THE UNIVERSITY OF CHICAGO

PLAYING THE HIGGS FIELD: USING THE HIGGS BOSON TO SEARCH FOR NEW PHYSICS WITH THE ATLAS DETECTOR

A DISSERTATION SUBMITTED TO
THE FACULTY OF THE DIVISION OF THE PHYSICAL SCIENCES
IN CANDIDACY FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

DEPARTMENT OF PHYSICS

BY
JESSICA L. RAINBOLT

CHICAGO, ILLINOIS
AUGUST 2026

ORCID iD: https://orcid.org/0000-0001-7394-0464

For my cat, Petra, who has kept me company during many long days and late nights.

TABLE OF CONTENTS

LIST OF FIGURES . . . vii

LIST OF TABLES . . . x

ACKNOWLEDGMENTS . . . xi

ABSTRACT . . . xii

1 INTRODUCTION . . . 1

2 THEORETICAL BACKGROUND . . . 3
2.1 Forces and gauge theories . . . 4
2.1.1 The electromagnetic force . . . 4
2.1.2 The strong force . . . 5
2.1.3 The weak force . . . 6
2.2 The elementary particles . . . 6
2.2.1 Fermions . . . 7
2.2.2 Bosons . . . 10
2.3 The Higgs mechanism and the Higgs field . . . 12
2.4 Shortcomings of the SM . . . 13

3 THE EXPERIMENT . . . 16
3.1 The Large Hadron Collider . . . 18
3.2 The ATLAS detector . . . 20
3.2.1 Coordinate system . . . 21
3.2.2 Inner detector . . . 23
3.2.3 Calorimeters . . . 27
3.2.4 Other detectors . . . 30
3.2.5 Trigger system . . . 31

4 COMMON ANALYSIS INGREDIENTS . . . 33
4.1 The ATLAS Run 2 dataset . . . 33
4.2 Signal, background, and simulation . . . 35
4.3 Particle reconstruction and physics objects . . . 37
4.3.1 Primary vertex . . . 38
4.3.2 Electrons . . . 39
4.3.3 Photons . . . 40
4.3.4 Muons . . . 41
4.3.5 Jets . . . 42
4.3.6 Missing transverse momentum . . . 44
4.4 Jet tagging . . . 45
4.4.1 Vertex tagging . . . 45

4.4.2 Flavor tagging 46

5 SEARCH FOR INVISIBLE AND PARTIALLY INVISIBLE HIGGS BOSON DECAYS 49
5.1 Analysis strategy 51
5.2 Triggers, physics objects, and event selection 53
5.2.1 Triggers and physics objects 53
5.2.2 SR event selection 55
5.2.3 CR event selection 57
5.3 Background estimation 58
5.3.1 Jet smearing method for $\gamma +$ jet background 60
5.4 Systematic uncertainties 65
5.4.1 Theoretical uncertainties 65
5.4.2 Experimental uncertainties 66
5.4.3 Background modeling uncertainties 67
5.5 Statistical analysis and results 68
5.5.1 Search for EW $Z(\rightarrow \nu\nu)\gamma$jj 68
5.5.2 Search for $H \rightarrow$ inv. 71
5.5.3 Search for $H \rightarrow \gamma\gamma_{\mathrm{d}}$ 72
5.6 Conclusion 75

6 SEARCH FOR NEW SCALAR RESONANCES PRODUCED IN ASSOCIATION WITH A HIGGS BOSON 77
6.1 Analysis strategy 79
6.2 Triggers 82
6.2.1 Integrated luminosity for b jet triggers 83
6.3 Physics objects and event selection 84
6.3.1 Resolved channel event selection 85
6.3.2 Mixed channel event selection 86
6.4 Background estimation 88
6.4.1 GP + NF method 89
6.4.2 Gaussian process regression 90
6.4.3 Normalizing flows 92
6.4.4 Background method validation 97
6.5 Systematic uncertainties 98
6.5.1 Signal systematics 98
6.5.2 Background systematics 99
6.6 Statistical analysis and validation 100
6.7 Results 102
6.8 Conclusion 109

7 ATLAS TRIGGER SYSTEM UPGRADE FOR THE HL-LHC 111
7.1 Testing of high-speed track-fitting hardware 111
7.1.1 Lab setup 112
7.1.2 Tuning transmitter and receiver settings 114
7.1.3 Bit error rate test 118
7.1.4 Conclusion 119
7.2 High-level synthesis implementation of track-fitting firmware 119
7.2.1 Firmware design 120
7.2.2 HLS implementation 122
7.2.3 Conclusion 124

8 CONCLUSION 126

BIBLIOGRAPHY 128

LIST OF FIGURES

2.1 Table of elementary particles in the SM, organized by type [8]. 8
2.2 The potential of the Higgs field before and after it acquires a nonzero vacuum expectation value v [12]. 12

3.1 Overall view of the LHC [15]. 17
3.2 Plan for the LHC and HL-LHC, updated in January 2025 [16]. 17
3.3 Schematic layout of the LHC [18]. 19
3.4 Cutaway view of the ATLAS detector [24]. 21
3.5 ATLAS coordinate system [26]. 22
3.6 Pseudorapidity η for various values of the polar angle θ [28]. 23
3.7 Cutaway view of the ATLAS inner detector [30]. 24
3.8 Paths of two charged particles, one with $\eta = 1.4$ and one with $\eta = 2.2$, through the inner detector [25]. 25
3.9 Cutaway view of the ATLAS calorimeter system [34]. 28
3.10 Schematic diagrams of ATLAS calorimeter modules, showing the arrangement of active and passive layers [35, 36]. 29
3.11 Cutaway view of the ATLAS muon system [37]. 31

4.1 Cumulative luminosity versus time delivered to ATLAS, recorded by ATLAS, and classified as good for physics during stable beams for Run 2 pp collisions [41]. . 34
4.2 Luminosity-weighted distribution of the mean number of interactions per crossing for Run 2 pp collisions [44]. 35
4.3 A slice of the ATLAS detector with visualizations of how different particles are detected [46]. 39

5.1 Feynman diagrams representing the three signal processes [2]. 50
5.2 Feynman diagrams representing the major background processes [2]. 52
5.3 ATLAS display of the VBF + $E_{\mathrm{T}}^{\mathrm{miss}} + \gamma$ signature in an event consistent with the signal $H \to$ inv. and $H \to \gamma\gamma_{\mathrm{d}}$ processes [2]. 54
5.4 Background composition in the $H \to$ inv. SR. 60
5.5 Background composition in the $H \to \gamma\gamma_{\mathrm{d}}$ SR. 60
5.6 Fits of jet energy response R in simulated γ+jet events for each range of truth jet p_{T}. "Biased" represents the raw fit result, and "Corrected" represents the result after correcting the distribution so it peaks at $R = 1$. 62
5.7 Distributions of truth $E_{\mathrm{T}}^{\mathrm{miss}}$ and of $\Delta\phi$ between the truth $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ and the nearest jet in simulated γ + jet events. 63
5.8 Normalized distributions of $E_{\mathrm{T}}^{\mathrm{miss}}$, jet p_{T}, and m_{jj} for smeared and original (i.e., unsmeared) γ + jet events. 64
5.9 Post-fit results for all m_{jj} SR and CR bins in the EW $Z(\to \nu\nu)\gamma$jj cross section measurement. The lower panel shows the ratio of data to the sum of all predicted contributions and a comparison of the pre-fit and post-fit predictions [2]. 69

5.10 Post-fit results for all DNN SR and CR bins in the search for $H \to$ inv. The $H \to$ inv. signal is overlaid, scaled to a 37% $\mathcal{B}(H \to \text{inv.})$. The lower panel shows the ratio of data to the sum of all predicted contributions and a comparison of the pre-fit and post-fit predictions [2]. 73
5.11 Post-fit results for the ten $(m_{\text{jj}}, m_{\text{T}})$ bins in the SR for the $H \to \gamma\gamma_{\text{d}}$ search, and the CRs shown inclusively. The $H \to \gamma\gamma_{\text{d}}$ signal is overlaid for two mass hypotheses. The lower panel shows the ratio of data to the sum of all predicted contributions and a comparison of the pre-fit and post-fit predictions [2]. 74
5.12 Post-fit m_{T} distribution in the inclusive $H \to \gamma\gamma_{\text{d}}$ SR with two example mass hypotheses overlaid. Events with $m_{\text{T}} > 500\,\text{GeV}$ are included in the final bin [2]. 75
5.13 Expected and observed 95% CL upper limits on the product of the VBF production cross section and the $\gamma\gamma_{\text{d}}$ branching fraction for a scalar mediator of various masses. Theoretical predictions for a 5% branching fraction are overlaid [2]. . . 76

6.1 Feynman diagram of the $X \to SH \to bbbb$ process, initiated by gluon–gluon fusion [4]. 78
6.2 Diagrams representing the methods used to reconstruct the $S \to bb$ and $H \to bb$ decays in each channel. 80
6.3 The sensitivity of the resolved and mixed channels to the signal mass hypotheses, shown as points in the (m_S, m_X) plane [4]. 81
6.4 Mixed channel selection efficiencies for example signal samples with fixed $m_S = 750\,\text{GeV}$, as a function of the m_X hypothesis. The numbers are total efficiencies for consecutively applied selection steps. The "preselection" step requires one of the large-R jet, 2b1j, or 2b2j triggers to fire [4]. 87
6.5 Background composition in the mixed channel signal region. 88
6.6 Overview of the GP + NF method for background estimation [4]. 90
6.7 GP regression fit to the m_H^{reco} distribution for the mixed channel background estimate. The lighter colors indicate the CR included in the fit, while the darker colors indicate the SR not seen during the training [4]. 93
6.8 Selected kinematic distributions predicted by the GP + NF method for the mixed channel. The observable m_S^{reco} is directly predicted by the NF as $\ln\left(m_S^{\text{reco}}\right)$, while $\Delta\eta_{SH}^{\text{reco}}$ and m_X^{reco} are reconstructed from the GP + NF ensemble predictions [4]. 96
6.9 Distribution of m_X^{reco} in the mixed channel 1b VR. 98
6.10 Signal m_S^{reco} distributions for each m_S hypothesis assigned an exclusive region in the mixed channel, with fitted skew normal distributions overlaid [4]. 101
6.11 Nonclosure of the GP + NF background prediction in the mixed channel 1b VR for two m_S regions. Plots are pre-fit and the ratio is the nonclosure uncertainty assigned to the SR [4]. 103
6.12 Relative variations of background systematic uncertainties with respect to the nominal background prediction for two m_S regions in the mixed channel. The inclusive m_S region is used to derive the nonclosure uncertainty for the $m_S = 3\,\text{TeV}$ fits. Both panels show the non-closure uncertainty together with the correlated shape uncertainties derived from the ensemble of background predictions [4]. . 104

6.13 Example post-fit m_X^{reco} distribution in the mixed channel Loose GN2X validation region under the background-only hypothesis. The lower panel shows the ratio of the observed data to the background prediction [4]. 104
6.14 Example post-fit m_X^{reco} distributions in two mixed channel m_S regions under the background-only hypothesis. Example signal templates are overlaid and scaled by the corresponding expected upper limit. The lower panels show the ratio of the observed data to the background prediction [4]. 106
6.15 Expected and observed 95% CL upper limits on the product of the X production cross section and the $X \to SH \to 4b$ branching ratio as a function of m_S for each m_X hypothesis. Limits from the resolved and mixed channels are shown separately. For visual clarity, limits corresponding to different m_X hypotheses are multiplied by the factor indicated in the figure. Predictions from the TRSM are overlaid for selected m_X values [4]. 107
6.16 Assignment of the resolved and mixed channels to the final results for all signal mass hypotheses, shown as points in the (m_S, m_X) plane [4]. 108
6.17 Expected and observed 95% CL upper limits on the product of the X production cross section and the $X \to SH \to 4b$ branching ratio as a function of m_S for each m_X hypothesis. For mass hypotheses covered by both the resolved and mixed channels, only the result from the channel with the lower expected upper limit is retained. Predictions from the TRSM are overlaid for selected m_X values [4]. . 109

7.1 Lab setup for TFM serial link tests, with relevant components labeled. 114
7.2 Comparison of eye diagrams using the default and tuned transceiver settings for both TFM serial links. 117
7.3 Block diagram of trackfitter design. 121
7.4 Workflow for HLS trackfitter implementation. 123

LIST OF TABLES

2.1 The four fundamental forces and their relative strengths, listed in decreasing order [5]. . . . 3

4.1 Integrated luminosities for the Run 2 pp dataset, with results from both the 2019 and 2022 luminosity calibrations. The 2019 calibration reports a combined value for 2015 and 2016 [23, 42]. . . . 34

5.1 Object requirements for $E_{\mathrm{T}}^{\mathrm{miss}}$, photons, central and forward jets, electrons, and muons in the SR. . . . 55
5.2 Contributions from different groups of systematic uncertainties to the 1σ uncertainty bands of the best-fit $\mu_{Z\gamma_{\mathrm{EW}}}$ value and on the 95% CL upper limits on $\mathcal{B}(H \to \mathrm{inv.})$ and $\mathcal{B}(H \to \gamma\gamma_{\mathrm{d}})$ [2]. . . . 70

6.1 Jet p_{T} thresholds and online b tagging efficiencies for 2b1j and 2b2j triggers. . . . 82
6.2 Integrated luminosity of b jet HLT GRLs and standard GRLs for each year of Run 2 data collection. . . . 84
6.3 Object requirements for small-R and large-R jets. Except where noted explicitly, the small-R jet requirements are the same for both channels. The large-R jet requirements are for the mixed channel only. . . . 85
6.4 Signal and control region boundaries for each channel. . . . 85
6.5 Summary of the impact of the different sources of systematic uncertainty on the search sensitivity. . . . 110

7.1 Range of values for each transmitter and receiver setting [115]. . . . 115
7.2 Tuned transceiver settings for both TFM serial links. . . . 116
7.3 Dimensions of eye openings for diagrams generated using default and tuned transceiver settings. . . . 117

ACKNOWLEDGMENTS

I would like to thank all of the people in my life, past and present, who helped me make it to where I am today.

I would especially like to thank all of my collaborators, peers, and mentors at the University of Chicago for creating a lively and supportive research environment. I will always treasure the years I spent here with you.

I would like to thank Cecilia Tosciri for her tutelage and companionship during the long $SH \to 4b$ analysis journey.

I would like to thank Mel Shochet for his mentorship in both my hardware and analysis work.

Finally, I would like to thank my advisor, Young-Kee Kim, for her invaluable guidance, advice, support, and understanding.

ABSTRACT

This dissertation presents the Higgs boson as a tool to search for physics beyond the Standard Model through its interactions with hypothetical particles. Two searches are presented, both using proton–proton collision data collected by the ATLAS detector during Run 2 of the Large Hadron Collider. First, the Higgs boson's couplings to dark matter candidates are investigated by analyzing fully and partially invisible decay channels. Second, interactions between the Higgs and new high-mass scalar bosons are probed to explore extensions of the scalar sector. By constraining these couplings, this work aims to reveal signatures of new physics and guide future experimental searches at high-energy colliders.

CHAPTER 1

INTRODUCTION

After the ATLAS and CMS experiments announced the discovery of a "Higgs-like" particle with mass approximately 125 GeV on July 4, 2012 [1], their first order of business was to determine whether this particle was, in fact, *the* Higgs boson described by the Standard Model of particle physics. In the fourteen years since, both experiments have found increasingly convincing evidence that it was, and the community at large has become increasingly comfortable referring to the 125 GeV particle as "the Higgs." But the case is not yet closed. The properties of the Higgs boson that made it difficult to find also make it difficult to study. Current measurements of its fundamental properties, particularly its couplings to itself and other particles, are still relatively imprecise.

In other words, it is still entirely possible—even after over a decade of careful study—that the 125 GeV Higgs boson interacts with new, undiscovered particles. Such particles, if they exist, may address known shortcomings of the Standard Model and provide answers to outstanding questions in fundamental physics. The properties of these hypothetical particles, as well as the nature of their interactions with the Higgs boson, are unknown. Countless theoretical models of physics beyond the Standard Model suggest a vast array of possibilities.

This work uses the Higgs boson itself as a tool to search for these new particles. It describes two searches, both performed using proton–proton collision data collected by the ATLAS detector at the CERN Large Hadron Collider (LHC) during its second run from 2015 to 2018. Together, these searches target a variety of hypothetical particles that may interact with the Higgs boson. Both searches feature innovations in design and methodology that pave the way for future use of the Higgs boson as a discovery tool.

The first search, published in ref. 2, is a search for Higgs boson decays to particles that are invisible to the detector. These may be Standard Model neutrinos or exotic, undiscovered particles like dark matter. Its key innovation is the use of a Higgs production mechanism that includes an additional photon, which not only improves discrimination between signal and background but also opens up new avenues for interpretation; in particular, Higgs boson decays to a photon and a so-called dark photon. This search also reports the first observation of one of the backgrounds to the invisible Higgs decay signals: the electroweak production of a Z boson and a photon produced in association with two jets, with the Z boson decaying invisibly. This observation has been followed by a second ATLAS measurement in ref. 3.

The second search, published in ref. 4, is a search for an exotic Higgs production mechanism involving two new scalar particles. These new particles may be part of an "extended Higgs sector," which is postulated by many popular extensions of the Standard Model. In the target process, one of the new particles decays to a second new particle and a Higgs boson, and both the Higgs boson and the second new particle decay to a bottom quark–antiquark pair. This search's key innovation is the use of a novel machine learning method to estimate contributions from background processes. The method, which combines Gaussian process regression and normalizing flows, simulates background contributions via sampling from a probability density learned from real data. This search reports upper limits on this process for a wide range of hypothetical particle masses.

This dissertation is organized as follows. Chapter 2 provides a brief introduction to the Standard Model and the Higgs field, focusing on the aspects most relevant to the searches. Chapter 3 is an overview of the LHC and the ATLAS detector. Chapter 4 describes features of ATLAS analyses that are common to both searches. The search for invisible and partially invisible Higgs boson decays is presented in chapter 5, and the search for new scalar particles produced in association with a Higgs boson is presented in chapter 6. Chapter 7 describes two upgrade projects for the ATLAS trigger system in anticipation of the High-Luminosity LHC, during which the size of the ATLAS dataset will increase by an order of magnitude. Finally, chapter 8 concludes the work.

CHAPTER 2

THEORETICAL BACKGROUND

The Standard Model (SM) of particle physics describes the most fundamental components of the known universe. It describes the elementary particles that make up matter and how they interact with one another via three of the four fundamental forces: the electromagnetic, strong, and weak forces. In general, the SM's particles have mass and therefore are also subject to the fourth fundamental force—the gravitational force—but it is not present in our current picture of the SM. At the elementary particle scale, the gravitational force is so weak in comparison to the other three that the SM is self-consistent without it. Table 2.1 lists the four fundamental forces and their relative strengths.

This chapter presents a brief description of the SM and the Higgs field to motivate the searches presented in this dissertation. Section 2.1 introduces the three fundamental forces included in the SM and the gauge theories that describe them. Section 2.2 describes each of the SM's elementary particles. Section 2.3 describes the Higgs field and the Higgs mechanism. Finally, section 2.4 describes some of the shortcomings of the SM and motivations for new physics searches.

Table 2.1: The four fundamental forces and their relative strengths, listed in decreasing order [5].

Force	Relative strength
Strong	10
Electromagnetic	10^{-2}
Weak	10^{-13}
Gravitational	10^{-42}

2.1 Forces and gauge theories

The fundamental forces in the SM are described by gauge theories. A gauge theory is a field theory that is invariant under local transformations of "gauge," a word which has the meaning of "length" or "calibration" [6], but in quantum field theory has come to refer to redundant degress of freedom. The stipulation that the Lagrangian characterizing any gauge theory must be invariant under local gauge transformations gives rise to gauge fields that are quantized as gauge bosons, described in the next section, that carry the theory's force. Gauge transformations are symmetry operations defined by mathematical groups, called Lie groups [7]. Each symmetry group is associated with a "charge," which is a quantum number that is conserved in all SM interactions.

This section gives a brief overview of each of the three fundamental forces described by the SM to lay the groundwork for the following sections. The electromagnetic force is described in section 2.1.1, the strong force in section 2.1.2, and the weak force in section 2.1.3.

2.1.1 The electromagnetic force

The electromagnetic (EM) force is described by the gauge theory of quantum electrodynamics (QED). The EM force acts on particles with electric charge Q, which takes positive and negative scalar values. Electric charge generates the symmetry group of QED, the unitary group $U(1)$, which is the set of all complex numbers with magnitude 1. This means QED is invariant under phase rotations $\psi \to e^{i\alpha}\psi$. Its gauge field is the electromagnetic field, whose physical manifestation is the photon. Because the $U(1)$ group is "abelian," meaning all of its elements commute, the photon itself does not carry electric charge and therefore does not interact with itself [7].

The EM force is responsible for virtually all physical phenomena we observe in our daily lives, aside from the gravitational force that anchors us to the Earth. It is responsible for our sense of sight, our sense of touch, and so on, all the way down to the electric impulses that link our metaphysical consciousness to physical reality. It is also responsible for the structure of everyday matter, since it attracts negatively charged electrons to positively charged atomic nuclei, and also assembles the resulting atoms into molecules.

2.1.2 The strong force

The strong force is described by the gauge theory of quantum chromodynamics (QCD). It acts on particles with color charge, which has three components: redness (r), blueness (b), and greenness (g). The three colors are the key ingredients to symmetry under the special unitary group $SU(3)$, which is the set of all unitary matrices with determinant 1. $SU(3)$ can be represented by the eight 3×3 Gell-Mann matrices, which do not commute with one another. As a result, QCD is nonabelian, and its gauge field, the gluon field, carries color charge and interacts with itself [7].

The phenomenon of “color confinement” is related to the nonabelian nature of QCD and the great strength of the strong force. Only color-neutral particles can exist in isolation. Particles with nonzero color charge cannot; when they are pulled apart, the strong force continuously links them together (rather than falling off, the like EM force does at long distances), and it is more energentically favorable for new color-charged particles to be created in the space between them than for them to continue moving apart. QCD alone does not conclusively imply color confinement, but this phenomenon has real physical and experimental consequences nonetheless [5].

The strong force is less accessible at the human scale than the EM force, but no less ubiquitous. It binds together the constituent particles of protons and neutrons, and indirectly binds protons and neutrons themselves within atomic nuclei. The strong force creates virtually all of the mass of macroscopic objects due to the strength of its interaction.

2.1.3 The weak force

The weak force is more difficult to describe than the EM and strong forces because it arises from a broken symmetry. The weak and EM forces are, in fact, both consequences of another gauge theory: electroweak (EW) theory. Above the Fermi scale of 246 GeV, the EM and weak forces merge into the singular EW force, which has $SU(2) \times U(1)$ symmetry. $SU(2)$ can be represented by the three 2×2 Pauli matrices, and its generator is "weak isospin" T, which has three components. The generator of the $U(1)$ symmetry is given the name "weak hypercharge," Y, and has one component. The result is a total of four gauge fields.

Below the Fermi scale, the EW symmetry is spontaneously broken via the Higgs mechanism as described in section 2.3, and the only unbroken remnant of the $SU(2) \times U(1)$ symmetry is the $U(1)$ symmetry of QED, whose charge Q is a linear combination of Y and one of the components of T [7]. The gauge fields of the weak force arise from a linear combination of the $SU(2) \times U(1)$ generators. They are associated with two different kinds of interactions: charged- and neutral-current interactions. The weak force is nonabelian, and its gauge fields, which manifest as the W and Z bosons, interact with themselves and with one another.

The most familar physical expression of the weak force is radioactive decay; for example, beta decay is a charged-current interaction. The weak force also plays an important role in the proton–proton chain of nuclear fusion within stars, when pairs of hydogen nuclei—that is, protons—convert to the proton–neutron pairs that themselves pair up to become helium-4 nuclei. The weak force is, threfore, no less critical to our continued existence than the others.

2.2 The elementary particles

The SM includes, depending on how one counts, at least 17 unique fundamental particles. No matter how one counts, they may be divided into two broad, mutually exclusive categories: fermions and bosons. Fermions comprise matter, and bosons mediate the fermions' interactions. Particles of both types are excitations or "waves" in quantum fields, which are formally described by quantum field theory. Whether it is more appropriate to think of particles or fields is a matter of perspective. In the context of high-energy LHC collisions, the discrete particle perspective is the most natural.

All particles have antimatter counterparts, called antiparticles. Particle–antiparticle pairs annihilate—they are waves that, upon meeting, cancel one another out. To mentally transform a matter particle into its antimatter counterpart, simply multiply its quantum numbers by -1. All other properties, including mass, remain unchanged. Particles whose quantum numbers are all 0 are their own antiparticles.

Figure 2.1 organizes the minimal 17 particles into a table, which may be used as a reference for the discussion in this section. The table omits not only antiparticles, but also the different color varieties of the strongly interacting particles. In the picture of the SM we are able to resolve with the ATLAS detector, described in later chapters, it is impossible to differentiate between particles of different colors and even between many particle–antiparticle pairs, so the simplified table is perhaps appropriate.

This section describes the SM's elementary particles and their basic properties. Section 2.2.1 describes the fermions. Section 2.2.2 describes the bosons, including a brief description of the Higgs boson. The Higgs field is discussed in further detail in section 2.3.

2.2.1 Fermions

Fermions are responsible for the small-scale structure of matter. Their intrinsic angular momentum, called "spin," takes on values that are half-integer multiples of the reduced Planck constant, $\hbar$. Because of this property, they are subject to the Pauli exclusion principle, which states that no two fermions can occupy the same state. This restriction stratifies the electrons surrounding atomic nuclei into distinct orbitals, which in turn dictate the properties

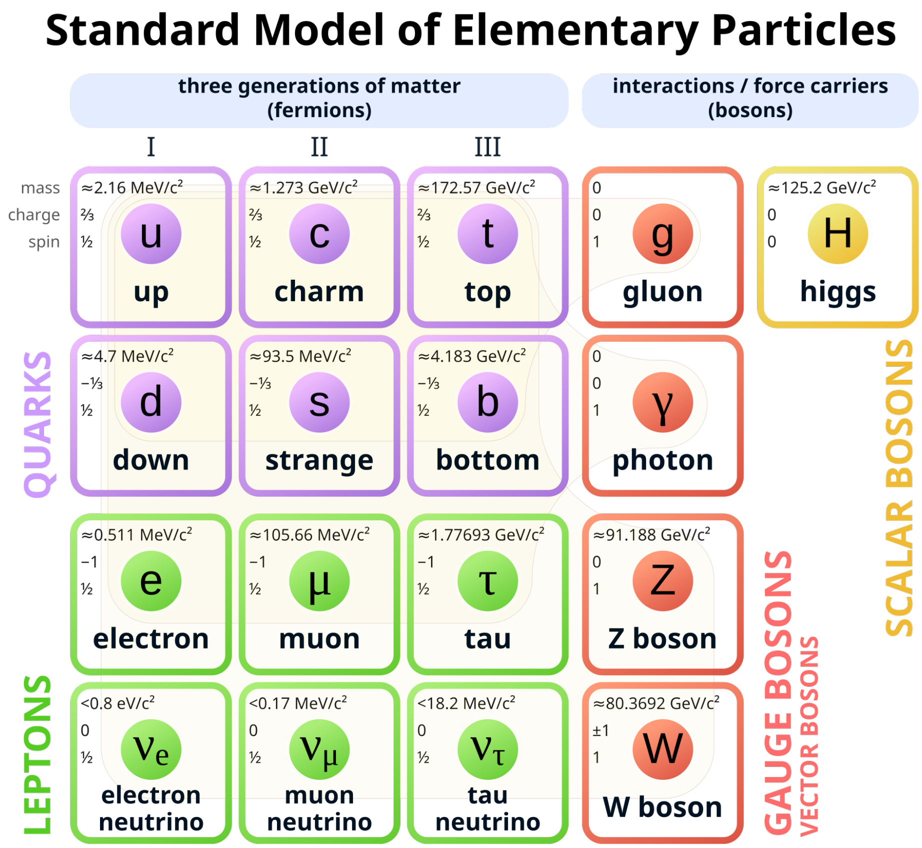

Figure 2.1: Table of elementary particles in the SM, organized by type [8].

of atomic bonds, which in turn dictate the structure of molecules, and so on. The SM contains two groups of elementary fermions, both taking spin quantum numbers $\pm\frac{1}{2}$.

Quarks

The first group of fermions, the "quarks," form composite particles called "hadrons" that are bound together by the strong force. By far the most famous hadrons are the proton and the neutron, which comprise atomic nuclei. Both are composed of up (u) quarks and down (d) quarks, with the proton consisting of two u quarks and one d quark, and the neutron consisting of one u quark and two d quarks. The quark content of the proton and neutron are responsible for their electric charge. The u quark has charge $+\frac{2}{3}$, and the d

quark has charge $-\frac{1}{3}$, with charge represented in units of the elementary charge e of the proton. Adding up the charges of the constituent quarks results in a total charge of +1 for the proton and 0 for the neutron, as expected.

The u and d quarks are only two of the six quark flavors, and as a pair they form the first of three quark "generations." Each of the remaining two generations also consists of an "uplike" quark with charge $+\frac{2}{3}$ and a "downlike" quark with charge $-\frac{1}{3}$. The second generation consists of the uplike charm (c) quark and the downlike strange (s) quark, and the third of the uplike top (t) quark and the downlike bottom (b) quark. The generations are ordered according to their mass, with the lightest first-generation quarks being the most stable and thus comprising the most stable hadrons in the universe. The neutron's lifetime is about 15 minutes outside of a nucleus, and even an isolated proton's lifetime is longer than 10^{29} years—a number nearly 20 orders of magnitude greater than the age of the universe [9].

Thanks to color confinement, quarks cannot be isolated and studied outside of hadrons. However, one quark, the t quark, is unique in that it is too heavy to form hadrons (although recent experimental evidence [10] suggest an extremely short-lived t quark–antiquark bound state has been observed). With mass 173 GeV, the t quark is the heaviest fundamental particle, weighing as much as a gold atom [11]. The t quark's large mass makes it very unstable. Its decay width is about 1.4 GeV [9], which corresponds to a lifetime of about 5×10^{-25} s.

Leptons

The second group of fermions, the "leptons," is home to the most familiar elementary particle: the electron. The leptons as a group were named originally for the electron's light weight in comparison to particles we now know are not elementary, such as the proton and the neutron. The name stuck, although the leptons are now characterized by their lack of participation in strong interactions rather than by their mass [5].

Leptons are subdivided into two groups of three. The charged leptons—the electron (e), the muon (μ), and the tau (τ)—interact via the weak and electromagnetic forces. They all have electric charge -1. The electron is stable, while the substantially more massive muon and τ have lifetimes of about $2\,\mu s$ and $0.3\,ps$, respectively. The neutral leptons are called neutrinos (ν), which interact only via the weak force. Neutrinos are incredibly light; they are actually massless within the SM, but experimental evidence suggests they have mass, and that the different flavors of neutrino have different masses [9]. The lack of mass for neutrinos is just one of the shortcomings of the SM that will be discusssed in section 2.4.

Each charged lepton has a partner neutrino with which it forms a generation. The leptons in each generation have a lepton number that is conserved in interactions: the electron and electron neutrino have electron number of 1, and likewise for the other generations. Correspondingly, antileptons have negative lepton number. Decays that produce a single charged lepton, such as the neutron's decay to a proton and an electron, conserve lepton number through the production of an antineutrino. In this example, $n \to p^+ e^- \overline{\nu}_e$.

2.2.2 Bosons

Bosons mediate the interactions between fermions. They have integer spin and do not obey the Pauli exclusion principle. In other words, they control rather than comprise matter. The SM includes bosons of two types: the gauge bosons and the Higgs boson.

Gauge bosons

When two or more fermions interact via one of the three fundamental forces in the SM, they exchange a "gauge" boson corresponding to that force. The name comes the fact that these bosons are physical manifestations of the gauge theories that describe the interactions they mediate, as explained in the previous section. The gauge bosons in the SM are vector bosons, so called because their corresponding fields are vector fields. They have spin 1, meaning their spin can take the values ± 1 and 0. The gauge bosons in the SM are the gluons, the photon, and the W and Z bosons.

Gluons mediate the strong interaction. They are massless and electrically neutral, and interact only via the strong force. Gluons come in eight varieties, each carrying a linear combination of color and anticolor charges; for example, $(r\bar{b} + \bar{r}b)/\sqrt{2}$. The linear combinations make gluons their own antiparticles. Gluons are stable since they are massless, but because of color confinement, they are very interactive.

The photon mediates the electromagnetic interaction. It is colorless, meaning it is unreactive to the strong force, as well as massless and electrically neutral. These properties make it its own antiparticle. It is also perfectly stable.

The W and Z bosons mediate the charged- and neutral-current weak interactions, respectively. Both are colorless. There are two types of W boson, W^+ and W^-, which have electric charge $+1$ and -1, respectively, and form a particle–antiparticle pair. The Z boson is electrically neutral and its own antiparticle. Unlike the other gauge bosons, the W and Z are massive, and quite so, with masses similar to rubidium and molybdenum atoms, both of which weigh close to 100 GeV. As such, they are unstable. Both have decay widths close to 2 GeV, giving them each a lifetime of about 3×10^{-25} s.

Higgs boson

The Higgs boson is the only SM boson that is not a gauge boson. It is a scalar boson, which means it has spin 0 and its corresponding field, the Higgs field, is a scalar field. It interacts with all massive particles. Its mass is, rather famously, 125 GeV—about the same as a cesium atom. Its decay width is about 3 MeV, making its lifetime about 2×10^{-22} s. It is electrically neutral, colorless, and its own antiparticle.

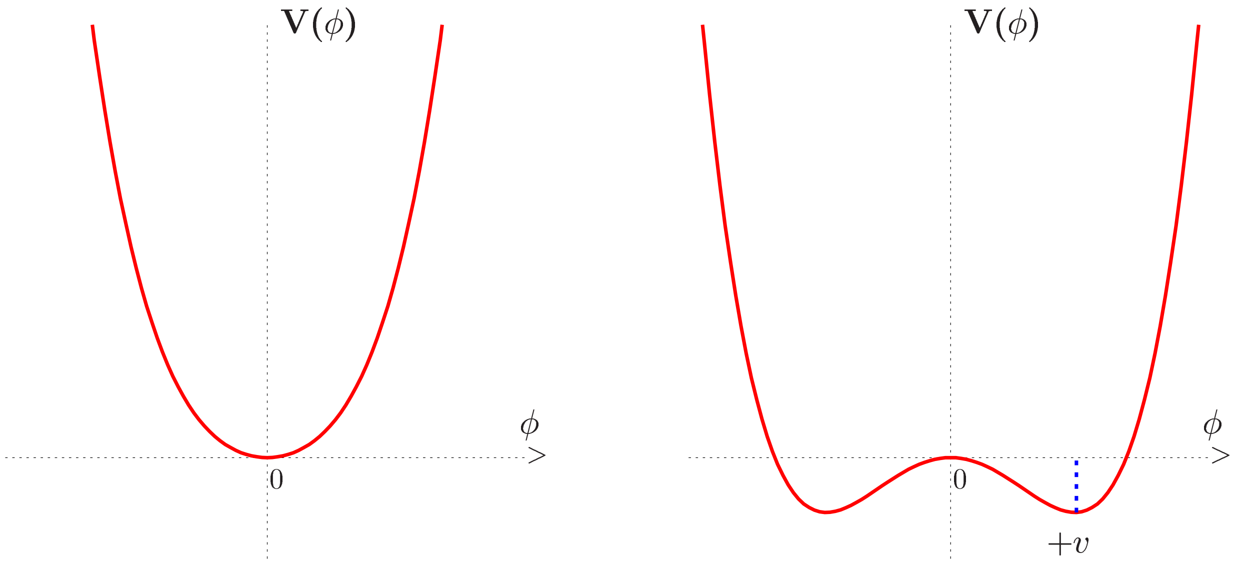

(a) Above Fermi scale

(b) Below Fermi scale

Figure 2.2: The potential of the Higgs field before and after it acquires a nonzero vacuum expectation value v [12].

2.3 The Higgs mechanism and the Higgs field

As anticipated in section 2.1.3, the Higgs mechanism breaks the EW gauge symmetry. This phenomenon is important not only because it gives rise to the familiar EM and weak forces, but also because it creates mass. Without the Higgs field, none of the elementary particles in the SM would have mass, and the physical manifestation of the SM would look very different from the one we are familiar with.

The Higgs field has a potential whose shape depends on its energy scale. At high energy scales, the potential has a stable minimum, and it is most energetically favorable for the Higgs field's vacuum expectation value v to be zero. But as the energy scale decreases, the shape of the potential changes, and the stable minimum becomes an unstable local maximum. In other words, v becomes nonzero; specifically, it becomes 246 GeV [9]. This shift is illustrated diagrammatically in fig. 2.2. The low-energy potential is often called a "sombrero" potential thanks to its shape.

The Higgs field's nonzero vacuum expectation value has serious consequences for EW theory. Since the Higgs field is an $SU(2)$ doublet, each of its two components is a complex scalar field. It has nonzero weak isospin and weak hypercharge, so it interacts with the EW gauge bosons. As the Higgs field acquires v, it effectively chooses a preferred direction in $SU(2)$ space. This breaks the EW gauge symmetry. The $U(1)$ symmetry of QED is preserved because the electric charge Q is a linear combination of the symmetrical $SU(2)$ direction and the unbroken EW $U(1)$ symmetry [12].

The Goldstone theorem states that, whenever a continuous symmetry is spontaneously broken, one massless, spin 0 "Goldstone" boson is created for each broken generator. Goldstone bosons add additional degrees of freedom to the Lagrangian. But the EW symmetry is a local gauge symmetry, so these degrees of freedom, like so many others in the SM, can be eliminated via a gauge transformation. In the unitary gauge, the W^+, W^-, and Z bosons each absorb one of the three Goldstone bosons created from the three broken EW symmetry generators, and the added degrees of freedom manifest as their masses. The same gauge transformation extends to the Yukawa term of the SM Lagrangian, which describes how the fermions couple to the Higgs field, and generates the fermion masses [12].

2.4 Shortcomings of the SM

The SM very successfully describes many properties of the elementary particles and three of the four fundamental forces. But even this description implies that at least something is lacking, and in fact there are many aspects of the SM that are unsatisfactory. This section is an inexhaustive list of some of its shortcomings, which motivate the searches for new physics presented later in this dissertation.

The SM includes mass, which determines the strength of gravitational interactions, but has nothing to say about gravity itself. This is almost stranger than a theory of completely massless particles—some, but not all, of the ingredients are there. Hypothetically, the gravitational force is mediated by a gauge boson with spin 2, but current theories including gravitons cannot be renormalized, so they cannot be added to the neat framework of the SM [7].

Another point of fault is that some parameters of the SM can only be determined empirically—that is, they must be measured and are not predicted by the theory itself. Examples include the values of the fermion and Higgs boson masses. Additionally, the SM does not provide any explanation for the nonzero neutrino masses. Perhaps neutrinos acquire their mass via the Higgs field, like the other fermions, or perhaps they do not. The mass of the Higgs boson itself is a particularly puzzling quandary. The "hierarchy problem" in particle physics asks why the Higgs boson mass is only 125 GeV when it should be senstive to quantum corrections from physics at energy scales all the way up to the point at which new physics, such as gravity, enters the picture.

Related to the topic of gravitation and mass is the SM's lack of dark matter, which is hypothesized to make up 85% of the universe's mass. Dark matter is so called because it is electromagnetically inert. Its primary interaction with light matter seems to be gravitational, but weak interactions have not been conclusively ruled out. Additionally, since dark matter does have mass, it is natural to ask where this mass comes from. If it comes from the Higgs field, dark matter must interact with the Higgs boson. This question is addressed in chapter 5, which describes the first search presented in this dissertation.

Finally, the SM does not explain the matter–antimatter asymmetry in the observable universe. Why is it made up almost entirely of matter, as opposed to antimatter? In the SM, matter and antimatter are always created and destroyed at equal rates, so one would expect the early universe to contain equal amounts of both. If this was the case, what happened to all of the antimatter between the early universe and today? If it was not the case, and the universe has always been asymmetrical in this sense, then why? The SM cannot provide answers for either of these questions.

One of the primary goals of experimental particle physics is find answers to these and other questions in the form of physics beyond the SM (BSM). The searches presented in this dissertation are concrete steps toward that goal.

CHAPTER 3

THE EXPERIMENT

The ATLAS detector is one of the four major experiments at the Large Hadron Collider (LHC), a particle accelerator operated by the European Organization for Nuclear Research, which is known by the historical acronym CERN [13]. The CERN laboratory is located on the border of France and Switzerland, near Geneva, and the LHC lies in a tunnel about 100 km (300 ft) below. The tunnel is about 27 km (17 mi) in circumference [14], making its outer diameter about 8.6 km (5.3 mi). Figure 3.1 is a diagram of the LHC showing its geographic location and basic dimensions. Also shown are three smaller accelerators (that deliver protons to the LHC), the ATLAS detector, and the other three major detectors: CMS, LHCb, and ALICE. CMS is another general-purpose detector that competes with ATLAS, while LHCb and ALICE each target different, more specific physics goals.

The LHC is exactly what its name suggests—a large machine that collides hadrons. The hadrons in question are predominantly protons, although the LHC also has operational periods dedicated to colliding heavier ions, like lead nuclei. The ATLAS detector is a general-purpose particle detector that is designed primarily to record proton–proton (pp) collisions, although it records heavy ion collisions as well. However, only pp collisions are relevant to the searches presented in this dissertation.

The experimental timeline of the LHC alternates between "run" periods, during which collision data are recorded, and "long shutdown" periods, during which the machines are maintained and upgraded. The first phase of the LHC spanned three runs, aptly named Run 1, Run 2, and Run 3. The second phase of the LHC, the High-Luminosity LHC (HL-LHC), will begin after the long shutdown following Run 3, and will span Run 4 and Run 5. A sketch of the overall timeline for the LHC, including the HL-LHC phase, is shown in fig. 3.2. The current forecast for the total pp integrated luminosity is $3000\,\text{fb}^{-1}$ by the end of 2041.

