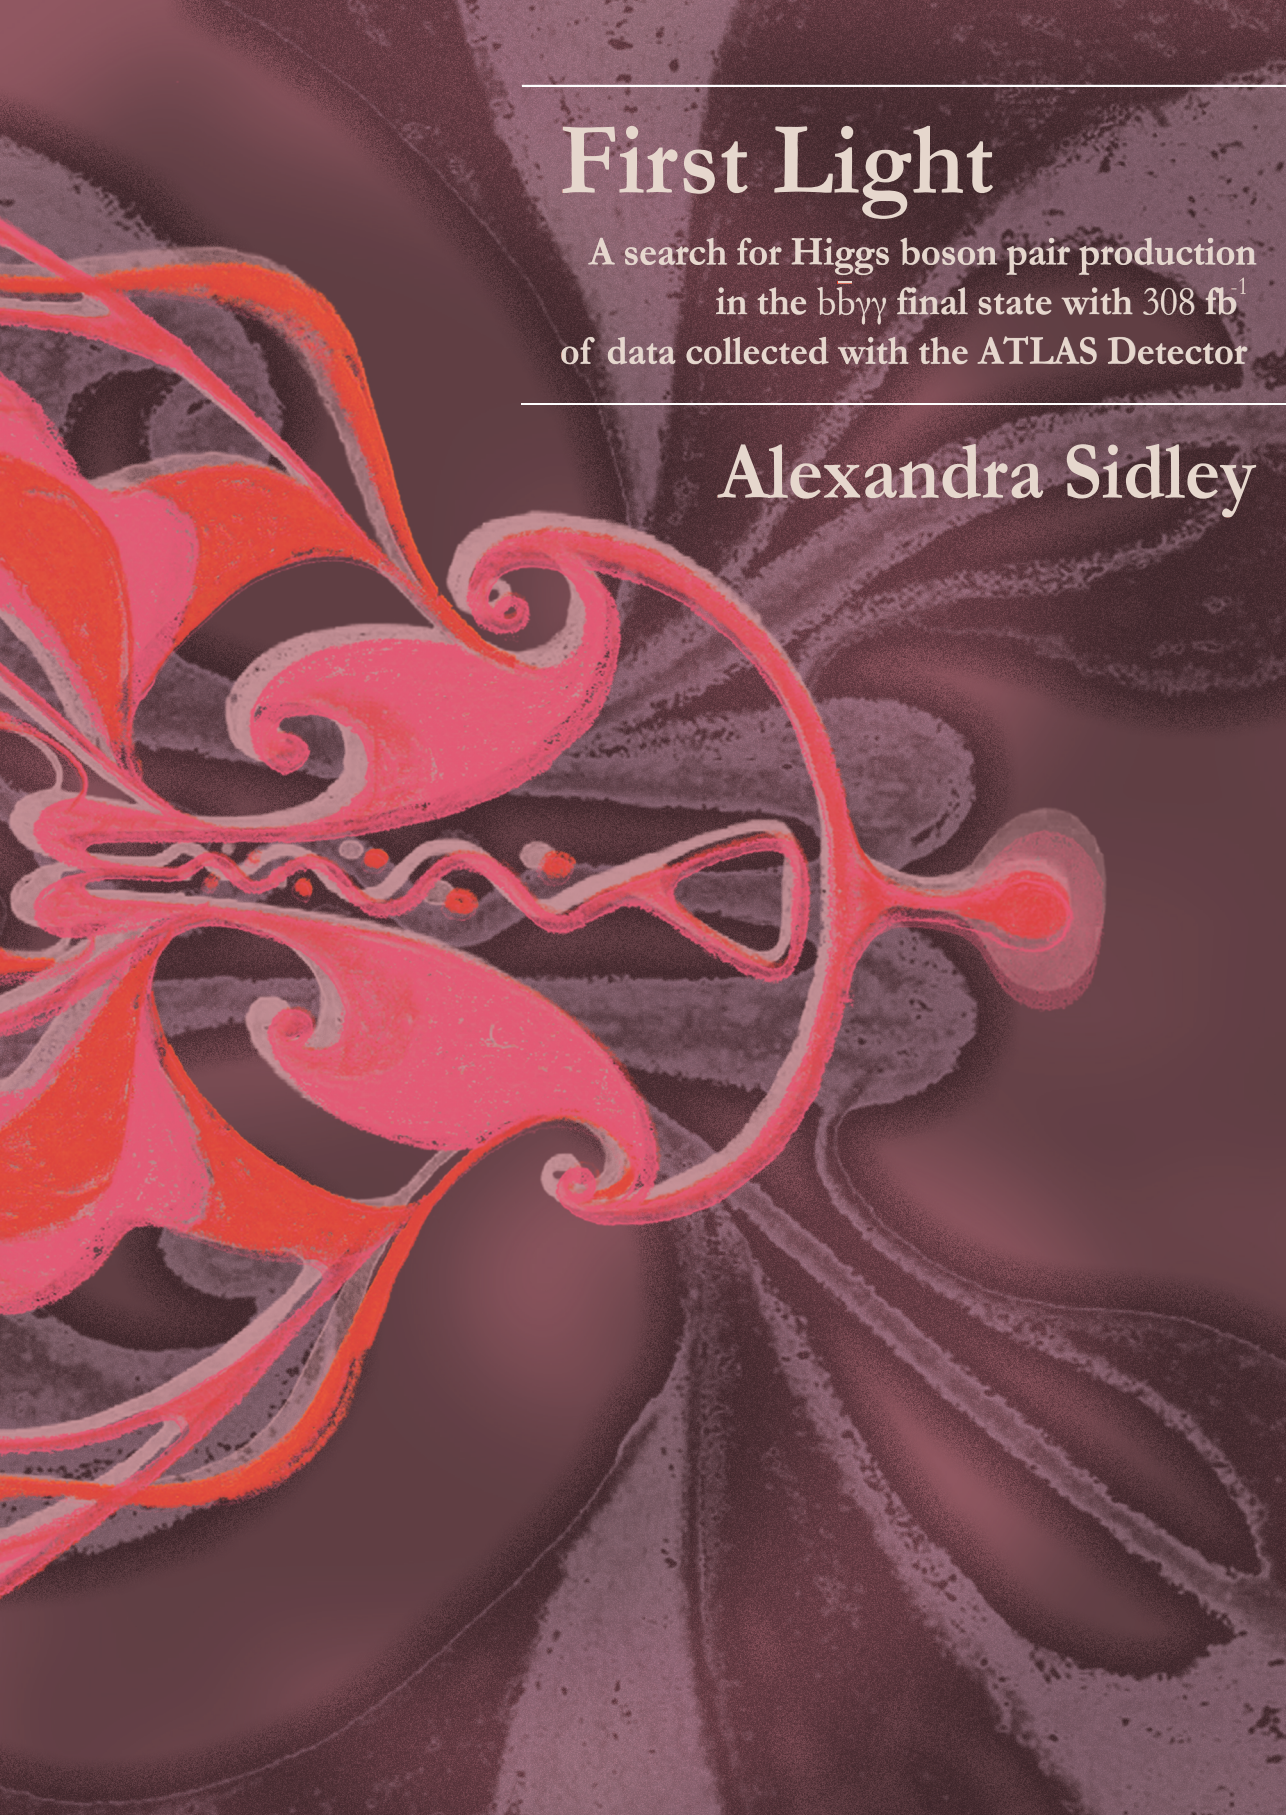

First Light

A search for Higgs boson pair production in the $b\bar{b}\gamma\gamma$ final state with 308 fb^{-1} of data collected with the ATLAS Detector

Alexandra Sidley

FIRST LIGHT

A search for Higgs boson pair production in the $b\bar{b}\gamma\gamma$ final state with 308 fb^{-1} of data collected with the ATLAS Detector

Alexandra Sidley

FIRST LIGHT

A search for Higgs boson pair production in the $b\bar{b}\gamma\gamma$ final state with 308 fb^{-1} of data collected with the ATLAS Detector

DOCTORAL THESIS

to obtain
the degree of doctor at the University of Twente,
on the authority of the rector magnificus,
prof.dr.ir. A. Veldkamp,
on account of the decision of the Doctorate Board,
to be publicly defended
on Tuesday the 8th of December 2026 at 12.30 hours

by

Alexandra Claire Sidley

This doctoral thesis has been approved by:

the promotor:
prof.dr. T.A. du Pree

the co-promotor:
dr. L. Brenner

Cover design: Théo Lange
Printed by: Proefschriftmaken
Layout: Alexandra Sidley
ISBN (print): 978-90-365-7474-7
ISBN (digital): 978-90-365-7475-4
URL: https://doi.org/10.3990/1.9789036574754

© 2026 Alexandra Sidley, The Netherlands. All rights reserved. No parts of this thesis may be reproduced, stored in a retrieval system or transmitted in any form or by any means without permission of the author. Alle rechten voorbehouden. Niets uit deze uitgave mag worden vermenigvuldigd, in enige vorm of op enige wijze, zonder voorafgaande schriftelijke toestemming van de auteur.

MANUSCRIPT COMMITTEE:

Chair/secretary:	prof.dr. J.J.L.M. Cornelissen (University of Twente)
Promotor:	prof.dr. T.A. du Pree (University of Twente)
Co-promotor:	dr. L. Brenner (Nikhef)
Members:	prof.dr.ir. A. Brinkman (University of Twente)
	dr. J.A. de Vries (Universiteit Maastricht)
	dr. L. Franco (DESY)
	prof.dr. W. Verkerke (Universiteit van Amsterdam)
	prof.dr.ir. S. Vanapalli (University of Twente)

Contents

1 Introduction **1**
1.1 Thesis Outline and Personal Contribution 2

2 Theoretical Foundations **5**
2.1 Foundations of the Standard Model 6
2.1.1 Fundamental Particles and Forces 6
2.1.2 Gauge Invariance and Quantum Electrodynamics 7
2.1.3 Quantum Chromodynamics 9
2.1.4 The Weak Interaction and Electroweak Theory 10
2.2 Electroweak Symmetry Breaking and the Higgs Mechanism . . . 12
2.2.1 The Higgs Field and its Potential 12
2.2.2 Spontaneous Symmetry Breaking 13
2.2.3 Gauge Boson Masses 14
2.2.4 Fermion Masses 15
2.2.5 The Physical Higgs Boson and the Self-Coupling 16
2.3 Limitations of the Standard Model 17
2.3.1 Observational Incompleteness 18
2.3.2 Theoretical Tensions 18
2.3.3 The Shape of the Higgs Potential 19

3 The ATLAS Experiment at the Large Hadron Collider **21**
3.1 The Large Hadron Collider 22
3.1.1 Luminosity and Pile-Up 24
3.1.2 LHC Run Periods 25
3.2 The ATLAS Detector 28
3.2.1 The ATLAS Coordinate System 28
3.2.2 The Inner Detector 30

3.2.3 Calorimeter System . . . 31
3.2.4 Muon Spectrometer . . . 34
3.2.5 Trigger and Data Acquisition System . . . 36
3.3 Upgrades for the High-Luminosity LHC . . . 38

4 Object Reconstruction **43**
4.1 Tracks and Vertices . . . 44
4.2 Topological Clusters . . . 45
4.3 Electrons and Photons . . . 46
4.4 Jets . . . 49
4.4.1 b-jets . . . 51
4.5 Muons . . . 54
4.6 Missing Transverse Energy . . . 55
4.7 Large-R Jets . . . 55
4.7.1 Calibration of the GN2X Tagger . . . 56

5 Event Generation **59**
5.1 Phenomenology of proton-proton collisions . . . 59
5.1.1 Parton distribution functions and factorisation . . . 60
5.1.2 Matrix element Calculation . . . 62
5.1.3 Parton Showering and Hadronisation . . . 62
5.2 Event Generators . . . 63
5.3 Detector Simulation . . . 64

6 Higgs Boson Phenomenology at the LHC **65**
6.1 Higgs Boson Production Mechanisms . . . 66
6.2 Higgs Boson Decay Channels . . . 68
6.3 Higgs Boson Couplings . . . 70
6.4 Higgs Boson Pair Production . . . 74
6.4.1 State of the Art of Higgs Boson Pair Production Searches with the ATLAS Experiment . . . 77
6.4.2 Triple Higgs . . . 80

7 The Search for Non-Resonant Higgs Boson Pair Production in the $b\bar{b}\gamma\gamma$ Final State **83**
7.1 Motivation . . . 83
7.2 Overview . . . 84
7.3 Data and Simulation Samples . . . 85
7.4 κ_λ and κ_{2V} Parameterisation . . . 89

7.5 Object-Level Selection 94
7.6 Kinematic Fit 97
7.6.1 Kinematic Fit likelihood construction 98
7.6.2 Kinematic Fit Second Iteration 104
7.7 Event-level Selection 107
7.8 Event Categorisation 108
7.9 Signal and Background Modelling 112
7.9.1 Signal and single-Higgs background 112
7.9.2 Continuum Background 114
7.10 Run 2 and Run 3 Correlation Strategy 115
7.11 Systematic Uncertainties 116
7.11.1 Experimental Uncertainties 117
7.11.2 Theoretical uncertainties 120
7.12 Statistical Analysis 124
7.13 Results 126
7.13.1 $HH \to b\bar{b}\gamma\gamma$ main results 126
7.13.2 Kinematic Fit Sensitivity Study 134
7.14 Summary 141

8 Prospects for a future Higgs pair and triple-Higgs combination 143
8.1 Indirect Probes of Coupling Parameters 144
8.2 Towards a Simultaneous κ_λ, κ_4 constraint in ATLAS 147

9 Conclusion and Outlook 153

10 Summary 157

11 Samenvatting 167

A Simulation and Event Selection 177
A.1 Signal Cutflows 177
A.2 Monte Carlo Yields 179
A.3 Yields by Analysis Category 180

B Signal and Background Modelling 183
B.1 Signal Model Fits 183
B.2 Background Template Validation 188
B.3 Combined Fits to the Diphoton Invariant Mass 194

C Kinematic Distributions of the κ_3, κ_4 Samples **197**
C.1 Distributions by Grid Point 197
C.2 Validation Against the Nominal Samples 201

Bibliography **203**

Acknowledgements **219**

Introduction 1

The last century of scientific inquiry into the physical world at its smallest scale has led us to the Standard Model – an astonishingly successful theory that governs the fundamental particles that make up matter and the forces between them. The Standard Model has displayed remarkable predictive power and held up against numerous experimental tests. The most recent and powerful example of this is the 2012 discovery of the Higgs boson by the ATLAS and CMS experiments [1, 2], confirming predictions made several decades earlier and ushering in a new era of research in collider physics.

This new era has been concerned with mapping out the properties of the Higgs boson and its interaction with the Standard Model. Among these, the least constrained is the Higgs boson's self-interaction, leaving open the question: what is the structure of the Higgs potential? This question is critical not only for better understanding this recently-discovered particle, but for uncovering the shortcomings of the Standard Model itself. It is widely known that this theory, despite its successes, is incomplete. It offers no explanation for several observed phenomena, including the presence of dark matter in galaxies and the matter-anti-matter asymmetry of the universe. For some of these phenomena, an explanation lies in modifications to the Higgs potential.

The shape of the Higgs potential is set by the Higgs boson self-coupling pa-

rameter λ, which is directly accessible only through Higgs boson pair production, a rare process in the Standard Model. Of the possible Higgs boson pair production final states accessible at the Large Hadron Collider (LHC), the one containing a pair of photons and a pair of (anti-)$b-$quarks, the $HH \to b\bar{b}\gamma\gamma$ channel, is the most sensitive to λ due to its unique trade-off between the large branching ratio of the $H \to b\bar{b}$ decay and the good mass resolution of the $H \to \gamma\gamma$ decay.

This thesis presents the most recent search for Higgs boson pair production in the $HH \to b\bar{b}\gamma\gamma$ channel. Importantly, it uses data from the ATLAS detector taken over the full Run 2 period and, for the first time in a Higgs pair production search, part of the Run 3 period. This larger dataset provides a significant increase in sensitivity, boosted even further by the use of an improved machine learning algorithm to identify b-jets and a kinematic fit to improve the mass resolution of the $H \to b\bar{b}$ system. Finally, this thesis looks forward to the possibility of a first simultaneous constraint on the trilinear and quartic couplings using ATLAS data.

1.1 Thesis Outline and Personal Contribution

The work presented in this thesis was performed within the ATLAS Collaboration and in close collaboration with members of the $HH \to b\bar{b}\gamma\gamma$ analysis team. Moreover, it rests on countless contributions from members of the ATLAS Collaboration at every step of the data-taking and reconstruction process, from hardware to computing and beyond. This section outlines the content of each chapter of the thesis and highlights my personal contributions to the work within this collaboration.

Chapter 2 provides a theoretical background to and motivation for the search performed in this thesis. The Standard Model and its particles and forces are introduced, along with the concept of gauge invariance, electroweak symmetry breaking and the Higgs boson. The limitations of the Standard Model are briefly discussed.

Chapter 3 sets the physical stage for the search performed in this thesis: the ATLAS Experiment at the Large Hadron Collider. The ATLAS subdetectors, trigger and data acquisition systems are described. My contributions to hard-

ware work for the upgrade to the High-Luminosity LHC are discussed, namely the replacement of ELMB2 modules in the muon spectrometer and the alcohol flushing of ITk cooling manifolds.

Chapter 4 describes how physics objects suitable for analysis are reconstructed from detector signals. I contributed to the first calibration of GN2X, the transformer-based graph neural network algorithm that identifies and classifies Large-Radius jets according to their flavour origin.

Chapter 5 describes the phenomenology of proton-proton collisions and the generation of events using Monte Carlo simulation.

Chapter 6 introduces the phenomenology of the Higgs boson and Higgs boson pair production at the LHC. Both single-Higgs and Higgs boson pair production and decay modes are discussed.

Chapter 7 contains the main body of this work, the $HH \to b\bar{b}\gamma\gamma$ analysis. Each step of the analysis chain is described in detail. I contributed to the technical implementation and validation of the kinematic fit: this includes developing the kinematic fit algorithm to function within the ntuple-processing software, testing it on the new ntuples and implementing it within the sequence of algorithms in the analysis framework. I also performed a sensitivity study on the kinematic fit, performing the full analysis chain on inputs with and without the kinematic fit to validate its effects and compare the results: the estimated improvements in sensitivity resulted in the adoption of the kinematic fit for the final analysis strategy. I performed several steps of supporting and validation work in other steps of the analysis, including new ntuple validation through cutflows and variable checks, background validation studies, spurious signal checks and theoretical systematic uncertainty estimation.

Chapter 8 provides an outlook on a novel combination of Higgs boson pair production and triple-Higgs measurements to constrain the trilinear and quartic self-coupling simultaneously. In collaboration with a team of three people, I generated and validated simulation samples of Higgs boson pair production in the $b\bar{b}\gamma\gamma$ final state that account for new second order electroweak calculations. This allows for the production within the ATLAS software framework of Higgs boson pair production samples that are sensitive to both coupling modifiers. Although this study could not be completed within the time constraints of my thesis, the prospects and suggested next steps are discussed.

Theoretical Foundations 2

Our present understanding of fundamental physics is the result of a rich interplay between theory and experiment conducted over centuries. As technology progresses to allow for more sophisticated experiments, the space of what can be tested empirically grows, and theories expand and evolve to accommodate this new space. In parallel, when novel and compelling ideas are brought forth by theorists, experiments are developed to attempt to test and falsify these ideas. Over the last half of the last century, this interplay has led to the development of a coherent theoretical framework of fundamental physics known as the Standard Model (SM) [3–5]. The SM, which describes all known fundamental particles and interactions except for gravity, is our most successful theory to date. It has withstood experimental tests of many kinds and made numerous predictions subsequently confirmed by experiments.

Despite its long-standing success, however, the SM is widely believed to be incomplete, failing to address several key questions and observations about the universe. This chapter provides an overview of the theoretical foundations of the Standard Model and its limitations, setting the stage and providing context for the subsequent discussion of Higgs boson pair production. Section 2.1 introduces the Standard Model, the fundamental particles and forces it describes and its underlying principles. Section 2.2 discusses the mechanism of electroweak symmetry breaking and the role of the Higgs boson. Finally, Sec-

tion 2.3 outlines some of the limitations of the Standard Model, motivating the search for new physics beyond it.

2.1 Foundations of the Standard Model

2.1.1 Fundamental Particles and Forces

The Standard Model describes elementary particles which cannot be broken down into more fundamental constituents. These particles are categorised into two groups: fermions and bosons. These groups are distinguished by their spin, a quantum number analogous to a particle's intrinsic angular momentum. Fermions are the building blocks of matter and are characterised by a half-integer spin. Bosons are the mediators of forces between particles and are characterised by an integer spin. Figure 2.1 shows a schematic of the Standard Model particles and their properties.

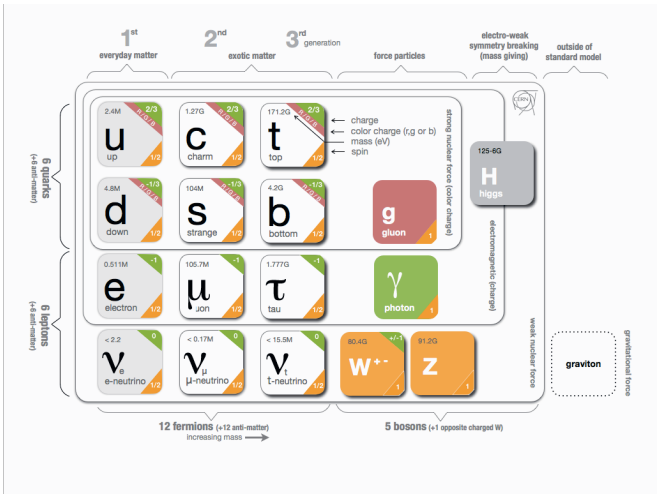

Figure 2.1: The fundamental particles of the Standard Model. The fermions are shown on the left, grouped by generation, and the bosons on the right. Each particle is shown with its mass and relevant quantum numbers. Taken from [6].

The bosons include the photon, a massless and chargeless spin-1 particle

that mediates the electromagnetic force; the charged $W^{\pm}$ and neutral Z bosons, massive spin-1 mediators of the weak force; the gluons, massless spin-1 particles that mediate the strong force; and the Higgs boson, the only scalar spin-0 particle in the SM, responsible for the generation of the masses of all other particles. The electromagnetic force is governed by the formalism of Quantum Electrodynamics (QED) and the strong force is described by Quantum Chromodynamics (QCD), both of which are described in subsequent sections.

The fermions in the Standard Model have a further structure, divided into two types of particles of three generations each. The leptons consist of the electron, muon and tau, along with their corresponding neutrinos. The quarks consist of the up and down quarks in the first generation, charm and strange quarks in the second generation and bottom and top quarks in the third generation. Each particle has a corresponding anti-particle with an opposite electromagnetic charge. The three generations of quarks and leptons carry identical charges and are distinguished primarily by their masses. The first generation fermions (the electron and the up and down quarks) are light and stable and make up the ordinary matter seen around us. The second and third generation fermions are heavier and less stable, and tend to decay into the lighter first generation particles. Leptons only interact via the electroweak force, while quarks interact via both the strong and electroweak force.

2.1.2 Gauge Invariance and Quantum Electrodynamics

The Standard Model is a quantum field theory in which particles manifest as excitations of underlying quantum fields that permeate all of spacetime. The dynamics of these fields characteristically follow the principle of gauge invariance, which requires that fundamental interactions be invariant under local transformations of certain symmetry groups. In particular, the SM is based on the non-abelian gauge group $SU(3)_C \times SU(2)_L \times U(1)_Y$, where $SU(3)_C$ corresponds to the strong interaction that acts on the colour charge carried by quarks and gluons and $SU(2)_L \times U(1)_Y$ to the electroweak interaction, discussed further in Section 2.1.4.

Quantum Electrodynamics (QED), which describes the interactions of electromagnetically charged particles, is the most straightforward manifestation of the principle of gauge invariance and serves as a prototype for the more complex QCD and electroweak theories. In QED, a free fermion is represented by

the Dirac Lagrangian

$$\mathcal{L}_{\text{free}} = \bar{\psi}(i\gamma^\mu \partial_\mu - m)\psi \tag{2.1}$$

where ψ is the fermion field, $\bar{\psi} = \psi^\dagger \gamma^0$, γ^μ are the Dirac gamma matrices and m is the mass of the field. This Lagrangian is invariant under a global U(1) transformation

$$\psi \to e^{i\alpha}\psi \tag{2.2}$$

where α is a constant phase. By Noether's theorem, this invariance implies a conserved quantity, in this case the electric charge q. However, the SM is based on the stronger requirement of *local* gauge invariance, in which the phase α is a function of spacetime. The transformation becomes

$$\psi \to e^{i\alpha(x)}\psi. \tag{2.3}$$

The Dirac Lagrangian in Equation 2.1 is not invariant under this local transformation. To restore invariance, the derivative is replaced by the covariant derivative

$$D_\mu = \partial_\mu - iqA_\mu, \tag{2.4}$$

where A_μ is a gauge field that transforms as $A_\mu \to A_\mu + \frac{1}{q}\partial_\mu \alpha(x)$ and q is the electromagnetic charge. Substituting this, the Lagrangian becomes

$$\mathcal{L} = \bar{\psi}(i\gamma^\mu D_\mu - m)\psi. \tag{2.5}$$

Expanding the covariant derivative in Equation 2.5 gives the free term $\mathcal{L}_{\text{free}}$ plus a second term that represents the interaction between the field ψ and the gauge field A_μ. This is a natural consequence of the imposition of local gauge invariance. The gauge field A_μ in fact represents the photon, the mediator of the electromagnetic force. Finally, a kinematic term for the gauge field itself is added to the Lagrangian, giving the full QED Lagrangian

$$\mathcal{L}_{\text{QED}} = \bar{\psi}(i\gamma^\mu D_\mu - m)\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu}, \tag{2.6}$$

where $F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu$ is the electromagnetic field strength tensor. Importantly, to add a mass term $\frac{1}{2}m_\gamma^2 A_\mu A^\mu$ for the photon would break the gauge invariance. It then follows that the photon must be massless.

The QED Lagrangian, with the principle of gauge invariance imposed, now encodes the dynamics of electromagnetically charged particles and forces the

existence of a spin-1 gauge boson. This principle underpins the physics of the full Standard Model. In a similar way, requiring local gauge invariance under transformations of the SU(3)$_C$ and SU(2)$_L$ $\times$ U(1)$_Y$ symmetry groups leads to the formulation of Quantum Chromodynamics and the electroweak theory, respectively, described in the following sections.

2.1.3 Quantum Chromodynamics

Quantum Chromodynamics (QCD) is the non-abelian gauge theory that describes the strong interaction between quarks and gluons. It is based on the SU(3)$_C$ symmetry group, where the subscript C denotes the colour charge carried by quarks and gluons. Quarks transform as triplets under SU(3)$_C$. This means that they carry one of three colour charges (red, green, or blue) and anti-quarks carry the corresponding anti-colour charge. They are represented by the quark fields q_f, where $f = u, d, c, s, t, b$ denotes the quark flavour. Each q_f is a three-component colour triplet, with the colour index left implicit throughout this section.

Gluons, the massless gauge bosons of QCD, are represented by the gluon fields G^a_μ. The indices $a = 1, 2, ...8$ correspond to the eight generators of the SU(3) group $T^a = \frac{\lambda^a}{2}$, where λ_a are the Gell-Mann matrices [7]. Gluons carry a combination of color and anti-colour charge, which allows them to interact with themselves as well as with quarks.

The QCD Lagrangian is given by

$$\mathcal{L}_{\text{QCD}} = \sum_f \bar{q}_f \left(i\gamma^\mu D_\mu - m_f\right) q_f - \frac{1}{4} G^a_{\mu\nu} G^{\mu\nu}_a \tag{2.7}$$

where $D_\mu = \partial_\mu - i g_s G^a_\mu T^a$ is the covariant derivative introduced to enforce local SU(3)$_C$ gauge invariance, g_s is the strong coupling constant and m_f is the mass of the quark of flavour f. The first term describes the free propagation of each quark, together with its interaction with the gluon field, while the second term describes the gluon three- and four-point self-interactions. Unlike the electroweak case discussed in Section 2.1.4, left- and right-handed quarks transform identically under SU(3)$_C$, so this explicit mass term does not violate colour gauge invariance.

The non-abelian nature of $SU(3)_C$ gives rise to the phenomena of confinement and asymptotic freedom. Confinement [8] refers to the fact that free quarks and gluons cannot be observed in isolation. Indeed, despite the overwhelming evidence for the existence of quarks, they have never been observed as free particles. They are instead confined to colourless bound states known as hadrons, which include quark-anti-quark pairs called mesons and three-quark particles called baryons. Asymptotic freedom [9] describes the behaviour of quarks as free, non-interacting particles at asymptotically high energies or short distances.

Both confinement and asymptotic freedom are consequences of the running coupling of QCD. In general, the energy dependence of coupling constants arises from the procedure of renormalisation, by which divergences that appear in loop integrals within perturbative quantum field theory are systematically absorbed into redefinitions of the theory's parameters [10]. The strong coupling $\alpha_s(Q^2) = \frac{g_s(Q^2)}{4\pi}$ is therefore not a true constant but is dependent on the energy scale Q^2 of the interaction. At low energies, $\alpha_s(Q^2)$ becomes large, leading to confinement. At high energies, $\alpha_s(Q^2)$ becomes small, allowing quarks to behave as quasi-free particles. The energy threshold between these regimes is known as the QCD scale $\Lambda_{\text{QCD}} \approx 200\,\text{MeV}$, a fundamental parameter of QCD. At energies below Λ_{QCD}, the strong coupling becomes too large for perturbation theory to be applicable, and non-perturbative methods such as lattice QCD are required to study the properties of hadrons and their interactions. At energies above Λ_{QCD}, which are achieved at hadron colliders such as the LHC, perturbative QCD can be used to make predictions.

2.1.4 The Weak Interaction and Electroweak Theory

The weak interaction governs the interactions between quarks and leptons through the exchange of the massive gauge bosons $W^{\pm}$ and Z. It is responsible for processes such as beta decay, in which a neutron decays into a proton, an electron and an anti-neutrino. The exchange of $W^{\pm}$ bosons allows quarks and leptons to change flavour via charged current. Interactions involving the Z boson, known as neutral current, do not change the flavour of the particles involved at tree-level. Flavour-changing decays are governed by the Cabibbo-Kobayashi-Maskawa (CKM) matrix which encodes the relative probability of a quark of a given flavour transitioning into another [11, 12].

The symmetry group associated with the weak interaction is $\mathrm{SU}(2)_L$. The subscript L denotes that only fermions with left-handed chirality transform under this group. Right-handed fermions are represented as singlets and therefore do not transform under $\mathrm{SU}(2)_L$. This chiral structure encodes the parity violation of the weak interaction, the possibility of which was raised by Lee and Yang in 1956 [13] and which was demonstrated experimentally the following year by Wu et al. [14].

Gauge invariance can be applied naïvely to the $\mathrm{SU}(2)_L$ group in a similar way to the QED case. This leads to the introduction of three gauge fields W_μ^a corresponding to the three generators of $\mathrm{SU}(2)_L$, known as the weak isospin, $T^a = \sigma^a/2$ ($a = 1, 2, 3$), where σ^a are the Pauli matrices and g is the associated coupling constant. Note that T^a here denotes the generators of $\mathrm{SU}(2)_L$ in the fundamental representation, and should not be confused with the $\mathrm{SU}(3)_C$ generators of the same name introduced in Section 2.1.3. These gauge fields, however, do not directly correspond to the physical $W^\pm$ and Z bosons for the following reasons:

- The Z boson couples to both left-handed and right-handed fermions, while the W_μ^3 gauge field couples only to left-handed fermions due to the chiral coupling structure of $\mathrm{SU}(2)_L$.
- It is known from experiment that the $\mathrm{W}^\pm$ and Z bosons are massive, but the gauge fields W_μ^a are massless by consequence of gauge invariance. This problem is resolved by the Higgs mechanism described in Section 2.2.

The solution to the former is electroweak unification, which extends the gauge group to $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$, unifying the weak interaction with the electromagnetic interaction. The $\mathrm{U}(1)_Y$ group has its own associated gauge field, B_μ, with coupling constant g', playing an analogous role for $\mathrm{U}(1)_Y$ to the one that W_μ^a and g play for $\mathrm{SU}(2)_L$. The subscript Y denotes the weak hypercharge, which was introduced by Glashow, Weinberg and Salam [15–17]. The third component of the weak isospin, T_3, and the weak hypercharge Y are related to the electromagnetic charge Q through the Gell-Mann-Nishijima relation

$$Q = T_3 + \frac{Y}{2}. \tag{2.8}$$

The covariant derivative for electroweak theory is given by

$$D_\mu = \partial_\mu - igW_\mu^a T^a - ig'\frac{Y}{2}B_\mu, \tag{2.9}$$

where g and g' are the coupling constants for $\mathrm{SU}(2)_L$ and $\mathrm{U}(1)_Y$, respectively. The physical $\mathrm{W}^{\pm}$ bosons are obtained through linear combinations of the W_μ^1 and W_μ^2 gauge fields:

$$\begin{pmatrix} W_\mu^+ \\ W_\mu^- \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & -i \\ 1 & i \end{pmatrix} \begin{pmatrix} W_\mu^1 \\ W_\mu^2 \end{pmatrix}. \tag{2.10}$$

The neutral Z boson and photon are recovered by introducing a neutral gauge field B_μ that mixes with the W_μ^3 field as follows:

$$\begin{pmatrix} Z_\mu \\ A_\mu \end{pmatrix} = \begin{pmatrix} \cos\theta_W & -\sin\theta_W \\ \sin\theta_W & \cos\theta_W \end{pmatrix} \begin{pmatrix} W_\mu^3 \\ B_\mu \end{pmatrix}, \tag{2.11}$$

where θ_W is the Weinberg angle defined by the relation $\tan\theta_W = \frac{g'}{g}$.

The theory now describes the weak interaction with the correct chirality, with the physical $\mathrm{W}^{\pm}$, Z and photon states recovered from the gauge fields W_μ^a and B_μ through the mixing relations above. The photon is indeed observed to be massless, while the $\mathrm{W}^{\pm}$ and Z bosons have measured, non-zero masses.

2.2 Electroweak Symmetry Breaking and the Higgs Mechanism

The electroweak theory described in Section 2.1.4 successfully unifies the electromagnetic and weak interactions under the gauge group $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$, but it does so at a cost: gauge invariance forbids explicit mass terms for the $W^{\pm}$ and Z bosons. Since the masses of these particles are established experimental facts, a mechanism to generate them while respecting the gauge invariance of the Lagrangian is required. The solution, proposed independently by Brout and Englert [18], Kibble, Hagen and Guralnik [19], and Peter Higgs [20] in 1964, is known as electroweak symmetry breaking (EWSB) or the Higgs mechanism. It exploits spontaneous symmetry breaking, in which the Lagrangian itself retains full gauge symmetry but the ground state of the theory does not.

2.2.1 The Higgs Field and its Potential

The Higgs mechanism introduces a new complex scalar field, the Higgs field, which transforms as a doublet under $\mathrm{SU}(2)_L$ with weak hypercharge $Y = 1$:

$$\Phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}. \tag{2.12}$$

The most general renormalisable Lagrangian for Φ that is invariant under the full $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ gauge symmetry is

$$\mathcal{L}_{\text{Higgs}} = (D_\mu \Phi)^\dagger (D^\mu \Phi) - V(\Phi), \tag{2.13}$$

where D_μ is the electroweak covariant derivative from Equation 2.9. The Higgs potential $V(\Phi)$ takes the form

$$V(\Phi) = \mu^2 \Phi^\dagger \Phi + \lambda \left(\Phi^\dagger \Phi\right)^2, \tag{2.14}$$

where μ^2 and λ are free parameters of the theory. The requirement that the potential be bounded from below imposes $\lambda > 0$, ensuring that the field energy cannot decrease without limit. The qualitative behaviour of $V(\Phi)$ depends on the sign of μ^2. For $\mu^2 > 0$, the potential has a single minimum at $\Phi = 0$, the symmetry of the ground state matches that of the Lagrangian, and no symmetry breaking occurs. For the Higgs mechanism to work as intended, $\mu^2 < 0$ is required. The shape of the potential for $\mu^2 < 0$ is shown in Figure 2.2. It has a local maximum at the origin and an infinite set of degenerate minima arranged on a circle, giving the potential its characteristic "Mexican hat" shape. The radius of the circle of minima defines the vacuum expectation value (VEV) v of the field,

$$v = \sqrt{-\frac{\mu^2}{\lambda}}, \tag{2.15}$$

which is a real, positive quantity.

2.2.2 Spontaneous Symmetry Breaking

Spontaneous symmetry breaking is characterised by a non-zero VEV: the ground state of the theory selects a particular point on the circle of minima. This is analogous to the ball rolling down the hill in Figure 2.2. By convention, the VEV is oriented in the neutral component of the doublet,

$$\langle \Phi \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}, \tag{2.16}$$

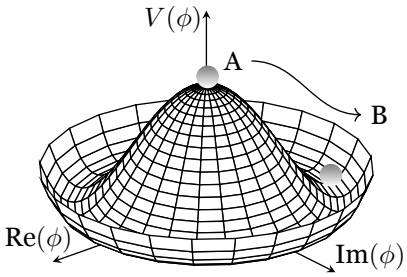

Figure 2.2: The shape of the Higgs potential $V(\Phi)$ for $\mu^2 < 0$ in the complex ϕ^0 plane. Adapted from [21].

a choice that breaks $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ down to $\mathrm{U}(1)_Q$. Three of the four generators of $\mathrm{SU}(2)_L \times \mathrm{U}(1)_Y$ are broken, while the combination $Q = T_3 + Y/2$ remains unbroken, keeping the photon massless and preserving the $\mathrm{U}(1)_Q$ symmetry of electromagnetism.

According to Goldstone's theorem [22], the spontaneous breaking of a continuous symmetry produces one massless scalar boson for each broken generator. Here, three such Goldstone bosons appear. The field can then be transformed by a chosen gauge called the unitary gauge such that the Goldstone components are set to zero. The Higgs field is left as a scalar fluctuation about the vacuum $H(x)$:

$$\Phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix}. \tag{2.17}$$

The three Goldstone degrees of freedom are not lost but rather absorbed into the longitudinal polarisation modes of the three massive gauge bosons, providing each with the third polarisation state of a massive spin-1 particle. This is the central result of the Higgs mechanism: the Goldstone bosons are "eaten" by the $W^{\pm}$ and Z bosons, which acquire mass in the process.

2.2.3 Gauge Boson Masses

The masses of the gauge bosons arise from the kinetic term $(D_\mu \Phi)^\dagger (D^\mu \Phi)$ in Equation 2.13. Substituting the field expansion of Equation 2.17 yields quadratic

mass terms for the gauge fields. Using the electroweak covariant derivative from Equation 2.9 and the mixing relations of Equation 2.11, these mass terms are

$$m_W = \frac{gv}{2}, \tag{2.18}$$

$$m_Z = \frac{v}{2}\sqrt{g^2 + g'^2} = \frac{m_W}{\cos\theta_W}, \tag{2.19}$$

$$m_\gamma = 0. \tag{2.20}$$

The photon remains massless, as required by the unbroken $\mathrm{U}(1)_Q$ symmetry. The ratio of the W and Z masses is fixed at tree level by the relation $\rho \equiv m_W^2/(m_Z^2 \cos^2\theta_W) = 1$. The measured masses $m_W \approx 80.4\,\mathrm{GeV}$ and $m_Z \approx 91.2\,\mathrm{GeV}$, together with the Fermi constant $G_F = 1/(\sqrt{2}\,v^2)$ measured in muon decay, fix the VEV to

$$v = \left(\sqrt{2}\,G_F\right)^{-1/2} \approx 246\,\mathrm{GeV}. \tag{2.21}$$

This value sets the scale of the particle masses.

2.2.4 Fermion Masses

The $\mathrm{SU}(2)_L$ gauge symmetry also prohibits explicit Dirac mass terms for fermions. The Higgs mechanism resolves this through gauge-invariant *Yukawa couplings* between the fermion fields and the Higgs doublet. For a fermion of flavour f, the Yukawa term is

$$\mathcal{L}_{\text{Yukawa}} = -y_f \left(\bar{\psi}_L\, \Phi\, \psi_R + \text{h.c.}\right), \tag{2.22}$$

where ψ_L and ψ_R are the left- and right-handed fermion fields and y_f is the Yukawa coupling, a free parameter for each fermion. Substituting the field expansion of Equation 2.17, the VEV generates a mass term for the fermion,

$$m_f = \frac{y_f\, v}{\sqrt{2}}, \tag{2.23}$$

while the coupling of the physical Higgs field H to the fermion is $-i \cdot m_f/v$. A critical consequence of this relationship is that the Higgs coupling to each fermion is directly proportional to the fermion mass. The top quark, with $m_t \approx 173\,\mathrm{GeV}$, therefore has a Yukawa coupling close to unity, while the electron

Yukawa coupling is of order 10^{-6}. The measurement of Higgs couplings to fermions thus constitutes a direct test of the mass-generation mechanism and is discussed further in Chapter 6.

2.2.5 The Physical Higgs Boson and the Self-Coupling

Substituting Equation 2.17 into the Higgs potential and using the minimum condition $\mu^2 = -\lambda v^2$, the potential expressed in terms of the physical field H becomes

$$V(H) = \lambda v^2 H^2 + \lambda v H^3 + \frac{\lambda}{4} H^4 + \text{const.} \tag{2.24}$$

Following electroweak symmetry breaking, the Higgs field manifests as a neutral spin-0 scalar particle called the Higgs boson. The quadratic term in Equation 2.24 identifies the Higgs boson mass,

$$m_H^2 = 2\lambda v^2 = -2\mu^2. \tag{2.25}$$

Since λ and μ^2 are free parameters of the Standard Model, the Higgs boson mass is not fixed by the theory and must instead be determined experimentally. Following the discovery of the Higgs boson by the ATLAS and CMS experiments in 2012 [1, 2], its mass has been measured with increasing precision; the most recent ATLAS combination gives $m_H = 125.11 \pm 0.11$ GeV [23]. Throughout this thesis, theoretical predictions and simulated samples are evaluated at the reference value $m_H = 125$ GeV, following the convention of the LHC Higgs Cross Section Working Group [24]; the difference with respect to the measured value has a negligible effect on the results presented here. Combined with the value of v obtained in Equation 2.21, this measurement fixes the value of λ,

$$\lambda = \frac{m_H^2}{2v^2} \approx 0.13. \tag{2.26}$$

The second and third terms in Equation 2.24 are self-interaction terms of the Higgs field which give rise to vertices involving three and four Higgs bosons. The cubic term corresponds to a trilinear Higgs self-interaction whose coupling strength is denoted by

$$\lambda_3 = \frac{3m_H^2}{v} = 6\lambda v, \tag{2.27}$$

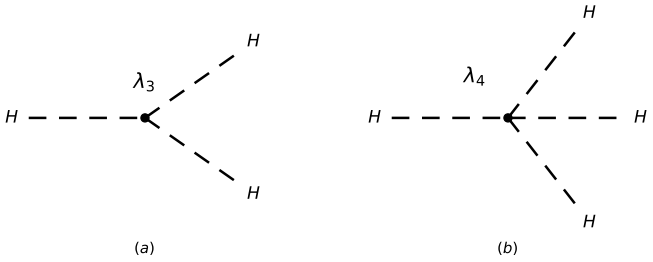

Figure 2.3: Feynman diagrams for the trilinear (a) and quartic (b) Higgs boson self-interactions.

while the quartic term corresponds to a quartic Higgs self-interaction with coupling

$$\lambda_4 = \frac{3m_H^2}{v^2} = 6\lambda. \tag{2.28}$$

Both vertices are shown in Figure 2.3.

The 2012 discovery of the Higgs boson confirmed the existence of the final particle predicted by the Standard Model, completing its particle content. This does not mean, however, that the electroweak symmetry breaking sector is fully understood: the measurement of m_H fixes λ only under the assumption that the potential in Equation 2.14 takes its minimal, renormalisable form. Whether this is true, and by extension, whether the values of λ_3 and λ_4 predicted above are correct, has not yet been established experimentally, as discussed in Section 2.3.3. The trilinear self-coupling λ_3 is accessible through Higgs boson pair production, providing a direct experimental test of the shape of the Higgs potential. This is the central topic of this work and is discussed in detail in Chapters 6 and 7.

2.3 Limitations of the Standard Model

Despite the remarkable success of the Standard Model in describing the known fundamental particles and their interactions, it is known to be an incomplete theory. A number of well-established experimental observations cannot be accommodated within its framework, and it contains some theoretical tensions.

These shortcomings motivate the search for physics beyond the Standard Model (BSM).

2.3.1 Observational Incompleteness

The most direct evidence that the SM is incomplete comes from cosmological and astrophysical observations. Galaxy rotation curves [25] and other astrophysical measurements indicate the presence of a form of matter that does not interact electromagnetically and is therefore invisible to conventional detectors. This dark matter accounts for approximately 27% of the energy content of the universe [26], yet the SM contains no candidate particle(s) for it. Similarly, the observed predominance of matter over antimatter in the universe requires departures from matter-antimatter symmetry beyond what the SM can provide [27]. While the SM does contain a source of CP-violation in the quark sector through the CKM matrix [11, 12] and, analogously in the lepton sector, through the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix [28], which encodes the mixing between neutrino flavour and mass eigenstates, it is insufficient by many orders of magnitude to generate the observed baryon asymmetry. Observation of neutrino oscillations [29, 30] demonstrates that neutrinos have non-zero masses, despite the SM requiring them to be strictly massless. Finally, the SM also provides no description of gravity, which becomes relevant at the Planck scale $M_{\text{Pl}} \approx 1.2 \times 10^{19}$ GeV.

2.3.2 Theoretical Tensions

Beyond these observational gaps, the SM also contains a feature that is theoretically unsatisfying: the Higgs boson mass is unusually sensitive to physics at very high energies. Unlike the masses of the gauge bosons and fermions, which are protected by symmetries of the SM, no such symmetry protects the Higgs mass. As a result, quantum loop corrections from any particle that couples to the Higgs field push its mass towards the highest energy scale of the theory. If the SM is valid up to a scale Λ (for instance the Planck scale), these corrections are of order $\delta m_H^2 \sim \Lambda^2/16\pi^2$, far larger than the measured value $m_H^2 \sim (125\,\text{GeV})^2$. Reproducing the small observed mass would require a bare mass parameter to cancel this correction to a precision of about one part in 10^{34}. This is a remarkably precise and seemingly unnatural coincidence. This tension between the small observed scale of electroweak physics and the much larger scale at which new physics is expected to appear is known as the hierarchy problem [31], and it strongly motivates the existence of new physics at

or below the TeV scale that would stabilise the Higgs mass without such an extreme cancellation.

2.3.3 The Shape of the Higgs Potential

Particularly relevant to this thesis is the fact that the Higgs potential introduced in Section 2.2 is not experimentally established in its full form. The potential $V(\Phi) = \mu^2\Phi^\dagger\Phi + \lambda(\Phi^\dagger\Phi)^2$ represents the simplest renormalisable choice consistent with gauge invariance, or a minimal Higgs potential. However, there is no fundamental principle that prevents additional higher-dimensional terms such as $(\Phi^\dagger\Phi)^3/\Lambda^2$ from contributing if the SM is an effective field theory valid below some cutoff scale Λ. Such terms would modify the shape of the potential and shift the Higgs self-couplings λ_3 and λ_4 away from their SM predictions without necessarily altering m_H or v. Additionally, such a term does not shift the two couplings equally. Expanding $(\Phi^\dagger\Phi)^3$ about the vacuum generates a correction to λ_4 approximately six times larger than the corresponding correction to λ_3 [32]. A non-minimal potential may therefore be substantially more visible in the quartic coupling than in the trilinear one. As shown in Section 2.2, both self-couplings are uniquely determined by the single parameter λ in the SM, and any deviation from the predicted values would constitute direct evidence of a non-minimal Higgs potential.

A further consequence concerns the long-term stability of the electroweak vacuum. The value of λ is in fact not constant but evolves with energy scale under the renormalisation group, similar to the running of the strong coupling constant described in Section 2.1.3. This running depends on the top Yukawa coupling y_t and the Higgs boson mass m_H. In particular, λ becomes negative for large y_t. With the current measured values of m_H and the top quark mass $m_t \approx 173\,\text{GeV}$, λ becomes negative at approximately 10^{10}–10^{12} GeV [33]. This renders the electroweak vacuum *metastable* rather than absolutely stable, which means that the minimum v is in fact a local minimum. Tunneling to the global minimum, which sits at a much higher energy scale, has a non-zero probability. Whether the true vacuum is absolutely stable, metastable, or unstable is highly sensitive to the precise values of m_H, m_t, and λ itself.

The experimental determination of the Higgs self-coupling λ_3 is therefore one of the most important outstanding measurements in particle physics. It provides sensitivity to BSM contributions that would otherwise be inaccessible through single-Higgs measurements alone. The phenomenology of Higgs

boson pair production, which is the most sensitive experimental probe to λ_3, is discussed in the Chapter 6.

The ATLAS Experiment at the Large Hadron Collider

3

The previous chapter formed the theoretical background to much of the experimental high-energy physics that is performed today and introduced the importance of the Higgs boson self-coupling. This chapter will describe the experimental apparatus capable of probing the Standard Model, and in particular, generating and detecting Higgs boson pair-production; namely, the Large Hadron Collider (LHC) [34, 35] and the ATLAS (A Toroidal LHC ApparatuS) experiment [36]. To study Higgs boson pair production experimentally, two requirements must be met:

- A large enough number of particle collisions with high enough energy to produce a statistically significant number of Higgs bosons
- Detectors sophisticated enough to reconstruct and identify their decay products

The former requirement is fulfilled by the LHC at CERN, the world's largest and highest-energy particle accelerator. The latter is fulfilled by the ATLAS experiment, one of the two general-purpose detectors on the LHC. The following sections contain an overview of the LHC and relevant concepts, followed by a description of the ATLAS detector and its subsystems.

3.1 The Large Hadron Collider

The LHC is a circular particle accelerator with a circumference of 26.7 km, located approximately 100 m underground beneath the French-Swiss border near Geneva. It is designed to collide protons at a centre-of-mass energy of up to 14 TeV. Heavy-ions are also collided at lower energies and intensities, but only the proton-proton collisions are relevant for this work.

Hydrogen atoms are stripped of their electrons to create protons which are then accelerated in multiple stages before being injected into the LHC beam pipe as 2 beams traveling in opposite directions. The acceleration chain, (see Figure 3.1) begins with a linear accelerator, which accelerates the protons to 50 MeV. They are then injected into the Proton Synchrotron Booster (PSB) and accelerated to 1.4 GeV. Next, they enter the Proton Synchrotron (PS), where they are accelerated to 25 GeV, before being injected into the Super Proton Synchrotron (SPS) and accelerated to 450 GeV. Finally, they are injected into the LHC into two separate beam pipes traveling in opposite directions, where they are accelerated to their final energy of up to 7 TeV per beam.

Collisions occur at four interaction points (IPs) around the LHC. These correspond to four experiments: ATLAS [36], CMS [38], LHCb [39] and ALICE [40]. The two largest, ATLAS and CMS, are general purpose experiments designed to deliver a large collision rate. Both cover an extensive range of SM and BSM searches and measurements. The purpose of having two independent general-purpose experiments is to be able to validate results with orthogonal, independent datasets. Additionally, the different detector designs of the ATLAS and CMS experiments and their independent collaborations allow for different areas of strength alongside the significant overlap in their physics programs.

LHCb and ALICE are designed for more specialised measurements with a smaller collision rate. The LHCb experiment is a single-arm forward spectrometer designed for B-physics and other heavy-flavour measurements. The ALICE experiment is designed to study the quark-gluon plasma using heavy-ion collisions.

The geometry of the LHC consists of eight arc sections and four straight sections, following the design of the former LEP collider [41], whose tunnel and infrastructure was reused to house the LHC. The arcs contain 1232 superconduct-

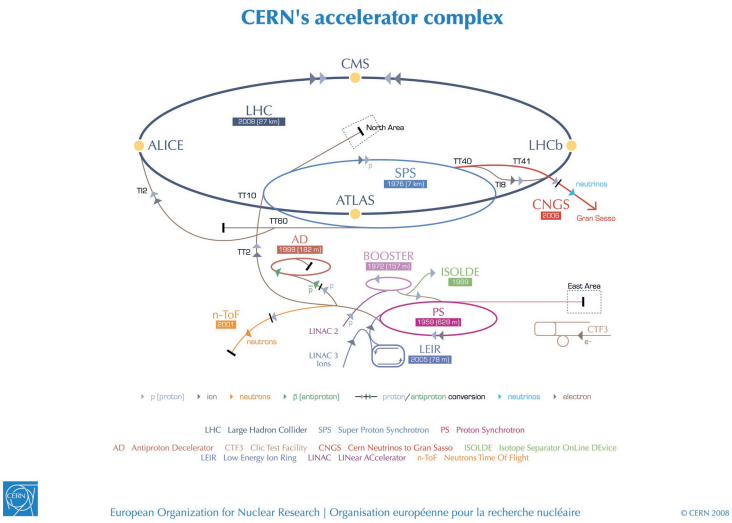

Figure 3.1: Schematic representation of the CERN accelerator complex, showing the different stages of acceleration and the locations of the experiments [37].

ing dipole magnets with a field strength of 8.33 T made from niobium-titanium coils to bend the beam around the ring, keeping it on a circular trajectory. The coils operate with a current of 11700 A and are cooled with superfluid helium down to 1.9 K. The straight sections serve as either experimental or utility regions. Acceleration of the beams occurs in RF cavities in the straight sections. A schematic of the full acceleration chain and the locations of the different experiments is shown in Figure 3.1.

The proton beams are grouped into 2808 bunches of approximately 1×10^{11} protons per bunch, with a nominal bunch spacing of 25 ns. Near the IPs, the proton bunches are focused by superconducting quadrupole magnets with a field gradient of approximately 233 T/m until the beam width is approximately $16\,\mu$m. In total, the LHC contains 858 quadrupole magnets used for beam focusing and over 6800 higher-order multipole magnets used as corrector magnets for beam distortion and other effects. Once the beams are accelerated and focused, they collide at the IPs, resulting in so-called bunch crossings at a rate of 40 MHz.

