



UNIVERSITÀ DEGLI STUDI DI MILANO

Scuola di Dottorato in Fisica, Astrofisica e Fisica Applicata

Dipartimento di Fisica

Ciclo XXXVII

Shedding light on Higgs boson self-interactions in the $b\bar{b}\gamma\gamma$ channel.

**Photon and b -jet calibrations,
and searches for Higgs boson pairs with the ATLAS experiment.**

Supervisor: Prof. Leonardo Carminati
Dr. Ruggero Turra
Dr. Stefano Manzoni

Coordinatore: Prof. Aniello Mennella

Tesi di Dottorato di:
Elena Mazzeo

Anno Accademico 2023/2024

Commission of the final examination:

Prof. Pamela Ferrari

Dr. Mauro Donegà

Dr. Fabio Cerutti

Referees:

Prof. Pamela Ferrari

Dr. Mauro Donegà

Final examination:

19 December 2024

Università degli Studi di Milano, Dipartimento di Fisica, Milano, Italy

Settore Scientifico Disciplinare:

PHYS/01-A

PACS:

07.05.Kf, 07.05.Mh, 14.80.Bn, 14.80.Cp

Cover illustration:

L'Empire des lumières, René Magritte.

To Momo.

Abstract

The work presented in this thesis is based on proton-proton collision data of the Large Hadron Collider (LHC), recorded by the ATLAS detector between 2015 and 2018 (Run 2), amounting to a total integrated luminosity of 140 fb^{-1} . This work focuses on two analyses, one in the Standard Model (SM) context, and one related to probing new physics scenarios, whose common theme is the study of the Higgs boson self-interactions through the search for Higgs boson pair production in the final state with two photons and two bottom quarks ($b\bar{b}\gamma\gamma$).

The first analysis is a search for SM Higgs boson pairs (HH) in the $b\bar{b}\gamma\gamma$ final state. No excess of events with respect to the SM background is observed. An observed upper limit of $\mu_{\text{HH}} < 4.0$ is set at 95% confidence level (CL) on the di-Higgs signal strength, corresponding to the di-Higgs production cross-section normalised to its SM prediction. The corresponding expected upper limit under the assumption of no HH production (SM HH production) is 5.0 (6.4). The observed 95% confidence intervals for the coupling modifiers are $-1.4 < \kappa_\lambda < 6.9$ and $-0.5 < \kappa_{2V} < 2.7$, while the corresponding expected results are $-2.8 < \kappa_\lambda < 7.8$ and $-1.1 < \kappa_{2V} < 3.3$, yielding the most stringent expected constraints on κ_λ achieved by a single analysis to date. This analysis is combined with all the other HH searches in various final states, thus providing the most complete picture on HH production with the Run 2 dataset. The combined observed (expected) upper limit on the signal strength is $\mu_{\text{HH}} < 2.9$ (2.4) at 95% CL. The observed 95% confidence intervals for the coupling modifiers are $-1.2 < \kappa_\lambda < 7.2$ and $0.6 < \kappa_{2V} < 1.5$, with corresponding expected intervals $-1.6 < \kappa_\lambda < 7.2$ and $0.4 < \kappa_{2V} < 1.6$.

The second analysis is a search for a new heavy scalar resonance X , decaying into a Higgs boson and another scalar particle S ($X \rightarrow SH$), in the final state where the Higgs boson decays into two photons and the S particle decays into two bottom quarks. This is the first search for asymmetric Higgs boson pair production in the $b\bar{b}\gamma\gamma$ channel using ATLAS data. It searches for a $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal within a wide range of masses for the two resonances X and S (denoted m_X and m_S), covering $170 \leq m_X \leq 1000 \text{ GeV}$ and $15 \leq m_S \leq 500 \text{ GeV}$. No significant excess above the expected background is found and 95% CL upper limits are set on the cross section times branching ratio, ranging from 39 fb to 0.09 fb. The largest deviation from the background-only expectation occurs for $(m_X, m_S) = (575, 200) \text{ GeV}$ with a local (global) significance of 3.5 (2.0) standard deviations.

Contents

Introduction	ix
1 The Standard Model and Higgs boson physics	1
1.1 The Standard Model	2
1.1.1 Elementary particles	2
1.1.2 Fundamental interactions	2
1.2 The Higgs mechanism	11
1.2.1 The Standard Model Higgs boson	13
1.3 Higgs boson physics at hadron colliders	17
1.3.1 Higgs boson production mechanisms	18
1.3.2 Higgs boson decay modes	20
1.3.3 Measurements of Higgs boson properties at the LHC	22
2 Constraining the Higgs boson potential at colliders	29
2.1 Higgs boson pair production in the SM	30
2.1.1 Higgs boson pair production modes	30
2.1.2 Dependence of the Higgs boson pair production cross section from κ_λ	33
2.1.3 Dependence of the Higgs boson pair production cross section from κ_{2V} and κ_V	35
2.1.4 Searches for SM Higgs boson pair production at the LHC	37
2.2 Beyond the Standard Model potential: new physics in Higgs boson pair production	39
2.2.1 Shortcomings of the Standard Model	39
2.2.2 Asymmetric Higgs boson pair production at the LHC	45
3 The Large Hadron Collider	49
3.1 Proton-proton interactions at hadron colliders	49
3.2 Operation of the LHC and the accelerator complex	51
3.3 Luminosity and pileup	53

4	The ATLAS experiment	59
4.1	The Inner Detector	61
4.1.1	The Pixel detector	61
4.1.2	The Semiconductor Tracker	62
4.1.3	The Transition Radiation Tracker	63
4.2	The calorimeters	64
4.2.1	The Electromagnetic Calorimeter	69
4.2.2	The Hadronic Calorimeter	70
4.2.3	The Forward Calorimeter	71
4.3	The Muon Spectrometer	71
4.4	The magnet system	72
4.5	The trigger and data acquisition systems	73
5	Photon reconstruction	77
5.1	Electron and photon reconstruction	78
5.1.1	EM Calorimeter clusters formation	80
5.1.2	Track reconstruction and track-cluster matching	81
5.1.3	Supercluster growth	84
5.1.4	Preparation of electron and photon candidates for physics analysis	85
5.2	Electron and photon energy calibration	86
5.2.1	Energy reconstruction in the EM Calorimeter	90
5.2.2	LAr layer calibration	90
5.2.3	In-situ measurement of the global energy scale and resolution from $Z \rightarrow ee$ events	93
5.2.4	Energy scale uncertainties	94
5.3	Photon identification	103
5.4	Photon isolation	109
5.5	Electrons	111
5.6	Muons	111
6	b-jets reconstruction	113
6.1	Jets	113
6.1.1	Particle flow jets	113
6.1.2	Energy calibration	116
6.2	b -jets	118
6.2.1	Jet flavour tagging	119
6.2.2	b -jet identification efficiency measurement	123
6.2.3	b -jet energy corrections	130
6.3	Missing transverse momentum	133
7	Search for Higgs boson pair production in the $HH \rightarrow b\bar{b}\gamma\gamma$ channel	135
7.1	Data	137
7.2	Simulated samples	138
7.2.1	HH samples	138
7.2.2	Background samples	144

7.3	Event selection	146
7.3.1	$b\bar{b}\gamma\gamma$ preselection	146
7.3.2	Categorisation	149
7.4	Signal and background modelling	158
7.4.1	Signal and resonant background model	158
7.4.2	Continuum background model	160
7.5	Systematic uncertainties	162
7.5.1	Experimental systematic uncertainties	163
7.5.2	Theoretical systematic uncertainties	167
7.6	Statistical model	169
7.7	Results	172
7.7.1	Observed data	172
7.7.2	Exclusion limits on the HH signal strength	174
7.7.3	Constraints on κ_λ and κ_{2V}	176
7.8	Combination with the other HH final states	179
7.8.1	Overview of the HH searches	179
7.8.2	Combined results	180
8	Search for resonant production of Higgs bosons plus new scalars in the $b\bar{b}\gamma\gamma$ final state	185
8.1	Data and simulated events	186
8.1.1	Signal samples	187
8.2	Event selection	190
8.3	PNN discriminant	194
8.4	Analysis strategy and background estimation	197
8.4.1	SM HH and single Higgs boson background	197
8.4.2	Continuum background	197
8.4.3	Analysis regions	199
8.5	Systematic uncertainties	201
8.5.1	Experimental systematic uncertainties	202
8.5.2	Theoretical systematic uncertainties	202
8.5.3	Signal interpolation uncertainty	204
8.5.4	Continuum background modelling uncertainty	206
8.6	Statistical model	206
8.6.1	Test statistic for discovery	210
8.7	Results	211
8.8	Look-elsewhere effect for the global significance	215
8.8.1	Generation of toy data	216
8.8.2	Statistical results for background only toys	219
8.8.3	Global significance	222
9	Conclusions	227
	Appendices	

A	Modeling the single Higgs boson production plus bottom quarks	231
A.1	An accurate simulation of the $b\bar{b}H$ process	233
A.1.1	Comparison with the current simulation	234
A.2	Impact of the new $b\bar{b}H$ modelling on the HH searches	235
B	An alternative background modelling for the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis	239
B.1	BDT discriminants	240
B.2	Analysis strategy	241
B.3	Signal and background modelling	242
B.4	Systematic uncertainties	246
B.4.1	Experimental systematic uncertainties	247
B.4.2	Theoretical systematic uncertainties	247
B.5	Expected results	251
	Bibliography	255
	Acknowledgements	273

Introduction

The discovery of the Higgs boson (H), announced by the ATLAS [1] and CMS [2] collaborations in July 2012, marked the final missing piece of the Standard Model (SM) of particle physics. The SM, which provides a comprehensive and self-consistent description for both the electroweak and strong phenomena, has been extensively tested through numerous experimental measurements, across a wide range of energies. In the decade following this discovery, impressive effort has been dedicated to studying the properties of this long-sought particle, including its interactions with the W and Z bosons, the heaviest quarks, and charged leptons. To date, all measurements of the Higgs boson's properties report an excellent agreement with the SM predictions [3, 4].

However, the strengths of the interaction of the Higgs boson with other Higgs bosons are still largely experimentally unconstrained. According to the SM, the Higgs boson *self-interactions* (described by the *scalar potential*) emerge from the spontaneous breaking of the electroweak gauge invariance (EWSB) realized by the *Higgs mechanism*, which is responsible for the mass generation of elementary particles, including the Higgs boson itself. Thus, experimental measurements of the Higgs boson *self-couplings*, particularly the *trilinear self-coupling* λ_{HHH} , is crucial for understanding the exact structure of the scalar potential. This, in turn, offers critical insights into the EWSB mechanism and the process of mass generation.

More precisely, the SM requires a specific relationship between the trilinear Higgs boson self-coupling λ_{HHH} , the Higgs boson mass ($m_H \approx 125$ GeV), and the vacuum expectation value of the Higgs field ($v \approx 246$ GeV). Since the latter two parameters are known from experimental measurements [5, 6], the SM provides a well-defined prediction for the self-coupling λ_{HHH} , corresponding to $\lambda_{HHH}^{\text{SM}} \approx 0.13$. The measurement of any deviation of λ_{HHH} from its SM prediction would profoundly challenge our current understanding of EWSB and provide a strong hint of new physics beyond the SM (BSM). In experimental measurements, this corresponds to testing for a deviation of its *coupling modifier* $\kappa_\lambda = \lambda_{HHH} / \lambda_{HHH}^{\text{SM}}$ from unity.

The only direct way to probe κ_λ is through the search for Higgs boson pair production (HH), which is influenced by the trilinear self-interaction vertex in both the dominant gluon-gluon fusion (ggF) and subdominant vector boson fusion (VBF) production modes, as shown in Figure 1. Up to now, the observation of SM double Higgs boson production has been out of reach for the ATLAS and CMS experiments at the Large Hadron Collider (LHC) [4, 7]. Therefore, a precise measurement of the trilinear Higgs boson self-coupling remains elusive.




Figure 1 – Feynman diagrams contributing to the total amplitude of HH production via (a) ggF and (b) VBF with terms proportional to the coupling modifier κ_λ .

The work presented in this thesis is based on proton-proton collision data of the LHC [8], recorded by the ATLAS detector [9] between 2015 and 2018 (corresponding to the LHC Run 2 data-taking period). The ATLAS detector is one of the two general-purpose experiments along the LHC ring: its main goal is to study the broadest range of signals as possible, either coming from SM processes or from new physics. Both these approaches are pursued in this thesis. My work focuses on two analyses, whose common theme is the search for Higgs boson pair production in the final state with two photons and two bottom quarks ($b\bar{b}\gamma\gamma$).

The first analysis is a search for production of SM Higgs boson pairs in the $b\bar{b}\gamma\gamma$ channel [10]. This final state where one Higgs boson decays to two bottom quarks and the other decays to a photon pair is one of the “golden channels” of di-Higgs searches [11]. This final state benefits from the high branching fraction of the $H \rightarrow b\bar{b}$ decay (59% for a SM Higgs boson). On the other hand, the low branching ratio for the $H \rightarrow \gamma\gamma$ decay (at just 0.23%) is compensated by the excellent trigger and reconstruction efficiency for photons with the ATLAS detector and the excellent resolution of the invariant mass of the photon pair (1-2 GeV), leading to a clear and distinguishable signal. Di-Higgs boson production is an extremely rare process, and only 0.26% of Higgs boson pairs generate a two photon plus two bottom quarks final state. According to the SM, only ~ 12 $HH \rightarrow b\bar{b}\gamma\gamma$ events are expected to be present within the full Run 2 dataset, which corresponds to a six and three order-of-magnitude suppression with respect to inclusive photon pair production and single $H \rightarrow \gamma\gamma$, respectively. The former makes up the primary source of background for the $HH \rightarrow b\bar{b}\gamma\gamma$ search. Hence, this analysis relies on the ability of selecting good candidates for the $H \rightarrow \gamma\gamma$ and $H \rightarrow b\bar{b}$ decays, to help separating the rare signal events from the overwhelming backgrounds.

Compared to a previous di-Higgs search in the $b\bar{b}\gamma\gamma$ channel using ATLAS Run 2 data [12], this new analysis exploits an improved optimisation strategy, incorporating both the dominant ggF production mode and the subdominant VBF production mode. Although Higgs boson pair production via VBF is incredibly rare (with a rate about 20 times smaller than ggF according to the SM), searching for the VBF HH process is particularly interesting. This production mode is characterized by two highly energetic hadronic

jets, reconstructed in the forward region of the ATLAS detector, in addition to the $b\bar{b}\gamma\gamma$ signature. This distinctive topology helps to isolating the Higgs boson pairs produced via VBF, thus enhancing the overall signal-to-background ratio. Additionally, beyond providing further constraints on the coupling modifier κ_λ when combined with ggF, di-Higgs production via VBF offers a unique probe to the quartic interaction between a Higgs boson pair and two vector bosons (W or Z). The strength of this interaction, denoted as g_{HHVV} , is quantified by the coupling modifier $\kappa_{2V} = g_{HHVV}/g_{HHVV}^{SM}$ and is closely connected to the EWSB sector of the SM. Measuring κ_{2V} through Higgs boson pair production via VBF would help to form a clearer picture of the EWSB and the Higgs mechanism.

To maximise sensitivity to Higgs boson pair production, the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis is combined with other HH searches across different final states [7], including the two other *golden channels*, $b\bar{b}b\bar{b}$ and $b\bar{b}\tau\tau$ [13–15]. Among these, the $b\bar{b}\gamma\gamma$ final state stands out for its best sensitivity to κ_λ . The HH combination offers the most comprehensive insight into Higgs boson self-interactions using ATLAS Run 2 data.

Despite the progress made with the new HH analyses with Run 2 data, which brings us closer to the sensitivity needed to detect di-Higgs production in a SM-like scenario, the Higgs boson self-interaction sector remains significantly unconstrained. As a result, new physics could still be hidden within the scalar sector [16, 17], which could potentially give rise to anomalous trilinear Higgs boson self-coupling values. New physics scenarios can be explored by searching for extensions of the scalar sector, which can lead to asymmetric Higgs boson pair production, where a heavy scalar resonance (X) decays into a SM Higgs boson and a lighter scalar particle (S). The second analysis presented in this thesis is a search for the $X \rightarrow SH$ process in the final state where the Higgs boson decays into two photons and the scalar S decays into two bottom quarks [18]. This is the first search for asymmetric Higgs boson pair production in the $b\bar{b}\gamma\gamma$ channel using ATLAS data. It searches for a $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal within a broad range of masses for the two resonances X and S (denoted m_X and m_S), reaching previously uncharted regions of the phase space.

The search for $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ exploits the presence of three resonances in the final state - $H \rightarrow \gamma\gamma$, $S \rightarrow b\bar{b}$, and $X \rightarrow b\bar{b}\gamma\gamma$ - which are crucial for isolating the signal from competing SM background processes mimicking the $b\bar{b}\gamma\gamma$ signature, where at least one of these resonances is absent. This analysis was designed to have a fully continuous sensitivity across the explored region in the plane defined by the masses of the two resonances. In absence of a statistically significant excess observed above the expected backgrounds, it sets upper limits on the cross section times branching fraction for the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal.

Summary of the manuscript and personal contributions

This thesis includes seven Chapters, which are organized as follows. The first chapter (Chapter 1) summarizes the formulation of the SM, focusing on the electroweak symmetry breaking and the Higgs mechanism. This chapter also includes a description of the Higgs boson physics at hadron colliders, and a short review of the main measurements

of the Higgs boson properties performed by the ATLAS and CMS collaborations with Run 2 data.

Chapter 2 focuses on the phenomenology of Higgs boson pair production, including its connection to the trilinear Higgs boson self-coupling λ_{HHH} , and to the quartic Higgs boson-vector boson coupling g_{HHVV} . This chapter also addresses the main shortcomings of the SM, and includes a discussion on how new physics models, involving extended scalar sector with multiple scalar particles, could address many of its open issues and are thus interesting to probe.

Chapter 3 and Chapter 4 describe the LHC accelerator complex and the ATLAS detector. In Chapter 4 special emphasis is given to the Liquid Argon electromagnetic calorimeter, which is one of the main sub-detectors exploited by the measurements described in this thesis, especially for the reconstruction of photons. During my doctorate I worked within the ATLAS Liquid Argon Calorimeter group, involved on both the day-to-day operation of the calorimeter and to the upgrade of the readout electronics foreseen for the High-Luminosity LHC phase. For the former activity, I have been one of the on-call experts for the system that delivers the high voltage to the calorimeter electrodes. My duties ranged from solving hardware issues happening during the data-taking, as well as hardware and software maintenance during the periods without proton beams circulating in the LHC. For the latter activity, I tested various powering solutions for the main readout boards of the calorimeter, which collect and process the ionization signals from each calorimeter cells.

Chapters 5 details the reconstruction, identification, and calibration of photons. Similarly, Chapter 6 describes the reconstruction, identification, and energy calibration of jets originating from bottom quark hadronization. Both these particles are essential components of the two analyses discussed earlier. I contributed to various aspects of both photon and b -jets reconstruction. Specifically, I was involved in the electron and photon energy calibration, in particular on the final step of the procedure, which consists in the *in-situ* measurement of the electron and photon energy scale using Z boson decays to electron pairs [19]. This measurement is impacted by several sources of systematic uncertainties, whose control is crucial, in order to achieve a precise reconstruction of the $H \rightarrow \gamma\gamma$ decay. My work included evaluating the systematic uncertainty related to photon conversion mismodelling and propagating the effects of all other systematic uncertainties to the electron and photon energy scale as functions of the pseudorapidity and transverse momentum of electrons and photons.

On the b -jets side, I contributed to measuring the efficiency of a new b -jets identification algorithm based on Graph Neural Networks (GNN) [20]. This algorithm was recently developed by the ATLAS collaboration for analyses using data from the ongoing Run 3 data-taking period of the LHC, and is expected to bring sizable improvements in the next generation of HH searches involving b -jets in the final state.

Chapter 7 and Chapter 8 present the two analyses, $HH \rightarrow b\bar{b}\gamma\gamma$ [10] and $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ [18], respectively. I contributed to most aspects of the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis. I worked on the analysis optimisation, on the signal and background modelling, on the evaluation of the systematic uncertainties, and on the statistical interpretation and the extraction of the final results. In addition, I also served as editor of the Internal sup-

porting documentation, detailing the full analysis. Chapter 7 also includes the description of the combination of HH final states with Run 2 data. My contribution to the HH combination involved studying the modelling of HH production via ggF as a function of the coupling modifier κ_λ , and extracting the final results, focusing especially on the constraints on the di-Higgs signal strength and on the coupling modifiers κ_λ and κ_{2V} . My contributions to the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ analysis involved studying the background modelling, developing one of the statistical frameworks adopted by the analysis, and extracting the results. In addition, I evaluated the look-elsewhere effect for the calculation of the global significance across the full explored (m_X, m_S) plane.

In the context of the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis, I also developed an alternative background modelling strategy, based on changing the final discriminant variable, from the diphoton invariant mass to BDT discriminants. This work, described in Appendix B, demonstrated the effectiveness of this approach, which was later adopted by the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ analysis. Finally, I contributed to a phenomenological study, detailed in Appendix A, which aims to improve the modeling of single Higgs boson production with emission bottom quark emissions, which constitutes a crucial background for HH searches [21]. Some minor contributions to other publications of the ATLAS Collaboration are not included in this thesis, such as the preparation of the statistical model for the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis [12] for the combination of HH and single Higgs boson final states [22].

The Standard Model and Higgs boson physics

The Standard Model (SM) of particle physics offers a unified framework to describe both strong and electroweak phenomena, while categorizing all known elementary particles, including force carriers and matter particles. The SM was developed over the course of a few decades between the 1950s and 1970s [23]. It has shown remarkable success in explaining the subatomic world up to the highest energy scales accessible by current experiments, and has survived every experimental test up to the present day.

Each elementary particle in the SM corresponds to a quantum field with specific transformation properties under the Poincaré group (which determine their *spin* and *parity*), and their interactions are governed by a common principle: *local gauge invariance* [24, 25]. The formulation of the SM as a *gauge theory* ensures its *renormalizability* [26, 27]. This property guarantees that the SM remains predictive at all energy scales, and that the theory can be consistently quantized. The SM integrates the electroweak theory developed by Glashow, Weinberg, and Salam [28, 29] with Quantum Chromodynamics (QCD) [30], the fundamental theory of strong interactions.

A cornerstone of the SM is the Higgs mechanism, which elegantly resolves the problem of incorporating mass terms for elementary particles into the SM Lagrangian without compromising the electroweak gauge invariance. This mechanism implies the existence of a new scalar field, known as the *Higgs field*, whose excitations are identified as *Higgs bosons*. The discovery of a particle resembling the Higgs boson by the ATLAS and CMS collaborations in July 2012 [1, 2] provided strong evidence of the Higgs mechanism, significantly advancing our comprehension of the origin of mass of elementary particles.

The Higgs mechanism realizes the spontaneous breaking of electroweak gauge invariance (EWSB), thereby generating mass terms for elementary particles through their interactions with the Higgs field. Concurrently, it introduces self-interaction terms for the Higgs field itself. The study of these Higgs field self-interactions is closely linked to experimental efforts to observe the production of Higgs boson pairs. Detecting pairs of SM Higgs bosons would offer definitive validation of the Higgs mechanism as described in the SM, while any deviations from the expectation would provide a strong hint of new physics.

This Chapter provides an overview of the SM (Section 1.1), with a particular focus on the Higgs mechanism, discussed in Sections 1.2. A phenomenological portrait of the Higgs boson physics at hadron colliders (including a brief description of the measurements of the Higgs boson properties at the LHC) follows in Section 1.3.

1.1 The Standard Model

The SM is a gauge quantum field theory based on the Lie group $SU(3)_C \times SU(2)_L \times U(1)_Y$, where $SU(3)_C$ and $SU(2)_L \times U(1)_Y$ are the gauge groups of the QCD and electroweak interactions respectively. The SM gauge symmetry governs the fundamental interactions of all elementary particles, by assigning the corresponding fields to specific representations of the gauge group. The electroweak gauge invariance $SU(2)_L \times U(1)_Y$ is spontaneously broken (via the Higgs mechanism) to $U(1)_{EM}$, which describes the electromagnetic interactions. The SM does not include a description of *gravity*. Nevertheless, at the subatomic scale, this force is the weakest among the fundamental interactions, and can be safely neglected at the energy scales probed by the experiments at hadron colliders.

1.1.1 Elementary particles

The SM incorporates 12 spin- $\frac{1}{2}$ *fermionic fields* that represent the fundamental constituents of matter, known as *quarks* and *leptons*. Leptons only interact through electroweak force, while quarks are subject to both electroweak and strong forces, thus engaging in all SM interactions. Quarks and leptons are organized into three *generations* (or *families*), which share the same properties under electroweak and strong forces (namely, they have the same *quantum numbers*) but have progressively larger masses. Each fermion has its own anti-particle, with identical mass and opposite quantum numbers. In the SM, all the matter fields are massive, except for the neutrinos, which are considered massless. Nevertheless, the observation of neutrino oscillations has confirmed the non-zero value of neutrino masses [31], which are, however, extremely small compared to the masses of other fermions ($m_\nu < 0.8$ eV [32]).

Local gauge invariance in the SM implies the existence of spin-1 *gauge fields* related to the generators of the SM gauge group, whose excitations, known as *gauge bosons*, mediate the fundamental interactions. The photon (γ) is the gauge boson for electromagnetic interactions, while the charged $W^\pm$ bosons and the neutral Z boson mediate weak interactions. The gluon (g) is the gauge boson for the strong force. Among these, only the photon and the gluon are massless. The masses of the weak bosons are acquired through the Higgs mechanism. The Higgs boson, the final component of the SM, is crucial for generating the masses of both the weak vector bosons and all massive fermions, as well as its own mass.

Figure 1.1 provides an overview of the elementary particles in the SM and their properties.

1.1.2 Fundamental interactions

As already mentioned, the SM provides a consistent description of the electroweak and strong forces at all energy scales, by enforcing local gauge invariance under the group $SU(3)_C \times SU(2)_L \times U(1)_Y$. In gauge theories, the form of the interaction is naturally set after defining the symmetry group, requiring gauge invariance, and determining the properties of the matter fields under gauge transformations.




Figure 1.1 – Elementary particles described by the SM. The fermion fields (on the left) are depicted with their *quantum numbers* (electric charge, color, and spin), dictating their interaction properties.

The simplest example of a gauge theory is Quantum Electrodynamics (QED), where the electromagnetic interaction is derived from gauge invariance under the Abelian symmetry group $U(1)$. This method of constructing QED as a gauge theory (summarized below) can be extended to the more complex framework of the SM.

Let us begin with the following Lagrangian:

$$\mathcal{L} = i\bar{\psi}\gamma^\mu\partial_\mu\psi - m\bar{\psi}\psi, \quad (1.1)$$

which describes a free (non-interacting) spin- $\frac{1}{2}$ fermion field $\psi(x)$, with $\bar{\psi} = \psi^\dagger\gamma^0$ as its adjoint. This Lagrangian remains invariant under the global transformations of the symmetry group $U(1)$, which are typically represented with the unitary operator $U = e^{i\theta Q}$. Here, θ is a constant phase (i.e. $\partial_\mu\theta = 0$), and Q is the generator of the group, which acts on the fermion field as $Q\psi = \psi$. The transformation laws for ψ and its adjoint field under global $U(1)$ transformations are:

$$\begin{aligned} \psi(x) &\longrightarrow e^{i\theta Q}\psi(x) = e^{i\theta}\psi(x), \\ \bar{\psi}(x) &\longrightarrow e^{i\theta Q}\bar{\psi}(x) = e^{-i\theta}\bar{\psi}(x), \end{aligned} \quad (1.2)$$

To introduce the electromagnetic interaction, we apply the gauge principle, converting the global $U(1)$ symmetry into a *local* gauge invariance. This means that the fundamental laws of the theory (the Lagrangian) must remain invariant under local phase transformations, where $\theta(x)$ is an arbitrary function of space-time coordinates x . While the mass term in the original Lagrangian is invariant under local $U(1)$ transformations, the kinetic term is not:

$$\partial_\mu\psi \longrightarrow \partial_\mu\left[e^{i\theta(x)}\psi\right] = e^{i\theta(x)}\left[\partial_\mu\psi + i\partial_\mu\theta\right]. \quad (1.3)$$

Local gauge invariance is restored by replacing the partial derivative in the kinetic term with a *covariant derivative*, which includes an additional vector (spin-1) field $A_\mu(x)$, characterized by the same quantum numbers as the generator Q . The properties of $A_\mu(x)$ under local phase transformations are defined as:

$$A_\mu(x) \longrightarrow A_\mu(x) - \frac{1}{e}\partial_\mu\theta(x), \quad (1.4)$$

where e is a constant.

The covariant derivative $D_\mu(x)$, acting on the fermion field as

$$D_\mu(x)\psi(x) = \left[\partial_\mu + ieA_\mu(x)\right]\psi(x), \quad (1.5)$$

restores the gauge invariance of the resulting Lagrangian. The transformation law of the vector field A_μ (termed the *gauge field*) cancels out the unwanted term containing the partial derivative of $\theta(x)$. To complete the QED Lagrangian, we need a gauge-invariant kinetic term for the vector field A_μ to describe its free propagation. This kinetic term,

$\mathcal{L}_{kin}^{QED} = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu}$, is based on the field strength

$$F^{\mu\nu} = \partial^\mu A^\nu - \partial^\nu A^\mu \quad (1.6)$$

which is manifestly gauge-invariant. The resulting QED Lagrangian is:

$$\begin{aligned} \mathcal{L} &= -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} + i\bar{\psi}\gamma^\mu D_\mu\psi - m\bar{\psi}\psi \\ &= -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} + i\bar{\psi}\gamma^\mu\partial_\mu\psi - m\bar{\psi}\psi - eA_\mu\bar{\psi}\gamma^\mu\psi \end{aligned} \quad (1.7)$$

where the last term couples the fermion field to the gauge field A_μ , identified with the electron and the photon (the carrier of the electromagnetic force) respectively. The constant $e = 1.602176 \times 10^{-19}$ C is interpreted as the electric charge of the electron. The photon remains massless, and a mass term $-\frac{1}{2}M^2 A^\mu A_\mu$ is forbidden by gauge invariance. In addition, no self-interaction terms for the photon are present. The A_μ field is neutral under electromagnetic interactions: this stems naturally from the Abelian nature of the U(1) symmetry group.

The interaction term in the QED Lagrangian involving the fermion field and the photon, specifically $\mathcal{L}_{int}^{QED} = -eA_\mu\bar{\psi}\gamma^\mu\psi$, results in a cubic interaction vertex, as illustrated in Figure 1.2. This vertex introduces a factor of e in the calculation of scattering amplitudes.




Figure 1.2 – Feynman diagram representing the cubic interaction vertex between two fermions (f) and the photon (γ) in QED.

From an experimental point of view, the physical quantity of interest is the interaction cross section, which is related to the probability of a particular scattering process to occur. The cross section is proportional to the square of the amplitude's modulus, making it dependent on e^2 . Therefore, the strength of the electromagnetic interaction is set by the fine structure constant $\alpha = \frac{e^2}{4\pi\epsilon_0\hbar c}$, rather than e itself. Here, $\epsilon_0 = 8.854187 \times 10^{-12}$ Fm $^{-1}$ is the vacuum dielectric constant, $\hbar = 1.054571 \times 10^{-34}$ Js is the reduced Planck constant, and $c = 299792458$ m/s is the speed of light in the vacuum. The quantization of this Lagrangian has been remarkably successful in describing electromagnetic interactions even at very high energies.

A similar procedure is used to construct the strong and electroweak Lagrangians follow-

ing the gauge principle, though additional complexities arise due to the non-Abelian nature of the symmetry group $SU(3)_C \times SU(2)_L \times U(1)_Y$.

Quantum Chromodynamics

The complex dynamics of strong interactions are described at a fundamental level by Quantum Chromodynamics (QCD). QCD is a gauge theory based on the $SU(3)_C$ symmetry group, where the subscript C denotes the color charge carried by particles influenced by the strong force. The color charge in QCD plays the same role as the electric charge in QED. However, there is a fundamental difference between the electric charge and the color charge: while there is a single electric charge, the strong interaction phenomenology suggests that quarks possess three types of color: red, green, and blue. Consequently, each quark field of a specific flavor (u, d, c, s, t , or b) is a color triplet $q_f(x) = (q_f^{\text{red}}(x), q_f^{\text{green}}(x), q_f^{\text{blue}}(x))$ under the transformations of $SU(3)_C$. Considering all quark flavors f , the QCD Lagrangian is formulated as:

$$\mathcal{L} = \sum_f i q_f D_\mu \gamma^\mu q_f - \frac{1}{4} \sum_{a=1}^8 F^{\mu\nu,a} F_{\mu\nu}^a, \quad (1.8)$$

where:

- The first term parallels QED, incorporating both the free propagation of quarks and their interaction with the gluons (the strong force carriers). The covariant derivative is defined as:

$$D_\mu(x) q_f(x) = \left[\partial_\mu + i g_s \sum_{a=1}^8 G_\mu^a(x) t^a \right] q_f(x), \quad (1.9)$$

including the eight gluon fields $G_\mu^a(x)$ - each corresponding to a $SU(3)_C$ generator - and the generators t^a acting on quark triplets. The interaction strength between the gluons and the quarks is governed by the parameter g_s . This quantity is unique across all quark flavors and is typically expressed in terms of $\alpha_s = \frac{g_s^2}{4\pi}$, which defines the *strong coupling constant*. α_s plays the same role in QCD as the fine structure constant in QED: it appears in the interaction cross section involving the cubic interaction vertex between quarks and gluons, thereby setting the strength of the strong force.

- The second term is the gluons' kinetic term, with:

$$F^{\mu\nu,a} = \partial^\mu G^{\nu,a} - \partial^\nu G^{\mu,a} - g_s \sum_{b,c=1}^8 f^{abc} G^{\mu,b} G^{\nu,c}, \quad (1.10)$$

with f^{abc} being the structure constants of $SU(3)_C$. The third term in the field strength - absent in QED - arises from the gauge principle and the non-Abelian nature of $SU(3)_C$. Since $F^{\mu\nu,a}$ is quadratic in gauge fields, $\frac{1}{4} \sum_{a=1}^8 F^{\mu\nu,a} F_{\mu\nu}^a$ includes cubic and quartic interaction terms (proportional to g_s and g_s^2 , respectively) involving gluons.

In QED, the photon interacts with electrically charged fermions but remains electrically neutral itself. In contrast, experimental evidence shows that gluons interact strongly with each other [33]. Therefore, gluons are represented as gauge fields that carry color charge.

Renormalizability in QCD manifests as a *running* coupling constant α_s , varying with the momentum transfer during strong interactions. For large momentum transfer Q , the strong coupling is:

$$\alpha_s(Q^2) \approx \frac{12\pi}{(33 - 2n_f) \cdot \ln\left(\frac{Q^2}{\Lambda_{QCD}^2}\right)}, \quad (1.11)$$

where n_f is the number of quark flavors ($n_f = 6$ in the SM) and $\Lambda_{QCD} \approx 200$ MeV sets the energy scale for perturbative QCD predictions. This expression holds in the perturbative regime, where the energy scale of the strong process significantly exceeds Λ_{QCD} . As $Q^2 \rightarrow \infty$, the strong coupling logarithmically approaches zero, leading to *asymptotic freedom*: the quarks act as free particles at high energies. Conversely, as the energy scale nears Λ_{QCD} , the running coupling diverges: the interaction between quarks grows stronger, and perturbative calculations fail to provide reliable predictions for strong processes. This behavior hints at another important property of QCD, namely, the *color confinement*. While this property is not yet proved mathematically, it manifests itself in experimental observations: particles carrying color charge (namely, quarks and gluons) cannot be observed isolated. Instead they group together to form color-neutral composite states, known as *hadrons*. Color confinement explains the observation of an abundance of color-singlet composite particles, including *mesons* and *baryons* (corresponding to $q\bar{q}$ and qqq bound states), and even more exotic hadronic states such as *tetraquarks* and *pentaquarks* (all predicted by QCD). Due to color confinement, quarks or gluons produced in high energy collisions undergo *hadronization*. The primary particle involved in the collision starts to emit gluons and quark-antiquark pairs with progressively lower energy. The secondary quarks and gluons group themselves in hadronic states, thus resulting in collimated sprays of hadrons termed *hadronic jets*. A notable exception to this behaviour is the heaviest quark, the *top quark*. Due to its extremely short lifetime, the top quark decays into a bottom quark and a W boson before hadronization can occur, thus allowing for the direct observation of its properties.

Electroweak interactions

A unified description of the electromagnetic and weak forces was achieved by Glashow, Weinberg, and Salam in the 1960s, leading to formulation of the electroweak interactions as a gauge theory of the Lie group $SU(2)_L \times U(1)_Y$. The subscript L denotes *left-handed* fermions, indicating that only the left chiral component $\psi_L = \frac{1}{2}(1 - \gamma^5)\psi$ of quarks and leptons - collectively indicated here using ψ - interacts with the three gauge fields $W_\mu^i(x)$ (with $i = 1, 2, 3$) associated to the $SU(2)_L$ symmetry group. Conversely, the right chiral component of a fermion field $\psi_R = \frac{1}{2}(1 + \gamma^5)\psi(x)$ is a singlet under $SU(2)_L$, and do not couple to its gauge fields. The three generators of $SU(2)_L$, T^i with $i = 1, 2, 3$, are

termed *weak isospin* operators. $U(1)_Y$ is not the gauge group of the electromagnetic interactions: its generator, Y , is termed *weak hypercharge* and, contrary to the electric charge Q , it affects the left-handed and the right-handed components of fermion fields differently. Left-handed fermions are represented as doublets under $SU(2)_L$, while, as mentioned above, right-handed fermions transform trivially (namely, $T^i \psi_R = 0$). The transformation properties of the left-handed fermions under $U(1)_Y$ maintain the $SU(2)_L$ structure, as the upper and lower components of each doublet share the same weak hypercharge quantum number.

The arrangement of leptons and quarks within each family into left-handed doublets and right-handed singlets is detailed below:

Table 1.1 – Doublet and singlet structure for the left and right-handed components of the fermion fields under the $SU(2)_L \times U(1)_Y$ symmetry group.

	Leptons	Quarks
Left-handed	$\begin{pmatrix} \nu_{e,L} \\ e_L \end{pmatrix}, \begin{pmatrix} \nu_{\mu,L} \\ \mu_L \end{pmatrix}, \begin{pmatrix} \nu_{\tau,L} \\ \tau_L \end{pmatrix}$	$\begin{pmatrix} u_L \\ d_L \end{pmatrix}, \begin{pmatrix} c_L \\ s_L \end{pmatrix}, \begin{pmatrix} t_L \\ b_L \end{pmatrix}$
Right-handed	e_R, μ_R, τ_R	$u_R, c_R, t_R, \\ d_R, s_R, b_R$

In each left-handed lepton doublet, neutrinos are the upper-row partners of the corresponding charged leptons. Right-handed neutrinos do not exist in the original formulation of SM. The quarks whose left-handed components occupy the upper row of the weak isospin doublets - up (u), charm (c), and top (t) quarks - are called up-type quarks. Conversely, the bottom components correspond to the left chiral projections of the down-type quarks, namely the down (d), strange (s), and bottom (b) quarks.

The gauge-invariant electroweak Lagrangian is given by

$$\mathcal{L} = i\bar{\psi}_L \gamma^\mu D_\mu \psi_L + i\bar{\psi}_R \gamma^\mu D_\mu \psi_R - \frac{1}{4} \sum_{i=1}^3 W^{\mu\nu,i} W_{\mu\nu}^i - \frac{1}{4} B^{\mu\nu} B_{\mu\nu}, \quad (1.12)$$

where sums are implied over all left-handed doublets ψ_L and right-handed singlets ψ_R . The last two terms in Equation 1.12 include the field strengths

$$W^{\mu\nu,i} = \partial^\mu W^{\nu,i} - \partial^\nu W^{\mu,i} - g \sum_{j,k=1}^3 \epsilon^{ijk} W^{\mu,j} W^{\nu,k} \quad (1.13)$$

and

$$B^{\mu\nu} = \partial^\mu B^\nu - \partial^\nu B^\mu, \quad (1.14)$$

constructed from the gauge fields $W^{\mu,i}$ associated with the three weak isospin generators, and B^μ , the gauge field for the weak hypercharge. The structure constants of the $SU(2)_L$ symmetry group appearing in Equation 1.13 are the Levi-Civita tensor components ϵ^{ijk} , while g (g') are the weak charges for $SU(2)_L$ ($U(1)_Y$).

The electroweak covariant derivative (present in the first and second terms of Equation 1.12), acts differently on the left- and right-handed components of the fermion fields. Given a $SU(2)_L$ doublet $\psi_L = \begin{pmatrix} u_L \\ d_L \end{pmatrix}$ whose weak hypercharge quantum number is $Y_L/2$, the covariant derivative takes the following form:

$$\begin{aligned}
 D_\mu(x)\psi_L(x) &= \left[\partial_\mu + ig \sum_{i=1}^3 W_\mu^i(x) T^i + ig' B_\mu(x) Y \right] \psi_L \\
 &= \partial_\mu \begin{pmatrix} u_L \\ d_L \end{pmatrix} + i \frac{g}{2} \begin{pmatrix} W_\mu^3 & W_\mu^1 - iW_\mu^2 \\ W_\mu^1 + iW_\mu^2 & -W_\mu^3 \end{pmatrix} \begin{pmatrix} u_L \\ d_L \end{pmatrix} + ig' B_\mu \frac{Y_L}{2} \begin{pmatrix} u_L \\ d_L \end{pmatrix}_L \\
 &= \partial_\mu \begin{pmatrix} u_L \\ d_L \end{pmatrix} + \frac{i}{2} \begin{pmatrix} gW_\mu^3 + g'Y_L B_\mu & g(W_\mu^1 - iW_\mu^2) \\ g(W_\mu^1 + iW_\mu^2) & -gW_\mu^3 + g'Y_L B_\mu \end{pmatrix} \begin{pmatrix} u_L \\ d_L \end{pmatrix}
 \end{aligned} \tag{1.15}$$

where the $SU(2)_L$ generators T^i where replaced with their doublet representations $\sigma^i/2$, with σ^i being the Pauli matrices.

On the other hand, the covariant derivative acting on the right-handed fields ψ_R is:

$$\begin{aligned}
 D_\mu(x)\psi_R(x) &= \left[\partial_\mu + ig \sum_{i=1}^3 W_\mu^i(x) T^i + ig' B_\mu(x) Y \right] \psi_R(x) \\
 &= \partial_\mu \psi_R + i \frac{g'}{2} B_\mu Y_R^{u(d)} \psi_R
 \end{aligned} \tag{1.16}$$

where ψ_R represents either the up-type or the down-type quark (namely, u_R or d_R), and $Y_R^{u(d)}/2$ is their weak hypercharge quantum number.

The interactions between the electroweak gauge fields and the fermion fields are therefore encoded in the following Lagrangian, obtained by replacing the covariant derivatives of Equation 1.15 and 1.16 in Equation 1.12:

$$\begin{aligned}
 \mathcal{L}_{int}^{EW} &= -\frac{1}{2} (\bar{u}_L \bar{d}_L) \gamma^\mu \begin{pmatrix} gW_\mu^3 + g'Y_L B_\mu & g(W_\mu^1 - iW_\mu^2) \\ g(W_\mu^1 + iW_\mu^2) & g'Y_L B_\mu - gW_\mu^3 \end{pmatrix} \begin{pmatrix} u_L \\ d_L \end{pmatrix} \\
 &\quad - \frac{g'}{2} \bar{u}_R Y_R^u \gamma^\mu B_\mu u_R - \frac{g'}{2} \bar{d}_R Y_R^d \gamma^\mu B_\mu d_R
 \end{aligned} \tag{1.17}$$

Linear combinations of the two gauge fields $W_\mu^{1,2}$ yield

$$W_\mu^\pm(x) = \frac{1}{\sqrt{2}} \left[W_\mu^1(x) \mp W_\mu^2(x) \right], \tag{1.18}$$

whose physical states correspond to the two charged $W^\pm$ bosons. The weak interactions mediated by $W^\pm$ - whose fields appear in the off-diagonal terms in the matrix that describes the interaction with the left-handed fermions - connect the up and the down

component of the weak isospin doublets (namely, u_L and d_L). The electromagnetic interaction is recovered by performing a rotation in the two-dimensional space spanned by the two remaining gauge fields of the $SU(2)_L \times U(1)_Y$ group, W_μ^3 and B_μ :

$$\begin{pmatrix} W_\mu^3 \\ B_\mu \end{pmatrix} = \begin{pmatrix} \cos \theta_W & \sin \theta_W \\ -\sin \theta_W & \cos \theta_W \end{pmatrix} \begin{pmatrix} Z_\mu \\ A_\mu \end{pmatrix} \quad (1.19)$$

This transformation defines the *weak mixing angle* θ_W . Identifying the gauge field A_μ with the photon requires the following relationship between the electron charge e and the weak charges g and g' :

$$e = g \sin \theta_W = g' \cos \theta_W \quad (1.20)$$

Moreover, in the context of the electroweak theory, the electric charge operator is defined as

$$Q = T^3 + \frac{Y}{2} \quad (1.21)$$

Alongside with the photon, another gauge boson enters the neutral sector of the electroweak interactions: the Z boson. The Z bosons mediates weak neutral interactions between quarks and leptons of the same flavor, with different coupling constant across left-handed and right-handed chiral components.

Applying the latter notations in Equation 1.17 recovers the electroweak Lagrangian in terms of the physical weak and electromagnetic currents:

$$\begin{aligned} \mathcal{L}_{int}^{EW} = & -e A^\mu \bar{\psi} \gamma_\mu Q \psi \\ & - \frac{g}{\sqrt{2}} \bar{\psi} \gamma_\mu \left[W^{+\mu} (T^1 - iT^2) + W^{-\mu} (T^1 + iT^2) \right] \psi \\ & + \frac{e}{2 \sin \theta_W \cos \theta_W} \bar{\psi} \gamma_\mu Z^\mu \left[(T^3 - 2Q \sin^2 \theta_W) - T^3 \gamma^5 \right] \psi \end{aligned} \quad (1.22)$$

where the first, second, and third lines include the *electromagnetic current* J_μ^{em} , the *weak charged currents*¹ $J_\mu^\pm$, and the *weak neutral current* J_μ^0 .

The full electroweak Lagrangian in Equation 1.12 is therefore able to describe charged and neutral interactions associated with electromagnetic and weak processes: it incorporates QED and contemplates self-interactions between the gauge bosons. However, while it describes the photon as a massless boson (in agreement with experimental findings), it does not account for the experimentally observed large masses of the weak bosons $W^\pm$ and Z [34]. In fact, mass terms of the form $m_W^2 W^{\pm,\mu} W_\mu^\mp$ and $\frac{1}{2} m_Z^2 Z^\mu Z_\mu$ would explicitly violate the gauge invariance.

¹The weak charged currents involve the raising and lowering weak isospin operators $T^\pm = T^1 \pm iT^2$, acting non-trivially only on the left-handed component of the fermion fields. Indeed, considering $\psi_L = \begin{pmatrix} u_L \\ d_L \end{pmatrix}$ as a representative case, $J_\mu^+ = \bar{\psi}_L \gamma_\mu T^+ \psi_L = \bar{u}_L \gamma_\mu d_L$, while $J_\mu^- = \bar{d}_L \gamma_\mu u_L$.

1.2 The Higgs mechanism

As noted in Section 1.1.2, the gauge principle prohibits explicit mass terms for gauge bosons in the electroweak Lagrangian. Nonetheless, experimental observations show that, of the four electroweak bosons, only the photon is massless, while the $W^\pm$ and Z bosons are massive [6]:

$$\begin{aligned} m_W &= 80.377 \pm 0.012 \text{ GeV} \\ m_Z &= 91.1876 \pm 0.0021 \text{ GeV} \end{aligned} \quad (1.23)$$

The Higgs mechanism (formulated by Higgs, Brout, Englert, and others [35–37]) solves this apparent tension between Glashow’s electroweak theory and the experimental data without abandoning the gauge principle, which is fundamental for the internal consistency of the SM.

Weak massive gauge bosons acquire mass *dynamically*, via interaction with a scalar field $\Phi(x)$. The dynamics of this new field generate a spectrum of physical states which does not reflect the underlying symmetry of the theory: this phenomenon is termed *spontaneous symmetry breaking*. In the SM, the spontaneous breaking of the electroweak gauge group $SU(2)_L \times U(1)_Y$ induces mass terms for the three vector fields $W_\mu^\pm$ and Z_μ , while maintaining the gauge invariance of the electroweak Lagrangian.

Let us consider a simple (unphysical) example to show how the Higgs mechanism works. Given a complex scalar field $\phi(x)$, the following Lagrangian describes its self-interactions via the potential $V(\phi)$ and its couplings with a (massless) gauge field $A_\mu(x)$:

$$\mathcal{L} = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} + (D^\mu\phi)^\dagger (D_\mu\phi) - V(\phi) \quad (1.24)$$

where the covariant derivative is $D_\mu(x)\phi(x) = [\partial_\mu + ieA_\mu(x)]\phi(x)$ and $F^{\mu\nu} = \partial^\mu A^\nu - \partial^\nu A^\mu$ is the usual field strength. The scalar potential contains a quadratic and a quadrilinear self-interaction term for the scalar field ϕ :

$$V(\phi) = \mu^2\phi^\dagger\phi + \lambda(\phi^\dagger\phi)^2 \quad (1.25)$$

The above Lagrangian in Equation 1.24 is invariant under local phase transformation of the $U(1)$ symmetry group, affecting both ϕ and A_μ as follows:

$$\begin{aligned} \phi(x) &\longrightarrow e^{i\theta(x)}\phi(x) \\ A_\mu(x) &\longrightarrow A_\mu(x) - \frac{1}{e}\partial_\mu\theta(x) \end{aligned} \quad (1.26)$$

The shape of the scalar potential $V(\phi)$ is the keystone of the Higgs mechanism. While $\lambda > 0$ in Equation 1.25 guarantees the presence of global minima in $V(\phi)$, the sign of μ^2 discriminates between a unique minimum (with $\mu^2 > 0$ at $\phi_0 = 0$) or a degeneracy




Figure 1.3 – Shape of the scalar potential with $\mu^2 < 0$ in two dimensions as a function of the real and imaginary field components. This shape is typically referred to as “Mexican hat” potential.

of equivalent minima (with $\mu^2 < 0$). In the latter case, as shown in Figure 1.3, there are infinite configurations of the scalar field that minimize the potential, given by:

$$\phi_0 = e^{i\alpha} \sqrt{\frac{-\mu^2}{2\lambda}}, \text{ for } 0 < \alpha < 2\pi \quad (1.27)$$

To proceed with the quantization of the theory, ϕ will acquire one of these equivalent configurations as ground (*vacuum*) state. Although the Lagrangian in Equation 1.24 remains gauge-invariant under U(1) transformations, the symmetry does not manifest in the spectrum of physical states after choosing a particular vacuum: namely, the underlying U(1) symmetry is *hidden* or *spontaneously broken*. For simplicity, let us adopt the choice of vacuum where ϕ_0 is real and positive (i.e. $\alpha = 0$):

$$\phi_0 = \sqrt{\frac{-\mu^2}{2\lambda}} \equiv \frac{v}{\sqrt{2}}, \quad (1.28)$$

where v is referred to as *vacuum expectation value* (*v.e.v.*). After the choice of vacuum, the Lagrangian is reorganized in terms of the oscillations $\phi_1(x)$ and $\phi_2(x)$ around the ground state. Substituting $\phi = \frac{v}{\sqrt{2}} + \frac{1}{\sqrt{2}} [\phi_1 + i\phi_2]$ in Eq. 1.24, new interesting features

emerge:

$$\begin{aligned}
\mathcal{L} = & -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} + \frac{1}{2}e^2v^2A^\mu A_\mu \\
& + \frac{1}{2}\left(\partial_\mu\phi_1\right)^2 - \frac{1}{2}(-2\mu^2)\phi_1^2 - \lambda v\phi_1^3 - \frac{\lambda}{4}\phi_1^4 \\
& + e^2v^2A^\mu A_\mu \frac{\phi_1}{v} + e^2v^2A^\mu A_\mu \frac{\phi_1^2}{v^2} \\
& + \frac{1}{2}\left(\partial_\mu\phi_2\right)^2 + ev(\partial_\mu\phi_2)A_\mu + (\text{others})
\end{aligned} \tag{1.29}$$

A mass term for the gauge field $A_\mu(x)$ appears, with $m_A = ev$. The real scalar field ϕ_1 , referred to as the *Higgs field*, is massive with $m_1 = \sqrt{-2\mu^2}$ and interacts with itself with trilinear and quadrilinear couplings, whose strengths (equal to λv and $\lambda/4$ respectively) are related to its mass via the *v.e.v.*, as described in Eq. 1.28. The coupling strengths for the cubic and quartic vertices involving the Higgs and the gauge fields all include factors e^2v^2 , and are thus proportional to m_A^2 . A second (massless) scalar field $\phi_2(x)$ (termed *Goldstone boson* [38]) and its interaction terms also appear in the above rearrangement of the Lagrangian.

The underlying gauge invariance allows to clarify the particle content of the theory. The function $\theta(x)$ in Equation 1.26 can be chosen such that the transformed ϕ field is real, by *rotating away* the complex phase². With this gauge choice, termed *unitary gauge*, the Goldstone boson disappears from the Lagrangian, which unambiguously describes a massive vector boson and a massive scalar boson (the Higgs boson)³.

1.2.1 The Standard Model Higgs boson

The Higgs mechanism was incorporated in the electroweak sector of the SM by Weinberg and Salam, through a weak isospin doublet of complex scalar fields $\Phi(x)$:

$$\Phi(x) = \begin{pmatrix} \phi^+(x) \\ \phi^0(x) \end{pmatrix} \tag{1.30}$$

The self-interactions of $\Phi(x)$ are described by the scalar potential

$$V(\Phi) = \mu^2\Phi^\dagger\Phi + \lambda(\Phi^\dagger\Phi)^2, \tag{1.31}$$

entering the Lagrangian $\mathcal{L} = (D^\mu\Phi)^\dagger(D_\mu\Phi) - V(\Phi)$, which is manifestly gauge-invariant under $SU(2)_L \times U(1)_Y$, provided that the upper and lower component of $\Phi(x)$ have the

²After the choice of vacuum, the complex scalar field ϕ may be parametrized as $\phi(x) = e^{-i\frac{\chi(x)}{v}}(v + \phi_1(x))$. By enforcing the gauge invariance under $U(1)$, the χ degree of freedom can always be rotated away using the transformation law $\phi(x) \rightarrow e^{i\frac{\chi(x)}{v}}\phi(x)$, leaving only the real component ϕ_1 .

³A colourful description of the Higgs mechanism states that the spurious degree of freedom represented by the Goldstone field ϕ_2 is *eaten up* by the gauge field A_μ , which, as a result, acquires mass.

same weak hypercharge quantum number ($Y_\Phi = 1$). Choosing $\mu^2 < 0$, the potential reaches its minimum for an infinite set of configurations of the complex doublet, all satisfying $\Phi^\dagger \Phi = \frac{-\mu^2}{2\lambda}$. The particular choice of vacuum

$$\Phi_0 = \begin{pmatrix} 0 \\ v/\sqrt{2} \end{pmatrix} \quad (1.32)$$

with $v = \sqrt{\frac{-\mu^2}{\lambda}}$ ensures that the photon remains massless, by preserving the unbroken symmetry under $U(1)_{EM}$, whose generator Q is the electric charge. Hence, the spectrum of physical states still manifests the QED gauge invariance. After spontaneous symmetry breaking the weak isospin doublet is conveniently parametrized as

$$\Phi(x) = \frac{1}{\sqrt{2}} e^{i \sum_{i=1}^3 \theta^i(x) T^i} \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix}, \quad (1.33)$$

where $\theta^i(x)$, $i = 1, 2, 3$ and $H(x)$ are real scalar fields. The gauge-invariance under $SU(2)_L$ allows to rotate away the spurious degrees of freedom $\theta^i(x)$, $i = 1, 2, 3$. In the unitary gauge,

$$\Phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix}, \quad (1.34)$$

where $H(x)$ is the Higgs field. The mass terms for the weak bosons $W^\pm$ and Z emerge from substituting Equation 1.34 in the kinetic term of the scalar Lagrangian:

$$\begin{aligned} (D^\mu \Phi)^\dagger (D_\mu \Phi) &= \frac{1}{2} (\partial_\mu H)^2 \\ &+ \frac{g^2 v^2}{4} W^{+, \mu} W_\mu^- + \frac{1}{2} \frac{g^2 v^2}{4 \cos^2 \theta_W} Z^\mu Z_\mu \\ &+ \frac{g^2 v}{2} W^{+, \mu} W_\mu^- H + \frac{g^2}{4} W^{+, \mu} W_\mu^- H^2 \\ &+ \frac{g^2 v}{4 \cos^2 \theta_W} Z^\mu Z_\mu H + \frac{g^2}{8 \cos^2 \theta_W} Z^\mu Z_\mu H^2 \end{aligned} \quad (1.35)$$

Mass terms for the weak bosons $W^\pm$ and Z appear, with

$$m_W = \frac{gv}{2}; \quad m_Z = \frac{gv}{2 \cos \theta_W} = \frac{m_W}{\cos \theta_W} \quad (1.36)$$

Moreover, both $W^\pm$ and Z interact with the Higgs field via cubic and quartic vertices, whose couplings - typically expressed as $g_{HVV} \cdot v$ and g_{HHVV} , where both g_{HVV} and g_{HHVV} are dimensionless quantities - are proportional to square of the weak bosons masses. In fact, using the definitions of m_W and m_Z given in Equation 1.36, the third and

fourth lines of Equation 1.35 can be rewritten as:

$$\begin{aligned}\mathcal{L}_{int}^{EWSB} = & \frac{2m_W^2}{v^2} \cdot v \cdot W^{+, \mu} W_{\mu}^{-} H + \frac{m_W^2}{v^2} \cdot W^{+, \mu} W_{\mu}^{-} H^2 \\ & + \frac{m_Z^2}{v^2} \cdot v \cdot Z^{\mu} Z_{\mu} H + \frac{m_Z^2}{2v^2} \cdot Z^{\mu} Z_{\mu} H^2\end{aligned}\quad (1.37)$$

When calculating scattering amplitudes, the interaction terms in Equation 1.37 give rise to cubic and quartic vertices involving the weak bosons and the Higgs boson, as illustrated in Figure 1.35.




Figure 1.4 – Feynman diagrams illustrating the cubic (left) and quartic (right) interaction vertices between the W or Z bosons and the Higgs bosons. The quartic vertex involving two W bosons and two Higgs bosons (HHW^+W^-) includes an additional factor of 2 in its coupling constant compared to the $\frac{m_W^2}{v^2}$ term in the Lagrangian (Equation 1.37), which arises from a combinatorial factor due to the two identical Higgs bosons, either of which can occupy the two legs of the interaction. Similarly, combinatorial factors of 2 also appear in the two bottom diagrams, resulting from the presence of either two identical Z bosons or two identical Higgs bosons.

The coupling constants for these interactions are related to each other and to the weak

boson masses by the following relation:

$$g_{HVV} = g_{HHVV} = \frac{m_V^2}{v^2} \quad (1.38)$$

This relationship is a direct result of the weak doublet structure of the Higgs field and the process of electroweak symmetry breaking. Specifically, both the interaction terms and the weak boson mass terms in Equation 1.35 are derived from expanding the covariant derivative applied to the Higgs field Φ , as well as the $(v + H)^2$ term, which originates from the bottom component of the doublets Φ and $\Phi^\dagger$ following the electroweak symmetry breaking.

The structure of the scalar potential unveils the Higgs boson self-interactions: by substituting Equation 1.34 in Equation 1.31, one finds

$$V(\Phi) = \frac{1}{2}(-2\mu^2)H^2 + \lambda v H^3 + \frac{\lambda}{4}H^4 \quad (1.39)$$

Hence, oscillations of the Higgs field correspond to a massive scalar particle (namely, the *Higgs boson*) with mass $m_H = \sqrt{-2\mu^2}$. The couplings of the Higgs boson self-interactions are controlled by the $v.e.v.$ and the Higgs boson mass:

$$\lambda = \frac{-\mu^2}{v^2} = \frac{m_H^2}{2v^2} \quad (1.40)$$

The $v.e.v$ can be measured via muon decay ($v = \sqrt{1/\sqrt{2}G_F}$, where $G_F = 1.1663788(6) \cdot 10^{-5} \text{ GeV}^{-2}$ [6] is the Fermi coupling constant, yielding $v \approx 246 \text{ GeV}$), while $m_H \approx 125 \text{ GeV}$. Therefore, once the $v.e.v.$ and the Higgs boson mass are known, the Higgs boson SM potential and its self-couplings are fully determined. This result is important from a phenomenological point of view: as will be explained in Chapter 2, the trilinear self-interaction term contributes to Higgs boson pair production with a well defined dimensionless SM coupling:

$$\lambda_{HHH}^{SM} = \lambda = \frac{m_H^2}{2v^2} \approx 0.13 \quad (1.41)$$

The observation of Higgs boson pairs would provide an independent measurement for λ_{HHH} , thus shedding light on the exact structure of the Higgs boson potential close to its minima.

Mass generation for fermions

The chiral nature of the electroweak gauge group forbids explicit mass terms for fermions: $-m\bar{\psi}\psi = -m(\bar{\psi}_L\psi_R + \bar{\psi}_R\psi_L)$ is not gauge-invariant, because $SU(2)_L \times U(1)_Y$ affects the left-handed and right-handed chiral components of the fermion field differently. However, experimental data confirm that all fermions are massive: again, the solution comes

from the interaction with $\Phi(x)$ and its charge conjugate $\Phi^C(x) = -iT^2\Phi(x)$ and spontaneous symmetry breaking.

Let us consider a single quark family. Gauge-invariant Yukawa interaction terms involving the weak isospin doublet $\begin{pmatrix} u_L \\ d_L \end{pmatrix}$ and the right-handed components u_R and d_R are included in the SM Lagrangian:

$$\mathcal{L} = y_u (\bar{u}_L \bar{d}_L) \cdot \Phi^C u_R + y_d (\bar{u}_L \bar{d}_L) \cdot \Phi d_R + \text{h.c.} \quad (1.42)$$

After spontaneous symmetry breaking, the Yukawa Lagrangian gives rise to mass terms for both the up- and down-type quarks ($-m_u \bar{u}u$ and $-m_d \bar{d}d$), and interaction terms between the massive fermions and the Higgs field. By expanding the Higgs doublet Φ in terms of its vacuum expectation value (*v.e.v.*) and the real scalar field H , as shown in Equation 1.34, the Yukawa Lagrangian takes the form:

$$\begin{aligned} \mathcal{L}^{Yukawa} = & -\frac{y_u v}{\sqrt{2}} \bar{u}u - \frac{y_d v}{\sqrt{2}} \bar{d}d \\ & -\frac{y_u}{\sqrt{2}} \bar{u}uH - \frac{y_d}{\sqrt{2}} \bar{d}dH \end{aligned} \quad (1.43)$$

The fermion masses $m_{u(d)} \equiv \frac{y_{u(d)} v}{\sqrt{2}}$, defined in the first line of Equation 1.43, are directly proportional to the *Yukawa couplings* $y_{u(d)}$, which dictate the interaction strengths between the fermions and the Higgs boson, as shown in the second line of Equation 1.43.

1.3 Higgs boson physics at hadron colliders

As shown in Section 1.2, the fermions and the weak gauge bosons acquire mass dynamically, via their interactions with the Higgs field. Hence, the strength of the interaction between the Higgs bosons and the other SM particle is driven by their mass: given a fermion (weak boson) with mass m , its coupling with the Higgs boson is proportional to m (m^2). Therefore, the SM privileges the interactions between the Higgs boson and the heaviest particles, namely top quark and the weak bosons, while the Higgs boson couplings to the lighter particles (such as the leptons or the lighter quarks) is suppressed. In addition, there is no direct coupling between the Higgs and the gluon fields or the photon fields, which is a direct consequence of the manifest gauge invariance of the SM under $SU(3)_C \times U(1)_{EM}$. Indeed, both the color and the electromagnetic gauge symmetries are unbroken, and the corresponding fields remain massless.

Hence, while the Higgs boson production and decay processes are primarily influenced by its couplings to the top quarks and weak bosons, given that most SM particles have mass, the Higgs boson phenomenology is rich and varied. In addition, Higgs boson production cross-sections and branching fractions are governed by the Higgs boson mass m_H , whose value is not predicted by the theory. A precise measurement of m_H is therefore fundamental for confirming the SM description of the Higgs boson interactions.

1.3.1 Higgs boson production mechanisms

Production of Higgs bosons at hadron colliders can occur through various mechanisms. The production mode with the largest cross section is gluon-gluon fusion (ggF), followed by vector boson fusion (VBF), vector boson associated production (VH, including WH and ZH), top-antitop quark pair associated production ($t\bar{t}H$), bottom-antibottom quark pair associated production ($b\bar{b}H$), and single top quark associated production (tHW and tHq). Figure 1.5 reports tree-level diagrams contributing to the main Higgs boson production modes at the LHC. Figure 1.6 shows the behaviour of the Higgs boson production cross sections through all the different production mechanisms available at hadron colliders, as a function of the center of mass energy of the proton-proton (pp) collisions (while fixing $m_H = 125$ GeV), while Table 1.2 quotes the Higgs boson production cross sections assuming $m_H = 125$ GeV and pp collision at the center of mass energy of $\sqrt{s} = 13$ TeV.

Production mode	Cross section [pb]	QCD Scale [%]	PDF + α_S [%]
ggF	48.58	+4.56 -6.72	3.20
VBF	3.797	+0.4 -0.3	2.1
WH	1.3728	+0.5 -0.7	1.1
qq $\rightarrow$ ZH	0.7612	+3.9 -3.0	1.6
gg $\rightarrow$ ZH	0.1226	+3.9 -3.0	1.6
$t\bar{t}H$	0.5071	+5.8 -9.2	3.6
$b\bar{b}H$	0.4880	+20.2 -23.9	
tHq (t-channel)	7.425×10^{-2}	+6.5 -14.9	3.7
tHq (s-channel)	2.879×10^{-3}	+2.4 -1.8	2.2
tHW	1.517×10^{-2}	+4.9 -6.7	6.3

Table 1.2 – Standard Model production cross sections [40] for a Higgs boson of $m_H = 125$ GeV in pp collision at $\sqrt{s} = 13$ TeV. Their relative theoretical uncertainties on QCD scale and PDF + α_S are reported. For $b\bar{b}H$, only the total uncertainty is provided.

Despite having no direct coupling to gluons, the ggF (Figure 1.5a) is the dominant Higgs boson production mode at the LHC, accounting for about 87% of the total cross section. Two gluons emerging from the colliding protons couple with the Higgs boson via a quark loop, whose main contribution comes from top quarks (due to their large mass), while the contribution of lighter quarks are suppressed. The state of the art of the theoretical calculation of the ggF production cross section reach next-to-next-to-next-to-leading order (N^3 LO) accuracy in QCD and involve EW corrections at next-to-leading order (NLO), and are performed in the limit of $m_t \rightarrow \infty$. The prediction for a 125-GeV Higgs boson, including the central value as well as its theoretical uncertainty (theory) and the uncertainty propagated from the Parton Distribution Functions (PDF) and the value of the strong coupling constant α_S (PDF + α_S), is

$$\sigma_{\text{ggF}}(H) = 48.58^{+2.22}_{-3.27} \text{ pb (theory)} \pm 1.56 \text{ pb (PDF + } \alpha_S) \text{ [41]} \quad (1.44)$$




Figure 1.5 – Leading order Feynman diagrams for SM Higgs boson production at LHC [3].




Figure 1.6 – Cross sections for the SM Higgs boson production modes [39] as a function of the center of mass energy of pp collisions. The Higgs boson mass is fixed at $m_H = 125$ GeV.

The VBF mechanism (whose Feynman diagram is visible in Figure 1.5b) is the second most abundant Higgs boson production mode at the LHC, contributing approximately 7% to the total cross section, as shown in Table 1.2. It occurs through the scattering of two vector bosons (either W or Z bosons), emitted by two quarks within the two colliding protons. The two vector bosons fuse together and generate a Higgs boson. The two quarks involved in the VBF process will then hadronize into two highly energetic jets, which are typically observed in the forward and backward regions of the detector. The next most relevant Higgs boson production mode is the associated production with a vector boson (either W or Z). Figure 1.5c and 1.5d show two tree-level Feynman diagram contributing to this production mechanism. As for the VBF, this process occurs via the Higgs boson coupling with a vector boson, which is typically emitted via the annihilation of two quarks from the colliding protons. The Z boson associated production could also occur via fusion of two gluons into a Z (via a quark loop), which then radiates a Higgs boson. This process is typically referred to as $gg \rightarrow ZH$ (Figure 1.5d). The VH processes are tagged by the final state particles produced by the W/Z bosons decay. Finally, the $b\bar{b}H$, $t\bar{t}H$, and tH production modes (whose Feynman diagrams are represented in Figures 1.5e, 1.5f, 1.5g, and 1.5h) are the ones with the lowest cross sections. These processes are tagged by the presence of multiple jets in the final state, and, in particular, the presence of jets from the hadronization of bottom quarks (termed b -jets). In $t\bar{t}H$, $b\bar{b}H$, and tH processes, bottom quarks are either produced in association with the Higgs boson, or are created via the top quark decay.

1.3.2 Higgs boson decay modes

After being produced in pp collisions, the Higgs boson lives for an extremely short time (approximately 1.6×10^{-22} s according to the SM [40]) before decaying into other parti-

cles. Figure 1.7 presents the main decay modes for the SM Higgs boson, and shows their Branching fractions ($\mathcal{BR}$ s) as a function of the Higgs boson mass in a window between 120 GeV and 130 GeV. The $\mathcal{BR}$ for a specific Higgs boson decay is defined as its partial width divided by the total Higgs boson width:

$$\mathcal{BR}(H \rightarrow X_i) = \frac{\Gamma(H \rightarrow X_i)}{\sum_i \Gamma(H \rightarrow X_i)} \quad (1.45)$$

The leading decay mode for a SM Higgs boson with $m_H = 125$ GeV is $H \rightarrow b\bar{b}$, with




Figure 1.7 – Branching fractions [39] for the main decay modes of the SM Higgs boson as a function of m_H varying in the window between 120 GeV and 130 GeV.

a $\mathcal{BR}$ of 58.24%. In principle, all the heavy fermions can constitute final states for a Higgs boson decay (except for the top quark, being $m_t \approx 172.5$ GeV [42]), and the corresponding $\mathcal{BR}$ is proportional to the fermion masses. The ATLAS and CMS experiments at the LHC observed the Higgs boson decays to the third generation fermions (bottom quarks and τ leptons) in pp collision data delivered by the LHC during its first and second data-taking periods (Run 1 and Run 2) [43–46], and are now seeing hints of Higgs boson couplings to second generation fermions (namely, muons) [47, 48]. The second highest $\mathcal{BR}$ is provided by the decay into a pair of weak gauge bosons (WW^* or ZZ^* , with one of them being off-shell). The $H \rightarrow ZZ^* \rightarrow 4\ell$ decay, where two lepton pairs originate from each Z boson, constituted one of the two *golden channels* for the discovery of the Higgs boson [1, 2]: this is due to the clear signals left by electrons and muons in the ATLAS and CMS detectors, the precise invariant mass resolution of the four leptons in the final state, and the low background from other SM processes that can mimic the same four-lepton signature, allowing a signal-to-background ratio close to unity.

Decays into massless gauge bosons (namely, gluons and photons) are also possible, and occur via a loop of heavy quarks or $W^\pm$ bosons (only in the case of photons). The de-

cay into gluons has a larger $\mathcal{BR}$ ($\sim 9\%$), but it is indistinguishable from the huge QCD background produced in a hadron collider. Conversely, the decay into a pair of photons has a tiny $\mathcal{BR}$ (0.23% for a SM Higgs boson), but is nevertheless able to provide an extremely clear signature. In addition, the excellent mass resolution of the diphoton system allows to efficiently isolate Higgs boson decays from competing backgrounds. For these reasons, the $H \rightarrow \gamma\gamma$ decay mode was the second golden channel for Higgs boson discovery.


1.3.3 Measurements of Higgs boson properties at the LHC

The observation of the Higgs boson announced in 2012 by the ATLAS and CMS experiments at the LHC [1, 2] prompted a decade-long (currently still ongoing) campaign of measurements of the properties of this long-sought particle, all reporting, up to this date, an excellent agreement with the SM predictions. At the end of Run 2 (2015-2018), the LHC delivered to the ATLAS and CMS experiments a huge amount of pp collision data at $\sqrt{s} = 13$ TeV, corresponding to an integrated luminosity of approximately 140 fb^{-1} [49]. This dataset offered a unique opportunity to improve our understanding of the Higgs boson, and allowed to reach an unprecedented precision in the measurements of its mass, couplings, and quantum numbers. The following Section provides a short review of the most up-to-date combination of measurements of the Higgs boson properties performed by the ATLAS and CMS collaborations using the LHC Run 2 data.