This chapter gives a brief overview of the LHC and the ATLAS detector, focusing on the

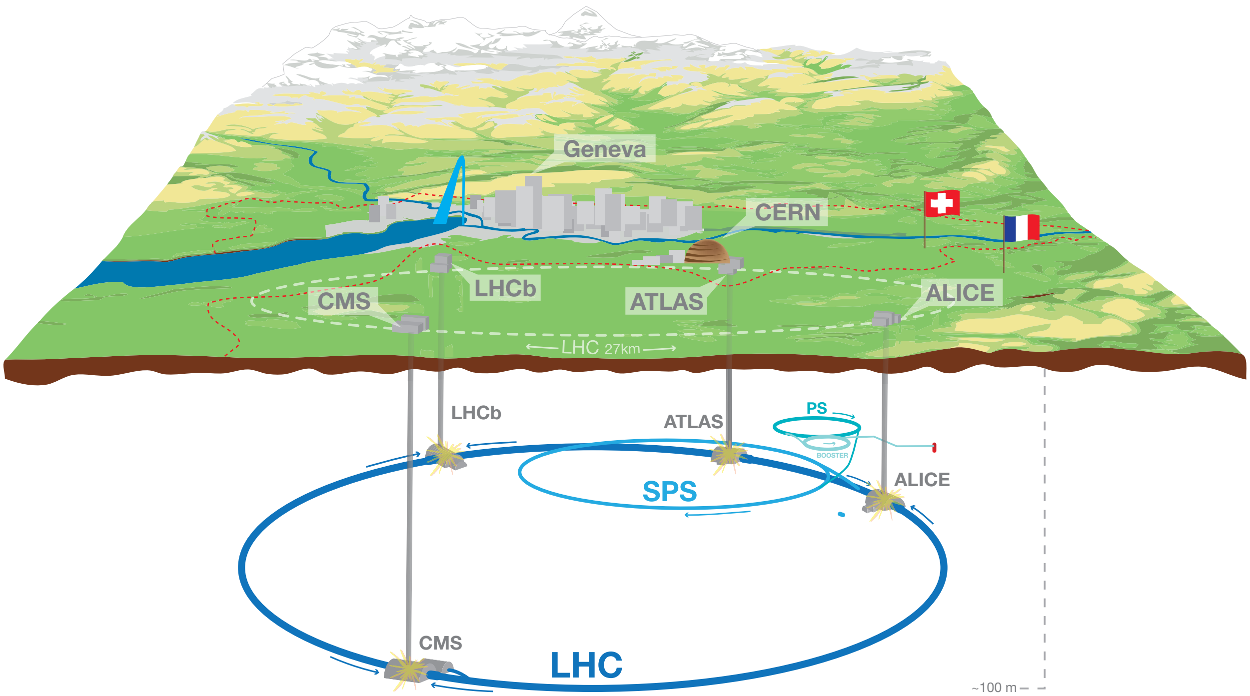

Figure 3.1: Overall view of the LHC [15].

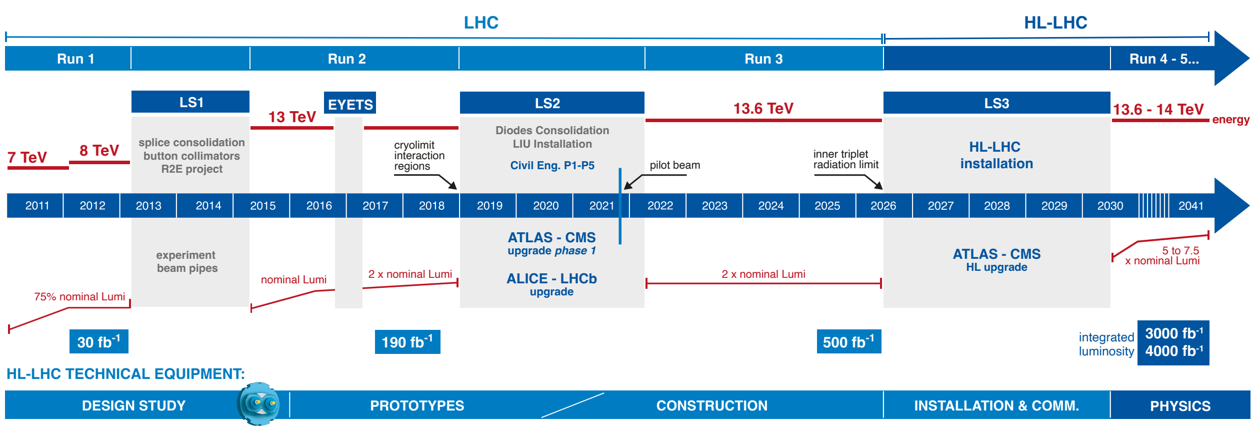

Figure 3.2: Plan for the LHC and HL-LHC, updated in January 2025 [16].

aspects most relevant to the work presented in this dissertation. The operation of the LHC is summarized in section 3.1, and the ATLAS detector is described in section 3.2.

3.1 The Large Hadron Collider

The LHC consists of two rings that accelerate beams of protons in opposite directions [14]. The tunnel is made of eight arcs and eight straight sections that are divided into "octants," numbered clockwise in accordance with the tunnel's eight access points, which have shafts connecting it to the surface [17]. Each octant has a straight section, called an insertion (IR), which is outfitted for a particular purpose [18]. Four of the IRs have an interaction point (IP) where the beams cross, coinciding with the locations of the four major detectors. The ATLAS detector is located at IP1, which is just below the CERN laboratory. The ALICE, CMS, and LHCb detectors are located at IP2, IP5, and IP8, respectively. IR2 and IR8 also include the injection systems for the two beams. The remaining IRs are used to maintain the beams. IR3 and IR7 house cleaning systems that are used to remove protons that have strayed from the beams, IR4 houses the system of radiofrequency (RF) cavities that accelerate the beams, and IR6 houses the beam dumping system that is used to remove the beams from the rings. A diagram of the LHC's layout, including the ring design and the function of each of the IRs, is shown in fig. 3.3.

The LHC's protons begin their journey in a tank of hydrogen gas [19] and travel though a series of smaller accelerators that increase their energy in stages until they reach the LHC, a process called the "injection chain" [14]. Three of the four accelerators used in the injection chain are shown in fig. 3.1. The injection chain has evolved over the LHC's lifetime, and the following description represents the injection chain used during Run 2. First, hydrogen molecules are released from the tank and passed through an electric field that strips them of their electrons [20]. The resulting protons are accelerated to an energy of 50 MeV via CERN's Linear accelerator 2 (Linac2), which is omitted from fig. 3.1. Next, the protons enter the Proton Synchrotron Booster, which accelerates them to 1.4 GeV and passes them to the Proton Synchrotron (PS), which accelerates them to 25 GeV [14]. The final step in the injection chain is the Super Proton Synchrotron (SPS), a synchrotron with a circumference

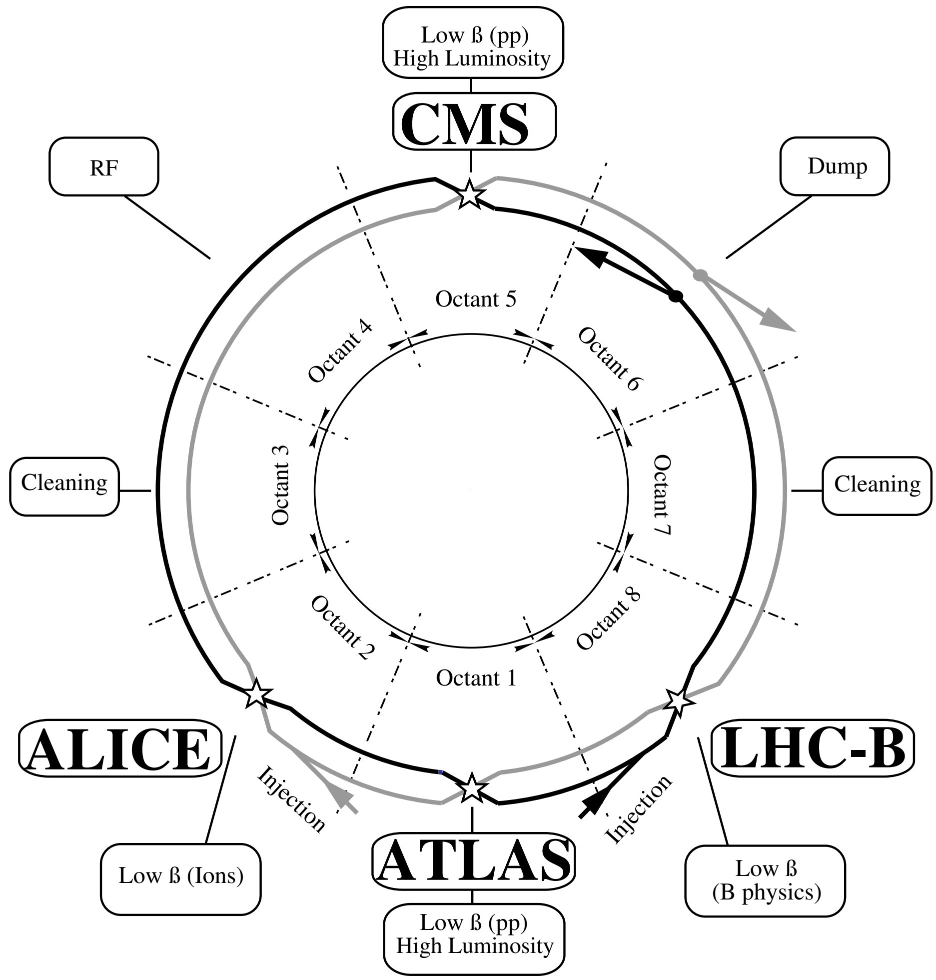

Figure 3.3: Schematic layout of the LHC [18].

of nearly 7 km (4.3 mi), that accelerates the protons to 450 GeV and injects them into the LHC [21]. The entire injection process takes about 16 minutes [14].

The LHC and all of the accelerators in its injection chain use RF cavities to accelerate the protons [14, 22]. RF cavities are metallic chambers that contain oscillating electric fields. Charged particles, like protons, receive an electric impulse from the field as they pass through it, which accelerates them. The frequency of the oscillations groups the protons into discrete "bunches." The bunches that arrive from the PS are spaced 25 ns apart, and they maintain this spacing as they pass through the LHC's 400 MHz RF cavities. By repeatedly passing through them, the protons are accelerated from 450 GeV to 6.5 TeV—a speed which is only a handful of miles per hour below the speed of light—over the course of about 20 minutes. Once the protons have reached this maximum energy, they circulate at a constant velocity until they collide inside one of the detectors or are dumped from the rings.

The proton beams inside the LHC are steered and maintained using a variety of magnets [14]. The primary magnets are superconducting dipole magnets that steer the beams through the rings; over 1,000 are required to cover the entire circumference of the LHC. They have a "twin bore" design, meaning they surround both rings in a figure eight shape, which enables a single magnet to provide a field for both beams. This design choice was necessitated by the tunnel's small size—it was built to house the Large Electron–Positron (LEP) collider, which required only one ring that was shared between the electron and positron beams [17]. A host of other magnets, some with as many as twelve poles, are used to focus the beams and correct their trajectory. Near the IPs, triplet assemblies consisting of three quadrupole magnets are used to constrict the beams to prepare them for collision. A narrower beam has a higher density of protons, increasing the instantaneous luminosity and potential for proton–proton collisions. The highest instantaneous luminosity inside the ATLAS detector during Run 2 was about $20 \times 10^{33}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ [23], or $0.02\,\mathrm{pb}^{-1}\,\mathrm{s}^{-1}$.

3.2 The ATLAS detector

ATLAS is a general purpose particle detector, meaning it is designed to detect as many particles as possible. Its design does not emphasize any particular particle or class of particles. It is composed of a host of complementary detector systems that, together, provide a very complete picture of every collision it records. Its geometry is cylindrical and symmetric about the LHC's beam line, which passes through its axis. At about 44 m (140 ft) long and 25 m (82 ft) in diameter, it is the size of a large, several-story building. Figure 3.4 is a diagram of the detector showing its overall dimensions and humans for scale, with labels indicating its main components.

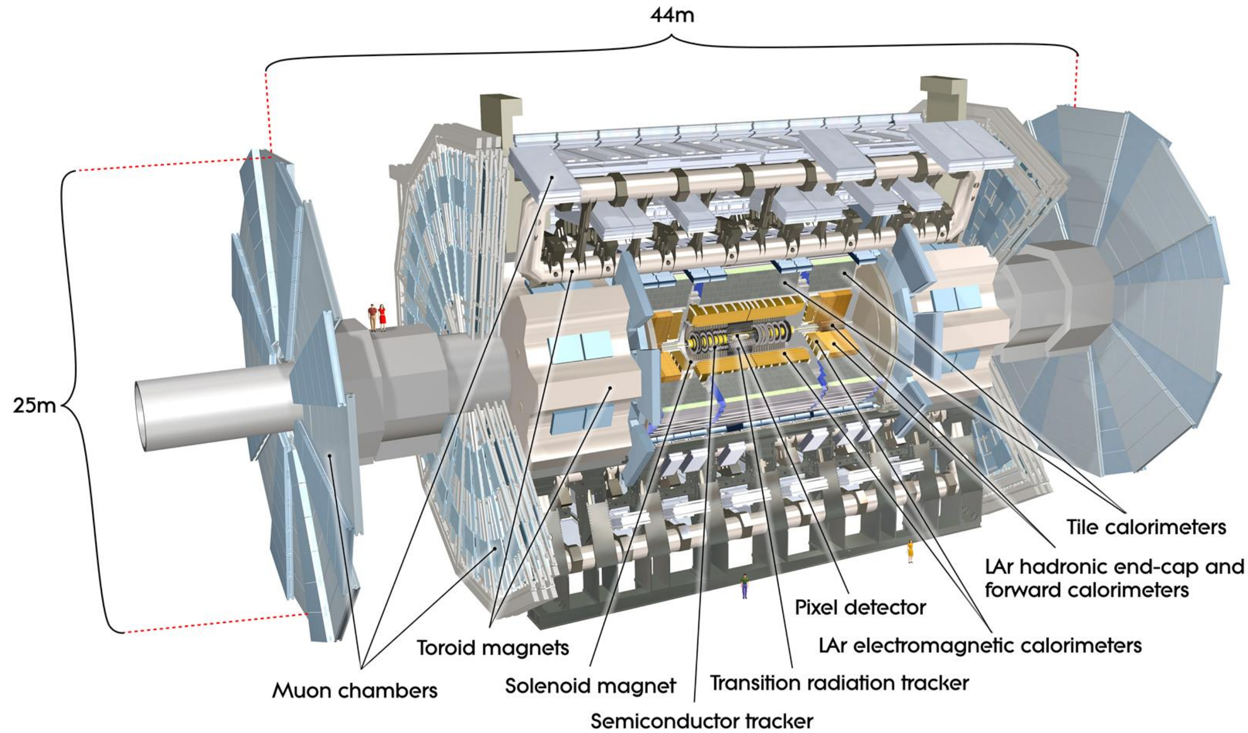

Figure 3.4: Cutaway view of the ATLAS detector [24].

This section describes the ATLAS detector as it was configured for Run 2 of the LHC, focusing on the elements most relevant to the work presented in this dissertation. It begins with a description of the detector's coordinate system in section 3.2.1. This is followed by a discussion of each of the detector's subsystems, beginning with those closest to the collision point and moving outward, mirroring the trajectory of a particle created in a collison. The tracking system, called the inner detector, is described in section 3.2.2. The calorimeter system is described in section 3.2.3. The muon spectrometer and the LUCID-2 detector for luminosity measurement are described in section 3.2.4. Finally, the trigger system, which uses information from multiple systems, is described in section 3.2.5.

3.2.1 Coordinate system

A coordinate system is defined to describe locations within the ATLAS detector. The origin coincides with the interaction point (IP) at the center of the detector where collisions occur.

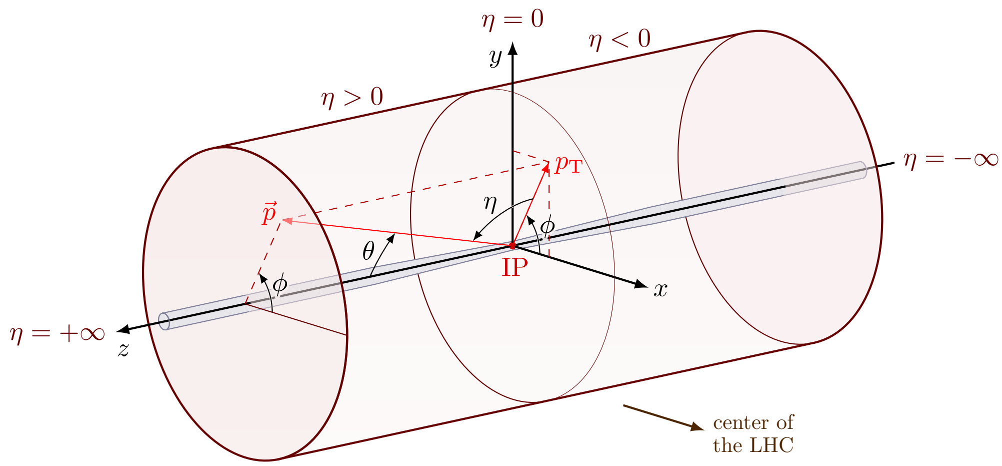

Figure 3.5: ATLAS coordinate system [26].

The beam line runs along the z axis. The xy plane, often called the "transverse" plane, is perpendicular to the beam line at the center of the detector. Many particles created at the IP will have trajectories that lie on or near the transverse plane. The azimuthal angle ϕ and polar angle θ are defined in the typical way [25]. Both spherical and Cartesian coordinates are shown in fig. 3.5, where $\vec{p}$ represents the momentum of a particle originating from the IP. The particle's transverse momentum, p_{T}, is the component of its momentum in the transverse plane. This quantity is often used as a proxy to a particle's total momentum.

Figure 3.5 also shows the pseudorapidity, η, which describes the distance between a particle's trajectory and the transverse plane. It is defined as

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

Accordingly, $\eta = 0$ coincides with $\theta = \pi/2$, while η approaches ∞ as θ approaches 0. The values of θ and η for several angles are shown in fig. 3.6. In the pp collision context, η is preferable to θ because, in the relativistic limit, differences in η are independent of boost in the z direction. Incoming quarks and gluons, collectively called "partons," will not in general have exactly the same momentum when they collide, meaning their center-of-mass frame will

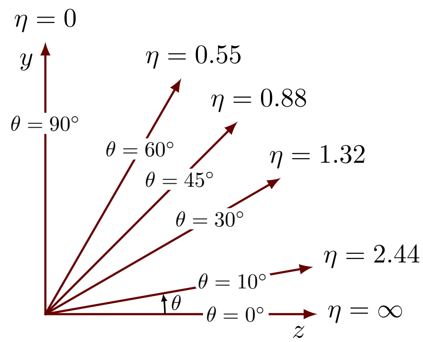

Figure 3.6: Pseudorapidity η for various values of the polar angle θ [28].

be boosted along the z axis with respect to the lab frame. The distribution of $\Delta\eta$ between objects from many such collisions will not be distorted by this boost [27].

Combining these definitions, the momentum of a particle (or any object) in the detector is typically written as

$$\vec{p} = (p_{\mathrm{T}}, \eta, \phi). \tag{3.2}$$

Another useful quantity is the angular separation of two objects, defined as

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}, \tag{3.3}$$

which describes a cone originating from the IP. Like $\Delta\eta$, ΔR is also invariant under boosts along the z direction in the massless approximation.

3.2.2 Inner detector

The ATLAS inner detector (ID) is used to record the tracks of charged particles. It is located at the center of the detector and surrounded by a superconducting solenoid magnet that creates a 2 T field within [25]. The magnetic field points along the z direction and curves the trajectories of charged particles moving through it. The radius of curvature r is given by

$$r = \frac{p_{\mathrm{T}}}{|Q|B}, \tag{3.4}$$

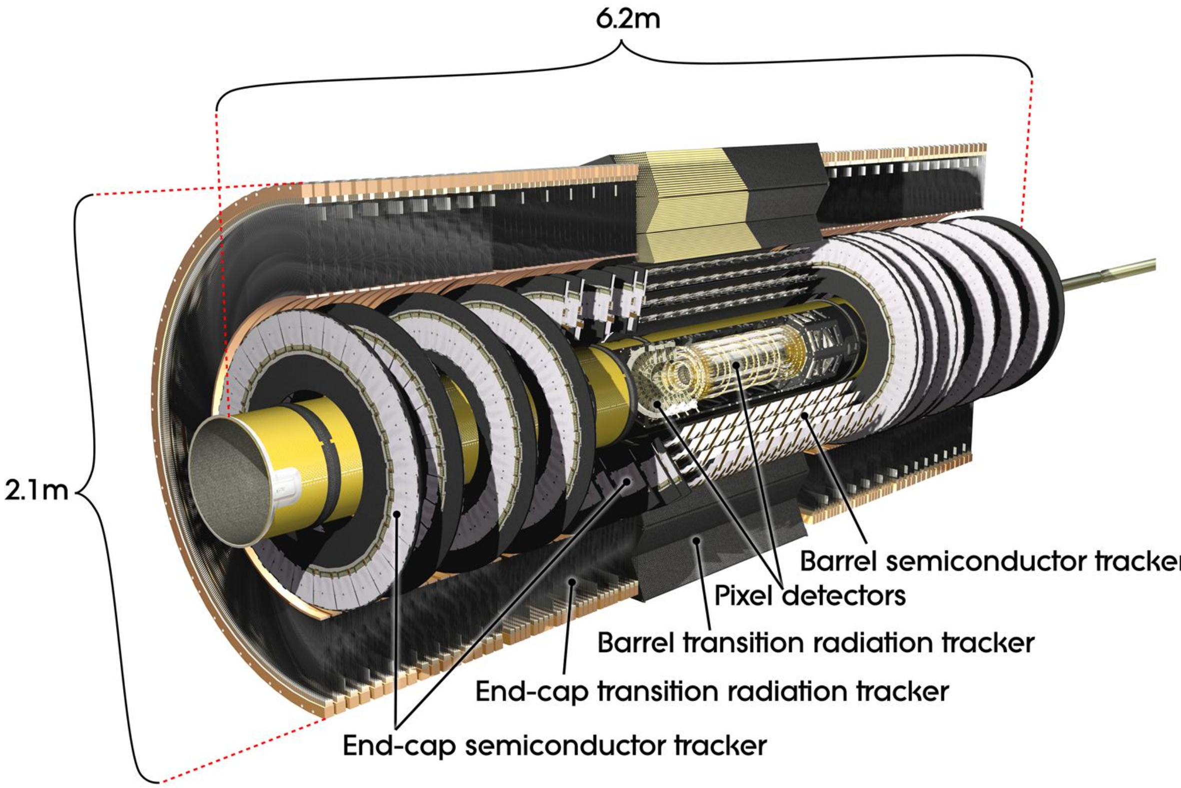

Figure 3.7: Cutaway view of the ATLAS inner detector [30].

where p_{T} is the particle's transverse momentum, Q is its charge, and B is the strength of the magnetic field [29]. The radius and direction of curvature can be used to measure a particle's momentum and determine whether it is positively or negatively charged.

The ID consists of three layers: the pixel detector, the semiconductor tracker (SCT), and the transition radiation tracker (TRT). All three are separated into barrel regions, which cover the center of the detector, and endcap regions, which are on either side of the barrel region. The total tracking coverage provided by the ID is $|\eta| < 2.5$. Figure 3.7 is a diagram of the entire ID, and fig. 3.8 shows the paths of two charged particles through a single radial section.

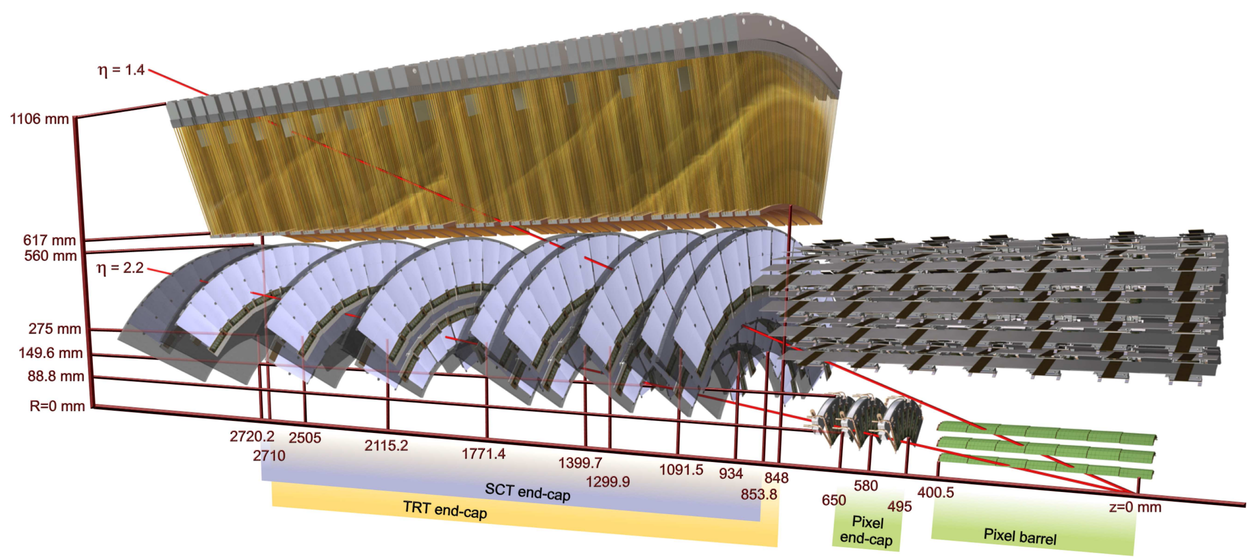

Figure 3.8: Paths of two charged particles, one with $\eta = 1.4$ and one with $\eta = 2.2$, through the inner detector [25].

Pixel detector

The pixel detector consists of layers of silicon pixels that provide precise tracking information in the region $|\eta| < 2.5$. Each pixel is a tiny semiconductor diode with an applied potential difference, called a "bias voltage." When a charged particle passes through the silicon, it leaves electron–hole pairs in its wake that drift toward the electrodes on either end of the pixels and are read out as signals [31]. The original ID design has three pixel layers in the barrel and three in each of the endcaps. A fourth layer, the insertable B layer, was added before the start of Run 2 and became the innermost layer of the pixel barrel [32]; it is not shown in fig. 3.7 or fig. 3.8. The pixels are 250 µm thick and between $50 \times 250\,\mu\text{m}^2$ and $50 \times 600\,\mu\text{m}^2$ in area, arranged with the long edge parallel to the beam line. The small size of the pixels enables tracking precision of about 10 µm in the radial direction and about 100 µm in the z direction. The entire pixel detector is composed of over 90 million pixels, and its usable radius is about 15 cm (6 in) [25].

Semiconductor tracker

The SCT covers the $|\eta| < 2.5$ region surrounding the pixel detector and is made up of paired layers of silicon strips. The operating principle of the silicon strips is the same as the pixels; the primary difference is their shape and arrangement. The strips are about 280 µm thick, 80 µm wide, and 12 cm long, and arranged parallel to the beam line. There are three SCT layers in the barrel and nine in each of the endcaps, composed of about 6 million total strips. Each layer is composed of "stereo pairs" of strips, which are offset at a small angle so the location of a hit can be triangulated in three dimensions. The SCT's precision is about 20 µm in the radial direction and about 600 µm in the z direction. Its usable radius extends to about 56 cm (22 in) [25].

Transition radiation tracker

The TRT surrounds the SCT and covers $|\eta| < 2.0$. It is made up of 4 mm diameter drift tubes, called "straw" tubes, that are sandwiched between layers of other materials. Drift tubes are proportional counters. An anode wire runs down the center of the tube, and the tube's surface acts as a cathode. When a charged particle passes through the gas in the tube, it ionizes, and the free electrons are drawn to the wire. When the signal is read out, the known drift time of electrons can be used to pinpoint how close the particle passed to the wire [31]. The layered structure of the TRT also creates its namesake transition radiation: photons that are produced when a particle moves across the boundary between different media. The quantity and energy of the photons is related to the mass of the particle and can be used to identify its type [33]. The TRT has about 300,000 tubes in total, and a typical track passes through 36 of them. The TRT only provides measurements in the radial direction, and its precision is about 150 µm. Its usable radius extends to about 1.1 m (3.5 ft) [25].

3.2.3 *Calorimeters*

The ATLAS calorimeter system surrounds the solenoid magnet and ID. It is used to absorb most particles created in collisions, preventing them from escaping the detector, and measuring the energy they deposit. The calorimeter system is complementary to the tracker—while the relative resolution of a momentum measurement from a track worsens with increasing momentum, the relative resolution of an energy measurement from a calorimeter improves with increasing energy [33]. Calorimeters are also sensitive to neutral particles that do not leave tracks in the ID.

The inner layers of the ATLAS calorimeter system are EM calorimeters that target electrons and photons, while the outer layers are hadronic calorimeters that target quarks, gluons, and composite hadrons like pions. Most of the calorimeters use liquid argon (LAr) as the detection medium, except the hadronic tile calorimeter, which uses plastic scintillator. The calorimeters may be divided into three major groups: the EM calorimeters, the tile calorimeter, and the LAr hadronic and forward calorimeters. The barrel and endcap calorimeters provide coverage over $|\eta| < 3.2$, and the forward calorimeters extend coverage to $|\eta| < 4.9$. Figure 3.9 is a diagram of the ATLAS calorimeter system.

Electromagnetic calorimeters

The EM calorimeter system has a barrel section that covers $|\eta| < 1.475$ and endcap sections that cover $1.375 < |\eta| < 3.2$. Both sections consist of alternating layers of lead and LAr arranged in an accordion pattern, shown in fig. 3.10a, to maximize the number of layers a particle passes through [25]. The “passive” lead layers serve to slow down electrons and photons via their interaction with the dense material. This interaction produces a “shower” of lower-energy electrons that stop in the “active” LAr layers when they ionize the LAr. The electic current from the ionization is read out, and these measurements are combined to determine the energy of the original particle. The ATLAS EM calorimeter is about

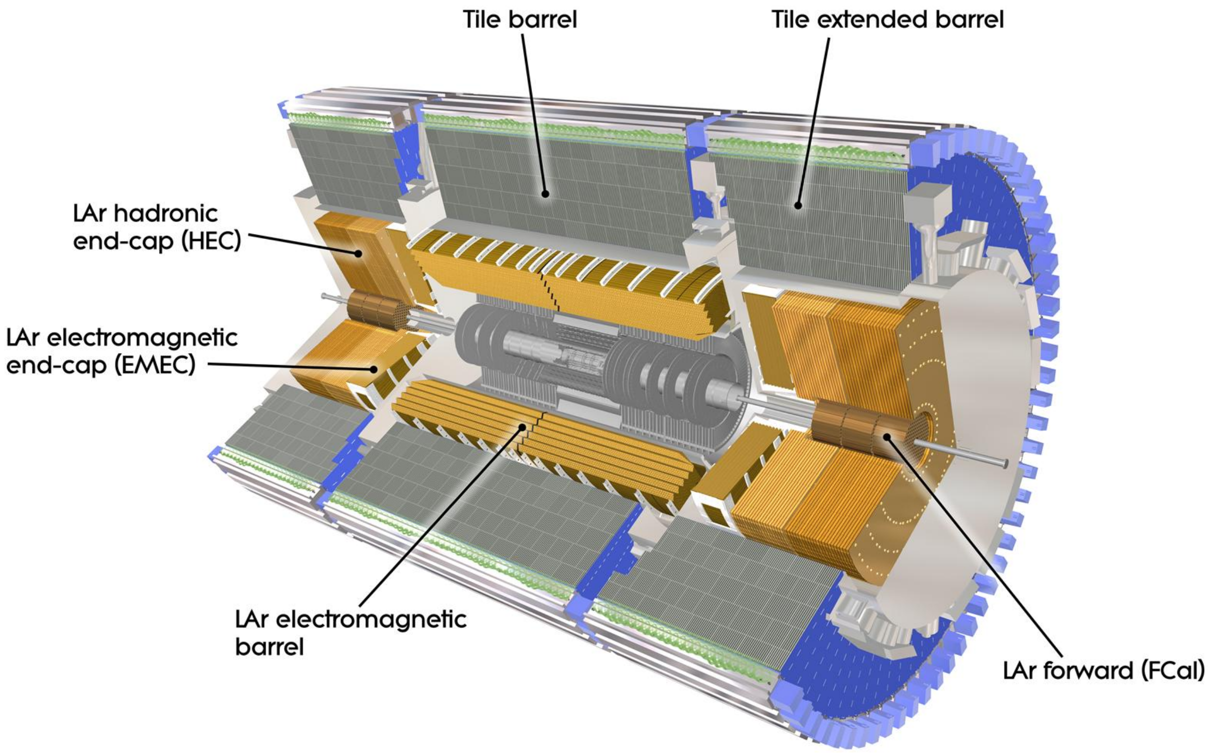

Figure 3.9: Cutaway view of the ATLAS calorimeter system [34].

1.2 m (4 ft) thick in the radial direction. A useful way to report this is in terms of "radiation length" X_0, the mean distance an electron must travel to lose $1 - e^{-1}$ (about 63%) of its energy [33]. The ATLAS EM calorimeter is more than $22X_0$ thick in the barrel and more than $24X_0$ thick in the endcaps, providing excellent energy absorption [25].

Tile calorimeter

The tile calorimeter (TileCal) has a barrel region that covers $|\eta| < 1.0$ and extended barrel regions that cover $0.8 < |\eta| < 1.7$. It is made of steel, the passive material, interspersed with tiles of plastic scintillator, the active material, laid perpendicular to the beam line as shown in fig. 3.10b to avoid gaps in azimuthal coverage. The TileCal's operation is analogous to the EM calorimeter, but for hadrons: they shower hadronically in the steel, and the shower particles ionize the plastic, which emits scintillation light [31, 33]. The scintillation photons

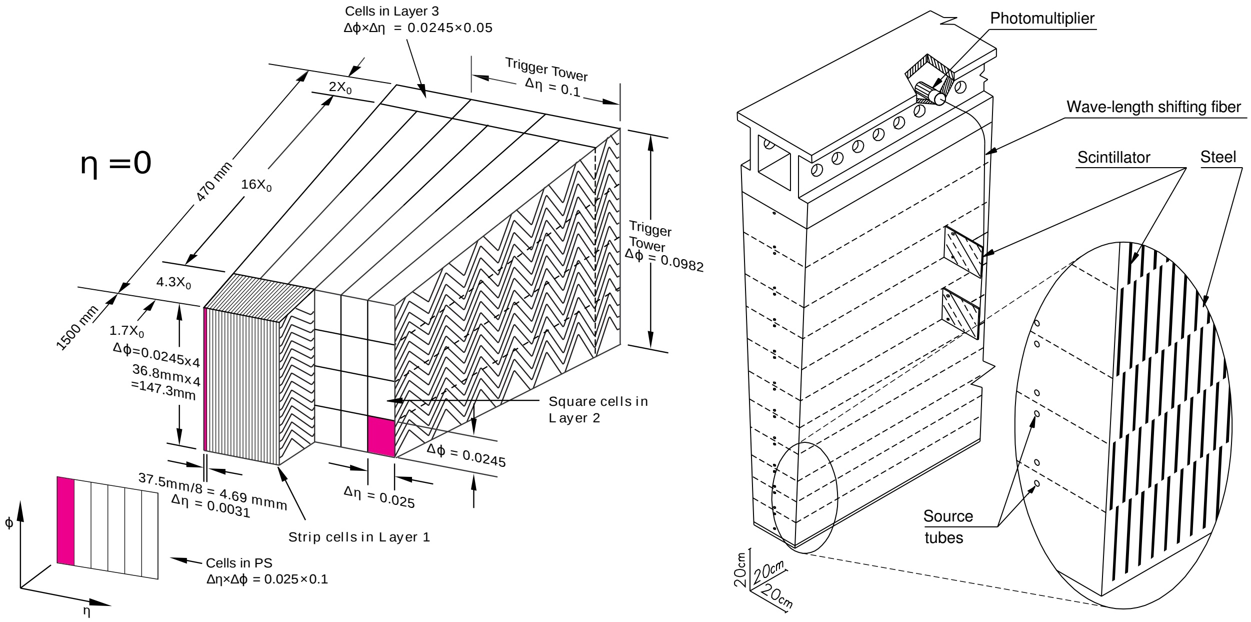

(a) LAr barrel module (b) TileCal module

Figure 3.10: Schematic diagrams of ATLAS calorimeter modules, showing the arrangement of active and passive layers [35, 36].

are collected using photomultiplier tubes, which provide energy measurements. The TileCal is about 2 m (6.6 ft) thick. The hadronic analogue to the radiation length is the "interaction length" λ, which is the mean distance a hadron travels before it interacts [33]. The TileCal is about 7.4λ thick, and contains over 400,000 scintillator tiles [25].

LAr hadronic and forward calorimeters

Each hadronic endcap calorimeter (HEC) is made of copper plates interspersed with layers of liquid argon. The HECs are composed of two "wheels," each with a radius of about 2 m (6.6 ft) and a thickness of about 90 cm (3 ft). Together, both wheels cover the range $1.5 < |\eta| < 3.2$ [25].

The forward calorimeter (FCal) system provides EM and hadronic calorimetry in the "forward" region, characterized by $3.1 < |\eta| < 4.9$. This is located between the beam pipe and the endcap calorimeters. Each FCal is composed of three modules: one made of copper and LAr to target EM calorimetry, and two made of tungsten and LAr to target hadronic calorimetry. Each module is made up of concentric metal rods and tubes, with LAr filling the gap between them, arranged parallel to the beam line. The modules are all about 45 cm (1.5 ft) in radius and depth. The EM modules are located closest to the IP and are greater than $27X_0$ thick. All three modules together are about 9.9λ thick [25].

3.2.4 Other detectors

Other ATLAS detector systems that are less relevant to the work presented in this dissertation, but still critical to its overall function, are described here. These systems measure muons and luminosity.

Muon spectrometer

Muons are much heavier than electrons and also have a relatively long lifetime, so they require dedicated detector systems. The ATLAS muon spectrometer (MS) is the outermost layer of the detector, which is responsible for its overall dimensions. It used to identify muons and measure their momenta in the range $|\eta| < 2.7$, using principles similar to the tracking system in the ID. It uses three toroid magnets, one in the barrel and one in each endcap, to create a magnetic field with strength ranging from about 0.5 to 1 T. It has four types of chamber, each using a different technology: monitored drift tube chambers, cathode strip chambers, resistive plate chambers, and thin gap chambers. They are labeled in fig. 3.11, which is a diagram of the MS. All of these technologies are gaseous ionization detectors that provide complementary coverage and tracking measurements. Monitored drift tubes provide the most precise measurements, with resolution as good as 35 µm [25].

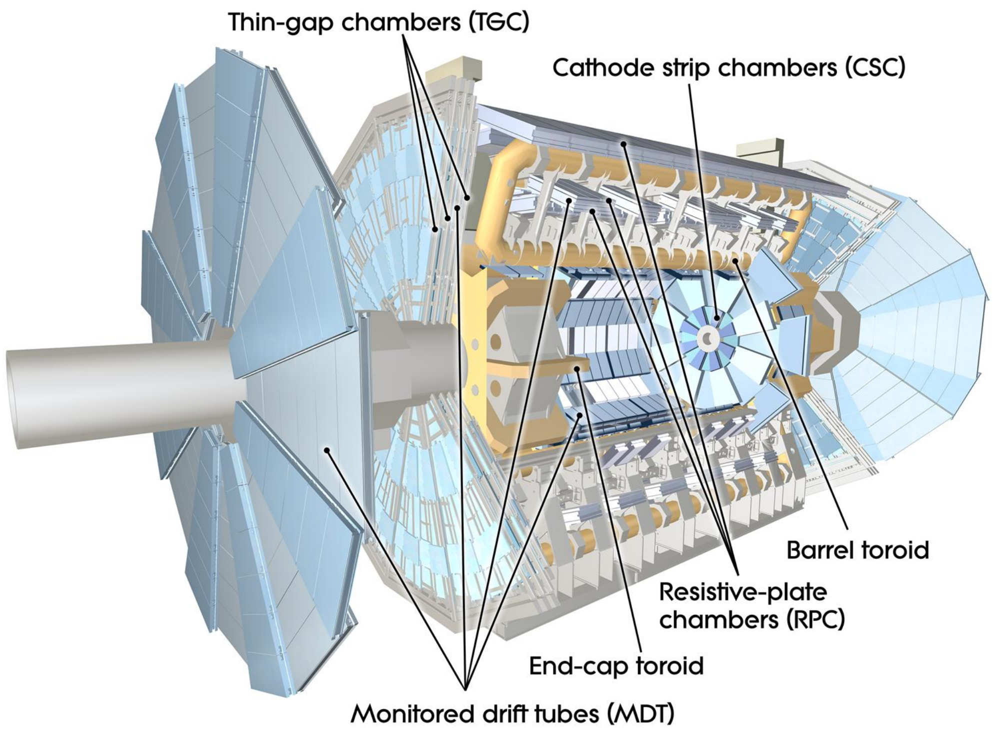

Figure 3.11: Cutaway view of the ATLAS muon system [37].

LUCID-2 detector

The luminosity of the ATLAS Run 2 dataset was measured using the LUCID-2 detector, which consists primarily of several groups of photomultiplier tubes (PMTs) positioned around both sides of the beam pipe, about 17 m (56 ft) from the IP [25]. The photomultipliers measure the Cherenkov light that is emitted when high-energy charged particles produced in *pp* collisions pass through the quartz windows of the PMTs [31]. The luminosity measurements from LUCID-2 are combined with measurements from the ID and calorimeters to improve their precision [38].

3.2.5 Trigger system

The LHC's 25 ns bunch spacing is equivalent to a data rate of 40 MHz, but the ATLAS electronics system can only record event data at a rate of about 1.2 kHz [39]. The ATLAS trigger system is designed to quickly identify and store only events with interesting physics signatures in order to make the most of the available data rate. The trigger system has two stages: the Level 1 (L1) trigger and the high-level trigger (HLT).

The L1 trigger operates in hardware and uses low-granularity information from the calorimeters and limited information from the MS to quickly seek out, typically, objects with high p_{T}. It has 2.5 μs to decide whether an event should be retained and must reduce the data rate to the maximum readout rate of 100 kHz [39]. If it identifies an object that passes criteria listed in the "trigger menu," it retains the event and passes information about the object's approximate location, called a region of interest (RoI), to the HLT [25].

The HLT operates in software and analyzes RoI data from all available detector systems. This includes high-granularity calorimeter information, precision measurements from the MS, and tracking information from the ID [40]. It performs "online" reconstruction of objects using fast versions of the algorithms used for the "offline" reconstruction that is carried out after an event has been permanently stored, a process that is described in section 4.3. The HLT has a few hundred milliseconds to decide whether to retain an event, which it does based on more finely tuned criteria than the L1 trigger. The average output rate from the HLT is the aforementioned 1.2 kHz [39].

