

Search for the production of a pair of highly Lorentz-boosted Higgs bosons via vector-boson fusion in final states with four bottom quarks with the ATLAS experiment

Dissertation
zur Erlangung des Doktorgrades
an der Fakultät für Mathematik, Informatik und Naturwissenschaften
Fachbereich Physik
der Universität Hamburg

vorgelegt von
Marcus Vinicius Gonzalez Rodrigues

Hamburg
2025

Gutachter/innen der Dissertation:

Prof. Dr. Elisabetta Gallo
Dr. Janna Katharina Behr

Zusammensetzung der Prüfungskommission:

Prof. Dr. Peter Schleper
Dr. Janna Katharina Behr
Prof. Dr. Arwen Pearson
Prof. Dr. Elisabetta Gallo
Dr. Jürgen Reuter

Vorsitzende/r der Prüfungskommission:

Prof. Dr. Arwen Pearson

Datum der Disputation:

25.07.2025

Vorsitzender des Fach-Promotionsausschusses PHYSIK:

Prof. Dr. Wolfgang J. Parak

Leiter des Fachbereichs PHYSIK:

Prof. Dr. Markus Drescher

Dekan der Fakultät MIN:

Prof. Dr.-Ing. Norbert Ritter

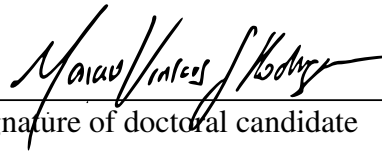
Declaration on oath

I hereby declare and affirm that this doctoral dissertation is my own work and that I have not used any aids and sources other than those indicated.

If electronic resources based on generative artificial intelligence (gAI) were used in the course of writing this dissertation, I confirm that my own work was the main and value-adding contribution and that complete documentation of all resources used is available in accordance with good scientific practice. I am responsible for any erroneous or distorted content, incorrect references, violations of data protection and copyright law or plagiarism that may have been generated by the gAI.

05.05.2025

Date



Signature of doctoral candidate

Acknowledgements

Working on this doctoral thesis at the DESY-ATLAS group has been for sure an incredible and rewarding experience. A great mixture of learning many new things about physics, having my first experiences in speaking daily in english and german, and adapting to a completely different culture.

To start with, I would like to thank my supervisor Katharina Behr for all the valuable learnings in the past years. You have always been present, always had an answer to whatever question about physics I had, and have been the whole time precise in knowing how to instigate me to pursue new ideas and techniques. I'm extremely glad to have had such a fantastic person to supervise me along this journey.

The search presented in this thesis was conducted in collaboration with many other researches. I would like to give a special thanks to Kunlin Run for all the help when I was still learning everything we need to know in an analysis. I also would like to thank Sebastien Rettie and Rachel Hyneman. I learned a lot from both of you and I hope we keep working together in future analyses.

Diogo Boito, I thanked you in my master's thesis, and now I want to thank you here again. Even though I'm no longer your student, you still care about my next-steps, believe in me, and remind me of how physics is cool. For me, you are a model of how any professor and supervisor should be.

I also would like to thank Peter Schleper, Arwen Pearson, Elisabetta Gallo, and Jürgen Reuter for accepting to be part of my PhD committee.

Life in Hamburg would be impossible without the close friends I made here. Henrique, Debora, Otavio, Lauren, Gilson, Arilene, Clara, Gustavo, Nany, and Larissa: you will always have a place in my heart. Dany, thanks for all these years.

Ich möchte auch meinem Deutschlehrer danken. Bruno, du hast keine Ahnung, wie sehr du mir mein Leben in Deutschland erleichtert hast.

Para finalizar, é claro que preciso agradecer à pessoa mais importante da minha vida. Mãe, espero que você saiba que nada disso seria possível sem o seu suporte, mesmo tendo um oceano inteiro nos separando. Obrigado.

Abstract

A search for the production of a pair of Higgs bosons (HH) via vector-boson fusion (VBF) in the Lorentz-boosted kinematic regime with subsequent decay of the Higgs bosons into four bottom quarks is presented. The analysis is based on 140 fb^{-1} of proton-proton collision data collected with the ATLAS detector at the Large Hadron Collider (LHC), at a center-of-mass energy of $\sqrt{s} = 13 \text{ TeV}$.

The results were interpreted in terms of both non-resonant and resonant production of Higgs-boson pairs. Non-resonant VBF HH production is predicted by the Standard Model (SM); the production cross-section is, however, significantly enhanced for anomalous, non-SM values for the trilinear Higgs-boson self-coupling and the quartic coupling between two Higgs bosons and two vector bosons, parametrised as κ_λ and κ_{2V} , respectively. Resonant VBF HH production is predicted only by extensions of the SM. In this thesis, the resonant HH production via the decay of a new heavy scalar particle originating from the vector-boson scattering was considered.

No significant deviation of the data from the expected background was observed. Hence, in the non-resonant search, limits on the κ_{2V} and κ_λ couplings were set at 95% confidence level. In the resonant search, upper limits on the signal cross-section were derived for intermediate masses ranging from 1.0 TeV to 5.0 TeV under the assumptions of a narrow- and broad-resonance widths. In the model-dependent broad-width scenario, a composite Higgs benchmark has been considered.

In addition to these results, which have been published in 2024, the relevant details of a successive on-going ATLAS analysis using additionally data collected in 2022 and 2023 at $\sqrt{s} = 13.6 \text{ TeV}$ and using more modern tools to identify the $H \rightarrow b\bar{b}$ decays is discussed here.

Finally, a proof-of-principle study using a transformer-based machine-learning algorithm to discriminate signal from background events in triple Higgs-boson searches is also shown in this thesis. The even rarer production of three Higgs bosons has been gaining significant interest since it provides unique access to the quartic Higgs-boson self-coupling.

Zusammenfassung

In dieser Arbeit wird eine Suche nach der Produktion eines Higgs-Bosonen-Paares (HH) durch Vektor-Bosonen-Fusion (VBF) im kinematisch Bereich mit hohen Lorentz-Boosts der Higgs-Bosonen und deren anschließendem Zerfall in vier Bottom-Quarks vorgestellt. Die Suche basiert auf 140 fb^{-1} Proton-Proton-Kollisionsdaten, die mit dem ATLAS-Detektor am Large Hadron Collider (LHC) bei einer Schwerpunktsenergie von $\sqrt{s} = 13 \text{ TeV}$ gesammelt wurden.

Die Ergebnisse der Suche wurden sowohl im Rahmen der nicht-resonanten als auch der resonanten Produktion von Higgs-Bosonen-Paaren interpretiert. Die nicht-resonante VBF HH -Produktion wird durch das Standardmodell (SM) vorhergesagt; der Produktionsquerschnitt ist jedoch bei anomalen Werten der trilineare Higgs-Bosonen-Selbstkopplung und der quartischen Kopplung zwischen zwei Higgs-Bosonen und zwei Vektorbosonen, parametrisiert als κ_λ bzw. κ_{2V} , deutlich erhöht. Die resonante VBF HH -Produktion wird ausschließlich durch Erweiterungen des SM vorhergesagt. In dieser Arbeit wurde die resonante HH -Produktion über den Zerfall eines neuen schweren Skalarteilchens, das aus der Vektor-Bosonen-Streuung stammt, untersucht.

Es wurde keine signifikante Abweichung der Daten vom erwarteten Hintergrund beobachtet. Daher wurden in der nicht-resonanten Suche Grenzwerte auf die Kopplungen κ_{2V} und κ_λ mit einem Konfidenzniveau von 95% festgelegt. In der resonanten Suche wurden obere Grenzen für den Wirkungsquerschnitt für Resonanzmassen im Bereich von 1.0 TeV bis 5.0 TeV unter der Annahme sowohl einer schmalen als auch einer breiten Resonanzbreite abgeleitet. Im modellabhängigen Szenario mit breiter Resonanzbreite wurde ein repräsentatives Modell mit einem Composite Higgs-Boson zugrunde gelegt.

Zusätzlich zu diesen Ergebnissen, die 2024 veröffentlicht wurden, werden hier die relevanten Details einer laufenden ATLAS-Analyse vorgestellt, die zusätzliche Daten aus den Jahren 2022 und 2023 bei $\sqrt{s} = 13.6 \text{ TeV}$ verwendet und modernere Algorithmen zur Identifikation von $H \rightarrow b\bar{b}$ -Zerfälle einsetzt.

Schließlich wird in dieser Arbeit auch eine Proof-of-Principle-Studie vorgestellt, bei der ein auf Transformatoren basierender Algorithmus für maschinelles Lernen zur Unterscheidung von Signal- und Hintergrundereignissen in der Suche nach simultaner Produktion von drei Higgs-Bosonen eingesetzt wird. Die noch seltenere Produktion von drei Higgs-Bosonen hat

großes Interesse geweckt, da sie einen einzigartigen Zugang zur quartischen Higgs-Bosonen-Selbstkopplung ermöglicht.

Contents

1	Introduction	1
2	Theoretical Foundations and Motivation	5
2.1	The Standard Model of Particle Physics	5
2.1.1	Quantum Chromodynamics	6
2.1.2	Electroweak Theory	8
2.1.3	Spontaneous Symmetry Breaking and the Higgs boson	11
2.1.4	The Higgs-boson self-couplings	14
2.2	Motivations for Beyond Standard Model Theories	17
2.2.1	Two-Higgs-doublet models (2HDMs)	18
2.2.2	Composite Higgs models	21
2.3	Production and Decay of Higgs-boson pairs	22
2.3.1	Calculation of production cross-sections	22
2.3.2	Overview of Higgs-boson pair production modes	24
2.3.3	VBF non-resonant production of Higgs-boson pairs	25
2.3.4	VBF resonant production of Higgs-boson pairs	26
2.3.5	Decay modes of the Higgs bosons	27
3	Experimental Setup and Simulation of Collision Events	29
3.1	The Large Hadron Collider	29
3.1.1	Luminosity	31
3.1.2	Pile-up	31
3.2	The ATLAS detector	33
3.2.1	General design	33
3.2.2	The Inner Detector	36
3.2.3	The Electromagnetic and Hadronic Calorimeters	37
3.2.4	The Muon Spectrometer	40
3.2.5	Trigger system	41
3.3	Dataset and simulation of collision events	42
3.3.1	Simulation of physics processes	42
3.3.2	Non-resonant production of $HH \rightarrow b\bar{b}b\bar{b}$	44

3.3.3	Resonant production of $HH \rightarrow b\bar{b}b\bar{b}$	47
3.3.4	QCD multijets and $t\bar{t}$ background	47
4	Object Definition and Event Reconstruction	49
4.1	Tracks	49
4.1.1	Track reconstruction	50
4.1.2	Vertex reconstruction	54
4.2	Hadronic jets	55
4.2.1	Jet clustering algorithm	55
4.2.2	Small- R jets	58
4.2.3	Large- R jets	60
4.2.4	Jet calibration	61
4.2.5	Single and double b -tagging algorithms	63
4.3	Leptons and Photons	67
4.3.1	Electrons and Photons	67
4.3.2	Muons	68
4.4	Missing transverse energy	69
5	Analysis Strategy	71
5.1	Event selection and reconstruction	72
5.1.1	Triggers	73
5.1.2	Jet cleaning	74
5.1.3	Higgs-bosons candidates	74
5.1.4	VBF signature	74
5.1.5	Resolved non-resonant analysis veto	75
5.1.6	Kinematic regions	76
5.1.7	Event selection performance	78
5.2	Data-driven background estimation	80
5.3	Multivariate analysis	82
5.4	Systematic uncertainties	89
5.4.1	General	91
5.4.2	Jet uncertainties	91
5.4.3	Double b -tagger	91
5.4.4	Background modelling	93
5.4.5	Signal modelling	94
6	Results and Statistical Interpretation	97
6.1	Unblinded distributions	97
6.2	Statistical analysis	98
6.3	Non-resonant interpretation	105
6.4	Resonant interpretation	112
6.5	Discussion	118

7	Boosted VBF $HH \rightarrow b\bar{b}b\bar{b}$ search with Run-3 data	121
7.1	The next-generation of the double b -tagger	122
7.2	UFO jets	124
7.3	Event selection and analysis regions	124
7.4	Background estimation	128
7.5	Multivariate analysis	128
7.6	Preliminary expected results	130
8	Transformers for triple-Higgs searches	133
8.1	HHH production in the SM	134
8.2	The ATLAS $HHH \rightarrow b\bar{b}b\bar{b}b\bar{b}$ search	134
8.3	The transformer-based architecture	136
9	Conclusions	139
A	Data-driven background modelling via reweighting	145
B	Naming conventions for the nuisance parameters	147
	Bibliography	149
	List of Figures	163
	List of Tables	173

Introduction

The Standard Model (SM) of Particle Physics [1–3] is to date one of the most successful theories in the history of science. It describes the interactions between elementary particles within the quantum field theory framework and was heavily influenced by the earlier works of Feynman [4], Schwinger [5], and Tomonaga [6] in the context of quantum electrodynamics. Since its consolidation in the 1970s [7–10], the SM has successfully foretold the existence of new particles and predicted a series of measurements with astonishing precision. Prominent examples are the discovery of the weak vector bosons Z and W^\pm in 1983 at CERN’s Super Proton Synchrotron [11, 12], and the top quark, observed for the very first time in 1995 at the Tevatron collider [13, 14]. The anomalous magnetic moment of the electron, today known at the impressive precision of 0.20 parts per billion [15], is considered one of the most precise determinations of a physical quantity in the history of physics.

One of the greatest triumphs of the SM came in 2012 after the ATLAS and CMS experiments of the Large Hadron Collider (LHC) observed the existence of a new scalar particle [16, 17] compatible with the Higgs boson H . Predicted 50 years prior to its discovery [18–20], the Higgs boson is a core element in the electroweak sector of the SM to spontaneously break the vacuum symmetry and give rise to the masses of the weak vector bosons Z and W^\pm . In fact, any massive particle in the SM acquires its mass through its interaction with the Higgs field.

Nevertheless, the SM has not been able to explain various phenomena. To cite a few, in the SM, neutrinos are described as massless particles, although neutrino oscillation experiments indicate they have mass [21]; it does not provide an answer for the baryon asymmetry, the observed dominance of matter over antimatter in the universe [22]; and it does not include a description of the gravitational force. In this context, many extensions of the SM, referred to as beyond SM (BSM) theories, aim to address these unanswered questions. Their validity can be tested from direct searches for the existence of new particles predicted by these models and/or indirectly through precision measurements of SM physical quantities, such as particle production and decay rates, which are expected to deviate from their SM values in the presence of BSM effects.

The fact that the Higgs boson is the only fundamental scalar observed so far and the most recently discovered elementary particle has made the Higgs sector a natural focus point for several BSM scenarios. Interesting examples are the inclusion of weakly interacting massive particles to account for the dark matter in the universe that couple to the Higgs sector [23]; the extension of the Higgs sector with additional Higgs fields, as motivated by several BSM theories [24, 25], that leads to additional Higgs bosons and a more complex structure for the Higgs potential [26]; and others. Measuring all Higgs-boson properties is therefore crucial for obtaining a complete picture of the Higgs sector and either exclude or provide hints on BSM phenomena.

Although several parameters of the Higgs mechanism are nowadays measured with impressive precision in single-Higgs searches [27, 28], some are still not well constrained due to the statistical limitations of the much rarer processes involving the production of two (or even three) Higgs bosons, through which they are accessed. This is the case of the trilinear Higgs-boson self-coupling and the coupling between two Higgs bosons and two vector bosons, both accessible via Higgs-boson pair (HH) production, and the quartic Higgs-boson self-coupling, which requires triple Higgs-boson (HHH) searches. Whereas measuring the coupling between two Higgs bosons and two vector bosons probes the doublet structure of the Higgs field and provides an important test of the SM Higgs mechanism, determining the Higgs-boson self-couplings is crucial for probing the shape of the Higgs potential, whose form directly determines the spontaneous symmetry breaking paradigm and influences the evolution and stability of the universe.

In terms of their coupling modifiers $\kappa_C \equiv \lambda_C/\lambda_C^{\text{SM}}$, such that $\kappa_C = 1$ corresponds to a coupling λ_C matching its SM prediction λ_C^{SM} , the most stringent constraints on the trilinear Higgs-boson self-coupling κ_λ and the coupling between two Higgs bosons and two vector bosons κ_{2V} are currently $-1.2 < \kappa_\lambda < 7.2$ [29] and $0.67 < \kappa_{2V} < 1.38$ [28] at 95% confidence level (CL). The most stringent constraint to date on the quartic Higgs-boson self-coupling is $-230 < \kappa_4 < 240$ [30] at 95% CL; it was derived only recently in the first ATLAS search for the extremely rare HHH production.

The main production mode for Higgs-boson pairs is via gluon-gluon fusion (ggF), followed by vector-boson fusion (VBF). Although the VBF channel has a cross-section ten times smaller than the ggF channel, it provides unique access to κ_{2V} , in addition to κ_λ — which is also accessible via the ggF channel. As a consequence, the VBF production mode is of paramount importance for any reliable determination of κ_{2V} .

The analysis presented in this thesis is a search for the production of Higgs-boson pairs via VBF using proton-proton (pp) collision data collected with the ATLAS detector [31] at a center-of-mass energy of 13 TeV during the second data-taking period of the LHC (Run-2). The details of an on-going successive analysis using Run-3 data are also presented in this thesis. Both searches target highly Lorentz-boosted Higgs bosons decaying into four bottom quarks. While this decay mode has the largest branching ratio, the large background rates arising from multijet production poses a greater challenge compared to other, non-fully hadronic, HH final states.

In the SM, Higgs-boson pairs originate from non-resonant production and the cross-section is critically dependent on the values of κ_{2V} and κ_λ . Alternatively, in many BSM scenarios, intermediate heavy resonances originating from the vector-boson scattering may also decay into HH , leading to a resonant production of Higgs-boson pairs. Both the non-resonant and resonant interpretations are considered in this analysis.

Among the most important aspects of this search are the use of a dedicated machine-learning based algorithm to detect the $H \rightarrow b\bar{b}$ decays, the estimation of the predicted background using data-driven techniques, and the use of multivariate analysis to enhance the overall signal over background sensitivity. This analysis provides the best κ_{2V} constraint in ATLAS to date and expands the VBF HH resonant search to mass ranges never explored before. The results obtained in this analysis are included in the legacy ATLAS combination of HH searches using full Run-2 data [29].

This thesis is structured as follows. In Chapter 2, an overview of the theoretical foundations of the SM and the motivations of BSM physics with special attention on some selected benchmark models for HH production is given. This chapter also introduces the production mechanism of Higgs-boson pairs in hadron colliders and their decay modes. The LHC, the ATLAS detector, the simulation of pp collision events, and the datasets used in the analysis are all described in Chapter 3. In Chapter 4, the reconstruction of physics objects in the ATLAS experiment, with a special emphasis on the objects relevant for this search, is presented. The details on the analysis strategy to search for the VBF production of Higgs-boson pairs decaying into four bottom quarks are provided in Chapter 5. An overview of the statistical tools required to interpret the results of this analysis and the results themselves are all presented in Chapter 6.

Chapter 7 is dedicated to the work being done in the context of an on-going search of VBF $HH \rightarrow b\bar{b}b\bar{b}$ using data from the on-going LHC Run-3 and using more modern tools to identify the Higgs-boson decay into a pair of bottom quarks, a key element that affects critically the analysis sensitivity.

Finally, in Chapter 8 studies into how to improve the search for the yet little explored triple Higgs-boson production are presented. These studies are based on the first and recent HHH search targeting final states with six bottom quarks [30]. The goal is to use the knowledge acquired here on a novel ATLAS $HHH \rightarrow b\bar{b}b\bar{b}\tau^+\tau^-$ analysis.

All the work and results presented in this thesis are summarised in Chapter 9.

Author's contributions

The ATLAS collaboration consists of many physicists, engineers, technicians, students, and support staff, with approximately 6000 members and 3000 scientific authors. All the collaborative work for maintaining and upgrading the detector, developing the software to reconstruct physics observables from raw data, as well as analysing the data and interpreting the results, is of fundamental importance for the work presented in this thesis. In this context, the author's specific contributions are summarised here.

The author was one of the key analysers in the search for Higgs-boson pairs via VBF using Run-2 data. The author contributed to all aspects of the resonant interpretation: sample production, analysis strategy optimisation, construction of the pBDT, background estimation, and estimation of several systematic uncertainties. The author was also one of the main contributors to the framework to perform the statistical analysis, which was used for both the non-resonant and resonant interpretations of the analysis.

Regarding the on-going equivalent non-resonant VBF HH analysis using additionally Run-3 data, the author has also been a key contributor. The author is one of the main developers of the framework to reconstruct the Higgs-boson candidates and VBF jets, one of the main developers of the statistical framework to interpret the results, and is contributing to all aspects of the analysis strategy. The author is the editor of the internal supporting document of this new VBF HH analysis.

In the context of the triple Higgs-boson studies, while the $HHH \rightarrow b\bar{b}b\bar{b}b\bar{b}$ event reconstruction strategy is based on the referenced paper, the studies involving transformers were all carried out by the author. Moreover, the author is now co-contact in the new ATLAS $HHH \rightarrow b\bar{b}b\bar{b}\tau^+\tau^-$.

Finally, the author is one of the experts and has made several contributions to the framework used to produce the data format necessary for studies involving tracking in dense hadronic environments.

Theoretical Foundations and Motivation

In this chapter, the theoretical foundations and the motivation for the analysis presented in this thesis are introduced. First, in Section 2.1 an overview of the Standard Model of Particle Physics is given, followed by a motivation for the search of Higgs-boson pair production within and beyond the Standard Model (Section 2.2). Lastly, in Section 2.3, a description of the production and decay modes of Higgs-boson pairs is presented.

2.1 The Standard Model of Particle Physics

The universe at the smallest scales can be described by the elementary particles and their interactions. In the framework of quantum field theory, particles are treated as excited states of their underlying quantum fields, and the interactions between them emerge from the exchange of mediating particles. Each elementary particle belongs to a group based on its quantum properties, like electric charge and spin. First, there are the fermionic spin-1/2 particles (and their anti-particles) which make up matter. They are divided into leptons and quarks, where only quarks carry colour charge. Leptons are further divided into electrically charged leptons — electrons, muons, and taus —, and neutrinos, which have neither electric nor colour charge. There are six species, or flavours, of leptons and another six of quarks. The combinations of a charged lepton with its neutrino partner and pairs of quarks of different flavours further separate all fermions into families. Photons, gluons, and the massive W^\pm and Z are spin-1 bosons responsible for the electromagnetic, strong, and weak interactions, respectively. Lastly, there is the neutral spin-0 Higgs boson, whose underlying field gives rise to the masses of the W^\pm and Z gauge bosons and matter particles. The properties of each of the elementary particles are summarised in Table 2.1.

The current best approach to date to describe the known elementary particles and their interactions is the Standard Model of Particle Physics (SM), a gauge field theory based on the symmetry group $SU(3)_C \otimes SU(2)_L \otimes U(1)_Y$ [1–3]. The theory of strong interactions, described by quantum chromodynamics (QCD), is incorporated in the SM symmetry group

Table 2.1: Properties of the elementary particles described by the Standard Model. The second, third, and fourth columns (in blue) show the properties of the fermions. The upper half shows the electron (e), muon (μ), tau (τ), together with their corresponding neutrinos ν . The lower half describes the quarks up (u), down (d), strange (s), charm (c), bottom (b), and top (t). The last column, in red, shows the properties, respectively, of the bosons W^\pm , Z , gluon (g), and photon (γ). The very last row in the table (in orange) contains the properties of the Higgs boson (H). Uncertainties and additional properties can be found in Ref. [32].

	e	ν_e	μ	ν_μ	τ	ν_τ	W^\pm	Z
Mass [MeV]	0.511	$< 0.225 \cdot 10^{-3}$	105.7	< 0.19	$1.78 \cdot 10^3$	< 18.2	$80.4 \cdot 10^3$	$91.2 \cdot 10^3$
Charge [e]	-1	0	-1	0	-1	0	± 1	0
Spin	1/2	1/2	1/2	1/2	1/2	1/2	1	1
	u	d	s	c	b	t	g	γ
Mass [MeV]	2.16	4.70	93.5	$1.27 \cdot 10^3$	$4.18 \cdot 10^3$	$172.6 \cdot 10^3$	0	$< 10^{-24}$
Charge [e]	$+2/3$	$-1/3$	$-1/3$	$+2/3$	$-1/3$	$+2/3$	0	0
Spin	1/2	1/2	1/2	1/2	1/2	1/2	1	1
Higgs boson H : mass = 125.2 GeV, neutral, spin-0								

via the non-Abelian $SU(3)_C$ gauge group, and is built upon the assumption of conservation of the quantum number colour [33], which realises itself in nature with $N_C = 3$ different values. The strong force is mediated by an octet ($N_C^2 - 1$) of vector boson gauge fields, the gluons, and is present only between particles that carry colour charge. The electroweak interaction (EW) [3, 8, 34], described by the remaining product group $SU(2)_L \otimes U(1)_Y$, is mediated by a triplet of gauge fields denoted $W_\mu^{a=1,2,3}$ of the non-Abelian $SU(2)_L$ group that couples only to left-handed (“ L ”) fermions, in addition to the singlet field B_μ originating from the unitary group $U(1)_Y$ that ensures hypercharge (“ Y ”) conservation. As will be shown later, the physical massive bosons W^\pm and Z , as well as the massless photon field, are described as superpositions of the gauge fields W_μ^a and B_μ .

2.1.1 Quantum Chromodynamics

The core of the theory of the strong interaction relies on the assumption of a new quantum number called *colour* that can take three possible values. The existence and phenomenology of the strong interaction is based on ample experimental evidence, notably for instance from electron-positron annihilation data that shows a higher cross-section into hadrons compared to muons that is compatible with the prediction where quarks carry another quantum number that leptons do not [35]. There are also other critical theoretical perspectives; a prominent example is the Δ^{++} (uuu) boson that would violate Pauli’s exclusion principle¹ if no colour quantum number existed.

The starting point to mathematically construct the theory of quantum chromodynamics is the free Dirac Lagrangian of a massless spin-1/2 quark field

$$\mathcal{L}_0 = \bar{q}(x)(i\gamma^\mu \partial_\mu)q(x), \quad (2.1)$$

¹ According to Pauli’s exclusion principle, two identical fermions can not occupy simultaneously the same quantum state within a quantum system.

where the quark field (just one flavour is assumed here for simplicity) is a colour triplet $q = (q^1, q^2, q^3)^T$, γ^μ are the 4×4 Dirac matrices, and $\bar{q} \equiv q^\dagger \gamma^0$. A quadratic mass term $m_q \bar{q} q$ has been omitted as it violates gauge invariance under the electroweak force discussed in Section 2.1.2.

The free Lagrangian \mathcal{L}_0 is invariant under global $SU(3)_C$ transformations represented by the unitary matrix

$$U = \exp(it^a \theta_a), \quad (2.2)$$

where t^a ($a = 1, 2, \dots, 8$) denote the 3×3 generators of the $SU(3)_C$ group in the fundamental representation and θ_a are real arbitrary parameters. The so-called “gauge principle” requires that the invariance should also hold locally, i.e., when the arbitrary parameters are dependent on the space-time point x , $\theta_a \rightarrow \theta_a(x)$. This is achieved if the quark derivative $\partial_\mu q(x)$ is replaced by its covariant counterpart

$$D_\mu(x) q(x) \equiv [\partial_\mu + ig_s t_a G_\mu^a(x)] q(x), \quad (2.3)$$

with g_s being a coupling and $G_\mu^a(x)$ eight boson fields that are identified as the gluons. The gluon fields are finally incorporated in the theory as real propagating fields by including in the Lagrangian the $SU(3)_C$ -invariant kinematic term

$$\mathcal{L}_{\text{kin}} = -\frac{1}{4} G_a^{\mu\nu} G_{\mu\nu}^a, \quad (2.4)$$

with the field strengths $G_a^{\mu\nu}$ defined as

$$G_a^{\mu\nu} \equiv \partial^\mu G_a^\nu - \partial^\nu G_a^\mu - g_s f^{abc} G_b^\mu G_c^\nu, \quad (2.5)$$

and f^{abc} being the structure constants that dictate the commutation relations of the $SU(3)_C$ generators.

Due to the presence of non-zero f^{abc} , the final QCD Lagrangian has terms that couple gluons to gluons, a crucial element responsible for the peculiar behaviour of the QCD coupling, later identified as α_s . In particular, the strong coupling parameter increases for large distances (or, equivalently, low energies), favouring the confinement of quarks as bound states known as hadrons. If two quarks from a bound state, say $q\bar{q}$, attempt to be separated, the strong interaction becomes so large that the energy stored in the gluon field is sufficient to create a new $q\bar{q}$ pair that will bind with the original quarks. This characteristic of the QCD coupling thus connects well with the long-standing observation of quark confinement [36]. On the other side of the energy spectrum, at small distances, the α_s tends towards zero and the phenomena of asymptotic freedom [9, 10] emerges in the theory of the strong interaction. In this scenario, quarks inside hadrons can be approximately treated as free and the use of perturbation theory is valid.