3.1.1 Luminosity and Pile-Up

Particle colliders are characterised by their centre-of-mass energy $\sqrt{s}$ and their instantaneous luminosity L. The centre-of-mass energy determines the maximum energy available to produce new particles in the collisions, while the instantaneous luminosity determines the collision rate at a given time. The instantaneous luminosity is defined as:

$$L = \frac{N_{\text{bunch}}^2 n_{\text{beam}} f_{\text{rev}} \gamma_{\text{r}}}{4\pi\epsilon_{\text{n}}\beta^*} F, \tag{3.1}$$

where N_{bunch} is the number of protons per bunch, n_{beam} is the number of bunches per beam, f_{rev} is the revolution frequency, γ_{r} is the relativistic gamma factor, ϵ_{n} is the normalised transverse beam emittance, β^* is the beta function at the collision point and F is a geometric factor accounting for the crossing angle of the beams at the IP. The LHC delivers a revolution frequency of $f_{\text{rev}} = 11.2\,\text{kHz}$, with a nominal design β^* of $0.55\,\text{m}$. At the nominal bunch size, spacing and population mentioned above, the LHC was designed to operate at a peak instantaneous luminosity of $10^{34}\ \text{cm}^{-2}\text{s}^{-1}$ [34]. In practice this design value was exceeded by roughly a factor of two, as discussed in Section 3.1.2.

The number of events of a given process with cross section σ produced at the LHC in a time period t is given by:

$$N_{\text{events}} = \sigma\mathcal{L} = \sigma \int L\text{d}t, \tag{3.2}$$

where $\mathcal{L}$ is the integrated luminosity over this time period. This quantity corresponds to the size of the collected dataset, a high value of which is crucial to collect a statistically significant number of events. A distinction is made between the integrated luminosity *delivered* by the LHC, that *recorded* by ATLAS, which is smaller due to trigger dead-time and detector readout inefficiency, and the subset certified as *good for physics*, for which all sub-detectors were fully operational and correctly calibrated. It is the last of these that defines the dataset available for analysis, and all luminosities quoted in this thesis refer to it unless stated otherwise.

Luminosity can be increased either by increasing the bunch revolution frequency, the number of bunches per beam or the number of protons per bunch; or by decreasing the total area of collision by further focusing the beams in

the transverse plane. This in turn increases the probability of interaction and therefore the total number of protons interacting in a single bunch. The integrated luminosity delivered by the LHC and recorded by ATLAS in each data-taking period is shown in Figure 3.3 and discussed in Section 3.1.2.

While high luminosity is advantageous for collecting large datasets, it can also negatively impact the performance of detectors through an effect known as pile-up. Pile-up refers to the presence of multiple proton-proton interactions occurring in the same or nearby bunch crossings. Inelastic proton-proton collisions occurring in the same bunch crossing as the hard-scattering process of interest produce additional particles that complicate the event and can obscure the process of interest. Multiple interactions occurring in the same bunch crossing is known as *in-time* pile-up, and accounts for the majority of pile-up. Additionally, interactions occurring in previous and subsequent bunch crossings can also contaminate the event due to the finite response time of the detector. This is known as *out-of-time* pile-up.

Pile-up is quantified by the mean number of interactions per bunch crossing. The number of interactions per bunch crossing follows a Poisson distribution with mean μ, where μ is given by:

$$\mu = \frac{L_b \sigma_{\text{inel}}}{f_{\text{rev}}}, \tag{3.3}$$

where L_b is the instantaneous luminosity per bunch and σ_{inel} is the inelastic proton-proton scattering cross section. The mean number of interactions per bunch crossing is shown in Figure 3.2 for different years during the second and third data-taking periods Run 2 and Run 3 of the LHC. The average pile-up $\langle\mu\rangle$ for each year is also shown, as well as the most probable value (μ_{MPV}) per year in the case of Run 3. During Run 2 the luminosity-weighted mean pile-up rose steadily, from $\langle\mu\rangle = 13.4$ in 2015 to 36.1 in 2018, giving an average of 33.7 over the full period. Run 3 conditions are harsher, with an average of $\langle\mu\rangle = 56$ and a most probable value of $\mu_{\text{MPV}} = 63$ across 2022–2025, almost double that of Run 2.

3.1.2 LHC Run Periods

The LHC has been operational since 2009 and has alternated between distinct periods of data-taking and long shutdowns. The long shutdowns are used for

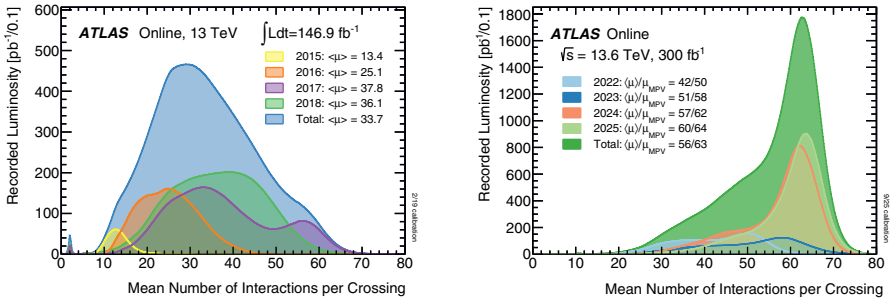

(a) (b)

Figure 3.2: Number of interactions per bunch crossing for different years during Run 2 (a) and Run 3 (b) of the LHC, weighted by luminosity. The average pile-up $\langle\mu\rangle$ for each year is also shown, as well as the most probable value (μ_{MPV}) per year in the case of Run 3. Taken from [42] and [43] for Run 2 and Run 3 respectively.

maintenance and upgrades to the LHC and its experiments. The first data-taking period, Run 1, took place between 2010 and 2012. During this time, the LHC operated at a centre-of-mass energy of $\sqrt{s} = 7$ TeV in 2010 and 2011 and $\sqrt{s} = 8$ TeV in 2012. During these periods, ATLAS recorded $5\,\text{fb}^{-1}$ and $21\,\text{fb}^{-1}$ of data respectively.

Long Shutdown 1 (LS1) took place between 2013 and 2015, during which the LHC underwent significant upgrades to increase the centre-of-mass energy and reach its design luminosity. The experiments also underwent upgrades to accommodate the increased luminosity and improve their performance.

Run 2 took place between 2015 and 2018, with the LHC operating at a centre-of-mass energy of $\sqrt{s} = 13$ TeV and reaching a peak instantaneous luminosity of $\text{L} = 2.1 \times 10^{34}\ \text{cm}^{-2}\text{s}^{-1}$ [42], approximately twice the design value. During this time ATLAS recorded $147\,\text{fb}^{-1}$ of data, as shown in Figure 3.3, of which $140\,\text{fb}^{-1}$ passed the data-quality requirements and is available for physics analysis.

Long Shutdown 2 (LS2) followed between 2019 and 2022. During this period, the LHC underwent further upgrades to increase the centre-of-mass en-

ergy to $\sqrt{s} = 13.6\,\text{TeV}$. The experiments also underwent further upgrades to their detectors and in preparation for the higher luminosity and pile-up conditions of Run 3 and beyond. Run 3 took place from 2022 to mid-2026, with the LHC operating at a centre-of-mass energy of $\sqrt{s} = 13.6\,\text{TeV}$. ATLAS recorded $332\,\text{fb}^{-1}$ of data during Run 3 [43], as shown in Figure 3.3.

Long Shutdown 3 (LS3) is planned for 2026-2030, during which the LHC will undergo major upgrades in preparation for the High-Luminosity LHC (HL-LHC) period from 2030 onwards. The HL-LHC will operate at a centre-of-mass energy of up to $\sqrt{s} = 14\,\text{TeV}$ and an instantaneous luminosity of $\text{L} = 5 \times 10^{34}\,\text{cm}^{-2}\text{s}^{-1}$, expected to produce a dataset of approximately $3000\,\text{fb}^{-1}$ in total. All experiments will undergo significant upgrades to their detectors to be able to manage the increased luminosity and pile-up conditions.

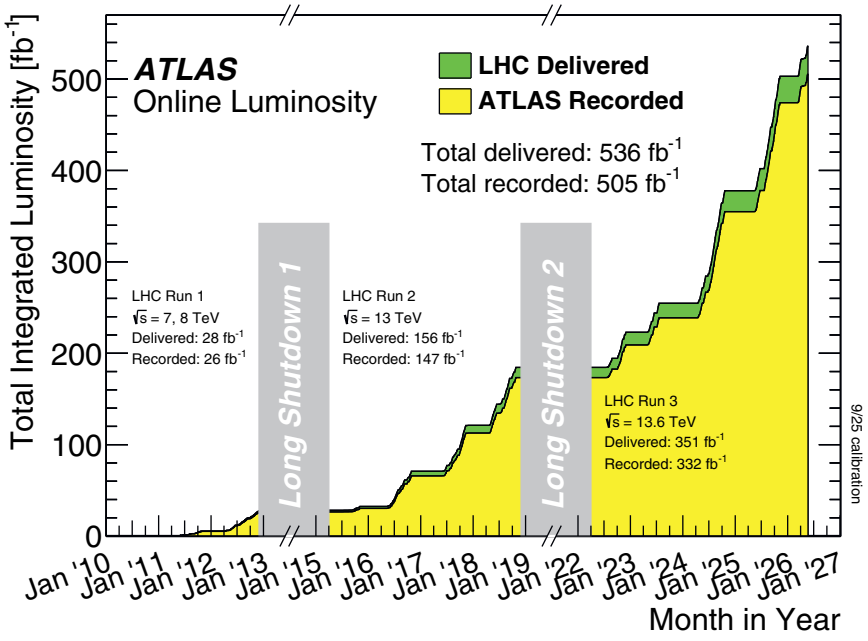

Figure 3.3: The integrated luminosity delivered by the LHC and recorded by ATLAS during each data-taking period of the LHC. Taken from [43].

3.2 The ATLAS Detector

The ATLAS detector is a general-purpose particle detector designed to study a wide range of both SM and BSM physics processes at the LHC. It is a cylindrical detector with a length of 44 m and a diameter of 25 m, weighing approximately 7000 tonnes, making it the largest particle detector in the world. The detector is composed of several subsystems, each designed to perform specific functions in the detection and measurement of particles produced in proton-proton collisions. The subdetectors are layered concentrically cylindrically outwards, with the IP at their centre.

The main subdetector systems consist of the inner detector (ID), the calorimeter system and the muon spectrometer (MS). A schematic representation of the ATLAS detector in its Run 3 configuration is shown in Figure 3.4. The different subsystems and detector functions are described in the following sections, with particular emphasis placed on the subsystems that are relevant to the reconstruction of physics objects in the $HH \to b\bar{b}\gamma\gamma$ analysis, namely photons and jets.

Object reconstruction is described fully in Chapter 4. Since this analysis uses both Run 2 and Run 3 data, the configuration of the detector over both data-taking periods is important. Upgrades applied to detector subsystems between Run 2 and Run 3 are highlighted where relevant.

3.2.1 The ATLAS Coordinate System

Physics variables within the ATLAS detector are described using a right-handed coordinate system with its origin at the nominal interaction point (IP) in the centre of the detector. The z-axis is defined along the beam pipe, the x-axis points from the IP to the centre of the LHC ring and the y-axis points upwards.

Cylindrical coordinates (r, ϕ) are used in the transverse (xy) plane, where r is the radial distance from the beam line and ϕ is the azimuthal angle around the beam pipe starting from the x-axis. The polar angle θ is defined as the angle in the yz plane. The rapidity y of a particle is commonly used instead of θ, defined as

$$y = \frac{1}{2} \ln \frac{E + p_z}{E - p_z}, \tag{3.4}$$

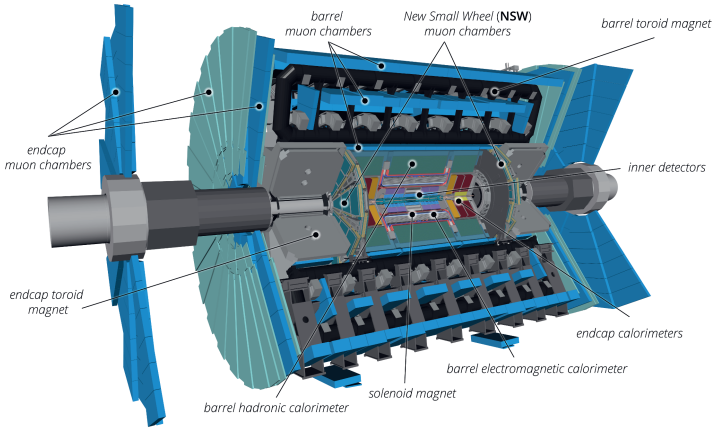

Figure 3.4: Cut-away view of the ATLAS detector in its Run 3 configuration and the locations of the subsystems [44].

where E is the energy of the particle and p_z is the component of its momentum along the beam axis. In the limit of massless particles, the rapidity is reduced to the pseudorapidity η, defined as

$$\eta = -\ln \tan \frac{\theta}{2}. \tag{3.5}$$

The geometric distance between two particles is expressed as ΔR, defined as

$$\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2}. \tag{3.6}$$

Because the particles produced in collisions have a centre of mass that is longitudinally boosted, quantities that are invariant under such boosts are preferred. Transverse momentum p_T, the component of momentum in the transverse plane, is invariant under longitudinal boosts. The transverse momentum is defined as

$$p_T = \sqrt{p_x^2 + p_y^2}, \tag{3.7}$$

where p_x and p_y are the x and y components of the momentum respectively. Similarly, the transverse energy E_T, representing the energy in the transverse plane, is often used. The ATLAS coordinate system on the LHC can be seen in Figure 3.5a and the distribution of η in the yz plane is shown in Figure 3.5b.

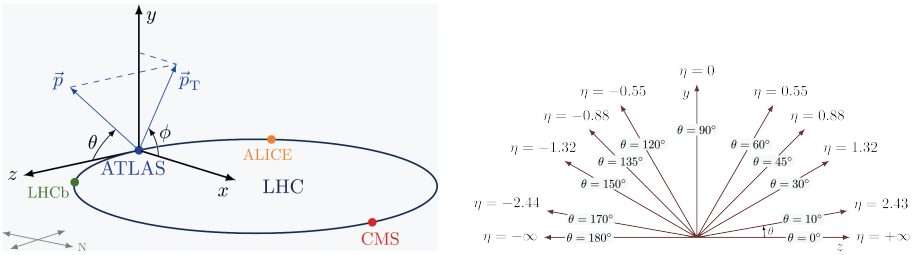

(a) (b)

Figure 3.5: The ATLAS coordinate system with the LHC and other experiments (a) and the distribution of pseudorapidity η in the yz plane (b) [45].

3.2.2 The Inner Detector

The inner detector is the innermost subsystem of the ATLAS detector, located close to the beam pipe and enclosing the IP, occupying a cylindrical volume of radius 1.15 m and length 7 m [36]. It is enveloped by a uniform axial magnetic field of 2 T generated by a central solenoid magnet of length 5.3 m and diameter 2.5 m, wound from an aluminium-stabilised NbTi superconductor chosen to minimise the material in front of the electromagnetic calorimeter [36]. The primary function of the ID is to measure the trajectories and momenta of charged particles produced in proton-proton collisions, which it is able to do accurately thanks to the magnetic field. The curvature of charged particles' trajectories is characterised by high-resolution position measurements known as hits, which are used to reconstruct tracks and to determine the positions of primary and secondary vertices with high spatial resolution. The ID is particularly important in the reconstruction and identification of jets originating from b-quarks, known as b-jets, as well as for the reconstruction of photon conversions occurring in the detector material, both of which are important for the $HH \to b\bar{b}\gamma\gamma$ analysis discussed in Chapter 7. A detailed description of the reconstruction of these objects can be found in Chapter 4.

The ID consists of a barrel region, covering a pseudorapidity range of $|\eta| < 1.2$, and two end-cap regions extending the coverage to $|\eta| < 2.5$. It is composed of three main subdetectors: the pixel detector, the Semiconductor Tracker (SCT) and the Transition Radiation Tracker (TRT). A schematic representation of the ID and its subdetectors is shown in Figure 3.6.

The Pixel Detector

The Pixel Detector is the innermost subdetector of the ID, located closest to the beam pipe. It is composed of four layers of modules. The first layer, closest to the beam pipe, is the Insertible B-layer (IBL), installed before Run 2. Positioned at a distance of $r = 33\,\text{mm}$ from the centre of the beampipe, the IBL consists of planar pixel sensors with a pixel size of $50 \times 250\,\mu\text{m}^2$ in the $r - \phi$ and z direction. The IBL has the best resolution of the four layers. The subsequent layers are composed of identical pixel sensors with a pixel size of $50 \times 400\,\mu\text{m}^2$.

The Semiconductor Tracker (SCT)

The SCT is a silicon microchip tracker composed of eight strip layers providing four space points per track. The layers are arranged cylindrically with a relative angle of $40\,\text{mrad}$ between strips in order to provide two-dimensional position measurements. The SCT covers a pseudorapidity range of $|\eta| < 2.5$ and provides high-precision position measurements with a resolution of $17\,\mu\text{m}$ in the $r - \phi$ direction and $580\,\mu\text{m}$ in the z direction.

The Transition Radiation Tracker (TRT)

The TRT is the outermost subdetector of the ID, consisting of layers of thin-walled drift tubes (straws) filled with a gas mixture. The straws provide both tracking information and electron identification through the detection of transition radiation. The TRT covers a pseudorapidity range of $|\eta| < 2.0$ and provides position measurements with a resolution of $130\,\mu\text{m}$ in the $r - \phi$ direction.

3.2.3 Calorimeter System

The calorimeter system is located outside the ID central solenoid and is enclosed by the muon spectrometer, extending in the barrel to an outer radius of $4.25\,\text{m}$ [36]. Its primary function is to measure the energies and positions of electromagnetically or strongly interacting particles. It is therefore divided into electromagnetic calorimeter (ECAL) and hadronic calorimeter (HCAL) systems. Energy measurements are obtained by absorbing and stopping incoming particles such that their energy is deposited in the detector.

The calorimeter system provides a nearly hermetic coverage of $|\eta| < 4.9$ and is composed of several layers of different materials designed to optimise the

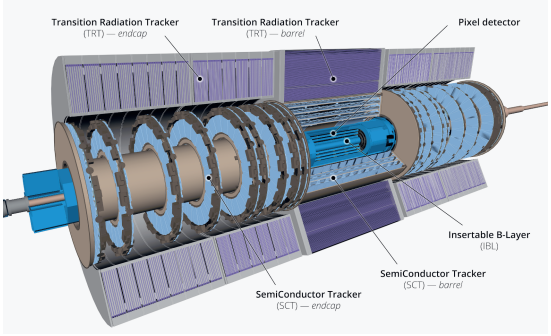

Figure 3.6: Cut-away view of the ATLAS ID and its subdetectors [44].

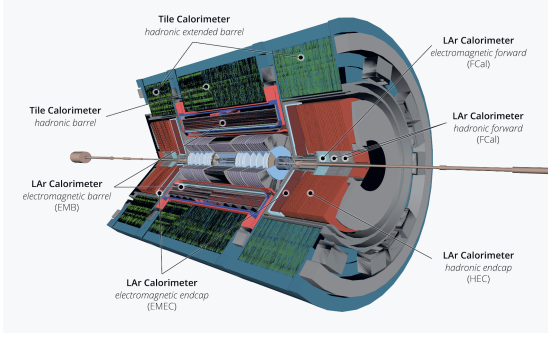

Figure 3.7: Cut-away view of the ATLAS calorimeter system and its subdetectors [44].

energy measurement of different types of particles. The ATLAS calorimeters are sampling calorimeters, which means that high-density absorbing materials are layered with active sections that measure particle energies. The calorimeter system is made up of two different sampling calorimeter technologies: the Liquid Argon (LAr) calorimeter and the Tile calorimeter (TileCal). The energies of electromagnetic particles, both neutral and charged, are measured by the LAr calorimeter. Hadronic showers are measured by both the LAr and the TileCal. These calorimeters are described further in the following. A schematic representation of the calorimeter system is shown in Figure 3.7.

Liquid Argon Calorimeter

The LAr Calorimeter system consists of several subsystems, namely the electromagnetic barrel (EMB) and endcap (EMEC) calorimeters, a hadronic endcap calorimeter (HEC) and a forward calorimeter (FCal). A thin presampler detector is placed in front of the electromagnetic calorimeter in regions of significant upstream material in order to correct for energy losses before the calorimeter.

The EMB, EMEC and part of the FCal make up the ECAL system, designed to measure the energies of electrons and photons. Muons are not absorbed by the ECAL system and are instead measured by the MS system described in the following section. The EMB and EMEC cover a pseudorapidity region of $\eta < 3.2$. They are composed of lead absorber plates interleaved with liquid argon active layers in an accordion geometry with full azimuthal symmetry. Electrons and photons that enter the calorimeter produce electromagnetic showers in the lead plates which then cause ionisation of the liquid argon. Electrodes capture the charges produced and measure the corresponding energy of the shower.

To accurately measure the energy of an electromagnetic shower, the full shower must be contained within the calorimeter. The EMB is between 22 and 30 radiation lengths (X_0) thick, while the EMEC is between $26\ X_0$ and $36\ X_0$ thick, which is sufficient to contain most electromagnetic showers.

The hadronic subsystems of the LAr calorimeter are the HEC and the FCal. The HEC is located directly behind the EMEC and is designed to measure hadronic showers in the endcap region. It consists of alternating layers of copper absorber plates and liquid argon active layers. Each endcap contains two HEC modules covering the region $1.5 < |\eta| < 3.2$ and overlapping slightly with the TileCal.

The FCal is located in the forward region of the detector, covering the pseudorapidity range $3.1 < |\eta| < 4.9$. Its role is to extend the calorimetric coverage as close to the beam line as possible [36]. It measures the energy of both electromagnetic and hadronic showers and consists of three modules. The first module is optimised for electromagnetic showers and is composed of copper absorber plates and liquid argon active layers. The second and third modules are optimised for hadronic showers and are composed of tungsten absorber plates and liquid argon active layers. The tungsten absorbers provide a high

density, which is necessary to minimise the lateral spread of hadronic showers.

Tile Calorimeter

The TileCal is a hadronic calorimeter located directly outside the LAr calorimeter, covering the pseudorapidity range $|\eta| < 1.7$ and extending radially from an inner radius of 2.28 m to the outer edge of the calorimeter system. It is composed of steel absorber plates interleaved with plastic scintillator tiles as active material. When hadrons enter the TileCal, they produce hadronic showers in the steel absorbers, which in turn cause the scintillator tiles to emit UV light. This light is then collected and shifted into the visible range and measured by photomultiplier tubes (PMTs) to determine the energy of the shower. The TileCal is approximately 10 nuclear interaction lengths thick, which is sufficient to contain most hadronic showers.

3.2.4 Muon Spectrometer

The MS is the outermost subsystem of the ATLAS detector, located outside the calorimeter system and designed to measure the trajectories and momenta of muons. It is enclosed by a large air-core toroidal magnet system, comprising a barrel toroid and two end-cap toroids, each built from eight superconducting coils arranged with eight-fold azimuthal symmetry about the beam axis [36]. The end-cap toroids are rotated by $22.5°$ with respect to the barrel in order to provide radial overlap and to optimise the bending power in the transition region between them. The coils are wound from an aluminium-stabilised NbTi/Cu conductor operating at 20.5 kA and 4.6 K, and the field is oriented azimuthally, bending muon trajectories in the $r - z$ plane with a bending power of 1.5 to 5.5 Tm in the barrel and approximately 1 to 7.5 Tm in the end-caps. The field strength range in each region is $0.2 - 2.5$ T and $0.2 - 3.5$ T respectively.

Muons are minimum ionising particles that are not stopped by the calorimeters and lose very little energy traversing them. Additionally, due to their mean lifetime of $\tau \approx 2.2\,\mu$s, the majority of muons travel completely through the ATLAS detector before decaying. This means that a signal in the MS is very unlikely to originate from any other SM particle. However, muons with a p_T lower than 3.0 GeV are unable to reach the MS due to energy loss in the detector materials.

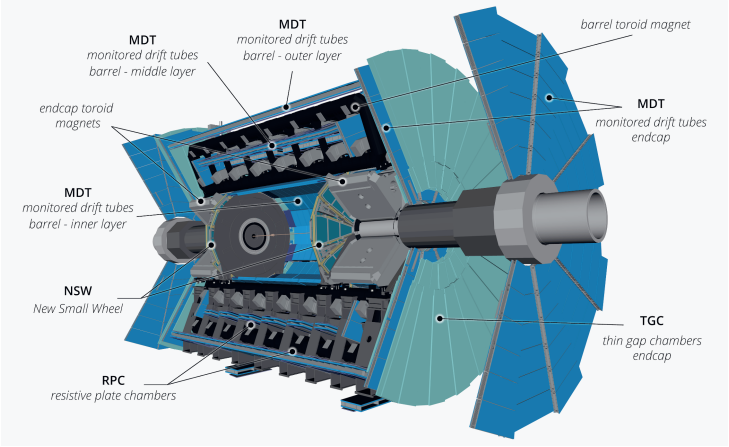

Figure 3.8: Cut-away view of the ATLAS Muon Spectrometer in Run 3 [44].

The MS provides a coverage of $|\eta| < 2.7$ and consists of a barrel and an endcap region with different detection technologies optimised for their respective η ranges. In the barrel, the chambers are arranged in three cylindrical layers at radii of approximately 5 m, 7.5 m and 10 m, while in the endcaps they form wheels perpendicular to the beam axis at $|z| \approx 7.4, 10.8, 14$ and 21.5 m [36]. The MS therefore defines the overall dimensions of the ATLAS detector. In the barrel region, the MS comprises multilayered Monitored Drift Tubes (MDTs) for precision tracking and Resistive Plate Chambers (RPCs) for triggering. The MDTs consist of layers of gas-filled aluminium drift tubes that provide high spatial resolution in the bending plane. The RPCs provide fast timing information with a resolution of 15 - 25 ns for triggering and contribute to the measurement of the azimuthal coordinate of muon tracks.

The endcap region underwent a significant upgrade between Run 2 and Run 3. In the Run 2 configuration, each endcap contained a Small Wheel covering $2.0 < |\eta| < 2.7$ and a Big Wheel extending the coverage to $1.0 < |\eta| < 2.7$. The Small Wheel is composed of MDTs and Cathode Strip Chambers (CSCs) for precision tracking, along with Thin Gap Chambers (TGCs) for triggering. For Run 3, the Small Wheel was replaced by the New Small Wheel (NSW) covering the same η range. The NSW features Micromegas (MM) detectors and small-strip TGCs (sTGCs), both of which provide precision tracking and trig-

gering capabilities. This upgrade delivers improved spatial and angular resolution, enhanced rate capability, and better timing resolution, which are essential for maintaining trigger selectivity and tracking performance in the high-luminosity conditions of Run 3. The Big Wheel, containing MDTs, CSCs, and TGCs, remains unchanged. The MS in its Run 3 configuration is shown in Figure 3.8.

3.2.5 Trigger and Data Acquisition System

Collisions at the LHC occur at a rate of $40\,\text{MHz}$, which is far too high to be stored on disk in its entirety. While the vast majority of events produced in proton-proton collisions are soft inelastic collisions, events of interest occur in hard-scattering processes. The ATLAS Trigger and Data Acquisition (TDAQ) system [46] is designed to select interesting events for recording while rejecting the large majority of uninteresting events. The full chain of the TDAQ system reduces the readout rate from $40\,\text{MHz}$ to approximately $1 - 3\,\text{kHz}$, which can then be stored for offline analysis. A diagram of the full Run 3 TDAQ system is shown in Figure 3.9.

The TDAQ system consists of a hardware-based Level-1 (L1) trigger and a software-based High-Level Trigger (HLT). The L1 trigger makes the first decision to record or discard an event based on reduced-granularity information from the subdetectors. It consists of two independent trigger systems: L1Calo, which receives information from the calorimeter system, and L1Muon, which receives information from the MS. The L1 topological processor (L1Topo) applies topological selections based on objects reconstructed by the two systems.

The L1 trigger performs selections that are coarse and computationally inexpensive relative to the HLT trigger. Events are selected based on event-level quantities (such as the total energy in the calorimeter), multiplicities of objects above thresholds (such as the number of muon tracks above a certain p_T) or topological quantities (such as invariant masses). Additionally, the L1 trigger identifies interesting features in events and records the region in η and ϕ in which they occur, known as the Region of Interest (RoI), which is then passed to the HLT in the event of a positive L1 trigger decision. The L1 trigger performs the initial rate reduction from $40\,\text{MHz}$ to $100\,\text{kHz}$ with a latency of $2.5\,\mu\text{s}$.

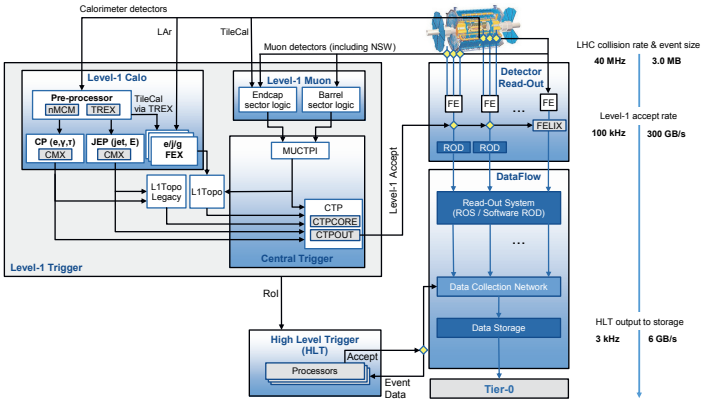

Figure 3.9: Schematic representation of the ATLAS TDAQ system in Run 3 [47].

For each event accepted by the L1 trigger, data is read out by front-end electronics in all detector systems and buffered by the ReadOut System (ROS). The ROS then sends the data to the second stage of the trigger, the HLT. The HLT makes use of finer-granularity reconstruction algorithms to make a final decision. Feature-extraction algorithms are applied within the RoI, making use of more precise calorimeter data and tracking information from the ID and MS. Processing is performed by over 100000 CPU cores comprising the HLT computing farm. Data recorded by the HLT is comprised of different physics trigger groups each subject to specific criteria, making up the HLT trigger menu. Each group is designed to select a different physics signature, such as photons or b-hadrons. Events may contain more than one of these signatures, therefore the groups overlap. As an example, the trigger group rates from a specific period in Run 2 are shown in Figure 3.10.

The HLT software was upgraded for Run 3 to make use of the full detector for the reconstruction of hadronic tracks. This improves the performance, but also increases the time to process events. To accommodate this upgrade

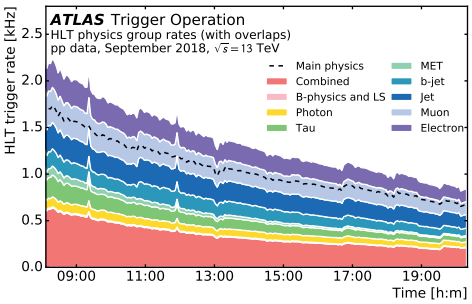

Figure 3.10: HLT physics trigger group rates as a function of time taken in September 2018 with a peak luminosity of $L = 2.0 \times 10^{34}$ cm^{-2}s^{-1} and a peak average number of interactions per crossing of $\langle \mu \rangle = 56$. Each group includes single object and multi-object triggers. Image taken from [48].

and keep resource usage within the allocated range, the HLT farm and reconstruction algorithms underwent upgrades. The HLT reduces the total readout rate to 3 kHz (increased from 1 kHz in Run 2) of data to be stored and used for offline reconstruction. Accepted events are compressed from their original size of approximately 3 MB to below 2 MB before being written to permanent storage, corresponding to an output bandwidth of approximately 6 GB/s [47]. The algorithms used to reconstruct physics objects from the stored data are described in the following chapter.

3.3 Upgrades for the High-Luminosity LHC

As mentioned in Section 3.1.2, the LHC will undergo its Phase II upgrade in anticipation of the High-Luminosity LHC, increasing its instantaneous luminosity by a factor of more than five and collecting a dataset on the order of ten times larger than the entire LHC data-taking period [49]. The dramatic increase in collision rate and total radiation dose poses fundamental challenges to essentially every subsystem of the ATLAS detector. In anticipation of this, the ATLAS detector itself is undergoing a comprehensive upgrade program to many of its subsystems. Upgrades include changes to the TDAQ system, a completely new inner tracker, a new silicon timing detector and extensive modifications to the MS. This section describes two of the major upgrades and my

contribution to them.

Inner Tracker Upgrade

The most significant change to the tracking system is the full replacement of the current ID with the ATLAS Inner TracKer (ITk). The entire tracking system described in Section 3.2.2 will be replaced by a new all-silicon pixel and strip tracker. The motivation for this complete overhaul is driven by the fact that the radiation levels, luminosity, and pile-up conditions at the HL-LHC will render the current ID inoperable. The ITk must tolerate the expected dose of radiation while maintaining efficient tracking.

The ITk consists of an inner pixel detector surrounded by a silicon strip detector. The strip system consists of four barrel layers and two endcaps with six disks each. In total, the ITk will house over 27000 modules made up of sensors and read-out electronics and contain more than 5 billion read-out channels [49].

The integration of the strip system is carried out at three independent sites: the Barrel section at CERN and the two endcap sections at DESY and Nikhef. Upon their completion, each endcap will be transported to CERN. The basic mechanical support structures of the endcaps are assembled at Nikhef. They require a two-phase CO_2 cooling system to maintain the silicon sensors at approximately -35 °C throughout the HL-LHC lifetime, protecting them from the effects of radiation damage [50]. The low temperature is required because radiation damage increases the leakage current of the silicon sensors, which in turn dissipates additional heat; a low temperature suppresses the leakage current and provides a margin against thermal runaway [51].

Before the pipes that distribute coolant (cooling manifolds) can proceed to pressure and leak testing, they must be cleaned to remove manufacturing residues and impurities that could obstruct the narrow cooling channels or compromise the leak-tightness of the circuit. I contributed to this quality control process by flushing the manifolds with isopropyl alcohol to remove contaminants, ensuring the integrity of the cooling circuits prior to integration into the endcap structures.

Muon Spectrometer Upgrade

The muon spectrometer also requires significant upgrades for the high luminosity phase. Along with the addition of the New Small Wheel mentioned in Section 3.2.4, a large fraction of the front-end electronics for the RPC, TGC, and MDT chambers will be replaced to make them compatible with the higher trigger rates and longer latencies necessary for the new level-0 trigger [52]. A particular concern for the MDT system is the Detector Control System (DCS), which is responsible for monitoring and controlling the operational parameters of the drift tubes, including gas pressure, temperatures, and high voltages. The DCS communicates with the MDT chambers via Embedded Local Monitor Board (ELMB) modules mounted directly on the chambers.

The ELMB2 is a general purpose plug-on I/O module for the monitoring and control of sub-detector front-end components, and serves as a replacement for the obsolete ELMB, tolerating higher levels of radiation than its predecessor while remaining fully compatible with the existing infrastructure. The original ELMB modules were designed for the radiation environment of the original LHC programme and were not designed to withstand the radiation levels expected during HL-LHC operation. As a result, all ELMB modules on the muon spectrometer must be replaced with ELMB2 modules before the start of Run 4.

Approximately 500 ELMB modules are accessible during the year-end technical stops (YETS) of Run 3 and represent an opportunity to progress with the replacement before LS3. I assisted with carrying out this replacement campaign on-site, physically accessing the MDT chambers and swapping the original ELMB modules for their ELMB2 successors. Figure 3.11a shows an example of a replacement on an MDT chamber.

Replacing as many modules as possible during the YETS reduces the volume of work required during LS3, when competing demands from the many simultaneous upgrades across all ATLAS subsystems place significant constraints on access time and personpower in the cavern. These efforts contribute to the readiness of the next generation of the ATLAS detector, ensuring that it can continue to perform optimally in the challenging conditions of the HL-LHC and maintain its ability to explore new frontiers in particle physics.



(a) (b)

Figure 3.11: Upgrade work performed over the course of the thesis: replacing an ELMB2 module on an MDT chamber in the ATLAS MS during the YETS of 2025 (a) and flushing the ITk cooling manifolds (b).

Object Reconstruction 4

The previous chapter described how the various subsystems of the ATLAS detector are optimised to identify and measure properties of the different particles produced in proton-proton collisions. This chapter describes the algorithms that reconstruct so-called *physics objects* from the raw detector information so that they can be used in physics analyses.

Figure 4.1 shows a cross section of the ATLAS detector and its subsystems with the signatures of different particles. Charged particles follow a helical trajectory through the inner detector (ID), while neutral particles travel in a straight line. Muons travel all the way to the muon spectrometer (MS). Photons and electrons deposit their energy in the electromagnetic calorimeter (ECAL). Hadrons such as protons and neutrons deposit their energy in the hadronic calorimeter (HCAL). Neutrinos do not interact with any of the detector subsystems and their energy escapes the detector. This energy is reconstructed as missing transverse energy (E^{T}_{miss}). In addition to the particles shown in Figure 4.1, collimated sprays of particles are produced in the hadronisation of quarks and gluons from the hard scattering events as discussed briefly in Chapter 2: these are known as jets and they are also reconstructed as distinct physics objects.

The reconstruction algorithms described in this chapter output a set of physics

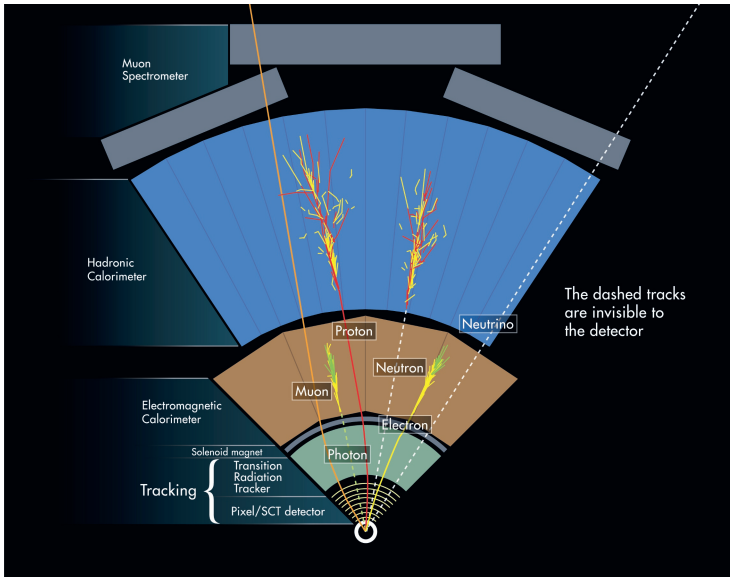

Figure 4.1: A cross section slice of the ATLAS detector showing the signatures left by different particles in the detector subsystems. The helical trajectories of charged particles are shown as curved in this 2D projection. Image taken from [53].

objects and their corresponding four-momenta. These objects are then calibrated to correct for detector effects. The central elements of object reconstruction in ATLAS are tracks, vertices and energy clusters. The reconstruction of these observables and how they are used to build physics objects are described in the following sections, with particular focus placed on the objects most relevant to the $HH \to b\bar{b}\gamma\gamma$ selection – namely, photons, jets and particularly b-jets, which are jets originating from a b-quark.

4.1 Tracks and Vertices

Reconstructed tracks form the basis of the reconstruction of physics objects in ATLAS. Track reconstruction is performed by a sequence of algorithms that use information from the inner detector. The first step is an *inside-out* algo-

rithm: track seeds are formed out of three points in space created from clusters of hits in the Pixel and SCT layers of the ID. Track candidates are built from the seeds by a combinatorial Kalman filter algorithm [54] that iteratively adds compatible hits moving away from the IP. Fake tracks and overlaps of space points between tracks are resolved by a dedicated ambiguity resolution algorithm. The tracks are then extended to the TRT.

Next, an *outside-in* algorithm starts with hits in the TRT and works inwards through the silicon layers towards the IP to account for tracks missed by the inside-out algorithm (often from secondary decays) in a process known as back-tracking. Surviving tracks then undergo a χ^2 fit and the parameters d_0, z_0, ψ, θ and q/p are assigned at a reference point. These represent the longitudinal and transverse distance from the track to the reference point, the azimuthal and polar angles of the momentum, and the (signed) electric charge over the momentum magnitude, respectively.

From the reconstructed tracks, it is possible to reconstruct the primary vertex (PV) at which the hard scatter event occurred, and other secondary vertices. An iterative vertex finding and fitting procedure [55] is used that takes the abovementioned quantities as input. Tracks compatible with the vertex position are assigned to the vertex and the vertex is in turn updated by weighting tracks according to their compatibility with it. Tracks associated with a given vertex are then refit under the assumption of originating from that vertex. Tracks that are incompatible with the vertex are used to reconstruct the next vertex. The procedure is repeated until no more compatible tracks are found. The PV is taken to be the vertex with the highest $\sum p_T^2$ of associated tracks out of all vertices containing at least two associated tracks with $p_T > 0.5\,\text{GeV}$.

4.2 Topological Clusters

The energy deposited by particles in the electromagnetic and hadronic calorimeters is clustered together into three-dimensional objects known as topological clusters (topo-clusters). The procedure, described fully in Ref. [56], begins with the identification of seed calorimeter cells and is based on the signal-to-noise ratio of the cells. This is defined as

$$\varsigma_{\text{cell}}^{\text{EM}} = \frac{E_{\text{cell}}^{\text{EM}}}{\sigma_{\text{noise,cell}}^{\text{EM}}}, \tag{4.1}$$

where $E_{\text{cell}}^{\text{EM}}$ is the cell signal and $\sigma_{\text{noise,cell}}^{\text{EM}}$ is the average expected noise for that cell calculated for each data-taking year. Seed cells with $|\varsigma_{\text{cell}}^{\text{EM}}| > 4$ are identified. Neighbouring cells with a $|\varsigma_{\text{cell}}^{\text{EM}}| > 2$ are iteratively added to the cluster until no more cells can be added. A final layer of cells with a $|\varsigma_{\text{cell}}^{\text{EM}}| > 0$ is added to the cluster. If more than one local maximum with an energy of above $500\,\text{MeV}$ is found, the cluster is split further using dedicated weighting algorithms.

Topo-clusters form the basis of the reconstruction of energetic particles and particularly of jets. In practice, single jets may deposit their energy in multiple topo-clusters or a single topo-cluster may contain the energy of multiple particles or jets. The reconstruction of jets is discussed further in Section 4.4.

4.3 Electrons and Photons

Photons are one of the two main objects reconstructed in the final state of the $HH \to b\bar{b}\gamma\gamma$ analysis described in Chapter 7. Both electrons and photons produce similar showers in the electromagnetic calorimeter, but electrons, due to their charge, are identified by one or more tracks in the inner detector as well as by ECAL clusters. Photons are identified primarily by their energy deposits in the ECAL, with associated tracks present only when the photon converts, as described below. Due to their importance to the $HH \to b\bar{b}\gamma\gamma$ analysis, photons are described in more detail.

Two significant effects complicate the reconstruction of electrons (or positrons) and photons: bremsstrahlung and photon conversion. Bremsstrahlung is the emission of photons by electrons as they interact with the material of the detector, causing the electrons to lose energy. Photon conversion is the process by which a photon converts into an electron-positron pair in the material of the detector. A photon is defined as converted if its energy deposit is matched to a conversion vertex in the ID. Two-track conversion vertices are reconstructed from two opposite-charge tracks forming a vertex consistent with that of a massless particle, while single-track vertices are formed from tracks with no hits in the innermost sensitive layers, where the two conversion tracks are not

separately resolved. Unconverted photons are matched to neither a conversion vertex nor an electron track.

Reconstruction from Superclusters

A dynamic, variable-sized cluster known as a supercluster is used to reconstruct electrons and photons and to correct for bremsstrahlung and photon conversion [57]. The supercluster consists of a topo-cluster, as the seed, and other nearby satellite clusters: seed candidates must have a minimum transverse energy E_T of 1 GeV for electrons and 1.5 GeV for photons. Electrons must also have a matched track containing at least four hits in the silicon layers of the ID.

Next, satellite clusters are identified within windows centred on the seed cluster barycentre. For both electrons and photons, a topo-cluster is considered a satellite if it falls within $\Delta\eta \times \Delta\phi = 0.075 \times 0.125$ of the seed, as such clusters typically arise from secondary electromagnetic showers initiated by the same particle. For electrons, a cluster in the larger window $\Delta\eta \times \Delta\phi = 0.125 \times 0.300$ is additionally taken as a satellite if its best-matched track is also the best-matched track of the seed cluster. For converted photons whose conversion vertices are built only from tracks with silicon hits, a cluster is added as a satellite if its best-matched track belongs to the conversion vertex matched to the seed. These track-based criteria serve to distinguish distant radiative photons and conversion electrons from pile-up and other unrelated clusters. Superclusters account for bremsstrahlung radiation by including satellite clusters around the electron cluster corresponding to the radiated photons.

After the superclusters are built, the measured energy of electron and photon candidate objects is calibrated using a multivariate regression algorithm. Photon and electron objects used in analysis are built from the superclusters. In some cases, the same seed cluster may produce both an electron and a photon object. In these cases several distinguishing criteria are checked – for example, a cluster with no good track attached is a photon – but in some cases, these do not bring a determinate result and the object is marked as explicitly ambiguous. The classification is then determined at the analysis level. For most physics analyses, reconstructed electron and photon candidates are required to satisfy *identification* and *isolation* criteria to improve the purity of

the selected objects.

Identification

Identification separates real electrons and photons from fakes. Fake photons and electrons have several sources, most often arising from hadrons misidentified as electrons or photons, or from hadron decays that produce electrons or photons. The dominant source of fake photons is $\pi^0 \to \gamma\gamma$ decays from a jet in which the two photons are highly collimated and are misidentified as a single photon. These decays can also mimic an electron if a charged hadron from the same jet provides an overlapping track. Photon conversions can also be misidentified as electrons since the electrons produced in the conversions leave similar calorimeter signatures.

Photon identification is performed by a cut-based selection that uses shower-shape variables from the calorimeters. Variables associated with the first layer of the EM calorimeter have particular discriminating power for prompt photons against $\pi^0 \to \gamma\gamma$ decays.

Three photon identification working points - Loose, Medium and Tight - are defined at increasing levels of background rejection. Loose and Medium are used for trigger selections, while Tight is the primary identification selection used in analysis. All working points are defined in separate bins of $|\eta|$ to account for the variation of shower-shape variables with the geometry of the calorimeter. The Tight working point is also optimised in bins of E_T and performed separately for converted and unconverted photons due to the differences in shower shapes.

Electron identification is based on a multivariate likelihood discriminant that takes as input shower shape variables, track quality variables and track-cluster matching variables. Loose, Medium and Tight electron identification working points are defined. The full procedure is described in Ref. [57].

Isolation

Isolation is used to separate prompt photons and electrons – those originating from the hard scatter event – from non-prompt photons and electrons that

typically originate from jets and have a large amount of surrounding activity. Prompt photons and electrons are expected to be well isolated from nearby particles. Two types of isolation are performed: calorimeter isolation and track-based isolation.

Calorimeter isolation sums the E_T of topo-clusters within a cone of size $\Delta R = 0.2$ around the candidate and subtracts the energy of a fixed 'core' around the candidate supercluster. The energy of the reconstructed electron or photon candidate is not perfectly contained within this core, so an additional correction is applied to correct for any energy leakage from the candidate into the cone. Corrections are also performed for pile-up and the underlying event.

Track-based isolation sums the p_T of tracks within a cone centred around the electron track or the direction of the photon cluster. Tracks matched to the electron or (converted) photon are subtracted. Several isolation working points are defined, based on selections on the calorimeter and track-based isolation variables whose thresholds scale with the candidate E_{T}. For photons these are Loose, Tight and TightCaloOnly, the last of which uses calorimeter isolation alone. A larger set of working points is available for electrons.

4.4 Jets

The first step in jet reconstruction is the Particle Flow (PFlow) algorithm [58], which uses information from both the ID and the calorimeters. For charged particles, selected tracks are matched to topo-clusters; to avoid double counting the energy from both the track and topo-cluster, the more precise momentum measurement from the track is used and the energy of the cluster is subtracted. Topo-clusters from neutral particles that leave no tracks are left unchanged. This collection of neutral clusters and charged particle tracks comprise the PFlow object.

Jets are built from PFlow objects using the anti-k$_t$ algorithm [59]. PFlow objects are clustered into jets based on a distance parameter

$$d_{ij} = \min\left(\frac{1}{p_{T,i}^2}, \frac{1}{p_{T,j}^2}\right)\frac{\Delta R_{ij}^2}{R^2}, \tag{4.2}$$

where $p_{T,i}$ and $p_{T,j}$ are the transverse momenta of objects i and j, ΔR_{ij} is the distance between the two objects in the $\eta - \phi$ plane and R is a fixed parameter that controls the size of the jet. The anti-k$_t$ algorithm iteratively merges pairs of PFlow objects based on a minimum value of d_{ij} and then recomputes the combined four-momenta until no more pairs can be merged.

Jets with a radius parameter of $R = 0.4$ are defined as small-radius (small-R) jets and jets with a radius parameter of $R = 1.0$ are defined as large-radius (large-R) jets. Large-R jets typically arise from highly boosted objects decaying into a pair of quarks that are too collimated for their resulting jets to be well reconstructed separately. The subsequent discussion applies only to Small-R jets, as these are the objects relevant to the $b\bar{b}\gamma\gamma$ final state; Large-R jets are discussed in the last section of this chapter.

After reconstruction, jets are calibrated; the jet calibration chain for Small-R jets is summarised in Figure 4.2. Pile-up effects are corrected for in two consecutive steps. First, a subtraction based on the estimated per-event pile-up contribution is performed: this is based on the median p_T density of a jet with area A in the $y - \phi$ plane. Residual pile-up not accounted for by the p_T density estimation (typically in the forward region of the calorimeter or in high occupancy regions of high-p_T jets) is then corrected for. This residual correction accounts for the difference between truth and reconstructed jet p_T dependent on pile-up sensitive variables.

After the initial pile-up corrections, the jet energy scale (JES) calibration [60] corrects the four-momenta of reconstructed jets to that of particle-level *truth-jets* built from stable particles in MC simulation. This procedure corrects for energy loss in the passive detector materials, out-of-cone effects and bias in the jet η reconstruction due to transitions between different calorimeter regions.

Residual dependencies of the JES on the flavour and energy distribution of constituent particles, their distribution within the shower and the initiating particle are corrected for by a global sequential calibration (GSC). The GSC uses global jet observables such as information from the trackers and muon chambers and the longitudinal shape of the calorimeter energy deposits, applied as a series of corrections designed to improve the jet resolution without altering the average jet energy response.

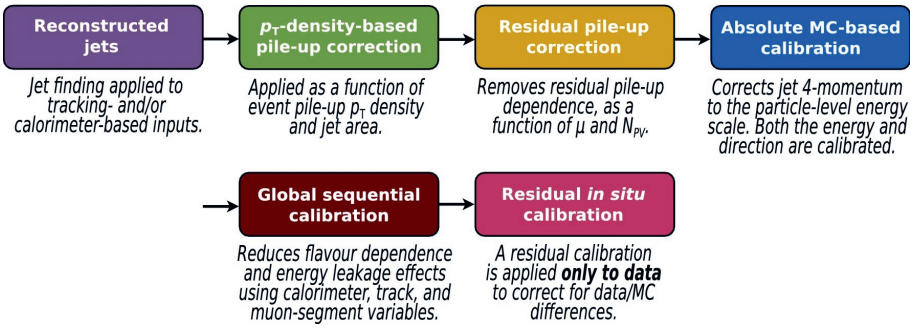

Figure 4.2: The jet energy scale calibration chain. Image taken from Ref. [60].

The final step in the calibration of jets accounts for the differences between data and simulation in the jet response due to imperfect modelling of the physics processes and detector effects. This 'in-situ' calibration is applied only to data and is performed using well-calibrated reference objects to estimate the ratio of the jet p_T to that of the reference object.

The jet vertex tagger (JVT) [61] is applied to further suppress pile-up jets beyond the initial pile-up corrections. The JVT is a multivariate algorithm that combines track-based variables with vertex-based variables to discriminate between jets originating from the hard scatter vertex and jets originating from pile-up. The JVT is only applicable to jets within the coverage of the tracker ($|\eta| < 2.5$). For jets outside of this region, a dedicated forward jet vertex tagger (fJVT) [62] is used, which matches forward jets to the missing transverse momentum computed from the tracks of each pile-up vertex.

4.4.1 b-jets

To perform physics analysis on final states with b-quarks, such as $HH \to b\bar{b}\gamma\gamma$, it is crucial to identify *b-jets* and distinguish them from jets originating from c-quarks, light quarks/gluons and τ leptons. Typically, b-jets contain a b-hadron that travels some distance away from the PV before further decaying. This results in a secondary vertex that is displaced from the primary vertex, and in tracks from the b-hadron decay with large transverse impact parameters d_0 relative to the primary vertex. An illustration of a typical b-jet is shown in Figure 4.3.

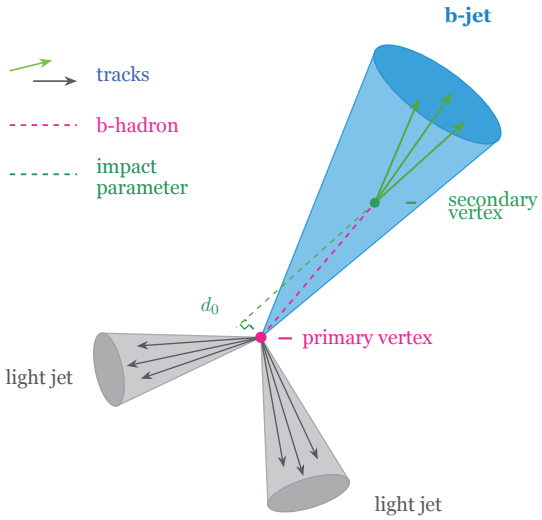

Figure 4.3: Schematic illustration of the signatures exploited for b-jet identification. A b-hadron produced at the primary vertex travels a distance (pink dashed line) before decaying at a displaced secondary vertex, producing a collimated jet of charged tracks. Extrapolating such a track back towards the beamline (light green) shows that it misses the primary vertex by the transverse impact parameter d_0, in contrast to the tracks of light-flavour jets, which originate directly from the primary vertex.

b-tagging algorithms use this characteristic feature and other variables to distinguish b-jets from other jets. The most recent and best performing b-tagging algorithm used in the ATLAS experiment at this time is the GN2 [63] algorithm, based on a transformer architecture [64]. Unlike previous b-tagging algorithms, GN2 does not rely on the higher level output of optimised intermediate algorithms, but instead learns directly on low-level track and jet information such as the number of hits in different layers of the inner detector, track-related variables and jet kinematics.