Higgs boson mass, width, spin, and parity

The golden channels for the measurement of the Higgs boson mass are the $H \rightarrow \gamma\gamma$ and $H \rightarrow ZZ^* \rightarrow 4\ell$ decays. Both final states are fully reconstructable, and benefit from the excellent resolution for the energy and momentum of electrons, muons, and photons with both the ATLAS and CMS detectors. The ultimate precision from a single experiment is reached by combining the two measurements in the di-photon and four-lepton channels. A combined legacy measurement of the Higgs boson mass, using data collected by the ATLAS and CMS detectors during Run 1, determined the mass to be $m_H = 125.09 \pm 0.21(\text{stat.}) \pm 0.11(\text{syst.}) \text{ GeV}$ [5]. Both ATLAS and CMS also measured the Higgs boson mass using Run 2 data. Although a combined ATLAS and CMS result relying on Run 2 data is not yet available, both experiments have independently performed Higgs boson mass measurements in the $H \rightarrow \gamma\gamma$ and $H \rightarrow ZZ^* \rightarrow 4\ell$ channels. The ATLAS combination of these channels, incorporating both Run 1 and Run 2 data [50], represents the most precise measurement to date, as shown in Figure 1.8a. CMS has also performed separate measurements in the di-photon [51] and four-lepton [52] final states, combining Run 1 and Run 2 data, whose results are shown in Figures 1.8b and 1.8c, respectively. The CMS measurement in the $H \rightarrow ZZ^* \rightarrow 4\ell$ channel is the most precise single-channel determination of the Higgs boson mass to date.


The decay mode into a pair of Z bosons is also essential for determining the natural decay width of the Higgs boson, Γ_H . This parameter affects only the production cross section of an on-shell Higgs boson, while the production of an off-shell Higgs boson is independent of Γ_H . Therefore, by measuring the ratio $\sigma^{\text{on-shell}} / \sigma^{\text{off-shell}}$ in the ZZ^*



(a)



(b)



(c)

Figure 1.8 – Summary of Higgs boson mass measurements from the individual analyses as well as the available combinations for both ATLAS ([50]) and CMS ([51] and [52]), using the Run 1 and Run 2 datasets. The statistical component to the total uncertainty (black error bar) is highlighted in yellow. The ATLAS results also highlight the systematic contribution to the total uncertainty (purple band).

final state and assuming no potential new physics contributions differently affect the on-shell and off-shell cross sections, both the ATLAS and CMS experiments successfully measured for the first time with Run 2 data the total Higgs boson decay width. ATLAS found $\Gamma_H = 4.5_{-2.5}^{+3.3}$ MeV [53] and CMS determined $\Gamma_H = 3.2_{-1.7}^{+2.4}$ MeV [54]: both measurements are consistent with the SM prediction ($\Gamma_H = 4.1$ MeV [40]).

Finally, the spin (J), and parity (P) of the Higgs boson can be measured exploiting its bosonic decay modes (namely, $H \rightarrow \gamma\gamma$, $H \rightarrow ZZ^*$, and $H \rightarrow WW^*$). Using the angular distributions of the particles in the final state (either photons or leptons from the W and Z bosons subsequent decays), the ATLAS and CMS experiments were able to exclude several non-SM spin-parity configurations of the Higgs boson with a confidence level larger than 99.9% already using Run 1 data [55, 56]. These results strongly support the $J^P = 0^+$ configuration, consistent with the SM prediction.

Higgs boson signal strengths and couplings

A combined measurement of the Higgs boson production cross sections in different decay channels can be harnessed to extract a portrait of the Higgs boson that is as complete as possible. The combined model is typically expressed in terms of *signal strengths*, whose measurements quantify the agreement between the observed data and the SM predictions. The signal strength (typically referred to as μ) modifies the expected SM cross section (or branching fraction) into the measured value. In other words, for a given process $Y \rightarrow H \rightarrow X$, one can define a combined signal strength as well as individual signal strengths for the production and the decay, corresponding to:

$$\begin{aligned}\mu_X^Y &= \frac{\sigma(Y \rightarrow H) \times \mathcal{BR}(H \rightarrow X)}{\sigma^{SM}(Y \rightarrow H) \times \mathcal{BR}^{SM}(H \rightarrow X)} \\ \mu^Y &= \frac{\sigma(Y \rightarrow H)}{\sigma^{SM}(Y \rightarrow H)} \\ \mu_X &= \frac{\mathcal{BR}(H \rightarrow X)}{\mathcal{BR}^{SM}(H \rightarrow X)}\end{aligned}\quad , \quad (1.46)$$

where $\sigma(Y \rightarrow H)$ ($\mathcal{BR}(H \rightarrow X)$) represents the measured cross section (branching fraction), while $\sigma^{SM}(Y \rightarrow H)$ ($\mathcal{BR}^{SM}(H \rightarrow X)$) is the corresponding SM prediction.

Figure 1.9 shows the individual signal strengths of the different Higgs boson production processes and decay branching fractions measured by CMS using Run 2 data. Conversely, Figure 1.10 reports the signal strengths for all the different combinations between Higgs boson production and decay modes extracted from ATLAS Run 2 data. All these measurements show an excellent agreement with the SM predictions over several order of magnitude.

Both production cross sections and decay rates are governed by the couplings of the Higgs boson to the massive fermions and the weak bosons, and are expected to respond consistently to any deviation of these couplings from the SM predictions. Such deviations are quantified using the κ framework, which requires that the Higgs boson cross




Figure 1.9 – Summary of signal strengths for various Higgs boson production modes (left) and decay processes (right), measured by CMS using Run 2 data [4]. The signal strengths for the production cross sections μ^Y are measured assuming branching fractions equal to the SM predictions and vice-versa. The uncertainties are expressed in terms of $1(2)\sigma$ confidence intervals, represented by the thick (thin) black lines. The statistical and systematic contributions to the 1σ uncertainty are highlighted using blue and red bands respectively.




Figure 1.10 – Signal strengths for different combinations of Higgs boson production and decay modes observed measured by ATLAS using Run 2 data [3]. The central values are shown together with the total uncertainty and its systematic component, highlighted as a coloured band.

sections and decay widths are parametrized in terms of *coupling modifiers* (termed κ), which change the strength of the couplings of the Higgs boson to other particles compared to the SM expectations. For instance, the branching fraction of the Higgs boson to bottom quarks is regulated by the Yukawa bottom coupling and is proportional to y_b^2 : hence, within the κ framework, $\mathcal{BR}(H \rightarrow b\bar{b})$ is parametrized as $\mathcal{BR}(H \rightarrow b\bar{b}) = \kappa_b^2 \cdot \mathcal{BR}^{SM}(H \rightarrow b\bar{b})$. A combined measurement of the Higgs boson cross sections and decay rates allows a simultaneous extraction of κ modifiers, and quantify the agreement of the observed couplings of the Higgs boson to a given particle with the SM. The κ framework allows to explore interference effects of the Higgs boson couplings between production and decay processes. This result is shown in Figure 1.11 as a function of the mass of the probed particles, using data collected the ATLAS (left) and CMS (right) experiments during the Run 2 of the LHC [3, 4]. In both cases, the absolute coupling strength modifiers - defined as $\sqrt{\kappa_V} \frac{m_V}{v}$ for the vector bosons and $\kappa_F \frac{m_F}{v}$ for the fermions, where m_V (m_F) is the mass of the vector boson (fermion) and v is the $v.e.v.$ - are shown, in order to keep proportionality to the particle mass (as predicted by the SM and described in Section 1.2.1). Figure 1.11 establishes a striking agreement of the measured couplings with the SM description of the Higgs mechanism across several order of magnitude in the masses of the probed particles. However, this picture is not yet complete. The couplings of the Higgs boson to the first generation fermions (electrons and lighter quark), that would populate the low mass frontier of Figure 1.11a and 1.11b, is out of reach for the LHC experiments. In addition, a prominent absence can be noticed in the high range of the SM particle mass spectrum: as explained in Section 1.2, the SM predicts that the Higgs boson couples with two other Higgs bosons via the trilinear self-coupling λ_{HHH} , whose modifier is defined as:

$$\kappa_\lambda = \lambda_{HHH} / \lambda_{HHH}^{SM} \quad (1.47)$$

which remains, to this day, unknown. Observing an agreement of the absolute trilinear self-coupling of the Higgs boson (namely, $\kappa_\lambda \cdot v$) with the diagonal lines in Figure 1.11 would provide the ultimate validation of the SM Higgs mechanism.




Figure 1.11 – Summary of the coupling strength modifiers for the Higgs boson to different particles measured by ATLAS (1.11a) and CMS (1.11b) using Run 2 data. In both cases, the absolute coupling strength modifiers are shown as a function of the mass of the probed particles, and the SM prediction is represented by the diagonal line.

Constraining the Higgs boson potential at colliders

The Higgs boson potential plays a central role in Higgs boson physics. It is responsible for the electroweak symmetry breaking (EWSB), and for the generation of mass of all the massive elementary particles in the SM (namely, the massive fermions and the weak bosons). The EWSB induces trilinear and quadrilinear self-interaction terms for the Higgs field, whose coupling strengths govern the shape of the scalar potential. Specifically, the trilinear Higgs boson self-coupling λ_{HHH} (already introduced in Section 1.2.1) determines the structure of the Higgs boson potential close to its minima. Its behaviour has profound implications for our understanding of particle physics, impacting both the TeV scale (probed by the experiments at the LHC) and much higher energy scales, such as the Planck scale, where quantum gravity effects may come into play.

For instance, the running of λ_{HHH} with the energy scale controls the stability of the vacuum [57, 58]. Current precision measurements of the Higgs boson mass and the top quark mass (influencing radiative corrections to λ_{HHH}) point to a second minimum in the Higgs boson potential, situated at a lower energy compared to our current vacuum state. This indicates that this vacuum may be metastable, and a tunnel transition to this new minimum could be possible, with a timescale set by the behaviour of λ_{HHH} in the ultraviolet region. Additionally, λ_{HHH} drives the nature of the electroweak phase transition [59, 60], a critical moment in the early Universe when the Higgs field acquired a non-zero vacuum expectation value ($v.e.v$) and spontaneously broke electroweak gauge invariance.

However, the Higgs boson self-coupling is still largely experimentally unconstrained. The only way to directly measure λ_{HHH} at the LHC is through the production of multiple Higgs bosons. The trilinear self-coupling term in the scalar potential generates an interaction vertex involving three Higgs bosons. This allows a single Higgs boson to decay into two additional Higgs bosons, therefore contributing to the production of Higgs boson pairs. According to the SM, the trilinear self-coupling influences the production of Higgs boson pairs through a scattering amplitude proportional to λ_{HHH}^{SM} , which is fully determined by the Higgs boson mass m_H and the $v.e.v$. An independent measurement of λ_{HHH} (or, equivalently, of its coupling modifier $\kappa_\lambda = \lambda_{HHH}/\lambda_{HHH}^{SM}$) from Higgs boson pair production in agreement with the SM expectations would strongly support our current view of the Higgs mechanism. Conversely, any deviation of κ_λ from unity emerging from Higgs boson pair production measurements would challenge our current

understanding of the EWSB and inevitably open the door to new physics scenarios. However, Higgs boson pair production is an extremely rare process, with a cross section three orders of magnitude smaller than that of a single Higgs boson. Although no Higgs boson pair production has been observed at the LHC so far, experimental sensitivity is gradually approaching the levels required to detect SM-like scenarios. Investigating and ultimately observing Higgs boson pair production is a key goal of the LHC physics program, especially with the increased statistics from the Run 3 dataset and the upcoming High Luminosity LHC phase. In particular, the sensitivity of the HH analyses performed by the ATLAS and CMS Collaborations is improving at a rate that outpaces the increase in collected data. If this trend continues, the combined data from the ATLAS and CMS experiments collected during the Run 2 and Run 3 periods could reveal the first hint of SM HH production.

2.1 Higgs boson pair production in the SM

2.1.1 Higgs boson pair production modes

The production mechanisms for SM Higgs boson pairs at hadron colliders mirror those of a single Higgs boson, already discussed in Section 1.3. The most abundant production mechanism is gluon-gluon Fusion (ggF HH), followed by vector boson fusion (VBF HH). In the SM scenario, these two dominant production modes account for more than 95% of the total SM di-Higgs cross section, and thus have been the primary focus of both theoretical calculations and experimental searches at the LHC. The next relevant production mode is top quark pair associated production ($t\bar{t}$ HH), followed by vector boson associated production (VHH, including both $W^\pm$ HH and ZHH), and single top quark associated production (t jHH). Figure 2.1 shows the di-Higgs production cross sections for the different mechanisms as a function of the center of mass energy of pp collisions, while fixing $m_H = 125$ GeV and all the couplings affecting HH production to the SM expectations.

Similarly to the single Higgs boson production, Higgs boson pair production via ggF occurs with two gluons emerging from the colliding protons, coupling with the Higgs bosons via a heavy quark loop (predominantly top quarks). At tree level, the ggF HH production mode proceeds via two amplitudes, shown in Figure 2.2, typically referred to as *box diagram* (Figure 2.2a) and *triangle diagram* (Figure 2.2b). The first diagram does not involve the Higgs boson self-couplings and is solely driven by the Higgs-top quark Yukawa interaction, characterized by the coupling modifier $\kappa_t = y_t/y_t^{SM}$. In contrast, the triangle diagram depends on both κ_t and κ_λ . The contributions of each amplitude and their interference to the double Higgs invariant mass distribution m_{HH} are illustrated in Figure 2.3. The destructive interference between the box and triangle diagrams generates an extremely small cross section:

$$\sigma_{\text{ggF}}(\text{HH}) = 31.05^{+1.86(+6\%)}_{-7.14(-23\%)} \text{ fb (scale+} m_t) \pm 0.71 (\pm 2.3\%) \text{ fb (PDF + } \alpha_s) \quad (2.1)$$




Figure 2.1 – Cross sections for double Higgs production at hadron colliders via gluon gluon fusion, vector boson fusion, vector boson associated production, and double Higgs production with top quarks ($t\bar{t}HH$ and $tjHH$) as a function of the center of mass energy of the pp collisions [11]. The perturbative order of the state-of-the-art theoretical predictions for each production mode is also indicated. Specifically, the ggF HH cross section is calculated at NNLO accuracy in QCD, incorporating finite top quark mass effects (FTapprox), while the VBF HH cross section is evaluated at $N^3\text{LO}$ accuracy.

This prediction, evaluated for SM $m_H = 125$ GeV Higgs boson pairs generated in pp collisions at $\sqrt{s} = 13$ TeV, achieves NNLO accuracy in QCD and incorporates top quark mass effects [61]. Its dominant uncertainty (“scale+ m_t ”), evaluated in [62], arises from the finite order of the perturbative expansions and the definition of the top quark mass in QCD corrections, while the uncertainty propagated from the PDFs and value of α_S (“PDF+ α_S ”) is subleading [63].




Figure 2.2 – Tree-level diagrams contributing to double Higgs production via ggF.




Figure 2.3 – Differential ggF HH cross section as a function of the di-Higgs invariant mass distribution [11]. The contribution to the ggF HH cross section (solid black line) from the box diagram, triangle diagram, and their interference are highlighted using dashed blue, red, and green lines respectively.

The SM ggF HH cross section in Equation 2.1 illustrates just how rare Higgs boson pair production is. Over the entire Run 2 data-taking period, during which the ATLAS detector recorded pp collisions at a center-of-mass energy of 13 TeV corresponding to an integrated luminosity of 140 fb^{-1} , the SM predicts the production of only about 4300 Higgs boson pairs. By comparison, the cross section for single Higgs boson production predicts more than eight million Higgs bosons in the same dataset - over three orders of magnitude greater than the number of Higgs boson pairs.

The VBF HH production mode involves the scattering of two weak bosons ($W^\pm$ or Z) radiated from two valence quarks inside the colliding protons. At tree level, the VBF HH cross section is driven by three diagrams, shown in Figure 2.4, scaling with $\kappa_\lambda \kappa_V$ (Figure 2.4a), κ_{2V} (Figure 2.4b), and κ_V^2 (Figure 2.4c) respectively, where κ_V and κ_{2V} describe the coupling strengths of the HVV and $HHVV$ interaction vertices, normalized to their SM predictions. Hence, double Higgs production via VBF depends from the trilinear Higgs boson self-coupling modifier κ_λ and simultaneously provides a unique probe to the electroweak symmetry breaking dynamics, involving cubic and quartic interactions between vector bosons and Higgs bosons. The state of the art for the theoretical calculation of the SM VBF HH cross section reaches $N^3\text{LO}$ accuracy in QCD and involves EW corrections at NLO. The prediction for pp collisions at $\sqrt{s} = 13$ TeV, assuming $m_H = 125$ GeV and all the other couplings affecting VBF HH production fixed to the SM expectations, is:

$$\sigma_{\text{VBF}}(\text{HH}) = 1.726^{+0.0005(+0.03\%)}_{-0.0007(-0.04\%)}(\text{scale}) \pm 0.036(\pm 2.1\%) \text{fb}(\text{PDF} + \alpha_s) \quad [63\text{--}66] \quad (2.2)$$

where the dominant contribution to the uncertainty arises from the $\text{PDF} + \alpha_s$ component, while the scale uncertainty (related to missing higher orders in the QCD perturbative expansion) is subleading. The defining feature of the final state in the VBF HH production mode is the presence of two highly energetic hadronic jets, typically reconstructed in the forward and backward regions of the particle detector. These two VBF *jets* originate from the hadronization of the two valence quarks that are involved in the hard scattering process. This signature, alongside the double Higgs final state, helps to efficiently discriminate between signal and competing background processes. For this reason, despite its extremely small SM cross section (almost 20 times smaller than the ggF HH cross section), the experimental study of Higgs pairs produced via VBF is particularly interesting.

2.1.2 Dependence of the Higgs boson pair production cross section from κ_λ

Both the dominant production modes for Higgs boson pair production at the LHC, ggF HH and VBF HH, are influenced by the Higgs boson trilinear self-coupling modifier κ_λ , as demonstrated by the tree-level Feynman diagrams in Figures 2.2b and 2.4a. Both the ggF HH and VBF HH cross sections exhibit a quadratic dependence on κ_λ , which is evident from Figure 2.5. The prediction for ggF HH combines the full NLO accuracy in QCD (including top quark mass effect) and the higher NNLO accuracy in the limit of $m_t \rightarrow \infty$ [67]: this framework is used to extract the ggF HH cross section for any κ_λ value. On the other hand, predictions for the VBF HH cross section for non-SM κ_λ values are available only at LO precision in QCD. Hence, the VBF HH prediction is based on the SM calculation (performed at $N^3\text{LO}$ accuracy in QCD), and on LO predictions for non-SM κ_λ values extracted using the MADGRAPH5_AMC@NLO event generator [68] (which are then rescaled using the SM $N^3\text{LO}$ to LO ratio). A quadratic fit is then used to interpolate the available points. As evident from Figure 2.5, non-SM κ_λ values can signif-




Figure 2.4 – Tree-level diagrams contributing to double Higgs production via VBF.




Figure 2.5 – Cross sections for double Higgs production at hadron colliders via gluon gluon fusion and vector boson fusion as a function of the Higgs boson trilinear self-coupling modifier κ_λ . These prediction assume pp collisions at $\sqrt{s} = 13 \text{ TeV}$, the Higgs boson mass equal to 125 GeV, and all the other couplings affecting HH production fixed to the SM expectations. The uncertainty bands represent the theoretical uncertainty affecting the predictions (namely, the scale+ m_t and the PDF+ α_s component for ggF HH, and the scale and the PDF+ α_s component for VBF HH).

icantly enhance the di-Higgs cross section for both the ggF and VBF production modes. Therefore, new physics scenarios with κ_λ values deviating from unity can be constrained with less data compared to what is required to observe SM HH production. Although VBF HH is less abundant than ggF HH, it experiences a greater enhancement with non-SM κ_λ values. Therefore, when investigating such new physics scenarios, considering HH production via VBF in addition to the main production mode helps to maximize the sensitivity.

Figure 2.6 presents the distributions of m_{HH} for Higgs boson pairs produced via ggF (Figure 2.6a) and VBF (Figure 2.6b) for the SM case and various κ_λ values. For both production modes, larger absolute values of κ_λ (positive or negative) result in softer m_{HH} spectra compared to the SM scenario, where the m_{HH} distribution peaks around 400 GeV. A distinctive double peak in the m_{HH} distribution for ggF HH appears at $\kappa_\lambda = 2.4$, corresponding to the maximum destructive interference between the triangle and box diagrams, which coincides with the minimum ggF HH cross section (see Figure 2.5).




Figure 2.6 – Distribution of the Higgs boson pair invariant mass (m_{HH}) for Higgs boson pairs produced via ggF (left) VBF (right) in the SM and for various anomalous κ_λ values.

2.1.3 Dependence of the Higgs boson pair production cross section from κ_{2V} and κ_V

As previously mentioned in Section 1.2.1, the SM Higgs mechanism induces a well-defined relation between the dimensionless couplings of the Higgs boson with vector bosons (g_{HVV} and g_{HHVV}) and the vector boson masses:

$$g_{HVV}^{SM} = g_{HHVV}^{SM} = \frac{m_V^2}{v^2} \quad (2.3)$$

Therefore, investigating the interactions between Higgs bosons and vector bosons provides valuable insights into the nature of the Higgs mechanism and the electroweak symmetry breaking. While κ_V is already measured with high accuracy from single Higgs

boson measurements (see Section 1.3.3), only Higgs boson pair production via VBF is directly sensitive to κ_{2V} (as illustrated in the Feynman diagram in Figure 2.4b). The VBF HH cross section as a function of κ_{2V} is presented in Figure 2.7. The prediction for the VBF HH cross section as a function of κ_{2V} follows the same approach as that for κ_λ . It is based on a quadratic fit to VBF HH cross sections for various non-SM κ_{2V} values, calculated at leading order (LO) [68] and then rescaled using the SM $N^3\text{LO}$ -to-LO ratio. This method leverages the $N^3\text{LO}$ accuracy available for the SM calculation, while non-SM VBF HH cross section predictions remain available only at LO. The VBF HH




Figure 2.7 – Cross section for double Higgs production via vector boson fusion as a function of the Higgs boson quartic coupling modifier κ_{2V} . The prediction assumes pp collisions at $\sqrt{s} = 13$ TeV, the Higgs boson mass equal to 125 GeV, and all other couplings affecting HH production fixed to the SM expectations, including $\kappa_V = 1$. The SM ggF HH cross section is also shown as a reference. The uncertainty bands represent the theoretical uncertainty affecting the prediction (namely, the scale and the PDF+ α_S components)

cross section reaches a minimum when κ_{2V} is exactly 1, corresponding to the SM expectation. This occurs due to the cancellation between the two diagrams in Figure 2.4b and Figure 2.4c, caused by $\kappa_{2V} = \kappa_V^2$ in the SM, which is a direct consequence of the weak isospin doublet structure of the Higgs field in the SM Higgs mechanism [69]. As a result, even small deviations from unity in the κ_{2V}/κ_V^2 ratio lead to significant enhancements in the VBF HH cross section. In addition, anomalous κ_{2V}/κ_V^2 couplings violate perturbative unitarity in the VBF HH processes, inducing a quadratic dependence of the VBF HH cross section with the center of mass energy of the quarks involved in the hard scattering process (namely, $\sigma_{\text{VBF}}(\text{HH}) \propto \hat{s}$) [11]. Consequently, such scenarios favor VBF HH production with highly energetic and central Higgs bosons, populating the

invariant mass (m_{HH}) region above 1 TeV, as shown in Figure 2.8. Therefore, studying VBF HH production, especially in the high- m_{HH} regime, is crucial for testing the nature of the electroweak symmetry breaking mechanism and constraining κ_{2V} .




Figure 2.8 – Distribution of the Higgs boson pair invariant mass (m_{HH}) for Higgs boson pairs produced via VBF for the SM case and various (κ_{2V}, κ_V) values. The normalisation and the shape of the m_{HH} distribution are driven by the ratio κ_{2V}/κ_V^2 , rather than the individual values of the two couplings.

2.1.4 Searches for SM Higgs boson pair production at the LHC

Due to the rich variety of Higgs boson decay modes discussed in Section 1.3.2, a wide set of signatures can be exploited when searching for Higgs boson pair production. The branching fractions for all the possible di-Higgs final states are summarized in Figure 2.9.

Higgs boson pair production is an extremely rare process, with a SM cross section that is a thousand times smaller than that of single Higgs boson production (refer to Section 2.1.1 and Section 1.3.1 for the numerical values). Hence, the most sensitive final states for studying HH production and its properties involve at least one Higgs boson decaying into two bottom quarks, leveraging its large $\mathcal{BR}$ (equal to 58.24% for a SM 125-GeV Higgs boson). Among these, three *golden channels* have been identified as significantly more sensitive than others: $HH \rightarrow b\bar{b}b\bar{b}$, $HH \rightarrow b\bar{b}\gamma\gamma$, and $HH \rightarrow b\bar{b}\tau\tau$.

- The $HH \rightarrow b\bar{b}b\bar{b}$ process has a fully reconstructable final state and benefits from the largest SM branching fraction among all the HH channels (equal to 34%). However, it suffers from overwhelming background due to multijet production. This background is increasingly suppressed in kinematic regions with Lorentz-boosted Higgs bosons, resulting in highly energetic central b -jets in the final state. Additionally, targeting the VBF HH production mode, characterized by two extra forward jets, helps further suppress the multijet background. Especially when combining the boosted topology with the VBF HH signature, this channel provides the

	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%

Figure 2.9 – $\mathcal{BR}$ s of different final states for double Higgs decay processes.

strongest sensitivity to non-SM κ_{2V}/κ_V^2 scenarios, which indeed favor high-energy central Higgs bosons (as detailed in Section 2.1.3).

- The $HH \rightarrow b\bar{b}\gamma\gamma$ final state has the smallest branching ratio among all the golden channels (just 0.26%). However, it benefits from the excellent trigger and reconstruction efficiency for photons in both the ATLAS and CMS detectors, as well as the excellent resolution for the di-photon invariant mass ($m_{\gamma\gamma}$). Requiring two isolated photons in the final state provides a strong trigger strategy, allowing the threshold for the photon transverse energy to be set as low as $E_T = 25$ GeV. Thanks to this, the $HH \rightarrow b\bar{b}\gamma\gamma$ final state is particularly sensitive in kinematic regions with softer Higgs bosons and lower di-Higgs invariant mass (e.g., $m_{HH} < 400$ GeV). This makes it the primary channel for sensitivity to SM and non-SM κ_λ scenarios, which favor HH production with softer m_{HH} spectra (see Section 2.1.2). This thesis focuses on the search for Higgs boson pair production in the $b\bar{b}\gamma\gamma$ final state.
- Finally, the $HH \rightarrow b\bar{b}\tau\tau$ channel strikes a balance between branching ratio and background contamination, offering the best sensitivity to SM HH production. It has a higher branching ratio (7.3%) compared to the $b\bar{b}\gamma\gamma$ channel and relatively low background from other SM processes. However, the reconstruction and identification of τ leptons, whether through their hadronic or leptonic decays, present significant challenges compared to photon reconstruction.

Searches for Higgs boson pair production in various final states, including the three golden channels, were conducted by the ATLAS and CMS Collaborations using up to 140 fb^{-1} of data collected during the Run 2 of the LHC. These searches primarily targeted the two leading production modes: ggF HH and VBF HH. No evidence of HH production was observed by either experiment in Run 2 data. The results are generally presented as

95% confidence level (CL) upper limits [70] on the di-Higgs signal strength:

$$\mu_{HH} = \frac{\sigma(HH)}{\sigma^{SM}(HH)} \quad (2.4)$$

and constraints on the coupling modifiers influencing HH production, specifically κ_λ and κ_{2V} ¹.

The best sensitivity to HH production and its properties is achieved by harnessing the statistical power of all the individual HH searches in different final states, and extracting the combined results. A combination of HH searches was performed by both the ATLAS and CMS collaboration, relying on data collected during the Run 2 of the LHC. While the ATLAS HH combination will be discussed in Chapter 7 of this thesis, I will anticipate here the results obtained by CMS. The combination of HH analyses with Run 2 data performed by CMS yields an observed 95% CL upper limit on μ_{HH} of 3.4, where the expected upper limit evaluated under the hypothesis of no HH production is 2.5 [4]. The observed and expected upper limits on μ_{HH} from the individual analyses and from the combination are shown in Figure 2.10. Similarly, the CMS HH combination allows to set the tightest constraints on the coupling modifiers κ_λ and κ_{2V} using CMS Run 2 data. CMS excludes κ_λ values outside the interval of $-1.7 < \kappa_\lambda < 7.0$ at 2σ (corresponding to 95.4% CL), where the expectation in the SM assumption is $-2.3 < \kappa_\lambda < 8.0$ [71]. Finally, the combined HH searches with Run 2 data exclude non-SM κ_{2V} values outside the interval $0.67 < \kappa_{2V} < 1.48$ at 95% CL, thereby rejecting the $\kappa_{2V} = 0$ hypothesis with a significance of more than 5σ [4]. Figure 2.11 illustrates the combined constraints on κ_λ and κ_{2V} from Higgs boson pair production as determined by CMS using Run 2 data.

2.2 Beyond the Standard Model potential: new physics in Higgs boson pair production

2.2.1 Shortcomings of the Standard Model

The Standard Model is exceptionally successful in describing the fundamental constituents of matter and their interactions at the subatomic scale. The SM has been extensively tested through precision measurements [72] and searches for new physics phenomena, and it has been experimentally verified at an excellent level of accuracy. Despite its remarkable success, there remain unresolved questions, pointing towards the existence of a more general underlying theory describing the whole of particle physics phenomenology across all energy scales.

Perhaps the most pressing limitation of the SM is the so-called *dark matter problem*. Several astrophysical and cosmological observations, such as galaxy rotation curves [73], gravitational lensing [74], and measurements of the cosmic microwave background [75], have provided compelling evidence of a significant excess of non-luminous, non-fermionic matter over ordinary matter. This form of matter, typically referred to as *dark matter*, has

¹In addition to κ_λ and κ_{2V} , Higgs boson pair production is driven, at tree level, by the coupling modifiers κ_t and κ_V . These are, however, determined with much higher precision from single Higgs boson measurements (see Section 1.3.3).




Figure 2.10 – 95% CL upper limits on the signal strength μ_{HH} (corresponding to $\sigma(pp \rightarrow HH)/\sigma_{\text{Theory}}$) obtained by the CMS Collaboration searching for Higgs boson pairs in different final states and from their combinations [4]. The green and the yellow bands identify the 68% and 95% confidence intervals on the limit on μ_{HH} expected under the background only hypothesis.




Figure 2.11 – Negative log-likelihood contours [70] corresponding to 68% and 95% CL in the $(\kappa_\lambda, \kappa_{2V})$ plane.

never been directly measured. According to the Lambda-CDM model of cosmology [76–78] and measurements of the cosmic microwave background by Planck [75], dark matter constitutes 27% of the total energy in the Universe, while fermionic ordinary matter, described by the SM, accounts for less than 5%. The remaining 68% is attributed to *dark energy*, a constant energy density that permeates the Universe and is thought to explain its accelerating expansion. The SM does not include a viable candidate for dark matter (nor does it account for dark energy), making its nature one of the most significant unresolved questions in particle physics and cosmology.

Another clear shortcoming of the Standard Model is its inability to reconcile general relativity, the current macroscopic theory of gravity, with the principles of gauge quantum field theory. Consequently, it fails to offer a consistent description of gravitational interactions at the subatomic scale. The inability to incorporate gravity is not critical at TeV energy scales, because gravity is extremely weak compared to the strong and electroweak interactions, and its effect can be safely neglected when studying particle physics phenomenology at colliders. However, at higher energies, particularly at the *Planck scale* (around $M_P \sim 10^{19}$ GeV), gravitational quantum effects become significant. The several order of magnitude difference between the electroweak scale ($M_{EW} \sim 100$ GeV) and the Planck scale, and the corresponding difference between the strengths of the electroweak interaction and gravitational interaction is considered a hint of the *unnaturalness* of the SM. This unnaturalness effect also manifests when considering the quantum corrections to the Higgs boson mass, coming from the virtual effects of every SM massive particle. The resulting physical mass m_H can be written as:

$$m_H^2 = m_{H,\text{bare}}^2 - \mathcal{O}(\lambda, g^2, y^2)\Lambda^2 + \mathcal{O}(\ln(\Lambda^2)) \quad [79] \quad (2.5)$$

Here, $m_{H,\text{bare}}$ represents the tree-level Higgs boson mass (also known as the *bare* mass), while λ , g , and y denote the trilinear Higgs boson self-coupling, electroweak couplings, and Yukawa couplings, respectively. Λ is the scale at which the SM ceases to be valid, and new physics is expected to emerge. The bare mass of the Higgs boson can be fine-tuned to match the observed Higgs boson mass of approximately 125 GeV. However, if the SM remains valid up to the Planck scale ($\Lambda = M_P$), the cancellation between $m_{H,\text{bare}}$

and $\mathcal{O}(\lambda, g^2, y^2)\Lambda^2$ must span several orders of magnitude: $\frac{(m_H^2)^2}{(M_P^2)^2} = \frac{(10^2)^2}{(10^{18})^2}$. From a

theoretical standpoint, this extreme fine-tuning seems highly unnatural and is known as the *hierarchy problem* [80].

The *hierarchy problem* hints at the possibility of a more natural theory, where the fine-tuned cancellation between the bare Higgs boson mass and the quadratic divergences from quantum corrections stems naturally from a first principle.

Over the past decades, numerous models have been proposed to extend the SM and address these unresolved points. One of the most popular and theoretically appealing beyond the SM (BSM) extensions is *supersymmetry* (SUSY) [81]. SUSY, at a minimum, doubles the particle content by introducing fermionic and bosonic *superpartners* for each SM particle, differing by half a unit of spin and known as *sparticles*. The presence of sparticles allows for the exact cancellation of quadratic divergences in the quantum corrections to the Higgs boson mass, thus elegantly solving the hierarchy problem.

A self-consistent SUSY theory that adheres to the gauge principle requires at least two Higgs doublets. The SUSY theory with exactly two Higgs doublets and the minimum number of additional sparticles is called the Minimal Supersymmetric Model (MSSM). The MSSM introduces an additional discrete symmetry, referred to as *R-parity* conservation, where the *R-parity* of a particle is defined as:

$$P_R = (-1)^{3(B-L)+2s} \quad (2.6)$$

In Equation 2.6, B is the baryon number, L is the lepton number, and s is the spin of the particle. All SM particles have a positive *R-parity* ($P_R = +1$), while their superpartners have negative *R-parity* ($P_R = -1$). The conservation of *R-parity* forbids the decay of the proton (in agreement with experimental data [82]). Furthermore, this symmetry prevents sparticles from decaying exclusively into SM particles and vice versa. The lightest supersymmetric particle (LSP), which should be neutral, cannot decay into other sparticles (being the lightest) or to SM particles due to *R-parity* conservation, making it stable and a strong candidate for dark matter. Therefore, supersymmetric theories, even in their minimal forms such as the MSSM, offer elegant solutions to several persistent issues within the SM. These features have made SUSY one of the most extensively studied extensions of the SM, drawing significant theoretical and experimental interest for decades. Both the ATLAS and CMS Collaborations at the LHC have conducted comprehensive searches for supersymmetric particles throughout Run 1, Run 2, and the ongoing Run 3 data-taking periods [83]. These searches exploit the rich phenomenology of SUSY, targeting various potential signals, including the production of *gluinos*, *squarks*, and *neutralinos*, among other sparticles. Despite this monumental experimental effort, no direct

evidence for supersymmetric particles has been observed so far. As a result, stringent exclusion limits have been placed on the allowed parameter space of SUSY models, significantly constraining the masses and properties of sparticles. For example, Figure 2.12 shows combined exclusion limits derived from Run 1 and Run 2 data, constraining the masses of two sparticles: the *gluino* ($\tilde{g}$), the superpartner of the SM gluon, and the lightest *neutralino* ($\tilde{\chi}_1^0$), a fermionic state resulting from the mixing of the superpartners of the SM weak gauge bosons and the Higgs boson - the *winos*, *bino*, and *higgsino*. These exclusion limits underscore the increasing challenge of detecting SUSY, as current experimental bounds continue to rule out a growing fraction of the parameter space within minimal SUSY model assumptions.

In recent years, growing interest has focused on more advanced supersymmetric models, such as the Next-to-Minimal Supersymmetric Standard Model (NMSSM), as well as broader extensions of the SM that feature an expanded Higgs sector with multiple scalar particles. Among these predicted particles, only the 125-GeV Higgs boson has been observed so far, while others remain undiscovered. The NMSSM, for instance, is a Next-to-Minimal Two-Higgs Doublet Model (N2HDM), which extends the SM scalar sector by introducing an additional complex doublet and a real scalar singlet [16]. Simpler alternatives are offered by the Two Real Scalar Singlet Model (TRSM) [17] or the complex scalar singlet extension of the scalar sector (cxSM) [84], where the SM is enhanced by adding two real scalar singlets or a complex scalar singlet respectively. A major strength of these models is their ability to naturally offer viable dark matter candidates [85, 86].

These models are particularly relevant for this thesis, as they allow for the asymmetric production of multiple Higgs bosons, including the $X \rightarrow SH$ phenomenology, where H represents the SM 125-GeV Higgs boson, and X and S are additional scalar particles. Indeed, Chapter 8 presents a search for a new heavy scalar particle X decaying into the SM Higgs boson and another lighter scalar particle S in the final state with two photons and two bottom quarks, using data collected by the ATLAS detector during the Run 2 of the LHC.

One of the simplest extensions of the SM scalar sector that accommodates the $X \rightarrow SH$ phenomenology is the TRSM [17]. This model introduces only two additional real singlets to the SM scalar sector, altering the SM scalar potential $V^{SM}(\Phi)$ by adding new terms involving the two extra degrees of freedom, ϕ_1 and ϕ_2 :

$$V(\phi_1, \phi_2, \Phi) = V^{SM}(\Phi) + V_{\text{singlets}}(\phi_1, \phi_2, \Phi) \quad (2.7)$$




Figure 2.12 – Exclusion limits in the $(m(\tilde{g}), m(\tilde{\chi}_1^0))$ plane, based on 13 TeV pp collision data recorded by the ATLAS detector, are shown for various simplified models where the gluino decays to the lightest supersymmetric particle (either the lightest neutralino or gravitino). These decays can occur directly or via a cascade involving other SUSY particles with intermediate masses. Each line represents a different gluino decay mode, as indicated in the legend, with the assumption that the decay occurs with a 100% branching fraction. Some limits are further influenced by assumptions regarding the masses of intermediate states, as detailed in the references cited in the plot [83].

The simplest² renormalizable, gauge-invariant scalar potential involving the new singlet fields of the TRSM model is:

$$\begin{aligned}
 V_{\text{singlets}}(\phi_1, \phi_2, \Phi) = & \mu_1^2 \phi_1^2 + \lambda_1 \phi_1^4 \\
 & + \mu_2^2 \phi_2^2 + \lambda_2 \phi_2^4 \\
 & + \lambda_{\Phi 1} \phi_1^2 (\Phi^\dagger \Phi) + \lambda_{\Phi 2} \phi_2^2 (\Phi^\dagger \Phi) + \lambda_{12} \phi_1^2 \phi_2^2
 \end{aligned} \tag{2.8}$$

where all the coefficients are real.

Since the new fields are singlets, they do not interact with gauge bosons, and their kinetic term is straightforward: $\mathcal{L}_{\text{kin}}^{\text{singlets}} = \sum_{i=1}^2 \partial^\mu \phi_i \partial_\mu \phi_i$. Moreover, it is impossible to construct gauge-invariant and renormalizable interactions between a scalar singlet and any of the SM fermions. Consequently, the singlets interact only with the SM complex scalar Φ through the couplings in the scalar potential described in Equation 2.8. If these singlets acquire a non-zero vacuum expectation value (*v.e.v.*), their physical states will mix with the SM Higgs field, inheriting some of its gauge and Yukawa couplings. Such symmetry-breaking scenarios are particularly intriguing because they produce three neutral scalar states, typically labeled H , X , and S , where H indicates the SM 125-GeV Higgs boson. Depending on the mixing parameters connecting the scalar fields with the physical states, the new particles S and X can inherit Higgs-like couplings to SM particles, allowing decays to fermion pairs, for instance. However, in TRSM-like scenarios, the leading interactions of these new scalar particles are with the 125-GeV Higgs boson, while their couplings with SM fermions or gauge bosons are subdominant. This makes (asymmetric) *Higgs boson pair production* (such as $X \rightarrow SH$) one of the most interesting experimental signature to probe when searching for new physics in the scalar sector.

A detailed review of the TRSM model and the resulting phenomenology at the LHC can be found in [17].

It is important to note that the phenomenology discussed here can also generically occur in other BSM models that feature three or more Higgs-like physical states. However, the couplings between the scalar particles and with SM particles may differ from those in the TRSM. Therefore, it is crucial to design searches that provide model-independent constraints (or measurements) of the new signal rate.

2.2.2 Asymmetric Higgs boson pair production at the LHC

The ATLAS and CMS experiments at the LHC cover a broad range of experimental searches for new physics in the scalar sector. However, they have primarily focused on symmetric Higgs boson pair production. This includes searches for processes such as $X \rightarrow HH$ [87, 88], where H is the SM 125 GeV Higgs boson and X is a new heavy scalar particle, and $H \rightarrow aa$ [89, 90], where a is a new light (pseudo-)scalar. Asymmetric Higgs boson pair production (namely, $X \rightarrow SH$) has only recently been explored as a viable

²To simplify the extended scalar sector, the TRSM model introduces two additional discrete symmetries, $\mathbb{Z}_2^1$ and $\mathbb{Z}_2^2$. These symmetries leave the SM unchanged, while $\mathbb{Z}_2^{1(2)}$ inverts the sign of $\phi_{1(2)}$ and acts trivially on $\phi_{2(1)}$.

signature, paving the way for more complex high-dimensional search spaces, involving the mass spectra of the two scalars X and S .

The production of a heavy scalar particle X could occur at hadron colliders via a gluon-gluon fusion mechanism, as illustrated by the Feynman diagrams in Figure 2.13.




Figure 2.13 – Feynman diagram for the production of a heavy scalar particle X decaying into a SM Higgs boson H and a lighter scalar S via gluon-gluon fusion.

Searches for asymmetric Higgs boson pair production at the LHC were conducted by both the ATLAS and CMS experiments using the Run 2 dataset. These searches considered various final states resulting from the decays of the SM Higgs boson H and the new scalar S to SM particles. The considered decay channels include $b\bar{b}\gamma\gamma$ [91, 92], $VV\tau\tau$ [93], $b\bar{b}\tau\tau$ [94], $\gamma\gamma\tau\tau$ [95], and $b\bar{b}b\bar{b}$ [96]. The decays of the SM Higgs boson (or the new particle S) to photons or τ leptons provide clear signatures for triggering and isolating the signal from the SM background, while the SM Higgs boson decay to bottom quarks exploits the highest $\mathcal{BR}$.

The decays of the particle S , assuming they are SM Higgs-like, depend on its mass: for $m_S < 130$ GeV, the dominant decay is $S \rightarrow b\bar{b}$, while for higher m_S , the decay to vector bosons (either $W^\pm$ or Z) becomes the primary mode.

Figure 2.14 illustrates the parameter space explored by the $X \rightarrow SH$ searches, defined by the masses of the two new scalars X and S . The searches complement each other by enhancing sensitivity across different (m_X, m_S) regions. Hence, the asymmetric Higgs boson pairs search program at the LHC covers a wide range of the masses of the new scalars, spanning up to two orders of magnitude. Specifically, the parameter space includes m_X values from 250 GeV to 4 TeV and m_S values from 15 GeV to 2 TeV.

Despite some emerging small excesses, no significant deviation from the SM background has been observed in any of these searches. Thus, no clear evidence of new physics in the scalar sector has been detected so far.




Figure 2.14 – Parameter space probed by different $X \rightarrow SH$ searches in various decay channels, performed by both ATLAS and CMS Collaborations [91].

The Large Hadron Collider

The Large Hadron Collider (LHC) [8] is the largest and highest energy particle accelerator in the world. The LHC accelerates and collides protons (pp) or heavy ions, in a 27 km-long tunnel located 100 m below the ground near CERN (Geneva). The LHC was designed to achieve a center-of-mass energy of $\sqrt{s} = 14$ TeV for pp collisions, with an instantaneous luminosity of $\mathcal{L} = 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. This thesis utilizes data from pp collisions delivered by the LHC during its second data-taking period - Run 2 - when the LHC operated at a center-of-mass energy of $\sqrt{s} = 13$ TeV and provided an integrated luminosity of $L = 140 \text{ fb}^{-1}$ over the course of four years (from 2015 to 2018).

3.1 Proton-proton interactions at hadron colliders

The LHC was conceived as a hadron accelerator and collider, in order to overcome the main limitation of accelerating electrons in a circular trajectory, namely the large energy losses via *synchrotron radiation emission*.

Charged particles emit synchrotron radiation when accelerated along a curved trajectory, losing energy according to the following relationship:

$$\frac{dE}{dt} \propto \frac{E^4}{m^2} \text{ [97]}, \quad (3.1)$$

where E and m represent the energy and mass of the particles, respectively. Because of the $\propto m^{-2}$ dependence, the radiative energy losses for an electron beam are much higher compared to proton beams with identical properties (i.e., $m_e^{-2}/m_p^{-2} \sim 10^9$). Consequently, achieving high center-of-mass energies for electron-positron (e^+e^-) collisions is far more challenging.

However, a hadron collider faces a different set of challenges. Protons are bound states of quarks and gluons (collectively known as partons), and their collisions can be categorized into two types:

- **Soft collisions** are long-distance, low momentum-transfer interactions between two protons within the two colliding proton beams. The particles in the final state generally have low transverse momentum and scatter at small angles. These events are referred to as *minimum bias*.

- **Hard collisions** involve large momentum-transfer interactions between two partons within the incoming protons. These interactions can produce final state particles with high transverse momentum and large mass. Such events are rare compared to soft collisions.

The latter are the most interesting from a physics perspective. The composite structure of protons introduces additional complexities when studying such interactions. In a hard collision, any two partons within the incoming protons can interact, each carrying only a fraction of the proton's momentum. Consequently, the actual center-of-mass energy for the hard scattering between the two partons (typically referred to as $\sqrt{\hat{s}}$) is lower than the nominal center-of-mass energy of the pp collision and, most importantly, it is not exactly known. This quantity can be expressed as:

$$\sqrt{\hat{s}} = \sqrt{x_a x_b s}, \quad (3.2)$$

where x_a and x_b are the momentum fractions of the partons within the two protons, and s is the center-of-mass energy of the pp collision. Therefore, the cross section for a hard scattering interaction $pp \rightarrow X$ is given by:

$$\sigma_{pp \rightarrow X} = \sum_{a,b} \int dx_a dx_b f_a(x_a, Q^2) f_b(x_b, Q^2) \cdot \hat{\sigma}_{ab \rightarrow X}(x_a, x_b), \quad (3.3)$$

where $\hat{\sigma}_{ab \rightarrow X}(x_a, x_b)$ is the *parton-level cross section*, and $f_{a(b)}(x_{a(b)}, Q^2)$ represents the probability of finding a parton carrying a fraction $x_{a(b)}$ of the proton's momentum at a given momentum transfer scale Q^2 . The integral spans all possible momentum fractions $x_{a(b)} \in [0, 1]$, and the sum covers all possible parton pairs a and b within the two protons. While the parton-level cross section $\hat{\sigma}_{ab \rightarrow X}$ can be calculated to any desired accuracy using perturbative QCD expansions, the *parton distribution functions* $f_{a(b)}(x_{a(b)}, Q^2)$ (typically referred to as PDFs) are influenced by the non-perturbative strong interactions of the partons within the protons. Their dependence on the momentum fractions $x_{a(b)}$ must be extracted from experimental data. Fortunately, once their behavior at a certain reference momentum transfer scale Q^2 is known, their evolution at all scales can be determined using the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) evolution equations [98]. Figure 3.1 shows the PDFs for the quarks and the gluons within a proton, evaluated for momentum transfer scales $\mu^2 = 10 \text{ GeV}^2$ and $\mu^2 = 10^4 \text{ GeV}^2$.

In summary, one of the main drawbacks of hadron colliders is the inability to directly measure the cross section of the underlying partonic interaction, happening at an unknown center-of-mass energy $\hat{s}$. Instead, only the hadronic cross section is experimentally accessible. While the partonic cross section can be calculated using perturbative QCD, ideally to any desired order of accuracy, computing the hadronic cross section requires incorporating the PDFs, which introduces additional uncertainties into the theoretical predictions.

Additionally, the partons within the colliding protons that are not involved in the hard scattering undergo hadronization or secondary multi-parton interactions, creating a dense environment and a substantial QCD background (known as the *underlying event*), which

pollute the event, thus making it more challenging to isolate the interesting scattering process.

Furthermore, hard and soft collisions occur simultaneously within the same proton bunch-crossing, resulting in background noise from minimum bias events.

The detectors positioned at the collision points along the LHC ring must extract a signal as clean as possible to distinguish the interesting processes from background events.

Soft collisions do not always constitute a background for hard scattering events; they too can be recorded and studied to gain a better understanding of QCD processes in hadron colliders.




Figure 3.1 – Parton distribution functions (PDFs) multiplied by the momentum fraction x , evaluated at $\mu^2 = 10 \text{ GeV}^2$ (left) and $\mu^2 = 10^4 \text{ GeV}^2$ (right) for quarks and gluons within a proton. The PDFs are extracted from the NNPDF3.1 NNLO set of PDFs [99].

3.2 Operation of the LHC and the accelerator complex

The pp (or heavy ion) interactions happen in four collision points along the LHC ring, where the four main experiments are placed: ATLAS, CMS, LHCb, and ALICE. LHCb [100] and ALICE (A Large Ion Collider Experiments) [101] are optimized for specific purposes: LHCb focuses on bottom quark physics and performing precision measurements targeting CP-violating processes, while ALICE is dedicated to study heavy-ion collisions and characterizing quark-gluon plasma. ATLAS (A Toroidal LHC Apparatus) [9] and CMS (Compact Muon Solenoid) [102] are general purpose experiments: they are designed to detect the broadest possible range of signals, whether they are generated from Standard

Model processes or are products of new physics.

The two proton beams circulating in the LHC travel in opposite directions within two adjacent beamlines. They are kept on their circular path by 1232 niobium-titanium superconducting dipole magnets (Figure 3.2a), which produce a magnetic field of up to 8.4 T and operate at a temperature of 1.9 K. Beam focusing in the horizontal and vertical directions is achieved using 392 superconducting quadrupole magnets, whose cross section is shown in Figure 3.2b. Higher-order magnets (sextupoles, octupoles, and decapoles) are employed to maintain beam stability by correcting minor non-uniformities in the magnetic field geometry.




Figure 3.2 – Cross section of the LHC ring, with a dipole magnet (left) [103] and a quadrupole magnet (right) [104]. The two proton beams travel within the two adjacent beampipes that can be noticed at the center of the ring’s cross sections.

When proton beams enter the detectors along the LHC ring, they are squeezed together to maximize the probability of collisions between the two incoming beams. This beam squeezing is achieved by a set of three quadrupole magnets, collectively known as the *inner triplet*. There are eight inner triplets positioned around the LHC ring, with two located at each of the four main experiments (ATLAS, CMS, LHCb, and ALICE).

Each proton beam is accelerated by eight superconducting radio-frequency (RF) cavities. These cavities provide a combined maximum accelerating voltage of 16 MV for each proton beam (2 MV per cavity) and oscillate at a frequency of 400 MHz. Depending on their exact energy, protons receive slightly different accelerating kicks. A proton that reaches the nominal energy of $\sqrt{s}/2$ at the ideal arrival time will not be accelerated when passing through the cavities. In contrast, protons with lower or higher energy arrive later or earlier than the ideal time and are thus accelerated or decelerated, bringing their energy closer to the nominal value. This mechanism helps to sort the proton beams into well separated *bunches*, each containing approximately 10^{11} protons.

The protons injected into the LHC start with an energy of 450 GeV. The RF cavities then ramp up the beam energy to $\sqrt{s}/2$ in about 30 minutes. These 450 GeV protons are delivered by a series of accelerators. The process begins with protons being extracted from a

hydrogen gas bottle by an electric field, which strips hydrogen atoms off their electrons. The protons are then accelerated to 50 MeV by a linear accelerator (LINAC 4). Next, the protons are injected into the Proton Synchrotron Booster (PSB) and subsequently into the Proton Synchrotron (PS), where their energy is increased to 1.4 GeV and 25 GeV, respectively. The final step in the accelerator chain is the Super Proton Synchrotron (SPS), which accelerates the proton beam to 450 GeV and delivers it to the LHC. During the Run 2 data-taking period, the maximum energy reached by the protons circulating in the LHC was 6.5 TeV, corresponding to a center-of-mass energy of $\sqrt{s} = 13$ TeV for pp collisions. In Run 3 (ongoing since 2022), this energy increased to 6.8 TeV, yielding a center-of-mass energy of $\sqrt{s} = 13.6$ TeV.

The entire CERN accelerator complex is illustrated in Figure 3.3.




Figure 3.3 – The CERN accelerator complex [105]. Protons are extracted from a hydrogen gas bottle and accelerated through a series of accelerators before being injected into the LHC.

3.3 Luminosity and pileup

In a particle collider, the rate of occurrence of a particular process is related to its cross section σ by the equation

$$\frac{dN}{dt} = \mathcal{L} \cdot \sigma \quad (3.4)$$

$\mathcal{L}$ is the *instantaneous luminosity*, and can be written as a function of the collider beam parameters as follows:

$$\mathcal{L} = \frac{n_{\text{beam}} N_{\text{bunch}}^2 f_{\text{rev}} \gamma F}{4\pi\epsilon_n \beta^*} \quad (3.5)$$

In Equation 3.5:

- n_{beam} is the number of proton bunches in a proton beam;
- N_{bunch} is the number of protons in a bunch;
- f_{rev} is the revolution frequency;
- γ is the relativistic Lorentz factor for the protons;
- ϵ_n is the normalized transverse beam emittance, representing the spread of the protons' positions and momenta in the two directions perpendicular to the beam trajectory.
- β^* represents the beta function (connecting the beam emittance with the geometrical size of the beam along its trajectory), evaluated at the collision point;
- F is the geometric luminosity reduction factor, which accounts for the decrease in luminosity when the two proton bunches do not collide head-on (i.e., when the crossing angle differs from zero). This factor is always less than 1.

Table 3.1 presents the LHC beam parameters during the four data-taking years in Run 2, compared with the design values. Figure 3.4 shows the peak instantaneous luminosity

Parameter	Design	2015	2016	2017	2018
Energy [TeV]	7.0	6.5	6.5	6.5	6.5
Revolution frequency f_{rev} [kHz]	11.2	11.2	11.2	11.2	11.2
Number of bunches n_{beam}	2808	2244	2220	2556 - 1868	2556
Bunch population, N_{bunch} [10^{11} p]	1.15	1.2	1.25	1.25	1.1
β^* [cm]	55	80	40	40 $\rightarrow$ 30	30 $\rightarrow$ 27 $\rightarrow$ 25
Normalized emittance ϵ_n [$\mu\text{m rad}$]	3.75	2.6 - 3.5	1.8 - 2	1.8 - 2.2	1.8 - 2.2
Crossing angle θ_c [μrad]	285	370	370 $\rightarrow$ 280	300 $\rightarrow$ 240	320 $\rightarrow$ 260
RMS bunch length σ_z [cm]	7.55	9	9	8	8
Relativistic γ for the proton	7462	6929	6929	6929	6929
Geometric luminosity loss F [%]	84	84	65	72	61
Peak luminosity [$10^{34} \text{ cm}^{-2} \text{ s}^{-1}$]	1.0	< 0.6	1.5	2.0	2.1

Table 3.1 – Summary of beam and machine parameters during the four years of Run 2, compared to the design values [106–108].

delivered by the LHC during stable beams the four data-taking years of Run 2. The peak value exceeds twice the nominal luminosity of $\mathcal{L} = 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ in both 2017 and 2018 runs (see Figure 3.4c and Figure 3.4d). During Run 3, the LHC has consistently operated at more than twice its nominal luminosity, reaching a peak of $\mathcal{L} = 2.5 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ in 2022 [110].

The luminosity integrated over time, known as *integrated luminosity* (L), measures the size of the collected dataset and relates the number of events observed for a given process to its cross section ($N = L \cdot \sigma$). Figure 3.5 displays the integrated luminosity delivered by the LHC and recorded by ATLAS over time during the four years of data-taking




Figure 3.4 – Peak instantaneous luminosity delivered by the LHC during stable beams for each fill in the four data-taking years of Run 2: 2015 (a), 2016 (b), 2017 (c), and 2018 (d) [109].

in Run 2. The total integrated luminosity recorded by the ATLAS experiment in the full Run 2, certified to include good quality data, amounts to $140.1 \pm 1.2 \text{ fb}^{-1}$ [49].




Figure 3.5 – Total integrated luminosity delivered by the LHC to the ATLAS experiment as a function of time [109], during the full Run 2 of the LHC. The recorded dataset (yellow) are superimposed to the total delivered luminosity (green).

In particle colliders, the instantaneous luminosity is not constant, but, for each fill, it decreases over time due to the degradation of both beam intensity (i.e. less number of protons in the bunches) and the emittance (more spread out momenta and positions of the protons within the bunches). The available time for stable beam data-taking is optimized by accounting for this luminosity decrease and the duration required for the accelerator chain to deliver a new beam. The primary factor contributing to the luminosity decay is the pp collisions occurring at the LHC interaction points. The relationship between luminosity and time in Run 2 is described by the following equation:

$$\mathcal{L}(t) = \mathcal{L}_0 e^{-t/\tau} \quad \text{with} \quad \tau = \frac{N_{\text{bunch},0}}{k\sigma_{pp}\mathcal{L}_0} \quad (3.6)$$

where $\mathcal{L}_0$ is the initial luminosity at the beginning of the fill, $N_{\text{bunch},0}$ is the initial number of protons in a bunch, σ_{pp} is the inelastic pp cross section, and k is a normalization factor that accounts for the number of pp interaction points along the beam ($k = 4$ for the LHC).

The LHC delivers such a high instantaneous luminosity that multiple pp interactions typically occur simultaneously at each bunch crossing within the interaction points. This is quantified by the average number of inelastic interactions per bunch crossing (μ), known as *pileup*. The pileup events, which primarily result from soft collisions, overlap with the hard interactions of interest and are generally treated as background in physics analyses. Pileup negatively impacts the reconstruction of the event of interest, degrading the energy resolution of physics objects such as jets, leptons, photons, and missing

transverse energy (E_T^{miss}). Additionally, it complicates the reconstruction of tracks and vertices by creating a denser environment within the particle detector's tracker. The pileup is directly proportional to the instantaneous luminosity:

$$\mu = \frac{\mathcal{L}}{\sigma_{pp}} \cdot \frac{1}{n_{\text{beam}}} \cdot \frac{1}{f_{\text{rev}}} \quad (3.7)$$

where σ_{pp} is the cross section for pp inelastic collisions. As shown in Figure 3.6a, the average number of interactions per bunch crossing during Run 2 was 33.7, and increased year by year, consistently with the increase of instantaneous luminosity delivered by the LHC. Hence, it is important to find a good compromise between pushing for higher (integrated) luminosities that correspond to higher data statistics, and worsening of pileup conditions that can degrade the quality of the data. In Run 3, the β^* luminosity lev-




Figure 3.6 – Luminosity-weighted distribution of the average number of interactions per bunch crossing (μ) in pp collisions recorded by ATLAS in the four data-taking years of Run 2 [109] (left) and in the first two data-taking years of Run 3 [110] (right).

elling technique [111] was adopted by the LHC to manage the pileup at the pp interaction points for the ATLAS and CMS experiments. This technique involves adjusting the beam squeezing to regulate the peak instantaneous luminosity, maintaining it at a constant level for several hours, delaying the onset of exponential decay described by Equation 3.6. By adjusting the instantaneous luminosity via the β^* parameter, as shown in Figure 3.7, the LHC ensures a stable target pileup level for longer during each fill, thus optimizing the data-taking conditions for the experiments and maximizing the integrated luminosity. Figure 3.6b illustrates the average number of interactions per bunch crossing during Run 3, which is significantly higher than in Run 2, reaching an average of 47.7.




Figure 3.7 – Pileup and β^* evolution at the ATLAS interaction point with β^* levelling during an LHC fill in 2022. The target pileup was $\mu = 54 \pm 2.5$. The final β^* value, marking the onset of the exponential decay of the instantaneous luminosity, was $\beta^* = 30$ cm and was reached after 5.2 hours of levelling.

The ATLAS experiment

ATLAS (A Toroidal LHC Apparatus) [9, 112] is one of the two general purpose experiments at the LHC, designed to explore the broadest possible range of physics phenomena arising from pp collisions, taking full advantage of the unprecedented energy levels achieved at the LHC. ATLAS is the largest particle detector ever built for a collider. It has a cylindrical structure, with diameter of 25 meters, a length of 46 meters, and a weight of approximately 7,000 tonnes. It is located in a cavern 100 meters below ground at one of the four collision points along the LHC ring (Interaction Point 1).




Figure 4.1 – A schematic view of the ATLAS detector.

The ATLAS detector consists of three major components, as illustrated in Figure 4.1: the **Inner Detector**, the **Calorimetric System**, and the **Muon Spectrometer**. These sub-detectors are arranged concentrically around the interaction point, forming an onion-like

structure. As particles produced by pp collisions move outward from the interaction point, they encounter each of these layers. The three sub-detectors complement one another: the Inner Detector tracks charged particles, the calorimeters measure the energy of photons, electrons, and hadrons, and the Muon Spectrometer handles tracking and momentum measurements for highly penetrating muons. Each sub-detector is designed to detect specific types of particles and measure certain properties, while the complete reconstruction of an event depends on the combined data from all three. Additionally, two distinct **magnet systems** generate the magnetic fields required to determine the momentum of charged particles by bending their trajectories: a 2-T solenoid magnet encircles the Inner Detector, and an 8-coil toroidal magnet system serves a similar purpose for the Muon Spectrometer.

A particle's position within the ATLAS detector is described by the coordinates (x, y, z) , where the origin is set at the collision point. The z axis runs along the beamline, while the (x, y) plane, known as the transverse plane, is perpendicular to the proton beams. The x axis points toward the center of the LHC ring, and the y axis points upwards.

Several key quantities are defined in the transverse plane, such as p_T (transverse momentum) and E_T^{miss} (missing transverse energy). The particle's position in this plane is often expressed using the modified polar coordinates, termed η and ϕ . Here, ϕ is the azimuthal angle, measured around the beam direction relative to the x axis, while η , known as *pseudorapidity*, is defined by the relation:

$$\eta = -\ln \left(\tan \frac{\theta}{2} \right), \quad (4.1)$$

where θ is the polar angle, measured from the z axis. Positive and negative values of η correspond to the positive and negative directions along the z axis, with $\eta = 0$ representing the transverse plane at the interaction point, as illustrated in Figure 4.2. Particles with large absolute values of η are referred to as *forward*, while those with smaller values are termed *central*.

The η coordinate can also be expressed in terms of the particle's momentum, $\mathbf{p}$, and its component along the beam axis, p_z :

$$\eta = \frac{1}{2} \ln \left(\frac{|\mathbf{p}| + p_z}{|\mathbf{p}| - p_z} \right) \quad (4.2)$$

The advantage of using pseudorapidity over the standard polar angle is that the difference in pseudorapidity between two points remains invariant under Lorentz boosts along the z axis. This feature is particularly useful when analyzing events from pp collisions, as the hard interactions between partons can be asymmetric and create boosted final states.

The angular distance ΔR between two points in the transverse plane is given by:

$$\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2} \quad (4.3)$$




Figure 4.2 – Polar coordinates used to identify particles' position within the ATLAS detector (left), and relation between the pseudorapidity η and the standard polar angle θ (right).

4.1 The Inner Detector

The ATLAS Inner Detector (ID) [113, 114], is the closest sub-detector to the interaction points. It consists of a barrel and two end-cap regions, extending its coverage up to $|\eta| < 2.5$. Its primary functions are providing high-precision tracking and high-resolution momentum measurements for charged particles, and reconstructing both primary and secondary vertices. To achieve these goals, the ID is designed with minimal material to reduce multiple scattering and energy loss, as well as to lower the probability of photon conversions. Additionally, it must have a strong radiation resistance to sustain the high radiation levels resulting from its proximity to the interaction point. The combined performance of the high-granularity pixel and silicon strip detectors (located closer to the beamline where track density is highest) and the Transition Radiation Tracker at larger radii ensures that these requirements are met. The complete layout of the ID, along with its three main components—the **Pixel detector**, the **Semi-Conductor Tracker (SCT)**, and the **Transition Radiation Tracker (TRT)**—is illustrated in Figure 4.3.

4.1.1 The Pixel detector

The Pixel Detector [116, 117] is the first system encountered by particles emerging from the pp interaction. In the barrel region, the detector includes 1736 silicon pixel modules arranged in four concentric layers at radii of 3.2 cm, 5.1 cm, 8.9 cm, and 12.3 cm from the beamline. The first of these layers, known as the **Insertable B-Layer (IBL)** [118], was added for the Run 2 data-taking to improve the tracking performance by reducing the distance between the interaction point and the first detector layer. As illustrated in Figure 4.4, the addition of the IBL led to a significant improvement in the resolution of




Figure 4.3 – Schematic layout of the ATLAS Inner Detector [115].

both the longitudinal and transverse impact parameters of the tracks, with resolutions improving by up to a factor two from Run 1 to Run 2. Each of the two end-caps contains 288 modules, organized into three disks. The segmentation of the pixel modules develops both in the R - ϕ plane and along the z axis, with a minimum pixel size of $50 \times 250 \mu\text{m}^2$ in the IBL and $50 \times 400 \mu\text{m}^2$ in the outer layers (with the longer side along the z coordinate). The detector achieves an intrinsic spatial resolution of $8 \mu\text{m}$ in the R - ϕ plane within the IBL, slightly degrading to $10 \mu\text{m}$ across both the outer barrel and end-cap regions. Along the z axis, the resolution is $40 \mu\text{m}$ in the IBL, degrading to $115 \mu\text{m}$ in the z (R) direction for the outer barrel (end-cap) layers. Typically, the Pixel Detector provides between three and four measurement points (*hits*) for track reconstruction.

4.1.2 The Semiconductor Tracker

Following the Pixel Detector, the Semiconductor Tracker (SCT) [120] provides between four and nine additional hit measurements for each track at medium radii (ranging from 30 cm to 51 cm) within the ID volume. The SCT is composed of eight concentric barrel layers of silicon microstrips organized into four cylinders, while the end-cap regions are composed of nine disks each, covering a pseudorapidity range of $|\eta| < 2.5$. In total, the SCT contains 4088 modules. The barrel modules are rectangular, measuring 64.0×63.6 mm with a strip pitch of $80 \mu\text{m}$, while the end-cap modules are trapezoidal, with strip pitches varying from $56.9 \mu\text{m}$ to $94.2 \mu\text{m}$. In the barrel, the strips run parallel to the z axis, while in the end-caps they run radially. To extract z coordinate information in the barrel region, a 40 mrad stereo angle is set between the two layers in each cylinder. A similar strategy is employed in the end-caps to provide measurements along the R direction. The silicon microstrips in the barrel (end-cap) regions offer single-hit spatial resolutions of $17 \mu\text{m}$ in the R - ϕ plane and $580 \mu\text{m}$ along the z (R) coordinate.




Figure 4.4 – Resolution of the transverse impact parameter (left) and longitudinal impact parameter (right) measured with 2015 (Run 2) data using the Inner Detector with the IBL, as a function of p_T for $0.0 < |\eta| < 0.2$. These results are compared to measurements taken with 2012 (Run 1) data [119].

4.1.3 The Transition Radiation Tracker

The Transition Radiation Tracker (TRT) [121] provides a large number of hits per track, typically around 36, although with lower spatial resolution ($130 \mu\text{m}$ in the R - ϕ plane). The TRT functions both as a drift chamber and as a transition radiation detector, thus allowing it to contribute to particle identification. The barrel region consists of 50,000 drift tubes, each 144 cm long, aligned parallel to the beamline, while an additional 250,000 shorter straws, each 39 cm long, are arranged radially in the end-caps. When a particle passes through a straw, the Xenon-Argon gas mixture inside becomes ionized. A fixed voltage causes the resulting ions and electrons to move towards the straw's outer wall and the central anode wire, respectively, producing a detectable signal. The position of the hit is determined by converting the electron drift time to the wire into a distance measurement. Finally, in addition to tracking, the TRT provides particle identification information in order to discriminate between electrons and charged pions. To generate transition radiation (TR), the straw tubes are interleaved with polymer fibers in the barrel and foils in the end-caps. When a particle crosses the boundary between different materials, it can emit TR. The intensity of the TR photons is influenced by the particle's relativistic Lorentz factor $\gamma = E/m$. This allows effective discrimination between electrons and pions of similar energies (within a range between 1 and 200 GeV), exploiting their mass difference ($m_e \approx 0.51 \text{ MeV}$ while $m_{\pi^\pm} \approx 139.57 \text{ MeV}$).

The overall ID momentum resolution (achieved in Run 1 so before the IBL insertion), was measured to be:

$$\frac{\sigma(p)}{p} = (4.83 \pm 0.16) \cdot 10^{-4} \text{ GeV}^{-1} \times p_T \quad [114] \quad (4.4)$$




Figure 4.5 – Track reconstruction efficiency of the as a function of the mother particle’s p_T for charged particles produced in the decay of a ρ meson, τ lepton, or a B^0 meson [122].

An idea of the tracking performance achieved by the ID during Run 2 is given in Figure 4.5, which shows the track reconstruction efficiency for charged particles produced in the decay of a ρ meson, τ lepton, or a B^0 meson, as a function of the mother particle’s p_T .

4.2 The calorimeters

The ATLAS calorimeters are designed to measure the energy of all charged and neutral particles produced in a pp interaction, with the exception of highly-penetrating muons and weakly-interacting neutrinos.

The ATLAS calorimetric system, depicted in Figure 4.6, consists of three distinct sub-detectors, each designed with specific purposes and utilizing different technologies. The **Electromagnetic (EM) Calorimeter** is optimized for measuring the energy of electrons and photons. The **Hadronic (HAD) Calorimeter** is dedicated to detecting and measuring incoming hadrons. Lastly, the **Forward Calorimeter (FCal)** is designed to measure both electromagnetic and hadronic showers in the forward regions of the detector.

The energy measurement performed by calorimeters is a destructive process. When a particle interacts with the detector material via electromagnetic or strong forces, it initiates a cascade (*shower*) of secondary particles, each with progressively lower energy. These secondary particles are fully absorbed, and the energy they deposit is converted into a measurable quantity, such as e.g. scintillation light, which is then used to determine the energy of the original particle.

ATLAS calorimeters are *sampling calorimeters* [123], which are constructed from alternating *absorbing layers* and *active layers*. The absorbing material (typically a high-density ma-




Figure 4.6 – Schematic layout of the ATLAS calorimetric system. Components highlighted in gold utilize Liquid Argon as the active medium, while those in gray employ scintillating tiles.

material such as lead or copper) reduces the energy of incoming particles, creating showers of secondary particles, while the active material collects the deposited energy and produces a measurable signal. In contrast, *homogeneous calorimeters* consist of a single material that both absorbs energy and generates the measurable signal. In homogeneous calorimeters, the entire energy deposited is collected and measured, while in sampling calorimeters the energy deposited in the absorber layers is lost. Therefore, homogeneous calorimeters typically offer better energy resolution than sampling calorimeters. An example of a homogeneous calorimeter is the CMS electromagnetic calorimeter [124], which uses PbWO_4 crystals that combine high density with adequate scintillating properties.

The size and geometry of the calorimeters are designed to contain the showers generated by incoming particles. The development of these showers depends on the particle's interactions with the detector material and varies significantly between primary electrons or photons and primary hadrons.

For example, high-energy electrons and photons produced in a pp collision interact with the detector material via bremsstrahlung and electron-positron pair production, respectively, generating secondary electrons and photons. These secondary particles, in turn, produce additional electrons and photons through the same processes, sustaining the development of the *electromagnetic shower*. The number of particles in the shower increases as it propagates through the material, until the average energy of the secondary particles falls below the *critical energy* ϵ_c , which corresponds to the threshold where ionization and photoelectric interactions become the dominant energy loss mechanisms for electrons and photons, respectively. reaching a maximum depth, t_{max} , which depends

on the initial energy E of the primary electron or photon, and on the *critical energy* ϵ_c . The critical energy depends on the properties of the absorber material (in particular, from its atomic number Z), and can be parameterized as:

$$\epsilon_c = \frac{610 \text{ MeV}}{Z + 1.24} \quad (4.5)$$

The shower depth at which the average secondary particles energy reaches ϵ_c is referred to as the *maximum depth* of the electromagnetic shower, and is given by:

$$t_{\max} = \ln \left(\frac{E}{\epsilon_c} \right) + C_j \quad (4.6)$$

where C_j equals -0.5 for a primary electron and +0.5 for a primary photon. At this stage, the showering process stops, and the remaining particles are fully absorbed by the material. The longitudinal development of the electromagnetic shower is characterized by the *radiation length*, X_0 , which represents the average distance an electron travels through the material to reduce its energy by a factor of e^{-1} (i.e., $\langle E(x) \rangle = E_0 e^{-x/X_0}$). The radiation length depends on the absorber material and is given by:

$$X_0 = \frac{716.4 \text{ g/cm}^2 \times A}{Z(Z + 1) \ln(287/\sqrt{Z})}, \quad (4.7)$$

where Z and A correspond to the atomic number and atomic mass of the absorber material, respectively. The depth required to contain 95% of the shower, in units of X_0 , can be expressed as:

$$t_{95\%} = t_{\max} + 0.08Z + 9.6 [X_0] \quad (4.8)$$

For example, for Liquid Argon and copper (used in the ATLAS electromagnetic calorimeter as the active medium and absorber material, respectively), the critical threshold and radiation length are:

$$\begin{aligned} \epsilon_c^{\text{LAr}} &= 32.84 \text{ MeV}, & X_0^{\text{LAr}} &= 14.00 \text{ cm} \\ \epsilon_c^{\text{Cu}} &= 19.42 \text{ MeV}, & X_0^{\text{Cu}} &= 1.436 \text{ cm} \end{aligned} \quad (4.9)$$

Hence, for a primary electron with an initial energy of $E = 100 \text{ GeV}$, the 95% longitudinal containment of the electromagnetic shower is achieved using approximately 260 cm of Liquid Argon or 20 cm of copper.

