

Searches for Higgs pair production in the 4 b-jet final state with the ATLAS Detector at the LHC

by

Christopher Gubbels

B.Sc., McMaster University, 2017

A DISSERTATION SUBMITTED IN PARTIAL FULFILLMENT
OF THE REQUIREMENTS FOR THE DEGREE OF

Doctor of Philosophy

in

THE FACULTY OF GRADUATE AND POSTDOCTORAL STUDIES
(Physics)

The University of British Columbia
(Vancouver)

July 2022

© Christopher Gubbels, 2022



The following individuals certify that they have read, and recommend to the Faculty of Graduate and Postdoctoral Studies for acceptance, the thesis entitled:

Searches for Higgs pair production in the 4 b-jet final state with the ATLAS Detector at the LHC

submitted by **Christopher Gubbels** in partial fulfillment of the requirements for the degree of **Doctor of Philosophy** in **Physics**.

Examining Committee:

Oliver Stelzer-Chilton, Adjunct Professor, Physics & Astronomy, UBC
Co-Supervisor

Maximilian Swiatlowski, Associate Research Scientist, TRIUMF
Co-Supervisor

Colin Gay, Professor, Physics & Astronomy, UBC
Supervisory Committee Member

Christopher Hearty, Professor/IPP Scientist, Physics & Astronomy, UBC
University Examiner

Takamasa Momose, Professor, Chemistry, UBC
University Examiner

Toyoko Orimoto, Associate Professor, Physics, Northeastern University College of Science
External Examiner

Additional Supervisory Committee Members:

Gary Hinshaw, Professor, Physics & Astronomy, UBC
Supervisory Committee Member

Gordon Semenoff, Professor, Physics & Astronomy, UBC
Supervisory Committee Member

Abstract

The Standard Model of Particle Physics is the prevailing theory for describing the interactions of all observed fundamental particles and three of the four known fundamental interactions. However, despite its profound success, the Standard Model fails to explain some observations, such as dark matter and matter-antimatter asymmetry. Additionally, incorporating Einstein's theory of general relativity has proven difficult. Many proposed extensions to the Standard Model resolve these and other open questions through modifications of the Higgs potential or through new particles which interact with the Higgs boson.

This dissertation presents a search for Higgs boson pair production in the signal rich 4 b -jet final state using 136 fb^{-1} of data collected by the ATLAS detector at the Large Hadron Collider. This search consists of two separate analysis strategies targeting resonant and non-resonant Higgs boson pair production. The resonant search provides sensitivity to the resonant decay of new massive scalar and spin-2 particles predicted by many extensions to the Standard Model. The non-resonant search is directly sensitive to the Higgs potential via the Higgs boson self-coupling and the four-point Higgs to vector boson couplings.

No significant excesses are observed in the resonant search. World-leading upper limits are placed on cross-sections for new massive scalar and spin-2 resonances decaying to two Higgs bosons. These limits exclude a wide range of Kaluza-Klein graviton models and provide leading sensitivity amongst the ATLAS Higgs pair searches for masses above 600 GeV.

In the non-resonant search, no evidence for Standard Model Higgs production is found. An observed (expected) upper limit of 5.49 (6.73) times the Standard Model cross-section is set. Limits are placed on the Higgs self-coupling and the four-point Higgs to vector boson couplings, with observed (expected) constraints of $\kappa_\lambda \in [-4.7, 12.2]$ ($[-4.1, 10.4]$) and $\kappa_{2V} \in [0.02, 2.05]$ ($[0.01, 2.06]$), respectively. These results are then reinterpreted in the more general Higgs Effective Field Theory and Standard Model Effective Field Theory frameworks. The resulting limits are found to provide leading sensitivity within the HH channel for several parameters and benchmark signals.

Lay Summary

The Standard Model of Particle Physics is the prevailing theory for describing the interactions of all known fundamental particles and forces, besides gravity. However, despite its profound success, the Standard Model fails to explain observations of dark matter and the abundance of matter within the universe. Additionally, incorporating Einstein's theory of gravity into the Standard Model has proven difficult. The Higgs boson serves a central role in the Standard Model, providing masses to all known particles. Many extensions of the Standard Model predict modifications of the rate at which Higgs bosons should be produced in pairs. This dissertation analyzes the second dataset collected by the ATLAS detector at the LHC to search for this pair production. Although no evidence of this process is observed, strong constraints are placed on the possible solutions to these open questions, a substantial contribution to the advancement of our understanding of the universe.

Preface

This thesis is based upon research done with the LHC and ATLAS detector as part of the ATLAS collaboration. The ATLAS collaboration has over 3000 research scientists from 181 institutions worldwide, representing 41 countries. The work presented here would not have been possible without the contribution of all of these members. The members of the ATLAS collaboration contribute scientific expertise, computational resources, and their time to ensure the smooth and efficient operation of the ATLAS detector as well as the central reconstruction of much of this data.

Figures with no references in the caption have been produced by me, at times in collaboration with other members of the collaboration. Figures with “**ATLAS**”, “**ATLAS Simulation**”, “**ATLAS Preliminary**”, “**ATLAS Simulation Preliminary**”, or “**ATLAS Online**” labels are public results released by the ATLAS collaboration. Those figures without any label are not approved by the collaboration and are prepared by me. The results presented in this thesis, and summarized in Chapter 6, led to the publications listed below.

The search for resonant pair production of Higgs bosons in the $b\bar{b}b\bar{b}$ final state, in particular the resolved channel of this search, is presented in Chapter 6. This search is performed using 139 fb^{-1} proton-proton collision data collected with the ATLAS experiment at the LHC. The analysis team consists of approximately 50 members. I personally contributed significantly to the mid-to-late stages of this analysis. I performed extensive studies of the background systematic uncertainties for this analysis. In particular, along with S. Gasiorowski, I studied the bootstrap uncertainty approximation methodology. Along with L. Borgna and expanding on suggestions from R. Teixeira De Lima, I then performed background estimate validation studies in newly defined additional validation regions. I also validated the background estimate by showing closure through an alternative mapping of the reweighting procedure. Finally, I contributed significantly to studies of the signal contamination within the validation region. In particular, I investigated the option of addressing this signal contamination by redefining the signal region. In public results, this work was

combined with a related analysis which focused on the same production process, but in the boosted channel (see Chapter 6). These combined results led to a conference note, a paper, and to a combined conference note with the other ATLAS Higgs boson pair production channels.

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.* [ATLAS-CONF-2021-035](#), 2021.

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.

ATLAS Collaboration. *Combination of searches for non-resonant and resonant Higgs boson pair production in the $b\bar{b}\gamma\gamma$, $b\bar{b}\tau\tau$, and $b\bar{b}b\bar{b}$ decay channels using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector.* [ATLAS-CONF-2021-052](#), 2021.

The search for non-resonant pair production of Higgs bosons in the $b\bar{b}b\bar{b}$ final state with 139 fb^{-1} of data is also presented in Chapter 6. This search was performed by a group of around 48 people, and I was part of the core team of analyzers for this work. I was involved in the development of a large percentage of the central software for this analysis. My most substantial software contribution was the production of a central data handling package utilized by the other major analysis frameworks. Additionally, I contributed substantially to the central limit framework, which was developed largely by R. Hyneman and myself, with additional contributions from J. Grundy and others. I also helped to maintain and further develop existing software used for background estimation and production of plots. In addition to this software development, I performed extensive studies to optimize the analysis regions. In particular, along with R. Teixeira De Lima, I developed and validated the quadrant control regions. And, along with L. Borgna, N. Hartman, and Y. Zhang, I studied and validated the various categorization and signal region options for this analysis. My most substantial and independent contribution was the initial implementation and studies of the impacts of SMEFT on Higgs pair production (in general), the development of an signal initial reweighting method, and for the implementation and limit setting to perform fits to the corresponding SMEFT and HEFT parameters. Thanks also goes to C. Vernieri and J. Pearkes for the frequent and useful discussions. The SMEFT reweighting method was later updated to the method described in this thesis. This method was developed primarily by T.

Ingebret-Carlson, with additional support from myself, and building on the HEFT reweighting tool developed by L. Pereira Sánchez and C. Dimitriadi, and using the methodology of [1]. The HEFT to SMEFT translation was initially suggested as a solution for k-factors by myself, with eventual implementation by M. Belfkir and L. Pereira Sánchez using the derivation from [1, 2]. The HEFT implementation and limits within this analysis were also done by me. This implementation made substantial use of the work of L. Pereira Sánchez and C. Dimitriadi who developed the reweighting and performed initial studies. L. Pereira Sánchez also provided me guidance throughout this process. These results led to a conference note, and to a combined conference note with the other ATLAS single and double Higgs boson pair production channels. A paper is also currently under internal review. Note that these public (and pending public) results were updated from those included in this dissertation to include theoretical uncertainties and other minor changes.

ATLAS Collaboration. *Search for non-resonant pair production of Higgs bosons in the $b\bar{b}b\bar{b}$ final state in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector.* [ATLAS-CONF-2022-035](#), 2022.

ATLAS Collaboration. *Constraining the Higgs boson self-coupling from single- and double-Higgs production with the ATLAS detector* [ATLAS-CONF-2022-050](#), 2022.

Table of Contents

Abstract	iii
Lay Summary	iv
Preface	v
Table of Contents	viii
List of Tables	xii
List of Figures	xiv
Glossary	xxv
Acknowledgments	.xxviii
1 The Standard Model of Particle Physics	1
1.1 Particles, fields, and symmetries	1
1.2 Quantum Chromodynamics	3
1.2.1 Running of the coupling and non-perturbativity	4
1.2.2 Parton model	5
1.3 Electroweak theory	7
1.4 Higgs potential and electroweak symmetry breaking	8
1.5 Quantum anomalies	12
1.5.1 Sphalerons and instantons	13
1.6 Successes and limitations	15
2 Higgs Pair Production: Phenomenology and Beyond the Standard Model Physics	19

2.1	Electroweak baryogenesis and non-resonant Higgs boson pair production	20
2.1.1	Electroweak baryogenesis	23
2.2	Resonant Higgs pair production	26
2.2.1	Additional scalars and the Higgs portal	27
2.2.2	Two-Higgs-doublet models	28
2.2.3	Theories with extra dimensions	29
2.3	Effective field theory in Higgs boson pair production	30
2.3.1	Formulating an effective field theory	31
2.3.2	Standard Model Effective Field Theory	33
2.3.3	Higgs Effective Field Theory	35
3	The ATLAS Detector and the Large Hadron Collider	38
3.1	The Large Hadron Collider	38
3.1.1	Proton acceleration	39
3.1.2	Luminosity and pile-up	40
3.2	The ATLAS detector	41
3.2.1	Coordinate system	42
3.2.2	Inner Detector	44
3.2.3	Calorimeter System	47
3.2.4	Muon Spectrometer	50
3.2.5	Trigger and data acquisition	52
4	Simulation	55
4.1	Event generation	55
4.1.1	Hard process	57
4.1.2	Parton showers	57
4.1.3	Hadronization	58
4.1.4	Underlying event	58
4.1.5	Decays	59
4.2	ATLAS detector simulation	59
4.3	Fast detector simulations	60
5	Object Reconstruction	62
5.1	Tracks, vertices and topological-clusters	62
5.1.1	Tracks	63
5.1.2	Vertices	65

5.1.3	Topological-clusters	67
5.2	Jets	68
5.2.1	Particle flow objects	68
5.2.2	Anti- k_t jet clustering algorithm	69
5.2.3	Jet calibration	70
5.2.4	Jet-vertex-tagger	72
5.3	b -tagging	73
5.3.1	Algorithm inputs	73
5.3.2	Algorithms	76
5.3.3	Calibration	78
5.3.4	Working points	79
5.3.5	b -jet energy correction	79
6	Higgs Pair Production in the $HH \rightarrow b\bar{b}b\bar{b}$ Final-state	81
6.1	Dataset and simulation	83
6.1.1	Dataset	83
6.1.2	Monte Carlo samples	83
6.1.3	Signal reweighting	85
6.2	Trigger	88
6.2.1	Trigger buckets	89
6.2.2	Trigger scale factors	91
6.3	Object definitions	91
6.4	Event selection and reconstruction	92
6.4.1	b -tagging categories	94
6.4.2	Non-resonant ggF and VBF channels	95
6.4.3	Higgs candidate pairing	95
6.4.4	Background reduction	97
6.4.5	Kinematic region definitions	100
6.4.6	Signal acceptance and cut-flows	104
6.4.7	Invariant mass correction	104
6.5	Background estimation	106
6.5.1	Methodological primer	107
6.5.2	Neural network reweighting methodology	108
6.5.3	Reweighting performance	111
6.6	Analysis categorization and binning	113

6.6.1	Non-resonant categorization strategy	114
6.7	Uncertainties	116
6.7.1	Statistical uncertainties	118
6.7.2	Background systematics from alternative models	122
6.7.3	Background modeling non-closure systematic	126
6.7.4	Signal uncertainties	127
6.7.5	Summary of uncertainties	129
6.8	Background Validation	131
6.8.1	Reverse $ \Delta\eta_{HH} $ validation	131
6.8.2	Validation in shifted analysis regions	133
6.9	Statistical analysis and results	137
6.9.1	Resonant results	137
6.9.2	Non-resonant results	142
7	Conclusions and Outlook	149
	Bibliography	151
A	Neural network reweighting performance in additional variables	165
A.1	Resonant	165
A.2	Non-resonant	165
A.2.1	ggF	165
A.2.2	VBF	165

List of Tables

Table 1.1	Standard model fermions and their corresponding electric charges, weak isospins, and weak hypercharges. Doublet pairings for Q_L and L_L are also shown.	8
Table 5.1	Selection cut and c -, τ - and light jet rejections corresponding to different b -jet efficiency working points as determined from evaluation on $t\bar{t}$ events as described in Section 5.3.3. [137]	79
Table 6.1	The relevant b -tagging region definitions and corresponding purpose within the analysis. The 2b and 4b are used in both the resonant and non-resonant analyses. The 3b1f region is used only for the non-resonant analysis. Here b -jets are defined as those which Loose b -jets are defined as those which fail the 77% efficiency working point and pass the 85% working point while,	94
Table 6.2	Set of input variables used for the reweighting in the ggF and VBF channels. Note that the trigger bucket index was not used for the resonant analysis.	111
Table 6.3	Logarithmic binning parameters for the non-resonant ggF and VBF channels.	114
Table 6.4	Magnitudes of background uncertainties for the resonant analysis in the 4b Signal Region	129
Table 6.5	Magnitudes of background uncertainties for the non-resonant ggF channel in the 4b Signal Region.	130
Table 6.6	Magnitudes of background uncertainties for the non-resonant VBF channel in the 4b Signal Region.	131

Table 6.7	The observed and expected upper limit on the combined Standard Model (SM) ggF and VBF HH production cross-section $\sigma_{\text{ggF+VBF}}$ at the 95% confidence limit. Expected limits are also shown for the ± 1 and $\pm 2\sigma$. Limits are normalized by the predicted cross-section of SM ggF and VBF HH production.	142
Table 6.8	Limits on the Standard Model Effective Field Theory (SMEFT) Warsaw coefficients, at the 95% confidence level, obtained from fits to SR 4b data in the ggF channel. These limits are those set when allowing only the indicated coefficient to vary from its SM value. The c_H limits are obtained using the relation between κ_λ and c_H that exists at this order of calculation.	146
Table 6.9	List of benchmark signals from [152] and corresponding Higgs Effective Field Theory (HEFT) coefficients.	147

List of Figures

Figure 1.1	A chart showing the particles of the SM, as well as their spins, masses, (electric) charges, and colors. Fermion anti-particles are not included directly, but these anti-particles have the same mass and spin, but opposite electric and color charges. The loops in the background enclose the particles which interact with each given gauge boson. Adapted from [10] and [11].	2
Figure 1.2	Feynman diagrams showing the three vertices representing the three core interactions of Quantum Chromodynamics (QCD). From left to right, a four-point gluon vertex, a three-point gluon vertex and a quark-gluon vertex. The first vertex is proportional to g_s^2 while the last two are proportional to g_s	4
Figure 1.3	Parton momentum distributions for a proton, with $Q^2 = (100\text{GeV})^2$, from the MSHT20 Parton Distribution Function (PDF) set. The gluon distribution is divided by a factor of 10. Figure taken from [19].	6
Figure 1.4	The Higgs potential as a function of $ \Phi $ for $\mu^2 < 0$ (solid black) and $\mu^2 > 0$ (dashed red).	9
Figure 1.5	Feynman diagrams showing the vertices of EW theory after Electroweak Symmetry Breaking (EWSB), where γ is the photon, $X^\pm$ is any charged particle, m is any massive SM particle, and m_B is any massive SM boson. Comma separated labels indicate multiple possibilities, but these possibilities are matched across the same diagram. e.g. in the second diagram in the top row, the fermion lines can represent either $\ell^\mp$ and ν_ℓ or u_i and d_i , but not $\ell^\mp$ and u_i . Labels separated by a slash can be used interchangeably.	11

Figure 1.6	Feynman triangle diagram for first order $SU(2)_L$ and $U(1)_Y$ anomalies. The cross indicates insertion of the axial vector current $\gamma^\mu \gamma^5$. The loop includes all fermions which couple to the relevant external gauge fields.	12
Figure 1.7	A one-dimensional schematic showing the free energy of the electroweak gauge fields as a function of the Chern-Simons number. The height of the energy barrier for a sphaleron transition is shown as $\sim 10\text{TeV}$, corresponding to the value at $T = 0$.	13
Figure 1.8	A summary of the standard model cross-section predictions and corresponding ATLAS' experiment measurements for a series of processes [40].	16
Figure 2.1	Total production cross sections for HH production within the SM in each of the four main production channels.[69]	20
Figure 2.2	Leading diagrams for non-resonant gluon-gluon fusion (ggF) Higgs pair production.	21
Figure 2.3	The inclusive cross sections for ggF and VBF HH production, as a function of the Higgs self-coupling (κ_λ). The ggF cross section is calculated to Next-to-Next-to-Leading Order (NNLO) using approximate top mass effects (FTapprox) [70] and the uncertainties are taken from [71]. The VBF cross section is calculated to N3LO with uncertainties from the α_s and PDF variations [72]. The SM predictions for both channels are indicated with stars.	21
Figure 2.4	The differential cross section for HH production as a function of the invariant mass of the HH system, m_{HH} , for several values of the Higgs self-coupling, κ_λ . Figure taken from [69].	22
Figure 2.5	Leading diagrams for non-resonant vector boson fusion (VBF) HH production.	23
Figure 2.6	Higgs potential at temperatures around the critical temperature for (a) a SM-like smooth crossover and (b) a first order EW phase transition. Note that in (b) the high temperature minimum ($\phi = 0$) becomes meta-stable at the critical temperature.	24
Figure 2.7	Feynman diagrams for resonant Higgs pair production facilitated by a scalar (left) or Kaluza-Klein graviton (right).	27

Figure 2.8	Maximum branching ratio for $h_2 \rightarrow h_1 h_1$ as a function of the mass, m_2 , of the new scalar for a simple singlet scalar extension with no additional Z_2 symmetry. Figure adapted from [69], originally from [79].	28
Figure 2.9	Schematic of the anti-de Sitter space arising from the Randall-Sundrum metric with the Higgs field localized to the infrared (IR) brane and fermions and gauge bosons in the bulk. Here the exponential suppression is shown in blue and parameterized by w , which runs from 0 at the ultraviolet (UV) brane to $ \pi $ at the IR brane, with the curvature and length scale parameterized as k and r . Figure adapted from [84].	30
Figure 2.10	Some of the leading order HH diagrams in SMEFT with insertions shown for c_H (pink), c_{tH} (turquoise), c_{HG} (green), and c_{tG} (yellow). Insertions for the operator corresponding to $c_{H\Box}$ modifies all Higgs couplings and so, these modifications are not shown above. Figure adapted using a setup from [93].	34
Figure 2.11	Leading order diagrams in HEFT, with insertions shown for c_{hhh} (pink), c_t (turquoise), c_{ggh} (navy), c_{gggh} (green), and c_{tt} (yellow). Figure adapted from [93].	36
Figure 3.1	The portions of the CERN accelerator complex used to accelerate proton beams for eventual collision in the Large Hadron Collider (LHC). Figure adapted from [98].	39
Figure 3.2	Yearly distributions of the mean number of interactions per bunch crossing for data recorded by the ATLAS experiment at the LHC [99].	41
Figure 3.3	Yearly distributions of the mean number of interactions per bunch crossing for data recorded by the ATLAS experiment at the LHC [99].	42
Figure 3.4	Computer generated image of the ATLAS detector including descriptions of the major subsystems [100].	43
Figure 3.5	Diagram of particle paths within the ATLAS detector [101].	44
Figure 3.6	ATLAS coordinate system overlaid on the LHC ring. The yellow point marks the location of the ATLAS experiment. Inspired by figure from [102].	45
Figure 3.7	A computer generated image of the ATLAS Inner Detector, showing the Pixel Detector, Semiconductor Tracker, and Transition Radiation Tracker [103].	46

Figure 3.8	A computer generated image of the ATLAS Calorimeter, showing the LAr electromagnetic, LAr hadronic, and tile calorimeters [106].	47
Figure 3.9	A computer generated image of the ATLAS Muon sub-system, showing the location of each of the main components [109].	51
Figure 3.10	Distributions of total CPU time (over $\sim 400\,000$ processing units) used per event for two different High-Level Trigger (HLT) algorithms [113].	54
Figure 4.1	A schematic showing the various stages of event generation. The hard interaction is indicated by a dark red circle. Additional final-state (red) and initial state (blue) is shown. Parton showering is also shown in the many branchings of both initial and final state partons. Hadronization is shown as light green blobs, resulting in the conversion of a single parton into multiple particles. A secondary interaction is indicated by a purple oval, decay of short-lived states is shown by small red circles, and decay of particles produced during hadronization is shown in dark green circles. Finally, the emission of photons and electrons is shown in yellow. Figure taken from [117].	56
Figure 4.2	Comparison of ATLFast-II (red squares) and GEANT4 (yellow) jet p_T (left) and rapidity (right) distributions for jets with $p_T > 30\text{ GeV}$ and $ \eta < 2.8$. [127].	61
Figure 5.1	Illustrations of (a) several single particle clusters where particles have produced a signal in multiple pixels and (b) a merged cluster, where multiple particles have contributed to the same cluster of pixels. Particle trajectories are shown by the arrows, with corresponding energy deposits shown in the correspondingly colored cells. [114]	64
Figure 5.2	Average jet energy response as a function of reconstructed jet (a) detector η and (b) energy [133].	71
Figure 5.3	Distributions of (a) corrJVF, (b) R_{pT} , and (c) jet-vertex-tagger (JVT) for simulated pile-up and hard-scatter jets with $20 < p_T < 30\text{ GeV}$ and $ \eta < 2.4$. [135]	74
Figure 5.4	Fake rates from pile-up as a function of hard-scatter jet efficiency for JVF, corrJVF, R_{pT} , and JVT. JVT can be seen to give consistently better performance than the other discriminants. [135]	75

Figure 5.5	An illustration of an interaction producing two light flavored jets and a b -jet. The secondary vertex and corresponding displaced tracks are labeled, as are the transverse decay length (L_{xy}) and impact parameter (d_0), both of which are relevant for b -tagging. [115]	76
Figure 6.1	Branching ratio for Higgs pairs produced from collisions at LHC energy scales (13 TeV).	82
Figure 6.2	Comparison of m_{HH} distributions obtained for $\kappa_\lambda=10$ using (a) the truth level linear combination and truth level simulation, and (b) fully simulated $\kappa_\lambda=10$ and reweighted SM for the 2018 data period.	87
Figure 6.3	Comparison of m_{HH} distributions obtained for $\kappa_\lambda=10$ using fully simulated $\kappa_\lambda=10$ and SM reweighted to $c_H = -19.23$. Perfect agreement is not expected for the SMEFT reweighting, as it is derived at Leading Order (LO) while the full simulation $\kappa_\lambda = 10$ sample is simulated at Next-to-Leading Order (NLO).	88
Figure 6.4	Comparison of reconstructed m_{HH} distributions for fully simulated VBF samples and those constructed using the linear combination method for (a) $\kappa_{2V}=0$ and (b) $\kappa_\lambda=0$. Note that $\kappa_\lambda = 0$ is used here as a comparison because $\kappa_\lambda = 10$ is part of the basis and so the linear combination yields a perfect match in that case.	89
Figure 6.5	Trigger bucketing procedure for the resonant analysis.	90
Figure 6.6	Trigger bucketing procedure for the non-resonant analysis.	91
Figure 6.7	Event selection flowchart for the resonant analysis.	93
Figure 6.8	Event selection flowchart for the non-resonant analysis channels.	93
Figure 6.9	Mass planes of 2b data from the non-resonant (a) ggF channel, (b) VBF channel, and (c) the resonant analysis. A mass plane is defined as the two-dimensional distribution of Higgs Candidate (HC) masses. Here the leading HC mass (m_{H1}) and sub-leading HC mass (m_{H2}) are displayed on the horizontal and vertical axes, respectively.	98

Figure 6.10	Distribution of $ \Delta\eta_{HH} $ for 4b Monte Carlo (MC) signals and data for (a) the resonant analysis and (b) the non-resonant ggF channel. The ggF non-resonant distribution includes only that 4b data falling outside the signal region (thus, referred to as blinded in the caption). The resonant distribution includes the 4b data from the 2017 validation region (defined in Section 6.4.5). A vertical line indicating selection cut of $ \Delta\eta_{HH} < 1.5$ is shown on both figures. Events to the right of these lines are discarded.	99
Figure 6.11	Flowchart for the derivation of the top veto discriminant X_{Wt} . Note that for the resonant analysis the additional jet used to reconstruct the top candidate is not required to be a b -jet.	99
Figure 6.12	Distribution of X_{Wt} for $t\bar{t}$ and signal MC for the non-resonant (a) ggF and (b) VBF channels, and for (c) the resonant analysis. The non-resonant distributions include all simulated 2018 4b events passing prior cuts. The resonant analysis distributions include simulated events from the 2017 4b signal region (defined in Section 6.4.5). A vertical line indicating selection cut of $X_{Wt} > 1.5$ is shown on both figures. Events to the left of these lines are discarded.	100
Figure 6.13	Regions used for the non-resonant VBF and ggF channels drawn over the 2b data and 4b MC signal mass planes for each channel.	102
Figure 6.14	Regions used for the non-resonant VBF and ggF channels drawn over the 2b data and 4b SM MC signal mass planes for each channel.	104
Figure 6.15	Acceptance times efficiency for the (a) resonant analysis spin-0 and (b) spin-2 signal models as a function of the mass, and the (c) non-resonant signals and data in the ggF channel and (d) VBF channel.	105
Figure 6.16	Effects of the mass correction on the m_{HH} distributions of 280 GeV, 400 GeV, and 700 GeV scalar signal signal samples.	106
Figure 6.17	Pictorial diagram of the reweighting procedure, represented as a traditional ABCD method. The light turquoise squares indicate unblinded background-rich regions (CR 2b, CR 4b, and SR 2b) while the dark pink square indicates the blinded signal region (SR 4b). The black arrow indicates that a function is derived for the density ratio using the 2b and 4b data in a control region while the grey arrow indicates the application of this density ratio to 2b data to estimate 4b data.	108

Figure 6.18	Distributions of corrected m_{HH} (a) before and (b) after the CR derived reweighting for the resonant analysis in CR data.	112
Figure 6.19	Distributions of m_{HH} (a) before and (b) after the CR1 derived reweighting for the non-resonant ggF channel in CR1 data.	112
Figure 6.20	Distributions of m_{HH} (a) before and (b) after the CR1 derived reweighting for the non-resonant VBF channel in CR1 data.	113
Figure 6.21	SR4b (a) $ \Delta\eta_{HH} $ and (b) X_{HH} distributions for reweighted 2b background and scaled MC signal for SM and $\kappa_\lambda=10$ variations. The purple vertical lines represent the category cuts. Clear discrimination power can be seen from the varying ratios of signal and background in each sliced region of the distribution.	115
Figure 6.22	SR4b invariant mass distributions in each ggF category for reweighted 2b data and scaled MC signal for SM and $\kappa_\lambda=10$ variations. The lower panel shows the $S/\sqrt{B}$ for each bin with the signal scaled by a value, α , as indicated in the legend of the upper panel.	116
Figure 6.23	SR4b $ \Delta\eta_{HH} $ distributions for reweighted 2b background and scaled MC signal for SM, $\kappa_{2V}=0$, and $\kappa_\lambda=10$ variations. The purple vertical line represents the category cut. Clear discrimination power can be seen from the varying ratios of signal and background on either side of the line.	117
Figure 6.24	SR4b invariant mass distributions in each VBF $ \Delta\eta_{HH} $ category for reweighted 2b data and scaled MC signal for SM, $\kappa_{2V}=0$, and $\kappa_\lambda=10$ variations. The lower panels show the $S/\sqrt{B}$ for each bin.	117
Figure 6.25	Sequence of plots depicting the approximate construction of the bootstrap error band, with each component shown as a ratio to the nominal estimate. The light grey lines indicate the predictions of each neural network in the ensemble. Figure 6.25a shows the event-level IQR histogram (blue) constructed from $w_M \pm w_I$, which approximately matches the shape of the true variation in the individual estimates (grey). Figure 6.25b shows the normalization correction (blue) given by $\pm\alpha_I H(\alpha_* w_M)/2$ which approximately matches the magnitude of the variation in the individual estimates (grey). Figure 6.25c shows the final band (solid black) produce from these two components and Figure 6.25d shows a comparison of this band (solid black) to the true histogram-level IQR (dashed black).	121

Figure 6.26	Bootstrap uncertainties (pink) for the non-resonant (a) ggF and (b) VBF channels. These uncertainties correspond to the standard deviations of the 100 predicted histograms (light grey) obtained from the neural network ensemble.	122
Figure 6.27	Decomposition of the background shape uncertainty for the ggF resonant analysis for 2018 SR 4b.	123
Figure 6.28	Visualization of the four SR quadrants used to define the uncertainty decomposition.	123
Figure 6.29	Quadrant decomposition of the background shape uncertainty in the non-resonant ggF channel for 2018 data. Q_N , Q_W , Q_E , Q_S are defined from the four quadrants of the signal region shown in Figure 6.28. . . .	124
Figure 6.30	Quadrant decomposition of the background shape uncertainty in the non-resonant VBF channel for the high and low $ \Delta\eta_{HH} $ categories. Q_N , Q_W , Q_E , Q_S are defined from the four quadrants of the signal region shown in Figure 6.28.	125
Figure 6.31	Performance of the (a) ggF (2018 only) and (b) VBF (all years) reweighting model predictions vs SR 3b1f data. The VBF predictions are found to be in good agreement with the data while the ggF predictions exhibit a small non-closure.	127
Figure 6.32	Non-closure uncertainty (red) and ratio of SR 3b1f data to the reweighted 2b background estimate (black) in each category of 2018 data for the ggF non-resonant channel.	127
Figure 6.33	Performance of the background estimation procedure in the reversed $ \Delta\eta_{HH} $ region ($ \Delta\eta_{HH} > 1.5$) SR 4b 2017 data, for the resonant analysis. The quoted p -value uses the χ^2 test statistic. Based on this p -values, there is no evidence that the background estimate differs from the data. . . .	132
Figure 6.34	Performance of the background estimation procedure for the non-resonant ggF channel in the 2018 reversed $ \Delta\eta_{HH} $ region ($ \Delta\eta_{HH} > 1.5$) SR 4b for the discriminating variables (a) m_{HH} , (b) $ \Delta\eta_{HH} $, and (c) X_{HH} . Agreement is seen to be good in all cases, with all $\chi^2/ndf \lesssim 1.1$ and $p \gtrsim 0.3$	133

Figure 6.35	Shifted regions overlaid on the ggF HC mass plane. The solid red curve in the center represents the true SR. The shifted regions are labeled in clockwise order: upper left (pink), upper center (blue), upper right (green), center right (light-blue), lower right (orange). The shifted region definitions are identical for the VBF and ggF analysis channels and the underlying ggF data is used here only for the sake of reference.	134
Figure 6.36	Reweighted 2018 SR 2b, and SR 4b data m_{HH} distributions for the shifted regions in the ggF channel. The upper regions and the center right region use CR1 derived weights, but the lower right region uses CR2 derived weights. The total uncertainty in the reweighted 2b estimate includes all background uncertainties except the 3b1f non-closure uncertainty. Good agreement is seen across the 5 regions.	135
Figure 6.37	Reweighted SR 2b and SR 4b data m_{HH} distributions for the shifted regions in VBF channel. The upper regions and the center right region use CR1 derived weights, but the lower right region uses CR2 derived weights. The total uncertainty in the reweighted 2b estimate includes all background uncertainties. Good agreement is seen across the 5 regions.	136
Figure 6.38	Results of a background only profile likelihood fit to the SR 4b m_{HH} spectrum for the resonant analysis. Scalar signals with masses of 280 GeV (black), 400 GeV (red), 600 GeV (blue), and 1000 GeV (orange) are overlaid.	138
Figure 6.39	Pulls and impacts for the background uncertainties for four representative scalar signals. Several H_T systematics are pulled to $\sim +1\sigma$, indicating that the VR model better modeled the data in this case. This is somewhat expected as the VR is closer to the SR and should provide a better model.	139
Figure 6.40	Upper limits on the cross-section for the resonant scalar signal model. Contributions from the boosted channel are included in the main expected lines. Those contributions arising only from the resolved channel detailed here are shown in pink.	140

Figure 6.41	Upper limits on the cross-section for the resonant spin-2 signal model. Contributions from the boosted channel are included in the main expected lines. Those contributions arising only from the resolved channel detailed here are shown in pink. RS Gravitons with $k/\bar{M}_{\text{Pl}} = 1$ under the analysis signal model are excluded by the resolved analysis between 300 and 1400 GeV.	141
Figure 6.42	Results of a background only profile-likelihood fit to SR 4b data in the non-resonant ggF channel for each year. Signal variations with $\kappa_\lambda=1$ (red) and 10 (turquoise), $c_{tH}=6$, and HEFT Benchmark 7 are overlaid.	143
Figure 6.43	Results of a background-only profile-likelihood fit to SR 4b data in the non-resonant VBF channel. Signal variations with $\kappa_\lambda=1$ (red), $\kappa_\lambda=10$ (turquoise), and $\kappa_{2V}=0$ (black) are overlaid.	144
Figure 6.44	Limits on the HH cross-section as function of the Higgs boson self-coupling obtained using the combination of ggF and VBF channels. Observed (expected) limits are $\kappa_\lambda \in [-4.7, 12.2]$ ($[-4.1, 10.4]$).	144
Figure 6.45	Limits on the HH cross-section as function of the HHVV coupling obtained using the combination of ggF and VBF channels. Observed (expected) limits are $\kappa_{2V} \in [0.02, 2.05]$ ($[0.01, 2.06]$).	145
Figure 6.46	One-dimensional limits on (a) c_{tH} (b) c_{tG} , (c) $c_{H\Box}$, and (d) c_{HG} from the non-resonant ggF channel only. At this order of calculation, c_H is directly related to κ_λ via linear transformation.	147
Figure 6.47	Upper cross-section limits on the 7 HEFT benchmarks $BM1 - 7$ at the 95% confidence level. The observed (expected) limits are included as text on the plot, in units of fb. The limit on the SM ggF cross-section is also included for reference.	148
Figure A.1	(Training variables) Distributions of ΔR between the closest Higgs Candidate jets, ΔR between the other two, and X_{Wt} before and after CR 1 derived reweighting for the 2016 Control Region 1.	166
Figure A.2	(Non-training variables) Distributions of p_T of the 1st and 3rd leading Higgs Candidate jets before and after CR 1 derived reweighting for the 2016 Control Region 1.	167
Figure A.3	(Non-training variables) Distributions of $ \Delta\eta_{HH} $, X_{HH} , and of m_{HH} before and after CR 1 derived reweighting for the 2016 Control Region 1.	168

Figure A.4	(Training variables) Distributions of ΔR between the closest Higgs Candidate jets, ΔR between the other two, and X_{W_t} before and after CR 1 derived reweighting for the 2017 Control Region 1.	169
Figure A.5	(Non-training variables) Distributions of p_T of the 1st and 3rd leading Higgs Candidate jets before and after CR 1 derived reweighting for the 2017 Control Region 1.	170
Figure A.6	(Non-training variables) Distributions of $ \Delta\eta_{HH} $, X_{HH} , and of m_{HH} before and after CR 1 derived reweighting for the 2017 Control Region 1.	171
Figure A.7	(Training variables) Distributions of ΔR between the closest Higgs Candidate jets, ΔR between the other two, and X_{W_t} before and after CR 1 derived reweighting for the 2018 Control Region 1.	172
Figure A.8	(Non-training variables) Distributions of p_T of the 1st and 3rd leading Higgs Candidate jets before and after CR 1 derived reweighting for the 2018 Control Region 1.	173
Figure A.9	(Non-training variables) Distributions of $ \Delta\eta_{HH} $, X_{HH} , and of m_{HH} before and after CR 1 derived reweighting for the 2018 Control Region 1.	174
Figure A.10	(Training variables) Distributions of the minimum and maximum di-jet masses, and the top veto variable, X_{W_t} , before and after CR 1 derived reweighting for the all years inclusive Control Region 1.	175
Figure A.11	(Non-training variables) Distributions of ΔR between the other two HCs and the number of jets before and after CR 1 derived reweighting for the all years inclusive Control Region 1.	176
Figure A.12	Distributions of $ \Delta\eta_{HH} $, the p_T of the HH system ($p_{T,HH}$), and m_{HH} before and after CR 1 derived reweighting for the all years inclusive Control Region 1.	177

Glossary

AdS	anti-de Sitter space
AFII	ATLFAST-II
BDT	Boosted Decision Tree
BSM	Beyond the Standard Model
CKM	Cabibbo–Kobayashi–Maskawa
CP	charge conjugation and parity symmetry
CSC	Cathode Strip Chamber
CTP	Central Trigger Processor
DGLAP	Dokshitzer–Gribov–Lipatov–Altarelli–Parisi
ECAL	Electromagnetic Calorimeter
EFT	Effective Field Theory
EM	Electromagnetic
EW	Electroweak
EWBG	Electroweak Baryogenesis
EWSB	Electroweak Symmetry Breaking
FCAL	Forward Calorimeter
FSR	final-state radiation
ggF	gluon-gluon fusion
GUT	Grand Unified Theory
HC	Higgs Candidate
HCAL	Hadronic Calorimeter

HEFT	Higgs Effective Field Theory
HLT	High-Level Trigger
HSV	hard-scatter vertex
ID	Inner Detector
IR	infrared
ISR	initial-state radiation
JES	jet energy scale
JVT	jet-vertex-tagger
KK	Kaluza-Klein
LHC	Large Hadron Collider
LO	Leading Order
L1	Level-1
MC	Monte Carlo
MDT	Muon Drift Tube
MS	Muon Spectrometer
MSSM	minimal supersymmetric extension of the Standard Model
NLO	Next-to-Leading Order
NN	Neural Network
NNLO	Next-to-Next-to-Leading Order
N3LO	Next-to-Next-to-Next-Leading Order
PDF	Parton Distribution Function
QCD	Quantum Chromodynamics
QED	Quantum Electrodynamics
RDO	Raw Data Object
RNN	recurrent neural network
RPC	Resistive Plate Chamber
RS	Randall-Sundrum
SCT	Semiconductor Tracker

SM	Standard Model
SMEFT	Standard Model Effective Field Theory
SPD	Silicon Pixel Detector
SSVF	single secondary vertex finding algorithm
SUSY	Supersymmetry
TDAQ	Trigger and Data Acquisition
TGC	Thin Gap Chamber
TRT	Transition Radiation Tracker
UV	ultraviolet
VBF	vector boson fusion

Acknowledgments

When I first enrolled at UBC, I insisted I was only staying for the MSc. Needless to say, things did not go that way, but I don't regret it in the least.