The discriminant of GN2 is defined as

$$D_{\text{GN2}} = \ln\left(\frac{p_b}{f_c p_c + f_\tau p_\tau + (1 - f_c - f_\tau) p_u}\right), \tag{4.3}$$

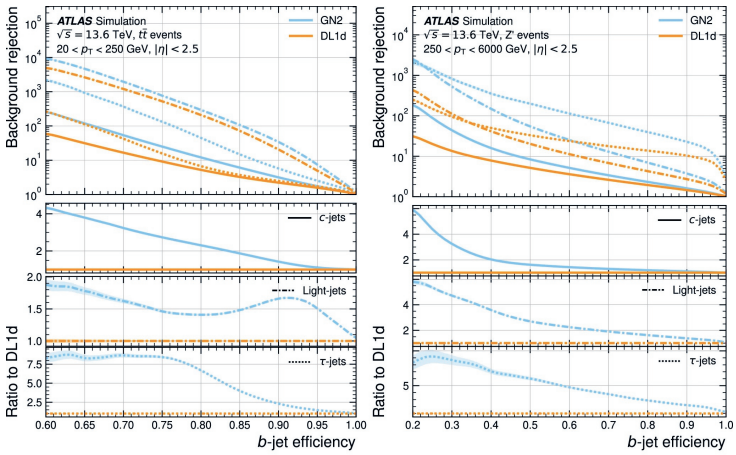

Figure 4.4: The background rejection of the GN2 b-tagging algorithm for b-jets against different backgrounds as a function of b-jet efficiency. Image taken from Ref. [63].

where p_b, p_c, p_τ and p_u are the output flavour prediction probabilities of a jet being a b-jet, c-jet, τ-jet or light jet respectively, and $f_{c(\tau)}$ is the relative weight between $p_{c(\tau)}$ and p_u optimised in simulated $t\bar{t}$ events. Several working points are defined based on the cut value of D_{GN2} for which true b-jets pass with efficiencies 60%, 70%, 77% and 85%. Physics analyses select the optimal working point for their particular needs based on the trade-off between signal efficiency and background rejection. Jets that pass the chosen working point are considered b-tagged jets. Jets are also assigned a pseudo-continuous b-tagging (PCBT) score based on a discretisation of the D_{GN2} discriminant to give more granular b-tagging information about each jet.

The background rejection of GN2 for b-jets against different backgrounds (c-jets, light-jets and τ-jets) as a function of b-jet efficiency is shown in Figure 4.4. Here, the predecessor to GN2, DL1d [65], is used as a reference from which to measure the performance. Figure 4.5 shows the improvement in background rejection over time for several flavour tagging algorithms in ATLAS for a 70% working point.

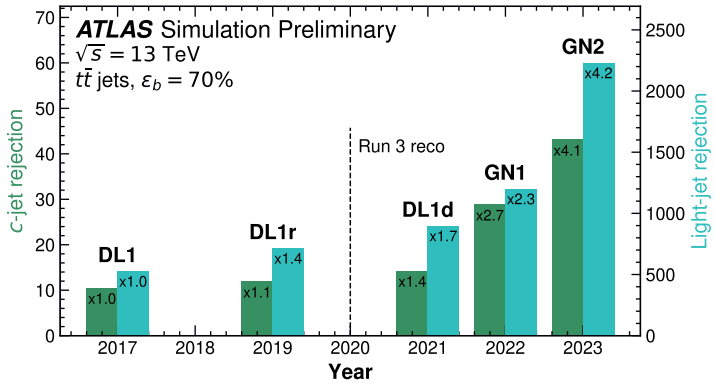

Figure 4.5: The c-jet and light jet rejection of different ATLAS b-tagging algorithms over time in simulated $t\bar{t}$ jets at 70% signal efficiency. Image taken from Ref. [66].

4.5 Muons

Muons are reconstructed using information from the muon spectrometer systems, the ID and the calorimeter [67]. Short, straight-line segments are created by hits in the muon chambers and combined into tracks. Track candidates are required to meet certain selection criteria based on the χ^2 of the track fit. Muons are classified into four types based on the information used in their reconstruction:

Combined: MS tracks and ID tracks are fitted to a single combined track. This category has the highest muon efficiency.
Segment-tagged: ID tracks are extrapolated to a segment in the MS. This category recovers low-p_T muons or those crossing only a single layer of the MS.
Calorimeter-tagged: ID tracks are matched to calorimeter deposits associated with the signature of a minimum ionising particle. This category recovers muons in the region around $|\eta| \approx 0$ and has the lowest efficiency.
Extrapolated: Muons are reconstructed from MS hits only, with a track loosely matched to the PV. This category recovers muons outside the ID coverage.

Similarly to photons and electrons, muons are required to meet identifica-

tion criteria to separate prompt muons from muons originating from hadron decays. These requirements are based on the quality of the muon track fit. Five working points are defined: the Tight, Medium, Loose working points correspond to increasing signal efficiency and decreasing background rejection; the high-p_T and low-p_T working points enhance the identification of muon tracks above 100 GeV and in the range $3 - 10$ GeV respectively.

4.6 Missing Transverse Energy

Missing transverse energy ($\mathrm{E}_T^{\mathrm{miss}}$) functions as a proxy object representing the transverse momentum of undetected particles such as neutrinos. The reconstruction of $\mathrm{E}_T^{\mathrm{miss}}$ is based on the conservation of transverse momentum and involves two components: signatures from the hard scattering event involving fully reconstructed objects and particles (e.g. photons, jets and muons), and *soft-event* signals consisting of reconstructed charged particle tracks originating from the PV but not associated with any hard object [68]. $\mathrm{E}_T^{\mathrm{miss}}$ is reconstructed as the negative vector sum of the transverse momenta of all reconstructed hard-scatter objects and soft-event signals, with the undetected particles accounting for the imbalance.

4.7 Large-R Jets

Very high p_T (boosted) particles produce collimated decay products that are reconstructed as a single Large-R jet. Large-R jets are reconstructed using the anti-k_t algorithm with radius parameter $R = 1.0$. The $HH \to 4b$ boosted channel [69] selects events in which both $H \to b\bar{b}$ decays are reconstructed with a Large-R jet, contributing significantly to the sensitivity to κ_{2V}. While Large-R jets are not used in the $HH \to b\bar{b}\gamma\gamma$ analysis described in Chapter 7, the exploration of the region in which the $H \to b\bar{b}$ decay is boosted is a potential topic of exploration for future searches.

The GN2X tagger [70] follows a similar transformer architecture and training as GN2, but is optimised for distinguishing both boosted $H \to b\bar{b}$ and $H \to cc$ jets from background processes. Large-R jets are classified based on their origin, leading to four classes: $H \to b\bar{b}$, $H \to cc$, top (arising from fully hadronic top-quark decays) and QCD (or multijet processes).

4.7.1 Calibration of the GN2X Tagger

Since GN2X is fully trained on simulated jets, it is necessary to calibrate the tagger to account for differences in the algorithm's performance on simulation and data. Both signal and background efficiency corrections are derived in the form of scale factors:

$$SF = \frac{\epsilon_{\text{data}}}{\epsilon_{\text{MC}}} = \frac{\frac{N^{\text{data}}_{\text{passed}}}{N^{\text{data}}_{\text{total}}}}{\frac{N^{\text{MC}}_{\text{passed}}}{N^{\text{MC}}_{\text{total}}}} = \frac{\frac{N^{\text{data}}_{\text{passed}}}{N^{\text{MC}}_{\text{passed}}}}{\frac{N^{\text{data}}_{\text{total}}}{N^{\text{MC}}_{\text{total}}}} = \frac{\mu_{\text{post-tag}}}{\mu_{\text{pre-tag}}}. \tag{4.4}$$

My contribution covers the first version of the $g \to bb$ background calibration. The $g \to bb$ multijet process, in which a gluon splits into a pair of b-quarks and hadronises into jets, is kinematically similar to boosted $H \to bb$ and a significant source of QCD background. The calibration is performed on Run 2 simulation samples of di-jet events enriched in $g \to bb$ by requiring the presence of a muon from the semileptonic decay of one of the b-hadrons and Run 2 data. At least one Large-R jet is required in each event, categorised into six flavour templates (BB, BC, BL, CC, CL, LL) based on the truth flavour labels of constituent variable-radius (VR) track jets that are ghost-associated to the Large-R jet. At least two VR track-jets are required.

A fit is performed over the six flavour templates to extract the normalisation factors and subsequent scale factors in Equation 4.4, which are then applied to tagged $g \to bb$ MC events. The chosen fitting variable is the mean secondary vertex (SV) mass of all VR subjets ghost-associated with the Large-R jet; this is a powerful discriminant because boosted $H \to bb$ tend to have a higher SV mass due to the decay of the b-hadron inside the jet. The mean SV mass for Large-R jets in the six flavour templates is shown in Figure 4.6.

Large-R jets are defined as tagged if they pass the 80% working point of the GN2X tagger and anti-tagged if not. The pre-fit and post-fit distributions, without systematics, of tagged and anti-tagged Large-R jets are shown in Figure 4.7. The agreement between data and Monte Carlo improves after the fit.

This represents the first iteration of this calibration setup for GN2X. Without separate scale factors for each working point, systematic uncertainties and binning in p_T, the scale factors are not yet optimised for use in boosted analy-

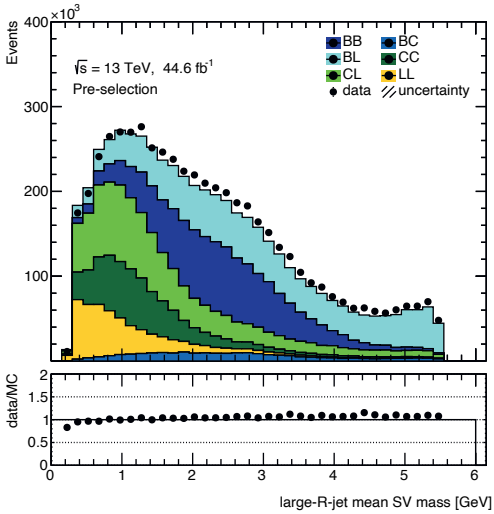

Figure 4.6: Large-R jet mean SV mass distribution of each MC flavour template and data.

ses. Subsequent and ongoing work extends and improves on this methodology, including with the use of a multivariate algorithm to select $g \to bb$ events. Additionally, *flat mass* working points are being adopted in newer versions of the tagger to avoid mass sculpting effects, defined such that efficiency is flat as a function of Large-R jet mass.

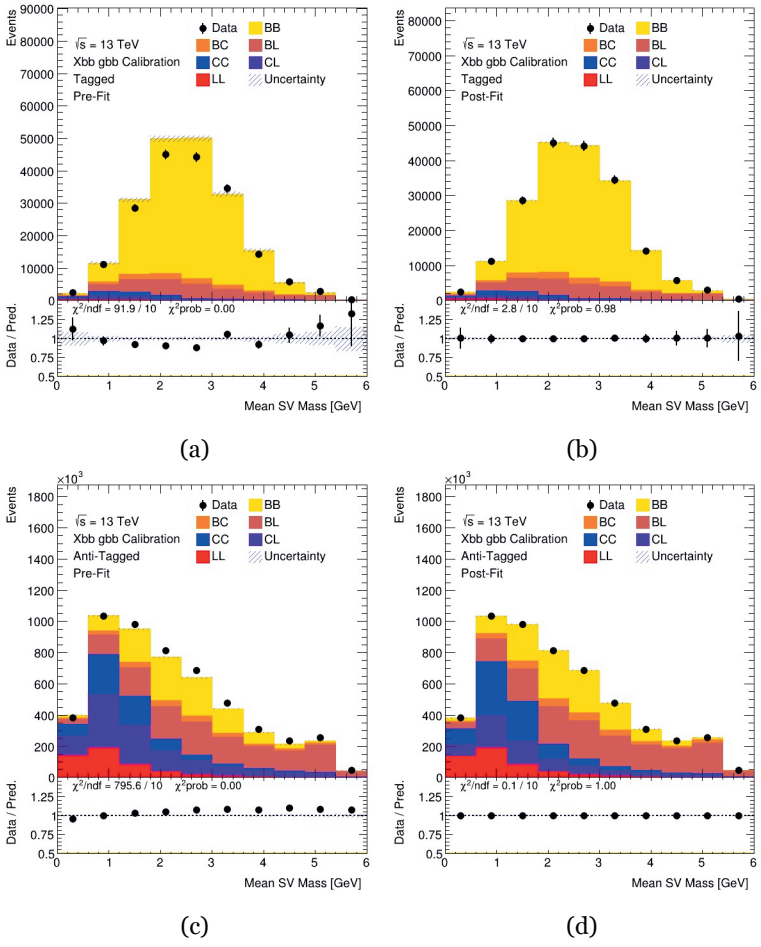

Figure 4.7: Distribution of the mean SV mass of Large-R jets in the six flavour templates for (a) tagged and (c) anti-tagged jets before the fit and (b) tagged and (d) anti-tagged jets after the fit.

Event Generation 5

Experimental searches and measurements by the ATLAS experiment rely on theoretical predictions to quantify the data's agreement with or deviation from the Standard Model. This requires accurate modelling of proton-proton collisions and a detailed understanding of the proton itself. Monte Carlo (MC) event generation is used widely in collider experiments as a tool to connect theoretical predictions to experimentally observable quantities by simulating proton-proton collisions and building distributions for different variables. This chapter briefly introduces the theoretical foundations underlying the phenomenology of proton-proton collisions at the LHC. Next, the different steps in the MC simulation chain are described. Finally, a discussion on MC generators and detector simulation is presented.

5.1 Phenomenology of proton-proton collisions

The proton is a composite particle composed of a bound state of three valence quarks (u, u, d) and a sea of gluons and quark-antiquark pairs that arise from quantum fluctuations in the vacuum. The valence quarks make up only a fraction of the proton's total mass, while the rest is dominated by QCD interactions between the quarks and gluons inside the proton. During a proton-proton collision, the constituent partons inside the protons interact with each other.

The probability of a given parton participating in an inelastic collision is determined using the parton distribution functions (PDFs) introduced in Section 5.1.1.

The physical processes involved in proton-proton collisions span a wide range of energy scales and are generally calculated and simulated separately. In a typical collision, two of the partons undergo a hard scattering process with a large momentum transfer. This corresponds to the process of interest, such as Higgs boson pair production. While the physics of the proton and its constituents occurs at long distances (low momentum transfer) and is therefore non-perturbative due to the running of the coupling constant described in Chapter 2, hard scattering occurs at short distances, in the perturbative regime of QCD, and can be calculated via the matrix element (ME) described in Section 5.1.2. Additional partons may be emitted by the incoming partons before the hard scattering (initial state radiation, or ISR) or by the outgoing partons (final state radiation, or FSR). This so-called parton showering still occurs in the perturbative regime.

At lower energies (~ 1 GeV), the process enters the non-perturbative regime in which partons hadronise into colour-neutral hadrons. This hadronisation leads to the formation of collimated sprays of particles known as jets. In addition to the partons involved in the hard scattering process, other 'spectator' partons from the involved protons may interact with each other and with the resulting partons of the hard scattering. Multiple parton interactions (MPI) may also occur in the same proton-proton collision. These effects are collectively known as the underlying event (UE). A schematic of a typical proton-proton collision with the processes described above is shown in Figure 5.1.

5.1.1 Parton distribution functions and factorisation

The probability of finding a parton of flavour i carrying a fraction x of the proton momentum at an energy scale μ_F is described by the parton distribution function (PDF) $f_i(x, \mu_F^2)$.

The factorisation theorem [72] provides the framework for calculating hadronic cross sections from partonic interactions. The cross section for a process $pp \to X$ can be written as the convolution of the PDFs of the relevant partons with the perturbatively calculable partonic cross section $\hat{\sigma}_{ij\to X}$:

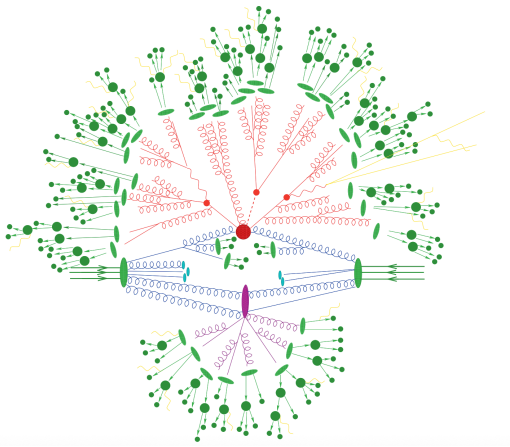

Figure 5.1: Schematic representation of the processes involved in a proton-proton collision in ATLAS and the corresponding steps in the MC simulation chain: the hard scattering process (red), PS and hadronisation (green), ISR (blue), UE (purple) and FSR (yellow) [71].

$$\sigma(pp \to X) = \sum_{i,j} \int_0^1 \mathrm{d}x_1 \int_0^1 \mathrm{d}x_2 \; f_i(x_1, \mu_F^2)\, f_j(x_2, \mu_F^2)\, \hat{\sigma}_{ij\to X}\left(\hat{s}, \mu_R^2, \mu_F^2\right), \tag{5.1}$$

where i, j are the flavours of the partons and $\hat{s} = x_1 x_2 s$ is the partonic centre-of-mass energy squared. The factorisation scale μ_F separates the non-perturbative proton structure from the perturbative hard-scattering process, while the renormalisation scale μ_R defines the scale at which the strong coupling constant α_s is evaluated. In practice, both scales are chosen to be characteristic of the energy scale of the hard-scattering process.

The PDFs cannot be calculated perturbatively and are instead determined from global fits to experimental data, including deep inelastic scattering measurements at HERA [73]. Their evolution with μ_F is described by the DGLAP equations [74–76], which account for parton radiation and splitting processes.

Several PDF sets are commonly used in LHC calculations, including PDF4LHC21 [77] and NNPDF30NLO [78] for the matrix element calculation described in Sec-

tion 5.1.2 and NNPDF23LO [79] for the parton showering described in Section 5.1.3. Non-perturbative parameters are often 'tuned' so that their predictions match experimental data. Different tunes are built for specific PDF sets. For example, the A14 tune is generally used with NNPDF23LO, both of which are used in the $HH \to b\bar{b}\gamma\gamma$ analysis described in this thesis. Uncertainties associated with the PDFs are propagated to predicted cross sections and acceptances in ATLAS analyses.

5.1.2 Matrix element Calculation

The partonic cross section in Equation 5.1, $\hat{\sigma}_{ij\to X}$ is calculated perturbatively as a power series in the strong coupling α_s:

$$\hat{\sigma}_{ij\to X} = \alpha_s^n \left[\hat{\sigma}^{\text{LO}} + \alpha_s\,\hat{\sigma}^{\text{NLO}} + \alpha_s^2\,\hat{\sigma}^{\text{NNLO}} + \cdots\right], \tag{5.2}$$

where n is the power of α_s at which the process first contributes. At leading order (LO), only tree-level Feynman diagrams are included. At next-to-leading order (NLO), one-loop virtual diagrams and real-emission diagrams with one additional final-state parton are added. Each successive order includes higher order loop and virtual corrections, providing a more accurate prediction of the cross section and reducing its dependence on the unphysical scales μ_R and μ_F. The variation of the result under changes to these scales is used as an estimate of the uncertainty due to missing higher-order terms. NLO accuracy ME calculation is the current standard for most processes simulated at the LHC. For some processes, such as ggF Higgs and di-Higgs production, NNLO calculations are available and are used to obtain the most precise predicted cross sections. Simulated samples are normalised to these cross sections.

5.1.3 Parton Showering and Hadronisation

Fixed-order ME calculations accurately describe hard emissions at large angles but become unreliable for lower energy or smaller angle emissions. The regions of phase space where these effects become important are known as the soft and collinear limits. At these limits, logarithmic enhancements in the perturbative series grow large and must be resummed to all orders. Parton showering addresses this by evolving the partonic final state from the hard scale down to a cutoff scale of order 1 GeV, a few times Λ_{QCD}, below which perturbation theory is no longer reliable. The evolution uses an approximation that is valid in the soft and collinear regime. The parton shower therefore

complements the ME by modelling the accompanying ISR and FSR where the fixed-order expansion is insufficient.

The transition between the two descriptions requires a matching procedure to avoid double-counting the first hard emission. The POWHEG method [80] generates the hardest emission at NLO accuracy within the ME before passing the event to the shower for subsequent softer radiation. This ensures that NLO accuracy is preserved while the shower fills in the remaining soft and collinear activity.

Once partons reach the non-perturbative scale, they hadronise into colour-neutral particles. In the Lund string model [81], the colour field between quarks is represented as a string of uniform energy density that fragments iteratively into hadrons as the quarks separate. The UE is modelled as part of the shower evolution.

5.2 Event Generators

In practice, the simulation chain described in Section 5.1 is implemented in software packages that specialise in one or more steps of the process. Dedicated ME generators such as POWHEG-BOX [82] and MadGraph5_aMC@NLO [83] compute the hard scatter at fixed perturbative order and interface to a separate shower program. Pythia8 [84] is widely used for parton showering, hadronisation, and the UE, and implements the POWHEG matching scheme when interfaced to a POWHEG-based ME calculation. For processes involving the decay of heavy-flavour hadrons, EVTGEN [85] provides dedicated decay tables with accurate branching fractions and spin correlations. Alternatively, multi-purpose generators such as Sherpa [86] handle ME generation and parton showering within a single framework, making them particularly suited to the simulation of multi-parton backgrounds such as diphoton production in association with jets. In the $HH \to b\bar{b}\gamma\gamma$ analysis described in Chapter 7, POWHEG-BOX is used for the ME calculation in the generation of the signal samples, Pythia8 is used for showering and hadronisation and EVTGEN is used for heavy flavour decay. The specific generators, PDF sets, and tunes used for each process in the $HH \to b\bar{b}\gamma\gamma$ analysis are detailed in Section 7.3.

5.3 Detector Simulation

The output of the simulation chain described above is a collection of final state particles and their four momenta with decay chain information. In order to have a meaningful comparison to real data collected by the ATLAS experiment, these events must be passed through a simulation of the ATLAS detector. Detector simulation software models the interactions of final state particles with the detector materials and the response of the detector to these interactions. Detector simulation is typically performed with GEANT4 [87]. Faster simulation software such as AtlFast3 [88] is also used in certain cases, including in the $HH \to b\bar{b}\gamma\gamma$ analysis. After this step, the same object reconstruction algorithms described in Chapter 4 are applied to the simulated events.

Higgs Boson Phenomenology at the LHC 6

Chapter 2 introduced the Higgs mechanism and the Higgs boson self-coupling, arguing that the measurement of the trilinear self-coupling parameter λ_3 is central to understanding the Higgs potential and providing a probe for BSM scenarios. λ_3 is experimentally accessible through Higgs boson pair production, commonly abbreviated to di-Higgs production. The experimental capabilities described in Chapters 3, 4 and 5 allow for the study of this process at the LHC with the ATLAS experiment.

Although this thesis is concerned with the production of Higgs boson pairs, an understanding of single-Higgs boson production and decay is a necessary prerequisite: the same mechanisms and decay channels underlie Higgs boson pair production and single-Higgs production constitutes one of the main backgrounds to the search presented in Chapter 7. This chapter therefore provides an overview of the phenomenology of the Higgs boson at the LHC, before turning to Higgs boson pair production. Sections 6.1 and 6.2 review the main production and decay modes of Higgs bosons at the LHC. Section 6.3 summarises the latest measurements for the Higgs boson's coupling to fermions and gauge bosons. Finally, Section 6.4 focuses on Higgs boson pair production, describing the main di-Higgs production and decay channels, recent results of di-Higgs searches at the LHC, and a brief mention of the prospects for triple Higgs production and the parameter λ_4.

6.1 Higgs Boson Production Mechanisms

The Higgs boson couples directly to all massive SM particles, encoded by different terms in the Higgs sector Lagrangian. Coupling to fermions proceeds via the Yukawa interaction, while both the coupling to gauge bosons and the self-coupling proceed via the Higgs mechanism. From the Lagrangian and the measured mass of the Higgs boson, the cross sections of different Higgs boson production modes from proton-proton collisions at a given energy can be calculated.

The most common Higgs boson production modes at the LHC are shown in Figure 6.1. These are, in order of decreasing cross section: gluon-gluon fusion (ggF), vector boson fusion (VBF), associated production with a vector boson (VH), associated production with a top or bottom quark pair (ttH/bbH), and associated production with a single top quark (tH). The cross section for the main production modes for a 125 GeV Higgs boson as a function of the centre-of-mass energy of the proton-proton collisions is shown in Figure 6.2.

The ggF process, in which gluon-gluon interactions create a Higgs boson through a heavy quark loop, dominates Higgs boson production at the LHC, accounting for approximately 87% of the total production rate. The second most common production mode is VBF, in which two (anti-)quarks scatter off of each other and exchange a W or Z boson, producing a Higgs boson and two energetic forward jets from the scattered quarks. The forward jets are a distinctive feature of the VBF production process and can be exploited to distinguish it from background. VBF accounts for about 7% of the total production rate.

In VH, which accounts for approximately 4%, a Higgs boson is radiated by a W or Z boson produced in a quark-anti-quark interaction. The ttH and bbH processes produce a Higgs boson via radiation from a top or bottom quark. These both account for approximately 1% of the total production rate. Finally, the tH process, in which a Higgs boson is radiated from a top quark produced in association with a single top quark, has the smallest cross section of the main production modes, at around 0.1% of the total production rate.

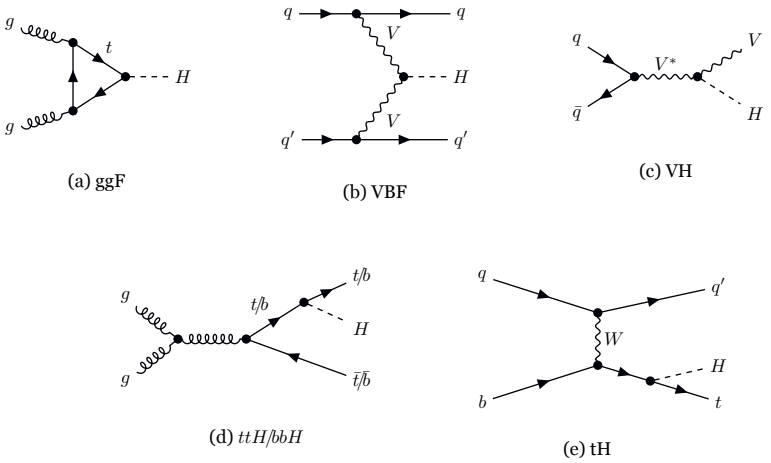

Figure 6.1: Main Higgs boson production modes at the LHC: (a) gluon–gluon fusion (ggF), (b) vector boson fusion (VBF), (c) associated production with a vector boson (VH), (d) associated production with heavy quarks (ttH/bbH), and (e) single top associated production (tH).

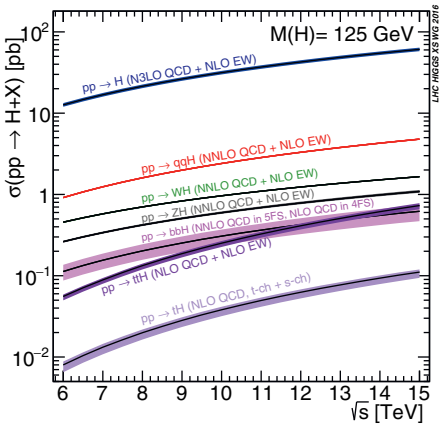

Figure 6.2: Cross section of Higgs boson production in proton-proton collisions as a function of centre-of-mass energy for the main Higgs production modes. Here, $pp \to H$ refers to ggF production and $pp \to qqH$ refers to VBF production. Taken from [24].

6.2 Higgs Boson Decay Channels

Once produced, the Higgs boson decays promptly into a variety of possible final states, with a total decay width of $\Gamma_H \approx 4.1\,\text{MeV}$ at $m_H = 125\,\text{GeV}$ [24], corresponding to a mean lifetime of approximately $1.6 \times 10^{-22}\,\text{s}$. The branching ratio (BR) into a given final state is determined by the partial width of that decay relative to the total width. The partial widths in turn depend on the Higgs coupling strengths and the available phase space. The predicted branching ratios as a function of m_H are shown in Figure 6.3, evaluated at $m_H = 125\,\text{GeV}$ for the remainder of this section.

The dominant decay mode is $H \to b\bar{b}$, with a branching ratio of approximately 58%, followed by $H \to WW^*$ at 21%. The general pattern reflects the fact that the Higgs field couples more strongly to heavier particles, as discussed in Section 6.3. However, the correspondence between particle mass and branching ratio is not straightforward. The b-quark, with $m_b \approx 4.2\,\text{GeV}$, is considerably lighter than the W boson at $m_W \approx 80.4\,\text{GeV}$, yet $H \to b\bar{b}$ is the dominant decay mode. This is due to the fact that at $m_H = 125\,\text{GeV}$, the Higgs boson is below the threshold for on-shell WW production. $H \to WW^*$ then proceeds via one off-shell W^*, which significantly suppresses the partial width. The $H \to b\bar{b}$ decay, by contrast, is fully on-shell and additionally benefits from a factor of three enhancement from quark colour. The remaining decay modes are $H \to gg$ at 8.6% (proceeding via a top-quark loop), $H \to \tau\tau$ at 6.3%, $H \to c\bar{c}$ at 2.9%, $H \to ZZ^*$ at 2.6%, and the rare loop-induced decay $H \to \gamma\gamma$ at 0.23%.

Representative Feynman diagrams for the two decay modes relevant to this thesis are shown in Figure 6.4. The $H \to b\bar{b}$ decay proceeds through a single tree-level Yukawa vertex, a topology shared by all of the other fermionic decay modes and, with the corresponding gauge vertex, by $H \to WW^*$ and $H \to ZZ^*$. The $H \to \gamma\gamma$ decay, by contrast, has no tree-level contribution: the Higgs boson does not couple directly to massless photons, and the decay proceeds instead through loops of charged particles, dominated by the W boson and the top quark. The two loop amplitudes interfere destructively, and it is this loop suppression that accounts for the small branching ratio.

A large branching ratio, however, is not sufficient to guarantee experimental sensitivity to a given decay channel. The sensitivity is also determined by the severity of the backgrounds and the efficiency and resolution with which

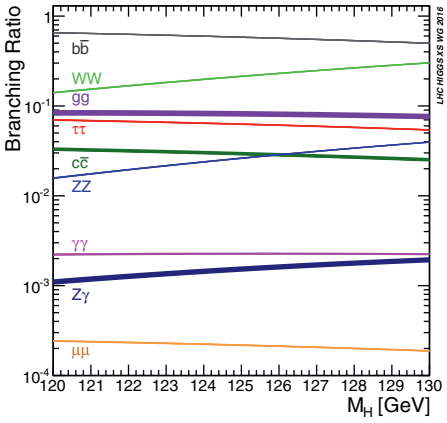

Figure 6.3: Branching ratios of the Higgs boson decay modes as a function of the Higgs boson mass. Taken from [24].

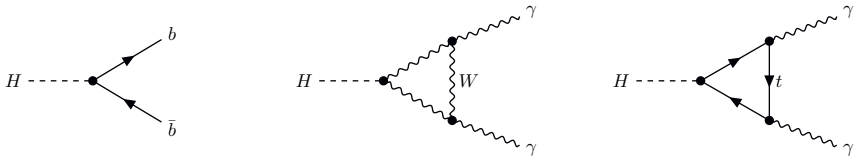

(a) (b) (c)

Figure 6.4: Feynman diagrams for the two Higgs boson decay modes relevant to this thesis: the tree-level decay $H \to b\bar{b}$ (a), and the loop-induced decay $H \to \gamma\gamma$ proceeding via a W boson loop (b) and a top quark loop (c). The topology of (a) is shared by the other fermionic decay modes, with the b-quark replaced by the corresponding fermion.

the final state can be reconstructed. The $H \to b\bar{b}$ channel illustrates this well: despite having the highest branching ratio by far, it is experimentally challenging due to the large QCD multijet background and the difficulty of triggering on purely hadronic final states. Conversely, the $H \to \gamma\gamma$ channel, with a branching ratio below 0.3%, was one of the two primary discovery channels for the Higgs boson in 2012 alongside $H \to ZZ^* \to 4\ell$ [1, 2]. Both channels benefit from fully reconstructed final states with excellent invariant mass resolution, which allows the signal to be isolated as a narrow peak above a smooth,

well-characterised background. The $H \to ZZ^* \to 4\ell$ channel in particular has a signal-to-background ratio close to unity in the signal region. By contrast, channels such as $H \to WW^*$ suffer from irreducible missing energy due to the neutrinos in the leptonic decay of the W, preventing a sharp mass peak and significantly complicating the separation of signal from background. These considerations are central to the choice of final state in Higgs boson pair production searches, which are discussed in detail later in this chapter.

6.3 Higgs Boson Couplings

As discussed in the previous section, the Higgs boson's couplings to fermions and gauge bosons arise from different mechanisms. Consequently, the mass dependence of the couplings is different for fermions and gauge bosons. For fermions, the coupling scales proportionally with the mass

$$g_{Hff} \propto m_f, \tag{6.1}$$

while for gauge bosons it scales with the mass squared

$$g_{HVV} \propto m_V^2. \tag{6.2}$$

The measurement of these couplings is a crucial test of the SM and has been central to the physics program at the LHC since the discovery of the Higgs boson itself.

Conventionally, deviations from the Standard Model are expressed in terms of coupling modifiers κ, defined such that κ_j^2 scales the partial width or production cross section of the corresponding process relative to its Standard Model prediction [89]. For example, κ_t modifies the top Yukawa coupling with respect to its Standard Model value, and a value of unity corresponds to exact agreement with the Standard Model. This framework will be used throughout the remainder of this thesis.

Roughly 7 million Higgs bosons were produced in proton-proton collisions during Run 2. Although a much smaller fraction of these were reconstructed and identified, this dataset is large enough to perform precise measurements of the Higgs boson's coupling to the heavier SM particles. The coupling to the Z boson is measured in the $H \to ZZ^* \to 4\ell$ discovery channel and the coupling

to the $W^{\pm}$ boson is measured through both the $H \to WW^* \to \ell\nu\ell\nu$ channel and via the VBF and VH production channels.

Among the fermion couplings accessible through Higgs boson decays, the experimentally most accessible are those probed by the $H \to b\bar{b}$ and $H \to \tau\tau$ channels, which are sensitive to the third generation Yukawa couplings and were measured from Run 2 data [69, 90–92]. The first and second generation fermion couplings are more challenging to measure due to the significantly lower BR of the corresponding decays. Notably, the first evidence for the Higgs boson coupling to muons was obtained in reference [93] despite the experimental challenges.

The measurement of the Higgs boson's coupling to the top quark is of particular interest, as the top quark is the heaviest SM particle with a predicted Yukawa coupling of close to unity. The top Yukawa coupling is also present in the dominant ggF production mode and, as mentioned in Chapter 2, it plays a role in the energy dependence of the self-coupling λ, which is relevant for the vacuum stability of the Higgs potential. Unlike analogous measurements such as $H \to b\bar{b}$, there exists no on-shell $H \to t\bar{t}$ process due to the large mass of the top quark ($m_t \approx 173\,\text{GeV}$). However, the top Yukawa coupling is directly accessible at tree-level through the $t\bar{t}H$ production mode. This measurement, despite being experimentally challenging due to the small cross section and complex final state topology, was achieved with the Run 2 dataset by both the ATLAS and CMS experiments [94].

Reference [95] performs a combination of individual coupling measurements from the ATLAS experiment with Run 2 data, and extracts the κ parameters using a global fit. In this result, limits on the charm Yukawa coupling to κ_c are obtained, although no significant observation of this coupling has been made yet. The resulting coupling values, or square root of the coupling for gauge bosons, as a function of the particle mass are shown in Figure 6.5. All measured values are found to be compatible with their SM predictions. These measurements represent the most precise tests of Higgs coupling universality to date. Coupling measurements for the first generation fermions and the strange quark are still unattainable at the current experimental precision.

After these, the final SM particle for which the coupling is not measured is the Higgs boson itself. In contrast to the remaining fermion couplings, κ_λ is directly related to the shape of the Higgs potential and therefore constitutes

a direct test of the Higgs mechanism itself. A deviation from the SM prediction for the fermion couplings may point to some BSM scenario that provides an additional source of mass beyond the VEV of the Higgs field. A deviation from the SM prediction for κ_λ is a more fundamental consequence because it points to a non-minimal Higgs potential. This would provide insight into the theoretical tensions described in Chapter 2, such as the vacuum stability. The measurement of κ_λ through Higgs boson pair production is discussed in the following section.

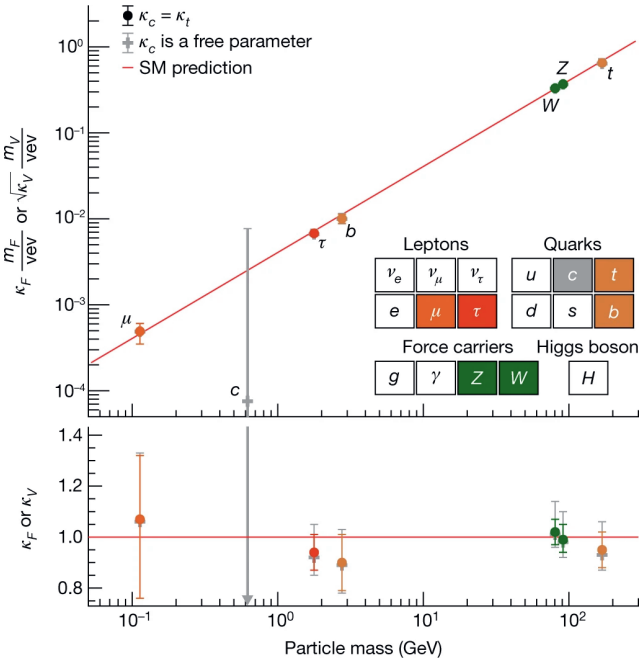

Figure 6.5: Summary of measurements of the Higgs boson couplings to fermions and gauge bosons, expressed as a ratio to the SM prediction. Measurements were done with κ_c fixed to its SM prediction (coloured points) and left to float (grey points). The red line corresponds to the SM expectation. Taken from [96].

6.4 Higgs Boson Pair Production

Higgs boson pair production is any process that creates a final state of two Higgs bosons. The inclusive cross section of Higgs boson pair production at the LHC is approximately 1000 times lower than that of single-Higgs production. Like single-Higgs, the dominant mode of Higgs boson pair production at the LHC is gluon-gluon fusion (ggF HH), accounting for approximately 90% of the total cross section. This process has two main contributing diagrams at leading order, shown in Figure 6.6.

The first is the triangle diagram, in which a single off-shell Higgs boson is created in a top quark loop and then decays into two on-shell Higgs bosons. This diagram depends on both κ_λ and κ_t. The second is the box diagram, in which the two Higgs bosons are created in a top quark loop, containing only κ_t vertices and therefore not sensitive to κ_λ. The amplitudes of the box and the triangle diagram interfere destructively due to the opposite signs of their form factors. This has several consequences for the cross section and kinematics of the ggF HH process, and these effects depend on the value of the self-coupling.

For $\kappa_\lambda = 0$, only the box diagram contributes and the cross section is enhanced relative to the SM. For $\kappa_\lambda \approx 2.45$, the cancellation is maximal, corresponding to the minimum cross section. For the SM scenario $\kappa_\lambda = 1$, not far from the maximum cancellation, the cross section remains low. For extreme (low or high) values of κ_λ, the cross section grows once the triangle diagram dominates. The cross section as a function of κ_λ for different Higgs boson pair production modes is shown in Figure 6.7.

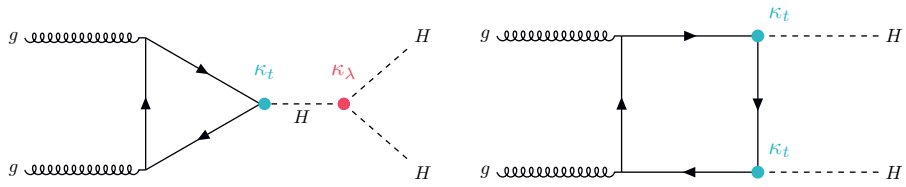

(a) (b)

Figure 6.6: The ggF HH production process, consisting at leading order of the triangle (a) and box (b) diagrams.

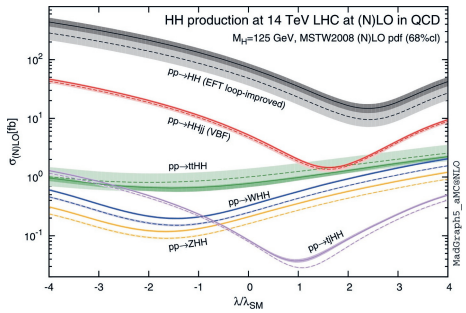

Figure 6.7: The cross section of different Higgs boson pair production modes as a function of κ_λ at $\sqrt{s} = 14\,\text{TeV}$. For ggF HH, a minimum can be seen at $\kappa_\lambda \approx 2.45$. Taken from [97].

The interference also has consequences for the kinematics of the ggF HH process, in particular for the di-Higgs invariant mass (m_{HH}) spectrum. The triangle diagram has a virtual Higgs propagator, therefore its amplitude peaks near the production threshold of $m_{HH} = 2m_H = 250\,\text{GeV}$. The box diagram has no resonant propagator and therefore a much broader m_{HH} spectrum that peaks at higher values of m_{HH}, falling off after $400\,\text{GeV}$. The relative contributions of both diagrams and their interference can be seen in Figure 6.8. Consequently, the $\kappa_\lambda = 0$ scenario produces a harder m_{HH} spectrum and more boosted Higgs bosons, while for $\kappa_\lambda \gg 1$ the spectrum is softer and enhanced near $2m_H$. These effects give the m_{HH} distribution a characteristic shape that can be used to discriminate between different values of κ_λ.

Like for single-Higgs production, the subdominant Higgs boson pair production mode is vector boson fusion (VBF HH), in which two $W^\pm$ or Z bosons fuse and produce two Higgs bosons. This can proceed via three diagrams at leading order. In Figure 6.9a, the vector bosons produce a single off-shell Higgs boson that decays into two on-shell Higgs bosons, providing direct sensitivity to κ_λ. In Figure 6.9b, a third vector boson is exchanged and the two Higgs bosons are produced through VVH vertices. The relevant coupling in this case is κ_V. The third VBF HH diagram is shown in Figure 6.9c, in which the two Higgs bosons are produced directly through a VVHH vertex, which is sensitive to the coupling between two vector bosons and two Higgs bosons, κ_{2V}. The VBF HH process has a smaller cross section than ggF HH, at around 5% of the total HH cross section. However, like the corresponding single-Higgs

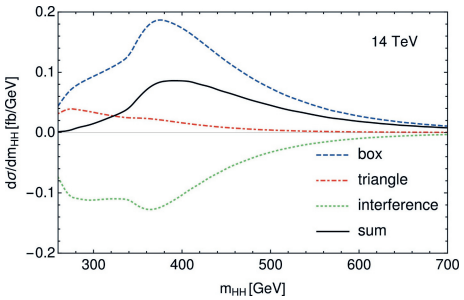

Figure 6.8: The m_{HH} spectrum of the box and triangle diagrams, their sum and their interference. Taken from [98].

production mode, it has a distinctive signature of two energetic forward jets, which can be exploited to separate it from background and from ggF HH. It is therefore typically considered alongside ggF HH as a signal in most Higgs boson pair production searches.

Other modes of Higgs boson pair production include associated production with a vector boson (VH HH) and associated production with a top quark pair (ttHH), but these, like the corresponding single-Higgs production modes, have much smaller cross sections than ggF HH and VBF HH and are therefore not discussed in detail.

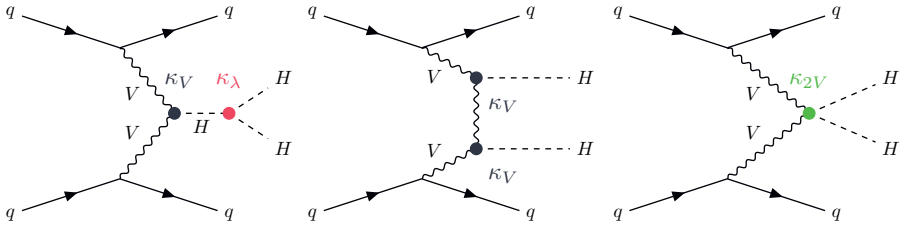

Figure 6.9: The VBF HH production process.

6.4.1 State of the Art of Higgs Boson Pair Production Searches with the ATLAS Experiment

Searches for Higgs boson pair production target different final states depending on how the Higgs pairs decay. Table 6.1 shows the branching fractions of the main di-Higgs decay modes which follow a similar trend as the corresponding single-Higgs processes. The dominant channel is $HH \to 4b$, followed by $HH \to bbWW$. However, as discussed in Section 6.2, the sensitivity of a given channel to both HH production and κ_λ depends not only on the branching ratio of the relevant process, but also on the amount and reducibility of the background and the resolution and reconstruction efficiency of the objects in the final state.

	bb	WW	$\tau\tau$	ZZ	$\gamma\gamma$
bb	34%				
WW	25%	4.6%			
$\tau\tau$	7.3%	2.7%	0.39%		
ZZ	3.1%	1.1%	0.33%	0.069%	
$\gamma\gamma$	0.26%	0.10%	0.028%	0.012%	0.0005%

Table 6.1: Branching ratios for Higgs boson pair production final states across the dominant decay modes bb, WW, $\tau\tau$, ZZ, and $\gamma\gamma$, computed from [24].

Using the full Run 2 dataset, the ATLAS experiment has performed legacy measurements in numerous Higgs boson pair production channels, obtaining both upper limits on the Higgs boson pair production signal strength

$$\mu_{HH} = \frac{\sigma_{HH}}{\sigma_{SM}}, \tag{6.3}$$

where σ_{HH} is the measured Higgs boson pair production cross section and σ_{SM} is the SM prediction, and 95% confidence intervals (CI) on κ_λ and κ_{2V}. The three most sensitive channels are $HH \to 4b$, $HH \to bb\tau\tau$ and $HH \to b\bar{b}\gamma\gamma$.

The $HH \to 4b$ channel is characterised by its large branching ratio and busy final state and therefore suffers from a large background from QCD multijet

processes. Both boosted [99] and resolved [100] topologies are explored to cover a larger region of the phase space. The VBF boosted channel in particular is extremely sensitive to κ_{2V}.

The $HH \to bb\tau\tau$ channel [101] includes the semi-leptonic and fully hadronic decay modes of the τ leptons. Its branching ratio is lower, but due to having only two b-jets in the final state, its backgrounds are manageable compared to the $4b$ channel.

The $HH \to b\bar{b}\gamma\gamma$ channel has the lowest branching ratio of the three most sensitive channels. However, due to the ATLAS detector's excellent photon reconstruction and mass resolution, the $H \to \gamma\gamma$ decay has a clean detector signature and allows for competitive sensitivity. Chapter 7 is dedicated to this channel.

The two other Higgs boson pair production channels considered by ATLAS consist of multiple less sensitive final states combined. The $HH \to bb\ell\ell + \text{MET}$ channel [102] covers final states with two $b-$quarks, two leptons and neutrinos. This can arise from $HH \to bbWW$, $HH \to bbZZ$ and the fully leptonic final state of $HH \to bb\tau\tau$. Finally, the $HH \to$ multileptons channel [103] consists of a combination of multiple final states with different numbers of leptons, hadronically decaying τ leptons and diphotons.

These final states are combined in reference [96] to form the ATLAS Run 2 legacy Higgs boson pair production combination, allowing for more stringent limits to be set on μ_{HH} and κ_λ than in any individual channel. Figure 6.10a shows the observed and expected upper limit on μ_{HH} for each individual channel and their combination, with a combined expected limit of 2.4. Figures 6.10b and 6.10c show the likelihood curves for κ_λ and κ_{2V} respectively for the individual channels and their combination. Of the individual channels, $HH \to bb\tau\tau$ is most sensitive to μ_{HH}, $HH \to b\bar{b}\gamma\gamma$ is most sensitive to κ_λ and $HH \to 4b$ is most sensitive to κ_{2V}.

This combination represents the state of the art of Higgs boson pair production searches with the ATLAS experiment, setting the most stringent limits on Higgs boson pair production and the self-coupling to date with the full Run 2 dataset. With the addition of Run 3 data, along with improvements in reconstruction and analysis techniques, the sensitivity to these parameters is expected to increase significantly. The first result to include Run 3 data is

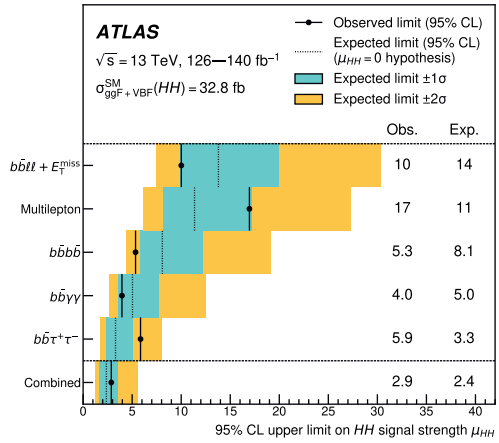

(a)

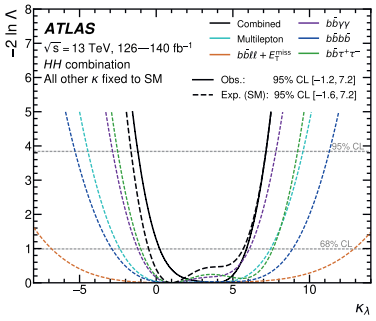

(b)

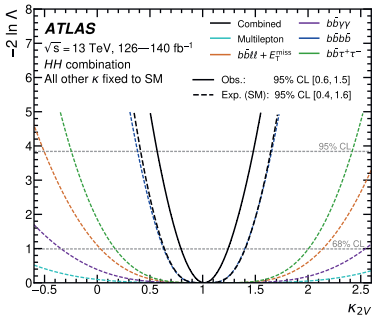

(c)

Figure 6.10: Results of the ATLAS Run 2 legacy Higgs boson pair production combination. (a) The observed and expected upper limits on μ_{HH} (a), the likelihood curve for κ_λ (b) and the likelihood curve for κ_{2V} (c) are shown. Individual channels are shown along with their combination. Taken from [96].

the $HH \to b\bar{b}\gamma\gamma$ channel in reference [104], discussed extensively in the next chapter, while new results of searches in other channels, and finally their combination, are expected to follow after the end of the Run 3 data taking period.

6.4.2 Triple Higgs

Although Higgs boson pair production and the trilinear self-coupling are the main focus of this thesis, it is worth noting that triple Higgs production, which is sensitive to the quartic self-coupling $\kappa_4 = \frac{\lambda_4}{\lambda_4^{SM}}$, has also become a process of interest at the LHC. The dominant mode of triple Higgs production at the LHC is shown in Figure 6.11.

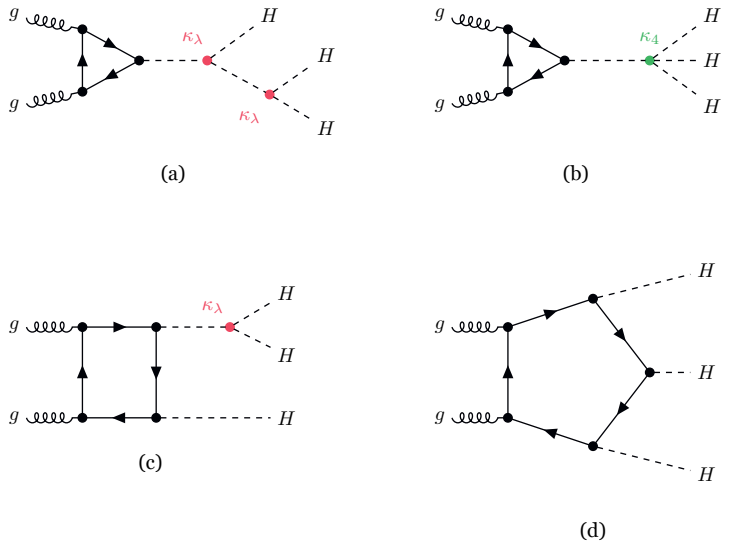

Figure 6.11: Leading order Feynman diagrams for triple Higgs production via gluon-gluon fusion. Diagrams (a) and (b) are sensitive to the quartic self-coupling, while diagrams (c) and (d) are not.

It is a rare process even compared to Higgs boson pair production, with a cross section on the order of 0.1 fb, placing it beyond the reach of current experimental sensitivity. However, ATLAS has placed the first limits on triple Higgs production in the $6b$ channel with the full Run 2 dataset, obtaining an

upper limit of $\sigma_{HHH} < 59\,\text{fb}$ [105]. This result is the first experimental probe of triple Higgs production and the quartic self-coupling, showing the feasibility of such measurements in preparation for the higher energy and luminosity.

The Search for Non-Resonant Higgs Boson Pair Production in the $b\bar{b}\gamma\gamma$ Final State 7

7.1 Motivation

The measurement of Higgs boson pair production provides direct access to the Higgs boson self-coupling, a fundamental parameter of the Higgs potential. Deviations from the SM prediction for Higgs boson pair production could indicate new physics beyond the Standard Model, while an observation of SM Higgs boson pair production would constitute a milestone in our experimental understanding of the Higgs sector. The study of Higgs boson pair production is therefore a key component of the LHC physics program.

This chapter focuses on the $HH \to b\bar{b}\gamma\gamma$ channel. As mentioned in Chapter 6, SM Higgs boson pair production occurs at a rate 1000 times lower than that of single-Higgs production, making it an extremely rare process. Furthermore, the $HH \to b\bar{b}\gamma\gamma$ channel makes up only 0.26% of the total Higgs boson pair production branching ratio. Its experimental signature, however, allows for a unique trade-off between the large branching ratio of the $H \to b\bar{b}$ decay and the excellent photon trigger and reconstruction efficiency of the $H \to \gamma\gamma$ decay in the ATLAS detector. Thanks to this trade-off, $HH \to b\bar{b}\gamma\gamma$ is one of the most sensitive channels to Higgs boson pair production and the most sen-

sitive channel to the Higgs boson trilinear self-coupling κ_λ in ATLAS as of the Run 2 combination [96]. The most recent search for Higgs boson pair production in the $HH \to b\bar{b}\gamma\gamma$ channel using the Run 2 dataset with a luminosity of $140\,\text{fb}^{-1}$ was performed by the ATLAS Experiment in Reference [106]. This analysis is hereafter referred to as the Run 2 legacy analysis.