The characteristic transverse size of the electromagnetic shower is described by the Molière radius, ρ_M , which represents the dispersion of the deflection of an electron from the shower axis after traveling a distance equal to X_0 . Approximately 95% of the shower energy is contained within a distance of twice the Molière radius. The Molière radius depends only on the properties of the detector material and is independent of the initial

particle energy. It is given by:

$$\rho_M = \frac{21 \text{ MeV}}{\epsilon_c} \times X_0 \quad (4.10)$$

For Liquid Argon and copper, the Molière radii are:

$$\rho_M^{\text{LAr}} = 8.95 \text{ cm}, \quad \rho_M^{\text{Cu}} = 1.55 \text{ cm} \quad (4.11)$$

Electromagnetic showers, initiated by incoming electrons or photons, consist exclusively of secondary electrons and photons. In contrast, hadronic showers - triggered by incoming hadrons - produce many different species of particles (e.g., protons, pions, neutrons) through nuclear interactions with the absorber nuclei. A key process within hadronic showers is the decay of neutral hadrons (π^0 , η , etc.), primarily through the $\gamma\gamma$ channel, where the resulting secondary photons generate additional electromagnetic showers in the detector. The energy deposited by this electromagnetic component is known as the electromagnetic fraction (f_{em}), which depends from the initial hadron energy and typically accounts for 30% to 60% of the total energy of the hadronic shower. The remaining energy is associated with the hadronic component (f_{had}), which includes processes such as slow neutron production and spallation of protons. Unlike the electromagnetic component, some of the hadronic energy does not contribute to the detector signals due to the loss of the *invisible energy*, which comes primarily from the binding energy of nucleons released during nuclear reactions. This invisible energy can represent up to 40% of the total non-electromagnetic energy. Consequently, the signal produced in a hadronic calorimeter by an incoming hadron is smaller than that produced by an electron or photon of the same energy. This difference in response to electromagnetic and hadronic showers is referred to as *non-compensation*.

The longitudinal development of a hadronic shower is governed by the interaction length, λ_{int} , which represents the average distance a particle travels within the absorber material before initiating a nuclear interaction. This interaction length can be approximately parameterized as:

$$\lambda_{\text{int}} \approx 35 \cdot A^{1/3} \text{ g cm}^{-2} \quad (4.12)$$

For a given material, λ_{int} is typically much larger than the radiation length X_0 . For example, in copper, the interaction length is 15.32 cm, over ten times larger than X_0^{Cu} , while in Liquid Argon, $\lambda_{\text{int}}^{\text{LAr}} = 85.77 \text{ cm}$, which is approximately six times X_0^{LAr} . The 95% longitudinal containment of a hadronic shower, which depends on the initial energy E of the primary hadron, is described in units of λ_{int} :

$$t_{95\%} = t_{\text{max}} + 2\lambda_{\text{int}} E^{0.13} [\lambda_{\text{int}}], \quad (4.13)$$

where t_{max} , the maximum shower depth, is given by:

$$t_{\text{max}} = 0.2 \ln \left(\frac{E}{1 \text{ GeV}} \right) + 0.7 [\lambda_{\text{int}}] \quad (4.14)$$

For instance, the 95% longitudinal containment of a hadronic shower initiated by 100-GeV pions requires approximately 450 cm of Liquid Argon and 90 cm of copper. Consequently, hadronic calorimeters - typically sized at 7 to 10 times λ_{int} - are substantially larger than electromagnetic calorimeters. In collider experiments, hadronic calorimeters are always sampling calorimeters using high atomic number materials as absorbers to allow for a more compact detector volume.

Event-by-event fluctuations in the calorimeter response influence the resolution of energy measurements. The energy resolution of a calorimeter depends on the initial energy of the incoming particle and can be expressed as follows:

$$\frac{\sigma(E)}{E} \approx \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c \quad (4.15)$$

The first term, a , is called *stochastic term* and describes the contributions to the resolution given by the statistical fluctuations in the shower development. In a homogeneous calorimeter, this term is primarily influenced by fluctuations in the number of signal quanta contributing to the calorimeter response. For example, in a calorimeter where the active medium generates an ionization signal, this corresponds to the number of ionization tracks created by secondary electrons in the shower. Since this number is typically very large, the resulting statistical fluctuations have a minimal impact on energy resolution. However, in calorimeters that rely on a Cherenkov signal, where the light yield is around 30 photoelectrons per GeV of deposited energy, these fluctuations can significantly affect the resolution. For a homogeneous calorimeter using ionization signals, the stochastic term is usually around $\text{few } \%/ \sqrt{E}$. In electromagnetic sampling calorimeters, the main contribution to the stochastic term comes from sampling fluctuations, which arise due to event-by-event variations in the total energy deposited by shower particles in the active material. These fluctuations are influenced by the sampling fraction (i.e., the proportion of active material with respect to the absorber layers) and the sampling frequency, which refers to the number of active layers traversed by the shower along its longitudinal development. For these calorimeters, the stochastic term typically ranges from 5% to $20\% / \sqrt{E}$.

Hadronic calorimeters present a more complex scenario. In addition to the fluctuations described above, event-by-event variations in the electromagnetic component and invisible energy induced by hadronic showers introduce further contributions to the energy resolution. These fluctuations exhibit a non-Gaussian distribution, causing their impact on energy resolution to be less dependent on the initial hadron energy than the standard $\propto 1/\sqrt{E}$ behavior. In non-compensating hadronic calorimeters, the stochastic term can range from 50% to $100\% / \sqrt{E}$.

The second term in Equation 4.15, b , is known as the *noise term*. It accounts for contributions from electronic noise induced by readout circuits, which is generally negligible within the energy range studied by the ATLAS detector. The final term, c , is the *constant term*, which arises from factors such as detector non-uniformity, alignment, and uncertainties in electronic calibration. The constant term becomes the limiting factor for resolution at very high energies.

For a complete review of calorimetry for high energy physics see [125].

4.2.1 The Electromagnetic Calorimeter

After exiting the Inner Detector, particles produced in a pp collision pass through the ATLAS Electromagnetic (EM) Calorimeter [126–128]. The EM Calorimeter is a cylindrical detector, with an outer radius of 2.25 m and an overall length of 6.8 m. It is divided into a central barrel section (EMB), covering the $|\eta| < 1.475$ region, and two end-cap sections (EMEC), extending the coverage to $|\eta| < 3.2$. The transition areas between the barrel and end-caps are known as *crack regions* ($1.37 < |\eta| < 1.52$). The EM Calorimeter is a sampling calorimeter that uses 4-mm thick layers of Liquid Argon (LAr) as its active medium. Ionization charges produced by the cascade of secondary particles induced by the primary electron or photon are collected by copper electrodes immersed in the LAr layers. The absorber is made up of lead plates with varying thicknesses, ranging from 1.1 mm to 1.5 mm in the barrel, and from 1.7 mm to 2.2 mm in the end-caps. To avoid gaps along the ϕ coordinate due to the readout system, the copper electrodes and lead plates are arranged in an *accordion geometry*, as shown in Figure 4.7.




Figure 4.7 – Layout of the three accordion layers of the EM barrel calorimeter (left), and close detail of the LAr gap in the accordion calorimeter (right), showing also the electrode composition [126]. The size of the cells in the $\eta = 0$ plane for each longitudinal layer in the accordion is reported in 4.7a.

This design provides an effective length of $22 X_0$ in the barrel and $24 X_0$ in the end-caps. Each section of the EM Calorimeter is housed in its own cryostat to maintain the LAr at its operational temperature of 89 K.

The Electromagnetic Calorimeter is divided into 182,468 cells along the η and ϕ coordinates and consists of four longitudinal layers designed to precisely capture the longitudinal development of electromagnetic showers. The first layer is a presampler placed in front of the solenoid magnet, while the remaining three layers are arranged in the accordion geometry and constitute the so-called accordion calorimeter. The granularity of the

cells in both η and ϕ directions varies across the different layers and detector regions, as detailed below. The four layers are:

- **Presampler (PS).** A very thin module with a thickness of 1.1 cm, covering the $|\eta| < 1.8$ region. It is used to correct for energy losses in the material upstream of the calorimeter (1.7 X_0 at $|\eta| = 0$ and 2.2 X_0 at $|\eta| = 1.4$ [129]). The PS is segmented into cells of size 0.025×0.1 in $\Delta\eta \times \Delta\phi$.
- **Layer 1 (L1) or strips.** This layer accounts for up to 6 X_0 (including material in front) and has the finest segmentation in η , with $\Delta\eta = 0.0031$, to distinguish between prompt photons and π^0 mesons decaying into two nearly collinear photons. The cell size in the azimuthal direction is $\Delta\phi = 0.1$, while the granularity in the η direction becomes coarser, moving from the central to the forward regions: $\Delta\eta = 0.0031$ for $|\eta| < 1.8$, $\Delta\eta = 0.004$ for $1.8 < |\eta| < 2.0$, $\Delta\eta = 0.006$ for $2.0 < |\eta| < 2.5$, and $\Delta\eta = 0.1$ for $2.5 < |\eta| < 3.2$.
- **Layer 2 (L2) or middle.** This layer is 16 X_0 thick and is segmented into cells of size $\Delta\eta \times \Delta\phi = 0.025 \times 0.025$. It contains the majority of the energy from showers generated by photons and electrons.
- **Layer 3 (L3) or back.** This layer is designed to estimate the energy that leaks into the hadronic calorimeter. It has a thickness of 2 X_0 and is segmented into cells with a size of 0.050×0.025 in $\Delta\eta \times \Delta\phi$.

The nominal energy resolution of the EM Calorimeter is

$$\frac{\sigma(E)}{E} \approx \frac{10\%}{\sqrt{E [\text{GeV}]}} \oplus 0.7\% [126], \quad (4.16)$$

which is consistent with the typical performance of sampling electromagnetic calorimeters. The sampling term varies with $|\eta|$, while the noise term is negligible and is not expected to contribute for energies above 0.5 GeV.

4.2.2 The Hadronic Calorimeter

The ATLAS Hadronic Calorimeter is an 11.5 m-long cylinder with an outer radius of 4.25 m, covering a pseudorapidity range of $|\eta| < 4.9$. It is placed after the EM Calorimeter and is designed to fully contain and measure showers initiated by hadrons. The central barrel region, which covers $|\eta| < 1.7$, contains the Tile Calorimeter [130, 131], a sampling calorimeter that uses scintillating plates (tiles) as the active medium and steel as the absorber. It is segmented longitudinally into three concentric layers, providing a total thickness of $11\lambda_{\text{int}}$, with a granularity of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$.

The Hadronic End-Cap Calorimeter (HEC) [132, 133] is also a sampling calorimeter, but it uses copper or tungsten plates as absorbers, interleaved with radiation-hard LAr gaps as the active material. The HEC consists of two separate disks (one for each end-cap) located directly behind the EM End-Cap Calorimeter and shares the same cryostat. It covers the pseudorapidity range $1.5 < |\eta| < 3.2$, overlapping with the Tile Calorimeter in the central region and the FCal in the forward region.

The combined performance of the EM and HAD Calorimeters results in a global energy resolution of

$$\frac{\sigma(E)}{E} \approx \frac{50\%}{\sqrt{E [\text{GeV}]}} \oplus 3\% \quad [126] \quad (4.17)$$

for jet energy measurements.

4.2.3 The Forward Calorimeter

The ATLAS Forward Calorimeter (FCal) [134] is designed to measure the energy of both electromagnetic and hadronic showers in the forward regions, extending the pseudo-rapidity coverage of the ATLAS calorimetric system to $3.1 < |\eta| < 4.9$. Tolerance to radiation is fundamental in the forward region: therefore, LAr is used as an active material, with copper or tungsten rods acting as absorbers.

The typical energy resolution achieved by the FCal for incoming pions is given by:

$$\frac{\sigma(E)}{E} \approx \frac{(94.2 \pm 1.6)\%}{\sqrt{E}} \oplus (7.5 \pm 0.4)\% \quad [135], \quad (4.18)$$

while for electrons is

$$\frac{\sigma(E)}{E} \approx \frac{(28.5 \pm 1.0)\%}{\sqrt{E}} \oplus (3.5 \pm 0.1)\% \quad [135], \quad (4.19)$$

4.3 The Muon Spectrometer

Muons are, along with the weakly-interacting neutrinos, the only particles able to escape the ATLAS calorimetric system without being absorbed.

Due to their significantly higher mass compared to electrons, muons experience a drastically reduced energy loss through bremsstrahlung - by a factor of approximately $m_\mu^{-2}/m_e^{-2} \approx 10^{-6}$ - which greatly lowers the likelihood of initiating electromagnetic showers in the ATLAS calorimeters. In addition, muons do not interact strongly with the detector material. For these reasons, muons can escape the ID and the calorimeters with minimal energy loss. Hence, the Muon Spectrometer (MS) [136], the outermost component of the ATLAS detector, is responsible for tracking their trajectories and measuring their momenta.

The design of the MS is determined by the structure of the toroid magnet (see Section 4.4), which divides the sub-detector into octants. The toroid magnetic fields curve the trajectories of incoming muons in the (R, z) plane, thus allowing the measurement of their momenta.

Muons' trajectories, projected onto the bending plane, are reconstructed using **Monitored Drift Tubes (MDTs)** across most of the detector's pseudorapidity range ($|\eta| < 2.7$). The MDTs are aided by the **Cathode Strip Chambers (CSCs)** in the forward region ($2 < |\eta| < 2.7$). Each MDT consists of 3 to 8 layers of aluminum drift tubes filled with a high-pressure Ar/CO₂ gas mixture with a central 50 μm wire, achieving a spatial resolution of 35 μm per chamber in the bending plane. CSCs are multiwire proportional cham-

bers, whose anodes lie along the radial direction, while the cathode strips run parallel or perpendicular to the wires. This configuration provides a three dimensional measurement of the muons' positions, offering a resolution of $40\ \mu\text{m}$ in the bending plane and 5 mm in the azimuthal direction.

The MS also incorporates two additional types of muon chambers: **Resistive Plate Chambers (RPCs)** and **Thin Gap Chambers (TGCs)**, which provide fast triggering (15-25 ns) for muon tracking in the barrel ($|\eta| < 1.05$) and end-cap regions ($1.05 < |\eta| < 2.4$), respectively.

The ATLAS muon system, particularly the MDTs, offers a momentum resolution ranging from 2-3% at 10 GeV to approximately 10% at 1 TeV.

4.4 The magnet system

A charged particle moving through a magnetic field follows a curved trajectory, whose curvature radius depends on the particle's momentum and the intensity of the magnetic field. For a uniform magnetic field of strength B , the component of the particle's momentum perpendicular to the direction of the magnetic field is given by:

$$p_T [\text{GeV}] \approx 0.3 \cdot |q| \cdot B [\text{T}] \cdot R [\text{cm}], \quad (4.20)$$

where q is the particle's charge and R is the curvature radius. The ATLAS detector exploits this relationship to measure the momenta of charged particles, by reconstructing the curvature of their tracks. This is achieved by immersing the Inner Detector and the Muon Spectrometer in powerful magnetic fields generated by the ATLAS magnet system, which covers a volume of approximately $12,000\ \text{m}^3$ by combining a superconducting solenoid magnet with a toroidal magnet system.

The solenoid magnet [137] is a superconducting coil that produces a 2 T magnetic field aligned with the beam axis for the ID. It operates with an 8 kA electric current and is maintained at a temperature of 4.5 K. The cooling is achieved by situating the coil directly within the cryostat, which also contains the ATLAS calorimeter (see Section 4.2). The solenoid is 5.8 m in length, with an inner radius of 1.23 m and an outer radius of 1.28 m, corresponding to just 0.66 radiation lengths.

The ATLAS toroid magnet system [138, 139] is located outside the calorimetric system and consists of a barrel toroid, which generates a magnetic field ranging from 0.2 to 2.5 T, and two endcap toroids, each providing a field strength of up to 3.5 T. The barrel toroid, shown in Figure 4.8, is composed of eight superconducting coils arranged radially and symmetrically around the beam axis, with a current of 20.5 kA running through them. Each coil is housed in an independent cryostat. The barrel toroid has an inner diameter of 9.4 m, an outer diameter of 20.1 m and it is 25.3 m long. Two endcap toroids generate the magnetic field in the forward regions, ensuring nearly complete coverage up to $|\eta| < 2.7$. Each endcap toroid consists of eight superconducting coils (all enclosed in a single cryostat), measuring 5 meters in length and spanning a radial width from 1.65 to 10.7 meters.




Figure 4.8 – Photo of the ATLAS detector taken during construction: all eight coils of the toroid barrel magnet system are visible.

4.5 The trigger and data acquisition systems

The LHC beams collide at the interaction point at the center of the ATLAS detector with a frequency of 40 MHz, meaning a bunch crossing occurs every 25 ns. Each bunch crossing results in 30 to 60 simultaneous proton-proton (pp) collisions, collectively referred to as an *event*. The average size of each event is around 1.5 MB [140]. At a collision rate of 40 MHz, this would generate an overwhelming data rate of approximately 60 TB/s. Processing and storing this volume of data is infeasible, so a rapid decision must be made for each event to determine whether it should be kept for further analysis. Therefore, one of the key components of the ATLAS experiment is the trigger and data acquisition system, which handles online event selection and the readout of data for the chosen events.

The trigger system is responsible for reducing the data rate from 40 MHz to 1 kHz, achieved through the coordination of two subsystems: the **Level 1 (L1) trigger** and the **High-Level Trigger (HLT)** [141, 142].

The L1 trigger [143] uses low granularity information from the calorimeters and muon subsystems to select events based on event-level quantities (such as total energy deposited in the calorimeter or missing transverse momentum), or number of physics objects exceeding a certain p_T threshold, or, finally, selection on topological quantities such as invariant masses and angular distances. These computations are performed in real time, by custom hardware located on the ATLAS detector. The L1 trigger reduces the event rate from 40 MHz to 100 kHz, making its selection decision within 2.5 μ s.

Upon receiving an L1 trigger accept signal, full-granularity data from various sub-detectors is sent to the ReadOut Drivers (RODs) for initial processing and formatting, and then to the ReadOut Systems (ROSeS) for buffering. The L1 trigger also provides the (η, ϕ) coordinates of Regions of Interest (RoIs) in the detector, where interesting activity was detected, to guide further analysis by the next-level trigger.

The HLT [140] uses the full-granularity data associated with the RoIs identified by the L1 trigger to perform partial event reconstruction using more complex algorithms, exploiting also the information from the ID which in Run 2 consisted of 40000 processing units (increased to 60000 in Run 3), making a decision with a 40 ms latency. The HLT further reduces the event rate from 100 kHz to around 1 kHz, and makes the final decision on whether to store the event data for offline reconstruction. Upon a trigger accept, the full event data is transferred from the ROs to CERN's Tier-0 data processing center at a rate of 1.2 Gb/s for initial reconstruction, and later distributed to computing centers worldwide for offline analysis.

The L1 and HLT decisions are guided by *trigger menus* [144], which describe the types of events deemed valuable for offline analysis and assigns a predefined rate for each of these types. Figure 4.9 illustrates the total L1 trigger rate for a typical LHC fill in Run 2, along with the contributions from key physics signatures, which include preliminary information on object identification and isolation (e.g., EM22VHI represents isolated electromagnetic activity in the calorimeter system, associated with a transverse momentum of at least 22 GeV). Figure 4.10 displays the total HLT rate, with a breakdown of the rates by reconstructed object groups, such as electrons and photons.




Figure 4.9 – Example of L1 trigger rates as a function of time, for a fill of the LHC taken in September 2018, when the peak instantaneous luminosity was registered to be $\mathcal{L} = 2.0 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, and the mean peak number of pp interaction per bunch crossing was 56. Different colors represent ROIs based on different objects, such as electromagnetic clusters (EM), muon candidates (MU), jet candidates (J), missing transverse energy (XE) and tau candidates (TAU) [145].




Figure 4.10 – Example of HLT trigger rates as a function of time, for a fill of the LHC taken in September 2018, when the peak instantaneous luminosity was registered to be $\mathcal{L} = 2.0 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$, and the mean peak number of pp interaction per bunch crossing was 56. Different physics signatures (e.g. photons, jets, etc.) are visible [145].

Photon reconstruction

Each event recorded by the ATLAS detector undergoes offline reconstruction to extract the physics objects (such as photons, electrons, hadronic jets, muons, τ leptons, and missing transverse momentum) from the measured signals in the ATLAS sub-detectors. For instance, hits in the different layers of the Inner Detector or the Muon Spectrometer (induced by charged particles passing through their active layers), are combined to form *tracks*. On the other hand, energy deposits in the calorimeters from electromagnetic and hadronic showers generated by incoming electrons or photons and hadrons are used to create *calorimetric clusters*. These low-level objects are then combined to reconstruct the final physics objects corresponding to the activity observed in the ATLAS sub-detectors. This information is then used in physics analyses. A simple sketch of how information from different sub-detectors is combined to identify particles emerging from the pp collisions is reported in Figure 5.1.

The searches for Higgs boson pair production in the final state with two photons and two bottom quarks discussed in this thesis depend heavily on the robust and efficient reconstruction and identification of photons and of hadronic jets originating from bottom quark fragmentation (b -jets). This Chapter provides a description of the reconstruction algorithms and identification techniques employed for photons during the Run 2 data-taking period, along with a discussion of their performance. The reconstruction and identification of other physics objects (such as electrons and muons) is also briefly addressed.

The searches for Higgs boson pair production where one of the two Higgs bosons decays to two photons (and, more generally, all the $H \rightarrow \gamma\gamma$ analyses) critically rely on the ability to include good-quality photon pairs in the signal region. These photons appear as compact and isolated energy deposits in the EM Calorimeter. The dominant background arises from hadronic jets with a large electromagnetic component (originated by neutral hadrons decays to photon pairs). These jets can be incorrectly reconstructed and classified as photons, in which case they are referred to as *fake photons*. The need to minimize the contribution from fake photons while maintaining high efficiency in retaining genuine photons drives the optimization of reconstruction, identification and isolation algorithms, which are outlined in the following sections.




Figure 5.1 – Sketch of an ATLAS wedge, with sub-detectors and different physics objects highlighted [146]. Dashed tracks are invisible to ATLAS, and their presence is inferred only via a transverse momentum imbalance (see Section 6.3).

5.1 Electron and photon reconstruction

Photon reconstruction in the ATLAS detector runs parallel to electron reconstruction: both are based on collections of energy deposits almost fully contained within the EM Calorimeter and on tracks and secondary vertices reconstructed in the Inner Detector. The reconstruction begins by defining clusters from energy deposits in topologically connected calorimeter cells (termed *topoclusters*). Subsequent steps in the algorithm then differentiate between electrons and photons, which may or may not have converted into an electron-positron pair due to interactions with the ID material. Photons are classified as *converted* or *unconverted* based on whether or not a conversion is reconstructed. A diagram representing the full algorithm for photon and electron reconstruction is shown in Figure 5.2.

In the reconstruction process, tracks and candidate conversion vertices identified in the ID are matched with energy topoclusters in the EM Calorimeter. If a topocluster has no associated tracks or conversion vertices, it is reconstructed as an unconverted photon. Conversely, if the topocluster is matched with one or more conversion vertices, a converted photon candidate is created. Finally, a topocluster matched with a track (or tracks) is identified as an electron candidate. The reconstruction of candidate electrons or photons then involves building superclusters using the matched topoclusters as inputs. Superclusters are dynamic, variable-size energy clusters designed to incorporate additional clusters from bremsstrahlung photons into the original electron cluster or to




Figure 5.2 – Diagram of the algorithm for photon and electron reconstruction [147]. The three conceptual phases of the algorithm (preparation of topoclusters in the EM Calorimeter and tracks in the Inner Detector, growth of the electron and photon superclusters, and building of the final analysis objects) are highlighted with different colors.

merge electron and positron clusters from photon conversions. Once the superclusters are formed, tracks are matched to electron superclusters, and conversion vertices are matched to photon superclusters. Supercluster reconstruction runs parallelly for both electrons and photons: therefore, a single topocluster can seed both an electron and a photon supercluster. In such cases, both an electron and a photon are reconstructed, and these objects are flagged as ambiguous. The final physics objects used in the analysis are then constructed: the ambiguous objects are classified, the electrons and photon energies are calibrated, and, finally, all the necessary information for identification and isolation is computed. For a comprehensive overview of the electron and photon reconstruction algorithms, refer to [147].

5.1.1 EM Calorimeter clusters formation

The reconstruction algorithm for both electrons and photons is built upon selecting energy deposits measured in variable-size clusters, known as topoclusters, which are constructed from topologically connected cells in the EM Calorimeter [148].

The key quantity that governs the development of topoclusters is the *significance* of a calorimeter cell, defined as:

$$\zeta_{\text{cell}}^{\text{EM}} = \frac{|E_{\text{cell}}^{\text{EM}}|}{|\sigma_{\text{noise}}^{\text{EM}}|} \quad (5.1)$$

where $E_{\text{cell}}^{\text{EM}}$ represents the energy deposited in the cell, and $\sigma_{\text{noise}}^{\text{EM}}$ is its expected noise, calculated considering both the read-out electronics and an estimate of the pileup noise corresponding to the average instantaneous luminosity. The clustering algorithm scans the calorimeter map, selects seed cells, and grows clusters around them following the “4-2-0” scheme:

1. Seed cells with $\zeta_{\text{cell}}^{\text{EM}} \geq 4$ are identified, excluding those in the Presampler and L1 layer of the EM Calorimeter.
2. Neighbouring cells with $\zeta_{\text{cell}}^{\text{EM}} \geq 2$ are then collected around the seeds. If a cell with $\zeta_{\text{cell}}^{\text{EM}} \geq 2$ is shared between two clusters, those are merged together forming a single one.
3. A surrounding crown, made up of all nearest-neighbor cells with $\zeta_{\text{cell}}^{\text{EM}} \geq 0$ is included in the topocluster.
4. Finally, topoclusters with two or more local maxima (i.e. cells with $E_{\text{cell}}^{\text{EM}} > 500$ MeV with at least four neighbours on the condition that none of them have a larger signal) are split into different topoclusters.

The same clustering algorithm is used by ATLAS for the reconstruction of many physics objects (τ leptons, jets, etc.), therefore topoclusters may include cells from both the EM and the hadronic calorimeters. To specifically identify electromagnetic showers initiated by incoming photons or electrons, the EM component of the topocluster, defined as

$$f_{\text{EM}} = \frac{E_{\text{L1}} + E_{\text{L2}} + E_{\text{L3}}}{E_{\text{cluster}}}, \quad (5.2)$$

where E_{Li} is the cluster energy in the i^{th} layer of the EM Calorimeter and E_{cluster} is the total cluster energy, must be greater than 0.5. Clusters with $f_{\text{EM}} > 0.5$ are considered to originate from an electromagnetic shower, and cells from the hadronic calorimeter are removed. After excluding the hadronic cells, only topoclusters with EM energy larger than 400 MeV are retained. Figure 5.3a illustrates the distribution of f_{EM} for topoclusters in the EM Calorimeter generated from simulated electrons without pileup, compared to those originating from pileup activity. Similarly, Figure 5.3b shows the efficiency of the f_{EM} selection as a function of the f_{EM} threshold for simulated electrons without pileup and for pileup clusters. The chosen cut, $f_{\text{EM}} > 0.5$, is able to reject approximately 60% of pileup clusters while maintaining 100% efficiency in selecting true electron topoclusters.




Figure 5.3 – Distribution of f_{EM} (left) and reconstruction efficiency as a function of the f_{EM} selection threshold (right) for simulated true electron clusters (black) and pileup clusters (red) clusters.

5.1.2 Track reconstruction and track-cluster matching

The next step in the reconstruction process is selecting reconstructed tracks and reconstructing photon conversion vertices in the Inner Detector and matching them with topoclusters in the EM Calorimeter. This track-cluster matching allows the initial differentiation between electrons, unconverted photons, and converted photons.

Track-pattern reconstruction is carried out throughout the entire Inner Detector as part of the full event reconstruction process [122, 149]. Charged particles passing through the ID deposit a small fraction of their energy in the ID layers through ionization, producing small electrical signals known as hits. These hits are used to reconstruct the trajectories of charged particles, which follow helicoidal paths with radii inversely proportional to their momentum, due to the magnetic field permeating the ID (see Equation 4.20).

The track reconstruction algorithm begins by identifying primary tracks produced by charged particles with lifetimes larger than 3×10^{-11} s, originating from the hard scattering vertex, using an inside-out approach [149]. The process starts with identifying

groups of three hits in the Pixel detector and the SCT, known as track seeds. Primary track candidates are then constructed by adding hits from other ID layers that are compatible with the track seed, using a Kalman filter algorithm [150]. A dedicated ambiguity resolution procedure is applied to resolve overlaps between track candidates and reject unrelated combinations of hits, that are not fired by a single particle (namely, *fake tracks*). The set of tracks resulting from the ambiguity resolution procedure are then refitted using a global χ^2 algorithm to determine the final impact parameters. Next, an attempt to extending the tracks to the TRT is performed: TRT hits compatible with the extrapolated trajectory are added, without modifying the original silicon track.

In the second step of the track reconstruction algorithm, an outside-in approach is used: it begins by identifying regions of interest in the EM Calorimeter and attempts to match them with segments of hits in the TRT. These TRT segments are then paired with short silicon track seeds (composed of two hits), creating track candidates through a process similar to the primary track reconstruction described earlier. Finally, a dedicated ambiguity resolution procedure is applied to the outside-in tracks, and a global χ^2 re-fit yields the final track candidates.

Track segments reconstructed in the TRT that do not match any track candidates are used to reconstruct secondary vertices, including those from photon conversions, by extrapolating the tracks from the TRT to the silicon detectors. Finally, an ambiguity resolution procedure is applied to tracks that share hits. The reconstructed tracks are the basis for the vertex-finding algorithm [151], which identifies candidate primary and secondary vertices. This algorithm determines the vertex position and refits the tracks, applying the constraint that they originate from the reconstructed vertex. Among the primary vertices, the one with the highest sum of the squared transverse momenta of its associated tracks is selected as the *hard-scattering vertex*.

Within the electron and photon reconstruction process, the standard tracking procedure is further refined to take into account large energy loss via bremsstrahlung for electrons and to improve the identification of photon conversions. Fixed-size clusters in the EM Calorimeter compatible with electromagnetic showers are used to define regions of interest. If track seeds without a matched primary track are found within these regions, a modified pattern recognition algorithm is applied. This modified algorithm, still based on the Kalman filter, allows for up to 30% energy loss at each material intersection to account for large bremsstrahlung effects. Following this, all tracks are refitted using a global χ^2 algorithm, allowing for additional energy loss where the standard track fit fails.

Track-cluster matching

The following steps involve matching the selected tracks with the topoclusters in the EM Calorimeter.

A preliminary *loose match* between fixed-size EM clusters and reconstructed tracks with silicon hits is performed, based on spatial constraints on the η and ϕ coordinates of the cluster's barycentre and the track extrapolated to the L2 layer of the EM Calori-

meter. For a cluster to loosely match a track with reconstructed charge q , the following conditions must be met: $|\Delta\eta| < 0.05$ and $-0.20 < q \cdot (\phi^{\text{track}} - \phi^{\text{cluster}}) < 0.05$ when the track energy is used for extrapolation from the last Inner Detector hit, or $-0.10 < q \cdot (\phi^{\text{track}} - \phi^{\text{cluster}}) < 0.05$ when the cluster energy is used for extrapolation from the track perigee. The asymmetric requirement on $\Delta\phi$ accounts for the energy of radiated photons included in the calorimetric cluster but missed by the track reconstruction. Only the silicon tracks that satisfy this loose match are then refitted using the Gaussian Sum Filter (GSF) algorithm [152], which further improves track parameter estimation. This is particularly beneficial for low- p_T electrons or electrons undergoing significant energy losses via bremsstrahlung.

The refitted tracks are then used to perform the final matching with the topoclusters. A track with reconstructed charge q is considered matched to a topocluster if, when extrapolated from its perigee to the EM L2 layer, it satisfies $|\Delta\eta| < 0.05$ and $-0.10 < q \cdot (\phi^{\text{track}} - \phi^{\text{cluster}}) < 0.05$.

If multiple tracks are matched to the same topocluster, the preferred track is selected based on the presence of pixel hits and the smallest ΔR between the topocluster barycentre and the extrapolated track position in the L2 layer. The properties of the candidate electron are then determined from the highest-ranking track.

Reconstruction of photon conversions

Tracks that are loosely matched with fixed-size EM clusters are used as input for reconstructing photon conversion vertices in the ID. Two types of photon conversions are considered: *double-track* and *single-track* conversions:

- A *double-track* conversion vertex is formed from two oppositely charged tracks originating from the same point. Geometric criteria, such as the angle between the tracks, their separation at the point of closest approach, and their distance at the reconstructed vertex, ensure that the tracks are consistent with the decay of a massless particle. Based on the presence of hits in the silicon layers of the ID, double-track conversion vertices are classified as *double silicon*, *silicon-TRT*, or *TRT-TRT*, depending on whether both, one, or neither of the associated tracks has silicon hits.
- *Single-track* conversion vertices account for asymmetric photon conversions, where one electron carries most of the initial photon energy, or for conversions of highly energetic photons, where the electron tracks are highly collimated. Single-track vertices are particularly relevant for photon conversions occurring at larger radii, where the TRT has lower spatial resolution compared to the silicon layers, making it easier to miss tracks or fail to resolve closely spaced tracks. Therefore, single-track conversion vertices are reconstructed only when there are no hits in the innermost ID layer, thus reducing the likelihood of prompt electrons being misidentified as single-track photon conversions. In such cases, the innermost track hit is assigned as the conversion vertex.

Conversion vertices are then matched to topoclusters by extrapolating the tracks to the EM L2 layer and comparing their distance to the topocluster barycentre in the η and ϕ

coordinates. For conversion vertices with tracks that have silicon hits, a match is considered valid if the distance between the tracks and the topocluster is within $\Delta\eta < 0.05$ and $\Delta\phi < 0.05$ after extrapolation. For conversion vertices composed solely of TRT tracks, the matching criteria depend on whether the first track is in the TRT barrel or endcaps. In the barrel, a match requires that the distance to the topocluster barycentre is within $\Delta\eta < 0.35$ and $\Delta\phi < 0.02$. In the endcap, stricter criteria are applied, requiring $\Delta\eta < 0.2$ and $\Delta\phi < 0.02$.

If multiple conversion vertices are matched to a single topocluster, preference is given to double-track conversions with silicon hits, followed by other double-track conversions and then single-track conversions. Within each category, the vertex with the smallest conversion radius is selected as the most likely candidate.

5.1.3 Supercluster growth

The loosely matched topoclusters and tracks serve as inputs to the *superclustering* algorithm [153], which operates in two distinct stages, described below.




Figure 5.4 – Pictorial representation of the superclustering algorithm for electrons and photons. Seed clusters are shown in red, satellite clusters in blue [147].

Supercluster seeding

The first stage involves scanning the reconstructed topoclusters to identify those suitable for seeding the growth of *superclusters*. The criteria for selecting seed clusters differ for

candidate electron and photon superclusters: the superclustering algorithm proceeds independently for electrons and photons. For electrons, seed topoclusters with $E_T > 1$ GeV that are matched to a track with at least four silicon hits in the ID are selected. In contrast, seed topoclusters with no matched track and $E_T > 1.5$ GeV are chosen for reconstructing photon superclusters.

Adding satellite topoclusters

In the second stage, the algorithm searches for satellite topoclusters within a window of $\Delta\eta \times \Delta\phi = 0.075 \times 0.125$ centered on the seed topocluster's barycentre. Only topoclusters with E_T lower than the seed are considered as potential satellite clusters. For electrons, additional satellite clusters are accepted within a larger window of $\Delta\eta \times \Delta\phi = 0.125 \times 0.3$ around the seed barycentre, provided they have at least one matched track with silicon hits in the ID and share the same best-matched track as the seed. For converted photons, satellite clusters are selected if they are matched to the same conversion vertex as the seed or to one of the tracks associated with that vertex.

Figure 5.4 illustrates the requirements for satellite clusters applied to candidate electron and photon superclusters.

Supercluster formation

Finally, the candidate electron and photon superclusters are formed by merging all the calorimeter cells from the seed topoclusters and their associated satellites. The size of each constituent topocluster is constrained to a maximum width of 0.075 in the η direction around its barycentre in the barrel (0.125 in the endcaps) of the EM Calorimeter. This constraint improves the estimation of systematic uncertainties related to the linearity of the energy response.

5.1.4 Preparation of electron and photon candidates for physics analysis

After the supercluster growth process is completed, tracks and conversion vertices are matched to electron and photon superclusters using the same procedures as for topoclusters.

Since electron and photon superclusters are constructed independently, a single topocluster might serve as the seed for both an electron and a photon supercluster. To address this, an ambiguity resolution algorithm is applied to finalize the classification of the candidates, ensuring that each particle is correctly identified as an electron, a photon, or both. The simplest scenario involves a photon supercluster whose seed is not also matched to an electron supercluster: in this case, only a photon is reconstructed. Similarly, only an electron is reconstructed when an input electron supercluster does not share its seed with a photon. Otherwise, the following ambiguity resolution process (also summarized in Figure 5.5) is used to address the ambiguous cases:

- A given initially ambiguous object is classified only as a **photon** if its associated electron track has no hits in the silicon trackers, or, if that is the case, if it has no hits in the Pixel detector and is part of a matched double silicon conversion vertex.

- The object is flagged only as an **electron** if its track has four hits in the silicon trackers and is not part of a conversion vertex. Conversely, if a conversion vertex is found, the candidate particle is still reconstructed as an electron if the electron track has an innermost hit¹ and the matched conversion vertex is not a double silicon vertex or, if that is the case, if at least one of its conversion tracks does not have an innermost hit. Finally, if the latter condition fails, an electron is reconstructed if its track is part of a silicon-TRT or a TRT-TRT conversion vertex whose radial distance with the innermost hit of the track is above 40 mm.
- **Both an electron and a photon** are reconstructed from the same initially ambiguous object if none of the previously described requirements are fulfilled or if the track p_T is smaller than 2 GeV or if $E^{\text{cluster}}/p^{\text{track}} > 10$. In such cases, the procedure still marks the initial candidate particle as ambiguous and each analysis is responsible for the final classification of such objects, depending on their specific requirements.

The construction of the final analysis objects is completed after their energy calibration and the calculation of the discriminating variables for particle identification and isolation criteria. These steps are summarized in the following Sections.




Figure 5.5 – Flowchart showing the logic of the ambiguity resolution for particles initially reconstructed both as electrons and photons. An *innermost hit* is a hit in the functioning pixel nearest to the beam-line along the track trajectory, E/p is the ratio of the supercluster energy to the measured momentum of the matched track, R_{conv} is the distance of the conversion vertex from the interaction point, and R_{firstHit} is the smallest radial position of a hit in the track or tracks that constitute the conversion vertex.

5.2 Electron and photon energy calibration

The accurate reconstruction and calibration of electron and photon energy are essential for physics analyses that involve these particles. This is true also for the analyses involving the $H \rightarrow \gamma\gamma$ decay (including the $HH \rightarrow b\bar{b}\gamma\gamma$ and the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ searches, described in this thesis), which depend on the precise measurement of the invariant mass

¹ An innermost hit is a hit in the functioning pixel nearest to the beamline along the track trajectory.

of the two photons in the final state, $m_{\gamma\gamma}$. The $m_{\gamma\gamma}$ variable is a powerful discriminator between the $H \rightarrow \gamma\gamma$ signal and the background. In fact, the signal exhibits a narrow peak in the $m_{\gamma\gamma}$ spectrum centered around $m_H \approx 125$ GeV, whose width is typically 1-2 GeV and is determined entirely by detector resolution effects, given that the natural width of the Higgs boson is approximately 4 MeV according to the SM (see Section 1.3.3). On the other hand, the dominant background processes produce a smoothly decreasing $m_{\gamma\gamma}$ distribution with no peaking structures. Accurate photon energy reconstruction and calibration lead to a sharper $m_{\gamma\gamma}$ peak for $H \rightarrow \gamma\gamma$ resonant processes, a more precise determination of the peak position, and a reduction in systematic uncertainties, affecting both the $m_{\gamma\gamma}$ peak and width. Improving the photon energy reconstruction and calibration enhances the signal-to-background ratio in $H \rightarrow \gamma\gamma$ analyses and increases the precision of measurements related to the Higgs mass and properties in the $H \rightarrow \gamma\gamma$ channel. Finally, the electron calibration is also a fundamental ingredient for measuring the Higgs mass and properties in the $H \rightarrow ZZ^* \rightarrow 4\ell$ channel and for precision electroweak measurements involving the W and Z bosons.

The energy reconstruction process begins by converting the electronic signal read out from each EM Calorimeter cell within the electron or photon supercluster into a raw energy value deposited by the incident particle in that cell [154, 155].

An overview of the full calibration procedure is presented in Figure 5.6, and the steps are described in more detail below and in the following Sections [19].

- **Step 1: LAr layers calibration.** Since the EM Calorimeter is longitudinally segmented, the response of each layer is calibrated separately. The raw energy of reconstructed electrons and photons in data is calibrated, to ensure that the response of each EM layer to incoming electrons or photons matches that in the simulation. This involves determining an energy scale for the Presampler, α_{PS} , and the relative responses of the first and second accordion layers (L1 and L2), α_{12} . A global scale, α_{acc} , is applied as an overall correction to the energy measured in all three accordion layers (L1, L2, and L3). After applying these corrections on data, the overall raw energy, E_{raw} , is obtained by summing the cell energies within the electron or photon supercluster.
- **Step 2: Simulation-based calibration.** The next stage of the energy calibration procedure is to apply a *simulation-based calibration* to the raw energy values in both data and Monte Carlo (MC) simulations. The simulation-based calibration corrects the raw energy, E_{raw} , which does not account for the electromagnetic shower energy deposited outside the cluster or for energy lost in inactive material upstream or within the calorimeter, to reflect as accurately as possible the true energy of the incoming particle. A Boosted Decision Tree (BDT)-based regression algorithm is trained separately for electrons, unconverted photons, and converted photons, in intervals of $|\eta|$ and transverse energy, using, as inputs, various quantities describing the longitudinal and lateral development of the EM shower, such as the ratio of raw energies across different layers, the η coordinate of cell edges, the periodicity of the lead absorber in each region, and the radius of photon conversions.




Figure 5.6 – Flowchart showing the electron and photon energy calibration procedure in ATLAS.

The BDT model is trained on simulated single-particle samples to predict the ratio $E_{\text{true}}/E_{\text{raw}}$ for incoming particles, where E_{true} denotes the true energy. This trained BDT model is then applied at the reconstruction level to both data and MC simulations. By evaluating the BDT on reconstructed electrons and photons, a predicted correction factor $E_{\text{true}}/E_{\text{raw}}$ is obtained, which is then multiplied by E_{raw} to yield the calibrated energy, E .

- **Step 3: Uniformity and ADC non-linearity corrections.** Following the simulation-based calibration, additional corrections are applied to data to account for geometric inhomogeneities and any residual non-linearity in the EM Calorimeter's electronic readout.
- **Steps 4 and 5: In-situ calibration for scale and resolution.** A final calibration step (the *in-situ calibration*) relies on $Z \rightarrow ee$ events. Global scale factors α are extracted and applied to data, in order to align the Z resonance peak position observed in data with that in the simulation. In the same measurement, smearing factors, which account for differences in the resolution of the $Z \rightarrow ee$ peak between data and simulation, are extracted. These smearing factors are applied to the simulation to match the slightly worse resolution observed in data. The presence of passive material upstream of and within the calorimeter, as well as the EM layer calibration procedure, introduces systematic uncertainties affecting the final calibrated energy E for electrons and photons. These uncertainties vary with the η coordinate and E_T of the candidate particle. To address this, $Z \rightarrow ee$ events are used to measure any residual dependence of the overall electron energy scale on the particle's E_T (*linearity fit*), allowing further calibration adjustments and providing additional constraints on associated uncertainties.
- **Step 6: Photon-specific corrections.** The calibration factors derived from $Z \rightarrow ee$ events are assumed to reflect the intrinsic response of the calorimeter and are applied identically to both electrons and photons. However, photon-specific corrections are necessary to account for differences in the lateral development of electron and photon showers.
- **Step 7: Validation.** Finally, the entire calibration chain is validated using low- E_T electron candidates from $J/\psi \rightarrow ee$ decays and photon candidates from radiative $Z \rightarrow \ell\ell\gamma$ decays.

My contribution to the electron and photon calibration is the extraction of the sensitivity of the total calibrated electron energy (E) to the energy measured in the accordion calorimeter (E_{acc}) from simulation, which is a key ingredient in the evaluation of the accordion scale α_{acc} . In addition, I evaluated, from simulation, the sensitivity of the total electron and photon energy to the raw energies in different layers and to the sum of the raw cell energies in the EM L1 or L2 layers read out either with the High, Medium, or Low gains of the EM Calorimeter electronics. These quantities are crucial for modelling the systematic uncertainties affecting the electron and photon energy across the E_T spectrum of the particle. Finally, I evaluated the systematic uncertainty on the photon energy scale related to the mismodelling of the photon conversions reconstruction between data and

simulations. This work is integral to the final electron and photon energy calibration of the ATLAS Collaboration for LHC Run 2 data.

5.2.1 Energy reconstruction in the EM Calorimeter

As electrons and photons pass through the EM Calorimeter, they interact with the absorber material, generating electromagnetic showers. These showers ionize the Liquid Argon in the active layers, creating an ionization current that is proportional to the deposited energy. The electrons and ions are then driven by a constant electric field toward the copper electrodes immersed in the Liquid Argon and placed between the active layers and the absorber layers. The ionization current forms a triangular pulse, read by the Front End Boards (FEBs) of the LAr calorimeter, which are responsible for amplifying and shaping the signal, using a CR-RC² filter. The shaper is equipped with three overlapping amplification gains (Low, Medium, and High), to accommodate the huge energy range of incoming electrons and photons. Upon a L1 trigger accept, the signal is digitized using the most suitable gain, and transmitted off-detector to the high-level trigger and data-acquisition systems.

The samplings of the shaped signal are converted to the energy deposited in the cell ($E_{\text{raw}}^{\text{cell}}$) using a linear combination of the first four digitized samples in the chosen gain, after subtracting the pedestal contribution (calculated from calibration runs). The linear coefficients for this combination are determined using an optimal filtering algorithm [156, 157], which maximizes energy and timing resolution. This value is then corrected using various calibration constants, derived from calibration runs, simulations, and test beam data, including cell gain, conversion factors from the digital signal counts to energy, and corrections for the different cell responses to ionization and calibration pulses.

5.2.2 LAr layer calibration

The Presampler energy scale, denoted as α_{PS} , is determined using $W \rightarrow \mu\nu$ and $Z \rightarrow \mu\mu$ events in a dedicated low-pileup data sample. Muon final states are ideal for this measurement because they allow α_{PS} to be independent from the influence of material in front of the calorimeter, which has minimal impact on muon energy deposits. The value of α_{PS} , shown in Figure 5.7a, is determined as the ratio of the muon energy peak position in data with that in simulation.

The intercalibration of the EM L1 and L2 layers, represented by the scale factor α_{12} , is determined using Z boson decays in both the di-muon and di-electron channels. For the $Z \rightarrow ee$ channel, α_{12} is derived by analyzing the mean values of the E/p and m_{ee} distributions across bins of the electrons' $|\eta|$ and E_1/E_2 in both data and simulation. Ideally, the ratio of the E/p and m_{ee} estimators between data and simulation should be constant as a function of E_1/E_2 , indicating perfect intercalibration between the L1 and L2 layers. Any observed slope in this ratio indicates a miscalibration, and α_{12} is adjusted to minimize the discrepancy of the data-to-MC ratio with respect to a constant. The final value of α_{12} is obtained by combining the results from both the di-muon and di-electron channels, shown in Figure 5.7b.




Figure 5.7 – (a) Measured presampler energy scale α_{PS} as a function of $|\eta|$, together with their total uncertainties (yellow bands) and their statistical uncertainties (error bars). (b) Relative calibration scale factor α_{12} of the first and second EM Calorimeter layers as a function of $|\eta|$. The results from the $Z \rightarrow \mu\mu$ and $Z \rightarrow ee$ analyses are shown in blue and red, respectively. The final combination combining electron and muon results, are shown as the black solid lines.

Accordion scale calibration

While α_{12} adjusts the relative response between the first two EM layers (L1 and L2), an additional scale factor is required to calibrate their combined response. Given that the third accordion layer (L3) contributes negligibly to the energy measurement for electrons from Z boson decays, the only remaining parameter to be calibrated is the overall energy scale of the accordion calorimeter, denoted as α_{acc} . α_{acc} adjusts the total energy measured in the three accordion layers (L1, L2, L3) in data, by applying the correction $E_{corr}^{acc} = E_{raw}^{acc} / (1 + \alpha_{acc})$. This calibration is performed in bins of the electron's η coordinate. The determination of α_{acc} is based on evaluating global energy scale factors α from $Z \rightarrow ee$ events, as detailed in Section 5.2.3.

Once the global scale factors α are obtained, the accordion scales are given by:

$$\alpha_{acc} = \frac{\alpha}{\frac{\partial \Delta E_{bias}(x)}{\partial x}}, \quad (5.3)$$

where $\partial \Delta E_{bias}(x) / \partial x$ is determined from simulation, and represents the sensitivity of the calibrated energy E to variations x in the raw accordion energy value. The calculation of $\partial \Delta E_{bias}(x) / \partial x$ involves introducing a known shift x into the raw accordion energy, E_{raw}^{acc} , and reapplying the entire calibration process. This produces a new calibrated energy, $E_{bias}(x)$, biased by the factor x . The relative difference between the original calibrated energy E and $E_{bias}(x)$, given by $\Delta E_{bias}(x) = E_{bias}(x) / E - 1$, changes linearly with x . The slope of the line $(x, \Delta E_{bias}(x))$ corresponds to $\partial \Delta E_{bias}(x) / \partial x$, and represents the impact

of varying $E_{\text{raw}}^{\text{acc}}$ on the calibrated energy.

Assuming a simple scenario where the total calibrated energy is approximately the sum of the raw energies from all calorimeter layers, $E = E_{\text{raw}} = E_{\text{raw}}^{\text{acc}} + E_{\text{raw}}^{\text{PS}}$ ², the biased calibrated energy for a variation x in $E_{\text{raw}}^{\text{acc}}$ is:

$$E_{\text{bias}}(x) = E_{\text{raw}}^{\text{acc}}(1 + x) + E_{\text{raw}}^{\text{PS}} \quad (5.4)$$

This yields:

$$\frac{\partial \Delta E_{\text{bias}}(x)}{\partial x} = \frac{E_{\text{raw}}^{\text{acc}}}{E_{\text{raw}}} = f^{\text{acc}}, \quad (5.5)$$

where f^{acc} is the fraction of the total raw energy deposited in the accordion layers. This first-order approximation of the *accordion sensitivity* $\partial \Delta E_{\text{bias}}(x)/\partial x$ as f^{acc} explains its behavior as a function of the electron's $|\eta|$, as shown in Figure 5.8. The sensitivity is less than one because a fraction of the total energy is deposited in the Presampler (for $|\eta| < 1.8$), and it can drop to around 0.5 in the transition region between the barrel and endcap calorimeters ($1.4 < |\eta| < 1.6$).




Figure 5.8 – Sensitivity of the calibrated electron energy to the energy measured in the accordion calorimeter, as a function of $|\eta|$ for electrons from Z boson decays.

As already mentioned, the accordion scales α_{acc} are applied to correct data. After applying the simulation-based calibration to both data and MC simulations, the $Z \rightarrow ee$ measurement is repeated, to extract residual global scales α_{residual} , which should be close

² Assuming that the total calibrated energy corresponds simply to the sum of the raw energies in the Presampler and in the three accordion layers corresponds to ignoring out-of-cluster effects and energy losses upstream, within, and after the EM Calorimeter, which are typically corrected by the simulation-based calibration.

to unity, confirming the convergence of the calibration procedure. The values of α_{acc} for the four data-taking years in Run 2, together with the residual scales α_{residual} , are shown in Figure 5.9.




Figure 5.9 – Energy scale calibration factors α_{acc} and α_{closure} as a function of $|\eta|$ for electrons from Z boson decays. The α_{acc} factors are shown separately for each data-taking year in Run 2, while the α_{closure} factors are integrated over the full Run 2 data. The bottom panel shows the difference between α_{acc} measured in a given data-taking period and the measurements using 2018 data [19].

5.2.3 In-situ measurement of the global energy scale and resolution from $Z \rightarrow ee$ events

After applying the layer calibration scale factors α_{PS} and α_{12} , along with the simulation-based calibration, a potential mismatch between the energy scales in data and simulation may still exist. This discrepancy is corrected using the accordion scale α_{acc} , as described in Section 5.2.2.

The calculation of the accordion scales is based on an *in-situ measurement* of global energy scale and resolution factors, referred to as α_i and c_i , using $Z \rightarrow ee$ events as a candle. The energy scale factors α_i are used to align the energy response in data with the MC

simulation, and are defined as:

$$E_i^{\text{data}} = E_i^{\text{MC}} / (1 + \alpha_i), \quad (5.6)$$

where the i index represents the bin of the electron's η coordinate. Similarly, the energy in MC simulations is smeared using a constant term c_i , such that the electron energy resolution is corrected as:

$$\left(\frac{\sigma(E)}{E} \right)_i^{\text{data}} = \left(\frac{\sigma(E)}{E} \right)_i^{\text{MC}} \oplus c_i, \quad (5.7)$$

where $\oplus$ denotes a quadrature sum, and c_i is also η -dependent.

For a pair of electrons from Z boson decays falling into the η bins i and j , the miscalibration of the di-electron invariant mass can be written in terms of the global scales α_i and the global smearing factors c_i as:

$$\begin{aligned} (m_{ee}^{\text{data}})_{ij} &= (m_{ee}^{\text{MC}})_{ij} / (1 + \alpha_{ij}), \quad \alpha_{ij} \approx \frac{\alpha_i + \alpha_j}{2} \\ \left(\frac{\sigma(m_{ee})}{m_{ee}} \right)_{ij}^{\text{data}} &= \left(\frac{\sigma(m_{ee})}{m_{ee}} \right)_{ij}^{\text{MC}} \oplus c_{ij}, \quad c_{ij} \approx \frac{c_i \oplus c_j}{2} \end{aligned} \quad (5.8)$$

The factors α_{ij} and c_{ij} are determined to best match the di-electron invariant mass distributions in data and simulation, separately for all (i, j) bins in the η coordinate of the two electron candidates. A fit across these bins is then used to extract the global scale factors α_i and smearing factors c_i . The global scale factors α_i are subsequently converted into accordion scales $\alpha_{\text{acc},i}$ (see Section 5.2.2) and applied to correct the accordion energy.

The final step of the calibration procedure consists in repeating the Z -based calibration fit, to test the convergence of the calibration chain. The resulting residual energy scale factors, $\alpha_{\text{residual},i}$, also shown in Figure 5.9, are always below 10^{-4} , except in the transition regions ($1.37 < |\eta| < 1.52$), where Tile scintillators contribute to the electron energy measurement but are not considered in this calibration procedure. This residual non-closure is applied as a final correction to the reconstructed energy in data.

The resolution terms c_i (shown in Figure 5.10) are typically less than 1% in the barrel and between 1% and 2% in the endcaps. This term corrects the energy resolution in simulation to match the slightly worse resolution observed in data.

Figure 5.11 presents the invariant mass distribution of $Z \rightarrow ee$ candidates in data and simulation, after applying the energy scale correction to the data and the resolution correction to the simulation. The uncertainty band reflects the propagation of uncertainties from the $\alpha_{\text{residual},i}$ and c_i factors. At the end of the calibration chain, a fair agreement between data and simulation is observed.

5.2.4 Energy scale uncertainties

Several sources of systematic uncertainty affect the energy scale and resolution for electrons and photons.




Figure 5.10 – Energy resolution calibration factors c as a function of η for electrons from Z boson decays. The bottom panel shows the difference between c measured in a given data-taking period and the measurements using 2018 data [19].




Figure 5.11 – Comparison of the invariant mass distributions of the electron pair in the selected $Z \rightarrow ee$ candidates in data and simulation, after the calibration and resolution corrections are applied. The total number of events in the simulation is normalized to that in data. The ratio between data and simulation is shown in the bottom panel, where the uncertainty band represents the impact of the uncertainties in the calibration and resolution correction factors [19].

- *In-situ $Z \rightarrow ee$ calibration.* The Z-based calibration fit used to extract the energy scale and resolution carries systematic uncertainties related to the measurement itself (for instance, from the event selection and from the choice of the fitting range in the m_{ee} spectrum). The impact is generally below 0.05% across most η regions but can reach 0.5% in the transition region.
- *LAr layer calibration.* The measurements of α_{pS} and α_{12} involve systematic uncertainties, such as those related to the selection of candidate electrons or muons, material uncertainties, and the choice of the estimators for $\langle E_i \rangle$, $\langle m_{ee} \rangle$, and $\langle E/p \rangle$. The uncertainty from α_{12} can impact the energy scale by up to 0.2%, while the effect of the presampler scale α_{pS} is generally below 0.1%.
- *Non-linearity of LAr readout electronics.* Energy reconstruction in the EM Calorimeter cells involves a linear conversion from ADC counts using a calibration factor derived from calibration runs. The conversion from ADC counts and injected current is not perfectly linear: the residual non-linearity is evaluated cell by cell, and is used to correct raw energies. The impact of the ADC non-linearity on the electron or photon energy is estimated by re-evaluating the non-linearity corrections for all the cells within the cluster, and computing the relative difference with the original cluster energy. This difference is reparametrized to be zero for electrons at $E_T = 40$ GeV, and affects the energy scale by up to 0.4% for low E_T electrons and photons and below 0.2% at higher E_T .
- *Intercalibration of High to Medium Gain.* The global scale is calibrated using electrons from $Z \rightarrow ee$ decays, with an average E_T of around 40 GeV, that are typically read out with the High gain. The transition to the Medium gain readout occurs around 60 GeV. The intercalibration between High and Medium gains is assumed perfect in simulation, and is measured in data using special data-taking runs where the Medium gain threshold in the L1 or L2 layers was lowered, such that the highest energy cells within electron clusters from $Z \rightarrow ee$ decays are read out with the Medium gain. Comparing $Z \rightarrow ee$ events in this special dataset with those reconstructed with nominal gain settings allowed to evaluate an uncertainty on the High to Medium gain intercalibration in the EM L1 or L2 layers: their impact is up to 0.1% in the barrel and 0.4% in the endcaps. A corresponding, independent, uncertainty is defined for the transition between the Medium and Low gain readouts, whose impact is non-negligible only at very high energies ($E_T > 400$ GeV). By isolating the High to Medium gain intercalibration uncertainty from the Medium to Low gain intercalibration uncertainty, the calibration fit described in the following Section constrains the former more effectively, since it is more relevant in the E_T range close to the average of $Z \rightarrow ee$ electrons.
- *Material modelling.* The electron and photon energy is affected by passive material in front of the EM Calorimeter. This is taken into account by the simulation-based calibration, which is applied to both data and MC simulations. However, differences between the simulation of the detector model and the actual detector lead to discrepancies in energy response between data and MC. For electrons with $E_T \approx 40$ GeV, such discrepancies are absorbed through Z-based calibration, but biases re-

main depending on particle type and energy. The impact of passive material uncertainties, parameterized as a function of particle type, E_T , and $|\eta|$, typically affects low E_T electrons and unconverted photons, reaching an impact of 0.1%-0.2%.

- *Photon conversions mismodelling.* The simulation-based calibration is optimized separately for unconverted photons and converted photons. Differences between data and simulation in classification rates of converted and unconverted photons generate biases in the photon energy scale. The conversion rates are measured using a sample of photons from radiative Z events. The uncertainty in the photon energy scale is evaluated by reweighting the conversion fractions (corresponding to the fraction of true photon conversions, photon conversions reconstruction efficiency, and fraction of reconstructed fake photon conversions) in simulated photons in single-particle samples to those observed in radiative Z samples in data and MC. The impact of the bias from the photon conversion fractions is assumed correlated among reconstructed converted and unconverted photons, across the full η range. However, the magnitude of the uncertainty varies between converted and unconverted photons and with their $|\eta|$ and E_T . For photons with $E_T = 60$ GeV, the uncertainty is about 0.02%-0.13% for unconverted photons and 0.01%-0.12% for converted photons, depending on the region. At lower E_T , these uncertainties increase significantly: for unconverted photons, they can be as large as 0.67%, while for converted photons they reach up to 1.31%.

Propagation of systematic uncertainties to the energy scale as a function of E_T

The Z-based calibration fit described in Section 5.2.3 is used to provide global calibration scale factors $\alpha_{acc,i}$ and $\alpha_{residual,i}$ as a function of the electrons' η coordinate, and fixes the energy scale and its uncertainty for electrons with E_T near the average produced in Z decays ($E_T \approx 40$ GeV). The impact of systematic uncertainties as a function of E_T for electrons and photons is propagated by evaluating the sensitivity of the total calibrated energy E to variations in the corresponding sources of uncertainty (for instance, the energy deposited in each EM layer E_{raw}^{Li}) as a function of E_T and particle type. The procedure for evaluating these sensitivities is similar to that used for the accordion sensitivity described in Section 5.2.2.

Consider, for instance, the systematic uncertainty affecting the energy scale of the Presampler. These uncertainties influence the modeling of E_{raw}^{PS} , the raw energy deposited in the presampler. Using simulated single-electron and single-photon events, E_{raw}^{PS} is artificially shifted by a known factor x , thus modifying it to $E_{raw}^{PS} \cdot (1 + x)$. The total calibrated energy is then recalculated by reapplying the full calibration procedure, producing a biased result due to the introduced variation x , referred to as $E_{bias}(x)$. The relative difference between the original calibrated energy E and $E_{bias}(x)$ is expressed as $\Delta E_{bias}(x) = E_{bias}(x)/E - 1$, and changes linearly with x . The slope of the line $(x, \Delta E_{bias}(x))$, denoted as $\partial \Delta E_{bias}(x) / \partial x$, quantifies the impact of varying E_{raw}^{PS} on the calibrated energy. This slope is called *presampler sensitivity* and is evaluated as a function of the particle's E_T and $|\eta|$. As an example, the presampler sensitivity curve for electrons and photons with two pseudorapidity values as a function of E_T is illustrated in Figure 5.12.




Figure 5.12 – Presampler sensitivity for electrons, converted photons, and converted photons as a function of the particle E_T for two pseudorapidity bins: (a) $0 < |\eta| < 0.1$ and (b) $1 < |\eta| < 1.1$.

The systematic uncertainty on the residual energy scale derived from the uncertainty in α_{PS} is determined using the in-situ Z-based measurement for electrons with E_T corresponding to the average for electrons from $Z \rightarrow ee$ decays. For electrons with lower or higher E_T , as well as for converted and unconverted photons, the impact is propagated using the presampler sensitivity curves as a function of E_T .

The same procedure is applied to evaluate the sensitivity of the total calibrated energy to variations in all the other ingredients of the calibration chain (such as the energy deposited in the EM layers L1 and L2 or in the full accordion, or the energy in the EM layers L1 or L2 readout with different gains).

This procedure allows to estimate the impact of the systematic uncertainty on the electron and photon energy scale across the full E_T range. The effects of the most significant uncertainties, already discussed in Section 5.2.4, are shown as a function of electron or photon E_T for two different pseudorapidity values in Figure 5.13.




Figure 5.13 – Relative energy scale calibration uncertainty for electrons, unconverted photons, and converted photons, as a function of E_T for $|\eta| = 0.3$ (left) and $|\eta| = 2.1$ (right) [19].

Constraining the energy scale as a function of E_T

As a first-order approximation, the energy scales derived from the in situ $Z \rightarrow ee$ calibration, as detailed in Section 5.2.2, are assumed to be valid across the full E_T spectrum. However, the E_T dependence of the energy scale corrections α_i is tested, by repeating the Z -based fit in bins of $|\eta|$ and E_T of the electron candidates. The calibrated energy can be

expressed as:

$$E_i^{\text{data}} = \frac{E_i^{\text{MC}}}{\left[(1 + \alpha_i)(1 + \alpha'_j)\right]}, \quad (5.9)$$

where α_i represents the energy scale factors determined from the initial Z-based calibration fit in bins of the electron's η (labeled by i), while α'_j describes a residual dependence of the energy scale on the electron's E_T . The index j categorises the electron's kinematics in a two-dimensional ($|\eta|, E_T$) space, with bin boundaries defined as 0, 0.6, 1.0, 1.37, 1.55, 1.82, and 2.47 in $|\eta|$ and 27, 33, 38, 44, 50, 62, 100, and $\sqrt{s}/2$ in E_T (in GeV). For this analysis, $Z \rightarrow ee$ events are classified into the j^{th} category if either of the two final-state electrons falls into the j^{th} ($|\eta|, E_T$) bin. A fit to the di-electron invariant mass distributions is performed to determine the values of α'_j that best align data and MC events. The resulting α'_j values are shown in Figure 5.14. These results indicate a significant residual E_T dependence of the energy scales, measurable with a precision better than 0.03% (rising to 0.3% in the transition region). The precision in measuring the α'_j factors surpasses that of the η -dependent scales α_i described earlier. Hence, this linearity fit, which quantifies the energy dependence of the energy scales, can provide an additional constraint on the systematic uncertainties affecting the energy scales across the entire E_T spectrum. While the behaviour of the systematic uncertainties as a function of E_T can be estimated from simulation (by evaluating the sensitivity of the reconstructed energy to each source of uncertainty), the linearity fit offers actual data points across various E_T regions (not just at $E_T \approx 40$ GeV, typical from electrons from Z decay) that can be used to constrain these uncertainties.

The impact of the linearity fit on the systematic uncertainties affecting the energy scales is illustrated in Figure 5.14, where the outer uncertainty band (pre-fit) is compared with the inner band (post-fit). The linearity fit reduces the systematic uncertainty on the energy scale by up to a factor of two for $E_T < 50$ GeV and up to a factor of three for $E_T \approx 150$ GeV.

The impact of the linearity fit on the electron energy scale uncertainties is further detailed in Figure 5.15a as a function of E_T and $|\eta|$. For electrons with $E_T \approx 40$ GeV, the precision remains largely unchanged, as these particles, typical of on-shell Z boson decays, are calibrated through the energy scale analysis that determines the accordion scales and the η -dependent residual scales. However, energy scale uncertainties for electrons with $E_T = 10$ GeV or $E_T = 1$ TeV are typically reduced by 30%-50%, varying from 0.2%-0.3% for $|\eta| < 1$ and $|\eta| > 1.8$, to between 0.5% and 1% for $1 < |\eta| < 1.8$. The impact of this analysis on photon calibration uncertainties is shown in Figure 5.15b for converted and unconverted photons at $E_T = 60$ GeV, which is the average for photons from Higgs boson decays. Uncertainties for converted photons, which are experimentally similar to electrons, are only marginally reduced by the linearity fit. However, for unconverted photons, the overall energy calibration uncertainty is typically reduced by 30% in the barrel and by up to a factor of two in the endcaps.




Figure 5.14 – Comparison of the measured values of $\alpha' j$ from $Z \rightarrow ee$ events (black dots) with the results extracted from the linearity fit, including their corresponding uncertainty band. The pre-fit value of α' (namely, $\alpha' = 0$ before the linearity constraint) and its associated uncertainty are also displayed. Additionally, the measured values of α' from $J/\psi \rightarrow ee$ events, used as a cross-check for the linearity model in the low E_T region, are shown (open dots) [19].




Figure 5.15 – Total relative systematic uncertainty in the energy scale as a function of $|\eta|$ for electrons with $E_T = 10$ GeV, 40 GeV or 1 TeV (left) photons with $E_T = 60$ GeV, after the constraints from the linearity fit. The bottom panels show the ratio of the post-fit to pre-fit uncertainties [19].

5.3 Photon identification

The photon identification criteria are optimized to efficiently select prompt and isolated photons while suppressing background from hadronic jets which are incorrectly classified as photons.

Three working points for photon identification - *Tight*, *Medium*, and *Loose* [158] - are defined by applying independent selection criteria on discriminating calorimetric variables, termed *shower shapes*, described in Table 5.1. Photons are expected to produce compact electromagnetic showers in the EM Calorimeter, giving rise to energy deposits characterized by a distinct maximum in its core, whereas showers initiated by jets tend to be deeper, wider, and the energy deposits are typically more spread out. Therefore, variables sensitive to the width of the shower in the η and ϕ directions, as well as those describing the energy leakage into the hadronic calorimeter, provide significant discriminative power between photon-induced and jet-induced energy deposits. In addition, some shower shape variables are specifically designed to take advantage of the fine segmentation EM L1 layer. These variables can distinguish cases where a neutral meson, such as a π^0 , carries most of the jet's p_T and decays into two photons, resulting in an electromagnetic shower with two distinct maxima that can be resolved within the L1 layer. The distributions of some of these variables for unconverted photons radiated from leptons in $Z \rightarrow \ell\ell\gamma$ decays are shown in Figure 5.16, for both data and MC samples. The simulation does not perfectly describe the shower shapes observed in data, particularly those related to the lateral development of the shower. To address this, dedicated corrections derived from data-to-MC comparisons are used to correct the MC

distributions, aligning them with the distributions observed in data [159].




Figure 5.16 – Examples of two shower shape variables, R_η (left) and $w_{s,3}$ (right), for unconverted photons from radiative Z decays. The black dots represent the data, which are compared to simulated events both before and after applying corrections to the shower shapes. The uncorrected MC distributions are shown in red, while the corrected distributions, adjusted by shifts to better align with the data, are displayed in blue [159].

The Loose identification requirement exploits variables such as R_{had} , R_{had_1} , R_η , and w_{η_2} and was used as for triggering events with photons during the 2015 and 2016 data-taking periods of Run 2. The Medium working point adds a loose cut on E_{ratio} and became the main photon trigger selection in early 2017 to maintain an acceptable trigger rate. The Tight working point is applied to a subset of candidate photons passing the Medium selection, introducing additional requirements on R_η and R_ϕ (describing the shower spread in the η and ϕ directions in the L2 layer), on $w_{s,3}$ and $w_{\text{tot}_{s1}}$ (describing the lateral shower width in the L1 layer), on f_{side} (the fraction of energy in the shower’s crown relative to its core), on ΔE (sensitive to the presence of double maxima within a single shower), and on f_1 (corresponding to the fraction of the photon energy deposited in the L1 layer).

The selection criteria for the three working points are optimized separately in different pseudorapidity bins to account for changes in calorimeter geometry and upstream materials with $|\eta|$, which affect the shower shape variables. Moreover, the selection optimization for the Tight working point is E_T -dependent. Compared to the previous E_T -independent optimization, this strategy enhances the Tight selection efficiency by up to 20% for photons with $E_T < 30$ GeV and simultaneously improves background rejection at high E_T [158]. Additionally, the 2-T solenoidal magnetic field imparts a deviation from a straight trajectory to the electron-positron pair created during a photon conversion in the ID, resulting in wider electromagnetic showers along the ϕ direction compared to showers from unconverted photons. The Tight identification criteria exploit this feature by distinguishing between unconverted and converted photons.

Category	Description	Name	Loose	Medium	Tight
Hadronic leakage	Ratio of E_T in the first sampling of the hadronic calorimeter to E_T of the EM cluster (used over the range $ \eta < 0.8$ or $ \eta > 1.37$)	R_{had_1}	✓	✓	✓
	Ratio of E_T in all the hadronic calorimeter samplings to E_T of the EM cluster (used over the range $0.8 < \eta < 1.37$)	R_{had}	✓	✓	✓
EM L1 layer	Ratio in η of cell energies in 3×7 cells over 7×7	R_η	✓	✓	✓
	Lateral width of the shower along η	w_{η_2}	✓	✓	✓
	Ratio in ϕ of cell energies in 3×3 cells over 3×7	R_ϕ			✓
EM L2 layer	Shower width calculated from three strips around the strip with maximum energy deposit	$w_{s,3}$			✓
	Total lateral shower width	$w_{\text{tot}_{s1}}$			✓
	Energy outside the core of the three central strips but within seven strips divided by energy within the three central strips	f_{side}			✓
	Difference between the energy associated with the second maximum in the strip layer and the energy reconstructed in the strip with the minimum value found between the first and second maxima	ΔE			✓
	Ratio of the energy difference associated with the largest and second largest energy deposits to the sum of these energies	E_{ratio}		✓	✓

Table 5.1 – Photon shower shape variables and identification criteria definitions [160].