CHAPTER 4

COMMON ANALYSIS INGREDIENTS

All ATLAS analyses use a common set of strategies and tools. This section presents information that is relevant to both of the searches presented in this dissertation. Section 4.1 gives a description of the Run 2 dataset used for both searches. Section 4.2 provides a brief overview of the general procedure for characterizing signal and background processes, with a focus on Monte Carlo simulation of events. Section 4.3 describes how particles are reconstructed with the ATLAS detector as physics objects used in analysis. Finally, section 4.4 describes tagging procedures used to select jets.

4.1 The ATLAS Run 2 dataset

The ATLAS Run 2 dataset is used for both of the searches presented in this dissertation. The center-of-mass energy, $\sqrt{s}$, of the Run 2 collisions was 13 TeV. Data collection began in 2015 and ended in 2018, and during this time period the LHC delivered $156\,\text{fb}^{-1}$ of collision data to the ATLAS detector. ATLAS recorded $147\,\text{fb}^{-1}$ of them, and 140 fb passed quality criteria that make them suitable for analysis [41]. The $140\,\text{fb}^{-1}$ value comes from a recalibration of the luminosity measurement that was published in 2022, and has an uncertainty of 0.83% [23]. Prior to 2022, the value was $139\,\text{fb}^{-1}$ and had an uncertainty of 1.7% [42]. Figure 4.1 shows the cumulative luminosity versus time during Run 2 for all three categories.

The integrated luminosity for each of the Run 2 data-taking years is listed in table 4.1 for both calibrations. The integrated luminosity for 2015 is much lower than for each of the other three years; for this reason, some Run 2 analyses omit the 2015 data. Some analyses may not use all of the data for other reasons, such as the choice of trigger. ATLAS data collection is divided into discrete periods called “runs” (not to be confused with the runs of the LHC) during which data are recorded continuously, and the runs during which analysis-

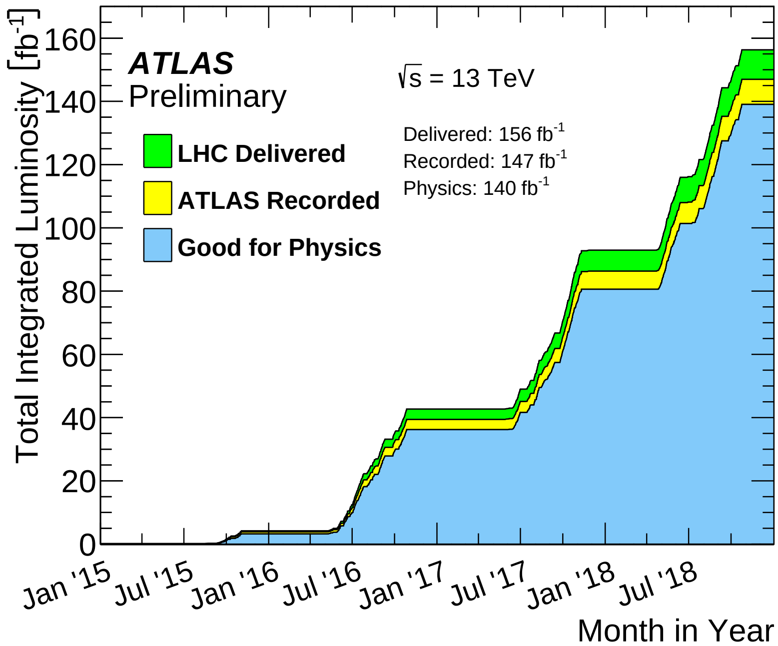

Figure 4.1: Cumulative luminosity versus time delivered to ATLAS, recorded by ATLAS, and classified as good for physics during stable beams for Run 2 *pp* collisions [41].

Table 4.1: Integrated luminosities for the Run 2 *pp* dataset, with results from both the 2019 and 2022 luminosity calibrations. The 2019 calibration reports a combined value for 2015 and 2016 [23, 42].

	Integrated luminosity (fb^{-1})	
Year	2019 calibration	2022 calibration
2015		3.2
2016	36.2	33.4
2017	44.3	44.6
2017	58.5	58.8
Total	139.0	140.1

ready data are taken are stored in a "good run list" (GRL) [43]. Most analyses use the so-called standard GRL, but other GRLs exist that are used with specific triggers depending on their needs. This is relevant for the search presented in chapter 6.

The LHC's peak instantaneous luminosity reached $20 \times 10^{33}\,\text{cm}^{-2}\,\text{s}^{-1}$ during Run 2, as mentioned in section 3.1. Accordingly, for each bunch crossing at the center of the AT-

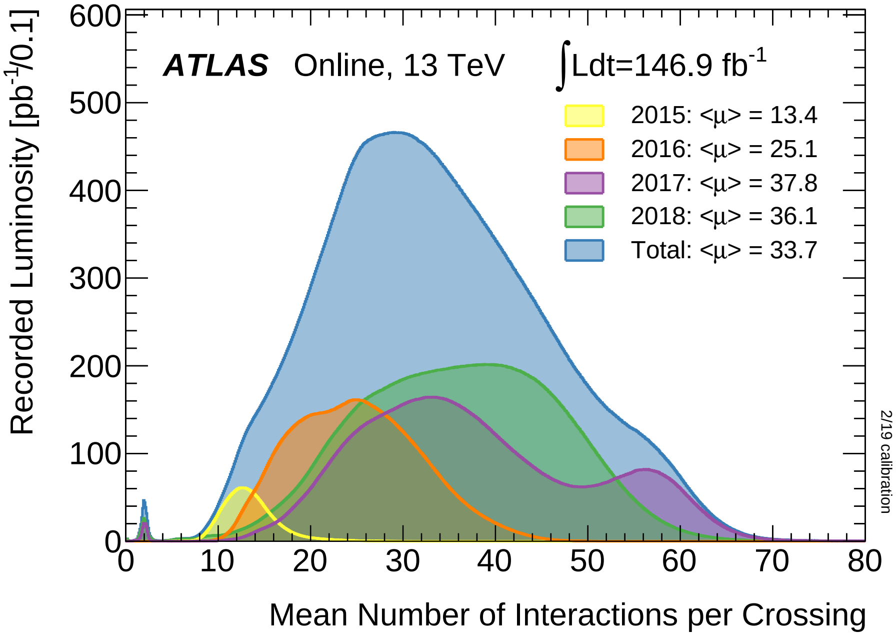

Figure 4.2: Luminosity-weighted distribution of the mean number of interactions per crossing for Run 2 pp collisions [44].

LAS detector, there are many simultaneous pp interactions. However, typically only one interaction per crossing results in interesting physics that we would like to study. The extraneous interactions, called "pileup," are a source of noise that must be mitigated during analysis. This mitigation becomes more challenging as the number of interactions per crossing increases. During Run 2, each bunch crossing resulted in an average of 34 interactions, although this number changed over the course of the run. Figure 4.2 shows the average number of interactions per bunch crossing, $\langle\mu\rangle$, for each year of data taking.

4.2 Signal, background, and simulation

The first step of any analysis is to identify the target physics process or processes, called the "signal." The second step is to identify the "backgrounds"—other processes that resemble the signal's detector signature. As the analysis develops, an event selection procedure is defined to choose events that likely originate from the signal process and, to the extent possible, reject events that likely originate from background processes. An effective event selection maximizes the ratio of signal to background events while keeping acceptance high enough for sufficient statistical precision.

It is difficult, if not impossible, to search for any process without a concrete idea of what one is looking for. For this reason, ATLAS analyses make heavy use of simulated physics events. These events are produced via a random sampling technique called Monte Carlo (MC) simulation. Due to the ubiquity of this technique, simulated events or samples are often referred to as simply "MC." The simulation process begins with an event generator, which models the momenta of incoming and outgoing particles in the hard scatter process of a collision event. Different MC generators are available that excel at generating events of different processes. The second step of the simulation process is "showering," which is the simulation of the hadronic activity that creates jets. The final steps are digitization and detector simulation, which recreate the conditions under which real particle collisions are recorded inside the ATLAS detector. The goal of the entire process is to simulate real collision data as closely as possible, but with the ability to isolate individual physics processes and access complete "truth-level" information about each interaction.

MC simulated events of the signal process are used to design and optimize the event selection for an analysis. MC simulated events of background processes are often used to inform the event selection and to estimate contributions from background processes. However, some processes are difficult to simulate accurately, and in these cases "data-driven" background estimation techniques are preferred. These techniques estimate background processes from real data events using an event selection that targets the background instead of the signal process. The searches presented in this dissertation use both methods to estimate background contributions.

In addition to the isolation of processes, another benefit of MC simulation is the ability to generate an arbitrary number of events. Typically, the MC events generated for a given process far outnumber the number of events expected from that process in real data. This way, simulated events are not limited by statistical precision. A scaling factor w is applied to MC events to scale them to the expected number of events:

$$w = \frac{L\sigma}{N}, \tag{4.1}$$

where L is the integrated luminosity of the target dataset, σ is the cross section of the simulated process, and N is the number of events generated. The impact of the uncertainty in the integrated luminosity is assessed by varying w within it.

Although MC simulated events bear very close resemblance to real data, there exist residual differences between MC and data that are assessed as systematic uncertainties. These systematic uncertainties are typically assessed using a "reweighting" procedure, in which MC events are multiplied by different constant factors when producing a histogram in order to match the shape of the same distribution in real data. For example, the $\langle\mu\rangle$ distribution in fig. 4.2 is not perfectly replicated when simulating the effects of pileup in MC events, so weights are applied to the MC events in order to make the simulated $\langle\mu\rangle$ distribution match the one from data events. Systematic uncertainties are assessed by varying these weights within their uncertainties. Similar procedures are used to assess systematic uncertainties from nearly all imaginable sources, including various aspects of the particle reconstruction methods detailed in section 4.3. Uncertainties in the theoretical calculations used to model the physics processes in MC event generation are also assessed as systematic uncertainties.

4.3 Particle reconstruction and physics objects

The ATLAS detector is designed to detect as many particles as possible. Particles with very short lifetimes—typically, W, Z, and Higgs bosons, τ leptons, and t quarks—decay immediately and must be reconstructed from their decay products. Particles with longer lifetimes—electrons, muons, photons, gluons, and the other five flavors of quark—can be detected directly. Neutrinos do not interact with the detector and cannot be detected at all, but their presence can be inferred from missing transverse momentum, discussed later in this section. Some composite particles can also be directly detected, but the ATLAS searches in this dissertation involve only elementary particles, so composite particles will not be discussed here.

Figure 4.3 is a schematic of a slice of the ATLAS detector with visualizations of how various particles are detected, which provides a visual reference for this section's discussion of particle reconstruction and physics objects. A "physics object" is a collection of measurements from various detector systems that, together, represent a reconstructed particle [45]. Physics objects are the main ingredients of any ATLAS analysis.

This section describes the general reconstruction methods for the physics objects used in the ATLAS searches presented in this dissertation. Primary vertex reconstruction is described in section 4.3.1. Electrons, photons, and muons are each reconstructed as distinct physics objects that share their names; they are discussed in section 4.3.2, section 4.3.3, and section 4.3.4, respectively. Quarks and gluons are reconstructed as objects called "jets," discussed in section 4.3.5. Missing transverse energy is described in section 4.3.6.

4.3.1 Primary vertex

A "vertex" is the reconstructed position of a single pp collision within the ATLAS detector. Each bunch crossing creates many pp interactions and therefore many vertices. The "primary vertex" (PV) is the vertex from which the physics of interest originates. For a vertex to be reconstructed, it must have at least two tracks in the ID, each with $p_{\mathrm{T}} > 0.5\,\mathrm{GeV}$, that originate from the interaction region in the xy plane. The PV is identified as the vertex with the largest sum of squared track momenta [47].

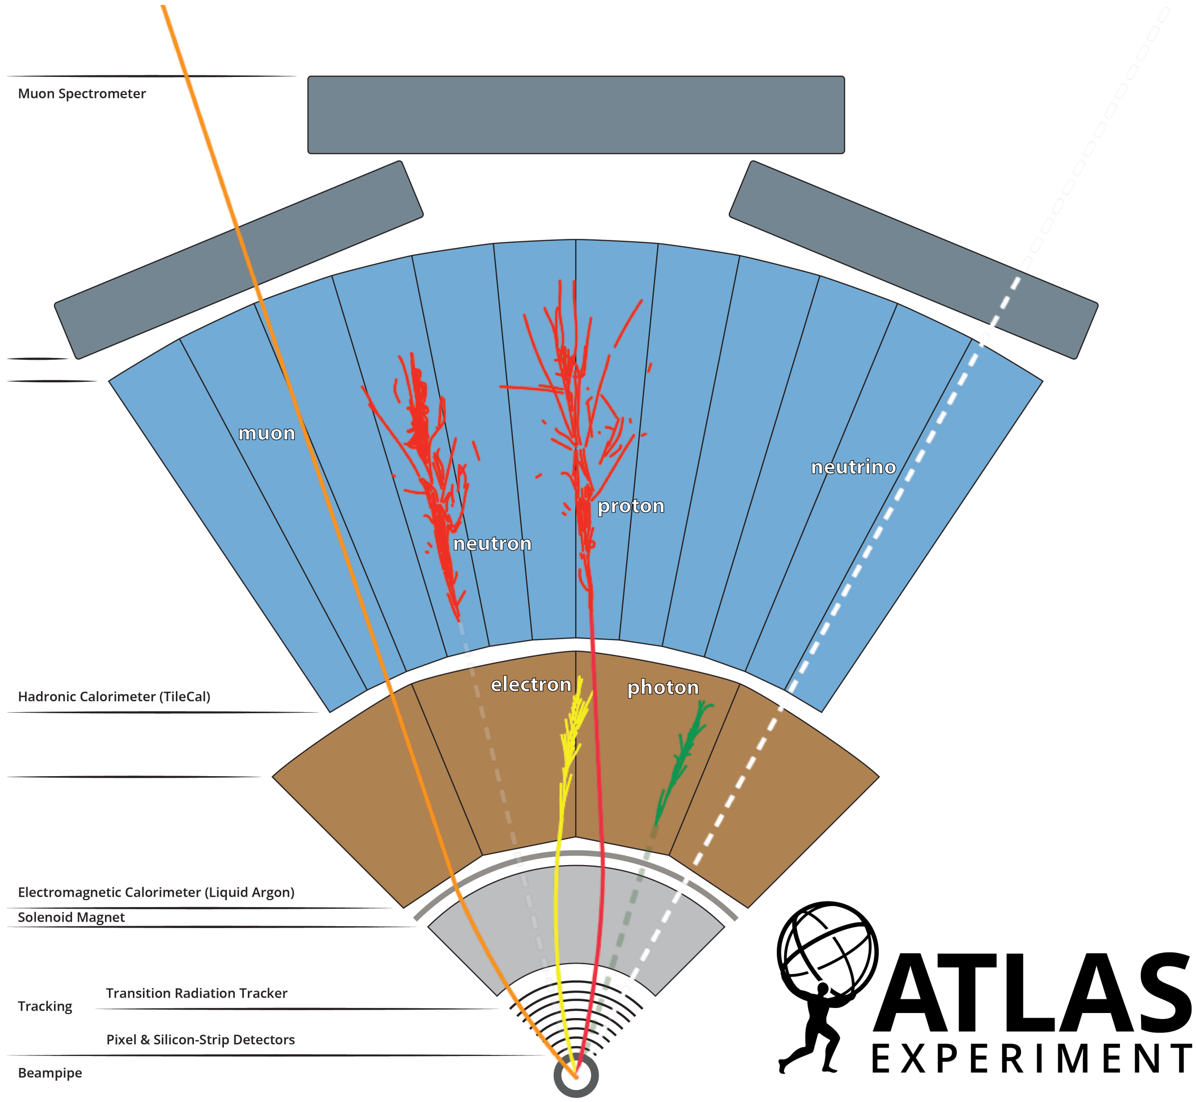

Figure 4.3: A slice of the ATLAS detector with visualizations of how different particles are detected [46].

4.3.2 Electrons

An electron consists of a "cluster" built from an energy deposit in the EM calorimeter that is matched to a track from the ID. The track maps out the curved path the electron takes through the magnetic field inside the ID, and the cluster originates from the shower it creates as it passes through the calorimeter. Electrons are typically reconstructed only within the "precision" area of the detector covered by the ID, which is $|\eta| < 2.47$. The range $1.37 < |\eta| < 1.52$ is called the "transition region" because it marks the transition from the barrel to endcap calorimeters. In this region, the clusters are supplemented with energy measurements from the presampler and the scintillator between the calorimeter cryostats [48].

The energy of an electron is estimated from the energy of the cluster, with a correction applied for energy lost outside the EM calorimeter. The energy resolution is lower in the transition region due to the greater amount of inert material, which causes a larger fraction of energy to be lost. Corrections are applied to the electron energy scale and resolution in MC simulation to account for observed differences between real and simulated $Z \to e^+e^-$ events [49]. The curvature of an electron's track is used to determine whether it has positive or negative charge [48].

The quality of an electron object, meaning the likelihood that its cluster and track originate from an electron as opposed to a jet or a photon, is assessed using identification criteria. A discriminant is constructed from the likelihoods of the electron object to originate from signal or background processes, based on probability distributions over many observables characterizing the properties of the cluster and track. These properties include, for example, the shape of the cluster and the number of tracking hits in the ID. "Loose," "medium," and "tight" working points are defined to broadly characterize the likelihood of an electron object to have originated from a true signal electron, with a tighter identification criterion corresponding to a higher likelihood. The efficiencies for a signal electron with $E_{\mathrm{T}} = 40\,\mathrm{GeV}$ to pass the loose, medium, and tight identification criteria are approximately 93%, 88%, and 80%, respectively [50].

4.3.3 Photons

Like an electron, a photon also deposits a cluster of energy in the EM calorimeter, but it does not leave a track because it is not electrically charged. However, about 30% of photons convert to an electron–positron pair within the ID [51], creating a pair of tracks called a "conversion vertex." Therefore, a photon is either a cluster matched to a conversion vertex or a cluster that is not matched to a cluster or to a vertex. The energy of a photon is estimated in the same way as the energy of an electron, and the same energy scale and resolution calibrations are applied, with additional photon-specific corrections derived from $Z \to e^+e^-\gamma$ and $Z \to \mu^+\mu^-\gamma$ events [48].

Photon identification and isolation are used to differentiate photons from electrons and jets. The identification criteria are cuts on variables that characterize the shape of the EM shower in the calorimeter. "Loose," "medium," and "tight" working points are defined. For a photon object with $E_\mathrm{T} = 40\,\mathrm{GeV}$, the tight working point has an efficiency above 90% and a misidentification rate below 40% [48]. Photon identification excludes the transition region of the EM calorimeter [51] because the information about shower shape in this region is limited. Photon isolation has loose and tight working points. The tight working point limits energy deposits in the calorimeter within a cone of size $\Delta R = 0.4$ around the photon's direction, and in the tracker within a cone of size $\Delta R = 0.2$. Calibrations derived from $Z \to e^+e^-\gamma$ and $Z \to \mu^+\mu^-\gamma$ events are used to correct observed differences between data and MC simulation in the shape of the calorimeter isolation distribution [48].

4.3.4 Muons

Muons are reconstructed from tracks in the MS and ID, supplemented with calorimeter measurements. A muon can be reconstructed using measurements from these three systems in a number of different ways, which are complementary in p_T and η. Tracks map out the trajectory of the muon through the ID and MS and are matched to recover the trajectory through the entire detector. Calorimeter deposits consistent with a minimum-ionizing particle are used to provide coverage primarily in the $|\eta| < 0.1$ region, where there is a gap in MS coverage. Muons are typically reconstructed within $|\eta| < 2.5$, which is the coverage of the ID, but this range can be extended to the MS coverage region of $|\eta| < 2.7$. A muon's momentum is measured by the curvature of its track in the magnetic fields within the MS and ID. Calibrations from $Z \to \mu^+\mu^-$ events are used to correct differences in momentum scale and resolution between data and simulation [52].

Prompt muons are distinguished from muons originating from hadrons with identification and isolation criteria. Muon identification criteria include the number of hits in the ID and MS, the quality of the track fit, and the compatibility of the ID and MS measurements. "Loose," "medium," and "tight" working points are defined by cuts on these quantities, with the loose working point having an efficiency of about 98% for muons and below 1% for misidentified hadrons with $p_{\mathrm{T}} > 20\,\mathrm{GeV}$ [53]. A set of "very loose" criteria is also defined as a superset of the loose criteria with even higher efficiency.

4.3.5 Jets

Quarks and gluons, collectively known as "partons," are reconstructed as objects called "jets." Because of the color confinement phenomenon described in section 2.1.2, partons cannot be observed in isolation. As they move away from one another, the potential energy between them increases rapidly, and it soon becomes energetically favorable for new quark–antiquark pairs to be created from the vacuum in the space between them [5]. The newly-created quarks immediately interact with one another and form hadrons. The result is a roughly cone-shaped "jet" of hadrons forming around what would have been the trajectory of the individual parton [45]. A "jet clustering" algorithm is used to group hadrons into jets. The most frequently used algorithm in ATLAS analyses, including the searches presented in this dissertation, is the anti-k_t algorithm, which takes as input a jet radius parameter R defined in (η, ϕ) space [54].

The hadrons within the jet create showers as they pass through the hadronic calorimeter, and charged hadrons also create tracks in the ID. The showers consist primarily of hadrons, but can also include leptons and photons. ATLAS uses a "particle flow" algorithm to combine measurements from the ID and calorimeter to form signals that, ideally, originate from individual particles created in the hard scatter process. The energy from the charged particles in the tracker is subtracted from the calorimeter measurements so they are not counted twice. The matched tracks and corrected calorimeter clusters are called "particle flow objects," which are used to reconstruct jets with $R = 0.4$ [55]. Jets with $R = 1.0$, called "large-R" jets, are also used in some cases. They are reconstructed from "unified flow objects," which are created from a hybrid of the particle flow algorithm and a similar approach that solves the double counting of charged particles via a reweighting procedure instead of subtraction [56, 57].

The energy of a jet is determined by combining the corrected energy of the clustered calorimeter deposits with the momenta of the matched tracks. The jet energy scale is calibrated to recover the energy of the single particle from which the jet originated. The calibrations include subtraction of energy originating from pileup interactions and addition of energy lost to passive detector material, which are derived from MC simulation. Additional corrections obtained from events in which a jet recoils against another jet or a different object, such as a photon, are applied as well. The jet energy resolution is also calibrated, using a resolution measured in dijet events and noise corrections derived from measuring residual energy deposits inside random cones equal in size to the jets [58, 59].

Unlike electrons, photons, and muons, jets have substantial mass. Large-R jets in particular can be very massive. The mass scale and resolution of large-R jets are calibrated using two methods. The first method uses tracking information to obtain an independent measurement of the jet mass scale and resolution. The second uses the mass scale and resolution of well-understood particles: W and Z bosons and t quarks. These particles are reconstructed from jets, and the positions and shapes of the peaks in the large-R jet mass spectrum are used to assess the mass scale and resolution [59, 60].

The concepts of identification and isolation do not extend to jets, since they are composite objects that can originate from different kinds of particles. However, various "tagging" methods exists to quantify jet quality. Because pileup interactions create many low-momentum jets in every collision event, care must be taken to ensure that the jets selected for analysis originate from the PV. This process is known as "vertex tagging" and is discussed in section 4.4.1. Although jets originating from different kinds of particles are largely indistinguishable from one another, jets originating from certain particles, like b and c quarks, have special properties that make it possible to differentiate them from jets originating from gluons and other flavors of quark. The process of identifying these jets is called "flavor tagging." The searches presented in this dissertation make use of b and bb tagging algorithms, which are presented in section 4.4.2.

4.3.6 Missing transverse momentum

Missing transverse momentum, $E_{\mathrm{T}}^{\mathrm{miss}}$, is reconstructed from the principle of momentum conservation. In the transverse plane, the momenta of all particles originating from a pp collision at the PV must sum to zero. This requirement can be used to infer the presence of particles like neutrinos that are invisible to the detector. The $E_{\mathrm{T}}^{\mathrm{miss}}$ is calculated as the sum of two terms: the "hard" term and the "soft" term. The hard term is computed from the momenta of all physics objects originating from the PV, meaning all leptons, photons, and jets. The soft term is computed from the momenta of all tracks that originate from the PV but are not associated with a physics object. The $E_{\mathrm{T}}^{\mathrm{miss}}$ exists only in the transverse plane, so it cannot be associated with a momentum vector in three-dimensional space, but its azimuthal angle, ϕ^{miss}, can be determined [61]. The shorthand $\vec{E}_{\mathrm{T}}^{\mathrm{miss}} = (E_{\mathrm{T}}^{\mathrm{miss}}, \phi^{\mathrm{miss}})$ is used to represent the $E_{\mathrm{T}}^{\mathrm{miss}}$ and its azimuthal angle.

4.4 Jet tagging

4.4.1 Vertex tagging

Jet vertex tagging is used to assess the likelihood that a given jet originated from the PV as opposed to a pileup interaction. The searches presented in this dissertation use three different vertex tagging algorithms for jets.

The jet vertex tagger (JVT) discriminant is built from two quantities that individually provide good discrimination between pileup and hard-scatter jets. The first quantity, $R_{p_{\mathrm{T}}}$, is the quotient of the scalar sum of the p_{T} of the tracks associated with the jet that originate from the PV and the jet's reconstructed p_{T}; this quantity peaks at zero for pileup jets, whose tracks do not primarily originate from the PV. The second quantity is the corrected jet vertex fraction—another ratio with the same numerator, but whose denominator is the scalar sum of the p_{T} of all tracks matched to the jet, regardless of their originating vertex, and corrected for the number of vertices in the event. It peaks near zero for pileup jets and near one for hard-scatter jets [62]. Because the JVT calculation relies on tracking information, JVT requirements are only imposed on small-R jets that fall within the ID region of $|\eta| < 2.5$. They are also only imposed on jets with $p_{\mathrm{T}} < 60\,\mathrm{GeV}$, which are more likely to originate from pileup vertices, and the JVT cut for these jets has an efficiency of about 93% for hard-scatter jets. The NNJVT algorithm is a next-generation version of JVT that uses a neural-network–based discriminant. Its efficiency ranges from 80% to 99% for hard-scatter jets [4].

The forward jet vertex tagger (fJVT) discriminant provides vertex tagging for jets with $2.5 < |\eta| < 4.5$, which are outside the ID volume. This discriminant is computed for each vertex by calculating the $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ at that vertex, projecting it onto the $\vec{p}_{\mathrm{T}}$ of the jet, and normalizing by the squared jet p_{T}. A cut is applied on largest fJVT value of all vertices. For pileup jets, fJVT tends toward 1, while for hard-scatter jets, it tends toward 0. This is because the p_{T} of a pileup jet is expected to be balanced by the $E_{\mathrm{T}}^{\mathrm{miss}}$ at its originating vertex [63]. The fJVT cut is applied to all jets with $|\eta| > 2.5$ and $p_{\mathrm{T}} > 50\,\mathrm{GeV}$. It accepts about 93% of hard-scatter jets and rejects about 58% of pileup jets [2].

4.4.2 Flavor tagging

Flavor tagging is used to identify jets originating from quarks of a particular flavor. The searches in this dissertation use b tagging to identify small-R jets from single b quarks and bb tagging to identify large-R jets from collimated $b\bar{b}$ decays. The presence of b quarks is identified thanks to the relatively long lifetime of b hadrons in comparison to hadrons composed of lighter quarks. The result is one or two decays within the jet that are slightly displaced from the PV, and the presence of these so-called secondary vertices is used to identify b and bb jets [64].

Single b tagging is the most widely used kind of flavor tagging, and several b tagging algorithms have been developed and used by ATLAS over the years. The MV2c10 algorithm was developed for Run 2 and was used in online b tagging; that is, b tagging in the HLT. It was also used for offline b tagging before more sophisticated algorithms were developed. The MV2c10 algorithm uses a discriminant built from a multivariate analysis using boosted decision trees. The input features include the presence of a secondary vertex and its properties, as well as the features of a reconstructed b hadron decay chain. MC simulated events of the $t\bar{t}$ process, which produces b jets, c jets, and "light" jets originating from gluons and u, d, and s quarks. The b jets are treated as signal and the others as background [64]. At the 77% efficiency working point, 77% of b jets are accepted, and the rejection factors for c jets and light jets are 6 and 134, respectively [2]. Here, a "rejection factor" is the reciprocal of the acceptance for a particular variety of background jet.

The newest generation of b tagging algorithms uses transformers. The GN2 algorithm uses a graph neural network to directly infer jet flavor from low-level track and jet inputs, without relying on intermediate reconstruction of secondary vertices. Training the network directly on the tracks allows it access to more finely-grained information and enhances its discrimination capabilities. The network is trained on MC simulated $t\bar{t}$ events to assess the probability that a particular jet originated from a b quark, from a c quark, from a τ lepton, or is a light jet; these probabilities are labeled p_b, p_c, p_τ, and p_u, respectively. A discriminant for b jet identification is built from these probabilities:

$$D_b = \ln\left(\frac{p_b}{f_c p_c + f_\tau p_\tau + (1 - f_c - f_\tau) p_u}\right), \tag{4.2}$$

where $f_c = 0.2$ and $f_\tau = 0.05$ are constant weight factors used to balance the rejection of the various background jet varieties. Compared to previous algorithms, the c and light jet rejection is improved by a factor of 2. The GN2 working points are defined by the efficiency to correctly tag b jets [65].

GN2X is another transformer-based algorithm, similar to GN2, that is used to identify large-R jets originating from $b\bar{b}$ decays. The target process is $H \to b\bar{b}$, but the working points are designed to provide a constant $b\bar{b}$ identification efficiency over the mass range $50 < m < 250\,\text{GeV}$. The network is trained on simulated $H \to b\bar{b}$, $H \to c\bar{c}$, $t\bar{t}$, and multijet events to assess the probability that a jet originated from a $b\bar{b}$ decay, from a $c\bar{c}$ decay, from a t quark, or from a multijet event; the respective probabilities are labeled p_{Hbb}, p_{Hcc}, p_{top}, and p_{QCD}. The discriminant for $b\bar{b}$ jet identification is constructed similarly to the GN2 discriminant:

$$D_{\text{Hbb}} = \ln\left(\frac{p_{\text{Hbb}}}{f_{\text{Hcc}} p_{\text{Hcc}} + f_{\text{top}} p_{\text{top}} + (1 - f_{\text{Hcc}} - f_{\text{top}}) p_{\text{QCD}}}\right), \tag{4.3}$$

where f_{Hcc} and f_{top} are weight factors used to balance the $c\bar{c}$ and t quark jet rejection. For the search presented in chapter 6 that uses the GN2X tagger, $f_{\text{Hcc}} = 0.02$ while $f_{\text{top}} = 0$, meaning that p_{top} is not included in the discriminant. This is to mitigate an observed dependence of p_{top} on the data-taking year. The resulting decrease in t quark rejection has a small effect on the analysis sensitivity. The GN2X working points are defined by the efficiency to incorrectly tag jets from multijet events as $b\bar{b}$ jets, called a mistag rate. At the 0.37% mistag rate working point, the efficiency to correctly tag $b\bar{b}$ jets is about 55% [66].

CHAPTER 5

SEARCH FOR INVISIBLE AND PARTIALLY INVISIBLE HIGGS BOSON DECAYS

Invisible decays of the Higgs boson, denoted $H \to$ inv., occur in the SM via the process $H \to ZZ^* \to \nu\bar{\nu}\nu\bar{\nu}$ with a branching fraction of about 0.1% [67]. This is two orders of magnitude smaller than the current experimental reach of the LHC—the best available upper limit on the $H \to$ inv. branching fraction, from ATLAS results combining the full Run 1 and Run 2 datasets, is 10.7% [68].

With this in mind, current searches for $H \to$ inv. are primarily motivated by BSM scenarios. Of particular interest are those involving dark matter, which is invisible by nature. Many theoretical models posit the existence of a "dark sector" containing one or more DM candidates [69]. In "Higgs portal" models, the Higgs boson is the sole link between this dark sector and the SM [70]. In $H \to$ inv. searches at the LHC, the notation $H \to \chi\chi$ is frequently used to represent a Higgs decay to DM particles, but the precise nature and number of invisible particles is open to interpretation.

A different class of DM candidates may be captured by partially invisible Higgs decays. One example is $H \to \gamma\gamma_{\mathrm{d}}$, where γ_{d} is a "dark photon." Dark photons may mediate long-range interactions in the dark sector, similarly to how light photons mediate electromagnetic interactions in the SM [71]. Other hypothetical particles may also decay to $\gamma\gamma_{\mathrm{d}}$, and searches for $H \to \gamma\gamma_{\mathrm{d}}$ can be interpreted in this context as well.

The first ATLAS analysis presented in this dissertation is published in ref. 2. It is a search for three different processes with the same final state, which consists of two jets with large pseudorapidity separation, large $E_{\mathrm{T}}^{\mathrm{miss}}$, and a central photon. This signature, called "VBF + $E_{\mathrm{T}}^{\mathrm{miss}} + \gamma$," was chosen to target the invisible decay of a Higgs boson produced via vector boson fusion (VBF) with a photon emitted during the initial state.

The VBF production mode of the Higgs boson occurs when two quarks from incoming

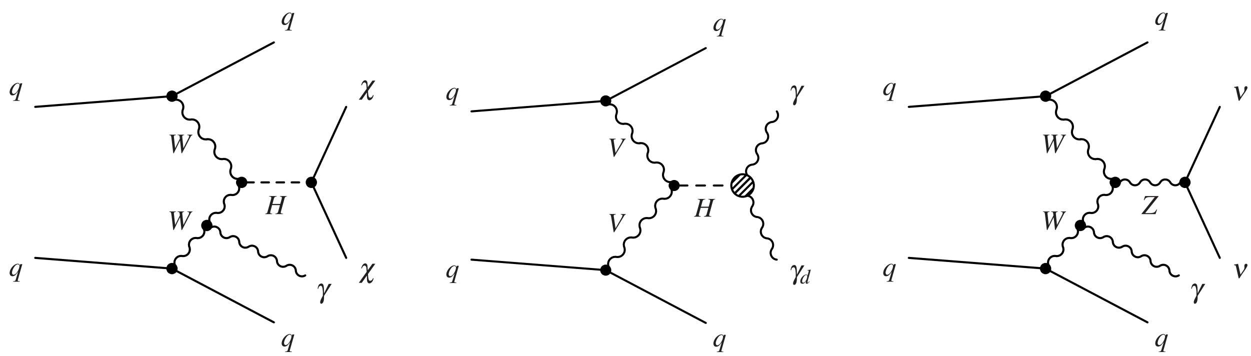

(a) $H \to$ inv. signal (b) $H \to \gamma\gamma_{\mathrm{d}}$ signal (c) EW $Z(\to \nu\nu)\gamma$jj signal

Figure 5.1: Feynman diagrams representing the three signal processes [2].

protons glance past one another and each emit a W or Z boson, which "fuse" and create a Higgs boson. When a photon is emitted in the initial state, both bosons must be W bosons, since Z bosons do not carry electric charge. This process, culminating in the $H \to$ inv. decay, is shown as a Feynman diagram in fig. 5.1a. The initial-state photon requirement decreases the cross section by about two orders of magnitude compared to VBF production with no such restriction [72], but dramatically reduces contributions from background processes.

The photon requirement also opens new avenues for interpretation. The $H \to \gamma\gamma_{\mathrm{d}}$ decay has a nearly identical signature, so long as the H is produced via VBF as shown in fig. 5.1b. In addition, one of the largest backgrounds in both the $H \to$ inv. and $H \to \gamma\gamma_{\mathrm{d}}$ interpretations is electroweak (EW) $Z(\to \nu\nu)\gamma$jj production, which had not been observed prior to this search. The leading Feynman diagram for this process, which accounts for about 85% of its cross section, is shown in fig. 5.1c; it is the same as the $H \to$ inv. diagram shown in fig. 5.1a, but with the $H \to \chi\chi$ portion replaced by $Z \to \nu\bar{\nu}$.

Both the ATLAS and CMS collaborations have published other searches for the $H \to$ inv. decay. Prior to the publication of ref. 2, both collaborations had published limits on the $H \to$ inv. branching ratio using combinations of several indepedent searches that used the full Run 1 data and the 2015–2016 Run 2 data [73, 74]. After the publication of ref. 2, ATLAS published an updated combination including this and other new Run 2 results [68].

CMS has also published updated results that include more recent searches [75], but to date has not published a search for $H \to$ inv. that includes a photon in the initial state.

Both collaborations have also published other results related to the $H \to \gamma\gamma_{\mathrm{d}}$ and EW $Z(\to \nu\nu)\gamma$jj interpretations. CMS published earlier limits on the $H \to \gamma\gamma_{\mathrm{d}}$ process in the ZH and VBF production modes [76, 77]. ATLAS later published a search targeting the ZH production mode [78], and another targeting the gluon–gluon fusion (ggF) production mode [79]. CMS published the first observation of the EW $Z\gamma$jj process using the $Z \to e^+e^-$ and $Z \to \mu^+\mu_-$ decay modes [80], but has yet to publish a search using the $Z \to \nu\overline{\nu}$ decay mode. ATLAS, on the other hand, has published a second search that includes a combination with the ref. 2 result and enhances the observation sensitivity [3].

This chapter presents the ATLAS VBF $+ E_{\mathrm{T}}^{\mathrm{miss}} + \gamma$ search and its results. It begins with a description of the overall analysis strategy in section 5.1. This is followed by a description of the trigger strategy, physics objects, and event selection in section 5.2. Section 5.3 describes the background processes and estimation methods. The systematic uncertainties are described in section 5.4. Section 5.5 presents the statistical analysis and results for all three interpretations. Finally, conclusions are presented in section 5.6.

5.1 Analysis strategy

All three searches follow the same basic strategy. For each signal process, a signal region (SR) is defined to target the VBF $+ E_{\mathrm{T}}^{\mathrm{miss}} + \gamma$ final state, in addition to four control regions (CR) that each target a different group of background processes. A simultaneous fit is performed to the signal and control regions to extract a signal strength while constraining the background contributions. The SR and CR definitions for the $H \to$ inv. and EW $Z(\to \nu\nu)\gamma$jj searches are the same, while those for the $H \to \gamma\gamma_{\mathrm{d}}$ search are slightly different to better represent the kinematics of that process.

For the signal $H \to$ inv. and $H \to \gamma\gamma_{\mathrm{d}}$ processes, both the VBF and ggF production

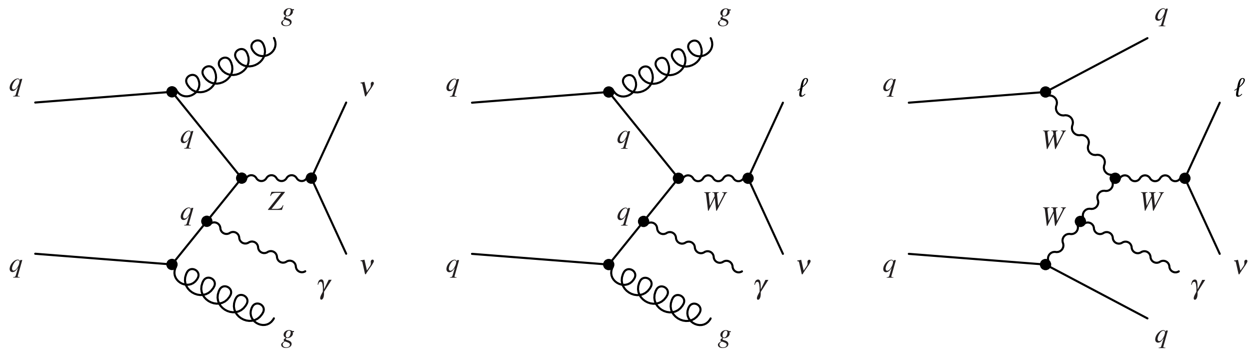

(a) Strong $Z(\to \nu\nu)\gamma$jj (b) Strong $W(\to \ell\nu)\gamma$jj (c) EW $W(\to \ell\nu)\gamma$jj

Figure 5.2: Feynman diagrams representing the major background processes [2].

modes are considered signal even though the search targets the VBF topology. The $H \to$ inv. signal process is modeled using SM $H \to ZZ^* \to \nu\bar{\nu}\nu\bar{\nu}$ events, which are kinematically indistinguishable from H decays to invisible BSM particles. The $H \to \gamma\gamma_\mathrm{d}$ process is modeled for a 125 GeV Higgs boson and for 13 additional masses ranging from 60 GeV to 2 TeV, representing hypothetical BSM scalars. The SR for the $H \to \gamma\gamma_\mathrm{d}$ search is slightly different from that of the $H \to$ inv. search mainly to increase the efficiency for the low- and high-mass BSM scalars.

All SR and CR definitons are informed by the dominant contributions to the VBF + $E_\mathrm{T}^\mathrm{miss} + \gamma$ final state. In addition to the EW $Z(\to \nu\nu)\gamma$jj process mentioned previously, "strong" $Z(\to \nu\nu)\gamma$jj production can occur when a Z boson is produced in a quark–antiquark collision if the quark and antiquark each emit a gluon and one also emits a photon. W bosons can also be produced via similar strong and EW processes. W boson decays can produce $E_\mathrm{T}^\mathrm{miss}$ via the $W^- \to \ell^-\bar{\nu}$ and $W^+ \to \ell^+\nu$ decay modes (where $\ell \in e, \mu, \tau$), which have a total branching fraction of about 33% [9]. It is possible to decrease contributions from $W(\to \ell\nu)\gamma$jj by vetoing events with an electron or a muon, but this is impossible if the lepton is "lost" because it falls outside the detector volume or is not reconstructed for other reasons. Feynman diagrams representing these backgrounds are shown in fig. 5.2.

The CRs are defined to constrain the contributions from $Z(\to \nu\nu)\gamma$jj and $W(\to \ell\nu)\gamma$jj by targeting the events from these processes that fall outside the SR. For $Z(\rightarrow \nu\nu)\gamma$jj, this means selecting events with a photon that is not produced centrally; these events are captured by the "$Z^{\gamma}_{\text{Rev.Cen}}$" CR, which is named for its reversal of the photon centrality requirement in the selection described in section 5.2. For $W(\rightarrow \ell\nu)\gamma$jj, this means selecting events that *do* have a reconstructed electron or muon. The CRs targeting these processes are called the "$W^{\gamma}_{e\nu}$" and "$W^{\gamma}_{\mu\nu}$" CRs. The $W^{\gamma}_{e\nu}$ CR in particular has a difficult-to-measure background arising from jets that are mistakenly reconstructed as electrons, which is constrained by a fourth CR, called the "Fake-e CR."