The Run 3 data-taking period has exceeded luminosity expectations, with $168\,\text{fb}^{-1}$ of data collected as of the end of 2024. The analysis presented in this chapter and published in Reference [104] combines the full Run 2 dataset with Run 3 data up to and including 2024.

7.2 Overview

This chapter presents the search for Higgs boson pair production in the $b\bar{b}\gamma\gamma$ final state using the full Run 2 and partial Run 3 data. This search is for *non-resonant* Higgs pair production, in which a Higgs boson (as opposed to a new scalar resonant particle) produced in a proton-proton collision decays into two Higgs bosons. The analysis strategy is broadly based on that of the ATLAS Run 2 legacy analysis. For this reason, the Run 2 legacy analysis is used as a reference from which to measure improvements. A main difference is the addition of the Run 3 dataset, taken at a higher collision energy and luminosity. The addition of Run 3 data provides $308\,\text{fb}^{-1}$ of data, more than doubling the total integrated luminosity with respect to the Run 2 legacy analysis and significantly enhancing the sensitivity to Higgs boson pair production. Furthermore, improvements in combined performance tools and analysis techniques developed during Run 2 are leveraged to further improve the sensitivity of the search. This makes the analysis presented in this chapter a timely and important update to previous Higgs boson pair production searches in the $b\bar{b}\gamma\gamma$ final state. The treatment and combination of the two data-taking periods are discussed later in this chapter. Important differences and improvements with respect to the Run 2 legacy analysis are highlighted in the text.

This chapter describes the $HH \to b\bar{b}\gamma\gamma$ analysis strategy, beginning with the data and simulation samples used for analysis in Section 7.3. The parameterisation of the signal in terms of the relevant couplings is described in Section 7.4. The object-level selection is described in Section 7.5. In Section 7.6, the use of a kinematic fit to improve the invariant mass resolution of the $H \to b\bar{b}$

system is detailed, followed by the event-level selection in Section 7.7. The event categorisation strategy is described in Section 7.8, including the use of a Boosted Decision Tree (BDT) algorithm to separate signal from background events. The signal and background modelling is discussed in Section 7.9, followed by the systematic uncertainties that affect the results in Section 7.11. The statistical analysis methods used to extract the signal from the data are presented in Section 7.12. Finally, the fit to the diphoton invariant mass is described, leading to the results of the search in Section 7.13. The results of a sensitivity study on the kinematic fit are also presented in Section 7.13.2.

During the development of the analysis strategy, in order to avoid bias, data is blinded in the signal region, which is defined as the diphoton invariant mass window around the Higgs mass $m_H = 125\,\text{GeV}$ at $m_{\gamma\gamma}\,\epsilon\,[120, 130]$ GeV. The optimisation of the selection and categorisation, as well as the signal and background modelling is therefore performed on MC simulated events and/or data events outside the signal region (referred to as the data sidebands). Once the analysis strategy is frozen, the data in the signal region is unblinded and the final fit is performed on the full unblinded dataset.

7.3 Data and Simulation Samples

Data Samples

This analysis is performed on proton-proton collision data collected by the ATLAS detector during Run 2 of the LHC at a centre-of-mass energy $\sqrt{s} = 13\,\text{TeV}$ between 2015 and 2018, and during Run 3 at a centre-of-mass energy $\sqrt{s} = 13.6\,\text{TeV}$ between 2022 and 2024. All collected data is subject to data-quality requirements that monitor the detector subsystems to identify issues in the integrity of the data and exclude data that may be compromised [107]. After such requirements are applied, the total integrated luminosity amounts to $140\,\text{fb}^{-1}$ for Run 2 data and $168\,\text{fb}^{-1}$ for the partial Run 3 data, with uncertainties of 0.83% and 4% respectively [108, 109].

Signal Samples

The signal MC samples include both ggF HH and VBF HH production modes. Other production modes have comparatively tiny branching ratios and are ignored. To incorporate both SM and BSM scenarios, samples are generated at different values of the coupling modifiers κ_λ, κ_{2V} and κ_V. The specific values

chosen are discussed further in Section 7.4.

The ggF HH samples are generated at NLO using the POWHEG-BOX-V2+FT event generator [80, 82] with full NLO corrections using a finite top quark mass. The parton distribution function (PDF) set PDF4LHC21 [77] is used for the matrix element. The decay of the Higgs bosons, parton showering, hadronisation and the underlying event are simulated with PYTHIA 8.309 [84].

The VBF HH samples are generated at LO using MADGRAPH5_aMC@NLO v2.6.1 [83] with the NNPDF30NLO PDF set [78]. Parton showering and hadronisation are again performed with PYTHIA 8.309 using the A14 NNPDF23LO tune [84, 110], and EVTGEN is used for bottom and charm hadron decays.

Finally, all signal events are passed through a fast simulation tool AtlFast3, based on a parametric description of the ATLAS detector response [88]. The generators, matrix element (ME) PDF sets and parton showering (PS) PDF sets and tunes are shown in Table 7.1 for SM and each BSM configuration of ggF HH and VBF HH samples.

Background Samples

The main backgrounds to Higgs boson pair production in the $b\bar{b}\gamma\gamma$ final state arise from processes that mimic the $b\bar{b}\gamma\gamma$ signature. The dominant background is non-resonant diphoton production in association with jets, or $\gamma\gamma+$jets. Production of events with either one photon and one jet (γj) or two jets (jj) in which one or both jets mimic the signature of one or 2 photons is also included in this background category. The diphoton invariant mass ($m_{\gamma\gamma}$) spectrum of the $\gamma\gamma$ + jets process is a smoothly falling distribution with no resonance peak. It is therefore also referred to as continuum or non-resonant background. Sherpa 2.2.14 [86] is used for both event generation and parton showering for the continuum background.

In addition to diphoton production in association with jets of any flavour ($\gamma\gamma+$jets), a filtered sample of diphoton production with two b-jets is also simulated ($\gamma\gamma+b\bar{b}$). This additional sample is only used for background modelling studies, as described further in Section 7.9. Continuum background events are also interfaced to AtlFast3. The use of fast simulation is justified by the large

Table 7.1: Summary of di-Higgs signal samples. Standard Model configurations are highlighted in bold.

Process	κ_λ	κ_{2V}	κ_V	Generator	PDF (ME)	PDF+Tune (PS)
ggF HH	**1**	–	–	Powheg + Pythia 8 + EvtGen	PDFLHC21	A14NNPDF23LO
	0	–	–			
	5	–	–			
	−1	–	–			
	2.5	–	–			
	10	–	–			
VBF HH	**1**	**1**	**1**	MadGraph5_aMC@NLO + Pythia 8 + EvtGen	NNPDF30NLO	A14NNPDF23LO
	0	1	1			
	2	1	1			
	10	1	1			
	1	0	1			
	1	0.5	1			
	1	1.5	1			
	1	2	1			
	1	3	1			
	1	1	0.5			
	1	1	1.5			
	0	0	1			
	−5	1	0.5			

number of events required to accurately model the continuum background.

The sub-dominant background is single-Higgs production, in which the Higgs boson decays into a photon pair. The $m_{\gamma\gamma}$ distribution therefore has a resonance peak centred at $m_H = 125\,\text{GeV}$, also referred to as resonant background. All main single-Higgs production modes (ggH, VBFH, WH, ZH, ttH, tH) are considered for the resonant background. Single-top associated production (tH) is split into the s-channel and t-channel contribution (tHjb) and the W-associated contribution (tHW). W-associated production (WH) is separated into W^+H and W^-H. ZH is also separated into quark-initiated ($qq \to ZH$) and gluon-initiated ($gg \to ZH$) production.

The ME and PS generators, PDF set and tune used for each single-Higgs process are summarised in Table 7.2. Detector simulation for the single-Higgs background is performed by GEANT4. In all simulations, the Higgs boson mass is assumed to be 125 GeV, which is compatible with the latest measurements [23].

Table 7.2: Summary of single-Higgs and continuum background samples.

Process	Generator	PDF (ME)	PDF+Tune (PS)
ggH	Powheg + Pythia 8 + EvtGen	NNPDF30NNLO PDF4LHC21	AZNLOCTEQ6 A14NNPDF23LO
VBF	Powheg + Pythia 8 + EvtGen	NNPDF30 PDF4LHC21	AZNLOCTEQ6 A14NNPDF23LO
W^+H	Powheg MiNLO + Pythia 8 + EvtGen	NNPDF30 PDF4LHC21	AZNLOCTEQ6 A14NNPDF23LO
W^-H	Powheg MiNLO + Pythia 8 + EvtGen	NNPDF30 PDF4LHC21	AZNLOCTEQ6 A14NNPDF23LO
$q\bar{q} \to ZH$	Powheg MiNLO + Pythia 8 + EvtGen	NNPDF30 PDF4LHC21	AZNLOCTEQ6 A14NNPDF23LO
$gg \to ZH$	Powheg MiNLO + Pythia 8 + EvtGen	NNPDF30 PDF4LHC21	AZNLOCTEQ6 A14NNPDF23LO
ttH	Powheg + Pythia 8 + EvtGen	NNPDF30 PDF4LHC21	A14NNPDF23LO
bbH	Powheg + Pythia 8 + EvtGen	NNPDF30 PDF4LHC21	A14NNPDF23LO
tHjb	MadGraph5_aMC@NLO + Pythia 8 + EvtGen	NNPDF30	A14NNPDF23LO
tHW	MadGraph5_aMC@NLO + Pythia 8 + EvtGen	NNPDF30	A14NNPDF23LO
$\gamma\gamma$ + jets	Sherpa 2.2.14	NNPDF30NNLO	–
$\gamma\gamma + b\bar{b}$	Sherpa 2.2.14	NNPDF30NNLO	–

7.4 κ_λ and κ_{2V} Parameterisation

The cross sections of ggF HH and VBF HH production depend on four coupling modifiers, as explained in Chapter 6: the trilinear Higgs boson self-coupling modifier κ_λ, the quartic HHVV coupling modifier κ_{2V}, the top-quark Yukawa coupling modifier κ_t and the HVV coupling modifier κ_V. In particular, ggF HH depends directly on κ_λ and κ_t, while VBF HH depends directly on κ_λ, κ_{2V} and κ_V. Only κ_λ and κ_{2V} are targeted in this search. Simulating κ_t variations would be computationally expensive, and since κ_t is already well-constrained by single-Higgs measurements [95], it is justified to fix it to its SM value of 1.

However, κ_V is less well-constrained and enters the VBF HH cross section in a more complex way, so this parameter cannot be fixed to its SM value. It is therefore varied in the parameterisation along with κ_λ and κ_{2V} and kept as a nuisance parameter in the statistical fit. At leading order, the total amplitude of VBF HH production can be expressed as a sum of the three diagrams in Figure 6.9:

$$\mathcal{A}(\kappa_\lambda, \kappa_{2V}, \kappa_V) = A_1 \cdot \kappa_\lambda \kappa_V + A_2 \cdot \kappa_{2V} + A_3 \cdot \kappa_V^2, \tag{7.1}$$

where A_1, A_2 and A_3 are the normalisation coefficients for each diagram. The differential cross section $\frac{d\sigma_{VBF}}{d\Phi}$ is proportional to the square of the amplitude, and can be expressed as:

$$\begin{aligned}
\frac{d\sigma_{\text{VBF}}}{d\Phi}(\kappa_\lambda, \kappa_{2V}, \kappa_V) &\propto \left|\mathcal{A}_{\text{VBF}}(\kappa_\lambda, \kappa_{2V}, \kappa_V)\right|^2 \\
&= \left|A_1 \cdot \kappa_\lambda \kappa_V + A_2 \cdot \kappa_{2V} + A_3 \cdot \kappa_V^2\right|^2 \\
&= |A_1|^2 \cdot \kappa_\lambda^2 \kappa_V^2 + |A_2|^2 \cdot \kappa_{2V}^2 + |A_3|^2 \cdot \kappa_V^4 \\
&\quad + 2\,\Re(A_1 A_2^*) \cdot \kappa_\lambda \kappa_V \kappa_{2V} \\
&\quad + 2\,\Re(A_1 A_3^*) \cdot \kappa_\lambda \kappa_V^3 + 2\,\Re(A_2 A_3^*) \cdot \kappa_{2V} \kappa_V^2 \\
&= x_1 \cdot \kappa_\lambda^2 \kappa_V^2 + x_2 \cdot \kappa_{2V}^2 + x_3 \cdot \kappa_V^4 \\
&\quad + x_4 \cdot \kappa_\lambda \kappa_V \kappa_{2V} + x_5 \cdot \kappa_\lambda \kappa_V^3 + x_6 \cdot \kappa_{2V} \kappa_V^2
\end{aligned} \tag{7.2}$$

where the x_i for $i = 1, ...6$ are the coefficients of each of the polynomial terms in the three couplings. Because the differential cross section is polynomial in the couplings, the coefficients can be solved exactly by a system of linear

equations with six basis samples generated with different values of κ_λ, κ_V and κ_{2V}. These basis samples and their respective cross sections are shown in Table 7.4. The choice of basis samples is made to include coupling values within a similar range to the experimental constraints on these couplings, while avoiding samples that have large weights or large negative weights.

In practice, this linear system is solved using the simulated cross section in each reconstructed m_{HH} bin for each basis sample, because the coupling changes not only the total cross section but the shape of the m_{HH} distribution. The differential VBF HH cross section then becomes:

$$\begin{aligned}
\frac{d\sigma_{\mathrm{VBF}}}{d\Phi}(\kappa_\lambda, \kappa_{2V}, \kappa_V) = &\left(\frac{68}{135}\kappa_{2V}^2 + \frac{20}{27}\kappa_{2V}\kappa_\lambda\kappa_V - 4\kappa_{2V}\kappa_V^2 + \frac{1}{9}\kappa_\lambda^2\kappa_V^2 - \frac{56}{27}\kappa_\lambda\kappa_V^3 + \frac{772}{135}\kappa_V^4\right) \\
&\cdot \frac{d\sigma_{\mathrm{VBF}}}{d\Phi}[1,1,1] \\
&+ \left(-\frac{4}{5}\kappa_{2V}^2 + 4\kappa_{2V}\kappa_V^2 - \frac{16}{5}\kappa_V^4\right) \\
&\cdot \frac{d\sigma_{\mathrm{VBF}}}{d\Phi}[1,1.5,1] \\
&+ \left(\frac{11}{60}\kappa_{2V}^2 - \frac{19}{24}\kappa_{2V}\kappa_\lambda\kappa_V + \frac{1}{3}\kappa_{2V}\kappa_V^2 - \frac{1}{8}\kappa_\lambda^2\kappa_V^2 + \frac{13}{6}\kappa_\lambda\kappa_V^3 - \frac{53}{30}\kappa_V^4\right) \\
&\cdot \frac{d\sigma_{\mathrm{VBF}}}{d\Phi}[2,1,1] \\
&+ \frac{1}{9}\left(-\frac{11}{60}\kappa_{2V}^2 + \frac{11}{24}\kappa_{2V}\kappa_\lambda\kappa_V + \frac{1}{8}\kappa_{2V}\kappa_V^2 - \kappa_\lambda^2\kappa_V^2 - \frac{5}{6}\kappa_\lambda\kappa_V^3 + \frac{13}{30}\kappa_V^4\right) \\
&\cdot \frac{d\sigma_{\mathrm{VBF}}}{d\Phi}[10,1,1] \\
&+ \left(\frac{88}{45}\kappa_{2V}^2 + \frac{4}{9}\kappa_{2V}\kappa_\lambda\kappa_V - \frac{16}{3}\kappa_{2V}\kappa_V^2 - \frac{4}{9}\kappa_\lambda^2\kappa_V^2 + \frac{152}{45}\kappa_V^4\right) \\
&\cdot \frac{d\sigma_{\mathrm{VBF}}}{d\Phi}[1,1,0.5] \\
&+ \left(\frac{8}{45}\kappa_{2V}^2 - \frac{4}{9}\kappa_{2V}\kappa_\lambda\kappa_V + \frac{4}{9}\kappa_\lambda\kappa_V^3 - \frac{8}{45}\kappa_V^4\right) \\
&\cdot \frac{d\sigma_{\mathrm{VBF}}}{d\Phi}[-5,1,0.5],
\end{aligned} \tag{7.3}$$

where $\frac{d\sigma_{\mathrm{VBF}}}{d\Phi}[1,1,1]$, $\frac{d\sigma_{\mathrm{VBF}}}{d\Phi}[1,1.5,1]$, ... are the differential cross sections per

m_{HH} bin for each basis sample.

The ggF HH parameterisation follows the same procedure of using basis samples to model the cross section at arbitrary values of κ_λ, using the differential cross section in each m_{HH} bin. The contributing diagrams to this process are the box and the triangle diagrams shown in Figure 6.6. The total amplitude at LO can therefore be expressed as:

$$\mathcal{A}_{\text{ggF}}(\kappa_\lambda, \kappa_t) = A_1 \cdot \kappa_\lambda \kappa_t + A_2 \cdot \kappa_t^2, \tag{7.4}$$

where A_1 and A_2 are again the normalisation coefficients for each diagram. In a similar expansion to the VBF HH case, the differential cross section can be expressed as:

$$\frac{d\sigma_{ggF}(\kappa_\lambda)}{d\Phi} \propto |\mathcal{A}_{ggF}(\kappa_\lambda, \kappa_t)|^2 = x_1 \cdot \kappa_\lambda^2 \kappa_t^2 + x_2 \cdot \kappa_\lambda \kappa_t^3 + x_3 \cdot \kappa_t^4 \tag{7.5}$$

To determine the three coefficients x_i, three basis samples with different values of κ_λ are required. The basis samples used are shown in Table 7.3 along with their respective cross sections. The specific basis is chosen to reduce the statistical uncertainties around the limits on κ_λ and around the maximum interference point of $\kappa_\lambda \approx 2.4$. Setting κ_t to its SM value of 1, the parameterised cross section becomes:

$$\begin{aligned} \frac{d\sigma_{ggF}}{d\Phi}(\kappa_\lambda) = \Bigg[& \left(1 + \frac{1}{5}\kappa_\lambda^2 - \frac{6}{5}\kappa_\lambda\right) \cdot \frac{d\sigma_{ggF}}{d\Phi}[0,1] \\ & + \left(\frac{5}{4}\kappa_\lambda - \frac{1}{4}\kappa_\lambda^2\right) \cdot \frac{d\sigma_{ggF}}{d\Phi}[1,1] \\ & + \left(\frac{1}{20}\kappa_\lambda^2 - \frac{1}{20}\kappa_\lambda\right) \cdot \frac{d\sigma_{ggF}}{d\Phi}[5,1] \Bigg] \end{aligned} \tag{7.6}$$

The parameterisation above defines how cross sections and therefore event yields are expected to vary under deviations of κ_λ and κ_{2V} from the Standard Model. With the samples used and the parameterisation of the signal established, the next section describes the steps to isolate the $HH \to b\bar{b}\gamma\gamma$ process from background in the data. Selections are applied at both the object level and the event level.

κ_λ	$\sigma_{13\,\text{TeV}}$ (fb)	$\sigma_{13.6\,\text{TeV}}$ (fb)
0	68.56	75.75
1	30.77	34.13
5	90.82	99.65
-1	127.50	140.70
2.5	13.67	15.13
10	641.20	704.20

Table 7.3: ggF HH samples produced and their Run 2 and Run 3 cross sections. The samples marked as **bold** are used as the basis for the cross section parameterisation. The non-basis samples are used for validation and some for training the BDT used in event categorisation.

κ_λ	κ_{2V}	κ_V	$\sigma_{13\,\text{TeV}}$ (fb)	$\sigma_{13.6\,\text{TeV}}$ (fb)
1	**1**	**1**	1.687	1.874
1	**1**	**0.5**	10.297	11.635
10	**1**	**1**	93.392	102.328
1	**1.5**	**1**	3.172	3.613
2	**1**	**1**	1.525	1.685
-5	**1**	**0.5**	6.127	6.943
1	2	1	13.631	15.504
1	0	1	25.642	28.851
1	0.5	1	9.159	10.268
1	3	1	61.372	69.626
0	1	1	4.427	4.880
1	1	1.5	62.793	70.270
0	0	1	34.991	39.171

Table 7.4: VBF HH samples produced and their Run 2 and Run 3 cross sections. The samples marked as **bold** are the ones used as the basis for the cross section parameterisation. The non-basis samples are used for validation and some for training the BDT used in event categorisation.

7.5 Object-Level Selection

Since photons can be triggered on efficiently in the ATLAS detector, diphoton triggers are used to select $HH \to b\bar{b}\gamma\gamma$ events. While trigger requirements differ slightly between years, all triggers require at least two photons to be present and pass trigger-level identification requirements. In particular, the 2015-2016 dataset was obtained with a diphoton trigger requiring a loose identification working point and HLT-level E_T thresholds of 35 and 25 GeV for the two legs of the trigger. To account for increased pile-up, the datasets of subsequent years used a medium identification working point and the same E_T thresholds. The efficiency of the 25 GeV leg of the diphoton trigger for each year of Run 2, for different values of transverse energy E_T, pseudorapidity η and pile-up $\langle\mu\rangle$ is shown in Figure 7.1. The trigger efficiency in 2022 for different thresholds at different values of E_T and η is shown in Figure 7.2. Efficiencies are above 95% for $\mathrm{E}_T > 30$ GeV

Events that pass the trigger requirements are fully reconstructed offline. Photons, leptons and jets are reconstructed as outlined in Chapter 4. The primary vertex used is the hardest vertex. This is defined as the vertex giving the highest sum of squares of the transverse momenta of contributing tracks. Beyond the trigger-level requirements, additional selections are applied to photons and jets to be consistent with the $HH \to b\bar{b}\gamma\gamma$ signature and reduce the fraction of misidentified particles. These are described in the following section.

Photons and Jets

The photon selection largely follows that of the $H \to \gamma\gamma$ analysis [112]. Selected events are required to contain two photons within the pseudorapidity range $|\eta| < 2.37$, excluding the transition region between the barrel and endcap calorimeters ($1.37 < |\eta| < 1.52$). Photons are required to satisfy a tight identification working point. To ensure that the photon candidates are well-isolated from other activity in the event, photons must also pass the tight working point of the calorimeter and track-based isolation [57] described in Chapter 4.

The two leading photon candidates that pass these criteria are then required to have a ratio $p_T/m_{\gamma\gamma} > 0.35$ (0.25) for the leading (subleading) photon, and the diphoton invariant mass is required to be within the range $105 < m_{\gamma\gamma} < 160$ GeV.

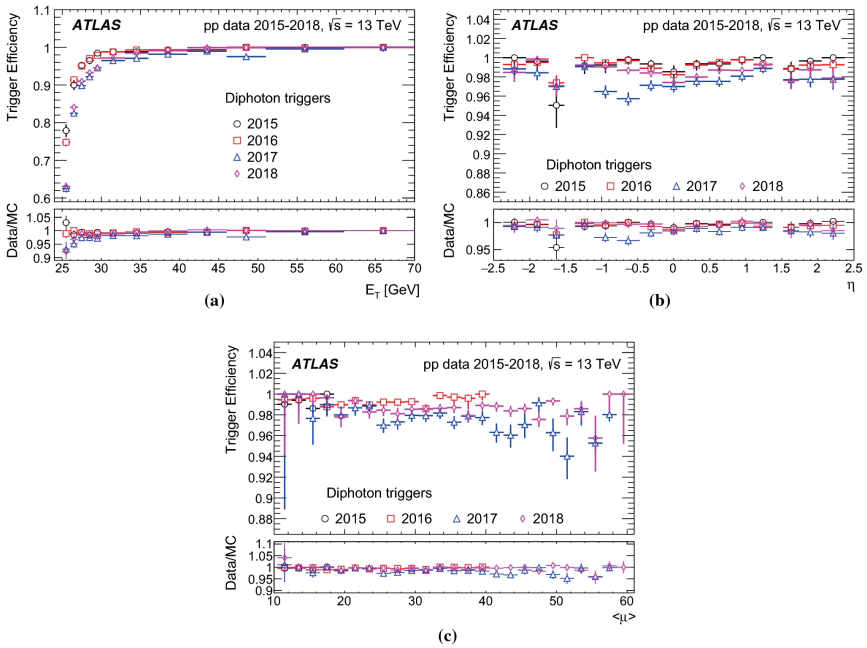

Figure 7.1: Trigger efficiency of the 25 GeV leg of the diphoton trigger for each year of Run 2 as a function of the transverse energy E_T (a), pseudorapidity η (b) and pile-up $\langle\mu\rangle$ (c), taken from Reference [111].

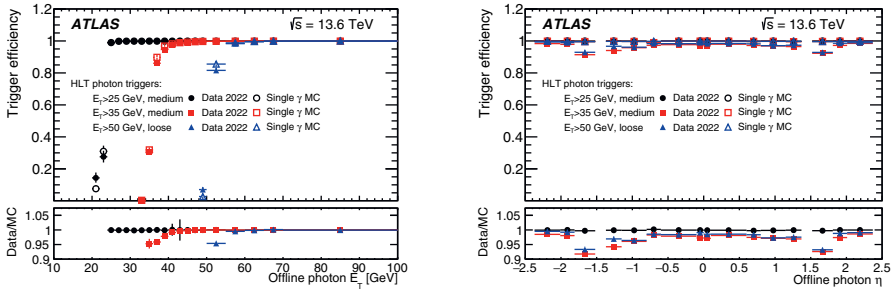

(a) (b)

Figure 7.2: Trigger efficiency (2022 data and MC) as a function of the transverse energy E_T (a) and pseudorapidity η (b) for the 25, 35 and 50 GeV legs of the diphoton trigger. Taken from Reference [47].

Jets are reconstructed from particle flow objects using the anti-k_t clustering algorithm [59] with a radius parameter $R = 0.4$ and calibrated to correct for detector effects and pile-up as discussed in Chapter 4. Jet candidates are required to have $p_T > 25$ GeV and $|\eta| < 4.5$. To ensure that jet candidates originate from the primary vertex, jets with $p_T < 60$ GeV and $|\eta| < 2.4$ are required to pass the jet vertex tagger (JVT) described in Chapter 4. Jets outside of this angular region are required to pass the forward jet-vertex tagger (fJVT). Only central jets (i.e. with $|\eta| < 2.5$ that fall within the inner detector's acceptance) are considered candidates for the $H \to b\bar{b}$ decay. Forward jets ($|\eta| > 2.5$) typically come from the VBF HH process.

b-jets

The GN2 algorithm [63] is used for the classification of jets originating from b-quarks, c-quarks, light-flavour quarks/gluons or taus, as described in Chapter 4. $H \to b\bar{b}$ candidate jets in this analysis must satisfy the 85% working point of GN2, which corresponds to a b-jet efficiency of 85% and misidentification rates of 0.01 for light-flavour jets and 0.2 for c-jets. This working point is chosen to maximise the sensitivity of the analysis: compared to the legacy analysis, the b-tagging working point is loosened due to the better background rejection of GN2 for the same signal efficiency compared to its predecessor.

b-jets in ATLAS are complex objects with imperfect reconstruction. For this reason, several calibrations of the b-jet energies are necessary. The first of these accounts for the energy contributions of muons from semileptonic b-hadron decays within *b*-jets. This is called a Muon-in-jet correction. It is applied by searching for muons within a cone of $\Delta R = 0.4$ around the jet axis. If a muon with a p_T above 4 GeV is found in this region, the muon four-momentum is added to that of the jet, and the jet energy is recalculated accordingly.

A second calibration called PtReco is applied to correct for the underestimation of the b-jet energy due to neutrinos from semileptonic b-hadron decays within the jet and out-of-cone effects. This correction is derived from simulated $H \to b\bar{b}$ events as a function of jet p_T and η, and applied to the *b*-jets after the Muon-in-jet correction. These two corrections combined improve the resolution of the $H \to b\bar{b}$ invariant mass $m_{b\bar{b}}$ by up to 17% in simulated $HH \to b\bar{b}\gamma\gamma$ events. However, there is still opportunity for further improvement of the $m_{b\bar{b}}$ resolution by employing more novel techniques. The following section intro-

duces the use of a kinematic fit to improve the $m_{b\bar{b}}$ resolution beyond the Muon-in-jet and PtReco corrections.

7.6 Kinematic Fit

The $HH \to b\bar{b}\gamma\gamma$ final state has a key asymmetry that can be exploited to improve the reconstruction of the $H \to b\bar{b}$ system: the photon four-momenta are measured with ∼2–3% precision by the electromagnetic calorimeter, while the b-jet four-momenta have only ∼15% resolution. A kinematic fit exploits this asymmetry by using the precisely-measured photon system as an anchor: imposing transverse momentum conservation across all reconstructed objects in the event constrains the jet four-momenta to be consistent with both the measured photon momenta and a model of the detector response. This yields a per-event correction to the b-jet kinematics and consequently an improved estimate of $m_{b\bar{b}}$. Figure 7.3 shows the momentum balance between the $H \to b\bar{b}$ and $H \to \gamma\gamma$ systems.

The main difference between this and the b-jet calibrations described in the previous section is that it is a per-event correction. The Muon-in-jet and PtReco corrections apply simulation-derived scale factors as a function of jet p_T and η, correcting for systematic shifts. The kinematic fit instead optimises the jet four-momenta event-by-event using the full event topology, and is therefore sensitive to per-event fluctuations that uniform corrections do not capture.

While $m_{\gamma\gamma}$ is chosen as the discriminating variable in the final fit, $m_{b\bar{b}}$ enters the analysis indirectly as a BDT input variable – and in fact the input variable with the highest discriminatory power. Improving its resolution therefore has a non-negligible effect on the analysis sensitivity, justifying the development of this novel per-event correction method.

The kinematic fit builds on the method described in References [113, 114], and is further developed for use in the Run 2 legacy analysis by F. Montereali in Reference [115]. In this implementation, an additional constraint is added to the baseline likelihood to improve, alongside $m_{b\bar{b}}$, the resolution of the four-body invariant mass $m^*_{b\bar{b}\gamma\gamma}$. This extension is discussed in Section 7.6.2, while the baseline kinematic fit is introduced below.

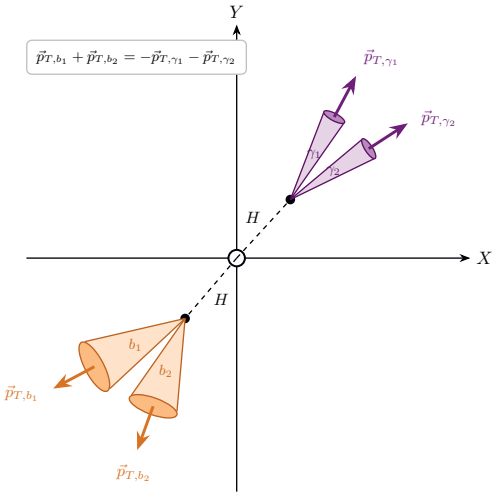

Figure 7.3: Schematic of the transverse momentum balance in the $HH \to b\bar{b}\gamma\gamma$ transverse plane. The kinematic fit imposes momentum conservation: the jet momenta (orange) are corrected using the precisely-measured photon momenta (purple). The cone width is indicative of the momentum resolution.

7.6.1 Kinematic Fit likelihood construction

The kinematic fit operates on the principle that the reconstructed four-momenta of the jets in an event are imperfect measurements of their true values, smeared by the finite resolution of the calorimeter and other effects. If a model can be constructed for how the detector smears these quantities, it is then possible to determine which set of "true" jet four-momenta is most consistent with both the observed measurements and the known physics of the event - for instance, that the total transverse momentum of all particles should sum to zero. The kinematic fit determines exactly this set of four-momenta.

It takes as inputs the reconstructed four-momenta of the candidate $HH \to b\bar{b}\gamma\gamma$ photons and b-jets, as well as all additional jets in the event, since these also contribute to the total momentum balance. The jet four-momenta are the free parameters of the fit, while the photon four-momenta are fixed. The fit then returns updated jet four-momenta that are most compatible with both the detector response model and the physical constraints imposed on the event, yielding

an improved estimate of $m_{b\bar{b}}$.

This is implemented as a maximum-likelihood estimation. For each event, a likelihood function is constructed that encodes the probability of observing the set of reconstructed quantities $\mathrm{y}_i^{\mathrm{obs}}$ given a set of underlying "true" values $\mathrm{y}_i^{\mathrm{fit}}$, where the index i represents the value of a given variable for a given object in the event. The likelihood is expressed as the product of probability density functions over all terms $f(y_i^{\mathrm{obs}}|y_i^{\mathrm{fit}})$:

$$\mathcal{L}(y_1^{\mathrm{obs}}, ..., y_N^{\mathrm{obs}}|y^{\mathrm{fit}_1}, ..., y_N^{\mathrm{fit}}) = \prod_i^N f(y_i^{\mathrm{obs}}|y_i^{\mathrm{fit}}) \tag{7.7}$$

The terms in the product fall into two categories. The first are transfer functions (TFs). These are per-object, per-variable probability density functions that model the detector response, or equivalently, how the detector transforms the true value of a kinematic quantity into its reconstructed value. These apply, in principle, to the variables that parameterise the four-momentum: energy, transverse momentum, pseudorapidity and azimuthal angle. A transfer function for the energy of a jet, for example, encodes the probability of reconstructing a jet with energy E^{obs} given that the true energy of the jet is E^{fit}. The derivation of the TFs is described in more detail in Section 7.6.1.1.

The second category of terms are global physical constraints $f_{\mathrm{constr.}}$, which impose event-level conditions on the fit parameters based on known properties of the collision. In the baseline kinematic fit, this consists of a transverse momentum conservation term. This is further discussed in Section 7.6.1.2, and the addition of a mass constraint term beyond the baseline likelihood expression is described in Section 7.6.2. Together, the TFs and global physical constraints fully define the structure of the likelihood.

In practice it is more convenient to minimise the negative log-likelihood (NLL), $-2\ln\mathcal{L}$, which transforms the product in Equation 7.7 into a sum. The NLL can then be expressed as the sum of contributions from transfer functions for the relevant variables $v = (\mathrm{E}, p_T, \eta, \phi)$ of objects $j = (\mathrm{jets}, \gamma)$ and global physical constraints:

$$-2\ln\mathcal{L} = \sum_{v=E,p_T,\eta,\phi} \sum_{j=\mathrm{jets},\gamma} [-2\ln \mathrm{TF}_v^j(v_j^{\mathrm{fit}}, v_j^{\mathrm{obs}})] - 2\ln f_{\mathrm{constr.}} \tag{7.8}$$

From this generalised form of the likelihood function, some assumptions and simplifications can be made. Firstly, the photon four-momenta are measured with ∼2–3% precision, an order of magnitude better than the jets. A photon transfer function would therefore be an extremely narrow distribution, contributing only a near-constant term to the NLL with negligible effect on the minimisation. In this limit, the impact of a photon transfer function on the fit result is negligible, and the photons can be approximated as fixed at their reconstructed values rather than treated as free parameters of the fit. Transfer functions are therefore not derived for photon variables.

Secondly, the effect of including jet η and ϕ transfer functions was tested on simulated data. These angular transfer functions were found to provide negligible improvement to the performance of the kinematic fit, while increasing its complexity significantly. Transfer functions are therefore included only for the p_T and E of the jets in the event. The likelihood expression after these simplifications becomes:

$$-2\ln\mathcal{L} = \sum_{v=E,p_T}\sum_{j=\text{jets}}[-2\ln \text{TF}_v^j(v_j^{\text{fit}}, v_j^{\text{obs}})] - 2\ln f_{\text{constr.}} \tag{7.9}$$

7.6.1.1 Transfer functions

Transfer functions (TFs) must accurately encapsulate detector smearing effects in an analytical, normalisable form suitable for a probability density function. In particular, they should describe the distribution of the relative difference between the values of the true and reconstructed variable, i.e. $\frac{v_{\text{True}}-v_{\text{Reco}}}{v_{\text{True}}}$. After standard jet calibrations are applied, the residual deviations between reconstructed and true jet kinematics still exhibit asymmetric, non-Gaussian tails, arising from effects such as semi-leptonic b-hadron decays and out-of-cone radiation that are not fully captured by average scale factor corrections. Standard analytical forms, including the Bukin and Double-Sided Crystal Ball functions, were tested in Reference [116] but found to model these tails inadequately. A systematic method for constructing flexible, normalizable asymmetric bell-shaped PDFs was therefore developed, yielding a function referred to as *TaOGaTa*. It consists of a Gaussian core with ascending and descending arctangent step-functions to allow for asymmetric tails. Its full form for a given variable *x* is given by:

$$
\begin{aligned}
TaOGaTa(x) = N \times & (\frac{1}{\pi}(\arctan(a(x-m)) + \frac{\pi}{2}))^{\alpha} \\
& \times \exp(-\frac{(x-\mu)^2}{2\sigma^2}) \\
& \times (\frac{1}{\pi}(\arctan(-b(x-n)) + \frac{\pi}{2}))^{\beta}
\end{aligned}
\tag{7.10}
$$

where a, b, m, n, α, β and μ are shape parameters and N is a normalisation factor. These parameters are extracted by fitting the *TaOGaTa* function to the distribution of the relative difference between the true and reconstructed quantities ($\frac{E_{\text{true}} - E_{\text{Reco}}}{E_{\text{True}}}$ and $\frac{p_T^{\text{true}} - p_T^{\text{Reco}}}{p_T^{\text{True}}}$) in simulated Run 2 $HH \to b\bar{b}\gamma\gamma$ events.

To account for the energy dependence of these distributions, they are binned in intervals of p_T. Furthermore, the energy resolution depends on the region of the detector in which the particle is reconstructed. Therefore, separate TFs are derived for jets in the central barrel region ($|\eta| < 1.37$), the forward end-cap region ($1.52 < |\eta| < 2.5$), the crack region between the barrel and the end-cap, where resolution is expected to be poor ($1.37 < |\eta| < 1.52$), and the region closest to the beam-pipe, where tracks cannot be reconstructed at all ($2.5 < |\eta| < 4.4$). Both energy and p_T transfer functions exhibit similar energy and angular dependence, thus the same p_T and detector region splits are used for both variables. The total number of transfer functions then amounts to: 4 detector regions $\times$ 6 p_T bins $\times$ 2 variables $\times$ 2 data-taking periods $=$ 96. The energy and p_T transfer functions for the two main detector regions, the barrel and the end-cap, are shown for the different p_T bins in Figure 7.4.

7.6.1.2 Global physical constraints

The second component of the kinematic fit likelihood is the global physical constraints $f_{\text{constr.}}$. These impose physical conditions on the fit parameters based on known properties of the event. In the baseline kinematic fit likelihood construction, the only constraint applied is that the total $\vec{p_T}$ of all the particles in a given event should be zero within the detector resolution. The conservation also holds for the separate components in the X and Y directions, p_X and p_Y.

To enforce $\sum p_X, \sum p_Y = 0$ directly is too strong a constraint given that the truth-level p_X and p_Y distributions, even before detector effects are considered, have a finite resolution around zero. This is due to effects such as missing transverse energy (MET) arising from semi-leptonic b-hadron decays, soft or

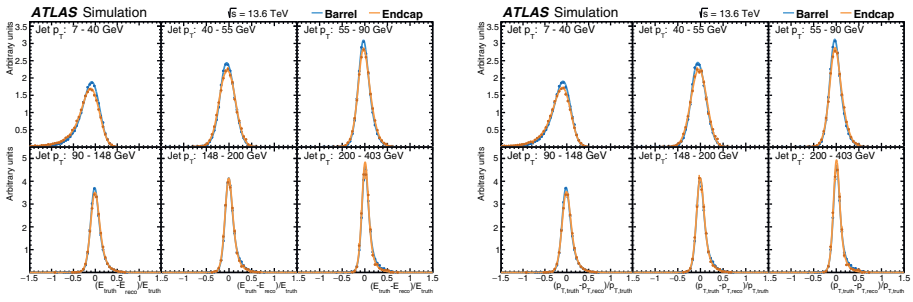

(a) (b)

Figure 7.4: Transfer functions for energy (a) and p_T (b) in the barrel and endcap regions for 6 p_T bins. The parameters are extracted by fitting the *TaOGaTa* function to the distribution of the relative difference between the true and reconstructed variables in simulated events.

out-of-cone radiation and forward particles beyond the jet η acceptance. Although the global transverse momentum is conserved, these effects contribute non-negligibly to the momentum balance of the candidate $HH \to b\bar{b}\gamma\gamma$ objects and additional jets.

The chosen approach therefore is to derive probability density functions to model the true distribution of p_X and p_Y. These are not well-modelled by Gaussian distributions because the effects mentioned previously are not purely random fluctuations. They may also have asymmetric and/or heavier tails, which renders a Gaussian model insufficient. The p_X and p_Y distributions are therefore modelled using a custom function that takes these considerations into account. It consists of an ascending and descending arctangent step function, similar to the *TaOGaTa* function but without the Gaussian core. The function, called *TaOTa*, is derived and discussed in references [115, 116]. It has the following form:

$$\begin{aligned} TaOTa(x) = N \times &(\frac{1}{\pi}(\arctan(a(x-m)) + \frac{\pi}{2}))^{\alpha} \\ &\times (\frac{1}{\pi}(\arctan(-b(x-n)) + \frac{\pi}{2}))^{\beta} \end{aligned} \tag{7.11}$$

where a, b, m, n, α, β are shape parameters and N is a normalisation factor. The parameters are obtained by fitting the function to the distributions of the

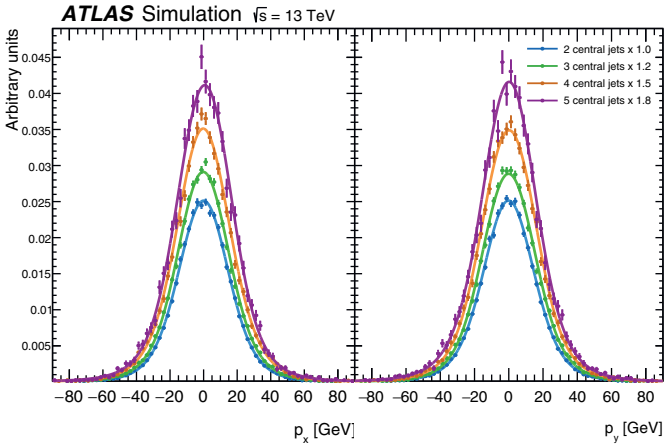

Figure 7.5: Distributions and PDFs of the x and y components of the total transverse momentum. Arbitrary units are used for visualisation purposes.

total truth-level p_X and p_Y of candidate jets and photons and additional jets in simulated Run 2 $HH \to b\bar{b}\gamma\gamma$ events. The truth-level jets are matched to the corresponding reconstructed jets by requiring the angular distance (ΔR) between them to be less than 0.1. This ensures that the truth-level distributions indeed reflect those of the full reconstructed event.

The presence of additional jets beyond the two candidate jets broadens the constraint distributions. These are therefore fitted separately for each possible count of additional *central* jets: 0, 1, 2, or 3, leading to a total of eight distributions. Forward jets ($|\eta| > 2.5$) are also included in the momentum sum, but the choice of distribution is indexed by the count of additional central jets only, since central jets are the dominant source of shape variation in the p_X and p_Y distributions. This count is bounded to at most three by the event-level requirement of fewer than six central jets per event (Section 7.7). The fitted p_X and p_Y PDF distributions for the different jet configurations are shown in Figure 7.5.

To allow the fit to give more or less weight to the constraint terms relative to the TFs, a scalar multiplier λ is investigated that multiplies the p_X and p_Y constraint terms. This term controls the strength of the constraints. A scan of

a large range of values yields $\lambda = 3.05$ as the optimal value that minimises the $m_{b\bar{b}}$ resolution in simulated $HH \to b\bar{b}\gamma\gamma$ events. The final form of the baseline kinematic fit likelihood function becomes:

$$\begin{aligned} -2\ln(\mathcal{L}) = \sum_{j=\text{jets}} \Bigg[& -2\ln\left(\text{TF}_E\left(\frac{E^j_{\text{Fit}} - E^j_{\text{Reco}}}{E^j_{\text{Fit}}}\right)\right)\Bigg] \\ & -2\ln\left(\text{TF}_{p_T}\left(\frac{p^j_{T,\text{Fit}} - p^j_{T,\text{Reco}}}{p^j_{T,\text{Fit}}}\right)\right) \\ & -2\lambda\ln[f_{p_X}(p_X)] - 2\lambda\ln[f_{p_Y}(p_Y)]\ . \end{aligned} \tag{7.12}$$

The first two terms are the energy and p_T transfer functions for all jets in the event. The last two terms are the momentum conservation constraint terms in the X and Y directions, represented by the PDFs discussed above and multiplied by the strength parameter λ.

7.6.2 Kinematic Fit Second Iteration

The baseline kinematic fit is run for each event and outputs updated jet four-momenta from which the corrected $m_{b\bar{b}}$ is calculated. This is the primary function of the kinematic fit. As mentioned previously, an additional term is added to the likelihood that performs the secondary function of improving the precision of the four-body invariant mass of the candidate jets and photons $m^*_{b\bar{b}\gamma\gamma}$(defined in Equation 7.15). This variable, like $m_{b\bar{b}}$, is used in the BDT classifiers and therefore plays a role in discrimination $HH \to b\bar{b}\gamma\gamma$ signal from background. The effect of this improvement is in fact most prominent not in $m^*_{b\bar{b}\gamma\gamma}$ directly, but in the difference between the reconstructed and the true value, $m^{*,Reco}_{b\bar{b}\gamma\gamma} - m^{*,True}_{b\bar{b}\gamma\gamma}$. Nonetheless, the performance of the kinematic fit can be seen to improve due to this constraint.

The additional term added to the likelihood penalises deviations of the fitted $m_{b\bar{b}}$ from the Higgs boson mass m_H, implemented as a penalty term $(m_{b\bar{b}} - m_H)^2$ added to the negative log-likelihood, weighted by a strength parameter λ_m described below. This constraint leverages the known mass of the $H \to b\bar{b}$ system to further correct the jet kinematics. The enforcement of this constraint would significantly bias the $m_{b\bar{b}}$ distribution, effectively forcing it to a narrow

peak at the Higgs boson mass. The $m_{b\bar{b}}$ distribution is therefore updated only using the output of the baseline kinematic fit. This effectively results in two iterations of the kinematic fit per event. The first iteration is performed by minimising the NLL in Equation 7.12 using the MIGRAD algorithm from ROOT's Minuit class [117] and updating the $m_{b\bar{b}}$ variable according to the fitted jet four-momenta. The second iteration minimises the NLL again in the same way but independently, this time including the Higgs boson mass constraint:

$$\begin{aligned} -2\ln(\mathcal{L}) = \sum_{j=\text{jets}} \Bigg[& -2\ln\left(\text{TF}_E\left(\frac{E^j_{\text{Fit}} - E^j_{\text{Reco}}}{E^j_{\text{Fit}}}\right)\right) \Bigg] \\ & -2\ln\left(\text{TF}_{p_T}\left(\frac{p^j_{T,\text{Fit}} - p^j_{T,\text{Reco}}}{p^j_{T,\text{Fit}}}\right)\right) \\ & -2\lambda\ln[f_{p_X}(p_X)] - 2\lambda\ln[f_{p_Y}(p_Y)] + \lambda_m(m_{b\bar{b}} - m_H)^2. \end{aligned} \tag{7.13}$$

The output jet four-momenta of the second iteration are then used to update all other variables that depend on the jet four-momenta, including the $m^*_{b\bar{b}\gamma\gamma}$ variable. Similarly to the strength parameter λ described above, λ_m is a scalar multiplier that controls the strength of the Higgs boson mass constraint term. A scan of possible values shows $\lambda_m = 0.1$ to give the optimal resolution of $(m^{*,Reco}_{b\bar{b}\gamma\gamma} - m^{*,True}_{b\bar{b}\gamma\gamma})$.

While the Higgs mass constraint term would bias the $m_{b\bar{b}}$ distribution severely, rendering it unsuitable for use, the bias effects on $m^*_{b\bar{b}\gamma\gamma}$ are less severe. Since $m^*_{b\bar{b}\gamma\gamma}$ is kinematically correlated with $m_{b\bar{b}}$, the constraint also biases $m^*_{b\bar{b}\gamma\gamma}$ to some degree. However, the effect of the constraint is to remove a source of resolution degradation in the $m^*_{b\bar{b}\gamma\gamma}$ variable, namely the mismatch between the true and reconstructed $m_{b\bar{b}}$. The effect of adding this constraint was checked on both signal and background simulated events. The shape of the $m^*_{b\bar{b}\gamma\gamma}$ distribution was not significantly altered in either case and no background sculpting effects were observed. Furthermore, the notably smaller optimal value of λ_m compared to λ means that the constraint is relatively weak compared to the momentum conservation constraints, which also helps to mitigate bias effects. The addition of this term was therefore found to be justified for the purpose of improving the $m^*_{b\bar{b}\gamma\gamma}$ resolution.

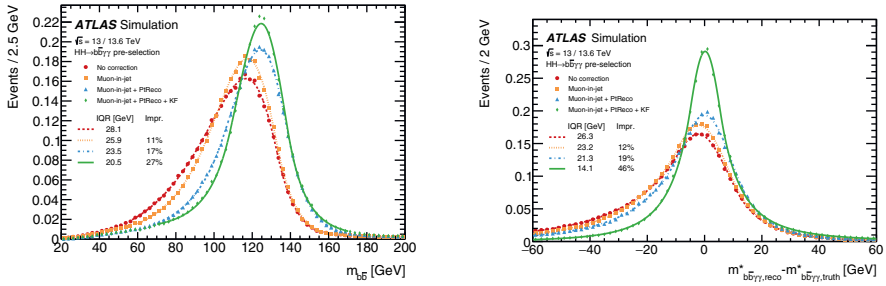

(a) (b)

Figure 7.6: $m_{b\bar{b}}$ (a) and $m^{*,Reco}_{b\bar{b}\gamma\gamma} - m^{*,True}_{b\bar{b}\gamma\gamma}$ (b) distributions for cases with no jet corrections (red), Muon-in-jet only (orange), Muon-in-jet and PtReco (blue) and Muon-in-jet, PtReco and the kinematic fit (green). $m_{b\bar{b}}$ distributions are fit to a Bukin function and $m^{*,Reco}_{b\bar{b}\gamma\gamma} - m^{*,True}_{b\bar{b}\gamma\gamma}$ is fit to a Double-Sided Crystal Ball. The inter-quartile range of the fitted function is used to quantify the resolution.

The effect of the kinematic fit is compared to different configurations of the other b-jet corrections in Figure 7.6. Figure 7.6a shows the $m_{b\bar{b}}$ distribution and 7.6b shows the $m^{*,Reco}_{b\bar{b}\gamma\gamma} - m^{*,True}_{b\bar{b}\gamma\gamma}$ distribution. The inter-quartile range is used as a proxy for the resolution of the fitted distributions. The kinematic fit improves the $m_{b\bar{b}}$ resolution by 27% and the $m^{*,Reco}_{b\bar{b}\gamma\gamma} - m^{*,True}_{b\bar{b}\gamma\gamma}$ by 46% compared to the case where no correction is applied.

Since $m_{b\bar{b}}$ and $m^{*}_{b\bar{b}\gamma\gamma}$ enter the final fit only indirectly through BDTs, improvements in the resolution of these variables do not scale directly into sensitivity to the final analysis targets μ_{HH}, κ_λ and κ_{2V}. A detailed study into the effect of the kinematic fit on the analysis sensitivity is shown in Section 7.13.2. Nonetheless, these improvements are promising. The kinematic fit is applied at this stage – after object-level selection but before event-level cuts – so that the corrected jet four-momenta and improved $m_{b\bar{b}}$ are available as inputs to the subsequent BDT-based event categorisation. The following section describes the event-level selection.

7.7 Event-level Selection

On top of the trigger and object requirements described above, cuts are applied at an event level to select $HH \to b\bar{b}\gamma\gamma$ candidate events and reduce backgrounds. Firstly, events must contain at least two b-jets that pass the 85% working point of the GN2 tagger, described in Chapter 4. Both the multiplicity and the working point of this cut are loosened with respect to the Run 2 legacy analysis thanks to the increased performance of GN2 over the previous tagger.

The b-jets in each event are ordered by their PCBT score. The first two b-jets that satisfy the 85% GN2 working point form the two $H \to b\bar{b}$ candidate jets. If two jets have the same PCBT score, they are then ordered by p_T.

$t\bar{t}H$ is a significant source of background to the $HH \to b\bar{b}\gamma\gamma$ process. It can produce both hadronically and leptonically decaying final states (or a combination of both). To reduce background from $t\bar{t}H$ with hadronically decaying top quarks, less than six central jets are required in each event. This aims to remove events arising from hadronic $t\bar{t}H$ background, which typically have a high jet multiplicity. To reduce the $t\bar{t}H$ background with leptonically decaying top quarks, events with any number of electrons or muons are rejected.

The distribution of the fitting variable $m_{\gamma\gamma}$ in the data sidebands after preselection is shown in Figure 7.7 compared with SM HH signal and backgrounds. The determination of the relative fractions of the simulated $\gamma\gamma+b\bar{b}$ and $\gamma\gamma+$jets in the continuum background and its normalisation is described in Section 7.9. The yields of data events remaining after the pre-selection are 2794 for Run 2 and 3392 for Run 3. The expected MC event yields for a given process are determined by the sum over all events of a per-event weight that takes into account the cross section of the process, the integrated luminosity and other generator and detector related effects, as well as calibration and trigger scale factors. The weight is evaluated as:

```
weight = generatorWeight * PileupWeight * Luminosity
         * Xsection * kFactor * FilterEff * SF / sumOfWeights
SF = btag_SF * jvt_SF * photon1_SF * photon2_SF
```
(7.14)

where `generatorWeight` is the weight assigned to the event by the MC generator, `PileupWeight` is the pile-up reweighting factor and `Luminosity` is the integrated luminosity for the respective data-taking period. `Xsection` is the theoretical

cross section for the given process, taken from the respective PDF set. `kFactor` corrects the theoretical cross section to a higher order calculation where available. `FilterEff` is a reweighting factor to account for the efficiency of any generator-level filters applied. `SF` is the product of calibration scale factors to account for data-MC differences. These include b-tagging, JVT and photon identification and isolation scale factors. Finally, `sumOfWeights` is the sum of the generator weights of all filtered simulated events for the given process, to normalise the total yield correctly. The MC signal yields (SM ggF and VBF HH combined) after pre-selection are 2.03 in Run 2 and 2.63 in Run 3. The full breakdown of MC yields per process is given in Appendix A.