The Loose photon identification is optimized to achieve a 99% efficiency for photons with $E_T > 40$ GeV, with a corresponding background rejection factor of about 1000 [161, 162]. The Tight selection, in contrast, achieves 85% efficiency for photons with $E_T > 40$ GeV, increasing to around 92% for unconverted photons and 95% for converted photons when $E_T > 100$ GeV, as shown in Figures 5.17 and 5.18. The background rejection factor for the Tight selection is approximately 5000 [161, 162].

Photon identification efficiency is measured in data using three different methods: the first involves selecting events with single isolated photons from data collected using single photon triggers with large prescale factors, and estimating the background contamination from jets using control regions (*matrix method*); the second utilizes photons radiated from leptons in $Z \rightarrow \ell\ell\gamma$ decays (radiative Z method); and the third employs electrons from $Z \rightarrow ee$ decays, applying a technique to transform electron shower shapes to resemble those of photons (*electron extrapolation method*). Figures 5.17 and 5.18 present the efficiency measurements for the Tight identification requirement, performed with these three methods using Run 2 data [163]. These methods are complementary, covering different ranges of photon E_T : up to 100 GeV for the radiative Z method, from 25 to 200 GeV for the electron extrapolation method (corresponding to the E_T spectrum of electrons from Z boson decay), and up to very high energies (up to 1 TeV) for the matrix method. Data-to-MC scale factors are derived from these methods, with their uncertainties considered as a systematic uncertainty in photon identification. The photon identification scale factors, visible in the bottom panels in Figure 5.17 and Figure 5.18, differ from unity by up to 3%-4% at $E_T = 10$ GeV and by at most 1%-2% at $E_T > 40$ GeV, with uncertainty decreasing from approximately 10% at $E_T = 10$ GeV to less than 1%-2% at higher E_T . To correct the photon identification efficiency in simulation to match data, the scale factors from the three methods are combined using a weighted average in each $|\eta|$ and E_T bin of the photon.




Figure 5.17 – Tight photon identification efficiency for unconverted photons, along with the data-to-MC efficiency ratio, as a function of E_T in four distinct $|\eta|$ regions. Efficiencies are derived using three different measurement methods. The bottom panels also provide the combined scale factor, calculated as a weighted average of the individual scale factors, together with its total uncertainty (represented by the shaded area).




Figure 5.18 – Tight photon identification efficiency for converted photons, along with the data-to-MC efficiency ratio, as a function of E_T in four distinct $|\eta|$ regions. Efficiencies are derived using three different measurement methods. The bottom panels also provide the combined scale factor, calculated as a weighted average of the individual scale factors, together with its total uncertainty (represented by the shaded area).

Working point	Calorimeter isolation	Track isolation
Loose	$E_T^{\text{cone}20} < 0.065 \times E_T$	$p_T^{\text{cone}20} / E_T < 0.05$
Tight	$E_T^{\text{cone}40} < 0.022 \times E_T + 2.45 \text{ GeV}$	$p_T^{\text{cone}20} / E_T < 0.05$
TightCaloOnly	$E_T^{\text{cone}40} < 0.022 \times E_T + 2.45 \text{ GeV}$	-

Table 5.2 – Definition of the photon isolation working points [147].

5.4 Photon isolation

A typical photon from a $H \rightarrow \gamma\gamma$ decay is expected to be surrounded by a region of minimal hadronic activity: isolation requirements on candidate photons are used to further reduce the background from fake photons.

The hadronic activity surrounding the photon candidate is quantified from nearby energy deposits in the calorimeters or from close-by tracks reconstructed in the ID. The idea is to measure the signals detected within a fixed distance from the photon candidate and require that all the surrounding activity is limited to a fraction of the energy deposited by the candidate photon itself. Two classes of variables are used to evaluate the photon isolation:

- **Calorimetric isolation.** $E_T^{\text{cone}0.2}$ is defined as the scalar sum of the transverse energy for all the topo-clusters whose barycentre falls within a cone of radius $\Delta R = \sqrt{(\Delta\eta)^2 + (\Delta\phi)^2} = 0.2$ centered around the candidate photon supercluster barycentre. The contribution of the photon energy itself is removed by subtracting the energy deposited in a core rectangular region of size $\Delta\eta \times \Delta\phi = 0.125 \times 0.175$. The isolation energy is also corrected to account for photon energy leakage outside the core window and contributions from pile-up and underlying events are removed. Other calorimetric isolation variables are similarly defined using cones of different radii, such as $E_T^{\text{cone}0.4}$.
- **Track isolation.** Similarly, the track-based isolation variable $p_T^{\text{cone}0.2}$ is computed as the scalar sum of the transverse momentum for all the selected tracks falling within a cone of radius $\Delta R = 0.2$ centered around the photon cluster direction. Tracks matched to the photon conversion are excluded, in case of a converted photon candidate. Selected tracks included in the sum are required to have $p_T > 1 \text{ GeV}$, $|\eta| < 2.5$, and to be compatible with originating from the primary vertex.

Three photon isolation operating points are defined based on specific requirements for calorimeter and track isolation variables, as outlined in Table 5.2. The photon isolation efficiency is evaluated using two methods: radiative Z boson decays and inclusive-photon production data, similar to the approach used for photon identification. The results from these methods are combined to produce a single set of scale factors for each working point, data-taking year, and photon conversion status. Figure 5.19 shows the evolution of isolation efficiency from the $Z \rightarrow \ell\ell\gamma$ measurement as a function of E_T , η , and conversion status, along with the data-to-MC efficiency ratio.




Figure 5.19 – Efficiency of the photon isolation working points, evaluated using photons from $Z \rightarrow \ell\ell\gamma$ decays, for unconverted (left) and converted (right) photons as a function of photon η (top) and E_T (bottom). The lower panel shows the ratio of the efficiencies measured in data and in simulation. The total uncertainty is shown, including the statistical and systematic components [147].

5.5 Electrons

As explained in Section 5.1, electron reconstruction is performed by matching a supercluster in the EM calorimeter to a track in the Inner Detector compatible with originating from the primary vertex. The four-momentum of the electron candidate is determined by combining information from the EM calorimeter and the tracker: the energy is calculated from the electron supercluster, while the η and ϕ coordinates are obtained from the matched track. In cases where the track has hits only in the TRT, both the energy and angular coordinates are derived exclusively from the calorimeter.

To distinguish prompt, isolated electrons from those originating from hadron decays, identification criteria are optimized using a multivariate likelihood function. Three identification operating points - *Tight*, *Medium*, and *Loose* - are defined based on whether the likelihood assigns a more electron-like or background-like probability to the candidate. Electron identification relies on both calorimetric shower shape variables and track properties [163, 164].

To further distinguish signal electrons from background, electrons are required to be well isolated from surrounding hadronic activity. Electron isolation is quantified using both calorimetric and track-based variables, for instance, $E_T^{\text{cone0.2}}$ and $E_T^{\text{varcone0.2}}$. The calorimetric isolation, $E_T^{\text{cone0.2}}$, is calculated as the scalar sum of the energy from all topoclusters within a cone of radius $\Delta R = 0.2$ around the electron supercluster barycentre, corrected for the electron's own energy and pile-up contributions. The track isolation, $p_T^{\text{varcone0.2}}$, sums the transverse momenta of all nearby tracks originating from the primary vertex, excluding the electron's track. A variable-sized cone, dependent on the electron's transverse momentum, is used to calculate $p_T^{\text{varcone0.2}}$: this allows to take into account decays of heavy particles (such as the top quark), where high- p_T electrons are produced in close proximity to other decay products.

Further details on electron identification and isolation algorithms can be found in [147, 163, 164].

5.6 Muons

Muons are reconstructed independently in the Inner Detector, Muon Spectrometer, and calorimeters. The information from each sub-detector is combined to form a complete muon track. Depending on which subsystem contributes most to the reconstruction, four types of muon candidates are defined:

- A *combined muon* is created when high-quality matching tracks are reconstructed in both the ID and MS, with hits from both sub-detectors refitted into a single muon track.
- An *inside-out combined muon* is formed when a high-quality track from the ID is extrapolated to the MS and matched to at least three loosely aligned MS hits. The ID track, the energy loss in the calorimeter, and the MS hits are then refitted into a single unified muon track.

- A *segmented muon* is formed when a track in the ID extrapolates to a local track segment in the MS.
- A *calorimeter-tagged muon* is recovered from energy deposits in the calorimeters that are consistent with a minimum ionizing particle. This extends muon reconstruction acceptance to the $|\eta| < 0.1$ region, where the MS is only partially instrumented to accommodate cabling and services for the calorimeters and ID.
- An *extrapolated Muon* is reconstructed only in the MS, with a loose association to the primary vertex.

To minimize contamination from muons originating in hadron decays, identification criteria are applied to select only prompt muons. This identification is based on variables that assess the quality of the track fit, as secondary muons from charged meson decays typically exhibit poorer track quality compared to prompt muons. Three standard identification working points are defined—*Tight*, *Medium*, and *Loose*—which offer increasing efficiency but decreasing purity with respect to background contamination. Additionally, two specialized working points, *Low- p_T* and *High- p_T* , are introduced to recover efficiency for low- p_T muons (typically $3 < p_T < 5$ GeV) and high- p_T muons (above 1 TeV) with improved momentum resolution that meet the Medium selection criteria. The reconstruction and identification of muons in the ATLAS detector is further discussed in [165].

As discussed in Chapter 5, one of the key physics objects in the Higgs boson pair production searches presented in this thesis is the jets arising from the hadronization of bottom quarks (*b*-jets). This chapter details the reconstruction algorithms and identification techniques used for *b*-jets during the Run 2 data-taking period, along with an overview of their performance. A summary of the measurement of *b*-jet identification efficiency for a new algorithm based on Graph Neural Networks [20], to which I made personal contributions, is also presented. The reconstruction and identification of the missing transverse momentum is also briefly addressed.

6.1 Jets

Quarks and gluons are produced abundantly in *pp* collisions at the LHC, either through underlying event activity or directly from the hard-scattering process. However, due to *color confinement*, these particles cannot be detected individually in the final state. Instead, the fragmentation and hadronization of a colored parton lead to the formation of a collimated spray of hadrons traveling in the direction of the initial parton, referred to as a *jet* [166].

Jets originating from the fragmentation and hadronization of bottom quarks are referred to as *b*-jets. An efficient identification of *b*-jets againsts background from jets originating from other quark flavors or gluons, as well as precise calibration of the *b*-jets' energy, is critical for Higgs boson pair production searches. Indeed, the most sensitive channels to HH production (including the $b\bar{b}\gamma\gamma$ final state explored in this thesis), involve at least one Higgs boson decaying into two bottom quarks (see Section 2.1.4).

This section covers the reconstruction of hadronic jets and their standard energy calibration, followed by a discussion on the *b*-jet identification and on the specific corrections applied to *b*-jet energy measurements.

6.1.1 Particle flow jets

The collimated spray of hadrons forming a jet consists of both charged and neutral particles. Charged hadrons leave tracks in the Inner Detector, while both charged and neutral

hadrons interact with the ATLAS calorimeters, inducing hadronic showers and depositing their energy in the calorimetric system. Jet reconstruction in the ATLAS detector follows the *particle flow* strategy [167, 168]: the main idea is to surpass the single physics object reconstruction involving the specifically-designed sub-detector (i.e. jet reconstruction using energy deposits in the calorimeters, without attempting to resolve the individual particles) and to pursue an holistic approach instead. Information from both the calorimeters and the ID is combined, and energy clusters and tracks are used to reconstruct all individual stable particles in the event (e.g., electrons, photons, pions). The reconstructed *particle flow objects*, which ideally originated from the hadronization of the same parton, form the basis for jet clustering and reconstruction.

The main advantage of the particle flow approach consists in the ability to fully leverage the strengths of both the calorimeters and the ID and to use the two sub-systems to compensate each other's short-comings:

- For low-energy charged particles, the momentum resolution provided by the tracker is significantly better than the energy resolution provided by the calorimeters. Conversely, the calorimeters offer a superior energy resolution at higher energies, have a larger pseudorapidity acceptance with respect to the tracker, and are able to detect neutral particles.
- The ID provides a better angular resolution for measuring the charged particles' direction, catches low-energy tracks which may be swept out of the jet cone by the solenoidal magnetic field before reaching the calorimeters, and allows to match a track with the primary interaction vertex. The latter feature is crucial for pileup rejection when clustering the hadronic jet.

The main downside of the particle flow strategy is the need to subtract calorimeter signals for particles already reconstructed using tracker information (e.g., low-energy charged hadrons) to prevent double-counting their energy. The ability to correctly identify and subtract all of the energy from a single particle without affecting the energy deposits of any other particle is the key challenge of the particle flow algorithm.

Particle flow algorithm

Topoclusters, built from energy deposits in the ATLAS calorimeters using the “4-2-0” scheme described in Section 5.1, along with tracks reconstructed in the ID, serve as inputs for the particle flow algorithm.

The first step of the particle flow algorithm is to select tracks for reconstructing (charged) particle flow objects. Reconstructed tracks must meet several criteria: they are required to have $p_T > 0.5$ GeV, $|\eta| < 2.5$, at least nine hits in the silicon layers of the ID and no missed hits in the Pixel Detector. No requirement on the association with the hard-scattering vertex is placed. In addition, since energetic particles are typically found close to the jet core, their tracks in the ID are often surrounded by nearby activity, making it challenging to subtract the corresponding calorimeter energy to avoid double-counting their contribution. To address this, tracks with $p_T > 100$ GeV are excluded from the selection, and calorimetric information is used instead for reconstructing high- p_T particles. Furthermore, tracks associated with medium-quality electrons or muons are masked, as

the particle flow algorithm is optimized to reconstruct hadronic showers.

Selected tracks are matched to reconstructed topoclusters by requiring the ratio between the topocluster energy and the track momentum to be $E^{\text{cluster}}/p^{\text{track}} > 0.1$. All matched tracks are then ranked based on the angular distance $\Delta R'$ between the topocluster's barycentre and the track's extrapolation to the EM L2 layer. This quantity is defined as:

$$\Delta R' = \sqrt{\left(\frac{\Delta\eta}{\sigma_\eta}\right)^2 + \left(\frac{\Delta\phi}{\sigma_\phi}\right)^2}, \quad (6.1)$$

where σ_η and σ_ϕ represent the lateral widths of the topocluster along the η and ϕ directions, computed as the standard deviation of the cells' positions relative to the topocluster barycentre. The best-matched track is the one with the lowest $\Delta R'$.

Once a track with momentum p^{track} is matched to a topocluster, the expected energy $\langle E_{\text{dep}} \rangle$ deposited in the calorimeters by the particle associated with the track is calculated as:

$$\langle E_{\text{dep}} \rangle = p^{\text{track}} \cdot \langle E_{\text{ref}}^{\text{cluster}} / p_{\text{ref}}^{\text{track}} \rangle \quad (6.2)$$

where $\langle E_{\text{ref}}^{\text{cluster}} / p_{\text{ref}}^{\text{track}} \rangle$ is the average energy-to-momentum ratio for topoclusters within a cone of radius $\Delta R = 0.4$ around a track, derived from simulated single-pion samples without pileup. The spread of the expected energy deposit, $\sigma(E_{\text{dep}})$, is determined from the standard deviation of the $E_{\text{ref}}^{\text{cluster}} / p_{\text{ref}}^{\text{track}}$ distribution in these single pion samples.

The expected deposited energy $\langle E_{\text{dep}} \rangle$ is then compared with the measured energy in the matched topocluster, E^{cluster} . If $\langle E_{\text{dep}} \rangle - E^{\text{cluster}}$ is positive and larger than the spread $\sigma(E_{\text{dep}})$, it is assumed that the particle that created the track deposited its energy in more than one topoclusters: additional nearby topoclusters (within a cone of radius $\Delta R = 0.2$ from the extrapolation of the track to the EM L2 layer) are therefore considered to be matched to the track.

Once a set of topoclusters corresponding to the track has been selected, the expected energy deposited by the particle that created the track is subtracted, in order to avoid double-counting the energy left in the tracker and in the calorimeter by the same particle. This subtraction is only performed for tracks and matched topoclusters with

$$\frac{E^{\text{cluster}} - \langle E_{\text{dep}} \rangle}{\sigma(E_{\text{dep}})} < 33.2 \times \log_{10} \left(40 \text{ GeV} / p_{\text{T}}^{\text{track}} \right) \quad (6.3)$$

This requirement ensures that calorimetric energy subtraction is applied for tracks with p_{T} below 40 GeV, except in cases where the calorimeter activity (indicated by E^{cluster}) is particularly high, such as in dense environments where subtraction accuracy decreases.

Jet clustering and reconstruction

Particle flow objects, defined using the topoclusters that survive the subtraction process and tracks associated with the hard-scattering vertex, are used as input for jet reconstruction. ATLAS employs the *anti- k_t algorithm* [169] for jet clustering, which ensures

that the resulting jets are *infra-red and collinear safe*, namely, the number of jets remains stable despite the emission of soft or collinear particles. This property allows to have finite predictions of cross sections for processes involving jets across all perturbative orders.

The clustering of particle flow objects into jets is governed by two distance parameters, d_i and d_{ij} , defined as:

$$d_i = \frac{1}{p_{T_i}^2}, \quad d_{ij} = \min \left(\frac{1}{p_{T_i}^2}, \frac{1}{p_{T_j}^2} \right) \frac{\Delta_{ij}^2}{R^2} \quad (6.4)$$

where i and j are indices of particle flow objects, $\Delta_{ij} = \sqrt{(\eta_i - \eta_j)^2 + (\phi_i - \phi_j)^2}$ represents their angular separation, and R is a fixed distance parameter used in jet reconstruction (for the $HH \rightarrow b\bar{b}\gamma\gamma$ and $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ analyses discussed in this thesis, the selected jets are reconstructed using $R = 0.4$). The algorithm starts by calculating the distances between all pairs of particles. If $d_{ij} < d_i$, objects i and j are merged into a new object, which replaces them in the input set. If $d_{ij} > d_i$, object i is identified as a jet and removed from the input pool. This process is repeated, with the distances d_{ij} and d_i being recalculated after each step, until no further combinations are possible.

6.1.2 Energy calibration

Jets are initially calibrated at the electromagnetic scale, which is measured to ensure that the deposited energy of electrons and photons matches their true energy. However, the response of the calorimeter to incoming jets is different to the one for electrons and photons: therefore, a recalibration is needed to adjust the jet energy scale to that of *truth-level jets*, namely, jets reconstructed at the particle level using stable final-state particles from MC simulations, without incorporating detector effects.

The complete jet energy scale (JES) calibration procedure [168], illustrated in Figure 6.1, involves several correction stages that adjust the four-momentum of the jets, scaling the jet p_T , energy, and mass.

The process begins with pileup corrections, to remove the excess energy from pileup activity, due to additional pp interactions occurring within the same (in-time) or nearby (out-of-time) bunch crossings. These corrections include:

1. A first correction based on the *jet area* (which measures the sensitivity of the jet to pileup activity) and on the median transverse momentum density (termed ρ) evaluated as a function of the (η, ϕ) region in the detector, which estimates the pileup activity;
2. A residual correction, parametrized using MC simulations, which adjusts for the residual dependence of the jet's p_T on the average number of pp collisions per bunch crossing (μ) and the number of primary vertices in the event (N_{PV}).

Next, the absolute JES calibration, or MCJES, aligns the reconstructed jet energy and direction with truth jets from di-jet MC events. This step addresses factors such as non-




Figure 6.1 – Flowchart of the jet energy scale calibration procedure [168].

compensating calorimeter response, energy losses in passive material, energy deposited outside of the cone of radius R used for jet clustering (*out-of-cone effects*), and biases in jet η reconstruction, which arise due to transitions between different calorimeter technologies and changes in calorimeter granularity.

Despite the MCJES calibration, variations in the energy response within the same kinematic region (defined by the jet's η and p_T) may still occur. These variations result from differences in flavor composition, energy distribution of constituent particles, their transverse spread, and fluctuations in the hadronic shower development in the calorimeter. The global sequential calibration (GSC) applies a series of multiplicative corrections to minimize these fluctuations and enhance jet resolution without changing the average jet energy response. Corrections are sequentially applied to the jet's four-momentum to remove dependencies on tracking, calorimeter, and muon spectrometer observables.

Finally, the η -intercalibration and in-situ calibration address the residual discrepancies between data and MC simulations, due to imperfect detector modeling or inaccuracies in the simulation of the physics processes. These discrepancies are quantified by comparing jets against well-measured reference objects [168]. For each in-situ estimate, corrections are derived by numerically inverting the ratio $R_{\text{data}}^{\text{in-situ}} / R_{\text{MC}}^{\text{in-situ}}$, where $R^{\text{in-situ}}$ represents the jet energy response, defined as the average ratio of jet p_T (or η , for the η intercalibration) to reference object p_T (or η), binned according to the reference object's p_T (or η). The data-to-MC ratio and associated systematic uncertainties from the Z -jet, γ -jet, and multi-jet calibrations are combined across overlapping p_T regions, and the resulting correction is applied to data, to match the jet energy scale to the one observed in simulations.

The full JES calibration introduces 125 systematic uncertainty terms from the various calibration stages. These uncertainties arise from event selection choices during in-situ analyses, event topology, MC mismodeling, and statistical limitations. Figure 6.2 provides an overview of these uncertainties as a function of p_T and η . The largest contributions are from pileup subtraction for low- p_T jets (reaching up to 4%) and jet flavor response (averaging 1.5% in the medium- p_T range), while in-situ calibration uncertain-

ties dominate at higher p_T .




Figure 6.2 – The relative impact of systematic uncertainties on the jet energy scale for anti- k_t $R = 0.4$ jets is illustrated as a function of η for jets with a fixed p_T of 60 GeV (left) and as a function of p_T for jets at $\eta = 0$ (right). The total uncertainty, depicted as the shaded region topped by a solid black line, is determined by summing all uncertainty components in quadrature. The primary contributions to the overall systematic uncertainty are represented by the colored lines. The “in-situ JES” contributions correspond to the final in-situ calibration, while uncertainties related to flavor composition, response, and “punch-through” are associated with the global sequential calibration. Flavor-dependent components account for the fact that the di-jet samples used in several calibration steps predominantly consist of gluon jets, although these corrections are applied to all jet flavors [168].

Similarly, discrepancies in jet energy resolution (JER) between data and MC simulation can arise from the mismodelling of detector effects and physics processes in the simulation. To understand and correct these differences, the JER is measured directly in data using di-jet events, where the momenta of the two jets in the final state are expected to be well balanced. Additionally, data collected from random unbiased triggers are used to quantify fluctuations in calorimeter energy deposits caused by pileup activity, providing an independent estimate of the noise term in the energy resolution (see Equation 4.15). These measurements are then combined to produce the final JER results, shown in Figure 6.3. A correction factor is derived and applied to the simulation to match the slightly degraded resolution observed in the data.

6.2 *b*-jets

As already mentioned, all the most sensitive channels for searching for Higgs boson pair production (namely, $b\bar{b}\gamma\gamma$, $b\bar{b}\tau\tau$, and $b\bar{b}b\bar{b}$) involve at least one $H \rightarrow b\bar{b}$ decay. Similarly, many searches for new physics in the scalar sector (among which the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ search) consider bottom quarks in the final state, exploiting the enhanced coupling of the scalar particles to bottom quarks (assuming that they have SM Higgs-like couplings, and their mass is lower than 130 GeV). The hadronization of the bottom quark pair produces two *b*-jets, and their accurate identification is critical to suppress the overwhelming background from jets originating from the hadronization of lighter quarks (u , d , s),




Figure 6.3 – Jet energy resolution measured in data for fully calibrated jets (left), and relative impact of the associated systematic uncertainty (right). The main contributions to the total uncertainty (represented as a shaded area topped with a black solid line), propagated from the two in-situ measurements using di-jet events or random unbiased data, are shown using colored lines.

charm quarks, or gluons. Additionally, the invariant mass of the b -jet pair, $m_{b\bar{b}}$, is a key discriminating variable between resonant $H \rightarrow b\bar{b}$ or $S \rightarrow b\bar{b}$ decays and background processes, where the emission of bottom quarks is non-resonant. To enhance the invariant mass resolution, specific corrections to the energy calibration of b -jets are applied in most searches for Higgs boson pair production, including the $HH \rightarrow b\bar{b}\gamma\gamma$ and the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ analyses.

6.2.1 Jet flavour tagging

The hadronization of bottom quarks results in the production of b -flavored hadrons, with over 90% of these hadrons being B mesons, as illustrated in Figure 6.4a. These B mesons decay via weak interactions, generating a collimated bunch of particles (mainly c - or light-flavored hadrons and leptons), which form the reconstructed b -jet.

A key feature that distinguishes b -jets is the relatively long lifetime of B mesons, approximately 1.6 ps (corresponding to $c\tau \approx 500 \mu\text{m}$ [6]). For a B meson inside a jet with a transverse momentum of 50 GeV, the average distance traveled before decaying is:

$$l = \beta\gamma c\tau \approx 3 \text{ mm} \quad (6.5)$$

As a result, the charged particles produced by B meson decay create tracks whose point of origin is not compatible with the primary vertex, where the B meson was produced. The presence of a displaced secondary vertex within the jet, corresponding to the B meson decay point, is thus a key signature of b -jets. In many cases, B meson decay products include charm hadrons, which decay with shorter lifetimes, creating a tertiary vertex displaced from both the primary and secondary vertices. Moreover, the particles originating from B meson decays are typically much lighter than the B mesons themselves




Figure 6.4 – (a) Production fractions of weakly decaying B hadrons and (b) multiplicity of charged stable particles produced in the decay of a B^0 mesons, estimated using simulated $t\bar{t}$ events produced in $\sqrt{s} = 8$ TeV pp collisions according to different MC generators [170].

(which have masses of around 5 GeV), meaning they tend to have only a few GeV of momentum in the B meson rest frame, often with a large transverse component relative to the jet axis. This causes b -jets to be broader and less collimated compared to jets of different flavour.

To identify b -jets, ATLAS uses advanced b -tagging algorithms [171], which exploit typical features of a B meson decay, such as:

- **Secondary vertices:** a secondary vertex displaced by several millimeters from the primary interaction point is a strong indicator of a B meson decay;
- **Large impact parameters:** tracks from the B meson decay tend to have a large distance of closest approach (*impact parameter*) relative to the primary vertex, as shown in Figure 6.5;
- **Soft leptons:** approximately 35% of B mesons decay semileptonically, meaning that a low-energy, non-isolated lepton (typically a muon) can often be found within the jet cone, providing another identifying feature;
- **High jet constituent multiplicity:** B meson decay chains typically produce more stable hadrons than those from charm or light hadrons, contributing to the higher particle multiplicity within b -jets.

These features are exploited by advanced multivariate techniques, which generate a final discriminant variable to differentiate b -jets from jets produced by charm quarks (c -jets), light quarks, or gluons (collectively referred to as light jets). First, low-level algorithms analyze individual charged particle tracks within the jet cone or reconstruct secondary vertices and B meson decay chains, which help to isolate b -jets from background jets. These low-level outputs are then passed to high-level multivariate classifiers, to maximize b -tagging performance.

As shown in Figure 6.5, tracks from b -hadron decays are displaced from the primary vertex. The IP2D and IP3D algorithms exploit histograms of transverse and longitudinal




Figure 6.5 – Illustration of a *b*-jet, with tracks originating from the fragmentation of the original bottom quark created at the primary vertex and the tracks from the *b*- and *c*-flavored hadron decay at the secondary and tertiary vertex, respectively. The large impact parameter of the tracks originated by the *b*-flavored hadron decay is also evident, when extrapolating the tracks to the primary vertex position.

nal impact parameter significances for each track in the jet, creating likelihood-based discriminants for *b*-tagging. However, these algorithms treat each track independently, limiting their ability to capture correlations between different tracks. To address this, a recurrent neural network (RNN)-based classifier, RNNIP, was developed. The RNN is able to exploit the correlations between track parameters across all tracks in the jet, and provides three output nodes, interpreted as probabilities that the jet is a *b*-jet, *c*-jet, or light jet. These probabilities are then combined into a final discriminant. The SV1 algorithm reconstructs a single secondary vertex within the jet by iteratively combining two-track vertex candidates and excluding tracks from photon conversions or long-lived light hadrons. Similarly, the JetFitter algorithm reconstructs the full *b*- and *c*-hadron decay chain using a Kalman filter, identifying a common line along which the primary and *b*-hadron decay vertices lie, effectively approximating the *b*-hadron flight path and vertex positions.

The outputs of these low-level algorithms are combined using high-level multivariate classifiers. The *b*-tagging algorithm used in ATLAS analyses for the full Run 2 dataset is based on a fully connected deep feed-forward neural network called DL1r, which is trained on information provided by IP2D, IP3D, RNNIP, SV1, and JetFitter. This network produces multiple output nodes that assign probabilities for a jet to be a *b*-jet, *c*-jet, or light-flavor jet, which are then combined into a single *b*-tagging discriminant variable:

$$D_{\text{DL1r}} = \ln \left(\frac{p_b}{f_c \cdot p_c + (1 - f_c) \cdot p_{\text{light}}} \right), \quad (6.6)$$

where p_b , p_c , and p_{light} represent the output probabilities of the network, and f_c is a tunable parameter, that can be optimized to prioritize the rejection of *c*-jets with respect to light jets or vice-versa. The value of $f_c = 0.018$ was chosen as a compromise between

having good rejection capabilities for both c - and light jets over a broad range of b -jet efficiencies, and has been optimized taking into account several ATLAS analyses, including those studying $H \rightarrow b\bar{b}$ decays. Figure 6.6 shows the distribution of D_{DL1r} for b -, c -, and light-flavor jets from simulated $t\bar{t}$ events. The performance of DL1r is compared to previous high-level b -tagging algorithms in ATLAS in Figure 6.7: the DL1 tagger is similar to DL1r but does not use the RNNIP discriminant as input for the training, while MV2 is a BDT-based version of DL1.




Figure 6.6 – Distributions of the D_{DL1r} discriminant for b -jets, c -jets and light jets in simulated $t\bar{t}$ events.

Selection criteria for the discriminant D_{DL1r} can be adjusted to establish optimal working points for balancing b -jet efficiency and background rejection. For Run 2 analyses, these working points are chosen using simulated $t\bar{t}$ samples to achieve fixed inclusive b -jet efficiencies of 60%, 70%, 77%, or 85%. Additionally, a pseudo-continuous b -tagging option is available, where each jet is assigned an integer score corresponding to the most stringent efficiency threshold it meets. For instance, jets passing the tightest operating point (corresponding to an inclusive b -jet efficiency of 60%) have a b -tagging probability between 60% and 0% and are given a score of 5. Jets that do not meet this threshold but pass the next highest (with a b -tagging probability between 70% and 60%) receive a score of 4. This pattern continues across the efficiency bins, with jets failing even the loosest working point (within the 100% to 85% efficiency range) being assigned a score of 1. The operating point used in the $HH \rightarrow b\bar{b}\gamma\gamma$ and $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ searches corresponds to the 77% inclusive b -jet efficiency, which offers a rejection factor of 1/130 for light jets and 1/4.9 for c -jets [172]. The $HH \rightarrow b\bar{b}\gamma\gamma$ analysis also leverages the pseudo-continuous b -tagging to further distinguish the HH signal, which ideally contains well-identified true b -jets, from background processes which include a large fraction of c - or light jets, incorrectly identified as b -jets.




Figure 6.7 – Light jet and *c*-jet rejection factors as a function of the *b*-jet efficiency for the high-level *b*-taggers DL1r, DL1, and MV2, evaluated using jets from simulated $t\bar{t}$ events. The two bottom panels, show the ratio of the light jet rejection and the *c*-jet rejection of the algorithms to MV2. For the same *b*-jet efficiency, the DL1r tagger enhances the light and *c*-jet rejection power of up to a factor two.

6.2.2 *b*-jet identification efficiency measurement

The *b*-jet efficiency of the *b*-tagging working points described in the previous section, along with the corresponding rejection factors for *c*- and light-flavor jets, is initially calculated using jets from simulated $t\bar{t}$ events. However, imperfect modeling of the detector response and of physical processes introduce potential biases in the *b*-tagging efficiency in MC simulations with respect to data. To account for these discrepancies, the *b*-tagging working points must be calibrated by measuring the corresponding *b*-jet efficiency directly from data [172]. The *c*- and light jet rejection factors offered by the *b*-tagging working points are also calibrated, using dedicated analyses [173, 174].

Measuring the *b*-jet efficiency requires to extract a pure sample of *b*-jets from data. This is achieved by selecting $t\bar{t}$ events, which contain two *b*-jets in the final state via the $t \rightarrow Wb$ decay (whose branching fraction approaches 100%). A clean sample of $t\bar{t}$ events is obtained by selecting the dileptonic decay channel, where both *W* bosons from the top quark pair decay leptonically ($W \rightarrow \ell\nu_\ell$). The topology of this process, referred to as

dileptonic $t\bar{t}$, is shown in Figure 6.8. This decay mode results in a final state characterized by two well-identified and isolated leptons (electrons and muons), along with two jets.




Figure 6.8 – Feynman diagram for a dileptonic $t\bar{t}$ event in the transverse plane. The central dot marks the interaction point from which the two top quarks emerge and subsequently decay. The decay chain of the two top quarks results in two leptons and two b -jets.

Event selection

The two leptons in the final state provide a reliable trigger strategy. A combination of single-lepton triggers is employed, whose p_T thresholds ensure to maintain a flat, near 100% efficiency for offline leptons with $p_T > 27$ GeV. The primary background arises from *semileptonic $t\bar{t}$ events*, where only one W boson decays leptonically while the other decays hadronically. *Fake lepton* contributions can originate from several sources, such as non-prompt leptons from decays of b - or c -hadrons, photon conversions (in the case of electrons), or jets misidentified as electrons. Other subdominant backgrounds fall into two categories: those with two real leptons in the final state and those where one lepton candidate is fake. In the first category, the main contributions are single top production in association with a W boson (tW), where the leptons are produced from the top decay and the W decay; diboson production (WW , WZ , ZZ), with leptons from the decays of Z or W bosons; and Z boson production with additional jets, where the Z decays leptonically. Backgrounds involving fake leptons, including, for instance, single top production and W boson production with additional jets, are heavily suppressed.

The following selection criteria are applied on top of the trigger requirements, to select a clean sample of dileptonic $t\bar{t}$ events with minimal background contamination. To eliminate backgrounds from fake or non-prompt leptons, exactly two well reconstructed and identified leptons with opposite charges are required¹. To suppress contributions

¹To be included in this analysis, electron candidates must meet the following criteria: $E_T > 28$ GeV, $|\eta| < 2.47$, be isolated, and pass the Tight identification working point. Similarly, candidate muons must have $p_T >$




Figure 6.9 – Distribution of the p_T (left) and η (right) of the leading jet for events passing the selection in data and simulation. Simulated events are split into physics process contributing to the analysis region. The bottom panels show the data-to-simulation ratio and the fraction of $t\bar{t}$ events among the simulated events.

from the Z +jets background where the Z boson decays leptonically, the two leptons are required to have different flavors: one must be an electron and the other a muon. Oppositely charged, opposite-flavor leptons can also arise from Z +jets events in which the Z boson decays to τ leptons, which subsequently decay leptonically. Such events typically populate the low dilepton invariant mass region. Therefore, this background is mitigated by requiring $m_{e\mu} \geq 50$ GeV. Additionally, contributions from $t\bar{t}$ events that involve high- p_T light jets (alongside the two b -jets from the top quark decay chain), arising from initial or final state radiation or W boson decays, are suppressed by requiring exactly two jets in the event with $p_T > 20$ GeV. After applying the selection criteria, the purity of $t\bar{t}$ events reaches 85%, as visible in Figure 6.9.

Event categorization

The goal of the measurement described in this Section is to extract the b -tagging efficiency for true b -jets. To achieve this, events entering the signal region are categorised as bb , bl , lb , or ll based on the true flavor of the leading and subleading selected jets. Specifically, bb (ll) events are those where both (or neither) of the selected jets are true b -jets. In contrast, bl (lb) events have only the leading (or subleading) jet as a true b -jet, with the

28 GeV, $|\eta| < 2.5$, be isolated, and meet the Tight identification working point requirements.

other being a *c*- or light-flavor jet. According to simulation, almost 60% of the selected events contain two true *b*-jets. The majority of *bb* events arise from $t\bar{t}$ production, with a smaller contribution from Wt events.

Events are further categorised to define a signal region enriched in *bb* events, which is used to extract the *b*-jet efficiency, and three control regions enriched in *lb*, *bl*, and *ll* events. These control regions help estimate the background contamination from events where one or both jets are not true *b*-jets. The estimated fractions of *lb*, *bl*, and *ll* events from the control regions are propagated to the signal region via a simultaneous log-likelihood fit (described below). The categorisation is based on the jet-lepton invariant masses, $m_{j_1\ell}$ and $m_{j_2\ell}$, where j_1 (j_2) refers to the leading (subleading) selected jet, and ℓ represents one of the two selected leptons. These variables are constructed by pairing the two jets and two leptons using the configuration that minimizes $(m_{j_1\ell}^2 + m_{j_2\ell}^2)$. This condition is imposed to avoid highly asymmetric scenarios, where the two invariant masses differ significantly or are both very large. If the paired lepton and jet originate from the same top quark decay, $m_{j\ell}$ is a proxy for the top quark mass, as it captures all its decay products except the missing transverse momentum from the neutrino. Thus, $m_{j\ell} \leq m_{\text{top}} \approx 172.5$ GeV. After pairing, both jet-lepton invariant masses are required to be above 20 GeV, eliminating events with very soft leptons or where the leptons and jets are very close, which are poorly modeled by simulation. Figure 6.10 shows the distribution of $m_{j_1\ell}$ and $m_{j_2\ell}$ after selection criteria are applied, with contributions from *bb*, *bl*, *lb*, and *ll* events highlighted. As mentioned, for events with true *b*-jets, a correctly paired $m_{j\ell}$ should have an upper limit near m_{top} but can be significantly smaller due to the undetected neutrino. In contrast, for events with *c*- or light jets, $m_{j\ell}$ has a much harder spectrum.

The event categorisation process, outlined in Figure 6.11, follows these steps:

1. **Jet p_T binning.** Selected events are first binned based on the p_T of the leading and subleading jets into a two-dimensional grid (p_{T_1}, p_{T_2}) , with bin boundaries at 20, 30, 40, 60, 85, 110, 175, 250, and 400 GeV.
2. **Separation into signal and control regions.** For each (p_{T_1}, p_{T_2}) bin, events are further divided into a *bb*-enriched signal region (SR) and three control regions: CR_{LL} , CR_{BL} , and CR_{LB} , enriched in *ll*, *bl*, and *lb* events, respectively. This separation is achieved by applying selections on the variables $m_{j_1\ell}$ and $m_{j_2\ell}$. The specific selections defining the SR and control regions are detailed in Table 6.1.
3. **Binning into pseudo-continuous *b*-tagging working points.** Events in the SR are further classified based on the pseudo-continuous *b*-tagging scores of the leading and subleading jets, denoted as w_1 and w_2 . As explained above, these values represent the tightest *b*-tagging operating point met by each jet or indicate if the jets are not *b*-tagged. There are five possible *b*-tagging categories for each jet, amounting to a total of 25 possible (w_1, w_2) pairs.

This categorisation results in a total of 1260 orthogonal categories.




Figure 6.10 – Distribution of the jet-lepton invariant masses for the leading (left) and sub-leading (right) selected jets for events passing the selection criteria, shown both in data and simulation. Simulated events are classified in terms of the flavour composition of the two jets. The bottom panels show the data-to-simulation ratio and the fraction of bb events among the simulated events.




Figure 6.11 – Schematic illustration of the event categorisation.

Analysis region	Description	Selection
SR	bb -enriched	$m_{j1\ell} < 175$ GeV and $m_{j2\ell} < 175$
CR_{LL}	ll -enriched	$m_{j1\ell} > 175$ GeV and $m_{j2\ell} > 175$
CR_{BL}	bl -enriched	$m_{j1\ell} < 175$ GeV and $m_{j2\ell} > 175$
CR_{LB}	lb -enriched	$m_{j1\ell} < 175$ GeV and $m_{j2\ell} < 175$

Table 6.1 – Definition of signal and control regions for the simultaneous estimation of the b -tagging probabilities and the true flavor fractions bb , ll , bl , lb .

Maximum likelihood fit

The *b*-jet efficiency for each pseudo-continuous *b*-tagging operating point is determined through a simultaneous fit to data across signal and control regions, spanning all jet p_T bins. This approach allows for a simultaneous estimation of the *b*-tagging probabilities for true *b*-jets and of the jet flavor composition in the selected data. The fit relies on a binned log-likelihood function, expressed as:

$$\log \mathcal{L}(\nu_{\text{tot}}, \Theta) = \sum_{i=1}^N n_i \log \nu_i(\nu_{\text{tot}}, \Theta), \quad (6.7)$$

where ν_{tot} is the total expected number of events (based on simulation), and $\Theta = (\theta_1, \theta_2, \dots, \theta_m)$ represents the set of free parameters in the fit, including both the parameters of interest (POIs) and nuisance parameters (NPs). The sum in Equation 6.7 runs over the $N = 1260$ categories, indexed by i , where n_i and ν_i represent the observed and expected number of events in the i^{th} category, respectively.

The POIs are the *b*-tagging probabilities, termed $\mathcal{P}_b(O^k|T^m)$, which represent the conditional probability that a *b*-jet with p_T in the bin labeled by (T^m) $m = 1, \dots, 9$ satisfies the pseudo-continuous operating point $O_{k=1, \dots, 5}^k$. They are directly related to the inclusive *b*-jet efficiency for a specific working point with efficiency $X\%$, which is defined as the sum of all *b*-tagging probabilities of pseudo-continuous bins above the threshold $X\%$:

$$\epsilon_b(X\%|T^m) = \sum_{O^k > X} \mathcal{P}_b(O^k|T^m) \quad (6.8)$$

In the control regions, the expected number of events depends on the p_T bins of the leading jet (T^m) and the subleading jet (T^n), and it incorporates contributions from all flavor fractions: *bb*, *bl*, *lb*, and *ll*. The equation is given by:

$$\nu_{\text{CR}}(T^m, T^n) = c_{bb}^{m,n} \nu_{\text{CR},bb}^{m,n} + c_{bl}^{m,n} \nu_{\text{CR},bl}^{m,n} + c_{lb}^{m,n} \nu_{\text{CR},lb}^{m,n} + c_{ll}^{m,n} \nu_{\text{CR},ll}^{m,n} \quad (6.9)$$

Here, $\nu_{\text{CR},ij}^{m,n}$ represents the expected number of events where the leading and subleading jets are of flavor i and j , respectively, as estimated from simulation. The factors $c_{ij}^{m,n}$ are unconstrained NPs in the fit and represent (p_{T_1}, p_{T_2}) -dependent correction factors for the *bb*, *bl*, *lb*, and *ll* fractions relative to the simulation.

In the signal region, which is used to extract the *b*-tagging probabilities, the expected number of events depends on the p_T bins of the leading and subleading jets (T^m and T^n), as well as the pseudo-continuous *b*-tagging probabilities for these jets, termed O^k and O^p . Therefore, the expected number of events in the SR can be written as follows:

$$\begin{aligned} \nu_{\text{SR}}(T^m, T^n, O^k, O^p) = & c_{bb}^{m,n} \nu_{\text{SR},bb}^{m,n} \cdot \mathcal{P}_b(O^k|T^m) \cdot \mathcal{P}_b(O^p|T^n) \\ & + c_{bl}^{m,n} \nu_{\text{SR},bl}^{m,n} \cdot \mathcal{P}_b(O^k|T^m) \cdot \mathcal{P}_l(O^p|T^n) \\ & + c_{lb}^{m,n} \nu_{\text{SR},lb}^{m,n} \cdot \mathcal{P}_l(O^k|T^m) \cdot \mathcal{P}_b(O^p|T^n) \\ & + c_{ll}^{m,n} \nu_{\text{SR},ll}^{m,n} \cdot \mathcal{P}_l(O^k|T^m) \cdot \mathcal{P}_l(O^p|T^n) \end{aligned} \quad (6.10)$$

In Equation 6.10, $\nu_{\text{SR},ij}^{m,n}$ is defined similarly as $\nu_{\text{CR},ij}^{m,n}$ but it refers to the SR rather than the CRs. The $c_{ij}^{m,n}$ factors are shared between the control and signal regions. The probability $\mathcal{P}_l(O^k|T^m)$ represents the conditional likelihood that a c - or light jet is b -tagged and satisfies the pseudo-continuous b -tagging bin O^k . These *light jet probabilities* are fixed parameters in the fit and are derived from the calibration of light jet mistagging efficiencies [174].

This measurement is subject to various sources of systematic uncertainties, related to both detector effects and physics modeling. These uncertainties include the jet energy scale and resolution, the modeling of signal $t\bar{t}$ events and background processes (with a specific uncertainty assigned to the background from fake leptons), and the uncertainty related light jet mis-tagging probabilities, which propagated from the corresponding measurement. To assess the impact of each source of uncertainty on the b -tagging probability, the corresponding parameter in the simulation or fit input is varied by one standard deviation. The analysis is then repeated with this modified parameter, and the difference between the biased b -tagging probability and the nominal results is used to determine the contribution of each uncertainty source.

Results for the DL1r tagger

This measurement was performed for the first time using data from the first three years of Run 2 (2015, 2016, and 2017) to determine the b -jet efficiency for the operating points optimized for the MV2 and DL1 taggers [172]. The analysis was then extended to include the full Run 2 dataset to evaluate the b -jet efficiency for the operating point of the DL1r tagger, which is presented here. Both the b -tagging probabilities and b -jet efficiencies obtained from the maximum likelihood fit to data are compared to those derived from MC samples. The resulting scale factors are then used to adjust the simulation so that it aligns with the observed data. Figure 6.12a displays the scale factors for the b -tagging probabilities for each pseudo-continuous operating point of the DL1r tagger, for jets with p_T in the range $85 < p_T < 110$ GeV. The b -jet efficiencies are then calculated from the b -tagging probabilities using Equation 6.8. Figure 6.12b shows the b -jet efficiency scale factors for the 70% operating point as a function of the jet's p_T .

Results for the GN2 tagger

The ATLAS Collaboration has recently developed a new generation of flavor tagging algorithms based on Graph Neural Networks (GNN) [20]. These advanced algorithms utilize information from a variable number of charged particle tracks within a jet to directly predict the jet flavor, bypassing the need for intermediate low-level algorithms. The state-of-the-art GNN-based algorithm currently employed by ATLAS is the GN2 tagger. This algorithm significantly improves upon the performance of previous-generation taggers, offering up to a fourfold increase in c -jet rejection and double the light jet rejection, while maintaining the same b -jet efficiency. This enhanced performance is illustrated in Figures 6.13 and 6.14.

The adoption of the GN2-based b -tagging algorithm is expected to significantly enhance




Figure 6.12 – (a) Data-to-MC scale factors (SFs) for the *b*-tagging probabilities in each pseudo-continuous bin of the DL1r tagger, for jet p_T in the range $85 < p_T < 110$ GeV, and (b) *b*-jet efficiency SFs for the 70% operating point of DL1r as a function of the jet's p_T .

the sensitivity of new searches for Higgs boson pair production, as the most sensitive channels involve at least two *b*-jets in the final state.

Five fixed-efficiency working points were established by setting thresholds on the GN2-based discriminant, corresponding to inclusive *b*-jet selection efficiencies of 65%, 70%, 77%, 85%, and 90%. These fixed-efficiency working points, along with the associated pseudo-continuous *b*-tagging probabilities, were calibrated using the same *t* $\bar{t}$ -based in-situ measurement outlined in Section 6.2.2. I have contributed to the in-situ calibration of the GN2 tagger using Run 2 and Run 3 data (collected in 2022 and 2023). The results obtained with the Run 2 dataset are summarised below².

Figure 6.15a displays the scale factors for the *b*-tagging probabilities for each pseudo-continuous operating point of the GN2 tagger, for jets with p_T in the range $85 < p_T < 110$ GeV. The corresponding *b*-jet efficiencies are calculated from the *b*-tagging probabilities using Equation 6.8. Figure 6.15b shows the *b*-jet efficiency scale factors for the 70% operating point as a function of the jet's p_T .

6.2.3 *b*-jet energy corrections

As illustrated in Figure 6.3a, the typical jet energy resolution for jets reconstructed in the ATLAS detector varies between 25% for $p_T \approx 20$ GeV to almost 5% at extremely high p_T values (above 1 TeV). This directly impacts the resolution of the dijet invariant mass

²The two searches for Higgs boson pairs described in this thesis, in Chapter 7 and Chapter 8, select *b*-jets in their signal region using the DL1r tagger. However, the new analyses, which use both the complete Run 2 and partial Run 3 datasets, are replacing DL1r with the GN2 tagger to take advantage of its enhanced *b*-tagging performance. Consequently, the *b*-jet calibration results presented here are a critical ingredient of the ongoing $HH \rightarrow b\bar{b}\gamma\gamma$ and $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ searches based on Run 2 and Run 3 data.




Figure 6.13 – The light jet and *c*-jet rejection factors as a function of *b*-jet efficiency are shown for the *b*-taggers GN2, GN1, and DL1d, evaluated using jets from simulated $t\bar{t}$ events. The bottom two panels display the ratio of the light jet and *c*-jet rejection rates of the algorithms relative to DL1d. Here, GN1 represents a GNN-based algorithm, while DL1d is a traditional high-level *b*-tagging algorithm, similar to DL1r.




Figure 6.14 – The c -jet and light-jet rejection of different flavour tagging algorithms developed by the ATLAS Collaboration over time, for an inclusive 70% b -jets efficiency from simulated $t\bar{t}$ events. The rejections are evaluated inclusively in the jets' p_T . Between the DL1r and DL1d, a transition from Run 2 to Run 3 reconstruction took place.




Figure 6.15 – (a) Data-to-MC scale factors (SFs) for the b -tagging probabilities in each pseudo-continuous bin of the GN2 tagger, for jet p_T in the range 85 < p_T < 110 GeV, and (b) b -jet efficiency SFs for the 70% operating point of GN2 as a function of the jet's p_T .

from $H \rightarrow b\bar{b}$ decays, leading to a resolution of $m_{b\bar{b}}$ between 10% and 20%. As a result, resolving the $H \rightarrow b\bar{b}$ mass peak is significantly more challenging compared to the $H \rightarrow \gamma\gamma$ peak, where the diphoton invariant mass resolution can be as precise as 1% to 2%. Beyond the intrinsically poorer energy resolution of jets with respect, for instance, to photons, measuring the energy of b -jets presents additional challenges compared to light jets. B meson decays frequently produce muons and neutrinos, whose energy escapes the ATLAS calorimeters, complicating the energy calibration. Furthermore, b -jets have broader cones and contain softer particles than light jets, which means that the standard jet energy calibration procedure, derived for light jets, is less effective for b -jets. To improve the resolution of the b -jet pair invariant mass, $m_{b\bar{b}}$, which is crucial for distinguishing the $H \rightarrow b\bar{b}$ (or $S \rightarrow b\bar{b}$) peak in double Higgs boson events from backgrounds, dedicated corrections are applied to the four-momentum reconstruction of b -jets [11]. These corrections include two components:

- **μ -in-jet.** When a non-isolated muon with $p_T > 5$ GeV is found within $\Delta R = 0.4$ of the jet axis, its four-momentum is added to the jet's, and the small amount of energy deposited by the muon in the calorimeters is subtracted. If more than one muon is matched to the b -jet, the closest muon to the jet axis is used.
- **PtReco.** MC simulations are employed to determine the expected response of reconstructed b -jets compared to truth jets, which are clustered using all stable hadrons, while also including non-isolated muons and neutrinos. A correction is then applied to account for residual differences in response, caused by energy leakage outside the b -jet cone (*out-of-cone* leakage) and energy carried away by undetected neutrinos. Separate corrections are derived for b -jets containing B mesons that decay semileptonically or fully hadronically.

A more detailed discussion of these b -jet energy corrections can be found in [175]. Implementing the μ -in-jet and PtReco corrections improves the $m_{b\bar{b}}$ resolution in $HH \rightarrow b\bar{b}\gamma\gamma$ events by approximately 20% compared to the standard jet energy calibration, as illustrated in Figure 6.16.

6.3 Missing transverse momentum

The presence of neutrinos or potential new beyond-the-SM particles that interact weakly with detector material can only be inferred through momentum imbalance in the transverse plane. Since the total transverse momentum in a pp collision is expected to be zero, any undetectable particles produced would result in a missing transverse momentum, E_T^{miss} , measured within the ATLAS detector.

The E_T^{miss} is an event-level quantity, calculated as the negative vector sum of the transverse momenta of all reconstructed objects in an event [176]. The E_T^{miss} components along the x and y axes are defined as:

$$E_{x(y)}^{\text{miss}} = - \sum_{\text{electrons}} E_{x(y)}^e - \sum_{\text{photons}} E_{x(y)}^\gamma - \sum_{\text{taus}} E_{x(y)}^\tau - \sum_{\text{jets}} E_{x(y)}^{\text{jet}} - \sum_{\text{muons}} E_{x(y)}^\mu - E_{x(y)}^{\text{soft}}, \quad (6.11)$$




Figure 6.16 – $m_{b\bar{b}}$ distribution for the two leading jets in simulated $HH \rightarrow b\bar{b}\gamma\gamma$ events, constructed using the standard jet energy calibration (blue) and after applying the specific b -jet energy corrections (pink).

where $E_{x(y)}^i$ represents the calibrated energy of the corresponding physics object projected along the x (y) axis, and $E_{x(y)}^{\text{soft}}$ refers to contributions from tracks with $p_T > 0.5$ GeV and $|\eta| < 2.5$, originating from the primary vertex but not associated with any reconstructed object.

The magnitude of the missing transverse momentum and its azimuthal angle are given by:

$$E_T^{\text{miss}} = \sqrt{(E_x^{\text{miss}})^2 + (E_y^{\text{miss}})^2}, \quad \phi^{\text{miss}} = \arctan\left(\frac{E_y^{\text{miss}}}{E_x^{\text{miss}}}\right) \quad (6.12)$$

Search for Higgs boson pair production in the $HH \rightarrow b\bar{b}\gamma\gamma$ channel

Double Higgs boson production is regulated by the trilinear Higgs boson self-coupling modifier $\kappa_\lambda = \lambda_{HHH}/\lambda_{HHH}^{SM}$. The search for Higgs boson pairs in the $b\bar{b}\gamma\gamma$ final state provides a powerful handle to κ_λ , which, in turn, would give crucial insights into the exact shape of the scalar potential close to its minima, which is currently unknown. This Chapter presents a search for Higgs boson pair production in the $b\bar{b}\gamma\gamma$ final state using 140 fb^{-1} of 13 TeV pp collision data recorded by the ATLAS detector during the Run 2 data-taking period [10].

As discussed in Section 2.1.1, SM double Higgs boson production is an extremely rare process, with the cross section for the dominant production mode (gluon fusion, ggF HH) suppressed by three orders of magnitude compared to single Higgs boson production rates. Additionally, the branching fraction for the $b\bar{b}\gamma\gamma$ channel is the lowest among the *golden channels* for di-Higgs searches, at just 0.26%. Therefore, to enhance the signal-to-background ratio, which is critical for this analysis, it is essential to thoroughly understand and model all competing SM processes that could mimic the $b\bar{b}\gamma\gamma$ signature.

In this analysis, HH production via ggF and VBF in the $b\bar{b}\gamma\gamma$ final state is considered the signal, while the primary irreducible background arises from diphoton production accompanied by the emission of two bottom quarks. In diphoton production, the invariant mass of the photon pair, $m_{\gamma\gamma}$ ¹, is expected to follow a continuous smoothly falling distribution, which does not exhibit any peaking structures: hence, the $\gamma\gamma+b\bar{b}$ background source corresponds to the dominant contribution of the so-called *continuum background*. Another sizable contribution to the continuum background is represented by diphoton production plus additional c - or light jets, that are incorrectly classified as b -jets by the b -tagging requirements. This component represents more than 30% of the continuum background in the phase space of the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis, as will be explained in more detail in this Chapter. A smaller, reducible contribution to the continuum background arises from γj and $j j$ production, where one or two jets are misidentified as photons. This *fake photon* background is suppressed thanks to placing stringent requirements on the quality of the photons selected in the signal region. Despite the cross section for dijet

¹Since photons are massless, their invariant mass can be expressed in terms of their energies E^{γ_1} and E^{γ_2} , and the angle α between their directions: $m_{\gamma\gamma} = \sqrt{2E^{\gamma_1}E^{\gamma_2}(1 - \cos\alpha)}$.

production being three orders of magnitude larger than the one for diphoton production [177–179], the reducible γj and $j j$ components accounts for only 20% of the overall continuum background in the $HH \rightarrow b\bar{b}\gamma\gamma$ phase space.

Production of a single Higgs boson decaying to two photons represents another subdominant background for the $HH \rightarrow b\bar{b}\gamma\gamma$ search. The single Higgs production modes providing the largest contributions are ggF, Z boson-associated production and $t\bar{t}H$. In fact, the ggF process is the most abundant single Higgs production mode, and can be accompanied by the emission of two or more (b -)jets in the final state, thus mimicking the $b\bar{b}\gamma\gamma$ signature. ZH production can also result in a $b\bar{b}\gamma\gamma$ signature, when the Z boson decays to two bottom quarks. Finally, the decay chain of the top quarks always results in the emission of two b -jets, thus making the $t\bar{t}H(\rightarrow \gamma\gamma)$ process one of the dominant single Higgs backgrounds for the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis. The two photons from the decay of the Higgs boson will appear as a resonant narrow peak centered around $m_H \approx 125 \text{ GeV}$ ² in the $m_{\gamma\gamma}$ spectrum: single Higgs boson production forms the *resonant* background.

Figure 7.1 illustrates the composition of the background in the $m_{\gamma\gamma}$ spectrum within the phase space of the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis. The di-Higgs signal is visible as a small additional contribution to the narrow peak from the $H \rightarrow \gamma\gamma$ resonance, which sits on top of the smoothly decaying $m_{\gamma\gamma}$ shape generated by the continuum background.




Figure 7.1 – Distributions of the diphoton invariant mass for events in data (dots with error bars) compared with the sum of the expected signal and backgrounds (histograms) in one of the signal regions of the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis, referred to as “High Mass 1”. In the plot, the relative fractions of the $\gamma\gamma+b\bar{b}$ and the $\gamma\gamma$ + other jets samples are highlighted. The reducible γj and $j j$ components of the continuum background (estimated using a data-driven technique) are also shown.

Section 7.1 and 7.2 provide a description of the dataset and all simulated samples for the

²This statement is true for both signal events (i.e. double Higgs production in the $b\bar{b}\gamma\gamma$ channel) and single Higgs background events, where $H \rightarrow \gamma\gamma$.

signal and background processes used in the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis. The definition of physics objects and the signal region, which includes events with a pair of good-quality photons and two b -jets, is detailed in Section 7.3.1. Section 7.3.2 outlines the event categorisation, while Section 7.4 illustrates the modeling of the diphoton invariant mass spectrum for both signal and background processes. Section 7.5 addresses the systematic uncertainties affecting the analysis results. Finally, all these elements are combined to construct the statistical model, as described in Section 7.6, leading to the final results outlined in Section 7.7.

Additionally, Section 7.8 presents the combination between the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis and the other di-Higgs searches conducted using Run 2 data in various final states, including the other two *golden channels*, namely $b\bar{b}b\bar{b}$ and $b\bar{b}\tau\tau$, plus two additional less sensitive channels (involving multi-leptons and $b\bar{b}\ell\ell + E_T^{\text{miss}}$ final states). The goal of the HH combination is to set the most stringent constraints on Higgs boson pair production and on the trilinear Higgs boson self-coupling modifier κ_λ using the full Run 2 dataset.

7.1 Data

This analysis is based on pp collision data collected by the ATLAS detector throughout the entire Run 2 of the LHC, spanning from 2015 to 2018, with proton beams colliding at a center-of-mass energy of 13 TeV. After applying data quality requirements [180], the final dataset corresponds to an integrated luminosity of $140 \pm 1.2 \text{ fb}^{-1}$ [49].

The two photons in the final state provide a robust handle for the trigger strategy, shared between all the $H \rightarrow \gamma\gamma$ analyses. The data used in this analysis are collected using a set of photon triggers [181] to ensure an efficient selection of high-quality photon pairs in the signal region. Specifically, the diphoton triggers HLT_g35_loose_g25_loose and HLT_g35_medium_g25_medium_L12EM20VH are applied to data collected in 2015-2016 and 2017-2018, respectively. Both triggers require events with two photons, where the leading (subleading) photon has transverse energy $E_T > 35(25) \text{ GeV}$. The former trigger applies a *Loose* photon identification, while the latter enforces stricter *Medium* identification criteria and calorimetric isolation requirements. This is needed to keep trigger rates at manageable levels, given the higher instantaneous luminosities in 2017 and 2018, compared with early Run 2 conditions, as explained in Section 3.3. To address potential inefficiencies in selecting $H \rightarrow \gamma\gamma$ candidates, single-photon triggers are employed alongside diphoton triggers. These are relevant in scenarios where the two photons from the Higgs decay display a large asymmetry in transverse energy, with the leading photon being significantly more energetic than the subleading one. They are also important for recovering events where the Higgs boson is highly boosted, causing the two photons from its decay to become so collimated that they are reconstructed as a single photon. Two sets of single-photon triggers are used: HLT_g120_loose for data from 2015 and HLT_g140_loose for data from 2016 to 2018. Both triggers select events with a *Loose* photon and $E_T > 120 \text{ GeV}$ or 140 GeV , respectively. The inclusion of these single-photon triggers has a negligible effect on the expected number of SM HH signal events entering the signal region, with a relative impact of less than 0.1%. However, for non-SM HH

signals characterised by anomalous κ_{2V} values, where boosted Higgs boson production is enhanced, the inclusion of single-photon triggers becomes more relevant, with a relative impact of up to 1% on the expected yields of events entering the signal region of the analysis. The online trigger E_T thresholds are reflected in the offline diphoton selection. At the analysis level, events with leading (subleading) photon transverse energy $E_T > 0.35(0.25) \cdot m_{\gamma\gamma}$ (where $m_{\gamma\gamma}$ is the invariant mass of the photon pair) are selected. An additional cut on the diphoton invariant mass is applied, requiring $105 < m_{\gamma\gamma} < 160$ GeV. Combining the cuts on the relative transverse energies of the photons and the lower bound of the $m_{\gamma\gamma}$ window, the final selection criteria for the leading (subleading) photon transverse energy is $E_T > 36.75(26.25)$ GeV.

7.2 Simulated samples

Simulated events are crucial for accurately modeling both signal and background processes, including the response of the ATLAS detector. *Monte Carlo* (MC) samples are sets of simulated events widely used in ATLAS physics analyses. They play a crucial role in estimating selection efficiencies, constructing distributions for key discriminating variables, and building the statistical models under background-only or signal-plus-background assumptions, in order to test real data against either of these hypotheses.

The production of MC samples begins with the generation of physics events, using amplitudes calculated to a specific order in perturbation theory. This process is handled by an *event generator*. The subsequent emission of quarks and gluons from the partons involved in the primary interaction (known as *parton showering* [182]) and their hadronization [183, 184] are also simulated by an event generator. The latter may not necessarily be the same that takes care of the hard scattering event, and typically it is also involved in the simulation of the decays of unstable particles created in the hard-scattering interaction. MC samples also simulate the behaviour of the partons that are not involved in the hard scattering (i.e. the *underlying event*, see Section 3.1), and pileup events. The next step involves simulating the interactions of stable particles in the final state with the ATLAS detector and modeling the detector's response. This is achieved either through an accurate and detailed simulation of the ATLAS detector using GEANT4 [185], or through a faster method, known as *ATLAS Fast Simulation*, which reduces CPU demands by employing a simplified detector geometry and parametrisations of the interactions between particles and the detector material [186–188]. Finally, identical reconstruction algorithms are applied to both simulated events and real data.

MC samples are generated for both signal and background processes in the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis.

7.2.1 HH samples

Events from ggF HH production with coupling modifiers $\kappa_\lambda = 1$ and 10 (with κ_t set to unity) are generated using the POWHEG-BOX event generator [189–193] at NLO accuracy in QCD, incorporating parton distribution functions (PDFs) from PDF4LHC15 [194] and accounting for finite top-quark mass effects [195]. Parton showering, hadroniza-

tion, and the underlying event are simulated with PYTHIA 8.244 [196], while theoretical uncertainties related to parton showering are assessed using the alternative generator HERWIG 7.1.6 [197, 198]. The cross section used to normalize the SM ggF HH sample is calculated at NNLO accuracy in QCD, including finite top quark mass effects, as detailed in Section 2.1.1. On the other hand, the cross section for the $\kappa_\lambda = 10$ sample combines NNLO accuracy in QCD in the $m_{\text{top}} \rightarrow \infty$ limit and the full top quark mass dependence within the NLO calculation (see Section 2.1.2). The two predictions are shown in Table 7.1.

Table 7.1 – Combinations of coupling modifiers (κ_λ, κ_t) used to generate ggF HH samples for the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis. The corresponding cross sections, evaluated with the highest available accuracy, are also shown.

κ_λ	κ_t	Cross section [fb]
1	1	31.05
10	1	672.2

VBF HH events are generated at leading order (LO) using MADGRAPH5_AMC@NLO 2.6.0 [68, 69]. The PDF set from NNPDF3.0 [199, 200] is employed, while the parton showering is handled by PYTHIA 8.224. SM VBF HH events (with coupling modifiers ($\kappa_\lambda, \kappa_{2V}, \kappa_V$) all set to unity) are generated. Additionally, twelve non-SM samples are produced with various combinations of ($\kappa_\lambda, \kappa_{2V}, \kappa_V$), as detailed in Table 7.2. The SM VBF HH sample is normalized using the cross section calculated at $N^3\text{LO}$ in QCD, including NLO electroweak corrections, while non-SM VBF HH samples adopt LO cross sections, rescaled with the $N^3\text{LO}$ to LO scale factor extracted from the SM cross section, as explained in Section 2.1.3.

κ_λ and κ_{2V} parametrisation

As discussed in Sections 2.1.2 and 2.1.3, Higgs boson pair production via ggF and VBF is influenced by four coupling modifiers: κ_λ , κ_t , κ_{2V} , and κ_V . This analysis aims to measure κ_λ and κ_{2V} , which remain largely experimentally unconstrained (single Higgs boson production already provides precise measurements for κ_t and κ_V , as detailed in Section 1.3.3). To achieve this, the presence of a HH signal needs to be tested over a two-dimensional search space, spanned by the parameters κ_λ and κ_{2V} . However, simulating HH samples across a fine grid in the ($\kappa_\lambda, \kappa_{2V}$) plane would be prohibitively resource-intensive. Hence, only a handful of HH signal samples are simulated using carefully selected coupling values, and ggF HH and VBF HH production is then parametrised as a function of κ_λ and κ_{2V} by performing a linear combination of the available samples. At LO, the ggF HH process involves contributions from two amplitudes, represented by the Feynman diagrams in Figures 2.2a and 2.2b. These amplitudes, $\mathcal{A}_{\text{box}}$ and $\mathcal{A}_{\text{triangle}}$, scale with the coupling modifiers κ_t and $\kappa_\lambda \cdot \kappa_t$, respectively. The differential ggF HH cross section in an arbitrarily small region of the phase space $d\Phi$ is proportional to the square of the total amplitude, which can be expressed as a second-degree polynomial of

Table 7.2 – Combinations of coupling modifiers $(\kappa_\lambda, \kappa_{2V}, \kappa_V)$ used to generate VBF HH samples for the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis. The corresponding cross sections, evaluated with the highest available accuracy, are also shown.

κ_λ	κ_{2V}	κ_V	Cross section [fb]
1	1	1	1.726
0	1	1	4.608
2	1	1	1.559
10	1	1	98.059
1	0	1	27.069
1	0.5	1	9.634
1	1.5	1	3.356
1	2	1	14.515
1	3	1	65.431
0	0	1	36.965
1	1	0.5	10.958
1	1	1.5	66.071
-5	1	0.5	6.472

κ_λ and κ_t :

$$\begin{aligned}
 \frac{d\sigma_{\text{ggF}}(\kappa_\lambda, \kappa_t)}{d\Phi} &\propto |\mathcal{A}_{\text{box}} + \mathcal{A}_{\text{triangle}}|^2 \\
 &= a_1 \cdot \kappa_\lambda^2 \kappa_t^2 + a_2 \cdot \kappa_t^2 + a_3 \cdot \kappa_\lambda \kappa_t^2 \\
 &= \kappa_t^2 \cdot \left(a_1 \cdot \kappa_\lambda^2 + a_2 + a_3 \cdot \kappa_\lambda \right),
 \end{aligned} \tag{7.1}$$

where a_1 , a_2 , and a_3 are independent of κ_λ and κ_t . Higher-order QCD corrections do not introduce additional $t\bar{t}H$ or trilinear Higgs boson vertices (which scale with κ_t and κ_λ respectively) into the diagrams shown in Figure 2.2. As a result, the parametrisation of the ggF HH cross section presented in Equation 7.1 remains valid at any order in QCD, even when the amplitudes $\mathcal{A}_{\text{box}}$ and $\mathcal{A}_{\text{triangle}}$ are modified to account for higher-order QCD corrections. Equation 7.1 is valid for any point in the $(\kappa_\lambda, \kappa_t)$ parameter space. For example, in the SM case, the differential ggF HH cross section can be expressed in terms of the coefficients a_1 , a_2 , and a_3 as follows:

$$\frac{d\sigma_{\text{ggF}}(1, 1)}{d\Phi} = a_1 + a_2 + a_3 \tag{7.2}$$

This opens the possibility of solving a system of linear equations, where the unknowns are the terms a_1 , a_2 , and a_3 , while the differential ggF HH cross sections, derived from available ggF HH MC samples simulated with specific $(\kappa_\lambda, \kappa_t)$ values, are known. Since there are three unknowns a_i (with $i = 1, 2, 3$), three ggF HH reference samples - referred to as *basis samples* - are required, each generated with different values of the κ_λ coupling

modifier³. By inverting the 3×3 coefficient matrix of this linear systems system, the a_i quantities can be expressed in terms of the known $\frac{d\sigma_{\text{ggF}}(\kappa_\lambda, \kappa_t)}{d\Phi}$ distributions obtained from the three basis samples. Substituting the solutions for the a_i terms back into Equation 7.1 yields the distribution of any variable Φ for ggF HH events as a function of $(\kappa_\lambda, \kappa_t)$, in the form of a linear combination of the three reference samples. The linear coefficients of this combination depend on the values of κ_λ and κ_t . This approach provides a natural parametrisation of ggF HH events across the $(\kappa_\lambda, \kappa_t)$ plane, once the basis is chosen. The solution to Equation 7.1 requires three ggF HH reference samples, but only two fully simulated ggF HH MC samples are available, as summarized in Table 7.1. To overcome this limitation, three additional ggF HH samples were generated at *truth level* with $\kappa_\lambda = 1, 0$, and 20 (and κ_t set to unity)⁴. Solving the system of linear equations using these three truth-level samples as the basis, and substituting the solutions back into Equation 7.1, yields the following expression for the differential ggF HH cross section:

$$\begin{aligned} \frac{d\sigma_{\text{ggF}}(\kappa_\lambda, \kappa_t)}{d\Phi} = \kappa_t^2 \cdot & \left[\left(\kappa_t^2 + \frac{\kappa_\lambda^2}{20} - \frac{399}{380} \kappa_\lambda \kappa_t \right) \cdot \frac{d\sigma_{\text{ggF}}(0, 1)}{d\Phi} \right. \\ & + \left(\frac{40}{38} \kappa_\lambda \kappa_t - \frac{2}{38} \kappa_\lambda^2 \right) \cdot \frac{d\sigma_{\text{ggF}}(1, 1)}{d\Phi} \\ & \left. + \left(\frac{\kappa_\lambda^2 - \kappa_\lambda \kappa_t}{380} \right) \cdot \frac{d\sigma_{\text{ggF}}(20, 1)}{d\Phi} \right] \end{aligned} \quad (7.3)$$

The linear combination of the basis samples in Equation 7.3 was then employed to generate truth-level histograms of the invariant mass of the HH system, m_{HH} , for different κ_λ values in the range $-30 \leq \kappa_\lambda \leq 30$. Figure 7.2 shows the distributions of m_{HH} for ggF HH events simulated at truth-level with $\kappa_\lambda = 1, 0$, and 20. The distribution of any other observable at the *reconstruction level* - after considering parton showering, detector simulation, and event reconstruction - for any targeted κ_λ value within the range $-30 \leq \kappa_\lambda \leq 30$ are obtained by reweighting the fully simulated SM or $\kappa_\lambda = 10$ samples [201]. This reweighting involves applying an event-by-event weight, computed as the ratio of the binned m_{HH} distribution for the target κ_λ value to the binned m_{HH} distribution of the SM (or $\kappa_\lambda = 10$) sample. This reweighting procedure assumes that the complete kinematics of the ggF HH process are driven solely by m_{HH} , meaning that the differential ggF HH cross sections for other observables are expected to vary consistently with changes in the m_{HH} distributions. The validity of this reweighting approach was checked by comparing the fully simulated ggF HH signal with $\kappa_\lambda = 10$ to the corresponding sample obtained by reweighting the SM sample, and vice versa. Figures 7.3a and 7.3b display the distributions of the diphoton invariant mass, $m_{\gamma\gamma}$, and the invariant mass of the four final-state objects (two photons and two b -jets), $m_{b\bar{b}\gamma\gamma}$, using both

³There is no need to vary κ_t when choosing the basis samples, as κ_t only serves as an overall normalization factor in Equation 7.1

⁴These truth-level simulations are incomplete and not suitable for extracting reliable predictions for the ggF HH process to compare with real data: they do not include parton showering or pass the events through detector simulation.