Indeed, it was thanks to the people I met at UBC, TRIUMF, and CERN that I was convinced to extend my stay. The support I received from these outstanding communities cannot be overstated. First and foremost, I would like to thank my supervisors: Oliver Stelzer-Chilton and Maximilian Swiatlowski. Some people are lucky enough to have one great supervisor, but I was even luckier: I had two.

Oliver, I thank you for taking me on as your student and for reassuring me that I had the background necessary to fit in with the group. I also thank you for your constant efforts to work with me as I followed somewhat non-traditional routes, both within my early academic career and on my way to and from CERN. Your efforts to strengthen the sense of community within the Vancouver ATLAS group did not go unnoticed and some of my most memorable moments at UBC were the events you hosted. It was also thanks to you, that I was able to begin to work with Max and to find an analysis that truly suited me.

Max, thank you for the time and effort you put in to help me establish myself within the always hectic HH4b analysis. I appreciate the effort you put in to build that foundation in the first few weeks I had at CERN and also that you maintained that same effort even over the many months of large time-zone differences. Thank you for holding me accountable, but also for being understanding of my illogical and erratic work schedule. The long Skype discussions were a surprising highlight, but it was there that I learned you were always willing to chat or to explain some obscure ATLAS concept (clearly too!). Most of all, thank you for the constant support and mentorship; I wouldn't have accomplished half as much without it.

I would also like to thank my supervisory committee: Colin Gay, Gary Hinshaw, and Gordon Semenoff for the useful discussions and thought provoking questions. It was these questions that encouraged me to learn many of the things I have written about in this dis-

sertation. More broadly, thank you to all of the ATLAS Collaboration, without whom, this work would not have been possible. In particular, I would like to thank the HH4b analysis team members, as well as the broader HH and HH combination teams. While I still refuse to type my titles with double H's, I really appreciated the commitment and the community building. And of course, the HH color schemes and logos are the best. I also acknowledge the financial support I received from the Natural Sciences and Engineering Research Council of Canada, which supported me in these continued studies.

Finally, thank you to the UBC and TRIUMF ATLAS groups, and in particular, to my fellow grad students (and post-docs) in particular: Steffen, Alex, Sébastien, Vincent, Elham, Félix, Robin, purple haired Robin, Dominique, Zhengcheng, Evan, Juan, Colm, Russell, and Ryan. I will always be grateful for the way you made the office feel like home, even before it was. The list of reasons is both too long to remember and too memorable to forget: long discussions about anything and everything, camping trips, desk assemblies, pumpkin carving, baked goods, Glühwein nights, ...

Chapter 1

The Standard Model of Particle Physics

The Standard Model of Particle Physics (SM) is the prevailing theory for describing the interactions of all observed fundamental particles and three of the four known fundamental interactions. Within the field of subatomic physics, this theory has achieved unparalleled success. Its predictions have been confirmed repeatedly and with remarkable precision, including fruitful predictions of previously undiscovered particles [3–6]. This chapter begins with a brief outline of the constituent fields and relevant symmetries of the Standard Model (SM). It then describes quantum chromodynamics and electroweak theory, the two major components of the SM. Although the Higgs field may be considered a component of the electroweak theory, due to the importance of the Higgs boson to this thesis, a dedicated section is allocated to the discussion of the Higgs field, its potential, and the Electroweak Symmetry Breaking (EWSB) it induces. Finally, this chapter concludes with a summary of the successes and limitations of the SM.

1.1 Particles, fields, and symmetries

The SM [7–9] is a relativistic quantum field theory described by a Lagrangian which satisfies a global Lorentz symmetry and translational symmetry, as well as a product of local gauge symmetries. The Lorentz symmetry covers the standard rotational symmetry familiar from Newtonian mechanics, plus the additional boost symmetry of special relativity. Together the global translational and Lorentz symmetries form a global Poincaré symmetry. The elementary particles of the SM each have a spin which labels the irreducible representation of the Lorentz group under which that particle transforms. Notably, due to the spin-statistics




Figure 1.1: A chart showing the particles of the SM, as well as their spins, masses, (electric) charges, and colors. Fermion anti-particles are not included directly, but these anti-particles have the same mass and spin, but opposite electric and color charges. The loops in the background enclose the particles which interact with each given gauge boson. Adapted from [10] and [11].

theorem, this spin also dictates whether the particles follow Fermi-Dirac or Bose-Einstein statistics, with half-integer spins following the former and integer spins following the latter. Particles also have a chirality which dictates whether the particle transforms under the left- or right-handed representation of the Poincaré group.

The local gauge symmetry of the SM, $SU(3) \times SU(2) \times U(1)$, defines the interactions of the theory. In addition to its spin and chirality, each SM particle has a charge under each component of the local gauge symmetry. The force mediators are vectors in the adjoint representation of each of these local gauge symmetries. The $SU(3)$ component maps to the theory of quantum chromodynamics (QCD), which governs the strong force. Particles which interact via the strong force carry a charge known as “color”, while objects which do not are said to be “colorless”. The $SU(2) \times U(1)$ component maps to electroweak theory (EW) which governs electromagnetism and the weak force. The $SU(2)$ charge is weak isospin while the $U(1)$ charge is weak hypercharge. Following the breaking of this symmetry, detailed in Section 1.4, a mixture of these two charges forms the electric charge. Particles which do not interact via electromagnetism are said to be (electrically) neutral.

The SM consists of three kinds of elementary particles: quarks, leptons, the force mediators, and the Higgs boson [12, 13]. These particles are the excitations of the corresponding

fields within the field theory. The charges and masses of all of these particles are shown in Figure 1.1 [14]. Each listed particle also has an anti-particle with the same mass, but with opposite chirality and which carries the opposite charges under the gauge symmetries [12, 13]. Throughout the rest of this thesis, particles and their corresponding anti-particles are referred to by the name of the particle, unless a distinction is needed. For example, “up quark” is used to refer to both up and anti-up quarks unless otherwise specified.

The quarks and leptons are fermions, with spin $\frac{1}{2}$. These particles are divided into three generations of increasing mass, shown as columns in Figure 1.1. All quarks and leptons, except neutrinos, occur with both left and right-handed chirality, both carrying the same electric and color charges. Only left-handed chiral neutrinos exist in the SM. All of these particles interact via the electroweak force, though neutrinos do not have an electric charge and right-handed leptons and left-handed anti-leptons do not interact via the weak force. The quarks also carry a color charge under $SU(3)$. This charge can be expressed as a $\mathbb{C}^3$ vector, with basis components commonly labeled as “Red”, “Blue”, and “Green”.

The force mediators and the Higgs boson are all bosons. The Higgs boson is a scalar boson with spin 0, it is a by-product of EWSB and shall be discussed further in Section 1.4. The force-mediators are vector bosons with spin 1 and form the adjoint representations of the symmetry for which they mediate the interaction. These force mediators will be further discussed in the following sections, but include the gluons which mediate the strong force, the $W^\pm$ and Z bosons which mediate the weak force, and the photon which mediates electromagnetism.

1.2 Quantum Chromodynamics

Quantum Chromodynamics (QCD) [15–17] is the $SU(3)$ gauge field theory describing the strong interaction. The force carriers of this theory are the massless gluon fields, which are denoted as G^A where A runs from 1 to 8. These gluons are in the adjoint representation of $SU(3)$, meaning they are represented as the generators of the Lie group. They therefore carry two color indices each, i.e. they may also be written G_{ij}^A where i and j run over the three colors. These indices are typically suppressed, unless for some reason it is useful to show the contraction of two fields over their color indices explicitly. The field strength tensor of the gluon fields is

$$G_{\mu\nu}^A = \partial_\mu G_\nu^A - \partial_\nu G_\mu^A - g_s f_{ABC} G_\mu^B G_\nu^C, \quad (1.1)$$




Figure 1.2: Feynman diagrams showing the three vertices representing the three core interactions of QCD. From left to right, a four-point gluon vertex, a three-point gluon vertex and a quark-gluon vertex. The first vertex is proportional to g_s^2 while the last two are proportional to g_s .

where g_s is the strong coupling, and f_{ABC} are the SU(3) structure constants, which are anti-symmetric in their indices. The corresponding gauge covariant derivative is [12]

$$D_\mu = \partial_\mu + ig_s T_A G_\mu^A, \quad (1.2)$$

where the $T_A \equiv \lambda_A/2$ where λ_A are the eight Gell-Mann matrices. This covariant derivative is used to describe the propagation of color-charged material through the gluon fields. The quark fields are denoted as q_r^i with r indexing the 6 quark flavors and i indexing the color. By writing out the kinetic term for the quarks as well as the energy of the field strength tensor for the gluon fields, we arrive at the full QCD component of the SM Lagrangian:

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4} G_{\mu\nu}^A G_A^{\mu\nu} + i\bar{q}_r^i \gamma^\mu (D_\mu)_{ij} q_r^j. \quad (1.3)$$

It should be noted that in pure QCD, one can also add a quark mass term as $-i\bar{q}_r^i m \delta_{ij} \bar{q}_r^j$ where δ_{ij} is the Kronecker delta; however, in the SM this term would violate gauge invariance due to the flavor changing currents of SU(2). The mass term in the SM theory is instead obtained via EWSB. Figure 1.2 shows Feynman diagrams for the three kinds of vertices present in QCD, namely the three- and four-point gluon vertices and the quark-gluon vertex.

1.2.1 Running of the coupling and non-perturbativity

In quantum field theory, the interaction of force mediators with the vacuum can lead to screening and anti-screening effects which cause the effective coupling of the theory to vary at different length/energy scales. This variation of the effective coupling with scale is known as the running of the coupling and is well described by the process of renormalization. In QCD, both screening and anti-screening effects occur due to the self-interaction of the gluon. As a consequence of these effects, the effective coupling strength decreases asymptotically

with increasing energy or decreasing distance. This phenomenon, known as asymptotic freedom allows the perturbative calculation of QCD processes at high energies. Contrastingly, at large distances/low energies the interaction of gluon propagators with gluons from the vacuum leads to an amplification of the effective coupling due to anti-screening. This increased effective coupling means that QCD dynamics at low energies/large distances are non-perturbative. One notable non-perturbative effect is color confinement, in which isolated color-charged objects are forbidden. For this reason, quarks and gluons are observed only as the constituents of color-neutral composite particles, known as hadrons.

1.2.2 Parton model

While known to be composed of quarks and gluons, the precise substructure of protons and other hadrons is difficult to model due to the non-perturbativity of QCD at small scales. Thus, to understand scattering processes involving one or more such hadrons, a phenomenological model must be used. The parton model [18] is one such model which describes hadrons as an infinite sea of partons (quarks and gluons) that together carry the momentum of the hadron. The momentum fraction carried by a given type of parton is described by a corresponding Parton Distribution Function (PDF). These sets of PDFs are unique to each type of hadron and are determined experimentally.. A set of proton PDFs are shown in Figure 1.3 for interactions with a resolution scale of $Q^2 = (100 \text{ GeV})^2$. All PDFs, $f_i(x)$, satisfy the relation:

$$1 = \sum_i \int_0^1 dx x f_i(x) \quad (1.4)$$

where i indexes the quark, anti-quark and gluon PDFs and x is the momentum fraction. These PDFs also allow the description of the net quark content of a hadron, which is obtained from the following relation:

$$N_q = \int_0^1 dx (f_q(x) - f_{\bar{q}}(x)). \quad (1.5)$$

This net quark content, is used to describe the constituent quarks of a hadron, with hadrons described as having N_q valence q -quarks. For example, the proton has $N_u = 2$, $N_d = 1$, and all other $N_q = 0$, and so the quark is often said to be composed of two valence up quarks and a down quark.

The parton model is particularly useful for describing hadron collisions as it allows the factorization of the cross-sections. For example, the inclusive cross-section for a process




Figure 1.3: Parton momentum distributions for a proton, with $Q^2 = (100 \text{ GeV})^2$, from the MSHT20 PDF set. The gluon distribution is divided by a factor of 10. Figure taken from [19].

$p p \rightarrow X$ can be written

$$\sigma = \sum_{i,j} \int dx_i dx_j d\Phi f_i(x_i) f_j(x_j) \frac{d\hat{\sigma}_{ij \rightarrow X}}{d\Phi} \quad (1.6)$$

where the sum is over all possible pairs of constituent partons (i and j) of the incoming protons, Φ is the phase space element, and $\hat{\sigma}_{ij \rightarrow X}$ is the corresponding differential parton level cross-section. In this case, the PDFs are the proton PDFs. To do this, a factorization scale, μ_F , is defined which separates the hard-scatter perturbative physics from the soft non-perturbative parton level physics described by the PDFs. Physics occurring at energies above this scale is treated perturbatively via the differential cross-section, while physics below this scale is incorporated as a normalization component via the PDFs. Since the factorization scale alters the physics included in the PDFs, they also have a dependence on μ_F and so are sometimes written as $f_i(x_i, \mu_F)$. This dependence is described by a set of renormalization group equations, known as the Dokshitzer–Gribov–Lipatov–Altarelli–Parisi (DGLAP) evolution equations, which are obtained by absorbing the collinear infrared sensitivity of the parton model into a running parton density [20, 21]. The resulting equations describe the energy evolution in terms of parton splitting functions which describe the processes: $g \rightarrow gg$, $g \rightarrow qq$, and $q \rightarrow qg$. The DGLAP equations play a crucial role in describing the evolution

of parton showers, as will be described further in Section 4.1.2

1.3 Electroweak theory

Electroweak (EW) theory [7–9] provides a unified description of two SM interactions: electromagnetism and the weak interaction. EW theory is a non-abelian chiral gauge theory with $SU(2)_L \times U(1)_Y$ gauge symmetries. The subscript L is used to indicate the left-chirality of the $SU(2)$ gauge interaction, while the subscript Y denotes weak hypercharge and is used to distinguish this symmetry from the maximal unbroken $U(1)_Q$ symmetry obtained after EWSB. Particles which interact via the $SU(2)$ gauge interaction are said to carry weak isospin, T , and are labeled by the third component of this isospin, T_3 . Those particles which interact via the $U(1)$ gauge interaction carry hypercharge Y .

The force carriers of the $SU(2)_L$ and $U(1)_Y$ theories are the W^i bosons ($i = 1, 2, 3$) and B boson, respectively. Like the gluons of QCD, these bosons are all massless and are in the adjoint representation of the corresponding symmetry and are singlets under other gauge symmetries. The W^i bosons all have $T = 1$, with each taking a different value of T_3 ($-1, 0, +1$). The B boson carries hypercharge 0.

$$W_{\mu\nu}^i = \partial_\mu W_\nu^i - \partial_\nu W_\mu^i - g\epsilon_{jk}^i W_\mu^j W_\nu^k B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu \quad (1.7)$$

where ϵ_{jk}^i is the anti-symmetric Levi-Civita tensor which arises here from the structure constants of $SU(2)$. The corresponding covariant derivative is

$$D_\mu = \partial_\mu + ig'YB_\mu + i\frac{g}{2}t_L^i W_\mu^i \quad (1.8)$$

where g' is the coupling strength of the B boson, Y is the hypercharge, and t_L^i are the $SU(2)_L$ generators. These generators are 0 for the trivial representation of particles without weak isospin and proportional to the Pauli matrices for those with weak isospin. As the weak interaction is chiral, the left-handed fermions and right-handed anti-fermions form doublets under $SU(2)_L$ while the right-handed fermions and left-handed anti-fermions form singlets. These singlets and doublets are summarized in Table 1.1 along with the weak hypercharge, weak isospin, and electric charge of each particle.

Fermion Type	Field	Q	T_3	Y
Quarks	$Q_L = \begin{pmatrix} u_L \\ d_L \end{pmatrix}$	$\frac{2}{3}$	$\frac{1}{2}$	$\frac{1}{6}$
		$-\frac{1}{3}$	$-\frac{1}{2}$	$\frac{1}{6}$
	u_R	$\frac{2}{3}$	0	$\frac{2}{3}$
	d_R	$-\frac{1}{3}$	0	$-\frac{1}{3}$
Leptons	$L_L = \begin{pmatrix} \nu_L \\ e_L \end{pmatrix}$	0	$\frac{1}{2}$	$\frac{1}{6}$
		-1	$-\frac{1}{2}$	$\frac{1}{6}$
	e_R	-1	0	-1

Table 1.1: Standard model fermions and their corresponding electric charges, weak isospins, and weak hypercharges. Doublet pairings for Q_L and L_L are also shown.

Writing out the full electroweak portion of the Lagrangian one obtains the following:

$$\begin{aligned} \mathcal{L}_{\text{EW}} = & -\frac{1}{4}B_{\mu\nu}B^{\mu\nu} - \frac{1}{4}W_{\mu\nu}^i W^{i\mu\nu} \\ & + \bar{Q}_L \gamma^\mu D_\mu Q_L + \bar{L}_L \gamma^\mu D_\mu L_L + \bar{u}_R \gamma^\mu D_\mu u_R + \bar{d}_R \gamma^\mu D_\mu d_R + \bar{e}_R \gamma^\mu D_\mu e_R \end{aligned} \quad (1.9)$$

As discussed briefly in the previous section, it is due to this chiral nature of the SU(2) interaction that it is impossible to write down gauge invariant Dirac mass terms, since the left-handed and right-handed fields carry different charges.

1.4 Higgs potential and electroweak symmetry breaking

The final component of the SM is the Higgs field [22–24], a complex scalar field which is a singlet under SU(3), a doublet under SU(2)_L, and which has a hypercharge of 1/2. This field may be written as

$$\Phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}, \quad (1.10)$$

where ϕ^+ and ϕ^0 are labeled as such due to their charges under the as of yet to be realized U(1)_Q symmetry. The Higgs field also has a corresponding conjugate field given by $\tilde{\Phi} \equiv i\sigma^2 \Phi$, where σ^2 is the second Pauli matrix. This field experiences the Higgs potential:

$$V(\Phi) = -\mu^2 \Phi^\dagger \Phi + \lambda \left(\Phi^\dagger \Phi \right)^2, \quad (1.11)$$




Figure 1.4: The Higgs potential as a function of $|\Phi|$ for $\mu^2 < 0$ (solid black) and $\mu^2 > 0$ (dashed red).

where λ and μ are parameters of the Higgs potential. Note that to ensure the potential is bounded from below and has a well defined vacuum state, it must be that $\lambda \geq 0$ if $\mu^2 > 0$ or $\lambda > 0$ otherwise. This potential, plus a Higgs field kinetic term, and a series of interactions with the quarks and leptons form the Higgs component of the SM Lagrangian:

$$\begin{aligned} \mathcal{L}_{\text{Higgs}} = & (D_\mu \Phi)^\dagger (D^\mu \Phi) + \mu^2 \Phi^\dagger \Phi - \lambda (\Phi^\dagger \Phi)^2 \\ & - \sum_{i,j} \left(y_{ij}^u \bar{Q}_{Li} \tilde{\Phi} u_{Rj} + y_{ij}^d \bar{Q}_{Li} \Phi d_{Rj} + y_{ij}^l \bar{L}_{Lj} \Phi e_{Rj} + h.c. \right) \end{aligned} \quad (1.12)$$

where the indices i and j run from 1 to 3, and y_{ij} are the Yukawa couplings.

To fully understand the implications of this last addition, the vacuum structure must be determined. For $\mu^2 > 0$ and $\lambda > 0$, the potential has a minimum at zero; however, for the SM case with $\mu^2 < 0$ and $\lambda > 0$, the Higgs potential forms the so-called Mexican hat potential, shown in Figure 1.4. The minimum of this potential is clearly not zero; in fact, it occurs at

$$|\Phi|^2 = -\frac{\mu^2}{2\lambda},$$

which defines the vacuum expectation value (vev)

$$v \equiv \sqrt{\frac{\mu^2}{\lambda}}. \quad (1.13)$$

This ground-state is degenerate, occurring for all phases. To parameterize it we write the Higgs field as

$$\Phi = \exp(i t_L^i \xi^i / v) \begin{pmatrix} 0 \\ \frac{v+h}{\sqrt{2}} \end{pmatrix}, \quad (1.14)$$

where ξ^i and h are real scalar fields. In hyperspherical coordinates, the ξ fields can be thought of as parameterizing the Higgs 3 angular components while the h field corresponds to the radial component or Higgs boson. This parameterization yields a new form for the Higgs potential:

$$V(h) = \frac{1}{2} m_H^2 h^2 - \frac{m_H^2}{2v} h^3 - \frac{m_H^2}{8v^2} h^4 + \dots \quad (1.15)$$

parameterized now in terms of the measurable Higgs boson mass $m_H \equiv \sqrt{-2\mu^2} = \sqrt{2\lambda}v$ and v . Making a similar substitution into the Higgs kinetic term and making appropriate substitutions we find that

$$(D_\mu \Phi)^2 = \frac{1}{2} (\partial_\mu h)^2 + \frac{g^2 v^2}{4} W_\mu^+ W^{-\mu} \left(1 + \frac{h}{v}\right)^2 + \frac{g^2 v^2}{8 \cos^2 \theta_w} Z_\mu Z^\mu \left(1 + \frac{h}{v}\right)^2 \quad (1.16)$$

where $\theta_w \equiv \arctan(g'/g)$ is the weak mixing angle and the new mass eigenstates are

$$\begin{aligned} A_\mu &\equiv \sin \theta_w W_\mu^3 + \cos \theta_w B_\mu, \\ Z_\mu &\equiv \cos \theta_w W_\mu^3 - \sin \theta_w B_\mu, \\ W_\mu^\pm &\equiv \frac{W_\mu^1 \mp i W_\mu^2}{\sqrt{2}}. \end{aligned} \quad (1.17)$$

These eigenstates are the familiar photon, Z boson, and $W^\pm$ bosons, with masses $m_W = gv/2$, $m_Z = gv/2 \cos \theta_w$, and $m_A = 0$. The massless A_μ field is the photon field, which is the gauge boson of the maximal unbroken subgroup $U(1)_Q$ which we associate with Quantum Electrodynamics (QED). Making a similar substitution into the Yukawa coupling terms we obtain:

$$\mathcal{L}_{\text{Yukawa}} = - \sum_{i,j} \left(y_{ij}^u \bar{Q}_{Li} u_{Rj} + y_{ij}^d \bar{Q}_{Li} d_{Rj} + y_{ij}^l \bar{L}_{Lj} e_{Rj} + h.c. \right) \left(\frac{v+h}{\sqrt{2}} \right) \quad (1.18)$$

where the fermions now have mass matrices given by $v y_{ij}$. These mass matrices are not necessarily diagonal, but can be diagonalized by a set of unitary matrices, $V_f^{L/R}$. Finally, propagating the new boson and fermion mass eigenstates back into the EW fermion interac-




Figure 1.5: Feynman diagrams showing the vertices of EW theory after EWSB, where γ is the photon, $X^\pm$ is any charged particle, m is any massive SM particle, and m_B is any massive SM boson. Comma separated labels indicate multiple possibilities, but these possibilities are matched across the same diagram. e.g. in the second diagram in the top row, the fermion lines can represent either $\ell^\mp$ and ν_ℓ or u_i and d_i , but not $\ell^\mp$ and u_i . Labels separated by a slash can be used interchangeably.

tion terms in Equation 1.9 we find the following charged and neutral current interactions:

$$\begin{aligned} \mathcal{L}_{\text{EW,int}} = & -\frac{g}{\sqrt{2}} \left(\bar{u}_L^i V_{CKM} W_\mu^+ \gamma^\mu d_L^i - \bar{\nu}_L^i W_\mu^+ \gamma^\mu e_L^i + h.c. \right) \\ & - \bar{\psi}^i \left(g s_w A_\mu \gamma^\mu \left(t_L^3 P_L + Y \right) + \frac{g}{c_w} Z_\mu \gamma^\mu \left(c_w^2 t_L^3 P_L - s_w^2 Y \right) \right) \psi^i \end{aligned} \quad (1.19)$$

where ψ^i is used as a shorthand for all fermion terms, $V_{CKM} \equiv V_u^{L\dagger} V_d^L$ is the Cabibbo–Kobayashi–Maskawa (CKM) matrix, P_L is the left-chiral projection operator, $c_w \equiv \cos \theta_w$, and $s_w \equiv \sin \theta_w$. Note that the unitary matrix for the neutrinos is taken to be the same as for the corresponding charged lepton, which we are free to do since the neutrinos are massless in the SM. From this equation we find the relation between the electric charge of a fermion and its weak hypercharge and weak isospin, namely:

$$Q = Y + T_3 \quad (1.20)$$

The CKM matrix in Equation 1.19 describes the rates of quark flavor mixing within the SM. The relevant interaction vertices for electroweak theory after EWSB are shown in Figure 1.5.




Figure 1.6: Feynman triangle diagram for first order $SU(2)_L$ and $U(1)_Y$ anomalies. The cross indicates insertion of the axial vector current $\gamma^\mu \gamma^5$. The loop includes all fermions which couple to the relevant external gauge fields.

1.5 Quantum anomalies

Beyond the gauge symmetries from which it was constructed, the SM appears to contain several accidental global symmetries. Most notable among these are baryon number ($U(1)_B$) and lepton number ($U(1)_L$), where quarks have $B = \frac{1}{3}$ and $L = 0$ and leptons have $B = 0$ and $L = 1$. While these two symmetries exist in the tree-level Lagrangian, they are violated in the full quantum field theory. The non-conservation of these classical symmetries in the quantum theory is known as a quantum anomaly [13, 25]. These quantum anomalies arise because they are non-chiral, while the underlying electroweak gauge theory is chiral. The chirality of electroweak theory means that the contributions of the left- and right-handed fermions to the $U(1)_B$ and $U(1)_L$ currents are imbalanced when interacting with the electroweak gauge fields. To show this, we begin by writing out the current corresponding to $U(1)_B$:

$$j_B^\mu = \sum_i B_i \bar{\psi}_i \gamma^\mu \psi_i = \frac{1}{3} \delta_{ij} (\bar{Q}_{Li} \gamma^\mu Q_{Lj} + \bar{u}_{Ri} \gamma^\mu u_{Rj} + \bar{d}_{Ri} \gamma^\mu d_{Rj}) \quad (1.21)$$

where i and j index the fermion generations. Computing the expectation value of the divergence of this current to second order in perturbation theory, we find two relevant anomalies for this current both of which may be represented by the triangle diagram in Figure 1.6. These anomalies can be written in the usual way [13, 25] as $U(1)_B U(1)_Y^2$ and $U(1)_B SU(2)_L^2$. The total divergence of the current from these anomalies is found to be [25]:

$$\langle \partial_\mu j_B^\mu \rangle = \frac{3}{64\pi^2} \left(g^2 \epsilon_{\mu\nu\alpha\beta} B^{\mu\nu} B^{\alpha\beta} - g'^2 \epsilon_{\mu\nu\alpha\beta} W^{i\mu\nu} W^{i\alpha\beta} \right) \quad (1.22)$$

As this divergence is not strictly zero, the current associated with baryon number is not conserved and baryon number within an isolated system can change with time. The non-




Figure 1.7: A one-dimensional schematic showing the free energy of the electroweak gauge fields as a function of the Chern-Simons number. The height of the energy barrier for a sphaleron transition is shown as $\sim 10 \text{ TeV}$, corresponding to the value at $T = 0$.

conservation of lepton number proceeds similarly. The relevant anomalies for the $U(1)_L$ current:

$$j_L^\mu = \delta_{ij} (\bar{L}_{Li} \gamma^\mu L_{Lj} + \bar{e}_{Ri} \gamma^\mu e_{Rj}) \quad (1.23)$$

are $U(1)_L U(1)_Y^2$, $U(1)_L SU(2)_L^2$, and the gravitational anomaly $U(1)_L (\text{grav})^2$. Under linearized gravity, the resulting total divergence is found to be [25, 26]:

$$\partial_\mu j_L^\mu = \frac{3}{64\pi^2} \left(g^2 \epsilon_{\mu\nu\alpha\beta} B^{\mu\nu} B^{\alpha\beta} - g'^2 \epsilon_{\mu\nu\alpha\beta} W^{i\mu\nu} W^{i\alpha\beta} + \frac{1}{12\pi^2} \epsilon_{\mu\nu\alpha\beta} R^{\mu\nu\sigma\tau} R_{\sigma\tau}^{\alpha\beta} \right) \quad (1.24)$$

where R is the Riemann curvature tensor [27] from the gravitational contribution. Of note, one then has that the symmetry $U(1)$ symmetry defined by $B - L$ is conserved, so long as gravitational effects are negligible.

1.5.1 Sphalerons and instantons

While baryon number and lepton number are violated, the violation occurs through a total divergence term and so does not appear in perturbation theory. To find processes which violate these global symmetries, one must instead search for those non-perturbative solutions which are gauge equivalent to 0, but topologically stable [25]. Two such class of solutions are the sphalerons and instantons. These solutions correspond to transitions between the periodic degenerate vacuum configurations of the electroweak gauge fields, shown in Figure 1.7. These vacuum configurations are topologically distinct and stable states which can be characterized by their Chern-Simons winding number, N_{CS} . Changes in N_{CS} can be

written as [28]:

$$\Delta N_{CS} = \frac{1}{64\pi^2} \int_0^t dt' \int d^3x \epsilon_{\mu\nu\alpha\beta} W^{i\mu\nu} W^{i\alpha\beta}. \quad (1.25)$$

Comparing this equation to the integral of the second term in Equation 1.22 and Equation 1.24 we see that for SU(2) induced B and L violation:

$$\Delta B = \Delta L = 3\Delta N_{CS} \quad (1.26)$$

The stable minima of the gauge fields correspond to integer values of N_{CS} , and so transitions between these minima will produce 3 net baryons and 3 net leptons. Thus, sphalerons and instantons both violate $U(1)_B$ and $U(1)_L$, but preserve $U(1)_{B-L}$.

As shown in Figure 1.7, these various minima are separated by an energy barrier. Instanton solutions correspond to the process of tunneling through this barrier, while sphaleron solutions correspond to passing over the barrier. The instanton tunneling probability in electroweak theory is negligibly small at $\sigma_{inst} \lesssim 10^{-170}$ [29] and so the following paragraphs shall focus only on sphalerons.

The minimum energy required for a sphaleron transition to occur is equal to the height of the barrier between the minima of Figure 1.7. The height of this barrier is temperature dependent and can be approximated as [30–32]

$$E_{sph}(T) \cong \frac{8\pi v(T)}{g} f\left(\frac{m_H}{m_W}\right), \quad (1.27)$$

where $f(m_H/m_W)$ is a function of order unity and $v(T)$ is the Higgs vev. This energy leads to a sphaleron transition rate of [33, 34]:

$$\Gamma_{sph} \cong A \left(\frac{g^2}{4\pi}\right)^4 \left(\frac{E_{sph}(T)}{T}\right)^7 e^{-E_{sph}(T)/T} \quad (1.28)$$

where A is a constant. At the low temperatures of the modern universe, the sphaleron energy is $\sim 10\text{TeV}$ and the transition rate is negligible. This large energy and the suppression of the process due to the necessity of a coherent state, makes this process unlikely to be observed at current or future high energy colliders [35]. However, for higher temperatures such as those at and prior to EWSB, the sphaleron rate is much higher and transitions readily occur. Of additional note, for $T \gg m_W$, the Higgs vev vanishes and Equation 1.28 fails. At this point the rate is no longer exponentially suppressed and instead scales as $\Gamma \propto T^4$ [36, 37].

1.6 Successes and limitations

The SM is an extremely successful theory which has been tested and repeatedly verified over the span of several decades. It has correctly predicted several new particles, including the Higgs boson, and the relations between its internal parameters have proven internally consistent. However, there are also known issues and outstanding questions, which shall be briefly outlined in this section.

To understand some of these problems, it is important to reference back to the free parameters of the SM. In particular, the SM consists of a total of 19 free parameters:

- 10 particle masses: 1 Higgs boson mass and 9 fermion masses (6 quark and 3 charged lepton), with the latter being dictated by the Yukawa couplings
- 1 vacuum expectation value (of the Higgs field)
- 3 gauge coupling constants
- 3 mixing angles and 1 phase of the CKM matrix
- θ_{QCD} , which was not covered in the previous sections, but enters if one includes the CP violating anti-symmetric gluon field strength term.

Notably, while it is known that neutrinos have masses and undergo flavor mixing, these effects are not encoded in the SM. However, this is primarily because the origin of these effects are still unclear. In fact, it is straightforward to extend the SM to include these effects through the addition of a neutrino equivalent of the CKM matrix, known as the PMNS matrix and by introducing 3 additional Yukawa interactions [38, 39]. Doing so would lead to a total of 26 free parameters and would successfully explain most observations seen to date, including neutrino fluctuations. The agreement of several experimental measurements with the predictions of the SM are shown in Figure 1.8. Despite these profound successes; however, there are still several open questions surrounding the SM.

First, there are the theoretical challenges. Finding a complete theory of physics that combines Einstein's theory of general relativity with the SM has proven difficult. There are several attempts to do so, including loop quantum gravity theories [41], string theories [42], and supergravity theories [43]; however, to date these theories still face conceptual challenges.

The SM also fails to provide explanations for the wide-varieties of scales appearing in the problem. One such example is the large gap between the EW scale ($m_H \sim 10^2 \text{ GeV}$)




Figure 1.8: A summary of the standard model cross-section predictions and corresponding ATLAS' experiment measurements for a series of processes [40].

and the Planck scale of gravity ($M_{pl} \sim 10^{19}$ GeV). This large gap in scales leads to a fine-tuning of the Higgs boson mass. The mass of the Higgs boson receives corrections from fermion loops, which are of the order of the UV cutoff to the theory. This means that if the SM were extended to the Planck scale, these corrections would be expected to be of order M_{pl} . For the Higgs boson mass to take the much lower value m_H , there must be a precise cancellation of the SM loop corrections. There are numerous theories attempting to explain this cancellation. Most notable among these are Supersymmetry (SUSY) [44], theories with large extra dimensions [45, 46], and theories with additional scalar fields such as two Higgs doublet models [47], and composite Higgs models [48].

Another open theoretical question is the problem of vacuum stability within the SM. While the Higgs potential is stable at low energies, the Higgs self-coupling is driven to negative values at large energies due to the running of the coupling. This leads to a minimum with a lower potential energy than the current Higgs vev. This lower minimum could lead to spontaneous vacuum decay [49]. In the SM, the timescale of this vacuum decay is longer

than the current age of the universe; however, this could be altered by updated measurements of the top Yukawa coupling or by modifications to the Higgs potential itself [50]. Such modifications exist in many extensions of the SM [51].

There are also problems driven by experimental results from various fields of physics. First, the values of the neutrino masses are not well measured. At present, only mass splittings and upper thresholds (see Figure 1.1) have been measured [52]. Additionally, as already mentioned, the origin of these masses is still unclear. Besides the aforementioned PMNS and Yukawa method for introduction in which neutrinos are Dirac fermions, it is also possible that they could be Majorana fermions meaning they are equivalent to their anti-particles [53]. This would violate the accidental global lepton symmetry within the SM. Experiments are ongoing to confirm this experimentally, through the observation of neutrinoless double β -decay [54–56].

Another, potentially related, open question is the inability of the SM to explain the so-called dark matter seen in astronomical observations. Dark matter is weakly- or non-interacting non-baryonic matter which is observed through its gravitational effects on astronomical objects, most famously, the rotational curves of galaxies [57]. There have been several attempts to explain dark matter through theories of modified gravity as well as many theories which propose particle dark matter candidates as extensions to the SM. These SM extensions include such things as the addition of sterile neutrinos [58] or of a so-called dark-sector which is only connected to the known SM particles through a weakly interacting portal [59].

Finally, the excess of baryons relative to anti-baryons within the universe (baryon asymmetry) is unexplained by the SM. The direct production of such a baryon asymmetry via baryogenesis would require a series of conditions, known as the Sakharov conditions, to be satisfied [60]. The three necessary conditions are Baryon number violation, CP violation, and interactions out of thermal equilibrium. Baryon number violation occurs in the SM at very high energies, such as those found in the early universe, via the non-perturbative sphaleron process. CP violation also occurs in the SM via the weak interaction; however, additional CP violation may be required in some theories [61–65]. Finally, there is no known mechanism for interactions to be brought out of thermal equilibrium in the SM. However, there are several theories which introduce the necessary mechanisms, including Grand Unified Theory (GUT) baryogenesis and Electroweak Baryogenesis (EWBG) [66]. GUT baryogenesis [67, 68] provides for the production of excess baryons via the decay of heavy baryons during or shortly after the grand unification epoch, in which all three interactions of the SM were unified. EWBG [51] relies on modifications to the Higgs potential

which produce a strongly first-order phase transition during EWSB.

As history has shown, advances in physics are most often achieved through the direct study of the failings of past theories, whether theoretically or experimentally. Thus, just as the Standard Model forms the basis of particle physics, these open questions together form the foundation guiding further study within the field.