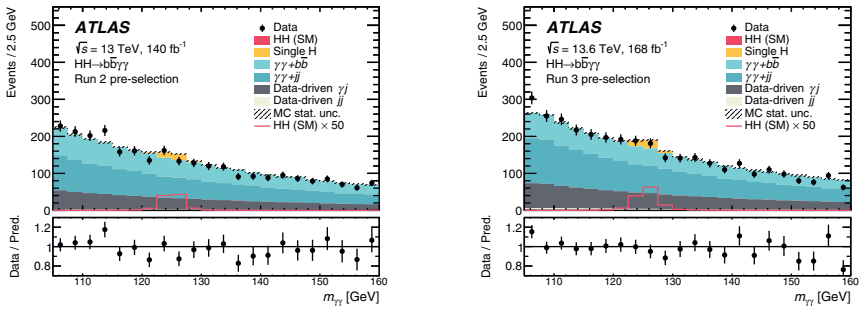

(a) (b)

Figure 7.7: Distributions of the diphoton invariant mass for events in data compared with the sum of the expected signal and background contributions after pre-selection for (a) Run 2 and (b) Run 3. The HH signal is shown as a filled histogram stacked on top of the backgrounds, normalised to the SM prediction. It is also shown unstacked as an unfilled histogram, scaled by a factor of 50 to enhance visibility.

7.8 Event Categorisation

To further reduce the background beyond the pre-selection cuts and improve the sensitivity of the analysis, events are divided into categories based on their kinematic properties and the output of a multivariate classifier. This approach not only further increases the signal-to-background ratio, but also allows for different categories to be optimised separately for SM or BSM signal hypotheses. One of the most sensitive variables to anomalous κ_λ values is the modified

four-body invariant mass $m^*_{b\bar{b}\gamma\gamma}$, defined in Equation 7.15. This variable peaks at lower values for κ_λ values different from 1. It is used for this reason as a basis for the event categorisation.

$$m^*_{b\bar{b}\gamma\gamma} = m_{b\bar{b}\gamma\gamma} - (m_{b\bar{b}} - 125\,\text{GeV}) - (m_{\gamma\gamma} - 125\,\text{GeV}) \tag{7.15}$$

A high-mass region (HM) is defined for $m^*_{b\bar{b}\gamma\gamma} > 350$ GeV that is most sensitive to SM ggF HH production (κ_λ close to 1) and both SM (κ_{2V} close to 1) and BSM (κ_{2V} different from 1) VBF HH production. A low-mass region (LM) is defined for $m^*_{b\bar{b}\gamma\gamma} < 350$ GeV that is most sensitive to BSM ggF HH production with κ_λ values different from 1, which enhance the cross section at low $m^*_{b\bar{b}\gamma\gamma}$ values.

A Boosted Decision Tree classifier (BDT), implemented with XGBoost [118], is trained separately in each of these regions to distinguish signal from background. Importantly, the Run 2 and Run 3 datasets are combined in the training of the BDTs in both regions instead of training a separate BDT for each data-taking period. This is hereafter referred to as the Run 2 and Run 3 correlation strategy and will be discussed further in Section 7.10. In both BDTs, all single-Higgs processes and $\gamma\gamma$ + jets are used as background. In the HM region, the BDT is trained using SM ggF HH and SM VBF HH, plus BSM VBF HH with (κ_λ, κ_V, κ_{2V}) = (0, 1, 1), (10, 1, 1), (1, 1.5, 1), (1, 3, 1), (-5, 1, 0.5) samples as signal. In the LM region, the BDT is trained using ggF HH samples with $\kappa_\lambda = 5$ and $\kappa_\lambda = 10$ and the same BSM VBF HH samples as in the HM region. The specific BSM samples used for training were inherited from the Run 2 legacy analysis and found to provide optimal coverage. The exact values have negligible impact on the training.

The hyperparameters of both BDTs were optimised in a grid-search method in [106] and have been re-used for the current analysis. The input variables used for the BDTs are summarised in Table 7.5. These variables are chosen for their discriminating power between signal and background. To avoid sculpting effects in the $m_{\gamma\gamma}$ distribution, $m_{\gamma\gamma}$ is not included in the BDT input variables, and other variables are decorrelated from $m_{\gamma\gamma}$ as much as possible, for example by dividing the p_T of the $H \to \gamma\gamma$ system by $m_{\gamma\gamma}$. It is not possible, however, to fully decorrelate all variables from $m_{\gamma\gamma}$. The output of the BDT is a score between 0 and 1 for each event, where events closer to 0 are more background-like and events closer to 1 are more signal-like.

Kinematic variables of candidate photons and *b*-jets
$p_T/m_{\gamma\gamma}$, η and ϕ of the candidate photons
p_T, η, ϕ and PCBT score of both candidate b-jets
p_T^{bb}, $\eta_{b\bar{b}}$, $\phi_{b\bar{b}}$ and $m_{b\bar{b}}$ of the $H \to b\bar{b}$ candidate
Variables to reduce $t\bar{t}H$ and ZH
Scalar sum of the transverse momenta of all jets: $H_T = \sum_{\text{jets}} p_T$
Single-topness $\chi_{Wt} = \min_{\text{jets}} \sqrt{\left(\frac{m_{j_1 j_2} - m_W}{m_W}\right)^2 + \left(\frac{m_{j_1 j_2} - m_t}{m_t}\right)^2}$ (where the minimum runs over all 3-jet combinations)
Missing transverse momentum E_T^{miss} and its azimuthal angle ϕ^{MET}
Additional jet variables
p_T, η, ϕ and PCBT score of the third and fourth jets if present
HH system variables
Modified four-body invariant mass $m^*_{b\bar{b}\gamma\gamma}$
Angular separation between photons $\Delta R(\gamma_1, \gamma_2)$
Angular separation between b-jets $\Delta R(b_1, b_2)$
VBF-specific variables
$\Delta\eta(j_1^{\text{VBF}}, j_2^{\text{VBF}})$ between the two VBF jets
VBF event-shape variables:
Di-jet invariant mass $m_{j_1^{\text{VBF}} j_2^{\text{VBF}}}$
Planar flow P_f and transverse sphericity $S_\perp$ (defined in refs)
p_T balance: $\frac{\left\|\sum_{i\in\{\gamma_1,\gamma_2,b_1,b_2\}} \vec{p_T^{\,i}}\right\|}{\sum_{i\in\{\gamma_1,\gamma_2,b_1,b_2\}} p_T^i}$

Table 7.5: Summary of BDT input variables.

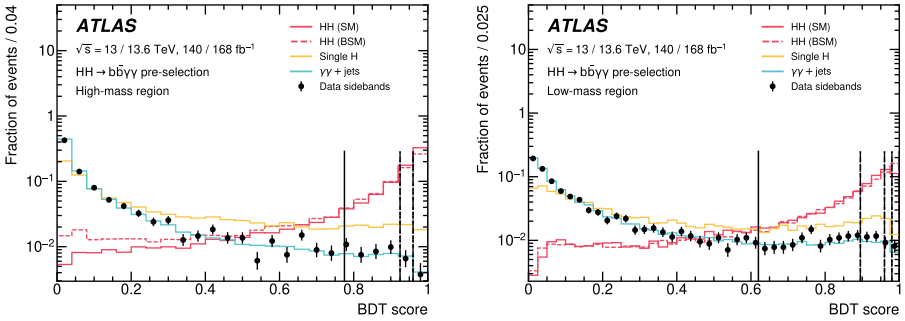

(a) (b)

Figure 7.8: BDT score distributions for (a) the HM region and (b) the LM region. The vertical lines indicate the category boundaries. HH (SM) and (BSM) correspond to ggF HH and VBF HH with the coupling values used in the training for SM and BSM described above.

The resulting BDT score distributions are used to define the analysis categories. Three categories are defined in the HM region and four in the LM region. The category boundaries are chosen to maximise the expected counting significance Z_{tot}. This is defined as:

$$Z_{bin} = \sqrt{2 \cdot [(S+B) \cdot \log(1+S/B) - S]}$$
$$Z_{tot} = \sqrt{\sum_{bin} Z_{bin}^2} \tag{7.16}$$

where S in the HM region is the expected yield of the SM ggF HH and SM VBF HH and S in the LM region is the expected yield of ggF HH ($\kappa_\lambda = 5$) and VBF HH ($\kappa_\lambda = 10$). B is the expected yield of the single-Higgs and continuum backgrounds. An additional requirement calls for at least 11 $\gamma\gamma$ + jets background events in the range $m_{\gamma\gamma}\ \epsilon[115, 160\,\text{GeV}]$ in each category. This ensures that the probability of having less than two events in the data sideband of any category is less than 0.1%, which is necessary to properly model the continuum background. The BDT score distributions with the category boundaries indicated are shown in Figure 7.8. The events with scores lower than the loosest category cut-off in each region are considered background and discarded. The categories are summarised in Table 7.6

Table 7.6: Thresholds in the $m^*_{b\bar{b}\gamma\gamma}$ variable and in the BDT score that define signal region categories.

Category	**Mass region**	**BDT score cuts (Run 2 + 3)**
High Mass 1	$m^*_{b\bar{b}\gamma\gamma} > 350\,\text{GeV}$	$0.775 \leq \text{BDT score} < 0.925$
High Mass 2	$m^*_{b\bar{b}\gamma\gamma} > 350\,\text{GeV}$	$0.925 \leq \text{BDT score} < 0.960$
High Mass 3	$m^*_{b\bar{b}\gamma\gamma} > 350\,\text{GeV}$	$\text{BDT score} \geq 0.960$
Low Mass 1	$m^*_{b\bar{b}\gamma\gamma} \leq 350\,\text{GeV}$	$0.620 \leq \text{BDT score} < 0.895$
Low Mass 2	$m^*_{b\bar{b}\gamma\gamma} \leq 350\,\text{GeV}$	$0.895 \leq \text{BDT score} < 0.960$
Low Mass 3	$m^*_{b\bar{b}\gamma\gamma} \leq 350\,\text{GeV}$	$0.960 \leq \text{BDT score} < 0.980$
Low Mass 4	$m^*_{b\bar{b}\gamma\gamma} \leq 350\,\text{GeV}$	$\text{BDT score} \geq 0.980$

7.9 Signal and Background Modelling

Due to its excellent resolution, the diphoton invariant mass $m_{\gamma\gamma}$ is able to discriminate well between the signal and the background. $m_{\gamma\gamma}$ is therefore used as the discriminating variable in the final fit. Both the di-Higgs signal and the much larger single-Higgs background appear as a narrow peak around the Higgs boson mass on top of the smoothly falling continuum background. For the fit to perform well, the $m_{\gamma\gamma}$ shape in signal and both background categories must be modelled accurately. The signal and background models are described in the following sections.

7.9.1 Signal and single-Higgs background

A Double-Sided Crystal Ball (DSCB) function is used to model the $m_{\gamma\gamma}$ distributions of both the di-Higgs signal and the single-Higgs background. The DSCB function, commonly used to describe the $H \to \gamma\gamma$ process, consists of a Gaussian core with two asymmetric power-law tails, and is completely determined by six parameters: μ_{CB} and σ_{CB} are the mean and width of the Gaussian core and $\alpha^{\text{Lo}}_{\text{CB}}$, $n^{\text{Lo}}_{\text{CB}}$, $\alpha^{\text{Hi}}_{\text{CB}}$, and $n^{\text{Hi}}_{\text{CB}}$ are shape parameters of the lower and higher tails. While well-reconstructed $H \to \gamma\gamma$ events are well-described by

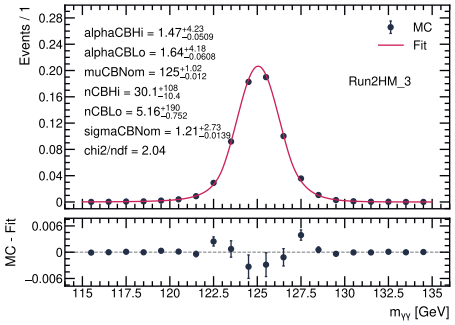

Figure 7.9: Fit of the DSCB function to the $m_{\gamma\gamma}$ distribution of the SM ggF + VBF HH MC sample in the Run 2 HM3 category.

the Gaussian core, the tails account for contributions from converted photons ($\gamma \to e^+e^-$) where at least one electron loses energy through bremsstrahlung radiation [119].

The nominal shape parameters of the DSCB are determined from an unbinned fit to full simulated SM HH samples in the range $m_{\gamma\gamma} \in [115, 135]$ GeV, with ggF HH and VBF HH modes weighted by their cross sections and summed. Separate parameters are derived for each analysis category. Within a given category, the same parameter values are used for both the di-Higgs signal and the single-Higgs background, since both processes produce a resonance at m_H. The bias introduced by applying the HH model parameters to the single-Higgs background is evaluated with a signal injection test, yielding a bias of less than 1%.

In the final fit to data, the tail parameters are fixed to their nominal values, while the mean and width of the Gaussian core are allowed to float within their uncertainties. The distribution for the Run 2 HM3 category with the best-fit function and parameters is shown in Figure 7.9. The fits for all categories can be found in Appendix B. The more sensitive categories to the SM signal (those with tighter BDT score boundaries) tend to contain more energetic and central diphoton events. As reflected by the σ_{CB} parameter of the DSCB shape, these categories exhibit a narrower resonance peak.

7.9.2 Continuum Background

The continuum background in each category is modelled by an exponential functional form whose parameters are left to float in the final fit to data. The choice of functional form may introduce a bias into the final fit. Too many degrees of freedom can result in an absorption of signal, while too few can introduce artificial signal. A spurious signal test is performed to evaluate the bias introduced by the choice of functional form [120].

Spurious signal is defined as the signal measured when fitting a signal + background model to a background-only template. The spurious signal test estimates the amount of spurious signal produced by a given functional form. The test requires the construction of a background-only template on which to perform the fit. The $\gamma\gamma$ + jets MC sample is an option for the background-only template, but this sample is left with few events after the pre-selection cuts, particularly the requirement for at least two b-jets, rendering it insufficient for a background-only template, which requires a large number of events to avoid large statistical fluctuations that may lead to inflated spurious signal. Therefore, a dedicated high-statistics $\gamma\gamma$ + heavy flavour MC sample, filtered for b-jets in the final state, is used for the background-only template. Plots validating several distributions of the $\gamma\gamma$ + heavy flavour sample against $\gamma\gamma$ + jets can be found in Appendix B.2. These were found to be in agreement.

The spurious signal test is performed over combined Run 2 and Run 3 categories (see Section 7.10) by fitting the background-only template with a signal + background model. The signal contribution is a DSCB function and the background model is the function under consideration. Both an exponential and a power law function are considered in this test. The fit is performed for different Higgs mass values ranging from 123 GeV to 127 GeV in intervals of 0.5 GeV. The spurious signal N_{SP} is defined as the number of fitted signal events with the largest magnitude out of all Higgs masses in this window.

Both the exponential and power-law functions perform similarly in the spurious signal test. The exponential function is chosen for its lower number of parameters. The acceptance criterion for the spurious signal is $N_{\text{SP}}/\sigma_{\text{bkg}} < 20\%$, where σ_{bkg} is the statistical uncertainty on the number of fitted signal events when fitting the signal-plus-background model to the background template. This ratio is used because the absolute spurious signal is expected to scale with the background yield, which varies across categories. The spurious signal sat-

isfies this criterion for all categories, with a maximum of 18% observed in the loosest low-mass category. N_{SP} is added as an uncertainty on the signal yield in the final fit.

7.10 Run 2 and Run 3 Correlation Strategy

The BDT hyperparameters and categorisation methodology are largely inherited from the Run 2 legacy analysis and remain optimal for this analysis. However, unlike the Run 2 legacy analysis, this analysis combines two data-taking periods with different centre-of-mass energies. This raises the question of how to combine data from Run 2 and Run 3.

The most conservative strategy is to treat Run 2 and Run 3 separately, training separate BDTs and optimising the categories independently for each data-taking period. The difficulty is that the most sensitive categories, HM3 and LM4, have by design the highest signal-to-background ratios and therefore the fewest background events. When Run 2 and Run 3 are treated independently, neither dataset alone has sufficient $\gamma\gamma$ + jets events in these categories to satisfy the 11-event requirement at the BDT score thresholds that maximise sensitivity. The category boundaries must therefore be loosened, reducing the sensitivity of the analysis.

To recover this sensitivity, an alternative strategy is adopted in which Run 2 and Run 3 samples are combined in both the BDT training and the category optimisation, as shown in Figure 7.8. The 11 $\gamma\gamma$ + jets requirement is then applied to each Run 2+Run 3 category pair (e.g. Run 2 HM3 + Run 3 HM3 combined), which provides sufficient statistics to satisfy the requirement at tighter BDT score thresholds than either dataset could achieve individually. Once the boundaries are defined, they are applied separately to Run 2 and Run 3 events, yielding 14 total analysis categories: 7 (3 HM + 4 LM) for each data-taking period with identical boundaries.

Additionally, the continuum background shape is correlated between Run 2 and Run 3 in the final fit. This is motivated by the observed compatibility between Run 2 and Run 3 background shape parameters and allows the full combined dataset to constrain the background shape simultaneously. The background normalisation is left uncorrelated and free to float independently for

each data-taking period, since it is determined directly from the data sideband and depends on the integrated luminosity, which differs significantly between Run 2 and Run 3. The HH signal and single-Higgs background modelling is based on MC and is likewise not correlated.

Tests indicate compatibility between Run 2 and Run 3 continuum background shape parameters. As a further check, the difference seen in simulation between Run 2 and Run 3 background shape parameters was enhanced by a factor of ten and no significant bias on the fitted signal is observed. The expected sensitivity to μ_{HH} with the Run 2 and Run 3 correlation strategy is increased by 8.7% compared to the conservative approach.

7.11 Systematic Uncertainties

Although the sensitivity to the $HH \to b\bar{b}\gamma\gamma$ signal is statistically limited, it is important to account for all the relevant sources of systematic uncertainties. Systematic uncertainties affect both the normalisation and shape of the $m_{\gamma\gamma}$ distributions of the signal and backgrounds used in the final fit. They are incorporated into the statistical analysis as nuisance parameters (NPs) and are profiled in the fit to data described in Section 7.12. Systematic uncertainties that modify the overall yield of a process without altering the shape of its $m_{\gamma\gamma}$ distribution are referred to as normalisation or yield uncertainties. Those that alter the shape of the $m_{\gamma\gamma}$ distribution itself are referred to as shape uncertainties.

Uncertainties are calculated separately for the ggF and VBF HH signal processes and for each single-Higgs background process. Uncertainties arising from the same source are treated as correlated across processes. Sources of systematic uncertainties fall into two categories: experimental uncertainties related to the reconstruction and calibration of physics objects, and theoretical uncertainties related to the modelling of signal and background processes. The continuum background, whose shape and normalisation are determined entirely from the fit, is not subject to theoretical or experimental systematic uncertainties from simulation. Its associated uncertainty is the spurious signal described in the previous section. The evaluation of experimental and theoretical uncertainties is described in the following sections.

7.11.1 Experimental Uncertainties

Experimentally determined quantities such as the photon identification efficiency or jet energy resolution have an associated uncertainty due to the limited precision of the measurement. Sources of experimental uncertainties and their evaluation are described below.

Luminosity: An overall normalisation uncertainty on the integrated luminosity is applied to all resonant processes, whose expected yields are derived from MC scaled by the integrated luminosity. The continuum background normalisation is a free parameter determined from data and is therefore unaffected. The luminosity uncertainty is 0.83% for Run 2 data and 4.0% for the partial Run 3 dataset, and is treated as uncorrelated between data-taking periods.

Pile-up: Uncertainties on the average number of interactions per bunch crossing $\langle\mu\rangle$ are evaluated per bin from data.

Diphoton trigger efficiency: Uncertainties associated with the diphoton trigger efficiency in Run 2 and 2022 are evaluated using control samples from radiative Z boson decays. In the absence of the same estimation for the 2023 and 2024 dataset, a conservative approach of doubling the 2022 uncertainty is taken.

Photon identification and isolation efficiency: Uncertainties on the photon identification and isolation efficiency in Run 2 are evaluated using control samples of prompt photons from photon+jet production, radiative $Z \to \ell\ell\gamma$ decays and electrons from $Z \to ee$ decays, in which electrons are used as a proxy for photons [57]. Corresponding uncertainties in Run 3 are extrapolated from Run 2.

Photon energy scale and resolution: Uncertainties on the photon energy scale and resolution in Run 2 are evaluated using $Z \to ee$ decays and radiative Z boson decays [57]. The corresponding uncertainties in Run 3 are extrapolated from Run 2 [121].

Jet energy scale and resolution: Uncertainties on the jet energy scale and resolution are evaluated using control samples of jets that recoil against well-calibrated reference objects, such as photons [60]. An additional uncertainty is introduced to account for the different responses of jets originating from b-quarks, light quarks and gluons.

Flavour tagging efficiency: Uncertainties in flavour tagging efficiencies and misidentification rates are evaluated for b- and c- jets using $t\bar{t}$ events and for light and gluon jets using Z+jets events [122–124].

JVT efficiency: Jet Vertex Tagger (JVT) efficiency uncertainties are evaluated using control samples of jets that recoil against well-calibrated objects

[61].
E_T^{miss} soft term resolution: The component of the missing transverse energy that is not associated with any well-reconstructed objects is called the soft term. The uncertainty on its resolution is derived using $Z \to \ell\ell$ events as a reference from which to measure the recoil.

Electron and muon scale and resolution uncertainties have a contribution due to the lepton veto applied in the pre-selection. However their contributions are extremely small, hence they are not discussed here.

Yield uncertainties

Each experimental systematic uncertainty listed above affects the expected yields of the signal and single-Higgs background processes. The effect is evaluated using additional MC samples in which the relevant quantity is varied up and down by its uncertainty. The variation in the expected yield of the i-th process in the c-th category is given by:

$$\delta N_c^i(\pm 1\sigma) = \frac{N_c^i(\pm 1\sigma)}{N_c^i(\text{nom})} - 1 \tag{7.17}$$

where $N_c^i(\text{nom})$ is the nominal expected yield and $N_c^i(\pm 1\sigma)$ is the expected yield of the process with the up or down variation. Each systematic uncertainty has a corresponding statistical uncertainty which is determined by the bootstrap method [125].

To reduce the number of nuisance parameters in the final fit, uncertainties may be either grouped together by source or neglected in certain cases: if both the up and down variations $\delta N_c^i(\pm 1\sigma)$ are less than 0.1%, they are assumed negligible and their values are set to zero. If the statistical uncertainty of $\delta N_c^i(\pm 1\sigma)$ is larger than the value of $\delta N_c^i(\pm 1\sigma)$ itself, the uncertainty is considered to be dominated by statistical fluctuations and is therefore neglected. In pathological cases in which the up and down variations have the same sign, the largest one is taken and the uncertainty is symmetrised.

Shape uncertainties

The photon energy scale and resolution affect not only the expected yield but also the $m_{\gamma\gamma}$ shape of the resonant processes. In particular, the photon energy scale affects the DSCB peak position μ_{CB}. The uncertainty on the photon energy scale for the i-th process in the c-th category is therefore propagated into a systematic uncertainty on μ_{CB} as follows:

$$\delta\mu_{\text{CB}}(\pm 1\sigma) = \frac{\langle m_{\gamma\gamma}\rangle_c^i(\pm 1\sigma)}{\langle m_{\gamma\gamma}\rangle_c^i(\text{nom})} - 1 \tag{7.18}$$

where $\langle m_{\gamma\gamma}\rangle_c^i(\text{nom})$ and $\langle m_{\gamma\gamma}\rangle_c^i(\pm 1\sigma)$ are the means of the nominal and varied $m_{\gamma\gamma}$ distributions respectively.

Similarly, the uncertainty on the photon energy resolution affects the DSCB peak width σ_{CB}. The uncertainty on the peak width is calculated as:

$$\delta\sigma_{\text{CB}}(\pm 1\sigma) = \frac{\text{IQR}_c^i(\pm 1\sigma)}{\text{IQR}_c^i(\text{nom})} - 1 \tag{7.19}$$

where $\text{IQR}_c^i(\text{nom})$ and $\text{IQR}_c^i(\pm 1\sigma)$ are the inter-quartile range of the nominal and varied $m_{\gamma\gamma}$ distributions respectively.

$\delta\mu_{\text{CB}}(\pm 1\sigma)$ and $\delta\sigma_{\text{CB}}(\pm 1\sigma)$ are calculated for the signal by merging the contributions of the ggF HH and VBF HH SM samples. Full simulation is used for the uncertainty extraction. The same uncertainties are applied to the single-Higgs background, since the same model is used.

Parameterisation uncertainties

An additional source of experimental uncertainty in the ggF HH and VBF HH expected yields arises from the disagreement between the yields obtained through the signal parameterisation procedure described in Section 7.4 and the yields obtained directly from simulated samples. The uncertainty is taken as the maximum deviation over all samples. These are on the order of 10% in the high-mass categories and notably higher in the low-mass categories, particularly in VBF HH (reaching up to 30%). Since VBF HH is a subdominant contribution to the signal, the sensitivity is not significantly affected.

7.11.2 Theoretical uncertainties

Theoretical uncertainties arise from several sources, including the limited precision of theoretical calculations and from the choice of parton distribution functions (PDFs) and Monte Carlo generators. Uncertainties due to missing higher order terms in perturbative QCD calculations are known as scale uncertainties. These affect the yields of the resonant processes and are evaluated by varying the renormalisation and factorisation scales μ_R and μ_F. Systematic uncertainties due to the choice of PDF and the value of the strong coupling constant α_S also affect the expected yields of resonant processes. They are combined into a single uncertainty (PDF + α_S). Alternative weights are stored in MC samples that encode variations in μ_R and μ_F, the PDF set and α_S. Uncertainties are calculated separately for each resonant process and for Run 2 and Run 3.

To calculate the scale uncertainty, the expected yield of the i-th (signal and resonant background) process in the c-th category is calculated for six variations of μ_R and μ_F. The scale uncertainty is then given by the envelope of the relative difference between the varied yields and the nominal yield.

PDF uncertainties are also evaluated using alternative MC weights. The ggF HH sample uses the PDF4LHC21 set, which includes 41 weights corresponding to the nominal PDF choice and 40 variations. The VBF HH sample is simulated using the NNPDF30NLO set. This PDF set provides 100 alternative weights corresponding to 100 variations in the PDF with respect to the nominal. The PDF uncertainty is taken as the square sum of the relative differences between the varied and nominal yields. Single-Higgs samples in Run 2 are generated using the PDF4LHC15 set, while Run 3 samples are generated with the PDF4LHC21 set, similarly to that of ggF HH. The former includes 30 variations, while the latter includes 40. The PDF uncertainty is calculated as described above for the signal.

Two variations in the value of α_S are provided in all resonant samples. The α_S uncertainty is taken as half the difference between the varied expected yields. The combined PDF+α_S uncertainty is given by the square sum of the PDF and α_S uncertainties.

Scale and PDF+α_S uncertainties also affect the Higgs boson pair production cross sections. Uncertainties in the branching fraction of the $H \to \gamma\gamma$ pro-

cess affect both signal and single-Higgs background, while uncertainties in the branching fraction of the $H \to b\bar{b}$ affect only the signal. These uncertainties are taken from the LHC Higgs Working Group [24].

Uncertainties due to the choice of parton shower generator are also taken into account. These are evaluated by comparing the expected yields of nominal HH and single-Higgs samples with alternative yields obtained from samples in which the parton showering is performed by a different generator. The parton shower uncertainty is taken as the relative difference between the alternative and nominal samples. For ggF HH, the envelope of the uncertainties over three basis samples (SM, $\kappa_\lambda = 0$, $\kappa_\lambda = 5$) is taken as the final uncertainty for the process. Alternative samples are only available for Run 3, however the resulting uncertainties are similar to the values seen in the Run 2 legacy analysis. The Run 3 uncertainties are therefore also applied to Run 2. Parton shower uncertainties for ggF HH are applied to VBF HH, since it was also seen in the legacy analysis that these were similar. All alternative samples are available for single-Higgs processes in both Run 2 and Run 3, so these are calculated individually as described above.

Due to known mismodelling in heavy-flavour production in association with $t\bar{t}$ [126] and W boson production in association with b-jets [127], a conservative uncertainty of 100% is applied to the ggH, VBFH and WH background processes.

Given the Run 2 and Run 3 correlation strategy described in Section 7.10, care must be taken in the treatment of systematic uncertainties. Uncertainties that are expected to be the same between Run 2 and Run 3 are correlated. These include all theoretical uncertainties and the spurious signal uncertainty (Section 7.9.2). Jet-related uncertainties are correlated with the exception of one component of the jet energy scale uncertainty. Luminosity, trigger, flavour-tagging and photon-related uncertainties are decorrelated between data-taking periods.

Uncertainties are propagated into nuisance parameters to be included in the likelihood function of the final fit via response functions. The response function of a given uncertainty corresponds to how the uncertainty affects the yield or shape. If the effect of the uncertainty is additive and symmetric (for example, the effect of the photon energy resolution on the signal peak width), a Gaussian function is used. For uncertainties that multiplicatively scale the

yield (for example, the luminosity uncertainty), a lognormal response function is used. The majority of uncertainties have an asymmetric effect and are propagated into nuisance parameters via an asymmetric function that allows for different sized up and down variations.

The relative impact of the main uncertainty sources on the HH signal strength is shown in Table 7.7.

As shown in Table 7.7, theoretical uncertainties are the dominant source of systematic uncertainty on μ_{HH}, with the QCD scale and heavy-flavour content uncertainties each exceeding 30%. Among experimental uncertainties, photon-related systematics – energy scale, resolution and efficiency – are the largest, as expected for an analysis that relies on the $m_{\gamma\gamma}$ resonance for signal extraction. The spurious signal uncertainty of $\pm 7.2\%$ is small, validating the choice of exponential functional form for the continuum background model. Nevertheless, as noted above, the analysis remains predominantly statistically limited: even the largest systematic uncertainties are small relative to the overall statistical uncertainty on μ_{HH}.

Source of systematic uncertainty	**Relative uncertainty [%]**	
	Up	**Down**
Experimental		
Photon energy scale	+22	-34
Photon energy resolution	+14	-7.5
Photon efficiency	+13	-2.5
Jet	+10	-6.9
Luminosity	+6.7	-1.2
Theory		
QCD scale + m_{top}, PDF+α_s	+35	-5.2
$\mathcal{B}(H \to \gamma\gamma, b\bar{b})$	+10	-2.1
Parton showering model	± 15	
Heavy-flavour content	± 30	
Background model		
Spurious signal	± 7.2	

Table 7.7: Relative effect of systematic uncertainties on the observed HH signal strength μ_{HH}. The (upper or lower) error on μ_{HH} is determined by a fit with only the respective nuisance parameter floating while all others are fixed to their best-fit values. The impact is calculated as the sum in quadrature of this error and the error when all NPs are fixed.

7.12 Statistical Analysis

The results of the search for $HH \to b\bar{b}\gamma\gamma$ are obtained through an unbinned maximum likelihood fit of the $m_{\gamma\gamma}$ distribution in the range $m_{\gamma\gamma} \in [105, 160]$ GeV performed simultaneously over all categories. The likelihood function is defined as:

$$\mathcal{L} = \prod_{c} \left(\text{Pois}(n_c | N_c(\boldsymbol{\theta})) \cdot \prod_{i=1}^{n_c} f_c(m_{\gamma\gamma}^i, \boldsymbol{\theta}) \cdot G(\boldsymbol{\theta}) \right), \tag{7.20}$$

where, for each category c, n_c and N_c are the observed and expected number of events respectively. For each event i in category c, f_c is the value of the probability density function of the $m_{\gamma\gamma}$ distribution, $\boldsymbol{\theta}$ is the vector of nuisance parameters (NPs) corresponding to systematic uncertainties and $G(\boldsymbol{\theta})$ is a constraint function for the NPs.

The observed number of events in a given category is modelled by a Poisson distribution with a mean equal to the expected number of events $N_c(\boldsymbol{\theta})$. The total expected number of events is a sum of the expected yields of the signal (N_c^{HH}) and single-Higgs background (N_c^{H}) processes (taken from MC), the expected number of continuum background events ($N_c^{\gamma\gamma}$) and the spurious signal contribution (N_c^{SS}):

$$N_c(\boldsymbol{\theta}) = \mu_{HH} \cdot N_{\text{c}}^{HH}(\boldsymbol{\theta}_{\text{yield}}^{HH}, \kappa_\lambda, \kappa_{2V}) + N_{\text{c}}^{\text{H}}(\boldsymbol{\theta}_{\text{yield}}^{\text{H}}) + N_{\text{c}}^{\text{SS}} \cdot \boldsymbol{\theta}_{\text{c}}^{\text{SS}} + N_{\text{c}}^{\gamma\gamma} \tag{7.21}$$

As described in Section 7.11, the expected yields of the signal and single-Higgs background are affected by systematic uncertainties. These are included as NPs $\boldsymbol{\theta}$, on which the yield terms in the likelihood depend for each respective process. The expected signal yield is also multiplied by the signal strength μ_{HH}, and depends on the coupling modifiers κ_λ and κ_{2V} which modify the expected signal yield as described in Section 7.4. μ_{HH}, κ_λ and κ_{2V} are the parameters of interest (POIs), which may either be fixed or floating in the fit, depending on which statistical interpretation is being performed. The expected yield of the continuum background is an unconstrained parameter in the fit and does not depend on any NPs. The spurious signal yield is multiplied with an associated NP $\boldsymbol{\theta_c^{SS}}$.

The probability density function f_c describes the shape of the signal + background model in the $m_{\gamma\gamma}$ spectrum. It is the sum of the DSCB functions that

model the signal and single-Higgs background and the exponential function that models the continuum background. The shape for each process is normalised by the corresponding expected event yields. The $m_{\gamma\gamma}$ shapes of the resonant processes are affected by the photon energy scale and resolution as described in Section 7.11.1. Hence, f_c depends on the corresponding nuisance parameters $\boldsymbol{\theta}$.

The final term in Equation 7.20 is a Gaussian constraint term $G(\boldsymbol{\theta})$ for the NPs. For each NP θ, the likelihood is multiplied by a Gaussian probability density function $G(a = 0 \mid \theta, 1)$, centred on θ with a width of 1, and a represents a global observable from an auxiliary measurement whose nominal value is 0. Hence, the constraint functions to penalise the fit for values of θ that deviate from the global observable a.

As discussed in Section 7.10, continuum background shape parameters are correlated between the respective Run 2 and Run 3 categories in the fit, but are otherwise unconstrained. The continuum background normalisation is completely unconstrained, since it is determined directly from the data sidebands.

The profile likelihood ratio [128] is used to measure the parameters of interest. This is defined as

$$\Lambda(\mu) = \frac{\mathcal{L}(\mu, \hat{\hat{\boldsymbol{\theta}}}(\mu))}{\mathcal{L}(\hat{\mu}, \hat{\boldsymbol{\theta}})} \tag{7.22}$$

where μ is the POI being measured (μ_{HH}, κ_λ or κ_{2V}). The denominator is the unconditionally maximised likelihood evaluated with the best-fit POI and NPs, $\hat{\mu}$ and $\hat{\boldsymbol{\theta}}$. The numerator is the conditionally maximised likelihood for a specific value μ under consideration, where $\hat{\hat{\boldsymbol{\theta}}}$ is the value of the NPs that maximises the likelihood for μ.

The unconditionally maximised likelihood gives the best-fit values of μ_{HH}, κ_λ or κ_{2V}. In the absence of signal, upper limits on the signal strength μ_{HH} are set at 95% CL using the CL_s method [129] and the profile likelihood ratio

test statistic $\tilde{q}_\mu$, defined as [128]:

$$\tilde{q}_\mu = \begin{cases} -2\ln\frac{\mathcal{L}(\mu,\hat{\hat{\boldsymbol{\theta}}}(\mu))}{\mathcal{L}(0,\hat{\hat{\boldsymbol{\theta}}}(0))} & \hat{\mu} < 0, \\ -2\ln\frac{\mathcal{L}(\mu,\hat{\hat{\boldsymbol{\theta}}}(\mu))}{\mathcal{L}(\hat{\mu},\hat{\boldsymbol{\theta}})} & 0 \leq \hat{\mu} \leq \mu, \\ 0 & \hat{\mu} > \mu. \end{cases} \tag{7.23}$$

By construction, large values of $\tilde{q}_\mu$ represent incompatibility of the data with the given signal strength μ. The probability distribution of $\tilde{q}_\mu$ can be derived analytically using the asymptotic approximation [128], which is required for the calculation of the CL_s. Since the asymptotic approximation may fail when the event count is low, a validation is performed using the frequentist approach. A series of toy experiments is run to obtain the distribution of $\tilde{q}_\mu$ and it was found to be compatible with the asymptotic approximation.

From the distribution of $\tilde{q}_\mu$, using the CL_s method, the upper limit on the signal strength μ_{HH} and the cross section σ_{HH} at the 95% CL can be calculated. Similarly, via a scan over the coupling modifiers κ_λ and κ_{2V}, confidence intervals can be obtained for these parameters.

7.13 Results

7.13.1 $HH \to b\bar{b}\gamma\gamma$ main results

Both expected and observed results are obtained via the statistical procedure described above. While observed results are extracted from a fit to data, expected results are obtained by performing the fit on an Asimov dataset [128] generated with the respective signal hypothesis. To determine the expected best-fit signal strength and significance, as well as the expected limits on the coupling modifiers κ_λ and κ_{2V}, the presence of SM HH signal is assumed. An Asimov dataset is generated with $\mu_{HH} = 1$ and the corresponding post-fit NPs and background shape parameters obtained from the conditionally maximised likelihood for $\mu_{HH} = 1$.

To calculate the expected upper limit on the signal strength in the absence of signal, an Asimov dataset is generated with $\mu_{HH} = 0$ and the corresponding conditionally maximised NP and background shape parameter values. It should be noted that this background-only hypothesis does not correspond to any physical coupling configuration: even setting $\kappa_\lambda = 0$ would still yield HH

production via the box diagram proportional to κ_t^4. However, it serves as the standard benchmark for expected limit computation and comparability with previous results. For the Asimov fits, the global observable a in the Gaussian constraint is also set to its post-fit value.

The observed $m_{\gamma\gamma}$ distributions in each category are shown alongside the signal + background, continuum background and total background in Figures 7.10 and 7.11. The continuum background includes the spurious signal. The total background is the sum of the continuum and single-Higgs backgrounds. The corresponding fits with Run 2 and Run 3 categories combined are shown in Appendix B.3.

The observed best-fit value for the signal strength is $\mu_{HH} = 0.9^{+1.3}_{-1.0}(\text{stat}) \pm 0.5(\text{sys})$, where the first set of uncertainties are statistical and the second systematic. The expected signal strength is $\mu_{HH} = 1^{+1.3}_{-1.0}$. The corresponding statistical significance for Higgs boson pair production is $0.78\,\sigma$, with an expected significance of $1.01\,\sigma$. The expected significance represents a sensitivity milestone for Higgs boson pair production searches, surpassing $1\,\sigma$ for the first time in any channel.

An observed upper limit on the signal strength at 95% CL assuming SM HH production is set at $\mu_{HH} < 3.7$, matching the expected limit under the same hypothesis. The expected upper limit assuming no HH production is $\mu_{HH} < 2.6$. These results are also derived for Run 2 and Run 3 separately. The significance and upper limit for Run 2 and Run 3 separately and combined are shown in Figure 7.12. These results are consistent with a slight deficit in Run 2 and an excess in Run 3. The combined result shows a small positive excess above the background-only expectation, though the observed significance of $0.78\,\sigma$ is slightly below the expected $1.01\,\sigma$ assuming SM HH production.

The expected and observed impacts of the first 20 systematic uncertainties on μ_{HH} are shown in Figure 7.13. Overall, the expected rankings are consistent with the observed and no systematics are significantly pulled or constrained by the data with the exception of the Run 3 photon energy scale. The larger impact of this systematic on the observed signal strength may be partially explained by the lower peak position of the data in the Run 3 HM3 category, as well as the larger uncertainties inherent in the extrapolation of Run 2 photon energy scale calibrations to Run 3 (Section 7.11).

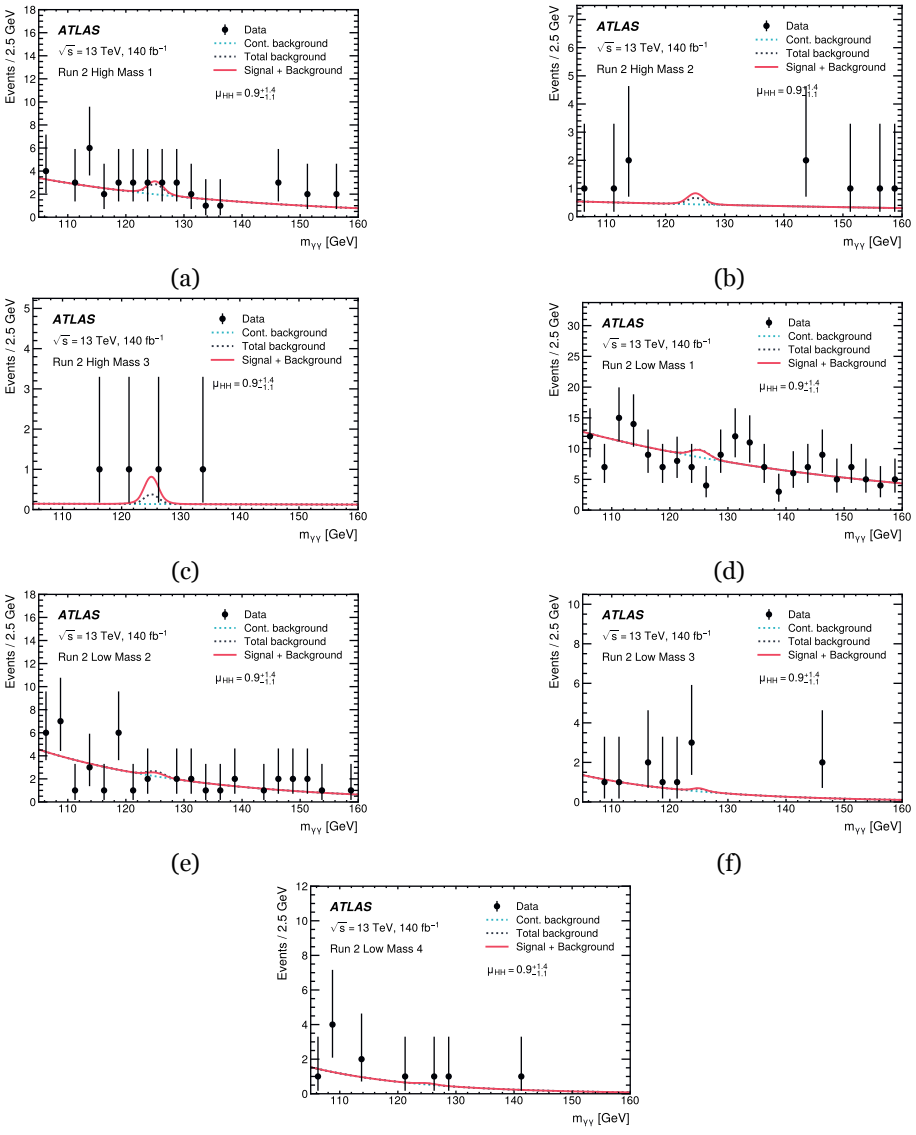

Figure 7.10: Observed $m_{\gamma\gamma}$ distributions in the Run 2 categories HM1 (a), HM2 (b), HM3 (c), LM1 (d), LM2 (e), LM3 (f), LM4 (g). The signal + background fit (red), single-Higgs + resonant background fit (black dotted) and the continuum background only fit (blue dotted) are also shown.

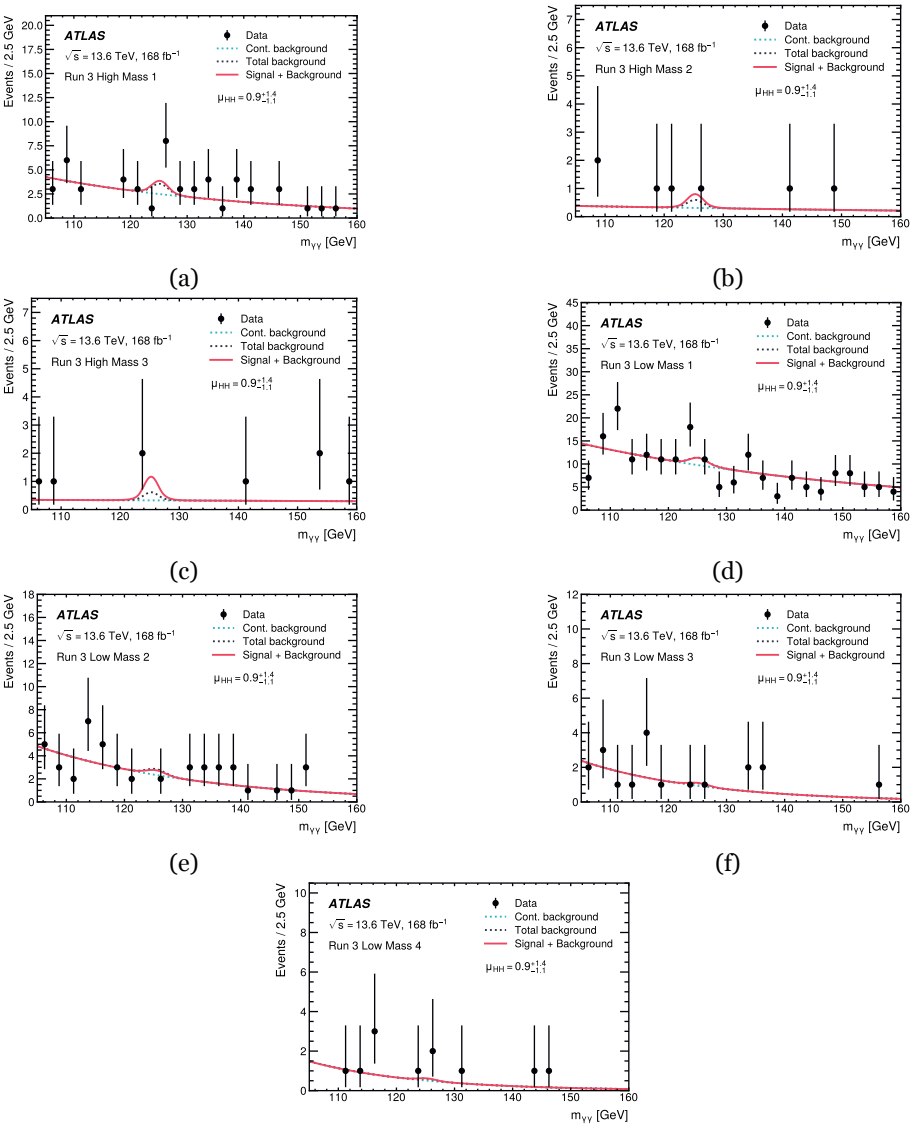

Figure 7.11: Observed $m_{\gamma\gamma}$ distributions in the Run 3 categories HM1 (a), HM2 (b), HM3 (c), LM1 (d), LM2 (e), LM3 (f), LM4 (g). The signal + background fit (red), single-Higgs + resonant background fit (black dotted) and the continuum background only fit (blue dotted) are also shown.

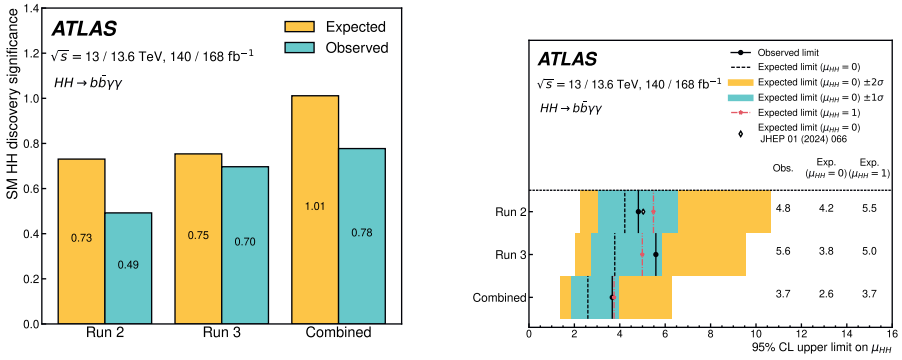

(a) (b)

Figure 7.12: Observed and expected significance (a) and upper limit on the signal strength μ_{HH} (b) for Run 2, Run 3 and combined. Expected limits for both the SM HH production and background-only hypotheses are shown.

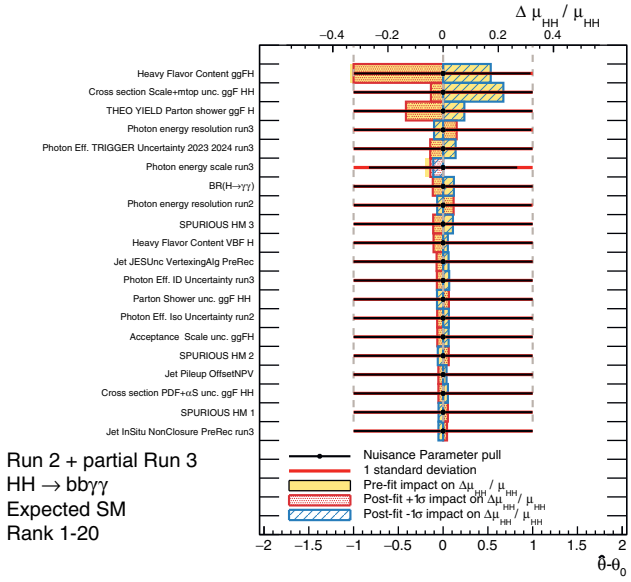

(a)

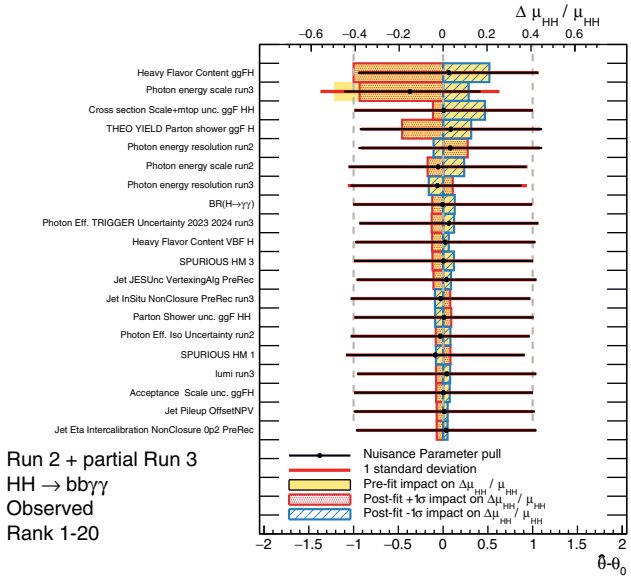

(b)

Figure 7.13: Expected (a) and observed (b) pre-fit and post-fit impact and pull of the first 20 systematic uncertainties on the μ_{HH} measurement.

Table 7.8 shows the impact of each group of systematic uncertainties on the expected ($\mu_{HH} = 0$) and observed upper limits on the signal strength. These are calculated by floating only the respective group of systematics and comparing the result to the stat-only case. The dominant uncertainty in both cases is the theoretical uncertainty on the HH cross section, contributing $\pm 5.2\%$ expected and $\pm 6.1\%$ observed; since μ_{HH} is defined relative to the predicted SM cross section, this uncertainty propagates directly onto the limit. The heavy-flavour content uncertainty, arising from the conservative 100% uncertainty assigned to the single-Higgs backgrounds (Section 7.11), is the next largest expected contribution but is much smaller on the observed limit ($\pm 0.98\%$ compared to $\pm 3.3\%$), since the observed fit constrains the resonant yield while the expected case leaves the corresponding nuisance parameter at its full prior width. The photon energy scale behaves in the opposite way, being negligible in the expected case but the second largest contribution to the observed limit, consistent with the large pull of the Run 3 photon energy scale seen in Figure 7.13. In all cases the total systematic contribution remains small compared to the statistical uncertainty.

Table 7.8: The breakdown of systematic uncertainty on the expected ($\mu_{HH} = 0$) and observed μ_{HH} upper limits, ordered by decreasing impact on the expected limit.

Uncertainty source	Expected $\Delta\sigma$ [%]	Observed $\Delta\sigma$ [%]
QCD Scale + m_{top}, PDF+α_S	± 5.2	± 6.1
Heavy-flavour content	± 3.3	± 0.98
Parton showering model	± 0.99	± 0.26
Jet	± 0.33	$< \pm 0.1$
Branching ratio	± 0.22	± 0.33
Photon Efficiency	± 0.22	± 0.27
Background model (spurious signal)	± 0.21	$< \pm 0.1$
Photon Energy Scale (PES)	$< \pm 0.1$	± 3.6
Photon Energy Resolution (PER)	$< \pm 0.1$	± 0.43
E_T^{miss}	$< \pm 0.1$	$< \pm 0.1$
Muon	$< \pm 0.1$	$< \pm 0.1$
Flavor Tagging	$< \pm 0.1$	$< \pm 0.1$
Pileup	$< \pm 0.1$	$< \pm 0.1$
Luminosity	$< \pm 0.1$	$< \pm 0.1$
All systematics	± 12	± 13

Finally, both 1D and 2D likelihood scans are performed to obtain confidence intervals for the coupling modifiers κ_λ and κ_{2V}. The 1D scans are performed by fixing one POI to its SM value and floating the other, while the 2D scan is performed by floating both simultaneously. The 1D likelihood scans and the resulting confidence intervals are shown in Figure 7.14. The observed and expected confidence intervals for κ_λ at 95% CL are $[-1.6, 6.6]$ and $[-1.8, 6.9]$ respectively. The observed and expected intervals for κ_{2V} are $[-0.5, 2.6]$ and $[-0.4, 2.6]$ respectively. The 2D likelihood scan is shown in Figure 7.15.

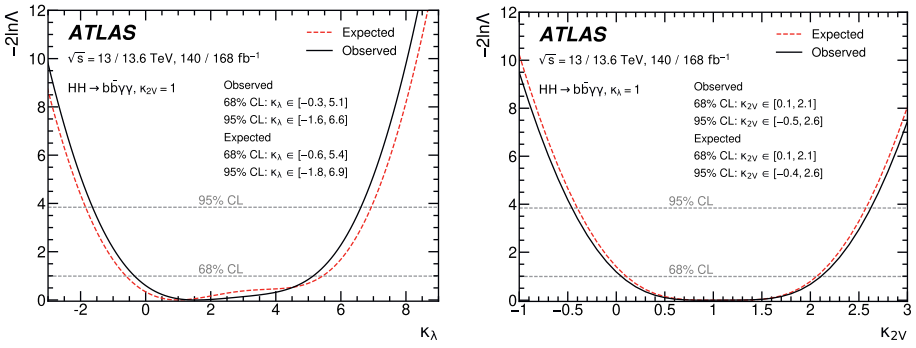

(a) (b)

Figure 7.14: Expected and observed 1D profiled likelihood scans of κ_λ (a) and κ_{2V} (b).