Figure 7.2 – Distributions of the invariant mass of the Higgs boson pair, m_{HH} , for ggF HH events simulated at truth-level with $\kappa_\lambda = 1, 0$, and 20 . The histograms are normalised to unity.

the fully simulated and reweighted samples. The agreement on the expected yields is typically within a factor of 10%.




Figure 7.3 – $m_{\gamma\gamma}$ (left) and $m_{b\bar{b}\gamma\gamma}$ (right) distributions obtained using simulated ggF HH events, with $\kappa_\lambda = 10$. The distributions extracted from the existing fully simulated sample (black) are compared to the corresponding results from the reweighted sample (pink).

To summarize, to parametrise the ggF HH process as a function of κ_λ and κ_t , a system

of three linear equations based on Equation 7.1 was solved using three ggF HH samples with $\kappa_\lambda = 1, 0$, and 20 as the basis, thus yielding Equation 7.3. The basis is built using the available fully simulated SM sample, while the two ggF HH samples with $\kappa_\lambda = 0$ and 20 are obtained by applying the m_{HH} -based reweighting described above to the existing simulated ggF HH samples with $\kappa_\lambda = 1$ and 10 respectively.

The parametrisation of the VBF HH process as a function of κ_λ , κ_{2V} , and κ_V follows a similar approach to that used for ggF HH. As illustrated in Figure 2.4, the total amplitude for Higgs boson pair production via VBF at LO is the sum of three diagrams, scaling with $\kappa_\lambda\kappa_V$, κ_{2V} , and κ_V^2 respectively:

$$\mathcal{A}_{\text{VBF}}(\kappa_\lambda, \kappa_{2V}, \kappa_V) = A_1 \cdot \kappa_\lambda\kappa_V + A_2 \cdot \kappa_{2V} + A_3 \cdot \kappa_V^2 \quad (7.4)$$

where A_1 , A_2 , and A_3 are normalization factors for each diagram. Therefore, the differential VBF HH cross section, $\frac{d\sigma_{\text{VBF}}(\kappa_\lambda, \kappa_{2V}, \kappa_V)}{d\Phi}$ (with $d\Phi$ representing the infinitesimal phase space element), can be expressed as a polynomial in the three coupling modifiers:

$$\begin{aligned} \frac{d\sigma_{\text{VBF}}(\kappa_\lambda, \kappa_{2V}, \kappa_V)}{d\Phi} &\propto |\mathcal{A}_{\text{VBF}}(\kappa_\lambda, \kappa_{2V}, \kappa_V)|^2 \\ &= b_1 \cdot \kappa_\lambda^2 \kappa_V^2 + b_2 \cdot \kappa_{2V}^2 + b_3 \cdot \kappa_V^4 + b_4 \cdot \kappa_\lambda \kappa_V \kappa_{2V} \\ &\quad + b_5 \cdot \kappa_\lambda \kappa_V^3 + b_6 \cdot \kappa_{2V} \kappa_V^2 \end{aligned} \quad (7.5)$$

where $b_1, b_2, \dots, b_6$ are constants. As with ggF HH, Equation 7.5 can be used to build a system of six linear equations with six unknown parameters, b_i , with $i = 1, \dots, 6$. The known terms are the differential VBF HH cross sections derived from MC samples simulated with different $(\kappa_\lambda, \kappa_{2V}, \kappa_V)$ values. To solve this system, six independent MC samples are required as the basis for the VBF HH parametrisation. Among the thirteen available VBF HH MC samples (summarized in Table 7.2), six are chosen as the basis for the parametrisation. These correspond to the following combinations of the coupling modifiers $(\kappa_\lambda, \kappa_{2V}, \kappa_V)$: $(1, 1, 1)$, $(0, 1, 1)$, $(10, 1, 1)$, $(1, 1.5, 1)$, $(1, 3, 1)$, and $(-5, 1, 0.5)$. This choice of basis was shown to be able to model VBF HH kinematics over a broad parameter space, extending to κ_λ , κ_{2V} , and κ_V values far beyond the SM. The differential

VBF HH cross section can then be expressed as:

$$\begin{aligned}
 \frac{d\sigma_{\text{VBF}}(\kappa_\lambda, \kappa_{2V}, \kappa_V)}{d\Phi} = & \left(\kappa_{2V}^2 - \frac{373\kappa_{2V}\kappa_\lambda\kappa_V}{594} - \frac{1150\kappa_{2V}\kappa_V^2}{297} + \right. \\
 & \left. - \frac{\kappa_\lambda^2\kappa_V^2}{9} + \frac{1033\kappa_\lambda\kappa_V^3}{594} + \frac{853\kappa_V^4}{297} \right) \cdot \frac{d\sigma_{\text{VBF}}(1, 1, 1)}{d\Phi} \\
 & + \left(\frac{2\kappa_{2V}\kappa_\lambda\kappa_V}{3} - \frac{2\kappa_{2V}\kappa_V^2}{3} + \frac{\kappa_\lambda^2\kappa_V^2}{10} - \frac{53\kappa_\lambda\kappa_V^3}{30} + \frac{5\kappa_V^4}{3} \right) \cdot \frac{d\sigma_{\text{VBF}}(0, 1, 1)}{d\Phi} \\
 & + \left(\frac{\kappa_{2V}\kappa_\lambda\kappa_V}{27} - \frac{\kappa_{2V}\kappa_V^2}{27} + \frac{\kappa_\lambda^2\kappa_V^2}{90} - \frac{13\kappa_\lambda\kappa_V^3}{270} + \frac{\kappa_V^4}{27} \right) \cdot \frac{d\sigma_{\text{VBF}}(10, 1, 1)}{d\Phi} \\
 & + \left(-\frac{4\kappa_{2V}^2}{3} - \frac{4\kappa_{2V}\kappa_\lambda\kappa_V}{33} + \frac{60\kappa_{2V}\kappa_V^2}{11} + \frac{4\kappa_\lambda\kappa_V^3}{33} - \frac{136\kappa_V^4}{33} \right) \cdot \frac{d\sigma_{\text{VBF}}(1, 1.5, 1)}{d\Phi} \\
 & + \left(\frac{\kappa_{2V}^2}{3} + \frac{5\kappa_{2V}\kappa_\lambda\kappa_V}{66} - \frac{10\kappa_{2V}\kappa_V^2}{11} - \frac{5\kappa_\lambda\kappa_V^3}{66} + \frac{19\kappa_V^4}{33} \right) \cdot \frac{d\sigma_{\text{VBF}}(1, 3, 1)}{d\Phi} \\
 & + \left(-\frac{16\kappa_{2V}\kappa_\lambda\kappa_V}{33} + \frac{16\kappa_{2V}\kappa_V^2}{33} + \frac{16\kappa_\lambda\kappa_V^3}{33} - \frac{16\kappa_V^4}{33} \right) \cdot \frac{d\sigma_{\text{VBF}}(-5, 1, 0.5)}{d\Phi}
 \end{aligned} \tag{7.6}$$

The VBF HH parametrisation method was validated by comparing distributions of observables from existing fully simulated MC samples (not already used to construct the basis) with those from the corresponding parametrised samples. As illustrated in Figure 7.4, which shows the $m_{\gamma\gamma}$ and $m_{b\bar{b}\gamma\gamma}$ distributions for both the existing and parametrised VBF HH samples, the parametrisation successfully reproduces the fully simulated samples with an agreement on the expected yields better than 10%.

7.2.2 Background samples

The background samples include simulated events of both single Higgs boson production, where the Higgs decays to two photons, and continuum diphoton production.

The single Higgs samples, shared across all $H \rightarrow \gamma\gamma$ analyses, are simulated as outlined in Table 7.3 [202]. All main single Higgs production modes are considered, including ggF, VBF, WH, ZH (split into the two $q\bar{q} \rightarrow ZH$ and $g\bar{g} \rightarrow ZH$ components), $t\bar{t}H$, $b\bar{b}H$, and tH , with the latter separated into s- and t-channel contributions (tHq) and W-associated production (tHW). For both single Higgs and HH production, a Higgs boson mass of $m_H = 125 \text{ GeV}$ was assumed. The inclusive cross sections for these single Higgs processes are normalised to the most precise available theoretical predictions, as detailed in Section 1.3.1.

The continuum background samples include diphoton production in association with jets of any flavor ($\gamma\gamma$ +jets), two additional b-jets ($\gamma\gamma$ + $b\bar{b}$), and top quark pairs ($t\bar{t}\gamma\gamma$). The $\gamma\gamma$ +jets process was simulated using the SHERPA 2.2.4 generator [203] at NLO accuracy in QCD for diphoton production with up to one additional parton, while up to three additional partons are simulated at LO accuracy. This event generation follows the *five-*




Figure 7.4 – $m_{\gamma\gamma}$ (left) and $m_{b\bar{b}\gamma\gamma}$ (right) distributions obtained using simulated VBF HH events, with $\kappa_\lambda = 0, \kappa_{2V} = 0, \kappa_V = 1$ (top) and $\kappa_\lambda = 1, \kappa_{2V} = 2, \kappa_V = 1$ (bottom). The distributions obtained using the fully simulated samples (black) are compared to the corresponding parametrised samples (pink).

flavor scheme, treating b -quarks as massless⁵. To improve the modeling of the $\gamma\gamma+b\bar{b}$ process, which is the dominant background in the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis, a dedicated sample was generated using SHERPA 2.2.12 [203] at NLO accuracy for the production of the two photons and the two b -quarks, employing the *four-flavor scheme*, where the b -quark is treated as massive. In both the $\gamma\gamma$ +jets and $\gamma\gamma+b\bar{b}$ samples, additional jets are produced via parton showering, which is also managed by the SHERPA event generator. While the $\gamma\gamma$ +jets sample includes the $\gamma\gamma+b\bar{b}$ component (although simulated with lower precision), the dedicated sample significantly reduces the statistical uncertainty on the simulation of the $\gamma\gamma+b\bar{b}$ background, compared to the inclusive $\gamma\gamma$ +jets sample, despite having 60 times fewer events, thanks to requirements placed on the b -quarks at generator level. Since the normalisation of the continuum background is derived directly from data, the $\gamma\gamma$ +jets, $\gamma\gamma+b\bar{b}$, and $t\bar{t}\gamma\gamma$ samples are normalised using the cross sections provided by the event generators, (quoted in Table 7.3), rather than relying on more accurate theoretical predictions or experimental measurements.

Table 7.3 – Summary of the single Higgs boson background and continuum background samples used in the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis. The generator used in the simulation is also provided. The final three columns list the accuracy in QCD of the event generator, the order in QCD of the calculated cross sections for the corresponding background process, and the cross section values. For the $\gamma\gamma$ +jets and $\gamma\gamma+b\bar{b}$ samples, a generator-level requirement of $95 \leq m_{\gamma\gamma} \leq 170$ GeV is applied to enhance efficiency in the $H \rightarrow \gamma\gamma$ phase space.

Process	Generator	Accuracy	Order of σ calculation	$\sigma \times BR$ [fb]
ggF H	POWHEG NNLOPS [189–191, 204, 205]	NNLO	N ³ LO	110.2766
VBF H	POWHEG-BOX v2 [189–191, 206]	NLO	NNLO	8.6192
WH	POWHEG-BOX v2 [189–191, 207]	NLO	NNLO	3.1163
qq $\rightarrow$ ZH	POWHEG-BOX v2 [189–191, 207]	NLO	NNLO	1.7279
gg $\rightarrow$ ZH	POWHEG-BOX v2 [189–191, 207]	LO	NLO	0.2783
t $\bar{t}$ H	POWHEG-BOX v2 [189–191, 208]	NLO	NNLO	1.1511
b $\bar{b}$ H	POWHEG-BOX v2 [189–191, 209]	NLO	NNLO	1.1078
tHq	MADGRAPH5_AMC@NLO [68]	NLO	NLO	0.1751
t $\bar{t}$ HW	MADGRAPH5_AMC@NLO [68]	NLO	NLO	0.0344
$\gamma\gamma$ +jets ($95 \leq m_{\gamma\gamma} \leq 170$ GeV)	SHERPA v2.2.4 [203]	$\gamma\gamma + 1$ (NLO), 2, 3 (LO)	$\gamma\gamma + 1$ (NLO), 2, 3 (LO)	$51.823 \cdot 10^3$
$\gamma\gamma+b\bar{b}$ ($95 \leq m_{\gamma\gamma} \leq 170$ GeV)	SHERPA v2.2.12 [203]	NLO	NLO	$1.528 \cdot 10^3$
t $\bar{t}\gamma\gamma$	MADGRAPH5_AMC@NLO [68]	NLO	NLO	10.5581

7.3 Event selection

7.3.1 $b\bar{b}\gamma\gamma$ preselection

A set of preselection criteria is applied to target events consistent with the $HH \rightarrow b\bar{b}\gamma\gamma$ signal, characterised by two photons and two b -jets in the final state, compatible with

⁵The five-flavor and four-flavor schemes represent different ways of handling bottom quarks in MC event generation. In the five-flavor scheme, bottom quarks are treated as massless partons and included in the proton structure, allowing them to participate as initial-state partons in hard scattering processes. This massless approximation is accurate at high energies and simplifies theoretical predictions significantly. Conversely, in the four-flavor scheme, bottom quarks are considered massive and are not included in the proton PDFs. Consequently, they can only appear as products of the hard scattering process, not in the initial state, and must be explicitly included in matrix element calculations. The four-flavor scheme is preferred in scenarios where bottom quark mass effects are significant, such as in heavy-flavor emissions.

originating from Higgs boson decays.

An initial data-quality selection is applied to ensure that events were recorded under optimal detector conditions, when all ATLAS sub-detectors were fully operational and free from data integrity errors or corruption due to some hardware failures (for instance, HV trips in the EM Calorimeter).

Each event must contain at least two photons that pass the *loose* identification criteria (see Section 5.3) with transverse energy $E_T > 25$ GeV and pseudorapidity $|\eta| < 2.37$, excluding the transition region between the barrel and endcap of the EM calorimeter ($1.37 \leq |\eta| \leq 1.52$). The two highest- E_T photons fulfilling the *loose* identification criteria are selected as photon candidates and used to reconstruct the $H \rightarrow \gamma\gamma$ decay. To accurately determine the primary vertex (PV) among all reconstructed hard-scattering vertices in the event, the two photon candidates are used. In processes like single Higgs production via ggF, in which, for example, the $H \rightarrow \gamma\gamma$ decay results into two unconverted photons and there is little additional activity in the tracker, selecting the vertex based on the highest sum of transverse momenta of associated tracks often does not yield the correct vertex. Instead, a neural network-based algorithm [210] is employed, combining photon trajectory information and tracks associated with each candidate vertex. In simulated ggF H events, the algorithm selects the correct diphoton vertex within 0.3 mm of the pp collision axis in 79% of cases, compared to 56% using the standard method. For HH production and other Higgs production modes, where additional jet activity is expected, the accuracy improves, ranging from 84% to 97%. For these processes, the accuracy of the NN over the standard hardest-vertex method improves only of a few percentage points.

To reduce the contribution of backgrounds involving fake photons, the two photon candidates are required to be matched to the objects that fired the triggers, to pass the *tight* identification criteria, and to meet the *Loose* isolation requirements (see Section 5.3 and 5.4). Isolation variables are calculated using only the tracks associated with the diphoton PV. The leading (subleading) photon must also satisfy a relative transverse energy requirement with respect to the diphoton invariant mass ($m_{\gamma\gamma}$), specifically $E_T/m_{\gamma\gamma} > 0.35$ (0.25). Finally the invariant mass of the photon pair must fall within the window $105 \leq m_{\gamma\gamma} \leq 160$ GeV. These photon selection criteria are shared across all $H \rightarrow \gamma\gamma$ analyses [211].

To suppress background from single Higgs boson production via $t\bar{t}H$, as well as continuum $t\bar{t}\gamma\gamma$ processes where the top quark decay chain produces leptons (electrons or muons), events are rejected if they contain one or more isolated lepton candidates with $p_T > 10$ GeV that meet the *medium* identification criteria (see Section 5.5 and 5.6).

Reconstructed jets must have $p_T > 25$ GeV and fall within the rapidity range $|y| < 4.4$. Jets with $p_T < 60$ GeV and within the acceptance region of the ID ($|\eta| < 2.5$) must be selected as consistent with originating from the PV using the *tight* operating point of the Jet Vertex Tagger (JVT) [212], to reject pileup jets contaminations.

Events must contain at least two *central* jets (i.e., jets within $|\eta| < 2.5$) that are eligible for flavor tagging. In order to target $H \rightarrow b\bar{b}$ decays, events must contain exactly two *b*-tagged jets, identified using the 77% efficiency WP of the DL1r tagger (see Section 6.2). These two *b*-jets are used to reconstruct the $H \rightarrow b\bar{b}$ candidate. Among the

selected b -jets, the leading jet is identified by the highest pseudo-continuous b -tagging score, while the subleading jet has the second-highest score. If there are ties in the b -tagging scores, jets are further ranked by their p_T . Additional jets are ranked similarly, first by their pseudo-continuous b -tagging score and then by p_T in case of ties. Events with six or more central jets are rejected to reduce background from $t\bar{t}H$ and $t\bar{t}\gamma\gamma$ events, where top quarks decay hadronically. Events with at least four jets may originate from HH production via VBF, where the hadronization of the two valence quarks from the hard scattering process produces two high-energy forward jets. These jets, referred to as *VBF jets*, provide a crucial handle for isolating HH events from background processes. To enhance the selection efficiency of VBF HH events, the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis employs a BDT-based algorithm, known as the *VBF jet tagger*, which identifies the two VBF jets among the remaining jets that are not already selected as the two b -jets from the $H \rightarrow b\bar{b}$ decay. The VBF jet tagger is trained on all possible jet pairs in an event, excluding the b -jets used to form the $H \rightarrow b\bar{b}$ candidate. The input features, listed in Table 7.4, include both jet-level (j) variables, such as p_T , η , and angular separation from the $b\bar{b}\gamma\gamma$ system, and jet-pair (j_1j_2) variables, like invariant mass, angular separations, and distances from the $b\bar{b}\gamma\gamma$ system. The BDT is trained on simulated SM VBF HH events, where jet pairs matched to the scattered quarks at the truth level are treated as signal, and all other possible jet pairs as background. After training, the VBF jet tagger is applied to all jet pair combinations in both data and simulated events: the jet pair with the highest tagger score identifies the VBF jet candidates. The VBF jet tagger correctly selects the VBF jet pair in 95% of simulated SM VBF HH events, compared to a (lower) 87% efficiency, reached when using a simpler method based on the selecting jet pair with the highest invariant mass.

Table 7.4 – Input variables used to train the VBF jet tagger algorithm.

Category	Variable
Jet-level kinematics	$p_T(j)$
	$\eta(j)$
	$\Delta\eta(j, b\bar{b}\gamma\gamma)$
	$\Delta R(j, b\bar{b}\gamma\gamma)$
Jet pair quantities	$m_{j_1j_2}$
	$\Delta\eta(j_1, j_2)$
	$\Delta\eta(j_1j_2, b\bar{b}\gamma\gamma)$
	$\Delta R(j_1j_2, b\bar{b}\gamma\gamma)$
Kinematics of the $j_1j_2 + b\bar{b}\gamma\gamma$ system	$p_T(j_1j_2b\bar{b}\gamma\gamma)$
	$\eta(j_1j_2b\bar{b}\gamma\gamma)$
	$m_{j_1j_2b\bar{b}\gamma\gamma}$
Event-level quantities	Scalar sum of all the jets' momenta, H_T

The $b\bar{b}\gamma\gamma$ preselection selects SM ggF HH and VBF HH events with an expected efficiency of 13% and 9% respectively. The number of events selected in data in this inclusive signal region is 1874 in the full Run 2 dataset.

7.3.2 Categorisation

Events entering the signal region of the analysis, described in Section 7.3.1, are split into mutually exclusive categories. The categorisation of events is one of the key ingredients of the analysis. Together with enhancing the signal over background ratio, an accurately optimised categorisation significantly impacts the analysis sensitivity to both SM HH production and to deviations of the coupling modifiers κ_λ and κ_{2V} from their SM values. The event categorisation is optimised on the two dominant HH production modes: via ggF and via VBF. It is based on the modified invariant mass of the $b\bar{b}\gamma\gamma$ system, $m_{b\bar{b}\gamma\gamma}^*$, defined as $m_{b\bar{b}\gamma\gamma}^* = m_{b\bar{b}\gamma\gamma} - (m_{b\bar{b}} - 125 \text{ GeV}) - (m_{\gamma\gamma} - 125 \text{ GeV})$ (where 125 GeV approximates the mass of the Higgs boson). With respect to $m_{b\bar{b}\gamma\gamma}$, $m_{b\bar{b}\gamma\gamma}^*$ is less sensitive to the potential degradation of diphoton and dijet energy resolution due to detector effects. As a result, using $m_{b\bar{b}\gamma\gamma}^*$ enhances the resolution of the four-object final state compared to the traditional $b\bar{b}\gamma\gamma$ invariant mass, as demonstrated in Figure 7.5.




Figure 7.5 – Distribution of the $b\bar{b}\gamma\gamma$ invariant mass, $m_{b\bar{b}\gamma\gamma}$ (dashed lines), compared with the the modified $b\bar{b}\gamma\gamma$ invariant mass, $m_{b\bar{b}\gamma\gamma}^*$ (solid lines), for simulated SM ggF HH and VBF HH events.

The $m_{b\bar{b}\gamma\gamma}^*$ variable acts as a proxy for the HH invariant mass (m_{HH}). As discussed in Section 2.1.2 and 2.1.3, the m_{HH} variable is highly sensitive to variations in the coupling modifiers affecting HH production. For instance, non-SM values of κ_λ favor softer m_{HH} spectra, typically peaking below 300 GeV, for both ggF and VBF production. On the other hand, deviations in κ_{2V}/κ_V^2 from unity lead to highly boosted HH production in the VBF mode, resulting in harder m_{HH} spectra, enhancing the m_{HH} region above 1 TeV. Finally, in the SM, HH production yields a moderately hard m_{HH} spectrum, peaking around 400 GeV for ggF and 350 GeV for VBF production. This behavior is mirrored

by the $m_{b\bar{b}\gamma\gamma}^*$ variable, whose distributions for simulated ggF HH and VBF HH events in both the SM scenario and various combinations of anomalous κ_λ and κ_{2V} values are shown in Figure 7.6. Hence, to optimise the analysis sensitivity to non-SM κ_λ scenarios




Figure 7.6 – Distributions of the modified four-object invariant mass $m_{b\bar{b}\gamma\gamma}^*$, after applying the analysis preselection, for simulated ggF HH (left) and VBF HH (right) events for several values of κ_λ and κ_{2V} . The threshold at $m_{b\bar{b}\gamma\gamma}^* = 350$ GeV used to define the Low Mass and High Mass regions is also shown.

and simultaneously maintain a strong sensitivity to both SM HH production and non-SM VBF HH production with κ_{2V} deviating from unity, two distinct $m_{b\bar{b}\gamma\gamma}^*$ regions are defined: the *Low Mass* region and the *High Mass* region. The Low Mass region corresponds to $m_{b\bar{b}\gamma\gamma}^* \leq 350$ GeV, while the High Mass region is defined by $m_{b\bar{b}\gamma\gamma}^* > 350$ GeV. The categories are ultimately constructed with the help of Boosted Decision Tree (BDT)-based algorithms, trained in each mass region to separate HH signal events from single Higgs backgrounds and continuum background from $\gamma\gamma$ +jets and $t\bar{t}\gamma\gamma$ processes. Different HH benchmarks are used as signals for the BDTs in the two mass regions:

- **High Mass Region.** The BDT was trained using SM ggF HH and VBF HH events, along with anomalous VBF HH samples corresponding to five combinations of the coupling modifiers $(\kappa_\lambda, \kappa_{2V}, \kappa_V)$: $(0, 1, 1)$, $(10, 1, 1)$, $(1, 1.5, 1)$, $(1, 3, 1)$, and $(-5, 1, 0.5)$. Together with the SM sample, these five variations serve as a basis for parametrising the VBF HH process as a function of the coupling modifiers $(\kappa_\lambda, \kappa_{2V}, \kappa_V)$ (as described in Section 7.2.1). Therefore, any $(\kappa_\lambda, \kappa_{2V}, \kappa_V)$ scenario can be modeled using a linear combination of the six VBF HH samples used in the training.
- **Low Mass Region.** The BDT was trained on a mixture of ggF HH events with $\kappa_\lambda = 5.6$ (reweighted from the SM sample) and $\kappa_\lambda = 10$, alongside the same anomalous VBF HH samples used for the High Mass BDT. The inclusion of this sample with $\kappa_\lambda = 5.6$ improved the expected constraints on both the HH signal strength and the coupling modifier κ_λ . Similar improvements were observed by repeating the training with ggF HH samples using slightly different κ_λ values (for instance $\kappa_\lambda = 6$ and $\kappa_\lambda = 7$). However, the $\kappa_\lambda = 5.6$ value was chosen as it lies one unit below the

upper exclusion limit on κ_λ set by the previous $HH \rightarrow b\bar{b}\gamma\gamma$ analysis using Run 2 data [12].

Each training sample is assigned a weight that reflects its importance in the BDT training. These weights, along with the BDT hyperparameters, are optimised using a bayesian algorithm [213] that maximises the expected number-counting significance, Z [70]. The significance is evaluated based on the expected yields of a benchmark HH signal and backgrounds from single Higgs production and continuum processes, within the $m_{\gamma\gamma}$ range of $120 < m_{\gamma\gamma} < 130$ GeV, as explained below. This optimisation procedure also determines the thresholds in the BDT discriminant for the High Mass and Low Mass regions, which are used to define the final categories. Optimising the training weights has a substantial impact, enhancing the analysis sensitivity to the HH signal strength and to κ_λ and κ_{2V} by up to 10%. In contrast, tuning the BDT hyperparameters has minimal effect, with changes in the expected constraints to μ_{HH} , κ_λ , and κ_{2V} staying below 1%. To minimise overtraining, an *early stopping* criterion was implemented in the optimisation process. This criterion applies both to the training stage and to the entire optimisation chain, including the tuning of thresholds in the BDT discriminants defining the analysis categories. At each iteration of the optimisation algorithm, the binary loss function adopted for the BDT training and the number-counting significance are evaluated on independent MC events, forming a validation dataset; the procedure halts if there is no improvement over 10 consecutive iterations. Finally, the full analysis is performed relying a third, independent set of MC events (termed the test dataset) that remains untouched throughout the optimisation phases.

Both BDTs are trained using the same set of input variables, detailed in Table 7.5. These input features include kinematic properties of the two photons and two b -jets in the final state, as well as b -tagging information for the two b -jets. The kinematic variables and the b -tagging status of the third and fourth jets (if present) are also included as input features for the training: these variables are particularly useful for isolating VBF HH events, which are typically characterized by four jets in the final state. Additional VBF-specific kinematic variables, such as the invariant mass and angular separation of the two VBF jets (identified using the VBF jet tagger), are included to enhance the BDTs' performance in isolating VBF HH production. The scalar sum of all jets' transverse momenta, H_T , is another input feature. This, along with the magnitude and direction of the missing transverse momentum (E_T^{miss} and ϕ^{miss}), as well as the *single topness* variable, is crucial for rejecting the $t\bar{t}H$ background. The single topness, defined as

$$\chi_{Wt} = \min_{j_1, j_2, j_3} \sqrt{\left(\frac{m_{j_1 j_2} - m_W}{m_W}\right)^2 + \left(\frac{m_{j_1 j_2 j_3} - m_t}{m_t}\right)^2}, \quad (7.7)$$

where the minimum is taken over all possible three-jet combinations, quantifies the likelihood that any combination of three jets in the event originate from a $t \rightarrow Wb \rightarrow q\bar{q}'b$ decay. The training also relies on event-level variables involving the two photons and two b -jets used to reconstruct the $H \rightarrow \gamma\gamma$ and $H \rightarrow b\bar{b}$ decays, such as the $m_{b\bar{b}\gamma\gamma}^*$ variable and angular separations between the photons and b -jets ($\Delta R(\gamma_1, \gamma_2)$ and $\Delta R(b_1, b_2)$), which are crucial for distinguishing the HH signal from backgrounds and for differentiating

between the ggF HH and VBF HH production modes. In particular, as explained earlier, the $m_{b\bar{b}\gamma\gamma}^*$ variable is also crucial for distinguishing anomalous κ_λ and κ_{2V} scenarios from SM HH production. Finally, three event-shape observables are employed: the transverse sphericity $S_\perp$ [214], planar flow P_f [215], and the transverse momentum balance of the two photons and two b -jets, defined as:

$$p_T^{\text{balance}} = \frac{|p_T(\gamma_1) + p_T(\gamma_2) + p_T(b_1) + p_T(b_2)|}{|p_T(\gamma_1)| + |p_T(\gamma_2)| + |p_T(b_1)| + |p_T(b_2)|}. \quad (7.8)$$

These variables capture the topology of the $b\bar{b}\gamma\gamma$ system, with back-to-back photons and b -jets expected in $HH \rightarrow b\bar{b}\gamma\gamma$ events, while background events tend to have a more isotropic distribution of final-state particles.

As previously mentioned, the final categories are defined by applying cuts to the BDT

Table 7.5 – Input variables used for training the BDT-based algorithm in both the High and Low mass regions.

Type	Variable	Description
Photon-related kinematic variables	$p_T/m_{\gamma\gamma}$	Transverse momentum of the leading and subleading photon, scaled with the di-photon invariant mass.
	η and ϕ	Angular variables of the leading and subleading photons.
Jet-related kinematic variables	b -tag status of the two b -jets	Pseudo-continuous b -tagging score of the two selected b -jets.
	p_T, η , and ϕ of the two b -jets	Transverse momentum and angular variables of the two selected b -jets.
	$p_T^{b\bar{b}}, \eta^{b\bar{b}}$, and $\phi^{b\bar{b}}$	Transverse momentum and angular variables of the di-jet object.
	$m_{b\bar{b}}$	Invariant mass of the two selected b -jets.
	b -tag status of the 3 rd and 4 th jets	Pseudo-continuous b -tagging score of the 3 rd and 4 th -ranked jets.
	p_T, η , and ϕ of the 3 rd and 4 th jets	Transverse momentum and angular variables of the 3 rd and 4 th -ranked jets.
	H_T	Scalar sum of the transverse momenta of all the jets reconstructed in the event.
	χ_{Wt}	Single topness.
VBF-targeting variables	m_{jj}^{VBF} and $\Delta\eta(j_1^{\text{VBF}}, j_2^{\text{VBF}})$	Invariant mass and angular separation of the two VBF jets.
Missing transverse momentum-related variables	E_T^{miss} and ϕ^{miss}	Missing transverse momentum and its azimuthal direction.
$b\bar{b}\gamma\gamma$ -related variables	$m_{b\bar{b}\gamma\gamma}^*$	Modified invariant mass of the $b\bar{b}\gamma\gamma$ system.
	$\Delta R(\gamma_1, \gamma_2)$ and $\Delta R(b_1, b_2)$	Angular separation between the two photons or the two b -jets.
Event-shape variables	$S_\perp, P_f$, and p_T^{balance}	Transverse sphericity, planar flow, and transverse momentum balance of the $b\bar{b}\gamma\gamma$ system.

discriminants in both the High Mass and Low Mass regions. Specifically, three categories are defined in the High Mass region (referred to as *High Mass i* , with $i = 1, 2, 3$) and four

categories in the Low Mass region (referred to as *Low Mass i* , with $i = 1, 2, 3, 4$). Higher category indices correspond to events with higher BDT discriminant values, which are more signal-like. Events with BDT scores below the threshold for the lowest-index category are discarded. Of the 1874 data events passing the $b\bar{b}\gamma\gamma$ preselection criteria (outlined in Section 7.3.1), only 340 are classified into one of these seven categories. The BDT score thresholds for each mass region are optimised to maximise the combined number-counting significance Z across all categories, defined as:

$$Z = \sqrt{\sum_{i=1}^N Z_i^2}, \quad (7.9)$$

where N represents the number of categories, which is 3 for the High Mass region and 4 for the Low Mass region. The number-counting significance Z_i for each category is calculated as:

$$Z_i = \sqrt{2 \cdot [(S_i + B_i) \cdot \log(1 + S_i/B_i) - S_i]}, \quad (7.10)$$

where S_i and B_i indicate the expected signal and background yields in each category i , in the diphoton invariant mass range $120 < m_{\gamma\gamma} < 130$ GeV. The expected background yields B_i include all single Higgs processes and the continuum background, estimated from the $\gamma\gamma$ +jets MC sample. The overall normalisation of the $\gamma\gamma$ +jets process after applying the $b\bar{b}\gamma\gamma$ preselection is adjusted to match data in the sideband regions ($105 < m_{\gamma\gamma} < 120$ GeV and $130 < m_{\gamma\gamma} < 160$ GeV) to account for the γj and jj contributions to the continuum background, which are not included in the $\gamma\gamma$ +jets simulation. As will be explained in Section 7.4.2, the reducible γj and jj backgrounds are estimated to contribute approximately as 20% of the continuum background, after applying the $b\bar{b}\gamma\gamma$ preselection. The expected signal yields S_i are calculated using different HH benchmark scenarios depending on the mass region, optimising for either SM-like or anomalous HH production. In the Low Mass region, which is more sensitive to non-SM κ_λ values, the number-counting significances are computed using ggF HH events with $\kappa_\lambda = 5.6$ and VBF HH events with $\kappa_\lambda = 10$. In contrast, the High Mass region prioritises sensitivity to SM HH production, so the signal yields S_i are evaluated using SM HH events.

The BDT discriminants for the Low and High Mass regions, displaying data, HH signals, and background distributions, are presented in Figure 7.7a and Figure 7.7b, respectively. The category boundaries are provided in Table 7.6.

The optimization of the BDT score thresholds is carried out by imposing a requirement to have at least two expected $\gamma\gamma$ +jets background events in each analysis category within the region $120 \leq m_{\gamma\gamma} \leq 130$. This constraint is necessary for modeling the continuum background in each category. Given the limited statistics of the $HH \rightarrow b\bar{b}\gamma\gamma$ final state, the boundaries of the most sensitive categories (High Mass 3 and Low Mass 4) are driven by this constraint on the number of expected $\gamma\gamma$ +jets events.

Figure 7.8 displays the expected purity of each analysis category within the $120 < m_{\gamma\gamma} < 130$ GeV range, using SM HH production as the signal. It also shows the contributions




Figure 7.7 – Distributions of the BDT discriminants for the Low Mass region (left) and High Mass region (right), extracted for ggF HH and VBF HH signal events (both in the SM and for anomalous κ_λ and κ_{2V} scenarios), single Higgs boson backgrounds, and $\gamma\gamma$ +jets events. Distributions extracted from data in the $m_{\gamma\gamma}$ sidebands ($105 < m_{\gamma\gamma} < 120$ GeV and $130 < m_{\gamma\gamma} < 160$ GeV) are also displayed. All distributions are normalised to unity. The vertical dashed lines indicate the thresholds defining the analysis categories. All events with BDT scores below the lowest threshold (marked by the thick dashed line) are discarded.

Table 7.6 – Definition of the seven analysis categories, established by applying cuts on the $m_{b\bar{b}\gamma\gamma}^*$ variable and BDT discriminants.

Category	Mass region	BDT discriminant
High Mass 1	$m_{b\bar{b}\gamma\gamma}^* > 350$ GeV	$0.545 < \text{BDT output} < 0.830$
High Mass 2	$m_{b\bar{b}\gamma\gamma}^* > 350$ GeV	$0.830 < \text{BDT output} < 0.905$
High Mass 3	$m_{b\bar{b}\gamma\gamma}^* > 350$ GeV	$\text{BDT output} > 0.905$
Low Mass 1	$m_{b\bar{b}\gamma\gamma}^* \leq 350$ GeV	$0.545 < \text{BDT output} < 0.785$
Low Mass 2	$m_{b\bar{b}\gamma\gamma}^* \leq 350$ GeV	$0.785 < \text{BDT output} < 0.890$
Low Mass 3	$m_{b\bar{b}\gamma\gamma}^* \leq 350$ GeV	$0.890 < \text{BDT output} < 0.950$
Low Mass 4	$m_{b\bar{b}\gamma\gamma}^* \leq 350$ GeV	$\text{BDT output} > 0.950$

from the main background processes, including single Higgs production via ggF, $t\bar{t}H$, and ZH, as well as the $\gamma\gamma$ +jets sample (representing the dominant contribution to the continuum background). The High Mass 3 category is the most enriched in SM HH signal, but its expected purity remains quite low, at just 20%.




Figure 7.8 – Expected composition of each analysis categories, in terms of HH signal (assuming SM HH production), single Higgs, and $\gamma\gamma$ +jets backgrounds.

Figure 7.9 illustrates the relative importance of ggF HH and VBF HH processes in each category, comparing the expected signal composition for SM HH production (Figure 7.9a), as well as scenarios with $\kappa_\lambda = 10$ (Figure 7.9b) and $\kappa_{2V} = 3$ (Figure 7.9c). In the SM scenario, where the ratio of VBF HH to ggF HH cross sections is at a minimum, the contribution of VBF HH events is negligible compared to the dominant ggF HH production mode, even in the High Mass 3 category. For $\kappa_\lambda = 10$, VBF HH contributions become more significant, particularly in the Low Mass categories optimised for anomalous κ_λ values. When setting $\kappa_{2V} = 3$, which does not affect the ggF HH process but strongly enhances the VBF HH cross section, the VBF HH signal becomes highly dominant with respect to ggF HH. Although this analysis benefits from optimising for VBF HH production - by incorporating VBF HH samples and VBF-targeting variables in BDT training and evaluating number-counting significances accounting for VBF HH signals - no dedicated VBF HH-enriched category is defined. This decision is based on the fact that separating into ggF-enriched and VBF-enriched categories does not improve sensitivity to non-SM κ_{2V} values, as the VBF HH signal is already dominant with respect to ggF HH production in these scenarios. On the other hand, while defining a VBF HH-enriched category could help to independently constrain SM ggF HH and VBF HH production, the sensitivity to SM VBF HH remains extremely low compared to ggF HH. In fact, only one VBF HH event in the $b\bar{b}\gamma\gamma$ final state is expected within the entire Run 2 dataset, according to the SM. Consequently, this analysis focuses on maximising overall sensitivity to SM HH production and to anomalous κ_λ and κ_{2V} values by defining categories that isolate HH events from backgrounds, rather than trying to explicitly trying separate ggF HH and VBF HH signals.




Figure 7.9 – Fraction of the two main HH production modes (via ggF and VBF) contributing to the overall signal yields in each analysis category in the (a) SM, (b) $\kappa_\lambda = 10$, and (c) $\kappa_{2V} = 3$ hypotheses.

Figures 7.10 and 7.11a illustrate the expected efficiency times acceptance for ggF HH and VBF HH production as a function of κ_λ in each analysis category. For both production modes, efficiency peaks around $\kappa_\lambda = 1$ (the SM value) in the High Mass categories, which are optimised to target SM HH production. In contrast, the Low Mass categories achieve maximum efficiency for anomalous κ_λ values. Finally, Figure 7.11 displays the expected efficiency times acceptance for VBF HH production as a function of κ_{2V} . The High Mass categories, particularly High Mass 3, have the strongest sensitivity to the VBF HH signal, both for the SM and for anomalous κ_{2V} values. The efficiency for VBF HH in the Low Mass categories is significantly lower, as expected.




Figure 7.10 – The acceptance times efficiency for the signal $ggF \text{ } HH$ process as a function of the coupling modifier κ_λ in each analysis category. The other coupling modifiers affecting $ggF \text{ } HH$ production are fixed to their SM predictions. The bands indicate the simulation statistical uncertainty.




Figure 7.11 – The acceptance times efficiency for the signal $VBF \text{ } HH$ process as a function of the coupling modifier κ_λ (left) and κ_{2V} (right) in each analysis category. The other coupling modifiers affecting $VBF \text{ } HH$ production (including κ_{2V} and κ_λ respectively) are fixed to their SM predictions. The bands indicate the simulation statistical uncertainty.

7.4 Signal and background modelling

The presence of a diphoton resonance in the final state is crucial in determining the choice of the final discriminant variable. As already mentioned, the diphoton invariant mass ($m_{\gamma\gamma}$) effectively distinguishes between the di-Higgs signal and the primary background source, corresponding to continuum diphoton production with additional jets. The signal manifests as a sharp peak around $m_H \approx 125$ GeV, whereas the background follows a smoothly decreasing distribution with no peak. Therefore, signal extraction involves modeling both the signal and background contributions in the $m_{\gamma\gamma}$ spectrum within each analysis category.

7.4.1 Signal and resonant background model

Both the HH signal and the single Higgs boson background involve the Higgs decaying into two photons, leading to a narrow resonance in the $m_{\gamma\gamma}$ distribution centered around the Higgs mass ($m_H \approx 125$ GeV). As a result, the $m_{\gamma\gamma}$ distribution for all resonant processes, including HH signals and single Higgs backgrounds, is modeled using a *Double-Sided Crystal Ball* (DSCB) distribution [210, 216]. This model combines a Gaussian core with asymmetric power-law tails and is characterized by six parameters: μ_{CB} , σ_{CB} , α_{Lo} , n_{Lo} , α_{Hi} , and n_{Hi} . The DSCB functional form is given by:

$$f_{\text{DSCB}}(x) = N \cdot \begin{cases} A_{\text{Lo}} \cdot \left(B_{\text{Lo}} - \frac{x - \mu_{\text{CB}}}{\sigma_{\text{CB}}} \right)^{-n_{\text{Lo}}} & \text{for } \frac{x - \mu_{\text{CB}}}{\sigma_{\text{CB}}} < -\alpha_{\text{Lo}} \\ A_{\text{Hi}} \cdot \left(B_{\text{Hi}} + \frac{x - \mu_{\text{CB}}}{\sigma_{\text{CB}}} \right)^{-n_{\text{Hi}}} & \text{for } \frac{x - \mu_{\text{CB}}}{\sigma_{\text{CB}}} > \alpha_{\text{Hi}} \\ \exp \left(-\frac{(x - \mu_{\text{CB}})^2}{2\sigma_{\text{CB}}^2} \right) & \text{for } -\alpha_{\text{Lo}} \leq \frac{x - \mu_{\text{CB}}}{\sigma_{\text{CB}}} \leq \alpha_{\text{Hi}} \end{cases} \quad (7.11)$$

Here, A_{Lo} , B_{Lo} , A_{Hi} , and B_{Hi} are normalization constants determined by:

$$A_k = \left(\frac{n_k}{|\alpha_k|} \right)^{n_k} e^{-\frac{\alpha_k^2}{2}}, \quad B_k = \frac{n_k}{|\alpha_k|} - |\alpha_k|, \quad (7.12)$$

where $k = \text{Lo}$ or $k = \text{Hi}$. These parameters ensure a smooth transition between the Gaussian core and the power-law tails.

The Gaussian core of the DSCB is defined by its mean (μ_{CB}) and width (σ_{CB}), while the parameters α_{Lo} and n_{Lo} (α_{Hi} and n_{Hi}) describe the low (high) energy tails. The α_k values represent the transition points, in units of σ_{CB} , between the Gaussian core transitions and the power-law tails, whose steepness is governed by the n_k parameters. The Gaussian core captures the mass of well-reconstructed $H \rightarrow \gamma\gamma$ candidates, while the tails account for photon energy miscalibration effects [217].

The central values of the DSCB shape parameters in each category are derived from a fit to simulated SM signal events, weighted according to their cross sections, efficiencies, and acceptances. Figure 7.12 shows the resonant $H \rightarrow \gamma\gamma$ model for the two most relevant categories: High Mass 3 and Low Mass 4.

The same functional form and shape parameters are also used to describe the single Higgs background processes. This assumption was validated by performing signal-plus-




Figure 7.12 – Diphoton invariant mass distributions for SM ggF HH and VBF HH simulated events, after applying the selection requirements for the (a) High Mass 3 and (b) Low Mass 4 categories. The results of the fit of the DSCB function to the $m_{\gamma\gamma}$ distributions are also shown (blue lines). The bottom panels display the residuals of the fit.

background fits on signal-plus-background templates built from MC samples, where negligible bias in single Higgs yields was observed when comparing MC distributions with the functional forms.

Moreover, the $m_{\gamma\gamma}$ distribution for HH signals shows negligible dependence on the coupling modifiers κ_λ and κ_{2V} . The behavior of the key shape parameters, μ_{CB} and σ_{CB} , across a range of κ_λ values within $-10 \leq \kappa_\lambda \leq 10$ exhibited no significant deviations, as demonstrated in Figure 7.13. Similarly, the κ_{2V} dependence within the range $-5 \leq \kappa_{2V} \leq 5$ was found to have no critical impact on μ_{CB} or σ_{CB} . Indeed, the effect of κ_λ and κ_{2V} dependence on the resonant shape parameters was determined to have a negligible influence on the expected 95% confidence intervals for both coupling modifiers, with an estimated impact below 0.5%.




Figure 7.13 – Shape parameters μ_{CB} (left) and σ_{CB} (right) as a function of κ_λ in the High Mass 3 category. The shaded bands represent the statistical uncertainties on the shape parameters extracted from the fit.

To construct the final statistical model, the resonant $m_{\gamma\gamma}$ shapes for both the HH signals and the single Higgs background processes must be correctly normalised using the expected yields for each process i in each analysis category c , N_c^i . The expected yield for process i in category c is given by:

$$N_c^i = L \cdot \sigma(i) \cdot \mathcal{BR}(i) \cdot (A \times \epsilon)_c^i, \quad (7.13)$$

where:

- L is the total integrated luminosity ($L = 140 \text{ fb}^{-1}$);
- $\sigma(i)$ represents the cross section of the considered resonant process;
- $\mathcal{BR}(i)$ is the branching fraction for the HH decay to two bottom quarks and two photons in case i stands for either the ggF HH or the VBF HH process, which is $\mathcal{BR}(\text{HH} \rightarrow b\bar{b}\gamma\gamma) = 0.26\%$ (see Section 2.1.4). On the other hand, single Higgs production processes, $\mathcal{BR}(i)$ represents the branching fraction for $H \rightarrow \gamma\gamma$, equal to 0.227% (as detailed in Section 1.3.2);
- $(A \times \epsilon)_c^i$ denotes the detector acceptance times the selection efficiency for the i^{th} resonant process in the c^{th} category, determined using MC simulations.

To investigate HH production for anomalous values of κ_λ and κ_{2V} , the expected signal yields are parameterised as functions of these coupling modifiers, as described in Section 7.2.1. This allows to construct a statistical model that is fully parametric in terms of these coupling modifiers. Since the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis primarily targets anomalous κ_λ and κ_{2V} scenarios, the other coupling modifiers that affect HH production (namely, κ_t and κ_V), are set to unity. Any impact on single Higgs background yields or on the $H \rightarrow \gamma\gamma$ branching fraction from κ_λ (due to electroweak corrections involving the trilinear Higgs boson vertex [22]) is neglected.

7.4.2 Continuum background model

As mentioned above, the continuum background consists mainly of diphoton plus jets ($\gamma\gamma$ +jets), with smaller contributions from γj and jj events, where one or more jets are misidentified as photons. Thanks to stringent photon identification and isolation requirements, the contribution from these reducible backgrounds is relatively small. The γj and jj background contributions, after the $b\bar{b}\gamma\gamma$ preselection and in each analysis category, are estimated using a data-driven two-dimensional matrix method [218]. This method aims to estimate the composition of $\gamma\gamma$, γj , and jj events in the signal region, where both reconstructed photons are required to be tight and isolated. To do this, 15 background-enriched control regions are used, where one or both photons fail the tight identification or isolation criteria but pass a looser identification requirement. The matrix method exploits the correlations between $\gamma\gamma$, γj , and jj events among the control and signal regions, and predict their yields in the signal region by solving a system of linear equations. After the $b\bar{b}\gamma\gamma$ preselection, the estimated fractions in the signal region are: $83.1 \pm 2.4\%$ (stat.) $^{+3.6\%}_{-0.6\%}$ (syst.) for $\gamma\gamma$, $16.4 \pm 2.4\%$ (stat.) $^{+0.6\%}_{-3.7\%}$ (syst.) for γj , and $0.5 \pm 0.2\%$ (stat.) $^{+0.1\%}_{-0.0\%}$ (syst.) for jj . Statistical uncertainties are derived from

data, while systematic uncertainties arise from varying the looser photon identification criteria in the control regions. Therefore, most of the continuum background consists of real photons. Comparisons of the $m_{\gamma\gamma}$ distribution between data, the $\gamma\gamma$ template (from $\gamma\gamma$ +jets simulated events), and the γj and jj templates (from control regions in data) showed compatible shapes across categories within statistical uncertainties. Within the $HH \rightarrow b\bar{b}\gamma\gamma$ phase space, the $\gamma\gamma$ +jets background includes exactly two b -tagged jets. These two jets can be either true b -jets, or c - or light-jets misidentified as b -jets. Truth-level information from the simulated samples was used to determine that only 54% of these events consist of two true b -jets. However, this fraction increases in the more sensitive analysis categories, mostly thanks to the use of b -tagging information in the BDT training. The $m_{\gamma\gamma}$ distribution in the data sidebands was compared to both the inclusive $\gamma\gamma$ +jets sample (which contains both true b -jets and mistagged c - or light-jets) and the $\gamma\gamma+b\bar{b}$ sample, which contains only the irreducible component where the two selected jets are true b -jets. These comparisons, presented in Figure 7.14 for the High Mass 3 and Low Mass 4 categories, show that both samples are consistent with each other and with the data, after the $b\bar{b}\gamma\gamma$ preselection and within each category. Therefore, the $\gamma\gamma+b\bar{b}$ sample can reliably model the $m_{\gamma\gamma}$ distribution of the continuum background.




Figure 7.14 – Diphoton invariant mass distribution in the data sidebands ($105 < m_{\gamma\gamma} < 120$ GeV and $130 < m_{\gamma\gamma} < 160$ GeV), compared to the $\gamma\gamma$ +jets and $\gamma\gamma+b\bar{b}$ samples, within the (a) High Mass 3 and (b) Low Mass 4 categories.

The continuum background is ultimately modeled using a functional form that is chosen as a compromise between having a small bias in each category and having a low number of degrees of freedom. The bias is evaluated through a signal-plus-background fit on a background-only template, which is constructed using the high-statistics and high-efficiency SHERPA 2.2.12 $\gamma\gamma+b\bar{b}$ sample. To achieve a more accurate representation of the continuum background, this sample is normalised to match the data in the sidebands. The signal yield is computed from the signal-plus-background fit at different Higgs boson mass values within the range $123 \leq m_H \leq 127$ GeV, in increments of 0.5

GeV. The largest signal yield in this window is termed the *spurious signal* (N_{sp}) [219], and corresponds to the bias of the chosen functional form. For an analytical function to be considered valid for modeling the continuum background, it must pass the *spurious signal test* by meeting one of the following criteria:

- $|N_{\text{sp}}| < 0.1 \cdot N_{S,\text{exp}}$, where $N_{S,\text{exp}}$ is the expected signal yield for SM HH events (evaluated using simulated samples);
- $|N_{\text{sp}}| < 0.2 \cdot \sigma_{\text{bkg}}$, where σ_{bkg} is the statistical uncertainty on the fitted number of signal events from the signal-plus-background fit on the background-only template.

If no functions satisfy these criteria, a modified *relaxed spurious signal test* is applied. In this case, a new variable, ζ_{sp} , is introduced:

$$\zeta_{\text{sp}} = \begin{cases} N_{\text{sp}} + 2\Delta_{\text{MC}} & \text{if } N_{\text{sp}} + 2\Delta_{\text{MC}} < 0 \\ N_{\text{sp}} - 2\Delta_{\text{MC}} & \text{if } N_{\text{sp}} - 2\Delta_{\text{MC}} > 0, \\ 0 & \text{otherwise} \end{cases} \quad (7.14)$$

where Δ_{MC} accounts for the statistical uncertainty in N_{sp} due to limited MC statistics in the background-only template. The spurious signal test is then applied using ζ_{sp} instead of N_{sp} . When multiple functions pass either the standard or relaxed test, the one with the fewest degrees of freedom is selected.

Ultimately, the exponential function was selected to model the $m_{\gamma\gamma}$ distribution in each category. The spurious signal values obtained for the exponential function in each category are shown in Table 7.7. In the final statistical model, both the normalization and shape parameters of the continuum background are left free to float and are determined directly from the data.

Table 7.7 – The spurious signal values for each category, corresponding to the final choice of using the exponential function $\exp(a \cdot m_{\gamma\gamma})$ to model the continuum background, are provided. For comparison, the expected yields of SM HH signal events are also reported.

Category	Spurious signal	Expected SM HH yield
High Mass 1	0.491	0.262
High Mass 2	-0.511	0.197
High Mass 3	-0.707	0.852
Low Mass 1	0.168	0.049
Low Mass 2	-0.179	0.039
Low Mass 3	-0.660	0.040
Low Mass 4	-1.040	0.033

7.5 Systematic uncertainties

The sensitivity of the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis is primarily constrained by statistical precision, due to the amount of data collected by the ATLAS detector during Run 2. However,

it is important to identify and evaluate all potential sources of systematic uncertainty that may affect the final results. Systematic uncertainties can arise from various factors, such as auxiliary measurements or theoretical predictions that carry their own uncertainties, or from specific modeling choices used to interpret the observed data. These uncertainties are classified based on their origin:

- **Experimental systematic uncertainties** stem from uncertainties in auxiliary measurements that are used to construct the statistical model (e.g., the total integrated luminosity collected by ATLAS during Run 2), as well as uncertainties in the reconstruction of physics objects with the ATLAS detector, or the impact of applying trigger requirements for events in the signal region. These uncertainties can affect both the event yields of resonant processes (namely, double and single Higgs production) in each category and their shape in the $m_{\gamma\gamma}$ spectrum.
- **Theoretical systematic uncertainties** impact the category yields of HH signal and single Higgs background processes. These may arise from mismodeling of cross sections or branching fractions, mismodelling in the simulation of parton showers in MC samples, or uncertainties in the heavy-flavor content of single Higgs boson production via ggF, VBF, and associated W boson production.

The only systematic uncertainty related to the continuum background modeling is the spurious signal, which is treated as an additional uncertainty on the number of HH signal events in each category. Typically, the impact of each systematic uncertainty, whether it affects the category yields or the $m_{\gamma\gamma}$ shape of the resonant processes, is quantified by propagating the uncertainty through the entire analysis. The impact is then expressed as the relative difference between the nominal and varied results. Both experimental and theoretical systematic uncertainties are evaluated separately for the two HH production modes (ggF and VBF) and the primary single Higgs backgrounds, which include ggF, $t\bar{t}H$, and ZH production mechanisms.

For the ggF HH process, the dependence of systematic uncertainties on κ_λ is accounted for by evaluating each source of uncertainty at $\kappa_\lambda = 1$ and $\kappa_\lambda = 10$. The final uncertainty, applied to the ggF HH process for any value of κ_λ , is taken as the envelope of these two results. Similarly, for the VBF HH process, any potential dependence on κ_λ , κ_{2V} , or κ_V is addressed by evaluating each uncertainty using the six basis samples introduced in Section 7.2.1, with coupling modifiers $(\kappa_\lambda, \kappa_{2V}, \kappa_V) = (1, 1, 1), (0, 1, 1), (10, 1, 1), (1, 1.5, 1), (1, 3, 1)$, and $(-5, 1, 0.5)$. The final systematic uncertainty for VBF HH is taken as their envelope and applies across all $(\kappa_\lambda, \kappa_{2V}, \kappa_V)$ values.

7.5.1 Experimental systematic uncertainties

Yield uncertainties

A 0.83% uncertainty on the integrated luminosity of the full Run 2 dataset is derived using a dedicated measurement [49] from the LUCID 2 detector [220].

A systematic uncertainty is assigned to the κ_λ -reweighting procedure used to parametrise the ggF HH expected yields as a function of κ_λ . The uncertainty in each category, shown in Table 7.8, is quantified by taking the maximum difference between the expected yields

Table 7.8 – Systematic uncertainties associated with the ggF HH and VBF HH parametrizations as functions of κ_λ and $(\kappa_\lambda, \kappa_{2V}, \kappa_V)$, respectively, are presented. Only significant results (where the statistical uncertainty is smaller than the systematic uncertainty) are included. The corresponding statistical uncertainties are also provided.

Category	Syst. unc. [%]	
	ggF HH parametrisation	VBF HH parametrisation
High Mass 1	-	-7.4 ± 6.3
High Mass 2	-	14.0 ± 8.7
High Mass 3	4.3 ± 2.1	15.2 ± 11.4
Low Mass 1	-5.8 ± 4.2	19.4 ± 18.9
Low Mass 2	-7.5 ± 4.5	-
Low Mass 3	7.7 ± 5.1	22.9 ± 17.4
Low Mass 4	-10.7 ± 4.4	-7.3 ± 7.1

obtained directly from the available ggF HH MC samples and those derived through the κ_λ -reweighting method. Only significant relative differences, where the statistical uncertainty on the relative difference is smaller than the difference itself, are considered. Similarly, an uncertainty related to the parametrisation of the VBF HH process as a function of $(\kappa_\lambda, \kappa_{2V}, \kappa_V)$ using the linear combination method is derived by comparing the expected yields from the parametrisation with those obtained from the available VBF HH MC samples (excluding those used as the basis for the parametrisation). Again, only significant relative differences are considered. These uncertainties are presented in Table 7.8. The uncertainties associated with the ggF HH and VBF HH parametrisations are treated as flat uncertainties, affecting the ggF HH and VBF HH yields uniformly across all values of κ_λ , κ_{2V} , and κ_V . The systematic uncertainties related to the ggF HH and VBF HH parametrisations can reach up to 10% and 20%, respectively. However, they are typically accompanied by large statistical uncertainties, often comparable in magnitude to the systematic uncertainties themselves. Therefore, the observed non-closure of the ggF HH and VBF HH parametrisations largely reflects statistical limitations in the HH MC samples rather than indicating a significant mismodeling of the signal processes as a function of the coupling modifiers κ_λ , κ_t , κ_{2V} , and κ_V .

The other sources of experimental systematic uncertainties affecting the event yields of the resonant processes include:

- photon energy scale and resolution;
- diphoton trigger efficiency;
- photon identification and isolation efficiencies;
- pileup modeling in MC samples;
- jet energy scale and resolution;
- Jet Vertex Tagger efficiency;
- Soft term for the missing transverse momentum (see Section 6.3);
- flavor-tagging efficiency for b -, c -, and light jets.

These uncertainties are evaluated using auxiliary samples, where the systematic variations are applied upstream. Asymmetric variations of the category yields are obtained for each HH signal process or single Higgs production mode: the relative difference with respect to the nominal result is the final systematic uncertainty. For example, the upper ($+1\sigma$) and lower (-1σ) systematic uncertainties on the event yields for the i^{th} resonant process in the c^{th} category are given by:

$$\delta N_c^i(\pm 1\sigma) = \frac{N_c^i(\pm 1\sigma)}{N_c^i} - 1, \quad (7.15)$$

where N_c^i and $N_c^i(\pm 1\sigma)$ are the nominal and varied yields, respectively.

Each systematic uncertainty is assigned a statistical uncertainty, calculated using the *bootstrap* method [221]. Three thousand pseudo-events are generated by fluctuating the event weights with multiplicative factors drawn from a Poisson distribution with mean 1: the multiplicative factors are kept consistent among the (real) MC events. The pseudo-events are processed through each step of the analysis, yielding three thousand replicaes of the category yields and their corresponding $\pm 1\sigma$ variations. The central value for $\delta N_c^i(\pm 1\sigma)$ is determined as the mean across all replicaes, while the statistical uncertainty is calculated as the standard deviation. The bootstrap method effectively accounts for potential correlations between events in both the nominal and varied samples when calculating the uncertainty on $\delta N_c^i(\pm 1\sigma)$.


The following pruning criteria are used to exclude systematic uncertainties with negligible effects, or whose impact is not statistically significant:

- Systematic uncertainties smaller than 0.1% are deemed negligible and set to 0.
- If the statistical uncertainty of a systematic uncertainty exceeds 100%, the systematic uncertainty is considered not statistically significant and set to 0.
- Systematic uncertainties with both upper and lower variations being negligible or not statistically significant are excluded from the statistical model.
- For pathological cases where both $\delta N_c^i(+1\sigma)$ and $\delta N_c^i(-1\sigma)$ have the same sign, the sign of the smaller contribution is reversed.


The experimental systematic uncertainties affecting the category yields for the ggF HH production mode are shown in Figure 7.15a and 7.15b for the High Mass 3 and Low Mass 4 categories. The largest contributions come from the ggF HH parametrisation (up to 10%), jet energy resolution (up to 6%), and flavor-tagging efficiencies (up to 3%).

Shape uncertainties

The position and resolution of the $m_{\gamma\gamma}$ resonance from the $H \rightarrow \gamma\gamma$ decay are affected by systematic uncertainties related to the photon energy scale and photon energy resolution. The impacts of these uncertainties are propagated to the $m_{\gamma\gamma}$ peak position and width, represented in the statistical model by the shape parameters μ_{CB} and σ_{CB} of the double-sided crystal ball, for each category and resonant process as follows.



(a) High Mass 3



(b) Low Mass 4

Figure 7.15 – The impact of experimental systematic uncertainties on the expected yields for the ggF HH process is shown for the (a) High Mass 3 and (b) Low Mass 4 categories. These uncertainties are determined by taking the envelope of the impacts evaluated for both the SM ggF HH sample and the ggF HH sample with $\kappa_\lambda = 10$, and are applied uniformly across the full κ_λ range.

- The uncertainty in the photon energy scale is applied upstream in dedicated samples, and introduces a systematic uncertainty affecting the DCSB parameter μ_{CB} . This uncertainty is quantified by calculating the relative difference between the mean of the nominal $m_{\gamma\gamma}$ distribution and its $\pm 1\sigma$ variations:

$$\delta\mu_{CB,c}^i(\pm 1\sigma) = \frac{\langle m_{\gamma\gamma}(\pm 1\sigma) \rangle_c^i}{\langle m_{\gamma\gamma} \rangle_c^i} - 1, \quad (7.16)$$

where $\langle m_{\gamma\gamma} \rangle_c^i$ and $\langle m_{\gamma\gamma}(\pm 1\sigma) \rangle_c^i$ represent the mean values of the nominal and varied $m_{\gamma\gamma}$ distributions, respectively, extracted from the nominal and auxiliary MC samples for the resonant process i , following the selection criteria of the category c .

- The uncertainty from the photon energy resolution impacts the DCSB parameter σ_{CB} , and is calculated by comparing the interquartile ranges (IQR) of the $m_{\gamma\gamma}$ distributions in the nominal case and the $\pm 1\sigma$ variations:

$$\delta\sigma_{CB,c}^i(\pm 1\sigma) = \frac{[\text{IQR}(\pm 1\sigma)]_c^i}{\text{IQR}_c^i} - 1, \quad (7.17)$$

where $[\text{IQR}(\pm 1\sigma)]_c^i$ and IQR_c^i represent the varied and nominal interquartile ranges of the $m_{\gamma\gamma}$ distribution for the resonant process i within the category c .

Therefore, although the DCSB parameters are the same across the resonant processes within each category, the magnitudes of the systematic uncertainties affecting these parameters vary by process. These uncertainties are calculated separately for the ggF HH, VBF HH⁶, and the main single Higgs backgrounds (ggF, t $\bar{t}$ H, and ZH).

The uncertainty on the photon energy scale was found to affect the $m_{\gamma\gamma}$ peak position with a relative effect of approximately 0.5%, consistent across all resonant processes and analysis categories. On the other hand, the uncertainty on the photon energy resolution had a more significant impact on the $m_{\gamma\gamma}$ peak width, with a relative effect around 10%, for all resonant processes and across all categories. The magnitudes of the shape systematic uncertainties related to photon energy scale and resolution for the ggF HH process in each category are illustrated in Figures 7.16a and 7.16b, respectively.

7.5.2 Theoretical systematic uncertainties

Systematic uncertainties arising from missing higher-order corrections in theoretical calculations (“scale” uncertainties), as well as uncertainties propagated from the PDFs and α_S (“PDF+ α_S ” uncertainties), impact the theoretical predictions for the double Higgs and single Higgs production cross sections, and are thus considered in the statistical model. For ggF HH production, the scale uncertainty is combined with the uncertainty related to the choice of renormalisation scheme for the top quark mass, resulting in the

⁶The dependence of the shape systematic uncertainties on the coupling modifiers κ_λ and κ_{2V} was neglected, as the nominal resonant $m_{\gamma\gamma}$ shape showed minimal sensitivity to these parameters, as shown in Section 7.4.1. Hence, the shape uncertainties for the HH signals were evaluated under the SM hypothesis and applied uniformly across all non-SM variations of κ_λ and κ_{2V} .




Figure 7.16 – The impacts of systematic uncertainties on the photon energy scale and resolution on the $m_{\gamma\gamma}$ peak position and $m_{\gamma\gamma}$ peak width respectively are shown for the SM ggF HH process in each analysis category.

“scale+ m_{top} ” uncertainty, as described in Section 2.1.1. The magnitudes of the scale and PDF+ α_s uncertainties impacting the HH cross sections are detailed in Section 2.1.1, while the corresponding uncertainties for single Higgs boson production are presented in Table 1.2. Additionally, these scale and PDF+ α_s uncertainties can influence the efficiencies of the HH and single Higgs processes within each category, as they may alter the shapes of discriminant variables, (e.g., $m_{b\bar{b}\gamma\gamma}^*$), potentially leading to event migrations between categories. This effect has also been accounted for, resulting in uncertainties affecting the HH and single Higgs yields in each category of up to 6%.

An uncertainty of $^{+2.90\%}_{-2.84\%}$ [40, 222] related to the theoretical calculation of the branching fraction for the $H \rightarrow \gamma\gamma$ decay is applied to all di-Higgs and single Higgs processes. Only the signal processes are impacted by the systematic uncertainty of magnitude $^{+1.70\%}_{-1.73\%}$ [40, 222] stemming from the theoretical prediction of $\mathcal{BR}(H \rightarrow b\bar{b})$ [222].

The effect of parton shower modeling on the HH and single Higgs processes was assessed by comparing the nominal category yields to those obtained from alternative MC samples that use a different generator for the simulation of the parton shower. All the nominal HH and single Higgs samples (listed in Table 7.3) are based on the PYTHIA 8 generator, while parton shower uncertainties are derived from comparing with samples relying on HERWIG 7. The parton shower uncertainties the yields of the resonant processes with a relative impact up to 10%.

The simulation of additional heavy-flavor jets in single Higgs boson production via ggF, VBF, and WH has to occur via the parton shower, and thus carries significant theoretical uncertainties. For this reason, a 100% uncertainty is assigned to these production modes. Adopting this uncertainty is supported by measurements of differential cross sections for single Higgs production in phase spaces that include at least one b -jet [223,

[224], as well as studies on $t\bar{t}$ production in association with heavy flavor jets [225, 226] and on W boson production accompanied by additional b -jets [227]. Despite this large uncertainty, its overall impact on the final results of the $HH \rightarrow b\bar{b}\gamma\gamma$ search remains limited, primarily because the analysis is strongly statistically limited, and the contribution of these single Higgs production mechanisms to the most sensitive analysis categories (High Mass 3 and Low Mass 4) is relatively small, as shown in Figure 7.8.

7.6 Statistical model

The results of the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis are extracted via a maximum likelihood fit to the $m_{\gamma\gamma}$ distribution in data, within the invariant mass range $105 \leq m_{\gamma\gamma} \leq 160$ GeV, performed simultaneously across all analysis categories. The likelihood function is defined as:

$$\mathcal{L} = \prod_c \left[\text{Pois}(n_c | N_c(\theta)) \cdot \prod_{i=1}^{n_c} f_c(m_{\gamma\gamma}^i | \theta) \right] \cdot G(\theta), \quad (7.18)$$

where c indexes the analysis categories, and i represents each data event within category c . The outer product runs over all analysis categories. In the likelihood function, n_c is the observed number of events, $N_c(\theta)$ is the expected number of events, and f_c is the probability density function modeling the $m_{\gamma\gamma}$ distribution in category c . The term $G(\theta)$ imposes constraints on the nuisance parameters (NPs), collectively denoted by θ , which affect both the expected yields and the $m_{\gamma\gamma}$ shape.

The observed event count n_c follows a Poisson distribution, centered around the expected number of events $N_c(\theta)$, which is the sum of the expected yields for the double and single Higgs processes, continuum background, and spurious signal in the c^{th} category:

$$\begin{aligned} N_c(\theta) = & \mu_{\text{HH}}^{\text{ggF}} \left(\mu_{\text{HH}}^{\text{ggF}} \cdot N_c^{\text{ggF HH}}(\theta_{\text{yield}}^{\text{ggF HH}}) + \mu_{\text{HH}}^{\text{VBF}} \cdot N_c^{\text{VBF HH}}(\theta_{\text{yield}}^{\text{VBF HH}}) \right) \\ & + \sum_{p \in \text{single H}} N_c^p(\theta_{\text{yield}}^p) \\ & + N_c^{\text{sp}} \cdot \theta_c^{\text{sp}} + N_c^{\text{continuum}} \end{aligned} \quad (7.19)$$

The yields of the resonant processes are subject to various systematic uncertainties, introduced into the model via the nuisance parameters θ . For example, each systematic uncertainty affecting the ggF HH signal yield modifies $N_c^{\text{ggF HH}}$ by a response function, $r(\theta)$, where θ is the NP associated to that systematic. The form of the response function depends on the type of constraint assigned to the specific NP θ . For Gaussian (or Log-Normal) constraints, $r(\theta)$ is defined as $(1 + \delta \cdot \theta)$ (or $\exp(\sqrt{(1 + \delta^2)\theta})$), where δ represents the magnitude of the systematic uncertainty on yields or shape parameters, estimated as described in Section 7.5. This results in the term $N_c^{\text{ggF HH}}(\theta_{\text{yield}}^{\text{ggF HH}}) = N_c^{\text{ggF HH}} \cdot \prod_j r_j(\theta_j)$, where j indexes the systematic uncertainties affecting the ggF HH signal yield. This term contributes to the expected category yield $N_c(\theta)$ in Equation 7.19, thus entering the likelihood function in Equation 7.18.

For each systematic uncertainty, the likelihood is multiplied by a Gaussian probability density function $G(a = 0|\theta, 1)$, centered in its NP θ with unitary width. The quantity a is termed *global observable*, and represents an auxiliary measurement (associated with the systematic uncertainty) whose nominal value is typically zero. Hence, the role of $G(a = 0|\theta, 1)$ is to provide a constraint for the NP θ . The product of all the gaussian constraints makes up the last term in Equation 7.18, namely $G(\theta)$.

As shown in Equation 7.19, only the continuum background yields $N_c^{\text{continuum}}$ (which are free parameters in the maximum-likelihood fit) do not suffer the influence of systematic uncertainties, and therefore do not depend from any nuisance parameter.

The overall expected yields for the HH signal are scaled by the di-Higgs signal strength, μ_{HH} , which quantifies the observed double Higgs production rate relative to SM predictions. Additionally, the individual contributions from ggF HH and VBF HH production are scaled by their respective signal strengths, $\mu_{\text{HH}}^{\text{ggF}}$ and $\mu_{\text{HH}}^{\text{VBF}}$. In the statistical model, μ_{HH} is the primary parameter of interest (POI) when measuring the HH signal strength or setting exclusion limits on it. Similarly, when analysing either the ggF HH or VBF HH production mode separately, the relevant signal strength ($\mu_{\text{HH}}^{\text{ggF}}$ or $\mu_{\text{HH}}^{\text{VBF}}$) becomes the POI, while the signal strength of the other production mode and the overall μ_{HH} factor are fixed to unity. Finally, when focusing on κ_λ or κ_{2V} measurements, $\mu_{\text{HH}}^{\text{ggF}}$ and $\mu_{\text{HH}}^{\text{VBF}}$ are reparameterized in terms of these coupling modifiers (as outlined in Section 7.2.1), which then serve as the new POIs, while μ_{HH} is fixed to unity.

The term $f_c(m_{\gamma\gamma}^i|\theta)$ encodes the shape information of the $m_{\gamma\gamma}$ spectrum. Specifically, it represents the overall probability density function f_c , evaluated on the $m_{\gamma\gamma}$ value for the event i . f_c is the sum of the DSCB functions (describing the resonant processes) and of the exponential shape (modelling the continuum background). Each component is normalised according to the relative contributions of the signal, resonant, and the continuum backgrounds. As described in Section 7.5.1, systematic uncertainties from photon energy scale and resolution impact the mean and width of the resonant shapes. These are modeled using Log-Normal response functions and Gaussian constraints, with associated NPs. The final maximum likelihood fit simultaneously constrains the signal strength, nuisance parameters, and the normalization and shape parameters of the continuum background, resulting in the fit constraining nearly 100 parameters.

The measurement of the parameters of interest in the model is performed using a statistical test based on the profile likelihood ratio (PRL), defined as:

$$\Lambda(\mu) = \frac{\mathcal{L}(\mu, \hat{\theta}(\mu))}{\mathcal{L}(\hat{\mu}, \hat{\theta})}, \quad (7.20)$$

Here, μ represent the parameters of interest (for instance, the HH signal strength, or the coupling modifiers $(\kappa_\lambda, \kappa_{2V})$), while θ represents the nuisance parameters. In the numerator, the nuisance parameters are set to their profiled values $\hat{\theta}(\mu)$, maximizing the likelihood for a fixed μ . In the denominator, both the parameter of interest and the nuisance parameters are set to their best-fit values, $\hat{\mu}$ and $\hat{\theta}$, respectively, which jointly maximize the likelihood. The best-fit $\hat{\mu}$ is the central value of measurement, while its uncertainty is determined from the 68% confidence interval [70], typically calculated

under the assumption of the asymptotic approximation, implying that the PRL follows a χ^2 distribution with number of degrees of freedom equal to the number of POIs in the model.