5.2 Triggers, physics objects, and event selection

The event selection aims to reconstruct the VBF + $E_{\text{T}}^{\text{miss}} + \gamma$ signature depicted in fig. 5.3, which consists of two forward jets in opposite detector hemispheres, large $E_{\text{T}}^{\text{miss}}$, and one central photon. This section describes the selection criteria. Section 5.2.1 lists the trigger and object requirements, section 5.2.2 describes the event selection for the SR, and section 5.2.3 describes the event selection for the CRs.

5.2.1 Triggers and physics objects

The principal triggers used in this search are $E_{\text{T}}^{\text{miss}}$ triggers with thresholds ranging from 70 to 110 GeV. Using $E_{\text{T}}^{\text{miss}}$ as opposed to photon triggers enables us to select photons with less energy than the photon trigger thresholds, which are about 140 GeV. Because this search targets forward jets, jet triggers have low efficiency and are therefore not used. Triggers requiring a single electron or muon are used for the $W^{\gamma}_{e\nu}$ and $W^{\gamma}_{\mu\nu}$ CRs. The lowest p_{T} thresholds for these triggers range from 20 to 26 GeV.

The main physics objects used in this search are jets, $E_{\text{T}}^{\text{miss}}$, and photons. Events with an electron or muon candidate are vetoed from the SR. Table 5.1 summarizes the object requirements for the SR. Overlap removal is performed among all objects using a detailed

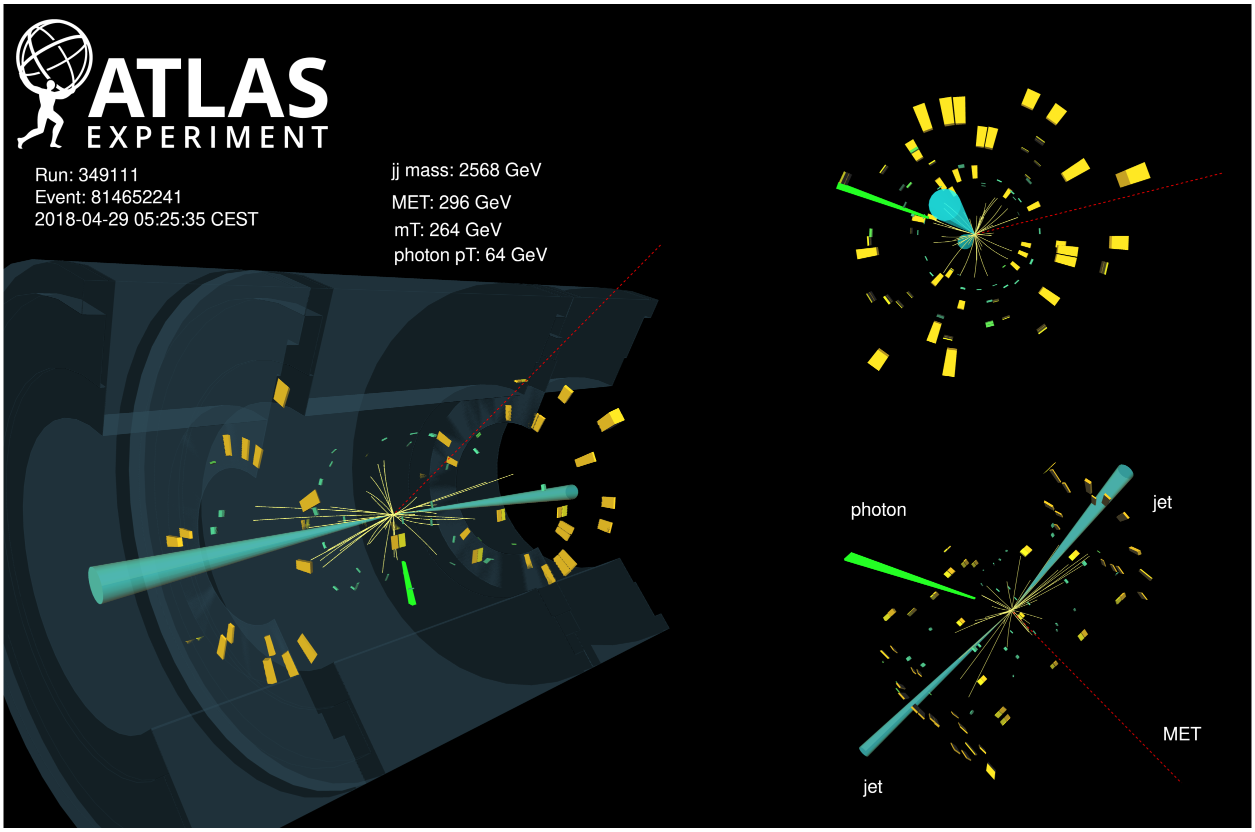

Figure 5.3: ATLAS display of the VBF + $E_{\mathrm{T}}^{\mathrm{miss}}$ + γ signature in an event consistent with the signal $H \to$ inv. and $H \to \gamma\gamma_{\mathrm{d}}$ processes [2].

set of criteria outlined in ref. 2.

The jets used are $R = 0.4$ jets as described in section 4.3.5. As the final state signature requires two jets with large separation, both central and forward jets are used. Central jets with $p_{\mathrm{T}} < 60\,\mathrm{GeV}$ must pass the JVT discriminant, while the leading two forward jets must pass the fJVT discriminant, both described in section 4.4.1. Only jets with $p_{\mathrm{T}} > 25\,\mathrm{GeV}$ are included, and higher p_{T} requirements are imposed on the forward jets during the event selection.

The definition of $E_{\mathrm{T}}^{\mathrm{miss}}$ is given in section 4.3.6. Crucially for this search, the $E_{\mathrm{T}}^{\mathrm{miss}}$ calculation does not include jets with $p_{\mathrm{T}} < 60\,\mathrm{GeV}$ that fail the JVT requirement, which means that "fake" $E_{\mathrm{T}}^{\mathrm{miss}}$ may be created by jets failing the requirement being omitted from the calculation. It is possible to remove many of these events by applying a cut to the

Table 5.1: Object requirements for $E_{\mathrm{T}}^{\mathrm{miss}}$, photons, central and forward jets, electrons, and muons in the SR.

Attribute	Requirement					
	$E_{\mathrm{T}}^{\mathrm{miss}}$	Photons	Forward jets	Central jets	Electrons	Muons
Number	—	1	2 or 3 total jets		0	0
Minimum E_{T}	150 GeV	—	—	—	—	—
Minimum p_{T}	—	15 GeV	50, 60 GeV	25 GeV	4.5 GeV	4 GeV
Minimum $\lvert\eta\rvert$	—	0	0	2.5	0	0
Maximum $\lvert\eta\rvert$	—	2.37	2.5	4.5	2.47	2.7
Vertex tagging	—	—	fJVT	JVT	—	—
Identification	—	Tight	—	—	Loose	Very loose
Isolation	—	Tight	—	—	None	None

quantity $E_{\mathrm{T}}^{\mathrm{jets,no\text{-}jvt}}$, which is the magnitude of the sum of all jet momenta, including jets that fail the JVT requirement so long as they have $p_{\mathrm{T}} > 20\,\mathrm{GeV}$. For the CRs with the lepton veto inverted, the quantity $E_{\mathrm{T}}^{\mathrm{miss,lep\text{-}rm}}$ is defined to estimate what the $E_{\mathrm{T}}^{\mathrm{miss}}$ would be if the lepton were lost; it is calculated from the sum of $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ and the lepton's momentum.

Photons, electrons, and muons are reconstructed as described in section 4.3.3, section 4.3.2, and section 4.3.4, respectively. Photons are required to fulfill the "tight" identification and isolation criteria. For the purposes of the SR veto, electrons and muons have very low p_{T} thresholds, are required to meet only "loose" or "very loose" identification criteria, and have no isolation requirements applied. For the $W_{e\nu}^{\gamma}$, $W_{\mu\nu}^{\gamma}$, and Fake-e CRs, $p_{\mathrm{T}} > 30\,\mathrm{GeV}$ is required for the lepton, and tighter identification and isolation requirements are applied. Full details on the requirements for CR leptons may be found in ref. 2.

5.2.2 SR event selection

Events must pass one of the $E_{\mathrm{T}}^{\mathrm{miss}}$ triggers described in section 5.2.1. After this trigger requirement, events must contain exactly two or three jets with $p_{\mathrm{T}} > 25\,\mathrm{GeV}$, exactly one photon, and no electrons or muons, using the object definitions summarized in table 5.1. The leading jet must have $p_{\mathrm{T}} > 60\,\mathrm{GeV}$ and the subleading jet must have $p_{\mathrm{T}} > 50\,\mathrm{GeV}$.

To ensure orthogonality with $t\bar{t}H \to$ inv. searches, no more than one b tagged jet may be present, using the 77% working point of the MV2c10 algorithm described in section 4.4.2. The b tagging veto has a negligible effect on this search's sensitivity.

The VBF signature must be present in each event. The leading two jets must have pseudorapidity separation $|\Delta\eta(\mathrm{j}_1, \mathrm{j}_2)| > 3.0$ and lie in opposite longitudinal hemispheres, meaning $\eta_{\mathrm{j}_1}\eta_{\mathrm{j}_2} < 0$. The dijet pair formed from these two jets must also have large invariant mass, $m_{\mathrm{jj}} > 250\,\mathrm{GeV}$. Finally, they must not be back-to-back in the transverse place, defined by a maximum allowed value of $\Delta\phi(\mathrm{j}_1, \mathrm{j}_2)$, which is 2.5 for the $H \to$ inv. SR and 2.0 for the $H \to \gamma\gamma_{\mathrm{d}}$ SR. This requirement reduces background from multijet events, which tend to have large azimuthal separation. The difference in the cut value between the SRs accounts for the different origin of the photon in the $H \to \gamma\gamma_{\mathrm{d}}$ decay.

For events with a third jet present, it must be consistent with final-state radiation from one of the two leading jets. The VBF signal process is characterized by limited hadronic activity in the gap between the two characteristic jets in contrast to the background QCD multijet and strong $Z(\to \nu\nu)\gamma\mathrm{jj}$ and $W(\to \ell\nu)\gamma\mathrm{jj}$ processes. This requirement is enforced via a cut on the centrality of the third jet, with centrality defined by

$$C_i = \exp\left(-\frac{4}{(\eta_{\mathrm{j}_1} - \eta_{\mathrm{j}_2})^2}\left(\eta_i - \frac{\eta_{\mathrm{j}_1} + \eta_{\mathrm{j}_2}}{2}\right)^2\right). \tag{5.1}$$

Here, C_i represents the centrality of some object i relative to the leading and subleading jets. C_i approaches 1 when the object approaches perfect centrality and e^{-1} when it approaches alignment with either jet. The centrality requirement for the third jet is $C_{\mathrm{j}_3} < 0.7$.

Next, requirements are imposed on the $E_{\mathrm{T}}^{\mathrm{miss}}$. The trigger thresholds influence the minimum $E_{\mathrm{T}}^{\mathrm{miss}}$ requirement, which is 150 GeV. This threshold ensures an overall trigger efficiency of 97% for the $H \to$ inv. signal. We also require $E_{\mathrm{T}}^{\mathrm{jets,no\text{-}jvt}} > 130\,\mathrm{GeV}$ to exclude multijet events with large "fake" $E_{\mathrm{T}}^{\mathrm{miss}}$ created by pileup or vertex tagging inefficiencies. The direction of the $E_{\mathrm{T}}^{\mathrm{miss}}$ must also be well separated from all jets in azimuth, meaning it must satisfy $\Delta\phi(\vec{E}_{\mathrm{T}}^{\mathrm{miss}}, \mathrm{j}) > 1$ for all selected jets, to further reduce backgrounds from fake $E_{\mathrm{T}}^{\mathrm{miss}}$ created by pileup jets.

Finally, the event must contain exactly one photon. The photon must be central, defined by $C_\gamma > 0.4$. To reduce contributions from pileup, its momentum must intersect with the beamline no more than 250 mm from the PV. A maximum p_{T} requirement is also imposed to reduce contributions from multijet events with a high-energy photon, described further in section 5.3. For the $H \to$ inv. SR, the requirement is $p_{\mathrm{T}} < 110\,\mathrm{GeV}$. For the $H \to \gamma\gamma_{\mathrm{d}}$ SR, it is relaxed to $p_{\mathrm{T}} < 0.733 m_{\mathrm{T}}$, where m_{T} is the transverse mass of the H candidate consisting of the photon and the $E_{\mathrm{T}}^{\mathrm{miss}}$, if $m_{\mathrm{T}} > 150\,\mathrm{GeV}$. This enhances the search sensitivity at large m_{T} and therefore allows us to probe more massive scalar mediators. In the $H \to$ inv. SR, we require $\Delta\phi(\vec{E}_{\mathrm{T}}^{\mathrm{miss}}, \gamma) > 1.8$ to ensure that the photon and $E_{\mathrm{T}}^{\mathrm{miss}}$ are back to back; this requirement is dropped in the $H \to \gamma\gamma_{\mathrm{d}}$ search to enhance the sensitivity at low m_{T}.

5.2.3 CR event selection

The event selection for the CRs is the same as the SR, but with one or two changes made to increase the selection efficiency for events of a particular background process.

$W_{\mu\nu}^{\gamma}$ CR

The SR lepton veto is replaced with the requirement that the event contain exactly one muon with $p_{\mathrm{T}} > 30\,\mathrm{GeV}$. Events with electrons are vetoed. The $E_{\mathrm{T}}^{\mathrm{miss}}$ requirement is replaced by $E_{\mathrm{T}}^{\mathrm{miss,lep\text{-}rm}} > 150\,\mathrm{GeV}$.

$W_{e\nu}^{\gamma}$ CR

The SR lepton veto is replaced with the requirement that the event contain exactly one electron with $p_{\mathrm{T}} > 30\,\mathrm{GeV}$. Events with muons are vetoed. The $E_{\mathrm{T}}^{\mathrm{miss}}$ requirement is replaced by $E_{\mathrm{T}}^{\mathrm{miss,lep\text{-}rm}} > 150\,\mathrm{GeV}$. An additional requirement on the proper $E_{\mathrm{T}}^{\mathrm{miss}}$ is enforced, $E_{\mathrm{T}}^{\mathrm{miss}} > 80\,\mathrm{GeV}$, to reduce contributions from events with fake electrons.

Fake-e CR

The SR lepton veto is replaced with the requirement that the event contain exactly one electron with $p_{\mathrm{T}} > 30\,\mathrm{GeV}$. Events with muons are vetoed. The $E_{\mathrm{T}}^{\mathrm{miss}}$ requirement is replaced by $E_{\mathrm{T}}^{\mathrm{miss,lep\text{-}rm}} > 150\,\mathrm{GeV}$. This CR targets events with fake electrons, so the requirement $E_{\mathrm{T}}^{\mathrm{miss}} < 80\,\mathrm{GeV}$ is imposed.

$Z^{\gamma}_{\mathrm{Rev.Cen}}$ CR

The photon centrality requirement is inverted: $C_{\gamma} < 0.4$. This CR targets the strong $Z(\rightarrow \nu\nu)\gamma\mathrm{jj}$ and $\gamma + \mathrm{jet}$ backgrounds.

5.3 Background estimation

The majority of the background contributions in both the $H \rightarrow$ inv. and $H \rightarrow \gamma\gamma_{\mathrm{d}}$ searches are the strong and EW $Z(\rightarrow \nu\nu)\gamma\mathrm{jj}$ and $W(\rightarrow \ell\nu)\gamma\mathrm{jj}$ processes mentioned in section 5.1. Together, these backgrounds make up about 90% of expected events in the $H \rightarrow$ inv. search and about 60% of expected events in the $H \rightarrow \gamma\gamma_{\mathrm{d}}$ search. These backgrounds are estimated using MC simulation and constrained using the $W^{\gamma}_{e\nu}$, $W^{\gamma}_{\mu\nu}$, and $Z^{\gamma}_{\mathrm{Rev.Cen}}$ CRs. They are abbreviated as "Strong $V + \gamma$" and "EW $V + \gamma$," where V can be either Z or W.

Much smaller bacgkround contributions arise from t quark processes, including $t\bar{t}$ production in association with a photon, and from W and Z production in association with two photons. These backgrounds are grouped together as "$t\bar{t}\gamma/V\gamma\gamma$" and are estimated using MC simulation. They account for about 2% of events in the $H \rightarrow$ inv. SR and about 4% of events in the $H \rightarrow \gamma\gamma_{\mathrm{d}}$ SR.

Like all analyses with jets in the final state, there is a background from events with multiple jets, which we call "$\gamma+$jet" to refer to the presence of the photon. This background is less than 1% in the $H \to$ inv. search and about 4% in the $H \to \gamma\gamma_{\mathrm{d}}$ search; this is a factor of ten smaller than the multijet background in other $H \to$ inv. searches, such as the ATLAS search presented in ref. 81, thanks to the photon requirement. It is relatively uncommon for multijet events to have substantial $E_{\mathrm{T}}^{\mathrm{miss}}$. However, if a jet's momentum is reconstructed incorrectly, the $E_{\mathrm{T}}^{\mathrm{miss}}$ may become large enough for the event to pass the selection. This effect is accounted for in the jet smearing procedure used to estimate the $\gamma+$jet background, which is described in section 5.3.1.

Reconstruction difficulties also underlie the remaining background contributions. In some cases, a jet or an electron may be misreconstructed as a photon. We refer to the backgrounds from these effects as "jet $\to \gamma$" and "$e \to \gamma$," respectively. They originate from events with a W or Z boson produced in association with jets; one of the jets or an electron from the W or Z decay may be reconstructed as a photon. Because MC simulation cannot reliably reproduce these effects, they are estimated using partially data-driven techniques. The jet $\to \gamma$ background is estimated using an $ABCD$ extrapolation method, and the $e \to \gamma$ background is estimated using a fake rate method. Descriptions of these methods may be found in ref. 2. The jet $\to \gamma$ background is about 2% in both SRs, while the $e \to \gamma$ background is about 7% in the $H \to$ inv. SR and about 33% in the $H \to \gamma\gamma_{\mathrm{d}}$ SR. The large difference in the $e \to \gamma$ background is due to the looser photon p_{T} requirement in the $H \to \gamma\gamma_{\mathrm{d}}$ SR and the removal of the $\Delta\phi(\vec{E}_{\mathrm{T}}^{\mathrm{miss}}, \gamma)$ cut.

The background composition of the $H \to$ inv. SR is shown in fig. 5.4, and the background composition of the $H \to \gamma\gamma_{\mathrm{d}}$ SR is shown in fig. 5.5. An additional background from jets misreconstructed as electrons, called "jet $\to e$," shows up only in the $W_{e\nu}^{\gamma}$ CR and in the Fake-e CR, the latter of which is designed to constrain these contributions. The actual processes contributing to this background include $\gamma+$ jet events with one jet faking an

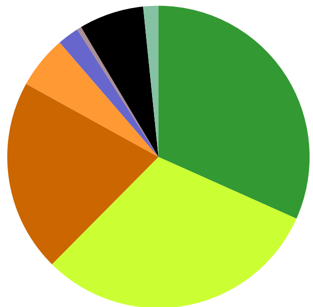

	Process	Fraction
	Strong $Z+\gamma$	31.69%
	EW $Z+\gamma$	30.75%
	Strong $W+\gamma$	20.60%
	EW $W+\gamma$	5.64%
	$t\bar{t}\gamma/V\gamma\gamma$	2.30%
	$\gamma+$jet	0.48%
	$e\to\gamma$	6.93%
	jet $\to\gamma$	1.62%

Figure 5.4: Background composition in the $H\to$ inv. SR.

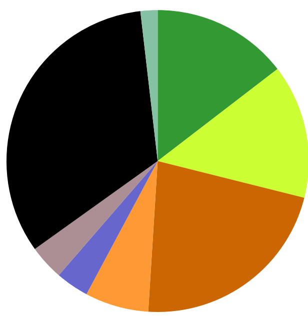

	Process	Fraction
	Strong $Z+\gamma$	14.56%
	EW $Z+\gamma$	14.35%
	Strong $W+\gamma$	22.08%
	EW $W+\gamma$	6.81%
	$t\bar{t}\gamma/V\gamma\gamma$	3.54%
	$\gamma+$jet	3.75%
	$e\to\gamma$	33.07%
	jet $\to\gamma$	1.83%

Figure 5.5: Background composition in the $H\to\gamma\gamma_{\mathrm{d}}$ SR.

electron and multijet events with one jet faking a photon and another faking an electron. This background is estimated using a fake rate method.

5.3.1 Jet smearing method for $\gamma+$jet background

In $\gamma+$jet events, true $E_{\mathrm{T}}^{\mathrm{miss}}$ can only arise from decays of "heavy flavor" hadrons within jets. Heavy flavor hadrons contain b or c quarks and can decay to final states including neutrinos. It is rare for the $E_{\mathrm{T}}^{\mathrm{miss}}$ created from these decays to be large. Studies of simulated $\gamma+$jet events suggest that 85% of them have truth $E_{\mathrm{T}}^{\mathrm{miss}}$ below 100 GeV. Therefore, the reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$ in selected $\gamma+$jet events originates primarily from mismeasurement of jet momenta.

One of the limitations of $\gamma+$jet MC, and multijet MC in general, is the number of events we are able to generate. The cross section for the $\gamma + \text{jet}$ process is huge, so even samples with tens of millions of events must have large event weights. The events in the samples used for this search had an average weight of about 20. These statistical limitations make it impossible to obtain accurate estimates of the number of $\gamma + \text{jet}$ events in the SR, must less draw any conclusions about the shapes of their distributions in any observable of interest.

This search uses a jet smearing technique to effectively increase the statistics of the $\gamma+$jet MC samples. The "smearing" consists of reproducing the effects of imperfect momentum reconstruction on truth jets by sampling p_{T}, ϕ, and η offsets from measured distributions. The truth $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ in each event is resampled as well. The events used for the smearing must pass a looser set of requirements than the SR event selection; in particular, the p_{T} requirements for the leading and subleading jets are lowered to 50 and 30 GeV, respectively, and no $E_{\mathrm{T}}^{\mathrm{miss}}$ requirement is imposed. The smearing is only applied to jets that pass certain isolation requirements and are matched to a truth jet within $\Delta R < 0.3$. Each event is smeared 20 times.

The jet p_{T} is smeared according to the jet energy response, which characterizes the accuracy and precision of the jet energy measurement. The jet energy response is defined by

$$R = \frac{p_{\mathrm{T}}^{\mathrm{reco}}}{p_{\mathrm{T}}^{\mathrm{truth}}}, \tag{5.2}$$

where $p_{\mathrm{T}}^{\mathrm{reco}}$ is the p_{T} of a reconstructed jet and $p_{\mathrm{T}}^{\mathrm{truth}}$ is the p_{T} of the closest truth jet. R is measured in bins of $p_{\mathrm{T}}^{\mathrm{truth}}$. The distributions tend to have a core that is roughly Gaussian and asymmetrical tails. They also tend to be biased, meaning that their peak is not at $R = 1$, especially in the low $p_{\mathrm{T}}^{\mathrm{truth}}$ bins, because reconstructed jets with $p_{\mathrm{T}}^{\mathrm{reco}} < 25\,\mathrm{GeV}$ were not retained. Each R distribution is fit to the sum of two Gaussians, and the result is adjusted so its peak is at $R = 1$. Figure 5.6 shows the biased and corrected R distributions for each of the five $p_{\mathrm{T}}^{\mathrm{truth}}$ bins in events simulated in accordance with 2018 detector conditions.

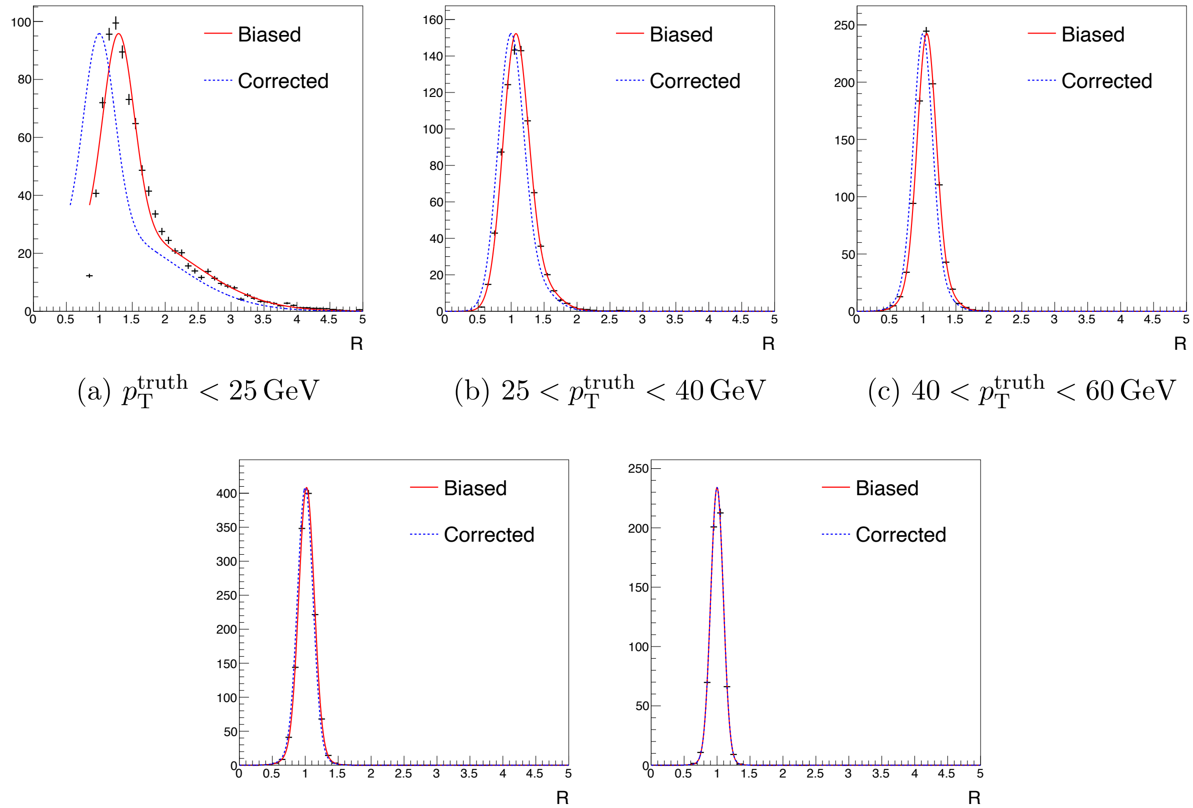

(a) $p_{\mathrm{T}}^{\mathrm{truth}} < 25\,\mathrm{GeV}$ (b) $25 < p_{\mathrm{T}}^{\mathrm{truth}} < 40\,\mathrm{GeV}$ (c) $40 < p_{\mathrm{T}}^{\mathrm{truth}} < 60\,\mathrm{GeV}$

(d) $60 < p_{\mathrm{T}}^{\mathrm{truth}} < 100\,\mathrm{GeV}$ (e) $p_{\mathrm{T}}^{\mathrm{truth}} > 100\,\mathrm{GeV}$

Figure 5.6: Fits of jet energy response R in simulated $\gamma+$jet events for each range of truth jet p_{T}. "Biased" represents the raw fit result, and "Corrected" represents the result after correcting the distribution so it peaks at $R = 1$.

The jet η and ϕ are smeared using Gaussian distributions consistent with the angular resolution of jet momenta in Run 2 data, and the jet mass is not smeared. The elements of a jet's smeared four-momentum are set as

$$p_{\mathrm{T}}^{\mathrm{smeared}} = \hat{R}\, p_{\mathrm{T}}^{\mathrm{truth}}, \tag{5.3}$$

$$\eta^{\mathrm{smeared}} = \eta^{\mathrm{reco}} + \hat{\eta}, \tag{5.4}$$

$$\phi^{\mathrm{smeared}} = \phi^{\mathrm{reco}} + \hat{\phi}, \tag{5.5}$$

$$m^{\mathrm{smeared}} = m^{\mathrm{reco}}, \tag{5.6}$$

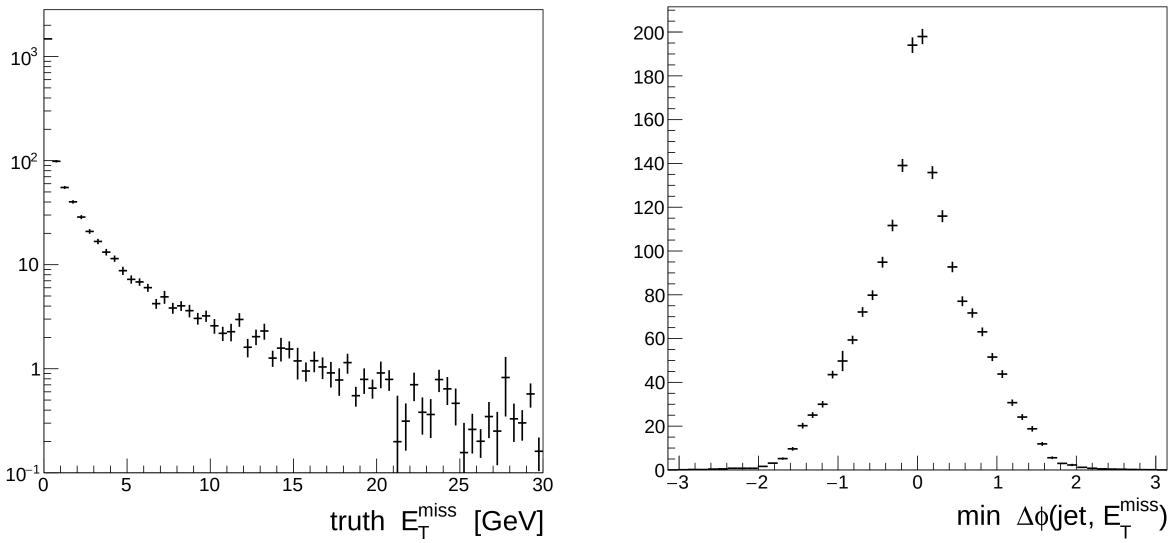

(a) Truth $E_{\mathrm{T}}^{\mathrm{miss}}$

(b) $\Delta\phi$ of truth $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ and nearest jet

Figure 5.7: Distributions of truth $E_{\mathrm{T}}^{\mathrm{miss}}$ and of $\Delta\phi$ between the truth $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ and the nearest jet in simulated $\gamma+$jet events.

where $\hat{R}$ is a value drawn from the R distribution appropriate for $p_{\mathrm{T}}^{\mathrm{truth}}$, and $\hat{\eta}$ and $\hat{\phi}$ are drawn from the angular resolution Gaussian distributions. The JVT and fJVT scores are recalculated for each smeared jet. In addition, smeared jets are removed if $p_{\mathrm{T}}^{\mathrm{smeared}} < 25\,\mathrm{GeV}$ or if they overlap with the photon.

The truth $E_{\mathrm{T}}^{\mathrm{miss}}$ in each smeared event is resampled by generating a random truth $E_{\mathrm{T}}^{\mathrm{miss}}$ value and direction. This is achieved by sampling from the truth $E_{\mathrm{T}}^{\mathrm{miss}}$ distribution and the distribution of $\Delta\phi$ between the $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ and the nearest jet. Both distributions are shown in fig. 5.7. The nearest jet is also randomly assigned. Then the reconstructed $E_{\mathrm{T}}^{\mathrm{miss}}$ is recalculated by subtracting the original truth $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$ and adding the sampled $\vec{E}_{\mathrm{T}}^{\mathrm{miss}}$, and is subsequently corrected for the changes in jet momenta.

Distributions of critical observables in smeared events are validated by comparing them to original $\gamma+$jet MC events passing a subset of the event selection criteria. Comparisons of the $E_{\mathrm{T}}^{\mathrm{miss}}$, jet p_{T}, and m_{jj} distributions are shown in fig. 5.8. The agreement is acceptable everywhere except at low and high jet p_{T}. Because the $\gamma+$jet background is small and is

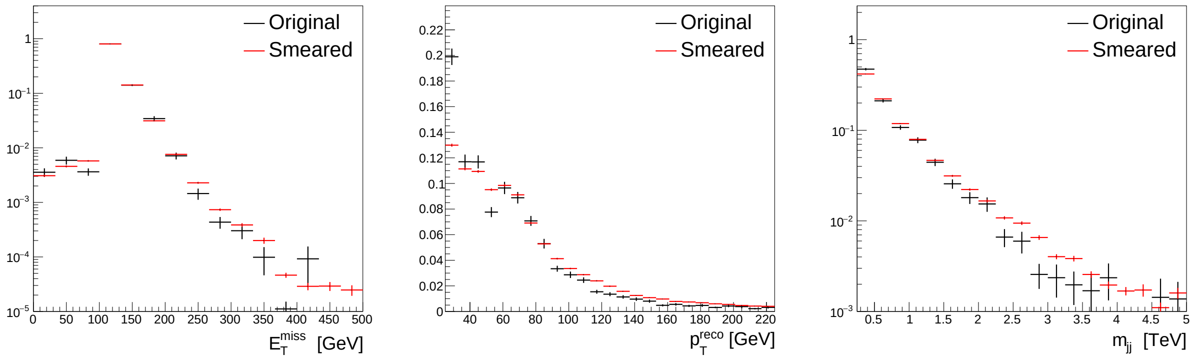

(a) $E_{\mathrm{T}}^{\mathrm{miss}}$ (b) Jet p_{T} (c) m_{jj}

Figure 5.8: Normalized distributions of $E_{\mathrm{T}}^{\mathrm{miss}}$, jet p_{T}, and m_{jj} for smeared and original (i.e., unsmeared) $\gamma+\mathrm{jet}$ events.

assigned a large uncertainty, we did not attempt to improve the modeling.

A "Low-$E_{\mathrm{T}}^{\mathrm{miss}}$" VR is defined to estimate the overall normalization of the $\gamma+\mathrm{jet}$ background. The selection criteria for this VR are identical to that of the SR except we require $110 < E_{\mathrm{T}}^{\mathrm{miss}} < 150\,\mathrm{GeV}$ and $0.25 < m_{\mathrm{jj}} < 1.0\,\mathrm{TeV}$, and we do not apply the $E_{\mathrm{T}}^{\mathrm{jets,no\text{-}jvt}}$ requirement or the upper bound on the photon p_{T}. The normalization factor is calculated from the ratio of observed and expected events after subtracting contributions from other backgrounds and found to be 0.91 ± 0.36, where the uncertainty is statistical only.

A systematic uncertainty is assessed by varying the smearing procedure in several ways, such as using the biased jet energy response functions, skipping the $E_{\mathrm{T}}^{\mathrm{miss}}$ resampling, and applying weights to correct the p_{T} distribution above 90 GeV. The spread of these variations is about 40%. An additional 75% uncertainty is applied to account for the extrapolation of the $E_{\mathrm{T}}^{\mathrm{miss}}$ trigger efficiency from the SR to the Low-$E_{\mathrm{T}}^{\mathrm{miss}}$ VR. Including the statistical component, the total uncertainty in the $\gamma+\mathrm{jet}$ background estimate is 94%.

5.4 Systematic uncertainties

Sources of systematic uncertainty in this search may be divided into three major groups: theoretical uncertaintites, experimental uncertainties, and background estimation uncertainties. Theoretical and experimental uncertainties affect both the background and signal processes, while the background modeling uncertainties affect only the background processes estimated using data-driven methods. This section describes all sources of systematic uncertainty and how each is assessed. Theoretical uncertainties are described in section 5.4.1, experimental uncertainties are described in section 5.4.2, and background modeling uncertainties are described in section 5.4.3.

5.4.1 Theoretical uncertainties

Theoretical uncertainties are most relevant for the $V+\gamma$ processes, in both the strong and EW production modes. QCD factorization and renormalization scale uncertainties are assessed by varying each by a factor of 0.5 or 2, separately and together, for a total of seven variations; the envelope of the overall difference in yield of each process ranges from about 25% to about 56% for the strong $V+\gamma$ processes, and from about 3% to about 11% for the EW $V+\gamma$ processes. Uncertainties from parton shower modeling are assessed using the scale variations for the strong $V+\gamma$ processes and by varying the shower modeling parameters for the EW $V+\gamma$ processes; they range from 4% to 15%. Parton distribution function (PDF) uncertainties are assessed using variations of the PDF set and are smaller than a few percent overall.

Some contributions to the $V+\gamma$ theory uncertainties only apply to specific processes. For the $Z(\rightarrow\nu\nu)\gamma$jj process, an additional uncertainty arises from interference between the strong and EW production modes, which is assigned to the EW $Z(\rightarrow\nu\nu)\gamma$jj process in the fit. This uncertainty is no more than 22% of the EW $Z(\rightarrow\nu\nu)\gamma$jj yield in each bin. Finally, for the strong $V+\gamma$ processes, an uncertainty due to the choice of MC generator is assessed by comparing the difference in shape and normalizaion between events from two different generators. This uncertainty is no larger than 20% in each bin.

Theoretical uncertainties for the VBF $H \to$ inv. signal samples are comparatively smaller. The envelope of the seven factorization and renormalization scale variations corresponds to an overall systematic uncertainty of about 1–2%. Uncertainties from parton shower modeling are assessed by comparing the differences between events showered using two different generators, and they range from 2% to 8%. Uncertainties from the choice of PDF are assessed using variations of the PDF set and are about 1–2%.

Theoretical uncertainties for the VBF $H \to \gamma\gamma_{\mathrm{d}}$ signals are slightly larger than for VBF $H \to$ inv. and depend on m_H, since this process is also modeled for BSM mediators. The factorization and renormalization scale uncertainties range from about 1% to 3%. The parton shower uncertainties are assessed by comparing two different generators, and range from 2% to 4%. The PDF uncertainties are no more than 2% in each bin. An additional 2% uncertainty from EW corrections at next-to-leading order is assessed for the $m_H = 125\,\mathrm{GeV}$ sample.

For both the $H \to$ inv. and $H \to \gamma\gamma_{\mathrm{d}}$ signals, the ggF production mode is included in the signal, although its contribution to the SR is subdominant. For the $H \to \gamma\gamma_{\mathrm{d}}$ process, the ggF contribution is about 20%. Theoretical uncertainties from the ggF $H \to$ inv. and $H \to \gamma\gamma_{\mathrm{d}}$ samples are dominated by the factorization and renormalization scale variations, which range from 20% to 30%.

5.4.2 Experimental uncertainties

Uncertainties are associated with the reconstruction of all physics objects listed in table 5.1. Jet energy scale, energy resolution, and vertex tagging uncertainties are assessed. Uncertainties in the electron and muon veto requirements are taken into account for events in the SR and $Z^{\gamma}_{\mathrm{Rev.Cen}}$ CR, while uncertainties in electron and muon identification, isolation, energy scale, and energy resolution are taken into account for events in the $W^{\gamma}_{e\nu}$ and $W^{\gamma}_{\mu\nu}$ CRs. In the SR and all CRs, photon identification, isolation, energy scale, and energy resolution are included. For the $E_{\mathrm{T}}^{\mathrm{miss}}$, the uncertainties from the reconstruction of the other objects are factored into the calculation, and additional uncertainties in the $E_{\mathrm{T}}^{\mathrm{miss}}$ scale and resolution are applied as well.

The $E_{\mathrm{T}}^{\mathrm{miss}}$ triggers described in section 5.2.1 are not perfectly efficient at the selection threshold of 150 GeV. Uncertainties are assessed to account for differences in efficiency between data and MC simulation. These differences range from about 3–6% near $E_{\mathrm{T}}^{\mathrm{miss}} = 150$ GeV to below 0.4% for $E_{\mathrm{T}}^{\mathrm{miss}} > 200$ GeV, and are applied as corrections to the simulated events. The uncertainties in the corrections range from 1–4% and originate from the statistical imprecision of the data used to derive them.

Additional sources of experimental uncertainty include pileup modeling and the luminosity measurement. The uncertainty in the integrated luminosity measurement for the combined 2015–2018 dataset was 1.7% at the time of this search. These uncertainties are small compared to those associated with object reconstruction.

5.4.3 Background modeling uncertainties

Dedicated uncertainties are assessed for the estimates of the backgrounds that arise from misreconstruction: $\gamma +$ jet, jet $\to \gamma$, $e \to \gamma$, and jet $\to e$. The uncertainty in the $\gamma +$ jet background estimate is assessed as described in section 5.3.1; the total uncertainty in this background estimate is 94%. The jet $\to \gamma$ background uncertainty is assessed by varying the photon identification and isolation requirements used to define the regions used for the $ABCD$ method extrapolation; the total uncertainty ranges from 80% to 90%. The $e \to \gamma$ contribution is assessed using a fake rate method, and the uncertainty in the fake rate, which is assessed by varying a number of factors in its estimation, ranges from 30% at $p_{\mathrm{T}} = 15$ GeV to 15% at $p_{\mathrm{T}} = 100$ GeV. The jet $\to e$ background contributes only to the $W^{\gamma}_{e\nu}$ and Fake-e

CRs; its uncertainty is about 80% in the $H \to$ inv. search and about 40% in the $H \to \gamma\gamma_{\mathrm{d}}$ search.

5.5 Statistical analysis and results

As mentioned in section 5.1, a simultaneous fit is performed to the signal and control regions in each search. In general, each background process is assigned a normalization factor that may be fixed to unity or allowed to float in the fit. Because the CRs are enriched in background events, their higher event yields and smaller theoretical uncertainties enable a much more precise background estimate than would be possible in the SR alone. All systematic uncertainties are correlated across the SR and CRs. Different discriminant variables are used in the fits for each signal process to optimize sensitivity.

This section describes the fit strategies used for each of the three signal processes and their results. The EW $Z(\to \nu\nu)\gamma\mathrm{jj}$ search is the simplest to describe, so it is presented first in section 5.5.1 as a segue to the other interpretations. Next, the $H \to$ inv. search is described in section 5.5.2. Finally, the $H \to \gamma\gamma_{\mathrm{d}}$ search is described in section 5.5.3.

5.5.1 *Search for EW $Z(\to \nu\nu)\gamma jj$*

As mentioned in section 5.1, the EW $Z(\to \nu\nu)\gamma\mathrm{jj}$ interpretation uses the same SR and CR definitions as the $H \to$ inv. search. The SR and CRs are binned in m_{jj}, the invariant mass of the two VBF jets, because this observable differentiates well between the signal EW and background strong $Z(\to \nu\nu)\gamma\mathrm{jj}$ processes. The normalization factors for the EW $Z(\to \nu\nu)\gamma\mathrm{jj}$ signal, the strong $Z(\to \nu\nu)\gamma\mathrm{jj}$ background, the total $W(\to \ell\nu)\gamma\mathrm{jj}$ background, and the jet $\to e$ background are allowed to float in the fit, while those for the other backgrounds are fixed.