Compared to the Run 2 legacy analysis, the sensitivity to the Higgs boson pair production signal strength μ_{HH} was improved by approximately 85%. Approximately 60% can be attributed to additional data (of which 50% comes from the partial Run 3 dataset and 10% comes from the optimisation of analysis categories). The use of the new b-tagging algorithm and looser working point improves the sensitivity by approximately 20% and the kinematic fit adds an improvement of 5%. The expected width of the 95% confidence intervals for κ_λ and κ_{2V} improved by approximately 20% and 30% respectively with respect to the Run 2 legacy analysis.

7.13.2 Kinematic Fit Sensitivity Study

A simplified version of the analysis described in this chapter was run in order to estimate the impact of the Kinematic Fit on both μ_{HH} and κ_λ sensitiv-

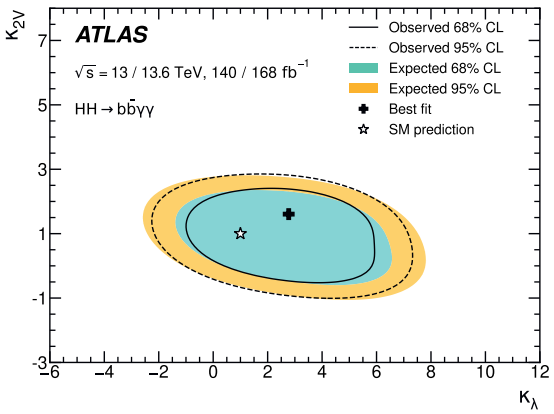

Figure 7.15: Expected and observed 2D profiled likelihood scan of κ_λ vs κ_{2V}. The inner and outer contours represent the 68% and 95% CL intervals respectively.

ity. This study was performed at an early stage in the analysis development to help determine whether the inclusion of the kinematic fit was feasible and justified. The study uses older MC samples that were later reproduced and an older version of the kinematic fit with transfer functions initially developed for use in the legacy analysis. The samples, analysis strategy and the kinematic fit itself were subsequently updated and improved after the completion of this study. The true sensitivity gain due to the kinematic fit in the final, unblinded results, shown in the previous section, is therefore larger than the gain estimated here. Nevertheless, these results - with the expectation that they would improve along with the development of the analysis - were sufficient to justify the use of the kinematic fit.

This study uses the baseline event selection described in Section 7.7. The BDT training described in Section 7.8 was performed twice independently: once with kinematic fitted input variables and once with the nominal input variables as they were calculated before the application of the kinematic fit. The affected input variables are those that depend on jet four-momenta: $\text{jet}_{pT}^{1,2,3,4}$, $\text{jet}_{\eta}^{1,2,3,4}$, $\text{jet}_{\phi}^{1,2,3,4}$, pT_{bb}, η_{bb}, ϕ_{bb}, m_{bb}, HT, single topness, m_{HH}, sphericityT, planarFlow, pT balance. Hereafter, **KF** is used to refer to results obtained using the kinematic fitted variables in the BDT and **nominal** is used to refer to results that

use the same variables without the kinematic fit. Apart from the training variables, all analysis steps are unchanged and performed independently for the KF and nominal cases.

BDT scores for signal and background processes in the HM and LM regions for Run 2 and Run 3 are shown in Figures 7.16 and 7.17 for the KF and nominal case. Optimised boundaries from the event categorisation step are also shown on the plots. The effect of the kinematic fit on the category boundaries in the separate regions and data-taking periods is not necessarily consistent - in the Run 2 HM region, for example, the kinematic fit allows for the category boundaries to be tighter, but in the Run 2 LM region, this is not the case.

A more robust initial evaluation can be done by calculating the SM and BSM benchmark significance for Run 2 and Run 3 combined. These benchmarks are defined in Section 7.8. The benchmark significance of the KF and nominal cases is shown in Table 7.9. Based on these benchmarks, the kinematic fit improves the SM significance by 3.1% and the BSM significance by 1% with respect to the nominal case.

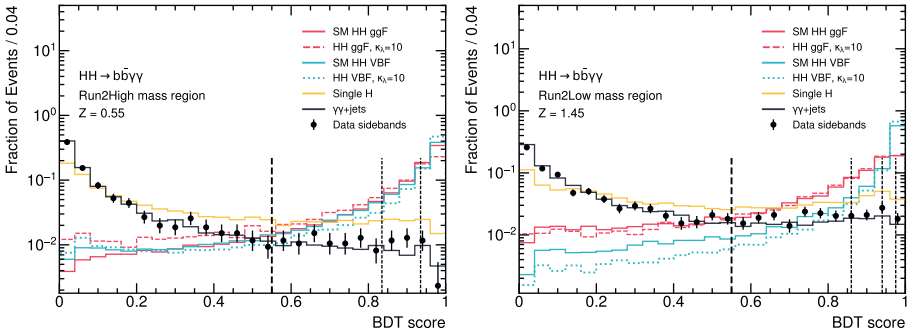

(a)

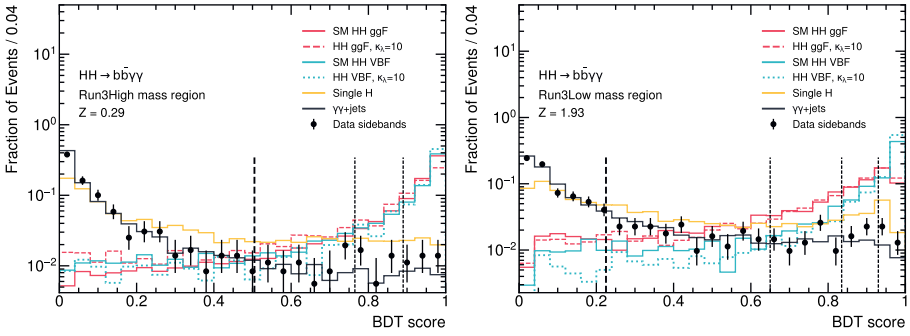

(b)

Figure 7.16: BDT scores and category boundaries for Run 2 (a) and Run 3 (b) in the high mass and low mass regions for the nominal case.

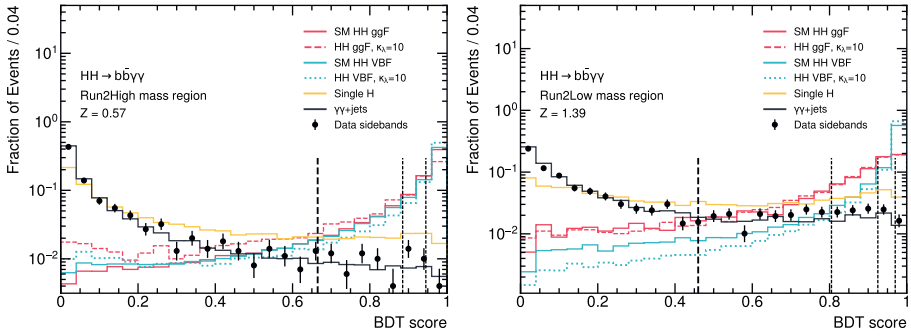

(a)

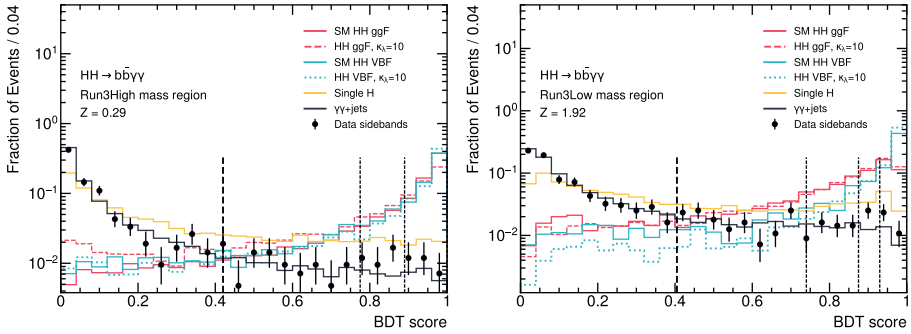

(b)

Figure 7.17: BDT scores and category boundaries for Run 2 (a) and Run 3 (b) in the high mass and low mass regions for the KF case.

	Z(SM)	Z(BSM)
Nominal Run 2	0.55	4.55
KF Run 2	0.57	4.60
Nominal Run 3	0.28	2.63
KF Run 3	0.29	2.65
Nominal Combined	0.62	5.25
KF Combined	0.64	5.30

Table 7.9: Z(SM) and Z(BSM) benchmark significance in Run 2, Run 3 and combined for the KF and nominal case.

To more precisely quantify the effect of the kinematic fit on the final results, estimates of the μ_{HH} upper limit, significance and κ_λ and κ_{2V} scans must be obtained using both KF and nominal quantities. A simplified version of the statistical analysis described in Section 7.12 was performed for both cases. To estimate the upper limit and significance on μ_{HH}, an unbinned maximum likelihood fit to the $m_{\gamma\gamma}$ distribution was performed without systematic uncertainties using Asimov data. The results of this fit are shown in Table 7.10 for Run 2 and Run 3 separately and combined for the KF and nominal cases. The Run 2 legacy results are also included for reference.

The kinematic fit improves the expected upper limit by 3.2% and increases the significance by 4.9% with respect to the nominal case. Likelihood scans for κ_λ and κ_{2V} are also performed for both cases. These are shown in Figure 7.18. Interestingly, the kinematic fit has almost no effect on these limits. Its influence is primarily on the SM signal sensitivity. This is not unexpected, as the transfer functions and momentum constraints that comprise the kinematic fit were developed exclusively on simulated SM signal.

These improvements are larger than the 3.1% estimated from the benchmark significance. This is expected, as the benchmark significance is an approximate counting-based figure of merit used to optimise the category boundaries, whereas the likelihood fit reflects the sensitivity of the analysis itself. The benchmark also underestimates the significance in absolute terms, giving 0.62 for the nominal combined case compared with 0.73 from the fit, and is

therefore best interpreted as a conservative proxy. In the final analysis, the kinematic fit contributes an improvement of approximately 5% to the expected μ_{HH} limit (Section 7.13), compared with the 3.2% estimated here, consistent with the updated samples, re-derived transfer functions and optimised analysis strategy.

	Expected μ_{HH} upper limit	Significance
Legacy Run 2	4.86	0.54
Nominal Run 2	4.04	0.67
KF Run 2	3.88	0.70
Nominal Run 3	9.14	0.30
KF Run 3	9.12	0.30
Nominal Combined	3.46	0.73
KF Combined	3.35	0.77

Table 7.10: Expected μ_{HH} upper limit and significance in Run 2, Run 3 and combined for the KF and nominal case.

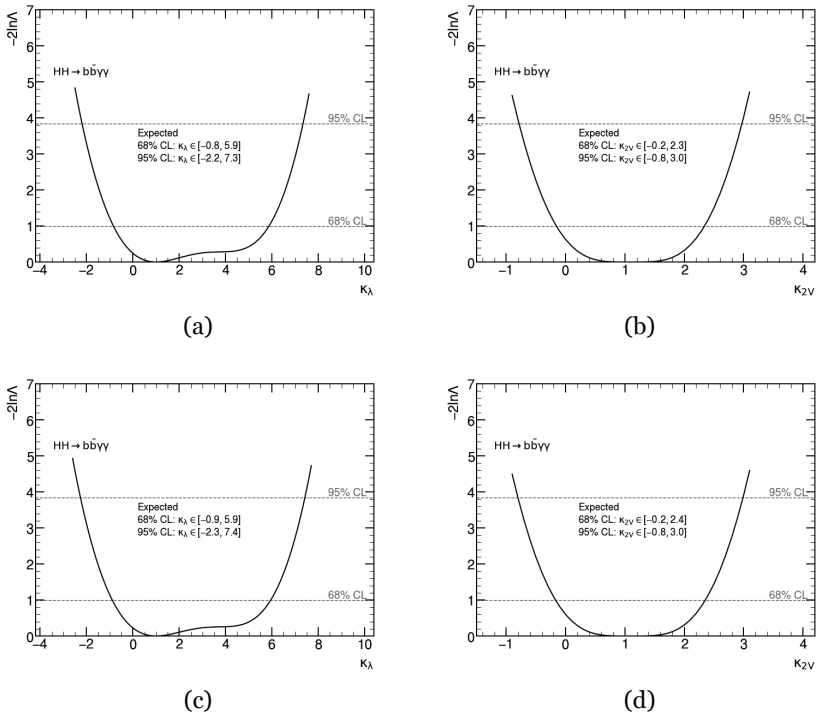

Figure 7.18: Expected 1D likelihood scans of κ_λ (a) and κ_{2V} (b) for the KF case and κ_λ (c) and κ_{2V} (d) for the nominal case.

7.14 Summary

This chapter has presented a search for Higgs boson pair production in the $b\bar{b}\gamma\gamma$ final state using the full Run 2 dataset together with a partial Run 3 dataset, corresponding to a total integrated luminosity of $308\,\text{fb}^{-1}$. Building on the Run 2 legacy analysis, the search incorporates the GN2 b-tagging algorithm at a looser working point, a re-optimised BDT-based event categorisation trained jointly across both data-taking periods, and a kinematic fit that improves the $m_{b\bar{b}}$ resolution by 27% through a per-event correction to the b-jet four-momenta.

The observed best-fit signal strength is $\mu_{HH} = 0.9^{+1.3}_{-1.0}\,(\text{stat}) \pm 0.5\,(\text{sys})$, cor-

responding to an observed significance of $0.78\,\sigma$ against an expected $1.01\,\sigma$ – the first time a single channel has exceeded $1\,\sigma$ expected sensitivity to SM Higgs boson pair production. An observed upper limit of $\mu_{HH} < 3.7$ is set at 95% CL, matching the expected limit under the same hypothesis. Observed 95% CL intervals of $[-1.6, 6.6]$ and $[-0.5, 2.6]$ are obtained for κ_λ and κ_{2V} respectively. The measurement remains statistically limited, with the total systematic contribution to the expected limit amounting to $\pm 12\%$.

As quantified in Section 7.13, these developments together improve the sensitivity to μ_{HH} by approximately 85% relative to the Run 2 legacy analysis, with roughly two-thirds of the gain deriving from the enlarged dataset and the remainder from improvements to the analysis techniques themselves. In surpassing $1\,\sigma$ expected sensitivity for the first time in any single channel, this result marks a milestone in the search for Higgs boson pair production and an important step towards its eventual observation.

Prospects for a future Higgs pair and triple-Higgs combination 8

The $HH \to b\bar{b}\gamma\gamma$ analysis described in Chapter 7 represents a sensitivity milestone for Higgs boson pair production and the Higgs boson trilinear self-coupling κ_λ. The High Luminosity LHC (HL-LHC) is projected to measure Standard Model Higgs boson pair production with a significance of $4.3\ \sigma$ in a baseline extrapolated scenario and to measure κ_λ with 15% precision [130]. Access to the quartic coupling κ_4, however, remains more distant. The first experimental probe of triple Higgs production in the ATLAS experiment [105] constrains κ_4 in the $HHH \to 6b$ channel to lie between -230 and 240 at 95% CL for a fixed $\kappa_\lambda = 1$. This bound is two orders of magnitude away from the Standard Model value, and therefore excludes only extreme deviations rather than testing the Standard Model prediction. Direct triple- Higgs production will remain far from the sensitivity required to probe $\kappa_4 = 1$ even with the full HL-LHC dataset.

The case for pursuing κ_4, despite how weakly it is currently constrained, rests on the fact that it cannot be estimated by measuring κ_λ, since the Higgs potential is uniquely dependent on both couplings. As discussed in Section 2.3.3, a higher-dimensional contribution to the Higgs potential shifts them by different amounts, generating a correction to the quartic coupling roughly six

times larger than to the trilinear. Constraining κ_4, even loosely, is therefore critical for capturing possible shifts in the potential that are not captured by measurements of κ_λ alone. This chapter discusses an alternative and complementary way to constrain κ_4 using indirect probes.

8.1 Indirect Probes of Coupling Parameters

For couplings that are only accessible at leading order through low cross section processes, but enter higher cross section processes through higher order loops, indirect probes can provide an alternative means of constraint and complement the sensitivity of the direct measurement. For example, Reference [131] proposes the use of single-Higgs production to constrain the trilinear self-coupling through next-to-leading order (NLO) electroweak loops. Examples of such diagrams are shown in Figure 8.1.

Anomalous coupling values impact the Higgs boson production cross section and kinematics through these electroweak effects. Sensitivity to κ_λ can therefore be increased by combining these κ_λ-dependent single-Higgs channels with selected Higgs boson pair production channels. Additionally, this method allows for other couplings – such as κ_t and κ_V – to be simultaneously constrained by the single-Higgs channel instead of fixing them to their SM value. This allows the assumption that new physics only enters through κ_λ to be relaxed and therefore reduces the model dependence of the constraint.

A recent application of this procedure within ATLAS is the combination of single-Higgs and Higgs boson pair production measurements with the full Run 2 dataset in Reference [132]. κ_λ-dependent NLO electroweak corrections are applied to the rates and kinematics of the single-Higgs measurements in the $\gamma\gamma$, ZZ^*, WW^*, $\tau\tau$ and $b\bar{b}$ decay channels, which are then combined with the direct constraints from the $HH \to 4b$, $HH \to bb\tau\tau$ and $HH \to b\bar{b}\gamma\gamma$ channels. This combination excludes values outside $-0.4 < \kappa_\lambda < 6.3$ at 95% CL assuming new physics will only affect κ_λ, while in a generic model with additional floating coupling modifiers the constraint is relaxed to $-1.4 < \kappa_\lambda < 6.1$. The inclusion of the single-Higgs channels improves the sensitivity to κ_λ compared to the Run 2 Higgs boson pair production combination in Reference [96].

The same logic can be applied to the measurement of the quartic coupling

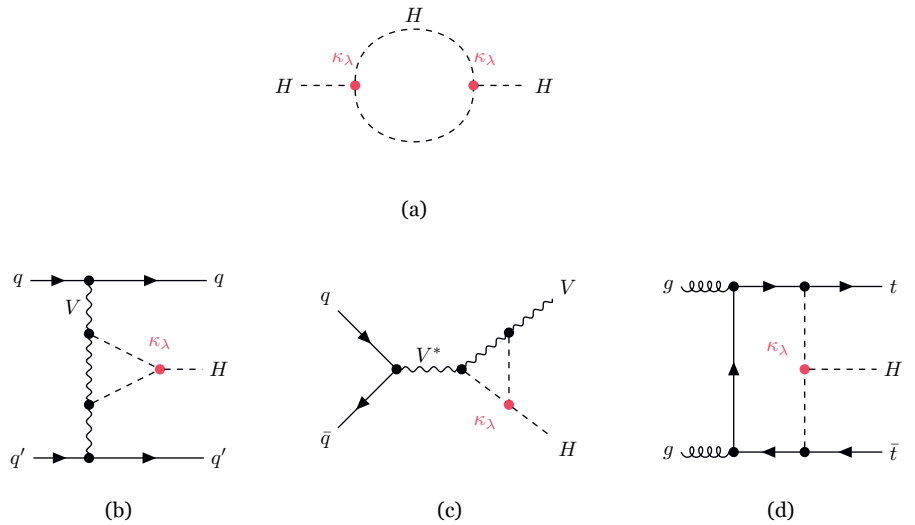

Figure 8.1: Example of one loop κ_λ-dependent diagrams for the Higgs boson self energy (a) and single-Higgs boson production in the VBF (b), VH (c), and $t\bar{t}H$ (d) modes. The self-coupling vertex is indicated in red.

κ_4, using Higgs boson pair production as an indirect probe and combining it with triple Higgs measurements. In Reference [133], calculations are performed of two-loop contributions to the amplitude and kinematics of the ggF HH process involving κ_4 dependent vertices. Examples of such processes are shown in Figure 8.2. These include box and crossed-box topologies dependent on the quartic coupling and a correction to the trilinear vertex with a quartic contribution. Here the top Yukawa coupling is set to its SM value.

While Reference [133] originally uses these calculations to estimate constraints on κ_λ and κ_4 for hypothetical future colliders with centre-of-mass energies of 27 and 100 TeV, Addendum [134] updates the results for centre-of-mass energies of 13, 13.6 and 14 TeV. The calculations are implemented in POWHEG-BOX [82], allowing for the generation of simulation samples at the aforementioned energies at differential values of κ_λ and κ_4.

Figure 8.3 shows the hypothetical constraints placed on κ_λ and κ_4 simultaneously from Higgs boson pair production and triple Higgs production for a

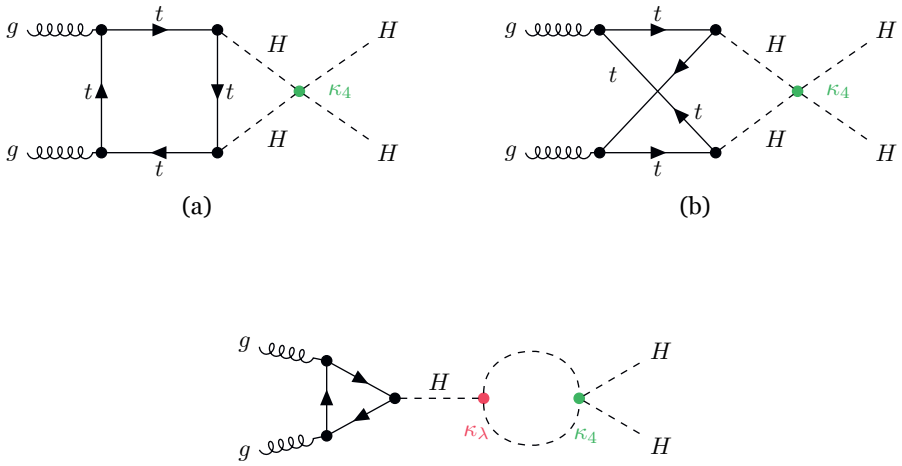

Figure 8.2: Examples of two-loop κ_4-dependent diagrams contributing to the ggF HH amplitude: the box (a) and crossed box (b) topologies, in which the quartic vertex couples directly to the top quark loop, and a correction to the trilinear vertex involving both the trilinear and quartic self-couplings (c). The trilinear and quartic self-coupling vertices are indicated in red and green respectively. Adapted from Ref. [133].

14 TeV collider such as the LHC. These projections assume an uncertainty on the Higgs boson pair production signal strength of 50%, which is representative of HL-LHC projections, and an upper limit for triple Higgs production of 20 times the Standard Model value. The two measurements are strongly complementary, with Higgs boson pair production constraining κ_λ to a narrow band but constraining κ_4 more loosely, and conversely, triple-Higgs constraining κ_λ only loosely. Their combination therefore leaves two allowed islands, one containing the Standard Model point and a second at $\kappa_\lambda \approx 3 - 4$. This arises from the degeneracy of the Higgs pair production rate in κ_λ.

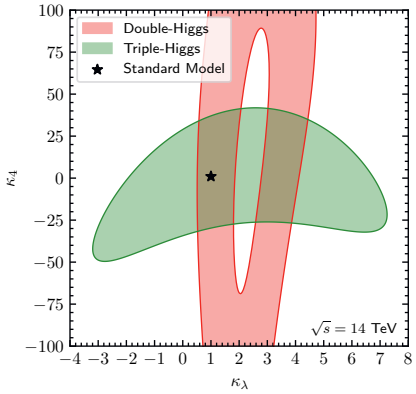

Figure 8.3: Simultaneous constraints on κ_λ and κ_4 from Higgs boson pair production (red) and triple Higgs production (green), adapted from [134]. The Standard Model prediction is shown as a star.

8.2 Towards a Simultaneous κ_λ, κ_4 constraint in ATLAS

The work described above is a phenomenological study and the first to perform this kind of indirect measurement of κ_4. The calculation implemented in POWHEG-BOX provides a useful basis from which to perform a more complete sensitivity study and apply this method using real ATLAS simulation and data. Such a study was not completed in this thesis due to time constraints. However, this section described the first steps already taken towards realising it, and the proposed future developments.

In order to obtain meaningful constraints at Run 2 and Run 3 energies, Monte Carlo simulations are required that include each step of the simulation chain described in Chapter 5: matrix element, parton showering, hadronisation, underlying event, detector simulation and reconstruction. An initial validation study was performed on truth-level events without reconstruction and detector simulation.

Using the POWHEG-BOX implementation of the calculation described above, truth-level ggF $HH \to b\bar{b}\gamma\gamma$ samples of 10000 events each were generated at

$\sqrt{s} = 13\,\text{TeV}$ for the following [κ_λ, κ_4] grid-points: [1,1], [1,0], [−1,0], [5,0], [5,100], [2,300], [2,−300], [−1,−100], [−1,−300], [−1,100], [−1,300]. The grid-points were chosen to cover a range of values both close to the Standard Model prediction and close to the experimental bounds on both couplings. Pythia8 was used for the parton showering, hadronisation and underlying event simulation.

Figure 8.4 shows the $m^*_{b\bar{b}\gamma\gamma}$ distribution for each grid point. This variable is chosen because the destructive interference between the triangle and box diagrams, and hence the dependence on the self-couplings, varies strongly with the invariant mass of the Higgs boson pair. The distributions of other kinematic variables are shown in Appendix C.

To validate the samples against the Run 2 ggF HH samples used in the $HH \to b\bar{b}\gamma\gamma$ analysis, the $m_{b\bar{b}}$ and jet multiplicity distributions are compared in Figures 8.5 and 8.6 for three $\kappa_4 = 0$ grid points against the equivalent κ_λ points. These variables are chosen for the validation because they are largely insensitive to the self-couplings, therefore any discrepancy would indicate a problem with the generator configuration or its interface to Pythia8 rather than a coupling effect. The p-values, χ^2 and ratios are also shown.

The $m^*_{b\bar{b}\gamma\gamma}$ distributions in Figure 8.4 show the expected behaviour for κ_λ variations, with a softer spectrum at higher values for κ_λ closer to unity and a narrower peak at lower values for κ_λ further from unity. Variations in κ_4 produce almost no change in the shape of the distributions, which are here normalised to unit area. The dependence on κ_4 enters predominantly through the overall normalisation. The generated samples are also statistically compatible with the nominal $HH \to b\bar{b}\gamma\gamma$ samples at the same values of κ_λ, as can be seen in Figures 8.5 and 8.6. Of the twelve comparisons performed across the variables shown here and in Appendix C, the p-values are uniformly distributed with the exception of the jet multiplicity at κ_λ = 5 (Figure 8.6c), for which $p = 0.048$. Given the number of comparisons, obtaining one p-value below 0.05 by chance is expected, and no systematic trend is observed. The POWHEG-BOX implementation is therefore taken to be valid for a full sensitivity study.

The next step in this study is to perform a full simulation of the generated samples within the ATLAS framework, including detector simulation and re-

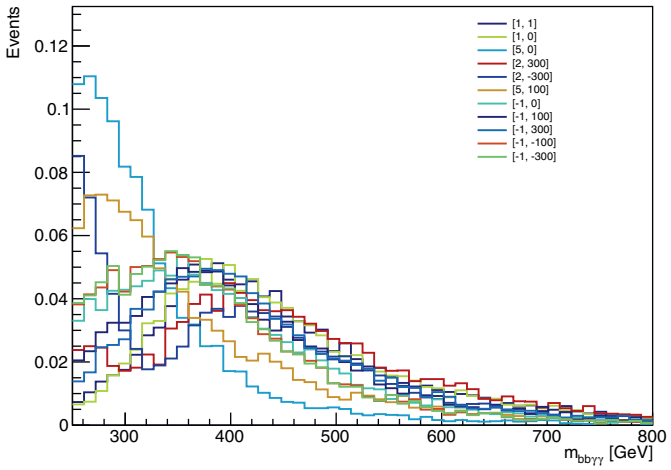

Figure 8.4: Distribution of $m^*_{b\bar{b}\gamma\gamma}$ for ggF $HH \to b\bar{b}\gamma\gamma$ samples generated at $\sqrt{s} = 13\,\text{TeV}$ for various $[\kappa_\lambda, \kappa_4]$ grid-points.

construction, and to generate the same κ_4-dependent samples in other Higgs boson pair production channels. A combination can then be performed with the triple-Higgs production measurement to obtain the first simultaneous constraint on κ_λ and κ_4 in ATLAS and produce a contour plot of the type shown in Figure 8.3 representing real experimental constraints.

The treatment of the other coupling modifiers would also require revisiting. In the $HH \to b\bar{b}\gamma\gamma$ analysis, κ_t is fixed to its Standard Model value on the grounds that it is well constrained by single-Higgs measurements, while κ_V, being less well constrained, is instead retained as a nuisance parameter (Section 7.4). The general criterion is whether the external constraint on a coupling modifier is tight compared to the sensitivity of the measurement itself. This is satisfied for the κ_λ measurement in $HH \to b\bar{b}\gamma\gamma$, but not for κ_4. The two-loop κ_4 contribution alters the ggF HH rate at the percent level, while the box contribution to the cross section scales as κ_t^4 and κ_t is currently known to around 10%. This is an order of magnitude larger than the effect being measured, and justifies the floating of κ_t in this measurement.

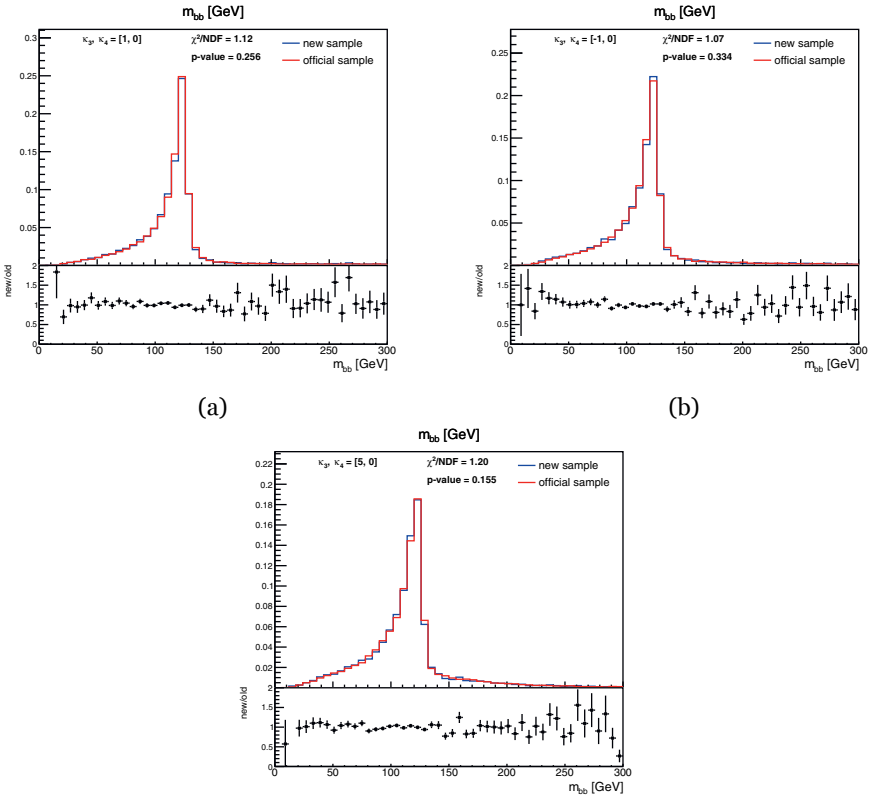

Figure 8.5: Distributions of $m_{b\bar{b}}$ for κ_4-sensitive ggF $HH \to b\bar{b}\gamma\gamma$ samples generated at $\sqrt{s} = 13\,\text{TeV}$ for grid-points (a) [1, 0], (b) [-1, 0], (c) [5, 0] compared to the equivalent $HH \to b\bar{b}\gamma\gamma$ κ_λ points. The p-values, χ^2 and ratios are also shown.

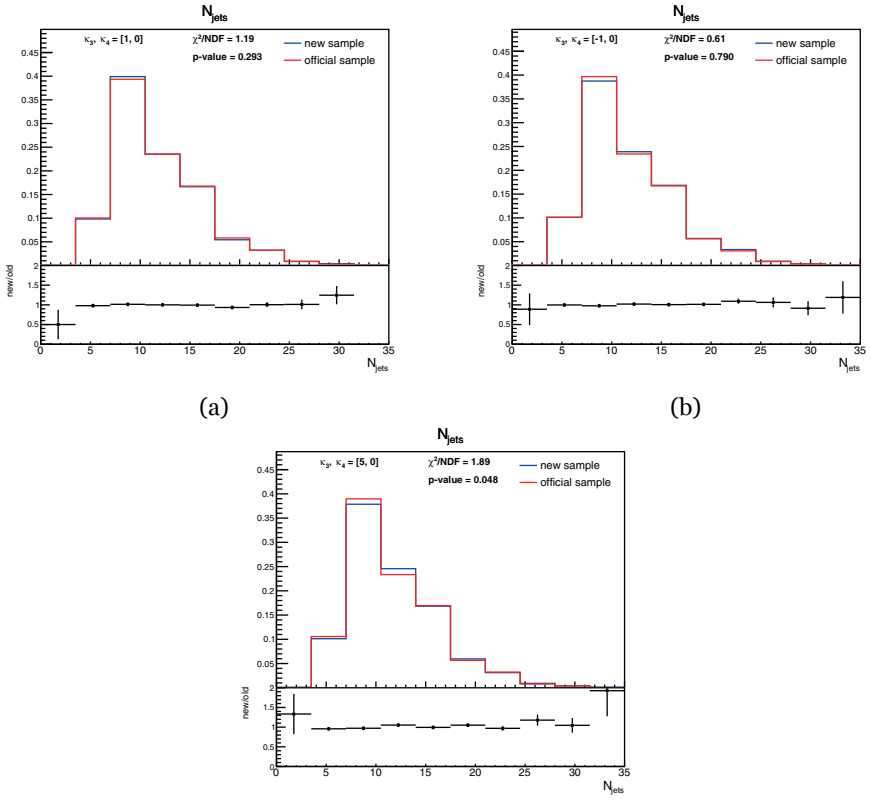

Figure 8.6: Distributions of number of jets for κ_4-sensitive ggF $HH \to b\bar{b}\gamma\gamma$ samples generated at $\sqrt{s} = 13\,\text{TeV}$ for grid-points (a) [1, 0], (b) [-1, 0], (c) [5, 0] compared to the equivalent $HH \to b\bar{b}\gamma\gamma$ κ_λ points. The p-values, χ^2 and ratios are also shown.

The assumption implicit in the κ-framework is that new physics acts only to rescale Standard Model couplings, as with the κ_λ and κ_{2V} constraints presented in this thesis. Interpretations in more general effective field theory frameworks, such as Standard Model Effective Field Theory (SMEFT) and Higgs Effective Field Theory (HEFT) as discussed in Reference [135], do not make this assumption and include other couplings. These frameworks are therefore suitable for this kind of measurement as a complement to the κ-framework.

While the prospect of constraining the quartic self-coupling may appear remote, it is worth recalling that Higgs boson pair production itself was long regarded as a measurement belonging to the HL-LHC era, if it was considered accessible at all. The analysis presented in Chapter 7 reaches an expected significance above $1\,\sigma$ in a single channel using a partial Run 3 dataset through a combination of increased luminosity and sustained improvement in reconstruction and analysis techniques. The quartic coupling today occupies the same position that the trilinear coupling occupied a decade ago: a parameter of the Higgs potential that is experimentally out of reach yet worth pursuing precisely because it is not yet known what will be found there. The samples and validation presented in this chapter represent a first step towards the first such constraint in ATLAS and towards a more complete experimental determination of the shape of the Higgs potential.

Conclusion and Outlook 9

The work presented in this thesis is concerned with uncovering the structure of the Higgs potential through the measurement of the Higgs boson trilinear self-coupling modifier κ_λ, accessible directly at the LHC through Higgs boson pair production. The final state most sensitive to κ_λ is $HH \to b\bar{b}\gamma\gamma$, due to the high branching ratio of the $H \to b\bar{b}$ decay ($\sim$58%) and the clean detector signature and good invariant mass resolution ($1-2$%) of the $H \to \gamma\gamma$ decay.

The search for Higgs boson pair production in the $b\bar{b}\gamma\gamma$ final state was performed with the ATLAS Experiment using both the Run 2 dataset at $\sqrt{s} = 13\,\text{TeV}$ and, for the first time, the partial Run 3 dataset at $\sqrt{s} = 13.6\,\text{TeV}$ up to and including 2024, amounting to a total integrated luminosity of $308\,\text{fb}^{-1}$. Novel analysis techniques were applied to enhance the sensitivity of the search, including the use of the GN2 b-tagging algorithm and a kinematic fit to improve the mass resolution of the $H \to b\bar{b}$ system.

This search yields an observed Higgs pair production signal strength of $\mu_{HH} = 0.9^{+1.4}_{-1.1}$, corresponding to an observed statistical significance of $0.78\,\sigma$. The expected significance for Standard Model Higgs pair production reached $1.01\,\sigma$, the first time any single channel has surpassed this sensitivity threshold. At 95% CL, limits of $[-1.6, 6.6]$ and $[-0.5, 2.6]$ are placed on the coupling modifiers κ_λ and κ_{2V} respectively.

Relative to the ATLAS Run 2 Legacy analysis [106], the expected sensitivity to μ_{HH} is improved by approximately 85%. Of this, roughly 60% is attributable to a larger dataset – 50% from the inclusion of Run 3 data and 10% from re-optimised analysis categories – with a further 20% from the adoption of the GN2 b-tagging algorithm and 5% from the kinematic fit. The expected 95% confidence intervals on κ_λ and κ_{2V} were narrowed by approximately 20% and 30% respectively.

Outlook

At the time of writing, Run 3 has seen its final collisions, marking the end of this era of the LHC. As the long shutdown and the preparations for the HL-LHC era are underway, it is worth reflecting on earlier projections for this era and how they have evolved in the course of Run 2 and Run 3. The prospects for Higgs boson pair production in particular serve as an illustration of this.

An early study, published in 2014, estimates a significance of $1.3\,\sigma$ for Standard Model Higgs boson pair production in the $b\bar{b}\gamma\gamma$ channel with the full HL-LHC dataset of $3000\,\text{fb}^{-1}$ [136]. This is of the same order as the expected significance of $1.01\,\sigma$ achieved with $308\,\text{fb}^{-1}$ of Run 2 and partial Run 3 data in the same channel. This projection is based on the expected performance of the upgraded detector and analysis techniques available at the time and has since been superseded by more optimistic predictions. A current ATLAS projection, combining six di-Higgs channels, predicts a significance of $4.3\,\sigma$ with the same dataset [130].

Although these projections involve different datasets and no direct comparison can therefore be made between them, they are nonetheless illustrative. The increasing optimism of these projections over the course of Run 2 and Run 3 illustrates how improvements in detector performance, object reconstruction and analysis techniques can make a process thought to be out of reach seem increasingly attainable. This history shows that it is not unrealistic to expect that, through continued improvements, seeing evidence ($\geq 3\,\sigma$) of Higgs pair production may already be possible with the full Run 3 dataset. As the HL-LHC delivers its much larger dataset, there is even more reason for optimism that discovery ($\geq 5\,\sigma$) of Higgs pair production may be achieved in

its lifetime.

This is also relevant when considering the prospects for probing the Higgs boson quartic self-coupling κ_4. This coupling is at present only very weakly constrained, but the evolution of the projections for Higgs boson pair production shows that, as experimental techniques continue to evolve, our experiments are able to achieve more than what was previously thought possible.

In this context, the simultaneous interpretation of Higgs pair production and triple-Higgs production offers a promising route towards constraining κ_λ and κ_4 together. Incorporating the second-order electroweak corrections to Higgs boson pair production introduces sensitivity to κ_4 into a process with a substantially larger production rate than triple-Higgs. A future combination of the two processes could therefore exploit the complementary dependence on the two self-couplings to increase sensitivity to κ_4 even before direct observation of triple-Higgs production becomes possible. The continued improvement of Higgs pair production sensitivity, together with the much larger HL-LHC dataset and ongoing advances in theoretical and experimental precision, makes this type of combination an especially interesting direction for future studies of the Higgs potential.

Summary 10

One of the most innate human qualities is curiosity. We see this clearly in young children, with their sense of wonder and endless repetitions of "Why?", peeling back layers like a Matryoshka doll until they reach the truth (or the end of their parents' knowledge and patience). While this curiosity often takes a back seat to adult life, I don't believe it ever truly leaves us. Sometimes it takes the form of scientific inquiry into the nature of the universe. Much like the child's Matryoshka doll of "Why?", you can look at the physical world around you as humans have always done, zooming in further and asking "What?": What is the sun? What is a hydrogen atom? What is a nucleus? What is a proton?

Scientists asking these questions for centuries have led us to understand that the universe as we know it is made up of a small number of elementary particles. These particles fall into two broad categories: fermions, which are the building blocks of matter, and bosons which carry forces between them. This set of particles and forces is the Standard Model of particle physics, and it encapsulates our current understanding of the makeup of the physical world.

The Standard Model has been remarkably successful in its description of nature and has predicted the existence of now-known particles long before we were able to observe them. But not everything that has been observed can be

adequately accounted for by the Standard Model. For example, the presence of dark matter in galaxies and the observed matter-antimatter asymmetry in the universe have no explanation in the Standard Model, among other tensions. For this reason, it is considered an incomplete theory, motivating the search for new physics beyond it.

What is responsible for mass in the Standard Model? This question is answered by the Higgs boson, hypothesised more than 60 years ago [18–20] and discovered by the ATLAS and CMS experiments in 2012 [1, 2]. Over the subsequent years, the properties of the Higgs boson and its interactions with the particles of the Standard Model have been measured with increasing precision. Such measurements provide an important test of the Standard Model: a measured value significantly different from the Standard Model prediction would be a clear sign of new physics beyond it. So far, the Standard Model has held up to these tests.

What kind of interaction does the Higgs boson have with itself? What is the Higgs potential? These questions are related by the Higgs boson's self-coupling, both the three-point (trilinear) and the four-point (quartic) self-coupling, which together directly determine the shape of the Higgs potential, a fundamental property of the Standard Model. This thesis is primarily concerned with the trilinear coupling, the more experimentally accessible of the two.

The Higgs potential, shown in Figure 10.1, affects the stability of the vacuum and consequently the evolution of the universe. Whether its true shape reflects its Standard Model form is a critical open question. A modified Higgs potential shape could also help resolve some of the Standard Model's open tensions: for instance, it could have permitted conditions in the early universe that allowed matter to dominate over antimatter, as mentioned above. For these reasons, the measurement of the Higgs boson self-coupling is a priority of experimental particle physics.

Our most powerful laboratory for such experimental endeavours is the Large Hadron Collider (LHC) at CERN, a circular particle accelerator with a circumference of 27 km, located 100 m underground in the Geneva region. The LHC accelerates and collides protons at unprecedented energies and luminosities, and because the proton is a composite particle, these collisions are capable of producing varied and rare final states. The 'camera' used to record these final states is the ATLAS detector, a general-purpose particle detector with a

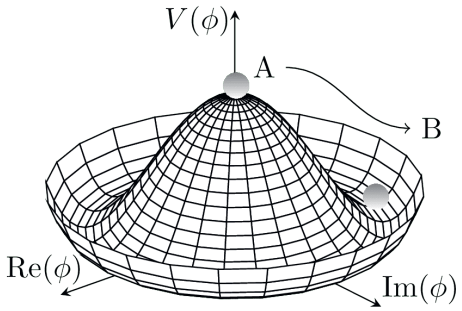

Figure 10.1: The shape of the Standard Model Higgs potential, taken from [21].

cylindrical geometry, designed to measure the energy and momentum of particles produced in LHC collisions. A schematic of the LHC with the locations of the ATLAS detector and the LHC's three other experiments, ALICE, CMS and LHCb, is shown in Figure 10.2

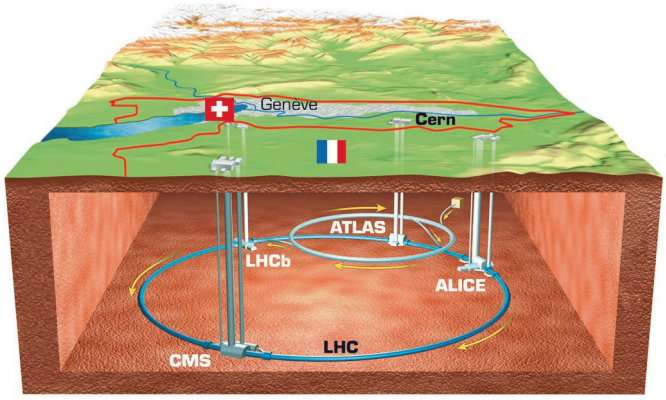

Figure 10.2: Schematic of the LHC underground facilities in the Geneva region and the caverns of the four experiments: ATLAS, ALICE, CMS and LHCb [137].

Higgs Boson Pair Production at the LHC

The Higgs boson self-coupling appears in exactly one of the rare processes that the LHC is capable of producing, namely Higgs boson pair production, occurring roughly 1000 times less than the production of a single Higgs boson. The dominant mode of Higgs pair production can occur through both the triangle and the box diagrams, illustrated in Figures 10.3a and 10.3b respectively, whose amplitudes interfere destructively. Of the two diagrams, only the triangle diagram contains the Higgs boson trilinear self-coupling, here denoted κ_λ. The subdominant mode of Higgs pair production is sensitive to κ_λ and additionally to the coupling between two Higgs bosons and two vector bosons, κ_{2V}.

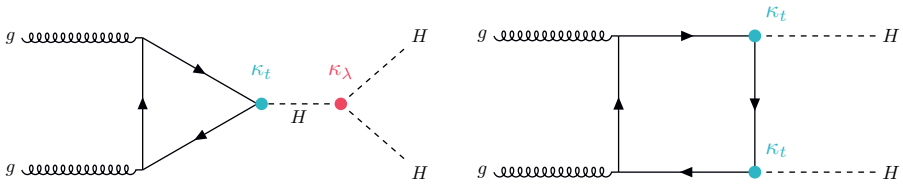

(a) (b)

Figure 10.3: The dominant Higgs pair production process, consisting of the triangle (a) and box (b) diagrams.

The Higgs boson can decay into a variety of final-states, giving a variety of possible signatures to Higgs boson pair production:

- The decay of both Higgs bosons into b-quarks, $HH \to 4b$, has the highest branching ratio of all the channels in ATLAS at 34%, but the final state of four b-jets is difficult to reconstruct and to separate from an overwhelming multi-jet background.

- The decay of one of the Higgs bosons into b-quarks and one into τ leptons has a lower branching ratio of 7.3% and more manageable background than the $HH \to 4b$ channel, but the presence of τ leptons still makes for challenging final-state reconstruction.

- The $HH \to b\bar{b}\gamma\gamma$ channel, in which one Higgs boson decays into two b-quarks and the other into two photons has a branching ratio of only 0.26%,

but the presence of the photons allows for a final state that is well reconstructed and easier to distinguish from background than the two higher branching ratio channels in ATLAS. It is this channel that is the focus of this thesis.

The search for $HH \to b\bar{b}\gamma\gamma$

This thesis details the search for Higgs boson pair production in the $HH \to b\bar{b}\gamma\gamma$ channel, performed using data collected by the ATLAS experiment during Run 2 at 13 TeV and Run 3 at 13.6 TeV and corresponding to an integrated luminosity of 140 fb^{-1} and 168 fb^{-1} respectively. The Run 3 dataset includes data taken in 2024, analyzed for the first time.

A large part of the effort made in physics analyses of this kind is the separation of signal from background. For a process as rare as $HH \to b\bar{b}\gamma\gamma$, this is especially important in order to obtain a meaningful result. The signal process is defined as Higgs boson pair production in the $b\bar{b}\gamma\gamma$ final state from the two dominant production processes. Both Standard Model Higgs boson pair production and Higgs boson pair production with coupling values different from the Standard Model prediction of one are considered.

The background processes are those which mimic the final state of the signal. Two types of background processes are considered: single-Higgs production, in which the Higgs boson decays into a photon pair and which displays a resonance peak in the diphoton invariant mass distribution at the Higgs boson mass of $m_H = 125$ GeV, and non-resonant diphoton production with jets in the final state, referred to as continuum background, which has a smoothly falling distribution in the diphoton invariant mass. Simulated data is generated for all signal and background processes in order to model their distributions such that the collected data can be compared to theoretical predictions.

To separate signal from background, a selection followed by a categorisation is performed on the data events. The selection includes the requirement that events contain two photons and at least two b-jets. The classification of b-jets against other jets is performed by a new transformer-based graph neural network algorithm called GN2, which rejects approximately four times more background than its predecessors. For this reason, a working point cor-

responding to an 85% b-jet efficiency is used, allowing for a higher signal efficiency than that of previous iterations in which a 77% working point was used.

The event categorisation is performed using a boosted decision tree (BDT), a machine-learning algorithm trained on the kinematic properties of simulated signal and background events, whose output is a score reflecting how 'signal-like' each event is. Two BDTs are trained in two separate regions of the final state invariant mass $m^*_{b\bar{b}\gamma\gamma}$ due to the dependence of this variable on the value of κ_λ. Values close to the Standard Model prediction of unity produce a sharper distribution peaking at higher values of $m^*_{b\bar{b}\gamma\gamma}$, while values far from unity produce a broader distribution that peaks at lower values of $m^*_{b\bar{b}\gamma\gamma}$. Splitting the events into mass categories allows for the exploitation of this variable to enhance sensitivity to κ_λ.

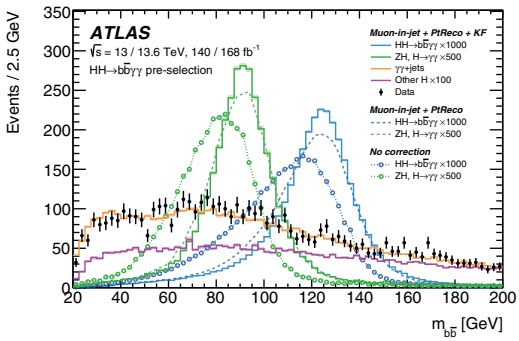

Figure 10.4: The $m_{b\bar{b}}$ distribution of the Higgs pair production signal (blue), single Higgs backgrounds (green and purple) and continuum background (orange) overlaid with data (black points). Distributions are shown for cases with no correction (dotted line), standard b-jet corrections only (dashed line) and standard b-jet corrections with the kinematic fit (solid line).

The highest-ranked variable in the BDT training is the invariant mass of the two b-jets in the final state, $m_{b\bar{b}}$, whose resolution is much worse than that of $m_{\gamma\gamma}$. A kinematic fit is performed that uses the good resolution of $m_{\gamma\gamma}$ and the conservation of transverse momentum to correct the four-momentum of jets in each event and subsequently improve the $m_{b\bar{b}}$ resolution. All inputs to the BDTs are based on the corrected four-momenta. The resulting $m_{b\bar{b}}$ distribu-

tion, compared to without the kinematic fit, is shown by the blue line in Figure 10.4. The $m_{b\bar{b}}$ resolution is improved by 27% with respect to no corrections.

Based on the output of the BDTs, events are split further into three categories in the high mass region and four categories in the low mass region. The category boundaries are determined by optimising the expected sensitivity to the signal and events that do not fall within any of the categories are rejected. An illustration of a real $HH \to b\bar{b}\gamma\gamma$ candidate event within the detector from Run 3 data left after the selection and categorisation steps is shown in Figure 10.5.

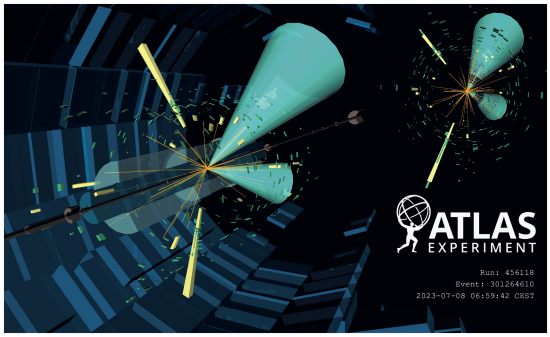

Figure 10.5: A $HH \to b\bar{b}\gamma\gamma$ candidate event from Run 3. The two final state b-jets (turquoise cones) and photons (yellow towers) are shown. Taken from [138].

The final results are obtained by performing a maximum likelihood fit, shown in Figure 10.6, over all categories to the diphoton invariant mass. The observed signal strength, defined as the ratio of the observed Higgs boson pair production rate to that predicted by the Standard Model, is $\mu_{HH} = 0.9^{+1.4}_{-1.1}$, compatible with the Standard Model prediction and corresponding to an observed statistical significance of $0.78\,\sigma$. Under the assumption of Standard Model Higgs pair production, this analysis was expected to observe it with a significance of $1.01\,\sigma$. This is the first time any single Higgs boson pair production channel has been expected to exceed a sensitivity of $1\ \sigma$: a marker of how far the field has advanced since the Higgs boson's discovery.

In the absence of a significant signal, an upper limit is set on the Higgs

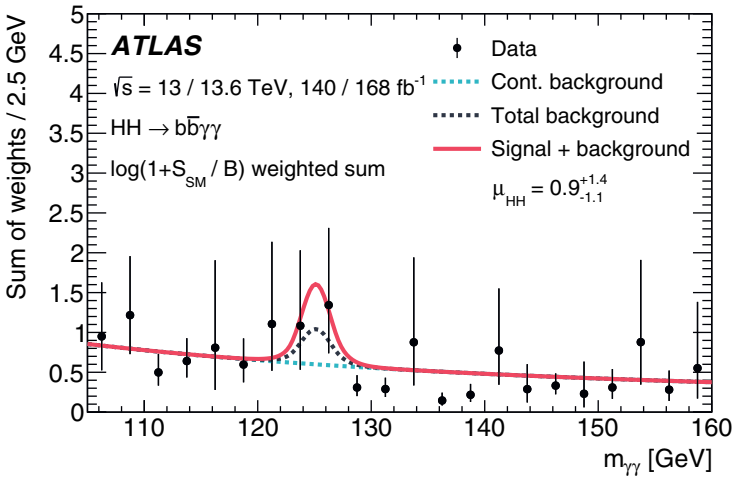

Figure 10.6: Weighted diphoton invariant mass distribution summed over all categories and the two data-taking periods. The events in each category are weighted by $\log(1 + S_{SM}/B)$. The points represent the data while the lines represent the fit results for the continuum background (light dotted), full background (black dotted) and full fit (solid).

boson pair production signal strength: at 95% confidence, the observed data excludes values above 3.7, in close agreement with the expected limit under the assumption of Standard Model Higgs boson pair production. Constraints are also placed on the couplings κ_λ and κ_{2V}, both of which are consistent with the Standard Model prediction of one, of $-1.6 < \kappa_\lambda < 6.6$ and $-0.5 < \kappa_{2V} < 2.6$, excluding coupling values outside these intervals with 95% confidence.