Chapter 2

Higgs Pair Production: Phenomenology and Beyond the Standard Model Physics

The Standard Model is an incredibly successful theory, but also an incomplete description of nature with several open questions. As the newest, least studied, and most-interacting constituent of the SM, the Higgs boson provides a natural window for further examination of this theory's weaknesses. Detailed characterization of the Higgs boson will provide strong constraints of the Beyond the Standard Model (BSM) theories, which attempt to extend the SM while resolving open questions. While many of the Higgs boson couplings have now been relatively well measured, the Higgs self-coupling remains largely unprobed. A measurement of this self-coupling would provide crucial information about the shape of the Higgs potential and could hint at the answers to many open questions. Higgs boson pair (HH) production provides a unique opportunity for the direct measurement of this coupling. In fact, HH production is highly relevant to many BSM theories. Many of the simplest well-motivated extensions to the SM modify the SM Higgs potential or introduce particles which decay favorably to pairs of Higgs bosons.

The rest of this chapter will provide a brief introduction to these theories and to their effects on HH production. In particular, this chapter will cover Electroweak Baryogenesis, real singlet scalar extensions, large extra dimensions, and two-Higgs-doublet models. This chapter will then conclude with an introduction to effective field theories, with an emphasis on their use as a tool for providing a more model-independent parameterization of BSM physics. A more expansive overview of most of these topics can be found in [69].




Figure 2.1: Total production cross sections for HH production within the SM in each of the four main production channels.[69]

2.1 Electroweak baryogenesis and non-resonant Higgs boson pair production

SM HH production occurs non-resonantly through four main production channels: gluon-gluon fusion (ggF), vector boson fusion (VBF), double Higgs-strahlung (WHH/ZHH), and double Higgs bremsstrahlung off top quarks (ttHH/tjHH). The total production cross-section for each of these channels is shown in Figure 2.1. Due to their larger cross sections and sensitivity to several parameters of interest, this thesis will focus only on ggF and VBF production. These channels provide excellent sensitivity to measurements of the Higgs self-coupling, which is highly relevant to many BSM theories, most notably to those predicting electroweak baryogenesis Section 2.2. VBF also provides sensitivity to the gauge structure of the Higgs sector.

Gluon-gluon fusion

Gluon-gluon fusion is the leading HH production channel with a SM rate over an order of magnitude larger than the other channels. At leading order, ggF HH production is facilitated by two diagrams: a box diagram and a triangle diagram Figure 2.2.

The Higgs self-coupling vertex is labeled and parameterized here by $\kappa_\lambda = \lambda_3/\lambda_{SM}$, that




Figure 2.2: Leading diagrams for non-resonant ggF Higgs pair production.




Figure 2.3: The inclusive cross sections for ggF and VBF HH production, as a function of the Higgs self-coupling (κ_λ). The ggF cross section is calculated to Next-to-Next-to-Leading Order (NNLO) using approximate top mass effects (FTapprox) [70] and the uncertainties are taken from [71]. The VBF cross section is calculated to N3LO with uncertainties from the α_s and PDF variations [72]. The SM predictions for both channels are indicated with stars.

is, the ratio of the true Higgs self-coupling to the SM prediction. The modification of this coupling can lead to amplification of the cross-section (Figure 2.3), but also substantial shape changes due to the destructive interference of these two diagrams. In particular, very large or small values of κ_λ will lead to a shift in the invariant mass distribution (m_{HH}) towards lower values, as shown in Figure 2.4. Thus, the Higgs self-coupling can be measured both through the HH cross-section and the shape of the m_{HH} distribution.




Figure 2.4: The differential cross section for HH production as a function of the invariant mass of the HH system, m_{HH} , for several values of the Higgs self-coupling, κ_λ . Figure taken from [69].

Vector boson fusion

Vector boson fusion is sensitive to the Higgs self-coupling, as seen from the Leading Order (LO) diagrams shown in Figure 2.5. It can provide some sensitivity to the measurement of κ_λ ; however it has a much lower cross-section than ggF (Figure 2.3) and shows comparatively less sensitivity to this parameter. However, VBF provides access to additional parameters of interest: the HHVV and HVV couplings. These couplings appear in the second and third diagrams in Figure 2.5 as κ_V (HVV) and κ_{2V} (HHVV) parameterized as ratio of the corresponding coupling to its SM value. Individually, these diagrams produce amplitudes which scale with the center of mass energy. More broadly, when summing these amplitudes together one has:

$$\mathcal{A} \sim \frac{\hat{s}}{v^2} (\kappa_{2V} - \kappa_V^2) + \dots \quad (2.1)$$

where $\hat{s}$ is the parton center of mass energy. In the SM and in BSM theories where the Higgs field is an SU(2) doublet, $(\kappa_{2V} - \kappa_V^2) = 0$; however, the κ_{2V} coupling is currently unconstrained by data with VBF HH production providing the only direct probe [69]. If $\kappa_{2V} \neq \kappa_V^2$, the combined amplitude would continue to scale with the energy of the collision. This violates perturbative unitarity bounds and requires unitarization by new states at some higher energy cut-off scale [73]. The unique sensitivity of VBF HH production to these




Figure 2.5: Leading diagrams for non-resonant VBF HH production.


effects make it a powerful probe of the gauge structure of the Higgs sector and of possible new high energy physics.

2.1.1 Electroweak baryogenesis

Electroweak baryogenesis provides a solution to the baryon asymmetry through the production of excess baryons at the electroweak phase transition. Three Sakharov conditions [60] must be met to allow for electroweak baryogenesis:

1. There must be a baryon number violating process: such a process is required to introduce the initial baryon asymmetry
2. There must be C and CP violation: otherwise the rate of a process that produces excess baryons must be balanced by the complementary process which produce excess anti-baryons
3. There must be non-equilibrium dynamics to shift the baryon number violating processes out of equilibrium in order to preserve the baryon asymmetry.

Baryon number violating processes are provided already in the SM, via the non-perturbative process known as sphalerons which convert three baryons (leptons) to three anti-leptons (anti-baryons) or vice-versa. These sphalerons are highly suppressed by the Higgs vev [74] and so do not appear with a measurable rate post-EWSB. C and CP violation are also provided within the SM, via the CKM phases. The SM effects may be too small in the SM [74], but there are numerous BSM models which provide the necessary additional CP violation, as will be discussed in the following sections. Finally, non-equilibrium dynamics are not provided in the SM, and so would need to be introduced. Such non-equilibrium dynamics occur naturally in many BSM theories via a strongly first-order electroweak phase transition [51].



(a) Smooth EW crossover (SM-like).

(b) First order EW phase transition.

Figure 2.6: Higgs potential at temperatures around the critical temperature for (a) a SM-like smooth crossover and (b) a first order EW phase transition. Note that in (b) the high temperature minimum ($\phi = 0$) becomes meta-stable at the critical temperature.

Electroweak phase transition

Due to the self-interaction of the Higgs field, as well as its interaction with other fields, the state which minimizes the energy of the Higgs field may not match with the potential in the Lagrangian. One may account for these effects by writing out the so-called effective potential, which accounts for these interactions. In doing so, one inherently introduces a dependence of the Higgs potential on the temperature of the surrounding bath of particles [74]. This leads to the following finite temperature effective Higgs potential at one-loop:

$$V(\phi, T) \cong D(T^2 - T_0^2)\phi^2 - ET\phi^3 + \frac{\lambda}{4}\phi^4 \quad (2.2)$$

where λ is a function which is slowly varying in T , D , E , and T_0 are constants. For the SM, these constants depend on the vev, and the masses of the top, Higgs boson, W and Z . The full SM form of these constants and of λ can be found in [74]. This potential clearly evolves with changing temperature, and the precise evolution depends on the constants and λ function.

During the early universe, the temperature was very high and the Higgs field was immersed in a hot bath of radiation making these temperature effects highly relevant. As the universe expanded and cooled, the Higgs potential evolved according to the temperature dependence of Equation 2.2. For the SM Higgs potential at zero temperature, the minimum is the normal vev, while at high temperatures, the minimum is zero due to the flip in sign and increasing contribution of the ϕ^2 . The transition between these temperatures is smooth

in the SM case [74], as shown in Figure 2.6a; however, this needn't be the case. To protect against non-perturbative effects at large temperatures, one can also add an additional term to the effective potential which accounts for the thermal mass corrections in the propagators via daisy resummation [32, 75]. This term depends on the masses of scalars and longitudinal vectors within the theory and the addition of such particles could lead to modifications of the cubic coefficient E . It is also possible for this coefficient to be modified via the addition of additional scalars which take a vev and mix with the Higgs field, leading to tree-level cubic terms [76]. Regardless of the source, such modifications could have profound consequences. For non-zero E , there is a first-order phase transition at EWSB. In this case, the high-temperature minimum becomes meta-stable at some critical temperature, T_c , and remains so for some range of temperature beyond this, as shown in Figure 2.6b. During this period of meta-stability, the Higgs field can tunnel to the true minimum within certain regions of the universe. This tunneling leads to the spontaneous production of bubbles in this lower energy, higher vev phase. If these bubbles are of sufficient size, such that the pressure induced by the surplus free energy within the bubble exceeds the surface tension, the bubble will begin to expand until all of local space-time has been converted to the lower energy post-EWSB phase [32, 74]. One of these expanded bubbles of the broken phase would then correspond to the current state of the observable universe.

CP violation and baryon asymmetry

A net baryon number may be produced during the electroweak phase transition through a combination of scattering, diffusion, and transfer reactions at the bubble walls [74]. Since sphaleron rates are suppressed by a large Higgs vev, rates are much lower in the broken phase inside the bubbles. This means that if a baryon asymmetry can be produced outside the bubble wall, it will be preserved inside. To place approximate bounds on the required first-order phase transition parameters, one can relate the energy required for a sphaleron to the temperature. At temperatures which are much larger than the sphaleron energy, sphalerons will proceed freely, while at energies much lower, this will not be true. As a first order approximation, knowing that the sphaleron energy is directly related to the Higgs vev, we can use ϕ_c/T_c to characterize the expected sphaleron rate in the broken phase at the start of the phase transition. Here ϕ_c is the minimum of the Higgs potential at the critical temperature. This ratio is often referred to as the strength of the transition. A strongly first order phase transition with $\phi_c/T_c > 1$ is required to ensure that the baryon asymmetry is not washed out by sphalerons in the broken phase [32].

Now, to produce the baryon asymmetry in the first place, we know that we must first

produce an excess of left-handed¹ leptons relative to their anti-particles which can then be converted to an excess baryon number by the sphalerons. This excess can be produced by CP-violating interactions at the bubble wall which proceed via a similar mechanism to the transmission and reflection familiar from quantum mechanics [32]. Once the excess is produced, it is rapidly converted to a baryon excess which must be rapidly swept up by the bubbles before it is washed away by the induced chemical potential for sphaleron processes. Thus, a net baryon number is left in the broken phase following the electroweak phase transition. This baryon number would remain today and would explain the observed baryon asymmetry [51].

While this mechanism is rather complex, the two components missing in the SM, additional CP-violating processes and a strongly first order phase transition, are common in many BSM theories [51]. In particular, such modifications are common in the two-Higgs-doublet models and simple scalar extensions which will be discussed in the following section. The measurement of HH production provides a prime avenue for investigating this phenomenon through the measurement of the critical Higgs field self-interaction term, which is directly linked to the cubic term that drives the first-order phase transition. Additionally, HH production provides sensitivity to a commonly theorized source of extra CP-violation, in which additional heavy scalars mix with the SM Higgs field [69]. Thus, HH production is a crucial channel for better understanding the Higgs potential, EWBG, and the baryon asymmetry problem.

2.2 Resonant Higgs pair production

The Higgs self-coupling takes a fixed and well-measured value within the SM, modifications of this value without modification of well-measured SM quantities would result in the violation of renormalizability and gauge invariance. Thus, such modifications would necessitate the addition of new BSM physics to restore these important attributes of the SM. A natural way to modify this value is through the addition of new couplings to the Higgs boson. These couplings would likely be to massive new particles, which could decay resonantly to pairs of Higgs bosons, as shown in Figure 2.7. This resonant HH production would lead to a sharp resonant peak in the invariant mass spectrum of the HH system.

¹Left-handed is all that is relevant here since sphalerons are chiral




Figure 2.7: Feynman diagrams for resonant Higgs pair production facilitated by a scalar (left) or Kaluza-Klein graviton (right).

2.2.1 Additional scalars and the Higgs portal

The simplest extension to the SM, common to many BSM theories, is the addition of one or more real singlet scalar fields. Such scalar fields can be used to connect the SM, via the Higgs sector, to a so-called hidden sector. These hidden sectors contain additional particles, which do not interact directly with the SM. This is a common proposal for solving the dark matter problem and is known as a Higgs portal. These scalar field extensions can also be used to drive inflation by providing additional flexibility to the scalar potentials during the early universe.

A benchmark model is provided by the most general renormalizable Lagrangian that includes the addition of a single such scalar field:

$$\mathcal{L}_{\text{BSM}} = \frac{a_1}{2} \Phi^\dagger \Phi S + \frac{a_2}{2} \Phi^\dagger \Phi S^2 + b_1 S + \frac{b_2}{2} S^2 + \frac{b_3}{3} S^3 + \frac{b_4}{4} S^4 \quad (2.3)$$

where S is the new scalar [69, 77, 78]. This scalar generically also takes some vacuum expectation value, v_s , with $S = v_s + s$. EWSB leads to a series of quadratic cross terms between the Higgs boson, h , and the scalar, s . Both particles are also singlets under the remaining gauge symmetry $\text{SU}(3)_C \times \text{U}(1)_Q$. There is then clearly mixing between these states, and h and s are not the mass eigenstates. Following the typical convention, we can label the mass eigenstates $h_{1,2}$ with corresponding masses $m_{1,2}$ with $m_2 \geq m_1$. Since only the Higgs boson originally interacted with the other SM particles, all SM couplings of the mass eigenstates would be suppressed by a factor related to the mass mixing angle and the production and decay rates would be suppressed by this factor squared. One of these mass eigenstates must be the observed Higgs boson, and so must have a mass of 125 GeV. Since no new scalar field has been observed, it is likely that if such a model exists it either has a substantially larger mass than the Higgs boson, it is degenerate with the Higgs boson, or the mixing angle is very small. The degenerate case allows for the modification of the self-coupling discussed in the previous section, but will not be discussed further here. In the




Figure 2.8: Maximum branching ratio for $h_2 \rightarrow h_1 h_1$ as a function of the mass, m_2 , of the new scalar for a simple singlet scalar extension with no additional Z_2 symmetry. Figure adapted from [69], originally from [79].

case that the unobserved eigenstate has a mass $m_2 \geq 2m_H$, it would allow on-shell decays to pairs of Higgs bosons. As shown in Figure 2.8, given no additional symmetries for the new scalar, the branching ratio for HH production ($h_2 \rightarrow h_1 h_1$) could be as high as 90% [79]. This large branching ratio, along with the straightforward resonant nature of the decay makes HH production an ideal location to search for these scalar extensions.

2.2.2 Two-Higgs-doublet models

Another set of useful benchmark models are the two-Higgs-doublet models (2HDMs). 2HDMs are distinguished from the SM by the presence of an additional SU(2) doublets composed of scalar fields with the same quantum numbers as the SM Higgs field. There are many motivations for such theories, first of all, by introducing additional vacuum structure they provide additional flexibility to the Higgs potential, easing constraints on EWBG and false vacuum decay [80]. 2HDM also provide additional CP-violation through the CP-odd Higgs boson produced during EWSB as a linear combination of the imaginary components of the neutral Higgs fields [81].

2HDM also share a close relation to theories of supersymmetry (SUSY) [82], where they are added to provide additional symmetry breaking effects, similarly to in the SM. SUSY explains the hierarchy problem by extending the Poincaré symmetry with fermionic generators that introduce a new superpartner for each SM particle. This new symmetry would

require that superpartners differ from their SM partner only by a half integer difference in spin. That is, the superpartners of bosons are fermions and vice-versa. Since fermions and bosons provide loop corrections of opposite sign to the Higgs mass, the corrections from these superpartners precisely cancel those of the SM particles, resolving the hierarchy problem. Of course, if such particles existed with the same masses they would have already been discovered. To resolve this, SUSY theories include supersymmetry that is broken softly, meaning broken only by terms with a coupling of positive mass dimension, m_{soft} . These additional SUSY violating terms lead to mass splittings of order m_{soft} between superpartners and a similarly sized correction to the Higgs mass [82]. In the minimal supersymmetric extension of the Standard Model (MSSM), and many other SUSY theories, this soft breaking is accomplished through an additional scalar doublet field [83].

Regardless of the origin of the 2HDM, whether it be a direct scalar extension or a by-product of SUSY, the general phenomenology will be similar. The 2HDM will produce 5 mixed scalars after EWSB: two neutral CP-even h_0/H_0 , two charged CP-even $H^\pm$, and one neutral CP-odd A^0 . The lighter neutral CP-even state, h_0 , is typically associated with the SM Higgs boson. Similarly to the singlet scalar extension, these mass mixings lead to both modifications of the potential (self-coupling) and to enhanced HH production via the resonant decay of the heavy neutral H^0 [47]. Again, the $H^0 \rightarrow h_0 h_0$ branching would be enhanced for a large range of mass mixing angles.

2.2.3 Theories with extra dimensions

The electroweak hierarchy problem can also be solved by taking the electroweak and gravitational interactions to be the same, while introducing a mechanism which dilutes the effective strength of gravity. This dilution is possible through the introduction of additional space-like dimensions. One promising option is the addition of a finite warped extra dimension, described by the Randall-Sundrum (RS) metric [84]. This metric is a solution to Einstein's equations in 5 dimensions with a bulk cosmological constant and provides a curvature to the space along the extra dimension. The resulting spacetime, an anti-de Sitter space (AdS), is shown in Figure 2.9, where the two boundaries are identified as the high-energy ultraviolet (UV) brane and low-energy infrared (IR) brane respectively. In such a theory, dimensionful quantities become exponentially suppressed as they move along this new dimension from the UV brane to the IR brane, with the precise rate determined by the curvature. This effective suppression provides a useful tool for solving the electroweak hierarchy problem and for explaining the various energy scales within the SM itself. Given a sufficient curvature and distance between the branes, this scaling can suppress the Planck




Figure 2.9: Schematic of the anti-de Sitter space arising from the Randall-Sundrum metric with the Higgs field localized to the IR brane and fermions and gauge bosons in the bulk. Here the exponential suppression is shown in blue and parameterized by w , which runs from 0 at the UV brane to $|\pi|$ at the IR brane, with the curvature and length scale parameterized as k and r . Figure adapted from [84].

scale (or a similar scale) to be the size of the electroweak scale at the IR brane. Thus, by localizing the Higgs field to the IR brane, the electroweak hierarchy can be resolved. Additionally, the large differences in particle masses can be explained by localizing fermions to different regions of the bulk.

These models also lead to new physics effects measurable at the LHC, notably, in such theories, fluctuations of the metric lead to towers of heavy spin-2 states known as Kaluza-Klein (KK) gravitons plus a light scalar state called a radion. Given a sufficiently general Lagrangian, this radion can mix with the SM Higgs boson, leading to a direct coupling between the Higgs boson and the KK gravitons [85]. Both the KK gravitons and the radion can then be expected to decay to HH with the branching ratio being determined by the various free parameters of the model. These decays would once again lead to resonant HH production at the LHC within the invariant mass range of 250 GeV up to the TeV scale.

2.3 Effective field theory in Higgs boson pair production

The previous two sections outlined the many BSM theories to which HH is sensitive, most of which predict new particles somewhere in the $\sim 10^2 \text{ GeV} - 10^3 \text{ GeV}$ range. To extend beyond these energy ranges, to those theories with physics at scales which cannot be directly probed at the LHC, a different approach must be considered. While the benchmark models described in the previous sections aim to reduce the enormous set of BSM theories to a more

manageable and general form, we can go further. The use of Effective Field Theory (EFT) helps us to address both of these issues.

EFT, as described here, is built on the idea that knowledge of physics at some small scale m does not depend on the detailed dynamics of the physics at some much larger scale Λ . That is, for example, the idea that the interactions of a very high energy particle can be replaced by a set of interactions consisting only of low energy fields. To do so, we can expand the full theory in powers of the small parameter m/Λ , producing a second perturbative expansion, this time in the energy itself. In doing so, we abandon the idea of a fully renormalizable theory. Instead, we accept theories which are renormalizable only order-by-order in this power counting series. In practice, there are many technical complications to constructing such theories, including the handling of UV divergences, the inconsistencies between different regularization schemes, handling of the running of couplings, and potentially non-trivial power counting [86, 87]. This is all that will be said about such technical issues here. The focus of the rest of this section will be on the conceptual development and structure of two important EFT extensions to the SM, namely the Standard Model Effective Field Theory and the Higgs Effective Field Theory.

2.3.1 Formulating an effective field theory

To formulate an effective field theory, one must first decide on a few important points: the relevant degrees of freedom, the symmetries (which can be more than the original theory), and the expansion parameter or scale for power counting. Together these components form the basis of an EFT. Given these building blocks, there are two main approaches: the top-down approach and the bottom-up approach.

Top-down formulations

In the top-down approach, we begin with a high energy theory for which we wish to find a corresponding theory at low energies. This can be useful as a calculational tool, particularly in cases where matching to a low energy theory is relatively easy. The various effective field theories of QCD [88, 89] are examples of this. To perform this conversion from high energy theory to low energy EFT, we begin by integrating out the high mass physics. This requires the first decision: the cut-off scale above which you are integrating. It is important to keep in mind that this cut-off scale will dictate the properties of the low energy operators and the mixing of the coefficients between those operators, but it will not influence the real low energy physics. This is the same idea as the cut-offs in regularization. The cut-off here is not a real physical parameter of the theory, and so it will not influence the physics, only the

mathematical manifestation of it. The end result is a new low energy effective field theory, with a new set of constants and operators consisting only of a set of low-energy fields.

Bottom-up formulations

One may also take the opposite approach, and begin with a low energy theory from which the effective field theory description is constructed. This is very useful when the true high energy theory is unknown, as is the case with the SM when we interpret it as the low energy version of a more complete theory. Two good examples of such EFTs are the Standard Model Effective Field Theory and the Higgs Effective Field Theory which will be discussed in the following sections. To go about constructing an EFT from a low energy theory, one simply writes down the most general possible set of operators which respect the given symmetries. Again, these operators needn't be renormalizable as we now only care about the theory being renormalizable order-by-order. This means, for example, that within the SM we may begin to write down operators which have mass dimensions larger than 4. As we do so, we naturally introduce a power counting. To maintain the mass dimension of the action, each of these operators must come along with a coefficient which carries the appropriate inverse mass dimensions so as to correctly cancel and yield a mass dimension 0 action. For example, in the four-dimensional SM, this means that operators of mass dimension 5 must have a coefficient with mass dimension -1 . This yields a natural suppression of higher order operators at low energies. For example taking the coefficients for a mass dimension n operator to be c/Λ^{n-d} , where c is naturally of order 1 and d is the space-time dimension, we see that the coefficients decrease by a factor proportional to the power of the mass scale of new physics. Such a formalism is a very natural one as non-renormalizable operators are expected to be suppressed as (cf. relevant/irrelevant and marginal operators in renormalization group). It is also well motivated based on history. As seen through the discovery of electroweak theory, these small non-renormalizable corrections can appear as the remnants of new massive particles or a new energy scale.

It is important to note the general nature of this approach. In formulating an EFT in this manner, we have made relatively few assumptions about the nature of the high energy theory. This allows us to map constraints placed on this low energy theory to a large variety of BSM theories within the equivalence class. There are however two notable assumptions. First, we have assumed that there is a clear delineation in scales, with the new physics² all above a single large energy scale, Λ , which is substantially larger than the scale being probed and then the scale of the SM. This assumption is built into our *power counting*. This

²At least that new physics which couples to the SM

excludes many new low-energy resonant phenomenon and so this should still be considered a complimentary search to a set of searches for resonant benchmark models. In selecting our symmetries we have also placed assumptions on the possible UV completions of our EFT. That is, we have placed some constraint on the set of BSM theories which can map to our EFT. This assumption will be briefly discussed in the next two sections with regards to model selection and a more thorough coverage can be found in [90, 91].

2.3.2 Standard Model Effective Field Theory

The SM is known to be the most general possible Lagrangian consisting of the SM fields and satisfying the SM gauge symmetries. The Standard Model Effective Field Theory (SMEFT) is a bottom-up approach that takes the perspective that the Higgs field is an SU(2) doublet and extends the structure of the SM by writing down all additional non-renormalizable operators. The Lagrangian can then be split into components based on the mass dimension:

$$\mathcal{L}_{\text{SMEFT}} = \mathcal{L}^{(0)} + \mathcal{L}^{(2)} + \mathcal{L}^{(4)} + \mathcal{L}^{(5)} + \mathcal{L}^{(6)} + \dots \quad (2.4)$$

where $\mathcal{L}^{(n)}$ is defined as

$$\mathcal{L}^{(n)} = \frac{1}{\Lambda^{n-4}} \sum_k c_k^{(n)} \mathcal{O}_k^{(n)} \quad (2.5)$$

where k sums over all operators $\mathcal{O}_k^{(n)}$ of mass-dimension n , with constant coefficients $c_k^{(n)}$ which are naturally of order one, but which are 0 in the SM. Notably, $\mathcal{L}^{(0)}$ contains the cosmological constant, $\mathcal{L}^{(2)}$ is the Higgs boson mass term and $\mathcal{L}^{(4)}$ includes all other SM terms. The first of the SMEFT additions are found in $\mathcal{L}^{(5)}$, which consists of only one operator per generation. These operators are closely related to neutrino masses, but violate lepton number and are therefore not generally relevant for the current precision of HH measurements. The dimension-7 operators are also ignored here, since they are tightly constrained by lepton and baryon number violation. Focusing now on the lowest even dimension operators, we note that given a new mass scale of $O(1 \text{ TeV})$ and $O(m_H)$ interaction energy, dimension-6 terms should provide corrections of a few %, while the dimension-8 operators provide corrections of $O(10^{-4})$. Thus, with current and near-future HH sensitivity, the only relevant operators will be those of dimension-6.

In grand total, there are 59 $\mathcal{O}^{(6)}$ operators per generation; however, most of these operators do not contribute to HH at LO and so, given their already small corrections, are not relevant to current HH precision ($\sim O(1)$). Taking only those operators which do contribute




Figure 2.10: Some of the leading order HH diagrams in SMEFT with insertions shown for c_H (pink), c_{tH} (turquoise), c_{HG} (green), and c_{tG} (yellow). Insertions for the operator corresponding to $c_{H\Box}$ modifies all Higgs couplings and so, these modifications are not shown above. Figure adapted using a setup from [93].

to LO HH, we find a reduced set of 6 operators:

$$\begin{aligned} \Delta\mathcal{L} = & \frac{c_{H\Box}}{\Lambda^2} \partial_\mu(\Phi^\dagger\Phi) \partial^\mu(\Phi^\dagger\Phi) + \frac{c_{tH}}{\Lambda^2} \left(\Phi^\dagger\Phi - \frac{v^2}{2} \right) \bar{Q}t\tilde{\Phi} + \frac{c_H}{\Lambda^2} \left(\Phi^\dagger\Phi - \frac{v^2}{2} \right)^3 \\ & + \frac{c_{HG}}{\Lambda^2} \left(\Phi^\dagger\Phi - \frac{v^2}{2} \right) G_A^{\mu\nu} G_{\mu\nu}^A + \frac{c_{tG}}{\Lambda^2} i g_s (\bar{Q}\tau^{\mu\nu} T_A t) \tilde{\Phi} G_{\mu\nu}^A + h.c. \quad (2.6) \end{aligned}$$

with using the Warsaw basis and the notation of the SMEFT@NLO model [92], but with slightly different naming conventions obtained by replacing ϕ with H in the subscripts. Note that here we ignore the strong CP-violating operator which in this naming convention would be called $\mathcal{O}_{H\tilde{G}}$. The diagrammatic contributions of these operators to LO ggF HH production are shown in Figure 2.10. From the above operators, $\mathcal{O}_{H\Box}$ and $\mathcal{O}_H$ induce a correction to the Higgs self-coupling, with the following relation:

$$\kappa_\lambda = 1 - 2 \frac{v^4}{m_H^2 \Lambda^2} c_H - 3 \frac{v^2}{\Lambda^2} c_{H\Box}. \quad (2.7)$$

More generally, the $\mathcal{O}_{H\Box}$ provides a universal scaling to all Higgs couplings, and for this reason it is well constrained by single Higgs boson measurements [69]. Notably, at LO the only contribution of $\mathcal{O}_H$ is to the Higgs self-coupling making c_H equivalent to a reparameterization of κ_λ at LO. $\mathcal{O}_{HG}$ provides direct gluon-Higgs interaction, allowing for direct gluon to Higgs boson production as well as tree-level gluon HH production. This operator is thus also highly constrained by single Higgs boson measurements. $\mathcal{O}_{tH}$ shifts the value of the top Yukawa coupling and also induces a new ttHH vertex, allowing for enhanced HH production. The relatively weak constraints of the top Yukawa coupling (relative to other Higgs couplings) and the enhanced HH production via ttHH makes HH an interesting location to search for this operator. Finally, $\mathcal{O}_{tG}$, known as the top-dipole operator, induces an anomalous color-dipole moment for the top quark, modifies the top-gluon coupling, and adds a new top-Higgs-gluon vertex [69]. This operator introduces several new LO diagrams relevant to HH production.

From the above discussion it is clear that high energy physics can enter HH production in many ways, some of which may provide counteracting effects. It is thus important to address all relevant operators simultaneously to properly understand and parameterize the unexcluded pockets of phase space or to fully characterize an anomaly. The use of SMEFT as a general model also simplifies direct combinations between various channels. As noted above, the various operators of SMEFT impact not only HH production, but also other SM processes, most notably single Higgs boson production. By performing global fits, or otherwise combining constraints across these channels, the degeneracies within any given channel may be resolved. This is particularly relevant for HH where uncertainty on $c_{H\Box}$ or c_{HG} would be expected to negate the unique sensitivity of the channel to c_H and c_{tH} , as has been observed previously in other generalized theory frameworks [69, 94]. Through the use of SMEFT, HH is able to take advantage of powerful constraints from single Higgs boson production and other SM measurements to remove these degeneracies, allowing constraints to be placed directly on those parameters to which it is uniquely sensitive.

2.3.3 Higgs Effective Field Theory

The Higgs Effective Field Theory (HEFT) takes the same bottom-up approach to EFT construction as SMEFT, but provides a different organization of the BSM physics by using a non-linear realization of $SU(2) \times U(1)$. In particular, motivated by the $SU(2) \times SU(2) \rightarrow SU(2)$ symmetry breaking seen in chiral symmetry breaking of QCD, HEFT preserves the unbroken custodial symmetry of the SM by taking the Higgs field to be an $SU(2) \times U(1)$ singlet. The power counting is more complex than SMEFT, serving as a double expansion in the




Figure 2.11: Leading order diagrams in HEFT, with insertions shown for c_{hhh} (pink), c_t (turquoise), c_{ggh} (navy), c_{gghh} (green), and c_{tt} (yellow). Figure adapted from [93].

loop order and the chiral dimension [69, 95]. The end result is the following additional Lagrangian terms relevant to current HH measurements:

$$\Delta\mathcal{L}_\chi = -m_t \left(c_t \frac{h}{v} + c_{tt} \frac{h^2}{v^2} \right) \bar{t}t - c_{hhh} \frac{m_H^2}{2v} h^3 + \frac{\alpha_s}{8\pi} \left(c_{ggh} \frac{h}{v} + c_{gghh} \frac{h^2}{v^2} \right) G_A^{\mu\nu} G_{\mu\nu}^A \quad (2.8)$$

where c_i are the new coefficients [69]. To lowest order, couplings c_t and c_{hhh} take SM values of 1, while c_{tt} , c_{ggh} and c_{gghh} take SM values of 0. The diagrammatic contributions of these operators to LO ggF HH production are shown in Figure 2.11. In particular, c_{hhh} modifies the Higgs self-coupling, c_t modifies the top Yukawa coupling, c_{ggh} induces a ggH vertex, c_{gghh} induces a ggHH vertex, and c_{tt} induces a ttHH vertex. Note that the equivalent to the SMEFT chromomagnetic operator $\mathcal{O}_{tG}$ does not appear until two-loop order in HEFT and so is ignored here. Regardless, as was the case for SMEFT, the cross-section and invariant mass distribution of HH production are highly dependent upon these parameters. It should be noted that HEFT provides a more general description of possible higher energy physics than SMEFT; however, much of this additional theory space is potentially constrained. This flexibility also comes with some downsides, as it removes useful relations between operators relevant to HH and those relevant to other common processes. This is particularly relevant

for the induced $t\bar{t}HH$ coupling which is an independent parameter in HEFT, while in SMEFT it is proportional to the $t\bar{t}H$ coupling. Due to its prioritization of new Higgs sector physics, it is also a less popular theory within the ATLAS collaboration, making it more difficult to use in combinations. There are clearly then both advantages and disadvantages to using HEFT over SMEFT or vice-versa. In this thesis, this issue will be avoided and limits will be set to both models, but further discussion on these comparisons can be found in [69, 90, 91].

Chapter 3

The ATLAS Detector and the Large Hadron Collider

As is evident by the prolonged search for the Higgs boson, it is a difficult particle to produce in quantities abundant enough to be observed and studied. At present, there is only one particle accelerator capable of producing a substantial number of Higgs bosons, the Large Hadron Collider (LHC) [96]. The work shown in this thesis relies on data obtained by the ATLAS detector [97] at the LHC, from 2015-2018, during what is known as Run-2. This chapter serves to outline those experimental facilities relevant to the work presented in this thesis.

3.1 The Large Hadron Collider

The LHC is a circular hadron accelerator which makes up one part of the CERN accelerator complex, located near Geneva, Switzerland. It is 26.7 km in circumference and located 100 m underground. It consists of two rings, with counter-circulating beams of charged particles (typically protons or heavy ions). These beams are made to cross and collide at four points around the ring. These four points, known as interaction points, correspond to the locations of the four largest LHC experiments: ATLAS, CMS, LHCb, and ALICE. The LHC is designed to provide proton collisions with a center-of-mass energy of $\sqrt{s} = 14 \text{ TeV}$ and instantaneous luminosity of $L = 10^{34} \text{ cm}^{-2} \text{ s}^{-2}$. During the Run-2 data taking period relevant to this thesis, the LHC was operated with a center-of-mass energy of $\sqrt{s} = 13 \text{ TeV}$.




Figure 3.1: The portions of the CERN accelerator complex used to accelerate proton beams for eventual collision in the LHC. Figure adapted from [98].

3.1.1 Proton acceleration

To reach the peak energy used during operation, the LHC relies on a series of smaller accelerators, shown in Figure 3.1. For proton beams, the process begins with Hydrogen atoms, which are stripped of their electrons by an electric field, before being accelerated by a linear accelerator (LINAC2) to an energy of 50 MeV. The resulting beam of protons is passed on to the circular Proton Synchrotron Booster which further accelerates them to an energy of 1.4 GeV. The beam is then injected into the Proton Synchrotron and accelerated to 25 GeV, before being passed along to the Super Proton Synchrotron which accelerates the beam to 450 GeV. It is at this point that the beam is finally injected into one of the two vacuum pipes of the LHC. The protons which reach the LHC arrive in bunches, separated by 25 ns. Each bunch is approximately 7.7 cm long and contains approximately 10^{11} protons at the start of a normal fill cycle. Once in the LHC, the bunches are further accelerated by a set of 8 radio frequency (RF) cavities per beam. Within each of these cavities, an electric field is made to fluctuate at 400 MHz. This timing is precisely maintained to ensure that the protons will always experience an accelerating voltage. This technique also works to stabilize the bunches, as protons which arrive too early or late relative to their bunch will naturally receive a compensatingly small or large acceleration due to their position within the Electromagnetic (EM) wave. As these proton beams continue to circulate within the LHC with increasing energy,

they are steered by a series of dipoles made of niobium-titanium magnets.

3.1.2 Luminosity and pile-up

When attempting to maximize the amount of usable data, an important quantity to consider is the luminosity, as the total number of expected events for a given process, N , is given by

$$N = L\sigma = \int dt \mathcal{L}\sigma, \quad (3.1)$$

where $\mathcal{L}$ is the instantaneous luminosity, σ is the cross-section of the process, and t is time. The integral of the instantaneous luminosity over time, L is called the integrated luminosity. Thus to maximize the number of events or integrated luminosity, both the up-time and the instantaneous luminosity must be maximized. In particular, the instantaneous luminosity is optimized by modifying a variety of beam characteristics which together determine the instantaneous luminosity as

$$\mathcal{L} = \frac{fn_1n_2n_b}{4\pi\sigma_x\sigma_y}F, \quad (3.2)$$

where f is the frequency of revolution of the beam, n_1 and n_2 are the number of protons in the two bunches, n_b is the number of bunches in the beam, σ_x and σ_y characterize the convolved beam width along each axis, and F is a form factor characterizing other geometric effects such as the angle of beam crossing. It is through the optimization of these quantities that instantaneous luminosities of upwards of $10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ are achieved at the LHC. However, this large instantaneous luminosity leads to a large number of collisions occurring within a short time window and affecting the detector simultaneously. At the LHC, there are typically tens of collisions per bunch crossing, the majority of which are low energy inelastic collisions. Those collisions, besides the collision of interest, are referred to as in-time pile-up. Additionally, due to the tight 25ns spacing of the bunches, there can also be out-of-time pile-up, where collisions from one of the neighboring bunch crossings affect the detector during the response window for the collision of interest. As shown in Figure 3.2, the pile-up generally increased in each year of data taking. This is due to the increase in instantaneous luminosity as LHC beam performance was further improved. Total up-time also improved each year. The combined effects of the increase in up-time and instantaneous luminosity may be observed by the larger integrated luminosities seen in Figure 3.3 for later years.




Figure 3.2: Yearly distributions of the mean number of interactions per bunch crossing for data recorded by the ATLAS experiment at the LHC [99].

3.2 The ATLAS detector

The ATLAS detector is a multi-purpose particle detector located at Interaction Point 1 on the LHC ring [97]. This general approach is reflected in the detectors design. As shown in Figure 3.4, the ATLAS detector is composed of several different detector subsystems and an array of magnets which together allow it to directly measure the mass, charge, momentum, and trajectory of particles produced at or near the interaction point. Each of the different subsystems, namely the Inner Detector, Calorimeter System, and Muon Spectrometer serve a unique role in measuring the variety of particles produced by the high energy collisions at the center of the detector. An overview of how different particles interact with each of these subsystems and their components is shown in Figure 3.5. These three detector subsystems and their components are detailed in Sections 3.2.2, 3.2.3 and 3.2.4.




Figure 3.3: Yearly distributions of the mean number of interactions per bunch crossing for data recorded by the ATLAS experiment at the LHC [99].