Intuitively, the profile likelihood ratio measures the agreement between the hypothesized value of μ and the observed data: $0 \leq \Lambda(\mu) \leq 1$, and $\Lambda(\mu)$ approaching unity implies a good agreement between the posited model and data.

In the absence of a signal (i.e., for $\mu_{\text{HH}} \approx 0$), the statistical results for the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis can be interpreted in terms of upper limits on the HH signal strength. The relevant test statistic for limits setting is $\tilde{q}_\mu$, defined as follows:

$$\tilde{q}_\mu = \begin{cases} -2 \ln \frac{\mathcal{L}(\mu, \hat{\boldsymbol{\theta}}(\mu))}{\mathcal{L}(0, \hat{\boldsymbol{\theta}}(0))} & \text{if } \hat{\mu} \leq 0 \\ -2 \ln \frac{\mathcal{L}(\mu, \hat{\boldsymbol{\theta}}(\mu))}{\mathcal{L}(\hat{\mu}, \hat{\boldsymbol{\theta}})} & \text{if } 0 < \hat{\mu} < \mu \\ 0 & \text{if } \hat{\mu} > \mu \end{cases} \quad (7.21)$$

For a given value of μ , its disagreement with the observed data is quantified by the p-value, p_μ :

$$p_\mu = \int_{\tilde{q}_{\mu, \text{obs}}}^{\infty} f(\tilde{q}_\mu | \mu), d\tilde{q}_\mu, \quad (7.22)$$

where $\tilde{q}_{\mu, \text{obs}}$ is the test statistic evaluated on the observed data, while $f(\tilde{q}_\mu | \mu)$ is the probability density function of the test statistic under the null hypothesis (i.e., the POI equals μ). As illustrated in Figure 7.17, a smaller p_μ indicates a greater incompatibility between the observed data and the hypothesized μ . Specifically, a low p_μ for $\mu > 0$ suggests regions where the signal hypothesis is rejected. Upper limits on the POI μ are established using the CL_s method [228]. The quantity

$$\text{CL}_s = \frac{p_\mu}{1 - p_b}, \quad (7.23)$$

is computed for various values of μ within the exclusion limit region. The numerator represents the standard p-value, while the denominator is defined as:

$$p_b = 1 - \text{CL}_b = \int_{\tilde{q}_{\mu, \text{obs}}}^{\infty} f(\tilde{q}_\mu | 0), d\tilde{q}_\mu, \quad (7.24)$$

where $f(\tilde{q}_\mu | 0)$ is the probability density of the test statistic under the alternate hypothesis of no signal. The upper limit on μ is the largest value for which $\text{CL}_s < \alpha$, where α indicates the confidence level (CL) for exclusion. For instance, with $\alpha = 0.05$, the presence of a signal is excluded for values above the upper limit at 95% CL. The distributions of $\tilde{q}_\mu$ under both null and alternate hypotheses are determined using asymptotic formulae derived from the asymptotic approximation. The accuracy of the asymptotic approximation was verified by generating pseudo-experiments to estimate the distributions of the test statistic $\tilde{q}_\mu$ under the null and alternative hypotheses, and repeating the entire statistical analysis. The discrepancy between these results and those based on the

asymptotic approximation was found to be less than 10%. For a comprehensive review of this likelihood-based limit-setting method and the derivation of the asymptotic formulae, see [70].




Figure 7.17 – Schematic description of a p -value obtained with an observed value $t_{\mu,\text{obs}}$ of the test statistic t_μ [70]. The same idea is valid for the test statistic $\hat{q}_\mu$.

7.7 Results

As discussed in the previous Section, the statistical results for the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis are obtained through a maximum likelihood fit to the $m_{\gamma\gamma}$ distribution in data, performed simultaneously across all analysis categories. As summarized in Table 7.9, no excess of events above the background-only expectations is observed. In fact, the measured HH signal strength is slightly negative ($\hat{\mu}_{HH} \approx -1.7$), indicating a small deficit of events observed in data, compared to the expected background. Therefore, 95% CL upper limits are established for the HH signal strength. Additionally, 95% confidence intervals are derived for the coupling modifiers κ_λ and κ_{2V} , assuming all the other coupling modifiers affecting HH production set to their SM value. Finally, a two-dimensional measurement in the $(\kappa_\lambda, \kappa_{2V})$ plane was performed, yielding a 95% CL exclusion region for these two parameters.

7.7.1 Observed data

Figure 7.18 presents the background-only fits to data for each analysis category. The contribution of the continuum background is highlighted, and the small resonant peak around $m_H \approx 125$ GeV observed in the fit results refers to single Higgs boson background events.




Figure 7.18 – Diphoton invariant mass distributions in data for each category within the mass range $105 \leq m_{\gamma\gamma} \leq 160$ GeV. The background-only fit results are also shown: both the contributions from the continuum background and single Higgs boson production are considered.

Table 7.9 – The expected number of events (estimated by using simulation) from HH signals with various κ_λ and κ_{2V} hypotheses and single Higgs boson production, and the expected number of events from the continuum background, evaluated in the $120 < m_{\gamma\gamma} < 130$ GeV window. For comparison, the number of observed data events is also shown. The uncertainties in the HH signals and single Higgs boson backgrounds include the systematic uncertainties discussed in Section 7.5. Asymmetric uncertainties arise primarily from the theory calculation of the SM ggF HH cross-section and the large uncertainty in the yield of single Higgs bosons produced in ggF events in association with heavy-flavour jets, parameterised by a lognormal distribution. The uncertainty in the continuum background is given by the sum in quadrature of the statistical uncertainty from the fit to the data and the spurious signal uncertainty.

	High Mass 1	High Mass 2	High Mass 3	Low Mass 1	Low Mass 2	Low Mass 3	Low Mass 4
SM HH ($\kappa_\lambda = 1$) signal	$0.26^{+0.03}_{-0.04}$	$0.194^{+0.021}_{-0.032}$	$0.84^{+0.10}_{-0.14}$	$0.048^{+0.007}_{-0.008}$	$0.038^{+0.004}_{-0.006}$	$0.039^{+0.004}_{-0.006}$	$0.032^{+0.004}_{-0.004}$
ggF	$0.25^{+0.03}_{-0.04}$	$0.188^{+0.021}_{-0.032}$	$0.81^{+0.10}_{-0.14}$	$0.046^{+0.007}_{-0.008}$	$0.036^{+0.004}_{-0.006}$	$0.037^{+0.004}_{-0.006}$	$0.025^{+0.004}_{-0.004}$
VBF [10^{-3}]	$7.9^{+0.6}_{-0.5}$	$5.3^{+0.5}_{-0.4}$	29^{+4}_{-3}	$1.98^{+0.28}_{-0.24}$	$1.71^{+0.16}_{-0.14}$	$1.96^{+0.21}_{-0.19}$	$7.4^{+0.6}_{-0.5}$
HH ($\kappa_\lambda = 10$) signal	$2.5^{+0.4}_{-0.3}$	$1.81^{+0.25}_{-0.20}$	$6.2^{+0.8}_{-0.6}$	$5.0^{+1.2}_{-0.9}$	$3.8^{+0.7}_{-0.5}$	$3.7^{+0.7}_{-0.6}$	$3.6^{+0.4}_{-0.4}$
ggF	$2.3^{+0.4}_{-0.3}$	$1.64^{+0.25}_{-0.19}$	$4.9^{+0.8}_{-0.6}$	$4.7^{+1.0}_{-0.8}$	$3.6^{+0.7}_{-0.6}$	$3.3^{+0.7}_{-0.5}$	$2.04^{+0.34}_{-0.27}$
VBF	$0.231^{+0.019}_{-0.017}$	$0.170^{+0.019}_{-0.017}$	$1.29^{+0.15}_{-0.14}$	$0.28^{+0.20}_{-0.11}$	$0.23^{+0.23}_{-0.12}$	$0.36^{+0.10}_{-0.08}$	$1.57^{+0.17}_{-0.16}$
VBF HH ($\kappa_{2V} = 3$) signal	$0.23^{+0.04}_{-0.04}$	$0.20^{+0.05}_{-0.04}$	$3.8^{+0.7}_{-0.6}$	$0.03^{+0.04}_{-0.02}$	$0.03^{+0.06}_{-0.02}$	$0.048^{+0.023}_{-0.015}$	$0.17^{+0.04}_{-0.03}$
Single Higgs background	$1.5^{+0.5}_{-0.3}$	$0.48^{+0.21}_{-0.10}$	$0.57^{+0.25}_{-0.14}$	$1.72^{+0.31}_{-0.19}$	$0.53^{+0.08}_{-0.06}$	$0.29^{+0.14}_{-0.07}$	$0.16^{+0.06}_{-0.03}$
ggF	$0.5^{+0.5}_{-0.2}$	$0.14^{+0.21}_{-0.09}$	$0.25^{+0.25}_{-0.12}$	$0.29^{+0.31}_{-0.15}$	$0.08^{+0.08}_{-0.04}$	$0.07^{+0.13}_{-0.06}$	$0.04^{+0.06}_{-0.03}$
ttH	$0.302^{+0.034}_{-0.032}$	$0.069^{+0.009}_{-0.008}$	$0.063^{+0.008}_{-0.007}$	$0.77^{+0.09}_{-0.08}$	$0.214^{+0.029}_{-0.026}$	$0.100^{+0.012}_{-0.012}$	$0.048^{+0.005}_{-0.005}$
ZH	$0.61^{+0.06}_{-0.05}$	$0.174^{+0.020}_{-0.016}$	$0.188^{+0.035}_{-0.029}$	$0.49^{+0.05}_{-0.04}$	$0.149^{+0.028}_{-0.025}$	$0.069^{+0.033}_{-0.023}$	$0.028^{+0.010}_{-0.007}$
Other	$0.17^{+0.08}_{-0.04}$	$0.089^{+0.030}_{-0.016}$	$0.07^{+0.04}_{-0.02}$	$0.181^{+0.030}_{-0.019}$	$0.089^{+0.016}_{-0.009}$	$0.046^{+0.007}_{-0.004}$	$0.036^{+0.008}_{-0.004}$
Continuum background	$11.3^{+1.5}_{-1.6}$	$3.2^{+0.8}_{-0.8}$	$2.8^{+0.8}_{-0.8}$	$37.2^{+2.9}_{-2.9}$	$10.8^{+1.5}_{-1.5}$	$4.4^{+0.9}_{-1.0}$	$1.1^{+0.5}_{-0.5}$
Total background	$12.8^{+1.6}_{-1.6}$	$3.7^{+0.9}_{-0.8}$	$3.4^{+0.8}_{-0.8}$	$38.9^{+2.9}_{-2.9}$	$11.3^{+1.5}_{-1.5}$	$4.7^{+0.9}_{-1.0}$	$1.3^{+0.5}_{-0.5}$
Data	12	4	1	29	8	5	4

7.7.2 Exclusion limits on the HH signal strength

No significant excess over the expected background is observed, leading the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis to set upper limits on Higgs boson pair production. The observed 95% CL upper limit on the total HH production signal strength, μ_{HH} , is 4.0. Under the assumption of no HH production, the expected 95% CL upper limit is 5.0, while assuming SM HH production, it is 6.4. These expected results are obtained by performing a maximum-likelihood fit on *Asimov data*⁷ [70], generated either assuming the background-only hypothesis or the SM HH signal-plus-background hypothesis. The 95% expected upper limits on μ_{HH} show an improvement of 12%, with respect to the previous $HH \rightarrow b\bar{b}\gamma\gamma$ analysis with Run 2 data [12], due to the new optimisation of the analysis categories and to the improved estimation of the spurious signal, as discussed below.

⁷The Asimov dataset is a toy dataset constructed to perfectly follow a probability density function without statistical fluctuations. By construction, when one evaluates the estimator for a parameter on the Asimov dataset (which is built using a certain model), the estimator returns the true value of the parameter. Asimov datasets are therefore ideal for studying the median sensitivity of a search.

Table 7.10 – Breakdown of the dominant systematic uncertainties in the expected μ_{HH} upper limit at 95% CL. The impact of the uncertainties corresponds to the relative variation of the expected upper limit when re-evaluating the profile likelihood ratio after fixing the nuisance parameter in question to its best-fit value, while all remaining nuisance parameters remain free to float. Only systematic uncertainties with an impact of at least 0.1% are shown.

Systematic uncertainty source	Relative impact [%]
Experimental	
Photon energy resolution	0.4
Photon energy scale	0.1
Flavour tagging	0.1
Theoretical	
Factorisation and renormalisation scale	4.8
$\mathcal{BR}(H \rightarrow \gamma\gamma, b\bar{b})$	0.2
Parton showering model	0.2
Heavy-flavour content	0.1
Background model (spurious signal)	0.1

Upper limits at 95% CL are also placed on ggF HH and VBF HH production separately, assuming that the other production mode is fixed to the SM expectation. The observed upper limit on the ggF HH signal strength, μ_{HH}^{ggF} is 4.1, with an expected limit of 5.3, assuming no ggF HH production. For VBF HH, the observed upper limit on μ_{HH}^{VBF} is 96, with an expected limit of 145 under the assumption of $\mu_{HH}^{\text{VBF}} = 0$. This shows that the analysis sensitivity to SM HH production is driven primarily by the dominant ggF HH production mode, while the sensitivity to SM VBF HH production remains limited.

The observed limits placed on μ_{HH} , μ_{HH}^{ggF} and μ_{HH}^{VBF} are tighter than the expected values, due to the slight deficit of events in the data compared to the background-only expectation.

Table 7.10 summarizes the relative impact of the systematic uncertainties discussed in Section 7.5 on the 95% CL expected upper limits on μ_{HH} . Overall, the effect of these uncertainties is small, with the “scale” uncertainty having the largest influence, reducing the analysis sensitivity by nearly 5%. This is primarily due to the large “scale+ m_{top} ” uncertainty of $^{+23\%}_{-6\%}$ affecting the SM ggF HH cross section. The impact of the spurious signal uncertainty is 0.1%, which corresponds to a large reduction with respect to the previous $HH \rightarrow b\bar{b}\gamma\gamma$ analysis using Run 2 data, where it reached 3%. This improvement is due to suppressed statistical fluctuations in the background-only template used to select the analytical function for background modeling, which otherwise lead to inflated spurious signal estimates. In the current analysis, the SHERPA 2.2.12 $\gamma\gamma+b\bar{b}$ sample, with much higher statistics in the analysis categories, replaces the SHERPA 2.2.4 $\gamma\gamma$ +jets sample used previously, leading to spurious signal values up to 50% lower.

Another perspective on the influence of systematic uncertainties in the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis is provided by Figure 7.19, which shows the impact of each systematic uncertainty on the expected measurement of μ_{HH} , performed under the assumption of SM HH production. The uncertainties are ranked by descending impact on the measurement: the leading one is the 100% uncertainty on the single Higgs background via ggF plus heavy-flavor emissions, which affects the μ_{HH} measurement by up to 35%. The “scale+ m_{top} ” uncertainty on the SM ggF HH cross section provides the second-to-largest impact on the μ_{HH} measurement (up to 20%), while all the other systematic uncertainties contribute to less than 10%.

7.7.3 Constraints on κ_λ and κ_{2V}

Figure 7.20a shows the profile log-likelihood ratio, $-2\ln(\Lambda)$, as a function of κ_λ , with all the other coupling modifiers affecting HH production (including κ_{2V}) set to their SM values. The value of κ_λ is constrained within the range $-1.4 < \kappa_\lambda < 6.9$ at 95% CL, compared to the expected range of $-2.8 < \kappa_\lambda < 7.8$, assuming the SM hypothesis ($\kappa_\lambda = 1$). The expected confidence interval has improved by 6% over the previous Run 2 analysis [12], thanks to the new optimised categorisation, which explicitly incorporates the VBF HH production mode, thus offering an improved sensitivity to κ_λ . However, the observed confidence interval for κ_λ is wider than in the previous analysis. In fact, the previous analysis observed a larger deficit of events in data compared to background-only expectations in the most sensitive analysis categories, leading to tighter observed constraints on κ_λ . A comparison of the two analyses, assessing the overlap of selected events and their correlations via a bootstrap technique [221], shows that the two observed κ_λ results are consistent within 0.3 standard deviations.

Similarly, Figure 7.20b shows the profile log-likelihood ratio for κ_{2V} , assuming all other coupling modifiers set to their SM values. The observed (expected) 95% confidence interval for κ_{2V} is $-0.5 < \kappa_{2V} < 2.7$ ($-1.1 < \kappa_{2V} < 3.3$). The expected range represents a 17% improvement over the previous Run 2 analysis, again due to the inclusion of the VBF HH production mode in the analysis optimisation, which provides the unique sensitivity to κ_{2V} .

As shown in Figure 7.20a, the expected profile log-likelihood ratio as a function of κ_λ reveals a partially resolved double minimum at $\kappa_\lambda = 1$ and approximately $\kappa_\lambda = 4$. This behavior arises because the cross section for ggF HH production, which predominantly drives the analysis sensitivity to κ_λ , is the same at both $\kappa_\lambda = 1$ and $\kappa_\lambda = 4$, with a value of around 30 fb. If the sensitivity to κ_λ were based only on the overall ggF HH cross section, the expected $-2\ln(\Lambda)$ curve would not distinguish between the two minima. However, by exploiting the differential information from binning in the $m_{b\bar{b}\gamma\gamma}^*$ variable, whose distribution differs between the $\kappa_\lambda = 1$ and $\kappa_\lambda = 4$ hypotheses, the analysis can resolve the two minima. In contrast, the observed $-2\ln(\Lambda)$ curve shows a single minimum at $\kappa_\lambda = 3$. This difference is driven by the deficit of observed events relative to the background-only expectation, making the κ_λ value that best fits the data the one associated with the minimum ggF HH cross section, around $\kappa_\lambda = 2.4$.

Finally, Figure 7.21 presents the observed and expected two-dimensional constraints at




Figure 7.19 – Fitted values of nuisance parameters associated to each systematic uncertainty (black markers) with their uncertainty (black lines), superimposed to a 100% uncertainty (red line). The pre-fit impacts on the expected μ_{HH} value (yellow boxes) are evaluated by fixing each NP, in turn, to its nominal value plus a $\pm 100\%$ uncertainty and performing an independent fit to measure μ_{HH} . The post-fit impacts, represented by the shaded blue or red boxes, are evaluated similarly, except that each NP is fixed to its fitted value plus its post-fit $\pm 1\sigma$ uncertainty. The impacts on μ_{HH} are measured assuming SM HH production. The nuisance parameters are ranked by decreasing impact on the measurement.




Figure 7.20 – Observed (solid line) and expected (dashed line) value of the profile log-likelihood ratio, $-2 \ln \Delta$, as a function of κ_λ (left) and κ_{2V} (right), with all the other coupling modifiers (including, respectively, κ_{2V} or κ_λ) are fixed to their SM predictions.

68% and 95% CL in the $(\kappa_\lambda, \kappa_{2V})$ plane, derived from the two-dimensional profile log-likelihood ratio, with all other coupling modifiers fixed to their SM values.




Figure 7.21 – Likelihood contours at 68% (solid line) and 95% (dashed line) CL in the $(\kappa_\lambda, \kappa_{2V})$ parameter space, with all the other coupling modifiers affecting HH production fixed to their SM predictions. The corresponding expected contours are shown by the inner and outer shaded regions. The SM prediction (corresponding to $(\kappa_\lambda, \kappa_{2V}) = (1, 1)$) is indicated by the star, while the best-fit value is denoted by the cross.

7.8 Combination with the other HH final states

As discussed in Sections 1.3.2 and 2.1.4, Higgs boson decays cover a diverse range of final states, providing a rich phenomenology for studying HH production. The most sensitive channels, referred to as the *golden channels*, include $b\bar{b}\gamma\gamma$, $b\bar{b}\tau\tau$, and $b\bar{b}b\bar{b}$. These final states are complementary, as they offer the best sensitivity to different regions of the m_{HH} spectrum (for example, $b\bar{b}\gamma\gamma$ is particularly effective at probing HH production in the softer m_{HH} range). To maximise the sensitivity to HH production and its properties, it is crucial to exploit the statistical power of all individual HH searches across various final states and extract the combined results. This Section presents the combination of multiple HH searches [7], offering the most complete insights into HH production and the trilinear Higgs boson self-coupling achieved using ATLAS Run 2 data.

7.8.1 Overview of the HH searches

The HH combination incorporates five input analyses. Along with the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis discussed in this chapter, the other two key searches included in the combination are the $HH \rightarrow b\bar{b}b\bar{b}$ [13, 14] and $HH \rightarrow b\bar{b}\tau\tau$ [15] analyses. Two additional analyses, with a limited sensitivity with respect to the three golden channels, contribute to the combination: the $HH \rightarrow$ multi-leptons [229] and the $HH \rightarrow b\bar{b}\ell\ell + E_T^{\text{miss}}$ [230] analyses, which are briefly presented below.

$HH \rightarrow b\bar{b}b\bar{b}$

The $HH \rightarrow b\bar{b}b\bar{b}$ analysis takes advantage of the highest HH branching fraction (34%), but suffers from a large SM background from multijet production. The analysis exploits both a resolved topology, where all four final-state jets are individually reconstructed and identified, and a boosted topology, where the two bottom quarks from the Higgs decay are merged into a single large-radius jet.

The $HH \rightarrow b\bar{b}b\bar{b}$ resolved analysis was performed using 126 fb^{-1} of data collected during the full Run 2. It selects events containing at least four b -tagged jets, which are further categorized into VBF HH- and ggF HH-enriched categories based on whether there are additional jets and on their characteristics. To mitigate contamination from $t\bar{t}$ production, a top-veto discriminant is applied to both categories, similar to the single-topness discriminant described in Section 7.3.2. The signal region in each category is defined within a two-dimensional plane of the Higgs boson candidates' masses, m_{H_1} and m_{H_2} , focusing on the area surrounding the expected Higgs boson mass peaks. The HH invariant mass is used as the discriminating variable in the maximum-likelihood fit to extract the final results. The background, primarily composed of QCD multijet and $t\bar{t}$ processes, is evaluated using a data-driven method. This approach involves extrapolating the m_{HH} distribution from control regions that are identical to the signal regions, except that they require exactly two b -tagged jets. To correct for differences in kinematic properties between events with four and two b -tagged jets, a neural network-based reweighting is applied.

The $HH \rightarrow b\bar{b}b\bar{b}$ boosted analysis uses the full 140 fb^{-1} dataset collected by ATLAS dur-

ing Run 2 and is specifically optimised to target the VBF HH production mode and scenarios with anomalous κ_{2V} . This analysis selects events with two large-radius jets, representing the Higgs boson candidates, and two small-radius jets with large angular separation and high invariant mass, identified as the VBF jets. The signal region is defined by requiring that both large-radius jets meet the tightest b -tagging working point, and that the masses of the two Higgs candidates, m_{H_1} and m_{H_2} , lie within a small area around the Higgs boson mass peaks in the (m_{H_1}, m_{H_2}) plane. A BDT-based discriminant, trained on MC samples using variables optimized for the HH system and the VBF topology, is used to extract the final results in the maximum likelihood fit. The primary background, composed of QCD multijet and $t\bar{t}$ events, is estimated using a control region defined similarly to the signal region, but with the condition that one of the large-radius jets fails the b -tagging criteria.

$HH \rightarrow b\bar{b}\tau\tau$

The $b\bar{b}\tau\tau$ final state has one of the highest branching fractions (7.3%) among the various HH decay channels and suffers from a moderate background contamination from competing SM processes. The $HH \rightarrow b\bar{b}\tau\tau$ search defines separate signal regions for the fully hadronic and semileptonic $\tau\tau$ final states, depending on the decay chains of the two τ leptons. In the semileptonic channel events must contain exactly one electron or muon and one hadronically decaying τ lepton, both with opposite charges. In the fully hadronic channel events are required to have exactly two hadronically decaying τ -leptons, also with opposite charges. Both channels require the presence of exactly two b -tagged jets.

The dominant backgrounds come from Z + heavy flavor jets and $t\bar{t}$ production, with additional contributions from single Higgs boson decays into $\tau\tau$, where $t\bar{t}H$, ggF , and ZH processes are the most significant. A BDT-based discriminant is trained independently in each signal region, and the fit is performed simultaneously across the signal regions, along with an additional control region that constrains the normalization of the Z + heavy flavor background.

In addition to the main analyses discussed earlier, two additional searches contribute to the HH combination, namely, the $HH \rightarrow \text{multi-leptons}$ and $HH \rightarrow b\bar{b}\ell\ell + E_T^{\text{miss}}$ analyses. The multi-leptons search focuses on HH decays into pairs of W or Z bosons, τ leptons, or photons, where the W and Z bosons, as well as the τ leptons, decay leptonically, resulting in multiple leptons in the final state. This channel has a total branching fraction of 6.5%. The $b\bar{b}\ell\ell + E_T^{\text{miss}}$ analysis targets HH final states where one Higgs boson decays into two bottom quarks, and the other decays into either a pair of Z or W bosons, or τ leptons, which subsequently decay into two opposite-sign, same-flavor leptons (either electrons or muons). This channel accounts for 2.9% of the total HH decays.

7.8.2 Combined results

The measured HH signal strength is $\mu_{HH} = 0.5_{-1.0}^{+1.2} = 0.5_{-0.8}^{+0.9} (\text{stat.})_{-0.6}^{+0.7} (\text{syst.})$, where “stat.” and “syst.” denote the statistical and systematic uncertainties, respectively, con-




Figure 7.22 – Fitted values of nuisance parameters associated to each systematic uncertainty (black markers) with their uncertainty (black lines), superimposed to a 100% uncertainty (red line). The pre-fit impacts on the expected μ_{HH} value (yellow boxes) are evaluated by fixing each NP, in turn, to its nominal value plus a $\pm 100\%$ uncertainty and performing an independent fit to measure μ_{HH} . The post-fit impacts, represented by the shaded blue or red boxes, are evaluated similarly, except that each NP is fixed to its fitted value plus its post-fit $\pm 1\sigma$ uncertainty. The impacts on μ_{HH} are measured assuming SM HH production. The nuisance parameters are ranked by decreasing impact on the measurement.

tributing to the 68% confidence interval. The expected result, assuming SM HH production, is $\mu_{HH} = 1.0^{+1.2}_{-1.0} = 1.0^{+1.0}_{-0.9} (\text{stat.})^{+0.7}_{-0.5} (\text{syst.})$. Among the systematic uncertainties, the largest contribution comes from the 100% uncertainty on the single Higgs plus heavy-flavor background, particularly affecting the $b\bar{b}\gamma\gamma$ and $b\bar{b}\tau\tau$ channels. This uncertainty impacts the expected value signal strength (measured assuming SM HH production) by approximately 25%, as shown in Figure 7.22, primarily from the $b\bar{b}\tau\tau$ channel. The observed (expected) discovery significance for the HH signal is just 0.4 (1.0) standard deviations [70], with respect to the hypothesis of no HH production.

Since no significant excess of signal is observed, upper limits are set on the HH signal strength. The observed 95% CL upper limit on the total HH production signal strength, μ_{HH} , is 2.9. Under the assumption of no HH production, the expected 95% CL upper limit is 2.4, while assuming SM HH production, it is 3.4. The observed limits are slightly less stringent than the expected under the $\mu_{HH} = 0$ hypothesis, due to a modest ex-

cess of HH events above the SM background (the best-fit value for the signal strength is $\mu_{HH} \approx 0.5$). This excess is primarily driven by the $HH \rightarrow b\bar{b}\tau\tau$ and $HH \rightarrow$ multi-leptons analyses, with observed significances of 1.4 and 1.0 standard deviations, respectively. The relative contributions of the individual HH channels to the overall sensitivity are illustrated in Figure 7.23. The $HH \rightarrow b\bar{b}\tau\tau$ analysis imposes the strongest expected constraint on μ_{HH} , with an upper limit of 3.3.




Figure 7.23 – Observed and expected 95% CL upper limits on the signal strength μ_{HH} , from the $b\bar{b}\tau\tau$, $b\bar{b}\gamma\gamma$, $b\bar{b}b\bar{b}$, multi-leptons, and $b\bar{b}\ell\ell + E_T^{\text{miss}}$ channels, and their combination. The expected results are calculated within the assumption of no HH production.

The self-coupling modifier κ_λ is explored in the ggF and VBF HH production processes. The impact of κ_λ on single-Higgs-boson background production and Higgs boson branching fraction is neglected. Assuming that other Higgs boson couplings are set to the SM predictions, a maximum likelihood fit to the data yields $\kappa_\lambda = 3.8^{+2.1}_{-3.6}$, which is compatible with the SM prediction, with a p -value of 0.53. The expected best-fit value of κ_λ is $1.0^{+4.7}_{-1.5}$ when assuming SM HH production. The observed (expected) 95% CL interval is $-1.2 < \kappa_\lambda < 7.2$ ($-1.6 < \kappa_\lambda < 7.2$). This result represents the best expected sensitivity to the trilinear Higgs boson self-coupling to date. The values of the profile-likelihood test statistic $-2 \ln \Lambda$ as a function of κ_λ are shown in Figure 7.24a for both the individual searches and their combination, highlighting the $b\bar{b}\gamma\gamma$ channel as the most sensitive.

Similarly, κ_{2V} is explored in the VBF HH production process. Assuming the SM predictions for other Higgs boson couplings, the observed (expected) value is $\kappa_{2V} = 1.02^{+0.22}_{-0.23}$ ($\kappa_{2V} = 1.00^{+0.40}_{-0.36}$). The observed (expected) 95% CL interval is $0.6 < \kappa_{2V} < 1.5$ ($0.4 < \kappa_{2V} < 1.6$). The values of $-2 \ln \Lambda$ as a function of κ_{2V} are shown in Figure 7.24b: the

$b\bar{b}b\bar{b}$ analysis is the most sensitive, thanks to the boosted VBF $HH \rightarrow b\bar{b}b\bar{b}$ analysis. A deficit of data events in this channel results in stronger constraints on κ_{2V} than expected. Finally, two-dimensional likelihood contours of $-2\ln\Lambda$ in the $(\kappa_\lambda, \kappa_{2V})$ plane are presented in Figure 7.25.




Figure 7.24 – Expected values of the profile log-likelihood ratio, $-2\ln\Lambda$, as a function of κ_λ (left) and κ_{2V} (right), are shown for all the single channels and for the combined data. The observed values from the combined data are depicted by solid black lines. These results are computed with the assumption that all other Higgs boson couplings follow the SM predictions.




Figure 7.25 – The expected 95% CL contours in the $(\kappa_\lambda, \kappa_{2V})$ plane, corresponding to the individual decay channels and their combination, are illustrated using dashed lines. The observed contour from the combined results is depicted by a solid black line. The SM prediction is marked by a star, and the combined best-fit value is indicated by a cross.

Search for resonant production of Higgs bosons plus new scalars in the $b\bar{b}\gamma\gamma$ final state

The self-interactions of the Higgs boson are governed by the scalar potential, whose structure is still largely unknown. New physics could hide in the scalar sector, potentially giving rise to anomalous trilinear Higgs boson self-coupling values. Models that extend the scalar sector to include three or more scalar states, of which the 125-GeV Higgs boson is the only one observed so far, are particularly interesting. Such models not only can address several shortcomings of the SM (including the hierarchy problem) but also offer viable dark matter candidates. These new physics scenarios can be explored by searching for asymmetric Higgs boson pair production, where a new heavy scalar particle X decays into the 125-GeV Higgs boson, H , and another scalar particle S . This kind of phenomenology is predicted in multiple beyond-the-SM frameworks, including the next-to-minimal supersymmetric model (NMSSM) [231] and the two-real-scalar-singlet model (TRSM) [17].

This Chapter describes a search for the $X \rightarrow SH$ process, where the Higgs boson decays into two photons and the new scalar particle S decays into two bottom quarks. This search is performed using 140 fb^{-1} of 13-TeV pp collision data recorded by the ATLAS detector during the Run 2 data-taking period [18]. This is the first search for asymmetric Higgs boson pair production in the $b\bar{b}\gamma\gamma$ final state performed with ATLAS data. The search for the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal is conducted across a wide range of masses for the new scalars (m_X and m_S), spanning $170 \leq m_X \leq 1000 \text{ GeV}$ and $15 \leq m_S \leq 500 \text{ GeV}$, reaching previously unexplored regions of the parameter space at low (m_X, m_S) values, as shown in Figure 2.14. The search space only includes (m_X, m_S) values where the $X \rightarrow SH$ decay is kinematically allowed, which requires that $m_X > m_S + m_H$.

For each (m_X, m_S) value, the key feature of the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal is the presence of three distinct resonances: in the invariant mass spectrum of the two photons ($m_{\gamma\gamma}$), the two b -jets ($m_{b\bar{b}}$), and the four-body $b\bar{b}\gamma\gamma$ system. The $H \rightarrow \gamma\gamma$ decay produces a narrow peak in the $m_{\gamma\gamma}$ distribution centered around $m_H \approx 125 \text{ GeV}$. The decay of the new scalar S into two bottom quarks results in a peak in the $m_{b\bar{b}}$ distribution, centered at m_S . Finally, the presence of the heavy resonance X manifests as a peak in the $m_{b\bar{b}\gamma\gamma}$ distribution, centered at m_X . These three resonances are crucial for distinguishing the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal from SM backgrounds, where at least one of these resonances is absent.

The backgrounds for the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ search largely overlap with those of the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis described in Chapter 7, since the two analyses share the same $b\bar{b}\gamma\gamma$ final state. The main background is the continuum background, whose irreducible component is diphoton production accompanied by additional b -jets. Reducible continuum backgrounds come from processes involving diphoton production with additional c - or light-flavor jets misidentified as b -jets, or from γj and jj processes where one or two jets are misidentified as photons. Smaller contributions come from continuum diphoton production plus $t\bar{t}$ or plus $Z \rightarrow b\bar{b}$. The latter background is especially relevant when m_S is near the Z boson mass ($m_Z \approx 91.2$ GeV). Other minor backgrounds include single Higgs boson production in the diphoton decay channel with additional jets and SM HH production in the $b\bar{b}\gamma\gamma$ final state. To suppress these backgrounds, a strategy similar to that used in the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis is employed: stringent photon identification and isolation criteria are applied to reduce fake photon contamination, and b -tagging is used to select events with b -jets in the final state.

Section 8.1 provides an overview of the dataset and simulated samples for the signal and background processes used in the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ search. Event selection is described in Section 8.2, while the final discriminant, based on a parametrized neural network (PNN) algorithm [232], is discussed in Section 8.3. The background estimation method and analysis strategy are detailed in Section 8.4. Section 8.5 outlines the systematic uncertainties, and the statistical model used for interpreting the results is described in Section 8.6. The results of the search for the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal are presented in Section 8.7. Finally, the evaluation of the look-elsewhere effect for the global discovery significance is presented in Section 8.8.

8.1 Data and simulated events

This analysis is based on 13 TeV pp collision data collected by the ATLAS detector during the Run 2 data-taking period, corresponding to an integrated luminosity of $140 \pm 1.2 \text{ fb}^{-1}$ [49]. The trigger strategy, driven by the presence of two photons in the final state, follows the same approach as that used in the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis, described in Section 7.1.

As will be detailed later in this Chapter, the background estimation for the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ search relies heavily on MC simulations. The continuum background is modeled using the $\gamma\gamma$ +jets, $t\bar{t}\gamma\gamma$, and $Z(\rightarrow q\bar{q}) + \gamma\gamma$ MC samples. The first two samples are the same as those discussed in Section 7.2.2. The $Z(\rightarrow q\bar{q}) + \gamma\gamma$ sample, which includes both the irreducible background from Z boson decays into two bottom quarks and the reducible background from Z decays into c - or lighter quarks, is generated using SHERPA 2.2.11 [203]. A generator-level requirement is placed on the diphoton invariant mass, $m_{\gamma\gamma} > 95$ GeV, to target the $H \rightarrow \gamma\gamma$ phase space. An alternative simulation for the $\gamma\gamma$ +jets process, used to evaluate the systematic uncertainty in the continuum background modeling, is based on MADGRAPH5_AMC@NLO [68]. This sample simulates diphoton production with up to two additional jets at NLO accuracy in QCD. The simulation of the parton shower is handled by PYTHIA 8.2. Compared to the nominal SHERPA 2.2.4 $\gamma\gamma$ +jets sample, the alternative MADGRAPH5_AMC@NLO sample contains much

less simulated events, and has therefore a much lower statistics in the $b\bar{b}\gamma\gamma$ final state. Since the $\gamma\gamma$ +jets sample does not account for fake photon contributions to the continuum background, its normalisation is adjusted to match data in the sideband regions ($105 \leq m_{\gamma\gamma} \leq 120$ GeV and $130 \leq m_{\gamma\gamma} \leq 160$ GeV).

The simulated samples used to describe the single Higgs boson backgrounds and the SM HH process are the same as those employed in the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis, outlined in Section 7.2, with inclusive normalisation is set to match the most precise theoretical predictions available for the production cross sections and branching fractions.

8.1.1 Signal samples

The $X \rightarrow SH$ signals are simulated for 165 mass points in the (m_X, m_S) plane using PYTHIA 8.2 for both event generation and parton showering. The signal samples only consider the decay channels $H \rightarrow \gamma\gamma$ and $S \rightarrow b\bar{b}$. The new scalar particles, X and S , are generated under the narrow width approximation (NWA) [233].

Signal interpolation

The aim of this analysis is to search for a $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal continuously across the entire explored (m_X, m_S) region, ensuring no gaps in sensitivity between the tested mass points. This requires testing for the presence of a signal against the background-only hypothesis throughout the (m_X, m_S) parameter space in a fully continuous way. However, generating $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ MC samples on an infinitely fine two-dimensional grid in the masses of the two resonances would be unfeasible. To address this challenge, a signal interpolation technique is used to estimate the expected signal shapes and efficiencies for (m_X, m_S) points that fall between the fully simulated mass points. The goal is to create a signal grid in the (m_X, m_S) plane with sufficient granularity so that the spacing between the grid points is smaller than the experimental widths of the signal in both the m_X and m_S directions. In this way, any $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal present between the available mass points (whether fully simulated or interpolated) would be detected by testing the presence of signal in the neighbouring available (m_X, m_S) points, effectively making the analysis sensitivity continuous across the (m_X, m_S) plane.

The interpolation procedure involves estimating the probability density of two key observables, $m_{b\bar{b}\gamma\gamma}^*$ ¹ and $m_{b\bar{b}}$, for each (m_X, m_S) mass point in the signal grid. This density, defined as:

$$\mathcal{P}(m_{b\bar{b}\gamma\gamma}^*, m_{b\bar{b}} | m_X, m_S) \quad (8.1)$$

is then used to build the final PNN-based discriminant for the signal in the considered (m_X, m_S) point (as will be described in Section 8.3). Given a target $(m_X^{\text{target}}, m_S^{\text{target}})$ point, the signal interpolation procedure involves starting from a reference $(m_X^{\text{ref}}, m_S^{\text{ref}})$ mass point, for which a fully simulated sample is available, and therefore can be used to extract the p.d.f. of the $(m_{b\bar{b}\gamma\gamma}^*, m_{b\bar{b}})$ variables. The interpolation procedure is used to es-

¹The modified four-body invariant mass $m_{b\bar{b}\gamma\gamma}^*$ is similar to the one described in Section 7.3.2, and is defined as $m_{b\bar{b}\gamma\gamma}^* = m_{b\bar{b}\gamma\gamma} - (m_{b\bar{b}} - m_S) - (m_{\gamma\gamma} - 125 \text{ GeV})$, where 125 GeV approximates the Higgs boson mass. This variable improves the resolution of the $b\bar{b}\gamma\gamma$ resonance compared to the traditional $b\bar{b}\gamma\gamma$ invariant mass.

timate $\mathcal{P}(m_{b\bar{b}\gamma\gamma}^*, m_{b\bar{b}} | m_X^{\text{target}}, m_S^{\text{target}})$ starting from $\mathcal{P}(m_{b\bar{b}\gamma\gamma}^*, m_{b\bar{b}} | m_X^{\text{ref}}, m_S^{\text{ref}})$, in two steps: a rescaling step, in which the peak of this distribution is scaled to take into account the distance between $(m_X^{\text{target}}, m_S^{\text{target}})$ and $(m_X^{\text{ref}}, m_S^{\text{ref}})$, and a reweighting step, that takes into account the different resolution of the $X \rightarrow b\bar{b}\gamma\gamma$ and $S \rightarrow b\bar{b}$ resonances at different mass points.

The rescaling step exploits the kinematics of the two-body decay $X \rightarrow SH$, which, in the rest frame of the mother particle, only depend from the masses of the three particles involved, m_X , m_S , and m_H . Indeed, in the rest frame of the X boson, the S scalar and the Higgs boson are generated back to back with the same momentum, whose magnitude is:

$$p^*(m_X, m_S) = \frac{\sqrt{[m_X^2 - (m_S + m_H)^2] \cdot [m_X^2 - (m_S - m_H)^2]}}{2m_X} \quad (8.2)$$

In the reference sample, the four-momenta of the S , H , and X particles are reconstructed for each event using the di-jet, diphoton, and $b\bar{b}\gamma\gamma$ objects in the final state, respectively. These four-momenta are then recomputed in the rest frame of the X particle, and the three-momentum components of the S and H bosons are rescaled using the ratio of the momentum magnitudes, computed with the target and reference masses:

$$R^* = \frac{p^*(m_X^{\text{target}}, m_S^{\text{target}})}{p^*(m_X^{\text{ref}}, m_S^{\text{ref}})} \quad (8.3)$$

Additionally, the mass of the S particle in the rest frame is rescaled by $m_S^{\text{target}}/m_S^{\text{ref}}$. The rescaled four-momenta are then boosted back to the laboratory frame, and the shifted $m_{b\bar{b}\gamma\gamma}^*$ and $m_{b\bar{b}}$ variables are recomputed. In this way, the modified $(m_{b\bar{b}\gamma\gamma}^*, m_{b\bar{b}})$ distribution is centered around the target $(m_X^{\text{target}}, m_S^{\text{target}})$ point.

The rescaling procedure adjusts the momenta of the X , S , and H particles but does not account for differences in experimental resolution (due to detector effects) of the $X \rightarrow b\bar{b}\gamma\gamma$ and $S \rightarrow b\bar{b}$ resonances at different mass points. In particular, the X and S resonances are expected to have a poorer resolution for lower m_X and m_S values, as shown in Figure 8.1. To address this, a reweighting is applied to the interpolated sample to ensure that the two-dimensional resolution of the $(m_{b\bar{b}\gamma\gamma}^*, m_{b\bar{b}})$ distribution matches the expected resolution at the target $(m_X^{\text{target}}, m_S^{\text{target}})$ point. Since detector resolution effects are much more significant for jets than for photons, the reweighting only accounts for the $m_{b\bar{b}}$ resolution. This is measured at fully simulated (m_X, m_S) points by fitting the $(m_{b\bar{b}} - m_S)/m_S$ distribution with a Bukin function [234]. The Bukin fit parameters are then interpolated across the (m_X, m_S) plane using Delaunay triangulation [235], thus creating two-dimensional maps of these parameters in the resonance masses. The interpolated Bukin parameters can be used to construct the resolution response for each event i with dijet invariant mass $m_{b\bar{b}}^i$, at any (m_X, m_S) point:

$$\text{Resolution response}^i(m_X, m_S) = \text{Bukin} \left(\frac{m_{b\bar{b}}^i - m_S}{m_S} \middle| m_X, m_S \right) \quad (8.4)$$




Figure 8.1 – Distributions of (a) $(m_{b\bar{b}\gamma\gamma}^* - m_X)/m_S$ and (b) $(m_{b\bar{b}} - m_S)/m_S$ for simulated $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signals with different resonance masses. All the distributions are normalised to unity.

Therefore, the final interpolated sample is obtained by applying an event-by-event weight, $w^{\text{res},i}$, which is calculated as the ratio of the Bukin functions evaluated at the target and reference (m_X, m_S) points, multiplied by the event's dijet resolution at the target $(m_X^{\text{target}}, m_S^{\text{target}})$ point:

$$w^{\text{res},i}(m_X, m_S) = \frac{\text{Bukin}\left(\frac{m_{b\bar{b}}^i - m_S^{\text{target}}}{m_S^{\text{target}}} \middle| m_X^{\text{target}}, m_S^{\text{target}}\right)}{\text{Bukin}\left(\frac{m_{b\bar{b}}^i - m_S^{\text{ref}}}{m_S^{\text{ref}}} \middle| m_X^{\text{ref}}, m_S^{\text{ref}}\right)} \cdot \frac{m_{b\bar{b}}^i - m_S^{\text{target}}}{m_S^{\text{target}}} \quad (8.5)$$

This interpolation procedure is considered robust only in regions of the (m_X, m_S) plane where the differences in resolution effects between neighboring fully simulated signal samples are small (typically, below 2-3%). Therefore, the interpolation is restricted to the region where $m_X > 300$ GeV and $m_S > 70$ GeV. The quality of the interpolation is evaluated by comparing the interpolated signal shapes with those from fully simulated ones, for (m_X, m_S) points where signal MC samples are available, and evaluating the expected 95% CL upper limits on the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ cross section times branching fraction in the two cases. In the region where the interpolation is applied, the difference in the expected analysis performance between using the interpolated signals and the simulated signals is found to be within 5%, while it can reach up to 10% at $m_S = 70$ GeV, at the boundary of the interpolation domain. To account for this difference, a dedicated systematic uncertainty related to the interpolation procedure is assigned to the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signals at (m_X, m_S) points where interpolation is applied, as described in Section 8.5.3.

In regions where the interpolation procedure is not applied, the signal grid consists exclusively of fully simulated samples.

The final signal grid, illustrated in Figure 8.2, consists of 373 mass points, with a spacing

of 5 GeV in the densest region (for low m_X and m_S values) and of up to 50 GeV in the less granular region, characterised by higher resonance masses and wider experimental widths. Figure 8.3 shows the $m_{b\bar{b}\gamma\gamma}^*$ and $m_{b\bar{b}}$ distributions for $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ samples at neighbouring (m_X, m_S) points in the signal grid. The distributions largely overlap in both the m_X and m_S directions, suggesting that the granularity of the signal grid is fine enough to have an effectively continuous sensitivity across the explored parameter space. The required grid granularity was ultimately determined through signal injection




Figure 8.2 – (m_X, m_S) points in the search space for which the presence of a $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal is tested. Simulations are available for the signals marked as black triangles and red squares, representing points analysed with two different analysis selections (that will be described in Section 8.2). For the points marked as blue circles, the results rely on interpolated signal shapes.

tion tests on background only simulated events, where an artificial $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal was introduced at a specific (m_X^*, m_S^*) point, with a cross section times branching fraction set to twice the expected limit under the background-only hypothesis. The injected signals consistently appeared as an excess with a discovery significance of at least 3 standard deviations in nearby (m_X, m_S) points (where no signal was injected). This confirms that the signal grid is sufficiently fine to ensure no gaps in sensitivity exist between the tested mass points.

8.2 Event selection

The object and event selection criteria for the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ search closely align with those used in the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis, described in Section 7.3.1. Both analyses share the same diphoton selection to target the $H \rightarrow \gamma\gamma$ signature. In addition, both analyses place a veto on events with at least one lepton, and require at least two, but fewer than six, central jets ($|\eta| < 2.5$) in the final state, to reduce contamination from the $t\bar{t}H$ and




Figure 8.3 – Distributions of (a) $m_{b\bar{b}\gamma\gamma}^*$ and (b) $m_{b\bar{b}}$ for simulated or interpolated $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal events, with neighbouring resonance masses (m_X, m_S) in the (a) m_X and (b) m_S directions in the search space. All the distributions are normalised to unity.

$t\bar{t}\gamma\gamma$ backgrounds. In addition, b -tagging requirements are applied, to target the $S \rightarrow b\bar{b}$ signature.

The topology of the $S \rightarrow b\bar{b}$ decay changes significantly depending on the region in the (m_X, m_S) parameter space. Specifically, in regions where $m_X \gg m_S + m_H$ and m_S is low, the S particle can acquire a very large momentum to compensate for this mass imbalance, becoming so boosted that the two bottom quarks from its decay are so collimated that they cannot be resolved and are reconstructed as a single small-radius b -jet. On the other hand, for $m_X \sim m_S + m_H$, the $S \rightarrow b\bar{b}$ decay typically results in two well-separated b -jets, arising from the hadronization of the two bottom quarks. Consequently, the b -tagging criteria differ from those used in the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis, and are adapted to reflect the distinct kinematics of the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal at various mass points.

Two b -tagging categories are defined, based on the expected topology of the $S \rightarrow b\bar{b}$ decay in different regions of the (m_X, m_S) plane:

- For (m_X, m_S) points where $m_S/m_X < 0.09$, the S particle is typically highly boosted, resulting in a single b -jet in the final state. For these point, the *1 b-tagged selection* is applied, requiring exactly one b -tagged jet in the final state, selected by the 77% efficiency working point of the DL1r tagger (introduced in Section 6.2).
- For (m_X, m_S) points where $m_S/m_X > 0.09$, a resolved topology of the $S \rightarrow b\bar{b}$ decay is expected. Therefore, when analysing these points, exactly two b -tagged jets are required in the final state, selected by the 77% efficiency working point of DL1r. This category is referred to as the *2 b-tagged category*.

The threshold of $m_S/m_X = 0.09$, used to distinguish between regions where a resolved topology is expected from those where a boosted topology is expected, is an empirical choice based on studies of the selection efficiency for the 1 b -tagged and 2 b -tagged categories for $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signals across the (m_X, m_S) plane. In fact, Figure 8.4 shows the fraction of signal events selected in the 2 b -tagged category at different (m_X, m_S)

points in the signal grid, with respect to the total signal yields considering both categories. The dashed line represents the empirical threshold of $m_S/m_X = 0.09$. Points above the threshold are identified as those where the $S \rightarrow b\bar{b}$ resonance typically has a resolved topology, and indeed, the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signals have a large fraction of events passing the 2 b -tagged selection (around 90%). On the other hand, for points below the threshold, the $S \rightarrow b\bar{b}$ resonance typically exhibits a boosted topology, and the selection efficiency for the 2 b -tagged category is very low (below 40%), meaning that most events (at least 60%) pass the 1 b -tagged selection.




Figure 8.4 – Fraction of signal events selected in the 2 b -tagged category, with respect to the overall signal yields in both the 1 b -tagged and 2 b -tagged category, as a function of the m_X and m_S masses. The dashed line represents the empirical threshold $m_S/m_X = 0.09$ that separates the (m_X, m_S) points analysed with the 2 b -tagged selection from those analysed with the 1 b -tagged selection.

Figure 8.5 shows the signal efficiency for the 1 b -tagged and 2 b -tagged category as a function of the m_X and m_S masses. The 1 b -tagged category selects $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal with resonance masses $m_S/m_X < 0.09$ with an efficiency ranging from 2% in the low m_X region to 12% in the high m_X region. Similarly, the 2 b -tagged category, which targets signals with $m_S/m_X > 0.09$, exhibits lower efficiency in the low m_X region, reaching up to 20% for higher m_X values. The reduced efficiency at lower m_X in the 2 b -tagged category is due to the softer p_T spectra of the selected b -jets compared to those at higher m_X , as illustrated in Figure 8.6. This leads to a larger fraction of b -jets falling below the threshold for jet reconstruction (25 GeV, as discussed in Section 7.3.1).




Figure 8.5 – Signal efficiency for the 1 b -tagged (left) and 2 b -tagged categories as a function of the (m_X, m_S) mass point in the search space.




Figure 8.6 – Distribution of the leading b -jet p_T for $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signals for different masses m_X and m_S , selected by the 2 b -tagged category. Object reconstruction criteria, outlined in Section 7.3.1, require the jet's p_T to be above 25 GeV.

8.3 PNN discriminant

The final discriminant for the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ analysis is based on the output of a parametrized neural network (PNN) [232]. A PNN is a neural network that simultaneously processes event features - represented by the vector $\mathbf{x} = (x_1, x_2, \dots, x_N)$, where each x_i denotes physical quantities such as kinematic variables of final-state objects - and the parameters of the search space, represented by $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_m)$. In this analysis, $\boldsymbol{\theta}$ corresponds to the masses of the two resonances, m_X and m_S . For each event j , characterized by its kinematic features $\mathbf{x}^j$, the PNN response is a continuous function of the search parameters $\boldsymbol{\theta}$. Figure 8.7 presents a simplified illustration of how the PNN operates. Given a point $\bar{\boldsymbol{\theta}}$ in the search space, the PNN evaluated at that point, $\text{PNN}(\bar{\boldsymbol{\theta}})$,




Figure 8.7 – Sketch illustrating the concept of how the PNN operates. For each event, the PNN response evaluated on its input features $\mathbf{x}$ is a continuous function of the search parameters $\boldsymbol{\theta}$.

is designed to provide the optimal separation between background events and signal events with the parameter $\bar{\boldsymbol{\theta}}$. Moreover, $\text{PNN}(\bar{\boldsymbol{\theta}})$ can also detect a signal with nearby parameters $\bar{\boldsymbol{\theta}}'$, although with a slightly reduced sensitivity. These properties of the PNN are particularly interesting for the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ analysis, which searches for a signal in a wide, two-dimensional space defined by the masses of the two resonances, m_X and m_S . The search is further complicated by the non-trivial dependence of the signal on the (m_X, m_S) masses. Ideally, the PNN allows an optimal signal-to-background separation at any mass point in the search space, with a single (parametric) discriminant. Furthermore, the ability of the discriminant evaluated at a specific mass point (m_X^*, m_S^*) , $\text{PNN}(m_X^*, m_S^*)$, to remain sensitive to signals with nearby masses m_X' and m_S' is crucial for ensuring continuous sensitivity across the entire search space, when combined with a sufficiently granular signal grid.

Two separate PNNs are trained for the 1 b -tagged and 2 b -tagged categories to isolate a mixture of benchmark $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signals from the dominant backgrounds. The

training procedure for each PNN is outlined below:

- 2 b -tagged category.** The PNN is trained using all simulated signals with $m_S > 15$ GeV, and the backgrounds include the $\gamma\gamma$ +jets sample and the main single Higgs processes (ggF, ttH, and ZH). The feature vector is $\mathbf{x} = (m_{b\bar{b}\gamma\gamma}^*, m_{b\bar{b}})$, while the parameter vector is $\boldsymbol{\theta} = (m_X, m_S)$. Since the mass parameters are irrelevant for the background processes, and to prevent the PNN from distinguishing signal from background purely based on the search parameters, background events are assigned random (m_X, m_S) values during training, following the same distribution as the signal events.
- 1 b -tagged category.** The PNN is trained on simulated signals with $m_S = 15$, plus the signals with (m_X, m_S) values equal to (400, 30), (500, 30), (750, 70), (1000, 70) GeV, which are also analysed using the 1 b -tagged category. The backgrounds include the $\gamma\gamma$ +jets sample and the same single Higgs processes considered for the 2 b -tagged category, plus single Higgs production via VBF and the SM HH processes. In the 1 b -tagged region, a single b -jet (incorporating both bottom quarks from the boosted S particle's decay) is reconstructed in the final state. Hence, the feature vector is changed to accomodate this particular topology, and includes the modified three-body invariant mass $m_{b\gamma\gamma}^* = m_{b\gamma\gamma} - (m_{\gamma\gamma} - 125 \text{ GeV})$ and the p_T of the selected b -jet: $\mathbf{x} = (m_{b\bar{b}\gamma\gamma}^*, p_T(b))$. The parameters' vector only includes the mass of the X resonance: $\boldsymbol{\theta} = (m_X)$. Like in the 2 b -tagged category, the background events are assigned random (m_X) values for the training, following the same m_X distribution as the signals.

Figure 8.8 shows the distribution of the 2 b -tagged PNN discriminant for two different mass points, comparing the corresponding $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signals with the continuum, single Higgs, and SM HH backgrounds. Since the PNN response for each event depends on the mass parameters, the background shapes also vary between the PNN spectra evaluated at different mass points. In Figure 8.9, the 2 b -tagged PNN discriminant, evaluated at $(m_X, m_S) = (250, 100)$ GeV, is shown for continuum, single Higgs, and SM HH backgrounds, as well as for three $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signals with resonance masses $(m_X, m_S) = (250, 100)$, (240, 100), and (300, 110) GeV. The discriminant $\text{PNN}(m_X = 250 \text{ GeV}, m_S = 100 \text{ GeV})$ effectively separates the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal at $(m_X, m_S) = (250, 100)$ GeV from the backgrounds, and also classifies the signal at the nearby mass points $(m_X, m_S) = (240, 100)$ GeV as signal-like, as its distribution peaks near 1. On the other hand, the discriminant is less sensitive to the signal with more distant mass values, such as $(m_X, m_S) = (300, 110)$ GeV, whose distribution in the $\text{PNN}(m_X = 250 \text{ GeV}, m_S = 100 \text{ GeV})$ spectrum appears to be more background-like. As already mentioned, this property of the $\text{PNN}(\boldsymbol{\theta}')$ discriminant, evaluated at a specific point $\boldsymbol{\theta}' = (m_X', m_S')$ in the search space, to retain sensitivity to signals with nearby values of resonance masses ensures that the analysis would not miss a signal whose values of m_X and m_S fall between the (m_X, m_S) points where the presence of a signal is actually tested. This allows to have a fully continuous sensitivity across the explored search space, even with a finite granularity of the signal grid.




Figure 8.8 – PNN output distributions for simulated signal and background events for different (m_X, m_S) values in the 2 b -tagged category. The $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signals are characterised by the same m_X and m_S values as the PNN discriminants. All distributions are normalised to unity.




Figure 8.9 – Discriminant distributions for the PNN at $\theta = (250, 100)$ GeV for several background and signal samples in the 2 b -tagged category. The distribution of the PNN(θ) output is shown for signals generated with $(m_X, m_S) = (250, 100)$ GeV, $(240, 100)$ GeV and $(300, 110)$ GeV. All distributions are normalised to unity.

8.4 Analysis strategy and background estimation

The use of PNN discriminants as the final observable requires a complete rethinking of the background modeling strategy, especially for the continuum background, compared to traditional $H \rightarrow \gamma\gamma$ analyses, which rely on the diphoton invariant mass $m_{\gamma\gamma}$ as the final discriminant. In the approach adopted by this analysis, both signal and background processes are modeled using histograms of the PNN discriminants derived from simulations.² A slightly different treatment is adopted for the continuum background, since the available $\gamma\gamma$ +jets sample only accounts for the true photon component, omitting the (subdominant) fake photon contribution. Special care is taken to ensure that the continuum background template extracted from the $\gamma\gamma$ +jets sample is reliable for modeling the entire continuum background, including the fake photon component.

8.4.1 SM HH and single Higgs boson background

The SM HH and single Higgs boson backgrounds are modeled using the PNN discriminant distributions obtained from the corresponding MC samples. The overall normalisation of these background processes is determined based on their SM cross sections, calculated using the state-of-the-art accuracy (quoted in Table 1.2 and Section 2.1.1), on the branching ratio of the $H \rightarrow \gamma\gamma$ decay for single Higgs events, or the $HH \rightarrow b\bar{b}\gamma\gamma$ decay for di-Higgs events (corresponding to 0.227% and 0.26% respectively), and on the integrated luminosity of the Run 2 dataset (equal to 140 fb^{-1}).

8.4.2 Continuum background

As previously mentioned, the continuum background is primarily composed of processes with two reconstructed photons in the final state, accompanied by additional jets. Subdominant contributions come from the $t\bar{t}\gamma\gamma$ and $Z + \gamma\gamma$ processes, where the Z boson decays hadronically. These two subdominant backgrounds are modeled using MC samples, and their normalisation is determined from the generator cross sections (quoted in Table 7.3 for $t\bar{t}\gamma\gamma$, and corresponding to 68.9 fb for $Z + \gamma\gamma$).

The continuum background involving two reconstructed photons plus additional jets includes the contributions from events with two real photon and from events with fake photons, arising from the γj and jj processes. The fractions of the irreducible true $\gamma\gamma$ component and the reducible γj and jj components in the 1 b -tagged and 2 b -tagged categories are measured using a two-dimensional matrix method, detailed in Section 7.4.2. These purity values, along with their statistical and systematic uncertainties, are shown in Table 8.1.

The SHERPA $\gamma\gamma$ +jets sample, which simulates the true photon component, is used to model the entire reconstructed diphoton background, including the reducible γj and jj components, which contribute around 15% in both 1 b -tagged and 2 b -tagged categories. This is justified because, as explained below, the shape of the PNN discriminants for the $\gamma\gamma$ +jets process was found to be consistent with those of the reducible diphoton

²For (m_X, m_S) points without fully simulated $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signals, the PNN discriminant is estimated using interpolated samples, as described in Section 8.1.

Table 8.1 – Fraction of $\gamma\gamma$, γj , and jj events contributing to the continuum background involving two reconstructed photons in the 1 b -tagged and 2 b -tagged categories. The central values are shown along with their statistical (“stat.”) and systematic (“syst.”) uncertainties.

Category	$\gamma\gamma$ [%]	γj [%]	jj [%]
1 b -tagged	86.8 ± 0.9 (stat.) $^{+2.9}_{-0.7}$ (syst.)	12.0 ± 0.9 (stat.) $^{+0.5}_{-2.5}$ (syst.)	1.2 ± 0.3 (stat.) $^{+0.2}_{-0.4}$ (syst.)
2 b -tagged	84.7 ± 2.9 (stat.) $^{+3.2}_{-1.5}$ (syst.)	13.9 ± 2.8 (stat.) $^{+2.1}_{-0.2}$ (syst.)	1.4 ± 0.8 (stat.) $^{+0.0}_{-0.5}$ (syst.)

backgrounds. To account for the difference in normalisation between the true photon component and the full reconstructed diphoton background, the $\gamma\gamma$ +jets template is scaled by a normalisation factor, $\mu_{\gamma\gamma}$, which is allowed to float freely in the final fit to data. Hence, the normalisation of the reconstructed photon background is data-driven, while the shape of the PNN discriminants is taken from MC simulations.

To ensure that the $\gamma\gamma$ +jets sample accurately describe the full reconstructed diphoton background, templates of the key discriminant variables (namely, $m_{b\bar{b}\gamma\gamma}^*$ and $m_{b\bar{b}}$ in the 2 b -tagged category, $m_{b\gamma\gamma}^*$ and $p_T(b)$ in the 1 b -tagged category, and all the PNN discriminants) are extracted for γj and jj events, and compared with the corresponding distribution from the $\gamma\gamma$ +jets sample. The γj and jj templates are built using the same two-dimensional matrix method used to measure the $\gamma\gamma$, γj , and jj purities, namely, by inverting the identification or isolation requirements of the two reconstructed photons in data. The fake photon background templates are found to be consistent within statistical uncertainties with the corresponding distributions from the $\gamma\gamma$ +jets sample, describing the true photons component. Therefore, applying the data-driven $\gamma\gamma$ normalisation factor effectively accounts for the reducible fake photon backgrounds, while any impact on the shape is considered small and is covered by the diphoton background modeling uncertainty, discussed in Section 8.5.4.

Finally, the accuracy of the continuum background model in describing the data is validated in a region where no signal or other resonant background is expected, and the background is dominated entirely by continuum processes. This region, defined by the $m_{\gamma\gamma}$ sidebands (SB) is used to compare the shapes of key discriminating variables ($m_{b\bar{b}\gamma\gamma}^*$, $m_{b\bar{b}}$, $m_{b\gamma\gamma}^*$, $p_T(b)$, and PNN discriminants) between the continuum background model and the data after performing a background-only fit to the data. The background-only fit is carried out exclusively in the $m_{\gamma\gamma}$ SB, using single-binned templates of the PNN output. The shape information from the PNN discriminants is not considered in this fit, and only the normalisation of the $\gamma\gamma$ +jets template is free to float. Figure 8.10 illustrates this comparison for $m_{b\bar{b}\gamma\gamma}^*$ and $m_{b\bar{b}}$ in the 2 b -tagged category, and $m_{b\gamma\gamma}^*$ and $p_T(b)$ in the 1 b -tagged category. Since $\text{PNN}(\vec{\theta})$ is a unique observable for each value of $\vec{\theta} = (m_X, m_S)$, this validation is performed for each (m_X, m_S) point in the signal grid. For example, Figure 8.11 shows the PNN discriminant distributions for two different mass points: $(m_X, m_S) = (250, 100)$ GeV in the 2 b -tagged category, and $(1000, 70)$ GeV in the 1 b -tagged category, comparing data and the continuum background model in the $m_{\gamma\gamma}$ SB after normalising the $\gamma\gamma$ +jets contribution with the factor extracted from the background-only fit described above. A good agreement between the data and the con-

tinuum background is observed across the entire search space, confirming the robustness of the background model.




Figure 8.10 – Distributions of (a) $m_{b\bar{b}\gamma\gamma}^*$, (b) $m_{b\bar{b}}$, (c) $m_{b\gamma\gamma}^*$, and (d) $p_T(b)$ in data and in the predicted model, in the $m_{\gamma\gamma}$ sidebands of the 2 b -tagged region (top) and 1 b -tagged region (bottom). The $\gamma\gamma$ +jets background is rescaled to its post-fit normalisation in a background-only fit in the SB region only. The $\gamma\gamma$ +jets template is used to model the full continuum diphoton background, including both the true and the fake photon components. The error band corresponds to the dominant modelling uncertainty, which will be discussed in Section 8.5.4, which arises from the non-resonant $\gamma\gamma$ +jets background.

8.4.3 Analysis regions

The presence of a $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal for a given (m_X, m_S) mass point is tested by performing a simultaneous fit to data in two orthogonal regions, defined using the diphoton invariant mass, $m_{\gamma\gamma}$, described below.

- A **Signal region (SR)** is defined by $120 \leq m_{\gamma\gamma} \leq 130$ GeV, corresponding to the window around the $H \rightarrow \gamma\gamma$ peak. Any potential $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal is expected to contribute to the $H \rightarrow \gamma\gamma$ resonance, centered at $m_H \approx 125$ GeV. Therefore, this region drives the sensitivity to the resonant signal. To enhance the signal-




Figure 8.11 – Distributions of the PNN discriminant for two choices of $\theta = (m_S, m_X)$ in data and in the predicted background model, in the $m_{\gamma\gamma}$ sidebands of the 2 b -tagged region (left) and 1 b -tagged region (right). The normalisation of the $\gamma\gamma$ +jets background is adjusted to match data, using a background only fit. The $\gamma\gamma$ +jets template models the full continuum diphoton background, whose uncertainty, represented by the shaded gray area, arises from the modelling of the $\gamma\gamma$ +jets template (as will be discussed in Section 8.5.4).

to-background ratio, the PNN templates in the SR are finely binned, especially in the high PNN score range where the signal is expected to peak.

- A **Sidebands region (SB)** is defined by $105 \leq m_{\gamma\gamma} < 120$ GeV plus $130 \leq m_{\gamma\gamma} \leq 160$ GeV, and is expected to include almost null contribution from any $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signals or from resonant backgrounds (namely, single Higgs or SM HH processes). The main role of the SB region is to constrain the normalisation of the continuum background, $\mu_{\gamma\gamma}$, in the final fit to data. Therefore, a single bin is used to construct the PNN templates in the SB.

The binning of the PNN templates in the SR is optimised individually for each point in the parameter space. The definition of the bin widths begins in the most signal-like region of the PNN spectrum, where the PNN score is near 1, and progressively moves toward the more background-like region, where the PNN score is closer to zero. The width of the most signal-like bin is set to ensure that at least one background event is expected in that bin. This choice was found to be sufficient to ensure the validity of the asymptotic approximation [70], adopted for the statistical interpretation of the results³. Subsequent bins are iteratively defined, with each bin containing one additional expected background event compared to the previous bin. This process continues until the signal-to-background ratio in a bin becomes lower than that of the unbinned SR. At that point, a single bin is created for the remaining of the distribution, at PNN output values close to zero.

³The accuracy of the asymptotic approximation was verified by generating pseudo-experiments to estimate the distributions of the test statistic required for determining the upper limits on the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ cross-section. The discrepancy between these results and those based on the asymptotic approximation was found to be less than 8% across the entire (m_X, m_S) plane.

8.5 Systematic uncertainties

Similar to the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis, the sensitivity of the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ search is limited by the finite amount of data collected by the ATLAS detector during Run 2, particularly within the 2 b -tagged category selection. Nevertheless, it is important to understand and evaluate all the sources of systematic uncertainties that may affect the final results.

Since the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ search and the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis share a very similar final state, and they both rely on MC simulations for modeling single Higgs and di-Higgs processes, many sources of systematic uncertainties are common to both analyses, including:

- the 0.83% uncertainty propagated from the measurement of the total integrated luminosity of Run 2 data [49];
- the experimental systematic uncertainties due to potential mismodelling of the detector response in simulation with respect to data, related to the modelling of photons, jets, flavor tagging efficiencies, and pileup;
- the theoretical uncertainties affecting predictions for the single Higgs and di-Higgs production cross sections, as well as the $H \rightarrow \gamma\gamma$ and $H \rightarrow b\bar{b}$ branching fractions (quoted in Table 1.2, Section 2.1.1, and Section 7.5.2);
- the 100% uncertainty related to the heavy-flavor content in single Higgs production via ggF, VBF, or WH (already explained in Section 7.5.2);
- the “scale” uncertainties, reflecting missing higher-order corrections in theoretical calculations, and the uncertainties arising from the PDFs and the α_s value (referred to as “PDF+ α_s ” uncertainties), affecting MC event generation (already discussed in Section 7.5.2);
- the parton shower modelling uncertainties in MC samples (introduced in Section 7.5.2).

These uncertainties may impact either just the normalisation or both the normalisation and shape of the PNN templates for both signal and background processes.

A key difference between the treatment of systematic uncertainties in the $HH \rightarrow b\bar{b}\gamma\gamma$ and in the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ analyses comes from the modelling of the continuum background. In the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ analysis, the modeling of this background is not fully data-driven. Instead, it relies on MC simulations to extract the PNN discriminant distribution, with the normalisation adjusted to data in the final fit using a dedicated control region (the $m_{\gamma\gamma}$ sidebands). Therefore, the experimental and theoretical uncertainties mentioned above, such as those related to detector response mismodelling, luminosity measurement, parton shower and event generator modelling, and the scale and PDF+ α_s uncertainties, also impact the continuum background processes.

In addition, a custom *modelling uncertainties* affect the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ search, related to the signal interpolation procedure described in Section 8.1.1. This *signal interpolation uncertainty* applies to (m_X, m_S) points where interpolated samples are used to construct the PNN template for the signal process.

The impact of the systematic uncertainties is typically evaluated by varying each source

by $\pm 1\sigma$ upstream in the analysis chain, and propagating these variations through the entire analysis. For each uncertainty source, varied PNN histograms are produced, reflecting differences from the nominal histograms either in normalisation or in both normalisation and shape. These varied histograms are incorporated in the statistical model to account for the systematic uncertainties. Since the PNN discriminant, $\text{PNN}(m_X, m_S)$, is effectively a different observable for different mass points in the search space, the systematic uncertainties are evaluated independently for each (m_X, m_S) point within the signal grid.

This procedure is more complicated when extracting varied PNN templates (reflecting the $\pm 1\sigma$ variations from theoretical or experimental uncertainties) for $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signals in mass points where an interpolated signal sample is used to estimate the nominal PNN template. In such cases, the interpolation procedure is applied to varied reference samples, with $\pm 1\sigma$ systematic variations introduced upstream. This results in varied interpolated target signals, which are then used to construct the corresponding varied interpolated PNN templates.

8.5.1 Experimental systematic uncertainties

The 0.83% luminosity uncertainty affects only the normalisation of the PNN templates for both signal and background processes.

The other experimental systematic uncertainties, arising from mismodelling of photon identification and isolation efficiencies, photon energy scale and resolution, Jet Vertex Tagger efficiency, jet energy scale and resolution, flavor tagging efficiencies, and pileup, affect both the shape and normalisation of the PNN templates for signal and background processes. These uncertainties cause variations in the overall yields and efficiency of the single Higgs and HH backgrounds by approximately 1% to 7%. For the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ process, the impact ranges from 1% to as much as 15%, observed for jet energy resolution uncertainties for signals with low m_S values. For instance, Figure 8.12 shows the nominal $\text{PNN}(\theta)$ templates alongside their $\pm 1\sigma$ variations due to flavor tagging uncertainties, for to the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal, single Higgs production via ggF, SM ggF HH production, and continuum background, in the mass point $(m_X, m_S) = (250, 100)$ GeV.

The overall impact of experimental systematic uncertainties on the expected limits for the cross section times branching fraction of the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal, $\sigma(pp \rightarrow X) \times \mathcal{B}(X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma)$, is minimal - below 1% - for high resonance masses ($m_X > 400$ GeV). However, in regions of the search space with lower resonance masses, the impact ranges from 2% to 20%. The largest effects (up to 20%) are observed at the (m_X, m_S) points analysed with the 1 b -tagged category selection, and are dominated by the uncertainties on the jet energy resolution.