2.1.2 Electroweak Theory

Many of the key properties of the electroweak theory were studied in low-energy experiments, such as in the seminal experiment of Wu [37] that provided the first evidences of parity violation. It has been found out that neutrinos (anti-neutrinos) have always left-handed (right-handed) chiralities and that only the left-handed (right-handed) fermions (anti-fermions) participate in the weak interaction. Here, a fermion field is said to be left- (L) or right-handed (R) if it is an eigenstate of the chirality operator $P_{L/R} = \frac{1}{2}(1 \mp \gamma^5)$, where γ^5 is the pseudoscalar Dirac matrix. Lepton and neutrino scattering data additionally provided fundamental insights into the existence of different types of neutrinos and lepton quantum numbers that are conserved in any process. While $\nu_e n \rightarrow e^- p$ for instance is a valid process observed in nature, $\nu_e p \rightarrow e^+ n$ is forbidden as the difference in the number of leptons and anti-leptons is not conserved. Measurements of muon decays $\mu^- \rightarrow e^- \bar{\nu}_e \nu_\mu$ were also of huge importance for the indirect evidence that the weak interaction involves a mediator coupling electrons and neutrinos with equal strength. Finally, the observation of neutrino-electron scattering processes [38] were also of great importance to confirm the need of three intermediate spin-1 bosons in the weak interaction, since processes like $\nu_e e^- \rightarrow \nu_e e^-$ are generated via charged boson exchange but $\nu_\mu e^- \rightarrow \nu_\mu e^-$ requires a neutral mediator [39]. All these arguments provide the fundamental ingredients for the theory of weak interactions.

Given that fermions from the same family like e and ν_e , or u and d , must have their left-handed component coupled with equal strength to the charged weak gauge bosons W^\pm , they should form a doublet under $SU(2)_L$. The right-handed components of the fermions form singlets. The existence of three weak gauge bosons is satisfied once the gauge principle is applied in the $SU(2)_L$ group, as already seen in the construction of QCD in the last section. For a unified electroweak theory, an additional $U(1)_Y$ symmetry should be included in order to provide an extra degree of freedom for the photon field [40].

The free spin-1/2 Lagrangian takes the form

$$\mathcal{L}_0 = \sum_{k=1}^3 \bar{\psi}_k(x) (i\gamma^\mu \partial_\mu) \psi_k(x), \quad (2.6)$$

where the fermion fields for one family in the quark sector are identified as

$$\psi_1(x) = \begin{pmatrix} u \\ d \end{pmatrix}_L, \quad \psi_2(x) = u_R, \quad \psi_3(x) = d_R, \quad (2.7)$$

and analogously for the lepton families. The free Lagrangian is invariant under $SU(2)_L \otimes U(1)_Y$ transformations

$$\psi_1(x) \rightarrow \exp(iy_1\beta) \exp\left(i\frac{\sigma_a}{2}\theta^a\right) \psi_1(x), \quad (2.8)$$

$$\psi_2(x) \rightarrow \exp(iy_2\beta) \psi_2(x), \quad (2.9)$$

$$\psi_3(x) \rightarrow \exp(iy_3\beta) \psi_3(x), \quad (2.10)$$

where σ_a ($a = 1, 2, 3$) are the Pauli matrices of the $SU(2)_L$ transformation being applied solely to the doublet $\psi_1(x)$, while the parameters y_k called hypercharges multiplied by β are the phase components of the $U(1)_Y$ transformation of each spinor.

Applying now the gauge principle, the invariance under local transformations with $\theta^a(x)$ and $\beta(x)$ is satisfied if the fermion derivatives $\partial_\mu \psi_k(x)$ are replaced by their covariant counterparts

$$D_\mu \psi_1(x) \equiv \left[\partial_\mu + ig \frac{\sigma_a}{2} W_\mu^a(x) + ig' y_1 B_\mu(x) \right] \psi_1(x), \quad (2.11)$$

$$D_\mu \psi_2(x) \equiv \left[\partial_\mu + ig' y_2 B_\mu(x) \right] \psi_2(x), \quad (2.12)$$

$$D_\mu \psi_3(x) \equiv \left[\partial_\mu + ig' y_3 B_\mu(x) \right] \psi_3(x), \quad (2.13)$$

with new boson fields $W_\mu^{a=1,2,3}$ and B_μ , and couplings g and g' . Mass terms like $m \bar{\psi}_k \psi_k$ in the Lagrangian are forbidden as otherwise the gauge invariance would not be satisfied.

Finally, the gauge-invariant kinetic terms for the gauge fields W_μ^a and B_μ are introduced in the theory via the Lagrangian term

$$\mathcal{L}_{\text{kin}} = -\frac{1}{4} B_{\mu\nu} B^{\mu\nu} - \frac{1}{4} W_{\mu\nu}^a W_a^{\mu\nu}, \quad (2.14)$$

with the field strengths defined as

$$B_{\mu\nu} \equiv \partial_\mu B_\nu - \partial_\nu B_\mu, \quad (2.15)$$

$$W_{\mu\nu}^a \equiv \partial_\mu W_\nu^a - \partial_\nu W_\mu^a - g \epsilon^{abc} W_\mu^b W_\nu^c, \quad (2.16)$$

where ϵ^{abc} is the Levi-Civita tensor.

When analysing the interactions between the bosonic gauge fields and the fermions in the final electroweak Lagrangian with the covariant derivatives, it is possible to relate W_μ^a and B_μ to the real fields W^\pm , Z , and γ that are observed in nature. In particular, the gauge invariant Lagrangian \mathcal{L} contains the following interaction terms:

$$\mathcal{L} \supset -g \bar{\psi}_1 \gamma^\mu \left(\frac{\sigma_a}{2} W_\mu^a \right) \psi_1 - g' B_\mu \sum_{k=1}^3 y_k \bar{\psi}_k \gamma^\mu \psi_k. \quad (2.17)$$

The first term containing the $SU(2)_C$ matrices generates the charged-current interactions between the W^\pm bosons and the families of leptons and quarks. More precisely, if defining $W_\mu \equiv \frac{1}{\sqrt{2}} (W_\mu^1 + i W_\mu^2)$ the charged-current term \mathcal{L}_{CC} reads

$$\mathcal{L}_{CC} = -\frac{g}{2\sqrt{2}} \left\{ W_\mu^\dagger [\bar{\nu}_e \gamma^\mu (1 - \gamma^5) e + \bar{u} \gamma^\mu (1 - \gamma^5) d] + \text{h.c.} \right\}, \quad (2.18)$$

and therefore it is possible to identify the real vector boson W^+ (W^-) as W_μ (W_μ^\dagger).

The second term in Equation 2.17 together with the diagonal elements in the matrix product $\sigma_a W_\mu^a$ gives rise to the neutral-current interactions. The naive assumption that B_μ and W_μ^3 would correspond directly to the photon and Z fields is forbidden as this would implicate a chirality dependent interaction between fermions and the electromagnetic field, which would be in disagreement with the observations e.g. from electron-proton scattering experiments [41]. Therefore, the real boson fields Z_μ and A_μ for the neutral Z boson and the photon are first written as a generic rotation from the unphysical B_μ and W_μ^3 , i.e.,

$$\begin{pmatrix} W_\mu^3 \\ B_\mu \end{pmatrix} = \begin{pmatrix} \cos \theta_W & \sin \theta_W \\ -\sin \theta_W & \cos \theta_W \end{pmatrix} \begin{pmatrix} Z_\mu \\ A_\mu \end{pmatrix}, \quad (2.19)$$

where θ_W is defined as a weak mixing angle. The neutral-current term \mathcal{L}_{NC} thus takes the form

$$\begin{aligned} \mathcal{L}_{NC} = & - \sum_{k=1}^3 \bar{\psi}_k \gamma^\mu A_\mu \left[g \frac{\sigma_3}{2} \sin \theta_W + g' y_k \cos \theta_W \right] \psi_k \\ & - \sum_{k=1}^3 \bar{\psi}_k \gamma^\mu Z_\mu \left[g \frac{\sigma_3}{2} \cos \theta_W - g' y_k \sin \theta_W \right] \psi_k. \end{aligned} \quad (2.20)$$

In order to obtain a description of the electromagnetic interaction between fermions and the photon field in the first part of the expression, the corresponding term in brackets should reduce to the electromagnetic coupling, i.e.,

$$\left[g \frac{\sigma_3}{2} \sin \theta_W + g' y_k \cos \theta_W \right] = e Q_k, \quad (2.21)$$

with Q_k being the electric charge (in units of e) of the corresponding fermion. This is only possible if the electromagnetic and weak couplings are unified via the relation

$$g \sin \theta_W = g' \cos \theta_W = e, \quad (2.22)$$

and the hypercharge y_k of each fermion is correlated with its electric charge and weak isospin quantum number T_3 by

$$y = Q - T_3, \quad (2.23)$$

where the weak isospin of a fermion field is the corresponding eigenvalue ($\pm 1/2$) when applying the isospin operator $\sigma_3/2$ to the field. It is equal to $+1/2$ ($-1/2$) for left-handed electrons and u -quarks (neutrinos and d -quarks), and zero for all right-handed fermions, as they form singlets.²

Hypercharge is by construction a conserved quantity due to the $U(1)_Y$ gauge group choice. The conservation of y and T_3 implies electric charge conservation, in accordance with

² This is expanded also for the other lepton and quark families.

experimental observations.

The remaining interactions between the neutral Z boson and the different fermions can also be easily derived by simply applying the constraints above in Equation 2.20.

2.1.3 Spontaneous Symmetry Breaking and the Higgs boson

As mentioned before, mass terms in the free Lagrangian would break the gauge invariance of the model, meaning that so far all fermions and gauge bosons in the theory are assumed to be massless, at odds with the experimental observations that the Z and W^\pm bosons, as well as quarks and charged leptons, are massive. The procedure to generate the masses of the weak gauge bosons is referred to as spontaneous symmetry breaking in a gauged Lagrangian of a scalar particle [18–20].

The starting point to understand this procedure is the free Lagrangian of a complex scalar field $\phi(x)$,

$$\mathcal{L}_0 = \partial_\mu \phi^\dagger \partial^\mu \phi - V(\phi), \quad (2.24)$$

where the potential $V(\phi)$ is given by

$$V(\phi) = \mu^2 \phi^\dagger \phi + \lambda |\phi^\dagger \phi|^2, \quad (2.25)$$

with the constraint that λ is a real, positive number in order to have a potential bounded from below. The scenario $\mu^2 < 0$ is as well of particular interest because in this case the potential exhibits minima that are not only obtained with the trivial free vacuum $|0\rangle$ expectation value $|\phi_0| \equiv |\langle 0 | \phi | 0 \rangle| = 0$. In particular, for $\lambda > 0$ and $\mu^2 < 0$ a minimum is obtained for

$$|\phi_0| = \sqrt{\frac{-\mu^2}{2\lambda}} \equiv \frac{v}{\sqrt{2}}. \quad (2.26)$$

Given the $U(1)$ invariance of the free Lagrangian, there exists an infinite number of vacuum degenerate states with minimum energy

$$\phi_0(x) = \frac{v}{\sqrt{2}} \exp(i\theta). \quad (2.27)$$

When breaking the symmetry by choosing one specific solution, say with $\theta = 0$, the excitations of the field around its minimum can be parametrized in terms of two real scalar fields $\varphi_1(x)$ and $\varphi_2(x)$ as

$$\phi(x) = \frac{1}{\sqrt{2}} [v + \varphi_1(x) + i\varphi_2(x)]. \quad (2.28)$$

An illustration of the potential $V(\phi)$ of a complex field in the ground state parametrised in terms of real fields is given in Figure 2.1. Under this parametrisation the potential becomes




Figure 2.1: Potential $V(\phi)$ of a complex scalar field $\phi(x)$ as given in Equation 2.25 for $\lambda > 0$ and $\mu^2 < 0$. Excitations of the field $\phi(x)$ from the ground state are parametrised in terms of the real fields $\varphi_1(x)$ and $\varphi_2(x)$ according to Equation 2.28.

$$V(\phi) = V(\phi_0) - \mu^2 \varphi_1^2 + \lambda v \varphi_1 (\varphi_1^2 + \varphi_2^2) + \frac{\lambda}{4} (\varphi_1^2 + \varphi_2^2)^2. \quad (2.29)$$

The presence of the quadratic term $\mu^2 \varphi_1^2$ in the potential without an equivalent term for φ_2 means that φ_1 can be understood as a massive scalar state, while φ_2 should be interpreted as massless. The appearance of a massless excitation in a scenario where there was a spontaneously symmetry breaking is part of a much more general result known as the Goldstone theorem [42–44]. More precisely, by breaking the symmetry of a Lagrangian invariant under a continuous symmetry group G , a massless scalar field (known as a Nambu-Goldstone boson) appears for each generator of G that was broken, i.e., does not belong to the new subgroup K of G under which the vacuum state is invariant.

The situation gets more intriguing when one considers a $SU(2)_L$ doublet of complex scalar fields

$$\phi(x) = \begin{pmatrix} \phi^{(+)}(x) \\ \phi^{(0)}(x) \end{pmatrix}, \quad (2.30)$$

with superscripts (+) and (0) indicating the electric charge of the corresponding field, in the $SU(2)_L$ gauged scalar Lagrangian

$$\mathcal{L} = (D_\mu \phi)^\dagger D^\mu \phi - V(\phi). \quad (2.31)$$

The covariant derivative is given in Equation 2.11 and the hypercharge y_ϕ is equal to $1/2$. Similarly to the $U(1)$ scenario, it is possible to parametrize the doublet as fluctuations of the vacuum state using four real fields $\theta^{a=1,2,3}(x)$ and $H(x)$ as

$$\phi(x) = \exp \left[i \frac{\sigma_a}{2} \theta^a(x) \right] \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix}, \quad (2.32)$$

and break the symmetry by choosing one particular solution $\theta^{a=1,2,3}(x) = 0$. In this scenario, the kinematic term of the $SU(2)_L$ gauged scalar Lagrangian reads

$$(D_\mu \phi)^\dagger D^\mu \phi = \frac{1}{2} \partial_\mu H \partial^\mu H + (v + H)^2 \left[\frac{g^2}{4} W_\mu^\dagger W^\mu + \frac{g^2}{8 \cos^2 \theta_W} Z_\mu Z^\mu \right]. \quad (2.33)$$

As a consequence, the Lagrangian now contains quadratic terms for the weak gauge bosons W^\pm and Z , which means they are now described as massive particles with masses m_W and m_Z given by

$$m_W = m_Z \cos \theta_W = \frac{1}{2} v g. \quad (2.34)$$

The relations between electroweak parameters as predicted in the equation above provide a large range of tests of the electroweak theory. First, it predicts a m_Z higher than m_W , in absolute agreement with the current experimental values $m_Z = 91.1880 \pm 0.0020$ GeV and $m_W = 80.369 \pm 0.013$ GeV [32]. From these numbers it is possible to extract the value of the electroweak angle θ_W , which is found to be in great accord with the values experimentally obtained for instance from neutrino-lepton scattering [45]. It is also possible to predict the vacuum expectation scale $v \approx 246$ GeV (commonly expressed in terms of the Fermi constant G_F as $(\sqrt{2}G_F)^{-1/2}$) and compare e.g. with results from precise lepton lifetime measurements [46].

A similar procedure can be applied using a Yukawa-like type of Lagrangian that couples the scalar complex doublet, the left-handed doublet of fermions, and right-handed fermion singlets, to generate the masses m_f of quarks and charged leptons. More specifically, the gauge invariant Yukawa-like Lagrangian \mathcal{L}_Y can be written as

$$\mathcal{L}_Y = - \left(1 + \frac{H}{v} \right) [m_d \bar{d}d + m_u \bar{u}u + m_e \bar{e}e], \quad (2.35)$$

and analogously for the second and third generations of fermions.

The process of spontaneously breaking the symmetry not only yields mass terms for the weak gauge bosons, but also introduces a new spin-0 field $H(x)$, called Higgs field, that couples to all massive particles and itself. The potential in Equation 2.25, now as a function of $H(x)$, reads

$$V(H) = \left(\frac{m_H^2}{2} \right) H^2 + \left(\frac{m_H^2}{2v^2} \right) v H^3 + \left(\frac{m_H^2}{8v^2} \right) H^4, \quad (2.36)$$

where the first term is a mass term, whereas the second and third terms refer to the trilinear and quartic Higgs-boson self-couplings. The mass of the Higgs boson is expressed as $m_H = \sqrt{2\lambda}v$ and to date is determined experimentally as $m_H = 125.20 \pm 0.11$ GeV [32]. Since the vacuum

expectation value v is already known from electroweak precision measurements, m_H is the only free parameter of the Higgs mechanism. Once it is determined, it is possible to predict all the production and decay modes of the Higgs boson.

In conclusion, SM particles acquire mass through their interaction with the Higgs field. Almost 50 years after its prediction, in 2012, a new spin-0 scalar particle compatible with the Higgs boson was observed by the ATLAS and CMS collaborations at the LHC [16, 17]. Since then, increasingly precise measurements of this new particle have shown that its characteristics, like production cross-sections and branching ratios, agree very well with those predicted by the SM Higgs mechanism within the current uncertainties [27, 28]. All these measurements, however, still rely on singly produced Higgs bosons.

2.1.4 The Higgs-boson self-couplings

As shown in Equation 2.36, the shape of the Higgs potential is closely related to the trilinear and quartic Higgs-boson self-couplings, defined respectively as

$$\lambda_{HHH}^{\text{SM}} \equiv \frac{m_H^2}{2v^2} \quad (2.37)$$

and

$$\lambda_{HHHH}^{\text{SM}} \equiv \frac{m_H^2}{8v^2}. \quad (2.38)$$

It is expected that modifications of the Higgs sector in beyond SM (BSM) theories (see Section 2.2), which may or may not introduce additional particles to the model, should affect the structure of the Higgs potential and therefore introduce deviations of the Higgs-boson self-couplings compared to those predicted by the SM. These modifications may, of course, not only affect the Higgs-boson self-couplings, but also the couplings of the Higgs boson to other massive SM particles. One particular example of interest discussed in this thesis is the coupling of a pair of Higgs bosons to a pair of vector bosons $V(= Z, W^\pm)$, λ_{2V}^{SM} , which depends on the vector boson masses via m_V^2/v^2 .

Given the potential effect of BSM phenomena on the Higgs-boson couplings, a great tour de force to measure them at the highest possible precision is pursued by the large LHC experiments. In the context of these measurements, the couplings are typically expressed in terms of κ -parameters defined as the ratio between an anomalous coupling value and the value predicted by the SM, i.e., $\kappa_C \equiv \lambda_C/\lambda_C^{\text{SM}}$. If no deviations to the SM exist, $\kappa_C = 1$. In Figure 2.2, a summary of the measurements of the Higgs-boson couplings to a pair of fermions or vector bosons performed by the ATLAS collaboration is presented, showing an excellent agreement between experiment and the SM within the current uncertainties [27].

The trilinear Higgs-boson self-coupling modifier κ_λ and the modifier for the coupling of a Higgs-boson pair to a pair of vector bosons κ_{2V} are, however, still not measured. This is mainly because these couplings are accessible from processes involving the yet unobserved production of Higgs-boson pairs, which typically have small SM cross-sections.




Figure 2.2: Measurements of the reduced couplings κ of the Higgs boson to fermions and vector boson performed by the ATLAS collaboration [27]. Two scenarios have been considered to cope with the low sensitivity to the charm coupling: in the first, the charm coupling modifier is set to be equal to the top coupling modifier (coloured circle markers); in the second, the charm coupling modifier is left as a free floating parameter in the fits (grey cross markers). Notice the y-axis is different for the vector bosons and is expressed in terms of $\sqrt{\kappa}$ instead of κ .

Nevertheless, due to large dependence of the κ -parameters on the production cross-section (see further details in Section 2.3.2), it is possible to constrain the values of these couplings at a certain confidence level (CL). The current most stringent constraint of κ_λ was recently obtained by the ATLAS collaboration combining the results of several HH searches and reads $-1.2 < \kappa_\lambda < 7.2$ at 95% CL [29]. The best limits for κ_{2V} were obtained by the CMS collaboration and reads $0.67 < \kappa_{2V} < 1.38$ at 95% CL [28]. The current upper limits on the HH production cross-section as a function of κ_λ and κ_{2V} are shown in Figures 2.4 and 2.3, respectively.³

³ The final ATLAS results for κ -constraints are actually based on a likelihood scan instead of using the cross-section upper limit approach, as discussed in further detail in Chapter 6. Here the cross-section plot is shown to align with what CMS has publicly available.




Figure 2.3: Upper limits at 95% CL on the HH production cross-section as a function of κ_λ obtained by the ATLAS collaboration [29]. The combined result (black) is the combination from several analyses. Regions where the predicted cross-section is higher than the observed cross-section are excluded at 95% CL.




Figure 2.4: Upper limits at 95% CL on the HH production cross-section as a function of κ_{2V} obtained by the CMS collaboration [28]. Regions where the predicted cross-section is higher than the observed cross-section are excluded at 95% CL.

2.2 Motivations for Beyond Standard Model Theories

Despite the great success of the SM in predicting numerous properties of the elementary building blocks of the universe, it is widely considered an incomplete theory. It can not satisfactorily describe the mass of the neutrinos; it does not include gravity in the theory, meaning that it attempts to describe only three out of the four forces of the universe; and it can not answer some of the big open questions of particle physics, such as the existence of dark matter and the observed excess of matter over anti-matter in the universe.

The concept of dark matter was first postulated based on cosmological observations that large astronomical objects, such as galaxies or even stars, move faster than the predicted kinematics expected from the gravitational pulls of other visible objects in the vicinity [47]. Further indications of dark matter arose from precision measurements of the cosmic microwave background power spectrum [22] and from gravitational lensing [48]. In fact, it is now suggested that only 15% of the matter in the universe is composed of ordinary matter described by the SM, while the remaining 85% are assumed to be composed of weakly interacting dark matter that interacts with visible matter mostly gravitationally [22]. The SM does not contain a particle candidate for dark matter, which makes BSM theories a must.

Another big open question of particle physics is the imbalance between the total number of baryons — hadrons with an odd number of at least 3 valence quarks⁴ — and anti-baryons in the universe. The hypothetical process of creating the excess of baryons over anti-baryons starting from a universe with equal numbers of both is typically referred as baryogenesis. This asymmetry may have arisen from an electroweak phase transition in the early universe after the Big Bang, but the SM electroweak symmetry breaking can not produce this phenomenon in a universe where $m_H \gtrsim 80$ GeV [49]. Alternative solutions for this problem should satisfy the Sakharov conditions for baryogenesis [50], in which the existence of charge conjugation and parity (CP) violation is mandatory, meaning that processes with opposite electric charge (particle \rightarrow anti-particle) and chirality ($L \rightarrow R$) should not necessarily occur at equal rates. While CP-violating processes do exist in the SM, as observed in kaon decays [51], they are not sufficient to account for the total baryon asymmetry present in the universe.

Many BSM theories also attempt to solve what is known as the hierarchy problem. Assuming that the SM is valid at the Planck scale $M_P \sim 10^{19}$ GeV, a point where gravity would start to contribute more significantly, the higher-order quantum-loop corrections of the Higgs-boson mass would be more than 30 orders of magnitude higher than the value m_H determined experimentally. This requires the bare Higgs-boson mass⁵ to be extremely fine-tuned to correctly cancel the large radiative corrections — a process that seems unnatural. Some solutions to the hierarchy problem normally invoke the existence of extra particles interacting with the Higgs boson that would reduce the size of the radiative corrections, with

⁴ Hadrons are not composed of a fixed number of quarks, but are rather a sea of interacting quarks and gluons. The so-called valence quarks of the hadron, similarly to what is seen in atoms, determine the properties of the particle.

⁵ The interactions modify some of the particle properties like charge and mass, leading to differences between the observed physical values and the bare values present in the Lagrangian.

supersymmetry being a famous example [24].

The following sub-sections are dedicated to explaining in greater detail two relevant BSM scenarios that attempt to solve some of these open questions, with special attention on how they can affect the Higgs sector.

2.2.1 Two-Higgs-doublet models (2HDMs)

One of the most studied, but still relatively simple, extensions of the SM are the two-Higgs-doublet models (2HDMs) [26, 52]. Motivations for an additional $SU(2)_L$ doublet in the Higgs sector appear in a series of proposed variations of the SM. To name a few examples, 2HDMs are an important component of some supersymmetric theories, in which each particle from the fermion group has a bosonic partner [24]; it is a key component in certain classes of axial models of dark matter [25]; and because of the introduction of new degrees of freedom, it is possible to provide extra sources of CP-violation [53].

2HDMs are known for their rich vacuum structure, with the most general potential having up to 14 parameters and allowing CP-symmetry conservation, CP-violation, or C-violation [26]. It is commonly imposed a softly broken Z_2 symmetry⁶ to avoid flavour-changing neutral currents at leading order — hypothetical interactions that change the flavour of a lepton or quark but not its electric charge.⁷ In this context, there are a few scenarios depending on how fermions couple to the two Higgs doublets: in the so-called type I 2HDM, all fermions couple to only one of the doublets; in type II 2HDMs, quarks with electric charge $-1/3$ and leptons couple to one of the doublets, while quarks with electric charge $+2/3$ couple to the second doublet. There are also two extra possibilities if a Z_2 symmetry is imposed depending on how leptons couple to the Higgs doublets.

In this thesis, only 2HDMs with a CP-conserving potential will be considered, which is one of the most common 2HDMs benchmarks, motivated e.g. by supersymmetric theories. Under this assumption and with a softly broken Z_2 discrete symmetry, the most general potential V for two $SU(2)_L$ doublets Φ_1 and Φ_2 contains 8 real parameters and reads

$$\begin{aligned}
 V = & m_{11}^2 \Phi_1^\dagger \Phi_1 + m_{22}^2 \Phi_2^\dagger \Phi_2 - m_{12}^2 [\Phi_1^\dagger \Phi_2 + \Phi_2^\dagger \Phi_1] + \frac{\lambda_1}{2} (\Phi_1^\dagger \Phi_1)^2 + \frac{\lambda_2}{2} (\Phi_2^\dagger \Phi_2)^2 \\
 & + \lambda_3 \Phi_1^\dagger \Phi_1 \Phi_2^\dagger \Phi_2 + \lambda_3 \Phi_1^\dagger \Phi_2 \Phi_2^\dagger \Phi_1 + \frac{\lambda_5}{2} [(\Phi_1^\dagger \Phi_2)^2 + (\Phi_2^\dagger \Phi_1)^2].
 \end{aligned} \tag{2.39}$$

The first tree terms proportional to m_{ij} are mass terms, with the softly Z_2 symmetry breaking coming from the cross-term proportional to a non-zero m_{12} . The remaining terms proportional to the couplings λ_i describe the interactions between both doublets.

In the same spirit of Section 2.1.3, after an electroweak symmetry breaking in the potential V , it is possible to expand the complex doublets around the minimum in terms of eight fields as follows:

⁶ Under a Z_2 symmetry the two $SU(2)_L$ doublets transform as $\Phi_1 \rightarrow \Phi_1, \Phi_2 \rightarrow -\Phi_2$.

⁷ In the SM this is only possible at next-to-leading order loop-induced processes.

$$\Phi_{a=1,2} = \begin{pmatrix} \phi_a^+ \\ (v_a + \rho_a + i\eta_a)/\sqrt{2} \end{pmatrix}, \quad (2.40)$$

where v_a denote the vacuum expectation values of the neutral components of the doublets satisfying the relation $v_1^2 + v_2^2 \equiv v^2 \approx (246 \text{ GeV})^2$. Three of those fields should be massless excitations that give rise to the masses of the vector bosons Z and W^\pm , as expected from the Goldstone theorem. The remaining fields are then translated into five physical Higgs bosons. Two of these fields correspond to charged Higgs bosons, two to scalars, and the last one to a pseudoscalar.⁸ The relevant rotations are given by [54]

$$\begin{pmatrix} \rho_1 \\ \rho_2 \end{pmatrix} = R_\alpha \begin{pmatrix} H_1 \\ H_2 \end{pmatrix}, \quad \begin{pmatrix} \eta_1 \\ \eta_2 \end{pmatrix} = R_\beta \begin{pmatrix} \eta \\ A \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} \phi_1^+ \\ \phi_2^+ \end{pmatrix} = R_\beta \begin{pmatrix} \phi^+ \\ H^+ \end{pmatrix}, \quad (2.41)$$

where R_x is the typical two-dimensional rotation matrix, α is defined as a CP-even Higgs mixing angle between the two scalars H_1 and H_2 , and β is written in terms of v_a as $\tan \beta = v_2/v_1$. In the equation above, the pseudoscalar A and the charged Higgs bosons H^\pm ($H^- = (H^+)^\dagger$) are already identified as physical fields, while η and ϕ^+ (and $\phi^- = (\phi^+)^\dagger$) are would-be Goldstone bosons that account for the masses of the SM weak gauge bosons. Regarding the scalars H_1 and H_2 , one of them can have its mass tuned via the parameters of the potential V to match the mass value of the already observed SM Higgs boson H^{SM} . Additionally, in the so-called alignment limit $\alpha = \beta - \pi/2$ [55] the scalar now identified as H^{SM} has its couplings at leading order identical to their SM values. The remaining scalar particle is then interpreted as a heavy Higgs boson \tilde{H} and, together with the additional pseudoscalar and charged Higgs bosons, is a new interacting particle predicted by 2HDMs that may be observed at particle colliders like the LHC. In these models, the soft-breaking scale $M^2 \equiv 2m_{12}^2/\sin 2\beta$, the angle $\tan \beta$, and the particle masses m_Σ for $\Sigma \in \{\tilde{H}, A, H^\pm\}$, which depend on the model couplings, M^2 , and v_a , are the free parameters that can be chosen to produce different 2HDM phenomenologies.

The incorporation of new particles in the theory may add a number of production and decay modes of SM particles, depending on their masses, as well as affect the strength and properties of their couplings. Even though the alignment limit in 2HDMs is chosen to explicitly reproduce the SM couplings of H^{SM} at leading order, loop corrections in the vertices of the theory introduce terms dependent on the 2HDM parameters, leading to effective couplings that deviate from their SM predictions. To give a prominent example, the effective trilinear Higgs-boson self-coupling at next-to-leading-order in aligned type I 2HDMs reads [56]

⁸ A pseudoscalar particle behaves as a scalar particle but it is not invariant under parity transformations.