3.2.1 Coordinate system

The coordinate system used to describe the detector is outlined in Figure 3.6. This is a right-handed system, originating at the interaction point, with the z-axis aligned with the beam pipe, the x-axis pointing towards the center of the LHC ring, and the y-axis pointing upwards. In polar coordinates, the azimuthal angle ϕ is defined in the x-y plane, originating on the x-axis, while the polar angle θ is defined in the y-z plane.

Due to the varying fraction of momentum carried by partons in hadronic collisions, the reference frames of such collisions are not consistently defined. It is therefore important for relevant kinematic variables to be Lorentz invariant. While differences in ϕ are Lorentz invariant, this is not true for θ and it is therefore rarely used. Instead, the polar angle is re-parameterized in either of two forms, rapidity or pseudo-rapidity. Differences in rapidity are Lorentz invariant and it is defined as:

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right), \quad (3.3)$$




Figure 3.4: Computer generated image of the ATLAS detector including descriptions of the major subsystems [100].

where E is the energy of a given reconstructed object and p_z is the z -component of the momentum. However, the more commonly used quantity is pseudo-rapidity:

$$\eta = -\ln \left(\tan \frac{\theta}{2} \right). \quad (3.4)$$

Differences in pseudo-rapidity are Lorentz invariant only for massless particles; however, this is a reasonable approximation for many high-energy scenarios. Pseudo-rapidity is often preferred because it does not rely on knowledge of the total momentum, which is difficult to measure for highly relativistic particles. Additionally, because total momentum in the plane perpendicular to the proton beam is approximately zero, momentum is typically dealt with only as a transverse quantity, p_T . In this case, transverse quantities are defined as the radial components of vectors and so $p_T = \sqrt{p_x^2 + p_y^2}$. Transverse momentum is also often used to define a so-called transverse energy, $E_T = \sqrt{m^2 + p_T^2}$. Finally, the angular distance between two objects is measured in units of $\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}$.




Figure 3.5: Diagram of particle paths within the ATLAS detector [101].

3.2.2 Inner Detector

The Inner Detector (ID), Figure 3.7, is the inner most subsystem of the ATLAS detector and is used for particle tracking with coverage over the $|\eta| < 2.5$ range. To accomplish this, the ID relies on three sub-detectors, each composed of several layers of sensitive components. Each of these sub-detectors is able to record the location that a charged particle interacts with the detector material. This recorded interaction point is known as a *hit*. As these sub-detectors are layered, they will record a sequence of hits as a particle passes through them. The combined positional information of these hits allow the reconstruction of the trajectory of a charged particle as it passes through the detector. Additionally, the ID is immersed in a 2 T magnetic field produced by the surrounding solenoid. This magnetic field is oriented along the beam axis and thus deflects charged particles in the ϕ direction. Through the measurement of this deflection, the ID enables the reconstruction of the p_T and charge of passing particles via the Lorentz force law.




Figure 3.6: ATLAS coordinate system overlaid on the LHC ring. The yellow point marks the location of the ATLAS experiment. Inspired by figure from [102].

Pixel Detector

The innermost sub-detector of the ID is known as the Silicon Pixel Detector (SPD). It is composed of 4 layers [104] of silicon pixel sensor arrays, which detect charged particles via the electron-hole pairs they produce as they pass through the sensor. To precisely measure the angular position of hits and to better distinguish nearby particles, the pixels are smallest in the innermost layer, where they measure $250\mu\text{m}$ in the z -direction and $50\mu\text{m}$ in the ϕ -direction. This inner layer was added prior to the start of Run-2 and is known as the Insertable B-Layer [104]. The pixels in the outer three layers of the SPD are larger, measuring $400\mu\text{m} \times 50\mu\text{m}$. The ϕ -direction is prioritized with higher resolution because this allows for more precise measurements of particle deflection and thus p_T . In total, these four layers contain around 90 million pixels.

Semiconductor Tracker

The next layers of the Inner Detector are part of the Semiconductor Tracker (SCT) subsystem. These layers are also composed of silicon sensors arrays; however, these sensors are




Figure 3.7: A computer generated image of the ATLAS Inner Detector, showing the Pixel Detector, Semiconductor Tracker, and Transition Radiation Tracker [103].

much longer in the z -direction than the ϕ -direction at $12\text{ cm} \times 80\mu\text{m}$. There are 4 layers of these sensor arrays in the barrel and 2 in each of the end-caps. The reduction in the z -resolution of the individual sensors is further compensated for by placing sensors in each layer at off-set z positions and with a 40 mrad stereo-angle between sensors on either side of a given segment.

Transition Radiation Tracker

The third and final sub-detector of the ID, the Transition Radiation Tracker (TRT), provides additional coverage to the $|\eta| < 2.0$ region. This sub-detector is composed of approximately 300 000 kapton straw tubes filled with a primarily Noble gas mixture. Each of these tubes contains a central wire which collects electron cascades induced from the ionization of the gas mixture by passing charged particles. The space between these straw tubes is filled with polypropylene. The large differences in the indices of refraction of these materials causes particles to emit transition radiation as they cross the boundaries. The amount of transition radiation produced by a particle of a given energy is related to its mass [105], which allows




Figure 3.8: A computer generated image of the ATLAS Calorimeter, showing the LAr electromagnetic, LAr hadronic, and tile calorimeters [106].

the distinction of light charged particles from heavier charged particles. Most notably, this allows electrons and positrons to be distinguished from pions.

3.2.3 Calorimeter System

The next subsystem in the ATLAS detector is the calorimeter system, Figure 3.8, which surrounds the solenoid magnet and inner detector, covering $|\eta| < 4.9$. This subsystem is composed of three separate calorimeters, the inner Electromagnetic Calorimeter (ECAL), the outer Hadronic Calorimeter (HCAL), and the Forward Calorimeter (FCAL). All of these calorimeters are sampling calorimeters, that is they are composed of alternating layers of absorbing and sensing materials. The ECAL and HCAL are further sub-divided into barrel and end-cap regions. There are two separate cryostats for the calorimeters, one containing the FCAL and end-cap calorimeters and the other containing the ECAL barrel section. The HCAL barrel section is not contained in a cryostat. These calorimeters work in conjunction to contain and measure the energy of electrons, photons, taus, and hadrons.

Electromagnetic calorimeter

The Electromagnetic Calorimeter (ECAL) covers the range $|\eta| < 3.2$ and focuses on absorbing energy via Electromagnetic (EM) showers. These showers develop in the lead absorber layers of the calorimeter, as incident electrons produce bremsstrahlung radiation and photons undergo pair-production. Both of these processes produce additional photons or electrons and the cycle continues, with the number of new particles in the chain increasing exponentially. This results in the energy of the incident particle being spread among an exponential number of shower constituents. These showers are characterized by the radiation length of the material, X_0 , which corresponds to the average distance over which bremsstrahlung reduces the energy of an electron by a factor of e . When the energies of shower constituents are sufficiently low the particles rapidly lose energy through the ionization of the liquid argon layers. The collection of this ionization energy by the active material leads to an electrical signal proportional to the energy of the original particle, allowing that original energy to be measured.

To distinguish the energy coming from the many particles that are produced during an event, the calorimeter is sub-divided into many segments. The lead absorber is folded into an ‘accordion’ shape which produces multiple effective layers of absorber and active material, providing further longitudinal segmentation of the calorimeter. The ECAL is also granular, with the highest granularity in the $|\eta| < 2.5$ region, at 0.025×0.025 in $\eta \times \phi$. In total, the ECAL has around 170 000 read-out channels.

With a total thickness of $\sim 20X_0$, the ECAL is highly effective at collecting energy from photons and electrons; however, it is not as effective for muons and hadrons. Muons, being much heavier than electrons, deposit a relatively small fraction of their energy in the ECAL because the rate of bremsstrahlung is inversely proportional to the square of the mass of the incident particle. Like electrons and other charged particles, muons also deposit energy via ionization. For a given particle and material, this energy loss depends only on the velocity of the particle and is described by the Bethe-Bloch equation (with corrections) [107]. Typical muons lose a few GeV to ionization, which is only enough to stop those with particularly low momentum. Hadrons on the other hand, do deposit a significant fraction of their energy in the ECAL, but they do so via nuclear interactions in addition to bremsstrahlung. These interactions are characterized by a length scale which is much larger than X_0 . Due to this longer length scale, most hadronic showers propagate beyond the ECAL and are further captured by the HCAL.

Hadronic calorimeter

The Hadronic Calorimeter (HCAL) is composed of three different sub-detectors, each using a different combination of active and passive material, but otherwise operating on the same principles. The central hadronic calorimeter covers the range $|\eta| < 1.7$ and is composed of passive steel absorber layers with plastic scintillator tiles as the active material. These scintillating tiles are read out by wavelength-shifting optical fibers which are connected to photo-multiplier tubes. The end-cap hadronic calorimeter, spanning $1.5 < |\eta| < 3.2$, uses a copper passive material and a liquid argon active material. These hadronic calorimeters operate upon the same principle as the ECAL, that is they measure the energy of a high energy particle by inducing particle showers and collecting the energy from the resulting low energy particles. The main difference, is that the hadronic calorimeter is optimized to induce hadronic showers via the strong interaction and ionization.

Before a shower begins, charged particles will have already begun to lose energy to ionization of the nuclei within the absorbers, with this ionization being measured by the active material as in the ECAL. The shower itself begins with a hard collision between an incident particle and a nucleus within the absorber material. This collision produces a large number of secondary particles, mostly pions and neutrons. These secondary particles collide with other nuclei, leading to a rapidly growing number of particles within the shower. The mean distance between these collisions is known as the nuclear interaction length, λ_{int} , and is material-specific. After several interaction lengths, the shower begins to terminate as the energies of the secondary particles fall below the threshold required for pion production. At this point, remaining particles begin to decay and nuclei de-excite, emitting electrons and photons which are measured by the active material.

The ATLAS HCAL has a depth of $\sim 10\lambda_{\text{int}}$ which is sufficient to completely stop most hadronic showers. However, unlike in EM showers, much of the energy of the captured hadronic showers will remain unmeasured. This is because much of the energy produced during these hadronic showers (usually 10-40%) is emitted through neutrons, neutrinos, heat, nuclear vibrations, and other processes which are not measured by the active material. Due to the large number of possible hadronic processes, there is much larger variation among hadronic showers than EM showers. This results in a correspondingly large variation in the fraction of energy deposited by hadronic showers, even for identical initiating particles. Hadronic showers exhibit even further variation beyond this based on the type and energy of the initiating particle. These discrepancies are accounted for during calibration, which is covered further in Section 5.2.3. As with the ECAL, most muons deposit very little of their energy in the HCAL, since they do not experience the strong interaction. They do, however,

deposit some energy via ionization (~ 3 GeV).

Forward calorimeter

The Forward Calorimeter (FCAL) [108] consists of two identical calorimeters covering the region $3.1 < |\eta| < 4.9$ at each end of the ATLAS experiment. In order to produce thin enough gaps to handle the high intensity radiation in the forward region, these calorimeters use modules composed of a metal matrix of tubes (oriented along the z -axis) instead of metal plates. These tubes contain inner rods made of the same metal, with liquid argon filling the ~ 0.3 mm gaps between the inner and outer tubes. Both the inner and outer tubes serve as absorbing material, while the liquid argon is the active material.

There are three modules in each calorimeter, with the inner most module being made of copper and the outer two of tungsten. The inner copper module is optimized for EM showers with a depth of $28X_0$ or $2.7\lambda_{\text{int}}$. The two tungsten modules are optimized for hadronic calorimetry with depths of $3.6 - 3.7\lambda_{\text{int}}$, providing improved containment of hadronic showers. This improved containment not only provides better measurements of hadronic showers, but also prevents leakage into the neighboring HCAL end-cap, where energies are much lower. Due to the high intensity radiation in the forward regions, the FCAL also includes an outer brass plate to stop any remaining hadronic showers from entering the muon system. However, as with the HCAL and ECAL, the majority of muons will pass through the FCAL having deposited only a relatively small fraction of their energy, by ionization.

3.2.4 Muon Spectrometer

As previously discussed, muons travel relatively unhindered through the calorimeters, providing the opportunity to more precisely measure and isolate their tracks. This is precisely the role of the Muon Spectrometer (MS), Figure 3.9, the outermost layer of the ATLAS detector. The MS is divided into three sections, one in each end-cap and one in the barrel, together covering the range $|\eta| < 2.7$. Like the Inner Detector, the MS is composed of several layers, 3 in this case, with each immersed in a magnetic field. In fact, each of the three MS sections is provided a magnetic field by a separate magnet. These three magnets provide an average magnetic field of 0.5 T, all pointing in the ϕ direction. The three layers allow the triangulation of the particles trajectory while the magnetic field imparts a deflection in the η direction. Again, just as in the ID, the deflection caused by this magnetic field allows for the measurement of the momentum. To accomplish these tracking goals, the MS takes advantage of two sets of detectors, the muon trigger system and the muon tracking system.




Figure 3.9: A computer generated image of the ATLAS Muon sub-system, showing the location of each of the main components [109].

Muon trigger system chambers

The muon trigger system is composed of two different types of detector systems: Resistive Plate Chambers (RPCs) for $|\eta| < 1.05$ and Thin Gap Chambers (TGCs) for $1.05 < |\eta| < 2.4$. These systems work together to provide high resolution timing information and triggering to the MS.

The RPCs are made from two parallel resistive plates lined with metallic strips on their outer surfaces and filled with a gas mixture. The strips are oriented in the ϕ direction on one plate and the η direction on the other. A 3.8 kV voltage is held between the plates. When muons or other charged particles pass through the central gas mixture, they ionize it, leading to an immediate avalanche of electrons which, though rapidly quenched by one of the gases, induces a signal on the metallic strips via capacitive coupling. Together these effects lead to a short signal pulse providing a high time resolution of approximately 1 ns [110]. This high resolution timing response is ideal for rapid triggering which can be used in conjunction

with the coarse measurement of the ϕ position provided to trigger selective readout of the neighboring high precision tracking components.

The TGCs are multi-wire proportional chambers based on similar principles to the RPCs, but are able to better handle the high background conditions of the forward regions. They are composed of graphite-coated cathode chambers filled with a gas mixture, with a plane of parallel high voltage wires along the length of the chamber. These chambers have a < 25 ns time resolution and provide ϕ coordinate information with a resolution of 5 mm.

Muon tracking system chambers

The muon tracking system also consists of two separate detector systems: Muon Drift Tubes (MDTs) for $|\eta| < 2.7$ and Cathode Strip Chambers (CSCs) for $2.0 < |\eta| < 2.7$. These detectors have slower read-out systems, but provide much higher resolution position information. To compensate for their slow read-out, they rely on the muon trigger system and only read-out after an initial trigger decision has been made.

The MDTs are 3 cm diameter aluminum drift tubes filled with a gas mixture and a central wire. These detectors work similarly to the TRT drift tubes, but without the transition radiation component. These tubes are oriented in the ϕ direction to ensure high-precision measurement of the η -deflection induced by the magnetic field, and thus of the muon momentum.

The CSCs are multi-wire proportional chambers and operate similarly to the TGCs. The cathodes on the outer surface are oriented in the ϕ direction on one side and the η direction on the other to provide both η and ϕ positional information. These chambers are used in the inner most layers of the forward region as the MDTs are unable to tolerate the higher rate requirements.

3.2.5 Trigger and data acquisition

With a bunch crossing rate of over 40 MHz, ~ 100 million read-out channels, and event read-out sizes of ~ 1 MB, it is impossible to record all the collisions which occur within ATLAS. For this reason, the ATLAS Trigger and Data Acquisition (TDAQ) system [111, 112] is used to select and store only those events which are likely to be relevant for pre-determined ATLAS physics goals. The end result is the greatly reduced output rate of 1 kHz, meaning data is only recorded for roughly 0.0025% of bunch crossings. To achieve this the ATLAS TDAQ system relies on a layered approach consisting of two sequential triggers: the hardware-based Level-1 (L1) trigger and the software-based High-Level Trigger (HLT). The L1 trigger serves as an initial filter. It operates at a much higher rate, but consequentially

makes only coarse simple decisions. The HLT receives the output from the L1 trigger, now $< 100\text{kHz}$, and is able to run more sophisticated algorithms to further reduce the data rate down to the final rate of 1kHz .

Level-1 trigger

The Level-1 trigger is a hardware-based synchronous trigger consisting of a Central Trigger Processor (CTP) and three trigger systems: L1Calo, L1Muon, and L1Topo. This trigger is allocated $2.5\mu\text{s}$ per bunch crossing, but due to data transfer time, only about $0.5\mu\text{s}$ of this is actually available for decision making. The decision making itself is based on a pre-programmed set of criteria known as a trigger menu.

The L1Calo trigger system uses dedicated lower-granularity analogue trigger signals from the calorimeter to search for signatures of high energy particles or large missing E_T consistent with conditions specified in the trigger menu. That is, the trigger system scans calorimeter regions for energy maxima which exceed certain menu thresholds. Each match is considered a Region of Interest (RoI). After this processing, the final number of RoI matching each menu item is transmitted to the CTP.

The L1Muon trigger system uses information from the RPC and TGC muon trigger chambers discussed in Section 3.2.4. It searches for coincident hits across these chambers which are generally consistent with a muon originating from the center of the detector. The trigger also calculates an estimate of the muon p_T from these hits. If this p_T exceeds a threshold in the trigger menu it is considered an RoI. As with L1Calo, L1Muon provides information on the number of RoI matching each menu item to the CTP.

The L1Topo trigger system receives trigger object information from the L1Calo and L1Muon systems and uses it to calculate angular separation, invariant mass, and the scalar sum of transverse momentum of objects in the event (H_T). This additional topological information is then matched against criteria in the trigger menu. These trigger menu items typically target events with large invariant mass, large H_T , low angular separation, or some combination of the above. After processing, this selection information is passed on to the CTP, as with the other trigger systems.

The CTP receives this information from the three trigger systems and makes a final decision before returning a single bit which indicates whether the information stored in the front-end electronics should be read out. In doing so, the CTP also handles deadtime, enforcing both a minimum time between accepts and a maximum number of accepts within a given time window. This deadtime is imposed to avoid exceeding detector read-out rates.




Figure 3.10: Distributions of total CPU time (over $\sim 400\,000$ processing units) used per event for two different HLT algorithms [113].

High-level trigger

The high-level trigger is a complex software based trigger which follows the Level-1 trigger. For every event accepted by L1, the data from the full detector is read-out and buffered. The HLT is then able to request this data as needed. However, the HLT typically takes advantage of initial fast trigger algorithms for early rejection. This often involves running algorithms over event data fragments from the RoI specified by L1, loading more data only as necessary. As shown in Figure 3.10, typical HLT algorithms process the vast majority of events very quickly. Thanks to these time saving early rejection elements and the 400 000 processing units dedicated to the HLT, later stages of HLT algorithms can be quite complex and often resemble those used in offline construction. Common HLT algorithms include tracking [114], b -tagging [115], and calorimeter clustering [116].

Chapter 4

Simulation

Experimental design and theoretical understanding both play central roles in the search for new physics at the LHC; however, we also require a method to connect these two domains. This is the role played by simulation. By using simulated physics samples, built upon the relevant underlying theories, we are able to better understand how best to search for new physics. For reasons which will become apparent in the following sections, these simulated samples are often simply called Monte Carlo (MC) samples within ATLAS. These MC samples consist of simulated events drawn from the expected distributions for a given theoretical model within a specified detector (ATLAS, in this case), under a certain set of assumptions. The process of generating these samples can be split broadly into two steps: event generation and detector simulation. The remainder of this chapter is dedicated to the description of these two steps and the underlying tools used for each.

4.1 Event generation

Event generation consists of the simulation of the interactions and decays beginning with the proton-proton interaction and ending just prior to the entry of particles into the sensitive material of the detectors. The event generation calculations rely on the parton model and are split not by the chronological ordering of particle interactions, but by factorization of the inclusive cross-section. More precisely, event generation is split into the following stages: simulation of the hard process, initial- and final-state radiation (parton showering), and hadronization. Figure 4.1 shows an illustration of this factorization procedure. The following sections describe the approximate procedures followed by many common event generators, covering each of these stages plus additional information about decays and multiple interactions. Further details of the specific event generators used in this thesis can be




Figure 4.1: A schematic showing the various stages of event generation. The hard interaction is indicated by a dark red circle. Additional final-state (red) and initial state (blue) is shown. Parton showering is also shown in the many branchings of both initial and final state partons. Hadronization is shown as light green blobs, resulting in the conversion of a single parton into multiple particles. A secondary interaction is indicated by a purple oval, decay of short-lived states is shown by small red circles, and decay of particles produced during hadronization is shown in dark green circles. Finally, the emission of photons and electrons is shown in yellow. Figure taken from [117].

found in the corresponding references for POWHEG Box v2 [118] with PYTHIA8 [119] and MADGRAPH5_AMC@NLO [120].

4.1.1 Hard process

Simulation of the hard process begins with the generation of the relevant parton-level matrix elements from a set of Feynman rules and, for Next-to-Leading Order (NLO), UV counter-terms and other cancellation rules. This matrix element is used to construct the corresponding parton-level differential cross-section. MC integration methods are then used to produce a set of representative events drawn from the distribution described by the differential cross-section. These methods are useful here due to the dimensionless scaling of the error and the ability to handle functions which may have divergent derivatives. Such attributes are important for these calculations as matrix elements can often involve high-dimensional phase-spaces and step-functions. Such methods proceed by drawing random samples, called events, from the phase-space of the differential cross-section and weighting them by the value of the function at that point. These events are then unweighted using acceptance-rejection sampling, whereby events are randomly accepted at a rate proportional to the probability density at that point [121]. The final output of this stage of event generation is a set of events containing kinematic information, spin, color connections, and flavor information for each particle/parton. The information provided with each of these events is known as the parton-level truth information.

4.1.2 Parton showers

During the interaction of the protons, gluons and quarks are undergoing acceleration and will therefore radiate other quarks and gluons. This radiation occurs for both the initial and final state particles and is known as initial-state radiation (ISR) or final-state radiation (FSR), respectively. The parton shower component of event generation consists of the simulation of these radiation effects. QCD radiation is enhanced in the infrared, and so ISR and FSR consist mainly of soft collinear emissions, which leads to divergences in the determination of the emission distribution. To avoid this, a lower cut off on the energy of the emissions is placed. This cut-off is justified because collinear partons are not resolvable and so can be safely treated as one parton in the showering process. With this in mind, Sudakov form factors are used to encode the probability density of emitting a resolvable parton of a given energy, given the energy of the original parton. For initial-state radiation, the parton distribution functions are also accounted for in this probability density function. An iterative MC algorithm is then used to determine the shower content for each parton by solving the DGLAP

equations (or other radiation functions) via a series of iterative branching steps which integrate the probability density function as a function of some shower evolution variable. It should be noted that particular choice of shower evolution variable, renormalization scale, radiation functions, and other aspects of the algorithm can lead to discernible differences in the produced shower [122].

4.1.3 Hadronization

Following showering, in addition to the colorless particles, the events contain a series of partons linked through a series of color connections; however, partons are not physically measurable particles. Due to QCD confinement, these partons will form hadrons, and this hadronization procedure must be simulated as well. There are two dominant models of hadronization: string fragmentation [123] and cluster fragmentation [124], only the former will be covered here.

In the Lund string fragmentation model, long-range QCD confinement interactions are represented as a series of strings. In this model, as a string grows and its energy increases, it begins to break apart via the quantum tunneling of new quark-antiquark pairs out of the vacuum. These new quarks form the new ends of the string at the break-point. This process proceeds until only on-mass-shell hadrons remain. To simulate this process, the color and particle history information from the previous steps is used to form a set of strings between partons. These strings are then split using an iterative algorithm which iterates along the length of the string splitting it based on the tunneling probabilities until the energy of all segments of string are lower than the hadronization energy. The result, known as truth particle-level information, is a set of particles each with corresponding kinematic, spin, flavor, and history information.

4.1.4 Underlying event

Event generators may also account for secondary parton collisions from the proton remnants during a bunch crossing by using information about the instantaneous luminosity. This information, along with information from the original hard-scatter process, is used to determine the expected rate of multiple parton-parton scatterings. These additional scatterings may then be calculated via similar methods to the above and added to the same event record.

4.1.5 Decays

Many of the states produced during the initial hard scatter and during subsequent hadronization will be unstable; however, the exact treatment of these states will depend on their lifetimes. Any states from the original hard scatter with particularly short lifetimes, less than the hadronization scale, are decayed as part of the hard process. States with lifetimes comparable to or longer than the hadronization scale, but not long enough to propagate through the detector, are decayed during hadronization as are any shorter lived states produced during hadronization. The final product of the entire event generation procedure is this final set of stable and long-lived particles, along with corresponding particle histories, kinematic information, and spins.

4.2 ATLAS detector simulation

Having generated a set of events, it is still necessary to understand the signatures left by the final-state particles during their propagation through the detector. This is done by propagating each final-state particle through a full model of the ATLAS detector using GEANT4 simulation toolkit [125] and some additional ATLAS software [126]. Interactions with this model detector are recorded as hits and the position, time, and energy deposited during collision are recorded for later digitization. Decays, interaction vertices, and tracks which are of particular interest to analyses (e.g. those of sufficient energy) are also saved as particle-level truth information. To simulate pile-up, additional hits from other sources are overlaid. These other sources can include things such as additional proton-proton interactions, cavern backgrounds, beam halo, and beam gas interactions. With all relevant hits now included, these hits are digitized by ATLAS digitization software [126]. This digitization software models the response of each different detector to a given energy deposit, including temporal and position effects which can lead to varying signal shapes. This simulation includes many fine-details such as cross-talk, electric noise, and channel dependent variations for each different detector sub-system and corresponding read-out electronics. The end result of digitization is a set of Raw Data Objects (RDOs). These RDOs are then passed through the ATLAS high-level trigger and standard object reconstruction. The input event plus additional simulation-level truth information is kept for use during calibration. The object reconstruction and calibration steps are detailed in Chapter 5.

4.3 Fast detector simulations

The full simulation procedure works well, but can be very computationally expensive, taking up to several minutes per event. For this reason, additional faster simulation methods are sometimes also used, most commonly ATLFAST-II (AFII). The majority of simulation time is spent in simulating the calorimeter, where a large number of shower particles are produced and must also be stepped through the detector. AFII, therefore, uses full-simulation for the tracking and muon systems, while for the calorimeter RDOs are directly simulated from event generator output using FastCaloSim [127]. FastCaloSim relies on fits to a large set of fully simulated (GEANT4) samples to parameterize the longitudinal and lateral calorimeter response. The lateral response shape is given by a radially symmetric function centered around the impact point of the particle in a given calorimeter layer, modified with additional asymmetries for non-perpendicular collisions, parameterized by the angle of incidence. Different function parameters are determined in different bins of particle type, energy, $|\eta|$, and shower depth. The longitudinal energy deposition is estimated from two-dimensional histograms of calibrated energy response and shower depth, and layer energy fraction and shower depth. Correlations between layers are preserved by a correlation matrix of layer energy fractions for each bin of shower depth. These correlations and histograms are performed separately for each bin of particle type, energy, and $|\eta|$. Simulation results are obtained by selecting the corresponding parameterization bin for that particle type, energy, and $|\eta|$, then randomly sampling a shower energy and shower depth. These values are rescaled to match the particle energy and then the energy fraction in each calorimeter layer is sampled from the correlation matrix for that shower depth. The combination of these two parameterizations leads to a reasonably accurate fast-simulation method, which is 20-40 $\times$ faster than full GEANT4 simulation, and which correctly models many observables, but notably, not jet substructure (see Section 5.2 and Section 5.3). Figure 4.2 shows a comparison of full and fast simulation.




Figure 4.2: Comparison of ATLAS-II (red squares) and GEANT4 (yellow) jet p_T (left) and rapidity (right) distributions for jets with $p_T > 30$ GeV and $|y| < 2.8$. [127].

Chapter 5

Object Reconstruction

After events of interest have been recorded as electrical signals in the detector, these signals are converted to both low-level and high-level physics objects through a series of processes known as object reconstruction. This chapter provides an overview of the reconstruction procedures used within this thesis. Section 5.1 details the reconstruction of several low-level physics objects used in the identification and reconstruction of particles and other high-level physics objects. Section 5.2 covers the reconstruction of jets, one of the core components of the analysis performed as part of this thesis. Finally, Section 5.3 covers the *b*-tagging algorithms used to identify jets containing b-hadrons.

5.1 Tracks, vertices and topological-clusters

By tracing a hypothetical particle out from the center of the detector, we can carefully consider each stage of the measurement. First, there is the point at which that particle originates. If at least one other particle also originated at this point, as is typical for high energy particles, this point may be identified as a vertex. Vertices associated with the original proton-proton collision are known as primary vertices. As the particle continues out further from the center of the detector, it will likely pass through the ID. If it does, and if it is charged, it will leave behind a trail of hits. Next, the particle will enter the calorimeter system, where most particles will begin to deposit energy, leaving a trail of energy deposits in the ECAL, HCAL, or both. These energy deposits can be clustered into a single object, known as a topological-cluster. Finally, if the particle makes it through the calorimeters, it will arrive at the MS where charged particles will leave additional hits in the MS chambers. Any additional such hits are also considered part of the aforementioned track. Together these tracks, vertices, and topological-clusters provide a useful source of information about a particles

path through the detector and, due to the magnetic fields spread throughout the detector, its momentum. More details about each of these objects and their reconstruction is provided in the following subsections.

5.1.1 Tracks

As discussed in Sections 3.2.2 and 3.2.4, charged particles leave hits on the ID and MS when passing through them. The particle trajectories reconstructed from these hits are known as tracks. However, there are a large number of particles produced during each event which means a large number of events, particularly when the pile-up is also large. It is therefore non-trivial to determine which hits are associated with the same track. In order to correctly reconstruct a track, track-finding algorithms [114] are used.

Space-points and clusters

This track finding procedure begins by reconstructing the three dimensional position of a hit. This is complicated by the fact that single particles typically leave hits in multiple sensing elements (e.g. neighboring pixels), as seen in Figure 5.1a. To resolve this, a connected components algorithm is used to group pixels which share a common edge or corner. The point at which the particle trajectory intersects with the sensor is then determined by interpolating from the distribution of ionization over the clustered pixels. The resulting position is known as a space point. Of course, it is possible that multiple particles could have interacted with the same region of the detector, meaning the resulting cluster contains contributions from all of these particles. An example of such a merged cluster is shown in Figure 5.1b.

Track finding and track candidates

The primary track finding algorithm is seeded by track seeds are formed from sets of three adjacent space points. To maximize purity, seeds are required to have a sufficiently small impact parameter, sufficiently large momentum and an additional space-point is required to be in good agreement with the predicted trajectory. These remaining high quality seeds are processed by a combinatorial Kalman filter [128] which builds track candidates by iteratively incorporating additional space-points which are consistent with the predicted trajectory of the prior iteration. This algorithm will produce all valid trajectories for each given seed, which could mean multiple trajectories per seed. These track candidates are then further processed by an ambiguity solver.




Figure 5.1: Illustrations of (a) several single particle clusters where particles have produced a signal in multiple pixels and (b) a merged cluster, where multiple particles have contributed to the same cluster of pixels. Particle trajectories are shown by the arrows, with corresponding energy deposits shown in the correspondingly colored cells. [114]

Ambiguity solving

To aid in resolving ambiguities and pruning of incorrect and overlapping tracks, track candidates are assigned a track score based on their quality. This score prioritizes tracks which are a good fit (low χ^2), have few missing space-points where the track crosses a detector (holes), have a large number of high resolution clusters, and have high p_T . Tracks with a bad score are immediately discarded, while the remainder are ordered by descending track score for further processing steps. The ambiguity solver then removes candidates which have more than two shared clusters or which contain a cluster shared by more than two candidates. When determining whether a cluster is shared, only the preceding candidates are considered. These removed candidates are further processed by a neural network to identify merged clusters. Candidates are not penalized for clusters which are identified as merged. If this further processing causes a track to now pass the previous criteria, it is accepted. Otherwise, clusters causing it to fail the criteria are removed and the candidate is added back to the queue. Those tracks which survive this ambiguity solving process are processed by a neural network to more precisely determine their positions and corresponding uncertainties. Finally, the tracks are re-fit using this updated information and the resulting tracks are taken as the final set of tracks for the event.

5.1.2 Vertices

Vertices are points at which particles have interacted or decayed and are identified as the point of origin of multiple tracks. These vertices are subdivided into two main classes: primary vertices resulting from the primary proton-proton collisions and secondary vertices occurring later in the event. The primary vertices may be further subdivided into hard-scatter vertices (HSVs) and pile-up vertices. Each triggered event has one HSV candidate corresponding to the highest energy collision in that event. All remaining primary vertices in the event are identified as pile-up vertices. Primary vertices play the important role of associating various objects with the same production event and providing a measure of the amount of pile-up. Furthermore, the HSV in particular is crucial for the reconstruction of event-level kinematic quantities and for pile-up suppression. Secondary vertices are of particular importance for flavor tagging and are typically only searched for within jets (Section 5.2). The role of these vertices in flavor-tagging will be discussed in Section 5.3. Due to the different natures of primary and secondary vertices, their surrounding environments, and downstream uses, they are reconstructed by separate algorithms.

Primary vertex reconstruction

To avoid using incorrectly reconstructed tracks or those coming from a secondary vertex, the primary vertex reconstruction algorithm [129] uses only a subset of tracks. These tracks must have sufficiently high p_T , since they should have originated from a hard-scatter process and because tracks are easier to reconstruct when they are sufficiently straight. To ensure the tracks originated near the beam line they must have a small and precisely measured impact parameter and at least one hit on the inner most layer of the pixel detector. Finally, these tracks must be generally high quality with lots of clusters and few shared clusters or holes.

Along with this high quality set of tracks, the vertexing algorithm also requires an initial vertex position as a seed. The transverse coordinate of this seed is taken to be the center of the beam-spot while the z -position is taken to be the mode of the z -positions of the tracks at their closest approach to the beam spot.

Given this seed and the tracks, a vertex is found by an iterative annealing procedure. At each iteration, the tracks which are incompatible with the vertex are down-weighted and then the vertex position is recalculated. After the algorithm has converged, the final vertex position is stored and tracks which are compatible with this vertex remain associated with it. The process is then repeated with the remaining tracks. This process terminates once no more vertices may be found. All vertices with more than one associated track are kept. The vertex with the highest sum of squared track p_T is identified as the HSV candidate, as it is

expected to correspond to the highest energy proton-proton collision.

While the HSV serves a central role in pile-up suppression and event reconstruction, the remaining primary vertices are useful for the complimentary information that they provide. Since all non-HSV primary vertices are considered pile-up vertices, the total number of primary vertices (N_{PV}) can be used as a measure of the amount of pile-up for a given event. Various downstream reconstruction algorithms (e.g. jet reconstruction) use N_{PV} as a proxy variable to measure and correct for the varying contribution of pile-up to reconstructed objects. The association of tracks and other objects to primary vertices also provides a handle for downstream algorithms to select only those objects associated with specific interactions or prompt decays.

Secondary vertex reconstruction in jets

There are far fewer constraints on the location of secondary vertices within the detector and on the underlying physics that produces them. For this reason, there are several different common vertexing algorithms within ATLAS. This section will focus on the single secondary vertex finding algorithm (SSVF) [130], which is used for flavor-tagging. The SSVF can make use of any tracks which were either not included in the original set for primary vertex reconstruction or which were not associated with a vertex during that stage. Of these tracks, it will only consider those which fall inside the cone of a jet. Jets will be discussed further in Section 5.2; for now, it is sufficient to consider them only as conical volumes of interest within the detector. SSVF relies on the fact that for an n track vertex, all tracks should be sufficiently close to a common space point, such that the vertex formed by any two of these tracks will also be close to the n track vertex. However, in a high-density tracking environment, tracks which originate at a primary vertex can sometimes become close enough to create apparent secondary vertices which could be incorrectly identified by such a simple algorithm. For this reason, an extra step is added which rejects many two-track vertices which are not compatible with the vertex of interest. SSVF performs this rejection by requiring that two-track vertices not have tracks with hits in layers closer to the center of the detector than the secondary vertex. Since the SSVF is focused on flavor-tagging, it also rejects two-track vertices which are compatible with background processes such as long-lived particle decays, hadronic interaction with the detector, or photo-conversion. The SSVF then combines the remaining two-track vertices into a multi-vertex by merging those vertices which are close in space. If the resulting vertex has an invariant mass $> 6\text{ GeV}$ or a high χ^2 , the track with the largest χ^2 of the track-vertex association is removed from the set and the process is repeated. This process is repeated until the vertex satisfies the criteria. If

the final vertex has only two tracks, the vertex cleaning steps are repeated again.

5.1.3 Topological-clusters

Thanks to the fine segmentation of the ATLAS calorimeter read-out system, it is possible to reduce the impact of background noise and pile-up on signal read-out in the calorimeter by grouping cells into larger clusters. These clusters are not necessarily expected to correspond to single particle showers, but can include merged particle responses, partial responses or some combination of the two. Instead, they should be thought of only as measures of energy flow within the calorimeter and as a tool for background suppression. To form these clusters, neighboring cells are grouped using the topological-cluster (topo-cluster) reconstruction algorithm [116].

The topo-cluster reconstruction algorithm makes use of the cell signal significance, $\zeta_{\text{cell}}^{\text{EM}}$, which is the ratio of the measured signal to the average expected background noise for that cell. The algorithm begins by creating seeds from all cells with sufficiently large signal significances, $|\zeta_{\text{cell}}^{\text{EM}}| \geq 4$. These seed cells form the initial proto-clusters. These seeds then add any neighboring cells with $|\zeta_{\text{cell}}^{\text{EM}}| \geq 2$ to the proto-cluster. This procedure is iteratively repeated with the neighbors of the newly added cells. If proto-clusters intersect, having added the same cell with $|\zeta_{\text{cell}}^{\text{EM}}| \geq 2$, they are merged into a single cluster. This iterative expansion ends once there are no more neighbors which meet the $|\zeta_{\text{cell}}^{\text{EM}}| \geq 2$ criteria. Next a single additional expansion step is made, adding all cells which neighbor the proto-cluster and which have $|\zeta_{\text{cell}}^{\text{EM}}| \geq 0$.