8.5.2 Theoretical systematic uncertainties

The theoretical systematic uncertainties from the cross sections of single Higgs and di-Higgs processes, as well as from the $H \rightarrow \gamma\gamma$ and $H \rightarrow b\bar{b}$ branching fractions, influence the normalisation of the PNN templates for single Higgs and SM HH backgrounds. The




Figure 8.12 – Nominal and varied $\text{PNN}(\theta)$ distributions for the (a) signal, (b) ggF H, (c) SM ggF HH, and (d) $\gamma\gamma$ +jets processes, in the mass point $(m_X, m_S) = (250, 100) \text{ GeV}$. The varied PNN templates refer to $\pm 1\sigma$ variations due to the flavor tagging uncertainties listed in the legend.

100% systematic uncertainty related to the heavy flavor simulation is assigned to the normalisation of the ggF, VBF, and WH single Higgs backgrounds.

Additionally, scale and PDF+ α_s uncertainties affect both the normalisation and shape of the PNN templates for the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signals, and for single Higgs, di-Higgs, $\gamma\gamma$ +jets, $t\bar{t}\gamma\gamma$, and $Z + \gamma\gamma$ backgrounds.

Parton shower uncertainties also impact both the normalisation and shape of the PNN templates. These are evaluated only for the signals and for the resonant single Higgs and SM HH backgrounds. For the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signals, the parton shower uncertainties are assessed by comparing the PNN templates extracted from nominal MC samples, relying on PYTHIA 8.2 for event generation and parton showering, with those obtained from alternative MC samples based on HERWIG 7.1. The alternative single Higgs and SM HH MC samples used to evaluate parton shower uncertainties rely on HERWIG 7, as already discussed in Section 7.5.2.

For the three single Higgs production modes via ggF, VBF, and WH, scale, PDF+ α_s , and parton shower uncertainties are not considered, as these effects are assumed to be covered by the large 100% uncertainty associated with the heavy-flavor content simulation. The overall impact of theoretical uncertainties on the analysis sensitivity is approximately 4%. However, in regions of the (m_X, m_S) plane characterized by low m_S , this impact can rise to as much as 10%, mostly due to the contribution of the 100% heavy flavor uncertainty assigned to the ggF single Higgs production.

8.5.3 Signal interpolation uncertainty

A dedicated uncertainty is assigned to the signal interpolation procedure, described in Section 8.1.1, which is used to estimate the PNN discriminant distribution for (m_X, m_S) points where fully simulated signal samples are not available. This systematic uncertainty is evaluated by modifying the interpolated samples, varying, in turn, each of the fit parameters of the Bukin functions (used to estimate the $m_{b\bar{b}}$ resolution response at the target $(m_X^{\text{target}}, m_S^{\text{target}})$ point) by $\pm 1\sigma$ within their fit uncertainties. These modified interpolated samples generate slightly different probability density functions $\mathcal{P}(m_{b\bar{b}\gamma\gamma}^*, m_{b\bar{b}} | m_X^{\text{target}}, m_S^{\text{target}})$, which are then used to construct corresponding varied PNN templates. The envelope of the differences between the nominal and varied PNN templates is adopted as a systematic uncertainty. This uncertainty is applied to signal samples at (m_X, m_S) points where the interpolation procedure is used. The interpolation uncertainty typically leads to bin-by-bin variations of up to 30% in the PNN templates, as illustrated in Figure 8.13. Furthermore, as shown in the same Figure, this uncertainty typically covers the non-closure observed between interpolated and fully simulated signals, verified for the points where fully simulated signals are available.

The signal interpolation uncertainty was found to have a maximum impact of 10% on the analysis sensitivity across all tested mass points.




Figure 8.13 – Comparison of simulation and interpolation for the PNN discriminant, for $m_X^{\text{target}} = 750 \text{ GeV}$ and $m_S^{\text{target}} = 200 \text{ GeV}$. The shaded green bands around the interpolated histogram shows the signal interpolation uncertainties. The contributions of statistical and experimental systematic uncertainty are also shown.

8.5.4 Continuum background modelling uncertainty

As mentioned earlier, experimental, scale and PDF+ α_s , and parton shower uncertainties all influence both the normalisation and shape of the PNN templates used to model the continuum background, derived from the $\gamma\gamma$ +jets, $t\bar{t}\gamma\gamma$, and $Z + \gamma\gamma$ MC samples.

A dedicated systematic uncertainty is assigned to the modelling of the continuum diphoton background, which is described using the $\gamma\gamma$ +jets sample. This uncertainty is evaluated by comparing the nominal PNN templates (extracted from the SHERPA sample) with those obtained from the alternative MADGRAPH5_AMC@NLO sample. However, as discussed in Section 8.1, the MADGRAPH5_AMC@NLO $\gamma\gamma$ +jets sample suffers from significantly lower statistics in the signal regions of the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ analysis with respect to the nominal SHERPA sample, particularly in the 2 b -tagged category. This results in an inflated magnitude of the modelling uncertainty due to large statistical fluctuations in the varied PNN templates. To mitigate the effect of low statistics in the MADGRAPH5_AMC@NLO sample for the 2 b -tagged category, the varied PNN templates were built after removing the b -tagging requirements, and reweighing the simulated events to match the jet flavor composition of the 2 b -tagged category. Despite this smoothing procedure, the statistical fluctuations in the varied PNN templates used to estimate the modelling uncertainty remain significant.

Indeed, the continuum background modelling uncertainty is the dominant source of uncertainty in this analysis, inducing bin-by-bin efficiency variations in the $\gamma\gamma$ +jets templates reaching up to 100%, as illustrated in Figure 8.14. The impact on the expected analysis sensitivity for (m_X, m_S) mass points analysed with the 2 b -tagging category selection ranges from 2% at higher resonance masses ($m_X > 600$ GeV) to as much as 20% at lower m_X values.

In the 1 b -tagged category, the b -tagging requirements are less stringent, reducing the statistical limitations of the $\gamma\gamma$ +jets sample. As a result, the impact of this modelling uncertainty on the expected sensitivity is generally lower, amounting to only a few percent in most cases. However, for the two mass points $(m_X, m_S) = (500, 30)$ GeV and $(230, 15)$ GeV, significant statistical fluctuations in the MADGRAPH5_AMC@NLO sample in signal-like bins of the PNN histograms result in an overestimated magnitude of this uncertainty, translating into an impact on the analysis sensitivity reaching 40%.

8.6 Statistical model

The presence of a $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal at a specific (m_X, m_S) mass point is tested using a binned maximum likelihood fit to the PNN distribution in data. This fit is performed simultaneously in SR and in the SB. The likelihood function is expressed as fol-




Figure 8.14 – Nominal $\text{PNN}(\theta)$ distribution for the continuum diphoton background for $(m_X, m_S) = (250, 100)$ GeV, extracted from the SHERPA $\gamma\gamma$ +jets sample. The nominal template is compared with the corresponding varied histograms obtained from the MADGRAPH5_AMC@NLO sample, both with and without the smoothing procedure. The shaded bands represent the statistical uncertainty associated to each template.

lows:

$$\begin{aligned}
 \mathcal{L} = & \text{Pois} \left(n_{SB} | \mu_{\gamma\gamma} \cdot N_{SB}^{\gamma\gamma}(\alpha, \gamma) + \sum_p N_{SB}^p(\alpha, \gamma) \right) \\
 & \times \prod_i \text{Pois} \left(n_{SR,i} | \mu \cdot N_{SR}^{\text{signal}}(\alpha) \cdot f_i^{\text{signal}}(\alpha, \gamma) + \mu_{\gamma\gamma} \cdot N_{SR}^{\gamma\gamma}(\alpha) \cdot f_i^{\gamma\gamma}(\alpha, \gamma) \right. \\
 & \quad \left. + \sum_p N_{SR}^p(\alpha) \cdot f_i^p(\alpha, \gamma) \right) \\
 & \times G(\alpha) \times \text{Pois}(\gamma)
 \end{aligned} \tag{8.6}$$

The binned likelihood in Equation 8.6 is a product of multiple Poisson terms, modelling the number of observed events in data in the SB region (n_{SB}) and in each bin i of the PNN templates in the SR ($n_{SR,i}$). Additionally, it includes Gaussian constraint terms for the nuisance parameters (NPs) associated with systematic uncertainties (α) and Poisson constraints for the statistical uncertainties (γ), reflecting the limited statistics in the MC samples used to generate the PNN templates.

The Poisson term modelling n_{SB} is centered around the expected number of background events contributing to the single-bin PNN templates in the SB. This includes the continuum diphoton background contribution, with its expected yield from the $\gamma\gamma$ +jets MC sample (indicated as $N_{SB}^{\gamma\gamma}(\alpha)$ in Equation 8.6), and all the other background processes, represented by the sum of expected yields $N_{SB}^p(\alpha)$ for each process p . As mentioned in Section 8.4.2, the expected yield for the $\gamma\gamma$ +jets background is scaled by the normalisation factor $\mu_{\gamma\gamma}$, which is a free parameter in the fit and adjusts the contribution of the continuum background in order to match the data in the SB. This normalisation factor is applied in both the SB and SR, allowing the continuum background normalisation to propagate across both regions.

For the SR, the number of observed events in each bin i of the PNN template, $n_{SR,i}$, is modeled using a Poisson function, centered around the expected event yield in that bin. The expected event count includes contributions from the $\gamma\gamma$ +jets background, the single Higgs, SM HH, and other continuum backgrounds, as well as the signal. The signal expected event yield, $N_{SR}^{\text{signal}}(\alpha)$, is scaled by the signal strength parameter μ , corresponding to the product of the cross section and the branching fraction of the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ process (namely, $\mu = \sigma(pp \rightarrow X) \times \mathcal{BR}(X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma)$), which represents the parameter of interest (POI) of the model.

The shape information of the signal and background processes in the PNN spectrum is encoded by the terms f_i^p , which represent the fractional contributions of each process p in bin i of the PNN histograms in the SR.

The NPs α and γ are associated to systematic and statistical uncertainties, respectively. The systematic uncertainties affect both the expected yields and the shapes of the signal and background processes, while the statistical uncertainties only impact the shape, by fluctuating the fractional contributions of the signal and background processes in each bin. Statistical uncertainties are treated as uncorrelated across different analysis regions and within each bin of the PNN templates in the SR: each bin i is assigned an indepen-

dent NP γ_i .

For instance, for a process p with expected yield $N_{SR}^{p,MC} \cdot f_i^{p,MC}$ in the bin i in the SR, carrying a statistical uncertainty δ_i^p , the NP γ_i^p modifies the fractional yield as:

$$f_i^p(\gamma_i^p) = \gamma_i^p \cdot f_i^{p,MC}. \quad (8.7)$$

The Poisson constraint term for γ_i^p is defined as:

$$\text{Pois}(\gamma_i^p) = \text{Pois} \left(m_i^p = \left(\frac{N_{SR}^{p,MC} \cdot f_i^{p,MC}}{\delta_i^p} \right)^2 \mid \gamma_i^p \cdot \tau_i^p \right), \quad (8.8)$$

where m_i^p represents a global observable whose nominal value is $\left(\frac{N_{SR}^{p,MC} \cdot f_i^{p,MC}}{\delta_i^p} \right)^2$, and τ_i^p is a fixed constant, set to $\left(\frac{N_{SR}^{p,MC} \cdot f_i^{p,MC}}{\delta_i^p} \right)^2$. Thus, the term in Equation 8.8 constrains γ_i^p to be close to unity, allowing for Poisson fluctuations reflecting the MC statistical uncertainty in bin i for the process p . The global Poisson term $\text{Pois}(\gamma)$ in Equation 8.6 is the product of all these constraints.

Systematic uncertainties that impact the normalisation of process p are modeled by introducing a NP α_j , which modifies the expected yields in the SB and SR, $N_{SB(SR)}^{p,MC}$, by multiplying them by response function $\eta(\alpha_j)$, defined as:

$$\eta(\alpha_j) = \begin{cases} \left(I_R^p(+1\sigma) / I_R^p(\text{nom}) \right)^{\alpha_j} & \text{if } \alpha_j > 0 \\ \left(I_R^p(-1\sigma) / I_R^p(\text{nom}) \right)^{-\alpha_j} & \text{if } \alpha_j < 0 \end{cases} \quad (8.9)$$

Here, $I_R^p(\text{nom})$ and $I_R^p(\pm 1\sigma)$ represent the nominal and varied expected yields for process p in region R . The overall normalisation for process p in the region R is then expressed as:

$$N_R^p(\alpha) = N_R^{p,MC} \cdot \prod_j \eta(\alpha_j), \quad (8.10)$$

where j indexes all the systematic uncertainties affecting the normalisation of process p . Systematic uncertainties affecting both the normalisation and the shape of the PNN template for the process p also modify its fractional contribution f_i^p in each bin i in the SR. The fractional yield, originally obtained from the MC simulation as $f_i^{p,MC}$, is modified as follows:

$$f_i^p(\alpha) = f_i^{p,MC} \cdot \left(1 + \sum_j \sigma(\alpha_j) \right) \quad (8.11)$$

The response function $\sigma(\alpha_j)$ in Equation 8.11 is defined as:

$$\sigma(\alpha_j) = \begin{cases} \left(I_{SR,i}^p(+1\sigma) - I_{SR,i}^p(\text{nom}) \right) \cdot \alpha_j & \text{if } \alpha_j > 0 \\ \left(I_{SR,i}^p(\text{nom}) - I_{SR,i}^p(-1\sigma) \right) \cdot \alpha_j & \text{if } \alpha_j < 0 \end{cases}, \quad (8.12)$$

where $I_{SR,i}^p(\text{nom})$ and $I_{SR,i}^p(\pm 1\sigma)$ are the nominal and varied expected yields for process p in bin i of the SR, derived from the nominal and varied PNN templates respectively. For each NP α_j , a new Gaussian constraint $G(a = 0|\alpha_j, 1)$ is added to the likelihood. The parameter a is a global observable with a nominal value of zero, thus penalising large deviations of α_j from zero. The combined Gaussian constraints for all NPs α_j form the global term $G(\alpha)$ in Equation 8.6.

As mentioned above, for each (m_X, m_S) point in the search space, a simultaneous binned maximum likelihood fit to the data in the SR and SB is performed. This fit estimates the signal strength parameter μ , the normalisation of the continuum diphoton background $\mu_{\gamma\gamma}$, along with all the NPs α and γ . In total, the fit constrains over 100 parameters.

One of my contributions to the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ search consisted in the implementation of one of the statistical frameworks adopted by the analysis, based on the TRExFitter package [236, 237].

8.6.1 Test statistic for discovery

The measurement of the signal strength μ is based on the profile likelihood ratio $\Lambda(\mu)$, defined in Equation 7.20. Specifically, to test for the presence of a $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal at specific resonance masses (m_X, m_S) , the level of disagreement between the observed data and the background-only hypothesis is quantified using the test statistic q_0 [70], given by:

$$q_0 = \begin{cases} -2 \ln \Lambda(0) & \text{if } \hat{\mu} \geq 0 \\ 0 & \text{if } \hat{\mu} < 0 \end{cases} \quad (8.13)$$

The test statistic assumes that a signal would manifest as a positive value of $\hat{\mu}$, indicating an excess of events over the background-only scenario ($\mu = 0$). A q_0 value near zero corresponds to a profile likelihood ratio $\Lambda(0)$ close to 1, suggesting that the background-only hypothesis closely matches the best-fit values of the signal strength and NPs $(\hat{\mu}, \hat{\theta})$ that jointly maximize the likelihood. In contrast, larger q_0 values indicate increasing disagreement between observed data and the background-only hypothesis, potentially hinting at the presence of a signal.

The confidence level associated with the presence of a signal is quantified by the p -value, defined as:

$$p_0 = \int_{q_{0,\text{obs}}}^{\infty} f(q_0|0) dq_0, \quad (8.14)$$

where $q_{0,\text{obs}}$ is the value of the test statistic in the observed data, and $f(q_0|0)$ represents the probability density function of q_0 under the background-only assumption. A smaller p_0 indicates a larger deviation from the background-only hypothesis, providing stronger confidence in the presence of a signal. To claim a discovery, the p_0 must reach a threshold of 2.87×10^{-7} [70].

The p_0 is also expressed in terms of the *discovery significance* Z , which corresponds to the number of standard deviations above the mean of a normal Gaussian distribution with the same upper-tail probability as p_0 . The relationship between p_0 and Z is defined as:

$$Z = \Phi^{-1}(1 - p_0), \quad (8.15)$$




Figure 8.15 – The standard Gaussian distribution $\phi(x) = (1/\sqrt{2\pi}) \cdot \exp(-x^2/2)$, showing the relation between the significance Z and the p -value, namely $p\text{-value} = \int_Z^\infty \phi(x) dx$.

where Φ^{-1} is the quantile (inverse of the cumulative distribution) of the standard Gaussian distribution. Figure 8.15 illustrates the relationship between Z and the p -value. The discovery threshold of $p_0 = 2.87 \times 10^{-7}$ translates into a discovery significance equal to 5.0σ .

To compute p_0 , the distribution of the test statistic q_0 under the background-only hypothesis, $f(q_0|0)$, is required, as shown in Equation 8.14. This distribution is typically evaluated using asymptotic formulae derived under the asymptotic approximation [70], which is expected to hold unless the data contain very few events. The validity of this approximation was confirmed by generating $f(q_0|0)$ through pseudo-experiments and comparing the expected analysis sensitivity to results derived using the asymptotic formulae. The approximation was found to be robust, with the analysis sensitivity varying within relative factors of 10% across the entire tested (m_X, m_S) plane.

In absence of a significant excess of events observed with respect to the background-only scenario, 95% CL upper limits are set on the POI $\mu = \sigma(pp \rightarrow X) \times \mathcal{BR}(X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma)$, using the test statistic $\tilde{q}_\mu$ and the CL_s method [228], described in Section 7.6.

8.7 Results

For each (m_X, m_S) mass point in the search space, a fit to the observed data is performed simultaneously in the SR and the SB, using the binned maximum likelihood method described in Section 8.6. The post-fit PNN distributions in the SR for two mass points $(m_X, m_S) = (250, 100)$ GeV and $(1000, 70)$ GeV are shown in Figure 8.16. These two mass points are analysed using the 2 b -tagged and 1 b -tagged category selection respectively.




Figure 8.16 – Post-fit distributions of the PNN discriminant output in the (a) 2 b -tagged signal region for $m_X = 250$ GeV and $m_S = 100$ GeV and (b) 1 b -tagged signal region for $m_X = 1000$ GeV and $m_S = 70$ GeV, after a background-only fit to data. The signals corresponding to the two PNN parameterisations, normalised to a 1 fb cross section times branching fraction, are illustrated for comparison (blue solid line). The error band corresponds to the total systematic uncertainty after fit.

For most (m_X, m_S) points, the fit results show a good agreement between the observed data and the SM background-only expectations. However, some deviations are observed at a few points. The most significant excess occurs at $(m_X, m_S) = (575, 200)$ GeV, where the local discovery significance reaches 3.5σ . Figure 8.17 shows the signal-plus-background fit to the PNN distribution in the SR at the mass point $(m_X, m_S) = (575, 200)$ GeV. The contribution from the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal at these resonance masses is depicted as a shaded gray histogram.

However, after accounting for the *look-elsewhere effect*, which arises from searching for a $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal across multiple resonance mass values, the global significance of the search over the entire (m_X, m_S) plane is reduced to 2.0σ . The procedure for evaluating the look-elsewhere effect is detailed in Section 8.8. The observed discovery significance across the tested region of the (m_X, m_S) plane is visualized in Figure 8.18, where the excess is indicated by a bright yellow spot centered at $(m_X, m_S) = (575, 200)$ GeV. Higher values of the discovery significance are also seen for neighboring mass points, reflecting the sensitivity of the PNN evaluated at a determined $\theta = (m_X, m_S)$ to a signal with different nearby masses $\theta' = (m'_X, m'_S)$.

A particularly interesting point is $(m_X, m_S) = (650, 90)$ GeV, where the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ search conducted by the CMS collaboration using Run 2 data [238] reported a deviation between observed data and the background-only expectation. The CMS search quoted a measured signal cross section times branching fraction of 0.35 fb for $(m_X, m_S) = (650, 90)$ GeV, with a local (global) significance of 3.8σ (2.8σ). The analysis presented in this Chapter demonstrated sensitivity to such a signal. In fact, a signal injection test with a 0.35 fb $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal at the resonance masses $(m_X, m_S) = (650, 90)$ GeV was




Figure 8.17 – Post-fit PNN distribution in the SR for the $(m_X, m_S) = (575, 200)$ GeV mass point, after a signal plus background fit. The observed data are shown as black points, while the signal contribution is indicated by the shaded gray histogram. The bottom panel presents the ratio of the observed data with respect to the expected background: a small excess is visible in the most signal-like region in the PNN spectrum.

found to produce an expected local significance of 2.7σ . However, the observed local significance at this mass point is only 1.2σ , and thus this analysis does not confirm the excess reported by the CMS search. In fact, this analysis sets a 95% CL observed upper limit on $\sigma(X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma)$ of 0.14 fb at $(m_X, m_S) = (650, 90)$ GeV, which would exclude the cross section of 0.35 fb measured by the CMS analysis at this mass point.

In the absence of significant excesses above the SM background expectations, 95% CL upper limits are set on the cross section times branching fraction for the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal throughout the (m_X, m_S) plane. Figure 8.19a shows the observed 95% CL upper limits on $\sigma(pp \rightarrow X) \times \mathcal{BR}(X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma)$, while the corresponding expected results, assuming no signal, are displayed in Figure 8.19b.

The observed (expected) limits range from 39 (25) fb at $m_X = 170$ GeV and $m_S = 30$ GeV, to 0.09 (0.14) fb at $m_X = 1000$ GeV and m_S between 250 and 300 GeV. The limits improve at higher masses, reflecting the fact that signals with larger m_X are easier to distinguish from SM processes, as the X and S resonances typically show better experimental resolution. On the other hand, the limits weaken at lower m_S , where the signal becomes highly boosted and thus more challenging to reconstruct, and at low m_X , where signal efficiency in the 2 b -tagged category is lower, due to the subleading b -jet having softer p_T spectra, as explained in Section 8.2.




Figure 8.18 – Local observed significance for signal discovery across the (m_X, m_S) search space. The points show where the significance was evaluated. The band at $m_S = 125$ GeV corresponds to a $X \rightarrow HH \rightarrow b\bar{b}\gamma\gamma$ signal, which is already probed by a previous search performed by the ATLAS Collaboration with Run 2 data [12], and is thus not shown here.




Figure 8.19 – Observed (left) 95% CL upper limits on the signal cross section times branching fraction for the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal, in the (m_X, m_S) plane. The corresponding expected results (right) are evaluated under the assumption of no signal. The points show where the limits were evaluated. The band at $m_S = 125$ GeV is not shown as those points are equivalent to those already probed in the previous $X \rightarrow HH \rightarrow b\bar{b}\gamma\gamma$ search performed with ATLAS Run 2 data [12].

8.8 Look-elsewhere effect for the global significance

As discussed in Section 8.7, the search for the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ signal, as a function of the resonance masses X and S , revealed an excess over background-only expectations near the point $(m_X, m_S) = (575, 200)$ GeV, with a peak local significance of 3.5σ . However, when probing for new physics across a search space spanned by some search parameters $\boldsymbol{\theta}$, quantifying the local p_0 for an excess at a specific local point $\bar{\boldsymbol{\theta}}$ is not sufficient to establish confidence in the actual presence of a signal.

In fact, a seemingly significant local excess may arise by chance, purely due to the extensive size of the explored search space. This is known as the *look-elsewhere effect* [239]. To account for this, the *global* p_0 (or *global significance*) must be evaluated across the entire parameter space.

The global p_0 represents the probability of observing a local deviation from the background expectation anywhere within the search space, assuming the background-only hypothesis holds true:

$$\text{Global } p_0 = \mathcal{P}(\max_{\boldsymbol{\theta} \in \mathcal{M}} q_0(\boldsymbol{\theta}) > q_{\text{obs}} | \mu = 0) \quad (8.16)$$

In Equation 8.16, $\boldsymbol{\theta} = (\theta_1, \theta_2, \dots, \theta_m)$ represents the search parameters (for the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ analysis, $\boldsymbol{\theta} = (m_X, m_S)$), and $\mathcal{M}$ denotes the full explored search space. The test statistic $q_0(\boldsymbol{\theta})$, defined in Section 8.6, depends on the search parameters, and its observed maximum across the search space is denoted by q_{obs} . Thus, the global p_0 quantifies the probability that, in the absence of a genuine signal but purely due to statistical fluctuations in the background, the observed data would exhibit deviations from the background-only expectations at any point $\boldsymbol{\theta}$ in the search space, with a maximum test statistic q_0 larger than the value q_{obs} actually observed in data.

As shown in Equation 8.16, calculating the global p_0 involves estimating the probability density function for the quantity $\max_{\boldsymbol{\theta} \in \mathcal{M}} q_0(\boldsymbol{\theta})$. This distribution could be sampled by generating a large number of pseudo-experiments (*toys*) under the background-only hypothesis, running the full analysis for each one, and determining the maximum test statistic for each toy. However, this method is computationally expensive, so an alternative approach was chosen. The global p_0 for the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ search accounts for the look-elsewhere effect using an asymptotic method [240], summarised as follows:

- The set of search parameters $\boldsymbol{\theta}$ where the test statistic $q_0(\boldsymbol{\theta})$ exceeds a certain threshold u is termed *excursion set*:

$$\mathcal{A}_u = \{\boldsymbol{\theta} \in \mathcal{M} | q_0(\boldsymbol{\theta}) > u\} \quad (8.17)$$

- Under the background-only hypothesis, the expectation value of the *Euler characteristic* of the excursion set, $\phi(\mathcal{A}_u)$ [241], asymptotically approaches the global p_0 for large enough thresholds u :

$$\mathbb{E}[\phi(\mathcal{A}_u)] \approx \mathcal{P}(\max_{\boldsymbol{\theta} \in \mathcal{M}} q_0(\boldsymbol{\theta}) > u | \mu = 0) \quad (8.18)$$

Therefore, evaluating the expectation value of the Euler characteristic of the excursion set for a threshold $u = q_{\text{obs}}$ provides an estimate of the global p_0 , as defined in Equation 8.16.

- In addition, for a two-dimensional search space, the expected Euler characteristic (referred in the following as simply $\phi(\mathcal{A}_u)$), under the background-only hypothesis, can be expressed as:

$$\phi(\mathcal{A}_u) = \mathcal{P}(\chi^2 > u) + e^{-u/2} \cdot (\mathcal{N}_1 + \sqrt{u} \cdot \mathcal{N}_2), \quad (8.19)$$

where $\mathcal{P}(\chi^2 > u)$ represent the cumulative probability of the χ^2 distribution with one degrees of freedom to exceed a threshold u , and $\mathcal{N}_1$ and $\mathcal{N}_2$ are empirical constants, valid for each value of u . Equation 8.19 is crucial for determining the global p_0 in the asymptotic approximation. One could, in principle, compute the expected Euler characteristic $\phi(\mathcal{A}_u)$ for convenient low enough thresholds u under the background-only hypothesis and use these values to estimate $\mathcal{N}_1$ and $\mathcal{N}_2$. Once these constants are determined, Equation 8.19 combined with Equation 8.18 allows for the calculation of the global p_0 as a function of the maximum test statistic u . Finally, evaluating Equation 8.19 for $u = q_{\text{obs}}$ yields the global p_0 for the observed data.

To reduce the impact of statistical uncertainty in determining the values of $\mathcal{N}_1$ and $\mathcal{N}_2$, it is helpful to generate a limited set of toys under the background-only hypothesis for calculating the Euler characteristic $\phi(\mathcal{A}_u)$ at selected thresholds u . For a given threshold, $\phi(\mathcal{A}_u)$ can be computed for each toy, and the values of $\mathcal{N}_1$ and $\mathcal{N}_2$ are extracted via a simultaneous fit to the $\phi(\mathcal{A}_u)$ values across all toys and thresholds. However, generating background-only toys for the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ search is non-trivial because both the event selection and the PNN discriminant vary with the resonance masses (m_X, m_S) within the search space. In other words, the background-only model depends on the search parameters (m_X, m_S), making it impossible to generate a single background-only toy from sampling the PNN templates for all mass points simultaneously, since the PNN discriminant itself is a function of the search parameters $\boldsymbol{\theta} = (m_X, m_S)$. To address this issue, an alternative procedure was developed, in order to generate toys directly from MC simulated events for the background processes. After generating the toys, the full $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ analysis is applied to each toy to produce two-dimensional maps of the test statistic q_0 as a function of the search parameters $\boldsymbol{\theta} = (m_X, m_S)$. These maps are then used to compute the Euler characteristics $\phi(\mathcal{A}_u)$, which is done by counting, for each toy, how many values of q_0 across the search space exceed the thresholds u . Once this is done, the values of $\mathcal{N}_1$ and $\mathcal{N}_2$ are determined, thus allowing the calculation of the global p_0 .

8.8.1 Generation of toy data

A total of 20 background-only toys were generated, starting from MC simulated events for the background processes contributing to the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ search. The goal of toy generation is to allow the full $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ analysis to be performed on each

toy. This is challenging because, as previously mentioned, different (m_X, m_S) points are analysed with different selection criteria and different final discriminants - specifically, the 1 b -tagged category applies to points with $m_X/m_S < 0.09$, while the 2 b -tagged category applies to points with $m_X/m_S > 0.09$, and the PNN discriminant is different for each (m_X, m_S) .

To overcome this, the toy generation starts with MC events passing a common preselection that applies to all (m_X, m_S) points. This preselection is formed by merging the requirements for the 1 b -tagged and 2 b -tagged categories, demanding at least one and at most two b -tagged jets identified using the 77% efficiency working point of the DL1r algorithm 6.2. For each background process b , the expected number of events, N_b , was estimated at the preselection level using the corresponding MC sample. The core idea behind toy generation is to combine randomly selected MC events from each background sample b , ensuring that the number of picked events follows a Poisson distribution centered around N_b . This method allows the contribution of each background process to fluctuate around N_b , reflecting the statistical uncertainty in the expected event counts from data.

Additional complexity arises from the fact that MC events are weighted⁴, and these weights may be negative. To handle this, the MC sample for each background process b is divided into two subsamples, $\text{Sample}(b, +)$ and $\text{Sample}(b, -)$, containing events with positive and negative weights, respectively. Two positive quantities, N_b^+ and N_b^- , are defined to satisfy the following conditions:

$$\begin{cases} N_b^+ - N_b^- = N_b \\ N_b^- = f_b \cdot (N_b^+ + N_b^-) \end{cases} \quad , \quad (8.20)$$

where f_b is the fraction of negatively weighted events in the sample b at the preselection level. For each toy i , the number of events picked from the positive and negative weight subsamples, $n_{b,i}^+$ and $n_{b,i}^-$, are sampled from two independent Poisson distributions centered around N_b^+ and N_b^- , respectively:

$$\begin{cases} n_{b,i}^+ \sim \text{Pois}(n|N_b^+) \\ n_{b,i}^- \sim \text{Pois}(n|N_b^-) \end{cases} \quad (8.21)$$

The events selected from $\text{Sample}(b, +)$ contribute to the toy with a weight of +1, while events from $\text{Sample}(b, -)$ contribute with a weight of -1. The selection of events from the subsamples is done randomly, with each event's probability of being picked proportional to its absolute MC weight. This approach ensures that the physical meaning of MC weights is respected in the generation of background-only toy datasets.

As an example, consider the ggF HH background. At the common preselection level, the expected event yield for the ggF HH process is $N_{\text{ggF HH}} = 3.37$ events. Additionally,

⁴Each event from the MC simulations is assigned a weight that reflects the underlying theoretical calculations involved in event generation and parton showering. These weights are essential to ensure that the distribution of physical observables derived from the simulated events have physical meaning. Furthermore, additional weights are applied to correct for discrepancies between the simulated detector response and actual data, in order to improve the modelling of the considered process.

5.8% of the simulated events in the ggF HH sample at this stage carry a negative weight: therefore, $f_{\text{ggF HH}} = 0.058$. Solving Equation 8.20 yields:

$$\begin{cases} N_{\text{ggF HH}}^+ = 3.59 \\ N_{\text{ggF HH}}^- = 0.22 \end{cases} \quad (8.22)$$

To account for the ggF HH background in the 20 generated toys, 20 integers are sampled from the two probability distributions $\text{Pois}(n|N_{\text{ggF HH}}^+ = 3.59)$ and $\text{Pois}(n|N_{\text{ggF HH}}^- = 0.22)$. The distributions of these integers, denoted as $n_{\text{ggF HH},i}^{+(-)}$, are shown in Figures 8.20a and 8.20b, respectively. The corresponding numbers of ggF HH simulated events with positive and negative weights are then randomly drawn from their respective subsamples.




Figure 8.20 – The distribution of the number of ggF HH simulated events selected for each toy, drawn from the subsample with (a) positive weights and (b) negative weights, is shown. The vertical lines represent the mean values of the Poisson distributions used for sampling the number of events, which are $N_{\text{ggF HH}}^+ = 3.59$ and $N_{\text{ggF HH}}^- = 0.22$, respectively.

The aim of this approach is to ensure that each toy i contains a number of simulated events, weighted either +1 or -1, from each background process b , with an average event count consistent with the expected value N_b . To verify this, the difference $n_{b,i}^+ - n_{b,i}^-$ was calculated for each toy i and averaged across all toys. As shown in Table 8.2, the average value $\langle n_{b,i}^+ - n_{b,i}^- \rangle$ is consistent with N_b for each background process b , confirming the robustness of this method. Before selecting events from the positive and negative weight subsamples, additional weights are applied to the $\gamma\gamma$ +jets simulated events, in order to match the overall normalisation of the continuum background to the data in the $m_{\gamma\gamma}$ SB region at the common preselection level and after applying the 1 b -tagged and 2 b -tagged category selections. Thus, the normalisation of the background-only toys is expected to align with the data in the SB region (where no signal is expected) for each analysis category.

Table 8.2 – Average difference $\langle n_{b,i}^+ - n_{b,i}^- \rangle$ where $n_{b,i}^{+(-)}$ represents the number of simulated events selected from the positive (negative) weight subsample for background process b in toy i . This average is compared with the expected yield, N_b , at the common preselection level. The values are computed across the 20 background-only toys, with the uncertainty given as the standard deviation of the mean.

Background process	$\langle n_{b,i}^+ - n_{b,i}^- \rangle$	Expected yield N_b
ggF HH	3.75 ± 0.37	3.37
VBF HH	0.20 ± 0.09	0.15
ggF H	52.65 ± 1.39	54.00
VBF H	8.70 ± 0.62	9.02
WH	6.10 ± 0.54	6.18
qq $\rightarrow$ ZH	8.50 ± 0.69	8.76
gg $\rightarrow$ ZH	2.55 ± 0.25	2.45
t $\bar{t}$ H	19.00 ± 1.00	20.00
tHq	3.80 ± 0.73	3.65
tHW	0.55 ± 0.32	0.70
b $\bar{b}$ H	3.25 ± 0.50	3.59
$\gamma\gamma$ +jets	18482.05 ± 25.66	18525.05
t $\bar{t}\gamma\gamma$	79.60 ± 2.04	78.48
Z + $\gamma\gamma$	80.65 ± 1.73	74.07

Using this procedure, 20 background-only toys were generated by combining simulated events drawn from each background sample b . The $m_{\gamma\gamma}$ distributions for the expected background contribution, evaluated from MC samples, for each toy, and for data in the $m_{\gamma\gamma}$ SB region are displayed in Figure 8.21. These distributions are shown for events passing the common preselection, as well as the 1 b -tagged and 2 b -tagged categories, respectively. Figure 8.21 also presents a comparison of the $m_{\gamma\gamma}$ distributions between the average of the 20 toys, the expected background, and the data.

8.8.2 Statistical results for background only toys

For each toy, the PNN discriminant was evaluated at each (m_X, m_S) point in the signal grid. The PNN distributions in the SR for the 20 toys, compared to the expected background templates, are shown in Figure 8.22 for the mass points $(m_X, m_S) = (250, 100)$ GeV and $(575, 200)$ GeV.

The full $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ analysis was repeated 20 times, once for each toy, resulting in 20 two-dimensional maps of the test statistic q_0 as a function of the resonance masses (m_X, m_S) . Figure 8.23 presents the distribution of the best-fit signal strength $\hat{\mu}$ for the 20 generated toys, compared to the result from an Asimov dataset generated under the background-only assumption, and to the fit result observed in data. This comparison is performed for the mass points $(m_X, m_S) = (250, 100)$ GeV (Figure 8.23a) and $(575, 200)$ GeV (Figure 8.23b). The test statistic q_0 , calculated for each toy, the background-only Asimov dataset, and for data, is shown in Figure 8.24. These are compared against the expected probability density function under the background-only hypothesis, $f(q_0|\mu = 0)$,




Figure 8.21 – Distributions of $m_{\gamma\gamma}$ for the expected background (solid black line, estimated from MC samples) and data (black markers), compared with those extracted from individual background-only toys (left) and the toy average (right). The error band on the toy average represents the standard deviation of the mean. These distributions are shown for events passing (a,b) the common preselection, (c,d) the 1 b -tagged category, and (e,f) the 2 b -tagged category.




Figure 8.22 – Comparison of the PNN output distributions for the total expected background (solid black line, estimated using MC samples) and the 20 background-only toys (faint blue lines) in the $m_{\gamma\gamma}$ SR. The PNN templates are shown for two mass points: $(m_X, m_S) = (250, 100)$ GeV (left) and $(m_X, m_S) = (575, 200)$ GeV (right).

for the same mass points.

For the mass point $(m_X, m_S) = (575, 200)$ GeV, which shows a local excess, the ob-




Figure 8.23 – Best-fit values of the signal strength $\hat{\mu}$ for the background-only toys (blue vertical lines and filled blue histogram), compared with the corresponding result for data (black marker) and the background-only Asimov dataset (solid red line, representing a Gaussian probability density, with mean and standard deviation equal to the best-fit value $\hat{\mu}$ and its uncertainty under the background-only hypothesis). The exact best-fit values in each case are quoted in the legend. Results are shown for two mass points: $(m_X, m_S) = (250, 100)$ GeV (left) and $(m_X, m_S) = (575, 200)$ GeV (right).




Figure 8.24 – Distribution of the test statistic q_0 for the background-only toys (blue vertical lines and filled blue histogram), compared with the result from data (solid black line) and the background-only Asimov dataset (dashed blue line). The integral over each bin of the expected distribution of the test statistic q_0 under the background-only hypothesis, $f(q_0|0)$, is shown as a solid red line. Results are provided for two signal points: $(m_X, m_S) = (250, 100)$ GeV (left) and $(m_X, m_S) = (575, 200)$ GeV (right).

served values of $\hat{\mu}$ and $q_{0,obs}$ exhibit a deviation from the distributions generated by the background-only toys. In contrast, at the mass point $(m_X, m_S) = (250, 100)$ GeV, where no local excess is observed, the observed $\hat{\mu}$ and $q_{0,obs}$ are consistent with the background-only toy distributions.

Finally, Figure 8.25 provides an example of a two-dimensional map of the test statistic q_0 as a function of (m_X, m_S) , generated for one of the background-only toys.

8.8.3 Global significance

Determination of the Euler characteristics

The two-dimensional maps of the test statistic $q_0(m_X, m_S)$ for each toy are used to evaluate the Euler characteristic, $\phi(\mathcal{A}_u)$, at four different thresholds: $u = 1.0, 2.0, 4.0$, and 6.0 . The Euler characteristic $\phi(\mathcal{A}_u)$ quantifies the number of connected regions in the (m_X, m_S) plane where the test statistic exceeds the chosen threshold u . Intuitively, $\phi(\mathcal{A}_u)$ reflects the number of “islands”, formed by nearest neighboring (m_X, m_S) points where $q_0(m_X, m_S)$ is larger than u . Figure 8.26 illustrates this for thresholds $u = 1.0$ and $u = 2.0$, displaying the $q_0(m_X, m_S)$ maps for three background-only toys. The mass points where $q_0(m_X, m_S) > u$ are highlighted in red: all the nearest neighbours (m_X, m_S) points with $q_0(m_X, m_S) > u$ form an island, and the number of islands corresponds to the Euler characteristic $\phi(\mathcal{A}_u)$.




Figure 8.25 – Two-dimensional map of the test statistic q_0 as a function of the (m_X, m_S) mass points for one background-only toy.




Figure 8.26 – Two dimensional maps of the test statistic q_0 in the (m_X, m_S) plane, for three toys. The red markers indicate the (m_X, m_S) mass points for which $q_0(m_X, m_S) > \bar{u}$, with $\bar{u} = 1.0$ (top) or $\bar{u} = 2.0$ (bottom). The Euler characteristics $\phi(\mathcal{A}_u)$ are quoted for each case.

Table 8.3 – Values of the empirical constants $\mathcal{N}_1$ and $\mathcal{N}_2$, extracted from a simultaneous fit to the Euler characteristics $\phi(\mathcal{A}_u)$ across the 20 background-only toys for four thresholds u , equal to 1.0, 2.0, 4.0 and 6.0.

$\mathcal{N}_1$	25.7 ± 1.1
$\mathcal{N}_2$	$0.0 \pm 8.3 \cdot 10^{-12}$

Determination of the $\mathcal{N}_1$ and $\mathcal{N}_2$ constants

A simultaneous fit to the distributions of the Euler characteristics $\phi(\mathcal{A}_u)$ from each toy at each threshold is employed to extract the two empirical constants, $\mathcal{N}_1$ and $\mathcal{N}_2$. For a given threshold $\bar{u}$, the Euler characteristics $\phi(\mathcal{A}_{\bar{u}})_i$ for each toy i are modeled using a Poisson function, with the mean value $\lambda_{\bar{u}}$ defined by Equation 8.19:

$$\lambda_{\bar{u}} = \mathcal{P}(\chi^2 > \bar{u}) + e^{-\bar{u}/2} \cdot (\mathcal{N}_1 + \sqrt{\bar{u}} \cdot \mathcal{N}_2) \quad (8.23)$$

The simultaneous fit of the Poisson functions to the toy data allows to determine of the values of $\mathcal{N}_1$ and $\mathcal{N}_2$, along with their uncertainties, which account for both the limited number of toys and the correlations among the Euler characteristics evaluated at different thresholds. The value of $\mathcal{N}_2$ was constrained to be non-negative. The fitted values of $\mathcal{N}_1$ and $\mathcal{N}_2$, along with their associated uncertainties, are summarised in Table 8.3. The values for $\mathcal{N}_1$ and $\mathcal{N}_2$ can be substituted into Equation 8.19 to calculate the expected Euler characteristic, $\mathbb{E}[\phi(\mathcal{A}_u)]$, as a function of the threshold u (interpreted as the maximum test statistic across the explored (m_X, m_S) plane). Figure 8.27, showing the fitted $\mathbb{E}[\phi(\mathcal{A}_u)]$ as a function of u , reports an excellent agreement between the fitted function and the four points obtained by averaging the values of $\phi(\mathcal{A}_u)_i$ across the 20 toys for each of the four selected thresholds u .

Determination of the global significance

The curve for the expected Euler characteristic, $\mathbb{E}[\phi(\mathcal{A}_u)]$, as a function of the maximum test statistic u , can be translated into a curve showing the global significance as a function of the maximum local significance. This is done by first using Equation 8.18 to convert the Euler characteristic into a global p_0 value, then applying Equation 8.14 to convert the test statistic u into a local p_0 value, and finally using Equation 8.15 to transform both the global and local p_0 values into significance levels. Figure 8.28 presents the global significance as a function of the maximum local significance. The global significance remains below the dashed line, which represents the global significance being equal to the maximum local significance, neglecting the look-elsewhere effect. Given the observed maximum local significance of 3.5σ at $(m_X, m_S) = (575, 200)$ GeV in the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ analysis, the corresponding global significance is $(1.992 \pm 0.021)\sigma$ (rounded to 2.0σ). The uncertainty arises from the fit uncertainties in the constants $\mathcal{N}_1$ and $\mathcal{N}_2$. The associated local and global p_0 values, along with the trial factor (the ratio of global to local p_0), are summarised in Table 8.4.




Figure 8.27 – Fitted curve of $\mathbb{E} [\phi(\mathcal{A}_u)]$ as a function of the threshold in test statistic q_0 . The uncertainty band is propagated from the fitted uncertainties on $\mathcal{N}_1$ and $\mathcal{N}_2$. The blue markers indicate the arithmetic averages of the $\phi(\mathcal{A}_u)$ values for the thresholds $u = 1.0, 2.0, 4.0$, and 6.0 , calculated across the 20 toys. The error bars on the markers correspond to the standard deviation of the mean.

Table 8.4 – Maximum of the local p_0 , global p_0 , and corresponding trial factor. The maximum of the local p_0 across the explored search space is found for $(m_X, m_S) = (575, 200)$ GeV.

Maximum local p_0	$2.33 \cdot 10^{-4}$
Global p_0	0.0232
Trial factor	120.17




Figure 8.28 – Global significance as a function of the maximum local significance for the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ analysis, shown with (solid blue line) and without (dashed green line) accounting for the look-elsewhere effect. The uncertainty band on the global significance is propagated from the simultaneous fit used to determine the constants $\mathcal{N}_1$ and $\mathcal{N}_2$. The white star marks the maximum local significance observed in the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ search, along with its corresponding global significance. The bottom panel presents the ratio of local to global significance.

Conclusions

This thesis presented two different analyses based on the proton-proton collision data collected by the ATLAS detector at $\sqrt{s} = 13$ TeV, corresponding to an integrated luminosity of 140 fb^{-1} . The common theme of these analyses is the search for Higgs boson pair production in the final state with two photons and two bottom quarks ($b\bar{b}\gamma\gamma$). The first analysis is the search for non-resonant production of two SM Higgs bosons (HH), while the second one is the search for asymmetric production of two different scalars (one SM Higgs boson and a new boson S), from the decay of a heavy scalar resonance X ($X \rightarrow SH$).

The non-resonant search for Higgs boson pair production did not show any significant excess over the background-only expectations. Hence, the analysis places a 95% confidence level (CL) upper limit on the di-Higgs signal strength, μ_{HH} , to be below 4.0 times the SM prediction. The corresponding expected limits, evaluated assuming either no HH production, or SM HH production, are 5.0 and 6.4, respectively. The analysis also sets limits on the two dominant HH production modes (via gluon-gluon fusion and vector boson fusion) separately. The observed (expected) 95% CL upper limit on the ggF HH signal strength is 4.1 (5.3). Similarly, the observed 95% CL upper limit on the VBF HH signal strength is 96, with a corresponding expected result of 145. The $HH \rightarrow b\bar{b}\gamma\gamma$ search also explores the trilinear Higgs boson self-coupling, placing 95% CL constraints on its coupling modifier κ_λ . The observed allowed range is $-1.4 < \kappa_\lambda < 6.9$, while the expected range, evaluated assuming SM HH production (namely, $\kappa_\lambda = 1$) is $-2.8 < \kappa_\lambda < 7.8$. These corresponds to the most stringent expected constraints on κ_λ placed by a single analysis. Thanks to considering the subdominant production mode, via VBF, in the analysis optimisation, the $HH \rightarrow b\bar{b}\gamma\gamma$ search is able to achieve competitive limits also to the strength of the interaction between two Higgs bosons and two vector bosons, parametrised using the coupling modifier κ_{2V} . The observed (expected) 95% CL constraints on κ_{2V} are $-0.5 < \kappa_{2V} < 2.7$ ($-1.1 < \kappa_{2V} < 3.3$). Finally, for the first time in the $b\bar{b}\gamma\gamma$ channel, this analysis also provides 95% CL two-dimensional observed and expected exclusion regions in the $(\kappa_\lambda, \kappa_{2V})$ parameter space.

The final and most complete picture on Higgs boson pair production obtained using ATLAS Run 2 data is achieved by combining all the HH searches, performed using various final states, including, alongside $b\bar{b}\gamma\gamma$, also $b\bar{b}\tau\tau$, $b\bar{b}b\bar{b}$, multi-leptons, and $b\bar{b}\ell\ell + \text{MET}$. Among these channels, $b\bar{b}\gamma\gamma$ stands out due to its best sensitivity to κ_λ . The observed combined 95% CL upper limit on the HH signal strength is 2.9 times the SM

prediction, with an expected limit of 2.4, calculated in the background-model assumption. The expected 95% CL upper limit on μ_{HH} , assuming SM HH production, is 3.4. The combined HH searches are able to set 95% CL constraints on κ_λ and κ_{2V} , to be $-1.2 < \kappa_\lambda < 7.2$ and $0.6 < \kappa_{2V} < 1.5$, respectively. The corresponding expected results, evaluated in the SM assumption, are $-1.6 < \kappa_\lambda < 7.2$ and $0.4 < \kappa_{2V} < 1.6$, thus yielding the tightest expected constraints on κ_λ to date.

The second analysis presented in this thesis, the search for $X \rightarrow SH$ production, is the first search for this process in the $b\bar{b}\gamma\gamma$ channel. The search is performed in a wide region in the parameter space spanned by the masses of the two resonances X and S , reaching m_X and m_S values as low as 170 GeV and 15 GeV, thus covering a previously unexplored regions of the parameter space. The signal plus background fit to data showed a good agreement with the background-only hypothesis, in most (m_X, m_S) points. The largest discrepancy with the SM assumption was found for resonance masses equal to $(m_X, m_S) = (575, 200)$ GeV, corresponding to a local (global) discovery significance equal to 3.5σ (2.0σ). Since no statistically significant excess was observed with respect, the analysis sets 95% CL upper limits on the product of the production cross section and the branching fraction of the $X \rightarrow SH \rightarrow b\bar{b}\gamma\gamma$ process, as a function of the resonance masses. The observed limits range from 39 fb at lower m_X values to 0.09 fb at higher m_X values.

The latest combined HH searches using the full Run 2 dataset have achieved an improvement of a factor four in sensitivity to SM Higgs pair production compared to the results obtained from the earlier partial Run 2 dataset (36.1 fb^{-1}), which exceeds of a factor two the improvement expected from the increase in the integrated luminosity alone. With the onset of Run 3 in 2022, which has already provided over 170 fb^{-1} of data and is expected to continue until 2026, there is a unique opportunity to push HH searches to new levels of sensitivity. This is being driven by improved reconstruction and identification techniques for key physics objects, innovative optimization strategies, and early efforts to control systematic uncertainties before they become a limiting factor, with the final goal to achieve the first hint of HH production with the combined data from both the ATLAS and CMS experiments by the end of Run 3. Following the conclusion of Run 3, the High-Luminosity LHC (HL-LHC) phase will begin, aiming to collect 3000 fb^{-1} of data by 2041 - 20 times the integrated luminosity achieved with Run 2. The definitive answer on Higgs boson pair production and the trilinear Higgs boson self-coupling is expected to emerge only at the end of the HL-LHC era, thanks to this unprecedented amount of data.

Appendices

Modeling the single Higgs boson production plus bottom quarks

One of the main irreducible backgrounds for Higgs boson pair production searches involving at least one $H \rightarrow b\bar{b}$ decay is the production of a single Higgs boson accompanied by additional bottom quarks. This process is referred to as $b\bar{b}H$. When considering higher-order QCD corrections, the $b\bar{b}H$ process arises from two distinct production mechanisms: gluon-gluon fusion (ggF) and bottom-quark annihilation ($b\bar{b}H$), with the corresponding Feynman diagrams shown in Figures A.1a and A.1b.

The first diagram is proportional to the Yukawa coupling of the Higgs boson to the top quark (y_t) and corresponds to ggF production, where two additional bottom quarks are emitted via QCD radiation. The second diagram, proportional to the Yukawa coupling of the Higgs boson to the bottom quark (y_b), represents the direct $b\bar{b}H$ production mechanism, as described in Section 1.3.1. Thus, the total $b\bar{b}H$ cross section, which comes from the modulus square of the combined amplitude, consists of three terms: one proportional to y_t^2 , another proportional to y_b^2 , and a third term arising from the interference of the two diagrams (Figures A.1a and A.1b), proportional to the product of y_t and y_b :

$$d\sigma(b\bar{b}H) \propto A_t \cdot y_t^2 + A_b \cdot y_b^2 + A_{\text{int}} \cdot y_t y_b \quad (1.1)$$

This background can be significant in the HH phase space. As illustrated in Figure 7.8, the contribution from single Higgs production via ggF is of the same order of magnitude as the SM HH signal, or it can even exceed the SM HH signal, in the signal regions of the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis. This ggF contribution includes not only the irreducible background characterised by the emission of two bottom quarks, but also a reducible background from ggF events accompanied by c - or light-flavor jets that are misidentified as b -jets.

In the HH searches performed by the ATLAS Collaboration, the component of the $b\bar{b}H$ background proportional to y_t^2 is modeled using MC simulated events from the inclusive POWHEG NNLOPS ggF sample [189–191, 204, 205]. This sample includes both the irreducible ggF + $b\bar{b}$ component and the reducible ggF + other jets background. As detailed in Section 7.2.2, this sample generates single Higgs boson events via ggF at NNLO accuracy in QCD: therefore, the ggF plus two additional jets component is simulated only at LO. Additionally, the generation employs the five-flavor scheme, treating the




Figure A.1 – Feynman diagrams contributing to the single Higgs boson production plus two additional bottom quarks. The blue markers represent QCD interaction vertices, whose contribution to the total cross amplitude is proportional to the square root of the strong coupling constant α_s . The red (green) markers represent the Yukawa interaction vertex between the Higgs boson and the top (bottom) quarks, whose contribution to the total amplitude is proportional to y_t (y_b).

bottom quark as a massless parton.

ATLAS HH analyses assign a 100% systematic uncertainty to the $ggF + b\bar{b}$ background due to challenges in accurately simulating heavy-flavor emissions. This is supported by multiple measurements showing discrepancies between data and simulations in phase spaces involving heavy-flavor jets [223–227]. Although most HH analyses are currently limited by the amount of data collected in Run 2, this 100% uncertainty contributes by 25% to the total uncertainty on the combined HH signal strength, μ_{HH} , as explained in Section 7.8.2. This uncertainty is already the dominant systematic affecting the measurement of the HH signal strength. Projections for the High-Luminosity LHC (HL-LHC) scenario indicate that it will significantly degrade both the HH discovery significance and the sensitivity to the trilinear coupling modifier, κ_λ , in HH analyses using the full 3000 fb^{-1} dataset [242, 243]. For this reason, a more precise simulation of the $b\bar{b}H$ background is essential to mitigate this significant source of uncertainty.

A simulation of the $b\bar{b}H$ process considering NLO corrections in QCD is already available, based on MADGRAPH5_AMC@NLO [244]. This simulation is based on the four-flavor scheme. It accounts for both the y_t^2 and y_b^2 components, as well as their interference. The y_t^2 component was found to be dominant with respect to the y_b^2 contribution, even at leading order. The dominance of the y_t^2 component becomes even more pronounced when NLO corrections are included, since the NLO/LO ratio for the y_t^2 contribution is larger than the one for the y_b^2 contribution. Specifically, the NLO/LO ratio for the y_t^2 component is around 2, whereas for the y_b^2 component it ranges between 1 and 1.5. In addition, both NLO/LO ratios exhibit a strong dependence on the differential distributions. Therefore, including NLO accuracy in QCD into MC simulations is crucial for a reliable modelling of the $b\bar{b}H$ process, particularly for the y_t^2 term. Furthermore, the scale uncertainties from missing higher-order QCD corrections are substantial for both components, with uncertainties around 50% on the inclusive cross sections.

The work presented in this Appendix introduces a new $b\bar{b}H$ simulation in the context of HH searches [21]. It builds on the existing NLO simulation of the $b\bar{b}H$ process, and matches it to the parton shower, which is essential for making this simulation usable in physics analyses. A dedicated study was performed, targeting the HH phase space, with the $b\bar{b}\gamma\gamma$ final state used as a benchmark. My contribution includes reproducing the MC simulations of HH and single Higgs production via ggF, which are currently used in ATLAS HH searches, for comparison with the new $b\bar{b}H$ simulation. I also studied the impact of parton showering on the existing POWHEG NNLOPS-based ggF simulation of single Higgs production and assessed how adopting the new $b\bar{b}H$ simulation, instead of the POWHEG NNLOPS sample, for modeling the ggF + $b\bar{b}$ background would affect HH analyses.

A.1 An accurate simulation of the $b\bar{b}H$ process

A NLO simulation of the $b\bar{b}H$ process was implemented using the MADGRAPH5_AMC@NLO generator, with event generation matched to the parton shower. The calculation was performed in the four-flavor scheme, treating the bottom quark as a massive particle. Only the components proportional to y_t^2 and y_b^2 were simulated, while the interference term was neglected¹. The simulated events were analysed at truth-level (without passing the events through the detector simulation), applying specific selection criteria defining a *fiducial region*. The fiducial region targets the HH phase space, using the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis as a representative case.

- The Higgs boson is required to decay into a pair of photons;
- The final-state photons are required to have $|\eta| < 2.37$.
- Jets are reconstructed using the anti- k_t algorithm with a radius parameter of $R = 0.4$, and are required to have $|\eta| < 2.5$ and $p_T > 25$ GeV.
- Events entering the fiducial region are required to contain at least two photons and exactly two b -jets (corresponding to jets containing at least one B -hadron among their constituents);
- The selected leading (subleading) photon is required to have $p_T/m_{\gamma\gamma} > 0.35(0.25)$ and $|\eta| < 2.37$, where $m_{\gamma\gamma}$ is the invariant mass of the two photons;
- The two selected b -jets were required to have an invariant mass in the range of $80 < m_{b\bar{b}} < 140$ GeV.

Table A.1 presents the cross sections for the $b\bar{b}H$ process in the fiducial region at both LO and NLO, with separate contributions from the y_t^2 and y_b^2 components, along with their total sum. The results are provided for parton shower matching with both PYTHIA 8 and HERWIG 7. The $b\bar{b}H$ cross sections are compared with those from the POWHEG NNLOPS ggF sample. For the POWHEG NNLOPS sample, cross sections are given for two configurations of the parton shower: one including the gluon splitting into bottom quarks

¹The results in [244] indicate that the interference contribution is subdominant compared to the y_t^2 and y_b^2 components, accounting for only 5% to 10% of the total $b\bar{b}H$ cross section, and, in particular, it is smaller than the scale uncertainties.

($g \rightarrow b\bar{b}$), and one where this process is turned off. In the same fiducial region, the cross section of the SM HH production in the $b\bar{b}\gamma\gamma$ channel was found to be 22.7 ab.

Table A.1 – Fiducial cross sections for the process $pp \rightarrow b\bar{b}H$ process, assuming the $H \rightarrow \gamma\gamma$ decay. The cross section values are in ab.

		LO	NLO	Scale unc.	Parton shower scale unc.	POWHEG NNLOPS (y_t^2 only at NLO)
y_b^2	PYTHIA 8	3.15	4.22	+15% -15%	+10% -4%	
	HERWIG 7	2.59	4.08		+8% -12%	
y_t^2	PYTHIA 8	8.24	18.1	+58% -34%	+10% -7%	29.9
	HERWIG 7	6.83	16.6		+4% -5%	No $g \rightarrow b\bar{b}$: 17.2
Sum	PYTHIA 8	11.4	22.3	+50% -30%	+10% -6%	
	HERWIG 7	9.42	20.7		+4% -6%	

The y_t^2 component is found to be more than twice larger with respect to the y_b^2 component in the fiducial region. NLO corrections significantly enhance both components, increasing the y_t^2 contribution by approximately 150% and the y_b^2 contribution by about 50%. Scale uncertainties, as shown in the fourth column of Table A.1, are still sizeable, reaching nearly 60% for the y_t^2 component (while they are around 15% for the y_b^2 component). The impact of the parton shower was evaluated by comparing results obtained with PYTHIA 8 and HERWIG 7 for parton shower matching, and by checking the uncertainty associated with the variation of the parton shower scale (presented in the fifth column of Table A.1). These shows a relatively minor impact with respect to scale uncertainties, with values up to 10%. Considering the fiducial $HH \rightarrow b\bar{b}\gamma\gamma$ cross section of 22.7 ab confirms that the $b\bar{b}H$ background is non-negligible within the HH phase space. The fiducial $b\bar{b}H$ cross section obtained from the POWHEG NNLOPS ggF sample (which only models the y_t^2 contribution) is found to be twice as large as the cross section extracted from the new NLO MADGRAPH5_AMC@NLO sample. This discrepancy could be traced back to the $g \rightarrow b\bar{b}$ splitting in the parton shower. When this process is disabled, the cross section from the POWHEG NNLOPS ggF sample is reduced by a factor of two, and becomes closer to the NLO $b\bar{b}H$ predictions.

A.1.1 Comparison with the current simulation

The new $b\bar{b}H$ simulation predicts fiducial cross sections that are half the size of those obtained from the POWHEG NNLOPS sample. However, this discrepancy disappears when the $g \rightarrow b\bar{b}$ splitting is turned off in the parton shower for the POWHEG NNLOPS sample. This suggests that the $g \rightarrow b\bar{b}$ splitting in the parton shower contributes to approximately half of the events in the fiducial region for the POWHEG NNLOPS sample (which are characterised by exactly two central b -jets having $p_T > 25$ GeV). There

are two possible explanations for this:

- The $g \rightarrow b\bar{b}$ splitting may generate soft, widely separated bottom quarks, which are clustered with other hard partons during jet reconstruction. This could result in two high- p_T b -jets, where the soft bottom quarks contribute only a small fraction of the jets' total momentum. This corresponds to $pp \rightarrow b\bar{b}H + qq'$ events, where the q and q' are high- p_T light-flavor partons. This topology (single Higgs boson production plus two jets) is accounted for at LO in the POWHEG NNLOPS event generation, while the bottom quarks are introduced through the parton shower.
- Alternatively, the $g \rightarrow b\bar{b}$ splitting in the parton shower might produce two energetic, well-separated bottom quarks, which directly form the two high- p_T b -jets observed in the final state.

The second possibility implies that the parton shower is operating beyond its intended range of validity. On the contrary, the parton shower should introduce additional QCD radiation under the soft and collinear approximation. These two possibilities are explored by examining the distributions of the invariant mass of the two highest- p_T B -hadrons, shown in Figure A.2a, and the number of additional light jets (shown in Figure A.2b) in the fiducial region for the POWHEG NNLOPS sample, both with and without the $g \rightarrow b\bar{b}$ splitting. The impact of the $g \rightarrow b\bar{b}$ splitting on the selected events can be determined by analyzing the differences between these distributions. Figure A.2a demonstrates that enabling the $g \rightarrow b\bar{b}$ splittings populates the harder end of the m_{BB} spectrum, indicating that this process is responsible for generating highly energetic and well-separated B -hadrons in the fiducial region. This supports the second hypothesis, suggesting that the parton shower is indeed operating beyond its intended validity. Further insight is provided by Figure A.2b, which shows the number of additional light jets in events after applying the fiducial region requirements. The $g \rightarrow b\bar{b}$ splittings contribute to about one-third of the fiducial events with no additional light jets. This rules out the first hypothesis, which assumed a $pp \rightarrow b\bar{b}H + qq'$ topology, where two additional hard, well separated light jets are obtained at the event generation level, while the $g \rightarrow b\bar{b}$ splittings would produce two soft bottom quarks via the parton shower. Therefore, the $g \rightarrow b\bar{b}$ splittings in the parton shower appear to introduce a potential double counting of the $pp \rightarrow b\bar{b}H$ (plus no additional jets) topology in the fiducial region, a contribution that should already be accounted for by the POWHEG NNLOPS event generation at LO. These findings suggest that the 100% uncertainty assigned by the ATLAS HH searches to the $ggF + b\bar{b}$ background is likely adequate.

A.2 Impact of the new $b\bar{b}H$ modelling on the HH searches

As summarised in Section A.1, the new $b\bar{b}H$ simulation, based on MADGRAPH5_AMC@NLO, provides NLO accuracy and four-flavor scheme prediction for the $b\bar{b}H$ process, covering both the y_t^2 and y_b^2 components. This represents an improvement over the current prediction for the y_t^2 component, which is simulated only at LO and adopting the five-flavor scheme. In addition, the new $b\bar{b}H$ simulation allows to have a more reliable uncertainty estimation for the $ggF + b\bar{b}$ background for HH searches, by providing scale uncertain-




Figure A.2 – Distribution in the (a) invariant mass of the two hardest B hadrons and in the (b) number of light jets in the POWHEG NNLOPS sample, with and without the contributions from $g \rightarrow b\bar{b}$ splittings in the parton shower, after applying the fiducial region requirements. The shaded area in the m_{BB} distribution corresponds to the invariant mass requirement on the b -jets in the fiducial region.

ties on the $b\bar{b}H$ prediction in the fiducial region.

The impact of the new $b\bar{b}H$ modelling for the y_t^2 -induced contribution was estimated for the HH searches in the $b\bar{b}\gamma\gamma$ and $b\bar{b}\tau\tau$ final states. The lower cross section predicted by the new $b\bar{b}H$ simulation are factored into the expected contribution of the $ggF + b\bar{b}$ background in the signal regions for the $HH \rightarrow b\bar{b}\gamma\gamma$ and the $HH \rightarrow b\bar{b}\tau\tau$ analyses. The background from single Higgs boson production via ggF plus additional jets consists of two parts. An irreducible component from $ggF + b\bar{b}$ production contributes by 80% to the total ggF background events. The remaining 20% arises from a reducible contribution from ggF events with additional c - or light-flavor jets that are mistakenly tagged as b -jets, contributing to approximately 80% and 20% of the total ggF background, respectively. The latter reducible component is not included in the new $b\bar{b}H$ simulation, and is thus left unchanged. The former component is instead scaled down by the ratio of the NLO $b\bar{b}H$ prediction to the POWHEG NNLOPS prediction for the y_t^2 component. Additionally, the 100% uncertainty assigned to the single Higgs ggF background in the current analyses is replaced by the uncertainty derived from scale variations of the y_t^2 -induced $b\bar{b}H$ component, as shown in Table A.1.

The impact of the new simulation was assessed by recalculating the expected upper limits on the HH signal strength for both the $HH \rightarrow b\bar{b}\gamma\gamma$ and $HH \rightarrow b\bar{b}\tau\tau$ analyses, revealing improvements of approximately 2% and 8%, respectively. A more significant effect of the new $b\bar{b}H$ simulation was observed when extrapolating the expected sensitivity of these analyses to the High-Luminosity LHC scenario. Under this scenario, the expected upper limit on the HH signal strength and the expected discovery significance improved by 10% for the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis and by 20% for the $HH \rightarrow b\bar{b}\tau\tau$ analysis. To leverage the improved predictions for the $ggF + b\bar{b}$ background provided by the new $b\bar{b}H$ simulation in HH analyses, it would be necessary to integrate the $b\bar{b}H$ simulation

with the current POWHEG NNLOPS sample, which accounts for the reducible ggF background component. One possible approach is to disable bottom quark emissions (in both the parton shower and matrix element calculations) in the POWHEG NNLOPS sample, allowing the new $b\bar{b}H$ simulation to cover these contributions. Full validation of this new $b\bar{b}H$ simulation will ultimately require a direct measurement of the $b\bar{b}H$ process in data, which is currently unavailable.

An alternative background modelling for the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis

This Appendix explores an alternative background modeling strategy for the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis. This approach uses simulation-based templates for modeling the continuum background and exploits machine-learning-based discriminants as the final observables, rather than the traditional $m_{\gamma\gamma}$. Although this method, originally proposed for the ATLAS $HH \rightarrow b\bar{b}\gamma\gamma$ analysis, was not adopted for that analysis as described in Chapter 7, the study demonstrated its effectiveness: this approach was then adopted by the search for asymmetric Higgs boson pair production in the $b\bar{b}\gamma\gamma$ final state, discussed in Chapter 8.

The $HH \rightarrow b\bar{b}\gamma\gamma$ analysis described in Chapter 7 adopts the same analysis strategy shared, traditionally, by all the $H \rightarrow \gamma\gamma$ analyses. The presence of two photons in the final state determines the selection of the diphoton invariant mass $m_{\gamma\gamma}$ as the final discriminant variable. In fact, $m_{\gamma\gamma}$ offers a strong discriminating power by clearly distinguishing between the resonant signal, which appears as an additional contribution to the narrow $H \rightarrow \gamma\gamma$ peak, and the continuum background, which exhibits a smoothly decaying shape in the $m_{\gamma\gamma}$ spectrum. The resonant processes (including the HH signals and the single Higgs backgrounds) are modeled using a double-sided Crystal Ball function, while the continuum background (mainly from diphoton production) is described by a decaying analytical function, such as a simple exponential. Statistical results are derived from a maximum likelihood fit to the $m_{\gamma\gamma}$ distribution, performed simultaneously across mutually exclusive analysis categories.

The definition of the analysis categories is typically based on choosing thresholds on BDT outputs, set in order to maximise the expected signal significance. The definition of the categories must be independent of $m_{\gamma\gamma}$ to avoid sculpting the continuum background through the category selection process. Indeed, a high correlation between the BDT outputs and the diphoton invariant mass could introduce an artificial peaking structure in the diphoton background within the most sensitive categories, which would degrade the discriminating power of $m_{\gamma\gamma}$ itself. Furthermore, to accurately model the continuum background using an analytical function, a threshold on the expected background events must be imposed for each category. This requirement was found to be the primary factor in defining the most sensitive analysis categories in the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis: the thresholds defining the High Mass 3 and Low Mass 4 categories are completely

driven by the requirement on the minimum number of background events, which is chosen empirically to ensure the fit stability.

To address these challenges, an alternative analysis strategy for the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis has been developed, centered around a different choice of the final discriminant variable. The main idea is to modify the observable used in the final fit. Instead of performing an unbinned maximum likelihood fit to the $m_{\gamma\gamma}$ distribution across BDT-based categories, the BDT output is directly used as the final discriminant in a binned maximum likelihood fit.

The BDT discriminants effectively exploit the strong discriminating power of $m_{\gamma\gamma}$, which is included as an input variable during training. Since $m_{\gamma\gamma}$ is not the final observable, the continuum background sculpting in the $m_{\gamma\gamma}$ spectrum due to the correlation with the BDT discriminant is not a problem. This new analysis strategy requires a complete rethinking of the signal and background modeling compared to the traditional approach. Instead of relying on analytical functions, the signal HH processes, the single Higgs backgrounds, and the continuum background are modeled using templates of the BDT outputs derived from the corresponding MC samples. In particular, since the continuum background is modeled using MC simulations, the spurious signal uncertainty (discussed in Section 7.4.2) does not play a role. With this analysis strategy, the systematic uncertainties affecting the continuum background modeling are incorporated into the statistical model by providing varied continuum background templates, where either theoretical or experimental systematic uncertainties are propagated.

B.1 BDT discriminants

Following the same strategy as the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis, two BDT algorithms are trained, using the LIGHTGBM package to distinguish between signal and background processes in the High Mass and Low Mass regions, as defined in Section 7.3.2. Both BDTs consider all single Higgs production modes as backgrounds, along with the $\gamma\gamma$ +jets sample for the continuum background. Different HH benchmarks are used as signals for the BDTs in the two mass regions:

- **High Mass Region.** The BDT is trained using SM ggF HH events alongside anomalous VBF HH samples that correspond to five combinations of the coupling modifiers $(\kappa_\lambda, \kappa_{2V}, \kappa_V)$: $(1, 0, 1)$, $(1, 0.5, 1)$, $(1, 1.5, 1)$, $(1, 2, 1)$, and $(1, 3, 1)$. This configuration privileges the sensitivity to both SM HH events and VBF HH processes with κ_{2V} values deviating from unity, which have $m_{b\bar{b}\gamma\gamma}^*$ spectra that peak in the High Mass region.
- **Low Mass Region.** The BDT is trained on a combination of ggF HH and VBF HH events with $\kappa_\lambda = 10$. This choice helps to maximise the analysis sensitivity to HH production with anomalous κ_λ values.

During training, each sample is assigned a weight to ensure that the total weights for all signal and background processes are balanced. The relative weights of VBF HH and ggF HH events within the signal, as well as the weights of single Higgs events and $\gamma\gamma$ +jets events within the total background, are determined based on their ex-

pected event yields in the High or Low Mass region, specifically within the $m_{\gamma\gamma}$ range of $123 < m_{\gamma\gamma} < 127$ GeV. Both BDTs are trained using the same set of input variables, which largely overlap with those used in the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis. A summary of these input variables, which include the diphoton invariant mass $m_{\gamma\gamma}$, is provided in Table B.1. The BDT discriminants for the Low and High Mass regions, displaying data

Table B.1 – Input variables used for training the BDT-based algorithms in both the High and Low mass regions.