Figure 2.5: Trilinear Higgs-boson self-coupling modifier κ_λ as a function of the pseudoscalar Higgs mass m_A at one-loop (dashed curve) and two-loop (solid curve) orders [58]. The mass of the charged Higgs boson is set to be equal to m_A , while the heavy Higgs-boson mass $m_{\tilde{H}}$ and the scale M are fixed to 600 GeV. Further improvements on the experimental side to constrain the values of κ_λ will decrease the range of allowed values of m_A that can be explored.

$$\begin{aligned} \lambda_{HHH}^{\text{eff}}(2\text{HDM}) = \frac{m_H^2}{2v^2} \left[1 - \frac{N_c m_t^4}{3\pi^2 v^2 m_{H\text{SM}}^2} + \frac{m_{\tilde{H}}^4}{12\pi^2 v^2 m_{H\text{SM}}^2} \left(1 - \frac{M^2}{m_{\tilde{H}}^2} \right)^3 \right. \\ \left. + \frac{m_A^4}{12\pi^2 v^2 m_{H\text{SM}}^2} \left(1 - \frac{M^2}{m_A^2} \right)^3 + \frac{m_{H^\pm}^4}{6\pi^2 v^2 m_{H\text{SM}}^2} \left(1 - \frac{M^2}{m_{H^\pm}^2} \right)^3 + \dots \right]. \end{aligned} \quad (2.42)$$

While the first term proportional to the top-quark mass m_t^4 is a correction to the effective self-coupling predicted by the SM, all the remaining terms arise from radiative corrections involving 2HDM particles. These extra BSM corrections are however suppressed depending on how the values of the new particles masses differ from the soft-breaking scale M .

Constraining and eventually measuring the Higgs-boson couplings is thus crucial to probe and potentially exclude different benchmark predicted by BSM theories. In Figure 2.5, the values of κ_λ at one-loop and two-loop orders corrections are shown as a function of m_A under the hypothesis $m_{H^\pm} = m_A$, $M = m_{\tilde{H}} = 600$ GeV, and $\tan \beta = 2$. This relation between the parameters is particularly interesting due to its connection with custodial symmetries, typically present in BSM theories that attempt to solve the hierarchy problem [57]. The current experimental upper limit on κ_λ is capable of probing this scenario only for $m_A \gtrsim 900$ GeV [58].

2.2.2 Composite Higgs models

An interesting approach to explain the origin of the electroweak sector and its symmetry breaking is the presence of a new strong dynamics in composite Higgs models [59], where the Higgs boson is treated as a bound state of this strongly interacting sector. This is similar to how pions appear in the context of QCD, in which pions emerge as pseudo Goldstone bosons resulting from the spontaneous breaking of an approximate chiral symmetry.

In the class of composite Higgs models, the Higgs boson itself is another Goldstone mode generated from the spontaneous breaking of a larger $SO(5)$ symmetry group into $SO(4)$ ⁹, as a consequence of the formation of a condensate due to the strong dynamics happening at a scale f . Three out of the four Goldstone modes generated from the $SO(5) \rightarrow SO(4)$ symmetry breaking are massless and give rise to the SM weak vector bosons, while the fourth, referred to as a pseudo Goldstone boson, acquires mass from small loop effects and is identified as the SM Higgs boson. In this context, the Higgs-boson mass is protected by an approximate global symmetry and generated exclusively via quantum effects, thus offering a solution to the hierarchy problem.

In the limit $f \rightarrow \infty$ with the SM vacuum expectation value v fixed at 246 GeV, the “composite sector” that delivers the Higgs boson as a bound state decouples from the low energy degrees of freedom associated to the model, recovering the SM and the Higgs boson as an elementary particle. In fact, the ratio $\xi = v^2/f^2$ works as proxy for the vacuum misalignment angle, which measures how much the vacuum is tilted toward the Higgs direction because of explicit symmetry-breaking effects and hence gives the Higgs field a non-zero vacuum expectation value. Effectively, ξ dictates the mechanism by which the Higgs boson can behave as if it was a composite bound state or an elementary particle, thus controlling the deviations of composite Higgs models from the SM. As an example, the Higgs-boson couplings κ_{2V} and κ_V are predicted as

$$\kappa_{2V} \approx 1 - 2\xi \quad (2.43)$$

and

$$\kappa_V \approx \sqrt{1 - \xi} \quad (2.44)$$

in the minimal $SO(5)/SO(4)$ models [60]. Here, κ_V stands for the coupling modifier between a pair of vector bosons and a single Higgs boson.

One natural signature of this new strong dynamics would be the presence of new resonances that couple to the Goldstone modes. A simple benchmark is the existence of a scalar η described by the Lagrangian [61]

$$\mathcal{L}_\eta = \frac{1}{2} \partial^\mu \eta \partial_\mu \eta - \frac{1}{2} m_\eta^2 \eta^2 + \frac{a_\eta}{f} \eta \partial^\mu \pi^b \partial_\mu \pi^b, \quad (2.45)$$

⁹ The SM electroweak gauge group can be embedded naturally within $SO(4)$, thus $SO(4)$ also provides the machinery to generate the electroweak sector.

where π^b is the quartet of Goldstone modes and a_η a dimensionless coupling of the theory. The Lagrangian is fully unitary by setting $a_\eta = 1$.

In principle, the scalar η can be produced in high-energy colliders via vector-boson scattering and decay into two Higgs bosons, a process that would result in the resonant production of Higgs-boson pairs. The decay width Γ_η of the new scalar can be narrow or broad depending on the particle mass and on the strong dynamics scale f . In particular [61],

$$\Gamma_\eta = \frac{a_\eta}{8\pi^2} \frac{m_\eta^3}{f^2}. \quad (2.46)$$

Assuming a scale f greater than 500 GeV, mainly in view of the current constraints on the Higgs-boson couplings (recall Equation 2.43), but also not too large to avoid the decoupling of the composite sector, it is expected that the new scalar at the TeV scale should have a width of order $\mathcal{O}(10\%)$ of its mass.

2.3 Production and Decay of Higgs-boson pairs

As described in the sections above, constraining and eventually measuring the Higgs-boson couplings can provide valuable probes of BSM phenomena. In addition, extensions of the Higgs sector with new particles can be directly probed by searching for hypothetical processes in which these particles are produced. Searching for Higgs-boson pair production allows accessing κ_λ and κ_{2V} , and provides sensitivity to new particles decaying into HH .

2.3.1 Calculation of production cross-sections

The process of Higgs-boson pair production in proton-proton collisions, $pp \rightarrow HH$, can be factorised into two sub-components. First, the hard-scattering between two partons a and b (in this case, quarks or gluons within the protons) producing the desired HH final state can be described by perturbative QCD; the calculated partonic cross-section $\hat{\sigma}_{ab \rightarrow HH}$ is then convoluted with the Parton Distribution Functions (PDFs) f_a and f_b , which represent the probability densities of finding a parton with a given momentum fraction x . Based on the factorisation theorem [62], the total $pp \rightarrow HH$ cross-section reads

$$\sigma_{pp \rightarrow HH} = \sum_{a,b=g,q,\bar{q}} \int_{\tau_1}^1 \int_{\tau_2}^1 dx_1 dx_2 f_a(x_1, \mu_F^2) f_b(x_2, \mu_F^2) \hat{\sigma}_{ab \rightarrow HH}(x_1 p_1, x_2 p_2, \mu_R^2, \mu_F^2, \mathcal{P}), \quad (2.47)$$

where the lower limits τ_1 and τ_2 in the integrals depend on the threshold energy for the production of the desired final state. The different components entering the cross-section calculation are schematically represented in Figure 2.6.

The partonic cross-section depends on model parameters \mathcal{P} , such as couplings and masses, on the parton momenta $x_i p_i$ written as fractions x_i of the parent proton momentum p_i , and




Figure 2.6: Schematic representation of the cross-section factorisation mechanism. A parton a with momentum fraction x_1 coming from a proton with momentum p_1 interacts with a parton b with momentum fraction x_2 coming from the second proton with momentum p_2 . The partonic cross-section $\hat{\sigma}$ depends on the process under study.

finally on the renormalisation and factorisation scales μ_R and μ_F . While μ_R is an energy scale arising from a choice of a renormalisation scheme to remove low-energy divergences from the finite physical objects, the scale μ_F factorises the proton non-perturbative dynamics from the perturbative hard cross-section. The independence of the choice of these arbitrary scales in physical results are assured on the side of μ_R with the renormalisation group equations [63] and with the DGLAP equations [64–67] for μ_F , although the incomplete knowledge of the truncated perturbative series on the theory side leaves some residual dependences, obligating one to choose some values for them motivated from typical energy scales in the process under study. An uncertainty for the predicted physics observable is typically evaluated by varying these scales by factors of 0.5 and 2 (see Section 5.4).

The exact formulas for the PDFs required to calculate the total cross-sections can not be calculated perturbatively and instead have to be estimated from fits to data.¹⁰ In the traditional approach, the quark and gluon PDFs are first parametrised in terms of a set of variables at a starting scale μ_{F0} where perturbation theory is valid. These parameters are then obtained from global fits to data (e.g. from lepton-proton deep inelastic scattering processes) and the extrapolation to other energy scales μ_F is derived from the DGLAP equations. Due to their crucial importance for analyses with LHC data, there is a great tour de force from several groups [68–70] to make independent precise extractions of the PDFs at the center-of-mass energies of the LHC. As an example, the PDF4LHC PDF sets [71] are displayed in Figure 2.7 as a function of the incoming proton momentum fraction x . With all the ingredients for Equation 2.47 at hand, the integrals can be carried on numerically.

¹⁰ Estimations of the PDFs from lattice QCD are in principle possible but to date can not produce competitive results with high precision.




Figure 2.7: PDF4LHC PDF sets [71] as a function of x , the incoming proton momentum fraction, for different choices of the factorisation scale [62].

2.3.2 Overview of Higgs-boson pair production modes

The main production mode of a pair of Higgs bosons in the SM is via virtual quark-loop induced gluon-gluon fusion (ggF), as depicted in the leading order Feynman diagrams in Figure 2.8. The first, “box”, diagram depends solely on the top-Higgs coupling¹¹, while the second, “triangle”, additionally depends on the trilinear Higgs-boson self-coupling. The couplings are quoted in the vertices of the diagrams in terms of their κ -modifiers for easier understanding. The total ggF cross-section in the SM at next-to-next-to-leading order is [72]

$$\sigma_{\text{ggF}}(pp \rightarrow HH) = 31.05 \pm 3\%(\text{PDF}+\alpha_s)^{+6\%}_{-23\%}(\text{scale}+m_t) \text{ fb} \quad (2.48)$$

at $\sqrt{s} = 13$ TeV center-of-mass energy, based on the finite top-quark mass m_t approximation, and a Higgs-boson mass of 125 GeV. Due to the destructive interference between the triangle and box diagrams, the predicted cross-section changes significantly under different assumptions of the Higgs-boson coupling values, reaching a maximum of interference in the choice $\kappa_\lambda = 2.4\kappa_t$. Considering that each diagram should give different contributions to the cross-section in different regions of the di-Higgs mass spectrum, m_{HH} , it is expected that different combinations of κ -values correspond to different m_{HH} shapes and thus can result in higher or lower fractions of events with highly Lorentz-boosted Higgs bosons. This makes possible to study BSM scenarios for the couplings and set constraints on their values [73, 74].

The sub-dominant HH production mode is via vector-boson fusion (VBF), in which the Higgs bosons interact directly with the vector bosons W^\pm or Z , as shown in the leading order Feynman diagrams in Figure 2.9. The VBF process leads to two extra quarks in the final state as a result of the interaction between the hard partons and the vector bosons. In the SM, the corresponding total cross-section is [76]

¹¹ Contributions coming from lighter quark loops are significantly smaller compared to the top-quark contribution and therefore are omitted in the discussion.




Figure 2.8: Feynman diagrams at leading order representing the gluon-gluon fusion process to produce a pair of Higgs bosons. The κ -modifiers at each Higgs vertex are included in the diagrams to mark the couplings for the different vertices [75].




Figure 2.9: Feynman diagrams at leading order representing the vector-boson fusion process to produce a pair of Higgs bosons. The κ -modifiers at each Higgs vertex are included in the diagrams to mark the couplings for the different vertices [75].

$$\sigma_{\text{VBF}}(pp \rightarrow HH) = 1.726 \pm 2.1\%(\text{PDF}+\alpha_s)^{+0.03\%}_{-0.04\%}(\text{scale}) \text{ fb} \quad (2.49)$$

at next-to-next-to-next-to-leading order for $\sqrt{s} = 13$ TeV and $m_H = 125$ GeV. Even though the VBF cross-section is over an order of magnitude smaller than the ggF cross-section, the VBF production mode provides unique access to κ_{2V} , the extraction of which is a key objective of the main analysis presented in this thesis. Here and in the following, VBF di-Higgs production via the diagrams shown in Figure 2.9 will be referred as the VBF *non-resonant* production mode, given that there is no intermediate on-shell heavy particle decaying into the Higgs-boson pair. The latter is a BSM-scenario only and will be discussed in Section 2.3.4.

2.3.3 VBF non-resonant production of Higgs-boson pairs

The small VBF production cross-section of Higgs-boson pairs is heavily influenced by the strong destructive interferences between the diagrams of Figures 2.9(a) and 2.9(b). As a first approximation, the amplitude \mathcal{A} of the first two diagrams can be calculated from the hard-scattering of two longitudinally polarised vector bosons $VV \rightarrow HH$ and behaves like [60]

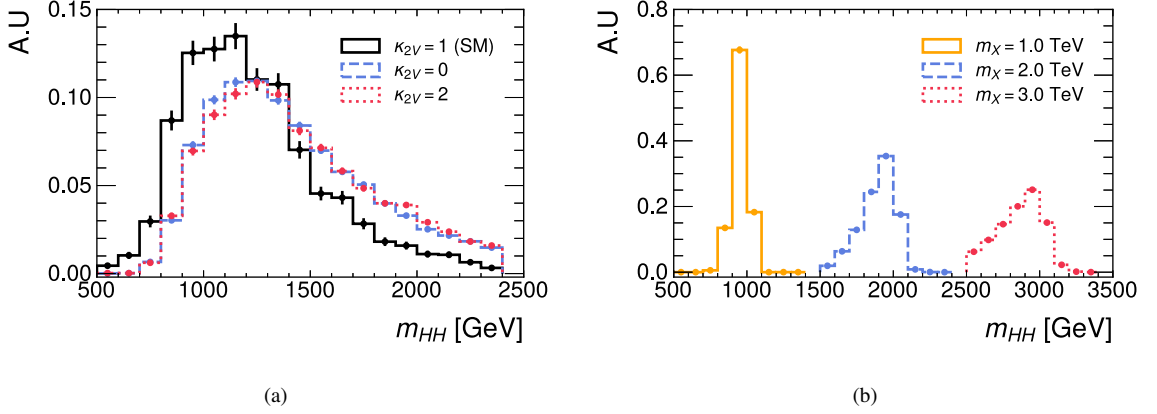


Figure 2.10: Reconstructed invariant mass distributions of the HH system. In (a) the pair of Higgs bosons are produced via non-resonant VBF under different assumptions for κ_{2V} . In (b) the BSM resonant VBF production mode for resonances with different masses m_χ . All histograms are normalised to unity.

$$\mathcal{A} \sim \frac{\hat{s}}{v^2} \left(\kappa_{2V} - \kappa_V^2 \right), \quad (2.50)$$

with sub-leading corrections depending on the masses of the bosons. In the equation above, \hat{s} is the partonic center-of-mass energy of the VV system that must be equal to the invariant mass m_{HH} of the reconstructed di-Higgs system due to energy-momentum conservation. For $\kappa_V = 1$, which is considered a safe approach given the stringent constraints already set from single Higgs-boson analyses [77, 78], it is possible to study the kinematic effects on the produced Higgs bosons when modifying the κ_{2V} coupling.

Since the total cross-section depends on the squared amplitude, it is expected that for all BSM scenarios with $\kappa_{2V} \neq 1$ the cross-section should be enhanced and the m_{HH} distribution should be shifted towards higher values, as exemplified in Figure 2.10(a) for a few values of κ_{2V} . Moreover, the positive correlation between m_{HH} and the outgoing momenta derived from kinematic relations implies that for $\kappa_{2V} \neq 1$ the reconstructed Higgs bosons should be much more energetic compared to the SM case. As will be discussed in greater detail in Section 5, exploiting an event reconstruction topology based on highly energetic Higgs bosons is of great advantage in the VBF analysis in order to have good discrimination power between signal and background events.

2.3.4 VBF resonant production of Higgs-boson pairs

The *resonant* production of a pair of Higgs-bosons via VBF is predicted by BSM-scenarios only, specifically those predicting the existence of a heavier particle that interacts both with vector bosons and the Higgs boson, and has a mass larger than $2m_H$.

In the analysis presented here, the focus are searches for resonant production of HH via




Figure 2.11: Leading order Feynman diagram for the resonant production of a pair of Higgs bosons via vector-boson fusion via an intermediate heavy X [75].

heavier scalar particles¹², as is the case for instance with the additional scalar heavy Higgs boson in the 2HDMs described in Section 2.2.1 and in the composite Higgs benchmark of Section 2.2.2. The corresponding leading order “ s -channel” Feynman diagram is depicted in Figure 2.11, where the new particle is conveniently labelled as X and has a mass m_X and a finite decay width Γ_X .

The main signature of the new particle is a localised peak around m_X in the m_{HH} distribution. This is well understood from the relativistic Breit-Wigner formula for the cross-section in the vicinity of a resonance [79]

$$\hat{\sigma}(\hat{s}) \propto \left| \frac{1}{\hat{s} - m_X^2 + im_X\Gamma_X} \right|^2, \quad (2.51)$$

valid for not very broad widths (typically up to 20% of the mass value). If Γ_X is too large, the term $m_X\Gamma_X$ should be replaced by a correct \hat{s} -dependent function encapsulating all corrections of the propagator of the particle X . The m_{HH} distributions reconstructed from detector level quantities for different choices of m_X and under the narrow width hypothesis $\Gamma_X \ll m_X$ (40 MeV) are presented in Figure 2.10(b).

2.3.5 Decay modes of the Higgs bosons

Once the value of the Higgs-boson mass is determined, all partial widths of the Higgs-boson decay modes can be predicted within the SM. Since the Higgs-boson couplings to massive fermions (bosons) depend linearly (quadratically) on the particle mass, the Higgs boson is more likely to decay into the heaviest ones, with the highest branching ratio of 58% being

¹² Considering other choices for the spin of the BSM heavy particle is also possible. The main differences compared to a scalar particle analysis would be on the angular kinematic distributions of the reconstructed Higgs bosons.

the Higgs-boson decay into a pair of bottom quarks. Decays to top quarks are kinematically forbidden given that $m_t > m_H$. Even though the weak vector bosons are more massive than the bottom quark, they should be produced in pairs with one being off-shell due to their large masses, and later decays into SM particles via electroweak interactions — reducing therefore the corresponding partial width. Decays of the Higgs boson to gluons or photons are also possible but significantly suppressed since they are only next-to-leading order loop-induced processes.

The branching ratios of the decay modes of the HH system are summarised in Figure 2.12. Fully hadronic final states, such as those from the dominant decays to four bottom quarks, suffer from a large background of QCD multijets, which is hard to simulate and instead needs to be estimated from data. Nevertheless, searches targeting VBF production benefits from the larger branching ratio of the $b\bar{b}b\bar{b}$ final state due to the low production cross-section. It is interesting to mention as well that ggF analyses targeting $b\bar{b}\tau\tau$ and $b\bar{b}\gamma\gamma$ final states are competitive to $b\bar{b}b\bar{b}$ ggF searches, despite the lower branching ratios, given the cleaner backgrounds. This is not the case for the VBF channel because of the much lower production cross-section.

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

Figure 2.12: Branching ratios of a pair of Higgs bosons with $m_H = 125$ GeV. The dominant decay mode is $HH \rightarrow b\bar{b}b\bar{b}$ and corresponds to almost 34% of the total HH decays. Numbers based on Ref. [32].

Experimental Setup and Simulation of Collision Events

Particle colliders are of fundamental importance to probe the SM and to search for phenomena predicted by BSM theories. From the collision of two particles, a whole new range of processes can be generated and reconstructed from the signals in specialised detectors. The generated processes post-collision of course depend on the energy and type of the colliding entities. When colliding leptons (e.g. in electron-positron colliders), there is a well defined center-of-mass energy determined by the beam energies that can be adjusted to enhance the production probability of specific processes. In contrast, hadron colliders open the possibility in exploring a wide range of phenomena without requiring any a priori knowledge regarding the masses and couplings values involved in the processes, making them ideal for model-independent searches. This is possible due to the composite nature of hadrons: the effective center-of-mass energy of two colliding partons depends proportionally on the momentum fraction x of each parton inside the hadron, with x ranging between 0 and 1 and its probability determined by the PDFs, as discussed in Section 2.3. Hence, hadron colliders offer a great environment to explore a vast range of energies under the same beam conditions.

The work presented in this thesis uses proton-proton collision data collected with the ATLAS detector at the Large Hadron Collider. They are introduced in Sections 3.1 and 3.2. In Section 3.3, the simulation of collision events for the relevant signal hypotheses is discussed.

3.1 The Large Hadron Collider

The Large Hadron Collider (LHC) is a 27 km circumference particle collider located at CERN, the European Organisation For Nuclear Research near Geneva, Switzerland, as depicted in Figure 3.1 [80]. The LHC tunnel was built below the Swiss-French border at depths ranging between 50 m and 175 m underground. It is designed to collide proton-proton (pp) beams or other nuclei at four different points around the accelerator ring. At these collision regions




Figure 3.1: Schematic illustration of the LHC complex [84].

are located the four main LHC experiments: ATLAS [31], CMS [81], LHCb [82], and ALICE [83]. Both ATLAS and CMS are multi-purpose experiments built to access a wide range of physics processes. Meanwhile, LHCb is specialised to b -hadron physics and ALICE is designed to study the physics of strongly interacting matter at extreme energy densities through heavy-ion collisions.

Protons are obtained through the ionisation of hydrogen atoms and accelerated in radio-frequency cavities distributed in straight segments along the accelerator complex. These cavities are metallic chambers containing an electromagnetic field adjusted to resonate at a specific frequency so that each stream of protons gains energy when passing through it. The proton beams are first accelerated to 50 MeV in the linear accelerator Linac4¹ before reaching the Proton Synchrotron Booster, which splits the beams into *bunches* and accelerates them simultaneously to 1.4 GeV in each synchrotron ring. The protons are then injected into the Proton Synchrotron ring to increase the beam energy to 25 GeV and in sequence into the Super Proton Synchrotron to further accelerate them to 450 GeV. Finally, the beams are injected into the main LHC ring in two separate streams, in opposite directions, where they reach the final energy. During the Run-2 period of the LHC, between 2015 and 2018, the final beam energy was set to 6.5 TeV, which corresponds to a total center-of-mass energy of $\sqrt{s} = 13$ TeV. After a 4-year shutdown, the LHC Run-3 program started in 2022 operating at $\sqrt{s} = 13.6$ TeV, slightly below the unprecedented maximum $\sqrt{s} = 14$ TeV for which the LHC is designed.

¹ Linac4 replaced Linac2 during the long shutdown between Run-2 and Run-3.

3.1.1 Luminosity

While the high center-of-mass energy in pp collisions provides access to a wide range of physics processes, its impact on the probabilities whether a process will occur or not is secondary. Hence, another key design parameter of particle colliders is the instantaneous luminosity

$$\mathcal{L}_i = f \frac{n_1 n_2}{A}, \quad (3.1)$$

where f is the bunch crossing frequency in the accelerator ring, n_1 and n_2 are the average number of particles in each colliding bunch, and A is the effective transversal area where the two bunches overlap. The rate in which a physics process should occur is given by the product between \mathcal{L}_i and its cross-section σ . The expected number of times N an interaction should take place during a period of time T is thus given by

$$N = \sigma \int_0^T \mathcal{L}_i(t) dt \equiv \sigma \mathcal{L}. \quad (3.2)$$

The value of the integrated luminosity \mathcal{L} works as a proxy for the total amount of data collected. In the particle physics community, it is commonly written in units of *barn*: $1 \text{ b} = 10^{-22} \text{ m}^2$.

Although the LHC was initially designed to operate with an instantaneous luminosity of approximately $10^{34} \text{ cm}^2 \text{ s}^{-1}$, this value has increased by about a factor of two during the Run-2, with the maximum peak of $2.1 \cdot 10^{34} \text{ cm}^2 \text{ s}^{-1}$ reached in May 2018, as shown in Figure 3.2(a). The integrated luminosity over time during the Run-2 is displayed in Figure 3.2(b). The total luminosity recorded by the ATLAS detector is lower than that delivered by the LHC mainly due to the time needed for the ATLAS sub-detectors to properly start for data-taking after the LHC has reached stable beams, as well as from momentaneous inefficiencies in the data acquisition system. The quality of the collected data also depends on whether all the detector components were not experiencing any major issues. During the whole Run-2 data-taking, about 96% of the data collected with the ATLAS detector were considered of good quality to be used in physics analyses. These datasets are part of the *Good Runs Lists* (GRLs) provided by the ATLAS Data Quality Group. Here, a *run* is the collection of data recorded with the ATLAS detector during a defined data-taking period, which typically corresponds to a fill of the LHC. These runs are further divided into small time intervals referred to as *luminosity blocks*, where the detector and beam conditions are expected to be approximately constant.

3.1.2 Pile-up

The high instantaneous luminosity comes at the cost of increasing pile-up, which refers to the extraneous pp interactions either within a single event or between multiple collision events. In-time pile-up encompasses all pp interactions occurring in a same bunch crossing besides the main, hard-collision of interest. Out-of-time pile-up is referred to the additional pp interactions originating from other bunch crossings with detector signals overlapping with

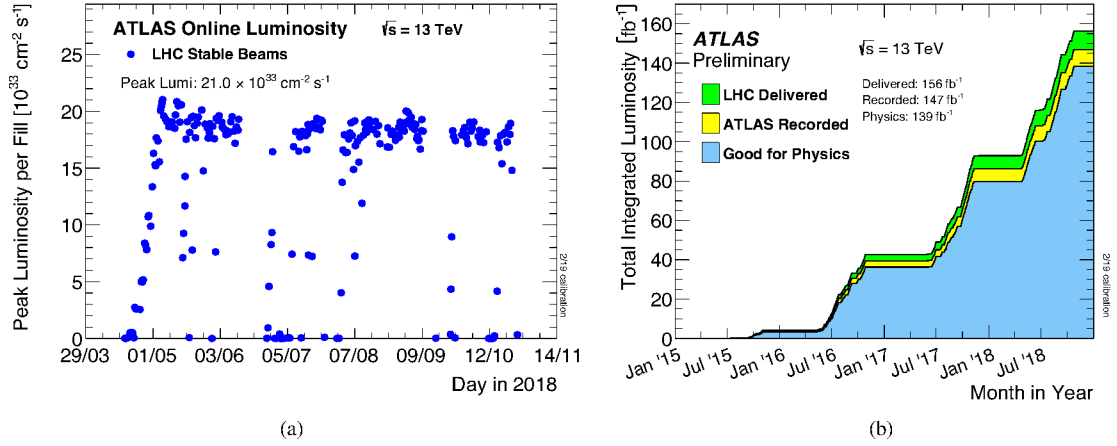


Figure 3.2: (a) Peak instantaneous luminosity delivered to the ATLAS detector during stable beams condition as a function of time in 2018. (b) Integrated luminosity as a function of the time during the Run-2 delivered by the LHC, recorded by ATLAS, and certified to be of good quality data [85].




Figure 3.3: Luminosity-weighted distribution of the average number of interactions per bunch crossing $\langle\mu\rangle$ for different years during the Run-2 data-taking [85].

the one of interest, which occur because the time elapsed between different bunch crossings is smaller than the readout time of some detector sub-systems.

The amount of pile-up in collision data is typically quantified by the average number of interactions per bunch crossing $\langle\mu\rangle$. During the Run-2 data-taking, the maximum value of $\langle\mu\rangle$ was reached in its final two years, as seen in Figure 3.3, mainly because of the continuously increasing instantaneous luminosity. In 2024, during the Run-3 data-taking, the average number of interactions per bunch crossing was about 61% higher compared to 2018 [86] — the last Run-2 year.

3.2 The ATLAS detector

A Toroidal LHC Apparatus, or in short words ATLAS, is a general purpose detector located at one of the four collision points of the LHC ring, suitable for both precision measurements and searches for a broad range of physics processes.

3.2.1 General design

Given the overall goal as a multi-purpose detector, the ATLAS design must meet an extensive list of requirements, the most important of which being a full 4π solid angle coverage around the collision point to assure that all particles resulting from the pp interactions pass through the detector and therefore can be reconstructed. Naturally, in practice this is not entirely possible because of the proton beams transversing the detector, which makes the reconstruction of the usually low-energy particles emitted at very small angles with respect to the beam line almost impossible.

The detector must allow for a precise reconstruction of different particle types, such as photons, electrons, muons, and hadrons.² The capability to identify different hadron species is also extremely desirable because of the rich hadronic environment generated from the pp collisions. The full four-momentum of these particles must also be reconstructed with high precision based on the energy deposits the particles leave on the detector elements when passing through them. Furthermore, having good control of the angular resolution is also crucial to be able to distinguish several close-by particles originating from the decay of massive particles — e.g., from the decay of highly energetic Higgs bosons.

The detector sub-systems must have fast read-out electronics to accompany the high bunch crossing frequency of the LHC without suffering from a high rate of out-of-time pile-up. In addition, because of the enormous amount of collisions happening every second, it is not possible to store permanently the data of all events. Hence, an efficient trigger system is fundamental to filter the interesting events that can be used in future analyses.