The proto-clusters built by this method fail to account for spatial structure within the cluster, inhibiting the detailed and unbiased reconstruction of future objects from these clusters. To account for this issue with the initial clustering procedure, clusters with multiple local maxima are split into multiple clusters. Cells within the original cluster are shared between the newly split clusters of the two closest local maxima. To avoid double counting, the contribution of a cell to each of these two clusters is determined by a geometric weight based on the cluster energies and the distance from the cell to the center of the two clusters. These two weights are chosen such that they always sum to unity. This splitting represents the final step, after which these proto-clusters, now topo-clusters, may be used for further reconstruction tasks, such as reconstruction of jets.

5.2 Jets

As discussed in Section 1.2, QCD does not allow for free colored states. Thus, quarks produced during an interaction undergo hadronization, producing a collimated shower of color neutral hadrons. This collection of particles is known as a jet. Due to the large number of hadrons in jets, jets are all but guaranteed to contain both electrically charged and neutral objects. For this reason, sufficiently high energy jets deposit energy in both the electromagnetic and hadronic calorimeters. Jets are reconstructed either by clustering topo-clusters (EMTopo jets) or so-called particle flow objects (PFlow jets), which are constructed from combinations of tracks and topo-clusters. Since all of the jets used in this thesis are PFlow jets, EMTopo jets will be ignored in further discussions and PFlow jets will be referred to simply as *jets*. Reconstructed jets, originally reconstructed at the EM scale of the calorimeter, are then calibrated to the particle level using a jet energy scale (JES) correction. Pile-up jets are then suppressed using the jet-vertex-tagger (JVT) discriminant.

5.2.1 Particle flow objects

The particle flow (PFlow) algorithm [131] takes advantage of the fact that charged jet constituents leave both tracks in the ID and energy deposits in the calorimeter. Using measurements from both of these objects together provides several advantages. Firstly, they provide complimentary measurements of energy. That is, the tracker has better resolution at low energies, while the calorimeter has better resolution at high energies. The ID is also able to reconstruct much lower energy objects due to lower noise occupancy. Additionally, lower energy jet constituents are often deflected out of the original jet cone by the magnetic field before reaching the calorimeter. The use of ID information allows these events to be located earlier and correctly associated with the jet cone. Another benefit is that the ID provides much better angular resolution than the calorimeter, while the calorimeter provides additional coverage in the forward regions. Finally, since tracks are associated to vertices, one is able to reject signals coming from pile-up vertices.

The particle flow algorithm begins by selecting tracks which are well measured, not matched to any previously reconstructed particles, and which have $0.5 \text{ GeV} < p_T < 40 \text{ GeV}$. The lower p_T threshold is used to remove background and to avoid wasting compute time on tracks which have little impact on the final energy, while the upper threshold is used to remove tracks which will not provide any improvement over the existing set of topo-clusters. Following track selection, track positions are extrapolated to the calorimeter. Tracks are then matched to the closest topo-cluster within $\Delta R < 1.64$ that has an energy which is at

least 10% of the track momentum. Tracks which are not matched are left in the list of tracks and ignored for further steps. To ensure that showers split across multiple topo-clusters are correctly reconstructed, tracks with energies significantly larger than the matched cluster are assigned to additional clusters until this is no longer the case. Next, to avoid double-counting energy, those tracks which were successfully matched to a topo-cluster must have their energy contribution subtracted from it. If the expected energy deposit of a track exceeds that of the assigned topo-clusters, they are simply removed. Otherwise, this subtraction is done moving outwards from the center of the projected energy density distribution subtracting rings of cells from each layer. For the final ring, where the energy of the cells exceeds the remaining track energy, the cell energies are scaled down to remove the remaining track energy. Remaining topo-clusters and cells are removed if their energy is consistent with a fluctuation in the expected energy deposit of the track. Following this matching and subtraction procedure, the selected tracks and remaining topo-clusters are considered fully reconstructed particle flow objects.

5.2.2 Anti- k_t jet clustering algorithm

Prior to jet reconstruction, particle flow tracks which do not pass close enough to the HSV are removed to reduce pile-up and topo-cluster positions are recalculated as being relative to the HSV. Following these adjustments, jet reconstruction proceeds using the anti- k_t jet clustering algorithm [132]. This algorithm begins by defining two important distance measures. The distance between two objects is defined as

$$d_{ij} = \min \left(\frac{1}{p_{T,i}^2}, \frac{1}{p_{T,j}^2} \right) \frac{\Delta_{ij}^2}{R^2}, \quad (5.1)$$

and the distance between the i 'th object and the beam is defined as:

$$d_{iB} = \frac{1}{p_{T,i}^2}, \quad (5.2)$$

where $\Delta_{ij} = \sqrt{(y_i - y_j)^2 + (\phi_i - \phi_j)^2}$ is the angular distance between the i 'th and j 'th objects, and R is the radius parameter, set to $R = 0.4$ for standard ATLAS jet reconstruction. These distances are calculated for all the relevant particle-flow objects. The algorithm then finds the smallest distance among all d_{ij} and d_{iB} . If the smallest distance is between two objects (d_{ij}), those objects are merged and the process is repeated. If the smallest distance is between some object and the beam d_{iB} , that object is called a jet and removed from the list.

This procedure continues until all objects have been removed from the list. The resulting output of this algorithm is a collection of jets. The anti- k_t algorithm has several beneficial properties. First of all, the algorithm is infrared safe, this follows from the fact that the algorithm clusters soft-particles with hard-particles before clustering them with each other. This resilience to soft-particles also reduces momentum resolution loss to pile-up [132]. Secondly, collinear particles are clustered immediately as they have $\Delta_{ij} = 0$, this means that if one particle were replaced by two collinear particles of half the original energy, the result would be the same (collinear safe). This algorithm also produces jets with a well defined shape, being conical unless they are overlapping in which case they may be clipped cones. Together these effects mean that this algorithm is both theoretically well-motivated and usually easier to calibrate.

5.2.3 Jet calibration

When measuring jet energy, the original measurements are made at the scale of the calorimeter (EM scale) not the energy scale of the original particle. These energy scales are different because the calorimeters only measure the fraction of the energy that interacts with the active material. The detector response also varies throughout the detector volume and in the different calorimeters, where the response is defined as the ratio of the measured energy to the true particle energy. Finally, the energy measured by the calorimeter can include contributions from pile-up which must be further corrected for. Thus, to obtain the correct particle level energy for a given jet, all of these differences must be accounted for. This is done through a procedure known as jet calibration [133].

Pile-up correction

The first two steps in the calibration procedure correct for pile-up. In particular, since pile-up from particles which left tracks has already been removed, this correction removes contributions from neutral and out-of-time pile-up, as well as those from the underlying event. To estimate the pile-up contribution, a set of k_t jets is formed from positive energy topo-clusters within the range $|\eta| < 2$ using the k_t algorithm [134] with $R = 0.4$. This algorithm is used because it naturally reconstructs a uniform background of soft pile-up jets. The average momentum density of these k_t jets in the $y - \phi$ plane, ρ , is then calculated and multiplied by the area, A , of each anti- k_t jet to estimate the contributions from pile-up. The four-momentum of each anti- k_t jet is then scaled to reduce the p_T by this amount. A residual correction is then applied to correct for differences in pile-up conditions between events. This correction is a function of the number of vertices, N_{PV} , and pile-up, μ , for that event. Separate residual




Figure 5.2: Average jet energy response as a function of reconstructed jet (a) detector η and (b) energy [133].

corrections are applied for several bins of truth-level p_T and detector $|\eta|$ to account for the varying impacts of these conditions on different regions of the calorimeter and for jets with different p_T .

Jet energy scale and η calibration

Next, the jet energy scale correction is applied to correct for the incomplete and varying response of the detector. This correction consists of two components: a correction to the energy and a correction to the jet η . Figure 5.2 shows the average ratio of the reconstructed and truth jet energies, the average jet energy response, as a function of reconstructed jet energy and η . Notably, the jet energy response varies significantly with energy and η . The variation with energy is caused by the increased fraction of the jet energy contained in the EM shower component for high energy particles. Whereas, the variation with η is a by-product of the varying detector geometry and calorimeter technologies. To correct for these effects, the jet energy response (jet η) is corrected by the average ratio of (difference in) the corresponding reconstructed and truth quantities found in MC dijet events. These corrections are applied in bins of reconstructed jet energy and η .

Global sequential calibration

Following these corrections, the jet energy response is still found to vary depending on the flavor and energy distribution of the jet constituents as well as other factors, such as the particle which initiated the jet. All of these factors can influence the resulting shower properties, most notably, the amount of the resulting shower which passes through the interface

between the ECAL and HCAL, the amount which makes it to the HCAL, and the amount that makes it through the HCAL. Since both the area after the HCAL and between the ECAL and HCAL consists of dead material and since the energy response of hadrons in the HCAL is lower, these effects can lead to an under-estimation of the true jet energy. The global sequential calibration consists of a series of sequential calibrations which correct the dependence of the response on:

- the jet energy fraction contributed by constituent tracks,
- the fraction measured in the last EM layer,
- the fraction measured in the first Tile layer, and
- the number of associated tracks in the MS.

The first component of this calibration provides information about the hadronic component of the jet, the next two about the amount of energy deposited in the dead region between the ECAL and HCAL, and the last about whether the jet made it through the HCAL.

***In situ* jet calibration**

Following the previous steps, jets have been calibrated to the particle level; however, differences still exist between data and MC. An *in situ* jet calibration is used to correct for these differences. Unlike in the previous steps where MC truth quantities could be used exclusively, to perform this calibration, jet response must be calculated in data. This measurement relies on jet balancing the p_T of a jet against that of a well-calibrated reference object or system. This reference objects p_T takes the place of the usual MC truth value and the response is calculated as the average ratio of the jet p_T to the reference object p_T , in bins of the reference object p_T . The ratio of the MC response to the data response is then used to correct the jet p_T in data following a numerical inversion procedure which converts this ratio from a function of reference object p_T to a function of jet p_T . To cover the full range of p_T and η , this calibration is implemented using three different *in situ* analyses: η intercalibration, Z/γ +jet and multijet balance. Full coverage of these analyses and their performance can be found in [133].

5.2.4 Jet-vertex-tagger

Many of the initially reconstructed jets come from pile-up, not from the hard-scatter event of interest. These pile-up jets are further suppressed by applying a cut on the jet-vertex-tagger


(JVT) discriminant. Two variables are used to formulate this discriminant: a corrected form jet vertex fraction (corrJVF) and R_{pT} . Jet vertex fraction (JVF) is defined as the fraction of the scalar sum of track p_T for that event which is associated with a given jet and originating from the HSV. The corrected jet vertex fraction adds to this an additional correction which accounts for the dependence of the sum on the number of primary vertices in that event. R_{pT} is defined as the fraction of the scalar sum of track p_T in a given jet coming from tracks associated to the HSV. Both of these variables are strong discriminants against pile-up (Figure 5.3), because pile-up jets are expected to consist mostly of low p_T tracks (low corrJVF) and tracks that are not associated with the HSV (low R_{pT}). The JVT discriminant itself is defined as the two-dimensional likelihood of corrJVF and R_{pT} specifying the relative probability for a jet to be a hard-scatter jet. The likelihood is obtained using a k-nearest neighbors algorithm trained on a set of simulated hard-scatter and pileup jets with $20 < p_T < 50 \text{ GeV}$ and $|\eta| < 2.4$ [135]. As shown in Figure 5.3c, the JVT discriminant provides good separation of pile-up and hard-scatter jets. The discriminating power of JVT allows for low pile-up jet fake rates across a wide range of hard-scatter jet efficiencies as shown in Figure 5.4.

5.3 b -tagging


Given a set of jets, it is typically useful to identify the flavor of the parton from which those jets originated. Since the Higgs boson decays most commonly to a pair of b -quarks and this thesis focuses on production of two Higgs bosons, the tagging of jets originating from b -quarks is of particularly importance. Such jets are called b -jets and this process of identifying them is known as b -tagging. An example of a b -jet is shown in Figure 5.5 with useful quantities of interest for tagging labeled.

5.3.1 Algorithm inputs


Since the b -mesons produced by the hadronization of b -quarks have relatively long lifetimes of $\sim 1.5 \text{ ps}$ [14], they typically travel a few millimeters before decaying. Thus, b -jets typically originate from a secondary vertex which is on average displaced by some similar (p_T dependent) amount from the production vertex. This secondary vertex is one of the core components in the identification of b -jets. There are also several other commonalities between the different ATLAS b -tagging algorithms [136]. These algorithms all require at least one primary vertex in the event and they frequently require the jet to have a minimum transverse impact parameter, measured relative to the HSV. In addition, in order to reject



(a) corrJVF



(b) R_{pT}



(c) Jet-vertex-tagger

Figure 5.3: Distributions of (a) corrJVF, (b) R_{pT} , and (c) JVT for simulated pile-up and hard-scatter jets with $20 < p_T < 30 \text{ GeV}$ and $|\eta| < 2.4$. [135]




Figure 5.4: Fake rates from pile-up as a function of hard-scatter jet efficiency for JVF, corrJVF, R_{pT} , and JVT. JVT can be seen to give consistently better performance than the other discriminants. [135]

background, ID tracks are required to have a $p_T > 500$ MeV, and jets with a $p_T < 120$ GeV are required to pass a cut based on the JVT discriminant. Track to jet matching for b -jets are made using the ΔR between the two objects, with a tighter cut for higher p_T since higher momentum jets are expected to be more collimated. The jet axis is also used to assign a sign to the track impact parameters, with positive impact parameters for tracks which cross the jet axis above the primary vertex and negative for those below. Finally, it is worth noting the way in which b -jets are labeled at truth level in MC simulation. Flavor labels are assigned




Figure 5.5: An illustration of an interaction producing two light flavored jets and a b -jet. The secondary vertex and corresponding displaced tracks are labeled, as are the transverse decay length (L_{xy}) and impact parameter (d_0), both of which are relevant for b -tagging. [115]

by matching jets to hadrons with $p_T > 5 \text{ GeV}$ within a cone of $\Delta R = 0.3$ around the jet axis. If the hadron is a b -hadron, the jet is identified as a b -jet. Jets matched to c -hadrons or τ -leptons are identified as c -jets and τ -jets, respectively. Finally, the remaining jets are labeled as light jets.

5.3.2 Algorithms

There are many b -tagging algorithms, all of which make use of vertexing information, track impact parameters and many of the same principles. This section will focus on those algorithms used within this thesis, namely MV2 and DL1r [137]. These algorithms are both trained/tuned using a hybrid MC sample composed of filtered $t\bar{t}$ and Z' (decaying into hadronic jet pairs) events. Additionally, they both make use of the outputs from several low-level b -tagging algorithms: IP2D and IP3D, RNNIP, SV1, and JETFITTER.

IP2D and IP3D

IP2D and IP3D [137] are complementary impact parameter-based algorithms. Both algorithms make use of the track signed transverse impact parameter, with IP3D also using the longitudinal impact parameter. They then fit probability density functions of the input variables for each flavor label. These functions are then used to construct discriminants for each possible pairing of flavor labels from the sum of the log probability ratios for each track for that pair of labels. This results in three final discriminants, one for discriminating b -jets from c -jets, one for distinguishing b -jets from light jets, and finally one for discriminating c -jets from light jets.

RNNIP

The IPD2 and IPD3 algorithms make the simplifying assumption that the flavor-dependent per-track likelihoods can be computed without considering correlations between tracks. This assumption is required to make computations tractable for these algorithms; however, there can be correlations between different tracks. For a b -jet, multiple charged particles can be produced at the secondary vertex corresponding to the b -hadron decay. These particles are correlated and would produce correlated tracks. This same effect would not exist for light jets, where there is no such decay. RNNIP [138] is an impact parameter based algorithm that uses an recurrent neural network (RNN) to account for these correlations. RNNs are Neural Networks (NNs) which, due to their graph structure and internal memory, are capable of capturing sequential-dependencies of arbitrary-length within the input. In this case, the RNN takes as input a variety of per-jet variables, including signed impact parameters, angular distance, and p_T information. Prior to input, the variable set for each track is ordered by the magnitude of the significance of the transverse impact parameter. This ordering is chosen to encode information about the relevant decay chains that may exist for some b -jets. The final output of the RNNIP algorithm is a set of b -jet, c -jet, τ -jet, and light jet probabilities for each jet.

SV1

The secondary vertex tagging algorithm relies on the vertex found for a given jet using SSVF, which was previously discussed in Section 5.1.2. The SV1 tagger uses 8 different discriminating variables based on relevant quantities related to the SSVF vertex, including: the invariant mass, SSVF vertex associated track multiplicity, and the three dimensional distance between this vertex and the PV. These 8 variables are all that are used in the MV2

and DL1r algorithms. A full list of variables and further details of how they are used in the original SV1 algorithm can be found in [130].

JETFITTER

JETFITTER [139] is a topological multi-vertex finding algorithm that exploits the topological structure of the weak decays of b - and c -hadrons for identification. This algorithm uses a modified Kalman filter to reconstruct the hadron flight direction by finding a common line between the PV and the b - or c -hadron vertex. Again eight discriminating variables may be produced from this algorithm, including vertex-associated track multiplicity, invariant mass from displaced vertices, and energy fraction of tracks associated with displaced vertices. As with SV1, these eight variables are all that is used in MV2 and DL1r. A full list of variables and further details on JETFITTER can be found in [139].

MV2

The MV2 algorithm [137] is a boosted decision tree (BDT) that combines the outputs of IP2D, IP3D, SV1, and JETFITTER. The output of this BDT is known as the MV2c10 score. The BDT is trained using b -jets as signal samples and background composed of 7% c -jets and 93% light jets. The algorithm is provided jet p_T and $|\eta|$, but during training the b -jets and c -jets are weighted to match the kinematic distributions of the light jets. This prevents the kinematic biases which could be introduced if the BDT optimized on the differences in these kinematic distributions.

DL1r

The DL1r algorithm [137] is a neural network (NN) which provides a multi-dimensional output corresponding to the probabilities for a jet to be a b -jet, c -jet, or light jet. The algorithm uses the same inputs as MV2, an additional set of c -tagging variables from JETFITTER, and the output from RNNIP.

5.3.3 Calibration

To ensure that tagger performance is consistent between data and MC, these b -tagging algorithms must be calibrated. To do this, b -tagging performance in data is estimated using b -, c -, and light jet enriched data samples [140]. The tagging efficiency in these data samples is then compared to that in MC, with the ratio of these efficiencies then used as a scale factor to correct the MC tagging efficiency. In order to correctly scale the MC tagging ef-

Working Point	Selection	Rejections		
		c -jet	τ -jet	light jet
60%	> 2.74	27	220	1300
70%	> 2.02	9.4	43	390
77%	> 1.45	4.9	14	130
85%	> 0.46	2.6	3.9	29

Table 5.1: Selection cut and c -, τ - and light jet rejections corresponding to different b -jet efficiency working points as determined from evaluation on $t\bar{t}$ events as described in Section 5.3.3. [137]

efficiency, these scale factors must depend on the flavor of the jet. For truth b -jets, they are simply set to the ratio mentioned above, while for other jets they are calculated as the ratio $(1 - \epsilon_{\text{data}})/(1 - \epsilon_{\text{MC}})$ where ϵ_{MC} and ϵ_{data} are the efficiencies in MC and data respectively. As with many calibration factors, these scale factors are found to vary for different jet p_T and η and so they are binned in these variables during derivation and application.

5.3.4 Working points

The DL1 and MV2 algorithms detailed in Section 5.3.2 provide a discriminant score as output, but make no decisions as to what cut should be placed on this output when labeling a b -jet. Instead, the cut may be determined such that a desired b -tagging efficiency is achieved. Several efficiencies, known as *working points*, are determined and corresponding scale factors are determined during calibration. At the analysis level, one of these working points is chosen and b -jets with a score greater than the working point threshold are labeled as b -jets. The working points used in this thesis are the 77% and 85% working points. Details of these working points as well as two other working points are provided for the DL1 algorithm in Table 5.1.

5.3.5 b -jet energy correction

The jet calibration described in Section 5.2.3 focuses on the more common light-quark and gluon-initiated jets. As a result, the calibration does not account for the effects of the b -hadron decay on the jet energy as this is relevant only for b -jets. In particular, b -hadron decays can produce neutrinos and muons and the energy from these decay products is not fully included in the jet energy. Additionally, b -hadron fragmentation produces a wider spread of hadrons which results in more energy falling out of the $R = 0.4$ jet cone. To

correct for the effect of missing muon energy, a search is done for muons within $\Delta R < \min(0.4, 0.04 + 10\text{GeV}/p_T^\mu)$ of the jet axis and which have $p_T > 4\text{GeV}$ and $|\eta| < 2.5$. If any muons are found within this range, the four-vector of the closest muon is added to the jet. The estimated energy contribution of any additional muons within this range is then subtracted from the jet energy. To correct for missing energy from neutrinos and out-of-cone energy, a jet p_T dependent correction factor is applied. This correction factor is derived in $t\bar{t}$ events.

Chapter 6

Higgs Pair Production in the $HH \rightarrow b\bar{b}b\bar{b}$ Final-state

This chapter builds upon previous chapters by describing the search for Higgs pair production at the LHC using data collected by the ATLAS experiment. In particular, the focus will be on searches in the $HH \rightarrow b\bar{b}b\bar{b}$ final-state. Due to the large branching ratio of $H \rightarrow b\bar{b}$, this final-state benefits from the largest expected signal (Figure 6.1); however, as a purely hadronic process it also contains difficult backgrounds which require an interesting set of new techniques to model.

The search itself is divided into two separate analyses: one focusing on resonant signals and one on non-resonant signals, including the SM signal. These analyses share many similarities, but were separately optimized due to the different properties of these two types of signals. The resonant analysis was the first analysis to be completed and so, served partially as a testing ground for the non-resonant analysis which would come later. These separate optimizations resulted in slightly different trigger strategies, pairing strategies, analysis regions, and final discriminants.

The differences arise primarily from the more peaked nature of the resonant signal and the focus on maintaining sensitivity to signals at low invariant masses (m_{HH}). As the background is peaked towards low m_{HH} (~ 300 GeV), sensitivity is typically lower in this region. For the non-resonant analysis, even those BSM signals which peak at lower m_{HH} still have relatively high m_{HH} tails which contribute substantially to sensitivity (cf large κ_λ values in Figures 2.2 and 2.4). For low mass resonant signals, however, there is no such high mass tail to provide sensitivity. Thus, for the non-resonant analysis a smoother background distribution and better high-mass signal acceptance can be prioritized over low mass signal

$b\bar{b}$	WW^*	gg	$\tau^+\tau^-$	$c\bar{c}$	ZZ^*	$\gamma\gamma$	
33.9%	24.9%	9.5%	7.3%	3.4%	3.1%	0.3%	$b\bar{b}$
	4.6%	3.5%	2.7%	1.2%	1.1%	< 0.1%	WW^*
		0.7%	1.0%	0.5%	0.4%	< 0.1%	gg
			0.4%	0.4%	0.3%	< 0.1%	$\tau^+\tau^-$
				< 0.1%	0.2%	< 0.1%	$c\bar{c}$
					< 0.1%	< 0.1%	ZZ^*
						< 0.1%	$\gamma\gamma$

Figure 6.1: Branching ratio for Higgs pairs produced from collisions at LHC energy scales (13 TeV).

efficiency, while in the resonant analysis compromises must be made to maintain signal efficiency for low mass resonances.

In addition to these optimization differences, the resonant analysis is complemented by an additional boosted resonant channel. The boosted channel targets high mass resonances (900 GeV to 5 TeV) where the Higgs bosons are produced with a large p_T . When such large p_T particles decay, the decay products will be highly collimated. This results in the two b -hadrons from each Higgs boson being reconstructed as a single large-radius jet instead of a pair of well isolated b -jets. This unique topology requires a different analysis approach, which is beyond the scope of this thesis. The resolved resonant channel discussed in this

thesis covers events where the 4 b -jets are well isolated, for signal masses up to 1.5 TeV. Further details on the boosted channel and the combination between these two channels can be found in [141]. The resonant analysis for the VBF channel was also performed separately [142] and will not be covered here.

These analyses build and expand upon previous work [94, 142, 143] with greater physics coverage, improved analysis strategies, and in some cases, more data (ggF previously used 27 fb^{-1} , a factor of $\sim 4.7\times$ less than the new analysis). In particular, object reconstruction has been improved through the use of the particle-flow jets described in Section 5.2. This analysis also benefits from the improved b -tagging performance of the DL1r algorithm relative to the MV2 algorithm used previously. The analysis strategy has been improved with a new neural-network driven background estimate, updated jet-pairing strategies, improved trigger scale factor methodologies, and a joint statistical treatment of the non-resonant VBF and ggF channels. Finally, the ggF non-resonant analysis has expanded substantially in scope with the inclusion of interpretations using two different EFT frameworks (SMEFT and HEFT).

6.1 Dataset and simulation

6.1.1 Dataset

Both the resonant and non-resonant analyses use the full ATLAS Run 2 dataset consisting of 126 fb^{-1} data collected from 2016 to 2018. Data collected during 2015 and $\sim 9 \text{ fb}^{-1}$ of data collected during 2016 is not used due to low trigger efficiency and b -jet trigger inefficiencies, respectively. The integrated luminosities for each year are 24.6 fb^{-1} in 2016, 43.7 fb^{-1} in 2017, and 57.7 fb^{-1} in 2018.

6.1.2 Monte Carlo samples

Both the resonant and non-resonant analyses use MC samples to model the signal processes. The backgrounds are not modeled using MC because the dominant background processes are QCD multi-jets processes. These processes are difficult to model with simulation due to the large number of relevant Feynman diagrams and the complexity of those diagrams. Instead, MC samples are used only as validation of the data-driven background estimation described in Section 6.5.

Non-resonant background validation samples

The background validation samples consist of a mixture of QCD di-jet samples and $t\bar{t}$ samples. The di-jet samples are generated using PYTHIA8 [119] with EvtGen, using the NNPDF2.3 LO PDF set [144]. The $t\bar{t}$ samples are generated using POWHEG Box v2 [118] with the NNPDF3.0 NLO PDF set [145], PYTHIA8 is used for parton showering and hadronization, and EvtGen is used for decays.

Non-resonant signal samples

The non-resonant analysis uses ggF HH samples produced for $\kappa_\lambda = 1$ (SM) and $\kappa_\lambda = 10$ and VBF samples produced for $(\kappa_\lambda, \kappa_{2V}, \kappa_V) = (1, 1, 1)$ (SM), $(1, 1.5, 1)$, $(2, 1, 1)$, $(10, 1, 1)$, $(1, 1, 0.5)$, $(-5, 1, 0.5)$ with additional samples for validation (see Section 6.1.3). The ggF events are generated using POWHEG Box v2 [118] at Next-to-Leading Order (NLO), with full NLO corrections with finite top mass, using the PDF4LHC PDF set [146]. VBF events are generated at Leading Order (LO) using MADGRAPH5_AMC@NLO [120] using the same PDF set. In both cases, PYTHIA8 [119] is used for the nominal parton showering and hadronization. Additional samples are also generated with HERWIG7 [147, 148] to assess parton shower uncertainties. All non-resonant samples utilize the full ATLAS detector simulation procedure described in Section 4.2. The ggF samples are then normalized to the NNLO cross-section calculated with approximate top-quark mass effects (FTapprox) [70]. The VBF samples are normalized to the Next-to-Next-to-Next-Leading Order (N3LO) QCD cross-section [149]. Non-resonant variations which are not simulated, including all SMEFT and HEFT variations, were obtained either through reweighting or as a linear-combination of simulated samples. These methods will be detailed in Section 6.1.3.

Resonant signal samples

Resonant signals are generated for two generalized models, a scalar model and a spin-2 model. In both cases, they are produced via gluon-gluon fusion and decayed resonantly to a pair of Higgs bosons. These models are produced for a range of signals from 251 GeV to 1.5 TeV. The lower mass threshold of 251 GeV is chosen as it corresponds to the minimum (whole number) resonance mass that can produce two Higgs bosons given $2m_H = (250.20 \pm 0.14)$ GeV. The scalar model is implemented as a 2HDM with the resulting heavy CP-even neutral scalar as the resonance and no other BSM particles included. The width of the scalar is set to be smaller than the detector resolution. The spin-2 model uses a KK graviton with $k/\bar{M}_{pl} = 1$, resulting widths vary with resonance mass from 3% for large masses to 13%

for low masses. In neither case is this negligible relative to the detector resolution which varies from $\lesssim 5\%$ at large masses to $\lesssim 15\%$ at low masses. Events are generated for both models at LO in α_s using MADGRAPH5_AMC@NLO and the NNPDF2.3 PDF set. For parton showering and hadronization, HERWIG7 is used with the MMHT2014 PDF set [150] for the scalar model and PYTHIA8 is used for the spin-2 model. Both models use EvtGen [151] for heavy flavor decays. AtlFast-II (AFII) is used for detector simulation for masses up to 900 GeV. Beyond this point, full simulation is used to maintain consistency with the boosted analysis, which required the jet substructure information that is poorly modeled in AFII (see Section 4.2).

6.1.3 Signal reweighting

As the number of signal variations being examined by the non-resonant analysis is quite large, it is not feasible to generate a sample for every variation. Instead, a combination of reweighting and linear-combination methods is used to model these variations using existing samples.

Reweighting ggF for κ_λ variations

To obtain additional variations of κ_λ , POWHEG Box v2 [118] is used to generate samples to truth parton level for $\kappa_\lambda = 0, 1, 10, 20$. First we note that, at LO, any given κ_λ sample is a linear combination of three basis components, namely the square amplitudes of the two LO diagrams (box and triangle in Figure 2.2), plus the interference cross-term. Writing out this general amplitude for a given value of κ_λ

$$|A(\kappa_\lambda)|^2 = |A_\square|^2 + \kappa_\lambda (A_\square^* A_\triangle + A_\square A_\triangle^*) + \kappa_\lambda^2 |A_\triangle|^2 \quad (6.1)$$

where $|A(\kappa_\lambda)|^2$ is the square amplitude for a given κ_λ and $A_\square$ and $A_\triangle$ are the amplitudes of the SM square and triangle diagram, respectively. Note that the interference cross-term cannot be found using the other two terms as the generated samples present square amplitudes which do not contain the necessary phase information to do this. Of course, we are free to choose any basis, and so we can generate additional variations by taking the appropriate linear combination of three of the generated samples. Coefficients for this linear combination can be found as usual in linear algebra, by solving the linear set of equations. Doing so for

the basis set $\kappa_\lambda = 0, 1, 20$ one obtains the following:

$$|A(\kappa_\lambda)|^2 = \frac{380 + 19\kappa_\lambda^2 - 399\kappa_\lambda}{380} |A(0)|^2 + \frac{400\kappa_\lambda - 20\kappa_\lambda^2}{380} |A(1)|^2 + \frac{\kappa_\lambda^2 - \kappa_\lambda}{380} |A(20)|^2 \quad (6.2)$$

To apply this equation, samples are used to form m_{HH} histograms, with 10 GeV bins and the equation is applied bin-wise. With this new truth level histogram for the κ_λ of interest, we can derive a reweighting as a function of m_{HH} by taking the ratio of the linear combination histogram to the histogram for $\kappa_\lambda=1$. This gives us our final truth level weights for that histogram. This process can be summarized as:

$$W(\kappa_\lambda, m_{HH}) = \frac{380 + 19\kappa_\lambda^2 - 399\kappa_\lambda}{380} R_{0,1}(m_{HH}) + \frac{400\kappa_\lambda - 20\kappa_\lambda^2}{380} + \frac{\kappa_\lambda^2 - \kappa_\lambda}{380} R_{20,1}(m_{HH}) \quad (6.3)$$

where $R_{i,j}(m_{HH})$ is the ratio of the square amplitude of $\kappa_\lambda = i$ to that of $\kappa_\lambda = j$ as a function of m_{HH} . These weights are applied to the fully simulated and reconstructed $\kappa_\lambda = 1$ sample as a function of the stored truth m_{HH} information to obtain the κ_λ variation of interest. These reweighted samples are then renormalized to the NLO FTapprox theory cross-section obtained using the equations from [70]. This process makes the assumption that the signal acceptance times efficiency can be well represented as a function of m_{HH} , with no other dependence on κ_λ besides the relative rates in any given m_{HH} bin. It also assumes that the LO equation is sufficient to parameterize a linear combination of these NLO samples. To confirm these assumption, these weights were cross-checked at truth level and reco-level vs the corresponding $\kappa_\lambda=10$ sample and good agreement was found as shown in Figure 6.2.

Reweightting ggF to obtain SMEFT and HEFT variations

For SMEFT and HEFT, a similar reweighting approach is used; however, the reweighting is derived in a different way. To obtain the reweighting for SMEFT, LO truth parton level samples are generated using MADGRAPH5_AMC@NLO with the SMEFT@NLO model [92] for a variety of values of the SMEFT coefficients. The ratio of each of these samples to the SM sample is then taken in bins of m_{HH} , similar to what was done for κ_λ . A complete second-order polynomial of the SMEFT coefficients is then fit to this ratio for each bin. This function can then be used to obtain a weight for any given combination of coefficient values. A similar function is used for HEFT; however, the coefficients are obtained from [1] where they are calculated at NLO. Taking advantage of the NLO HEFT cross-sections and additional LO HEFT cross-sections [1] and a SMEFT to HEFT translation [152], the SMEFT weights are scaled by the ratio of the NLO to LO HEFT cross-sections. Since the




Figure 6.2: Comparison of m_{HH} distributions obtained for $\kappa_\lambda=10$ using (a) the truth level linear combination and truth level simulation, and (b) fully simulated $\kappa_\lambda=10$ and reweighted SM for the 2018 data period.

chromomagnetic operator does not appear in HEFT at LO, the value of the corresponding SMEFT coefficient (c_{tG}) is ignored when performing the translation and applying these scalings. The final signal model is then obtained, for both HEFT and SMEFT, by using the corresponding function to reweight the fully simulated SM ggF sample as a function of truth m_{HH} . The performance of the SMEFT reweighting is checked by comparing to the ggF κ_λ reweighting and to the fully simulated $\kappa_\lambda=10$ sample using the equivalent c_H variation ($c_H = -19$, cf. Equation 2.7). These results are found to exhibit generally similar shapes, as seen in Figure 6.3. Perfect agreement is not expected for the SMEFT reweighting as it is derived at LO, while the full simulation $\kappa_\lambda = 10$ sample is simulated at NLO. Slight disagreement is observed for the HEFT reweighting in the last few bins; however this difference was found to have a negligible impact on the limits.

Linear combinations for κ_λ , κ_{2V} , and κ_V variations in VBF

While this reweighting method works well for ggF, the VBF acceptance effects cannot be modeled as well by m_{HH} because of the more complex kinematics, and so poor performance was found with a reweighting based setup. Instead, the set of basis samples listed in the previous section are fully simulated and then a linear combination is performed with these basis samples. The linear combination is derived using the same logic as the truth level ggF κ_λ linear combination; however, as there are 3 LO VBF diagrams there are 6 terms total and 6 basis samples are required. This linear combination is cross-checked using additional sam-




Figure 6.3: Comparison of m_{HH} distributions obtained for $\kappa_\lambda=10$ using fully simulated $\kappa_\lambda=10$ and SM reweighted to $c_H = -19.23$. Perfect agreement is not expected for the SMEFT reweighting, as it is derived at LO while the full simulation $\kappa_\lambda = 10$ sample is simulated at NLO.

ples generated at the following points $(\kappa_\lambda, \kappa_{2V}, \kappa_V) = (1, 0, 1), (1, 0.5, 1), (1, 2, 1), (1, 3, 1), (0, 1, 1), (0, 0, 1)$. Again, agreement is found to be good using this method as shown in Figure 6.4.

6.2 Trigger

This analysis uses a combination of multi- b -jet triggers which provide complimentary sensitivity across the m_{HH} spectrum. As discussed in Section 3.2.5, these triggers consist of an L1 trigger decision and an HLT decision. Due to the high-throughput required at L1, the initial trigger decisions can use only simple information like jet p_T and η , leading to lower acceptance. At the HLT stage, where more time is available, b -tagging is run using an online implementation of the MV2 algorithm (Section 5.3.2). This additional b -tagging informa-




Figure 6.4: Comparison of reconstructed m_{HH} distributions for fully simulated VBF samples and those constructed using the linear combination method for (a) $\kappa_{2V}=0$ and (b) $\kappa_\lambda=0$. Note that $\kappa_\lambda = 0$ is used here as a comparison because $\kappa_\lambda = 10$ is part of the basis and so the linear combination yields a perfect match in that case.

tion allows for higher efficiency trigger decisions to be made. The HLT trigger requirements used for this analysis are as follows:

2b + 2j requires two b -tagged jets and two additional jets, all with $p_T > 35 \text{ GeV}$

2b + 1j requires two b -tagged jets with $p_T > 55 \text{ GeV}$ and one additional jet with $p_T > 100 \text{ GeV}$ (in 2016) or 150 GeV (in 2017 and 2018)

2b + HT (2017 and 2018 only) requires two b -tagged jets with $p_T > 55 \text{ GeV}$ and that the value of the scalar sum of p_T across all jets in the event (H_T) is $> 300 \text{ GeV}$

1b requires one b -tagged jet with $p_T > 225 \text{ GeV}$ (in 2016) or 400 GeV (in 2017 and 2018)

All of these triggers were used in the resonant analysis; however, only the first two triggers were used in the non-resonant analysis as the additional triggers were found to contribute minimal sensitivity to these non-resonant signals.

6.2.1 Trigger buckets

While the use of multiple triggers provides improved sensitivity, it also complicates the application of the trigger scale factors used to account for differences in efficiency between MC and data. This complication arises because trigger efficiencies are calculated at the jet-level while scale factors must be applied at the event-level. These event-level scale factors

are calculated from the jet-level scale factors for each jet in the event, while accounting for the possibility that the same event has passed multiple triggers. To simplify this procedure, a bucketing strategy is used. This bucketing procedure splits events up into orthogonal categories by applying a series of offline selections mimicking the original trigger decisions. The events falling in each category are assigned to the trigger that category was designed to mimic. The corresponding scale factor is then applied to all events which pass the assigned trigger. Events which do not pass the assigned trigger are discarded. The precise cuts and variables used in the trigger selections were chosen to minimize the amount of discarded signal. The buckets used for the resonant and non-resonant analysis are shown in Figure 6.5 and Figure 6.6, respectively.




Figure 6.5: Trigger bucketing procedure for the resonant analysis.




Figure 6.6: Trigger bucketing procedure for the non-resonant analysis.