Compared to the previous ATLAS Run 2 legacy search in this channel [106], the sensitivity of this analysis improved by approximately 85%. While the larger dataset is responsible for the majority of this gain, a meaningful part also comes from analysis improvements developed as part of this thesis: the improved b-tagging algorithm described above, together with a kinematic fit that improves the resolution of the reconstructed $m_{b\bar{b}}$, contributed a further improvement of approximately 25%.

Future Prospects

The search described in this thesis was a milestone for Higgs boson pair production, but the future brings developments that will allow us to peel back the curtain even more. As the LHC moves beyond Run 3 and towards its high-luminosity upgrade, the dataset available to this measurement will continue to grow, and further improvements in detector performance and analysis technique will follow. History offers good reason for optimism: expectations for this measurement have repeatedly been surpassed, and there is little reason to think this trend will stop. It is not unreasonable to hope that evidence, and eventually discovery, of Higgs boson pair production lies within reach in the years ahead.

Looking further still, this same kind of search may one day be combined with future measurements of the even rarer production of three Higgs bosons, opening a path towards constraining not just the trilinear coupling but also its quartic counterpart, and bringing the full shape of the Higgs potential within experimental reach for the first time.

Not so long ago, the existence of the Higgs boson was itself an open question; today, we study how it interacts with itself, and with the world around it. This is, in many ways, the nature of scientific inquiry: each answer reached only uncovers the next layer of the Matryoshka doll. This thesis is a small contribution to that ongoing search, and, in the same spirit of curiosity with which this chapter began, there is every reason to expect this layer will not be the last.

Samenvatting 11

Nieuwsgierigheid is een van de meest wezenlijke menselijke eigenschappen. We zien dit duidelijk terug bij jonge kinderen, met hun gevoel van verwondering en eindeloze herhaling van “Waarom?”, laag voor laag afpellend als een matroesjkapop totdat ze de waarheid bereiken (of de grens van de kennis en het geduld van hun ouders). Hoewel deze nieuwsgierigheid in het volwassen leven vaak naar de achtergrond verdwijnt, verdwijnt ze naar mijn overtuiging nooit helemaal. Soms neemt ze de vorm aan van wetenschappelijk onderzoek naar de aard van het universum. Net als bij het matroesjkapopje van het kind met zijn “Waarom?”, kun je naar de fysieke wereld om je heen kijken zoals mensen dat altijd al hebben gedaan, steeds verder inzoomend en “Wat?” vragend: Wat is de zon? Wat is een waterstofatoom? Wat is een atoomkern? Wat is een proton?

Wetenschappers die deze vragen al eeuwenlang stellen, hebben ons doen begrijpen dat het universum zoals we het kennen is opgebouwd uit een klein aantal elementaire deeltjes. Deze deeltjes vallen uiteen in twee brede categorieën: fermionen, de bouwstenen van materie, en bosonen, die de krachten tussen deeltjes overdragen. Deze verzameling van deeltjes en krachten vormt het Standaardmodel van de deeltjesfysica, en het omvat ons huidige begrip van de opbouw van de fysieke wereld.

Het Standaardmodel is opmerkelijk succesvol gebleken in zijn beschrijving van de natuur en heeft het bestaan van inmiddels bekende deeltjes voorspeld lang voordat we ze konden waarnemen. Maar niet alles wat is waargenomen kan afdoende door het Standaardmodel worden verklaard. Zo hebben, naast andere spanningen, de aanwezigheid van donkere materie in sterrenstelsels en de waargenomen materie-antimaterie-asymmetrie in het universum geen verklaring binnen het Standaardmodel. Om die reden wordt het beschouwd als een onvolledige theorie, wat de zoektocht naar nieuwe natuurkunde voorbij het Standaardmodel motiveert.

Wat is verantwoordelijk voor massa in het Standaardmodel? Deze vraag wordt beantwoord door het Higgsboson, meer dan 60 jaar geleden voorspeld [18–20] en in 2012 ontdekt door de ATLAS- en CMS-experimenten [1, 2]. In de jaren daarna zijn de eigenschappen van het Higgsboson en zijn interacties met de deeltjes van het Standaardmodel met steeds grotere precisie gemeten. Dergelijke metingen vormen een belangrijke toets voor het Standaardmodel: een gemeten waarde die significant afwijkt van de voorspelling van het Standaardmodel zou een duidelijk teken zijn van nieuwe natuurkunde daarbuiten. Tot nu toe heeft het Standaardmodel deze toetsen doorstaan.

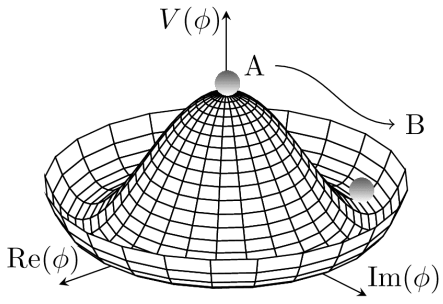

Figuur 11.1: De vorm van de Higgspotentiaal in het Standaardmodel, overgenomen uit [21].

Wat voor interactie heeft het Higgsboson met zichzelf? Wat is de Higgspotentiaal? Deze vragen hangen samen met de zelfkoppeling van het Higgsboson: zowel de drie-punts- (trilineaire) als de vier-punts- (quartaire) zelfkoppeling, die samen rechtstreeks de vorm van de Higgspotentiaal bepalen, een funda-

mentele eigenschap van het Standaardmodel. Dit proefschrift richt zich voornamelijk op de trilineaire koppeling, de experimenteel beter toegankelijke van de twee.

De Higgspotentiaal, weergegeven in Figuur 11.1, beïnvloedt de stabiliteit van het vacuüm en daarmee de evolutie van het universum. Of de werkelijke vorm ervan overeenkomt met de voorspelling van het Standaardmodel, is een belangrijke nog onbeantwoorde vraag. Een afwijkende vorm van de Higgspotentiaal zou ook een oplossing kunnen bieden voor enkele van de open spanningen binnen het Standaardmodel: zo zou deze bijvoorbeeld omstandigheden in het vroege universum mogelijk hebben gemaakt die ervoor zorgden dat materie de overhand kreeg op antimaterie, zoals hierboven genoemd. Om deze redenen is de meting van de zelfkoppeling van het Higgsboson een prioriteit binnen de experimentele deeltjesfysica.

Ons krachtigste laboratorium voor dergelijk experimenteel werk is de Large Hadron Collider (LHC) bij CERN, een ringvormige deeltjesversneller met een omtrek van 27 km, gelegen op 100 m diepte onder de grond in de regio Genève. De LHC versnelt en laat protonen botsen bij ongekende energieën en luminositeiten, en omdat het proton een samengesteld deeltje is, kunnen deze botsingen uiteenlopende en zeldzame eindtoestanden opleveren.

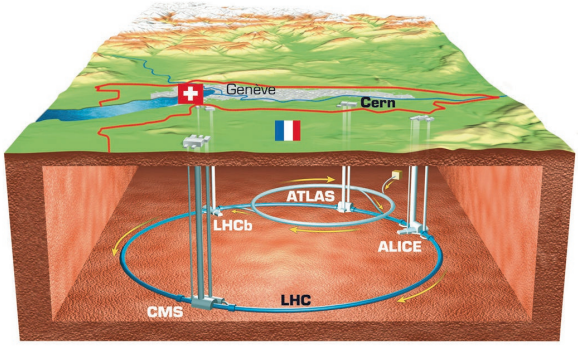

Figuur 11.2: Schematische weergave van de ondergrondse LHC-faciliteiten in de regio Genève en de cavernes van de vier experimenten: ATLAS, ALICE, CMS en LHCb [137].

Higgsboson-paarproductie aan de LHC

De 'camera' die deze eindtoestanden vastlegt, is de ATLAS-detector, een universele deeltjesdetector met een cilindrische geometrie, ontworpen om de energie en impuls te meten van deeltjes die ontstaan bij botsingen in de LHC. Een schematische weergave van de LHC met de locaties van de ATLAS-detector en de drie andere LHC-experimenten, ALICE, CMS en LHCb, is te zien in Figuur 11.2.

De zelfkoppeling van het Higgsboson komt voor in precies één van de zeldzame processen die de LHC kan produceren, namelijk Higgsboson-paarproductie, die ongeveer 1000 keer minder vaak voorkomt dan de productie van een enkel Higgsboson. De dominante productiewijze van Higgsboson-paren kan plaatsvinden via zowel het driehoeksdiagram als het boxdiagram, respectievelijk weergegeven in Figuur 11.3a en 11.3b, waarvan de amplitudes destructief interfereren. Van deze twee diagrammen bevat alleen het driehoeksdiagram de trilineaire zelfkoppeling van het Higgsboson, hier aangeduid als κ_λ. De subdominante productiewijze van Higgsboson-paren is gevoelig voor κ_λ en daarnaast voor de koppeling tussen twee Higgsbosonen en twee vectorbosonen, κ_{2V}.

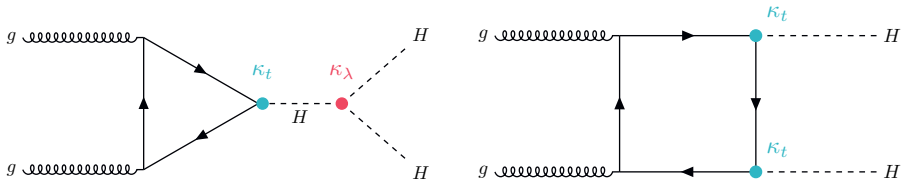

(a) (b)

Figuur 11.3: Het dominante productieproces van Higgsboson-paren, bestaande uit het driehoeksdiagram (a) en het boxdiagram (b).

Het Higgsboson kan vervallen naar uiteenlopende eindtoestanden, wat leidt tot verschillende mogelijke signaturen voor Higgsboson-paarproductie:

- Het verval van beide Higgsbosonen naar b-quarks, $HH \to 4b$, heeft de hoogste vervalskans van alle kanalen in ATLAS met 34%, maar de eindtoestand van vier b-jets is lastig te reconstrueren en te scheiden van een overweldigende multi-jet-achtergrond.

- Het verval van één Higgsboson naar b-quarks en het andere naar τ-leptonen heeft een lagere vervalskans van 7.3% en een beter beheersbare achtergrond dan het $HH \to 4b$-kanaal, maar de aanwezigheid van τ-leptonen maakt de reconstructie van de eindtoestand nog steeds lastig.

- Het $HH \to b\bar{b}\gamma\gamma$ kanaal, waarbij één Higgsboson vervalt naar twee b-quarks en het andere naar twee fotonen, heeft een vervalskans van slechts 0.26%, maar de aanwezigheid van de fotonen zorgt voor een eindtoestand die goed te reconstrueren is en beter te onderscheiden van de achtergrond dan de twee kanalen met een hogere vervalskans in ATLAS. Dit is het kanaal waarop dit proefschrift zich richt.

De zoektocht naar $HH \to b\bar{b}\gamma\gamma$

Dit proefschrift beschrijft de zoektocht naar Higgsboson-paarproductie in het $HH \to b\bar{b}\gamma\gamma$ kanaal, uitgevoerd met data verzameld door het ATLAS-experiment tijdens Run 2 bij 13 TeV en Run 3 bij 13.6 TeV, overeenkomend met een geïntegreerde luminositeit van respectievelijk 140 fb^{-1} en 168 fb^{-1}. De Run 3-dataset omvat data verzameld in 2024, hier voor het eerst geanalyseerd.

Een groot deel van de inspanning bij dit soort natuurkundige analyses gaat naar het scheiden van signaal en achtergrond. Voor een proces zo zeldzaam als $HH \to b\bar{b}\gamma\gamma$ is dit van bijzonder belang om tot een betekenisvol resultaat te komen. Het signaalproces wordt gedefinieerd als Higgsboson-paarproductie in de $b\bar{b}\gamma\gamma$ eindtoestand afkomstig van de twee dominante productieprocessen. Zowel Higgsboson-paarproductie volgens het Standaardmodel als Higgsboson-paarproductie met koppelingswaarden die afwijken van de Standaardmodel-voorspelling van één worden in beschouwing genomen.

De achtergrondprocessen zijn processen die de eindtoestand van het signaal nabootsen. Er worden twee soorten achtergrondprocessen beschouwd: enkelvoudige Higgsboson-productie, waarbij het Higgsboson vervalt naar een fotonpaar en die een resonantiepiek vertoont in de invariante-massaverdeling van het twee-fotonsysteem bij de Higgsbosonmassa van $m_H = 125$ GeV, en niet-resonante twee-fotonproductie met jets in de eindtoestand, aangeduid als continuümachtergrond, die een gelijkmatig dalende verdeling vertoont in de invariante twee-fotonmassa. Voor alle signaal- en achtergrondprocessen wordt

gesimuleerde data gegenereerd om hun verdelingen te modelleren, zodat de verzamelde data met de theoretische voorspellingen kan worden vergeleken.

Om signaal van achtergrond te scheiden, worden de data-events eerst geselecteerd en vervolgens gecategoriseerd. De selectie omvat de eis dat events twee fotonen en ten minste twee b-jets bevatten. De classificatie van b-jets ten opzichte van andere jets gebeurt met een nieuw, op een transformer gebaseerd graph neuraal-netwerkalgoritme genaamd GN2, dat ongeveer vier keer meer achtergrond verwijdert dan zijn voorgangers. Om die reden wordt een werkpunt gebruikt dat overeenkomt met een b-jet-efficiëntie van 85%, wat een hogere signaalefficiëntie mogelijk maakt dan bij eerdere versies, waarin een werkpunt van 77% werd gebruikt.

De categorisering van events gebeurt met behulp van een boosted decision tree (BDT), een machine-learningalgoritme dat getraind is op de kinematische eigenschappen van gesimuleerde signaal- en achtergrondevents, en waarvan de uitkomst een score is die weergeeft hoe 'signaalachtig' elk event is. Er worden twee BDT's getraind in twee afzonderlijke gebieden van de invariante massa van de eindtoestand, $m^*_{b\bar{b}\gamma\gamma}$, vanwege de afhankelijkheid van deze variabele van de waarde van κ_λ. Waarden dicht bij de Standaardmodelvoorspelling van één leveren een scherpere verdeling op die piekt bij hogere waarden van $m^*_{b\bar{b}\gamma\gamma}$, terwijl waarden ver van één een bredere verdeling opleveren die piekt bij lagere waarden van $m^*_{b\bar{b}\gamma\gamma}$. Door de events in massacategorieën te splitsen kan deze variabele worden benut om de gevoeligheid voor κ_λ te vergroten.

De hoogst gerangschikte variabele in de BDT-training is de invariante massa van de twee b-jets in de eindtoestand, $m_{b\bar{b}}$, waarvan de resolutie veel slechter is dan die van $m_{\gamma\gamma}$. Er wordt een kinematische fit uitgevoerd die gebruikmaakt van de goede resolutie van $m_{\gamma\gamma}$ en het behoud van transversale impuls om de viervector van jets in elk event te corrigeren en zo de $m_{b\bar{b}}$ resolutie te verbeteren. Alle invoer voor de BDT's is gebaseerd op de gecorrigeerde viervectoren. De resulterende $m_{b\bar{b}}$ verdeling, vergeleken met de verdeling zonder kinematische fit, is weergegeven door de blauwe lijn in Figuur 11.4. De $m_{b\bar{b}}$ resolutie verbetert met 27% ten opzichte van geen correcties.

Op basis van de uitkomst van de BDT's worden de events verder opgesplitst in drie categorieën in het hogemassagebied en vier categorieën in het lagemassagebied. De categoriegrenzen worden bepaald door de verwachte gevoelig-

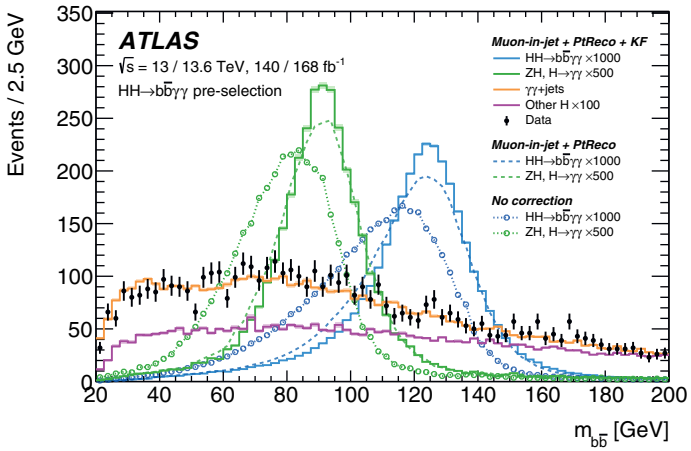

Figuur 11.4: De $m_{b\bar{b}}$ verdeling van het Higgsboson-paarproductiesignaal (blauw), de enkelvoudige-Higgsboson-achtergronden (groen en paars) en de continuümachtergrond (oranje), samen met de data (zwarte punten). De verdelingen zijn weergegeven voor de gevallen zonder correctie (stippellijn), met alleen standaard b-jetcorrecties (streepjeslijn) en met standaard b-jetcorrecties plus de kinematische fit (doorgetrokken lijn).

heid voor het signaal te optimaliseren, en events die in geen van de categorieën vallen worden verworpen. Een illustratie van een echt $HH \to b\bar{b}\gamma\gamma$ kandidaatevent in de detector, afkomstig uit Run 3-data en overgebleven na de selectie- en categoriseringsstappen, is te zien in Figuur 11.5.

De uiteindelijke resultaten worden verkregen door een maximum-likelihoodfit, weergegeven in Figuur 11.6, over alle categorieën op de invariante twee-fotonmassa uit te voeren. De waargenomen signaalsterkte, gedefinieerd als de verhouding tussen de waargenomen productiesnelheid van Higgsboson-paren en de voorspelling van het Standaardmodel, bedraagt $\mu_{HH} = 0.9^{+1.4}_{-1.1}$, in overeenstemming met de Standaardmodelvoorspelling en overeenkomend met een waargenomen statistische significantie van $0.78\,\sigma$. Onder de aanname van Higgsboson-paarproductie volgens het Standaardmodel werd verwacht dat deze analyse dit zou waarnemen met een significantie van $1.01\,\sigma$. Dit is de eerste keer dat voor een enkel Higgsboson-paarproductiekanaal een gevoeligheid van meer dan $1\,\sigma$ werd verwacht: een teken van hoezeer het vakgebied is gevorderd sinds de ont-

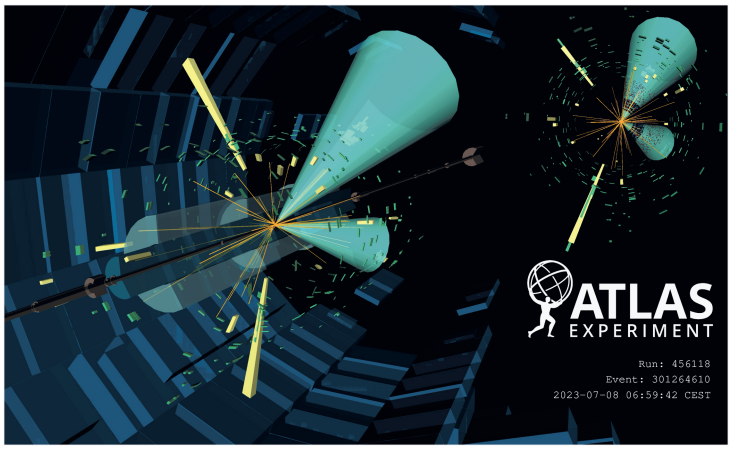

Figuur 11.5: Een $HH \to b\bar{b}\gamma\gamma$ kandidaatevent uit Run 3. De twee b-jets (turquoise kegels) en fotonen (gele torens) in de eindtoestand zijn weergegeven. Overgenomen uit [138].

dekking van het Higgsboson.

Bij afwezigheid van een significant signaal wordt een bovengrens gesteld aan de signaalsterkte van Higgsboson-paarproductie: bij een betrouwbaarheidsniveau van 95% sluiten de waargenomen data waarden boven 3.7 uit, in goede overeenstemming met de verwachte grens onder de aanname van Higgsboson-paarproductie volgens het Standaardmodel. Er worden ook grenzen gesteld aan de koppelingen κ_λ en κ_{2V}, die beide consistent zijn met de Standaardmodelvoorspelling van één, namelijk $-1.6 < \kappa_\lambda < 6.6$ en $-0.5 < \kappa_{2V} < 2.6$, waarbij koppelingswaarden buiten deze intervallen met 95% betrouwbaarheid worden uitgesloten.

Vergeleken met de eerdere ATLAS Run 2-legacyzoektocht in dit kanaal [106], is de gevoeligheid van deze analyse met ongeveer 85% verbeterd. Hoewel de grotere dataset verantwoordelijk is voor het grootste deel van deze winst, komt een significant deel ook voort uit analyseverbeteringen die als onderdeel van dit proefschrift zijn ontwikkeld: het verbeterde b-taggingalgoritme dat hierboven is beschreven, droeg samen met een kinematische fit die de resolutie van de gereconstrueerde $m_{b\bar{b}}$ verbetert, bij aan een verdere verbetering van ongeveer 25%.

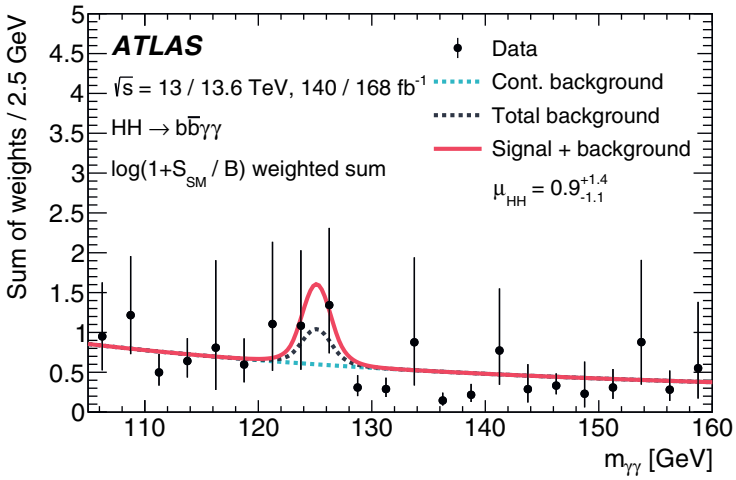

Figuur 11.6: Gewogen invariante-twee-fotonmassaverdeling, gesommeerd over alle categorieën en de twee dataverzamelingsperiodes. De events in elke categorie zijn gewogen met $\log(1+\mathrm{S_{SM}}/\mathrm{B})$. De punten stellen de data voor, terwijl de lijnen de fitresultaten weergeven voor de continuümachtergrond (lichte stippellijn), de volledige achtergrond (zwarte stippellijn) en de volledige fit (doorgetrokken lijn).

Toekomstperspectieven

De zoektocht die in dit proefschrift wordt beschreven, was een mijlpaal voor Higgsboson-paarproductie, maar de toekomst brengt ontwikkelingen die ons in staat zullen stellen nog verder achter het gordijn te kijken. Naarmate de LHC verder gaat dan Run 3 en toewerkt naar zijn high-luminosity-upgrade, zal de beschikbare dataset voor deze meting blijven groeien, en zullen verdere verbeteringen in detectorprestaties en analysetechnieken volgen. De geschiedenis geeft goede reden voor optimisme: de verwachtingen voor deze meting zijn herhaaldelijk overtroffen, en er is weinig reden om aan te nemen dat deze trend zal stoppen. Het is niet onredelijk om te hopen dat bewijs, en uiteindelijk ontdekking, van Higgsboson-paarproductie in de komende jaren binnen bereik ligt.

Verder vooruitkijkend zou dit soort zoektocht ooit gecombineerd kunnen worden met toekomstige metingen van de nog zeldzamere productie van drie Higgsbosonen, wat een weg zou openen om niet alleen de trilineaire koppeling maar ook zijn quartaire tegenhanger te beperken, en de volledige vorm van de Higgspotentiaal voor het eerst binnen experimenteel bereik te brengen.

Niet zo lang geleden was het bestaan van het Higgsboson zelf nog een open vraag; vandaag bestuderen we hoe het met zichzelf interageert, en met de wereld om zich heen. Dit is, in veel opzichten, de aard van wetenschappelijk onderzoek: elk bereikt antwoord onthult slechts de volgende laag van de matroesjkapop. Dit proefschrift is een kleine bijdrage aan die voortdurende zoektocht, en, in dezelfde geest van nieuwsgierigheid waarmee dit hoofdstuk begon, is er alle reden om te verwachten dat deze laag niet de laatste zal zijn.

Simulation and Event Selection A

This appendix collects supplementary material relating to the simulated samples and the event selection of the $HH \to b\bar{b}\gamma\gamma$ analysis presented in Chapter 7.

A.1 Signal Cutflows

Tables A.1 and A.2 give the cutflow for the Standard Model ggF and VBF HH signal samples through the pre-selection described in Section 7.7, for Run 2 and Run 3 respectively. For each selection step the expected yield, the efficiency relative to the preceding step and the cumulative efficiency are given. The yields correspond to the integrated luminosities quoted in Section 7.3, and all numbers are obtained from the AtlFast3 samples used as the nominal signal simulation in the analysis.

Both photons are required to satisfy the tight identification working point and the tight calorimeter and track isolation working points. These requirements retain 87.9% and 89.9% of ggF signal events respectively in Run 2, and 89.0% and 90.9% in Run 3, corresponding to a combined photon identification and isolation efficiency of approximately 79% and 81%. The requirement of at least two b-tagged jets is less efficient for VBF than for ggF production (35% compared to 49%).

Table A.1: Cutflow for the Standard Model ggF and VBF $HH \to b\bar{b}\gamma\gamma$ signal samples in Run 2. For each selection step the expected yield, the efficiency relative to the preceding step and the cumulative efficiency are given.

	ggF HH (SM)			VBF HH (SM)		
Selection	Yield	Rel. [%]	Cum. [%]	Yield	Rel. [%]	Cum. [%]
All events	11.40	100.00	100.00	0.62	100.00	100.00
Event cleaning and Dalitz filter	11.33	99.38	99.38	0.62	99.02	99.02
Pass trigger	7.00	61.78	61.39	0.35	55.96	55.41
Two loose photons	6.35	90.75	55.71	0.31	89.56	49.63
Trigger matching	6.31	99.42	55.39	0.31	99.27	49.27
Two tight ID photons	5.55	87.93	48.71	0.27	87.10	42.91
Two tight ID and isolated photons	4.99	89.88	43.78	0.24	89.34	38.34
Relative p_T cuts	4.58	91.87	40.22	0.22	91.01	34.89
$m_{\gamma\gamma}$ in $[105, 160]$ GeV	4.58	99.88	40.17	0.22	99.49	34.71
$N_{\text{lep}} = 0$	4.56	99.66	40.04	0.22	99.59	34.57
At least two jets	4.03	88.25	35.33	0.21	95.15	32.90
$N_{\text{jet}}^{\text{central}} < 6$	3.94	97.75	34.54	0.20	99.41	32.70
At least two b-jets (85% WP)	1.96	49.89	17.23	0.07	35.56	11.63

Table A.2: Cutflow for the Standard Model ggF and VBF $HH \to b\bar{b}\gamma\gamma$ signal samples in Run 3, with columns as in Table A.1.

	ggF HH (SM)			VBF HH (SM)		
Selection	Yield	Rel. [%]	Cum. [%]	Yield	Rel. [%]	Cum. [%]
All events	15.16	100.00	100.00	0.83	100.00	100.00
Event cleaning and Dalitz filter	14.96	98.68	98.68	0.82	98.27	98.27
Pass trigger	8.79	58.75	57.97	0.43	53.00	52.08
Two loose photons	8.11	92.30	53.51	0.40	91.28	47.54
Trigger matching	8.05	99.26	53.11	0.39	99.10	47.11
Two tight ID photons	7.17	89.00	47.27	0.34	87.82	41.37
Two tight ID and isolated photons	6.51	90.88	42.96	0.31	90.35	37.38
Relative p_T cuts	6.01	92.22	39.61	0.28	91.34	34.14
$m_{\gamma\gamma}$ in $[105, 160]$ GeV	6.00	99.89	39.57	0.28	99.48	33.97
$N_{\text{lep}} = 0$	5.98	99.67	39.44	0.28	99.63	33.84
At least two jets	5.28	88.27	34.81	0.27	94.84	32.09
$N_{\text{jet}}^{\text{central}} < 6$	5.16	97.71	34.01	0.27	99.39	31.90
At least two b-jets (85% WP)	2.54	49.18	16.73	0.09	34.41	10.98

A.2 Monte Carlo Yields

Table A.3 gives the full breakdown of the expected Monte Carlo yields per process after the pre-selection described in Section 7.7 for Run 2 and Run 3 separately. The yields are normalised according to the per-event weight given in Section 7.7 and correspond to the integrated luminosities quoted in Section 7.3.

Table A.3: Expected Monte Carlo yields per process after the pre-selection, for the Run 2 and Run 3 datasets.

Process	Run 2	Run 3
ggF HH (SM)	1.96	2.54
VBF HH (SM)	0.07	0.09
ggH	7.95	10.10
bbH	0.88	1.12
VBFH	0.87	1.03
qqZH	3.31	3.97
ggZH	1.02	1.31
W^+H	0.17	0.22
W^-H	0.13	0.17
tWH	0.19	0.09
tHjb	1.28	0.57
ttH	10.79	14.26
$\gamma\gamma$ + jets	1545.04	1893.71

A.3 Yields by Analysis Category

Tables A.4 and A.5 give the expected signal and background yields, together with the observed data counts, in each of the fourteen analysis categories defined in Section 7.8. The signal yields are quoted for the nine basis samples of the cross section parameterisation introduced in Section 7.4. The single-Higgs processes are grouped by production mode, with the W^+H, W^-H, $ggZH$ and $qqZH$ contributions collected into VH and those of $tHjb$ and tWH into tH. Events falling below the loosest category boundary in each region are discarded.

Table A.4: Expected signal and background yields and observed data counts in each analysis category for the Run 2 dataset. The signal rows are the basis samples from which the yield in each category is parameterised as a function of κ_λ and κ_{2V} (Section 7.4); VBF HH samples are labelled by $(\kappa_\lambda, \kappa_{2V}, \kappa_V)$. The continuum background is estimated from simulation with a scale factor $SF_{\text{Run 2}} = 1.79$ applied to the $\gamma\gamma$+jets sample, derived from the $m_{\gamma\gamma}$ sidebands at pre-selection level. The two continuum samples are listed separately rather than summed; their relative fractions in the background model are determined as described in Section 7.9. Uncertainties are statistical only and are omitted for readability.

Process	HM1	HM2	HM3	LM1	LM2	LM3	LM4
HH signal (parameterisation basis samples)							
ggF HH ($\kappa_\lambda = 0$)	0.769	0.506	1.026	0.263	0.202	0.087	0.059
ggF HH ($\kappa_\lambda = 1$, SM)	0.409	0.276	0.641	0.057	0.041	0.020	0.013
ggF HH ($\kappa_\lambda = 5$)	0.330	0.175	0.307	0.980	0.691	0.285	0.228
VBF HH $(1, 1, 1$, SM)	0.012	0.007	0.021	0.003	0.002	0.002	0.006
VBF HH $(1, 1, 0.5)$	0.069	0.054	0.569	0.010	0.010	0.009	0.036
VBF HH $(10, 1, 1)$	0.383	0.280	1.042	0.386	0.405	0.353	1.426
VBF HH $(1, 1.5, 1)$	0.020	0.015	0.188	0.003	0.003	0.002	0.009
VBF HH $(2, 1, 1)$	0.008	0.005	0.013	0.004	0.004	0.004	0.014
VBF HH $(-5, 1, 0.5)$	0.020	0.016	0.281	0.016	0.016	0.013	0.051
Single-Higgs background							
ggH	0.404	0.145	0.179	0.283	0.103	0.028	0.018
VBFH	0.068	0.019	0.023	0.041	0.013	0.006	0.004
VH	0.487	0.094	0.065	0.430	0.122	0.035	0.015
$t\bar{t}H$	0.224	0.045	0.028	1.137	0.396	0.101	0.036
tH	0.089	0.022	0.013	0.121	0.030	0.009	0.016
$b\bar{b}H$	0.026	0.003	<0.001	0.109	0.057	0.020	0.017
Total	1.298	0.330	0.308	2.121	0.722	0.198	0.106
Continuum background							
$\gamma\gamma + b\bar{b}$	31.7	8.0	5.8	92.4	30.8	9.5	7.0
$\gamma\gamma$ + jets	39.9	9.2	5.6	141.6	38.2	11.5	8.2
Data							
Sideband	29	9	2	145	39	7	8
Full range	41	9	4	173	44	11	11

Table A.5: Expected signal and background yields and observed data counts in each analysis category for the Run 3 dataset. The signal rows are the basis samples from which the yield in each category is parameterised as a function of κ_λ and κ_{2V} (Section 7.4); VBF HH samples are labelled by $(\kappa_\lambda, \kappa_{2V}, \kappa_V)$. The continuum background is estimated from simulation with a scale factor $SF_{\text{Run 3}} = 1.47$ applied to the $\gamma\gamma$+jets sample, derived from the $m_{\gamma\gamma}$ sidebands at pre-selection level. The two continuum samples are listed separately rather than summed; their relative fractions in the background model are determined as described in Section 7.9. Uncertainties are statistical only and are omitted for readability.

Process	HM1	HM2	HM3	LM1	LM2	LM3	LM4
HH signal (parameterisation basis samples)							
ggF HH ($\kappa_\lambda = 0$)	0.992	0.639	1.318	0.324	0.245	0.121	0.079
ggF HH ($\kappa_\lambda = 1$, SM)	0.539	0.348	0.821	0.072	0.052	0.025	0.017
ggF HH ($\kappa_\lambda = 5$)	0.427	0.223	0.398	1.250	0.856	0.388	0.310
VBF HH $(1, 1, 1$, SM)	0.015	0.009	0.026	0.003	0.003	0.002	0.008
VBF HH $(1, 1, 0.5)$	0.098	0.074	0.763	0.012	0.012	0.011	0.044
VBF HH $(10, 1, 1)$	0.514	0.366	1.208	0.498	0.554	0.440	1.830
VBF HH $(1, 1.5, 1)$	0.027	0.020	0.253	0.004	0.004	0.003	0.011
VBF HH $(2, 1, 1)$	0.011	0.005	0.016	0.005	0.006	0.004	0.017
VBF HH $(-5, 1, 0.5)$	0.027	0.022	0.380	0.021	0.018	0.016	0.062
Single-Higgs background							
ggH	0.493	0.156	0.185	0.595	0.155	0.066	0.040
VBFH	0.092	0.017	0.024	0.044	0.017	0.010	0.015
VH	0.617	0.128	0.099	0.517	0.149	0.043	0.018
$t\bar{t}H$	0.299	0.054	0.030	1.493	0.534	0.119	0.052
tH	0.032	0.015	0.008	0.045	0.009	0.009	0.006
$b\bar{b}H$	0.020	0.006	0.004	0.100	0.037	0.006	0.016
Total	1.554	0.376	0.351	2.794	0.902	0.254	0.148
Continuum background							
$\gamma\gamma + b\bar{b}$	39.3	10.5	7.5	116.5	40.1	13.6	9.6
$\gamma\gamma$ + jets	44.7	9.1	6.5	176.5	42.4	13.1	7.2
Data							
Sideband	38	5	6	156	43	18	8
Full range	54	7	8	202	47	20	11

Signal and Background Modelling B

This appendix collects the per-category signal and background model fits supporting the modelling strategy of Section 7.9 and the results of Section 7.13.

B.1 Signal Model Fits

The Double-Sided Crystal Ball (DSCB) function is fitted to the $m_{\gamma\gamma}$ distribution of the combined SM ggF and VBF HH simulation separately in each of the fourteen analysis categories, as described in Section 7.9. The fit for the Run 2 HM3 category is shown in Figure 7.9; the fits for all remaining categories are collected in Figures B.1–B.4. In each case the upper panel shows the simulated $m_{\gamma\gamma}$ distribution with the best-fit DSCB function overlaid together with the fitted parameters, and the lower panel shows the residual between simulation and fit.

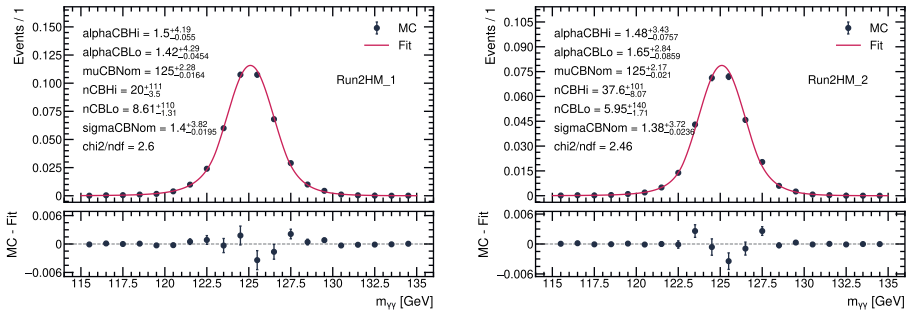

(a) (b)

Figure B.1: DSCB fits to the $m_{\gamma\gamma}$ distribution of the SM ggF + VBF HH simulation in the Run 2 HM1 (a) and HM2 (b) categories.

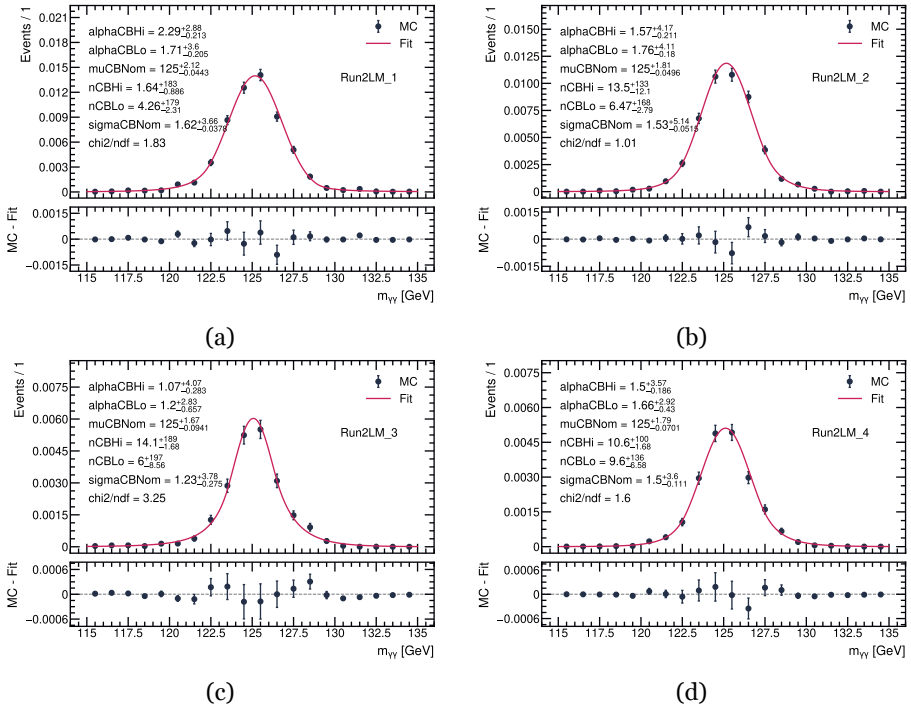

Figure B.2: DSCB fits to the $m_{\gamma\gamma}$ distribution of the SM ggF + VBF HH simulation in the Run 2 LM1 (a), LM2 (b), LM3 (c) and LM4 (d) categories.

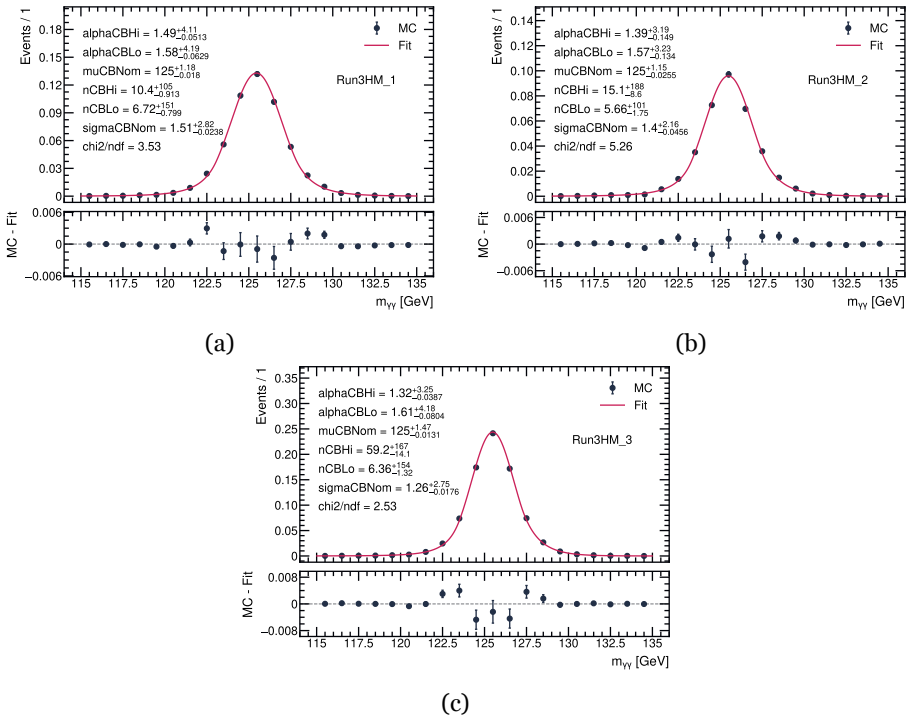

Figure B.3: DSCB fits to the $m_{\gamma\gamma}$ distribution of the SM ggF + VBF HH simulation in the Run 3 HM1 (a), HM2 (b) and HM3 (c) categories.

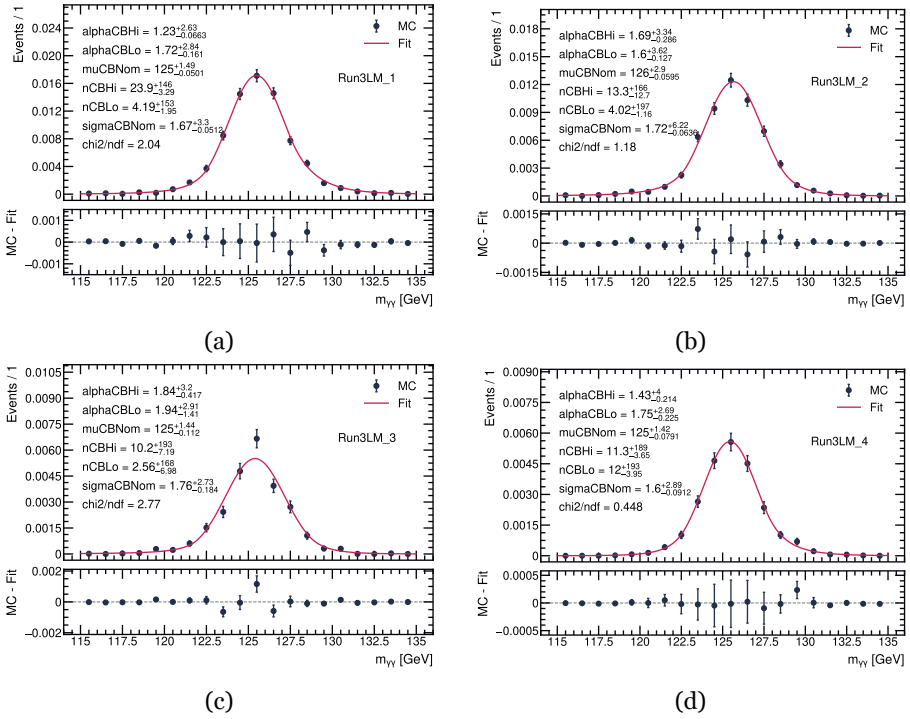

(a) (b)

(c) (d)

Figure B.4: DSCB fits to the $m_{\gamma\gamma}$ distribution of the SM ggF + VBF HH simulation in the Run 3 LM1 (a), LM2 (b), LM3 (c) and LM4 (d) categories.

B.2 Background Template Validation

As described in Section 7.9.2, the $\gamma\gamma$ + jets sample retains too few events after the pre-selection, particularly after the requirement of at least two b-jets, to serve as the background-only template of the spurious signal test, and a dedicated high-statistics $\gamma\gamma$+ heavy flavour sample is used in its place. This section presents the validation of that sample.

The relevant variable for the spurious signal test is $m_{\gamma\gamma}$. Figures B.5–B.8 show the $m_{\gamma\gamma}$ distribution of the two samples in each of the fourteen analysis categories. The two agree within the statistical precision of the $\gamma\gamma$ + jets sample across all categories, which supports the use of the heavy flavour sample as the template.

Distributions of the other variables entering the analysis were also compared, and are found to agree within the same statistical precision. Figures B.9 and B.10 show a representative selection for the Run 2 LM1 and HM3 categories, which are the loosest and tightest categories respectively: Some residual difference is visible in the pseudo-continuous b-tagging (PCBT) score, reflecting the different flavour composition of the two samples. This difference is expected, is within the statistical uncertainty of the $\gamma\gamma$ + jets sample, and does not affect the $m_{\gamma\gamma}$ shape; it therefore does not bias the spurious signal estimate. The remaining categories show the same behaviour.

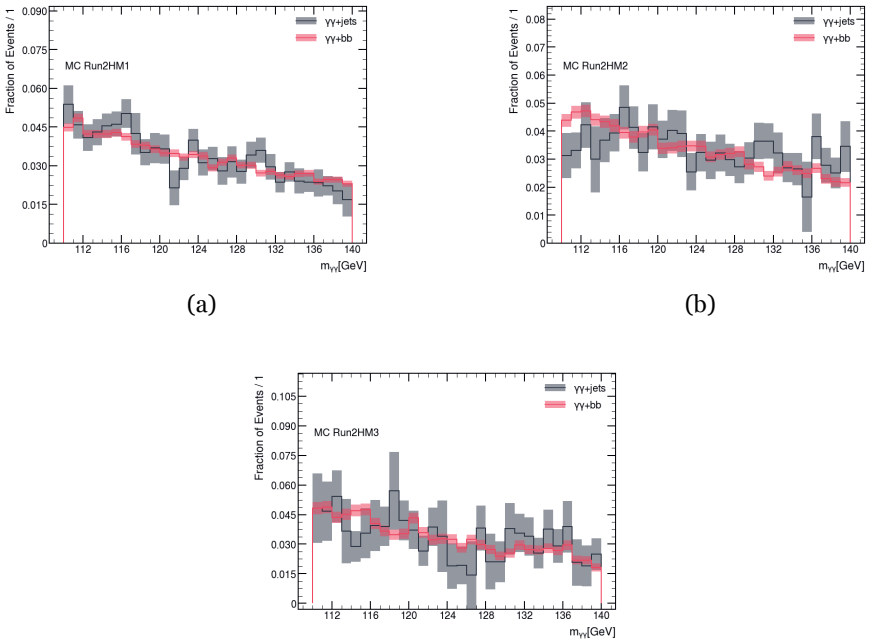

Figure B.5: Comparison of the $m_{\gamma\gamma}$ distribution between the $\gamma\gamma$ + jets sample (grey) and the dedicated $\gamma\gamma$+ heavy flavour sample (red) in the Run 2 HM categories: HM1 (a), HM2 (b), HM3 (c). Distributions are normalised to unit area and the shaded bands show the statistical uncertainty of each sample.

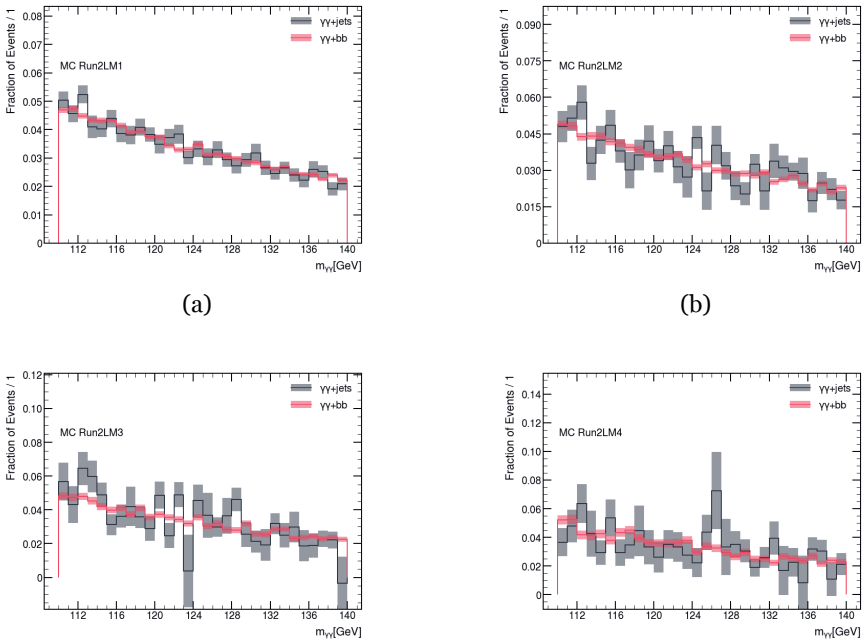

(a) (b)

(c) (d)

Figure B.6: Comparison of the $m_{\gamma\gamma}$ distribution between the $\gamma\gamma$ + jets sample (grey) and the dedicated $\gamma\gamma$+ heavy flavour sample (red) in the Run 2 LM categories: LM1 (a), LM2 (b), LM3 (c), LM4 (d). Distributions are normalised to unit area and the shaded bands show the statistical uncertainty of each sample.

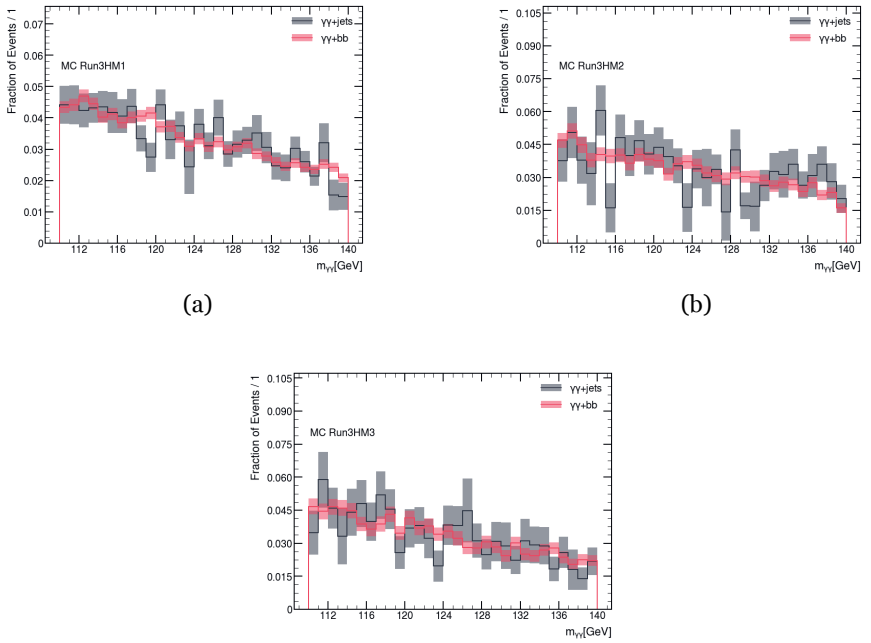

Figure B.7: Comparison of the $m_{\gamma\gamma}$ distribution between the $\gamma\gamma$ + jets sample (grey) and the dedicated $\gamma\gamma$+ heavy flavour sample (red) in the Run 3 HM categories: HM1 (a), HM2 (b), HM3 (c). Distributions are normalised to unit area and the shaded bands show the statistical uncertainty of each sample.

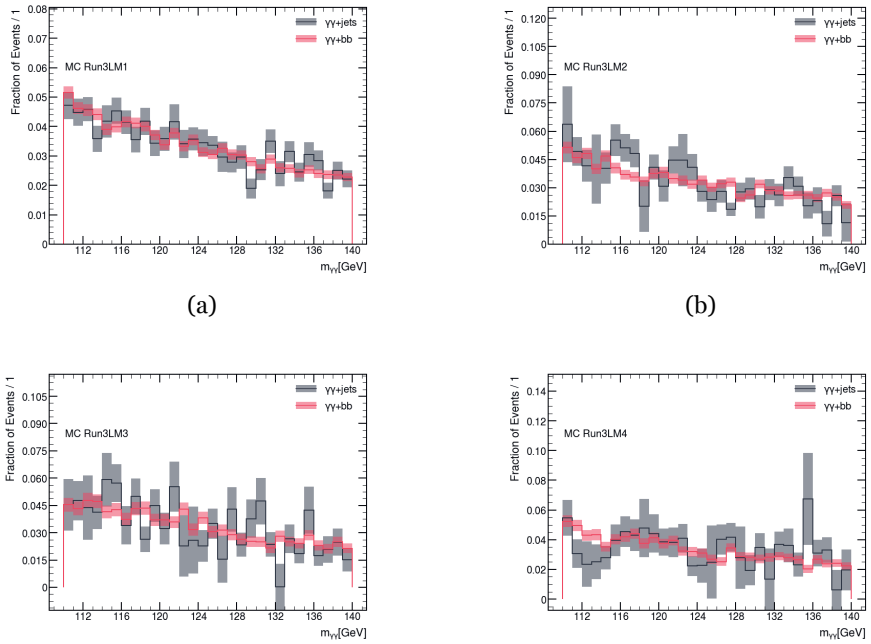

(a) (b)

(c) (d)

Figure B.8: Comparison of the $m_{\gamma\gamma}$ distribution between the $\gamma\gamma$ + jets sample (grey) and the dedicated $\gamma\gamma$+ heavy flavour sample (red) in the Run 3 LM categories: LM1 (a), LM2 (b), LM3 (c), LM4 (d). Distributions are normalised to unit area and the shaded bands show the statistical uncertainty of each sample.