Type	Variable	Description
Photon-related kinematic variables	$m_{\gamma\gamma}$	Invariant mass of the two selected photons.
	p_T, η and ϕ	Transverse momentum and angular variables of the leading and subleading photons.
	$p_T^{\gamma\gamma}, \eta^{\gamma\gamma}$, and $\phi^{\gamma\gamma}$	Transverse momentum and angular variables of the diphoton system.
Jet-related kinematic variables	b -tag status of the two b -jets	Pseudo-continuous b -tagging score of the two selected b -jets.
	p_T, η , and ϕ of the two b -jets	Transverse momentum and angular variables of the two selected b -jets.
	$p_T^{b\bar{b}}, \eta^{b\bar{b}}$, and $\phi^{b\bar{b}}$	Transverse momentum and angular variables of the di-jet object.
	$m_{b\bar{b}}$	Invariant mass of the two selected b -jets.
	H_T	Scalar sum of the transverse momenta of all the jets reconstructed in the event.
VBF-targeting variables	χ_{Wt}	Single topness.
	b -tag status of the two VBF jets	Pseudo-continuous b -tagging score of the two selected VBF jets.
	p_T, η , and ϕ of the two VBF jets	Transverse momentum and angular variables of the two selected VBF jets.
	m_{jj}^{VBF} and $\Delta\eta(j_1^{\text{VBF}}, j_2^{\text{VBF}})$	Invariant mass and angular separation of the two VBF jets.
	$\eta(j_1^{\text{VBF}}) \times \eta(j_2^{\text{VBF}})$	Product of the pseudorapidities of the two VBF jets.
Missing transverse momentum-related variables	$\min \Delta R(\gamma, j^{\text{VBF}})$ and $\min \Delta R(b, j^{\text{VBF}})$	Minimum angular distance between a photon and a VBF jet and a b -jet and a VBF jet.
	E_T^{miss} and ϕ^{miss}	Missing transverse momentum and its azimuthal direction.
$b\bar{b}\gamma\gamma$ -related variables	$m_{b\bar{b}\gamma\gamma}^*$	Modified invariant mass of the $b\bar{b}\gamma\gamma$ system.
Event-shape variables	$S_\perp, P_f$	Transverse sphericity and planar flow of the $b\bar{b}\gamma\gamma$ system.

in the $m_{\gamma\gamma}$ sidebands region, HH signals, and background distributions, are presented in Figure B.1b and Figure B.1a, respectively.

B.2 Analysis strategy

Signal extraction is performed through a simultaneous fit to the BDT discriminants in four orthogonal regions, defined using the modified $b\bar{b}\gamma\gamma$ invariant mass ($m_{b\bar{b}\gamma\gamma}^*$) and the diphoton invariant mass ($m_{\gamma\gamma}$), as described below.

First, events are categorized based on the $m_{b\bar{b}\gamma\gamma}^*$ variable into the High and Low Mass




Figure B.1 – Distributions of the BDT discriminants for the High Mass region (left) and Low Mass region (right), extracted for ggF HH and VBF HH signal events (both in the SM and for anomalous κ_λ and κ_{2V} scenarios), single Higgs boson backgrounds, and $\gamma\gamma$ +jets events. Distributions extracted from data in the $m_{\gamma\gamma}$ sidebands ($105 < m_{\gamma\gamma} < 120$ GeV and $130 < m_{\gamma\gamma} < 160$ GeV) are also displayed. All distributions are normalised to unity.

regions, separated by a threshold of $m_{b\bar{b}\gamma\gamma}^* = 350$ GeV. Within each mass region, events are further divided into two categories based on the $m_{\gamma\gamma}$ variable:

- A **Signal region (SR)** is defined by $120 \leq m_{\gamma\gamma} \leq 130$ GeV, corresponding to the window around the $H \rightarrow \gamma\gamma$ peak. Any potential HH signal is expected to contribute to the $H \rightarrow \gamma\gamma$ resonance, centered at $m_H \approx 125$ GeV. Therefore, this region drives the sensitivity to the resonant signal.
- A **Sidebands region (SB)** is defined by $105 \leq m_{\gamma\gamma} < 120$ GeV plus $130 \leq m_{\gamma\gamma} \leq 160$ GeV, and is expected to include almost null contribution from the HH signal or from single Higgs backgrounds. The main role of the SB region is to constrain the normalisation of the continuum background in the final fit to data.

The definition of the analysis regions is summarised in Figure B.2.

B.3 Signal and background modelling

Both the signal and background processes in each analysis region are modeled using templates of the two BDT discriminants. In the High Mass region, the High Mass BDT output is used as the final discriminant, while the Low Mass BDT output is used in the Low Mass region.

The BDT templates for the HH signals and single Higgs boson backgrounds are extracted from the corresponding MC samples (described in Section 7.2.1 and 7.2.2). These histograms are normalized according to the state-of-the-art theoretical predictions for inclusive cross sections and branching fractions, as described in Table 1.2 and Section 7.4.1. The effects of finite MC statistics when building the BDT templates are taken into account




Figure B.2 – Sketch of the four analysis regions.

in the statistical model by introducing dedicated nuisance parameters (as described in Section 8.6).

The continuum background consists of two reconstructed photons plus additional jets, and includes the contributions from events with two real photon and from events with fake photons, arising from the γj and jj processes. The fractions of the irreducible true $\gamma\gamma$ component and the reducible γj and jj components are determined using the data-driven two-dimensional matrix method, detailed in Section 7.4.2. The SHERPA $\gamma\gamma$ +jets sample, which simulates the true photon component, is used to model the entire continuum background, including the reducible components. As will be explained in detail later, the shape of the BDT discriminants for the $\gamma\gamma$ +jets process was found to be consistent with those from the reducible diphoton backgrounds. Nevertheless, a dedicated systematic uncertainty is assigned to this assumption. To correct for the difference in normalisation between the true photon component and the full reconstructed diphoton background, the $\gamma\gamma$ +jets templates are scaled by normalisation factors which are allowed to vary freely in the final fit to data. Thus, while the shape of the BDT discriminants is derived from MC simulations, the normalisation of the continuum background is data-driven. Two normalisation factors are defined separately for the High Mass and Low Mass regions ($\mu_{\gamma\gamma}^{\text{HM}}$ and $\mu_{\gamma\gamma}^{\text{LM}}$) and are mainly constrained by the data in the High and Low Mass SB regions.

This fitting strategy, where the continuum background normalisation is primarily constrained in the SB regions and propagated to the corresponding SRs, is validated by performing background-only fits to single-binned BDT templates in data in the SB regions. In these fits, only the normalisation factors $\mu_{\gamma\gamma}^{\text{HM}}$ and $\mu_{\gamma\gamma}^{\text{LM}}$ are allowed to float. The $m_{\gamma\gamma}$ distributions for the background samples - where the $\gamma\gamma$ +jets sample is normalised based on the post-fit values of $\mu_{\gamma\gamma}^{\text{HM}}$ and $\mu_{\gamma\gamma}^{\text{LM}}$ - are then compared with the $m_{\gamma\gamma}$ distributions observed in the data. This comparison, shown in Figure B.3, confirms that the normalisation of the continuum background extracted in the SB regions is consistent with the data in the full $m_{\gamma\gamma}$ spectrum.




Figure B.3 – Comparison of the $m_{\gamma\gamma}$ distributions in background-only simulated events and in data in the High Mass (left) and Low Mass (right) regions. The normalisation factors applied to the $\gamma\gamma$ +jets sample are extracted via a background-only fit to data in the SB regions. The $m_{\gamma\gamma}$ distributions for the SM HH signal are also shown. The bottom panel shows the ratio of the data to the background-only expectation in the SB regions.

The approach to modeling the continuum background in this analysis represents a significant change from the traditional method, where both the shape and normalisation of the continuum background are mainly data-driven. In the traditional approach, the only MC-based input is the choice of the particular analytical function used to model the continuum background in each category, which is selected based on a high-statistics $m_{\gamma\gamma}$ template. To ensure the reliability of the BDT templates derived from the SHERPA $\gamma\gamma$ +jets sample in describing the entire continuum background, the shape of these templates is compared to the data in the SB regions, as shown in Figure B.4. Additionally, since the SHERPA $\gamma\gamma$ +jets sample does not include the fake photon contributions from γj and jj processes, BDT templates for the reducible components are extracted from control regions in data enriched in fake photon background events, and compared with the $\gamma\gamma$ +jets-based histograms. The comparison between the fake photon templates and the SHERPA $\gamma\gamma$ +jets templates shows consistency within uncertainties, as displayed in Figure B.5.

Figure B.6 displays the BDT templates for the continuum background, HH signals, and single Higgs backgrounds in the High Mass and Low Mass SRs. The binning of these templates is selected to ensure at least one expected background event in the most signal-like bin (where the BDT output is close to unity). This binning strategy was found to be sufficient to preserve the validity of the asymptotic approximation used for extracting statistical results.

In Figure B.7, the BDT templates for the continuum background, HH signals, and single Higgs backgrounds in the High Mass and Low Mass SB regions are shown. Since the signal and resonant background contamination in the SB regions is negligible and their




Figure B.4 – Comparison of the BDT discriminant distributions from the $\gamma\gamma$ +jets process and from data in the High Mass and Low Mass SB regions. The bin edges are chosen to ensure at least one event in data for each bin. All distributions are normalised to unity. The bottom panel shows the ratio of the data to the SHERPA $\gamma\gamma$ +jets sample.




Figure B.5 – Comparison of the BDT templates for the γj and jj processes, built using control regions in data, with those derived from the SHERPA $\gamma\gamma$ +jets sample, in the High Mass SR (left) and Low Mass SR (right). The bin edges are chosen to ensure at least one event in data for each bin. All distributions are normalised to unity. The bottom panel shows the ratio of the data-driven γj and jj distributions to the SHERPA $\gamma\gamma$ +jets sample.




Figure B.6 – BDT templates for the SM HH signals, single Higgs backgrounds, and continuum backgrounds in the High Mass (left) and Low Mass (right) SRs. The normalisations of the continuum background templates are set from a background-only fit to data in the SB regions.

primary role is to constrain the continuum background normalisation to match data, single-binned BDT templates are chosen for these regions.




Figure B.7 – BDT templates for the SM HH signals, single Higgs backgrounds, and continuum backgrounds in the High Mass (left) and Low Mass (right) SB regions. The normalisations of the continuum background templates are set from a background-only fit to data in the SB regions.

B.4 Systematic uncertainties

The choice of the BDT-based final discriminant and the modelling strategy adopted by this analysis introduce a different implementation of the systematic uncertainties for both the signal and the backgrounds, compared with the traditional $H \rightarrow \gamma\gamma$ approach. In this method, systematic uncertainties that impact both the shape and normalisation

of signal and background processes are incorporated into the statistical model by providing alternative BDT templates, where the systematic variations are propagated. This also applies to the continuum background, where both experimental and theoretical uncertainties (described in Section 7.5) are accounted for. In addition, since the continuum background is modelled using MC simulations, there is no need to select a particular analytical function to describe the background shape: therefore, the spurious signal uncertainty is not present in this analysis.

Although the experimental and theoretical systematic uncertainties affecting the HH signal and single Higgs processes are also present in the traditional $H \rightarrow \gamma\gamma$ analysis (see Section 7.5) - and are expected to have a similar impact here - the treatment of the continuum background differs significantly in this approach. Therefore, the following sections focus on assessing the impact of experimental and theoretical systematic uncertainties on the continuum background.

As the normalisation of the continuum background is determined by the fit to data, even a very large systematic variation can be fully absorbed by the fit, as long as its impact is flat in the spectrum of the BDT discriminant and is similar in both the SRs and SB regions. Therefore, the systematic uncertainties affecting only the normalisation of the continuum background have a negligible effect on the statistical results. On the other hand, the systematic uncertainties that affect the shape of the BDT discriminants for the continuum background are expected to have a more noticeable effect on the statistical results and on the analysis sensitivity.

B.4.1 Experimental systematic uncertainties

A 0.83% uncertainty, related to the measurement of the total integrated luminosity of Run 2 data [49], is applied to the normalisation of the continuum background templates. In addition, systematic uncertainties due to potential mismodelling of the detector response in simulation with respect to data, related to the modelling of photons, jets, flavor tagging efficiencies, and pileup affect both the normalisation and the shape of the continuum background. These uncertainties lead to variations in the overall yield and efficiency of the continuum background, ranging from approximately 1% to 10%.

B.4.2 Theoretical systematic uncertainties

Scale and $\text{PDF}+\alpha_s$ uncertainties

Theoretical systematic uncertainties arising from missing higher-order corrections in the perturbative calculations of the $\gamma\gamma$ +jets process (referred to as “scale” uncertainties) affect the modeling of the continuum background. Similarly, uncertainties propagated from the PDFs and the value of the strong coupling constant α_s , collectively termed “ $\text{PDF}+\alpha_s$ ” uncertainties, impact both the normalisation and shape of the continuum background templates. The effects of scale and $\text{PDF}+\alpha_s$ uncertainties on the continuum background (ranging from approximately 1% for the $\text{PDF}+\alpha_s$ uncertainties to over 40% for the scale uncertainties) across the four analysis regions are shown in Figures B.8 and B.10.




Figure B.8 – Nominal and varied BDT distributions for the continuum background in the four analysis regions. The varied BDT templates refer to $\pm 1\sigma$ variations from the scale uncertainty on the $\gamma\gamma + \text{jets}$ process.




Figure B.9 – Nominal and varied BDT distributions for the continuum background in the four analysis regions. The varied BDT templates refer to $\pm 1\sigma$ variations from the PDF+ α_S uncertainty on the $\gamma\gamma + \text{jets}$ process.

Uncertainty on the reducible $\gamma j + jj$ component

As detailed in Section B.3, the SHERPA $\gamma\gamma$ +jets MC sample is used to model the entire continuum background, including the fake photon component. A dedicated systematic uncertainty, related to neglecting the γj and jj contributions when constructing the continuum background BDT templates, is introduced. An alternative BDT template is created by combining the $\gamma\gamma$, γj , and jj components according to their purities (listed in Section 7.4.2). The BDT template for the $\gamma\gamma$ component is derived from the SHERPA $\gamma\gamma$ +jets sample, while the γj and jj templates are built using control regions in data, as explained in Section 7.4.2. Figure B.10 illustrates the nominal and alternative BDT templates for the continuum background, along with the impact of this uncertainty, for both the High Mass and Low Mass SRs.




Figure B.10 – Nominal and alternative BDT distributions for the continuum background in the High Mass (left) and Low Mass (right) SRs. The alternative BDT templates represent the systematic uncertainty from neglecting the reducible diphoton component when building the continuum background histogram.

Matrix-element and parton shower uncertainty

An uncertainty related to the modelling of both the matrix element and the parton shower in the $\gamma\gamma$ +jets process is considered. This uncertainty is evaluated by comparing the nominal BDT templates (extracted from the SHERPA sample) with those obtained from an alternative $\gamma\gamma$ +jets sample, relying on MADGRAPH5_AMC@NLO for the description of the matrix element and on PYTHIA 8 for the modelling of the parton shower. However, the MADGRAPH5_AMC@NLO $\gamma\gamma$ +jets sample has significantly lower statistics in the signal regions of the $HH \rightarrow b\bar{b}\gamma\gamma$ analysis compared to the nominal SHERPA sample. This leads to inflated modeling uncertainties due to large statistical fluctuations in the alternative BDT templates. The impact of this uncertainty on the shape and

normalisation of the continuum background in the four analysis regions is shown in Figure B.11.




Figure B.11 – Nominal BDT distribution for the continuum diphoton background extracted from the SHERPA $\gamma\gamma$ +jets sample in the four analysis regions. The nominal template is compared with the corresponding varied histograms obtained from the MADGRAPH5_AMC@NLO sample. The shaded bands represent the statistical uncertainty associated to each template.

B.5 Expected results

The expected sensitivity of the analysis is determined through a binned maximum likelihood background-only fit to the BDT distributions in the SB regions. The statistical analysis allows to extract an expected upper limit on the HH signal strength μ_{HH} , un-

der the hypothesis of no HH production, at 95% CL. Similarly, 95% expected confidence intervals for the coupling modifiers κ_λ and κ_{2V} are obtained using the expected profile log-likelihoods $-2 \ln \Lambda$, based on the assumption of SM HH production, where $\kappa_\lambda =$ and $\kappa_{2V} = 1$. When deriving the expected constraints for κ_λ or κ_{2V} , all other coupling modifiers that influence HH production are kept fixed at their SM values.

The 95% CL expected upper limit on μ_{HH} is 4.97. This outcome is competitive with the sensitivity of the traditional $HH \rightarrow b\bar{b}\gamma\gamma$ analysis, which has a 95% CL expected upper limit on μ_{HH} of 4.88, evaluated including only the spurious signal uncertainty¹. The effects of the systematic uncertainties on the continuum background modeling are detailed in Table B.2, with the overall impact found to be below 2%.

Table B.2 – Breakdown of the dominant systematic uncertainties in the expected μ_{HH} upper limit at 95% CL. The impact of the uncertainties corresponds to the relative variation of the expected upper limit when re-evaluating the profile likelihood ratio after fixing the nuisance parameter in question to its best-fit value, while all remaining nuisance parameters remain free to float.

Systematic uncertainty source	Relative impact [%]
MC statistics	0.9
Experimental	< 0.1
Scale	< 0.1
PDF+ α_S	< 0.1
Matrix Element and Parton showering model	0.1
Fake photon component	0.7

The expected values of $-2 \ln \Lambda$ as functions of κ_λ and κ_{2V} are illustrated in Figures B.12a and B.12b, respectively. The value of κ_λ is constrained within the range of $-2.1 < \kappa_\lambda < 7.4$ at 95% CL, while the allowed range for κ_{2V} is $-0.7 < \kappa_{2V} < 2.9$ at 95% CL. While the sensitivity to the coupling modifier κ_λ aligns with that of the traditional $HH \rightarrow b\bar{b}\gamma\gamma$ analysis, the sensitivity to κ_{2V} has significantly improved. In fact, the width of the expected allowed κ_{2V} interval is reduced by approximately 15% compared to the traditional analysis.

¹The expected upper limit on μ_{HH} for the traditional $HH \rightarrow b\bar{b}\gamma\gamma$ analysis was determined by fitting data exclusively in the sidebands region, to allow for a fair comparison with the approach presented in this Appendix. Hence, the value quoted here differs from the expected upper limit on μ_{HH} reported in Chapter 7, where the limit was derived from a fit to data across the full $m_{\gamma\gamma}$ spectrum, spanning $105 \leq m_{\gamma\gamma} \leq 160$ GeV.




Figure B.12 – Expected value of the profile log-likelihood ratio, $-2\ln\Lambda$ as a function of κ_λ (left) and κ_{2V} (right), with all the other coupling modifiers (including, respectively, κ_{2V} or κ_λ) are fixed to their SM predictions. The value of $-2\ln\Lambda$ is shown either including (solid line) or neglecting (dashed line) the systematic uncertainties on the continuum background.

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*, *Phys. Lett. B* **716** (2012) 1, arXiv: [1207.7214 \[hep-ex\]](#) (cit. on pp. ix, 1, 21, 22).
- [2] CMS Collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, *Phys. Lett. B* **716** (2012) 30, arXiv: [1207.7235 \[hep-ex\]](#) (cit. on pp. ix, 1, 21, 22).
- [3] ATLAS Collaboration, *A detailed map of Higgs boson interactions by the ATLAS experiment ten years after the discovery*, *Nature* **607** (2022) 52, arXiv: [2207.00092 \[hep-ex\]](#) (cit. on pp. ix, 19, 25, 26).
- [4] CMS Collaboration, *A portrait of the Higgs boson by the CMS experiment ten years after the discovery.*, *Nature* **607** (2022) 60, arXiv: [2207.00043 \[hep-ex\]](#) (cit. on pp. ix, 25, 26, 39, 40).
- [5] ATLAS and C. Collaborations, *Combined measurement of the Higgs boson mass in pp collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS experiments*, *Phys. Rev. Lett.* **114** (2015) 191803, arXiv: [1503.07589 \[hep-ex\]](#) (cit. on pp. ix, 22).
- [6] R. L. Workman et al., *Review of Particle Physics*, *PTEP* **2022** (2022) 083C01 (cit. on pp. ix, 11, 16, 119).
- [7] 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** (2024) 101801, arXiv: [2406.09971 \[hep-ex\]](#) (cit. on pp. ix, xi, 179).
- [8] L. Evans and P. Bryant, *LHC Machine*, *JINST* **3** (2008) S08001 (cit. on pp. x, 49).
- [9] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, *JINST* **3** (2008) S08003, URL: <https://cds.cern.ch/record/1129811> (cit. on pp. x, 51, 59).
- [10] 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*, *JHEP* **01** (2024) 066, arXiv: [2310.12301 \[hep-ex\]](#) (cit. on pp. x, xii, 135).

- [11] B. Di Micco et al., *Higgs boson pair production at colliders: status and perspectives. Higgs boson potential at colliders: Status and perspectives*, *Rev. Phys.* **5** (2020) 100045, ed. by B. Di Micco, arXiv: [1910.00012](#), URL: <https://cds.cern.ch/record/2692014> (cit. on pp. x, 31, 32, 36, 133).
- [12] ATLAS Collaboration, *Search for Higgs boson pair production in the two bottom quarks plus two photons final state in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **106** (2022) 052001, arXiv: [2112.11876 \[hep-ex\]](#) (cit. on pp. x, xiii, 151, 174, 176, 214).
- [13] 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** (2023) 052003, arXiv: [2301.03212 \[hep-ex\]](#) (cit. on pp. xi, 179).
- [14] 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** (2024) 139007, arXiv: [2404.17193 \[hep-ex\]](#) (cit. on pp. xi, 179).
- [15] 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** (2024) 032012, arXiv: [2404.12660 \[hep-ex\]](#) (cit. on pp. xi, 179).
- [16] C.-Y. Chen et al., *Next-to-minimal two Higgs doublet model*, *Phys. Rev. D* **89** (2014) 075009, arXiv: [1312.3949 \[hep-ph\]](#) (cit. on pp. xi, 43).
- [17] T. Robens et al., *Two-real-scalar-singlet extension of the SM: LHC phenomenology and benchmark scenarios*, *Eur. Phys. J. C* **80** (2020) 151, arXiv: [1908.08554 \[hep-ph\]](#) (cit. on pp. xi, 43, 45, 185).
- [18] ATLAS Collaboration, *Search for a resonance decaying into a scalar particle and a Higgs boson in the final state with two bottom quarks and two photons in proton-proton collisions at a center of mass energy of 13 TeV with the ATLAS detector*, (2024), arXiv: [2404.12915 \[hep-ex\]](#) (cit. on pp. xi, xii, 185).
- [19] ATLAS Collaboration, *Electron and photon energy calibration with the ATLAS detector using LHC Run 2 data*, *JINST* **19** (2024) P02009, arXiv: [2309.05471 \[hep-ex\]](#) (cit. on pp. xii, 87, 93, 95, 96, 100, 102, 103).
- [20] *Graph Neural Network Jet Flavour Tagging with the ATLAS Detector*, tech. rep., CERN, 2022, URL: <https://cds.cern.ch/record/2811135> (cit. on pp. xii, 113, 129).
- [21] S. Manzoni et al., *Taming a leading theoretical uncertainty in HH measurements via accurate simulations for $b\bar{b}H$ production*, *JHEP* **2309** (2023) 179, arXiv: [2307.09992](#), URL: <http://cds.cern.ch/record/2875151> (cit. on pp. xiii, 233).
- [22] 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** (2023) 137745, arXiv: [2211.01216 \[hep-ex\]](#) (cit. on pp. xiii, 160).

- [23] S. Weinberg, *The making of the Standard Model*, [The European Physical Journal C - Particles and Fields](#) **34** (2004) 5, URL: <https://doi.org/10.1140/epjc/s2004-01761-1> (cit. on p. 1).
- [24] G. 't Hooft, *Under the spell of the gauge principle*, vol. 19, 1994 (cit. on p. 1).
- [25] B. de Wit, *Introduction to gauge theories and the Standard Model*, (1995), URL: <https://cds.cern.ch/record/292286> (cit. on p. 1).
- [26] G. 't Hooft, *Renormalizable Lagrangians for massive Yang-Mills fields*, [Nucl. Phys. B](#) **35** (1971) 167, ed. by J. C. Taylor (cit. on p. 1).
- [27] G. 't Hooft and M. J. G. Veltman, *Regularization and renormalization of gauge fields*, [Nucl. Phys. B](#) **44** (1972) 189 (cit. on p. 1).
- [28] S. L. Glashow, *Partial-symmetries of weak interactions*, [Nuclear Physics](#) **22** (1961) 579, ISSN: 0029-5582, URL: <https://www.sciencedirect.com/science/article/pii/0029558261904692> (cit. on p. 1).
- [29] S. Weinberg, *A Model of Leptons*, [Phys. Rev. Lett.](#) **19** (21 1967) 1264, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.19.1264> (cit. on p. 1).
- [30] F. Gross et al., *50 Years of quantum chromodynamics*, [The European Physical Journal C](#) **83** (2023) 1125, URL: <https://doi.org/10.1140/epjc/s10052-023-11949-2> (cit. on p. 1).
- [31] Y. Fukuda et al., *Evidence for oscillation of atmospheric neutrinos*, [Phys. Rev. Lett.](#) **81** (8 1998) 1562, URL: <https://link.aps.org/doi/10.1103/PhysRevLett.81.1562> (cit. on p. 2).
- [32] M. Aker et al., *Direct neutrino-mass measurement with sub-electronvolt sensitivity*, [Nature Physics](#) **18** (2022) 160, URL: <https://doi.org/10.1038/s41567-021-01463-1> (cit. on p. 2).
- [33] B. R. Stella and H.-J. Meyer, *$Y(9.46\text{ GeV})$ and the gluon discovery (a critical recollection of PLUTO results)*, [Eur. Phys. J. H](#) **36** (2011) 203, arXiv: 1008.1869 [hep-ex] (cit. on p. 7).
- [34] L. Di Lella and C. Rubbia, *The Discovery of the W and Z Particles*, [Adv. Ser. Dir. High Energy Phys.](#) **23** (2015) 137, URL: <https://cds.cern.ch/record/2103277> (cit. on p. 10).
- [35] F. Englert and R. Brout, *Broken symmetry and the mass of gauge vector mesons*, [Phys. Rev. Lett.](#) **13** (1964) 321, ed. by J. C. Taylor (cit. on p. 11).
- [36] P. W. Higgs, *Broken symmetries and the masses of gauge bosons*, [Phys. Rev. Lett.](#) **13** (1964) 508, ed. by J. C. Taylor (cit. on p. 11).
- [37] G. S. Guralnik et al., *Global conservation laws and massless particles*, [Phys. Rev. Lett.](#) **13** (1964) 585, ed. by J. C. Taylor (cit. on p. 11).
- [38] G. S. Guralnik, *Gauge invariance and the Goldstone theorem*, [Mod. Phys. Lett. A](#) **26** (2011) 1381, arXiv: 1107.4592 [hep-th] (cit. on p. 13).

- [39] S. Heinemeyer et al., *Handbook of LHC Higgs Cross Sections: 3. Higgs Properties: Report of the LHC Higgs Cross Section Working Group*, ed. by S. Heinemeyer, CERN Yellow Reports: Monographs, 2013, URL: <https://cds.cern.ch/record/1559921> (cit. on pp. 20, 21).
- [40] D. de Florian et al., *Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector*, CERN Yellow Reports: Monographs, Geneva: CERN, 2017, URL: <https://cds.cern.ch/record/2227475> (cit. on pp. 18, 20, 24, 168).
- [41] M. Bonvini et al., *On the Higgs cross section at $N^3\text{LO}+N^3\text{LL}$ and its uncertainty*, *Journal of High Energy Physics* **2016** (2016) 105, URL: [https://doi.org/10.1007/JHEP08\(2016\)105](https://doi.org/10.1007/JHEP08(2016)105) (cit. on p. 18).
- [42] ATLAS and CMS Collaborations, *Combination of measurements of the top quark mass from data collected by the ATLAS and CMS experiments at $\sqrt{s} = 7$ and 8 TeV*, *Phys. Rev. Lett.* **132** (2024) 261902, arXiv: 2402.08713 [hep-ex] (cit. on p. 21).
- [43] ATLAS Collaboration, *Observation of $H \rightarrow b\bar{b}$ decays and VH production with the ATLAS detector*, *Phys. Lett. B* **786** (2018) 59, arXiv: 1808.08238 [hep-ex] (cit. on p. 21).
- [44] ATLAS Collaboration, *Evidence for the Higgs-boson Yukawa coupling to tau leptons with the ATLAS detector*, *JHEP* **04** (2015) 117, arXiv: 1501.04943 [hep-ex] (cit. on p. 21).
- [45] CMS Collaboration, *Evidence for the Higgs boson decay to a bottom quark–antiquark pair*, *Phys. Lett. B* **780** (2018) 501, arXiv: 1709.07497 [hep-ex] (cit. on p. 21).
- [46] CMS Collaboration, *Observation of the Higgs boson decay to a pair of τ leptons with the CMS detector*, *Phys. Lett. B* **779** (2018) 283, arXiv: 1708.00373 [hep-ex] (cit. on p. 21).
- [47] ATLAS Collaboration, *A search for the dimuon decay of the Standard Model Higgs boson with the ATLAS detector*, *Phys. Lett. B* **812** (2021) 135980, arXiv: 2007.07830 [hep-ex] (cit. on p. 21).
- [48] CMS Collaboration, *Evidence for Higgs boson decay to a pair of muons*, *JHEP* **01** (2021) 148, arXiv: 2009.04363 [hep-ex] (cit. on p. 21).
- [49] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*, *Eur. Phys. J. C* **83** (2023) 982, arXiv: 2212.09379 [hep-ex] (cit. on pp. 22, 56, 137, 163, 186, 201, 247).
- [50] ATLAS Collaboration, *Combined measurement of the Higgs boson mass from the $H \rightarrow \gamma\gamma$ and $H \rightarrow ZZ^* \rightarrow 4\ell$ decay channels with the ATLAS detector using $\sqrt{s} = 7, 8$, and 13 TeV pp collision data*, *Phys. Rev. Lett.* **131** (2023) 251802, arXiv: 2308.04775 [hep-ex] (cit. on pp. 22, 23).
- [51] CMS Collaboration, *A measurement of the Higgs boson mass in the diphoton decay channel*, *Phys. Lett. B* **805** (2020) 135425, arXiv: 2002.06398 [hep-ex] (cit. on pp. 22, 23).

- [52] CMS Collaboration, *Measurement of the Higgs boson mass and width using the four leptons final state*, tech. rep., CERN, 2023, URL: <https://cds.cern.ch/record/2871702> (cit. on pp. 22, 23).
- [53] ATLAS Collaboration, *Evidence of off-shell Higgs boson production from ZZ leptonic decay channels and constraints on its total width with the ATLAS detector*, *Phys. Lett. B* **846** (2023) 138223, arXiv: 2304.01532 [hep-ex] (cit. on p. 24).
- [54] CMS Collaboration, *Measurement of the Higgs boson width and evidence of its off-shell contributions to ZZ production. First evidence for off-shell production of the Higgs boson and measurement of its width*, *Nature Phys.* (2022), arXiv: 2202.06923, URL: <https://cds.cern.ch/record/2801541> (cit. on p. 24).
- [55] ATLAS Collaboration, *Study of the spin and parity of the Higgs boson in diboson decays with the ATLAS detector*, *Eur. Phys. J. C* **75** (2015) 476, arXiv: 1506.05669 [hep-ex] (cit. on p. 24).
- [56] CMS Collaboration, *Constraints on the spin-parity and anomalous HVV couplings of the Higgs boson in proton collisions at 7 and 8 TeV*, *Phys. Rev. D* **92** (2015) 012004, arXiv: 1411.3441 [hep-ex] (cit. on p. 24).
- [57] D. Buttazzo et al., *Investigating the near-criticality of the Higgs boson*, *JHEP* **12** (2013) 089, arXiv: 1307.3536 [hep-ph] (cit. on p. 29).
- [58] G. Degrandi et al., *Higgs mass and vacuum stability in the Standard Model at NNLO*, *JHEP* **08** (2012) 098, arXiv: 1205.6497 [hep-ph] (cit. on p. 29).
- [59] D. S. Gorbunov and V. A. Rubakov, *Introduction to the theory of the early universe: Cosmological perturbations and inflationary theory*, 2011 (cit. on p. 29).
- [60] S. Dawson et al., “Report of the Topical Group on Higgs Physics for Snowmass 2021: The Case for Precision Higgs Physics”, *Snowmass 2021*, 2022, arXiv: 2209.07510 [hep-ph] (cit. on p. 29).
- [61] M. Grazzini et al., *Higgs boson pair production at NNLO with top quark mass effects*, *JHEP* **05** (2018) 059, arXiv: 1803.02463 [hep-ph] (cit. on p. 32).
- [62] J. Baglio et al., *gg → HH : Combined uncertainties*, *Phys. Rev. D* **103** (2021) 056002, arXiv: 2008.11626 [hep-ph] (cit. on p. 32).
- [63] LHC Higgs Cross Section Working Group, *Website of the LHC HH Cross Section Working Group*, URL: <https://twiki.cern.ch/twiki/bin/view/LHCPhysics/LHCHWGH> (cit. on pp. 32, 33).
- [64] F. A. Dreyer and A. Karlberg, *Fully differential Vector-Boson Fusion Higgs Pair Production at Next-to-Next-to-Leading Order*, *Phys. Rev. D* **99** (2019) 074028, arXiv: 1811.07918 [hep-ph] (cit. on p. 33).
- [65] F. A. Dreyer and A. Karlberg, *Vector-Boson Fusion Higgs pair production at N³LO*, *Phys. Rev. D* **98** (2018) 114016, arXiv: 1811.07906 [hep-ph] (cit. on p. 33).
- [66] F. A. Dreyer et al., *Precise predictions for double-Higgs production via vector-boson fusion*, *Eur. Phys. J. C* **80** (2020) 1037, arXiv: 2005.13341 [hep-ph] (cit. on p. 33).

- [67] S. Amoroso et al., “Les Houches 2019: Physics at TeV Colliders: Standard Model Working Group Report”, *11th Les Houches Workshop on Physics at TeV Colliders: PhysTeV Les Houches*, 2020, arXiv: [2003.01700 \[hep-ph\]](#) (cit. on p. 33).
- [68] 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*, [JHEP 07 \(2014\) 079](#), arXiv: [1405.0301 \[hep-ph\]](#) (cit. on pp. 33, 36, 139, 146, 186).
- [69] F. Bishara et al., *Higgs pair production in vector-boson fusion at the LHC and beyond*, [Eur. Phys. J. C 77 \(2017\) 481](#), arXiv: [1611.03860 \[hep-ph\]](#) (cit. on pp. 36, 139).
- [70] G. Cowan et al., *Asymptotic formulae for likelihood-based tests of new physics*, [Eur. Phys. J. C 71 \(2011\) 1554](#), arXiv: [1007.1727 \[physics.data-an\]](#) (cit. on pp. 39, 41, 151, 170, 172, 174, 181, 200, 210, 211).
- [71] CMS Collaboration, *Constraints on the Higgs boson self-coupling from the combination of single and double Higgs boson production in proton-proton collisions at $\sqrt{s} = 13$ TeV*, (2024), arXiv: [2407.13554 \[hep-ex\]](#) (cit. on p. 39).
- [72] *Standard Model Summary Plots June 2024*, tech. rep., CERN, 2024, URL: <https://cds.cern.ch/record/2903866> (cit. on p. 39).
- [73] V. C. Rubin and W. K. Ford Jr., *Rotation of the Andromeda Nebula from a Spectroscopic Survey of Emission Regions*, [Astrophys. J. 159 \(1970\) 379](#) (cit. on p. 39).
- [74] J. A. Tyson et al., *Detailed mass map of CL0024+1654 from strong lensing*, [Astrophys. J. Lett. 498 \(1998\) L107](#), arXiv: [astro-ph/9801193](#) (cit. on p. 39).
- [75] P. Collaboration, *Planck 2018 results. VI. Cosmological parameters*, [Astron. Astrophys. 641 \(2020\) A6](#), arXiv: [1807.06209 \[astro-ph.CO\]](#) (cit. on pp. 39, 41).
- [76] P. J. E. Peebles, *Tests of cosmological models constrained by inflation*, [Astrophys. J. 284 \(1984\) 439](#) (cit. on p. 41).
- [77] S. M. Carroll, *The Cosmological constant*, [Living Rev. Rel. 4 \(2001\) 1](#), arXiv: [astro-ph/0004075](#) (cit. on p. 41).
- [78] P. J. E. Peebles and B. Ratra, *The cosmological constant and dark energy*, [Rev. Mod. Phys. 75 \(2003\) 559](#), ed. by J.-P. Hsu and D. Fine, arXiv: [astro-ph/0207347](#) (cit. on p. 41).
- [79] S. P. Martin and D. G. Robertson, *Higgs boson mass in the Standard Model at two-loop order and beyond*, [Phys. Rev. D 90 \(2014\) 073010](#), arXiv: [1407.4336 \[hep-ph\]](#) (cit. on p. 41).
- [80] G. Altarelli, *The Higgs: so simple yet so unnatural*, [EPJ Web Conf. 71 \(2014\) 00005](#), ed. by L. Bravina et al. (cit. on p. 42).
- [81] S. P. Martin, *A Supersymmetry primer*, [Adv. Ser. Direct. High Energy Phys. 18 \(1998\) 1](#), ed. by G. L. Kane, arXiv: [hep-ph/9709356](#) (cit. on p. 42).
- [82] A. Takenaka et al., *Search for proton decay via $p \rightarrow e^+ \pi^0$ and $p \rightarrow \mu^+ \pi^0$ with an enlarged fiducial volume in Super-Kamiokande I-IV*, [Phys. Rev. D 102 \(2020\) 112011](#), arXiv: [2010.16098 \[hep-ex\]](#) (cit. on p. 42).

- [83] ATLAS Collaboration, *The quest to discover supersymmetry at the ATLAS experiment*, (2024), arXiv: [2403.02455 \[hep-ex\]](#) (cit. on pp. 42, 44).
- [84] V. Barger et al., *Complex Singlet Extension of the Standard Model*, [Phys. Rev. D **79** \(2009\) 015018](#), arXiv: [0811.0393 \[hep-ph\]](#) (cit. on p. 43).
- [85] I. Engeln et al., *The Dark Phases of the N2HDM*, [JHEP **08** \(2020\) 085](#), arXiv: [2004.05382 \[hep-ph\]](#) (cit. on p. 43).
- [86] A. Arhrib and M. Maniatis, *The two-real-singlet Dark Matter model*, [Phys. Lett. B **796** \(2019\) 15](#), arXiv: [1807.03554 \[hep-ph\]](#) (cit. on p. 43).
- [87] ATLAS Collaboration, *Combination of searches for resonant Higgs boson pair production using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS Detector*, [Phys. Rev. Lett. **132** \(2024\) 231801](#), arXiv: [2311.15956 \[hep-ex\]](#) (cit. on p. 45).
- [88] A. Hayrapetyan et al., *Searches for Higgs Boson Production through Decays of Heavy Resonances*, (2024), arXiv: [2403.16926 \[hep-ex\]](#) (cit. on p. 45).
- [89] *Summary of exotic Higgs boson decays from the ATLAS experiment*, tech. rep., CERN, 2021, URL: <https://cds.cern.ch/record/2758783> (cit. on p. 45).
- [90] CMS Collaboration, *Search for the exotic decay of the Higgs boson into two light pseudoscalars with four photons in the final state in proton-proton collisions at $\sqrt{s} = 13$ TeV*, [JHEP **07** \(2023\) 148](#), arXiv: [2208.01469 \[hep-ex\]](#) (cit. on p. 45).
- [91] ATLAS Collaboration, *Search for a resonance decaying into a scalar particle and a Higgs boson in the final state with two bottom quarks and two photons in proton-proton collisions at a center of mass energy of 13 TeV with the ATLAS detector*, (2024), arXiv: [2404.12915 \[hep-ex\]](#) (cit. on pp. 46, 47).
- [92] CMS Collaboration, *Search for a new resonance decaying into two spin-0 bosons in a final state with two photons and two bottom quarks in proton-proton collisions at $\sqrt{s} = 13$ TeV*, [JHEP **05** \(2024\) 316](#), arXiv: [2310.01643 \[hep-ex\]](#) (cit. on p. 46).
- [93] ATLAS Collaboration, *Search for a new heavy scalar particle decaying into a Higgs boson and a new scalar singlet in final states with one or two light leptons and a pair of τ -leptons with the ATLAS detector*, [JHEP **10** \(2023\) 009](#), arXiv: [2307.11120 \[hep-ex\]](#) (cit. on p. 46).
- [94] CMS Collaboration, *Search for a heavy Higgs boson decaying into two lighter Higgs bosons in the $\tau\tau b\bar{b}$ final state at 13 TeV*, [JHEP **11** \(2021\) 057](#), arXiv: [2106.10361 \[hep-ex\]](#) (cit. on p. 46).
- [95] *Search for the nonresonant and resonant production of a Higgs boson in association with an additional scalar boson in the $\gamma\gamma\tau\tau$ final state*, tech. rep., CERN, 2024, URL: <https://cds.cern.ch/record/2893031> (cit. on p. 46).
- [96] CMS Collaboration, *Search for a massive scalar resonance decaying to a light scalar and a Higgs boson in the four b quarks final state with boosted topology*, [Phys. Lett. B **842** \(2023\) 137392](#), arXiv: [2204.12413 \[hep-ex\]](#) (cit. on p. 46).
- [97] E. J. N. Wilson, *An introduction to particle accelerators*, 2001, ISBN: 978-0-19-852054-2 (cit. on p. 49).

- [98] G. Altarelli and G. Parisi, *Asymptotic freedom in parton language*, *Nucl. Phys. B* **126** (1977) 298 (cit. on p. 50).
- [99] R. D. Ball et al., *Parton distributions from high-precision collider data. Parton distributions from high-precision collider data*, *Eur. Phys. J. C* **77** (2017) 663, arXiv: 1706.00428, URL: <http://cds.cern.ch/record/2267455> (cit. on p. 51).
- [100] LHCb Collaboration, *The LHCb Detector at the LHC*, *JINST* **3** (2008) S08005, URL: <https://cds.cern.ch/record/1129809> (cit. on p. 51).
- [101] ALICE Collaboration, *The ALICE experiment at the CERN LHC*, *JINST* **3** (2008) S08002 (cit. on p. 51).
- [102] CMS Collaboration, *The CMS experiment at the CERN LHC. The Compact Muon Solenoid experiment*, *JINST* **3** (2008) S08004, URL: <https://cds.cern.ch/record/1129810> (cit. on p. 51).
- [103] A. Team, “Diagram of an LHC dipole magnet. Schéma d’un aimant dipôle du LHC”, 1999, URL: <https://cds.cern.ch/record/40524> (cit. on p. 52).
- [104] J.-L. Caron, “LHC quadrupole cross section.”, 1998, URL: <https://cds.cern.ch/record/841485> (cit. on p. 52).
- [105] E. Lopienska, *The CERN accelerator complex, layout in 2022. Complexe des accélérateurs du CERN en janvier 2022*, (2022), URL: <https://cds.cern.ch/record/2800984> (cit. on p. 53).
- [106] *9th LHC Operations Evian Workshop, Proceedings*, CERN, Geneva, 2019, URL: <http://cds.cern.ch/record/2706427> (cit. on p. 54).
- [107] R. Bruce et al., *Review of LHC Run 2 Machine Configurations*, (2019) 187, URL: <https://cds.cern.ch/record/2750415> (cit. on p. 54).
- [108] B. Salvachua, *Overview of Proton-Proton Physics during Run 2*, (2019) 7, URL: <https://cds.cern.ch/record/2750272> (cit. on p. 54).
- [109] ATLAS Collaboration, *Website of ATLAS Luminosity Public Results in Run 2*, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2> (cit. on pp. 55–57).
- [110] ATLAS Collaboration, *Website of ATLAS Luminosity Public Results in Run 3*, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun3> (cit. on pp. 54, 57).
- [111] M. Hostettler et al., *Operational β^* levelling at the LHC in 2022 and beyond*, *JACoW IPAC2023* (2023) MOPL045 (cit. on p. 57).
- [112] ATLAS Collaboration, *The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run 3*, *JINST* **19** (2024) P05063, arXiv: 2305.16623 [physics.ins-det] (cit. on p. 59).
- [113] ATLAS Collaboration, *ATLAS inner detector: Technical Design Report, 1*, Technical design report. ATLAS, Geneva: CERN, 1997, URL: <https://cds.cern.ch/record/331063> (cit. on p. 61).

- [114] ATLAS Collaboration, *The ATLAS Inner Detector commissioning and calibration*, *Eur. Phys. J. C* **70** (2010) 787, arXiv: 1004.5293 [physics.ins-det] (cit. on pp. 61, 63).
- [115] J. Pequena, “Computer generated image of the ATLAS inner detector”, 2008, URL: <https://cds.cern.ch/record/1095926> (cit. on p. 62).
- [116] N. Wermes and G. Hallewel, *ATLAS pixel detector: Technical Design Report*, Technical design report. ATLAS, Geneva: CERN, 1998, URL: <https://cds.cern.ch/record/381263> (cit. on p. 61).
- [117] ATLAS Collaboration, *Performance of ATLAS Pixel Detector and Track Reconstruction at the start of Run 3 in LHC Collisions at $\sqrt{s} = 900$ GeV*, tech. rep., CERN, 2022, URL: <https://cds.cern.ch/record/2814766> (cit. on p. 61).
- [118] M. Capeans et al., *ATLAS Insertable B-Layer Technical Design Report*, (2010) (cit. on p. 61).
- [119] ATLAS Collaboration, *Website of ATLAS Public Results on Tracking performance: Impact parameter resolution in Run 1 and Run 2*, URL: <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/IDTR-2015-007/> (cit. on p. 63).
- [120] J. N. Jackson, *The ATLAS semiconductor tracker (SCT)*, *Nucl. Instrum. Meth. A* **541** (2005) 89, ed. by T. Ohsugi et al. (cit. on p. 62).
- [121] A. Vogel, *ATLAS Transition Radiation Tracker (TRT): Straw Tube Gaseous Detectors at High Rates*, tech. rep., CERN, 2013, URL: <https://cds.cern.ch/record/1537991> (cit. on p. 63).
- [122] ATLAS Collaboration, *Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2*, *Eur. Phys. J. C* **77** (2017) 673, arXiv: 1704.07983 [hep-ex] (cit. on pp. 64, 81).
- [123] R. Wigmans, *Sampling calorimetry*, *Nucl. Instrum. Meth. A* **494** (2002) 277, ed. by G. V. Fedotovitch and B. Khazin (cit. on p. 64).
- [124] CMS Collaboration, *The CMS electromagnetic calorimeter project: Technical Design Report*, (1997) (cit. on p. 65).
- [125] C. W. Fabjan and F. Gianotti, *Calorimetry for particle physics*, *Rev. Mod. Phys.* **75** (2003) 1243 (cit. on p. 68).
- [126] ATLAS Collaboration, *ATLAS liquid argon calorimeter: Technical design report*, (1996) (cit. on pp. 69–71).
- [127] ATLAS Collaboration, *Readiness of the ATLAS Liquid Argon Calorimeter for LHC Collisions*, *Eur. Phys. J. C* **70** (2010) 723, arXiv: 0912.2642 [physics.ins-det] (cit. on p. 69).
- [128] M. Aleksa et al., *ATLAS Liquid Argon Calorimeter Phase-I Upgrade Technical Design Report*, (2013) (cit. on p. 69).
- [129] ATLAS Collaboration, *The Expected Performance of the ATLAS Inner Detector*, tech. rep., CERN, 2008, URL: <https://cds.cern.ch/record/1118445> (cit. on p. 70).

- [130] ATLAS Collaboration, *ATLAS tile calorimeter: Technical Design Report*, Technical design report. ATLAS, Geneva: CERN, 1996, URL: <https://cds.cern.ch/record/331062> (cit. on p. 70).
- [131] ATLAS Collaboration, *Readiness of the ATLAS Tile Calorimeter for LHC collisions*, *Eur. Phys. J. C* **70** (2010) 1193, arXiv: 1007.5423 [physics.ins-det] (cit. on p. 70).
- [132] ATLAS Collaboration, *Performance of the ATLAS Hadronic Endcap Calorimeter in beam tests: Selected results*, tech. rep., CERN, 2022, URL: <https://cds.cern.ch/record/2811731> (cit. on p. 70).
- [133] T. Barillari, *The ATLAS liquid argon hadronic end-cap calorimeter: Construction and selected beam test results*, *Nucl. Phys. B Proc. Suppl.* **150** (2006) 102, ed. by C. Bosio et al., arXiv: physics/0407026 (cit. on p. 70).
- [134] A. Artamonov et al., *The ATLAS Forward Calorimeter*, *JINST* **3** (2008) P02010, URL: <https://cds.cern.ch/record/1094547> (cit. on p. 71).
- [135] R. S. Orr, *The ATLAS Forward Calorimeter*, tech. rep., CERN, 2011, URL: <https://cds.cern.ch/record/1393395> (cit. on p. 71).
- [136] ATLAS Collaboration, *ATLAS muon spectrometer: Technical design report*, (1997) (cit. on p. 71).
- [137] ATLAS Collaboration, *ATLAS central solenoid: Technical design report*, (1997) (cit. on p. 72).
- [138] ATLAS Collaboration, *ATLAS barrel toroid: Technical design report*, (1997) (cit. on p. 72).
- [139] ATLAS Collaboration, *ATLAS end-cap toroids: Technical design report*, (1997) (cit. on p. 72).
- [140] ATLAS Collaboration, *ATLAS high-level trigger, data acquisition and controls: Technical design report*, (2003) (cit. on pp. 73, 74).
- [141] ATLAS Collaboration, *Performance of the ATLAS Trigger System in 2015*, *Eur. Phys. J. C* **77** (2017) 317, arXiv: 1611.09661 [hep-ex] (cit. on p. 73).
- [142] ATLAS Collaboration, *The ATLAS trigger system for LHC Run 3 and trigger performance in 2022*, *JINST* **19** (2024) P06029, arXiv: 2401.06630 [hep-ex] (cit. on p. 73).
- [143] ATLAS Collaboration, *ATLAS level-1 trigger: Technical Design Report*, (1998) (cit. on p. 73).
- [144] ATLAS Collaboration, *Operation of the ATLAS trigger system in Run 2*, *JINST* **15** (2020) P10004, arXiv: 2007.12539 [physics.ins-det] (cit. on p. 74).
- [145] ATLAS Collaboration, *Website of ATLAS Trigger Operation Public Results in Run 2*, URL: <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/TriggerOperationPublicResults> (cit. on pp. 74, 75).
- [146] S. Mehlhase, *ATLAS detector slice (and particle visualisations)*, (2021), URL: <https://cds.cern.ch/record/2770815> (cit. on p. 78).

- [147] ATLAS Collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data*, *JINST* **14** (2019) P12006, arXiv: 1908.00005 [hep-ex] (cit. on pp. 79, 80, 84, 109–111).
- [148] ATLAS Collaboration, *Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1*, *Eur. Phys. J. C* **77** (2017) 490, arXiv: 1603.02934, URL: <https://cds.cern.ch/record/2138166> (cit. on p. 80).
- [149] T. Cornelissen et al., *Concepts, Design and Implementation of the ATLAS New Tracking (NEWT)*, tech. rep., CERN, 2007, URL: <https://cds.cern.ch/record/1020106> (cit. on p. 81).
- [150] R. Frühwirth, *Application of Kalman filtering to track and vertex fitting*, *Nucl. Instrum. Meth. A* **262** (1987) 444 (cit. on p. 82).
- [151] W. Waltenberger et al., *Adaptive vertex fitting*, *Journal of Physics G: Nuclear and Particle Physics* **34** (2007) N343, URL: <https://dx.doi.org/10.1088/0954-3899/34/12/N01> (cit. on p. 82).
- [152] *Improved electron reconstruction in ATLAS using the Gaussian Sum Filter-based model for bremsstrahlung*, tech. rep., CERN, 2012, URL: <https://cds.cern.ch/record/1449796> (cit. on p. 83).
- [153] *Electron and photon reconstruction and performance in ATLAS using a dynamical, topological cell clustering-based approach*, tech. rep., CERN, 2017, URL: <https://cds.cern.ch/record/2298955> (cit. on p. 84).
- [154] M. Aharrouche et al., *Measurement of the response of the ATLAS liquid argon barrel calorimeter to electrons at the 2004 combined test-beam*, *Nucl. Instrum. Methods Phys. Res., A* **614** (2010) 400, URL: <https://cds.cern.ch/record/1273537> (cit. on p. 87).
- [155] M. Aleksa et al., *ATLAS Combined Testbeam: Computation and Validation of the Electronic Calibration Constants for the Electromagnetic Calorimeter*, tech. rep., CERN, 2006, URL: <https://cds.cern.ch/record/942528> (cit. on p. 87).
- [156] W. E. Cleland and E. G. Stern, *Signal processing considerations for liquid ionization calorimeters in a high rate environment*, *Nucl. Instrum. Meth. A* **338** (1994) 467 (cit. on p. 90).
- [157] D. Oliveira Damazio, *Signal Processing for the ATLAS Liquid Argon Calorimeter : studies and implementation*, tech. rep., CERN, 2013, URL: <https://cds.cern.ch/record/1630826> (cit. on p. 90).
- [158] ATLAS Collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data*, *JINST* **14** (2019) P12006, arXiv: 1908.00005, URL: <https://cds.cern.ch/record/2684552> (cit. on pp. 103, 104).
- [159] ATLAS Collaboration, *Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run 2 data collected in 2015 and 2016*, *Eur. Phys. J. C* **79** (2019) 205, arXiv: 1810.05087, URL: <https://cds.cern.ch/record/2643391> (cit. on p. 104).

- [160] ATLAS Collaboration, *Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run-1 data*, *Eur. Phys. J. C* **76** (2016) 666, 57 p, URL: <https://cds.cern.ch/record/2158117> (cit. on p. 105).
- [161] ATLAS Collaboration, *Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run-1 data*, *Eur. Phys. J. C* **76** (2016) 666, arXiv: 1606.01813 [hep-ex] (cit. on p. 106).
- [162] *Expected photon performance in the ATLAS experiment*, tech. rep., CERN, 2011, URL: <https://cds.cern.ch/record/1345329> (cit. on p. 106).
- [163] ATLAS Collaboration, *Electron and photon efficiencies in LHC Run 2 with the ATLAS experiment*, *JHEP* **05** (2024) 162, arXiv: 2308.13362 [hep-ex] (cit. on pp. 106, 111).
- [164] ATLAS Collaboration, *Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton-proton collision data at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **79** (2019) 639, arXiv: 1902.04655 [physics.ins-det] (cit. on p. 111).
- [165] ATLAS Collaboration, *Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **81** (2021) 578, arXiv: 2012.00578 [hep-ex] (cit. on p. 112).
- [166] G. Altarelli, *Collider Physics within the Standard Model: a Primer*, tech. rep., 2013, arXiv: 1303.2842, URL: <https://cds.cern.ch/record/1527372> (cit. on p. 113).
- [167] ATLAS Collaboration, *Jet reconstruction and performance using particle flow with the ATLAS Detector*, *Eur. Phys. J. C* **77** (2017) 466, arXiv: 1703.10485 [hep-ex] (cit. on p. 114).
- [168] ATLAS Collaboration, *Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **81** (2021) 689, arXiv: 2007.02645 [hep-ex] (cit. on pp. 114, 116–118).
- [169] M. Cacciari et al., *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063, arXiv: 0802.1189 [hep-ph] (cit. on p. 115).
- [170] ATLAS Collaboration, *Comparison of Monte Carlo generator predictions for bottom and charm hadrons in the decays of top quarks and the fragmentation of high p_T jets*, tech. rep., CERN, 2014, URL: <https://cds.cern.ch/record/1709132> (cit. on p. 120).
- [171] ATLAS Collaboration, *ATLAS flavour-tagging algorithms for the LHC Run 2 pp collision dataset*, *Eur. Phys. J. C* **83** (2023) 681, arXiv: 2211.16345 [physics.data-an] (cit. on p. 120).
- [172] ATLAS Collaboration, *ATLAS b -jet identification performance and efficiency measurement with $t\bar{t}$ events in pp collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **79** (2019) 970, arXiv: 1907.05120 [hep-ex] (cit. on pp. 122, 123, 129).
- [173] 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*, *Eur. Phys. J. C* **82** (2022) 95, arXiv: 2109.10627 [hep-ex] (cit. on p. 123).

- [174] 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 $\sqrt{s} = 13\text{ TeV}$* , *Eur. Phys. J. C* **83** (2023) 728, arXiv: 2301.06319 [hep-ex] (cit. on pp. 123, 129).
- [175] ATLAS Collaboration, *Evidence for the $H \rightarrow b\bar{b}$ decay with the ATLAS detector*, *JHEP* **12** (2017) 024, arXiv: 1708.03299 [hep-ex] (cit. on p. 133).
- [176] ATLAS Collaboration, *Performance of missing transverse momentum reconstruction with the ATLAS detector using proton–proton collisions at $\sqrt{s} = 13\text{ TeV}$* , *Eur. Phys. J. C* **78** (2018) 903, arXiv: 1802.08168 [hep-ex] (cit. on p. 133).
- [177] ATLAS Collaboration, *Measurement of the production cross section of pairs of isolated photons in pp collisions at 13 TeV with the ATLAS detector*, *JHEP* **11** (2021) 169, arXiv: 2107.09330 [hep-ex] (cit. on p. 136).
- [178] ATLAS Collaboration, *Measurement of the cross section for isolated-photon plus jet production in pp collisions at $\sqrt{s} = 13\text{ TeV}$ using the ATLAS detector*, *Phys. Lett. B* **780** (2018) 578, arXiv: 1801.00112 [hep-ex] (cit. on p. 136).
- [179] ATLAS Collaboration, *Measurement of inclusive jet and dijet cross-sections in proton–proton collisions at $\sqrt{s} = 13\text{ TeV}$ with the ATLAS detector*, *JHEP* **05** (2018) 195, arXiv: 1711.02692 [hep-ex] (cit. on p. 136).
- [180] ATLAS Collaboration, *ATLAS data quality operations and performance for 2015–2018 data-taking*, *JINST* **15** (2020) P04003, arXiv: 1911.04632 [physics.ins-det] (cit. on p. 137).
- [181] ATLAS Collaboration, *Performance of electron and photon triggers in ATLAS during LHC Run 2*, *Eur. Phys. J. C* **80** (2020) 47, arXiv: 1909.00761 [hep-ex] (cit. on p. 137).
- [182] S. Höche, “Introduction to parton-shower event generators”, *Theoretical Advanced Study Institute in Elementary Particle Physics: Journeys Through the Precision Frontier: Amplitudes for Colliders*, 2015 235, arXiv: 1411.4085 [hep-ph] (cit. on p. 138).
- [183] B. Andersson et al., *Parton Fragmentation and String Dynamics*, *Phys. Rept.* **97** (1983) 31 (cit. on p. 138).
- [184] B. R. Webber, *A QCD Model for Jet Fragmentation Including Soft Gluon Interference*, *Nucl. Phys. B* **238** (1984) 492 (cit. on p. 138).
- [185] S. Agostinelli et al., *GEANT4—a simulation toolkit*, *Nucl. Instrum. Meth. A* **506** (2003) 250 (cit. on p. 138).
- [186] W. Lukas, *Fast Simulation for ATLAS: AtIfast-II and ISF*, *J. Phys. Conf. Ser.* **396** (2012) 022031, ed. by M. Ernst et al. (cit. on p. 138).
- [187] ATLAS Collaboration, *The simulation principle and performance of the ATLAS fast calorimeter simulation FastCaloSim*, tech. rep., CERN, 2010, URL: <https://cds.cern.ch/record/1300517> (cit. on p. 138).
- [188] ATLAS Collaboration, *AtIfast3: The Next Generation of Fast Simulation in ATLAS*, *Comput. Softw. Big Sci.* **6** (2022) 7, arXiv: 2109.02551 [hep-ex] (cit. on p. 138).

- [189] P. Nason, *A New method for combining NLO QCD with shower Monte Carlo algorithms*, [JHEP 11 \(2004\) 040](#), arXiv: [hep-ph/0409146](#) (cit. on pp. 138, 146, 231).
- [190] S. Frixione et al., *Matching NLO QCD computations with Parton Shower simulations: the POWHEG method*, [JHEP 11 \(2007\) 070](#), arXiv: [0709.2092 \[hep-ph\]](#) (cit. on pp. 138, 146, 231).
- [191] S. Alioli et al., *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*, [JHEP 06 \(2010\) 043](#), arXiv: [1002.2581 \[hep-ph\]](#) (cit. on pp. 138, 146, 231).
- [192] G. Heinrich et al., *NLO predictions for Higgs boson pair production with full top quark mass dependence matched to parton showers*, [JHEP 08 \(2017\) 088](#), arXiv: [1703.09252 \[hep-ph\]](#) (cit. on p. 138).
- [193] G. Heinrich et al., *Probing the trilinear Higgs boson coupling in di-Higgs production at NLO QCD including parton shower effects*, [JHEP 06 \(2019\) 066](#), arXiv: [1903.08137 \[hep-ph\]](#) (cit. on p. 138).
- [194] J. Butterworth et al., *PDF4LHC recommendations for LHC Run II*, [J. Phys. G 43 \(2016\) 023001](#), arXiv: [1510.03865 \[hep-ph\]](#) (cit. on p. 138).
- [195] S. Borowka et al., *Higgs Boson Pair Production in Gluon Fusion at Next-to-Leading Order with Full Top-Quark Mass Dependence*, [Phys. Rev. Lett. 117 \(2016\) 012001](#), arXiv: [1604.06447 \[hep-ph\]](#) (cit. on p. 138).
- [196] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, [Comput. Phys. Commun. 191 \(2015\) 159](#), arXiv: [1410.3012 \[hep-ph\]](#) (cit. on p. 139).
- [197] M. Bahr et al., *Herwig++ Physics and Manual*, [Eur. Phys. J. C 58 \(2008\) 639](#), arXiv: [0803.0883 \[hep-ph\]](#) (cit. on p. 139).
- [198] J. Bellm et al., *Herwig 7.0/Herwig++ 3.0 release note*, [Eur. Phys. J. C 76 \(2016\) 196](#), arXiv: [1512.01178 \[hep-ph\]](#) (cit. on p. 139).
- [199] R. D. Ball et al., *Parton distributions with LHC data*, [Nucl. Phys. B 867 \(2013\) 244](#), arXiv: [1207.1303 \[hep-ph\]](#) (cit. on p. 139).
- [200] R. D. Ball et al., *Parton distributions for the LHC Run II*, [JHEP 04 \(2015\) 040](#), arXiv: [1410.8849 \[hep-ph\]](#) (cit. on p. 139).
- [201] ATLAS Collaboration, *Validation of signal Monte Carlo event generation in searches for Higgs boson pairs with the ATLAS detector*, tech. rep., CERN, 2019, URL: <https://cds.cern.ch/record/2665057> (cit. on p. 141).
- [202] ATLAS Collaboration, *Combined measurements of Higgs boson production and decay using up to 80 fb⁻¹ of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment*, [Phys. Rev. D 101 \(2020\) 012002](#), arXiv: [1909.02845 \[hep-ex\]](#) (cit. on p. 144).
- [203] E. Bothmann et al., *Event Generation with Sherpa 2.2*, [SciPost Phys. 7 \(2019\) 034](#), arXiv: [1905.09127 \[hep-ph\]](#) (cit. on pp. 144, 146, 186).
- [204] K. Hamilton et al., *MINLO: multi-scale improved NLO*, [JHEP 10 \(2012\) 155](#), arXiv: [1206.3572 \[hep-ph\]](#) (cit. on pp. 146, 231).

- [205] K. Hamilton et al., *Finite quark-mass effects in the NNLOPS POWHEG+MiNLO Higgs generator*, *JHEP* **05** (2015) 140, arXiv: [1501.04637 \[hep-ph\]](#) (cit. on pp. 146, 231).
- [206] P. Nason and C. Oleari, *NLO Higgs boson production via vector-boson fusion matched with shower in POWHEG*, *JHEP* **02** (2010) 037, arXiv: [0911.5299 \[hep-ph\]](#) (cit. on p. 146).
- [207] G. Luisoni et al., *$HW^\pm/HZ + 0$ and 1 jet at NLO with the POWHEG BOX interfaced to GoSam and their merging within MiNLO*, *JHEP* **10** (2013) 083, arXiv: [1306.2542 \[hep-ph\]](#) (cit. on p. 146).
- [208] H. B. Hartanto et al., *Higgs boson production in association with top quarks in the POWHEG BOX*, *Phys. Rev. D* **91** (2015) 094003, arXiv: [1501.04498 \[hep-ph\]](#) (cit. on p. 146).
- [209] B. Jager et al., *Higgs boson production in association with b-jets in the POWHEG BOX*, *Phys. Rev. D* **93** (2016) 014030, arXiv: [1509.05843 \[hep-ph\]](#) (cit. on p. 146).
- [210] ATLAS Collaboration, *Measurement of Higgs boson production in the diphoton decay channel in pp collisions at center-of-mass energies of 7 and 8 TeV with the ATLAS detector*, *Phys. Rev. D* **90** (2014) 112015, arXiv: [1408.7084 \[hep-ex\]](#) (cit. on pp. 147, 158).
- [211] ATLAS Collaboration, *Measurement of the properties of Higgs boson production at $\sqrt{s} = 13$ TeV in the $H \rightarrow \gamma\gamma$ channel using 139 fb^{-1} of pp collision data with the ATLAS experiment*, *JHEP* **07** (2023) 088, arXiv: [2207.00348 \[hep-ex\]](#) (cit. on p. 147).
- [212] ATLAS Collaboration, *Tagging and suppression of pileup jets with the ATLAS detector*, tech. rep., CERN, 2014, URL: <https://cds.cern.ch/record/1700870> (cit. on p. 147).
- [213] J. Snoek et al., “Practical Bayesian Optimization of Machine Learning Algorithms”, *Advances in Neural Information Processing Systems*, ed. by F. Pereira et al., vol. 25, Curran Associates, Inc., 2012, URL: https://proceedings.neurips.cc/paper_files/paper/2012/file/05311655a15b75fab86956663e1819cd-Paper.pdf (cit. on p. 151).
- [214] ATLAS Collaboration, *Measurement of event shapes at large momentum transfer with the ATLAS detector in pp collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J. C* **72** (2012) 2211, arXiv: [1206.2135 \[hep-ex\]](#) (cit. on p. 152).
- [215] L. G. Almeida et al., *Substructure of high- p_T Jets at the LHC*, *Phys. Rev. D* **79** (2009) 074017, arXiv: [0807.0234 \[hep-ph\]](#) (cit. on p. 152).
- [216] M. Oreglia, *A Study of the Reactions $\psi' \rightarrow \gamma\gamma\psi$* , Other thesis, 1980 (cit. on p. 158).
- [217] 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** (2018) 052005, arXiv: [1802.04146 \[hep-ex\]](#) (cit. on p. 158).

- [218] ATLAS Collaboration, *Measurement of isolated-photon pair production in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *JHEP* **01** (2013) 086, arXiv: 1211.1913 [hep-ex] (cit. on p. 160).
- [219] *Recommendations for the Modeling of Smooth Backgrounds*, tech. rep., CERN, 2020, URL: <https://cds.cern.ch/record/2743717> (cit. on p. 162).
- [220] G. Avoni et al., *The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS*, *JINST* **13** (2018) P07017, URL: <https://cds.cern.ch/record/2633501> (cit. on p. 163).
- [221] *Evaluating statistical uncertainties and correlations using the bootstrap method*, tech. rep., CERN, 2021, URL: <https://cds.cern.ch/record/2759945> (cit. on pp. 165, 176).
- [222] LHC Higgs Cross Section Working Group, *Website of the LHC Higgs Working Group*, URL: <https://twiki.cern.ch/twiki/bin/view/LHCPhysics/HiggsXSBR> (cit. on p. 168).
- [223] ATLAS Collaboration, *Measurements of the Higgs boson inclusive and differential fiducial cross-sections in the diphoton decay channel with pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **08** (2022) 027, arXiv: 2202.00487 [hep-ex] (cit. on pp. 168, 232).
- [224] ATLAS Collaboration, *Measurements of the Higgs boson inclusive and differential fiducial cross sections in the 4ℓ decay channel at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **80** (2020) 942, arXiv: 2004.03969 [hep-ex] (cit. on pp. 168, 232).
- [225] 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*, *JHEP* **04** (2019) 046, arXiv: 1811.12113 [hep-ex] (cit. on pp. 169, 232).
- [226] ATLAS Collaboration, *Study of heavy-flavor quarks produced in association with top-quark pairs at $\sqrt{s} = 7$ TeV using the ATLAS detector*, *Phys. Rev. D* **89** (2014) 072012, arXiv: 1304.6386 [hep-ex] (cit. on pp. 169, 232).
- [227] G. Aad et al., *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*, *JHEP* **06** (2013) 084, arXiv: 1302.2929 [hep-ex] (cit. on pp. 169, 232).
- [228] A. L. Read, *Presentation of search results: The CL_s technique*, *J. Phys. G* **28** (2002) 2693, ed. by M. R. Whalley and L. Lyons (cit. on pp. 171, 211).
- [229] 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*, *JHEP* **08** (2024) 164, arXiv: 2405.20040 [hep-ex] (cit. on p. 179).
- [230] ATLAS Collaboration, *Search for non-resonant Higgs boson pair production in the $2b + 2\ell + E_T^{\text{miss}}$ final state in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **02** (2024) 037, arXiv: 2310.11286 [hep-ex] (cit. on p. 179).
- [231] U. Ellwanger et al., *The Next-to-Minimal Supersymmetric Standard Model*, *Phys. Rept.* **496** (2010) 1, arXiv: 0910.1785 [hep-ph] (cit. on p. 185).

- [232] P. Baldi et al., *Parameterized neural networks for high-energy physics*, *Eur. Phys. J. C* **76** (2016) 235, arXiv: 1601.07913 [hep-ex] (cit. on pp. 186, 194).
- [233] D. Berdine et al., *Breakdown of the Narrow Width Approximation for New Physics*, *Phys. Rev. Lett.* **99** (2007) 111601, arXiv: hep-ph/0703058 (cit. on p. 187).
- [234] A. D. Bukin, *Fitting function for asymmetric peaks*, 2007, arXiv: 0711.4449 [physics.data-an], URL: <https://arxiv.org/abs/0711.4449> (cit. on p. 188).
- [235] B. Delaunay, *Sur la sphère vide*, French, Bulletin de l'Académie des Sciences de l'URSS. Classe des sciences mathématiques et na **1934** (1934) 793 (cit. on p. 188).
- [236] K. Cranmer et al., *HistFactory: A tool for creating statistical models for use with RooFit and RooStats*, tech. rep., New York U., 2012, URL: <https://cds.cern.ch/record/1456844> (cit. on p. 210).
- [237] M. Pinamonti et al., *TRExFitter package*, URL: <https://github.com/liboyang0112/TRExFitter/tree/master> (cit. on p. 210).
- [238] CMS Collaboration, *Search for a new resonance decaying into two spin-0 bosons in a final state with two photons and two bottom quarks in proton-proton collisions at $\sqrt{s} = 13$ TeV*, *JHEP* **05** (2024) 316, arXiv: 2310.01643 [hep-ex] (cit. on p. 212).
- [239] E. Gross and O. Vitells, *Trial factors for the look elsewhere effect in high energy physics*, *Eur. Phys. J. C* **70** (2010) 525, arXiv: 1005.1891 [physics.data-an] (cit. on p. 215).
- [240] O. Vitells, *Estimating the "look elsewhere effect" when searching for a signal*, (2011) 183, URL: <https://cds.cern.ch/record/2203251> (cit. on p. 215).
- [241] K. S. Brown, "Euler Characteristics", *Cohomology of Groups*, New York, NY: Springer New York, 1982 230, ISBN: 978-1-4684-9327-6, URL: https://doi.org/10.1007/978-1-4684-9327-6_10 (cit. on p. 215).
- [242] ATLAS Collaboration, *Updated projection of the sensitivity of searches for Higgs boson pair production in the $b\bar{b}\tau^+\tau^-$ final state from LHC Run 2 to the High Luminosity LHC with the ATLAS detector*, tech. rep., CERN, 2024, URL: <https://cds.cern.ch/record/2910850> (cit. on p. 232).
- [243] ATLAS Collaboration, *HL-LHC prospects for the measurement of Higgs boson pair production in the $b\bar{b}b\bar{b}$ final state and combination with the $b\bar{b}\gamma\gamma$ and $b\bar{b}\tau^+\tau^-$ final states at the ATLAS experiment*, tech. rep., CERN, 2022, URL: <https://cds.cern.ch/record/2841244> (cit. on p. 232).
- [244] N. Deutschmann et al., *Top-Yukawa contributions to bbH production at the LHC*, *JHEP* **1907** (2019) 054, arXiv: 1808.01660, URL: <http://cds.cern.ch/record/2633583> (cit. on pp. 232, 233).

Acknowledgements

There are a lot of people I want to thank as I arrive to the end of my Ph.D.

First and foremost, I am extremely grateful to my supervisors: Leonardo, Ruggero, and Stefano. Your guidance, encouragement, and endless patience have been invaluable, and I feel very lucky to have worked with you. Thank you for everything that you have taught me and for all the opportunities that you gave me. Your door was always open, and this thesis would not have even started without you.

I also want to extend my thanks to the entire Milano ATLAS group. I am grateful to all of you for creating such a warm and welcoming environment, and for making me feel at home from the start.

The work presented in this thesis would not have been possible without the contributions from the HH, e/γ , and flavor tagging communities. I am very grateful to the conveners, sub-conveners, and analysis contacts for their direction and support. I would also like to thank my fellow Ph.D. students, who shared this journey with me. A special thanks to Zihang and Alkaid: it was a pleasure to work with you, and I have many happy memories of our time together.

I would like to thank the members of the Liquid Argon team in Milano and at CERN. Working on the LAr upgrade and being part of the operations team was such a fun and rewarding experience. A special thanks to Francesco for his dedication, kindness, and for everything that he taught me.

A special thought should go to the friends that I made during my time in Milano and at CERN. Laura N., Sofia, Federica, Chiara L., Chiara M., Laura P., Davide, Stefano, Claudia, Alberto, and so many others: thank you for making me feel welcome in a new home. I am grateful for all the happy moments that we shared, and I hope that we can have many more.

Finally, I would like to thank my family and friends in Italy. Your constant support and encouragement have meant a lot to me, and I could not have completed this journey without you.