6.2.2 Trigger scale factors

Due to the rate requirements of the trigger system and the more static nature of trigger software, the jets reconstructed by the trigger (online jets) are not identical to those reconstructed by the downstream analysis software (offline jets). To resolve this discrepancy, offline jets are matched to the closest online jet within $\Delta R < 0.4$. Events are discarded if the matched jets do not pass the online b -tagging requirements. Per-jet b -tagging scale factors are then calculated for each offline jet using this matching information along with the online, conditional, and offline MC and data efficiencies [153]. The per-event b -tagging scale factor for each event is taken as the product of all of the per-jet b -tagging scale factors of the jets in that event. The full event scale factors are calculated as the product of these per-jet b -tagging scale factors are the normal kinematic jet scale factors described in Section 5.2.3. These scale factors are used as event weights for the MC events used throughout all remaining stages of this analysis.

6.3 Object definitions

Throughout this section, several important object definitions will be used, these are listed below:

Central jet A jet reconstructed using the anti- k_t algorithm with $R = 0.4$ and which satisfies the following cuts: $|\eta| < 2.5$ and $p_T > 40 \text{ GeV}$. These values are chosen to ensure the jet falls within the Inner Detector (necessary for b -tagging) and to avoid being too far onto the trigger turn-on curve.

Forward jet A jet reconstructed using the anti- k_t algorithm with $R = 0.4$ and which satisfies the following cuts: $|\eta| \geq 2.5$ and $p_T > 30 \text{ GeV}$. These values are chosen to improve the acceptance of the VBF process.

b -jet A jet which is b -tagged by the DL1r algorithm at the 77% efficiency working point.

Loose b -jet A jet which is b -tagged by the DL1r algorithm at the 85% efficiency working point, but not the 77% working point.

Higgs Candidate (HC) Represents a potential Higgs boson reconstructed from a pair of central jets. The jet pairings are selected by the analysis pairing algorithm.

VBF jets The pair of un- b -tagged jets with the largest reconstructed invariant mass (m_{jj}). These jets are only considered for the non-resonant analysis.

6.4 Event selection and reconstruction

In addition to the trigger bucket selections, a series of analysis selections are applied to reduce background while ensuring that HH signal events are included and correctly reconstructed. The steps in the event selections for each analysis are outlined in brief below, with further details in the following sections. Flowcharts detailing the event selection procedures are shown in Figure 6.7 (resonant) and Figure 6.8 (non-resonant).

Resonant event selection

1. Events must pass the trigger selection detailed in Section 6.2.1
2. Events must contain ≥ 4 central jets
3. At least 4 of these central jets must be b -tagged
- 4–5. Must pass top-veto, HC must be sufficiently collimated (Section 6.4.4)
6. Events must be in the kinematic signal region, $X_{HH} < 1.6$ (Section 6.4.5)




Figure 6.7: Event selection flowchart for the resonant analysis.

Non-resonant event selection

1. Events must pass the trigger selection detailed in Section 6.2.1
2. Events must contain ≥ 4 central jets
3. At least 4 of these central jets must be b -tagged
- 4–6. Split ggF and VBF channels with cuts indicating VBF topology (Section 6.4.2)
- 7, 8. **ggF** Must pass the top-veto; ggF HC must be sufficiently collimated (Section 6.4.4)
8. **VBF**, 9. **ggF** Events must be in the kinematic signal region, $X_{HH} < 1.6$ (Section 6.4.5)




Figure 6.8: Event selection flowchart for the non-resonant analysis channels.

6.4.1 b -tagging categories

Following the requirement of at least 4 central jets, both the resonant and non-resonant analysis place a cut on the number of central b -jets. This cut is natural when targeting the $b\bar{b}b\bar{b}$ topology and provides significant background reduction. However, additional b -tagging categories are also defined by altering this cut to allow for background rich control regions. In particular, the resonant analysis uses two different b -tagging categories: 2b and 4b. The non-resonant analysis uses these same two categories, plus a 3b1f category. The purpose and definition of each of these categories is further outlined in Table 6.1 and in the following paragraph.

Table 6.1: The relevant b -tagging region definitions and corresponding purpose within the analysis. The 2b and 4b are used in both the resonant and non-resonant analyses. The 3b1f region is used only for the non-resonant analysis. Here b -jets are defined as those which Loose b -jets are defined as those which fail the 77% efficiency working point and pass the 85% working point while,

Notation	Definition	Usage
2b	Exactly 2 b -jets	Background estimation and validation
4b	At least 4 b -jets	Signal region
3b1f	Exactly 3 b -jets and no loose b -jets	Background validation and systematic uncertainty derivation (non-resonant)

The 4b category, as already described, is generally more signal rich and includes the 4b kinematic signal region where final fits are performed. The 2b category is background-rich and is used for background estimation (Section 6.5). The 3b1f region is also background-rich, but has kinematics which more closely resemble the 4b category due to an increased presence of heavy flavor jets. The 3b1f region is used for validation and for deriving a systematic uncertainty Section 6.7.3. Events are discarded if they do not match any of these b -tagging criteria.

Note that events falling within the various b -tagging categories will be generally referred to in short form, for example: “2b data” for data events falling in the 2b category or “4b signal” for signal events falling in the 4b category. Additionally, the kinematic regions used in this analysis are defined for each b -tag category. That is, there will be a 4b signal region, a 3b1f signal region, and a 2b signal region; only the former of which is actually signal-rich.

6.4.2 Non-resonant ggF and VBF channels

The non-resonant analysis is split into two orthogonal channels with one optimized for VBF production and one for ggF production. These channels are defined by the following series of three cuts:

- Must have a total of ≥ 6 central and forward jets.
- The VBF jets must have $m_{jj} > 1 \text{ TeV}$ and $|\Delta\eta_{jj}| > 3$.
- The p_T of the sum of the two VBF jets and remaining 4 HC jets is $< 65 \text{ GeV}$.

Events which pass all of these cuts are placed in the VBF channel, while other events are placed in the ggF channel. These cuts are all characteristic of the VBF event topology. To understand the first two cuts, first note that VBF events typically include at least two additional high energy jets from the outgoing quarks which first emitted the vector bosons. Thus, at least 6 jets are required. The jets considered are expanded to include forward jets because the quark jets are typically more longitudinally separated and allowing for forward jets improves acceptance of these jets. These two high energy outgoing quarks are expected to correspond to the reconstructed VBF jets, hence the VBF jet cuts (high energy and well separated). The final cut is accounts for the larger expected p_T of the VBF system relative to the ggF system. The exact values used for the cuts were optimized to maximize signal sensitivity in the VBF channel and were cross-checked and found to have minimal impact on the ggF channel.

Following this series of cuts, any central VBF jets are removed from the set of central jets used for the reconstruction of the HH system and of the top veto discriminant, X_{Wt} .

6.4.3 Higgs candidate pairing

Following all of the above requirements, four central jets are selected for each event. First, the highest p_T b -jets for each category are selected. For 3b1f and 2b, the highest remaining p_T central jets are then selected (1 or 2, respectively) to obtain a total of four jets. These selected jets are paired to form HCs using a dedicated pairing algorithm.

Given any such set of 4 jets there are three possible pairings of those jets. The purpose of the pairing algorithm is to determine the pair most likely to correspond to the two jets produced by the decay of a real Higgs boson. The rate at which jets are paired correctly to form a HC can be determined from the truth information stored with MC simulation samples. However, it is not sufficient to optimize on this pairing efficiency alone. One must also account for the many data events which do not contain 4 selected central jets

coming from two real Higgs bosons. For these events there is no true pairing. This is critical because these events are severely sculpted by many naive pairing strategies, such as choosing the pairing with an invariant mass closest to 125 GeV. Such sculpting would lead to many background events having reconstructed invariant masses which resemble signal, causing a loss in sensitivity and making the background more difficult to estimate. Thus, when determining a pairing algorithm, the degree of background sculpting must also be considered.

Resonant pairing algorithm

The resonant analysis makes use of a Boosted Decision Tree (BDT) based pairing algorithm. This BDT takes as input the total angular separation, pseudo-rapidity separation, and azimuthal separation between the jets in each pair. The BDT is also parameterized by the truth value of the m_{HH} of the pairings. The BDT is trained on all of the correct and incorrect pairings for one quarter of the total AFII simulated scalar MC events. Since all correct and incorrect pairings are used in the training, any given truth m_{HH} value will have both correct and incorrect pairs (1:2 ratio) meaning that it is not possible for the BDT to determine the correct pairing based on this information alone. However, being provided this truth information will allow the BDT to fit different functions for different m_{HH} values. To improve training statistics for low m_{HH} events, the BDT training used a looser cut of $p_T > 35$ GeV for the four central jets. In order to have reliable pairing information available, the events used for training were also required to have four distinct jets that could be geometrically matched to the truth b -quarks. The BDT was trained using the Gradient-based One Sided Sampling (GOSS) algorithm and the hyperparameters were optimized using a Bayesian optimization procedure [154] and three-fold cross-validation. Events used in the training were removed from the samples used for the later stages of the analysis, to avoid bias.

Non-resonant pairing algorithm

While the resonant analysis BDT pairing focuses on accurately reconstructing the full spectrum of m_{HH} , the non-resonant analysis uses a much simpler pairing method to obtain a smooth unsculpted background distribution at the cost of lower efficiency at low m_{HH} . This sacrifice is acceptable because the heavier tails of the non-resonant distributions allow this sensitivity to be regained through better modeling of the mid to high m_{HH} regime. To achieve all of this, the non-resonant pairing exploits the fact that the decay products of the Higgs should be more collimated due to the Higgs initial momentum. The procedure itself is as follows:

1. For each pairing, identify the pair with the largest p_T as the leading Higgs Candidate for that pairing.
2. For each pairing, calculate the ΔR between the constituents of the leading Higgs candidate.
3. Select the pairing for which this leading Higgs Candidate ΔR is the lowest.

Higgs Candidate mass planes

The joint distribution of the reconstructed masses of these newly paired HCs defines a mass plane, as shown in Figure 6.9. The mass plane is particularly useful for distinguishing true HH signal from background because the signal is expected to peak near:

$$(m_{H1}, m_{H2}) \sim (125 \text{ GeV}, 125 \text{ GeV}),$$

where m_{H1} and m_{H2} are the reconstructed masses of the leading and sub-leading Higgs candidates, respectively. Contrastingly, the dominant QCD background would be expected to peak at low masses. In Figure 6.9a and Figure 6.9b, the background-rich 2b data is seen to peak at low masses for the non-resonant ggF and VBF channels, as expected. However, this is not true for the resonant analysis, where data sculpting caused by the BDT pairing algorithm leads to a large peak close to the true Higgs mass.

6.4.4 Background reduction

For further background reduction, additional cuts are placed which make use of these new HCs and their constituent jets. First, to remove QCD background events are required to have HCs with a maximum pseudorapidity separation of $|\Delta\eta_{HH}| < 1.5$. This cut targets the more separated topologies for which the ggF channels do not expect significant signal and is applied in the resonant analysis and in the non-resonant ggF channel. The cut is not applied in the non-resonant VBF channel because it would remove significant sensitivity to the more separated topologies expected for some VBF signal models (e.g $\kappa_{2V} = 0$) as will be shown in Section 6.6.1. Figure 6.10 shows the distributions of $|\Delta\eta_{HH}|$ for 4b data and MC signal and the applied cut in the resonant analysis and non-resonant ggF channel. The value of this cut was optimized to maximize sensitivity and can be seen to significantly suppress background with a far lesser reduction to the signal.

Next, to suppress background contributions from hadronic $t\bar{t}$ decays a dedicated top veto is applied. This veto removes events which are consistent with the dominant top quark




Figure 6.9: Mass planes of 2b data from the non-resonant (a) ggF channel, (b) VBF channel, and (c) the resonant analysis. A mass plane is defined as the two-dimensional distribution of Higgs Candidate (HC) masses. Here the leading HC mass (m_{H1}) and sub-leading HC mass (m_{H2}) are displayed on the horizontal and vertical axes, respectively.

decay process in which the top quark decays to a b -quark and a W . To do so, a W candidate is reconstructed from each possible pair of central jets within the event, including those not selected for HC reconstruction. For each of these W candidates, a top candidate is reconstructed for each of the remaining b -jets (non-resonant) or jets (resonant) not used in that W candidate. The top veto discriminant, X_{Wt} , is then defined for each event:

$$X_{Wt} = \min_{t \text{ cand.}} \sqrt{\left(\frac{m_W - 80.4 \text{ GeV}}{0.1 m_W}\right)^2 + \left(\frac{m_t - 172.5 \text{ GeV}}{0.1 m_t}\right)^2} \quad (6.4)$$




Figure 6.10: Distribution of $|\Delta\eta_{HH}|$ for 4b MC signals and data for (a) the resonant analysis and (b) the non-resonant ggF channel. The ggF non-resonant distribution includes only that 4b data falling outside the signal region (thus, referred to as blinded in the caption). The resonant distribution includes the 4b data from the 2017 validation region (defined in Section 6.4.5). A vertical line indicating selection cut of $|\Delta\eta_{HH}| < 1.5$ is shown on both figures. Events to the right of these lines are discarded.




Figure 6.11: Flowchart for the derivation of the top veto discriminant X_{Wt} . Note that for the resonant analysis the additional jet used to reconstruct the top candidate is not required to be a b -jet.

where m_W and m_t are the reconstructed masses of the W and t candidates. A flowchart showing the steps in this derivation is shown in Figure 6.11. Events are vetoed for $X_{Wt} < 1.5$. X_{Wt} distributions for $t\bar{t}$ MC, and signal MC are shown in Figure 6.12. The additional requirement that the jet used in the top-candidate be $\bar{t}$ -tagged was added for the non-resonant analysis as it was found to reduce the amount of systematic mismodeling in the background




Figure 6.12: Distribution of X_{Wt} for $t\bar{t}$ and signal MC for the non-resonant (a) ggF and (b) VBF channels, and for (c) the resonant analysis. The non-resonant distributions include all simulated 2018 4b events passing prior cuts. The resonant analysis distributions include simulated events from the 2017 4b signal region (defined in Section 6.4.5). A vertical line indicating selection cut of $X_{Wt} > 1.5$ is shown on both figures. Events to the left of these lines are discarded.

estimate.

6.4.5 Kinematic region definitions

For a true signal event, it is expected that the reconstructed masses of the HCs will be close to the true mass of the Higgs boson. Given this fact, we may define a series of analysis regions in the mass plane which include either the signal-rich area around this point or the background rich regions further from it. Due to the substantially different pairing algorithms used in the resonant and non-resonant analysis, the precise structure of the mass plane varies.

These two analyses therefore proceed in slightly different ways.

Signal region

For both the non-resonant and resonant analyses, a discriminant, X_{HH} , is defined that quantifies the difference between the expected HC masses for signal and the actual HC masses for that event. X_{HH} can be written as a function of the leading and sub-leading HC masses as follows:

$$X_{HH} = \sqrt{\left(\frac{m_{H,1} - m_{H,1}^{(0)}}{0.1 m_{H,1}}\right)^2 + \left(\frac{m_{H,2} - m_{H,2}^{(0)}}{0.1 m_{H,2}}\right)^2} \quad (6.5)$$

where $(m_{H,1}^{(0)}, m_{H,2}^{(0)}) = (124 \text{ GeV}, 117 \text{ GeV})$ for the non-resonant analysis and $(120 \text{ GeV}, 110 \text{ GeV})$ for the resonant analysis. These central masses are not the truth level $(125 \text{ GeV}, 125 \text{ GeV})$ that might be expected, but rather are slightly shifted. This shift accounts for energy loss in the reconstruction procedure, with the precise values determined from the median masses of the corresponding signal distributions. The kinematic signal region (SR) for both analyses is defined to be the region for which $X_{HH} < 1.6$. This cut forms an egg-shaped region within the mass plane and as shown in Figures 6.14 and 6.13 is found to include the majority of the signal peak. For both analyses, several alternative signal regions were explored. In particular, regions defined from contours of the signal significances were tested and found to improve the stat-only limit and signal region cut efficiency. However, these regions are found to provide worse limits overall, due to an increase in systematic background mismodeling. This increase is driven by the inclusion of additional events from the more difficult to model lower mass region where structure is complicated by the X_{Wt} cut. This further motivates the use of the X_{HH} shape for the signal region definition as it precisely avoids this region of the mass plane. Of note, “signal region” will be used here strictly to refer to the kinematic region defined by this cut. As discussed in Section 6.4.1, there are multiple different b -tagging categories for each kinematic region and of these, only the 4b signal region (SR 4b) is signal-rich.

Non-resonant control regions

The non-resonant analysis defines a pair of control regions to be used for background estimation: These two control regions are denoted: Control Region 1 (CR1) and Control Region 2 (CR2). CR1 is used to derive the nominal background estimate while CR2 is used to derive an alternative model, which is used to derive a systematic uncertainty (Section 6.7.2). Since it is ideal for the background model to be derived in a region as kinematically similar




Figure 6.13: Regions used for the non-resonant VBF and ggF channels drawn over the 2b data and 4b MC signal mass planes for each channel.

to the SR as possible, these regions are defined from quadrants of a ring surrounding the signal region. The inner boundary of this ring is the outer boundary of the signal region (i.e. $X_{HH} = 1.6$) and the outer boundary of this ring is given by:

$$\sqrt{\left(m_{H,1}^2 - 1.05 \cdot 124 \text{ GeV}^2\right)^2 + \left(m_{H,2}^2 - 1.05 \cdot 117 \text{ GeV}^2\right)^2} = 45 \text{ GeV}. \quad (6.6)$$

The center of this ring is shifted 5% up and to the right to avoid the irregular structure of the mass plane in the lower left region caused by the X_{Wt} cut (compare Figure 6.9a and Figure 6.13). The ring is then subdivided into quadrants by two perpendicular cuts which cross at the center of the SR. This quadrant splitting allows the CR1 and CR2 regions to be

equally close to the center, while also containing events further from the SR to reduce signal contamination. The opposing quadrants are paired rather than the adjacent quadrants to ensure that the m_{HH} distributions are more consistent with that of the SR. Due to the X_{Wt} cut, there are two bands within which the number of events is lower and the flavor composition can be slightly different. These bands can be seen as light strips in the mass plane stretching vertically and horizontally at $m_{H,i} \sim 80 \text{ GeV}$ on each axis. To ensure that these regions contribute equally to both CR1 and CR2, the quadrant cuts are made at an angle of 45° relative to the $m_{H,1}$ axis. This set of regions was found to improve performance in the alternative validation regions which will be discussed in Section 6.7. Signal contamination within these regions control regions is found to be consistently below 25% of the SR sensitivity, a level which has negligible impact on the background model.

Resonant control and validation region

In addition to the shifted (120 GeV, 110 GeV) SR center, the resonant analysis takes an alternative approach to defining additional regions. In this case, a kinematic control region (CR) and validation region (VR) are defined by a pair of circles which enclose the signal region. The CR is used to derive a background model for the SR, while the VR is used to estimate an extrapolation uncertainty. These two circles are given by the following two equations:

$$C_o = \sqrt{\left(m_{H,1}^2 - 1.05 \cdot 120 \text{ GeV}\right)^2 + \left(m_{H,2}^2 - 1.05 \cdot 110 \text{ GeV}\right)^2} \quad (6.7)$$

$$V_o = \sqrt{\left(m_{H,1}^2 - 1.03 \cdot 120 \text{ GeV}\right)^2 + \left(m_{H,2}^2 - 1.03 \cdot 110 \text{ GeV}\right)^2} \quad (6.8)$$

The VR is defined to be the region between the SR boundary and the inner circle given by the constraints $X_{HH} \geq 1.6$ and $V_o < 30$. The CR is then defined as the region between the two circles, satisfying the constraints $V_o \geq 30$ and $C_o < 45$. The shifts are chosen to approximately maintain first and second moments of the SR m_{HH} distribution. These region definitions originated with the previous analysis [143]. A modified set of regions based on a simple significance clustering algorithm and a slight shift to the CR and VR was tested and found to yield a 5% improvement as measured by the simple significance ($S/\sqrt{B}$). Since this gain was relatively small and discovered late in the analysis process, the original analysis region definitions were kept for the sake of simplicity.




Figure 6.14: Regions used for the non-resonant VBF and ggF channels drawn over the 2b data and 4b SM MC signal mass planes for each channel.

6.4.6 Signal acceptance and cut-flows

The acceptance times efficiency for each stage in the analysis selection is shown in Figure 6.15. Acceptance times efficiency is lower at low masses primarily due to trigger limitations (see Section 6.2), while for higher masses jets begin to merge and the jet cuts begin to fail. These effects can be observed in the non-resonant ggF analysis, where the SM signal has a higher acceptance times efficiency than the $\kappa_\lambda=10$ variation, because the latter is softer. In the resonant analysis, this effect is particularly pronounced for scalar signals because unlike the wider non-resonant and spin-2 signals, the scalars do not have a tail up towards moderate m_{HH} where acceptance is higher. The sharp increase in acceptance times efficiency for particularly low mass resonances is a by-product of the kinematic threshold imposed by the requirement that there must be at least 250 GeV to produce two on-shell Higgs bosons.

6.4.7 Invariant mass correction

In order to correct for some of these aforementioned selection effects and energy loss during reconstruction, the resonant analysis performs a correction to the invariant mass of the HCs. To calculate this correction, the HC four-momenta are scaled to match the reconstructed mass of the HC to 125 GeV. The sum of these scaled HC four-momenta is taken to obtain the four-momenta of the HH system, with the invariant mass of this system defined as the corrected m_{HH} . The rescaled HC momenta obtained during this procedure are not stored or used in other analysis calculations. This correction is found to significantly improve




Figure 6.15: Acceptance times efficiency for the (a) resonant analysis spin-0 and (b) spin-2 signal models as a function of the mass, and the (c) non-resonant signals and data in the ggF channel and (d) VBF channel.

the resolution of the resonant signal distributions while also re-centering them. Figure 6.16 shows a comparison of the distributions obtained for 280 GeV, 400 GeV, and 700 GeV scalar signal samples before and after applying this correction.

This correction was also tested for the non-resonant analysis, but it was not found to give a substantial improvement. The non-resonant analysis therefore uses the original reconstructed m_{HH} rather than this corrected version.




Figure 6.16: Effects of the mass correction on the m_{HH} distributions of 280 GeV, 400 GeV, and 700 GeV scalar signal samples.

6.5 Background estimation

Following event selection, there are two major sources of background: QCD multi-jets processes and $t\bar{t}$, with expected QCD contents of $\sim 80\%$ for ggF non-resonant, $\sim 90\%$ for VBF non-resonant, and $\sim 95\%$ for ggF resonant. QCD multi-jets processes are particularly difficult to model in simulation, this analysis therefore uses a fully data-driven background estimation procedure. This procedure sacrifices the distinction between the QCD and $t\bar{t}$ background components, but has been validated and found to model the $t\bar{t}$ component correctly and to yield a better model than using a mixture of $t\bar{t}$ MC and data-driven methods.

The background estimation method used for this analysis exploits the consistent kinematic differences seen between 2b and 4b data across local regions of the mass plane. In particular, a reweighting function serves as a mapping which morphs 2b data kinematic distributions to match the 4b kinematic distributions, including normalization. The nominal mapping for the non-resonant (resonant) analysis is derived in the CR1 (CR) for resonant,

while an alternate mapping is derived from the CR2 (VR) and used to derive a systematic uncertainty.

Due to the different analysis region definitions, the alternate mappings and corresponding systematic uncertainties for the resonant and non-resonant analysis have slightly different interpretations and cover slightly different systematic biases. In the resonant analysis, the alternative estimate is derived closer to the SR than the nominal estimate and so the systematic bias measured here is that due to extrapolation. In the non-resonant analysis, the orthogonal nature of the CR1 and CR2 quadrants, means that the alternative mapping is interpolated into the SR along an orthogonal axis and so measures the systematic bias arising from the different possible interpolations into the SR. The precise derivation and validation of these systematic uncertainties is discussed in Section 6.7.

6.5.1 Methodological primer

This reweighting procedure can be considered a more sophisticated version of the traditional data-driven ABCD method, as shown in Figure 6.17. In the ABCD method, one assumes that the ratio of some quantity (typically the yield) is equal in two pairs of regions: A and B , and C and D . The problem is typically framed such that the 4 regions can be considered to form a set of quadrants along a pair of axes. Taking these four regions to be oriented as in Figure 6.17, the assumption that $A/B = C/D$ corresponds to the assumption that changes along the vertical axis are independent of changes along the horizontal axis. This relation can then be used to estimate the unknown quantity of D from the known quantities of A , B , and C with the equation $D = C(B/A)$.

Now, as mentioned the reweighting procedure used here can be taken as the generalization of the ABCD method to the case of a *density* ratio (or more generally, some mapping). The regions A , B , C , and D correspond to $CR2b$, $CR4b$, $SR2b$, and $SR4b$, respectively. The assumption made is that the 4b/2b density ratio does not change as one shifts from the control region to the signal region. The resulting procedure can then be written (very roughly) as $SR4b = SR2b(CR4b/CR2b)$. Of note, a simple rearrangement of the above equation gives: $SR4b = CR4b(SR2b/CR2b)$. And so one would expect that if the methodology outlined in this way is valid, then one should be equally able to derive the SR/CR density ratio in the 2b regions and apply it in the 4b regions. In fact, empirical studies confirmed that this can be done and the resulting predictions are statistically equivalent. However, deriving the 4b/2b density ratio is preferred because the statistics in the 2b region are much larger ($n_{2b} \sim 100n_{4b}$). Splitting the two higher statistics regions ($SR2b$ and $CR2b$) between the training and prediction steps provides a more balanced use of the statistics and yields a final




Figure 6.17: Pictorial diagram of the reweighting procedure, represented as a traditional ABCD method. The light turquoise squares indicate unblinded background-rich regions (CR 2b, CR 4b, and SR 2b) while the dark pink square indicates the blinded signal region (SR 4b). The black arrow indicates that a function is derived for the density ratio using the 2b and 4b data in a control region while the grey arrow indicates the application of this density ratio to 2b data to estimate 4b data.

prediction with a lower statistical uncertainty.

6.5.2 Neural network reweighting methodology

As described above, the background estimation method works by deriving a reweighting function to map 2b data to match 4b kinematic distributions. Naively, one might attempt to derive this reweighting function by deriving weights in bins of the relevant kinematic variables, as is done for the truth- to reconstruction-level ggF signal models. Unfortunately, such an approach encounters several problems for the background estimate. Firstly, because of the relative size of the background, to ensure strong limits the background estimate is

required to have much better relative precision. In addition to this, the relation between 2b and 4b data is significantly more complex, showing differences in many variables and correlation effects. One could attempt to bin further in these additional variables to account for correlations; however, this leads to the curse-of-dimensionality. That is, phase space volume bins becomes exponentially less populated as the dimension increases. One is then forced to choose between including the correlation effects and using a more coarse binning which loses resolution or using a fine binning which will lead to overfitting and instability due to the low statistics. A BDT was used in a previous related analysis [155]; however, while BDTs provide a more adaptive binning, they still rely on a series of sequential tree decisions which split the phase-space into bins meaning that after several branches the statistics rapidly become limited. Thus, the depth of the tree is limited by the statistics which decreases the capacity of the BDT and consequentially the flexibility of the learned reweighting function.

Unlike these aforementioned methods, Neural Networks (NNs) perform truly multivariate and continuous calculations allowing them to directly learn the reweighting function. This allows them to bypass the binning issue and treat the problem as one of multivariate density estimation. In particular, the neural network learns the density ratio of 4b and 2b data:

$$w(x) = \frac{p_{4b}(x)}{p_{2b}(x)} \quad (6.9)$$

where $p_{4b}(x)$ and $p_{2b}(x)$ are the probability densities for 4b and 2b data across some input variables x . The approach used here is modified from those described in [156, 157] and makes use of the following loss function:

$$\mathcal{L}(Q(x)) = \mathbb{E}_{x \sim p_{2b}} [\exp(Q(x)/2)] + \mathbb{E}_{x \sim p_{4b}} [\exp(-Q(x)/2)] \quad (6.10)$$

where $Q(x)$ is an estimator dependent upon x and $\mathbb{E}_{x \sim p_{ib}}$ is the expectation value with respect to the 2b ($i = 2$) or 4b ($i = 4$) probability density. The use of an exponential encoding is used here to avoid the need to enforce direct constraints on positivity of the predicted function. Negative weights are undesirable because they could lead to unphysical negative background predictions. We note that the estimator which minimizes this loss can be found to satisfy the following relation:

$$Q_{\min}(x) \equiv \arg \min_Q \mathcal{L}(Q(x)) = \log \frac{p_{4b}(x)}{p_{2b}(x)} \quad (6.11)$$

and the desired reweighting function, the density ratio, can be obtained as $e^{Q_{\min}(x)}$.

The neural network is trained using variables which were determined to be sensitive to the differences between 2b and 4b data, excepting the discriminant variables themselves

(m_{HH} or corrected m_{HH}) to avoid potential biases or undesired fit correlations. These variables are listed in Table 6.2. The natural logarithm is used in some cases to reduce the dynamic range of those distributions exhibiting skewed power-law distributions (e.g. transverse momenta). The reweighting variables for the resonant and non-resonant ggF analyses are the same save for the change in X_{Wt} definition (Section 6.4) and the addition of a one-hot-encoding of the trigger bucket. One-hot encoding refers to the use of a group of binary variables, only one of which may be true, to represent each state of a (typically categorical) variable. The VBF analysis uses a different set of variables tuned to the VBF event kinematics. Prior to input to the neural network, all variables are standardized by subtracting the training sample mean and dividing by the standard deviation.

The neural networks for the ggF analyses consist of three densely connected hidden layers of 50 nodes each with ReLU activation functions and a single node linear output. The architecture used for the VBF channel is similar, but only 20 nodes are used due to the lower statistics this channel. These configurations were optimized and demonstrated good performance in the modeling of a variety of input and non-input variables, including (corrected) m_{HH} .

The precise performance of these neural networks after training can fluctuate due to the random initial conditions and the training statistics. To account for this variation between trainings and to produce a more stable estimate, an ensemble of 100 neural networks is used with each neural network trained on a different Poisson bootstrap resampling of the data [158]. The reweighted estimate for each of these networks is renormalized such that it matches the normalization of the target distribution (for that bootstrap). The nominal prediction for the resonant (non-resonant) analysis is then taken as the median (mean) prediction of the 100 resulting predictions. A final normalization is applied to correct the prediction to the un-resampled target distribution. This procedure, and its use for estimating the statistical uncertainty of the prediction, will be described further in Section 6.7.1. This approach was found to provide improved prediction stability and overall modeling performance.

The ensemble networks for the resonant (non-resonant) analysis are trained in the CR (CR1) region using the full set of 2b and 4b data. These networks are then used to predict a weight for each event in the background rich SR 2b region with these now weighted events used to estimate the background in the blinded SR 4b signal region. In addition, a separate ensemble of networks is trained in the VR (CR2) and the estimate of these networks is used to derive an uncertainty, detailed in Section 6.7.2. Kinematic differences exist between years due to the different data-taking conditions and triggers used. For this reason, the years are trained separately for the resonant and non-resonant ggF channels. For the lower statistics

Table 6.2: Set of input variables used for the reweighting in the ggF and VBF channels.
Note that the trigger bucket index was not used for the resonant analysis. .

ggF + Resonant	VBF
1. $\log p_T$ of the 4th leading HC jet	1. Maximum di-jet mass out of the possible pairings of the four HC jets
2. $\log p_T$ of the 2nd leading HC jet	
3. $\log \Delta R$ between the closest two HC jets	2. Minimum di-jet mass out of the possible pairings of the four HC jets
4. $\log \Delta R$ between the other two HC jets	3. Energy of the leading HC
5. Average $ \eta $ of the HC jets	4. Energy of the sub-leading HC
6. $\log p_T$ of the HH system	5. Second smallest ΔR between the jets in the leading HC (out of the three possible pairings for the leading HC)
7. ΔR between the two HCs	6. $\log X_{Wt}$,
8. $\Delta\phi$ between the jets in the leading HC	7. Average $ \eta $ of the HC jets
9. $\Delta\phi$ between the jets in the sub-leading HC	8. Trigger bucket index as a one-hot-encoder
10. $\log X_{Wt}$	9. Year index as a one-hot-encoder
11. Number of jets in the event	
12. Trigger bucket index as one-hot-encoder (not used for resonant analysis)	

non-resonant VBF channel, the years are trained inclusively because the increase in statistics outweighs the increased systematic bias from the kinematic differences.

6.5.3 Reweighting performance

The performance of this kinematic reweighting method in the training region (CR1 or CR) will be shown here for several important discriminating variables. Figure 6.18 shows the performance of the resonant analysis reweighting in corrected m_{HH} . Figures 6.19 and 6.20




Figure 6.18: Distributions of corrected m_{HH} (a) before and (b) after the CR derived reweighting for the resonant analysis in CR data.




Figure 6.19: Distributions of m_{HH} (a) before and (b) after the CR1 derived reweighting for the non-resonant ggF channel in CR1 data.

show the performance for the non-resonant ggF and VBF channel reweightings in m_{HH} . The reweighting is found to produce distributions which are in generally good agreement with the 4b data for most variables. Comparisons of the reweighted and unweighted distributions for additional variables can be found in Appendix A.




Figure 6.20: Distributions of m_{HH} (a) before and (b) after the CR1 derived reweighting for the non-resonant VBF channel in CR1 data.

6.6 Analysis categorization and binning

The remaining validation and statistical analysis examines the data in bins of corrected m_{HH} for the resonant analysis, and m_{HH} as well as additional “categories” (2-3 bins/variable) for the non-resonant analysis.

Resonant corrected m_{HH} binning In the resonant analysis, since the signals are narrow, it is important to maintain consistent resolution across the the corrected m_{HH} spectrum. For this reason the binning was optimized considering only the case of a consistent bin width. The optimization yielded 84 equal width bins from 250 GeV to 1450 GeV, with a bin width of 14.3 GeV. The last bin is treated as an overflow bin.

Non-resonant m_{HH} binning Contrastingly, the signals in the non-resonant analysis are broad and so a more flexible binning approach is used to adapt the bin size to the varying statistics across the m_{HH} spectrum. In practice, this is done by increasing the bin width for bins with lower statistics to ensure a statistical error in the background estimate that is consistently lower than 30%. To avoid over-optimizing the binning, the procedure used is a simple one known as a logarithmic binning, which uses only three free parameters. These three free parameters are the lowest bin edge, the upper threshold, and a parameter which determines the scaling rate, X_p . Beginning with the lowest bin edge, the second bin edge is set to $(100 + X_p\%) \times$ the lowest bin edge. This process is then repeated, with the edge used in the scaling always taken to be the one calculated during the previous step. This process

Table 6.3: Logarithmic binning parameters for the non-resonant ggF and VBF channels.

Category	Binning Parameters			
	Min [GeV]	Max [GeV]	X_p (%)	Precision [GeV]
ggF, All	280	950	9	1
VBF, low $ \Delta\eta_{HH} $	280	890	10	5
VBF, high $ \Delta\eta_{HH} $	290	1470	9	5

terminates when the calculated bin edge exceeds the upper threshold, with the previously calculated edge kept as the final edge. This yields a binning in which the bin edges are spaced logarithmically. To further maintain simplicity, bin edge values are rounded to the nearest 1 GeV for ggF or 5 GeV for VBF. The ggF channel includes underflow and overflow events in the first and last bin, respectively. The VBF channel includes overflow events in the last bin, but does not include an underflow bin due to poor modeling for m_{HH} less than the value of the lower bin edge. Different logarithmic parameters are used for the ggF and VBF channels, and in each category, as will be detailed in the following subsections.

6.6.1 Non-resonant categorization strategy

Beyond binning in the m_{HH} distribution, the non-resonant analysis makes use of further categorization to provide additional discriminatory power. The categorization variables are selected from the set of all common kinematic variables which satisfy several criteria. First and foremost, the variable should provide discriminatory power, that is the significance ($S/\sqrt{B}$) should vary significantly as a function of that variable. Second, the variable should be well modeled by the background estimation procedure. Variables which are distributed similarly in normalized 2b and 4b data are particularly well-motivated because the background estimation method (detailed in Section 6.5) is required to do less morphing. This means that these variables are simple to model and less likely to be subject to systematic uncertainties in the background modeling procedure. Finally, these variables should be similarly distributed within at least one (but as many as possible) of the additional validation region. This allows a consistent validation of these additional variables. One additional consideration during categorization is the statistics of the analysis. In a simple orthogonal bin categorization scheme like that used here, the number of events in each bin decreases exponentially with the number of categorization variables. This means that for the VBF channel, where statistics are lower, additional categorization variables or bins are required to provide even more




Figure 6.21: SR4b (a) $|\Delta\eta_{HH}|$ and (b) X_{HH} distributions for reweighted 2b background and scaled MC signal for SM and $\kappa_\lambda=10$ variations. The purple vertical lines represent the category cuts. Clear discrimination power can be seen from the varying ratios of signal and background in each sliced region of the distribution.

discriminatory power in order to provide a net gain in sensitivity.

ggF categorization

Keeping these criteria in mind, several variables were studied for the ggF channel categorization including X_{Wt} , several angular variables, number of jets, and X_{HH} . Of these variables, categorization in $|\Delta\eta_{HH}|$ and X_{HH} was found to yield the best sensitivity and modeling. Figure 6.21 shows that these variables lend discriminatory power, while the modeling before and after reweighting for these variables can be found in Section A.2.1. In Section 6.7, these variables are validated in alternative validation regions. Using these variables, the ggF channel is categorized into 3 evenly sized bins of $|\Delta\eta_{HH}|$ – $|\Delta\eta_{HH}| < 0.5$, $0.5 < |\Delta\eta_{HH}| < 1$, $1 < |\Delta\eta_{HH}| < 1.5$ – and into 2 bins of X_{HH} – $X_{HH} < 0.95$ and $X_{HH} > 0.95$. Evenly sized bins were chosen for $|\Delta\eta_{HH}|$ due to the flat background distribution while the X_{HH} boundary was selected to evenly distribute correctly paired signal events between the two bins. Taking all combinations of these bins yields a total of 6 categories. Since statistics and m_{HH} distributions are relatively similar across categories, the same m_{HH} binning is used for each category. Figure 6.22 shows the $S/\sqrt{B}$ for each bin and category.




Figure 6.22: SR4b invariant mass distributions in each ggF category for reweighted 2b data and scaled MC signal for SM and $\kappa_\lambda=10$ variations. The lower panel shows the $S/\sqrt{B}$ for each bin with the signal scaled by a value, α , as indicated in the legend of the upper panel.