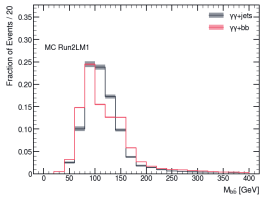

(a)

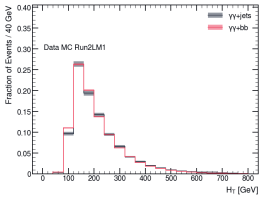

(b)

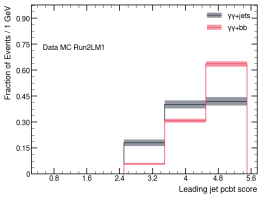

(c)

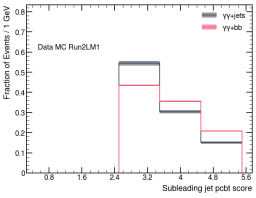

(d)

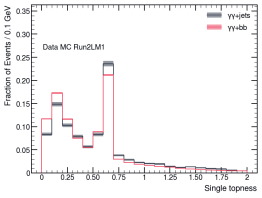

(e)

Figure B.9: Comparison of the $\gamma\gamma$ + jets sample (grey) and the $\gamma\gamma$+ heavy flavour sample (red) in the Run 2 LM1 category for $m_{b\bar{b}}$(a), H_T (b), leading jet PCBT score (c), sub-leading jet PCBT score (d), single-topness (e). Distributions are normalised to unit area and the shaded bands show the statistical uncertainty of each sample.

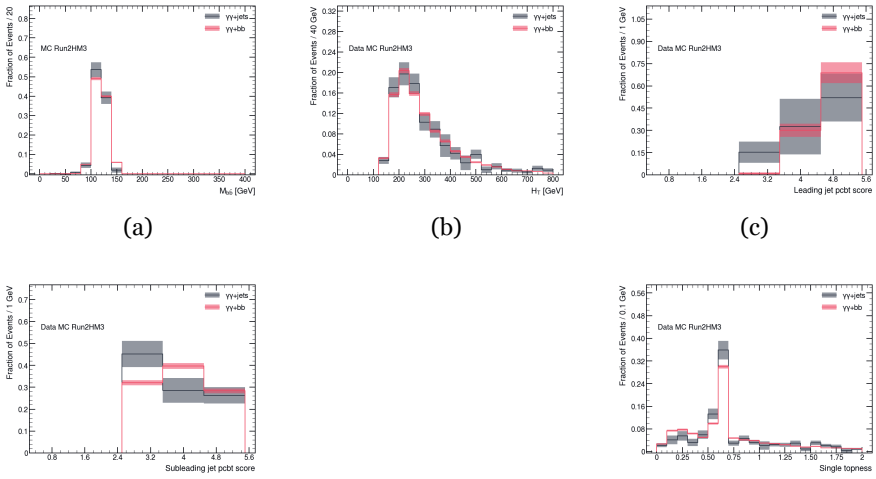

(a) (b) (c)

(d) (e)

Figure B.10: Comparison of the $\gamma\gamma$ + jets sample (grey) and the $\gamma\gamma$+ heavy flavour sample (red) in the Run 2 HM3 category for $m_{b\bar{b}}$(a), H_T (b), leading jet PCBT score (c), sub-leading jet PCBT score (d), single-topness (e). Distributions are normalised to unit area and the shaded bands show the statistical uncertainty of each sample.

B.3 Combined Fits to the Diphoton Invariant Mass

Figures 7.10 and 7.11 show the fits to the $m_{\gamma\gamma}$ distribution in each category separately for Run 2 and Run 3. For completeness, the fits obtained with the two data-taking periods combined are shown here for each of the seven analysis categories.

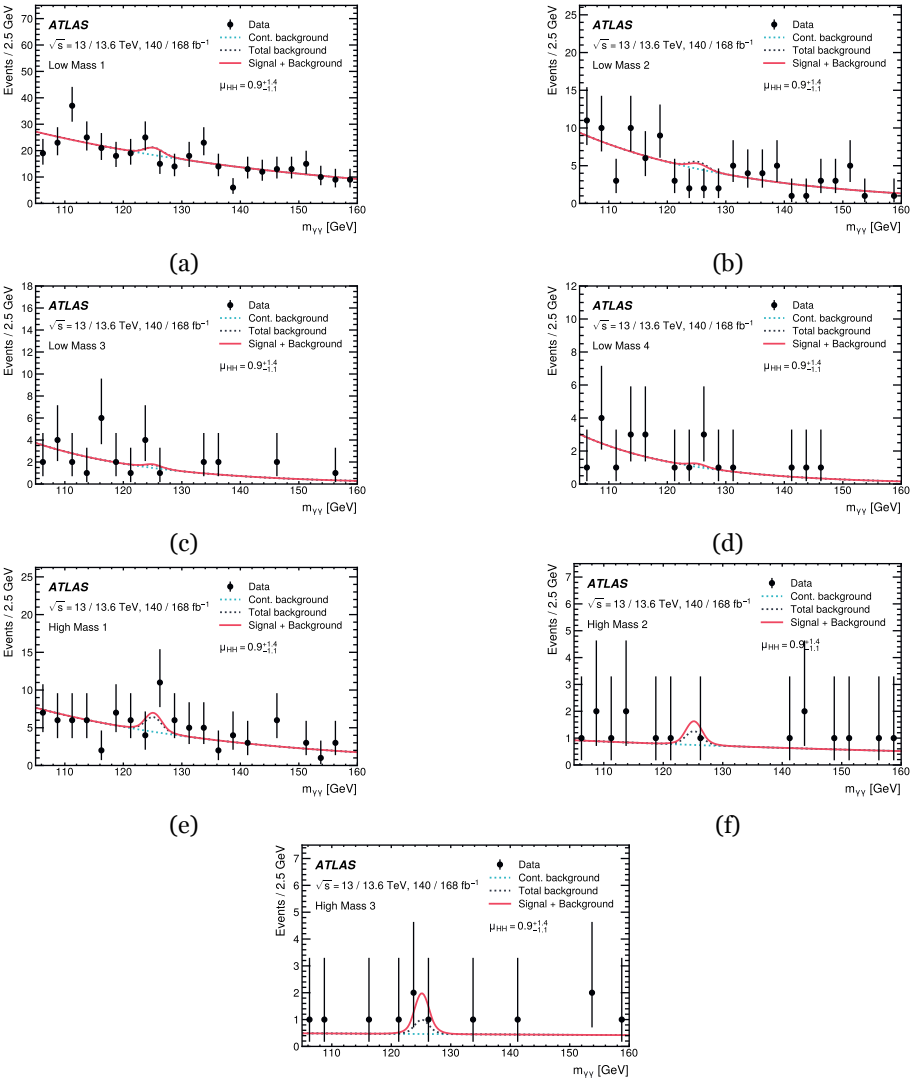

Figure B.11: Fits to the $m_{\gamma\gamma}$ distribution in each analysis category with the Run 2 and Run 3 datasets combined. The dashed blue line shows the continuum background, the dashed black line the total background including the single-Higgs contribution, and the solid red line the signal-plus-background model evaluated at the best-fit signal strength.

Kinematic Distributions of the κ_3, κ_4 Samples C

This appendix collects the remaining kinematic distributions of the truth-level ggF $HH \to b\bar{b}\gamma\gamma$ samples generated for the κ_3, κ_4 study of Chapter 8. Section C.1 shows the distributions for each generated [κ_λ, κ_4] grid point, and Section C.2 shows the remaining comparisons against the nominal $HH \to b\bar{b}\gamma\gamma$ samples used for validation.

C.1 Distributions by Grid Point

The $m^*_{b\bar{b}\gamma\gamma}$ distribution, which carries the bulk of the sensitivity to the self-couplings, is shown in the main text. The remaining kinematic variables show little dependence on either coupling and are collected here. All distributions are normalised to unit area.

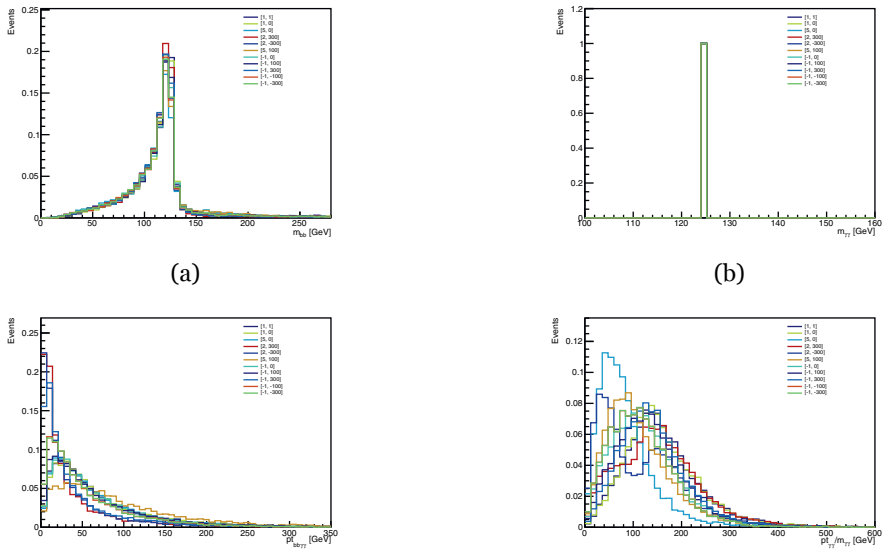

(a) (b)

(c) (d)

Figure C.1: Distributions of (a) $m_{b\bar{b}}$, (b) $m_{\gamma\gamma}$, (c) the transverse momentum of the $b\bar{b}\gamma\gamma$ system and (d) the transverse momentum of the diphoton system, for the ggF $HH \to b\bar{b}\gamma\gamma$ samples generated at $\sqrt{s} = 13\,\text{TeV}$ at each $[\kappa_\lambda, \kappa_4]$ grid point.

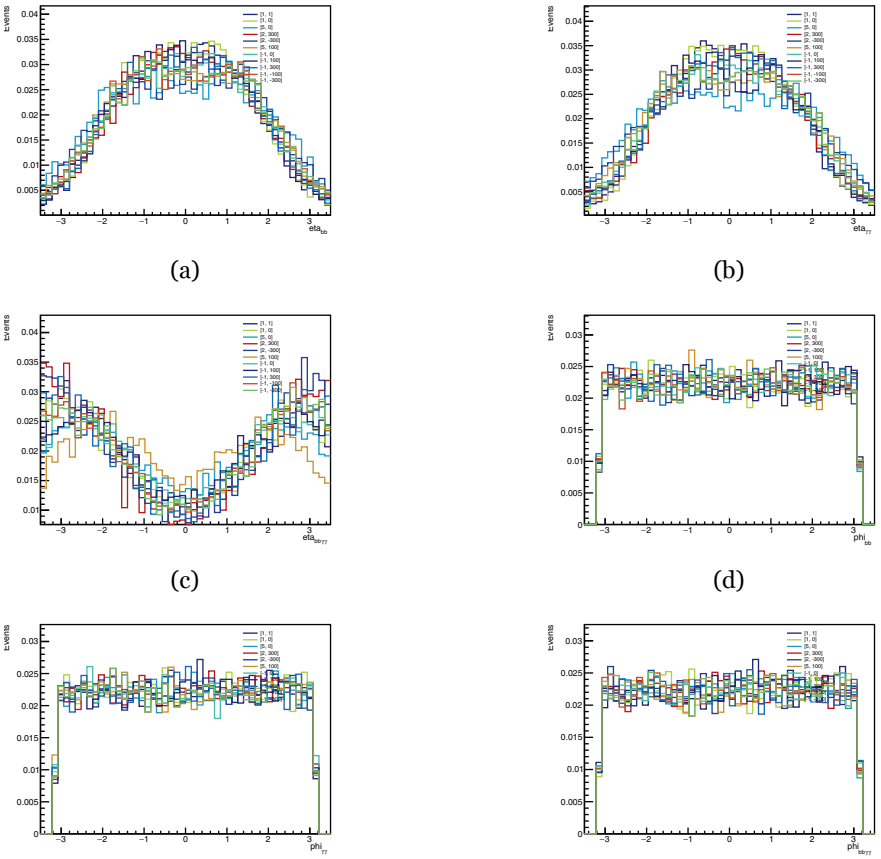

(a) (b) (c) (d) (e) (f)

Figure C.2: Pseudorapidity of the (a) $b\bar{b}$, (b) $\gamma\gamma$ and (c) $b\bar{b}\gamma\gamma$ systems, and azimuthal angle of the (d) $b\bar{b}$, (e) $\gamma\gamma$ and (f) $b\bar{b}\gamma\gamma$ systems, for the ggF $HH \to b\bar{b}\gamma\gamma$ samples generated at each [κ_λ, κ_4] grid point.

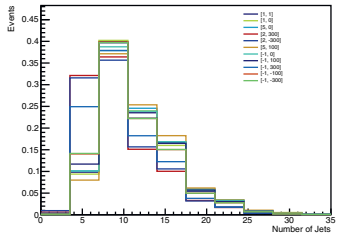

(a)

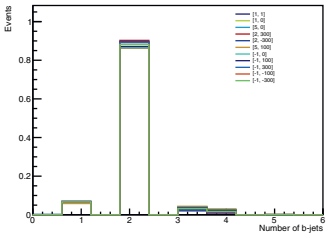

(b)

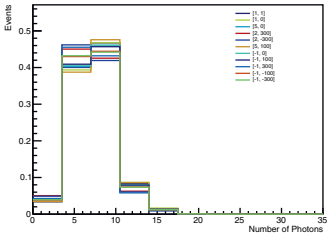

(c)

Figure C.3: Multiplicity of (a) jets, (b) b-jets and (c) photons for the ggF $HH \to b\bar{b}\gamma\gamma$ samples generated at each [κ_λ, κ_4] grid point.

C.2 Validation Against the Nominal Samples

The $m_{b\bar{b}}$ and jet multiplicity comparisons against the nominal $HH \to b\bar{b}\gamma\gamma$ samples are shown in the main text. The corresponding comparisons for $m^*_{b\bar{b}\gamma\gamma}$ and the b-jet multiplicity are collected here. In each case the newly generated $\kappa_4 = 0$ sample is compared to the nominal sample at the equivalent value of κ_λ, with the p-value, χ^2 and ratio shown.

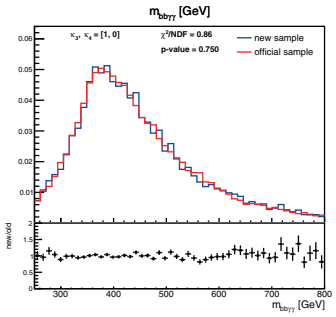

(a)

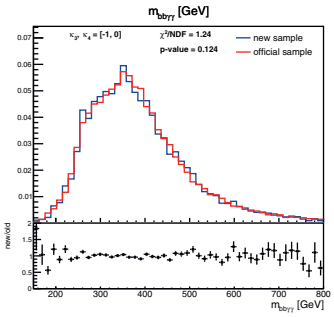

(b)

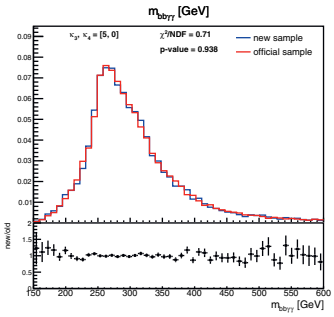

(c)

Figure C.4: Comparison of the $m^*_{b\bar{b}\gamma\gamma}$ distribution between the newly generated samples and the nominal $HH \to b\bar{b}\gamma\gamma$ samples for $\kappa_\lambda = 1$ (a), $\kappa_\lambda = -1$ (b) and $\kappa_\lambda = 5$ (c), all at $\kappa_4 = 0$.

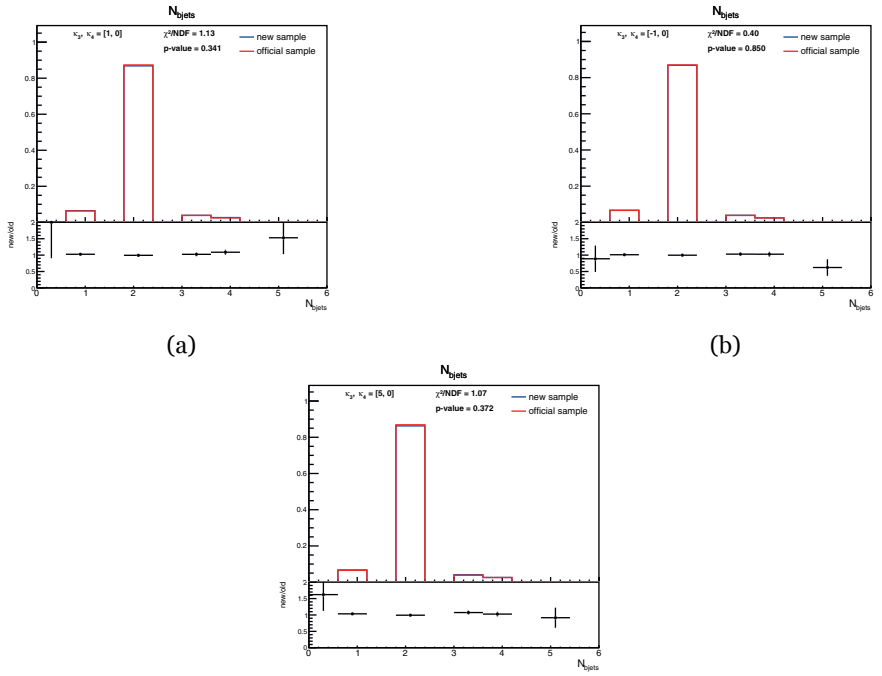

Figure C.5: Comparison of the b-jet multiplicity between the newly generated samples and the nominal $HH \to b\bar{b}\gamma\gamma$ samples for $\kappa_\lambda = 1$ (a), $\kappa_\lambda = -1$ (b) and $\kappa_\lambda = 5$ (c), all at $\kappa_4 = 0$.

Bibliography

[1] ATLAS Collaboration. Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC. *Physics Letters B*, 716(1):1–29, September 2012. doi: 10.1016/j.physletb.2012.08.020.

[2] CMS Collaboration. 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. doi: 10.1016/j.physletb.2012.08.021.

[3] David Griffiths. *Introduction to Elementary Particles*. Wiley-VCH, Weinheim, Germany, 2008.

[4] Mark Thomson. *Modern Particle Physics*. Cambridge University Press, 2013.

[5] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to Quantum Field Theory*. Addison-Wesley, Massachusetts, USA, 1995.

[6] CERN and Andrew Purcell. Go on a particle quest at the first CERN webfest, 2012. URL `https://cds.cern.ch/record/1473657`. Accessed April 2026.

[7] Murray Gell-Mann. Symmetries of Baryons and Mesons. *Phys. Rev.*, 125:1067–1084, Feb 1962. doi: 10.1103/PhysRev.125.1067.

[8] Kenneth G. Wilson. Confinement of quarks. *Phys. Rev. D*, 10:2445–2459, Oct 1974. doi: 10.1103/PhysRevD.10.2445.

[9] David J. Gross and Frank Wilczek. Ultraviolet Behavior of Non-Abelian Gauge Theories. *Phys. Rev. Lett.*, 30:1343–1346, Jun 1973. doi: 10.1103/PhysRevLett.30.1343.

[10] G. 't Hooft and M. Veltman. Regularization and Renormalization of Gauge Fields. *Nuclear Physics B*, 44(1):189–213, 1972. doi: 10.1016/0550-3213(72)90279-9.

[11] Nicola Cabibbo. Unitary Symmetry and Leptonic Decays. *Phys. Rev. Lett.*, 10:531–533, Jun 1963. doi: 10.1103/PhysRevLett.10.531.

[12] Makoto Kobayashi and Toshihide Maskawa. CP-Violation in the Renormalizable Theory of Weak Interaction. *Progress of Theoretical Physics*, 49(2):652–657, 02 1973. doi: 10.1143/PTP.49.652.

[13] T. D. Lee and C. N. Yang. Question of Parity Conservation in Weak Interactions. *Phys. Rev.*, 104:254–258, Oct 1956. doi: 10.1103/PhysRev.104.254.

[14] C. S. Wu et al. Experimental Test of Parity Conservation in Beta Decay. *Phys. Rev.*, 105:1413–1415, Feb 1957. doi: 10.1103/PhysRev.105.1413.

[15] Sheldon L. Glashow. Partial-symmetries of weak interactions. *Nuclear Physics*, 22(4):579–588, 1961. doi: 10.1016/0029-5582(61)90469-2.

[16] Steven Weinberg. A Model of Leptons. *Phys. Rev. Lett.*, 19:1264–1266, Nov 1967. doi: 10.1103/PhysRevLett.19.1264.

[17] Abdus Salam. Weak and Electromagnetic Interactions. In N. Svartholm, editor, *Elementary Particle Theory: Relativistic Groups and Analyticity (Nobel Symposium No. 8)*, pages 367–377, Stockholm, 1968. Almqvist & Wiksell.

[18] F. Englert and R. Brout. Broken Symmetry and the Mass of Gauge Vector Mesons. *Phys. Rev. Lett.*, 13:321–323, Aug 1964. doi: 10.1103/PhysRevLett.13.321.

[19] G. S. Guralnik et al. Global Conservation Laws and Massless Particles. *Phys. Rev. Lett.*, 13:585–587, Nov 1964. doi: 10.1103/PhysRevLett.13.585.

[20] Peter W. Higgs. Broken Symmetries and the Masses of Gauge Bosons. *Phys. Rev. Lett.*, 13:508–509, Oct 1964. doi: 10.1103/PhysRevLett.13.508.

[21] Riebesell, Janosh. Higgs potential – TikZ.net, 2022. URL `https://tikz.net/higgs-potential/`. Accessed April 2026.

[22] J. Goldstone. Field theories with «Superconductor »solutions. *Il Nuovo Cimento (1955-1965)*, 19(1):154–164, 1961. doi: 10.1007/BF02812722.

[23] ATLAS Collaboration. Combined measurement of the Higgs boson mass from the $H \to \gamma\gamma$ and $H \to ZZ^* \to 4\ell$ decay channels with the ATLAS detector using $\sqrt{s}$ = 7, 8 and 13 TeV pp collision data. *Phys. Rev. Lett.*, 131:251802, 2023. doi: 10.1103/PhysRevLett.131.251802.

[24] D. de Florian et al. *Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector*, volume 2 of *CERN Yellow Reports: Monographs*. CERN, Geneva, 2017. doi: 10.23731/CYRM-2017-002.

[25] Vera C. Rubin and Jr. Ford, W. Kent. Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions. *Astrophys. J.*, 159: 379–403, 1970. doi: 10.1086/150317.

[26] Planck Collaboration. Planck 2018 results. VI. Cosmological parameters. *Astron. Astrophys.*, 641:A6, 2020. doi: 10.1051/0004-6361/201833910.

[27] A. D. Sakharov. Violation of CP Invariance, C Asymmetry, and Baryon Asymmetry of the Universe. *JETP Lett.*, 5:24–27, 1967. doi: 10.1070/PU1991v034n05ABEH002497.

[28] Ziro Maki et al. Remarks on the Unified Model of Elementary Particles. *Prog. Theor. Phys.*, 28:870–880, 1962. doi: 10.1143/PTP.28.870.

[29] Super-Kamiokande Collaboration. Evidence for Oscillation of Atmospheric Neutrinos. *Phys. Rev. Lett.*, 81:1562–1567, 1998. doi: 10.1103/PhysRevLett.81.1562.

[30] SNO Collaboration. Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory. *Phys. Rev. Lett.*, 89:011301, 2002. doi: 10.1103/PhysRevLett.89.011301.

[31] Gerard 't Hooft. Naturalness, Chiral Symmetry, and Spontaneous Chiral Symmetry Breaking. In *Recent Developments in Gauge Theories*, volume 59 of *NATO Sci. Ser. B*, pages 135–157. Plenum Press, 1980. doi: 10.1007/978-1-4684-7571-5_9.

[32] Fabio Maltoni et al. Constraining the Higgs self-couplings at e^+e^- colliders. *JHEP*, 07:087, 2018. doi: 10.1007/JHEP07(2018)087.

[33] Giuseppe Degrassi et al. Higgs mass and vacuum stability in the Standard Model at NNLO. *JHEP*, 08:098, 2012. doi: 10.1007/JHEP08(2012)098.

[34] Lyndon Evans and Philip Bryant. LHC Machine. *Journal of Instrumentation*, 3(08):S08001, aug 2008. doi: 10.1088/1748-0221/3/08/S08001.

[35] Oliver Sim Brüning et al. *LHC Design Report*. CERN Yellow Reports: Monographs. CERN, Geneva, 2004. doi: 10.5170/CERN-2004-003-V-1.

[36] ATLAS Collaboration. The ATLAS Experiment at the CERN Large Hadron Collider. *Journal of Instrumentation*, 3(08):S08003, aug 2008. doi: 10.1088/1748-0221/3/08/S08003.

[37] Christiane Lefèvre. The CERN accelerator complex, 2008. URL https://cds.cern.ch/record/1260465.

[38] CMS Collaboration. The CMS experiment at the CERN LHC. *Journal of Instrumentation*, 3(08):S08004, aug 2008. doi: 10.1088/1748-0221/3/08/S08004.

[39] LHCb Collaboration. The LHCb Detector at the LHC. *Journal of Instrumentation*, 3(08):S08005, aug 2008. doi: 10.1088/1748-0221/3/08/S08005.

[40] ALICE Collaboration. The ALICE experiment at the CERN LHC. *Journal of Instrumentation*, 3(08):S08002, aug 2008. doi: 10.1088/1748-0221/3/08/S08002.

[41] Stephen Myers and Emilio Picasso. The design, construction and commissioning of the CERN large Electron–Positron collider. *Contemporary Physics*, 31(6):387–403, 1990. doi: 10.1080/00107519008213789.

[42] ATLAS Collaboration. ATLAS Public Luminosity Results for Run 2, 2026. URL https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2. Accessed 18 August 2026.

[43] ATLAS Collaboration. ATLAS Public Luminosity Results for Run 3, 2026. URL https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun3. Accessed 6 August 2026.

[44] ATLAS Collaboration. The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run-3. *Journal of Instrumentation*, 19(05):P05063, 2024. doi: 10.1088/1748-0221/19/05/P05063.

[45] Izaak Neutelings. CMS coordinate system (axis3d_cms), 2024. URL https://tikz.net/axis3d_cms/. Licensed under CC BY-SA 4.0. Accessed 12 March 2026.

[46] ATLAS Collaboration. Operation of the ATLAS trigger system in Run 2. *Journal of Instrumentation*, 15(10):P10004, oct 2020. doi: 10.1088/1748-0221/15/10/P10004.

[47] ATLAS Collaboration. The ATLAS trigger system for LHC Run 3 and trigger performance in 2022. *Journal of Instrumentation*, 19(06):P06029, jun 2024. doi: 10.1088/1748-0221/19/06/P06029.

[48] ATLAS Collaboration. Trigger Operation Public Results, 2026. URL https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TriggerOperationPublicResults. Accessed 9 February 2026.

[49] Laura Gonella. The ATLAS ITk detector system for the Phase-II LHC upgrade. *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, 1045:167597, 2023. doi: 10.1016/j.nima.2022.167597.

[50] ATLAS Collaboration. Technical Design Report for the ATLAS Inner Tracker Strip Detector. Technical Design Report CERN-LHCC-2017-005, ATLAS-TDR-025, CERN, Geneva, 2017.

[51] Graham Beck et al. Thermo-electrical modelling of the ATLAS ITk Strip Detector. *Nucl. Instrum. Meth. A*, 969:164023, 2020. doi: 10.1016/j.nima.2020.164023.

[52] ATLAS Collaboration. Technical Design Report for the Phase-II Upgrade of the ATLAS Muon Spectrometer. Technical Design Report CERN-LHCC-2017-017, ATLAS-TDR-026, CERN, Geneva, 2022.

[53] Joao Pequenao and Paul Schaffner. How ATLAS detects particles: diagram of particle paths in the detector. Accessed May 2026, 2013. URL https://cds.cern.ch/record/1505342.

[54] ATLAS Collaboration. Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2. *Eur. Phys. J. C*, 77: 673, 2017. doi: 10.1140/epjc/s10052-017-5225-7.

[55] ATLAS Collaboration. Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC. *Eur. Phys. J. C*, 77:332, 2017. doi: 10.1140/epjc/s10052-017-4887-5.

[56] ATLAS Collaboration. Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1. *Eur. Phys. J. C*, 77:490, 2017. doi: 10.1140/epjc/s10052-017-5004-5.

[57] ATLAS Collaboration. Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data. *Journal of Instrumentation*, 14(12):P12006, dec 2019. doi: 10.1088/1748-0221/14/12/P12006.

[58] ATLAS Collaboration. Jet reconstruction and performance using particle flow with the ATLAS Detector. *The European Physical Journal C*, 77 (7), July 2017. doi: 10.1140/epjc/s10052-017-5031-2.

[59] Matteo Cacciari et al. The anti-k_tjet clustering algorithm. *Journal of High Energy Physics*, 2008(04):063–063, April 2008. doi: 10.1088/1126-6708/2008/04/063.

[60] ATLAS Collaboration. Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *The European Physical Journal C*, 81(8):689, 2021. doi: 10.1140/epjc/s10052-021-09402-3.

[61] ATLAS Collaboration. Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector. *Eur. Phys. J. C*, 76:581, 2016. doi: 10.1140/epjc/s10052-016-4395-z.

[62] ATLAS Collaboration. Identification and rejection of pile-up jets at high pseudorapidity with the ATLAS detector. *Eur. Phys. J. C*, 77:580, 2017. doi: 10.1140/epjc/s10052-017-5081-5. [Erratum: Eur. Phys. J. C 77 (2017) 712].

[63] ATLAS Collaboration. Transforming jet flavour tagging at ATLAS. *Nature Communications*, 17(1):541, 2026. doi: 10.1038/s41467-025-65059-6.

[64] Ashish Vaswani et al. Attention Is All You Need. In *Advances in Neural Information Processing Systems 30*, pages 5998–6008, 2017. URL https://arxiv.org/abs/1706.03762.

[65] ATLAS Collaboration. Neural Network Jet Flavour Tagging with the Upgraded ATLAS Inner Tracker Detector at the High-Luminosity LHC. Technical Report ATL-PHYS-PUB-2022-047, CERN, Geneva, 2022. URL https://cds.cern.ch/record/2839913.

[66] ATLAS Collaboration. Jet Flavour Tagging With GN1 and DL1d. Generator dependence, Run 2 and Run 3 data agreement studies, 2023. URL https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/FTAG-2023-01/. ATLAS public plots FTAG-2023-01. Accessed 22 June 2026.

[67] ATLAS Collaboration. Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV. *The European Physical Journal C*, 81(7):578, 2021. doi: 10.1140/epjc/s10052-021-09233-2.

[68] ATLAS Collaboration. Performance of missing transverse momentum reconstruction with the ATLAS detector using proton–proton collisions at $\sqrt{s} = 13$ TeV. *The European Physical Journal C*, 78(11):903, 2018. doi: 10.1140/epjc/s10052-018-6288-9.

[69] ATLAS Collaboration. Constraints on Higgs boson production with large transverse momentum using $H \to b\bar{b}$ decays in the ATLAS detector. *Phys. Rev. D*, 105:092003, May 2022. doi: 10.1103/PhysRevD.105.092003.

[70] ATLAS Collaboration. Transformer Neural Networks for Identifying Boosted Higgs Bosons decaying into $b\bar{b}$ and $c\bar{c}$ in ATLAS. Technical Report ATL-PHYS-PUB-2023-021, CERN, Geneva, 2023. URL https://cds.cern.ch/record/2866601.

[71] Stefan Höche. Introduction to parton-shower event generators. In *Journeys Through the Precision Frontier: Amplitudes for Colliders (TASI 2014)*, pages 235–295. World Scientific, 2015. doi: 10.1142/9789814678766_0005.

[72] John C. Collins et al. Factorization of Hard Processes in QCD. *Adv. Ser. Direct. High Energy Phys.*, 5:1–91, 1989. doi: 10.1142/9789814503266_0001.

[73] Max Klein and Rik Yoshida. Collider physics at HERA. *Progress in Particle and Nuclear Physics*, 61(2):343–393, 2008. doi: 10.1016/j.ppnp.2008.05.002.

[74] V. N. Gribov and L. N. Lipatov. Deep inelastic e p scattering in perturbation theory. *Sov. J. Nucl. Phys.*, 15:438–450, 1972.

[75] Yuri L. Dokshitzer. Calculation of the Structure Functions for Deep Inelastic Scattering and e+ e- Annihilation by Perturbation Theory in Quantum Chromodynamics. *Sov. Phys. JETP*, 46:641–653, 1977.

[76] Guido Altarelli and Giorgio Parisi. Asymptotic Freedom in Parton Language. *Nucl. Phys. B*, 126:298–318, 1977. doi: 10.1016/0550-3213(77)90384-4.

[77] Richard D. Ball et al. The PDF4LHC21 combination of global PDF fits for the LHC Run III. *J. Phys. G*, 49(8):080501, 2022. doi: 10.1088/1361-6471/ac7216.

[78] Richard D. Ball et al. Parton distributions for the LHC run II. *Journal of High Energy Physics*, 2015(4):40, 2015. doi: 10.1007/JHEP04(2015)040.

[79] Richard D. Ball et al. Parton distributions with QED corrections. *Nucl. Phys. B*, 877:290–320, 2013. doi: 10.1016/j.nuclphysb.2013.10.010.

[80] Stefano Frixione et al. Matching NLO QCD computations with Parton Shower simulations: the POWHEG method. *JHEP*, 11:070, 2007. doi: 10.1088/1126-6708/2007/11/070.

[81] Bo Andersson. *The Lund Model*, volume 7. Cambridge University Press, 1998. ISBN 978-1-009-40129-6, 978-1-009-40125-8, 978-1-009-40128-9, 978-0-521-01734-3, 978-0-521-42094-5, 978-0-511-88149-7. doi: 10.1017/9781009401296.

[82] Simone Alioli et al. A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX. *Journal of High Energy Physics*, 2010(6):43, 2010. doi: 10.1007/JHEP06(2010)043.

[83] J. Alwall et al. The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton

shower simulations. *Journal of High Energy Physics*, 2014(7):79, 2014. doi: 10.1007/JHEP07(2014)079.

[84] Christian Bierlich et al. A comprehensive guide to the physics and usage of PYTHIA 8.3. *SciPost Phys. Codebases*, page 8, 2022. doi: 10.21468/SciPostPhysCodeb.8.

[85] David 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. doi: 10.1016/S0168-9002(01)00089-4. BEAUTY2000, Proceedings of the 7th Int. Conf. on B-Physics at Hadron Machines.

[86] Enrico Bothmann et al. Event generation with Sherpa 2.2. *SciPost Phys.*, 7:034, 2019. doi: 10.21468/SciPostPhys.7.3.034.

[87] S. Agostinelli 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. doi: 10.1016/S0168-9002(03)01368-8.

[88] ATLAS Collaboration. AtlFast3: The Next Generation of Fast Simulation in ATLAS. *Computing and Software for Big Science*, 6(1):7, 2022. doi: 10.1007/s41781-021-00079-7.

[89] André David et al. LHC HXSWG interim recommendations to explore the coupling structure of a Higgs-like particle. Technical Report CERN-LHCHXSWG-2012-001, CERN, 2012. URL `https://arxiv.org/abs/1209.0040`.

[90] ATLAS Collaboration. Measurement of Higgs boson decay into b-quarks in associated production with a top-quark pair in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Journal of High Energy Physics*, 2022 (6), June 2022. doi: 10.1007/jhep06(2022)097.

[91] ATLAS Collaboration. Measurements of Higgs bosons decaying to bottom quarks from vector boson fusion production with the ATLAS experiment at $\sqrt{s} = 13$ TeV. *The European Physical Journal C*, 81(6):537, 2021. doi: 10.1140/epjc/s10052-021-09192-8.

[92] ATLAS Collaboration. Measurements of Higgs boson production cross-sections in the $H \to \tau^+\tau^-$ decay channel in pp collisions at $\sqrt{s} = 13$

TeV with the ATLAS detector. *Journal of High Energy Physics*, 2022 (8), August 2022. doi: 10.1007/jhep08(2022)175.

[93] ATLAS Collaboration. Evidence for the Dimuon Decay of the Higgs Boson in pp Collisions with the ATLAS Detector. *Phys. Rev. Lett.*, 135: 231802, Dec 2025. doi: 10.1103/gzdh-p159.

[94] ATLAS Collaboration. Observation of Higgs boson production in association with a top quark pair at the LHC with the ATLAS detector. *Physics Letters B*, 784:173–191, September 2018. doi: 10.1016/j.physletb.2018.07.035.

[95] ATLAS Collaboration. A detailed map of Higgs boson interactions by the ATLAS experiment ten years after the discovery. *Nature*, 607(7917): 52–59, 2022. doi: 10.1038/s41586-022-04893-w.

[96] ATLAS Collaboration. Combination of Searches for Higgs Boson Pair Production in pp Collisions at $\sqrt{s} = 13$ TeV with the ATLAS Detector. *Phys. Rev. Lett.*, 133:101801, Sep 2024. doi: 10.1103/PhysRevLett.133.101801.

[97] R. Frederix et al. Higgs pair production at the LHC with NLO and parton-shower effects. *Physics Letters B*, 732:142–149, 2014. doi: 10.1016/j.physletb.2014.03.026.

[98] Biagio Di Micco et al. Higgs boson potential at colliders: Status and perspectives. *Reviews in Physics*, 5:100045, 2020. doi: 10.1016/j.revip.2020.100045.

[99] ATLAS Collaboration. Search for pair production of boosted Higgs bosons via vector-boson fusion in the $b\bar{b}b\bar{b}$ final state using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Lett. B*, 858:139007, 2024. doi: 10.1016/j.physletb.2024.139007.

[100] ATLAS Collaboration. Search for nonresonant 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. *Phys. Rev. D*, 108(5):052003, 2023. doi: 10.1103/PhysRevD.108.052003.

[101] ATLAS Collaboration. Search for the nonresonant production of Higgs boson pairs via gluon fusion and vector-boson fusion in the $b\bar{b}\tau^+\tau^-$ final state in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS

detector. *Phys. Rev. D*, 110(3):032012, 2024. doi: 10.1103/PhysRevD.110.032012.

[102] ATLAS Collaboration. Search for non-resonant Higgs boson pair production in the $2b + 2\ell + E_{\mathrm{T}}^{\mathrm{miss}}$ final state in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Journal of High Energy Physics*, 2024(2):37, 2024. doi: 10.1007/JHEP02(2024)037.

[103] ATLAS Collaboration. Search for non-resonant Higgs boson pair production in final states with leptons, taus, and photons in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Journal of High Energy Physics*, 2024(8):164, 2024. doi: 10.1007/JHEP08(2024)164.

[104] ATLAS Collaboration. Study of Higgs boson pair production in the $HH \to b\bar{b}\gamma\gamma$ final state with 308 fb^{-1} of data collected at $\sqrt{s} = 13$ TeV and 13.6 TeV by the ATLAS experiment. *Physics Letters B*, 876:140280, 2026. doi: 10.1016/j.physletb.2026.140280.

[105] ATLAS Collaboration. Search for triple Higgs boson production in the $6b$ final state using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Rev. D*, 111:032006, Feb 2025. doi: 10.1103/PhysRevD.111.032006.

[106] ATLAS Collaboration. Studies of new Higgs boson interactions through nonresonant HH production in the $b\bar{b}\gamma\gamma$ final state in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Journal of High Energy Physics*, 2024(1):66, 2024. doi: 10.1007/JHEP01(2024)066.

[107] ATLAS Collaboration. ATLAS data quality operations and performance for 2015-2018 data-taking. *JINST*, 15(04):P04003, 2020. doi: 10.1088/1748-0221/15/04/P04003.

[108] ATLAS Collaboration. Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC. *The European Physical Journal C*, 83(10):982, 2023. doi: 10.1140/epjc/s10052-023-11747-w.

[109] ATLAS Collaboration. Preliminary analysis of the luminosity calibration for the ATLAS 13.6 TeV data recorded in 2023. ATLAS PUB Note ATL-DAPR-PUB-2024-001, CERN, Geneva, 2024. URL `https://cds.cern.ch/record/2900949`.

[110] ATLAS Collaboration. ATLAS Pythia 8 tunes to 7 TeV data. Technical Report ATL-PHYS-PUB-2014-021, CERN, Geneva, 2014. URL https://cds.cern.ch/record/1966419.

[111] ATLAS Collaboration. Performance of electron and photon triggers in ATLAS during LHC Run 2. *The European Physical Journal C*, 80(1):47, 2020. doi: 10.1140/epjc/s10052-019-7500-2.

[112] ATLAS Collaboration. Measurement of the properties of Higgs boson production at $\sqrt{s} = 13$ TeV in the $H \to \gamma\gamma$ channel using 139 fb^{-1} of pp collision data with the ATLAS experiment. *Journal of High Energy Physics*, 2023(7):88, 2023. doi: 10.1007/JHEP07(2023)088.

[113] Peilong Wang. *Observation and Measurement of the Higgs Boson Produced in Association with a Vector Boson and Decaying into Bottom Quarks*. Phd thesis, CERN, Switzerland, 2024. URL https://inspirehep.net/literature/1841340. INSPIRE record 1841340.

[114] Mohamed Belfkir. *Search for Higgs pair production at LHC collider (CERN): The first measurement for Higgs potential and search for new physics*. Phd thesis, Université Savoie Mont Blanc, France, September 2021. CERN-THESIS-2021-221.

[115] Federico Montereali. *Search for Higgs boson pair production in the two photons plus two bottom quarks final state using partial Run-3 proton-proton collision data with the ATLAS detector*. PhD thesis, Università degli Studi di Roma Tre, February 2025. CERN-THESIS-2025-031.

[116] Romano Orlandini. Study of a kinematic fit algorithm for the $HH \to b\bar{b}\gamma\gamma$ channel reconstruction in the ATLAS experiment at the LHC. Master's thesis, University of Roma Tre, 2023.

[117] F. James and M. Roos. Minuit - a system for function minimization and analysis of the parameter errors and correlations. *Computer Physics Communications*, 10(6):343–367, 1975. doi: 10.1016/0010-4655(75)90039-9.

[118] Tianqi Chen and Carlos Guestrin. XGBoost: A Scalable Tree Boosting System. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, KDD '16, pages 785–794, New York, NY, USA, 2016. Association for Computing Machinery. doi: 10.1145/2939672.2939785.

[119] ATLAS Collaboration. Measurements of Higgs boson properties in the diphoton decay channel with 36 fb^{-1} of pp collision data at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Phys. Rev. D*, 98:052005, Sep 2018. doi: 10.1103/PhysRevD.98.052005.

[120] ATLAS Collaboration. Recommendations for the Modeling of Smooth Backgrounds. Technical Report ATL-PHYS-PUB-2020-028, CERN, 2020. URL https://cds.cern.ch/record/2743717.

[121] ATLAS Collaboration. Measurement of the $H \to \gamma\gamma$ and $H \to ZZ^* \to 4\ell$ cross-sections in pp collisions at $\sqrt{s} = 13.6$ TeV with the ATLAS detector. *The European Physical Journal C*, 84(1):78, 2024. doi: 10.1140/epjc/s10052-023-12130-5.

[122] ATLAS Collaboration. Measurement of the c-jet mistagging efficiency in $t\bar{t}$ events using pp collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS detector. *The European Physical Journal C*, 82(1):95, 2022. doi: 10.1140/epjc/s10052-021-09843-w.

[123] ATLAS Collaboration. Calibration of the light-flavour jet mistagging efficiency of the b-tagging algorithms with Z+jets events using 139 fb^{-1} of ATLAS proton–proton collision data at $\sqrt{s} = 13$ TeV. *The European Physical Journal C*, 83(8):728, 2023. doi: 10.1140/epjc/s10052-023-11736-z.

[124] ATLAS Collaboration. ATLAS flavour-tagging algorithms for the LHC Run 2 pp collision dataset. *The European Physical Journal C*, 83(7):681, 2023. doi: 10.1140/epjc/s10052-023-11699-1.

[125] Bradley Efron. Bootstrap Methods: Another Look at the Jackknife. *The Annals of Statistics*, 7(1):1–26, 1979. doi: 10.1214/aos/1176344552.

[126] ATLAS Collaboration. Measurements of inclusive and differential fiducial cross-sections of $t\bar{t}$ production with additional heavy-flavour jets in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector. *Journal of High Energy Physics*, 2019(4):46, 2019. doi: 10.1007/JHEP04(2019)046.

[127] ATLAS Collaboration. Measurement of the cross-section for W boson production in association with b-jets in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector. *Journal of High Energy Physics*, 2013(6):84, 2013. doi: 10.1007/JHEP06(2013)084.

[128] Glen Cowan et al. Asymptotic formulae for likelihood-based tests of new physics. *The European Physical Journal C*, 71(2):1554, 2011. doi: 10.1140/epjc/s10052-011-1554-0.

[129] A L Read. Presentation of search results: the CLs technique. *Journal of Physics G: Nuclear and Particle Physics*, 28(10):2693, sep 2002. doi: 10.1088/0954-3899/28/10/313.

[130] ATLAS Collaboration. Projected sensitivity of measurements of Higgs boson pair production with the ATLAS experiment at the HL-LHC. Technical Report ATL-PHYS-PUB-2025-006, CERN, Geneva, 2025. URL https://cds.cern.ch/record/2925853.

[131] G. Degrassi et al. Probing the Higgs self coupling via single Higgs production at the LHC. *Journal of High Energy Physics*, 2016(12), December 2016. doi: 10.1007/jhep12(2016)080.

[132] ATLAS Collaboration. Constraints on the Higgs boson self-coupling from single- and double-Higgs production with the ATLAS detector using pp collisions at $\sqrt{s} = 13$ TeV. *Phys. Lett. B*, 843:137745, 2023. doi: 10.1016/j.physletb.2023.137745.

[133] Wojciech Bizoń et al. Constraints on the quartic Higgs self-coupling from double-Higgs production at future hadron colliders. *JHEP*, 10:267, 2019. doi: 10.1007/JHEP10(2019)267. Addendum: JHEP 02 (2024) 170, arXiv:2402.03463.

[134] Wojciech Bizoń et al. Addendum to: Constraints on the quartic Higgs self-coupling from double-Higgs production at future hadron colliders. *Journal of High Energy Physics*, 2024(2):170, 2024. doi: 10.1007/JHEP02(2024)170.

[135] Lina Alasfar et al. Effective Field Theory descriptions of Higgs boson pair production, 2024. URL https://arxiv.org/abs/2304.01968.

[136] ATLAS Collaboration. Prospects for measuring Higgs pair production in the channel $H(\rightarrow \gamma\gamma)H(\rightarrow b\bar{b})$ using the ATLAS detector at the HL-LHC. Technical Report ATL-PHYS-PUB-2014-019, CERN, Geneva, 2014. URL https://cds.cern.ch/record/1956733.

[137] Philippe Mouche. Overall view of the LHC, 2014. URL https://cds.cern.ch/record/1708847. CERN photograph.

[138] ATLAS Collaboration. ATLAS Event Displays: Di-Higgs decays to pairs of b-quarks and photons, 2025. URL https://cds.cern.ch/record/2932627. ATLAS event display.

Acknowledgements

Although it feels surreal, I have finally completed my PhD, which somehow simultaneously passed by in a flash and lasted multiple lifetimes. Here I want to reflect on this four-year period and the people who were beside me during its ups and downs. Looking back at the last semester of my Master's in 2022, I remember being unsure what my next step would be and rather worried about it. After applying for a few positions both inside and outside of academia, I felt no closer to an answer, so I decided to stop stressing, focus on my exams and leave the decision for afterwards, when I would have a clearer head and more time to think. It was as soon as I stopped looking that I was contacted by Tristan to have a chat, and this position essentially found me.

Tristan, your passion and enthusiasm for Higgs physics, di-Higgs and for this project were immediately contagious, and from that very first chat I had an intuitive excitement about it too, something I hadn't felt deeply for the other positions I'd applied for so far. Throughout the four years of this work, your enthusiasm continued to radiate and provide a source of inspiration, reminding me not to lose sight of the bigger picture and the fact that this research is in fact extremely cool and important, even through the more stressful periods or the days spent bogged down in the details. Thank you for your guidance and for your very quick and insightful reviewing of my thesis in the last months. Thank you also for being a fellow appreciator of dad jokes and terrible puns, and for the many good laughs in the office.

Lydia, I don't think I would have made it without your advice, support and encouragement. And I don't think I would have made it *on time* without your very effective reviewing of my thesis. Our almost-weekly chats over the whole period were so valuable to me, whether they were about physics or otherwise. Thank you for being so present and engaged in my progress. You always had

time for my questions and concerns and were able to restore my confidence when I doubted myself. It's truly been a pleasure to have you as a co-promotor, colleague and friend.

Thank you to my committee members, for taking the time and effort to read and assess my thesis. Jacco, it feels like a perfect full-circle moment to have had you as my first particle physics lecturer during my bachelor's and now as a member of my PhD committee. You were a big part of the inspiration for me to pursue particle physics in the first place: both your courses at MSP and your introduction to CERN through the study trip and Summer Student Programme sowed the seeds of excitement in me about this world of research. Luca, you were an excellent mentor and collaborator during my first years of the PhD. Your depth of knowledge and clear, detailed explanations helped me understand both the physics and the analysis techniques on a much deeper level. Wouter, I am grateful to have had you as group leader during the time of my PhD. You made me feel welcome in the group and were always approachable and knowledgeable.

Mara, it was good to have you as my C3 member and as a friendly and familiar face at Nikhef. You showed real care for my progress and well-being and made me feel comfortable to share any concerns and questions.

Carlo, thank you for your guidance and mentorship through my AQP in the first year. I found our discussions insightful and learned a lot from them. Thank you also to Francesco and the FTag calibration team for your support throughout the project.

To the $HH \to b\bar{b}\gamma\gamma$ analysis team, it was a pleasure to work with you and learn from you. What we accomplished is no easy feat, and it took hard work and incredible collaboration of a large group to achieve. Knowledge was shared freely and enthusiastically by team members with a broad range of expertise, and every study, meeting and discussion played a role in this important result. I am immensely grateful to have been part of this team.

I believe that the PhD experience is determined at least as much by the group and office environment as it is by the content of the research. I feel incredibly lucky to have been part of the ATLAS group at Nikhef. In my completely unbiased opinion, you are the most welcoming, fun, talented and diverse group I could have asked for.

Osama, we have known each other since before I started and it was good to be on this parallel journey with you. I appreciate our insightful talks about every topic under the sun, your warm and calming presence and our piano duets. Thanks for being a wonderful office-mate, housemate at CERN and longtime Amsterdam friend. Zef, after becoming friends in our bachelor's, I was so happy to discover we'd be PhD colleagues too. It's been good to have you around these last years, and I hope your prophecy comes true and that maybe we'll be colleagues again one day.

Oliver, I'll miss being your neighbour, the office banter and getting demolished by you in ping-pong. Karel, how fortunate that you stayed on after your master's thesis. You're such a consistent presence in the group, always up for a coffee break or a post-lunch walk. I've had such fun times with both of you inside and outside the office and I look forward to many more still.

Robin, Liza and Lisa, I always appreciated your warm and kind presence, and thank you for being the backbone of our Nederlands Koffie – such a lovely ritual that I'll miss a lot. Ambre, the closest to me in PhD 'age', I enjoyed our talks, going around Marseille with you and sharing both the good and bad of our journeys together. Andrea and Andrea, your rivalry and fighting was an endless source of entertainment in the office even though we know that deep down there's only love there.

To the other (PhD) students and post-docs around me, both in the ATLAS group and at Nikhef: every chat, coffee, ping pong game, beer and outing contributed to making me feel like I belong here and made it a fun and warm experience. Thanks to Alejandro, Arturo, Bart, Bryan, Casper, Charlotte, Clara, Diana, Dorian, Duncan, Dylan, Edwin, Eugene, Greta, George, Jasper, Jochem, Jordan, Jordy, Jurjan, Lieke, Lisanne, Liya, Luca, Lucas, Malin, Maria, Marion, Marta, Marten, Mat, Melle, Naz-Pari, Niels, Oussama, Petja, Polina, Pranati, Rahul, Robin, Sacha, Sam, Siang-Yuan, Slava, Stefano, Varsha, Vincenzo, Vlad, Walter, Zhuoran and the rest of the Nikhef community.

To my wonderful friends at CERN: Rali, Asar, Cami, Alfonso, Mattias, Jedd and many more. Thank you for making me feel at home during my time there.

I would now like to thank the people close to me during these four years. Leo, how lucky I am to have found you 7 years ago at CERN. Together we went

through COVID, multiple questionable housing situations, break-ups, career changes and amazing home-cooked meals and our friendship evolved along with us. Thank you for being an amazing friend and paranymph. Carlo, from being my first friend in the master's to being each other's paranymphs, you've been such a constant by my side through all of this. I knew when we spontaneously drove to Italy that this would be a great and lasting friendship, and it continues to be. Cecile and Marta, our dinner group kept me sane and grounded and you have both been incredible friends over the last years.

Tal, thank you for your support, care and insightful advice during my PhD.

Thank you to the friends both old and new who have made Amsterdam feel like home: Anna, Adi, Gughi, Mees, Daria, Felix, Bo, Gabi, Maya, Marina, Merlijn, Kelvin, Lucile, Tomi, Theo, Teresa, Vedrana, Yassi and many more. Living abroad isn't easy and can get lonely, but the dinners, parties, laughs, cries, drinks and trips (with some of you to 3 different continents!) helped me realise I have a family here and I'm never alone.

To Dave, although I met you later in this journey, you helped me get through some real lows and I couldn't have finished it without you. I'm so grateful that I found you. You fully see and believe in me even through my doubts. Thank you for opening my world, for the beautiful time we've shared already. We have so much more to come.

Finally to my beloved family, without whom I would be nowhere. Joe, you are the best sibling I could ever ask for and I don't take it a single day for granted that you live here with me and that we remain such close friends. Emma, you are like a sister to me and such a light. I'm so happy to have you here in Amsterdam too. Mom and Dad, you raised me with an appreciation for arts and science, encouraged me to read widely and fostered my curiosity in many spheres. You encouraged me to take risks, pursue my interests, make my own mistakes and learn from them. Your support and love have made me feel like I can do anything and go anywhere and still be safe, stable and held by an incredible family. No words are sufficient for that, but thank you. This is for you.