Taking all these factors into account, the 44 m long and 25 m tall cylindrical ATLAS detector consists of several sub-detectors and covers almost the entire solid angle, as seen in Figure 3.4(a). Each of the detector sub-systems are optimised to identify aspects of the stable particles coming out of the collision, as exemplified in Figure 3.4(b). The tracking system, being the first layer of the ATLAS sub-systems, captures information from charged particles close to the interaction point. The energy of electrons, photons, and hadrons are measured using the electromagnetic and hadronic calorimeters. Finally, the outermost layer is dedicated to the reconstruction of muons.

The ATLAS detector is divided into two main regions based on its design: the barrel region is the central, cylindrical part that surrounds the beam axis and covers particles emitted at large angles relative the beam axis; the end-cap, forward part of the detector, has disc-shaped

² Neutrinos do not interact with detector elements.




Figure 3.4: (a) The ATLAS detector [31]. (b) Cross-sectional view of the ATLAS barrel to exemplify the interaction between the different particle types with the detector elements [87].

sections at both ends of the barrel and is built to detect particles moving close to the beam axis.

The details of each of these detector sub-systems and the ATLAS trigger system will be discussed in the following sections.

ATLAS coordinate system

ATLAS uses a right-handed coordinate system with the z -axis parallel to the beam line and the x - y plane referred as the transverse plane with the y -axis placed vertically. By convention, the origin is the nominal interaction point. In cylindrical coordinates, the azimuthal angle $\phi = \arctan(y/x)$ is measured around the beam axis. The polar angle θ from the beam axis is determined from the pseudorapidity η via

$$\eta = -\ln \left[\tan \left(\frac{\theta}{2} \right) \right], \quad (3.3)$$

with $\eta = 0$ indicating $\theta = 90^\circ$ and the limit $\eta \rightarrow \infty$ indicating $\theta \rightarrow 0^\circ$. For massless particles or in the relativistic limit $v/c \rightarrow 1$, with v being the particle velocity and c the speed-of-light, the pseudorapidity approaches the rapidity

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right), \quad (3.4)$$

where E is the particle energy and p_z the z -component of the particle three-momentum. Using the pseudorapidity instead of θ is particularly interesting because differences in η are Lorentz invariant under boosts along the beam axis.

Since the z -axis is oriented along the beam line, the p_z -component of the resulting particles reaching the detector at a hadron collider is unknown. This is because the interactions of interest are between the partons inside of the hadrons and they carry an a priori unknown momentum fraction. Hence, the relevant momentum variable for particle reconstruction is the transverse momentum

$$p_T = \sqrt{p_x^2 + p_y^2} = |\mathbf{p}| \sin \theta = \frac{|\mathbf{p}|}{\cosh \eta}. \quad (3.5)$$

At very good approximation, the transverse components of the colliding partons is zero, meaning that the vector sum of the transverse components of the final-state particles must also be zero. The vector sum of the transverse momenta of all reconstructed particles is thus used to determine the missing transverse momentum of an event (see Section 4.4).

Together with the energy, azimuthal angle, and pseudorapidity, the transverse momentum is used to write the four-momentum of reconstructed particles in the form of

$$(E, p_x, p_y, p_z) = (E, p_T \cos \phi, p_T \sin \phi, p_T \sinh \eta). \quad (3.6)$$

Finally, the angular separation between two reconstructed particles (or objects) i and j is measured based on the pseudorapidity and the azimuthal angle via the distance parameter

$$\Delta R(i, j) = \sqrt{(\eta_i - \eta_j)^2 + (\phi_i - \phi_j)^2}. \quad (3.7)$$

3.2.2 The Inner Detector

The Inner Detector is the innermost sub-system of the ATLAS detector, consisting of the Pixel Detector, Semiconductor Tracker (SCT), and the Transition Radiation Tracker (TRT), covering a pseudorapidity range up to $|\eta| = 2.5$ and spanning in a radii up to $R = 1082$ mm, as displayed in Figure 3.5. The entire Inner Detector is immersed in a 2 T magnetic field of surrounding solenoids pointing along the z -axis to bend the trajectory of charged particles and thereby enabling the measurement of their momenta. Being the closest sub-system to the interaction point, the Inner Detector is of fundamental importance for the reconstruction of the primary vertex, known as the point where the main, hard-collision is understood to have happened (see Section 4.1.2). Furthermore, all the information collected from secondary vertices related to the decays of particles at a distance from the interaction point plays crucial role for the machine learning algorithms used to identify b -hadron decays — one the most important tools for the analysis presented in this thesis.

The innermost part of the Inner Detector is the Pixel Detector. Its innermost layer referred to as insertable B-layer (IBL) is located at 33 mm of the interaction point and is surrounded by three cylindrical layers of silicon pixels. With more than 90 million pixel sensors, the IBL and the outermost silicon layers have a high granularity to allow precise measurements of charged particles. The IBL has the highest resolution with pixel sensors of pitch $50 \times 250 \mu\text{m}^2$ in the transverse and z -directions, while the outermost silicon layers have a resolution of $50 \times 400 \mu\text{m}^2$.

The SCT surrounds the Pixel Detector. It is a silicon-microstrip detector spanning radii from 299 mm to 514 mm, segmented in four cylindrical layers and about 6.3 million read-out channels. The SCT additionally has 9 silicon disks displayed at each end of the cylindrical layers to expand the SCT coverage to a pseudorapidity up to $|\eta| = 2.5$. The intrinsic resolution in the SCT is of $17 \times 580 \mu\text{m}^2$ in the transverse and z -directions.

Finally, the outermost layer of the Inner Detector is the TRT. It consists of about 50,000 straw tubes filled with a gas mixture of Xe, CO₂, and O₂, surrounded by a material that induces transition radiation when a charged particle passes through it. The TRT spans radii from 554 mm to 1082 mm and goes up to $|\eta| = 2.0$ in pseudorapidity, allowing for the reconstruction of charged particles far from the interaction point. The intrinsic resolution is lower compared to the Pixel Detector and SCT, with an accuracy of $170 \mu\text{m}$ in the transverse direction and no measurement in the z -direction. However, the number of hits left by charged particles is much higher in the TRT, which contributes significantly to the measurement of the particle momentum, and the larger radius at which it is located helps determining the particle curvatures more accurately. Furthermore, the transition radiation is enhanced when




Figure 3.5: Illustration of the ATLAS Inner Detector. The Pixel Detector is the closest to the beam axis. In sequence, there are the Semiconductor Tracker and the Transition Radiation Tracker. The IBL is not displayed in the figure as it was included in the Pixel Detector only after the end of the Run-1 [31].

induced by ultrarelativistic particles, providing therefore valuable information for particle identification (e.g., if it was an electron or a pion).

3.2.3 The Electromagnetic and Hadronic Calorimeters

The calorimeter system surrounds the Inner Detector and covers a region up to $|\eta| = 4.9$. It is composed of three main blocks: the electromagnetic calorimeter (ECAL), which is used to measure the energy of electrons and photons; the hadronic calorimeter (HCAL), meant to determine the energy of hadrons; and the forward calorimeter (FCAL), designed to detect particles emitted at small angles relative to the interaction point. A schematic representation of the calorimeter system is displayed in Figure 3.6. ATLAS utilises a sampling calorimeter design, meaning it is built using alternating layers with absorbing and sensitive materials.

The Electromagnetic Calorimeter

The ECAL is split into barrel (with coverage $|\eta| < 1.475$) and end-cap (with coverage $1.374 < |\eta| < 3.2$) regions, extending up to 2.25 m from the interaction point, and subdivided into three main layers in the radial direction consisting of cells with different




Figure 3.6: Illustration of the ATLAS calorimeter system [31]. The Electromagnetic Calorimeter surrounds the Pixel Detector and is overlayed by the Hadronic Calorimeter. The calorimeter layout and technologies are different in the barrel and end-cap regions. The main components are pointed in the figure.

itches, as illustrated in Figure 3.7. The innermost layer has the highest granularity of $\Delta\eta \times \Delta\phi = 0.003 \times 0.1$ to enhance the resolution of the η -profile and help discriminate showers initiated by single photons or, e.g., by neutral pion decays $\pi^0 \rightarrow \gamma\gamma$ with typically low η -separation between the final-state photons. The middle and outermost layers have granularities of $\Delta\eta \times \Delta\phi = 0.025 \times 0.0245$ and $\Delta\eta \times \Delta\phi = 0.05 \times 0.0245$, respectively.

The ECAL consists of lead and liquid argon (LAr) layers to play the role of absorbing and sensitive materials, respectively. The lead material induces electromagnetic showers via bremsstrahlung and e^+e^- pair-production when electrons and photons pass through it. The electromagnetic showers when passing through the LAr layers then induce signal ionisation that can be used to measure the shower energy. To ensure that the electromagnetic shower will be fully absorbed³ before reaching the HCAL, the thickness of the ECAL is chosen to be much larger than the radiation length X_0 of its materials, defined as the distance an electron has to travel inside the material before losing $1/e \approx 37\%$ of its incident energy. In the barrel region, the ECAL extends to $24 X_0$ (see again Figure 3.7).

The Hadronic Calorimeter

The HCAL surrounds the ECAL and is subdivided into a tile calorimeter and two end-cap calorimeters. The tile calorimeter uses as active material radially-oriented scintillating tiles and has iron absorber plates. It is subdivided in three cylinders, covering a pseudorapidity range up to $|\eta| = 1.7$. The end-cap hadronic calorimeters are situated in the pseudorapidity

³ The shower stops when either losses from other processes besides bremsstrahlung start to dominate or, in case of photons, when there is no more enough energy to produce e^+e^- pairs.




Figure 3.7: Schematic representation of the electromagnetic calorimeter. The thickness of each layer is given in terms of radiation length X_0 (see text) [31].

range $1.5 < |\eta| < 3.2$ and have copper plates as absorbing material and LAr as the sensitive material. The overall granularity in the HCAL is much coarser compared to the ECAL, with a typical cell pitch of the order of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$. While the principle of the shower energy measurement is similar to the ECAL using intercalated absorbing and sensitive layers, the hadronic shower in the HCAL is primarily initiated via the strong interaction between the incoming particles and the material atomic nuclei.

The hadronic shower is considerably more complex than the electromagnetic shower because the strong interaction between hadrons and the detector material can convert the incoming energy into forms of energy harder to be detected. Some examples include intra-nuclear cascades, in which the incoming hadron frees neutrons and protons from the material nuclei that might leave the calorimeter undetected; the generation of electromagnetic shower components, e.g. from $\pi^0 \rightarrow \gamma\gamma$ decays; and weak hadronic decays like $\pi^- \rightarrow \mu^- \nu_\mu$ that produce neutrinos — particles undetectable by the detector. As a result, the calorimeter response to hadrons is much lower compared to electrons and photons.

The Forward Calorimeter

The FCAL is also a hadronic calorimeter but situated in a region of large pseudorapidity, up to $|\eta| = 4.9$, to allow the reconstruction of particles emitted at small angles from the beam line. It also works as an additional shield to prevent particles other than muons from reaching the Muon Spectrometer.

3.2.4 The Muon Spectrometer

The Muon Spectrometer (MS) is the outermost sub-system of the ATLAS detector, has a coverage up to $|\eta| = 2.7$, and is dedicated to the reconstruction and identification of muons. Since the amount of energy lost via bremsstrahlung is inversely proportional to the square of the incoming particle mass, muons leave only a very small amount of energy in the calorimeter system. This allows for precision reconstruction of the muon trajectory and momentum using the MS standalone. It is possible however to increase the momentum resolution even more when using combined information from the Inner Detector. (Details about charged-particle reconstruction are given in Section 4.1.)

The structure of the MS serves two purposes: one is the precise muon trajectory reconstruction using the Monitored Drift Tubes (MDTs) and the Cathode Strip Chambers (CSCs); the other is to provide fast response to the ATLAS trigger system (see Section 3.2.5) using the dedicated Resistive Plate Chambers (RPCs) and Thin Gap Chambers (TGCs). An illustration of the MS is shown in Figure 3.8.

The MDTs consist of aluminium tubes filled with a gas mixture of Ar and CO₂ and arranged in three concentric cylindrical shells with radius of 5 m, 7.5 m, and 10 m from the interaction point. In the end-cap region the shells are arranged as large wheels transversal to the beam axis. In the forward $|\eta| > 2.0$ region of the detector, the cylindrical MDTs are replaced by the CSCs and arranged in four concentric disks around the beam axis. The CSCs are multiwire proportional chambers with cathode strips and have a much higher granularity than the MDTs. This allows having better control of muon reconstruction in the high η -region, where larger background rates from collisions or interactions other than the main hard-collision of interest are expected. The muon reconstruction principle in the MS is similar to the track reconstruction in the TRT of the Inner Detector: a magnetic field with average strength of 0.5 T is immersed in the MS to bend the muon trajectory and allow momentum measurement; the muons leave signals (or hits) in the MS by ionising the gas when passing through the tubes.

The trigger chambers are built with three layers of RPCs — consisting of narrow gas gaps between parallel resistive plates and further separated by insulating spacers — in the barrel $|\eta| < 1.05$ region, and three layers of TGCs in the end-cap $1.05 < |\eta| < 2.4$ regions, which are also multiwire proportional chambers. This system is designed to provide fast information on muon-tracks that can be used by the ATLAS trigger system to decide whether an event should be stored permanently or not. Additionally, the trigger chambers with its high time resolution can be used to distinguish different bunch crossings.




Figure 3.8: Schematic representation of the ATLAS Muon Spectrometer [31]. The Monitored Drift Tubes and Cathode Strip chambers are dedicated for precise muon trajectory reconstruction, while the Resistive Plate Chambers and Thin Gap Chambers are used primarily to provide fast response to the ATLAS trigger system.

3.2.5 Trigger system

With a bunch crossing frequency of 25 ns, 40 million collision events happen every second in the LHC during stable beams conditions, which would amount to an average of 40 TB of data to be stored per second, based on the event size for the current detector. The huge majority of these events, however, originate from non-interesting inelastic and elastic scattering. Hence, the ATLAS Trigger and Data Acquisition (TDAQ) system is designed to select the relevant events and reduce the data rate to be permanently stored for physics analyses to approximately 1 kHz. To be able to make a final decision on whether to keep or not an event based on a refined set of selection criteria and using as much detector information as possible, the ATLAS trigger system works in two consecutive stages: first, the hardware-based Level-1 (L1) trigger with a small decision time of $2.5 \mu\text{s}$ reduce the data rate from 40 MHz to 100 kHz; next, a software-based high-level trigger (HLT) employs a finer selection on the events passing the L1 trigger and further reduce the event rate to 1 kHz.

Level-1 Trigger

The first stage of the ATLAS trigger system is the L1 trigger, for which only $2.5 \mu\text{s}$ of processing time are allocated to make a decision on whether the event should be passed to the HLT. The L1 trigger does not perform any computationally expensive reconstruction of physics objects, but instead uses only hardware-based trigger systems in the calorimeter and

in the trigger chambers of the MS to detect potential signatures of highly energetic electrons, photons, hadrons, or muons. Regions in the calorimeter and in the MS where any of these objects activate the L1 trigger are defined as regions of interest (RoIs). If the event is accepted by the L1 trigger, these RoIs are additionally passed to the HLT system. The L1 trigger reduces the incoming data rate to an average of 100 kHz.

High-level Trigger

In case an event is accepted by the L1 trigger, the information from the remaining detector sub-systems are read out and sent to the software-based HLT system. It uses information from the calorimeter and MS with finer granularity and employs track reconstruction in the RoIs to reconstruct the relevant physics objects with increased resolution. If the event passes any of the several selection combinations (trigger chains), the event is accepted by the HLT system with probability $1/n$, where n is a prescale⁴ of the fired trigger chain to control the rate of accepted events.

The time allocated for HLT processing is, on average, approximately 500 ms. In the end, the rate of data stored permanently for physics analyses is about 1 kHz.

3.3 Dataset and simulation of collision events

The dataset used in the analysis presented here is the full Run-2 dataset collected with the ATLAS detector in pp collisions between 2015 and 2018, corresponding to an integrated luminosity of 140 fb^{-1} . All luminosity blocks in this dataset are part of the GRLs provided by the ATLAS Data Quality Group, ensuring that the data used for the physics analysis satisfy some quality criteria, such as not suffering from temporary sub-detector malfunctions.

Comparing the observations made in data to some expectation is the core of the physics analyses. These expectations encompass how the observations should look like either in the presence of only the SM hypothesis, or in the presence of alternative BSM signals. Hence, carefully simulating the particle interactions in pp collisions under different hypotheses is of fundamental importance.

3.3.1 Simulation of physics processes

The simulation of pp collision events relies on the Monte Carlo (MC) method, a statistical technique that generates possible outcomes based on random sampling from probability distributions. In the context of experimental particle physics, this is done as follows.

First, the hard interaction between two colliding partons is simulated through the calculation of the matrix-element (ME) using perturbation theory at a fixed order in the coupling parameters. This is only possible due to the large momentum transfers involved in the hard

⁴ Trigger chains with $n = 1$ are called *unprescaled triggers*.

interaction, meaning that the energy scales involved in the process are within the region where the use of perturbative QCD is applicable.

The ME process is then modified by accounting for additional processes, typically at softer scales, like gluon radiation and $q\bar{q}$ -pairs production. These processes can not be described in the context of perturbation theory, but rather should be calculated approximately. This stage is known as parton showering (PS) and, a priori, can be computed by general-purpose generators like Pythia [88] or Herwig [89] that are designed for both the ME and PS calculations. In LHC analyses, however, the ME calculation stage is typically done by dedicated ME generators such as MadGraph [90] and Powheg [91] and later injected into one of the PS generator softwares. The PS stage also handles the hadronisation of the partons into colour-neutral particles and performs the decay of (unstable) particles with a lifetime smaller than 10 ps that decay before reaching the detector elements.

The hard interaction is always accompanied by several softer interactions between the additional partons in the protons that do not take part in the main scattering process, while still being colour-connected to it. These interactions are referred as the underlying event (UE) and must also be included in the simulation. The UE is modelled phenomenologically and depends on a series of parameters of the ME generators. Choosing these parameters is known as *tune*, and this is accomplished by comparing simulated events to a set of data containing predominantly soft collision events.

Moreover, to include the effects from pile-up into the MC generation, simulated QCD multijets events are overlaid on top of the hard-scattering part of the simulation using a pile-up profile similar to the one observed in real data.

Finally, to mimic as much as possible how real collision data is captured by the detector, all final stable particles produced in the event generation are passed through a full simulation of the ATLAS detector [92] based on the Geant4 toolkit [93]. The simulated energy deposits left by the stable particles reaching the detector are digitised into information like voltages and currents equivalent to what is seen by the detector in real collision events. This allows processing real and simulated data using the same reconstruction software.

To account the differences each Feynman diagram contribute to the process under study when simulating events using beyond leading-order calculations, an event weight is assigned to each simulated event individually. These event weights should be further modified to include for instance the latest and most precise cross-section predictions; the integrated luminosity of the data period the MC simulation is meant to represent; pile-up weight corrections to account for different pile-up profiles observed in simulation and in real data; etc. All these factors must be included at analysis-level for a reliable comparison between data and simulation.

Different MC generators are better suited for different processes. In the following sections an overview of the setup to generate the relevant signal hypotheses for this analysis is presented. Expected background events were also simulated for initial studies, but the final background model was obtained using data-driven techniques, as explained in further details in Section 5.2.

3.3.2 Non-resonant production of $HH \rightarrow b\bar{b}b\bar{b}$

One of the key goals of this analysis is to test the agreement between data and the presence of a non-resonant production of Higgs-boson pairs decaying into four bottom quarks under different hypotheses for the values of some Higgs-boson couplings, namely κ_{2V} and κ_λ .

The relevant nominal ggF samples were simulated using the `Powheg Box v2` generator at next-to-leading order (NLO) in QCD, with full NLO corrections from finite top-mass effects⁵. The PDF4LHC PDF set [71] was used. The parton shower and hadronisation were simulated using `Pythia8.244` with the A14 set of tuned parameters. In order to estimate the systematic uncertainties due to the PS simulator choice, alternative samples were simulated interfacing the ME calculation with `Herwig7`. (For further details on the PS systematic uncertainties, see Section 5.4.)

VBF signal samples were simulated using `MadGraph 2.7.3` at LO in QCD interfaced with `Pythia8.244` and A14 tune, and using the NNPDF3.0NLO PDF set [94]. Alternative samples were also generated using `Herwig7` for the parton shower.

All samples used the `EvtGen1.7.0` program [95] to model the b - and c -hadron decays. The mass of the Higgs boson was set at 125 GeV for all samples.

A priori, it would be required to generate a signal sample for each combination of κ -values desired to be tested against observed data. While it is theoretically possible to generate simulated collision events for any choice of κ_{2V} and κ_λ (and the other Higgs-boson couplings modifiers), it is a computationally expensive and time-consuming task. Hence, only a handful signal samples are actually simulated and a reweighting technique is employed to model the signal hypotheses across the whole coupling modifiers space.

ggF reweighting

As a starting point, consider the ggF production mode of Higgs-boson pairs. The differential cross-section contributions can be expressed in terms of the box and triangle diagram contributions \mathcal{M}_Δ and \mathcal{M}_\square (see leading-order Feynman diagrams in Figure 2.8). Parametrised in the terms of κ_λ and the Higgs boson to top-quark pairs coupling modifier κ_t , the differential cross-section reads

$$\begin{aligned} \frac{d\sigma(\kappa_\lambda, \kappa_t)}{dx} &= \left| \kappa_t^2 \mathcal{M}_\square + \kappa_t \kappa_\lambda \mathcal{M}_\Delta \right|^2 \\ &= \kappa_t^4 a_1(x) + \kappa_t^2 \kappa_\lambda^2 a_2(x) + \kappa_t^3 \kappa_\lambda a_3(x), \end{aligned} \quad (3.8)$$

where x is some (kinematic) variable and a_1 , a_2 , and a_3 depend only on x . In this analysis, the most relevant variable x is the model output described in Section 5.3. With the knowledge of the differential cross-section for three different $(\kappa_\lambda, \kappa_t)$ points, it is possible to determine the functions a_i by solving a system of linear equations.

⁵ In many calculation the top-quark mass is set as much larger than other mass values. This simplifies significantly the calculations.

Since HH searches have much less sensitivity to κ_t compared to single-Higgs ttH analyses, in the following the top coupling modifier is kept fixed at $\kappa_t = 1$. The basis functions are then generated with $\kappa_\lambda = 0, 1, 20$, in accordance with other ATLAS HH analyses [73, 96–99]. The system of equations to be solved is

$$\begin{cases} \frac{d\sigma(0)}{dx} &= a_1(x) \\ \frac{d\sigma(1)}{dx} &= a_1(x) + a_2(x) + a_3(x) \\ \frac{d\sigma(20)}{dx} &= a_1(x) + 400a_2(x) + 20a_3(x) \end{cases}, \quad (3.9)$$

which leads to the reweighting formula for arbitrary κ_λ :

$$\frac{d\sigma(\kappa_\lambda)}{dx} = \left(1 - \frac{21}{20}\kappa_\lambda + \frac{1}{20}\kappa_\lambda^2\right) \frac{d\sigma(0)}{dx} + \frac{\kappa_\lambda(20 - \kappa_\lambda)}{19} \frac{d\sigma(1)}{dx} + \frac{\kappa_\lambda(\kappa_\lambda - 1)}{380} \frac{d\sigma(20)}{dx}. \quad (3.10)$$

VBF reweighting

The same approach is used to evaluate the VBF differential cross-section for arbitrary κ_{2V} , κ_λ , and κ_V coupling modifiers, which can be expressed in terms of the contributions from three Feynman diagrams:

$$\frac{d\sigma(\kappa_{2V}, \kappa_\lambda, \kappa_V)}{dx} = \left| \kappa_{2V}\mathcal{M}_{2V} + \kappa_V^2\mathcal{M}_V + \kappa_V\kappa_\lambda\mathcal{M}_\lambda \right|^2, \quad (3.11)$$

with diagram amplitudes identified accordingly to their coupling dependencies, as depicted in the leading-order Feynman diagrams in Figure 2.9. Expanding the equation above yields six independent functions $a_i(x)$ that do not depend on the coupling modifiers. Hence, six samples must be chosen to form the basis functions of the VBF reweighting technique. Even though any choice of six independent samples would theoretically work, it is important to choose samples that provide a good coverage in the whole kinematic space expected to be explored in the analysis to avoid mis-modelling with non-physical reweighted differential cross-section. The chosen basis of the parameter space reads $(\kappa_{2V}, \kappa_\lambda, \kappa_V) = \{(1, 1, 1), (1.5, 1, 1), (1, 2, 1), (1, 10, 1), (1, 1, 0.5), (1, -5, 0.5)\}$ and has already been used in previous ATLAS HH searches [73, 96, 97, 99]. This yields the following reweighting formula for the VBF differential cross-section:




Figure 3.9: Comparison of the m_{HH} distribution for the VBF $\kappa_{2V} = 0$ signal sample if using directly simulated events or the reweighting technique of Equation 3.12.

$$\begin{aligned}
 \frac{d\sigma(\kappa_{2V}, \kappa_\lambda, \kappa_V)}{dx} = & \left(\frac{68\kappa_{2V}^2}{135} - 4\kappa_{2V}\kappa_V^2 + \frac{20\kappa_{2V}\kappa_V\kappa_\lambda}{27} + \frac{772\kappa_V^4}{135} - \frac{56\kappa_V^3\kappa_\lambda}{27} + \frac{\kappa_V^2\kappa_\lambda^2}{9} \right) \frac{d\sigma(1, 1, 1)}{dx} \\
 & + \left(-\frac{4\kappa_{2V}^2}{5} + 4\kappa_{2V}\kappa_V^2 - \frac{16\kappa_V^4}{5} \right) \frac{d\sigma(1.5, 1, 1)}{dx} \\
 & + \left(\frac{11\kappa_{2V}^2}{60} + \frac{\kappa_{2V}\kappa_V^2}{3} - \frac{19\kappa_{2V}\kappa_V\kappa_\lambda}{24} - \frac{53\kappa_V^4}{30} + \frac{13\kappa_V^3\kappa_\lambda}{6} - \frac{\kappa_V^2\kappa_\lambda^2}{8} \right) \frac{d\sigma(1, 2, 1)}{dx} \\
 & + \left(-\frac{11\kappa_{2V}^2}{540} + \frac{11\kappa_{2V}\kappa_V\kappa_\lambda}{216} + \frac{13\kappa_V^4}{270} - \frac{5\kappa_V^3\kappa_\lambda}{54} + \frac{\kappa_V^2\kappa_\lambda^2}{72} \right) \frac{d\sigma(1, 10, 1)}{dx} \\
 & + \left(\frac{88\kappa_{2V}^2}{45} - \frac{16\kappa_{2V}\kappa_V^2}{3} + \frac{4\kappa_{2V}\kappa_V\kappa_\lambda}{9} + \frac{152\kappa_V^4}{45} - \frac{4\kappa_V^3\kappa_\lambda}{9} \right) \frac{d\sigma(1, 1, 0.5)}{dx} \\
 & + \left(\frac{8\kappa_{2V}^2}{45} - \frac{4\kappa_{2V}\kappa_V\kappa_\lambda}{9} - \frac{8\kappa_V^4}{45} + \frac{4\kappa_V^3\kappa_\lambda}{9} \right) \frac{d\sigma(1, -5, 0.5)}{dx}
 \end{aligned} \tag{3.12}$$

The bin-by-bin reweighting approach of Equation 3.12 is validated by comparing the reconstructed invariant mass of the HH system, i.e. $x = m_{HH}$, for $\kappa_{2V} = 0$ using simulated and reweighted events, as displayed in Figure 3.9. The details on the analysis strategy to reconstruct the HH system are given in Chapter 5.

3.3.3 Resonant production of $HH \rightarrow b\bar{b}b\bar{b}$

The resonant production of Higgs-boson pairs is simulated at LO in QCD using MadGraph5 version 2.6.1 interfaced with Pythia8.244 with its A14 tune and EvtGen1.7.0. The nominal NNPDF3.0NLO PDF set was used. The scalar resonance was modelled based on a type-II 2HDM (see Section 2.2.1), generating only the heavy scalar \tilde{H} for the signal process via VBF and forced to decay into a pair of SM Higgs bosons with 100% branching ratio. None of the other particles predicted 2HDMs are considered in the event generation. For convenience, in the following the scalar resonance will be referred as X .

In the nominal resonant search, the narrow-width approximation is assumed for the heavy scalar resonance. Hence, the width of the resonance Γ_X was set to 40 MeV — a value much smaller than the detector resolution of 5-6% of the particle mass. A total of 13 signal samples were generated, with the mass m_X of the heavy scalar ranging from 1.0 TeV to 5.0 TeV.

A few samples were also generated under a broad-width hypothesis. These were produced based on the composite Higgs benchmark of Section 2.2.2, with masses ranging from 1.2 TeV to 2.0 TeV and fixed relative width $\Gamma_X/m_X = 20\%$.

In summary, the following signal hypotheses were generated for the search of a resonant production of Higgs-boson pairs:

- 2HDM, narrow width $\Gamma_X = 40$ MeV: $m_X \in [1.0, 1.1, 1.2, 1.3, 1.4, 1.5, 1.6, 1.8, 2.0, 2.5, 3.0, 4.0, 5.0]$ TeV
- Composite Higgs model, $\Gamma_X = 20\% m_X$: $m_X \in [1.2, 1.5, 1.8, 2.0]$ TeV.

Additional samples with the alternative Herwig7 PS were also generated for a few mass hypotheses to estimate the systematic uncertainty related to the choice of the PS model.

3.3.4 QCD multijets and $t\bar{t}$ background

The main background of the analysis consists of QCD multijets and $t\bar{t}$ events. Here, the dominant QCD multijets, composing about 90% of the total background, stands for the events containing multi b -jets or jets that were mis-identified as b -jets, all originating from QCD interactions. The sub-dominant $t\bar{t}$ -background comprises all events where a $t\bar{t}$ -pair is produced and each top-quark decays into a b -quark and a W -boson that further decay into a b -quark and another light-quark.