VBF categorization

Following a similar procedure, categorization was tested for two variables: the invariant mass of the VBF jets (m_{jj}) and $|\Delta\eta_{HH}|$. Both variables were found to provide good discriminatory power; however, m_{jj} was not sufficiently well modeled to be reliably used as a categorization variable. For this reason, only categorization in $|\Delta\eta_{HH}|$ is used for VBF, with the two bins $|\Delta\eta_{HH}| < 1.5$ and $|\Delta\eta_{HH}| > 1.5$. The discriminatory power provided by this split on $|\Delta\eta_{HH}|$ is shown in Figure 6.23, while the modeling before and after reweighting can be found in Section A.2.2. Due to the varying length of the high m_{HH} tails in the two categories, different binning schemes are used for each with larger bins and a larger range for the high $|\Delta\eta_{HH}|$ category. The $S/\sqrt{B}$ for each bin and category can be found in Figure 6.24.

6.7 Uncertainties

This analysis uses a variety of nuisance parameters to parameterize the uncertainties associated with known biases in the underlying calibrations, methods, and objects. The largest two uncertainties are both associated with the background modeling procedure. The first of these is associated with the systematic bias from the derivation of the reweighting function outside of the SR and the subsequent interpolation/extrapolation. The second uncertainty accounts for the statistical variation of the neural network training. There are also additional smaller uncertainties associated with the mismodeling that exists in MC simulations used for




Figure 6.23: SR4b $|\Delta\eta_{HH}|$ distributions for reweighted 2b background and scaled MC signal for SM, $\kappa_{2V}=0$, and $\kappa_\lambda=10$ variations. The purple vertical line represents the category cut. Clear discrimination power can be seen from the varying ratios of signal and background on either side of the line.




Figure 6.24: SR4b invariant mass distributions in each VBF $|\Delta\eta_{HH}|$ category for reweighted 2b data and scaled MC signal for SM, $\kappa_{2V}=0$, and $\kappa_\lambda=10$ variations. The lower panels show the $S/\sqrt{B}$ for each bin.

the signal modeling and calibration. There are also Poisson statistical uncertainties which account for statistical variation in the SR 2b data used to predict the SR 4b background and in the MC simulation used to generate the signal models.

6.7.1 Statistical uncertainties

There are two sources of statistical variation for the background estimation procedure. The first of these is the Poissonian nature of the events collected in the SR 2b region and used to predict the SR 4b data. Once these events are weighted, the uncertainty in the counts of a given histogram bin, n_i , can be determined using the following formula:

$$\delta n_i = \sqrt{\sum_{j \in i} w_j^2} \quad (6.12)$$

where the sum is over all events j which fall into bin i , with w_j is the weight of the event. Note that for unweighted events ($w_j = 1$), this reduces to the usual $\delta n_i = \sqrt{n_i}$.

The second component of the uncertainty is that due to the limited training statistics and inherently stochastic nature of the neural network training.¹ In fact, since there is an order of magnitude more 2b data than 4b data and since the neural network training is relatively unstable, the dominant contribution to the statistical error arises from the training itself. A standard method for estimating such uncertainties is the bootstrap resampling technique. This method relies on the resampling of the original dataset, with replacement, to construct a series of statistically equivalent sets. These sets may be used to estimate the variance of estimators derived using the original sample. In this case, these samples are used to estimate the variance of the reweighting function (which is an estimator of the true density ratio) for any given set of input parameters. Since $\lim_{n \rightarrow \infty} \text{Binomial}(n, 1/n) = \text{Poisson}(1)$, for the large training sample used here the resampling procedure is approximately equivalent to assigning a set of weights drawn from a Poisson(1) to each element in the sample. This procedure is known as the Poisson bootstrap [158]. Thus, when training the ensemble of neural networks, for each network a different random set of weights drawn from Poisson(1) is assigned to each event. The initial neural network weights and optimizers are also given a different random initialization for each ensemble. The ensemble of trainings will then capture both the effects of the random initialization and those of the training sample itself. The variation in the predictions of this ensemble of neural networks is used to assign an uncertainty within the fit. The standard procedure for handling this ensemble would be to

¹This stochasticity arises from the stochastic batch gradient descent algorithm and from the standard procedure of randomizing the initial NN weights.

follow the full prediction procedure for each neural network. That is each network would be used to predict a weight for every SR2b event. These weighted events would then be used to form a histogram which represented the SR4b background estimate. Doing this for every network in the ensemble, one obtains 100 histograms which parameterize the variation in the background estimate. It is at this point where the treatment within the resonant and non-resonant analyses begins to diverge.

Resonant bootstrap uncertainty

The resonant analysis procedure treats these 100 variations as though they were uncorrelated across bins and so a bin-by-bin statistical uncertainty is assigned. Following the standard procedure, one would take the median prediction in each bin as the nominal prediction and some interquartile range (IQR) would be taken as the uncertainty. Unfortunately, such a prediction and corresponding uncertainty would not be generalizable across histograms in different variables or with different binnings. This means that a new set of histograms would frequently need to be regenerated and all original weights would need to be stored. Due to the large datasets being used, this procedure is thus impractical both due to computational time and storage.

These issues motivate the use of a few event level quantities (rather than 100) for the derivation of the predictions and uncertainties. However, since the median and IQR operations do not commute with summation, the median and IQR of histograms cannot be generated directly from the median and IQR of the event weights. Instead, we use an approach which has been empirically shown to produce results consistent with the histogram-level quantities. Note that while the mean operation does commute with summation, it is not as robust to outliers and is therefore not preferred. The IQR is also preferred over the standard deviation for this reason.

The components used in this derivation are as follows:

1. Bootstrap weight (w_n): weight predicted for a given event by a network on bootstrapped dataset n .
2. Bootstrap normalization correction (α_n): normalization factor for bootstrapped dataset n . This factor normalizes the predicted target distribution to the true target distribution for that bootstrapped dataset.
3. Median weight (w_M): median weight for a given event across all 100 datasets.
4. IQR of weights (w_I): IQR of weights, w_n , for a given event across all 100 datasets.

5. IQR of norms (α_I): IQR of the norms, α_n , for a given event across all 100 datasets.
6. Median normalization correction (α_M): normalization factor which normalizes the final prediction to the true target distribution for the original sample.

As discussed previously, the nominal estimate for a given weight is the product of the median weight and the median normalization correction, $\alpha_M \cdot w_M$. This prediction is found to agree with the histogram-level median to within $< 2\%$ percent with the agreement varying slightly with m_{HH} .

To estimate the histogram-level IQR, a set of varied weights is constructed: $w_M + w_I/2$. These weights are then used to construct the event-level IQR histogram. This event-level IQR does not correspond to the histogram-level IQR histogram because the weights for all events are being fluctuated up simultaneously, which leads to a correlated effect which is far too large in magnitude. However, it was found that this event-level IQR shares a very similar shape to the histogram-level IQR. It is thus used to give the “shape” component of the IQR, that is, the relative magnitude of the IQR within each bin as a fraction of the total expected fluctuation in the yields. To isolate the shape component of this histogram relative to the nominal, it is re-normalized to the nominal prediction and then the nominal histogram is subtracted from it. To estimate the total magnitude of the fluctuations, we rescale the nominal histogram by $\alpha_I/2$ and add this to the shape component to get the total uncertainty. This procedure can be summarized by the following equation:

$$\delta_{BS} = \left[\left(\sum \alpha_* w_M \right) \frac{H\left(w_M + \frac{w_I}{2}\right)}{\sum \left(w_M + \frac{w_I}{2}\right)} - H(\alpha_* w_M) \right] + \frac{\alpha_I}{2} H(\alpha_* w_M) \quad (6.13)$$

where $H(x)$ corresponds to the histogram of x and the sums are performed over all events (not just those in the given bin). The square brackets above correspond to the shape component while the second term corresponds to the correction to the magnitude. This uncertainty is referred to as the bootstrap uncertainty. The components of this uncertainty as well as a comparison of the final uncertainty to the histogram-level IQR is shown in Figure 6.25. The uncertainty is found to be in good agreement with the histogram-level IQR, the largest deviations occur in the tail where statistical effects are dominated by the normal Poisson uncertainty.

Non-resonant bootstrap uncertainty

Unlike the resonant analysis, the non-resonant analysis uses the simpler, but more computationally method of computing the uncertainty directly from the set of 100 predicted




Figure 6.25: Sequence of plots depicting the approximate construction of the bootstrap error band, with each component shown as a ratio to the nominal estimate. The light grey lines indicate the predictions of each neural network in the ensemble. Figure 6.25a shows the event-level IQR histogram (blue) constructed from $w_M \pm w_I$, which approximately matches the shape of the true variation in the individual estimates (grey). Figure 6.25b shows the normalization correction (blue) given by $\pm \alpha_I H(\alpha_* w_M)/2$ which approximately matches the magnitude of the variation in the individual estimates (grey). Figure 6.25c shows the final band (solid black) produce from these two components and Figure 6.25d shows a comparison of this band (solid black) to the true histogram-level IQR (dashed black).

histograms. This change was made because the approximation used in the resonant analysis was no longer found to be accurate, particularly for the VBF channel. In addition, since the yields in each bin were determined to be Gaussian distributed and to ensure consistency with the Gaussian nuisance parameter which parameterizes this uncertainty, the standard deviation is used instead of the IQR. To summarize, the non-resonant bootstrap uncertainty is taken as an uncorrelated bin-by-bin uncertainty determined by the standard deviation of the 100 histograms predicted by the NN ensemble. The resulting uncertainties for the VBF and ggF channels are shown in Figure 6.26




Figure 6.26: Bootstrap uncertainties (pink) for the non-resonant (a) ggF and (b) VBF channels. These uncertainties correspond to the standard deviations of the 100 predicted histograms (light grey) obtained from the neural network ensemble.

6.7.2 Background systematics from alternative models

To account for the systematic bias associated with deriving the reweighting in a control region and extrapolating to the signal region, an alternative background model is derived in an additional region. A corresponding uncertainty is derived from this alternative model by taking the difference between the alternative and nominal model predictions in the SR. This difference is then symmetrized about the nominal prediction to obtain a two-sided shape uncertainty on the m_{HH} spectrum.

For the resonant analysis, this shape uncertainty is decomposed into two components to provide additional flexibility to the fit. In particular, it is split into low H_T and high H_T components. The low H_T uncertainty is determined by assigning the alternate weight only to those events with $H_T \leq 300$, the remaining events maintain the nominal weight. The difference with the nominal is then taken and symmetrized as described previously. For the high H_T uncertainty, the method proceeds similarly, but this time those events with $H_T > 300$ are assigned the alternative model weight while the remainder receive the nominal weight. This split effectively decouples the uncertainty for low m_{HH} and high m_{HH} as H_T is highly correlated with m_{HH} ; however, it avoids the sharp cut-off in the uncertainty that one would obtain with an m_{HH} based cut. The resulting uncertainties are shown in Figure 6.27.

The non-resonant analysis also decomposes the background shape systematic; however, in this case, the quadrants give a natural decomposition. Signal region quadrants are defined using the same axes as the control region quadrants. The four resulting regions, shown in




Figure 6.27: Decomposition of the background shape uncertainty for the ggF resonant analysis for 2018 SR 4b.




Figure 6.28: Visualization of the four SR quadrants used to define the uncertainty decomposition.

Figure 6.28 are named according to their relative positions in the mass plane: Q_N , Q_S , Q_E , Q_W . In this case, alternative estimates are obtained by applying the CR2 model weights only within one given quadrant while the CR1 weights are applied in the remaining three. By repeating this for each quadrant, we obtain four alternative estimates. As before, the difference with the nominal is taken and then symmetrized to obtain four uncertainties. These quadrants were chosen to align with the control region quadrants to allow the shape systematics to naturally follow the kinematic similarity of the adjacent control region quadrant. The resulting uncertainties are shown in Figure 6.29 for ggF and in Figure 6.30 for VBF.




Figure 6.29: Quadrant decomposition of the background shape uncertainty in the non-resonant ggF channel for 2018 data. Q_N , Q_W , Q_E , Q_S are defined from the four quadrants of the signal region shown in Figure 6.28.




Figure 6.30: Quadrant decomposition of the background shape uncertainty in the non-resonant VBF channel for the high and low $|\Delta\eta_{HH}|$ categories. Q_N , Q_W , Q_E , Q_S are defined from the four quadrants of the signal region shown in Figure 6.28.

6.7.3 Background modeling non-closure systematic

Good performance of the background estimate in the training region is not sufficient for understanding the performance of the estimate in the SR. While the uncertainties outlined above help to quantify the potential extrapolation uncertainty, it is still necessary to validate this uncertainty. To test the coverage of the previous uncertainties and to determine if an additional non-closure systematic is required, the non-resonant analysis repeats the background estimation procedure is repeated with 3b1f as the target category. That is to say, instead of estimating the 4b/2b density ratio, the 3b1f/2b density ratio is estimated. Estimating the 3b1f/2b density ratio serves as an excellent simulation of the true reweighting problem because the 3b1f region is background-rich, but the kinematics are much more similar to that of 4b than 2b due to the more similar flavor composition (Section 6.4.1). Additionally the 3b1f category has larger statistics than 4b, allowing a more precise determination of any mismodeling. Aside from the alternate target, the background estimation procedure is identical to that described for estimating 4b: the same 2b data and neural network architecture are used, the target SR 3b1f is left blinded, and uncertainties are calculated in the same way. After performing the full background estimation procedure for 3b1f, the background-rich SR 3b1f region is “unblinded” and compared to the background estimate. Any differences not covered by the existing uncertainties are used to derive a non-closure systematic.

This region is a more recent development and so was not used for validation or derivation of a non-closure systematic for the resonant analysis. For the non-resonant VBF channel, good agreement is found when comparing SR 3b1f data to the background estimate (Figure 6.31b). Therefore, no non-closure systematic is applied for the VBF analysis. For the non-resonant ggF channel, some non-closure is observed between the background estimate and the SR 3b1f data; therefore, an additional systematic is assessed for this region

Derivation of the 3b1f non-closure systematic

The non-closure systematic is computed bin-by-bin in each category by subtracting from the SR 3b1f yield the closest point on the 1-sigma band of the background estimate. That is, if the uncertainty band covers the difference, the difference is taken as zero, otherwise it is taken as the difference between the band and the uncertainty. To remove statistical fluctuations, the uncertainty is smoothed using a mean filter with a width of 3 bins, with boundaries treated as zero. Similarly to the interpolation systematic described in the previous section, this non-closure systematic is symmetrized to obtain a two sided template which is then treated as a shape systematic on the m_{HH} spectrum. The largest non-closure is found in 2018 and the corresponding non-closure systematic is shown in Figure 6.32.




Figure 6.31: Performance of the (a) ggF (2018 only) and (b) VBF (all years) reweighting model predictions vs SR 3b1f data. The VBF predictions are found to be in good agreement with the data while the ggF predictions exhibit a small non-closure.




Figure 6.32: Non-closure uncertainty (red) and ratio of SR 3b1f data to the reweighted 2b background estimate (black) in each category of 2018 data for the ggF non-resonant channel.

6.7.4 Signal uncertainties

There are many additional smaller systematic biases which arise from mismodeling in the MC simulations used at various steps of the analysis. Since the background method is fully data-driven these uncertainties are only relevant for the signal models. These uncertainties can be divided into two broad categories, those uncertainties arising from detector modeling and reconstruction and those arising from theoretical predictions.

Detector modeling and reconstruction uncertainties account for the differences observed between MC simulation and data due to mismodeling of the detector and inconsistent performance of reconstruction algorithms for data and MC simulation (e.g. cf. Section 5.3.3).

More specifically these uncertainties account for uncertainty in the measured luminosity, the jet energy scale and resolution, b -jet tagging efficiency, and trigger efficiency. Luminosity uncertainties account for the uncertainty in the measured luminosity values used to scale our predicted signal cross-sections to obtain expected number of events in data. The potential inaccuracy in this measurement is accounted for by an uncertainty of 1.7% [159]. There are many different jet energy scale and resolution uncertainties, which are treated as outlined in [133]. The b -jet tagging efficiency uncertainties follow the official prescriptions outlined in [136]. The trigger uncertainties account for variation in the scale factors used to estimate the trigger efficiency differences between MC simulation and data (Section 6.2). The variation in these scale factors is taken as the uncertainty, with the precise uncertainty depending on the trigger which was matched to that event.

Theoretical uncertainties arise as a by-product of the various assumptions and simplifications made during theory calculations. These effects include uncertainties in predicted cross-sections due to finite order calculations and varying PDFs, variations in parton shower algorithms, and uncertainties arising from the signal reweighting procedures used. Uncertainties in the matrix element were considered for the resonant analysis, but were found to be negligible as were uncertainties in the PDFs. For the non-resonant analysis, uncertainties in the cross-section are used, following the scheme prescribed by [160] for ggF and [72] for VBF. To measure the variation in the signal model arising from parton showering, samples were generated using both PYTHIA8 and HERWIG7. These samples were compared and found to be in generally good agreement. A flat uncertainty of 5% for the resonant analysis. The non-resonant analysis uses an uncertainty parameterized by a linear function of m_{HH} . The uncertainties from the signal reweighting procedure for the resonant analysis are found to have a $< 1\%$ impact and are therefore not included.

The impact of variations in the b -jet energy corrections used during calibration, the acceptance times efficiency, and the underlying event simulation were also examined, and found to be negligible. Uncertainties are therefore not included for these effects.

6.7.5 Summary of uncertainties

Resonant uncertainties

The background uncertainties for the resonant analysis are summarized for 2016, 2017, and 2018 are summarized in Table 6.4, respectively. The signal uncertainties are much smaller in comparison and do not significantly impact overall sensitivity and are therefore not included in the table below.

Non-resonant uncertainties

The background uncertainties for the ggF and VBF channels of the non-resonant analysis are summarized in Table 6.5 and Table 6.6, respectively.

Table 6.4: Magnitudes of background uncertainties for the resonant analysis in the 4b Signal Region

2016 Signal Region 4b (Bkgd Yield = 9004.0)	Counts	Percent of Background (%)
Bootstrap Error	306.7	3.4 %
Poisson Error	99.1	1.1 %
Extrapolation Error	21.2	0.2 %
Total Error	322.9	3.6 %
2017 Signal Region 4b (Bkgd Yield = 10953.9)	Counts	Percent of Background (%)
Bootstrap Error	260.6	2.4 %
Poisson Error	112.2	1.0 %
Extrapolation Error	117.2	1.1 %
Total Error	306.9	2.8 %
2018 Signal Region 4b (Bkgd Yield = 19448.3)	Counts	Percent of Background (%)
Bootstrap Error	606.3	3.1 %
Poisson Error	145.8	0.7 %
Extrapolation Error	276.2	1.4 %
Total Error	682.0	3.5 %

Table 6.5: Magnitudes of background uncertainties for the non-resonant ggF channel in the 4b Signal Region.

2016 Signal Region 4b (Bkgd Yield = 3211.1)	Counts	Relative (%)
Poisson Error	106.0	3.3
Bootstrap Error	197.1	6.1
Shape Q _N	37.8	1.2
Shape Q _W	29.9	0.93
Shape Q _S	35.5	1.1
Shape Q _E	45.1	1.4
3b1f NC	7.3	0.23
Total Error	236.1	7.4

2017 Signal Region 4b (Bkgd Yield = 4492.7)	Counts	Relative (%)
Poisson Error	166.0	3.7
Bootstrap Error	256.7	5.7
Shape Q _N	88.8	2.0
Shape Q _W	66.1	1.5
Shape Q _S	69.8	1.6
Shape Q _E	64.6	1.4
3b1f NC	17.8	0.40
Total Error	339.2	7.6

2018 Signal Region 4b (Bkgd Yield = 7720.0)	Counts	Relative (%)
Poisson Error	218.3	2.8
Bootstrap Error	426.6	5.5
Shape Q _N	143.3	1.9
Shape Q _W	56.6	0.73
Shape Q _S	42.5	0.55
Shape Q _E	88.7	1.1
3b1f NC	18.0	0.23
Total Error	513.2	6.6

Table 6.6: Magnitudes of background uncertainties for the non-resonant VBF channel in the 4b Signal Region.

Signal Region 4b (Bkgd Yield = 354.6)	Counts	Relative (%)
Poisson Error	9.6	2.7
Bootstrap Error	30.8	8.7
Shape Q_N	7.9	2.2
Shape Q_W	5.3	1.5
Shape Q_S	4.4	1.2
Shape Q_E	7.1	2.0
Total Error	34.6	9.8

6.8 Background Validation

Additional background validation checks are performed by defining sets of additional validation regions with a similar structure to the nominal region set, but with a background-rich “signal region”. As was done for the 3b1f non-closure uncertainty, a full “dress rehearsal” of the reweighting procedure can then be performed in these regions, with the alternative background-rich pseudo-signal region then unblinded for validation. In this case, however, we are able to use 4b data as the target due to the lower signal content of these new shifted regions

6.8.1 Reverse $|\Delta\eta_{HH}|$ validation

Both the resonant and non-resonant ggF channels use a reverse $|\Delta\eta_{HH}|$ region, which includes the data with $|\Delta\eta_{HH}| > 1.5$, that is, with the original $|\Delta\eta_{HH}|$ cut reversed (cf. Section 6.4.4). This region is not usable as a validation region for the VBF channel, since no $|\Delta\eta_{HH}|$ cut is applied in that case. This region is orthogonal to the nominal signal region, is background rich, and displays generally similar kinematics. To perform the validation in this region, the nominal background estimation is followed using the newly defined CR/CR1 and VR/CR2 to derive a nominal and alternative estimate as normal, but entirely in the $|\Delta\eta_{HH}| > 1.5$ region. The nominal prediction in the reverse $|\Delta\eta_{HH}|$ signal region is then compared, including background uncertainties, to the unblinded background-rich data. The results of this validation procedure are shown in Figures 6.33 and 6.34. Good agreement is observed across the entire m_{HH} spectrum. The pulls are also in good agreement with a standard normal distribution, signifying that there is no significant global bias and that the

uncertainties are appropriately sized.




Figure 6.33: Performance of the background estimation procedure in the reversed $|\Delta\eta_{HH}|$ region ($|\Delta\eta_{HH}| > 1.5$) SR 4b 2017 data, for the resonant analysis. The quoted p -value uses the χ^2 test statistic. Based on this p -values, there is no evidence that the background estimate differs from the data.




Figure 6.34: Performance of the background estimation procedure for the non-resonant ggF channel in the 2018 reversed $|\Delta\eta_{HH}|$ region ($|\Delta\eta_{HH}| > 1.5$) SR 4b for the discriminating variables (a) m_{HH} , (b) $|\Delta\eta_{HH}|$, and (c) X_{HH} . Agreement is seen to be good in all cases, with all $\chi^2/ndf \lesssim 1.1$ and $p \gtrsim 0.3$.

6.8.2 Validation in shifted analysis regions

The non-resonant analysis, including the VBF channel, uses 5 additional validation regions referred to here as “shifted region”. The shifted regions are defined to be identical to the true signal and control regions, but located at shifted locations in the HC mass plane. The positions of these regions are chosen such that they are close to, but not overlapping with, the true signal region. Note that the corresponding control regions are not completely orthogonal to each other. The set of shifted regions is shown in Figure 6.35. These regions are advantageous because they allow us to directly test with 4b data which is still quite kinematically similar to the true SR. Additionally, they provide a set of estimates which can be treated as a (small) statistical sample for comparison. As was done for the $|\Delta\eta_{HH}|$ region, these regions are validated by running the nominal procedure, but using the corresponding new regions. This results in 5 different predictions in 5 different regions. Comparisons of these predictions to the data in the corresponding background-rich shifted signal regions are shown in Figures 6.36 and 6.37. Consistently good agreement is found in both channels for all 5 regions.




Figure 6.35: Shifted regions overlaid on the ggF HC mass plane. The solid red curve in the center represents the true SR. The shifted regions are labeled in clockwise order: upper left (pink), upper center (blue), upper right (green), center right (light-blue), lower right (orange). The shifted region definitions are identical for the VBF and ggF analysis channels and the underlying ggF data is used here only for the sake of reference.




Figure 6.36: Reweighted 2018 SR 2b, and SR 4b data m_{HH} distributions for the shifted regions in the ggF channel. The upper regions and the center right region use CR1 derived weights, but the lower right region uses CR2 derived weights. The total uncertainty in the reweighted 2b estimate includes all background uncertainties except the 3b1f non-closure uncertainty. Good agreement is seen across the 5 regions.




Figure 6.37: Reweighted SR 2b and SR 4b data m_{HH} distributions for the shifted regions in VBF channel. The upper regions and the center right region use CR1 derived weights, but the lower right region uses CR2 derived weights. The total uncertainty in the reweighted 2b estimate includes all background uncertainties. Good agreement is seen across the 5 regions.

6.9 Statistical analysis and results

The resonant and non-resonant analyses are used to set upper limits on the cross-sections for a variety of signal hypotheses. In both cases, this is done using the CL_s method [161]. The test statistic used is q_μ [162]:

$$q_\mu = \begin{cases} -2 \ln \left(\frac{\mathcal{L}(\mu, \hat{\theta}(\mu))}{\mathcal{L}(\hat{\mu}, \hat{\theta}(\mu))} \right) & \hat{\mu} \leq \mu \\ 0 & \hat{\mu} > \mu \end{cases} \quad (6.14)$$

where μ is the signal strength, θ represents the nuisance parameters, a single hat represents the maximum likelihood estimate of a parameter and a double-hat represents the conditional maximum likelihood estimate for a fixed argument. The CL_s is then be defined as the ratio of the p -values for the background only and combined signal+background hypothesis

$$CL_s = \frac{CL_{s+b}}{CL_b} = \frac{p(q_\mu \geq q_{\mu,obs}|s+b)}{p(q_\mu \geq q_{\mu,obs}|b)} \quad (6.15)$$

The CL_s is useful because it conditions the normal signal+background p -value on the p -value of the background only hypothesis to ensure that spurious deficits in the background do not lead to excluded signals. Since this analysis maintains sufficiently large statistics, the p -values are calculated using the asymptotic approximation [162].

The statistical implementation for both analyses follows the conventions and methodology of the RooStats framework [163]. In particular, the resonant analysis makes use of HistFactory [164] while the non-resonant analysis uses pure-Python HistFactory (pyhf) [165]. The background PDFs are created from the corresponding background estimates discussed in the previous sections and the signal PDFs are created from the relevant (reweighted) MC simulation. A one-sided upper limit is placed by finding the μ value for which the $CL_s = 0.05$. The uncertainties on this limit are obtained by by calculating the CL_s at $\pm 1\sigma$ and $\pm 2\sigma$ and finding where these values equal 0.05. Signal strengths which yield a $CL_s < 0.05$ are said to be excluded.

6.9.1 Resonant results

The resonant analysis performs fits in corrected m_{HH} due to its separation power between signal and background as well as between different signal hypotheses. The procedure begins with a profile-likelihood background-only fit to the SR 4b m_{HH} spectrum categorized by




Figure 6.38: Results of a background only profile likelihood fit to the SR 4b m_{HH} spectrum for the resonant analysis. Scalar signals with masses of 280 GeV (black), 400 GeV (red), 600 GeV (blue), and 1000 GeV (orange) are overlaid.

year. The results of this fit, with four representative scalar signals overlaid, are shown in Figure 6.38. The resulting distribution is clearly well modeled, and so the analysis proceeds with signal+background fits. Pulls and impacts from four such fits with representative scalar signals are shown in Figure 6.39. Several H_T systematics are pulled to $\sim +1\sigma$ for these fits. This indicates that the VR estimate better modeled the data in this case. Since the VR is closer to the SR in the mass plane, it has more similar kinematics. A model trained in a region with more similar kinematics to the SR should provide better modeling in the SR, and so it is unsurprising that the VR provides a slightly better model than the CR. This behavior was part of the motivation for using a more balanced set of control regions for the nominal and alternate estimate in the non-resonant analysis. The largest observed impacts are the background shape systematics, as expected. In particular, low H_T shape systematics have




Figure 6.39: Pulls and impacts for the background uncertainties for four representative scalar signals. Several H_T systematics are pulled to $\sim +1\sigma$, indicating that the VR model better modeled the data in this case. This is somewhat expected as the VR is closer to the SR and should provide a better model.

the largest impact for the 280 GeV signal while the high H_T shape systematic has the largest impact for the higher masses, including the displayed 600 GeV and 1000 GeV.

Since all initial checks are within expectation, modeling is assumed to be good and limits and significances are calculated. No globally significant excesses are observed for any signal mass points. The largest local excess gives a significance of 2.6σ for the scalar signal and 2.7σ for the spin-2 signal. Upper cross-section limits are then set for the scalar and spin-2 signal models. The resulting limits are shown in Figures 6.40 and 6.41. Contributions from the boosted channel are also included in the public limits and are therefore shown here. The contributions of the boosted channel begin to become significant at ~ 1.3 TeV, continuing to 5 TeV. The RS spin-2 graviton model considered here for $k/\bar{M}_{\text{Pl}} = 1$ is excluded from 300 GeV to 1.4 TeV by the resolved resonant channel. All other plots shown here use results from fits which only included the resolved channel.




Figure 6.40: Upper limits on the cross-section for the resonant scalar signal model. Contributions from the boosted channel are included in the main expected lines. Those contributions arising only from the resolved channel detailed here are shown in pink.




Figure 6.41: Upper limits on the cross-section for the resonant spin-2 signal model. Contributions from the boosted channel are included in the main expected lines. Those contributions arising only from the resolved channel detailed here are shown in pink. RS Gravitons with $k/\bar{M}_{\text{Pl}} = 1$ under the analysis signal model are excluded by the resolved analysis between 300 and 1400 GeV.

6.9.2 Non-resonant results

As discussed in Section 6.6, the non-resonant analysis performs fits in bins of m_{HH} with further categorization by year, X_{HH} , and $|\Delta\eta_{HH}|$. The general fitting procedure is similar to that used for the resonant analysis, first a profile-likelihood background-only fit is performed to the SR 4b data in both the ggF and VBF channel. The results of this fit are shown in Figure 6.42 for the ggF channel and in Figure 6.43 for the VBF channel, with representative signals overlaid. The resulting distribution are found to be well modeled with no significant pulls or constraints of the nuisance parameters and so the procedure continues. Signal and background fits are then run for the relevant signal models and 95% confidence interval limits are placed on the relevant parameters, as outlined in the following sections.

Higgs boson self-coupling and HHVV coupling limits

For the SM signal, limits are placed on the combined ggF and VBF signal cross-section, with the relative contributions of ggF and VBF production to the total cross section fixed to the SM prediction. The upper limit on the SM cross-section is shown in Table 6.7. Upper limits on the HH production cross-section as a function of κ_λ and κ_{2V} are shown in Figure 6.44 and Figure 6.45, respectively.

Table 6.7: The observed and expected upper limit on the combined SM ggF and VBF HH production cross-section $\sigma_{\text{ggF+VBF}}$ at the 95% confidence limit. Expected limits are also shown for the ± 1 and $\pm 2\sigma$. Limits are normalized by the predicted cross-section of SM ggF and VBF HH production.

	Observed	-2σ	-1σ	Expected	$+1\sigma$	$+2\sigma$
$\sigma_{\text{ggF+VBF}}$	5.49	3.61	4.85	6.73	9.37	12.61




Figure 6.42: Results of a background only profile-likelihood fit to SR 4b data in the non-resonant ggF channel for each year. Signal variations with $\kappa_\lambda=1$ (red) and 10 (turquoise), $c_{tH}=6$, and HEFT Benchmark 7 are overlaid.




Figure 6.43: Results of a background-only profile-likelihood fit to SR 4b data in the non-resonant VBF channel. Signal variations with $\kappa_\lambda=1$ (red), $\kappa_\lambda=10$ (turquoise), and $\kappa_{2V}=0$ (black) are overlaid.




Figure 6.44: Limits on the HH cross-section as function of the Higgs boson self-coupling obtained using the combination of ggF and VBF channels. Observed (expected) limits are $\kappa_\lambda \in [-4.7, 12.2]$ ($[-4.1, 10.4]$).




Figure 6.45: Limits on the HH cross-section as function of the HHVV coupling obtained using the combination of ggF and VBF channels. Observed (expected) limits are $\kappa_{2V} \in [0.02, 2.05]$ ($[0.01, 2.06]$).

SMEFT limits

SMEFT limits are set using only the ggF channel, with no VBF effects included. Upper cross-section limits are set at the 95% confidence level on each of the 4 relevant SMEFT parameters, excluding c_H . At this order of calculation, limits on c_H would be equivalent to those set on κ_λ under the linear transformation $c_H \cong -2.14(\kappa_\lambda - 1)$. As the SMEFT variations are calculated only to LO and do not include VBF (which is known to have a significant impact), the c_H limits are neglected in this case in favor of the more accurate description provided by the κ_λ limits. The resulting cross-section limits as a function of each parameter are shown in Figure 6.46. The 95% confidence level constraints for each coefficient are summarized in Table 6.8, including the constraints on c_H obtained using the κ_λ limits and the aforementioned transformation.

HEFT limits

Upper limits are set on the cross-sections of several HEFT benchmark signal models, again using the ggF channel data only. These benchmarks are intended to parameterize the possible HEFT signals by the general set of possible shapes they can produce in the invariant mass spectrum [152]. The coefficient values for these benchmark are shown in Table 6.9. The resulting limits are summarized in Figure 6.47.

Table 6.8: Limits on the SMEFT Warsaw coefficients, at the 95% confidence level, obtained from fits to SR 4b data in the ggF channel. These limits are those set when allowing only the indicated coefficient to vary from its SM value. The c_H limits are obtained using the relation between κ_λ and c_H that exists at this order of calculation.

Coefficient	Observed	Expected
c_{tH}	[-11.1, 6.4]	[10.3, 6.5]
$c_{H\Box}$	[-9.9, 15.1]	[-9.7, 14.2]
c_{tG}	[-1.09, 1.07]	[-0.90, 0.87]
c_{HG}	[-0.07, 0.06]	[-0.05, 0.05]
c_H (from κ_λ)	[-23.9, 12.2]	[-20.1, 10.9]




Figure 6.46: One-dimensional limits on (a) c_{tH} (b) c_{tG} , (c) $c_{H\Box}$, and (d) c_{HG} from the non-resonant ggF channel only. At this order of calculation, c_H is directly related to κ_λ via linear transformation.

Table 6.9: List of benchmark signals from [152] and corresponding HEFT coefficients.

Benchmark	κ_t	κ_λ	c_{tt}	c_{ggh}	c_{gghh}
SM	1	1	0	0	0
1	0.94	3.94	-0.333	0.5	0.333
2	0.61	6.84	0.333	0	-0.333
3	1.05	2.21	-0.333	0.5	0.5
4	0.61	2.79	0.333	-0.5	0.167
5	1.17	3.95	-0.333	0.167	-0.5
6	0.83	5.68	0.333	-0.5	0.333
7	0.94	-0.10	1	0.167	-0.167




Figure 6.47: Upper cross-section limits on the 7 HEFT benchmarks $BM1 - 7$ at the 95% confidence level. The observed (expected) limits are included as text on the plot, in units of fb. The limit on the SM ggF cross-section is also included for reference.

Chapter 7

Conclusions and Outlook

It has been one decade since the final particle in the Standard Model (SM), the Higgs boson, was discovered. Despite this observational “completion” of the theory, open questions still remain. The Higgs boson plays a central role in many of these questions and in the Beyond the Standard Model (BSM) theories which attempt to resolve them. A decade of study has yielded increasingly precise knowledge about the characteristics of this particle. Yet, to date, these studies have done little to resolve these questions and have not provided any clear evidence of new physics. Thus, it is crucial to not only increase precision, but to expand the scope and generality of the studies being performed.

One such avenue is the search for Higgs pair production. The study of this as of yet unobserved process provides access to a broad range of BSM scenarios. This includes complimentary sensitivity to both resonant and non-resonant processes. It also allows more precise constraints to be set on two poorly measured Higgs boson couplings: the Higgs boson self-coupling and the four-point Higgs to vector boson coupling.

This dissertation has outlined two searches for Higgs pair (HH) production in the 4 b -quark final state, one focusing on resonant production and one on non-resonant production. These studies were performed using 136 fb of data collected by the ATLAS experiment from LHC proton-proton collisions at $\sqrt{s} = 13$ TeV. These studies build and expand upon previous work [94, 142, 143] with greater physics coverage, improved analysis strategies, and in some cases, more data. In doing so, these searches have become uniquely broad in scope and have placed competitive, and in some cases world-leading, limits on several important coefficients, couplings and cross-sections.

In particular, the non-resonant analysis places observed upper limits on the SM HH production cross-section of 5.49 times the SM cross-section. These limits represent an

improvement of over 40% relative to the previous analysis (scaled to the current data set size). Improved limits have also been set to the Higgs boson self-coupling and the four-point Higgs to vector boson coupling (κ_{2V}). Of note, the observed limits on the latter coupling $\kappa_{2V} \in [0.02, 2.05]$ have improved over 50% relative to the previous κ_{2V} limits which were set using the same data set. These κ_{2V} limits are now the leading limits amongst similar analyses. The improved self-coupling limits of $\kappa_\lambda \in [-4.7, 12.2]$ should provide additional sensitivity to future combinations with the other HH decay channels. With this additional sensitivity, the combined HH limits should begin to exclude realistic models of EWBG. The physics scope has also been expanded significantly relative to past studies with the inclusion of interpretations within two common EFT frameworks: HEFT and SMEFT. These frameworks parameterize a broad spectrum of potential high-energy BSM physics and in doing so, provide a more general and robust framework for setting limits. This is particularly noteworthy as it is the first step towards incorporating HH production into global fits. In this combination, HH is expected to provide significant sensitivity to several coefficients, while helping to alleviate some of the degeneracies present within recent global fits [166].

The resonant analysis performed a search for new massive scalar and spin-2 resonances with masses between 251 GeV and 1.5 TeV decaying to a pair of Higgs bosons. No significant excess was observed, and upper limits on the cross-sections were set. These limits exclude RS Gravitons with $k/\bar{M}_{\text{Pl}} = 1$ between 300 GeV and 1.4 TeV and place world leading limits on general narrow-width scalar resonances over a wide range of masses.

Together the two analyses detailed in this dissertation represent a significant step forward in both scope and precision. With these improvements, the 4 b -quark final state is becoming increasingly competitive with other HH channels and is among the leaders in sensitivity. With the planned High-Luminosity LHC (HL-LHC) upgrades, when combined with the other HH channels, sensitivity to SM Higgs pair production is likely. Furthermore, one of the primary targets of this high-luminosity era is a high-precision global EFT measurement. With the newly expanded physics scope provided by the analyses outlined in this dissertation, Higgs pair production will be able to play an important role in these measurements.