Figure 5.9 shows the post-fit m_{jj} distribution in the SR and all CRs. The data and post-fit predictions have excellent agreement in every single bin, even within the data statistical

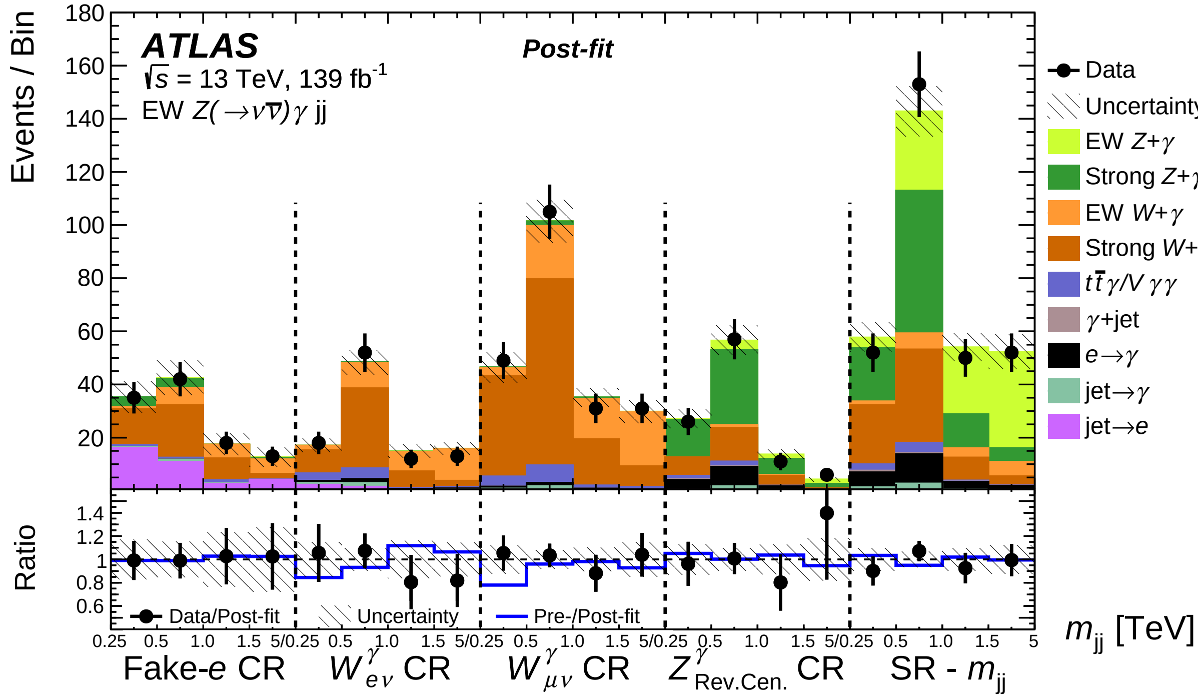

Figure 5.9: Post-fit results for all m_{jj} SR and CR bins in the EW $Z(\to \nu\nu)\gamma\text{jj}$ cross section measurement. The lower panel shows the ratio of data to the sum of all predicted contributions and a comparison of the pre-fit and post-fit predictions [2].

uncertainty. The best-fit normalization factors for the EW $Z(\to \nu\nu)\gamma\text{jj}$ signal, the strong $Z(\to \nu\nu)\gamma\text{jj}$ background, and the total $W(\to \ell\nu)\gamma\text{jj}$ background are all consistent with unity, while the best-fit normalization factor for the jet $\to e$ background is about 1.6. The difference between the pre-fit and post-fit distributions is shown as a ratio. The pre-fit agreement is also very good, and the difference between pre-fit and post-fit is small.

The best-fit value of the EW $Z(\to \nu\nu)\gamma\text{jj}$ normalization factor, called the "signal strength," is

$$\mu_{Z\gamma_{\text{EW}}} = 1.03 \pm 0.16(\text{stat.}) \pm 0.19(\text{syst.}) \pm 0.02(\text{lumi.}). \tag{5.7}$$

Compared to the background-only hypothesis of $\mu_{Z\gamma_{\text{EW}}} = 0$, the measured signal strength corresponds to an observation of the EW $Z(\to \nu\nu)\gamma\text{jj}$ process with an observed significance of 5.2σ. The expected significance is 5.1σ. Table 5.2 tabulates the contributions to the 1σ uncertainty in $\mu_{Z\gamma_{\text{EW}}}$. The data statistical uncertainty of 0.16 makes the largest contri-

Table 5.2: Contributions from different groups of systematic uncertainties to the 1σ uncertainty bands of the best-fit $\mu_{Z\gamma_{\mathrm{EW}}}$ value and on the 95% CL upper limits on $\mathcal{B}(H \to \mathrm{inv.})$ and $\mathcal{B}(H \to \gamma\gamma_{\mathrm{d}})$ [2].

	1σ uncertainty		
Source	$\mu_{Z\gamma_{\mathrm{EW}}}$	$\mathcal{B}(H \to \mathrm{inv.})$	$\mathcal{B}(H \to \gamma\gamma_{\mathrm{d}})$
Jet scale and resolution	0.076	0.045	0.0011
$V\gamma + \mathrm{jets}$ theory	0.067	0.044	0.0018
Pileup	0.040	0.021	0.0004
Photon	0.035	0.031	0.0011
$e \to \gamma$, jet $\to e$, and jet $\to \gamma$ backgrounds	0.035	0.034	0.0028
Lepton	0.027	0.003	0.0008
$E_{\mathrm{T}}^{\mathrm{miss}}$	0.023	0.018	0.0003
Signal theory shape	0.020	—	—
Signal theory acceptance	0.12	—	—
Data statistics	0.16	0.11	0.0056
$W\gamma + \mathrm{jets}$ and $Z\gamma + \mathrm{jets}$ normalization	0.073	0.013	0.0004
MC statistics	0.063	0.046	0.0026
Total	0.25	0.15	0.0073

bution, followed closely by theoretical uncertainties in the signal prediction, which have a combined uncertainty of about 0.12.

The best-fit value of $\mu_{Z\gamma_{\mathrm{EW}}}$ is used to extract a two-sided confidence interval for the fiducial cross section. The measured value is

$$\sigma^{\mathrm{fid.}}_{Z(\to\nu\nu)\gamma\mathrm{EW}} = 1.31 \pm 0.20(\mathrm{stat.}) \pm 0.20(\mathrm{syst.})\,\mathrm{fb}, \tag{5.8}$$

using the fiducial region defined in ref. 2. This is in good agreement with the theoretical cross section calculated by MadGraph [82], which is $1.27 \pm 0.17\,\mathrm{fb}$.

This observation of EW $Z(\to \nu\nu)\gamma jj$ is more significant than a dedicated ATLAS search published one year later in ref. 3. The dedicated search has an observed significance of 3.2σ and an expected significance of 3.7σ. The two searches are orthogonal in photon p_{T}, with the dedicated search relying on photon triggers and requiring a photon with $p_{\mathrm{T}} > 150\,\mathrm{GeV}$

to increase the search's sensitivity to anomalous quartic gauge couplings. This is the largest difference between the selection criteria of the two searches and is the primary driver for the difference in significance. The signal-to-background ratio of the EW $Z(\rightarrow \nu\nu)\gamma$jj SR in the VBF + $E_{\mathrm{T}}^{\mathrm{miss}}$ + γ search is approximately 30%, while in the dedicated search it is approximately 10%. The searches are complementary, and a combination was performed to increase the overall significance of the observation. The observed significance of the combination is 6.3σ, and its expected significance is 6.6σ.

5.5.2 Search for $H \rightarrow$ inv.

The $H \rightarrow$ inv. search uses a dense neural network (DNN) classifier to discriminate between signal and background processes. The classifier is trained on a mixture of labeled signal and background events that pass loosened selection criteria for the SR and CRs. The input features are a selection of kinematic observables constructed from the $E_{\mathrm{T}}^{\mathrm{miss}}$, the photon candidate, and the leading two jets. The trained classifier processes these observables for selected events and assigns each a "DNN score" from 0 to 1, with lower scores assigned to background-like events and higher scores assigned to signal-like events. Binning the SR and CRs using this DNN score provides an improvement in sensitivity of about 20% compared to the m_{jj} binning used for the EW $Z(\rightarrow \nu\nu)\gamma$jj search.

Another difference between the statistical analyses for the $H \rightarrow$ inv. and EW $Z(\rightarrow \nu\nu)\gamma$jj searches is which normalization factors are allowed to float. For the $H \rightarrow$ inv. search, the normalization factors for $W(\rightarrow \ell\nu)\gamma$jj and jet $\rightarrow e$ are allowed to float like in the EW $Z(\rightarrow \nu\nu)\gamma$jj search, but the normalization factors for both strong and EW $Z(\rightarrow \nu\nu)\gamma$jj are fixed to unity. For EW $Z(\rightarrow \nu\nu)\gamma$jj, this is because its DNN score distribution is so similar in shape to the $H \rightarrow$ inv. signal that allowing its normalization to float would have a detrimental effect on the signal strength. For strong $Z(\rightarrow \nu\nu)\gamma$jj, allowing it to float would slightly improve the constraints on the theoretical uncertainties associated with this background, but would also slightly degrade the signal strength. The benefit does not outweigh the cost, so it is fixed. The normalization factors for these backgrounds were found to be consistent with unity in the EW $Z(\to \nu\nu)\gamma\text{jj}$ fit.

Figure 5.10 shows the post-fit DNN score distribution in the SR and all CRs. The last two bins of the $Z^{\gamma}_{\text{Rev.Cen}}$ CR are merged into one because there are no background events in this CR with DNN score above 0.8. Agreement between data and the post-fit background-only prediction is good, with a slight discrepancy in only one of the 19 bins. The lack of an excess in the SR means that the $H \to$ inv. process is not observed. It is shown scaled to the observed 95% CL upper limit:

$$\mathcal{B}(H \to \text{inv.}) < 0.37\ (0.34^{+0.15}_{-0.10}\ \text{expected}). \tag{5.9}$$

The largest source of uncertainty is the data statistical uncertainty, which accounts for about 75% of the 1σ uncertainty. All uncertainty contributions are listed in table 5.2, which shows that the remaining uncertainty is due primarily to the MC statistical uncertainty, jet reconstruction systematics, and theoretical uncertainties in the $Z(\to \nu\nu)\gamma\text{jj}$ and $W(\to \ell\nu)\gamma\text{jj}$ background predictions.

5.5.3 Search for $H \to \gamma\gamma_d$

For the $H \to \gamma\gamma_{\text{d}}$ search, the SR and CRs are binned in the transverse mass m_{T} of the $E_{\text{T}}^{\text{miss}}$ and the photon, which is a proxy for the mass of the H or BSM mediator. The SR and all CRs are divided into two regions, one with $m_{\text{jj}} < 1\,\text{TeV}$ and one with $m_{\text{jj}} > 1\,\text{TeV}$, which are more sensitive to the ggF and VBF production modes, respectively. Both the ggF and VBF production modes are simulated for a 125 GeV Higgs mediator, while only the VBF mode is simulated for mediators of other masses. The same background normalization factors are fixed as for the $H \to$ inv. search, although allowing them to float does not affect

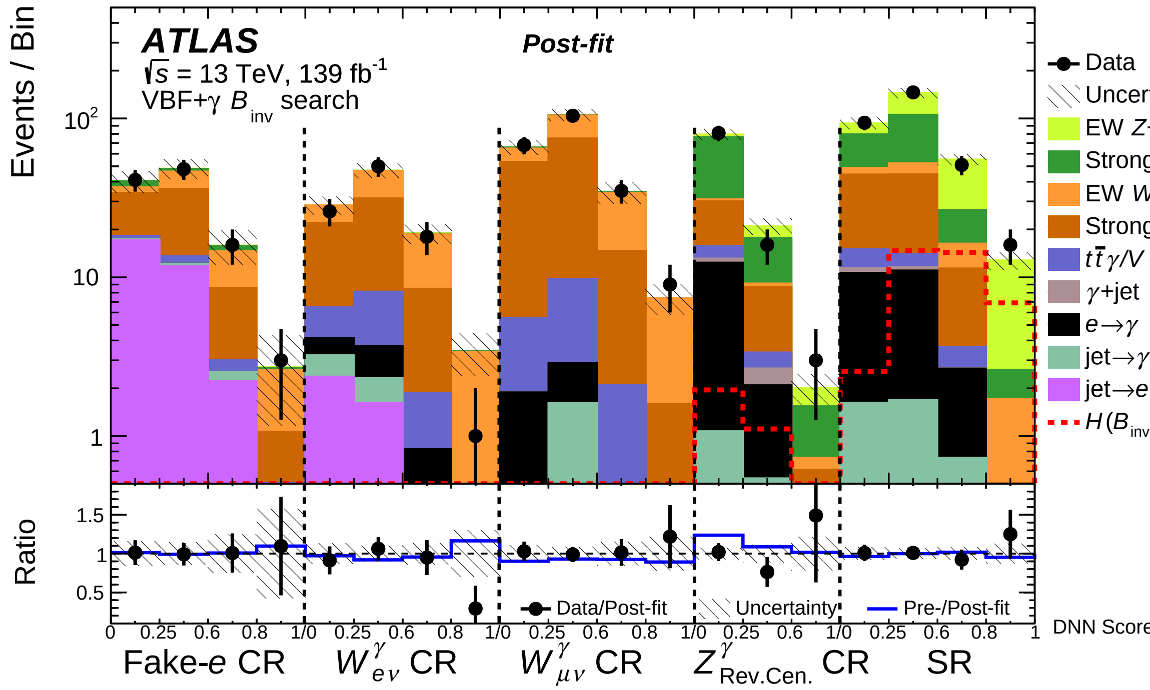

Figure 5.10: Post-fit results for all DNN SR and CR bins in the search for $H \to$ inv. The $H \to$ inv. signal is overlaid, scaled to a 37% $\mathcal{B}(H \to \text{inv.})$. The lower panel shows the ratio of data to the sum of all predicted contributions and a comparison of the pre-fit and post-fit predictions [2].

the fit results within two significant digits.

Figure 5.11 shows the post-fit m_{T} distributions in the $H \to \gamma\gamma_{\mathrm{d}}$ SR. The CRs are shown inclusively in the figure, although they use the same 10 bins as the SR in the fit. Agreement between data and the post-fit background-only prediction is reasonable overall, but the pre-fit to post-fit ratio suggests that the prediction in the high-m_{T} bins and the $Z^{\gamma}_{\mathrm{Rev.Cen}}$ CR is too high. The theoretical uncertainties in the $Z(\to \nu\nu)\gamma jj$ predictions are highest when m_{T} is large, which may contribute to the discrepancy. No evidence of $\gamma\gamma_{\mathrm{d}}$ production is observed for any scalar mediator mass. Signals for 125 GeV and 500 GeV mediators are shown as examples, scaled to the arbitrary branching fractions of 2% and 1%, respectively.

The same two mass hypotheses are shown in the inclusive SR in fig. 5.12, which demonstrates the deficit in the high-m_{T} bin. Because of this deficit, the observed exclusion power

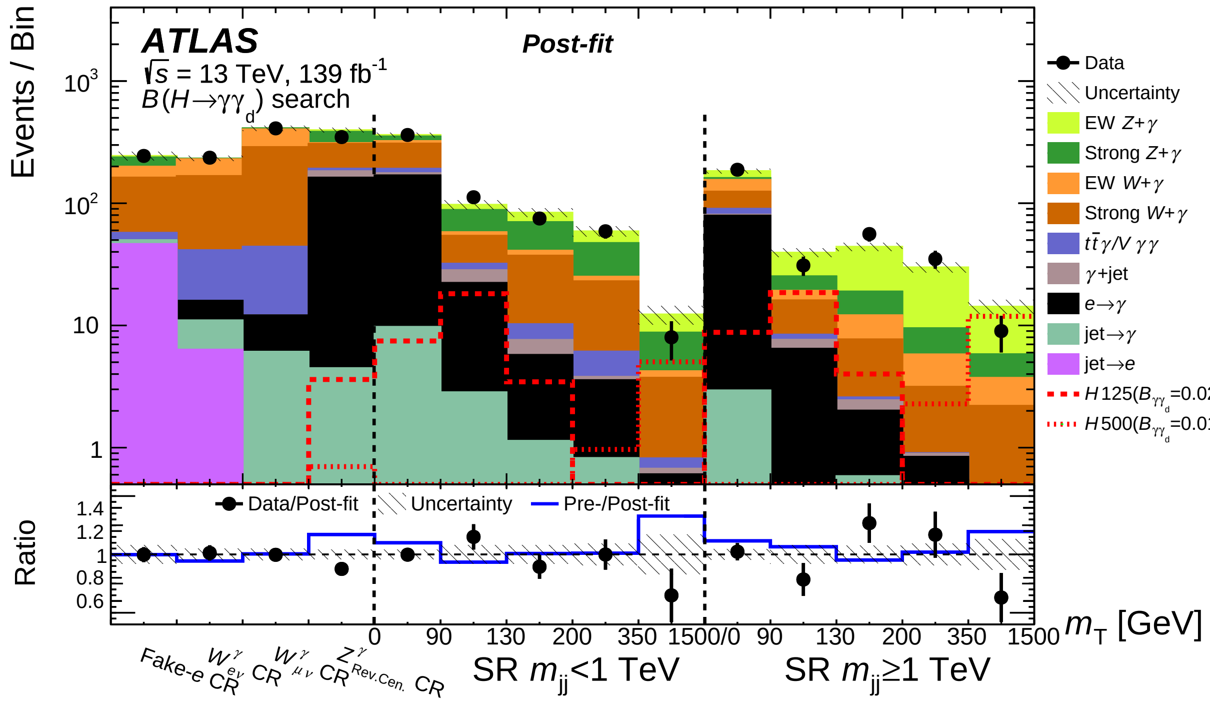

Figure 5.11: Post-fit results for the ten $(m_{\rm jj}, m_{\rm T})$ bins in the SR for the $H \to \gamma\gamma_{\rm d}$ search, and the CRs shown inclusively. The $H \to \gamma\gamma_{\rm d}$ signal is overlaid for two mass hypotheses. The lower panel shows the ratio of data to the sum of all predicted contributions and a comparison of the pre-fit and post-fit predictions [2].

is significantly better than the expected for mediator masses above the bin's low edge of 350 GeV. This is reflected in fig. 5.13, which shows observed and expected 95% CL upper limits on the VBF $H \to \gamma\gamma_{\rm d}$ cross section for m_H values ranging from 60 GeV to 2 TeV. Theoretical predictions for a 5% $\mathcal{B}(H \to \gamma\gamma_{\rm d})$ are overlaid to aid in interpreting these limits as a branching fraction. A 5% branching fraction is ruled out for all $m_H < 1.5$ TeV.

For $m_H = 125$ GeV, observed and expected 95% CL upper limits are set on $\mathcal{B}(H \to \gamma\gamma_{\rm d})$, considering both the VBF and ggF production modes:

$$\mathcal{B}(H \to \gamma\gamma_{\rm d}) < 0.018 \; (0.017^{+0.007}_{-0.005} \text{ expected}). \tag{5.10}$$

The data statistical uncertainty has the largest contribution to the 1σ uncertainty, with the remaining contributions arising primarily from the MC statistical uncertainty and the γ+jet

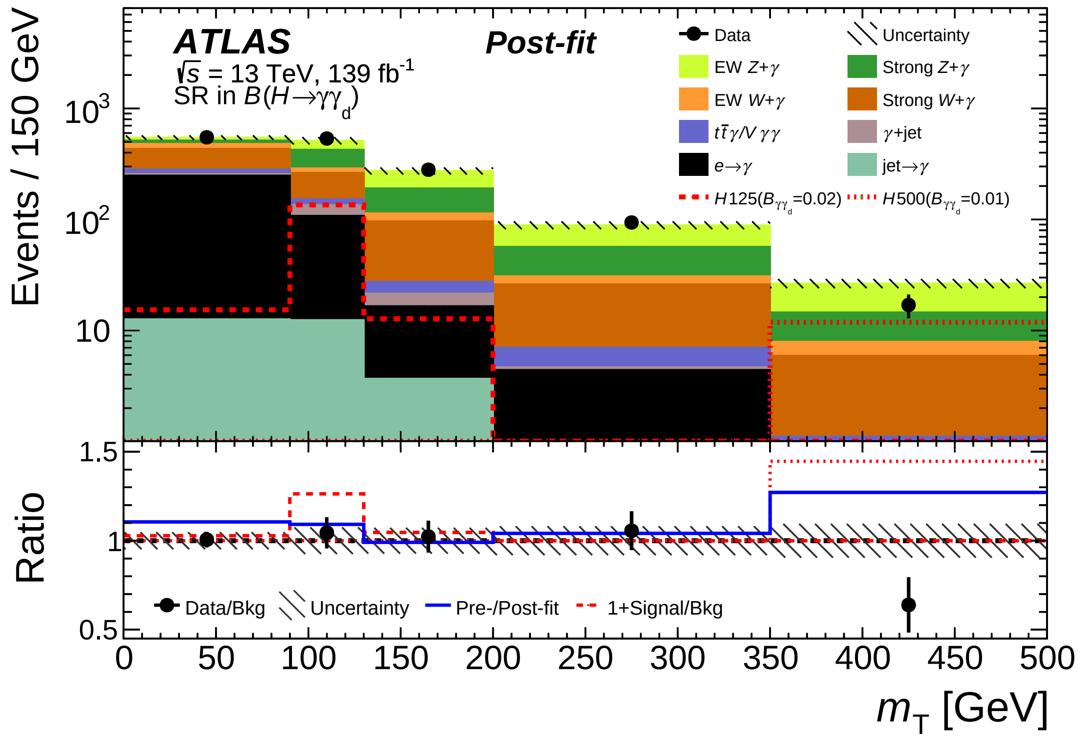

Figure 5.12: Post-fit m_T distribution in the inclusive $H \to \gamma\gamma_d$ SR with two example mass hypotheses overlaid. Events with $m_T > 500\,\text{GeV}$ are included in the final bin [2].

background estimation uncertainty. All contributions are shown in table 5.2. The search is far more sensitive to the VBF mode by design, which is illustrated by the 95% CL upper limits on the product of each cross section and the branching fraction:

$$\sigma(\text{VBF } H \to \gamma\gamma_d) > 0.064\,\text{pb}\ (0.064^{+0.030}_{-0.019}\,\text{pb expected}), \tag{5.11}$$

$$\sigma(\text{ggF } H \to \gamma\gamma_d) > 10.2\,\text{pb}\ (7.3^{+3.4}_{-1.9}\,\text{pb expected}). \tag{5.12}$$

5.6 Conclusion

This search, targeting the VBF + $E_T^{\text{miss}} + \gamma$ final state, uses $139\,\text{fb}^{-1}$ of data collected by the ATLAS detector at $\sqrt{s} = 13\,\text{TeV}$ during Run 2 of the LHC. The search primarily targets the $H \to$ inv. and $H \to \gamma\gamma_d$ processes, but also presents the first observation of the EW

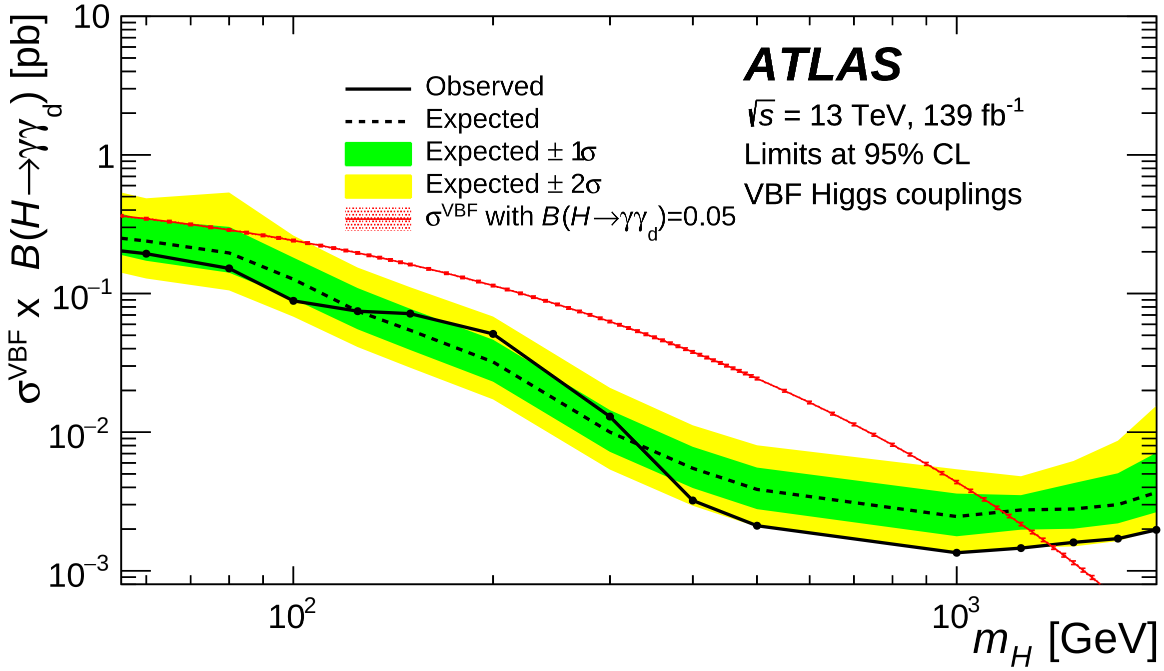

Figure 5.13: Expected and observed 95% CL upper limits on the product of the VBF production cross section and the $\gamma\gamma_\mathrm{d}$ branching fraction for a scalar mediator of various masses. Theoretical predictions for a 5% branching fraction are overlaid [2].

$Z(\rightarrow \nu\nu)\gamma$jj process. The $H \rightarrow$ inv. process is not observed, and a 95% CL upper limit on its branching ratio is set at 37%. The $H \rightarrow \gamma\gamma_\mathrm{d}$ process is not observed for any Higgs boson or BSM mediator mass between 60 GeV and 2 TeV. The search results exclude values of the product of the VBF H production cross section and the $H \rightarrow \gamma\gamma_\mathrm{d}$ branching ratio from 0.15 pb for $m_H = 60$ GeV to 3 fb for $m_H = 2$ TeV. For $m_H = 125$ GeV, a 95% CL upper limit on the $H \rightarrow \gamma\gamma_\mathrm{d}$ branching ratio is set at 1.8%. The EW $Z(\rightarrow \nu\nu)\gamma$jj process is observed with a significance of 5.2σ. A combination with a later ATLAS search [3] increases the significance to 6.6σ.

CHAPTER 6

SEARCH FOR NEW SCALAR RESONANCES PRODUCED IN ASSOCIATION WITH A HIGGS BOSON

Many proposed extensions of the SM include an "extended Higgs sector" that contains one or more BSM scalar fields. These models are intended to rectify specific limitations of the SM mentioned in section 2.4, such as the hierarchy problem that arises from the low mass of the Higgs boson [83]. An extended Higgs sector may arise by manually adding scalar fields to the SM Lagrangian, or as a natural consequence of more general SM extensions like supersymmetry. These models may also answer other questions the SM does not by virtue of containing, for example, dark matter candidates and explanations for the matter–antimatter asymmetry [84–86].

An extended Higgs sector may contain any number of BSM scalar or pseudoscalar bosons with properties that may, in general, be different from the SM Higgs. One such difference is the mass of the additional boson or bosons, but they may be even more exotic—for example, they may have nonzero charge. However, models with additional bosons that are identical to the 125 GeV Higgs in all properties except mass are arguably the simplest. In any case, the BSM bosons are expected to interact with the 125 GeV Higgs, so it can be used as a candle to illuminate search avenues that may lead to the discovery of an extended Higgs sector.

The second ATLAS analysis presented in this dissertation is published in ref. 4. It is a search for the $X \to SH \to bbbb$ process, where X and S are scalar resonances and H is the SM Higgs boson. This search identifies both S and H through their respective decays to a bottom quark–antiquark pair, $S \to b\bar{b}$ and $H \to b\bar{b}$. Because it is not possible to differentiate between jets originating from b and $\bar{b}$ in the ATLAS detector, in this context we make no distinction between the two and refer to the $b\bar{b}$ final state by the shorthand "bb." Figure 6.1 shows the Feyman diagram for the $X \to SH \to bbbb$ process, including the production of X via gluon–gluon fusion.

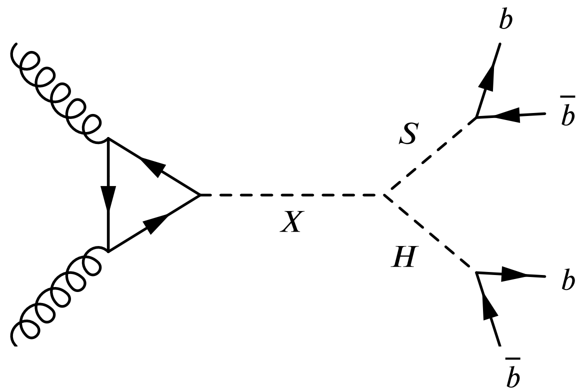

Figure 6.1: Feynman diagram of the $X \to SH \to bbbb$ process, initiated by gluon–gluon fusion [4].

A search for H in the bb final state has the advantage of a larger cross section than other decay modes, because H decays most frequently to $b\bar{b}$ with an expected branching ratio of about 58% [67]. The $S \to b\bar{b}$ branching fraction depends on both its mass and the theoretical model under consideration, but is always nonzero if S has the couplings expected of a Higgs boson.

This search does not target a particular theoretical model, but processes equivalent to $X \to SH$ are included in several noteworthy models with extended Higgs sectors. Examples include the SM extended by two real singlets (TRSM) [87] or a single complex singlet (CxSM) [84], the Next-to-Minimal Supersymmetric Standard Model (NMSSM) [85, 88], the complex Two-Higgs Doublet Model (2HDM) [89], and the 2HDM extended by a real or complex scalar singlet [86, 90, 91]. The results of this search may rule out certain configurations of these models.

The $X \to SH$ decay is kinematically feasible when $m_X > m_S + m_H$. This search targets m_X between 300 GeV and 6 TeV and m_S between 70 GeV and 5 TeV, but excludes m_S near 125 GeV. In ATLAS, the $m_S \approx m_H$ mass range is covered by dedicated searches for the $X \to HH$ process, which are optimized for the symmetry of that decay. Recent ATLAS searches for the $X \to HH$ process using the full Run 2 dataset include the $bbbb$ final state [92–94]. The CMS Collaboration has also published a dedicated $X \to HH$ search in this final state using their full Run 2 dataset [95].

This is the first ATLAS search for the $X \to SH \to 4b$ process, but searches for $X \to SH$ in various final states have already been published by the ATLAS and CMS experiments. ATLAS has published searches in the $VV\tau\tau$ [96], $bb\gamma\gamma$ [97, 98], $VV\gamma\gamma$ [99], and $HHH \to 6b$ [100] final states, and CMS has published searches in the $4b$ [101, 102], $bb\tau\tau$ [103], $bb\gamma\gamma$ [104, 105], $\gamma\gamma\tau\tau$ [106], and $VVbb$ [107] final states. The search presented here is competitive with the CMS $bbbb$ results.

This chapter presents the ATLAS $X \to SH \to bbbb$ search and its results. An overview of the search, highlighting its most important attributes and the overall analysis strategy, is given in section 6.1. This is followed by a description of the trigger strategy in section 6.2, then a description of the physics objects and event selection in section 6.3. The background estimation is explained in section 6.4. Section 6.5 describes the systematic uncertainties and how they are assessed, and a description of the statistical analysis follows in section 6.6. Finally, the results are presented in section 6.7, and the chapter is concluded in section 6.8.

6.1 Analysis strategy

Samples of MC simulated $X \to SH \to bbbb$ events are generated for each (m_S, m_X) pair shown in fig. 6.3. As explained previously, the range $100 < m_S < 170\,\text{GeV}$ is not included because it is covered by dedicated $HH \to 4b$ searches. For simplicity, only the gluon–gluon fusion production mode for X is simulated. The narrow-width approximation is used for both the X and S resonances, which means the widths of the reconstructed m_X and m_S distributions depend only on the detector resolution and the methods used to reconstruct them. The relative width of the reconstructed m_S distribution ranges from about 25% to 100%, while the relative width of the reconstructed m_X distribution ranges from 5% to 12%.

The kinematics of the $S \to bb$ and $H \to bb$ decays vary dramatically over the (m_S, m_X)

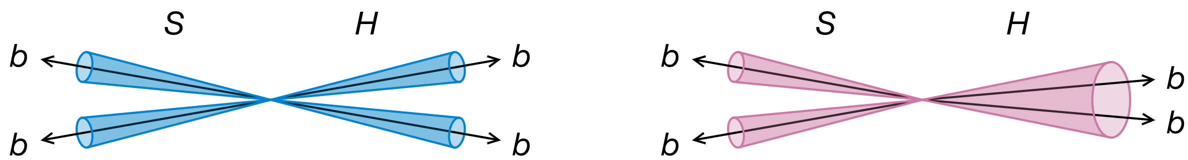

(a) Resolved channel (b) Mixed channel

Figure 6.2: Diagrams representing the methods used to reconstruct the $S \to bb$ and $H \to bb$ decays in each channel.

plane, and the preferred reconstruction method for a bb decay depends on its kinetic energy. A less energetic decay, with p_{T} below about 200 GeV, is best reconstructed as a pair of small-R or "resolved" jets, each identified as containing a single b quark. In a more energetic decay, the two b quarks are too collimated to distinguish (or "resolve") as individual small-R jets, so they are better reconstructed together as a single large-R or "boosted" jet.

This search includes two channels, "resolved" and "mixed," each corresponding to a different reconstruction method for the $H \to bb$ decay. In the resolved channel, both the $S \to bb$ and $H \to bb$ decays are reconstructed from a pair of small-R jets. In the mixed channel, the $S \to bb$ decay is reconstructed from a pair of small-R jets and the $H \to bb$ decay is reconstructed from a large-R jet. Both channels are represented as schematic diagrams in fig. 6.2. The resolved channel was previously explored by CMS in ref. 102. The mixed channel is unique to the present analysis. A "boosted" channel, in which both decays are reconstructed using boosted jets, was studied by CMS in ref. 101 and will also be studied in a future ATLAS search that will cover the the upper-left corner of the (m_S, m_X) plane.

Of the 83 simulated mass pairs, 14 are exclusive to the resolved channel, 51 are exclusive to the mixed channel, and 18 are accessible to both channels. Loosely speaking, the resolved channel is sensitive to mass pairs with small to intermediate m_X, while the mixed channel is sensitive to mass pairs with intermediate to large m_X. In the intermediate m_X range, the resolved channel excels when m_S/m_X is large. In addition, the mixed channel is insensitive to mass pairs with $m_S > m_H$ because the large-R jet is always identified with the H

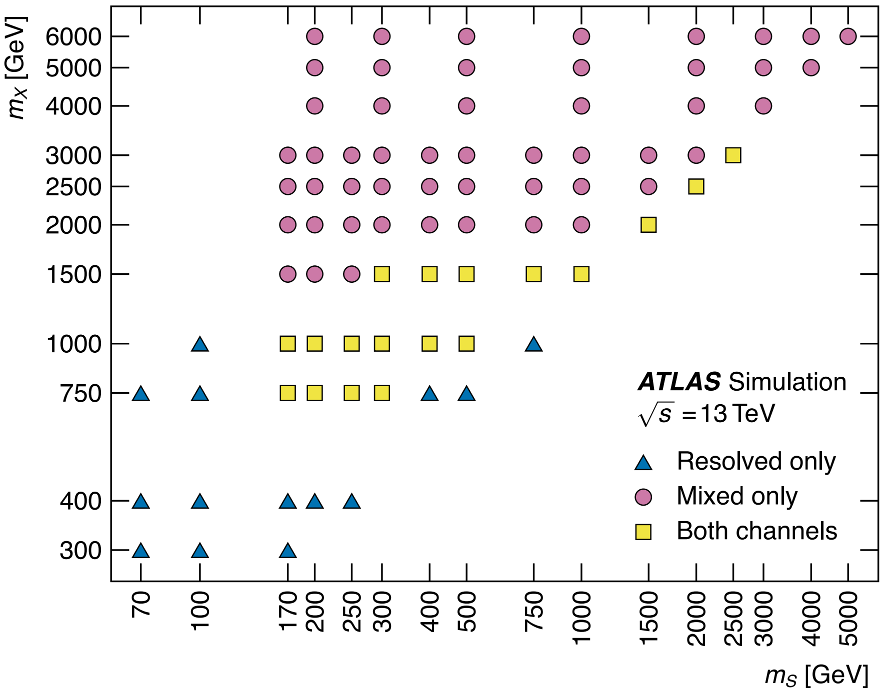

Figure 6.3: The sensitivity of the resolved and mixed channels to the signal mass hypotheses, shown as points in the (m_S, m_X) plane [4].

candidate, making it more boosted. Other restrictions on each channel's sensitivity come from kinematic requirements that are explained in section 6.3.

In both channels, an SR and a CR are defined based on the reconstructed H mass. The SRs target the H boson mass of 125 GeV, while the CRs are sidebands used to estimate background contributions. The search is performed by fitting the distribution of the reconstructed X mass and extracting a signal strength for each (m_S, m_X) hypothesis. To improve the signal-to-background ratio for each hypothesis, additional cuts on the reconstructed S mass are applied before performing the fit. The fits are performed independently for each channel, even for the mass hypotheses that are accessible to both. The event selection criteria for each channel are designed to optimize sensitivity to their dedicated regions of the (m_S, m_X) plane, so they are not orthogonal. For each mass hypothesis with shared sensitivity, the upper limit from the channel whose expected upper limit is lowest is chosen as the final result.

Table 6.1: Jet p_T thresholds and online b tagging efficiencies for 2b1j and 2b2j triggers.

		p_T threshold (GeV)		
Type	Year(s)	Untagged jet(s)	b tagged jets	b tagging efficiency (%)
2b1j	2016	100	55	60
	2017, 2018	150	55	70
2b2j	2016, 2018	35	35	60
	2017	35	35	40

6.2 Triggers

This search uses triggers of three types:

1. large-R jet triggers, which require one large-R jet,
2. 2b1j triggers, which require two b tagged small-R jets and one additional small-R jet, and
3. 2b2j triggers, which require two b tagged small-R jets and two additional small-R jets.

2b1j triggers are used in both channels, while large-R jet triggers are used only in the mixed channel and 2b2j triggers are used only in the resolved channel. The p_T threshold for all large-R jet triggers is 420 GeV. The p_T thresholds for the 2b1j and 2b1j triggers vary by data-taking year, as does the efficiency of the online tagger used to select the b jets. These details are summarized for all 2b1j and 2b2j triggers in table 6.1.

In general, when multiple triggers are used to select events from the same dataset, a dedicated calibration must be performed. To avoid this, and to simplify the trigger check in the event selection, we employ a "trigger bucketing" strategy. Events are divided among mutually exclusive "buckets" that are each assigned a trigger. Events that pass the trigger for their bucket are retained, and events that fail the trigger for their bucket are rejected.

In the resolved channel, events that contain at least one small-R jet with $p_T > 170$ GeV and at least two more with $p_T > 70$ GeV are assigned to the 2b1j trigger bucket, while events that do not meet these requirements are assigned to the 2b2j trigger bucket. These cuts are chosen to maximize the overall signal trigger efficiency.

In the mixed channel, events that contain at least one large-R jet with $p_{\mathrm{T}} > 520\,\mathrm{GeV}$ are assigned to the large-R jet trigger bucket, while events that do not meet this requirement are assigned to the 2b1j trigger bucket. This cut is chosen to ensure near-perfect efficiency of the large-R jet triggers, with the benefit that events in the large-R jet trigger bucket do not require scale factors.

6.2.1 Integrated luminosity for b jet triggers

Online b tagging relies on precise estimation of the beamspot position, which must be updated at the beginning of each run using new data [108]. Consequently, events recorded before the beamspot position was updated cannot reliably be used with b jet triggers. Therefore, these triggers require specialized GRLs, called "b jet HLT GRLs," that do not include these events and whose integrated luminosity is approximately 1.5% below that of the standard GRLs.

Additionally, after data collection began in 2016, a discrepancy was discovered between the online and offline algorithms used to locate the primary vertex. This discrepancy necessitated that events with a beamspot located further than 2 mm from the detector origin be removed from the affected portion of the 2016 b jet HLT GRL. As a result, the integrated luminosity of the 2016 b jet HLT GRL is approximately 25% below that of the standard GRL.

Because of their reliance on online b tagging, the integrated luminosity available to the 2b1j and 2b1j triggers is slightly reduced compared to that available to the large-R jet triggers. The integrated luminosities are compared in table 6.2. Accordingly, the total integrated luminosity for the resolved channel is $126\,\mathrm{fb}^{-1}$, while it ranges from $126\,\mathrm{fb}^{-1}$ to $136\,\mathrm{fb}^{-1}$ for the mixed channel. Signal events in the mixed channel analysis are scaled to the integrated luminosity appropriate for their trigger bucket. The 2015 data are omitted from

Table 6.2: Integrated luminosity of b jet HLT GRLs and standard GRLs for each year of Run 2 data collection.

	Integrated luminosity (fb^{-1})	
Year	b jet HLT GRL	Standard GRL
2016	24.6	33.0
2017	43.7	44.3
2018	57.7	58.5
Total	126.0	135.8

both channels because the b jet triggers for that year used different online b tagging algorithms that were not calibrated. The integrated luminosity of the 2015 dataset is approximately $3\,\text{fb}^{-1}$—so small that the effect of its exclusion on the search sensitivity is negligible.

6.3 Physics objects and event selection

Jets are the only physics objects used in this search, reconstructed as described in section 4.3.5. Small-R jets are used in both channels, while large-R jets are used only in the mixed channel. The GN2 and GN2X algorithms described in section 4.4.2 are used to identify jets originating from b quarks and from $H \to bb$ decays, respectively. Small-R jets with $p_{\text{T}} < 60\,\text{GeV}$ must pass the NNJVT jet vertexing algorithm mentioned in section 4.4.2. All object requirements for both types of jet are summarized in table 6.3.

Both channels define signal and control regions using the reconstructed H candidate mass, m_H^{reco}. The boundaries for these regions differ slightly between the two channels, and are listed in table 6.4. The nominal H mass of $125\,\text{GeV}$ is at or near the center of each SR, and each CR is composed of sidebands adjacent to the SR. The CRs are used for the background estimate described in section 6.4. The resolved channel SR and CR boundaries were chosen as the edges of a distorted ellipse in the $(m_H^{\text{reco}}, m_S^{\text{reco}})$ plane that is parameterized by the signal m_H^{reco} width, while the mixed channel boundaries were set using the signal m_H^{reco} width and rounded to the nearest $10\,\text{GeV}$.

Table 6.3: Object requirements for small-R and large-R jets. Except where noted explicitly, the small-R jet requirements are the same for both channels. The large-R jet requirements are for the mixed channel only.