While the final background estimate was obtained using a data-driven technique (see Section 5.2), simulated events were still used in the initial analysis studies to optimise the event selection. All background samples were generated using Powheg Box v2 interfaced with Pythia8.244 with its A14 tune and EvtGen1.7.0. The multijets ($t\bar{t}$) samples used the NNPDF2.3LO (NNPDF3.0NLO) PDF set.

Object Definition and Event Reconstruction

Using the information of all of the ATLAS sub-detectors, physics objects like leptons, photons, and hadronic jets can be reconstructed. They form the main set of inputs used downstream in physics analyses to identify the collision events with detector signatures consistent with the signal process of interest.

In Section 4.1, the reconstruction of charged-particle trajectories — also known as tracks — using information from the Inner Detector is presented. The introduction of algorithms to reconstruct hadronic jets and the identification of jets originating from b -quarks and b -quark pairs — all key components for this $HH \rightarrow b\bar{b}b\bar{b}$ search — are presented in Section 4.2. An overview of the identification of leptons and photons and the determination of the missing transverse energy in events are given in Sections 4.3 and 4.4, respectively.

4.1 Tracks

By definition, tracks represent the trajectories of charged particles that pass through the Inner Detector. They are reconstructed from the combined information of energy deposits left by particles in the Pixel Detector, SCT, and in the TRT extension. Each track is described by its momentum in polar coordinates, a pair of impact parameters, and a reference point, as depicted in Figure 4.1. The set of track parameters are:

- q/p : the ratio between the particle charge and reconstructed momentum. It quantifies the particle curvature inside the detector due to the magnetic field.
- θ : polar angle of the track in the longitudinal plane.
- ϕ : azimuthal angle of the track in the transverse plane.




Figure 4.1: Schematic perigee representation of the track parameters [100]. Here, the perigee representation describes the particle trajectory at its point of closest approach to the reference point.

- z_0 : longitudinal impact parameter, defined as the minimum distance between the particle trajectory and the reference point, in the longitudinal plane.
- d_0 : transverse impact parameter, defined as the minimum distance between the particle trajectory and the reference point, in the transverse plane.

Both the primary vertex of the collision event, which represents the point in which the main pp hard-scatter took place (see Section 4.1.2), and the center of the beam spot are considered good choices for the reference point. The former however is commonly used in most ATLAS analyses.

4.1.1 Track reconstruction

When an ionising particle passes through a silicon sensor in the Pixel or SCT detectors, it deposits an energy that may be interpreted as a signal. The higher the energy deposited, the longer the signal remains and the higher its intensity. If the time elapsed between the time t_0 at which the signal passes a certain threshold and the time t_1 at which the sensor returns to a base value exceeds a certain duration, the sensor is said to be activated by a charged particle. Since particles typically cross the Pixel and SCT layers with non-zero incident angle, in the vast majority of cases it will activate more than one single sensor, with an average of two sensors in the central part of the detector and going up to three sensors in the large- η end-barrels region [101]. Single activated sensors or adjacent activated sensors are then identified as clusters and represent three-dimensional space-points — with position and uncertainties depending on detector geometry and sensor pitch — indicating where charged particles traversed the detector layers. These space-points are interpreted as “hits” left by charged particles in the Inner Detector.




Figure 4.2: Illustration of the ATLAS track reconstruction steps. Each yellow dot represents a hit in the Inner Detector. The solid curves represent selected triplets of hits to form the initial track seeds. The dashed curve represents the adding of extra hits to the particle trajectory compatible with the original track seed. All particle trajectories passing a stringent set of requirements are then extended to the TRT to improve the resolution of the track momenta [102].

With all clusters hits defined, triplets of hits in the silicon layers are used as seeds for the track reconstruction and provide the first crude estimation of the track parameters. Additional hits from different layers are added to the track seed if they are compatible with the original particle trajectory, as shown by the solid and dashed blue curves in Figure 4.2. Moving forward in the track reconstruction algorithm, a combinatorial Kalman filter [101] is implemented to identify and include more space-points compatible with the first track seed and refine the determination of the track parameters. At this stage, there is a whole collection of track candidates, with each cluster potentially being used for the reconstruction of more than one single track. These shared clusters might exist either because of fake tracks, which appear as candidates during the reconstruction step but do not represent real charged-particle trajectories, or because the cluster indeed has been activated by more than one ionising particle. The differences between standard single-particle clusters and shared clusters can be visualised in Figure 4.3.

Although shared clusters likely indicate fake or low quality tracks, it is expected that tracks will share hits in detector regions with high particle densities — commonly referred to as a dense hadronic environment. Properly identifying these shared clusters is therefore a crucial task to reduce the number of fake tracks without affecting the reconstruction of real charged-particle trajectories.

Other track properties can also indicate if the track is more likely to represent a real particle

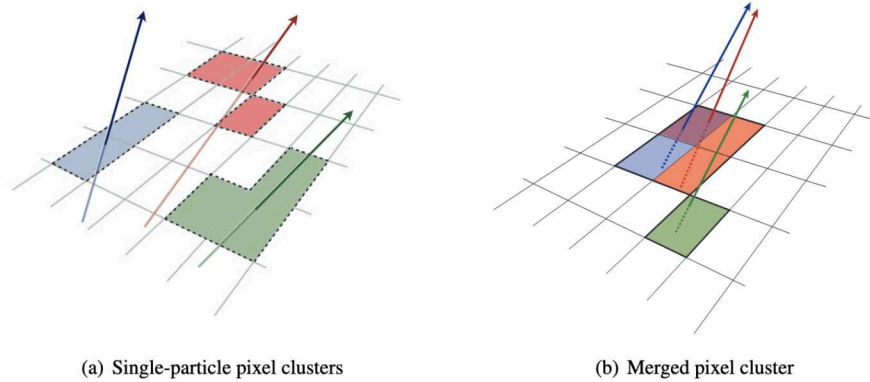


Figure 4.3: (a) Several clusters, each activated by a single ionising particle. (b) Cluster activated by several ionising particles. Different colors refer to the energy deposits of different particles [100].

trajectory or not. Hence, in the stage known as ambiguity solver, a *score* is assigned to each track candidate in order to select which tracks should be part of the final track collection, with the goal to reduce as much as possible the rate of fake tracks while maintaining great track reconstruction efficiency. The score favours more energetic tracks, penalises candidates containing holes in the silicon layers — intersections of the trajectory with a detector element without a matching cluster —, and reflects the quality of the initial estimations of the impact parameters. A shared-cluster-penalty is applied if the shared cluster is not identified as “merged” by a specialised multivariate algorithm. For this purpose, a set of dense neural networks [103] receiving information of the energy deposits in the clusters, incident angle of the ionising particles in the silicon layer, and relative position of the silicon sensors in the cluster, is used to identify if the cluster was activated by one, two, or three or more charged particles. If a cluster is identified as merged, i.e., activated by more than one particle, a mixed density network [104] is used to provide an estimation of the individual particles positions and uncertainties.

With the final track collection made of high-score tracks fulfilling all requirements at hand, all the particle trajectories are extended to the TRT and the compatible TRT hits are used to improve the resolution of the track momenta (see last layer in Figure 4.2). Finally, a final high-resolution χ^2 -fit is performed to obtain the final estimation of the track parameters.

The track reconstruction as described above is known as the inside-out algorithm, which relies heavily on the silicon layers of the Inner Detector and is optimised for the reconstruction of particles created close to the interaction point. However, tracks originating from the decay of a long-lived particle, such as kaons, typically have large impact parameters and leave only few hits in the silicon layers. Hence, an additional outside-in reconstruction algorithm which forms track seeds primarily from TRT hits is also applied to reconstruct such charged-particles trajectories [105]. The outside-in algorithm relies on detector hits that were not assigned to any tracks in the inside-out reconstruction.

Performance of track reconstruction in hadronic dense environments

The reconstruction of tracks composing highly energetic jets — group of close-by hadronic particles — is significantly more challenging because of the high particle densities in their cores, as the spatial separation between the several collimated charged particles becomes of the order of the granularity of the sensors in the silicon layers, which in turn increases the rate of shared clusters. The more boosted a jet is, the more collimated are its constituents. A good control of track reconstruction in these dense environments is fundamental for the search presented in this thesis, as the core of the analysis strategy is identifying Higgs bosons produced in the Lorentz-boosted regime.

To give a prominent example of how the specialised algorithms to identify shared clusters become even more relevant for highly energetic jets, in Figure 4.4(a) it is depicted the average number of pixel clusters of a track that have been identified as merged — split hits — as a function of the angular distance between the track and its associated jet, $\Delta R(\text{trk}, \text{jet})$, for two different intervals of jet transverse momentum. The closer the track is to the jet core (typically defined as $\Delta R < 0.02$), the higher the average number of clusters that have contributions from multiple particles and were identified as merged by the neural network. The rate of merged clusters per track also increases with increasing jet p_T .

In view that the algorithms to identify those multiple-particle clusters are not 100% efficient, it is expected that the overall track reconstruction efficiency should decrease as a function of jet p_T , since tracks sharing at least one cluster not correctly identified as merged are penalised during the ambiguity solver stage of the reconstruction algorithm and are likely lost. An interesting way of measuring the track reconstruction efficiency in real collision data without relying on simulated events is through the variable r_{trk} , defined as the ratio between the sum of all track transverse momenta inside a jet and the total jet p_T ,

$$r_{\text{trk}} = \frac{1}{p_T^{\text{jet}}} \sum_{t \in \text{Tracks}} p_T^t. \quad (4.1)$$

This method relies on the fact that at particle level the ratio between the number of charged and neutral particles within a jet should be constant as a function of the jet p_T . Therefore, as a first approximation considering that jets are mainly composed by pions, r_{trk} should be $2/3$ given that only two out of the three pions are charged. When considering additional effects of QED radiation and hadronic decays of resonances with equal number of charged and neutral particles (e.g., kaons), the ratio is found to be closer to $4/7$. Track reconstruction inefficiencies, which are larger at the core of highly energetic jets, introduce a dependence of r_{trk} on the jet p_T because of the loss of tracks. This can be seen for data collected during the Run-2 in Figure 4.4(b). The r_{trk} distribution has been normalised to the value at $p_T^{\text{jet}} = 500$ GeV to allow for an easier comparison of different p_T regimes. The value of r_{trk} decreases by about 20% when comparing jets with a p_T of 500 GeV and 2000 GeV.




Figure 4.4: (a) Average number of pixel clusters used to reconstruct a track that were identified as merged by the neural network as a function of $\Delta R(\text{trk}, \text{jet})$, for two different intervals of jet p_T . (b) Normalised average r_{trk} as a function of jet p_T . Both plots use data collected in 2018.

4.1.2 Vertex reconstruction

As explained in Section 3.1.2, multiple proton-proton interactions take place in the same bunch crossing. The positions at which each of these interactions occur are defined as vertices of the event. By construction, the primary vertex (PV) is defined as the collision vertex with the largest sum of the squared transverse momentum of all its composing tracks. It represents the point in which the main pp hard-scatter interaction took place [106].

To determine the PV candidates, only a sub-set of the reconstructed tracks passing a stringent selection is used as input. This helps reducing the contamination from poorly reconstructed tracks or tracks that did not originate from the primary interaction. Such requirements involve, for instance, a minimum number of silicon hits left by the particle in the Inner Detector and a maximum number of holes. The first seed position of the PV is selected in the transverse plane based on the beam spot position and in the longitudinal plane based on the z -coordinates of the selected tracks at their points of closest approach to the beam spot.

With a set of tracks and a seed position for the PV at hand, an iterative χ^2 minimisation fit is carried out, and in each iteration all tracks with low compatibility to the current vertex position are down-weighted. Once the first PV candidate position is determined, all tracks that are incompatible with it are removed and used to determine other vertex candidates. This procedure is repeated until all tracks have a well-defined compatible vertex or if no additional vertex can be found. The PV from the hard-scattering process is then determined as the PV candidate with largest $\sum_{t \in \text{tracks}} p_{T,t}^2$.

In many analyses, having an accurate knowledge of displaced secondary vertices (or even higher-order vertices) is crucial to detect the decay of particles with “longer” lifetime, such as b -hadron decays (see Section 4.2.5), or likewise to detect long-lived particles that may

arise in the context of BSM physics. The strategy to determine the position of secondary vertices depends on the signature the analysis is searching for (see e.g. Ref. [107]). As an example, the quality criteria applied to select the input reconstructed tracks varies depending on the lifetime of the decaying particle. The longer-lived the particle, for instance, the lower the expected number of hits it should leave in the first layers of the silicon detectors. The vertex finder algorithm is also adapted for such cases.

4.2 Hadronic jets

Hadronic jets are taken as the shower of particles reaching the detector that was initiated by a quark or a gluon. The kinematics of the jets hence approximate the kinematics of the original parton. Jets within each event collision, however, are not unique and depend on the *jet definition*, which relies i) on a jet clustering algorithm and ii) on the input constituents used for the clustering (such as Inner Detector tracks or calorimeter clusters). The choices for the jet definitions will depend on the analysis strategy and final state of the search. Hence, various *jet collections* can be defined for the same collision event.

A good jet definition must be robust enough to neither be affected by the inclusion of additional soft particles with energy close to zero (infrared safety) nor by the splitting of an input constituent into two softer ones travelling in the same direction (collinear safety). All jet definitions discussed here satisfy this infrared-collinear-safety.

4.2.1 Jet clustering algorithm

Over the years, many jet algorithms have been developed [108–111]. The most common ones used in LHC analyses relies on an iterative approach by which particles (or proto-jets) are combined based on two distance metrics. They are

$$d_{ij} = \min(p_{T,i}^{2n}, p_{T,j}^{2n}) \frac{\Delta R_{ij}^2}{R^2}, \quad (4.2)$$

$$d_i = p_{T,i}^{2n}, \quad (4.3)$$

where $\Delta R_{ij} = \sqrt{(y_i - y_j)^2 + (\phi_i - \phi_j)^2}$ is the angular distance between two input objects, $p_{T,i}$ is the transverse momentum of the i -th input object, n is an integer to govern the relative power of the energy versus the geometric scales, and R is a user-defined radius parameter representing the jet size. In ATLAS the two common choices are $R = 0.4$ and $R = 1.0$.

At each step in the inclusive clustering algorithm, d_{ij} and d_i are calculated for all input constituents in the event. If $d_i > d_{ij}$, the entities i and j are combined, otherwise i is identified as a final jet and removed from the list of input objects. The iterative procedure terminates when all inputs of the event are clustered.

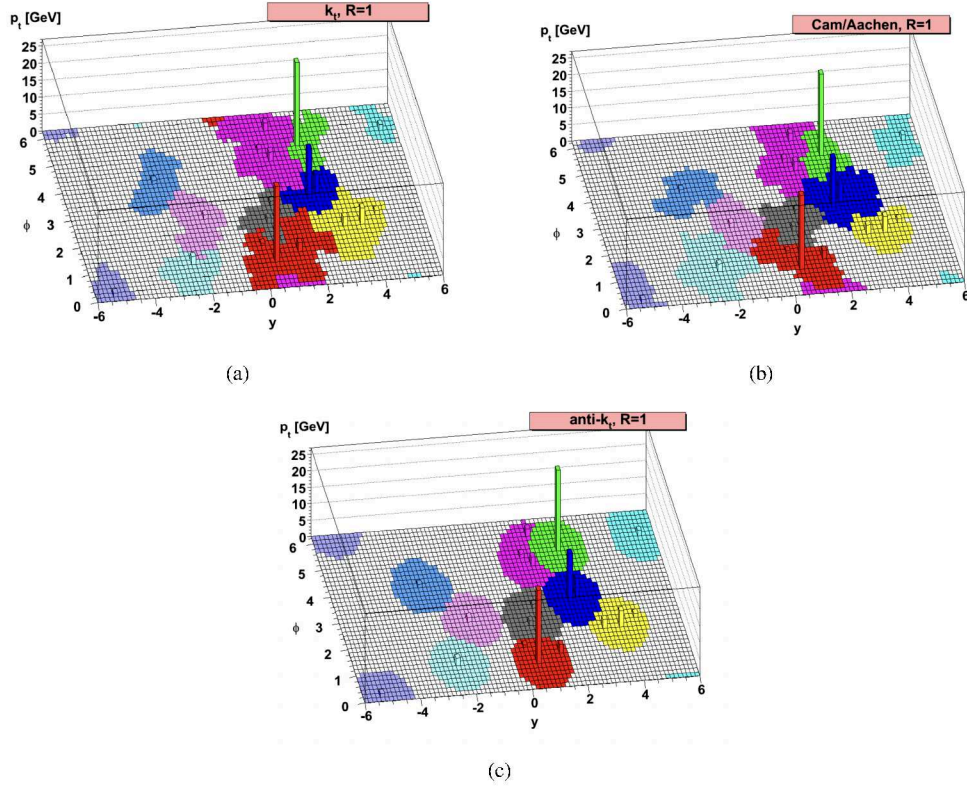


Figure 4.5: Illustrative examples of the shape of the reconstructed jets using (a) the k_t algorithm, (b) the Cambridge/Aachen algorithm, and (c) the anti- k_t algorithm for the same set of input particles. In all cases the radius parameter was set to 1.0 for simplicity. Different colors correspond to different jets [111].

The choice of n determines the criteria based on which the particles/proto-jets are clustered. The simplest choice $n = 0$ is known as the Cambridge/Aachen algorithm [109] and has the two metrics independent of the transverse momentum. In this algorithm, objects with small angular separation are clustered first and the recombination depends solely on ΔR_{ij} . For $n = 1$, known as the k_t -algorithm [108], the recombination prioritizes soft and close-by objects over harder, wide-angle emissions. This behaviour is generalised for any $n > 0$. Finally, for $n = -1$ (and with similar behaviour for any negative n) one has the anti- k_t algorithm [111], which builds jets starting from the hard, well-separated constituents, providing therefore the overall jet geometry already within the first steps in the iterative procedure. Unlike the Cambridge/Aachen and k_t algorithms, the resulting jets from the anti- k_t algorithm have an approximately circular shape in the $y \times \phi$ plane. The typical shape of the jets for these algorithms can be seen in the illustrative examples of Figure 4.5 using the same set of input particles.

The anti- k_t algorithm is the common choice in most ATLAS analyses and is also the jet algorithm used throughout this thesis, unless stated otherwise.

Calorimeter jets

One possible choice of input constituents for the jet algorithm are the signal clusters of the calorimeter cells, commonly referred to as topo-clusters, as they provide a good proxy of the energy deposits of the individual hadrons from each jet [112].

In order to obtain the collection of topo-clusters of an event, a “significance” ς_i is assigned for each calorimeter cell i . It is defined as the ratio of the energy deposit on the cell E_i to the average expected pile-up and electronic noise $\sigma_{\text{noise},i}$,

$$\varsigma_i = \frac{E_i}{\sigma_{\text{noise},i}}. \quad (4.4)$$

All cells with $\varsigma_i > 4$ are by construction used as seeds for the topo-clusters collection. The topo-cluster then grows by merging into it all adjacent cells with $\varsigma_i > 2$. Once all cells have been clustered, a scan over the energy distributions over all topo-clusters is performed and, if a cluster happens to have two regions of energy maxima, it is split into two.

The topo-clusters must also be calibrated to correct for effects of the non-compensating calorimeter response to hadrons, as part of the hadronic shower energy is dissipated via nuclear processes, for accidental signal losses when clustering the cells, and for energy lost in inactive material. The calibration strategy is based on the local hadronic cell weighting (LCW) scheme [112], where for each topo-cluster it is assigned a probability classifying it as either hadronic or EM based on the cluster shape. The energy of the topo-cluster is then corrected by a weighting factor derived from simulated single pion events and parametrised in terms of the above mentioned probability and energy values and density inside the cluster.

Particle flow jets

While the energy resolution of topo-clusters associated to highly energetic particles is very good, it is much lower for low- p_T jets. Fortunately, the ATLAS tracker allows for the reconstruction of low-energy particles with much better resolution thanks to the precise measurements of track curvature in the magnetic field. Hence, in order to maximally benefit from the performance of both ATLAS subsystems, an ensemble combining simultaneously tracks from the hard interaction and calorimeter energy clusters can be used. This set of input quantities is what is called particle flow objects [113]. To avoid a double-counting of energy and momentum measurement from the Inner Detector and calorimeter, an overlap removal algorithm to remove from the topo-clusters the energy contribution associated to matched tracks must be carried out. The track energy measurement is obtained from its p_T under the pion-mass assumption.¹

First, the tracks are selected based on a set of stringent quality criteria depending on the hits in the silicon detectors, and are required to have $|\eta| < 2.5$ (Inner Detector coverage)

¹ Tracks that are found to be associated to an electron are re-fitted under the electron-mass assumption. See Section 4.3.

and $0.5 \text{ GeV} < p_T < 40 \text{ GeV}$. The p_T -requirement assures only tracks with non-negligible contributions to the total jet- p_T are picked and removes energetic particles that are often poorly isolated from nearby activity in the detector and would deprecate the accuracy of the overlap removal algorithm. Each selected track is then matched to all topo-clusters close to it based on their proximity in the $\eta \times \phi$ plane and if the topo-cluster carries at least 10% of the track momentum. The overlap removal proceeds by subtracting the track energy from the topo-clusters energy, at a cell-by-cell basis, starting from the regions with highest energy density. In case the remaining energy in the topo-clusters cells associated to the track is consistent with the overall expected energy resolution, the remnant cells of the topo-clusters system are also removed. In the end, the particle flow objects are composed by a list of tracks, and modified and unmodified topo-clusters. All objects made of topo-clusters associated to tracks are attributed to charged hadrons, while topo-cluster-only objects are assumed to arise from neutral hadrons.

In most ATLAS Run-2 analyses, the particle flow objects and the calorimeter topo-clusters were the two main sets of input constituents to reconstruct jets. For ATLAS Run-3 analyses, another set of objects is now available. This is further explained in Chapter 7, where the work of an on-going VBF $HH \rightarrow b\bar{b}b\bar{b}$ analysis using also Run-3 data is presented.

4.2.2 Small- R jets

The standard jet collection used in ATLAS Run-2 analyses is based on the anti- k_t algorithm with radius parameter $R = 0.4$ and using particle flow objects as input to the jet algorithm. Throughout this thesis, $R = 0.4$ jets will be referred to as *small- R* jets.

In principle, small- R jets are meant to reconstruct the multiple hadrons originating from a single quark or gluon. In the search presented here, the presence of jets typically in the forward ($|\eta| > 2.5$) part of the detector plays a central role in identifying the signature of the VBF production of Higgs-boson pairs; the reconstruction of the Higgs bosons themselves, however, relies on a more complex jet collection, as discussed in Section 4.2.3.

Jet-vertex-tagger

The jet algorithms presented in the last sub-section are used to reconstruct jets regardless of whether they originate from a parton arising from the hard-scatter primary vertex or not. Additionally, some jets might be reconstructed by clustering hadrons that do not arise from one single vertex. Any jet originating from a parton not from the hard-scattering primary vertex and/or with a large fraction of pile-up contribution should be properly identified by a dedicated algorithm.

To decide whether a jet originates from a hard-scattering primary vertex or not, it is used the corrected Jet Vertex Fraction (JVF) variable that quantifies the relative impact of tracks from the main primary vertex PV_0 and tracks originating from primary vertices $PV_{n \geq 1}$ due to pile-up interactions in the same bunch crossing. It is given by [114]

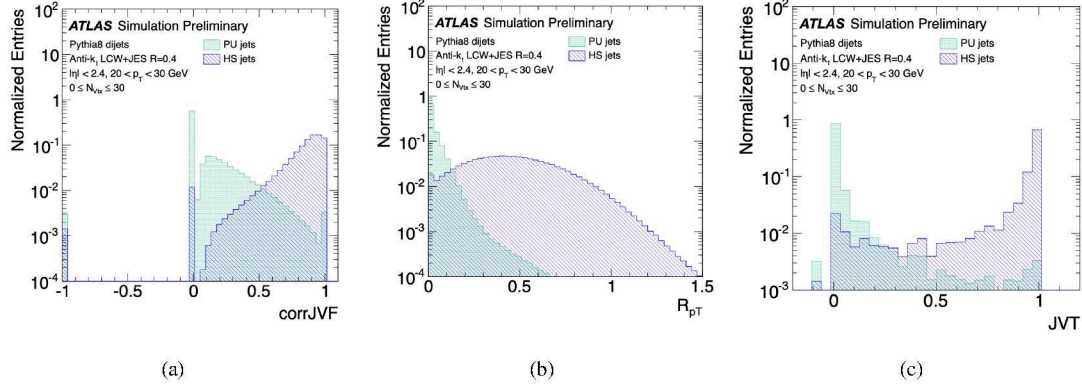


Figure 4.6: Comparison of hard-scatter and pile-up jets in the range $20 \text{ GeV} < p_T < 30 \text{ GeV}$ in the variables (a) corrJVF , (b) R_{p_T} , and (c) JVT [115].

$$\text{corrJVF} = \left(\sum_{t \in \text{Tracks in PV}_0} p_T^t \right) / \left(\sum_{t \in \text{Tracks in PV}_0} p_T^t + \frac{p_T^{\text{PU}}}{\alpha n_{\text{trk}}^{\text{PU}}} \right), \quad (4.5)$$

where the summation in the numerator and denominator is the scalar p_T sum of all tracks associated to the jet that originated from the main primary vertex PV_0 . The denominator further includes the contribution from all the additional tracks from pile-up vertices through

$$p_T^{\text{PU}} = \sum_{n \geq 1} \sum_{t \in \text{Tracks in PV}_n} p_T^t \quad (4.6)$$

and corrected by the factor $\alpha \cdot n_{\text{trk}}^{\text{PU}}$ to mitigate the approximate linear increase of p_T^{PU} with the total number of pile-up tracks $n_{\text{trk}}^{\text{PU}}$ per event. The value $\alpha = 0.01$ is approximated by the slope of the average $\langle p_T^{\text{PU}} \rangle$ with $n_{\text{trk}}^{\text{PU}}$.

The values of corrJVF ranges from 0 to 1, with larger values associated mostly to jets that originated from the hard-scatter. If it is equals to zero, there were no hard-scatter tracks associated to the jet. If no tracks at all were associated to the jet, it is assigned by default the value $\text{corrJVF} = -1$. The corrJVF distributions of hard-scatter and pile-up jets in simulated multijets events are shown in Figure 4.6(a).

Furthermore, as a measurement of the overall contribution of charged particles in the jet that originate from the primary vertex, the variable

$$R_{p_T} = \frac{1}{p_T^{\text{jet}}} \sum_{t \in \text{Tracks in PV}_0} p_T^t \quad (4.7)$$

is used. It is steeply falling for pile-up jets and peaks at zero if the jet has no track contributions from the primary vertex, as shown in Figure 4.6(b).

Finally, by combining information of both corrJVF and R_{p_T} in a single discriminant called




Figure 4.7: Schematic representation of how the angular separation between the decay products of the process $A \rightarrow B + C$ becomes smaller for increasing transverse momentum of the mother particle A .

Jet-Vertex-Tagger (JVT) using the k-nearest neighbours machine learning algorithm [114], it is possible to assign a probability of whether a jet is or is not coming from the hard-scatter collision or has a large pile-up contribution. The discriminant distributions of hard-scatter and pile-up jets in simulated multijets events are shown in Figure 4.6(c). A selection requirement is placed on the value of the JVT discriminant to select jets originating from the main hard-scatter vertex; different choices for these values correspond to different efficiencies in the overall number of jets being kept in each event. These choices are commonly referred to as working points (WP). Across ATLAS, it is typically recommended to use a JVT requirement corresponding to a 96% efficiency for hard-scatter jets.

4.2.3 Large- R jets

The decays of high- p_T heavy particles such as Higgs, W , and Z bosons into two particles are often reconstructed as a single large radius jet. This is well motivated by the fact that, as p_T increases, the decay products of a particle are increasingly collimated in the particle's direction in the laboratory frame. This behaviour is parametrised by the empirical formula (see e.g. [116]) relating the angular separation $\Delta R_{B,C}$ between the decay products of the two-body decay $A \rightarrow B + C$ of a particle A with mass m_A and transverse momentum p_T^A :

$$\Delta R_{B,C} \approx \frac{2m_A}{p_T^A}. \quad (4.8)$$

A schematic representation of the collimation behaviour is shown in Figure 4.7.

These jets, commonly known as large- R or boosted jets, are reconstructed in ATLAS analyses using the anti- k_t algorithm with jet radius parameter set to $R = 1.0$. In the Run-2 analyses, calorimeter topo-clusters were used as input objects to the large- R jet algorithm.

Due to the large area of calorimeter cells the large- R jets are expected to activate, the topo-clusters are susceptible to energy contaminations from the underlying pile-up event. Hence, to correct for this effect, the large- R jets are trimmed by re-clustering each $R = 1.0$ jet into several $R = 0.2$ sub-jets. If a sub-jet has a p_T lower than 5% of the total large- R jet p_T , it is removed from the large- R jet reconstruction [117].