Bibliography

- [1] G. Buchalla, M. Capozzi, A. Celis, G. Heinrich, and L. Scyboz. Higgs boson pair production in non-linear Effective Field Theory with full mt-dependence at NLO QCD. *Journal of High Energy Physics*, 2018(9), Sep 2018. → pages vii, 86
- [2] L. Alsafar. EFT Modelling at NLO. Feb 2021. → page vii
- [3] J. Beringer, J. Arguin, R. Barnett, K. Copic, O. Dahl, D. Groom, C. Lin, J. Lys, H. Murayama, C. Wohl, et al. Review of particle physics. *Physical Review D-Particles, Fields, Gravitation and Cosmology*, 86(1):010001, 2012. → page 1
- [4] P. Bagnaia, M. Banner, R. Battiston, P. Bloch, F. Bonaudi, K. Borer, M. Borghini, J.-C. Chollet, A. Clark, C. Conta, et al. Evidence for $Z^0 \rightarrow e^+ e^-$ at the CERN pp Collider. *Physics Letters B*, 129(1-2):130–140, 1983.
- [5] Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. 2012.
- [6] S. Chatrchyan, V. Khachatryan, A. M. Sirunyan, A. Tumasyan, W. Adam, E. Aguilo, T. Bergauer, M. Dragicevic, J. Erö, C. Fabjan, et al. Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC. *Physics Letters B*, 716(1):30–61, 2012. → page 1
- [7] S. L. Glashow. Partial-symmetries of weak interactions. *Nuclear physics*, 22(4):579–588, 1961. → pages 1, 7
- [8] S. Weinberg. A model of leptons. *Physical review letters*, 19(21):1264, 1967.
- [9] A. Salam. Weak and electromagnetic interactions. In *Selected Papers Of Abdus Salam: (With Commentary)*, pages 244–254. World Scientific, 1994. → pages 1, 7
- [10] C. Burgard. Example: Standard model of physics.
<https://texample.net/tikz/examples/model-physics/>. → pages xiv, 2
- [11] D. Galbraith. UX: Standard Model of the Standard Model.
<http://davidgalbraith.org/portfolio/ux-standard-model-of-the-standard-model/>. → pages xiv, 2

- [12] M. E. Peskin. *An introduction to quantum field theory*. CRC press, 2018. → pages 2, 3, 4
- [13] M. Srednicki. *Quantum field theory*. Cambridge University Press, 2007. → pages 2, 3, 12
- [14] P. D. Group, P. Zyla, R. Barnett, J. Beringer, O. Dahl, D. Dwyer, D. Groom, C.-J. Lin, K. Lugovsky, E. Pianori, et al. Review of particle physics. *Progress of Theoretical and Experimental Physics*, 2020(8):083C01, 2020. → pages 3, 73
- [15] M. Gell-Mann. A schematic model of baryons and mesons. *Phys. Letters*, 8, 1964. → page 3
- [16] H. Fritzsch, M. Gell-Mann, and H. Leutwyler. Advantages of the color octet gluon picture. *Physics Letters B*, 47(4):365–368, 1973.
- [17] C.-N. Yang and R. L. Mills. Conservation of isotopic spin and isotopic gauge invariance. *Physical review*, 96(1):191, 1954. → page 3
- [18] R. P. Feynman. *Photon-hadron interactions*. CRC Press, 2018. → page 5
- [19] MSHT pdfs. → pages xiv, 6
- [20] Y. L. Dokshitzer. Calculation of the structure functions for deep inelastic scattering and e^+e^- annihilation by perturbation theory in quantum chromodynamics. *Zh. Eksp. Teor. Fiz*, 73(1216):26, 1977. → page 6
- [21] V. N. Gribov and L. N. Lipatov. DEEP INELASTIC ep -SCATTERING IN A PERTURBATION THEORY. Technical report, Inst. of Nuclear Physics, Leningrad, 1972. → page 6
- [22] P. W. Higgs. Broken symmetries and the masses of gauge bosons. *Physical Review Letters*, 13(16):508, 1964. → page 8
- [23] F. Englert and R. Brout. Broken symmetry and the mass of gauge vector mesons. *Physical review letters*, 13(9):321, 1964.
- [24] G. S. Guralnik, C. R. Hagen, and T. W. Kibble. Global conservation laws and massless particles. *Physical Review Letters*, 13(20):585, 1964. → page 8
- [25] M. D. Schwartz. *Quantum field theory and the standard model*. Cambridge University Press, 2014. → pages 12, 13
- [26] L. Alvarez-Gaume and E. Witten. Gravitational anomalies. *Nuclear Physics B*, 234(2):269–330, 1984. → page 13
- [27] C. W. Misner, K. S. Thorne, and J. A. Wheeler. *Gravitation*. Macmillan, 1973. → page 13

- [28] G. t Hooft. Symmetry breaking through Bell-Jackiw anomalies. *Physical Review Letters*, 37(1):8–11, 1976. → page 14
- [29] V. A. Rubakov and M. E. Shaposhnikov. Electroweak baryon number non-conservation in the early universe and in high energy collisions. In *AIP Conference Proceedings*, volume 419, pages 347–412. American Institute of Physics, 1998. → page 14
- [30] N. Manton. Topology in the Weinberg-Salam theory. *Physical Review D*, 28(8):2019, 1983. → page 14
- [31] Y. Brihaye and J. Kunz. Electroweak bubbles and sphalerons. *Physical Review D*, 48(8):3884, 1993.
- [32] D. E. Morrissey and M. J. Ramsey-Musolf. Electroweak baryogenesis. *New Journal of Physics*, 14(12):125003, Dec 2012. → pages 14, 25, 26
- [33] P. Arnold and L. McLerran. Sphalerons, small fluctuations, and baryon-number violation in electroweak theory. *Physical Review D*, 36(2):581, 1987. → page 14
- [34] S. Y. Khlebnikov and M. Shaposhnikov. The statistical theory of anomalous fermion number non-conservation. *Nuclear Physics B*, 308(4):885–912, 1988. → page 14
- [35] V. V. Khoze and D. L. Milne. Suppression of electroweak instanton processes in high-energy collisions. *International Journal of Modern Physics A*, 36(05):2150032, 2021. → page 14
- [36] G. D. Moore. Measuring the broken phase sphaleron rate nonperturbatively. *Physical Review D*, 59(1):014503, 1998. → page 14
- [37] M. D’Onofrio, K. Rummukainen, and A. Tranberg. Sphaleron rate in the minimal standard model. *Physical Review Letters*, 113(14):141602, 2014. → page 14
- [38] B. Pontecorvo. INVERSE β PROCESSES AND NONCONSERVATION OF LEPTON CHARGE. *Zhur. Eksptl’. i Teoret. Fiz.*, 34, 1958. → page 15
- [39] Z. Maki, M. Nakagawa, and S. Sakata. Remarks on the unified model of elementary particles. *Progress of Theoretical Physics*, 28(5):870–880, 1962. → page 15
- [40] J. Kretzschmar. Standard Model physics at the LHC. pages 153–171. 19 p, Mar 2018. Contribution to the volume “From my Vast Repertoire - the Legacy of Guido Altarelli” : 24 pages, 13 figures. → pages xv, 16
- [41] C. Rovelli. Loop quantum gravity. *Living reviews in relativity*, 11(1):1–69, 2008. → page 15

- [42] O. Aharony, S. S. Gubser, J. Maldacena, H. Ooguri, and Y. Oz. Large N field theories, string theory and gravity. *Physics Reports*, 323(3-4):183–386, 2000. → page 15
- [43] H. P. Nilles. Supersymmetry, supergravity and particle physics. *Physics Reports*, 110(1-2):1–162, 1984. → page 15
- [44] N. Arkani-Hamed and S. Dimopoulos. Supersymmetric unification without low energy supersymmetry and signatures for fine-tuning at the LHC. *Journal of High Energy Physics*, 2005(06):073, 2005. → page 16
- [45] N. Arkani-Hamed, S. Dimopoulos, and G. Dvali. The hierarchy problem and new dimensions at a millimeter. *Physics Letters B*, 429(3–4):263–272, Jun 1998. → page 16
- [46] N. Arkani-Hamed, S. Dimopoulos, and G. Dvali. Phenomenology, astrophysics, and cosmology of theories with submillimeter dimensions and TeV scale quantum gravity. *Physical Review D*, 59(8), Mar 1999. → page 16
- [47] B. Hespel, D. López-Val, and E. Vryonidou. Higgs pair production via gluon fusion in the Two-Higgs-Doublet Model. *Journal of High Energy Physics*, 2014(9), Sep 2014. → pages 16, 29
- [48] K. Agashe, R. Contino, and A. Pomarol. The minimal composite Higgs model. *Nuclear Physics B*, 719(1-2):165–187, 2005. → page 16
- [49] G. Degrandi, S. Di Vita, J. Elias-Miro, J. R. Espinosa, G. F. Giudice, G. Isidori, and A. Strumia. Higgs mass and vacuum stability in the Standard Model at NNLO. *Journal of High Energy Physics*, 2012(8):1–33, 2012. → page 16
- [50] V. Branchina, E. Messina, and M. Sher. Lifetime of the electroweak vacuum and sensitivity to Planck scale physics. *Physical Review D*, 91(1):013003, 2015. → page 17
- [51] A. Riotto. Theories of baryogenesis. *Proceedings, summer School in High-energy physics and cosmology: Trieste, Italy*, pages 326–436, 1998. → pages 17, 23, 26
- [52] F. Capozzi, E. Lisi, A. Marrone, D. Montanino, and A. Palazzo. Neutrino masses and mixings: Status of known and unknown 3ν parameters. *Nuclear Physics B*, 908:218–234, 2016. → page 17
- [53] E. Ma. Verifiable radiative seesaw mechanism of neutrino mass and dark matter. *Physical Review D*, 73(7):077301, 2006. → page 17
- [54] M. Auger, D. Auty, P. Barbeau, E. Beauchamp, V. Belov, C. Benitez-Medina, M. Breidenbach, T. Brunner, A. Burenkov, B. Cleveland, et al. Search for neutrinoless double-beta decay in Xe 136 with EXO-200. *Physical Review Letters*, 109(3):032505, 2012. → page 17

- [55] E. Armengaud, C. Augier, A. Barabash, F. Bellini, G. Benato, A. Benoit, M. Beretta, L. Bergé, J. Billard, Y. A. Borovlev, et al. New Limit for Neutrinoless Double-Beta Decay of Mo 100 from the CUPID-Mo Experiment. *Physical review letters*, 126(18):181802, 2021.
- [56] D. Adams, C. Alduino, K. Alfonso, F. Avignone III, O. Azzolini, G. Bari, F. Bellini, G. Benato, M. Biassoni, A. Branca, et al. Improved Limit on Neutrinoless Double-Beta Decay in Te 130 with CUORE. *Physical review letters*, 124(12):122501, 2020. → page 17
- [57] G. Bertone, D. Hooper, and J. Silk. Particle dark matter: Evidence, candidates and constraints. *Physics reports*, 405(5-6):279–390, 2005. → page 17
- [58] S. Dodelson and L. M. Widrow. Sterile neutrinos as dark matter. *Physical Review Letters*, 72(1):17, 1994. → page 17
- [59] X. Chu, T. Hambye, and M. H. Tytgat. The four basic ways of creating dark matter through a portal. *Journal of Cosmology and Astroparticle Physics*, 2012(05):034, 2012. → page 17
- [60] A. D. Sakharov. Violation of CP-invariance, C-asymmetry, and baryon asymmetry of the Universe. In *In The Intermissions... Collected Works on Research into the Essentials of Theoretical Physics in Russian Federal Nuclear Center, Arzamas-16*, pages 84–87. World Scientific, 1998. → pages 17, 23
- [61] M. B. Gavela, P. Hernandez, J. Orloff, and O. Pène. Standard model CP-violation and baryon asymmetry. *Modern Physics Letters A*, 9(09):795–809, 1994. → page 17
- [62] P. Huet and E. Sather. Electroweak baryogenesis and standard model CP violation. *Physical Review D*, 51(2):379, 1995.
- [63] M. B. Gavela, P. Hernández, J. Orloff, O. Pène, and C. Quimbay. Standard model CP-violation and baryon asymmetry (II). Finite temperature. *Nuclear Physics B*, 430(2):382–426, 1994.
- [64] S. Bruggisser, T. Konstandin, and G. Servant. CP-violation for electroweak baryogenesis from dynamical CKM matrix. *Journal of Cosmology and Astroparticle Physics*, 2017(11):034, 2017.
- [65] W. Chao. CP violation at the finite temperature. *Physics Letters B*, 796:102–106, 2019. → page 17
- [66] A. Riotto and M. Trodden. Recent progress in baryogenesis. *Annual Review of Nuclear and Particle Science*, 49(1):35–75, 1999. → page 17
- [67] E. W. Kolb, A. Riotto, and I. I. Tkachev. GUT baryogenesis after preheating: Numerical study of the production and decay of X-bosons. *Physics Letters B*, 423(3-4):348–354, 1998. → page 17

- [68] D. Hooper and G. Krnjaic. GUT baryogenesis with primordial black holes. *Physical Review D*, 103(4):043504, 2021. → page 17
- [69] B. D. Micco, M. Gouzevitch, J. Mazzitelli, and C. Vernieri. Higgs boson potential at colliders: Status and perspectives. *Reviews in Physics*, 5:100045, Nov 2020. → pages xv, xvi, 19, 20, 22, 26, 27, 28, 35, 36, 37
- [70] S. Amoroso, P. Azzurri, J. Bendavid, E. Bothmann, D. Britzger, H. Brooks, A. Buckley, M. Calvetti, X. Chen, M. Chiesa, L. Cieri, V. Ciulli, J. Cruz-Martinez, A. Cueto, A. Denner, S. Dittmaier, M. Donegà, M. Dührssen-Debling, I. Fabre, S. Ferrario-Ravasio, D. de Florian, S. Forte, P. Francavilla, T. Gehrmann, A. G.-D. Ridder, L. Gellersen, E. W. N. Glover, P. Gras, C. Gwenlan, Y. Haddad, G. Heinrich, J. Hessler, T. J. Hobbs, M. Höfer, A. Huss, J. Huston, T. Ježo, S. P. Jones, S. Kallweit, M. Klasen, G. Knippen, A. Larkoski, M. LeBlanc, P. Loch, K. Long, D. Maître, S. Marzani, J. Mazzitelli, J. A. Mcfayden, E. Metodiev, J. K. L. Michel, M. M. Llácer, B. Nachman, P. Nadolsky, D. Napoletano, E. R. Nocera, C. Oleari, C. Pandini, M. Pellen, S. Pigazzini, J. Pires, S. Plätzer, S. Prestel, K. Rabbertz, E. Re, P. Richardson, F. Ringer, J. Rojo, J. Roloff, R. Röntsch, M. Schönherr, C. Schwan, F. Siegert, D. Soper, G. Soyez, M. Spira, M. R. Sutton, F. J. Tackmann, V. Theeuwes, S. L. Villani, J. Whitehead, H. T. Yang, and J. Zhou. Les Houches 2019: Physics at TeV Colliders: Standard Model Working Group Report, 2020. → pages xv, 21, 84, 86
- [71] LHC Higgs Cross Section HH Sub-group. → pages xv, 21
- [72] F. A. Dreyer and A. Karlberg. Vector-boson fusion higgs pair production at n3lo. *Physical Review D*, 98(11), Dec 2018. → pages xv, 21, 128
- [73] F. Bishara, R. Contino, and J. Rojo. Higgs pair production in vector-boson fusion at the LHC and beyond. *The European Physical Journal C*, 77(7), Jul 2017. → page 22
- [74] M. Quiros. Finite temperature field theory and phase transitions. *Proceedings, Summer school in high-energy physics and cosmology: Trieste, Italy*, 1999:187–259, 1998. → pages 23, 24, 25
- [75] R. R. Parwani. Resummation in a hot scalar field theory. *Physical Review D*, 45(12):4695, 1992. → page 25
- [76] P. Huang, A. J. Long, and L.-T. Wang. Probing the electroweak phase transition with Higgs factories and gravitational waves. *Physical Review D*, 94(7), Oct 2016. → page 25
- [77] V. Barger, P. Langacker, M. McCaskey, M. J. Ramsey-Musolf, and G. Shaughnessy. CERN LHC phenomenology of an extended standard model with a real scalar singlet. *Physical Review D*, 77(3), Feb 2008. → page 27

- [78] D. O’Connell, M. J. Ramsey-Musolf, and M. B. Wise. Minimal extension of the standard model scalar sector. *Physical Review D*, 75(3), Feb 2007. → page 27
- [79] I. Lewis and M. Sullivan. Benchmarks for double Higgs production in the singlet-extended standard model at the LHC. *Physical Review D*, 96(3), Aug 2017. → pages xvi, 28
- [80] A. Barroso, P. Ferreira, I. Ivanov, and R. Santos. Metastability bounds on the two Higgs doublet model. *Journal of High Energy Physics*, 2013(6):1–33, 2013. → page 28
- [81] J. M. Cline, K. Kainulainen, and M. Trott. Electroweak baryogenesis in two Higgs doublet models and B meson anomalies. *Journal of High Energy Physics*, 2011(11):1–49, 2011. → page 28
- [82] S. P. Martin. A supersymmetry primer. In *Perspectives on supersymmetry II*, pages 1–153. World Scientific, 2010. → pages 28, 29
- [83] S. Dimopoulos and H. Georgi. Softly broken supersymmetry and SU (5). *Nuclear Physics B*, 193(1):150–162, 1981. → page 29
- [84] L. Randall and R. Sundrum. Large Mass Hierarchy from a Small Extra Dimension. *Physical Review Letters*, 83(17):3370–3373, Oct 1999. → pages xvi, 29, 30
- [85] C. Csáki, M. L. Graesser, and G. D. Kribs. Radion dynamics and electroweak physics. *Physical Review D*, 63(6), Feb 2001. → page 30
- [86] A. V. Manohar. Introduction to effective field theories. 2018. → page 31
- [87] C. Burgess. An Introduction to Effective Field Theory. *Annual Review of Nuclear and Particle Science*, 57(1):329–362, Nov 2007. → page 31
- [88] S. Scherer. Introduction to chiral perturbation theory. *Advances in Nuclear Physics*, Volume 27, pages 277–538, 2003. → page 31
- [89] B. D. Serot and J. D. Walecka. Recent progress in quantum hadrodynamics. *International Journal of Modern Physics E*, 6(04):515–631, 1997. → page 31
- [90] T. Cohen, N. Craig, X. Lu, and D. Sutherland. Is SMEFT enough? *Journal of High Energy Physics*, 2021(3):1–83, 2021. → pages 33, 37
- [91] A. Falkowski and R. Rattazzi. Which EFT. *Journal of High Energy Physics*, 2019(10):1–16, 2019. → pages 33, 37
- [92] C. Degrande, G. Durieux, F. Maltoni, K. Mimasu, E. Vryonidou, and C. Zhang. Automated one-loop computations in the standard model effective field theory. *Physical Review D*, 103(9), May 2021. → pages 34, 86

- [93] HEFT interpretations of Higgs boson pair searches in $b\bar{b}\gamma\gamma$ and $b\bar{b}\tau\tau$ final states and of their combination in ATLAS. Technical report, CERN, Geneva, Apr 2022. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2022-019>. → pages xvi, 34, 36
- [94] ATLAS Collaboration. Constraints on the Higgs boson self-coupling from the combination of single-Higgs and double-Higgs production analyses performed with the ATLAS experiment. Technical report, CERN, Geneva, Oct 2019. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/CONFNOTES/ATLAS-CONF-2019-049>. → pages 35, 83, 149
- [95] G. Buchalla, O. Catà, and C. Krause. Complete electroweak chiral Lagrangian with a light Higgs at NLO. *Nuclear Physics B*, 880:552–573, Mar 2014. → page 36
- [96] L. Evans and P. Bryant. LHC machine. *Journal of instrumentation*, 3(08):S08001, 2008. → page 38
- [97] G. Aad, E. Abat, J. Abdallah, A. Abdelalim, A. Abdesselam, B. Abi, M. Abolins, H. Abramowicz, E. Acerbi, B. Acharya, et al. The ATLAS experiment at the CERN large hadron collider. *Journal of instrumentation*, 3:S08003, 2008. → pages 38, 41
- [98] E. Mobs. The CERN accelerator complex. Complexe des accélérateurs du CERN. Jul 2016. General Photo. → pages xvi, 39
- [99] ATLAS Collaboration. Luminosity Public Results. → pages xvi, 41, 42
- [100] J. Pequenaio. Computer generated image of the whole ATLAS detector. Mar 2008. → pages xvi, 43
- [101] J. Pequenaio and P. Schaffner. How ATLAS detects particles: diagram of particle paths in the detector. Jan 2013. → pages xvi, 44
- [102] G. Strong. On the impact of selected modern deep-learning techniques to the performance and celerity of classification models in an experimental high-energy physics use case. *Machine Learning: Science and Technology*, 1, 06 2020. → pages xvi, 45
- [103] J. Pequenaio. Computer generated image of the ATLAS inner detector. Mar 2008. → pages xvi, 46
- [104] B. Abbott, J. Albert, F. Alberti, M. Alex, G. Alimonti, S. Alkire, P. Allport, S. Altenheiner, L. S. Ancu, E. Anderssen, et al. Production and integration of the ATLAS Insertable B-Layer. *Journal of instrumentation*, 13(05):T05008, 2018. → page 45

- [105] J. D. Jackson. Classical electrodynamics, 1999. → page 46
- [106] J. Pequenaio. Computer Generated image of the ATLAS calorimeter. Mar 2008. → pages xvii, 47
- [107] M. Inokuti. Inelastic collisions of fast charged particles with atoms and molecules—the Bethe theory revisited. *Reviews of modern physics*, 43(3):297, 1971. → page 48
- [108] A. Artamonov, D. Bailey, G. Belanger, M. Cadabeschi, T.-Y. Chen, V. Epshteyn, P. Gorbounov, K. Joo, M. Khakzad, V. Khovanskiy, et al. The ATLAS forward calorimeter. *Journal of Instrumentation*, 3(02):P02010, 2008. → page 50
- [109] J. Pequenaio. Computer generated image of the ATLAS Muons subsystem. Mar 2008. → pages xvii, 51
- [110] ATLAS Collaboration et al. Performance of the ATLAS RPC detector and Level-1 muon barrel trigger at $\sqrt{s} = 13$ TeV. *arXiv preprint arXiv:2103.01029*, 2021. → page 51
- [111] M. Aaboud, G. Aad, B. Abbott, J. Abdallah, O. Abdinov, B. Abeloos, R. Aben, O. S. AbouZeid, N. L. Abraham, and et al. Performance of the ATLAS trigger system in 2015. *The European Physical Journal C*, 77(5), May 2017. → page 52
- [112] A. Ruiz-Martinez and ATLAS Collaboration. The Run-2 ATLAS Trigger System. Technical report, CERN, Geneva, Feb 2016. → page 52
- [113] T. M. and. Frameworks to monitor and predict rates and resource usage in the ATLAS High Level Trigger. *Journal of Physics: Conference Series*, 898:032007, oct 2017. → pages xvii, 54
- [114] M. Aaboud, G. Aad, B. Abbott, J. Abdallah, O. Abdinov, B. Abeloos, S. H. Abidi, O. S. AbouZeid, N. L. Abraham, and et al. Performance of the ATLAS track reconstruction algorithms in dense environments in LHC Run 2. *The European Physical Journal C*, 77(10), Oct 2017. → pages xvii, 54, 63, 64
- [115] ATLAS Collaboration. Configuration and performance of the ATLAS b -jet triggers in Run 2, 2021. → pages xviii, 54, 76
- [116] G. Aad, B. Abbott, J. Abdallah, R. Aben, M. Abolins, O. AbouZeid, H. Abramowicz, H. Abreu, R. Abreu, Y. Abulaiti, et al. Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1. *The European Physical Journal C*, 77(7):1–73, 2017. → pages 54, 67
- [117] T. Gleisberg, S. Höche, F. Krauss, M. Schönherr, S. Schumann, F. Siegert, and J. Winter. Event generation with SHERPA 1.1. *Journal of High Energy Physics*, 2009(02):007–007, Feb 2009. → pages xvii, 56

- [118] S. Alioli, P. Nason, C. Oleari, and E. Re. A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX. *Journal of High Energy Physics*, 2010(6):1–58, 2010. → pages 57, 84, 85
- [119] T. Sjöstrand, S. Ask, J. R. Christiansen, R. Corke, N. Desai, P. Ilten, S. Mrenna, S. Prestel, C. O. Rasmussen, and P. Z. Skands. An introduction to PYTHIA 8.2. *Computer Physics Communications*, 191:159–177, Jun 2015. → pages 57, 84
- [120] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer, H.-S. 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. *Journal of High Energy Physics*, 2014(7), Jul 2014. → pages 57, 84
- [121] J. Hammersley. *Monte carlo methods*. Springer Science & Business Media, 2013. → page 57
- [122] J. Alwall, C. Biscarat, S. Moretti, J. Rathsmann, and A. Sopczak. The $p\bar{p} \rightarrow tbH^\pm$ process at the Tevatron in HERWIG and PYTHIA simulations. *The European Physical Journal C-Particles and Fields*, 39(1):37–39, 2005. → page 58
- [123] B. Andersson, G. Gustafson, G. Ingelman, and T. Sjöstrand. Parton fragmentation and string dynamics. *Physics Reports*, 97(2-3):31–145, 1983. → page 58
- [124] A. Kupco. Cluster hadronization in HERWIG 5.9. *arXiv preprint hep-ph/9906412*, 1999. → page 58
- [125] S. Agostinelli, J. Allison, K. a. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand, et al. GEANT4—a simulation toolkit. *Nuclear instruments and methods in physics research section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 506(3):250–303, 2003. → page 59
- [126] G. Aad, B. Abbott, J. Abdallah, A. A. Abdelalim, A. Abdesselam, O. Abidinov, B. Abi, M. Abolins, H. Abramowicz, and et al. The ATLAS Simulation Infrastructure. *The European Physical Journal C*, 70(3):823–874, Sep 2010. → page 59
- [127] T. Yamanaka, ATLAS Collaboration, et al. The ATLAS calorimeter simulation FastCaloSim. In *Journal of Physics: Conference Series*, volume 331, page 032053. IOP Publishing, 2011. → pages xvii, 60, 61
- [128] R. E. Kalman. A new approach to linear filtering and prediction problems. 1960. → page 63
- [129] M. Aaboud, G. Aad, B. Abbott, J. Abdallah, O. Abidinov, B. Abeloos, R. Aben, O. S. AbouZeid, N. L. Abraham, and et al. Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC. *The European Physical Journal C*, 77(5), May 2017. → page 65

- [130] A. Note. Secondary vertex finding for jet flavour identification with the ATLAS detector. 2017. → pages 66, 78
- [131] M. Aaboud, G. Aad, B. Abbott, J. Abdallah, O. Abdinov, B. Abeloos, S. H. Abidi, O. S. AbouZeid, N. L. Abraham, and et al. Jet reconstruction and performance using particle flow with the ATLAS Detector, Jul 2017. → page 68
- [132] M. Cacciari, G. P. Salam, and G. Soyez. The anti-ktjet clustering algorithm. *Journal of High Energy Physics*, 2008(04):063–063, Apr 2008. → pages 69, 70
- [133] ATLAS Collaboration. Jet energy scale and resolution measured in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector, 2021. → pages xvii, 70, 71, 72, 128
- [134] S. D. Ellis and D. E. Soper. Successive combination jet algorithm for hadron collisions. *Physical Review D*, 48(7):3160–3166, Oct 1993. → page 70
- [135] K. Tomiwa. Performance of Jet Vertex Tagger in suppression of pileup jets and in ATLAS detector. In *Journal of Physics: Conference Series*, volume 802, page 012012. IOP Publishing, 2017. → pages xvii, 73, 74, 75
- [136] G. Aad, B. Abbott, D. C. Abbott, A. A. Abud, K. Abeling, D. K. Abhayasinghe, S. H. Abidi, O. S. AbouZeid, N. L. Abraham, and et al. ATLAS b-jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ tev. *The European Physical Journal C*, 79(11), Nov 2019. → pages 73, 128
- [137] F. A. Di Bello. Optimisation of the ATLAS b -tagging algorithms for the 2017-2018 LHC data-taking. Technical report, ATL-COM-PHYS-2017-1472, 2017. → pages xii, 76, 77, 78, 79
- [138] ATLAS Collaboration et al. Identification of jets containing b-hadrons with recurrent neural networks at the ATLAS experiment. *ATLAS note: ATL-PHYS-PUB-2017-003*, <http://cds.cern.ch/record/2255226>, 2017. → page 77
- [139] ATLAS Collaboration. Topological b -hadron decay reconstruction and identification of b -jets with the JetFitter package in the ATLAS experiment at the LHC. Technical report, CERN, Geneva, Oct 2018. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PUBNOTES/ATL-PHYS-PUB-2018-025>. → page 78
- [140] ATLAS Collaboration. Calibration of the ATLAS b -tagging algorithm in $t\bar{t}$ semi-leptonic events. Technical report, CERN, Geneva, Sep 2018. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/CONFNOTES/ATLAS-CONF-2018-045>. → page 78

- [141] 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. Technical report, CERN, Geneva, Aug 2021. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/CONFNOTES/ATLAS-CONF-2021-035>. → page 83
- [142] G. Aad, B. Abbott, D. C. Abbott, A. Abed Abud, K. Abeling, D. K. Abhayasinghe, S. H. Abidi, O. S. AbouZeid, N. L. Abraham, and et al. Erratum to: Search for the $HH \rightarrow b\bar{b}b\bar{b}$ process via vector-boson fusion production using proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Journal of High Energy Physics*, 2021(1), Jan 2021. → pages 83, 149
- [143] M. Aaboud, G. Aad, B. Abbott, O. Abidinov, B. Abeloos, S. H. Abidi, O. S. AbouZeid, N. L. Abraham, H. Abramowicz, and et al. Search for pair production of Higgs bosons in the $b\bar{b}b\bar{b}$ final state using proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Journal of High Energy Physics*, 2019(1), Jan 2019. → pages 83, 103, 149
- [144] R. D. Ball, V. Bertone, S. Carrazza, C. S. Deans, L. Del Debbio, S. Forte, A. Guffanti, N. P. Hartland, J. I. Latorre, J. Rojo, and M. Ubiali. Parton distributions with LHC data. *Nuclear Physics B*, 867(2):244–289, Feb 2013. → page 84
- [145] R. D. Ball, V. Bertone, S. Carrazza, C. S. Deans, L. Del Debbio, S. Forte, A. Guffanti, N. P. Hartland, J. I. Latorre, J. Rojo, and M. Ubiali. Parton distributions for the LHC run II. *Journal of High Energy Physics*, 2015(4), Apr 2015. → page 84
- [146] J. Butterworth, S. Carrazza, A. Cooper-Sarkar, A. D. Roeck, J. Feltesse, S. Forte, J. Gao, S. Glazov, J. Huston, Z. Kassabov, and et al. PDF4LHC recommendations for LHC Run II. *Journal of Physics G: Nuclear and Particle Physics*, 43(2):023001, Jan 2016. → page 84
- [147] M. Bähr, S. Gieseke, M. A. Gigg, D. Grellscheid, K. Hamilton, O. Latunde-Dada, S. Plätzer, P. Richardson, M. H. Seymour, A. Sherstnev, and et al. Herwig++ physics and manual. *The European Physical Journal C*, 58(4):639–707, Nov 2008. → page 84
- [148] J. Bellm et al. Herwig 7.0/Herwig++ 3.0 release note. *Eur. Phys. J. C*, 76(4):196, 2016. → page 84
- [149] F. A. Dreyer and A. Karlberg. Vector-boson fusion Higgs pair production at N3LO. *Physical Review D*, 98(11), Dec 2018. → page 84
- [150] L. A. Harland-Lang, A. D. Martin, P. Motylinski, and R. S. Thorne. Parton distributions in the LHC era: MMHT 2014 PDFs. *The European Physical Journal C*, 75(5), May 2015. → page 85

- [151] D. J. Lange. The EvtGen particle decay simulation package. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 462(1):152–155, 2001. BEAUTY2000, Proceedings of the 7th Int. Conf. on B-Physics at Hadron Machines. → page 85
- [152] M. Capozzi and G. Heinrich. Exploring anomalous couplings in Higgs boson pair production through shape analysis. *Journal of High Energy Physics*, 2020(3), Mar 2020. → pages xiii, 86, 146, 147
- [153] Configuration and performance of the ATLAS b -jet triggers in Run 2. *arXiv preprint arXiv:2106.03584*, 2021. → page 91
- [154] T. Head, MechCoder, G. Louppe, I. Shcherbatyi, fcharras, Z. Vinícius, cmmalone, C. Schröder, nel215, N. Campos, T. Young, S. Cereda, T. Fan, rene rex, K. K. Shi, J. Schwabedal, carlosdanielcsantos, Hvass-Labs, M. Pak, SoManyUsernamesTaken, F. Callaway, L. Estève, L. Besson, M. Cherti, K. Pfannschmidt, F. Linzberger, C. Caulet, A. Gut, A. Mueller, and A. Fabisch. scikit-optimize/scikit-optimize: v0.5.2, Mar. 2018. → page 96
- [155] Search for pair production of higgsinos in final states with at least three b -tagged jets in $\sqrt{s} = 13$ TeV pp collisions using the ATLAS detector. *arXiv preprint arXiv:1806.04030*, 2018. → page 109
- [156] T. Kanamori, S. Hido, and M. Sugiyama. A least-squares approach to direct importance estimation. *The Journal of Machine Learning Research*, 10:1391–1445, 2009. → page 109
- [157] G. V. Moustakides and K. Basioti. Training Neural Networks for Likelihood/Density Ratio Estimation, 2019. → page 109
- [158] N. Chamandy, O. Muralidharan, A. Najmi, and S. Naidu. Estimating uncertainty for massive data streams. 2012. → pages 110, 118
- [159] ATLAS Collaboration. Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC. Technical report, CERN, Geneva, Jun 2019. All figures including auxiliary figures are available at <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/CONFNOTES/ATLAS-CONF-2019-021>. → page 128
- [160] J. Baglio, F. Campanario, S. Glaus, M. Mühlleitner, J. Ronca, and M. Spira. $g g \rightarrow H H$: Combined uncertainties. *Physical Review D*, 103(5):056002, 2021. → page 128
- [161] A. L. Read. Presentation of search results: the CLs technique. *Journal of Physics G: Nuclear and Particle Physics*, 28(10):2693, 2002. → page 137

- [162] G. Cowan, K. Cranmer, E. Gross, and O. Vitells. Asymptotic formulae for likelihood-based tests of new physics. *The European Physical Journal C*, 71(2):1–19, 2011. → page 137
- [163] L. Moneta, K. Belasco, K. Cranmer, S. Kreiss, A. Lazzaro, D. Piparo, G. Schott, W. Verkerke, and M. Wolf. The roostats project. *arXiv preprint arXiv:1009.1003*, 2010. → page 137
- [164] K. Cranmer, A. Shibata, W. Verkerke, L. Moneta, and G. Lewis. HistFactory: A tool for creating statistical models for use with RooFit and RooStats. Technical report, 2012. → page 137
- [165] L. Heinrich, M. Feickert, G. Stark, and K. Cranmer. pyhf: pure-Python implementation of HistFactory statistical models. *Journal of Open Source Software*, 6(58):2823, 2021. → page 137
- [166] J. Ellis, C. W. Murphy, V. Sanz, and T. You. Updated global SMEFT fit to Higgs, diboson and electroweak data. *Journal of High Energy Physics*, 2018(6), Jun 2018. → page 150

Appendix A

Neural network reweighting performance in additional variables

The plots below show the performance of the resonant, and non-resonant ggF and VBF background models in several different variables.

A.1 Resonant

- ggF 2016 4b CR1: Figures A.1–A.3.
- ggF 2017 4b CR1: Figures A.4–A.6.
- ggF 2018 4b CR1: Figures A.7–A.9.

A.2 Non-resonant

A.2.1 ggF

- ggF 2016 4b CR1: Figures A.1–A.3.
- ggF 2017 4b CR1: Figures A.4–A.6.
- ggF 2018 4b CR1: Figures A.7–A.9.

A.2.2 VBF

- VBF all years 4b CR1: Figures A.10–A.12.




Figure A.1: (Training variables) Distributions of ΔR between the closest Higgs Candidate jets, ΔR between the other two, and X_{Wt} before and after CR 1 derived reweighting for the 2016 Control Region 1.




Figure A.2: (Non-training variables) Distributions of p_T of the 1st and 3rd leading Higgs Candidate jets before and after CR 1 derived reweighting for the 2016 Control Region 1.




Figure A.3: (Non-training variables) Distributions of $|\Delta\eta_{HH}|$, X_{HH} , and of m_{HH} before and after CR 1 derived reweighting for the 2016 Control Region 1.




Figure A.4: (Training variables) Distributions of ΔR between the closest Higgs Candidate jets, ΔR between the other two, and X_{Wt} before and after CR 1 derived reweighting for the 2017 Control Region 1.




Figure A.5: (Non-training variables) Distributions of p_T of the 1st and 3rd leading Higgs Candidate jets before and after CR 1 derived reweighting for the 2017 Control Region 1.




Figure A.6: (Non-training variables) Distributions of $|\Delta\eta_{HH}|$, X_{HH} , and of m_{HH} before and after CR 1 derived reweighting for the 2017 Control Region 1.




Figure A.7: (Training variables) Distributions of ΔR between the closest Higgs Candidate jets, ΔR between the other two, and X_{Wt} before and after CR 1 derived reweighting for the 2018 Control Region 1.




Figure A.8: (Non-training variables) Distributions of p_T of the 1st and 3rd leading Higgs Candidate jets before and after CR 1 derived reweighting for the 2018 Control Region 1.




Figure A.9: (Non-training variables) Distributions of $|\Delta\eta_{HH}|$, X_{HH} , and of m_{HH} before and after CR 1 derived reweighting for the 2018 Control Region 1.




Figure A.10: (Training variables) Distributions of the minimum and maximum di-jet masses, and the top veto variable, X_{Wt} , before and after CR 1 derived reweighting for the all years inclusive Control Region 1.




Figure A.11: (Non-training variables) Distributions of ΔR between the other two HCs and the number of jets before and after CR 1 derived reweighting for the all years inclusive Control Region 1.




Figure A.12: Distributions of $|\Delta\eta_{HH}|$, the p_T of the HH system ($p_{T,HH}$), and m_{HH} before and after CR 1 derived reweighting for the all years inclusive Control Region 1.