *Proceedings of the 27th International Conference on International Conference on Machine Learning*. ICML'10. Haifa, Israel: Omnipress, 2010, pp. 807–814. ISBN: 978-1-60558-907-7. URL: http://dl.acm.org/citation.cfm?id=3104322.3104425 (cit. on p. 249).

[221] Leo Breiman. "Random Forests". In: *Machine Learning* 45.1 (Oct. 2001), pp. 5–32. ISSN: 1573-0565. DOI: 10.1023/A:1010933404324. URL: https://doi.org/10.1023/A:1010933404324 (cit. on p. 249).

[222] Nitish Srivastava et al. "Dropout: A Simple Way to Prevent Neural Networks from Overfitting". In: *JMLR* 15 (2014), pp. 1929–1958. URL: http://jmlr.org/papers/v15/srivastava14a.html (cit. on pp. 249, 250).

[223] Diederik P. Kingma and Jimmy Ba. "Adam: A Method for Stochastic Optimization". In: *CoRR* abs/1412.6980 (2014). arXiv: 1412.6980. URL: http://arxiv.org/abs/1412.6980 (cit. on p. 251).

[224] CMS Collaboration. "Measurement of the cross section and angular correlations for associated production of a Z boson with b hadrons in pp collisions at $\sqrt{s} = 7$ TeV". In: *JHEP* 12 (2013), p. 039. DOI: 10.1007/JHEP12(2013)039. arXiv: 1310.1349 [hep-ex] (cit. on pp. 267, 279).

[225] ATLAS Collaboration. "Measurement of differential production cross-sections for a Z boson in association with b-jets in 7 TeV proton-proton collisions with the ATLAS detector". In: *JHEP* 10 (2014), p. 141. DOI: 10.1007/JHEP10(2014)141. arXiv: 1407.3643 [hep-ex] (cit. on pp. 267, 279).

[226] Pierre Baldi et al. "Parameterized neural networks for high-energy physics". In: *Eur. Phys. J.* C76.5 (2016), p. 235. DOI: 10.1140/epjc/s10052-016-4099-4. arXiv: 1601.07913 [hep-ex] (cit. on pp. 283, 284).

[227] Soo Ryu. "FELIX: The new detector readout system for the ATLAS experiment". In: *Journal of Physics: Conference Series* 898 (Oct. 2017), p. 032057. DOI: 10.1088/1742-6596/898/3/032057. URL: https://doi.org/10.1088%2F1742-6596%2F898%2F3%2F032057 (cit. on p. 290).

[228] Jonathan L. Feng. "Naturalness and the Status of Supersymmetry". In: *Ann. Rev. Nucl. Part. Sci.* 63 (2013), pp. 351–382. DOI: 10.1146/annurev-nucl-102010-130447. arXiv: 1302.6587 [hep-ph] (cit. on p. 290).

[229] Stephen P. Martin. "A Supersymmetry primer". In: (1997). [Adv. Ser. Direct. High Energy Phys.18,1(1998)], pp. 1–98. DOI: 10.1142/9789812839657_0001,10.1142/9789814307505_0001. arXiv: hep-ph/9709356 [hep-ph] (cit. on p. 291).

[230] ATLAS Collaboration. *SUSY October 2019 Summary Plot Update*. Tech. rep. ATL-COM-PHYS-2019-1317. Geneva: CERN, Oct. 2019. URL: https://cds.cern.ch/record/2693810 (cit. on p. 291).