The corresponding mass of boosted jets is, in many cases, a powerful tool to distinguish the decay of a heavy particle (often a signal) to some background. This mass is evaluated using information from both the Inner Detector tracks and the calorimeter clusters [118] for increased resolution. First, the calorimeter-based jet mass m^{calo} is defined as

$$m^{\text{calo}} = \sqrt{\left| \sum_{i \in \text{jet}} E_i \right|^2 - \left| \sum_{i \in \text{jet}} \mathbf{p}_i \right|^2}, \quad (4.9)$$

where E_i and \mathbf{p}_i are the energy and vector momentum of the i^{th} calorimeter cell composing the large- R jet. In addition, the track-assisted jet mass m^{TA} is calculated as

$$m^{\text{TA}} = \frac{p_T^{\text{calo}}}{p_T^{\text{tracks}}} \times m^{\text{tracks}}, \quad (4.10)$$

where p_T^{calo} is the transverse momentum of the calorimeter jet, and p_T^{tracks} and m^{tracks} are the transverse momentum and invariant mass of the four-vector sum of all tracks associated to the calorimeter jet. (For the calculation of the invariant mass of all tracks, the invariant mass of each track is set to the pion-mass value.) The ratio $p_T^{\text{calo}}/p_T^{\text{tracks}}$ is an important component to correct charged-to-neutral particles fluctuations, as m^{tracks} standalone is based only on the charged components of the jet and hence would underestimate the real jet mass. Finally, the combined mass is calculated using

$$m^{\text{comb}} = w \cdot m^{\text{calo}} + (1 - w) \cdot m^{\text{TA}}, \quad (4.11)$$

with the weight w determined independently for each large- R jet based on the resolution of the calibrated calorimeter and track-assisted mass terms [118].

4.2.4 Jet calibration

At detector level, only the energy deposits of particles interacting with detector elements can be measured. Due to detector imperfections and pile-up processes that interfere the signal in detector elements, extra corrections must be applied to correct the jet energy scale (JES) to that of the initial parton originating the jet. At the same time, it is important to optimise the overall jet energy resolution (JER) to allow for a precise reconstruction of the jet kinematic properties [119, 120].

In practice, the corrections in reconstructed jets are applied to match the kinematics of particle-level (truth jets), which are clustered with the anti- k_t algorithm using stable final state particles (except muons and neutrinos) from simulation. To match each truth jet to a reconstructed jet, the $\Delta R < 0.3$ matching between them is used.

First, a potential overestimation of the jet energy originating from pile-up is removed

primarily using information based on the jet area and transverse momentum density. An extra residual correction to the pile-up effects is derived from simulated events using a parametrized function of the average number of interactions per bunch crossing and the total number of reconstructed primary vertices in the event. In the case of large- R jets, a muon-in-jet correction must also be carried out to account for the potential energy missed from semi-leptonic b -decays — see Section 4.3.2.

The JES correction factor c_{JES} is then derived in bins of the reconstructed jet energy and pseudorapidity using the energy response function $E_{\text{reco}}/E_{\text{truth}}$, where E_{reco} is the reconstructed jet energy and E_{truth} the energy of the corresponding truth jet [121]. The JES corrections modify the jet energy, pseudorapidity, transverse momentum, and mass (in case of large- R jets only) as

$$E_{\text{reco}} = c_{\text{JES}} E_0, \quad (4.12)$$

$$\eta_{\text{reco}} = \eta_0 + \Delta\eta, \quad (4.13)$$

$$p_T^{\text{reco}} = c_{\text{JES}} |\mathbf{p}_0| / \cosh(\eta_0 + \Delta\eta), \quad (4.14)$$

$$m_{\text{reco}} = c_{\text{JES}} m_0 \quad (\text{if large-}R \text{ jet}) \quad (4.15)$$

where E_0 , η_0 , \mathbf{p}_0 , and m_0 are the jet properties prior to any calibration. The jet pseudorapidity is corrected with a shift $\Delta\eta$ to correct for any bias relative to truth jets in certain regions of the detector [120].

To account for imperfections in the simulation of the detector response in the JES calibration using truth jets, an extra *in situ* calibration is used to correct for the differences of calibrated objects in simulated and in real data. In the case of large- R jets, the method relies on p_T conservation in events where a large- R jet recoils against a well measured object. This is done using Z + jet events — where the recoiling system is a reconstructed $Z \rightarrow e^+e^-$ or $Z \rightarrow \mu^+\mu^-$ decay —, γ + jets events — where the recoiling system is the reconstructed photon —, and multijet events — where the balance system is formed by several low- p_T small- R jets [121]. From the p_T balance it is extracted normalisation factors to correct the p_T of reconstructed jets from real data to match simulation. The process is analogous for the in-situ calibration of small- R jets [119].

Specifically to large- R jets, jet mass scale (JMS) corrections must also be employed once the JES calibration has been applied. Similarly to the JES, the correction factor c_{JMS} is determined as a function of the jet E_{reco} , m_{reco} , and additionally $\log(m_{\text{reco}}/E_{\text{reco}})$, using the mass response function $m_{\text{reco}}/m_{\text{truth}}$. The large- R jet kinematic variables after the JES+JMS calibration becomes

$$E_{\text{reco}} = c_{\text{JES}} E_0, \quad (4.16)$$

$$\eta_{\text{reco}} = \eta_0 + \Delta\eta, \quad (4.17)$$

$$p_T^{\text{reco}} = c_{\text{JES}} \sqrt{E_0^2 - c_{\text{JMS}}^2 m_0^2} / \cosh(\eta_0 + \Delta\eta), \quad (4.18)$$

$$m_{\text{reco}} = c_{\text{JES}} c_{\text{JMS}} m_0. \quad (4.19)$$

An additional *in situ* mass calibration is also carried out to account for the differences in real data and simulation in the determination of the jet mass. The method utilises the measurements of the jet p_T using solely the calorimeter topo-clusters p_T^{calo} (which provides information from the full hadron shower) and the jet momentum reconstructed using only tracks p_T^{track} (which provides information only from the charged-particles in the hadron shower). The extra normalisation corrections to the jet mass are determined by comparing the average response function $p_T^{\text{calo}}/p_T^{\text{track}}$ in data and simulated multijet events [121].

As should be expected, a number of factors can affect the JES and JMS calibration, such as detector effects, event selections related to the in-situ measurements, modelling, and statistics. All of these effects are taken into account in physics analyses through the inclusion of systematic uncertainties during the statistical analysis stage (see Section 5.4).

Finally, the overall JER and the jet mass resolution (JMR) are evaluated once all energy- and mass-scale corrections have been applied to the large- R jets. In the case of large- R jets, the JER measurement relies on the calculation of the p_T balance between leading and sub-leading large- R jets in multijets events in different η regions, while the JMR is extracted from fits to the W -boson and top-quark mass peaks in high- p_T $t\bar{t}$ events with leptons+jets in the final state [121]. For small- R jets the JER is measured by studying the p_T asymmetry in events containing only two jets, since from momentum conservation in the transverse plane the transverse momenta of these two jets should be balanced [119].

4.2.5 Single and double b -tagging algorithms

The ability to identify whether a jet originated or not from a b -quark is of paramount importance for an extensive list of analyses conducted by, but not only, the LHC experiments. Dedicated tools are developed for such task, and they are referred to as b -tagging algorithms. Among them, the single b -tagging algorithm is the most common tool used across the physics analyses; it is used to *tag* jets with small radius parameter (e.g., the typical small- R jets already introduced) that should capture the decay of a single parton. For more complex topologies like the one presented in this thesis, where the Higgs-boson decay $H \rightarrow b\bar{b}$ is reconstructed as a single large- R jet, dedicated double b -tagging algorithms to identify jets composed by the hadronic decays of two close-by b -quarks are crucial. (Details on the analysis event reconstruction strategy are given in Chapter 5.)

In this section, the taggers used for Run-2 analyses, which includes the search presented here, are described. More efficient algorithms based on more modern machine learning tools are naturally becoming available for Run-3 analyses; this is further discussed in Chapter 7.

The single b -tagging algorithm

Because of the relatively long lifetime of b -hadrons of the order of 1.5 ps, they typically travel a few millimetres (up to centimetres!) before decaying into lighter hadrons. Hence, the decay vertex can be reconstructed as a displaced secondary vertex and all its associated tracks exhibit large impact parameters with respect to the primary vertex, as illustrated in Figure 4.8.




Figure 4.8: Illustration of an interaction producing one b -jet and two light-flavour jets. The b -hadron travels a distance L_{xy} before decaying and producing displaced tracks originated from a secondary vertex. Compared to prompt tracks originating from the primary vertex, the displaced tracks have a large transverse impact parameter d_0 [122].

The standard single b -tagging algorithm used in Run-2 analyses is called DL1r [123]. It is based on a sequence of deep learning algorithms that receives as input an extensive list of track, primary vertex, and secondary vertex features. First, a recursive neural network takes as input a variable number of tracks associated to the jet; the network outputs combined to other several reconstructed vertex properties are fed into a dense neural network that calculates the probabilities of the jet being originated from a b -, c -, or light-quark. The latter two are interpreted here as background jets. The three probabilities are then finally used to construct a one dimensional discriminant that is used to make a decision on the nature of the jet.

Other alternatives, such as the DL1 algorithm that does not make use of the recursive neural network and the MV2c10 algorithm based on boosted decision trees, have also been developed [124]. In Figure 4.9, the background rejection in simulated $t\bar{t}$ events as a function of the tagger efficiency in keeping b -jets when applying a requirement to the jet based on the tagger discriminant is shown. The DL1r algorithm outperforms the other alternatives.

The double b -tagging algorithm

The algorithm employed in this analysis to identify whether a large- R jet originated or not from the hadronic Higgs-boson decay $H \rightarrow b\bar{b}$ uses the kinematics of the jet itself combined with information from its underlying sub-jets — a type of jets with smaller radius that are spatially matched to the large- R jet.

The individual b -jets inside a large- R jet, considered here as a proxy candidate for a Higgs-boson decay, are reconstructed using Inner Detector tracks and a variable-radius (VR) anti- k_t jet algorithm [125]. Contrary to the fixed-radius jet algorithm, the effective radius of the sub-jets in the VR jet algorithm varies as a function of the large- R jet p_T in order to better describe the natural collimation of the sub-jets (recall Equation 4.8). The effective




Figure 4.9: Background rejection in simulated $t\bar{t}$ events as a function of the b -jet efficiency for the different single b -tagging algorithms. The two lower ratio-pads is with respect to the MV2c10 algorithm. In the labels, f_c is a parameter used to construct the tagger discriminants [123].

radius parameter is given by

$$R_{\text{VR}} = \frac{\rho}{p_T}, \quad (4.20)$$

with lower and upper bounds at $R_{\text{min}} = 0.02$ and $R_{\text{max}} = 0.4$, and a scaling parameter $\rho = 30$ GeV. All these parameters were optimised to maximise the matching efficiency of the jet axis and the b -hadrons [126]. Each of the reconstructed VR sub-jets are required to have $p_T > 7$ GeV, $|\eta| < 2.5$, and at least two track constituents satisfying some track quality criteria.

All reconstructed VR sub-jets of each large- R jet are then tagged using the above mentioned DL1r single b -tagging algorithm. The DL1r discriminants of the three p_T -leading sub-jets² are then combined with the p_T and η of the large- R jet³ and fed into a dense neural network. The double b -tagger network outputs three values p_{Higgs} , p_{top} , and p_{multijet} corresponding, respectively, to the probability of the jet to be originated from a $H \rightarrow b\bar{b}$ decay, top-decay, or

² In case a large- R jet has fewer than three VR sub-jets, the corresponding input values to the neural network are replaced by the mean discriminant values of the other VR sub-jets.

³ The mass of the large- R jet is not used because the mass distribution of Higgs-boson jets peaks close the Higgs-boson mass, while this is not true for background jets. This could introduce bias into the model learning. Although the jet- p_T could also introduce a similar bias, a downsampling technique to remove jets from the training samples so that the shape of the p_T distributions of all processes are similar is applied.




Figure 4.10: Multijet (a) and top (b) jet rejection as a function of large- R jet p_T for the 60% efficiency $H \rightarrow b\bar{b}$ WP. The double b -tagger performance is shown as blue dots and supersedes the alternative methods [127].

a light-quark or gluon jet. These three probabilities are used to define the double b -tagger discriminant

$$D_{Xbb} = \ln \frac{P_{Higgs}}{f_{top} \cdot p_{top} + (1 - f_{top}) \cdot p_{multijet}}, \quad (4.21)$$

where f_{top} is a parameter determining the fraction of top-decays background jets and optimised as $f_{top} = 0.25$ [127]. Higher values of D_{Xbb} indicate higher chances of a jet to be a $H \rightarrow b\bar{b}$ decay.

In this analysis, the large- R jets are identified as $H \rightarrow b\bar{b}$ jet if $D_{Xbb} > 2.44$, corresponding to a WP of 60% efficiency in keeping $H \rightarrow b\bar{b}$ jets. With this WP, the double b -tagger rejects approximately 80% (28%) of multijet (top-decay) jets with p_T up to 1 TeV. In Figure 4.10, the multijet (a) and top-decay (b) rejections as a function of the large- R jet p_T are displayed for the 60% WP. These plots also show the expected rejections if only b -tagging the VR sub-jets using the DL1r or MV2 algorithms and not using any information regarding the large- R jet kinematics. In this alternative approach, the large- R jet identification is based on the minimum single b -tagging discriminant value of the two highest p_T sub-jets.

The performance of the double b -tagging algorithm using the neural network supersedes the other methods significantly, especially in the high- p_T regime. Moreover, it is interesting to note how the performance when using fixed-radius $R = 0.2$ sub-jets (gray open circles) is significantly worse in the high- p_T regime. This is due the incapability of the fixed-radius algorithm in resolving the overlap of the very collimated b -hadrons inside of highly energetic large- R jets.

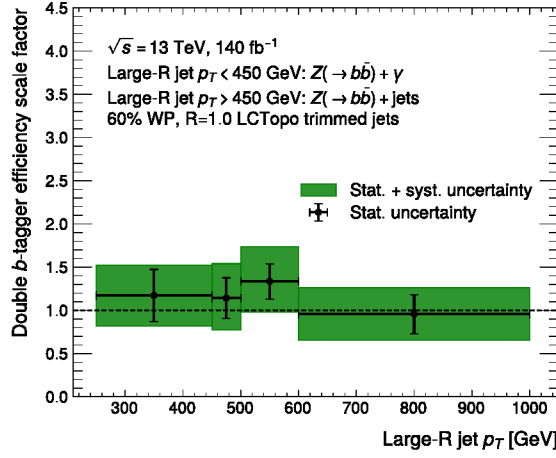


Figure 4.11: Signal efficiency scale factors for the double b -tagger at 60% WP in four large- R jet p_T bins. The black error bars indicate the statistical component of the scale factor uncertainties and the green bands indicate the combined statistical plus systematic uncertainties. The calibration was performed using $Z(\rightarrow b\bar{b}) + \gamma$ and $Z(\rightarrow b\bar{b}) + \text{jets}$ events. This figure is adapted from Ref. [128] based on the b -tagger version used in the analysis presented here.

Double b -tagger calibration

Since the double b -tagging algorithm is trained using only MC simulated collision events, the signal efficiency might not necessarily be identical to that in real data. Hence, the double b -tagger has to be calibrated and scale factors are derived to re-scale the event weights as a function of the large- R jet p_T to account for the tagger efficiency difference between data and simulation. The calibration has been performed using $Z(\rightarrow b\bar{b}) + \text{jets}$ and $Z(\rightarrow b\bar{b}) + \gamma$ events at high Z -boson transverse momenta. The scale factors are then defined as the ratio between the relative efficiencies in data and MC in detecting the $Z \rightarrow b\bar{b}$ signal events [128]. These scale factors come with statistical and systematic uncertainties that must be taken into account in the analysis (see Section 5.4).

The scale factors for the 60% WP are presented in Figure 4.11 in four p_T bins. If a jet has $p_T > 1.0$ TeV, the scale factor from the last p_T bin of the calibration is used.

4.3 Leptons and Photons

The reconstruction of photons, electrons, and muons only plays a small role in the analysis presented in this thesis. Nevertheless, a brief overview is given in this section.

4.3.1 Electrons and Photons

The most important ATLAS sub-detector used to reconstruct electrons and photons is the EM calorimeter. In the central region of the ATLAS detector $|\eta| \lesssim 2.5$, electrons are identified

by matching EM calorimeter clusters to Inner Detector tracks. For the forward, higher pseudorapidity regions outside of the Inner Detector coverage, the forward EM calorimeters must be used.

The EM calorimeter cluster is based on a topo-cluster reconstructed in the same way as for the calorimeter jets (see Section 4.2.1). These clusters are used as seed-clusters and are merged to close-by clusters that could have been created either from bremsstrahlung photons or from electrons arising from photon conversions. All Inner Detector tracks with a minimum p_T of 0.4 GeV and lying within a range of $\Delta R < 0.3$ from the barycentre of the cluster are matched to the cluster. Each resulting cluster-track pair is considered an electron candidate. The electron p_T and angular coordinates are determined by the track properties (due to higher precision) and the energy is determined by the calorimeter cluster. All tracks considered to belong to electrons are re-fitted under the electron-mass hypothesis and using an optimised Gaussian-sum filter [129] to include effects of energy loss in the track measurement (e.g. from bremsstrahlung).

In case no matching Inner Detector track exists, the object is identified as a photon. For cases that could be ambiguous, extra information based on the electromagnetic shower shape and the associated tracks is used to distinguish electrons and photons [130].

4.3.2 Muons

Because of their higher mass compared to electrons, muons do not lose a significant part of their energy in the calorimeter and thus reach the outermost part of the ATLAS detector, the MS. The design of the ATLAS detector allows for a stand-alone muon reconstruction using only information from the MS, but both the identification efficiency and precision on the particle parameters is significantly improved when using additional information from Inner Detector tracks.

First, MS tracks are reconstructed independently by searching for consecutive hits in one single muon station and then combining the consecutive particle trajectories from different muon stations while taking into account the magnetic field in the MS. Next, MS tracks are combined to matching Inner Detector tracks in a global χ^2 -minimisation accounting for the differences of the MS and Inner Detector track parameters and for the energy lost in the EM calorimeter. After the global fit, hit outliers are removed and the fit is redone for increased resolution on the final muon kinematics [131].

In this analysis, muons are used only in the context of muon-in-jet corrections to sub-jets of large- R jets (recall definition in Section 4.2.5) to account for the energy missed from semi-leptonic b -decays. If a muon satisfies a set of quality criteria — namely having $p_T > 4$ GeV, $|\eta| < 2.5$, and having at least three consecutive hits in the first muon station — and is spatially matched to one of the track jets inside a large- R jet, the muon is by construction assumed to be part of the jet and thus is added to the large- R jet four-vector.

4.4 Missing transverse energy

The missing transverse energy E_T^{miss} accounts for a potential momentum imbalance in the detector transverse plane caused by particles that do not interact with the detector elements. By definition, E_T^{miss} is given by the magnitude of the negative vectorial sum of all calibrated physics objects p_T in the event and all tracks that were reconstructed but did not match any calibrated object [132],

$$E_T^{\text{miss}} = \left| - \sum_{i \in \{\gamma, e, \mu, \tau, \text{jet}\}} \mathbf{p}_T^{(i)} - \sum_{i \in \{\text{unused tracks}\}} \mathbf{p}_T^{(i)} \right|. \quad (4.22)$$

Using already calibrated objects is of crucial importance to account for effects from mis-measurements due to geometric detector coverage, detector resolution, noise, and others, which may lead to false E_T^{miss} signatures.

In case of SM processes, non-zero E_T^{miss} are typically present only in events containing neutrinos. Other sources of E_T^{miss} could be from some BSM processes that predict the presence of new hypothetical particles that do not interact with the detector elements.

There is no E_T^{miss} requirement applied in the search presented in this thesis, as it involves only fully hadronic final states.

Analysis Strategy

The search for the VBF production of Higgs-boson pairs with four b -quarks in the final state in the kinematic regime characterised by larger Lorentz boosts of the Higgs bosons targets collision events with a detector signature of two large- R jets — one for each pair of b -quarks from the Higgs-bosons decays — and two small- R jets to fully reconstruct the additional two jets from the hard-collision that define the VBF topology. Reconstructing the Higgs bosons as boosted, merged jets is particularly relevant in BSM scenarios involving high-mass resonances decaying to Higgs-boson pairs or with $\kappa_{2V} \neq 1$ that produce high- p_T Higgs bosons, which in turn tend to have close-by decay products.

A schematic representation of the detector signature of signal candidate events is displayed in Figure 5.1. Based on the characteristics of the final state, the event selection is optimised to maximise the selection efficiency of signal-like events while rejecting background events from other SM processes with similar detector signatures. The reconstructed masses of the leading and sub-leading Higgs-bosons is further used to define the signal, validation, and control regions of the analysis. While the signal region is a signal-enriched phase-space region sensitive for the statistical inference, the validation and control regions are signal depleted and can be used to model the expected background. The analysis strategy is largely the same for both the non-resonant and resonant interpretations of this VBF $HH \rightarrow b\bar{b}b\bar{b}$ search; the only difference arises from an extra veto to remove events that lie inside the signal region of the previous non-resonant analysis [73]. The details on the event selection strategy are described in Section 5.1.

The background events that can mimic the signal signature arise mostly from non-resonant production of multijets initiated from heavy (b , t) quark flavours or from non-heavy quarks that were misidentified as a heavy quark. The multijet background is composed of approximately 90% of QCD produced multijets and 10% of $t\bar{t}$ events. Additional contributions to the total background from single Higgs-boson production and diboson (WW , ZZ , WZ) events are negligible. The fact that the QCD multijets background has a huge contribution from detector/reconstruction inefficiencies that lead to particle mis-identifications, besides




Figure 5.1: Schematic representation of the VBF production of pairs of boosted Higgs-bosons decaying into four b -quarks. Each large- R jet represents a Higgs boson decaying into a pair of b -quarks passing a double b -tagging algorithm criteria. The two extra small- R jets account for the VBF signature of the production mode.

other complex factors like high jet-multiplicity, makes its MC simulation an exceptionally challenging task. It often does not present a good closure with real data in regions where the presence of signal is expected to be negligible. Hence, as described in Section 5.2, a data-driven method is used to estimate the background in the signal region of the analysis. Simulated background events were, however, used in the early stages of the analysis to help designing the event selection strategy.

Further discrimination power between signal and background events in the analysis signal region is obtained using Boosted Decision Trees (BDTs) models. The BDTs response distributions of data and simulated events are used as observables in the statistical inference of this analysis (see Chapter 6). The details on the model training and on the BDTs performance both in the non-resonant and resonant searches are presented in Section 5.3.

The estimation and treatment of systematic uncertainties constitute an important task in any analysis. In Section 5.4, an overview of all the systematic uncertainties affecting both the non-resonant and resonant searches is presented.

5.1 Event selection and reconstruction

Reconstructed collision events (both from real data or MC simulated events) must pass a set of requirements in order to be considered a signal candidate event. First, a pre-selection based on trigger information and on the quality of the jets within each event is applied to pick the events that satisfy the minimum requirements to have a signal-like signature. The events passing these criteria are then further subject to object-related selection requirements, specific to the final state of the analysis, such as the presence of tagged $H \rightarrow b\bar{b}$ jets and VBF jets. Only events passing the full selection are used in the final statistical analysis.

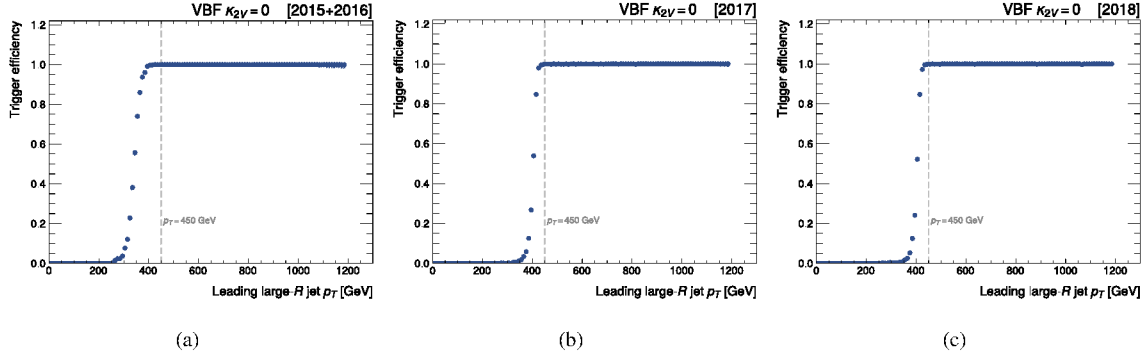


Figure 5.2: Trigger efficiency as a function of the leading large- R jet p_T for the trigger chains of (a) the combined 2015+2016, (b) the 2017, and (c) the 2018 datasets. The vertical dashed line indicates the plateau region at $p_T = 450$ GeV above which all trigger chains present a stable efficiency. The non-resonant VBF $\kappa_{2V} = 0$ signal has been used in this study.

The same event selection strategy is used in the non-resonant and resonant searches. In the kinematic regime with highly Lorentz-boosted Higgs bosons, the non-resonant search is expected to have enhanced sensitivity to BSM scenarios $\kappa_{2V} \neq 1$, and therefore the analysis strategy is designed to detect BSM values of κ_{2V} . Since the kinematic distributions for different BSM values of κ_{2V} are very similar (see Section 5.3 for further discussion), the signal $\kappa_{2V} = 0$ is chosen as a proxy representative sample to optimise the event selection. All selections are also verified to have good counting significance¹ in the resonant search. Each step of the event selection is described in the following subsections.

5.1.1 Triggers

In order to pre-select events that may potentially have signal-like characteristics, both collision data and MC simulated events are required to pass the lowest unscaled large- R jet trigger decisions which require the presence of at least one jet with a minimum transverse momentum. The minimum jet p_T is 360 GeV for data collected in 2015 and 420 GeV for data collected in 2016, 2017, and 2018; the triggers available for data collected in 2017 and 2018 have an additional requirement on the mass of the large- R , which is required to be greater than 35 and 40 GeV, respectively. The trigger chain efficiencies in the VBF $\kappa_{2V} = 0$ signal sample for the combined 2015+2016, 2017, and 2018 datasets as a function of the leading large- R jet p_T are presented in Figure 5.2. For events containing at least one large- R with p_T greater than 450 GeV, all trigger chains present a stable efficiency close to 100%.

¹ Here counting significance is understood as the ratio s/\sqrt{b} , where s and b are the signal and background event yields, respectively. Further details about statistical tools are given in Chapter 6.

5.1.2 Jet cleaning

Jets within each event produced in the pp collisions must be distinguished from misidentified jets of non-collision nature. Examples of such non-collision jets include jets originating from cosmic-ray showers that are produced in the atmosphere and reach the ATLAS detector, and secondary particle cascades originating from soft interactions from proton beam remnants with the cavern walls and detector shielding. The details on how these background jets are identified are described in Ref. [133]. If a single jet fails in passing the jet quality criteria, the entire event is vetoed.

5.1.3 Higgs-bosons candidates

Each event must have at least two large- R jets, denoted here and hereafter as J , corresponding to the decaying Higgs bosons. They must satisfy the kinematic requirements $250 \text{ GeV} < p_T^J < 3000 \text{ GeV}$, $50 \text{ GeV} < m_J < 600 \text{ GeV}$, and $|\eta_J| < 2.0$, which correspond to the region where the ATLAS large- R jet calibration is valid. Should more than two large- R jets be reconstructed in the event, the most energetic ones are chosen as the Higgs-bosons candidates, with the (sub-)leading jet in p_T representing the (sub-)leading Higgs boson. The leading Higgs-boson jet candidate must additionally satisfy the requirement $p_T > 450 \text{ GeV}$ to ensure full online trigger efficiency, as discussed in Section 5.1.1.

The double b -tagging algorithm based on a deep neural network as described in Section 4.2.5 is applied to the large- R jets to identify $H \rightarrow b\bar{b}$ decays. In this analysis, the Higgs-bosons candidates should pass the 60% efficiency WP of the double b -tagger. From now on, events where two or one large- R jets pass the tagger criteria are referred as 2Pass and 1Pass, respectively. The latter, although they do not satisfy the minimum requirements to be signal candidate events, are used for the data-driven background estimation (see Section 5.2).

The distribution of the reconstructed masses m_{H_1} and m_{H_2} of the leading and sub-leading Higgs-boson candidates are shown in Figure 5.3 for two different signal hypotheses, non-resonant VBF $\kappa_{2V} = 0$ and resonant narrow width $m_X = 2.0 \text{ TeV}$. The center of the distributions deviate from the known Higgs-boson mass of 125 GeV mainly due to detector effects, radiation outside the jet cone, energy lost due to neutrinos from semi-leptonic b -hadron decays, and final- and initial-state radiation that may lead to extra jets that are included in the Higgs-boson reconstruction.