	Requirement	
Attribute	Small-R jets	Large-R jets
Number	≥4 (resolved), ≥2 (mixed)	≥1 (mixed)
Minimum p_{T}	40 GeV (resolved), 55 GeV (mixed)	250 GeV
Maximum $\lvert\eta\rvert$	2.5	2.0
Mass range	—	70–200 GeV
Vertex tagging	NNJVT ($p_{\mathrm{T}} < 60$ GeV)	—
Flavor tagging	GN2 85% efficiency WP	GN2X 0.37% mistag rate WP

Table 6.4: Signal and control region boundaries for each channel.

	m_H^{reco} boundary (GeV)			
Channel	Lower CR	Lower SR	Upper SR	Upper CR
Resolved	59.3	107.8	148.8	190.7
Mixed	70	110	140	200

In both channels, events are selected by attempting to reconstruct S and H candidates from the relevant jet collections, subject to some additional criteria specific to each channel. Events that do not fall within the appropriate SR or CR are rejected. Section 6.3.1 describes the event selection for the resolved channel, and section 6.3.2 describes the event selection for the mixed channel.

6.3.1 Resolved channel event selection

After passing either a 2b1j or 2b2j trigger as described in section 6.2, resolved channel events must contain at least four small-R jets with $p_{\mathrm{T}} > 40$ GeV and $\lvert\eta\rvert < 2.5$. The p_{T} requirement is chosen to be slightly above the 2b2j trigger threshold. All jets that satisfy these requirements are sorted in order of decreasing p_{T}. The first four that pass the 85% efficiency working point of the GN2 tagger are selected to reconstruct the S and H candidates. If fewer than four pass this b tagging requirement, the event is rejected.

Once four b tagged jets have been selected, the S and H candidates are reconstructed using a boosted decision tree classifier, which is trained using kinematic observables constructed from each of the three possible pairings of the four jets in truth-matched events. During the selection, the trained classifier processes these observables for the event at hand and returns a probability for each of the three jet pairings. The pairing with the highest probability is chosen to reconstruct S and H, with the jet pair of larger angular separation $\Delta R(\mathrm{j},\mathrm{j}')$ identified as the more massive boson. This choice results in less mass sculpting than simply choosing the pair with larger invariant mass.

The event is rejected if the pseudorapidity separation between the reconstructed candidates does not satisfy $|\Delta\eta_{SH}^{\mathrm{reco}}| < 1.5$. This cut was chosen based on sensitivity studies performed for ATLAS $HH \to 4b$ searches, in which it reduces the background from multijet events. However, the $SH \to 4b$ process differs kinematically from the $HH \to 4b$ process, and while this cut improves the signal-to-background ratio for many (m_S, m_X) pairs, it is not optimal for all of them. Mass pairs for which the signal efficiency of this cut is 50% or below are excluded from the resolved channel analysis, and, as a result, it has limited sensitivity to the $1500 \leq m_X \leq 3000\,\mathrm{GeV}$ region of the (m_S, m_X) plane.

6.3.2 Mixed channel event selection

After passing either a large-R jet trigger or a 2b1j trigger as desribed in section 6.2, mixed channel events must contain at least two small-R jets with $p_{\mathrm{T}} > 55\,\mathrm{GeV}$ and $|\eta| < 2.5$, and at least one large-R jet with $p_{\mathrm{T}} > 250\,\mathrm{GeV}$ and $|\eta| < 2.0$. The small-R jet p_{T} requirement is chosen to meet the threshold of the b jet legs of the 2b1j triggers, while the large-R jet p_{T} requirement is set by the GN2X tagger. This requirement makes the mixed channel insensitive to mass pairs with $m_X - (m_H + m_S) < 250\,\mathrm{GeV}$.

The H candidate is reconstructed from the leading large-R jet in the event, provided it passes the 0.37% mistag rate working point of the GN2X tagger, which has an $H \to bb$

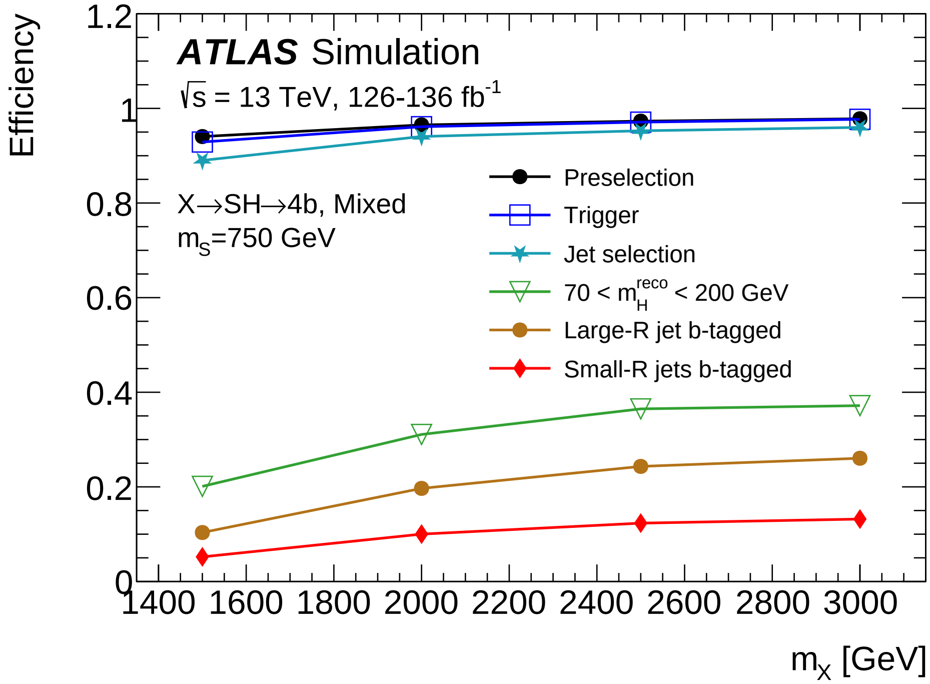

Figure 6.4: Mixed channel selection efficiencies for example signal samples with fixed $m_S = 750\,\text{GeV}$, as a function of the m_X hypothesis. The numbers are total efficiencies for consecutively applied selection steps. The "preselection" step requires one of the large-R jet, 2b1j, or 2b2j triggers to fire [4].

tagging efficiency of approximately 55%. If this jet does not pass the GN2X requirement, the event is rejected. After the H candidate has been identified, any small-R jets that overlap with it are removed prior to reconstructing the S candidate. This measure ensures that the S and H candidates are reconstructed from independent objects. Overlap in this case is defined as an angular separation $\Delta R(H_{\text{reco}}, \text{j}) < 1.0$ for any small-R jet j.

The S candidate is then reconstructed from the leading two small-R jets that remain, provided they both pass the 85% efficiency working point of the GN2 tagger. If they do not, or if fewer than two small-R jets remain after the overlap removal, the event is rejected. The S candidate must also satisfy the mass requirement $m_S^{\text{reco}} > 70\,\text{GeV}$, which reduces contributions from background processes. Figure 6.4 shows the efficiency of each stage of the mixed channel event selection for four mass hypotheses.

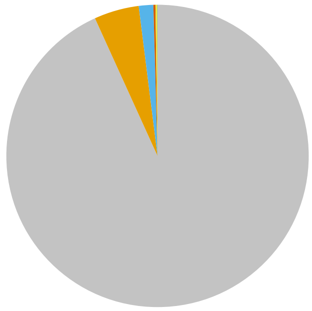

Process	Fraction
Multijet	93.21%
$t\bar{t}$	4.80%
$Z \rightarrow bb$	1.55%
$ZH \rightarrow qqbb$	0.22%
ggF $H \rightarrow bb$	0.20%
VBF $H \rightarrow bb$	0.02%

Figure 6.5: Background composition in the mixed channel signal region.

6.4 Background estimation

The abundance of jets produced at the LHC incurs large backgrounds in analyses with all-hadronic final states. In this search, the requirement that four b quarks are identified greatly reduces the multijet background compared to searches that require a similar number of jets but fewer b tags. Even with this reduction, multijet events are by far the largest background for this search, making up about 93% of expected events in the mixed channel.

Additional background contributions arise from other processes that produce b quarks. In particular, the $t\bar{t}$ process has a very large production cross section of about 800 pb [109] and always produces b quarks via the decays $t \rightarrow W^+b$ and $\bar{t} \rightarrow W^-\bar{b}$. In addition, the W branching fraction to hadrons is about 67% [9], which means either or both W bosons frequently decay to hadrons, creating even more jets that may pass the b-tagging requirements. The $t\bar{t}$ contribution to the mixed channel background is about 5%.

The smallest source of background is genuine $b\bar{b}$ pair production near the Higgs mass, which may arise from $Z \rightarrow b\bar{b}$ or $H \rightarrow b\bar{b}$ decays. These decays may occur separately alongside other jets, or together as the process $ZH \rightarrow qqbb$. Together, these processes comprise about 2% of the mixed channel background.

Figure 6.5 shows the relative contributions to the mixed channel signal region from all background processes.

6.4.1 GP + NF method

MC simulation of multijet events is notoriously challenging and imprecise in the phase-space regions of interest to most searches and measurements, so analyses with large multijet backgrounds usually estimate them using real data. These "data-driven" techniques typically use events from background-enriched CRs to approximate the background contribution to the SR. However, the kinematic and other properties of events in the CR differ from those in the SR by construction, so these estimates are often associated with large systematic uncertainties. In addition, the statistical precision of the background estimate is limited by the number of CR events.

This search uses a novel machine-learning approach to infer the kinematic properties of background events in the SR from interpolated CR data. The background estimate is performed in two steps using two machine learning techniques: Gaussian processes and normalizing flows. This approach, abbreviated "GP + NF," enables us to simulate the four-momenta of reconstructed S and H candidates in background events. The method provides a single background estimate that accounts for all of the processes listed in fig. 6.5.

The principle underlying the GP + NF method is that the kinematics of the S and H candidates in background events can be characterized by a joint probability distribution,

$$p(\vec{x}, m_H^{\text{reco}}) = p(m_H^{\text{reco}})\, p(\vec{x} \,|\, m_H^{\text{reco}}), \tag{6.1}$$

where m_H^{reco} is the H candidate mass, and the vector $\vec{x}$ describes all of the kinematic properties of the H and S candidates aside from m_H^{reco}:

$$\vec{x} = \left(p_{\text{T}}^{H_{\text{reco}}}, \eta_H^{\text{reco}}, p_{\text{T}}^{S_{\text{reco}}}, \eta_S^{\text{reco}}, m_S^{\text{reco}}, \Delta\phi_{SH}^{\text{reco}}\right). \tag{6.2}$$

It is useful to factorize $p(\vec{x}, m_H^{\text{reco}})$ this way because m_H^{reco} is used to define the SR and CR boundaries, so it requires special treatment.

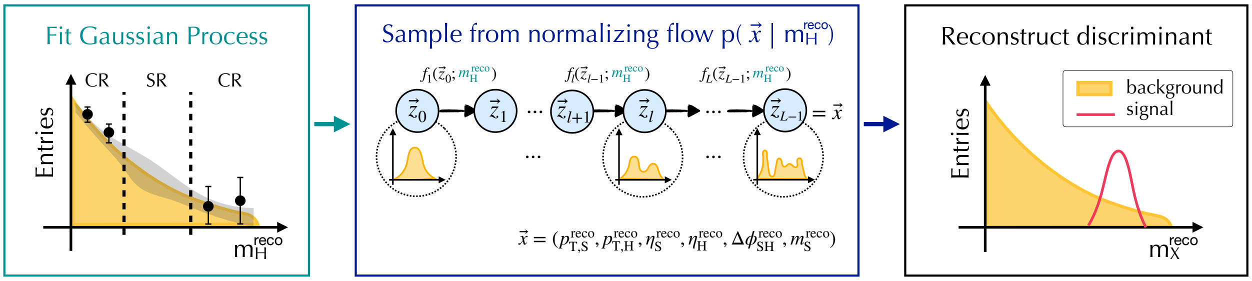

Figure 6.6: Overview of the GP + NF method for background estimation [4].

The two components of the GP+NF method correspond to the two functions on the right side of eq. (6.1). First, Gaussian process regression is used to interpolate the background m_H^{reco} distribution from the CR into the SR, which determines $p(m_H^{\text{reco}})$. Second, normalizing flows are trained to model the other kinematic quantities' dependence on m_H^{reco}, which determines $p(\vec{x}\,|\,m_H^{\text{reco}})$. Finally, the two components are combined by feeding the interpolated m_H^{reco} distribution into a trained NF model and generating events, which represent $p(\vec{x}, m_H^{\text{reco}})$. The GP + NF method is represented pictorially in fig. 6.6.

This section describes both steps of the GP + NF procedure and its validation. Section 6.4.2 describes Gaussian process regression and its use in the background estimation, while section 6.4.3 describes normalizing flows and their application. Validation regions used to test the method in the mixed channel are defined in section 6.4.4.

6.4.2 Gaussian process regression

Gaussian processes generalize Gaussian probability distributions to function space. Whereas a multivariate Gaussian probability distribution concerns random variables whose correlations are described by a covariance matrix, a Gaussian process concerns stochastic processes whose correlations are described by a covariance function [110]. Gaussian process (GP) regression, like any other method of linear regression, can be used to fit smooth a function to noisy data. Unlike other many other methods, however, GP regression is nonparametric, which means it does not require the user to explicitly specify the function's form. This is especially useful in the context of high-energy physics data, which are rarely well described by simple analytical functions. GP regression instead requires the specification of a covariance function or "kernel," which describes the correlations between data points.

Another difference between GP regression and more traditional fitting techniques is the form of the result, which is not a single function but a distribution over functions. The mean of this "posterior" distribution may be thought of as *the* result, while its variance may be used to assess an uncertainty. Both the mean and the variance may be determined by repeatedly sampling from the posterior distribution.

In this search, the *scikit-learn* library [111] is used to perform GP regression fits to m_H^{reco} distributions of CR data, represented as histograms. The fit results naturally interpolate m_H^{reco} through the SR and provide an estimate for the number of expected background events, calculated as the integral over the SR. For the covariance, we use the squared exponential function with an added noise term. Using the abbreviation $m \equiv m_H^{\mathrm{reco}}$, this is

$$\mathrm{cov}(m, m') = \exp\left(-\frac{m - m'}{2\ell^2}\right) + \sigma^2(m)\,\delta(m, m'), \tag{6.3}$$

where m and m' represent the centers of any two bins, ℓ is the "length scale," and $\sigma^2(m)$ is the statistical variance or "noise" in the bin with center m. The noise term informs the fit of the statistical precision of the input data. The only free parameter in eq. (6.3) is ℓ, which can be thought of as the minimum distance allowed between "wiggles" in the posterior functions. *Scikit-learn* optimizes ℓ during the fit.

The GP regression is performed separately for the mixed and resolved channels. Additionally, we perform a separate fit for each data-taking year to account for differences in trigger requirements and detector conditions. For each fit, we calculate the mean and variance from 25 samples, which we found to be sufficient for both to converge. The standard deviation in each bin is interpreted as an uncertainty. We also assess a 5% "normalization" uncertainty in the expected background yield, which is calculated as the integral of the mean fit result. The optimized values of ℓ are close to 30 GeV, which is the width of the SR. This reflects the expectation that the background m_H^{reco} distribution varies smoothly across the SR.

In fig. 6.7, the mixed channel GP fit results are compared to the CR and SR data. In each panel, "GP model" represents the mean of 25 samples from the posterior distribution, and "Uncertainty" represents the sum in quadrature of their standard deviation and the 5% uncertainty in the expected background yield. Since the fits were trained on CR data, the level of agreement in the CR illustrates how much the smooth GP regression result differs from the noisy training data. The SR data were not included in the training, so the SR agreement illustrates how well the actual data agree with the interpolation. Agreement is good for all three years, and we do not observe any significant shape or normalization differences between the SR data and the GP fit results. Figure 6.7d compares the data for all three years together to the sum of the fit results.

6.4.3 Normalizing flows

Normalizing flows (NF) are generative machine-learning models of multivariate probability distributions. Here, "generative" means a model can be used to sample from the probability distribution it has learned. Generative models are well suited to high-energy physics applications because sampled outcomes can be handled like MC simulated events, using familiar analysis tools and statistical techniques. In this sense, they provide similar utility to traditional data-driven procedures for background estimation using reweighted CR events. However, generative models are advantageous compared to traditional techniques in terms of statistical precision, because an arbitrary number of samples can be generated, meaning the statistical precision in the SR is unlimited. They can also be tailored to specific applications and phase space regions more readily than traditional MC simulation.

An NF model starts with a multivariate distribution of training data, called the "target

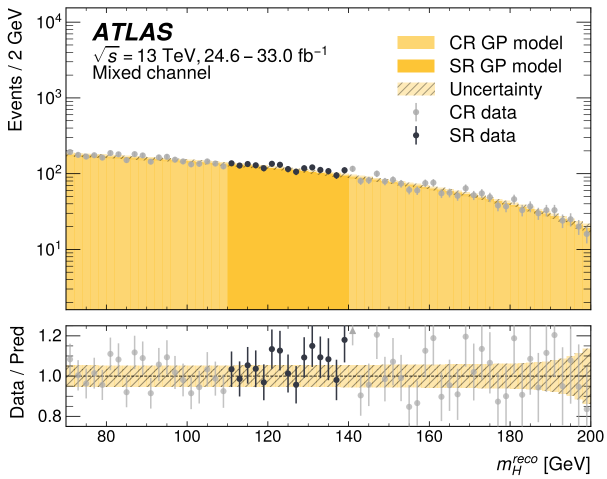

(a) 2016 data

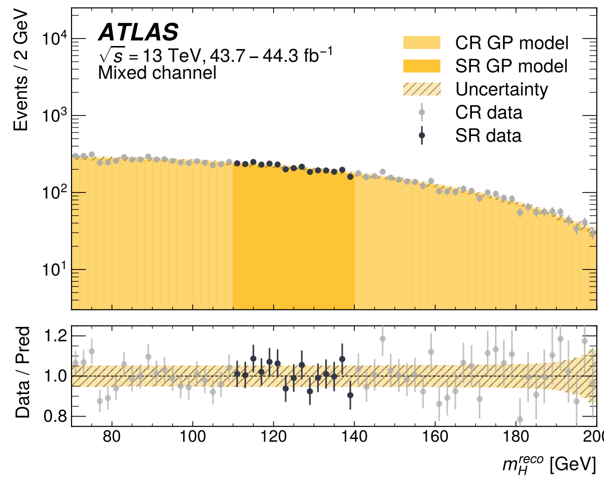

(b) 2017 data

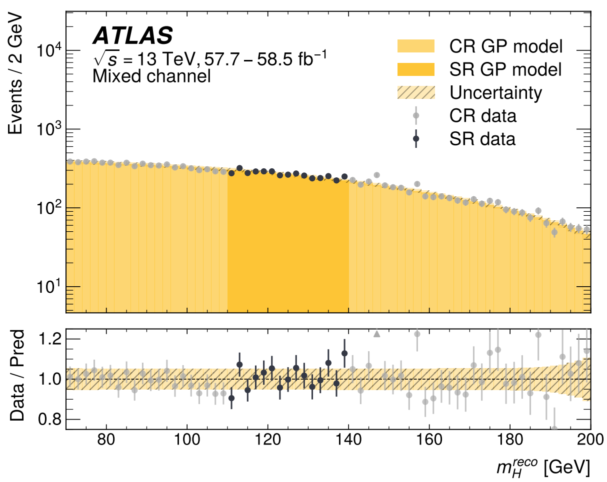

(c) 2018 data

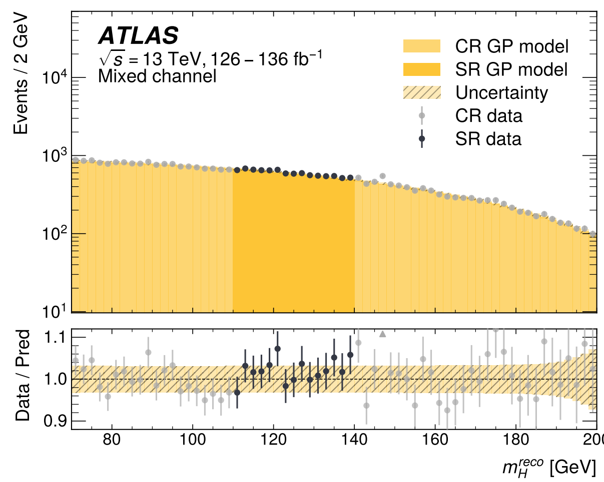

(d) Total

Figure 6.7: GP regression fit to the m_H^{reco} distribution for the mixed channel background estimate. The lighter colors indicate the CR included in the fit, while the darker colors indicate the SR not seen during the training [4].

distribution." The training process involves learning a series of invertible transformations that transform the target distribution into something very closely resembling a multivariate Gaussian distribution. The parameters of each transformation are learned using a neural network. In order for the the result to be a valid probability distribution, the transformations must all be differentiable as well as invertible [112]. Once the training is complete, the inverse transformations are applied to samples drawn from a multivariate Gaussian distribution, transforming them into samples from a very close approximation of the target distribution. The "flow" is the journey the samples take as they are transformed, and it is "normalizing" because the transformations applied during training transform the target distribution into a Gaussian.

For our background estimation, we use "rational quadratic spline coupling flows," which apply coupling transformations that are based on splines [113]. Coupling transformations split the inputs into two groups and use one group to parameterize a transformation for the other group. Transformations like this enable the NF to learn conditional probability distributions. A "spline" is a piecewise function whose pieces are all polynomials, and a "rational quadratic spline" has pieces that can each be expressed as a quotient of two quadratic polynominals. Splines are a good choice for modeling high-energy physics distributions because they are very general and can be used to describe complicated shapes. Rational quadratic splines can be made continuously differentiable, which is necessary for this application. They provide a good balance between flexibility and analytical tractability.

To learn the conditional probability $p(\vec{x}\,|\,m_H^{\mathrm{reco}})$ from eq. (6.1), we "condition" the NF on m_H^{reco}. In practice, we also condition on the data-taking year to account for the observed m_H^{reco} dependence on the data-taking year. This conditioning is achieved using the coupling transformations described previously. Our NF model has 10 total transformations and uses 4 "knots," or subdomain boundaries, for the splines. A neural network with 4 layers is used to learn the parameters of each spline transformation. Observables with sharp features are logarithmically transformed prior to training. Including these transformations, the observables learned directly in the training are

$$\ln\left(p_{\mathrm{T}}^{H_{\mathrm{reco}}}\right), \quad \eta_H^{\mathrm{reco}}, \quad \ln\left(p_{\mathrm{T}}^{S_{\mathrm{reco}}}\right), \quad \eta_S^{\mathrm{reco}}, \quad \ln\left(m_S^{\mathrm{reco}}\right), \quad \ln\left(\pi - \Delta\phi_{SH}^{\mathrm{reco}}\right). \tag{6.4}$$

After training, the inverse transformations are applied to recover the elements of $\vec{x}$. We train 25 NF models for each SR, each using a different initial seed. Independent models are trained for the resolved and mixed channels.

After the training, we sample from each model by feeding it one of the 25 samples from the GP posterior. The set of 25 coupled GP and NF predictions is called an "ensemble." The sampling is performed separately for each data-taking year, and for each year 100,000 "events" are sampled from each model across the CR and SR. After the sampling, events are discarded if they do not pass the kinematic cuts described in section 6.3, have m_X^{reco} outside the relevant range $200 < m_X^{\text{reco}} < 8000\,\text{GeV}$, or have azimuthal separation outside the physically meaningful range $0 < \Delta\phi_{SH}^{\text{reco}} < \pi$. Together, these restrictions remove less than 2% of sampled events in the mixed channel. The sampled events for each data-taking year are scaled to the estimated CR or SR yield, as appropriate, and summed. Each yield estimate is obtained by integrating the mean GP fit result over the relevant m_H^{reco} range as mentioned in section 6.4.2.

In fig. 6.8, three distributions learned by the GP+NF ensemble are compared to real data in the mixed channel SR. The models for all data-taking years are shown together, with the data superimposed. In each panel, "GP+NF model" represents the mean of 25 trainings and "1σ uncertainty" represents the sum in quadrature of their standard deviation and the 5% uncertainty in the expected background yield. The observable m_S^{reco} is transformed from the $\ln\left(m_S^{\text{reco}}\right)$ distribution learned by the NF models, while $\Delta\eta_{SH}^{\text{reco}}$ and m_X^{reco} are reconstructed from m_H^{reco} and the elements of $\vec{x}$. Agreement is good for the m_S^{reco} and $\Delta\eta_{SH}^{\text{reco}}$ distributions, with no overall trends. For the m_X^{reco} distribution, there is a trend: the model tends to underpredict at low m_X^{reco} and overpredict at large m_X^{reco}. This difference was anticipated based on GP + NF results from a validation region, described in section 6.4.4, and we assess a "nonclosure" uncertainty to account for it.

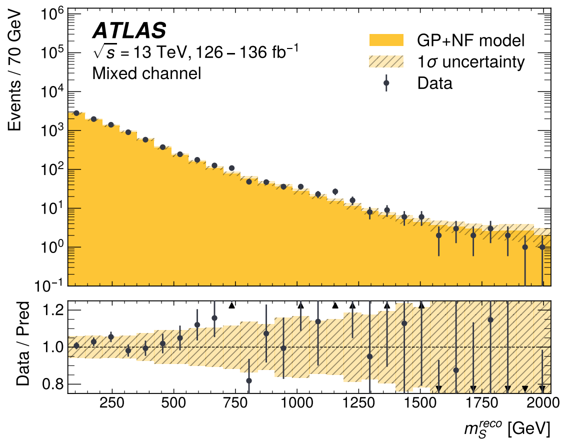

(a) m_S^{reco}

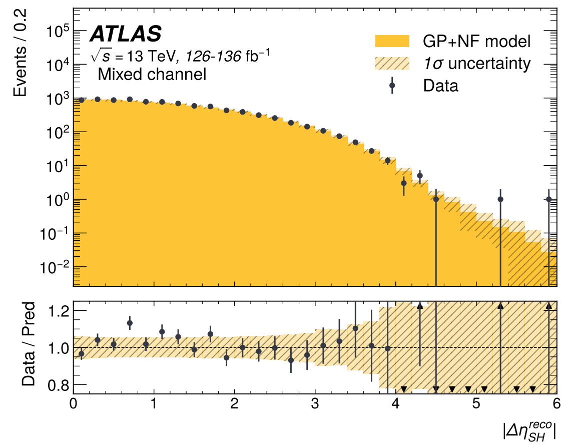

(b) $\Delta\eta_{SH}^{\text{reco}}$

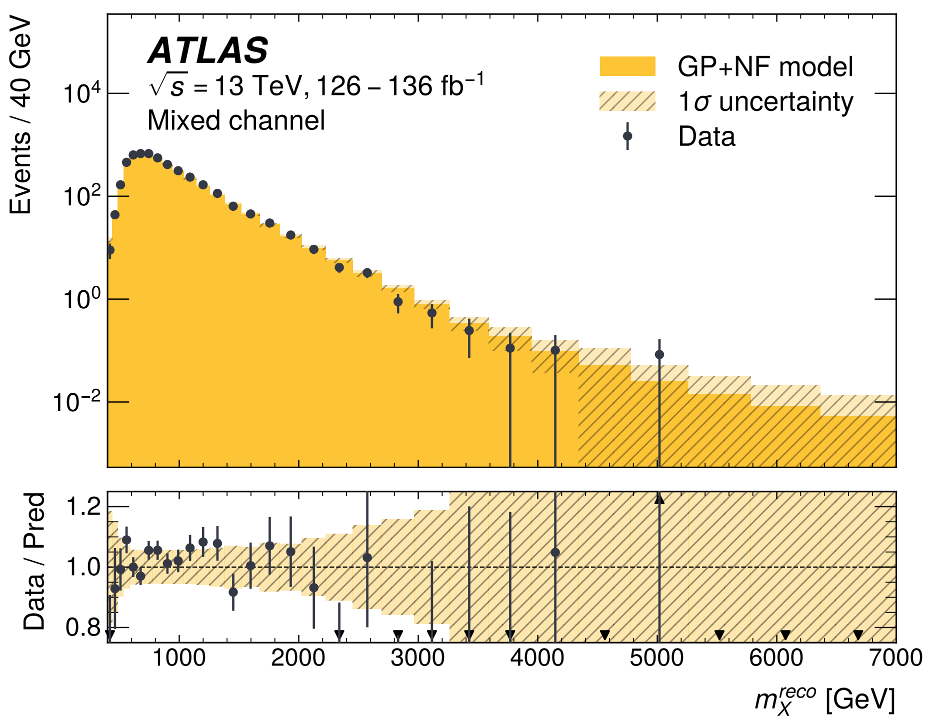

(c) m_X^{reco}

Figure 6.8: Selected kinematic distributions predicted by the GP+NF method for the mixed channel. The observable m_S^{reco} is directly predicted by the NF as $\ln\left(m_S^{\text{reco}}\right)$, while $\Delta\eta_{SH}^{\text{reco}}$ and m_X^{reco} are reconstructed from the GP + NF ensemble predictions [4].

6.4.4 Background method validation

To assess the performance of the GP + NF method, we used data from validation regions (VRs). These regions are designed to be kinematically similar to the SR but have a much lower signal-to-background ratio. The complete GP + NF method is performed in each VR using the same procedure as described above. The m_H^{reco} boundaries listed in table 6.4 are also used for these VR estimates; just like in the actual background estimate, the models are trained using "CR" data, and their performance is assessed using "SR" data.

For the mixed channel, two VRs are defined, which are used for different purposes. The event selection for each VR is the same as described in section 6.3.2, but with one requirement changed.

1b VR

As in section 6.3.2, the S candidate is reconstructed from the two leading small-R jets that remain after the overlap removal with the H candidate. However, instead of requiring that both of these jets pass the 85% efficiency working point of the GN2 tagger, we require that only one passes and the other fails. This VR has about ten times as many events as the SR and is used to derive the nonclosure uncertainty. Figure 6.9 shows the m_X^{reco} distribution from the GP + NF model compared to real data in this VR. The model tends to underpredict the number of events at low m_X^{reco} and overpredict in the tail. To account for this difference, we use the m_X^{reco} distribution in this VR to assess a nonclosure uncertainty as described in section 6.5.2.

Loose GN2X VR

As in section 6.3.2, the H candidate is reconstructed from the leading large-R jet in the event. However, instead of requiring that it pass the 0.37% mistag rate working point of the GN2X tagger, it must fail this working point and instead pass a looser working point with

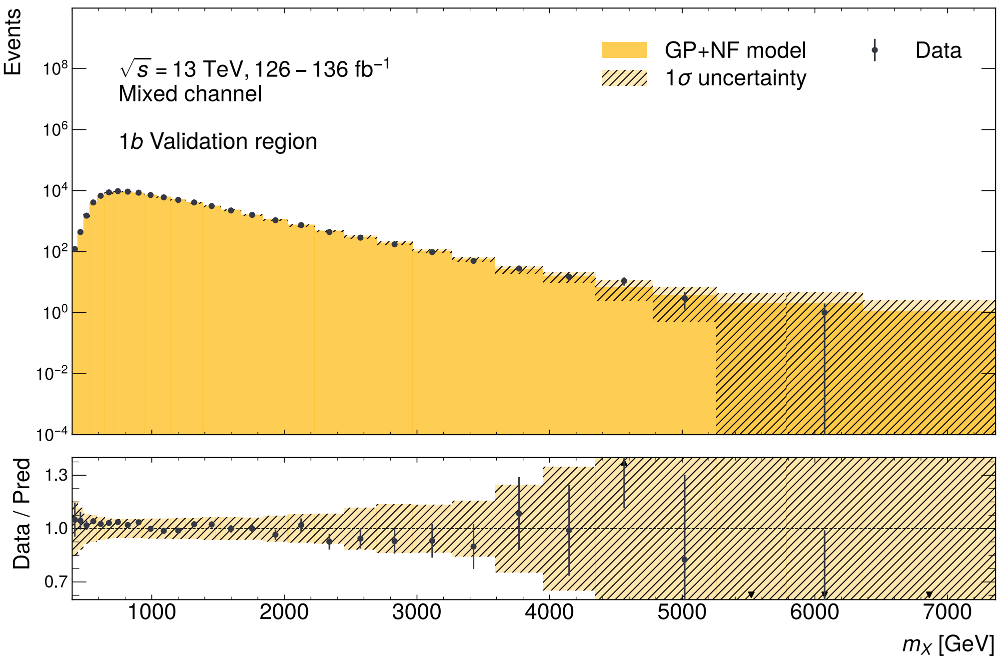

Figure 6.9: Distribution of m_X^{reco} in the mixed channel 1b VR.

a 1.55% mistag rate. This VR has about three times as many events as the SR and is used to test the statistical analysis as described in section 6.6.

6.5 Systematic uncertainties

Because the background estimation in this search is entirely data driven, theoretical uncertainties and uncertainties associated with the detector response are applied only to the signal. For the background, dedicated uncertainties developed for the GP + NF method are applied. This section describes the sources of systematic uncertainty for both signal and background, in section 6.5.1 and in section 6.5.2, respectively.

6.5.1 Signal systematics

Uncertainties associated with jet reconstruction and b tagging are assessed for both small-R and large-R jets. The uncertainties associated with the GN2X tagger used for large-R jets in the mixed channel range from 31–38%; they are by far the largest source of uncertainty affecting the signal in this search. Uncertainties from other aspects of jet reconstruction, including the GN2 tagger efficiency, are no more than a few percent each. They include uncertainties in the jet energy and mass scale and resolution, which account for differences between the fast simulation used to produce the signal samples and the full simulation used to produce the samples from which these uncertainties are derived.

Uncertainties in the trigger efficiency affect the signal as well. No uncertainty is assessed for the large-R jet triggers because they are used in the p_T region with perfect efficiency, as explained in section 6.2. For the b jet triggers, only preliminary calibrations are available for software configuration used to generate the signal samples. A 20% uncertainty is assessed in the 2b1j and 2b2j trigger efficiencies based on the available calibration data. Its effect on the search sensitivity is negligible for nearly all mixed channel mass hypotheses.

The remaining uncertainties affecting the signal come from theory, pileup modeling, the luminosity measurement, and the statistical precision of the MC samples. All sources of theoretical uncertainty, including the choice of factorization and renormalization scales, the choice of PDF, and the modeling of the parton shower and underlying event, are assessed together as a 5% normalization uncertainty based on the methodology in ref. 93. The uncertainty in the luminosity measured for the combined 2016–2018 dataset is 0.84%. The uncertainty in the statistical precision of the signal MC samples is very small compared to other sources.

6.5.2 Background systematics

In both the mixed and resolved channels, the mean of the 25-model GP + NF ensemble is used as the nominal background shape, computed as

$$H_\mathrm{nom} = \sum_{i=0}^{N-1} \frac{H_i}{N}, \tag{6.5}$$

where H_{nom} is the nominal background histogram, H_i is the histogram for model i, and $N = 25$ is the total number of models. All H_i have the same integral, which is computed from the GP regression as described in section 6.4.2.

We use the 25 histograms to assess a systematic uncertainty for the shape of the background m_X^{reco} distribution. Their spread characterizes the uncertainty associated with the statistical precision of the CR data used to train the models, as well as the models' ability to reliably reproduce the m_X^{reco} distribution from a random initial seed. In the mixed channel, this uncertainty is assessed as 25 sets of correlated up and down variations, $H_i^{\pm}$, which are each computed as

$$H_i^{\pm} = H_{\text{nom}} \pm \frac{H_{\text{nom}} - H_i}{\sqrt{N}}. \tag{6.6}$$

If any bin of $H_i^{\pm}$ has a negative value, it is set to 0. In this way, each of the 25 models is represented by one nuisance parameter.

In both channels, a nonclosure uncertainty is assessed using a VR. In the mixed channel, this is the 1b VR described in section 6.4.4. The up and down nonclosure variations, $H_{\text{noncl}}^{\pm}$, are calculated as

$$H_{\text{noncl}}^{\pm} = H_{\text{nom}} \pm \frac{\sum H_{\text{nom}}}{\sum H_{\text{VR}}^{\text{exp}}}(H_{\text{VR}}^{\text{obs}} - H_{\text{VR}}^{\text{exp}}), \tag{6.7}$$

where $H_{\text{VR}}^{\text{exp}}$ is the histogram of the expected m_X^{reco} distribution in the VR, $H_{\text{VR}}^{\text{obs}}$ is the observed distribution, and $\sum$ represents a histogram's integral.

Finally, as described previously, we also assess an uncertainty on the normalization of the background model. This uncertainty is ±5%, which is chosen to cover normalization differences observed in the VRs.

6.6 Statistical analysis and validation

As mentioned in section 6.1, a signal strength is extracted for each (m_S, m_X) hypothesis via a fit to the m_X^{reco} distribution. The fits are performed separately for the resolved and mixed

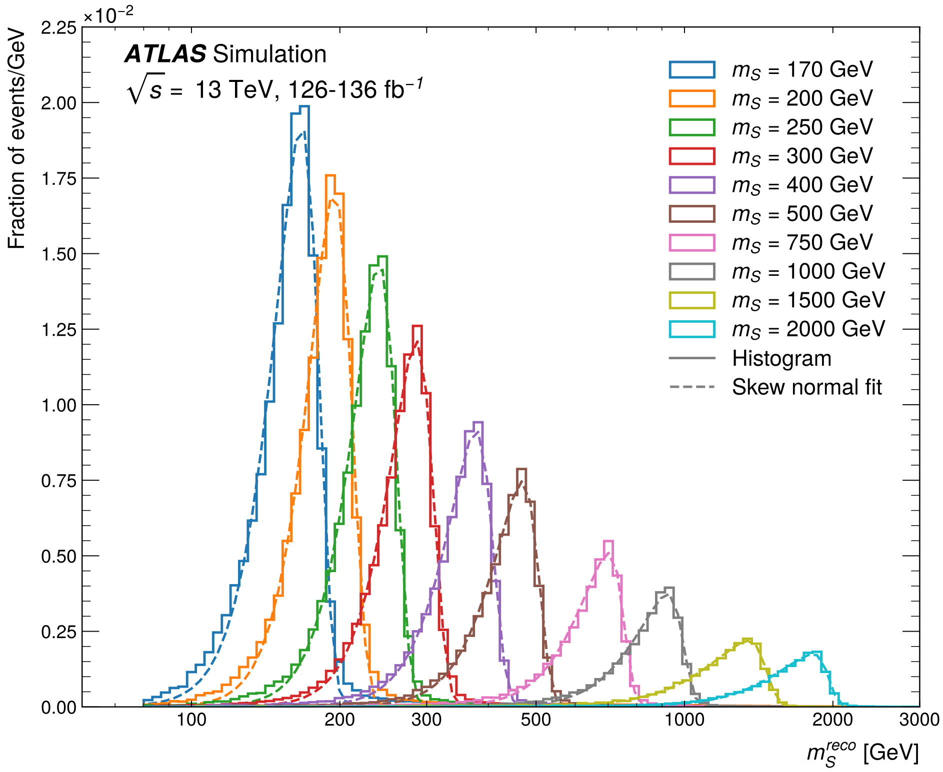

Figure 6.10: Signal m_S^{reco} distributions for each m_S hypothesis assigned an exclusive region in the mixed channel, with fitted skew normal distributions overlaid [4].

channels.

Before the fit, the signal-to-background ratio for each mass pair is improved by applying cuts that depends on m_S. In the resolved channel, this takes the form of a cut on a two-dimensional distance variable in the $(m_S^{\text{reco}}, m_X^{\text{reco}})$ plane. In the mixed channel, a one-dimensional m_S^{reco} window is defined based on the 99% probability interval of a skew normal distribution fit to the signal m_S^{reco} distribution. This cut is not applied for mixed channel mass hypotheses with $m_S \geq 2.5\,\text{TeV}$ or $m_X \geq 4\,\text{TeV}$ because, in these cases, it does not improve the sensitivity. The relevant m_S^{reco} distributions and the fitted skew normal distributions are shown in fig. 6.10.

After these m_S-dependent cuts are applied, the signal and background m_X^{reco} distributions are histogrammed using bins whose width increases logarithmically with m_X^{reco} to adequately capture the tail of the distribution. Any bins with fewer than 10 expected background events are merged with the bin to their left. This procedure ensures that all bins have sufficient statistical precision for the fit to converge, even in the event that one or more of them have substantially fewer observed than expected events.

Systematic uncertainties are assessed using the histograms after the m_S-dependent cuts and rebinning are applied. As explained in section 6.5.2, a nonclosure uncertainty is derived from the difference between the data and GP + NF model m_X^{reco} distributions in a VR. Figure 6.11 shows this comparison in the mixed channel $1b$ VR for the $m_S = 750\,\text{GeV}$ region and the inclusive m_S region used for all mass pairs with $m_S \geq 2.5\,\text{TeV}$ or $m_X \geq 4\,\text{TeV}$. Figure 6.12 shows the relative nonclosure uncertainty in the m_X^{reco} distribution in both regions, as well as the relative shape uncertainty from each of the 25 GP + NF models.

The Loose GN2X VR is used to validate the fit method in the mixed channel, with the number of events scaled down to match the expected SR yield. Figure 6.13 shows an example post-fit m_X^{reco} distribution for the $m_S = 750\,\text{GeV}$ region. The good agreement indicates that the background estimation method performs well in this VR and should, likewise, perform well in the SR. Signal injection tests were also performed in this VR, using the signal templates from the SR, and the injected signal was consistently recovered.

6.7 Results

As explained in section 6.1, separate fits are performed for the resolved and mixed channels. For each of the 18 mass hypotheses with shared sensitivity, shown in fig. 6.3, the upper limit from the channel with the best expected sensitivity is selected as the final result.

The post-fit m_X^{reco} distributions for two representative m_S regions in the mixed channel are shown in fig. 6.14. Agreement between data and the post-fit background-only prediction is good in both cases, and no significant deviations from the background-only hypothesis are observed. Example signal hypotheses are overlaid and scaled to their expected upper limit. The imperfect post-fit agreement and limited constraining of systematics suggest that the fit

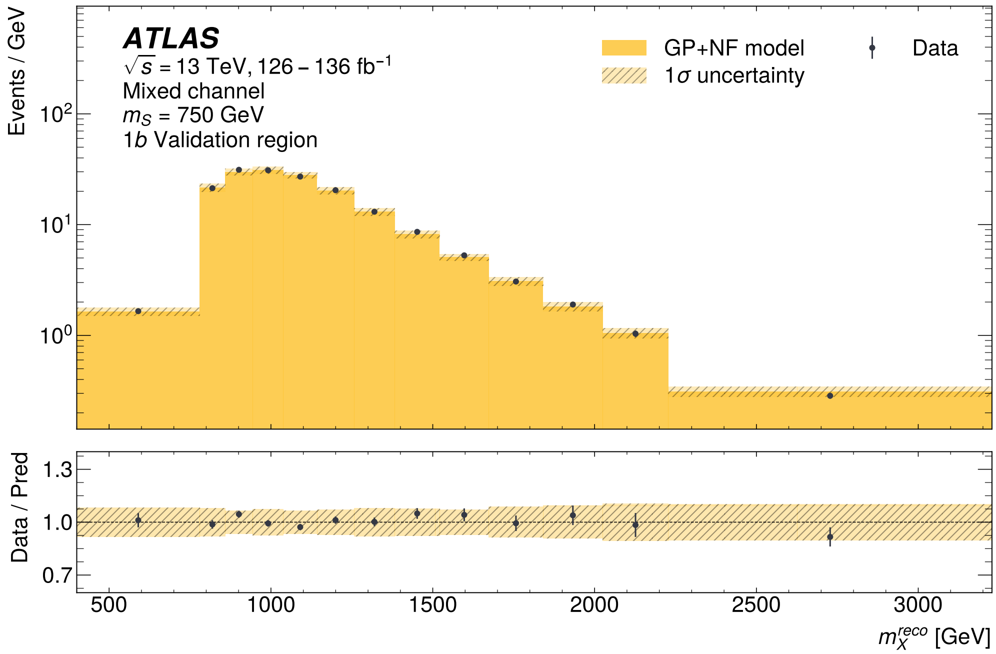

(a) $m_S = 750\,\text{GeV}$

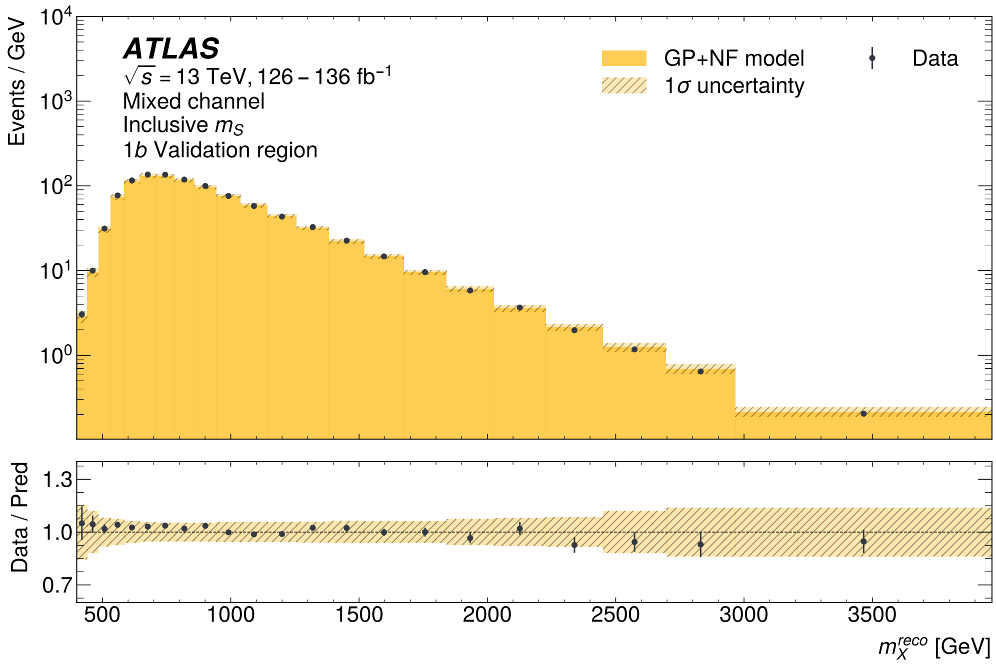

(b) Inclusive m_S

Figure 6.11: Nonclosure of the GP + NF background prediction in the mixed channel $1b$ VR for two m_S regions. Plots are pre-fit and the ratio is the nonclosure uncertainty assigned to the SR [4].

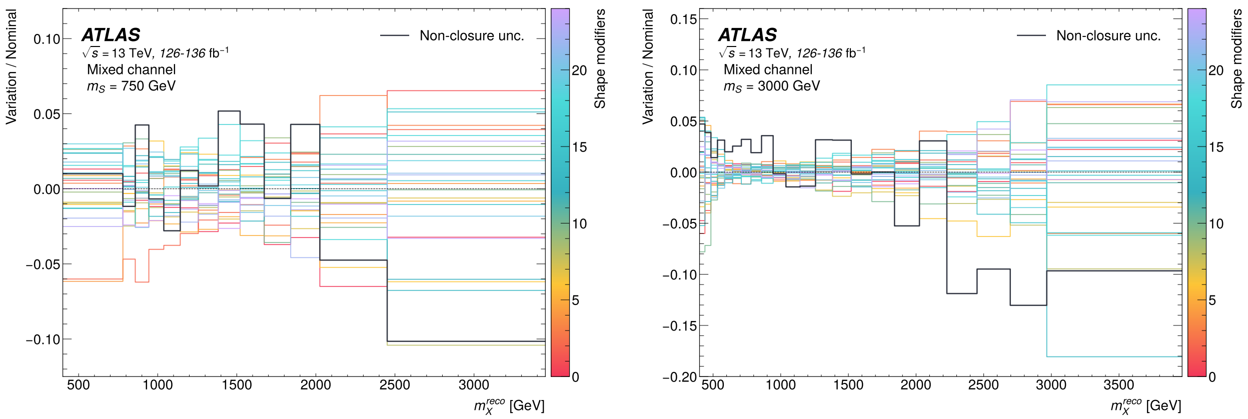

(a) $m_S = 750\,\mathrm{GeV}$ (b) Inclusive m_S

Figure 6.12: Relative variations of background systematic uncertainties with respect to the nominal background prediction for two m_S regions in the mixed channel. The inclusive m_S region is used to derive the nonclosure uncertainty for the $m_S = 3\,\mathrm{TeV}$ fits. Both panels show the non-closure uncertainty together with the correlated shape uncertainties derived from the ensemble of background predictions [4].

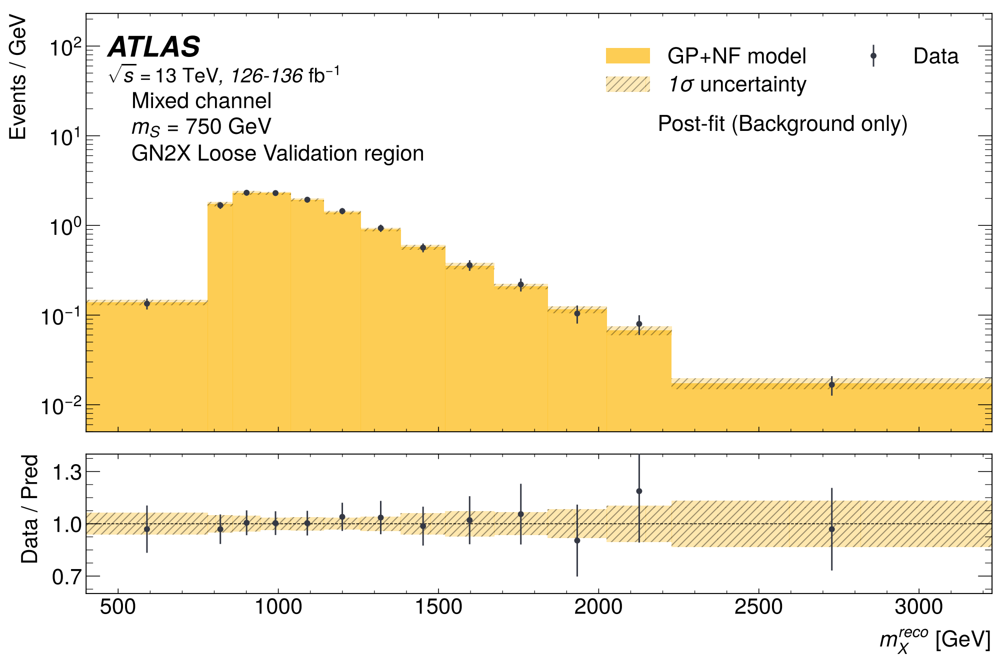

Figure 6.13: Example post-fit m_X^{reco} distribution in the mixed channel Loose GN2X validation region under the background-only hypothesis. The lower panel shows the ratio of the observed data to the background prediction [4].

is not overly flexible, which means it has good sensitivity; an overly flexible fit might absorb a potential signal.

Expected and observed upper limits at 95% CL on the product of the X production cross section and the $X \to SH \to 4b$ branching ratio are assessed for all mass hypotheses. The spacing in m_S of the tested mass hypotheses is smaller than the m_S^{reco} resolution, so a linear interpolation is performed between mass hypotheses that are adjacent in m_S, except across the gap containing $m_S = 125\,\text{GeV}$ that is left to dedicated $HH \to 4b$ searches. However, the spacing in m_X is larger than the m_X^{reco} resolution, so no interpolation is performed in the m_X direction.

Figure 6.15 shows the upper limits for the resolved and mixed channels separately, with each tested m_X hypothesis scaled to a unique factor of 10. Of the 18 mass hypotheses with shared sensitivity, the mixed channel expected upper limits are best for two of them: $(m_S, m_X) = (300\,\text{GeV}, 1.5\,\text{TeV})$ and $(400\,\text{GeV}, 1.5\,\text{TeV})$. The resolved channel limits are best for the remaining 16 shared hypotheses. The channel assignment for the final limits is shown for each mass hypothesis in fig. 6.16. Figure 6.17 shows the upper limits from only the most sensitive channel, with no scaling applied. Predictions from the TRSM, using the parameters presented in ref. 4 to maximize the cross section at each mass hypothesis, are also shown in both figures. None of the predictions are excluded, meaning the search is not yet sensitive to the TRSM; however, the observed limit at $(m_S, m_X) = (100\,\text{GeV}, 400\,\text{GeV})$ is very close to the prediction.

The relative impacts on the upper limits from all sources of uncertainty for both channels are listed in table 6.5. The impact of the data statistical uncertainty is evaluated as the expected upper limit obtained with all systematic uncertainties excluded, while the impact of each group of systematics is evaluated as the absolute difference between the expected upper limit obtained with all systematic uncertainties included and with that group of uncertainties excluded. All impacts are assessed relative to the expected upper limit, with all systematic

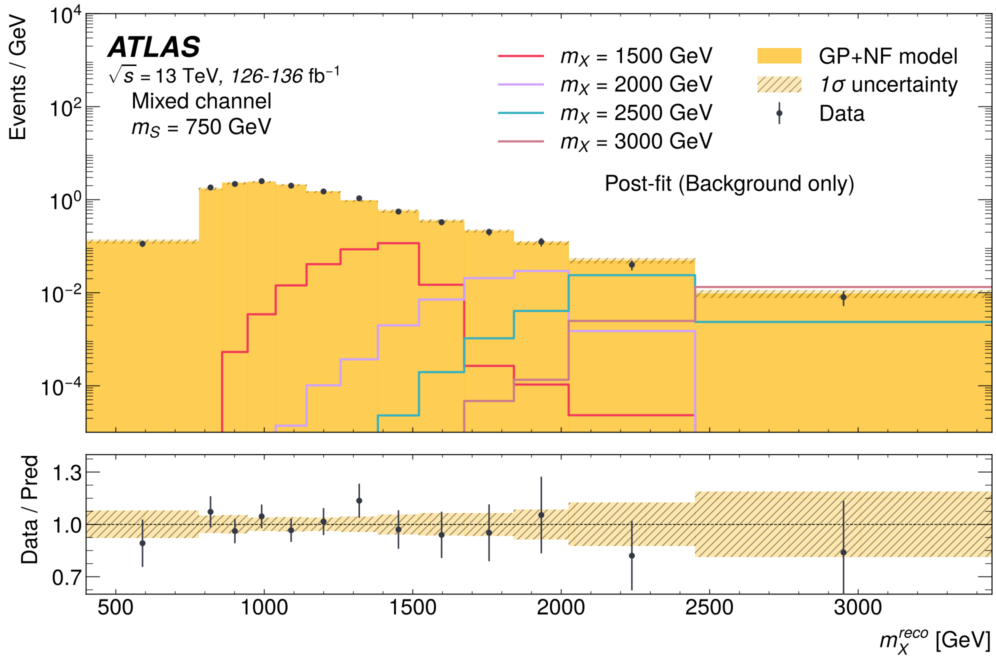

(a) $m_S = 750\,\text{GeV}$

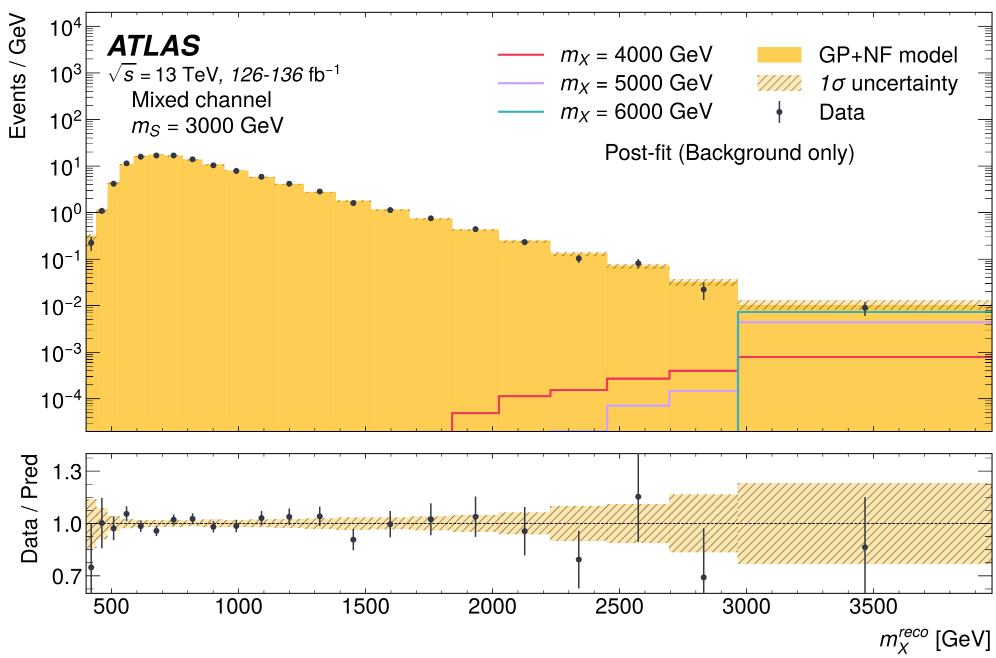

(b) $m_S = 3\,\text{TeV}$

Figure 6.14: Example post-fit m_X^{reco} distributions in two mixed channel m_S regions under the background-only hypothesis. Example signal templates are overlaid and scaled by the corresponding expected upper limit. The lower panels show the ratio of the observed data to the background prediction [4].

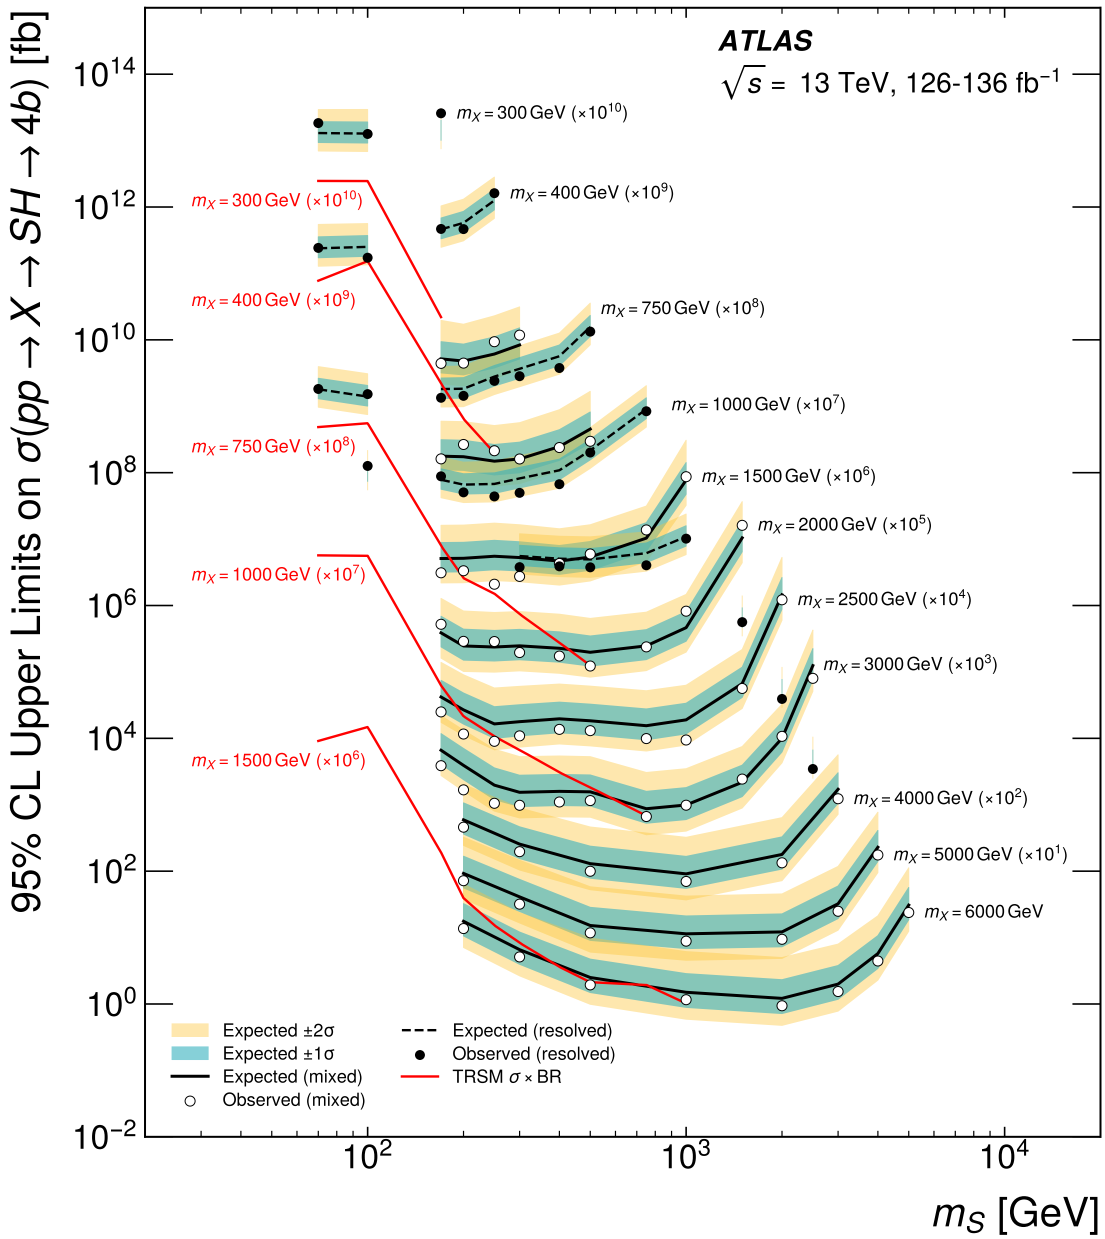

Figure 6.15: Expected and observed 95% CL upper limits on the product of the X production cross section and the $X \to SH \to 4b$ branching ratio as a function of m_S for each m_X hypothesis. Limits from the resolved and mixed channels are shown separately. For visual clarity, limits corresponding to different m_X hypotheses are multiplied by the factor indicated in the figure. Predictions from the TRSM are overlaid for selected m_X values [4].

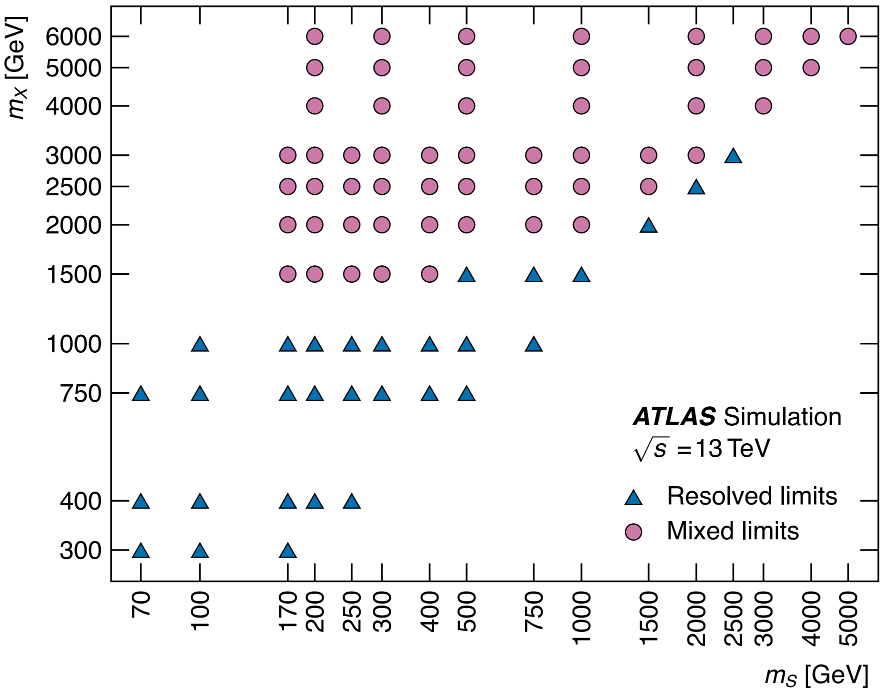

Figure 6.16: Assignment of the resolved and mixed channels to the final results for all signal mass hypotheses, shown as points in the (m_S, m_X) plane [4].

uncertainties included.

The mixed channel uncertainties are dominated by the data statistical uncertainty. The largest sources of systematic uncertainty in the mixed channel are associated with the background modeling and the GN2X tagger. The "flavor tagging" row also includes contributions from the GN2 tagger, but these contributions are only a few percent. The background systematic uncertainties are split into two categories. The first, "network retraining and CR stats," represents the shape uncertainties described in section 6.5, which are computed by retraining the NF models but are influenced by the limited statistical precision of the CR. The impacts of this category are much larger in the resolved channel than in the mixed channel because the two channels assess it using different methods; the method used in the resolved channel analysis involves uncorrelated per-bin modifiers that incur large impacts.

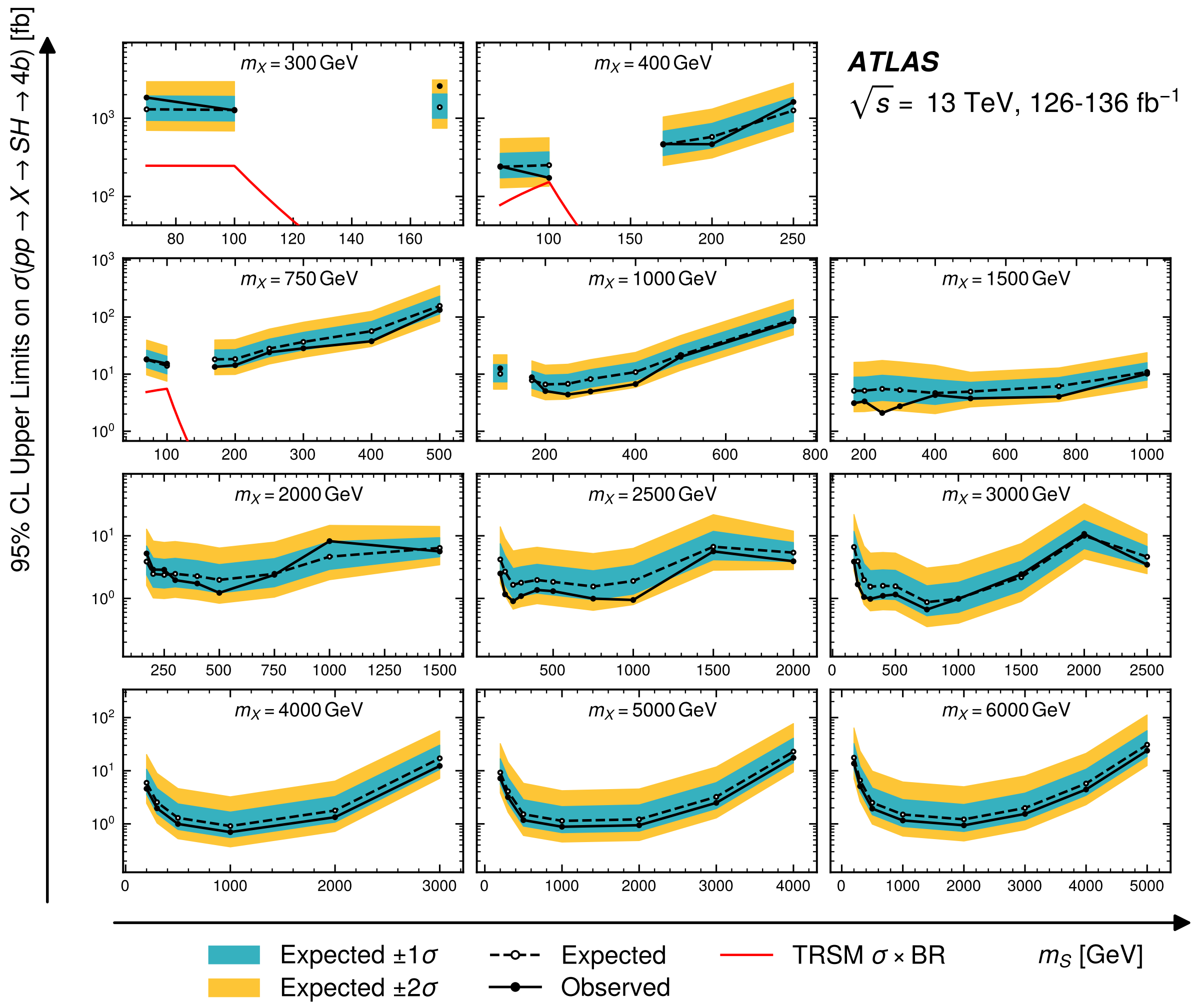

Figure 6.17: Expected and observed 95% CL upper limits on the product of the X production cross section and the $X \to SH \to 4b$ branching ratio as a function of m_S for each m_X hypothesis. For mass hypotheses covered by both the resolved and mixed channels, only the result from the channel with the lower expected upper limit is retained. Predictions from the TRSM are overlaid for selected m_X values [4].

6.8 Conclusion

This search for the $X \to SH \to 4b$ process uses the full Run 2 dataset collected by the ATLAS detector from 2016 to 2018. The search targets a wide range of masses for the hypothetical X and S resonances, ranging from $300 \leq m_X \leq 6000$ GeV and $70 \leq m_S \leq 5000$ GeV. It includes two channels, resolved and mixed, that target complementary regions of the (m_S, m_X) plane. The novel GP + NF method is used for background estimation in both

Table 6.5: Summary of the impact of the different sources of systematic uncertainty on the search sensitivity.

Uncertainty source	Relative impact (%) Resolved channel	Relative impact (%) Mixed channel
Data statistical uncertainty	24–51	55–74
All systematic uncertainties	49–76	26–45
Experimental	48–76	26–45
Detector response	9–15	24–33
Luminosity and pileup	<1	<1
Flavor tagging	0–5.4	22–32
Jet reconstruction	0.1–2.7	0.4–2.5
Trigger efficiency	8–9	<1
Background modeling	44–73	1–26
Network retraining and CR stats	36–63	0.1–23
Nonclosure and normalization	0–25	0.4–7.2
Theoretical	<1	<1
MC statistical	<1	<1

channels.

No significant excess is observed for any (m_S, m_X) mass hypothesis. The largest excess has a local significance of 1.8σ and is observed at the hypothesis $(170\,\text{GeV}, 300\,\text{GeV})$. Upper limits at 95% CL are set on the product of the X production cross section and the $X \to SH \to bbbb$ branching ratio. The observed limits range from 0.7 fb to 2.6 pb, with the lowest limits set by the mixed channel, which is more sensitive to mass hypotheses with large m_X. These results are the first ATLAS limits for the $X \to SH \to bbbb$ process and include the broadest range of mass hypotheses probed to date.

CHAPTER 7

ATLAS TRIGGER SYSTEM UPGRADE FOR THE HL-LHC

The HL-LHC is designed to attain an instantaneous luminosity as high as $7.5\times10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$, which is nearly four times the peak instantaneous luminosity reached during Run 2. With this increase in luminosity will come a corresponding increase in pileup. The Run 2 LHC configuration produced at most 38 pileup interactions per bunch crossing on average [23]; for the luminosity of HL-LHC, this number is expected to approach 200 [114].

The ATLAS trigger system must be enhanced to cope with this increase in pileup interactions. Specifically, the system's access to and utilization of tracking information must be improved. By quickly reconstructing particle tracks from hits in the Inner Tracker (ITk), the successor to the ID, the trigger system will be better able to associate decay products with the correct interaction vertices and make better decisions of which events to store. Typically, the reconstruction of particle tracks is a computationally expensive and rather lengthy process, so the Run 2 trigger system made only limited use of partial tracking information, even at the HLT stage. Very fast track reconstruction will be necessary to provide the future ATLAS trigger system with tracking information from the entire detector.

This section describes two projects related to the ATLAS trigger system upgrade for the HL-LHC. The first, described in section 7.1, is electronics testing for a proposed hardware tracking system. The second, described in section 7.2, is the implementation of tracking firmware for a hardware-accelerated Event Filter system roughly analogous to the HLT.

7.1 Testing of high-speed track-fitting hardware

The Hardware-based Tracking for the Trigger (HTT) system was proposed as a potential means of fast track reconstruction. This system would allow the trigger system to perform fast track reconstruction using hardware, bypassing software latency [114].

The Track Fitter Mezzanine (TFM) card, one of the many components of the HTT system, is a mezzanine card designed to perform second-stage track fitting. The TFM receives preliminary fits to the first 8 layers of the ITk and hits from the remaining 5, and uses these data to perform a fit to all 13 layers. In the HTT design, two TFMs are mounted on a Second-Stage Tracking Processor (SSTP) board, which facilitates communication between the TFMs and other HTT subsystems. The complete HTT design has 196 TFMs mounted on 96 SSTP boards. The principal component of the TFM is an Intel Stratix 10 FPGA, which performs the track fitting and controls data input and output. Twelve serial links send data between the TFM and the SSTP. It is critical to the performance of the HTT, and of the trigger system as a whole, that communication between the TFM and SSTP be fast and accurate.

Loopback tests were used to test the performance of the TFM serial link transceivers. A loopback test is performed by directly linking the transmitter and receiver of a single transceiver. This configuration does not introduce external factors and so is an effective method of testing the functionality of the TFM hardware. The signal leaves the Stratix 10, travels through the hardware, and returns. Therefore, the loopback test can be used to gauge the hardware performance.

This section describes the methodology and results of these tests. Section 7.1.1 describes the equipment and lab setup. Section 7.1.2 describes the procedure used to tune the transmitter and receiver settings to maximize the link performance. Section 7.1.3 describes a bit error rate test performed using the final configuration. Finally, section 7.1.4 section 7.1.4 discusses the relevance of the test results.

7.1.1 Lab setup

The TFM card was mounted on a Tracking Processor (TP) mechanical demonstrator board. The mechanical demonstrator has less functionality than the complete SSTP design, but it was sufficient to test the TFM serial links. In the following, the TP mechanical demonstrator will be referred to as simply the "TP" for brevity.

The TFM serial links are accessed via SMA connectors on the TP. Two links were tested, which are referred to as link 7 and link 15. The number of available links was limited by the design of the TP mechanical demonstrator (which did not implement all of the the links), the design of the TFM test firmware, and physical damage to one of the SMA connectors. The transmitters and receivers of each link were connected using 6 in SMA cables.

A labeled photo of the lab setup is shown in fig. 7.1. The TFM was mounted on the TP using a total of 13 screws, 8 of which are used to create even compression across the four Z-Ray interposers that interface between the TFM and TP. A heat sink was mounted over the Stratix 10, and a fan provided additional cooling. A power supply delivered 48.7 V to the TFM. Two of the transceivers were connected in a loopback configuration using the SMA connectors on the TP. A USB Blaster cable linked the Stratix 10's JTAG connector to a computer. Intel Quartus Prime software was used to program the Stratix 10.

Prior to the following tests of the serial links, a number of steps were taken to prepare the TFM for optimal performance. In particular, care was taken to ensure uniform compression along the TFM–TP interface. The Z-Ray interposers are small, thin strips of PCB with dozens of tiny compression contacts on one side. Each contact corresponds to one pin connection between the TFM and TP. Because the contacts are so small, uneven compression may result in incomplete pin connections, which would be detrimental to the signal quality and would, in turn, limit the performance of the serial links. A procedure was designed to tighten the 8 Z-Ray screws to create uniform compression within a tolerance of 0.1 mm across the entire interface, which was verified by measuring the distance between the TFM and TP using the 24 measurement holes included on the TP for this purpose. The Z-Ray interposers were also replaced before further testing; their compression contacts are fragile and sustained damage during the repeated unmounting and remounting of the TFM while

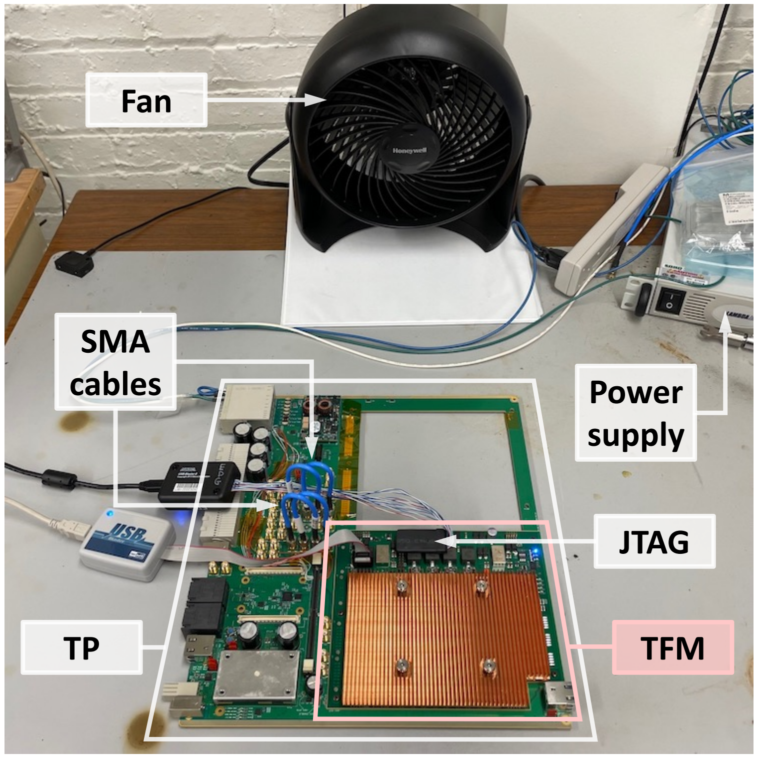

Figure 7.1: Lab setup for TFM serial link tests, with relevant components labeled.

the compression procedure was developed.

7.1.2 Tuning transmitter and receiver settings

The Stratix 10 has a number of equalization settings available for the transmitters and receivers. By default, the Quartus Prime software uses a set of default values for the transmitter settings ($V_{\mathrm{OD}} = 31$, pre-emphasis pre/post-tap = 0) and an adaptive mode for the receiver equalizers, but the serial link performance can be improved by tuning these settings. The settings and their possible values are listed in table 7.1.

For the transmitters, the adjustable settings are the differential output voltage, V_{OD}, and the pre-emphasis. The V_{OD} is the potential difference across the positive and negative transmitter pins, which controls the base amplitude of the transmitted signal. Its greatest possible value is $V_{\mathrm{OD}} = V_{\mathrm{CCT}}$, where V_{CCT} is the supply voltage to the transceiver. The transmitter pre-emphasis is a modification to the transmitted waveform shape that compen-

Table 7.1: Range of values for each transmitter and receiver setting [115].

Setting	Value range	Interpretation
V_{OD} control	17 to 31	0.55 to $1 \times V_{\mathrm{CCT}}$
Pre-emphasis post-tap	−15 to 15	6 to −6 dB
Pre-emphasis pre-tap	−24 to 24	14 to −14 dB
VGA DC gain	0 to 31	−5 to 5 dB
CTLE EQ gain	0 to 47	0 to 16 dB
CTLE AC gain	0 to 15	−5 to 10 dB

sates for possible interference between adjacent symbols or pulses. Two pre-emphasis taps are available on the Stratix 10: one pre-tap and one post-tap. The pre-tap emphasizes or de-emphasizes the bit just before the transition or edge, and the post-tap does the same for the bit just after the transition [116]. All of the transmitter settings can be adjusted manually.

For the receivers, variable gain amplification (VGA), continuous time linear equalization (CTLE), and decision feedback equalization (DFE) are available. The VGA DC gain, which can be manually or adaptively tuned, modifies the amplification of the incoming signal. Each receiver has a CTLE circuit that can be programmed to amplify the high-frequency part of the incoming signal. The CTLE AC gain can be adjusted manually or adaptively, as can the CTLE EQ gain. The EQ gain modifies the DC gain to de-emphasize the low-frequency part of the incoming signal. DFE is similar to CTLE, but has an improved signal-to-noise ratio, and can only be adjusted adaptively [115].

All settings were varied one at a time in the order specified by the Intel guide for transceiver high-speed link tuning [117] to determine the optimal configuration. The Stratix 10 Transceiver Toolkit software was used to generate eye diagrams to assess the serial link performance for each step of the tuning procedure. An eye diagram is a superposition of successive waveforms, and its shape resembles a human eye. The eye opening for an ideal digital signal is rectangular and represents a perfect square wave. However, actual signals have some degree of variation in signal amplitude, which affects the height of the eye open-

Table 7.2: Tuned transceiver settings for both TFM serial links.

Setting	Value	
	Link 7	Link 15
V_{OD} control	31	31
Pre-emphasis post-tap	0	0
Pre-emphasis pre-tap	0	0
VGA DC gain	0	0
CTLE EQ gain	15	15
CTLE AC gain	15	15

ing, and in signal timing, which affects the width of the eye opening. The area of the eye opening, computed as the product of its width and height, is a useful figure of merit; a larger area indicates a higher-quality signal [118].

The tuned values for all settings are listed for both links in table 7.2. Notably, the optimal transceiver settings are identical to the default values, and all values are the same for both links. The procedure was later repeated, and slightly different values (within ± 1) were obtained. However, the eye areas were unchanged, indicating that there is some tolerance for the transceiver settings. These are positive outcomes because they indicate that the same tuned transceiver settings can be used for both links. Therefore, it is likely that all of the transceivers on the TFM could be programmed with the same settings, rather than each one requiring individual tuning, which would be time intensive.

Eye diagrams were generated using the default and the tuned settings to compare their performance. They are shown for both serial links in fig. 7.2. Their dimensions are compared in table 7.3, which indicates that the eye areas are approximately 50% greater when the tuned settings are used. Based on the size of these eye openings, it was concluded that the signal quality of the TFM serial links was satisfactory.

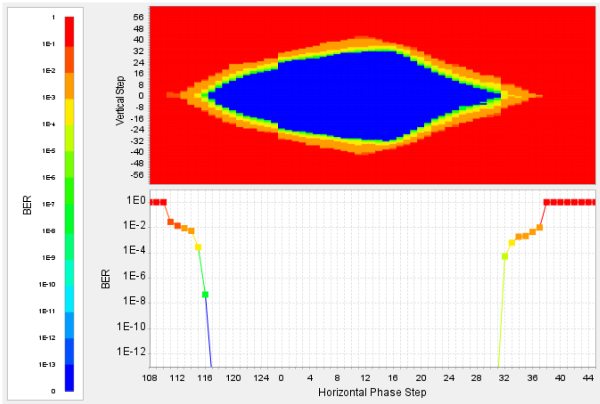

(a) Link 7, default settings

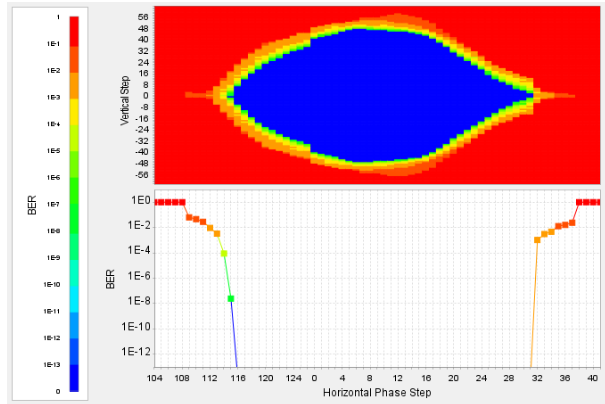

(b) Link 7, tuned settings

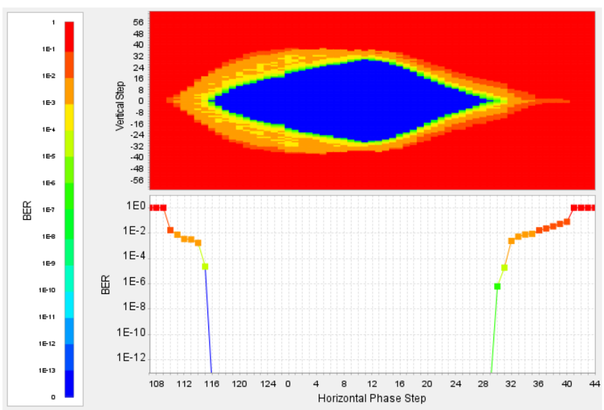

(c) Link 15, default settings

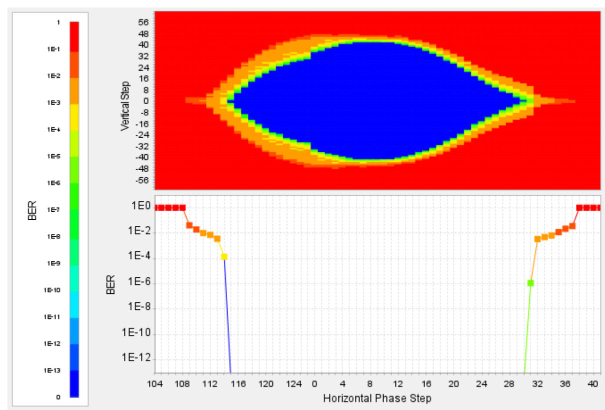

(d) Link 15, tuned settings

Figure 7.2: Comparison of eye diagrams using the default and tuned transceiver settings for both TFM serial links.

Table 7.3: Dimensions of eye openings for diagrams generated using default and tuned transceiver settings.