5.1.4 VBF signature

Two extra jets in the final state are the main signature of the VBF production mode of Higgs-boson pairs. They are reconstructed as small- R jets, denoted here and hereafter as j , and must have $p_T^j > 20 \text{ GeV}$ and $|\eta_j| < 4.5$, which correspond to the region where the ATLAS small- R jet calibration is valid. In order to mitigate pile-up effects, jets with $p_T < 60 \text{ GeV}$ and $|\eta_j| < 2.4$ should additionally satisfy a requirement based on the output of the JVT algorithm (see Section 4.2.2). To ensure no overlap between the VBF jet candidates

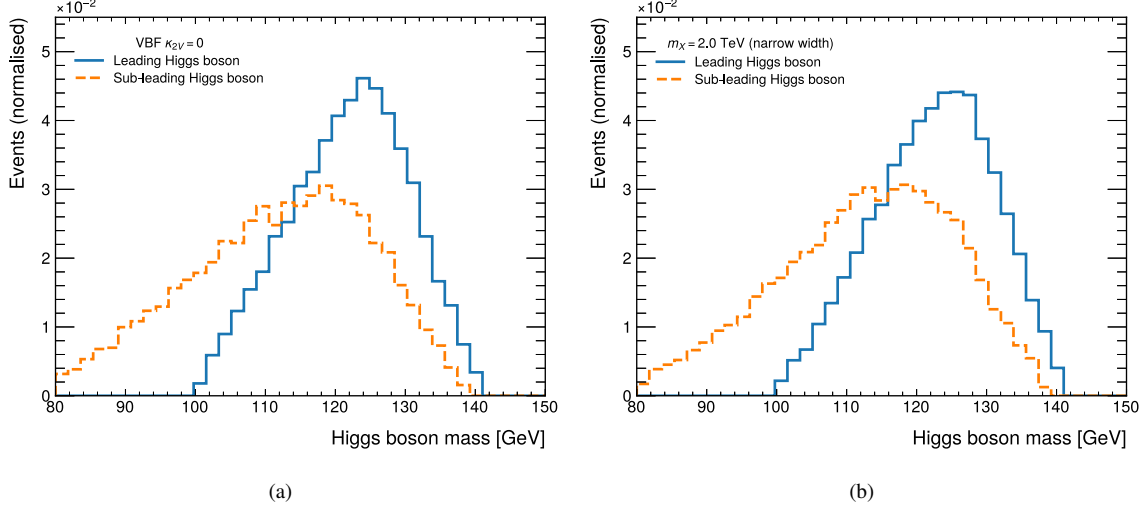


Figure 5.3: Reconstructed masses of the leading (solid line) and the sub-leading (dashed line) Higgs-bosons candidates for (a) the non-resonant VBF $\kappa_{2V} = 0$ and (b) resonant narrow width $m_X = 2.0$ TeV signals.

and the Higgs-bosons candidates, the angular distance between the small- R jets and the selected Higgs-bosons candidates must satisfy $\Delta R(J, j) > 1.4$. If an event has more than two small- R jets satisfying all these criteria, the two most energetic ones are selected as the VBF jets.

The kinematic distributions of the VBF jets provide valuable information to discriminate signal from background events, particularly the invariant mass m_{jj} and the pseudorapidity separation $\Delta\eta(j, j)$ of the VBF system. The kinematic distributions of these variables for the MC simulated background, the non-resonant VBF $\kappa_{2V} = 0$ signal, and the resonant narrow-width $m_X = 2.0$ TeV signal, are displayed in Figure 5.4. Signal events tend to have larger values of m_{jj} and $|\Delta\eta(j, j)|$. A scan over the values of m_{jj} and $|\Delta\eta(j, j)|$ was made in order to find the optimal selection requirements that provide a good counting significance of signal over background events while still keeping a reasonable expected background yield — important for numerical stability in the statistical analysis and to avoid huge statistical fluctuations. The assigned VBF jets thus have to satisfy the additional requirements

- $m_{jj} > 1.0$ TeV,
- $|\Delta\eta(j, j)| > 3.0$.

5.1.5 Resolved non-resonant analysis veto

In order to maximise the sensitivity to κ_{2V} and κ_λ , the search for non-resonant VBF production of $HH \rightarrow b\bar{b}b\bar{b}$ in the boosted regime, as described here, will be combined with the previous ATLAS non-resonant search using a resolved topology, where each of the b -jets were

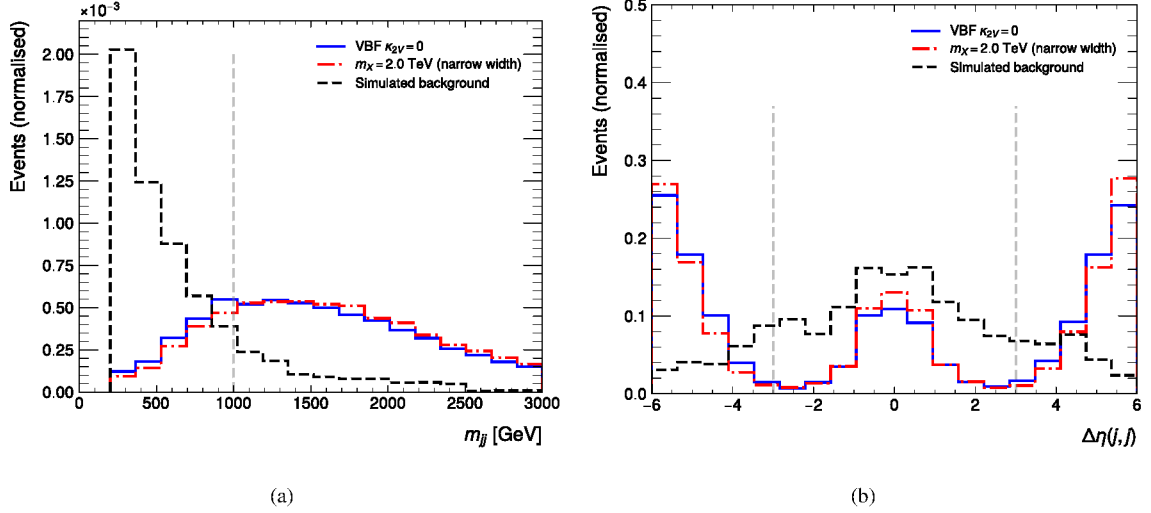


Figure 5.4: Distributions for (a) the invariant mass and (b) the pseudorapidity separation of the system of VBF jets for the MC simulated background, the non-resonant VBF $\kappa_{2V} = 0$ signal, and the resonant narrow-width $m_X = 2.0$ TeV signal. The vertical gray dashed lines indicate the thresholds where a selection requirement based on the variables is applied to select events with signal-like characteristics.

reconstructed individually as small- R jets [73]. (See Section 6.3 for the combined results.) To ensure full orthogonality between both analyses, all events passing the event selection of the resolved analysis are vetoed.² This veto is not applied in the resonant search.

5.1.6 Kinematic regions

After the aforementioned event selection, the 1Pass and 2Pass events are further classified into signal region (SR), validation region (VR), and control region (CR), depending on the following criteria defined in the mass-plane m_{H_1} - m_{H_2} . Events are classified as SR events if they satisfy

$$\sqrt{\left(\frac{m_{H_1} - 124 \text{ GeV}}{1500 \text{ GeV}/m_{H_1}}\right)^2 + \left(\frac{m_{H_2} - 117 \text{ GeV}}{1900 \text{ GeV}/m_{H_2}}\right)^2} < 1.6 \text{ GeV}, \quad (5.1)$$

where the values of 124 GeV and 117 GeV are chosen to match the approximate centers of the m_{H_1} and m_{H_2} distributions of simulated signal VBF HH events (recall Figure 5.3). The functional forms of the denominators, $1500 \text{ GeV}/m_{H_1}$ and $1900 \text{ GeV}/m_{H_2}$, as well as the closure boundary at 1.6 GeV, were optimised through a scan over the values of a_1 , a_2 , and r using the functional forms

² If both analysis topologies were being studied simultaneously, the choice on which analysis to apply the veto would depend on which offers the greatest sensitivity. Since the resolved analysis was conducted and published first, the veto has to be applied in this new, boosted analysis.

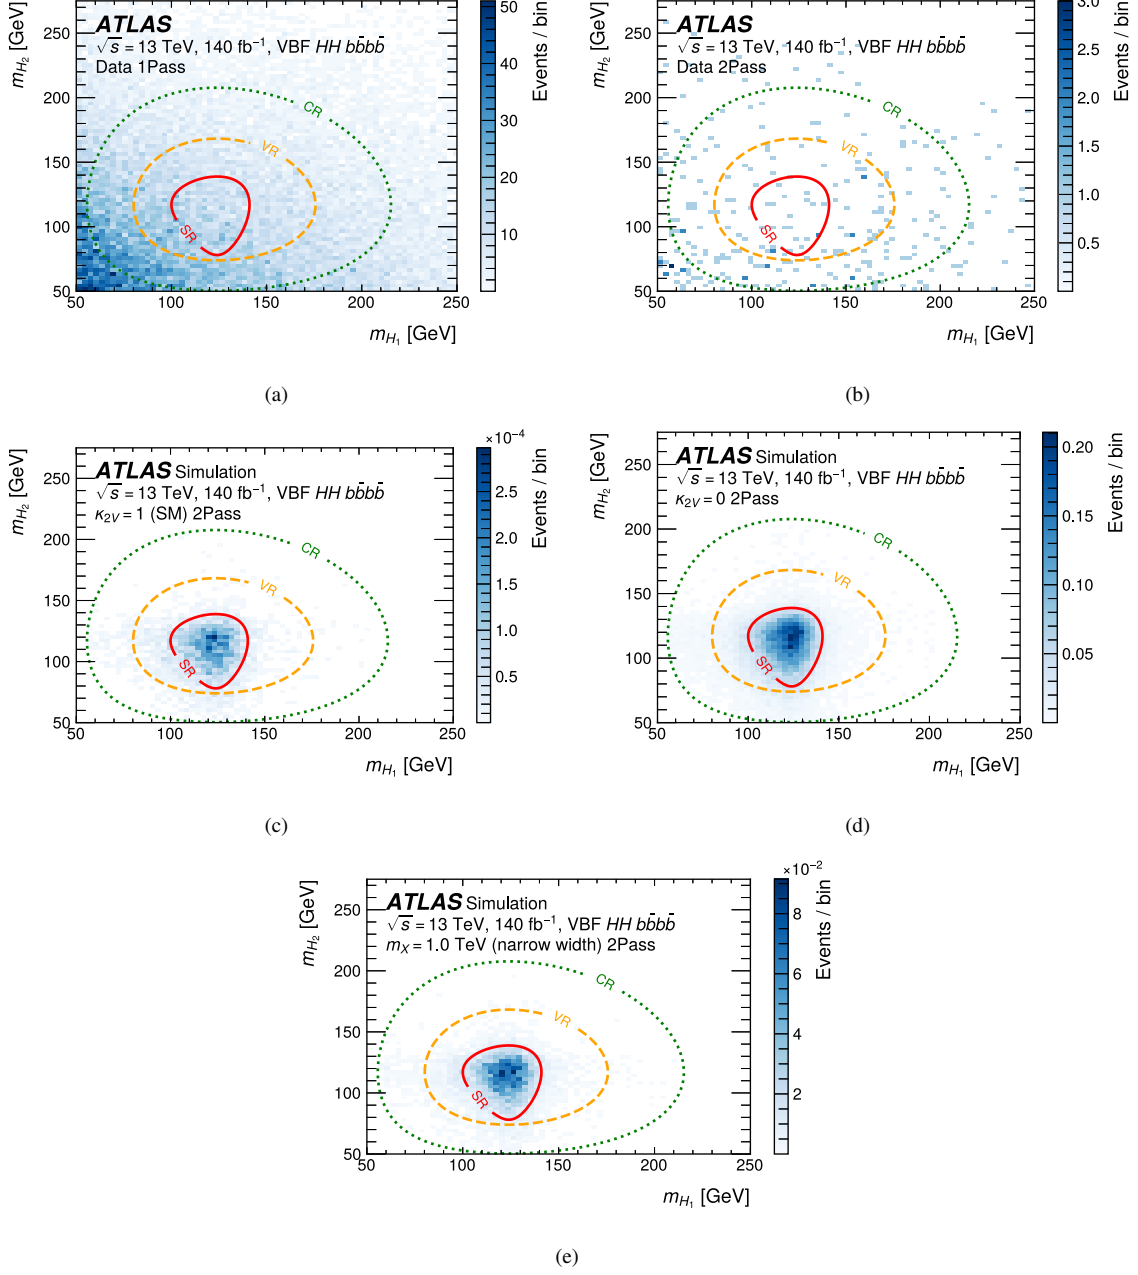


Figure 5.5: Reconstructed m_{H_1} - m_{H_2} plane for the (a) 1Pass and (b) 2Pass selections of the analysis, shown for the data events. The mass planes for the 2Pass selection of the analysis are also shown for the (c) VBF SM $\kappa_{2V} = 1$, (d) VBF $\kappa_{2V} = 0$, and (e) $m_X = 1.0$ TeV spin-0 narrow-width resonance HH samples. The SR is described by the continuous red line and captures most of the signal events, as seen in c), d), and e). The VR is defined as the region between the continuous red line and the dashed yellow line. The CR is delimited between the dashed yellow line and the dotted green line [75].

$$\sqrt{\left(\frac{m_{H_1} - 124 \text{ GeV}}{a_1/m_{H_1}}\right)^2 + \left(\frac{m_{H_2} - 117 \text{ GeV}}{a_2/m_{H_2}}\right)^2} < r, \quad (5.2)$$

$$\sqrt{\left(\frac{m_{H_1} - 124 \text{ GeV}}{2 \cdot a_1/(m_{H_1} + m_{H_2})}\right)^2 + \left(\frac{m_{H_2} - 117 \text{ GeV}}{a_2/m_{H_2}}\right)^2} < r, \quad (5.3)$$

$$\sqrt{\left(\frac{m_{H_1} - 124 \text{ GeV}}{a_1/m_{H_1}}\right)^2 + \left(\frac{m_{H_2} - 117 \text{ GeV}}{2 \cdot a_2/(m_{H_1} + m_{H_2})}\right)^2} < r, \quad (5.4)$$

and selecting the contour that maximises the overall counting significance of the VBF $\kappa_{2V} = 0$ signal over simulated QCD multijet and $t\bar{t}$ background. This SR definition was further verified to also provide optimal sensitivity in the resonant search.

Events are assigned to the VR if they fail the SR requirement in Equation 5.1 but pass the requirement

$$\sqrt{\left(\frac{m_{H_1} - 124 \text{ GeV}}{0.1 \ln(m_{H_1}/\text{GeV})}\right)^2 + \left(\frac{m_{H_2} - 117 \text{ GeV}}{0.1 \ln(m_{H_2}/\text{GeV})}\right)^2} < 100 \text{ GeV}. \quad (5.5)$$

Events failing both SR and VR requirements are assigned to the CR if they pass the requirement

$$\sqrt{\left(\frac{m_{H_1} - 124 \text{ GeV}}{0.1 \ln(m_{H_1}/\text{GeV})}\right)^2 + \left(\frac{m_{H_2} - 117 \text{ GeV}}{0.1 \ln(m_{H_2}/\text{GeV})}\right)^2} < 170 \text{ GeV}. \quad (5.6)$$

The SR, VR, and CR boundaries in the reconstructed m_{H_1} - m_{H_2} plane for data and several signal processes using the 2PASS and 1PASS selections of the analysis are presented in Figure 5.5. For signal processes, as expected, the majority of events are captured within SR boundary, although some of them fall outside due to resolution effects. The fraction of 2PASS HH signal events in the SR ranges from 55% – 78%, depending on the signal hypothesis.

The functional forms of the VR and CR help to reduce the contributions from multijets background events that typically populate the low-mass region of the m_{H_1} - m_{H_2} plane. A naive definition of the VR and CR using the same functional form as for the SR but with increased boundary radius r would not guarantee an unique and closed curve.

5.1.7 Event selection performance

Calculating the fraction of events that were visible to the detector (acceptance) times the fraction of events passing the analysis requirements (efficiency) is important to measure the performance of the analysis strategy in selecting the signal events of interest.

The cumulative acceptance times efficiency for the event selection of VBF non-resonant




Figure 5.6: Cumulative acceptance times efficiency as a function of κ_{2V} for each event selection step in the VBF non-resonant analysis. “Upstream selection” refers to selections applied when reconstructing objects and cleaning events [75].

samples as a function of κ_{2V} for each selection step in the final event selection strategy of this analysis is presented in Figure 5.6. The VBF SM $\kappa_{2V} = 1$ signal has the lowest selection acceptance times efficiency, with only 0.02% of the events residing inside the analysis SR; this reflects that most produced Higgs bosons are not highly energetic in the SM scenario. The performance is, however, significantly better for non-SM values of κ_{2V} , with the total acceptance times efficiency in the full selection chain ranging from 1% – 2%, depending on the value of κ_{2V} .

The analogous cumulative acceptance times efficiency curves of resonant signals for the narrow- and broad-width hypotheses as a function of m_X are shown in Figure 5.7. The total acceptance times efficiency is best in the mass window ranging between 1.3 TeV and 2.0 TeV, mainly because of the performance of the double b -tagging algorithm in this region. While for low masses the produced pairs of Higgs bosons are less energetic and thus are less likely to be fully contained within a single large- R jet, in the high mass hypotheses closer to 5.0 TeV the individual b -jets are so collimated that are hardly distinguishable inside the large- R jet, hence reducing the performance of the double b -tagger.

Numerical values of the event yields after each step in the event selection for data and some selected signal processes are displayed in Table 5.1.




Figure 5.7: Cumulative acceptance times efficiency in the resonant search as a function of m_X for each event selection step for the (a) narrow-width and (b) broad-width hypotheses. “Upstream selection” refers to selections applied when reconstructing objects and cleaning events [75].

Table 5.1: Event yields after each selection step for data and several representative signal hypotheses: non-resonant SM ggF, non-resonant SM VBF, non-resonant VBF $\kappa_{2V} = 0$, and resonant narrow width $m_X = 1.0$ TeV and $m_X = 5.0$ TeV signal samples. All signal yields except “Raw events” are MC generator-weighted and normalised to 140 fb^{-1} . Resonant samples were additionally normalised using $\sigma(pp \rightarrow X \rightarrow HH) = 1 \text{ fb}$. “Upstream selection” refers to selections applied when reconstructing objects and cleaning events. No “Veto resolved selection” is applied to resonant signals. The yields for the broad-width signal samples are similar to those of the narrow-width signal samples listed here [75].

Selection	Data	Non-resonant			Spin-0 resonant VBF	
		SM ggF	SM VBF	VBF $\kappa_{2V} = 0$	Narrow-width m_X 1.0 TeV	5.0 TeV
Raw events	16854036422	1480	82.0	1290	140	140
Trigger & upstream selection	63944638	20.9	1.15	235	70.7	126
≥ 2 large- R jets (η, m)	57510800	14.1	0.531	168	48.7	119
Double b -tagging	12875	5.35	0.131	77.4	25.2	24.9
≥ 2 small- R jets	5762	2.24	0.105	57.2	18.8	16.0
Large- R jets (p_T)	3902	1.41	0.0700	48.3	13.7	16.0
Small- R jets ($\Delta\eta(j, j), m_{jj}$)	314	0.148	0.0380	32.3	8.58	12.0
Signal region	23	0.0970	0.0290	24.5	6.68	6.59
Veto resolved selection	21	0.0590	0.0200	18.8	-	-

5.2 Data-driven background estimation

Fully hadronic final states come with the price of large rates of the hard-to-be-simulated QCD multijets background. In these cases, the background prediction in the analysis SR is often estimated from data itself using collision events from regions where the contributions from

signal processes are expected to be negligible.

The strategy employed in this analysis is based on the ABCD method (also known as matrix method), which was used for the first time in Ref. [134] and has been applied with modified versions in several LHC analyses [30, 73, 135]. Here, the method relies on the definition of orthogonal signal, validation, and control regions defined based on the number of Higgs-boson candidates passing the double b -tagger requirement. To estimate the background shape in the SR 2Pass as a function of a set \mathbf{x} of features characterizing the kinematics of each event, which is described by an unknown probability density function $p_{2\text{Pass}}(\mathbf{x})$, it is used the accessible probability density function of 1Pass data in the SR, $p_{1\text{Pass}}(\mathbf{x})$. The relation between both probability density functions is

$$p_{2\text{Pass}}(\mathbf{x}) = w(\mathbf{x}) \cdot p_{1\text{Pass}}(\mathbf{x}), \quad (5.7)$$

where $w(\mathbf{x})$ is a reweighting function from the 1Pass to the 2Pass regions and a priori may also depend on \mathbf{x} . The reweighting function is assumed to be the same regardless of the analysis region, and thus is obtained by learning the differences between data events in the CR 2Pass and CR 1Pass regions, where both probability density functions are known and the signal contamination is small. In fact, the signal contamination in the 1Pass selection is at most 1% depending on the signal considered and has an impact of at most 8% in the most signal-like bin of the discriminants used in the statistical analysis (see Section 5.3), which is a value covered by the statistical uncertainty in that bin. Typically, $w(\mathbf{x})$ would be determined with the help of machine learning classifiers. However, as shown in Figure 5.8, the shape of relevant kinematic variables are consistent in the 2Pass and 1Pass control regions within statistical uncertainties, and the impact of using a proper machine learning based reweighting has been found to be negligible, as discussed in Appendix A. This is mainly due to the very low statistics after applying the analysis event selection and due to the similarity of the kinematic distributions of the 2Pass and 1Pass events. Hence, an inclusive normalisation factor, independent of \mathbf{x} , is derived by simply calculating the ratio of the total number of events in the CR 2Pass and CR 1Pass regions:

$$w = 0.0081 \pm 0.0010 \quad (\pm 0.0011 \text{ resonant analysis}), \quad (5.8)$$

where the uncertainty is obtained using the VR (see Section 5.4.4 for details) and is different in the non-resonant and resonant analyses because of the extra veto applied in the non-resonant search only. Potential residual shape differences are also covered by the systematic uncertainties of the background modelling, as will be discussed later. Alternative definitions of the VR and CR, which split the nominal definitions of Equations 5.5 and 5.6 into quadrants, are found to yield values for the normalisation factor that are in good agreement with the nominal estimation. This demonstrates that the value provided in Equation 5.8 is largely independent on the choices of how the VR and CR were defined in this analysis.

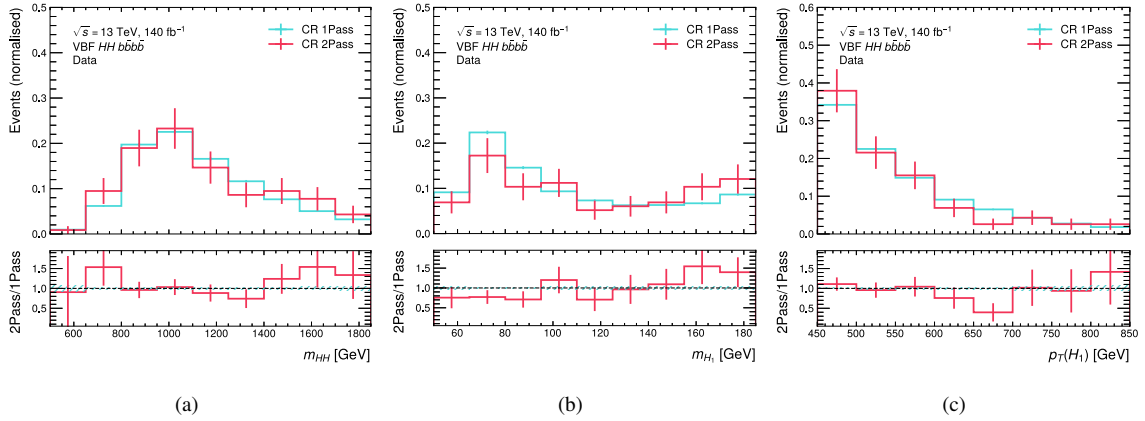


Figure 5.8: Data distributions in the CR 2Pass and CR 1Pass regions of (a) the invariant mass of the HH system, (b) the reconstructed mass of the leading Higgs boson, and (c) the reconstructed p_T of the leading Higgs boson. All uncertainties are of statistical nature.

5.3 Multivariate analysis

Although the event selection strategy in this analysis is designed to detect events with signal-like characteristics, its standalone application may not be optimal to have the best possible sensitivity, as simple kinematic requirements hardly capture all the kinematic nuances that make a collision event to be indeed a good signal event candidate. With the use of multivariate algorithms, it is possible to explore non-obvious relations and correlations between the kinematic properties of an event and use them to enhance the discrimination power between signal and background events in the analysis SR. In this analysis, a more robust version of the decision tree algorithm called Boosted Decision Trees (BDTs) [136] is used to capture more information about events that pass all the requirements of the event selection.

A decision tree is a supervised machine learning algorithm based on a hierarchical structured tree, typically (but not only) used for classification tasks [137]. As exemplified in Figure 5.9, a decision tree is composed primarily of a root node including all samples that are meant to be classified under a class of labels, such as whether an event should or should not be a signal event candidate. Each sample is described by a common set of features (e.g. selected kinematic variables of the event). Based on the property or value of one feature, all samples are divided into two subsets and passed to two internal nodes of the tree. The optimal feature and split point to divide the datasets in each node is found using elements of information-gain theory [137], with the objective to obtain nodes with the minimum possible entropy among the samples, i.e., reduce as much as possible the impurity of the dataset within a node. This process is repeated recursively, in a top-down manner, at each internal node until the majority of samples are classified or some stop condition is met. The nodes where no further split is applied are called leaves. The stop condition is particularly important to reduce an excessive fragmentation of the datasets used during training that leads to over-fitting — a state in which the model becomes too specialised in the training data and performs poorly on




Figure 5.9: Schematic representation of a decision tree. It starts with the root containing all samples meant to be classified. A series of decision requirements are applied to split the samples across different internal nodes.

new, unseen data.

The decision tree algorithm standalone may not be sufficient or may be susceptible to over-fitting if the samples meant to be classified have a large number of features with a high level of attribute complexity. In this case, the gradient tree boosting technique [136], the BDTs, provides a huge improvement in the performance of the machine learning model. In a nutshell, boosting refers to a combination of several, weak learner classifiers trees into a single strong classifier. A set of trees T is created recursively and the output of each tree is evaluated based on the score of each of its leaves. The combined trees provide a prediction \hat{y}_i of a sample i labelled with true value y_i . The ultimate goal is to minimise the objective function

$$\text{obj} = \sum_{i=1}^{N_{\text{samples}}} l(\hat{y}_i, y_i) + \sum_{t \in T} \Omega(f_t), \quad (5.9)$$

where $l(\hat{y}_i, y_i)$ is a loss-function that measures the distance between the prediction and the truth of sample i and $\Omega(f_t)$ is a regularisation function that penalises highly complex trees with tree structure f_t (e.g., with too many splits).

There are many boosting algorithms available that employ different strategies on how to define the loss and regularisation functions, as well as on how to iteratively increase the number of trees. In this analysis, the BDTs used to separate signal from background events in the SR are implemented with the XGBoost library [138]. A number of hyperparameters should be optimised in order to achieve simultaneously the best possible model performance while showing no indications of over-fitting. The main parameters are the number of boosted iterations (i.e., the number of trees created by the model); the maximum depth of each tree; and the learning rate that dictates how much the model is impacted by new iterations. Signal and background events are split into train and test datasets, with the former used to train the model and the latter to quantify the model performance using events not used during training. The performance of the model in discriminating signal from background events

Table 5.2: Kinematic variables used as input features for the BDTs models. The same set of variables is used in the non-resonant and resonant searches.

Physics objects	BDT input variables
Higgs-boson candidate ($H_i, i = 1, 2$)	$p_T^{H_i}, \eta_{H_i}$
Di-Higgs system (HH)	$p_T^{HH}, \eta_{HH}, m_{HH}$
VBF jets ($j_i, i = 1, 2$)	$p_T^{j_i}, \eta_{j_i}, E_{j_i}$

must be similar in both train and test datasets, otherwise the model is said to be biased to the train dataset and incapable of generalising to new data.

Since the BDT's task here is to perform binary classification, the model output is a single number ranging from 0 to 1, representing a probability; the common practice is to design the model such that the outputs correspond to the probability of the event being signal-like.

Different BDT models are employed in the non-resonant and resonant searches, since the kinematic distributions of non-resonant and resonant signal events differ significantly. The set of kinematic variables used in the models, however, are the same and is presented in Table 5.2. The masses of the Higgs-bosons candidates are not used to ensure the model performance will not be dependent on the mass-plane region — a crucial element for the validity of the data-driven background modelling of Section 5.2. For both the non-resonant and resonant BDTs models, the invariant mass m_{HH} and the transverse momentum of the Higgs bosons are found to be amongst the variables that provide the highest gains in discriminating signal from background events after each split in a tree node.³ The kinematic distributions of these variables for the data-driven background and non-resonant signal events under several hypotheses of κ_{2V} are depicted in Figure 5.10. The kinematic distributions of samples with $\kappa_{2V} \neq 1$ are very similar among each other and differs notably from the SM $\kappa_{2V} = 1$ scenario, which has kinematic distributions much more similar to those of the background processes. The analogous distributions comparing background and resonant signal events are presented in Figure 5.12. The shape of the distributions in the resonant search is heavily dependent on the value of the mass of the intermediate particle.

The details on how the BDT models for the non-resonant and resonant searches are constructed are given in the following sub-sections.

Non-resonant analysis optimisation

Since signal events with non-SM values of κ_{2V} have very similar kinematic distributions (see Figure 5.10), the $\kappa_{2V} = 0$ signal is chosen as a representative proxy for those samples and a BDT model is trained to separate $\kappa_{2V} = 0$ signal events from data-driven background

³ The importance of each feature is determined by the overall increase in model performance (or analogously, reduction of samples entropy) at each of the nodes the feature is used, weighted by the number of observations the node is responsible to discriminate.

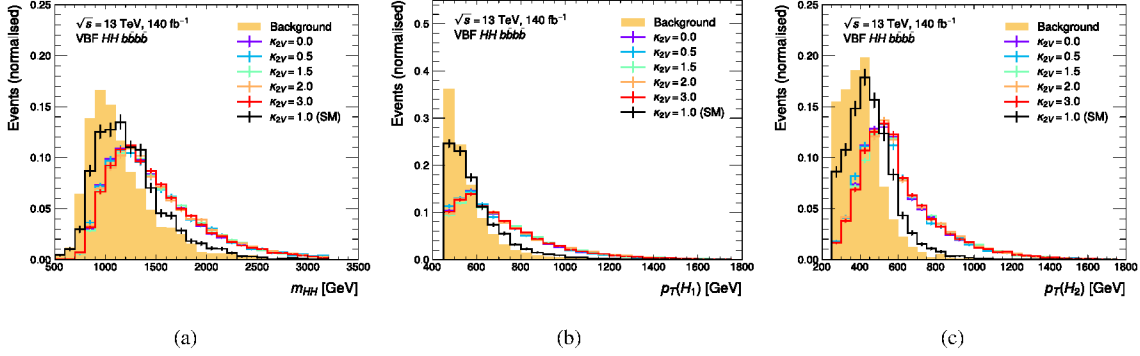


Figure 5.10: Distributions of (a) the invariant mass of the HH system and transverse momentum of (b) the leading and (c) the sub-leading Higgs-bosons candidates, for data-driven background and non-resonant signal events. The SM $\kappa_{2V} = 1$ distributions are shown in black and differ notably from the $\kappa_{2V} \neq 1$ distributions.

events. In order to maximise the sensitivity to the κ_{2V} coupling and in view of the low event selection efficiency of the VBF and ggF SM samples (see Table 5.1), which strongly limits that capability of the analysis in constraining the SM HH production cross-section, the SM samples during the BDT training — and only during training! — are treated as background together with the data-driven background estimation. The hyperparameters of the BDT model were optimised through a grid search and are set to 161 for the number of boosting iterations, to 2 for the maximum depth of each tree, and to 0.1 for the learning rate.