Link	Transceiver settings	Eye width	Eye height	Eye area
7	Default	43	64	2752
	Tuned	45	95	4275
15	Default	42	59	2478
	Tuned	44	83	3652

7.1.3 Bit error rate test

A bit error rate (BER) measurement is a key figure of merit in digital communications. A BER test (BERT) sends a stream of data from a transmitter to a receiver and compares the bits sent to those received. Any bit that differs between the transmitted sequence and the received sequence is marked as an error. The BER is defined as

$$\text{BER} = \frac{\text{number of error bits}}{\text{total number of bits transmitted}}. \tag{7.1}$$

Rigorous BERTs use a pseudorandom bit sequence (PRBS) as the data stream. The rigor of the test is affected by the length of the sequence, which is sent repeatedly. A PRBSn sequence is $2^n - 1$ bits long [119]. The BERTs for the TFM were performed using PRBS7.

The target BER for the TFM transceivers is below 10^{-15}. Approximately 3×10^{15} bits must be tested to claim a BER below 10^{-15} at 95% CL. This is a lengthy procedure. Even at the TFM's target data rate of 10 Gbps, such a test would take approximately 3.5 days to complete. It was not possible to perform BERTs of the necessary length until the final stages of the project due to concerns that the Stratix 10 may overheat. This concern was eventually resolved by the addition of a shutdown signal to the test firmware, which would power off the Stratix 10 if its temperature exceeded 50 °C. The typical operating temperature of the Stratix 10 was approximately 40 °C.

Two BERTs were performed, both using the tuned transceiver settings listed in table 7.2. The first BERT tested 2.84×10^{15} bits, and the second tested 3.39×10^{15} bits. Both tests measured a BER of 0 for both links. It can be concluded from these results that the BER is below the target 10^{-15} for both links tested.

7.1.4 Conclusion

Overall, the results of this project indicate that the TFM serial link performance is on target and that the optimal equalization settings are likely similar for all transceivers. These findings are significant because they would simplify the setup and maintenance of the complete HTT design, which includes 196 TFMs. Although ATLAS has elected not to use the HTT system for the trigger upgrade, understanding the capabilities of the TFM serial links will aid in future designs and plans for hardware tracking and the ATLAS trigger system.

7.2 High-level synthesis implementation of track-fitting firmware

After ATLAS elected not to use the HTT system for the trigger upgrade, a new plan was developed to implement tracking in the Event Filter (EF) system, which is roughly analogous to the HLT system used during Runs 1–3. When this decision was made, the plan for tracking at the EF level was to use one CPU and one FPGA accelerator per event. A proposal was made to repurpose the track-fitting firmware designed by the UChicago group for the HTT system to filter out fake tracks in the EF, which would require adapting it to the EF model of one FPGA used for an entire event. Previous iterations of the firmware considered only one RoI at a time.

Part of the proposal was to use high-level synthesis (HLS) to write a new implementation of the track-fitting firmware in order to improve the maintainability and adaptability of the design. In recent decades, the typical method of firmware design takes place at the register-transfer level (RTL), which describes how data are manipulated by hardware components, such as adders, as they are moved between registers [120]. Because this method requires manual translation of an abstract design to a single hardware system, it is hardware-dependent, tedious, and error-prone.

These shortcomings are addressed by HLS, which “synthesizes” an RTL design from an implementation in a higher-level programming language like C++. From the user’s perspective, the process is similar to compiling a C++ program into machine language for a specific computer. Because HLS allows designs to be finalized at a higher level of abstraction, it shortens development time, eases maintenance, and increases design reusability [121]. These considerations are particularly important for a large collaboration like ATLAS, in which there is frequent turnover for firmware development and maintenance. In the context of upgrade work, the design reusability feature of HLS in principle enables the same design to be implemented on a number of technologies being considered; in practice, reusability is limited by the proprietary nature of the FPGA technologies and corresponding HLS software.

The firmware design was implemented using the Intel HLS Compiler, which accepts designs written using a limited subset of C++. The target was the Stratix 10 FPGA, but the HLS compiler enables design testing on CPU using a simulation of the Stratix 10. The project was terminated before it reached the stage of testing on the physical Stratix 10.

This section describes the track-fitting firmware and the HLS implementation procedure. The firmware design is described in section 7.2.1. The HLS implementation process is described in section 7.2.2. Section 7.2.3 provides concluding thoughts on the project.

7.2.1 Firmware design

The trackfitter design was originally implemented for a trigger-level tracking system that was to be implemented before Run 3, and so assumes the ID configuration described in section 3.2.2. Since a complete RTL implementation of this design already existed, the HLS implementation was to replicate this design in order to compare the performance of the two implementations and use the results to draw a conclusion about the utility of HLS.

The trackfitter design, shown as a block diagram in fig. 7.3, consists of several discrete units. The first, the input decoder, receives and parses a stream of data that contains the coordinates of hits in the pixel detector and SCT. The input data may include up to 3 hits in any given layer, and up to one "empty" layer with 0 hits is allowed. The decoded hits are

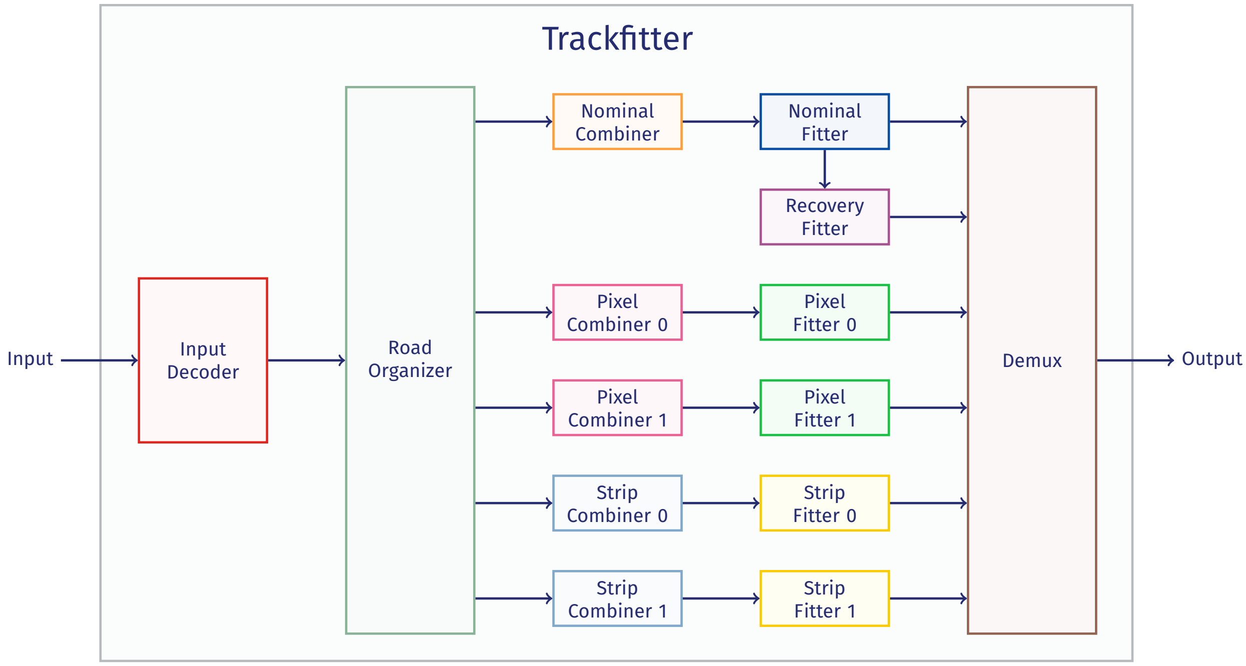

Figure 7.3: Block diagram of trackfitter design.

organized by layer into a data collection called a "road," which is passed to the next unit, the road organizer. The road organizer determines what type of fit is needed based on the location of the empty layer, if one exists. Roads without an empty layer need a "nominal" fit, while roads with an empty pixel layer need a "pixel" fit, and roads with an empty SCT layer need a "strip" fit. The different fit types account for the different number of parameters involved in the fit calculation, which differs based on whether a layer is empty and, if so, which kind it is.

The trackfitter design allows it to perform up to five fits in parallel: one nominal fit, two pixel fits, and two strip fits. This distribution is based on the expected frequency of roads with missing hits, which are more common than roads without them. After determining what kind of fit needs to be performed, the road organizer transfers the road to a "combiner" block that prepares it for the fit. The combiner assembles all possible tracks with one hit in each nonempty layer and passes them one by one to the "fitter" block, which performs a linear fit using matrix multiplication. The fit quality is assessed with a χ^2 calculation, and if it is below a certain threshold, the fit result is passed to the demux block, which formats the fit data and passes it on. If a nominal fit does not pass the χ^2 threshold, it is passed to the recovery fitter, which retries the fit with each layer removed. This makes it possible to recover tracks that may have erroneously been associated with an unrelated hit.

One of the challenges of the trackfitter design is arbitration. In general, each road contains a different number of hits in each layer and so correponds to a different number of track candidates. This means the time required to perform the fit may vary significantly between roads, so the various fitters and combiners will be busy for different lengths of time. Part of the road organizer's functionality is to check whether the various combiners are occupied; this process is called "arbitration" and requires communication between the road organizer and the combiners.

7.2.2 HLS implementation

The HLS implementation of the trackfitter was prepared in three major stages, illustrated in fig. 7.4. The first stage of development was a pure C++ implementation of the trackfitter design without regard to the limitations of the HLS compiler. This way, a correct implementation could be written quickly and used to verify that later iterations continued to produce correct results. A set of "test vectors," being artificial sets of hits used to test the various aspects of the trackfitter design, were created to test the pure C++ and later implementations.

The second stage was to adapt the C++ implementation to use only the subset of C++ supported by the Intel HLS compiler. The compiler does not support many basic features of C++, including, for example, dynamic memory allocation or recursive functions, because these constructs cannot be replicated in the FPGA's digital logic [122, 123]. For this step, the individual units of the trackfitter design were adapted one by one as standalone functions, and their functionality and performance were verified using a software testbench.

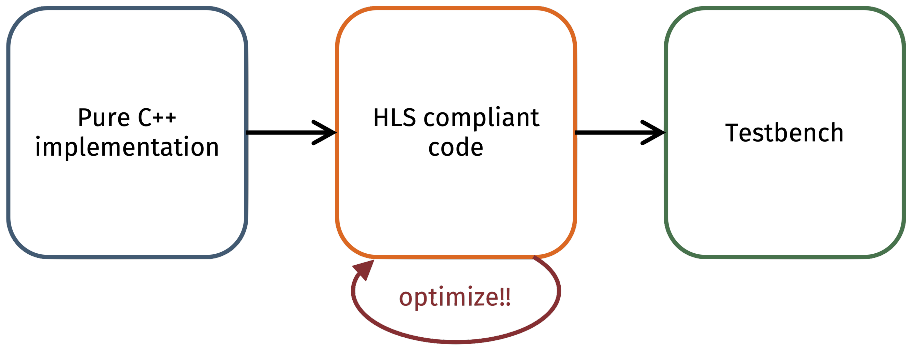

Figure 7.4: Workflow for HLS trackfitter implementation.

It is not sufficient for the HLS implementation of each unit to simply be written using C++ patterns supported by the HLS compiler; care must also be taken to optimize its performance. Without HLS-specific optimization, the RTL design synthesized by the compiler will be inefficient in terms of both “area,” which is hardware resource usage, and “latency,” which is the time delay between subsequent operations. Optimization can be achieved by restructuring the C++ code and by providing specific instructions to the HLS compiler using special preprocessor directives that tell it, for example, when a loop may be parallelized [124]. The optimization process requires many iterations. Each unit was optimized until its latency and area were small enough to run easily on the Stratix 10. As an example, the optimized road organizer used far less than 1% of the Stratix 10's area and took 59 clock cycles to complete for one test vector. Further optimization was deferred to a later stage of the project.

The combiners and fitters must run asynchronously, so they were implemented using a system of tasks. A system of tasks consists of a number of “task functions” that may run in parallel [122]. Essentially, a system of tasks emulates the behavior of a group of interconnected firmware components, but within a single component. While it is possible and often standard practice to use HLS to design individual components that are interfaced at RTL, for this project a single-component approach was attempted to benefit as much as possible from the advantages of HLS. However, designing a messaging system between the road organizer and combiners to facilitate the arbitration proved to be difficult if not impossible to achieve within the system of tasks, due to limitations on the number of read and write sites for data streams.

The final stage, the testbench, would implement the HLS trackfitter design on the actual Stratix 10 hardware. This stage would require a fully functional trackfitter design, including arbitration, but was not reached during the lifetime of the project. Had it been reached, the latency and area of both trackfitter designs would have been compared directly. Depending on these results, further optimization of the HLS implementation may have been attempted. If the HLS implementation's performance was found to be adequate, it would have been updated to meet the specifications of the upgraded ATLAS detector and the "one event, one board" model for the EF.

7.2.3 Conclusion

Before the testbench stage of the project was reached, Intel announced that it would discontinue support of its HLS compiler. As a result of this development, the project was shelved.

Although it was not completed, the project provided insights regarding the application of HLS that may inform future efforts. One lesson learned was the utility of the pure C++ implementation of the trackfitter design. While this step was not strictly necessary, it was a useful starting point; while an RTL implementation was available, the HLS-compliant code is still more similar to C++ than to hardware description language. The pure C++ implementation also provided a useful benchmark for accuracy, since it could be launched using a similar software testbench to the one written for the HLS. Another lesson concerns the difficulty of arbitration within the system of tasks. The utility of a system of tasks could be gauged in future projects by, for example, attempting to set up arbitration among a system of dummy tasks. This could expedite the process of learning whether a system of tasks is sufficient or whether the project would be better implemented as HLS components interfaced at RTL.

CHAPTER 8

CONCLUSION

This dissertation presents two searches for new particles that interact with the Higgs boson, using the Higgs boson itself as a tool for discovery. Both searches use the ATLAS Run 2 dataset, consisting of $140\,\text{fb}^{-1}$ of pp collisions with center-of-mass energy $\sqrt{s} = 13\,\text{TeV}$.

The first search, published in ref. 2, targets invisible and partially invisible decays of the Higgs boson. The $H \to \text{inv.}$ interpretation targets the VBF production mode of the Higgs boson with a radiated photon, and sets an upper limit on the $H \to \text{inv.}$ branching ratio of 37%. The $H \to \gamma\gamma_\text{d}$ interpretation provides experimental constraints on Higgs boson decays to a photon and a dark photon, with an upper limit on the $H \to \gamma\gamma_\text{d}$ branching fraction of 1.8%, as well as constraints on similar decays mediated by a BSM boson with mass between 60 GeV and 2 TeV. Finally, this search also presents the first observation of the EW $Z(\to \nu\nu)\gamma\text{jj}$ process, one of the backgrounds to the $H \to \text{inv.}$ and $H \to \gamma\gamma_\text{d}$ searches, with an observed significance of 5.2σ. A combination with a subsequent ATLAS search for the same process in ref. 3 increases the significance to 6.3σ.

The second search, published in ref. 4, targets the $X \to SH \to bbbb$ process. Here, X and S are hypothetical scalar bosons included in many models of new physics that include an extended Higgs sector. This search leverages the large $H \to b\bar{b}$ branching ratio and the similar $S \to b\bar{b}$ decay, and uses a novel GP + NF method to estimate background contributions from probability distributions learned from real data via machine learning techniques. A wide range of mass hypotheses, from 70 GeV to 5 TeV for m_S and from 300 GeV to 6 TeV for m_X, are considered. Upper limits ranging from 0.7 fb to 2.6 pb are set on the $pp \to X \to SH \to 4b$ cross section. The most stringent limits come from the mixed channel, in which the $H \to b\bar{b}$ decay is reconstructed from a single large-radius jet, which targets signal mass hypotheses with large m_X.

While neither search shows evidence for new physics, both are nonetheless worthwhile contributions to our understanding of the Higgs boson. The journey ahead, to an even better understanding, will require us to explore multiple paths—some obvious, some not so obvious. As a straightforward path, both of the searches presented in this dissertation can be expanded to include data from Run 3 of the LHC, which concluded in June 2026, shortly before this dissertation was completed. The Run 3 dataset is over $300\,\mathrm{fb}^{-1}$ in size and will improve the statistical precision of all results. This will be particularly relevant to the mixed channel of the $X \to SH \to 4b$ search, which is statistically limited. But systematic uncertainties must also be reduced; in particular, the theoretical uncertainties in the first search, the $H \to bb$ tagging uncertainties in the second, and the background estimation uncertainties in both.

In conclusion, future use of the Higgs boson as a discovery tool will require not only more data, but also more creativity and rigor. As we look toward the HL-LHC, we must also continue to make innovations in analysis strategy. Only time may tell whether the Higgs boson will lead us to the answers we seek about our universe. But the search is far from over, and so long as we press on, the future of Higgs physics will certainly be bright.

BIBLIOGRAPHY

[1] CERN, *CERN experiments observe particle consistent with long-sought Higgs boson*, press release, (July 2012) `https://home.web.cern.ch/news/press-release/cern/cern-experiments-observe-particle-consistent-long-sought-higgs-boson` (visited on 04/11/2026).

[2] ATLAS Collaboration, "Observation of electroweak production of two jets in association with an isolated photon and missing transverse momentum, and search for a Higgs boson decaying into invisible particles at 13 TeV with the ATLAS detector", Eur. Phys. J. C **82**, 105 (2022), arXiv:`2109.00925 [hep-ex]`, `https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PAPERS/EXOT-2021-17`.

[3] ATLAS Collaboration, "Measurement of electroweak $Z(\nu\bar{\nu})\gamma jj$ production and limits on anomalous quartic gauge couplings in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector", JHEP **06**, 082 (2023), arXiv:`2208.12741 [hep-ex]`.

[4] ATLAS Collaboration, "Search for new scalars via $X \to SH \to b\bar{b}b\bar{b}$ in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector", submitted to Phys. Rev. D (2026), arXiv:`2607.18484 [hep-ex]`, `https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PAPERS/HIGP-2024-31`.

[5] D. Griffiths, *Introduction to Elementary Particles*, 2nd ed. (Wiley-VCH, Portland, OR, 2008).

[6] L. O'Raifeartaigh and N. Straumann, "Gauge theory: Historical origins and some modern developments", Rev. Mod. Phys. **72**, edited by J.-P. Hsu and D. Fine, 1–23 (2000).

[7] M. E. Peskin and D. V. Schroeder, *An Introduction to Quantum Field Theory* (Perseus Books, Reading, MA, 1995).

[8] Wikimedia Commons, “Standard Model of Elementary Particles”, Sept. 2019, https://en.wikipedia.org/wiki/File:Standard_Model_of_Elementary_Particles.svg.

[9] Particle Data Group, “Review of particle physics”, Phys. Rev. D **110**, 030001 (2024), https://pdg.lbl.gov/.

[10] ATLAS Collaboration, “Observation of a cross-section enhancement near the $t\bar{t}$ production threshold in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector”, Rept. Prog. Phys. **89**, 057801 (2026), arXiv:2601.11780 [hep-ex].

[11] ATLAS Collaboration, *Boosting precision of top-quark mass measurement with ATLAS*, physics briefing, (Mar. 2025) https://atlas.cern/Updates/Briefing/Boosting-Top-Quark-Mass (visited on 07/02/2026).

[12] A. Djouadi, “The Anatomy of Electro-Weak Symmetry Breaking. I: The Higgs Boson in the Standard Model”, Phys. Rept. **457**, 1–216 (2008), arXiv:hep-ph/0503172.

[13] CERN, *The European Organization for Nuclear Research is born*, https://timeline.web.cern.ch/european-organization-nuclear-research-born (visited on 06/16/2026).

[14] “LHC Machine”, JINST **3**, edited by L. Evans and P. Bryant, S08001 (2008), https://cds.cern.ch/record/1129806.

[15] Service graphique, CERN, “Overall view of the LHC”, 2014, https://cds.cern.ch/record/1708849.

[16] CERN, *The HL-LHC project*, (Jan. 2025) https://hilumilhc.web.cern.ch/content/hl-lhc-project (visited on 06/17/2026).

[17] D. Brandt, H. Burkhardt, M. Lamont, S. Myers, and J. Wenninger, “Accelerator physics at LEP”, Rept. Prog. Phys. **63**, 939–1000 (2000), https://cds.cern.ch/record/448334.

[18] The LHC Study Group, *The Large Hadron Collider: Conceptual design*, tech. rep. (Oct. 1995), https://cds.cern.ch/record/291782/.

[19] CERN, *How an accelerator works*, https://home.cern/how-accelerator-works/ (visited on 06/17/2026).

[20] CERN, *Linear accelerator 2*, https://home.cern/science/accelerators/linear-accelerator-2/ (visited on 06/17/2026).

[21] CERN, *The Super Proton Synchrotron*, https://home.cern/science/accelerators/super-proton-synchrotron/ (visited on 06/17/2026).

[22] CERN, *Accelerating: Radiofrequency cavities*, https://home.web.cern.ch/science/engineering/accelerating-radiofrequency-cavities/ (visited on 06/17/2026).

[23] ATLAS Collaboration, "Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC", Eur. Phys. J. C **83**, 982 (2023), arXiv:2212.09379 [hep-ex].

[24] J. Pequenao, "Computer generated image of the whole ATLAS detector", 2008, https://cds.cern.ch/record/1095924.

[25] ATLAS Collaboration, "The ATLAS Experiment at the CERN Large Hadron Collider", JINST **3**, S08003 (2008), https://cds.cern.ch/record/1129811/.

[26] I. Neutelings, *CMS coordinate system*, (June 2017) https://tikz.net/axis3d_cms/.

[27] E. Daw, "Lecture 7 - Rapidity and Pseudorapidity", Mar. 2012, https://www.hep.shef.ac.uk/edaw/PHY206/Site/2012_course_files/phy206rlec7.pdf.

[28] I. Neutelings, *Pseudorapidity*, (Mar. 2026) https://tikz.net/axis2d_pseudorapidity/.

[29] H. D. Young and R. A. Freedman, *University Physics with Modern Physics*, 15th ed. (Pearson, 2020).

[30] J. Pequenao, "Computer generated image of the ATLAS inner detector", 2008, https://cds.cern.ch/record/1095926.

[31] G. F. Knoll, *Radiation Detection and Measurement*, 4th ed. (Wiley, 2010).

[32] ATLAS IBL Collaboration, "Production and Integration of the ATLAS Insertable B-Layer", JINST **13**, T05008 (2018), arXiv:1803.00844 [physics.ins-det].

[33] H. Kolanoski and N. Wermes, *Particle Detectors: Fundamentals and Applications* (Oxford University Press, 2020).

[34] J. Pequenao, "Computer Generated image of the ATLAS calorimeter", 2008, https://cds.cern.ch/record/1095927.

[35] ATLAS Collaboration, "Deep Generative Models for Fast Photon Shower Simulation in ATLAS", Comput. Softw. Big Sci. **8**, 7 (2024), arXiv:2210.06204 [hep-ex].

[36] ATLAS Collaboration, "Operation and performance of the ATLAS tile calorimeter in LHC Run 2", Eur. Phys. J. C **84**, 1313 (2024), arXiv:2401.16034 [hep-ex].

[37] J. Pequenao, "Computer generated image of the ATLAS Muons subsystem", 2008, https://cds.cern.ch/record/1095929.

[38] G. Avoni et al., "The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS", JINST **13**, P07017 (2018).

[39] ATLAS Collaboration, "Operation of the ATLAS trigger system in Run 2", JINST **15**, P10004 (2020), arXiv:2007.12539 [physics.ins-det].

[40] ATLAS Collaboration, "Performance of the ATLAS Trigger System in 2015", Eur. Phys. J. C **77**, 317 (2017), arXiv:1611.09661 [hep-ex].

[41] ATLAS Collaboration, *Data Quality Efficiency Plots*, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/DataQualityResults (visited on 06/17/2026).

[42] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, tech. rep. (CERN, Geneva, June 2019), https://cds.cern.ch/record/2677054.

[43] ATLAS Collaboration, "ATLAS data quality operations and performance for 2015–2018 data-taking", JINST **15**, P04003 (2020), arXiv:1911.04632 [physics.ins-det].

[44] ATLAS Collaboration, *Public ATLAS Online Luminosity Plots for Run-2 of the LHC*, https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2 (visited on 06/17/2026).

[45] ATLAS Collaboration, *ATLAS Open Data Documentation*, (2026) https://opendata.atlas.cern/docs/documentation/introduction/introductory_page (visited on 06/29/2026).

[46] S. Mehlhase, "ATLAS detector slice (and particle visualisations)", 2021, https://cds.cern.ch/record/2770815.

[47] ATLAS Collaboration, "Track and Vertex Reconstruction with the ATLAS Inner Detector", (2026), arXiv:2605.07585 [physics.ins-det].

[48] ATLAS Collaboration, "Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data", JINST **14**, P12006 (2019), arXiv:1908.00005 [hep-ex].

[49] ATLAS Collaboration, "Electron and photon energy calibration with the ATLAS detector using 2015–2016 LHC proton-proton collision data", JINST **14**, P03017 (2019), arXiv:1812.03848 [hep-ex].

[50] ATLAS Collaboration, "Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton-proton collision data at $\sqrt{s} = 13$ TeV", Eur. Phys. J. C **79**, 639 (2019), arXiv:1902.04655 [physics.ins-det].

[51] ATLAS Collaboration, "Electron and photon efficiencies in LHC Run 2 with the ATLAS experiment", JHEP **05**, 162 (2024), arXiv:2308.13362 [hep-ex].

[52] ATLAS Collaboration, "Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13$ TeV", Eur. Phys. J. C **76**, 292 (2016), arXiv:1603.05598 [hep-ex].

[53] ATLAS Collaboration, "Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV", Eur. Phys. J. C **81**, 578 (2021), arXiv:2012.00578 [hep-ex].

[54] M. Cacciari, G. P. Salam, and G. Soyez, "The anti-k_t jet clustering algorithm", JHEP **04**, 063 (2008), arXiv:0802.1189 [hep-ph].

[55] ATLAS Collaboration, "Jet reconstruction and performance using particle flow with the ATLAS Detector", Eur. Phys. J. C **77**, 466 (2017), arXiv:1703.10485 [hep-ex].

[56] ATLAS Collaboration, "Optimisation of large-radius jet reconstruction for the ATLAS detector in 13 TeV proton–proton collisions", Eur. Phys. J. C **81**, 334 (2021), arXiv:2009.04986 [hep-ex].

[57] ATLAS Collaboration, *Improving jet substructure performance in ATLAS using Track-CaloClusters*, tech. rep. (CERN, Geneva, 2017), https://cds.cern.ch/record/2275636.

[58] ATLAS Collaboration, "Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector", Eur. Phys. J. C **81**, 689 (2021), arXiv:2007.02645 [hep-ex].

[59] ATLAS Collaboration, "In situ calibration of large-radius jet energy and mass in 13 TeV proton–proton collisions with the ATLAS detector", Eur. Phys. J. C **79**, 135 (2019), arXiv:1807.09477 [hep-ex].

[60] ATLAS Collaboration, *Measurement of the ATLAS Detector Jet Mass Response using Forward Folding with 80 fb^{-1} of $\sqrt{s}$ = 13 TeV pp data*, tech. rep. (CERN, Geneva, 2020), https://cds.cern.ch/record/2724442.

[61] ATLAS Collaboration, "The performance of missing transverse momentum reconstruction and its significance with the ATLAS detector using 140 fb^{-1} of $\sqrt{s} = 13$ TeV pp collisions", Eur. Phys. J. C **85**, 606 (2025), arXiv:2402.05858 [hep-ex].

[62] ATLAS Collaboration, "Performance of pile-up mitigation techniques for jets in *pp* collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector", Eur. Phys. J. C **76**, 581 (2016), arXiv:1510.03823 [hep-ex].

[63] ATLAS Collaboration, *Forward jet vertex tagging using the particle flow algorithm*, tech. rep. (CERN, Geneva, 2019), https://cds.cern.ch/record/2683100.

[64] ATLAS Collaboration, "Measurements of b-jet tagging efficiency with the ATLAS detector using $t\bar{t}$ events at $\sqrt{s} = 13$ TeV", JHEP **08**, 089 (2018), arXiv:1805.01845 [hep-ex].

[65] ATLAS Collaboration, "Transforming jet flavour tagging at ATLAS", Nature Commun. **17**, 541 (2026), arXiv:2505.19689 [hep-ex].

[66] ATLAS Collaboration, *Transformer Neural Networks for Identifying Boosted Higgs Bosons decaying into $b\bar{b}$ and $c\bar{c}$ in ATLAS*, tech. rep. (CERN, Geneva, 2023), https://cds.cern.ch/record/2866601.

[67] LHC Higgs Cross Section Working Group, "Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector", CERN Yellow Rep. Monogr. **2**, 1–869 (2017), arXiv:1610.07922 [hep-ph].

[68] ATLAS Collaboration, "Combination of searches for invisible decays of the Higgs boson using 139 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment", Phys. Lett. B **842**, 137963 (2023), arXiv:2301.10731 [hep-ex].

[69] R. Essig et al., "Working Group Report: New Light Weakly Coupled Particles", in Snowmass 2013: Snowmass on the Mississippi (Oct. 2013), arXiv:1311.0029 [hep-ph].

[70] L. Lopez-Honorez, T. Schwetz, and J. Zupan, "Higgs portal, fermionic dark matter, and a Standard Model like Higgs at 125 GeV", Phys. Lett. B **716**, 179–185 (2012), arXiv:1203.2064 [hep-ph].

[71] S. Biswas, E. Gabrielli, M. Heikinheimo, and B. Mele, "Dark-Photon searches via Higgs-boson production at the LHC", Phys. Rev. D **93**, 093011 (2016), arXiv:1603.01377 [hep-ph].

[72] E. Gabrielli, F. Maltoni, B. Mele, M. Moretti, F. Piccinini, and R. Pittau, "Higgs Boson Production in Association with a Photon in Vector Boson Fusion at the LHC", Nucl. Phys. B **781**, 64–84 (2007), arXiv:hep-ph/0702119.

[73] ATLAS Collaboration, "Combination of searches for invisible Higgs boson decays with the ATLAS experiment", Phys. Rev. Lett. **122**, 231801 (2019), arXiv:1904.05105 [hep-ex].

[74] CMS Collaboration, "Combined measurements of Higgs boson couplings in proton–proton collisions at $\sqrt{s} = 13$ TeV", Eur. Phys. J. C **79**, 421 (2019), arXiv:1809.10733 [hep-ex].

[75] CMS Collaboration, "A search for decays of the Higgs boson to invisible particles in events with a top-antitop quark pair or a vector boson in proton-proton collisions at $\sqrt{s} = 13$ TeV", Eur. Phys. J. C **83**, 933 (2023), arXiv:2303.01214 [hep-ex].

[76] CMS Collaboration, "Search for dark photons in decays of Higgs bosons produced in association with Z bosons in proton-proton collisions at $\sqrt{s} = 13$ TeV", JHEP **10**, 139 (2019), arXiv:1908.02699 [hep-ex].

[77] CMS Collaboration, "Search for dark photons in Higgs boson production via vector boson fusion in proton-proton collisions at $\sqrt{s} = 13$ TeV", JHEP **03**, 011 (2021), arXiv:2009.14009 [hep-ex].

[78] ATLAS Collaboration, "Combination of searches for Higgs boson decays into a photon and a massless dark photon using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector", JHEP **08**, 153 (2024), arXiv:2406.01656 [hep-ex].

[79] ATLAS Collaboration, "Search for dark photons from Higgs boson decays in the gluon-gluon fusion channel in proton-proton collisions at $\sqrt{s} = 13.6$ TeV with the ATLAS detector", (2026), arXiv:2606.14362 [hep-ex].

[80] CMS Collaboration, "Measurement of the electroweak production of Zγ and two jets in proton-proton collisions at $\sqrt{s} = 13$ TeV and constraints on anomalous quartic gauge couplings", Phys. Rev. D **104**, 072001 (2021), arXiv:2106.11082 [hep-ex].

[81] ATLAS Collaboration, "Search for invisible Higgs-boson decays in events with vector-boson fusion signatures using 139 fb^{-1} of proton-proton data recorded by the ATLAS experiment", JHEP **08**, 104 (2022), arXiv:2202.07953 [hep-ex].

[82] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H. Shao, T. Stelzer, P. Torrielli, and M. Zaro, "The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations", JHEP **07**, 079 (2014), arXiv:1405.0301 [hep-ph].

[83] J. F. Gunion, "Extended Higgs Sectors", in 10th International Conference on Supersymmetry and Unification of Fundamental Interactions (SUSY02) (Dec. 2002), pp. 80–103, arXiv:hep-ph/0212150.

[84] V. Barger, P. Langacker, M. McCaskey, M. Ramsey-Musolf, and G. Shaughnessy, "Complex Singlet Extension of the Standard Model", Phys. Rev. D **79**, 015018 (2009), arXiv:0811.0393 [hep-ph].

[85] M. Maniatis, “The Next-to-Minimal Supersymmetric extension of the Standard Model reviewed”, Int. J. Mod. Phys. A **25**, 3505–3602 (2010), arXiv:0906.0777 [hep-ph].

[86] X.-G. He, T. Li, X.-Q. Li, J. Tandean, and H.-C. Tsai, “Constraints on Scalar Dark Matter from Direct Experimental Searches”, Phys. Rev. D **79**, 023521 (2009), arXiv:0811.0658 [hep-ph].

[87] T. Robens, T. Stefaniak, and J. Wittbrodt, “Two-real-scalar-singlet extension of the SM: LHC phenomenology and benchmark scenarios”, Eur. Phys. J. C **80**, 151 (2020), arXiv:1908.08554 [hep-ph].

[88] U. Ellwanger, C. Hugonie, and A. M. Teixeira, “The Next-to-Minimal Supersymmetric Standard Model”, Phys. Rept. **496**, 1–77 (2010), arXiv:0910.1785 [hep-ph].

[89] I. F. Ginzburg, M. Krawczyk, and P. Osland, “Two Higgs doublet models with CP violation”, in International Workshop on Linear Colliders (LCWS 2002) (Nov. 2002), pp. 703–706, arXiv:hep-ph/0211371.

[90] M. Muhlleitner, M. O. P. Sampaio, R. Santos, and J. Wittbrodt, “The N2HDM under Theoretical and Experimental Scrutiny”, JHEP **03**, 094 (2017), arXiv:1612.01309 [hep-ph].

[91] S. Baum and N. R. Shah, “Two Higgs Doublets and a Complex Singlet: Disentangling the Decay Topologies and Associated Phenomenology”, JHEP **12**, 044 (2018), arXiv:1808.02667 [hep-ph].

[92] ATLAS Collaboration, “Search for the $HH \to b\bar{b}b\bar{b}$ process via vector-boson fusion production using proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector”, JHEP **07**, [Erratum: JHEP 01, 145 (2021), Erratum: JHEP 05, 207 (2021)], 108 (2020), arXiv:2001.05178 [hep-ex].

[93] ATLAS Collaboration, "Search for resonant pair production of Higgs bosons in the $b\bar{b}b\bar{b}$ final state using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector", Phys. Rev. D **105**, 092002 (2022), arXiv:`2202.07288 [hep-ex]`.

[94] ATLAS Collaboration, "Search for pair production of boosted Higgs bosons via vector-boson fusion in the $b\bar{b}b\bar{b}$ final state using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector", Phys. Lett. B **858**, 139007 (2024), arXiv:`2404.17193 [hep-ex]`.

[95] CMS Collaboration, "Search for resonant pair production of Higgs bosons in the $\mathrm{b\bar{b}b\bar{b}}$ final state using large-area jets in proton-proton collisions at $\sqrt{s} = 13$ TeV", JHEP **02**, 040 (2025), arXiv:`2407.13872 [hep-ex]`.

[96] ATLAS Collaboration, "Search for a new heavy scalar particle decaying into a Higgs boson and a new scalar singlet in final states with one or two light leptons and a pair of τ-leptons with the ATLAS detector", JHEP **10**, 009 (2023), arXiv:`2307.11120 [hep-ex]`.

[97] ATLAS Collaboration, "Search for a resonance decaying into a scalar particle and a Higgs boson in the final state with two bottom quarks and two photons in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector", JHEP **11**, 047 (2024), arXiv:`2404.12915 [hep-ex]`.

[98] ATLAS Collaboration, "Search for a resonance decaying into a scalar particle and a Higgs boson in the final state with two bottom quarks and two photons with 199 fb^{-1} of data collected at $\sqrt{s} = 13$ TeV and $\sqrt{s} = 13.6$ TeV with the ATLAS detector", Phys. Lett. B **877**, 140425 (2026), arXiv:`2510.02857 [hep-ex]`.

[99] ATLAS Collaboration, "Search for a resonance decaying into a scalar particle and a Higgs boson in final states with leptons and two photons in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector", JHEP **10**, 104 (2024), arXiv:`2405.20926 [hep-ex]`.

[100] ATLAS Collaboration, "Search for triple Higgs boson production in the $6b$ final state using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector", Phys. Rev. D **111**, 032006 (2025), arXiv:2411.02040 [hep-ex].

[101] CMS Collaboration, "Search for a massive scalar resonance decaying to a light scalar and a Higgs boson in the four b quarks final state with boosted topology", Phys. Lett. B **842**, 137392 (2023), arXiv:2204.12413 [hep-ex].

[102] CMS Collaboration, "Search for a new heavy scalar resonance decaying into the Higgs boson and a new scalar particle in the $\mathrm{b\bar{b}b\bar{b}}$ final state using proton-proton collisions at $\sqrt{s} = 13$ TeV", (2026), arXiv:2605.02848 [hep-ex].

[103] CMS Collaboration, "Search for a heavy Higgs boson decaying into two lighter Higgs bosons in the $\tau\tau$bb final state at 13 TeV", JHEP **11**, 057 (2021), arXiv:2106.10361 [hep-ex].

[104] CMS Collaboration, "Search for a new resonance decaying into two spin-0 bosons in a final state with two photons and two bottom quarks in proton-proton collisions at $\sqrt{s} = 13$ TeV", JHEP **05**, 316 (2024), arXiv:2310.01643 [hep-ex].

[105] A. Hayrapetyan et al. (CMS), "Search for a new scalar resonance decaying to a Higgs boson and another new scalar particle in the final state with two bottom quarks and two photons in proton-proton collisions at $\sqrt{s} = 13$ TeV", JHEP **12**, 178 (2025), arXiv:2508.11494 [hep-ex].

[106] CMS Collaboration, "Search for the nonresonant and resonant production of a Higgs boson in association with an additional scalar boson in the $\gamma\gamma\tau\tau$ final state in proton-proton collisions at $\sqrt{s} = 13$ TeV", JHEP **05**, 186 (2026), arXiv:2506.23012 [hep-ex].

[107] CMS Collaboration, "Search for heavy scalar resonances decaying to Lorentz-boosted Higgs and Higgs-like bosons in the $\mathrm{b\bar{b}}$4q final state at $\sqrt{s} = 13$ TeV", (2026), arXiv:2602.00273 [hep-ex].

[108] ATLAS Collaboration, "Configuration and performance of the ATLAS b-jet triggers in Run 2", Eur. Phys. J. C **81**, 1087 (2021), arXiv:2106.03584 [hep-ex].

[109] ATLAS Collaboration, "Climbing to the Top of the ATLAS 13 TeV data", Phys. Rept. **1116**, 127–183 (2025), arXiv:2404.10674 [hep-ex].

[110] C. E. Rasmussen and C. K. I. Williams, *Gaussian Processes for Machine Learning* (MIT Press, 2006), https://gaussianprocess.org/gpml/.

[111] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, A. Müller, J. Nothman, G. Louppe, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and É. Duchesnay, "Scikit-learn: Machine Learning in Python", JMLR **12**, 2825–2830 (2011), arXiv:1201.0490 [cs.LG].

[112] G. Papamakarios, E. Nalisnick, D. J. Rezende, S. Mohamed, and B. Lakshminarayanan, "Normalizing Flows for Probabilistic Modeling and Inference", JMLR **22**, 1–64 (2021), arXiv:1912.02762 [stat.ML].

[113] C. Durkan, A. Bekasov, I. Murray, and G. Papamakarios, "Neural Spline Flows", in Advances in Neural Information Processing Systems, Vol. 32, edited by H. Wallach, H. Larochelle, A. Beygelzimer, F. d'Alché-Buc, E. Fox, and R. Garnett (2019), arXiv:1906.04032 [stat.ML].

[114] ATLAS Collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS TDAQ System*, tech. rep. (CERN, Geneva, Sept. 2017), https://cds.cern.ch/record/2285584.

[115] *L- and H-Tile Transceiver PHY User Guide*, Intel Corporation (Sept. 2021), https://www.intel.com/content/www/us/en/programmable/documentation/wry1479165198810.html.

[116] *Understanding the Pre-Emphasis and Linear Equalization Features in Stratix IV GX Devices*, Altera Corporation (Nov. 2010), https://www.intel.com/content/dam/www/programmable/us/en/pdfs/literature/an/an602.pdf.

[117] *Quick Guide for Intel Arria 10 and Intel Cyclone 10 GX Transceiver High-Speed Link Tuning*, Intel Corporation (Sept. 2018), https://www.intel.com/content/www/us/en/programmable/documentation/fbr1532995817619.html.

[118] E. Säckinger, "Eye Diagrams", in *Analysis and Design of Transimpedance Amplifiers for Optical Receivers* (Wiley, 2017) Chap. B, pp. 413–420, https://onlinelibrary.wiley.com/doi/pdf/10.1002/9781119264422.app2.

[119] J. Redd, "Explaining those BER testing mysteries", Lightwave Magazine (2004), https://www.lightwaveonline.com/home/article/16647704/explaining-those-ber-testing-mysteries.

[120] F. Vahid, "Register-Transfer Level (RTL) Design", in *Digital Design with RTL Design, VHDL, and Verilog*, 2nd ed. (Wiley, 2011) Chap. 5, pp. 247–316.

[121] M. Fingeroff, *High-Level Synthesis Blue Book* (Mentor Graphics, 2010), https://cse.usf.edu/~haozheng/teach/cda4253/doc/hls/hls_bluebook_uv.pdf.

[122] *Intel High Level Synthesis Compiler Pro Edition: Reference Manual*, Intel Corporation (Dec. 2021), https://docs.altera.com/r/docs/683349/21.4/intel-high-level-synthesis-compiler-pro-edition-reference-manual.

[123] *Intel High Level Synthesis Compiler Pro Edition: User Guide*, Intel Corporation (Dec. 2021), https://docs.altera.com/r/docs/683456/21.4/intel-high-level-synthesis-compiler-pro-edition-user-guide.

[124] *Intel High Level Synthesis Compiler Pro Edition: Best Practices Guide*, Intel Corporation (Dec. 2021), https://docs.altera.com/r/docs/683152/21.4/intel-high-level-synthesis-compiler-pro-edition-best-practices-guide.