The performance of the final BDT model in identifying signal events with $\kappa_{2V} \neq 1$ is evaluated using the ROC curve, acronym for receiver operating characteristic curve. It measures the signal efficiency selection against the background efficiency selection for several threshold settings based on the output of the BDT model. If the signal and background efficiencies increase at equal rates for varying thresholds, the classifier is understood to have no sensitivity in discriminating the samples and is generally making random identifications only. Ideally, the signal efficiency should be as high as possible at low background efficiency. The area under the ROC curve (AUC) provides a proxy of the model accuracy. The ROC curves for different non-SM values of κ_{2V} are presented in Figure 5.11. Although only the proxy $\kappa_{2V} = 0$ was used as a signal sample during training, the BDT performs equally well in identifying other values of κ_{2V} as signal events, with an overall AUC around 89%.

Resonant analysis optimisation

In the resonant analysis, there is no possibility in choosing only one proxy representative sample to identify as signal during the training of the BDT model because of the distinct kinematic behaviours depending on the value of m_X (see Figure 5.12). The approach used in the resonant search to accommodate multiple mass hypotheses in one single classifier model is to consider a mass-parametrised BDT (pBDT), inspired by parametrised neural networks [139]. In addition to the variables listed in Table 5.2, the pBDT includes an extra parametric input given by the true mass of the heavy intermediate particle. The set of signals




Figure 5.11: ROC curves and AUC evaluated for several non-SM VBF HH signal events and the background components using the BDT classifier. The ratio plot compares the performances of the signal samples with respect to the proxy $\kappa_{2V} = 0$ signal. The dashed gray line indicates the hypothetical ROC curve if the model makes random predictions only.

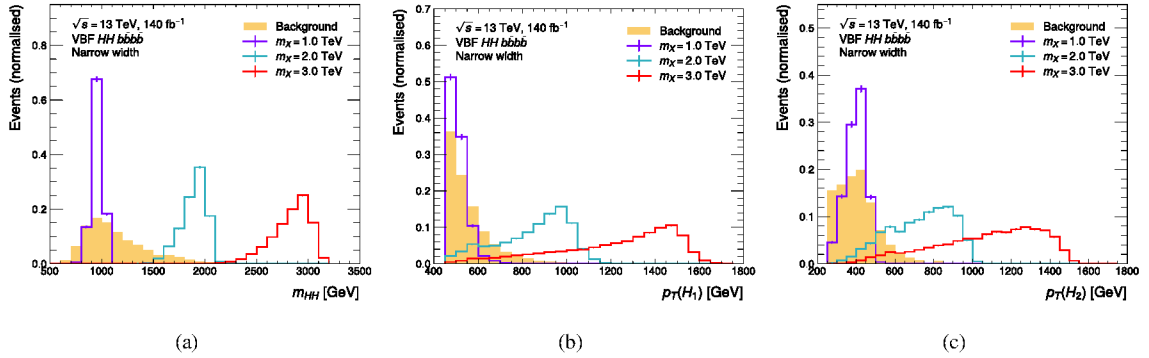


Figure 5.12: Distributions of (a) the invariant mass of the HH system and transverse momentum of (b) the leading and (c) the sub-leading Higgs-bosons candidates, for the data-driven background and resonant signal events under the narrow-width hypothesis. The kinematic distributions of signal events depend heavily on the value of m_X .

for training are then composed by all thirteen narrow-width MC samples with different values of m_X . (No broad-width signal has been used for the pBDT training). Given that for the data-driven background no true resonant mass exists, a random value from the available mass values considered in this analysis is taken for each background event. The pBDT model thus can be considered a function of the true mass of the heavy particle. The pBDT output when setting the true mass of the event to m_X is used to discriminate background from resonant signal events where there is an intermediate particle of mass m_X decaying into the HH




Figure 5.13: ROC curves and AUC evaluated for the background and several (a) narrow-width and (b) broad-width signals using the pBDT classifier. The dashed gray line indicates the hypothetical ROC curve if the model makes random predictions only.

system.

Finally, in order to ensure an adequate number of trainings statistics, some of the requirements of the event selection have to be relaxed, namely the removal of the VBF jets kinematic cuts and the double b -tagger WP loosened to 70%. The nominal event selection is reinstated after training. Analogously to the non-resonant BDT training, the pBDT model hyperparameters were optimised through a grid search and are set to 600 for the number of boosting iterations, to 4 for the maximum depth of each tree, and to 0.03 for the learning rate.

The ROC curves and AUC of all narrow-width and broad-width resonant signals are shown in Figure 5.13. In general, the pBDT model performs better for higher masses, an expected behaviour given that signal and background events become more distinct for high- m_X hypotheses, with almost no overlap between signal and background in some kinematic distributions for m_X close to 5.0 TeV. The model accuracy ranges from 90% to almost 100%. The pBDT also performs well in identifying broad width signals, although with lower accuracy compared to the narrow-width signals. This was foreseen due to the degradation of the kinematic distributions in signals under the broad-width hypothesis, which increases the rate of signal events with background-like characteristics. A comparison of the reconstructed m_{HH} distributions for the resonant $m_X = 2.0$ TeV signal under the narrow-width and broad-width hypotheses are shown in Figure 5.14.

Usage of the BDTs in the analysis

In this analysis, the statistical inference (see Chapter 6) is performed using the (p)BDT histogram distribution as an observable, rather than simply using the model output to apply a selection requirement in order to enhance the signal significance in the analysis SR.




Figure 5.14: Invariant mass m_{HH} distributions for the resonant $m_X = 2.0$ TeV signal under the narrow-width (solid line) and broad-width (dashed line) hypotheses.

The choice of the binning for the binned (p)BDT distributions affects the sensitivity of the search. For an ideal performance, it is desired to have some sensitive, signal-like bins, with sufficient statistics and good counting significance of signal over background. Hence, a binning optimisation is carried out using the “transformation D” method of the ATLAS fitting package `TRExFitter` [140]. The bins edges are optimised separately in the non-resonant and resonant analyses, and in the resonant search a different binning is used for each of the thirteen mass hypotheses the pBDT was trained on. In a nutshell, in the transformation D method the upper and lower edges of a first bin candidate starts at one — the highest possible output value of the BDT models; the lower bin edge is then reduced in small steps until the factor

$$z = z_b \frac{b_{\text{bin}}}{b_{\text{total}}} + z_s \frac{s_{\text{bin}}}{s_{\text{total}}}, \quad (5.10)$$

becomes greater than one. It quantifies how much of the total background and signal yields are contained within the bin. In the equation above, b_{bin} (s_{bin}) are the background (signal) yields in the bin candidate, b_{total} (s_{total}) are the total background (signal) yields, and z_b and z_s are generic positive integers that should be optimised. Once the first bin is defined, the process is repeated recursively downwards, starting from the lower edge of the previous bin, until reaching zero — the minimum (p)BDT output value. In addition to the $z > 1$ threshold to define a new bin, it is also required that each bin contains at least one background event to avoid numerical instabilities in the statistical analysis.

For each z_b and z_s pair, a different binning is generated. A grid search over the values of z_b and z_s is performed and the binning that provides the highest significance

$$Z = \sqrt{\sum_{i=1}^{N_{\text{bins}}} \left[2(s_i + b_i) \ln \left(1 + \frac{s_i}{b_i} \right) - 2s_i \right]} \quad (5.11)$$

is selected. The binnings used in the non-resonant and resonant searches are presented in

Table 5.3: Bins edges of the non-resonant BDT model and of resonant pBDTs for each signal hypothesis. The optimal binnings were obtained using the “transformation D” algorithm of the ATLAS fitting package `TRExFitter`. If the lowest bin during the optimisation had a width lower than 0.01, it has been merged to the following one.

BDT model	Binning
Non-resonant BDT	[0.000, 0.125, 0.302, 0.478, 0.685, 1.000]
pBDT with $m_X = 1.0$ TeV	[0.000, 0.212, 0.852, 0.948, 1.000]
pBDT with $m_X = 1.1$ TeV	[0.000, 0.333, 0.882, 0.945, 1.000]
pBDT with $m_X = 1.2$ TeV	[0.000, 0.028, 0.847, 0.945, 1.000]
pBDT with $m_X = 1.3$ TeV	[0.000, 0.658, 0.947, 1.000]
pBDT with $m_X = 1.4$ TeV	[0.000, 0.312, 0.947, 1.000]
pBDT with $m_X = 1.5$ TeV	[0.000, 0.470, 0.938, 1.000]
pBDT with $m_X = 1.6$ TeV	[0.000, 0.010, 0.108, 0.538, 0.928, 1.000]
pBDT with $m_X = 1.8$ TeV	[0.000, 0.030, 0.260, 0.892, 1.000]
pBDT with $m_X = 2.0$ TeV	[0.000, 0.050, 0.845, 1.000]
pBDT with $m_X = 2.5$ TeV	[0.000, 0.560, 1.000]
pBDT with $m_X = 3.0$ TeV	[0.000, 0.272, 1.000]
pBDT with $m_X = 4.0$ TeV	[0.000, 0.043, 1.000]
pBDT with $m_X = 5.0$ TeV	[0.000, 0.042, 1.000]

Table 5.3. Details on the formula for the significance are given in Chapter 6.

5.4 Systematic uncertainties

The estimation of the systematic uncertainties related to the reconstruction of different physics objects, the background modellings, MC simulation, luminosity measurements, and others, are of fundamental importance in the statistical analysis of the data. When searching for new physics processes, deviations of the data from the background estimation might actually be covered (or at least reduced) by systematic uncertainty variations. They affect the significance of deviations between the observed and expected exclusion limits, as well as between observed and expected confidence intervals in parameter estimations.

If data is compared to the background expectation based on a binned distribution, like the output of a multivariate model, the systematic uncertainties can affect the distributions in two ways: through normalisation effects, where the overall yields change but the shape of the distributions remain the same; or through shape effects, where the shape of the distributions

Table 5.4: Systematic uncertainties considered in the analysis and the effects on the fitting variable (normalization or shape effects).

Systematics	Type	Signal	Background
General			
Luminosity	Norm	✓	
Pile-up reweighting	Norm & Shape	✓	
Large- R jets			
Jet energy scale	Norm & Shape	✓	
Jet energy resolution	Norm & Shape	✓	
Jet mass scale	Norm & Shape	✓	
Jet mass resolution	Norm & Shape	✓	
Small- R jets			
Jet energy scale	Norm & Shape	✓	
Jet energy resolution	Norm & Shape	✓	
b -tagging			
Double b -tagger	Norm & Shape	✓	
Signal modelling			
Scale	Norm & Shape	✓	
PDF+ α_s	Norm & Shape	✓	
Branching ratio $H \rightarrow b\bar{b}$	Norm	✓	
SM HH cross-sections	Norm	✓	
Parton shower	Norm & Shape	✓	
Data-driven estimations			
Background modelling	Norm & Shape		✓
MC (data-driven) statistical errors		✓	✓

change but the overall yield is unaffected. The majority of the systematic uncertainties considered in this analysis have both normalisation and shape variations components. In this analysis, whenever a distribution has to be recreated using alternative event weights, the corresponding systematic uncertainty is estimated by comparing the nominal and the alternative (p)BDT output distributions of the samples.

A summary of the systematic uncertainties affecting the analysis presented in this thesis is displayed in Table 5.4. The systematic uncertainties are categorised accordingly to their sources and the information on how they affect the distributions and in which samples they are present is also given. Given that the background is estimated from data, all systematic uncertainties, except of course those from the background modelling, affect only the signal samples. In the following sub-sections, the meaning of each the systematic uncertainties and how they are estimated are presented.

5.4.1 General

Luminosity

The Run-2 dataset used in this analysis with all the ATLAS data quality requirements [141] has an integrated luminosity of $140.1 \pm 1.2 \text{ fb}^{-1}$ [142], i.e. with an overall normalisation uncertainty of 0.83%. The total luminosity and its uncertainty was measured using the LUCID-2 detector [143].

Pile-up reweighting

The pile-up weights to correct the pile-up profile of simulated events to match those seen in data, as discussed in Section 3.3, are affected by uncertainties, such as by varying the average number of interactions per bunch-crossing in their calculation. These weight variations thus have to be included as a systematic uncertainty in the analysis. By changing the contribution of each MC event, both the total yield and shape of distributions taken from an MC sample are affected. The uncertainty related to the pile-up modelling is however found to be extremely small and to have an overall impact of less than 0.1% in all bins of the final (p)BDT distributions for all MC signal samples.

5.4.2 Jet uncertainties

The jet energy scale (JES) and jet energy resolution (JER) uncertainties affect both large- R and small- R jets and account for the differences of jet reconstruction in data and simulated collision events, as well as for detector effects and reconstruction modelling [121, 144]. For the large- R jets, the combined JES and JER uncertainties include about 36 sources of independent components, while for small- R jets this number goes up to 43. The large- R jets have additional uncertainties arising from the jet mass scale (JMS) and jet mass resolution (JMR), with a total of 19 components.

All these uncertainties are provided by the ATLAS jet calibration teams. Each of the systematic components modifies (separately) the jets, and the variation in the distributions represents the corresponding normalisation plus shape uncertainty. The overall impact of the combined JES and JMS uncertainties in the large- R jets ranges from 1% – 10% depending on the signal sample; the combined JER and JMR uncertainties in the large- R jets are small with effects smaller than 1%; and the impact of the combined JES and JER uncertainties in the small- R jets are also found to lie between 1% and 10%, depending on the signal sample.

5.4.3 Double b -tagger

The dominant systematic uncertainty in this analysis stems from the double b -tagging algorithm to identify the $H \rightarrow b\bar{b}$ decays. The scale factors to correct the MC efficiencies to those observed in data are presented in Figure 4.11 in four large- R jet p_T bins, for the 60% WP of the double b -tagger used in this analysis. If a jet has $p_T > 1000 \text{ GeV}$, the scale

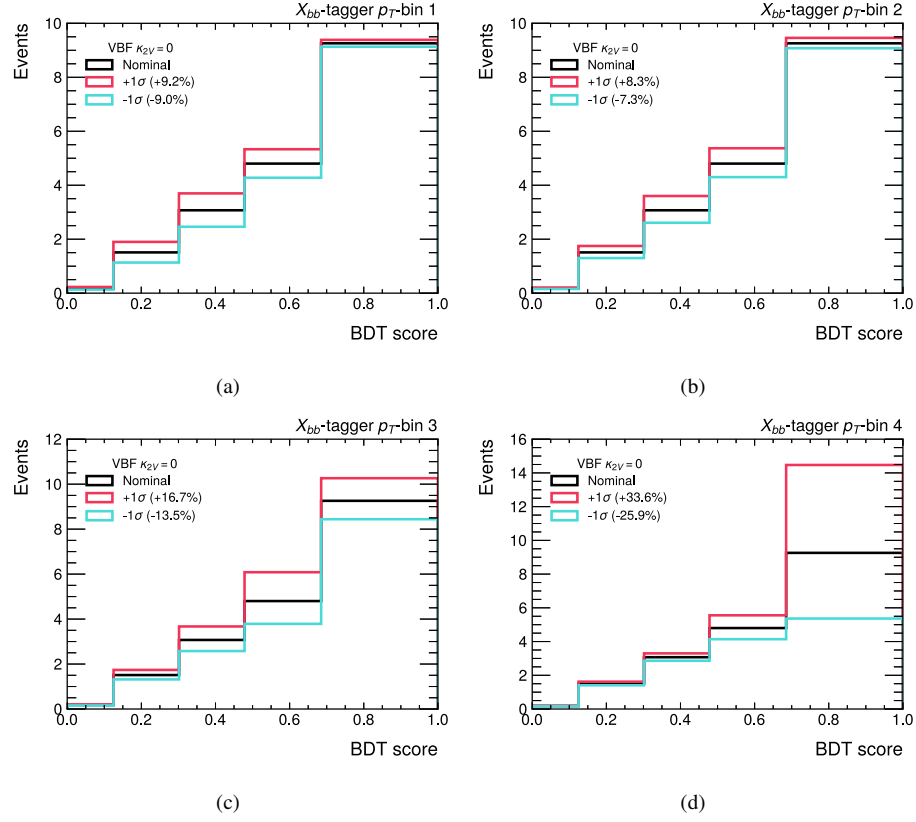


Figure 5.15: Impact of the systematic uncertainties of the double b -tagger scale factors per calibration p_T -bin on the VBF non-resonant $\kappa_{2V} = 0$ signal as a function of the BDT score.

factor from the last p_T bin of the calibration is used. As shown in the plot, the magnitude of the scale factor uncertainties are of the order of 30% in the whole p_T range. The dominant component of is the statistical uncertainty in the tagger calibration, which indicates that the scale factors from different p_T bins can be safely treated as uncorrelated. In this approach, the double b -tagger has four sources of systematic uncertainties, one for each p_T bin. The impact of each of the four tagger systematic uncertainty components on the VBF $\kappa_{2V} = 0$ sample as a function of the BDT score is show in Figure 5.15.

If an event has the two Higgs-bosons candidates with p_T values within the same bin of the calibration, the impact of the systematic uncertainty will be up to 70%, as the jets would be treated as fully correlated. This is the reason why the last p_T bin ($p_T > 600$ GeV) has the largest relative impact of the tagger uncertainties, as a large fraction of signal events has both Higgs-boson candidates in that bin.

The more conservative approach considering fully correlated jets regardless of their p_T bin would reduce the sensitivity of this analysis by about 25%.




Figure 5.16: BDT distributions in data in the VR 2Pass and VR 1Pass regions using (a) the optimal binning for the non-resonant analysis and (b) the two-bin smoothed version of the binning. The shape systematic uncertainties are estimated from the ratio pads [75].

5.4.4 Background modelling

The uncertainties related to the background modelling of Section 5.2 rank among the most impactful in this analysis. While the CR 2Pass and CR 1Pass data events are used to determine the normalisation weight w that is used to reweight the SR 1Pass data distributions to estimate the expected background in SR 2Pass, the VR is used to estimate the systematic uncertainties to account for the incomplete knowledge of how the shape of the distributions can differ in the different mass-plane regions and between 1Pass and 2Pass regions. First, a normalisation uncertainty Δw for w is estimated by comparing the yield ratios from the VR and CR via

$$\Delta w = 1 - \frac{N(\text{VR}, 2\text{Pass})/N(\text{VR}, 1\text{Pass})}{N(\text{CR}, 2\text{Pass})/N(\text{CR}, 1\text{Pass})}, \quad (5.12)$$

where $N(\text{CR}, 2\text{Pass})$ stands for the total data yield in the CR 2Pass region, and analogously for the others. The Δw is estimated as 12.3% for the non-resonant analysis and 14.4% for the resonant analysis⁴, producing the uncertainties already presented earlier in Equation 5.8.

Moreover, a shape systematic uncertainty is estimated to cover potential shape differences between the 2Pass and 1Pass regions, which were assumed to provide similar kinematic distributions. The shape uncertainties are derived comparing the data (p)BDT distributions in the VR 2Pass and VR 1Pass regions. This is illustrated in Figure 5.16(a) for the non-resonant analysis. The ratio pad indicates huge statistical fluctuations mainly because of the low yield

⁴ Recall the event selection of the non-resonant and resonant analyses differs on the veto of events passing the SR of the previous resolved analysis.

of VR 2Pass. In order to mitigate effects from limited statistics in the shape uncertainty estimation, a two-bin smoothing method is applied to reduce statistical fluctuations from the BDT distribution. This is done by merging the first three bins and last two bins of the BDT distribution, as seen in Figure 5.16(b). The shape uncertainties are then estimated by the ratio pad in the smoothed distributions and propagated to the original distributions with the optimal binning. In the non-resonant analysis, the shape uncertainties are obtained as 3% for the first three BDT distribution bins and 11% for the last two BDT distribution bins and treated as correlated. The process is completely analogous for the estimation of the shape uncertainties in the several pBDT distributions used for the resonant analysis. The two-bin smoothing method is applied for the mass points where the optimal binning has more than two bins. Considering all mass hypotheses and optimal binnings, the shape uncertainties in the pBDT distributions bins are found to range between 1% and 25%, with the large majority being below 10%.

5.4.5 Signal modelling

PDF+ α_s and scale variations

The combined PDF+ α_s uncertainty is a quadratic sum of the individual contributions from α_s variations and variations of the several different PDF parameters of the PDF set. The nominal PDF used in this analysis is the NNPDF set at next-to-leading order assuming $\alpha_s(m_Z) = 0.118$, NNPDF30_nlo_0118. The α_s uncertainty is estimated using the MC weights obtained with the NNPDF set evaluated at $\alpha_s(m_Z) = 0.117$ and $\alpha_s(m_Z) = 0.119$, re-creating the (p)BDT distributions with the alternative event weights and comparing to the nominal distributions. The process is analogous for the PDF variations, where the uncertainty is estimated from the largest difference between the nominal PDF set and its 100 variations with varied parameters values, corresponding e.g. to different fit methodologies, treatment of heavy-flavour quarks, etc. Since the PDF+ α_s variations change the MC event weights, its corresponding systematic uncertainty affects both the yield and the shape of the (p)BDT distributions. The overall impact was found to be of the order of 1% for all non-resonant and resonant samples.

The renormalisation and factorisation scales uncertainties account for the missing higher-order contributions in the perturbative series describing the process, as the dependence of the scales on physics processes are a consequence of the incomplete knowledge of the series. The two scales are varied by factors of 0.5 and 2.0 within their nominal values, thus generating seven alternative weight values for each MC event. The final uncertainties are obtained from the largest overall difference found between the nominal and alternative distributions. The overall impact of the renormalisation and factorisation scales range from 1% to 5% depending on the signal sample.

Branching ratio $H \rightarrow b\bar{b}$

The branching ratio of the Higgs boson decaying to a pair of b -quarks assuming $m_H = 125$ GeV is calculated as $0.582^{+3.4\%}_{-3.5\%}$ [145], where the uncertainty includes both theoretical errors arising from the missing higher-order terms in the perturbative series to derive the Higgs-boson partial widths, as well as parametric errors arising from the variations of QCD parameters in the calculation of the branching ratio, namely the strong coupling and the masses of the heavy quarks [146]. The theoretical and parametric uncertainties are added linearly. The total branching ratio uncertainty is applied as a normalisation variation for all signal samples.

SM HH cross-sections

In the non-resonant search, the theoretical uncertainties related to the HH cross-section must also be included. These uncertainties are provided centrally by the LHC Higgs Working Group and harmonised across all HH searches. It includes the relative uncertainties from PDF+ α_s , renormalisation and factorisation scale variations, and, in the case of the ggF production mode, the top-quark mass determinations. At $\sqrt{s} = 13$ TeV and for $m_H = 125$ GeV, the cross-section for the VBF production of Higgs-boson pairs has an uncertainty of $^{+0.03\%}_{-0.04\%}$ and $\pm 2.1\%$ arising from scale variations and PDF+ α_s , respectively. For the ggF production of Higgs-boson pairs, the uncertainty on the SM cross-section is $^{+6\%}_{-23\%}$ for scale and m_t variations, and $\pm 3.0\%$ for PDF+ α_s variations.

This uncertainty is not applicable in the resonant analysis, as in that case the cross-section itself is a parameter of interest when setting upper limits (see Chapter 6).

Parton shower

The systematic uncertainties due to the modelling of the parton shower are evaluated by comparing samples generated with Pythia 8 (nominal) and Herwig 7 (alternative). Because of the high CPU-cost in simulating collision events, only a fraction of the signal hypotheses used in this analysis has a sample generated with the alternative parton shower. Therefore, proxy representative samples and interpolations (in the case of the resonant search) are considered in order to estimate the uncertainty for all signals.

In the non-resonant search, the proxy $\kappa_{2V} = 0$ signal is used to estimate the uncertainty for all non-resonant VBF samples, including the VBF SM one. The BDT distributions of the nominal and alternative $\kappa_{2V} = 0$ samples are compared and a bin-by-bin uncertainty is derived. To ensure that normalisation uncertainties are coming exclusively from the difference in the event selection efficiencies for different choices of the parton shower settings, but not from possible differences in the total number of simulated events, each BDT histogram is first normalised to the initial total number of (weighted) events of its sample. The final estimation of the parton shower systematic uncertainties for non-resonant VBF samples ranges from 8% to 17%, depending on the bin of the BDT distribution. In the case of the SM ggF sample, a bin-by-bin estimation is severely affected by the low statistics of the sample after applying




Figure 5.17: Estimated parton shower normalisation uncertainties in the narrow-width resonant analysis. The blue points indicate estimations from the direct comparison between nominal and alternative parton shower samples. The red dashed line is obtained via linear interpolation. The same uncertainties are used for broad-width samples.

the whole event selection. Hence, only a normalisation uncertainty derived by the total yield difference of the nominal and alternative samples is considered. The value to be used for SM ggF is 21%.

In the resonant search, an interpolation is required to derive the uncertainties for all the mass points, as the very different kinematic distributions of the resonant signals do not allow the use of only one proxy sample. Moreover, a bin-by-bin estimation would be challenging as the binning of the pBDT distributions is optimised separately for each mass hypothesis (see Table 5.3), which makes an interpolation non-trivial. Thus, as done for the SM ggF sample, only a normalisation uncertainty derived from a comparison of the total yield difference of the nominal and alternative samples is considered. The estimated parton shower normalisation uncertainties for resonant narrow-width samples are depicted by the blue dots in Figure 5.17. The values range from 15% to 25%. The error bars indicate only statistical fluctuations and are not meant to represent any additional systematic uncertainties. For mass points for which no alternative sample is available, the uncertainty is estimated via linear interpolation. The same parton shower systematic uncertainties are used for the broad-width samples.

Results and Statistical Interpretation

The results of the analysis presented in this thesis and their statistical interpretation are presented in this chapter. In Section 6.1, the event yields and (p)BDT response distributions of observed data and expected background are shown. The statistical tools required for the interpretation of the results in terms of coupling constraints and cross-section upper limits are introduced in Section 6.2. The final results of the non-resonant and resonant searches are given in Sections 6.3 and 6.4, respectively. A discussion of the results obtained can be found in Section 6.5.

6.1 Unblinded distributions

The event reconstruction and selection for the non-resonant and resonant VBF HH analyses described in Chapters 4 and 5 are applied both to the MC generated signal samples and collision data. A comparison between the expected background and the observed data yields is presented in Table 6.1. The observed data yield is higher in the resonant search because of the extra veto applied in the non-resonant analysis discussed in Section 5.1.5; this veto also affects the background modelling when comparing 2Pass and 1Pass data distributions, which explains the different expected background yields in the non-resonant and resonant searches.

In the non-resonant search, the contributions from the SM VBF and ggF samples are negligible compared to the total background yield. The expected background and its total uncertainty in the resonant search is provided for three mass hypotheses, as the shapes of the systematic uncertainties in the data-driven background modelling depend on the value of m_X . In all cases the observed data yields agree well with the background expectations within uncertainties. The BDT distributions of observed data, expected background, SM VBF, SM ggF, and the BSM scenario $\kappa_{2V} = 0$ in the non-resonant search are shown in Figure 6.1, with the use of the optimal binning as earlier presented in Table 5.3. No data has been observed in the rightmost, signal-like bin. The pBDT distributions of observed data, expected background,

Table 6.1: Total yields of observed data, expected background, and SM VBF and ggF signals, in the non-resonant and resonant searches. The total uncertainty of the expected background in the resonant analysis depends on the mass hypothesis.

Search	Observed data	Expected background	SM VBF	SM ggF
Non-resonant	21	20.78 ± 1.66	0.02	0.06
Resonant	23	19.97 ± 2.94 (1.0 TeV)	-	-
		19.97 ± 2.90 (2.0 TeV)		
		19.97 ± 2.90 (5.0 TeV)		

and signal under the narrow-width assumption evaluated at three choices of m_X are presented in Figure 6.2. Similarly to what has been seen in the non-resonant search, there is a deficit of observed data in the rightmost signal-like bin for all mass points $m_X \geq 1.6$ TeV. These distributions are obtained before the profile likelihood fit to data discussed in the following sections.

In view of the good qualitative agreement between the observed data and the expected background in the (p)BDT distributions, the goal now is to set limits on the signal hypotheses (e.g., determining constrain intervals for the κ_{2V} and κ_λ couplings and/or calculating upper limits on the signal cross-section). All the relevant tools for such task are presented in the next section.

6.2 Statistical analysis

In many particle physics searches, the aim is to detect deviations from a predicted background originating either from a SM signal process that has not yet been observed, as was the case for the Higgs-boson discovery, or from BSM processes. The general procedure used to search for undiscovered phenomena is to first define what is called the null hypothesis, the scenario where the observed data is correctly described by the background expectation, and the alternative hypothesis, that additionally includes the signal process under study.¹ This signal can be either a SM process that has not yet been observed, as is the case of SM HH production, or some BSM process. The level of agreement between the observed data and a given hypothesis \mathcal{H} is quantified by calculating the so-called p -value, the probability that an observed incompatibility between data and \mathcal{H} does not arise merely from statistical fluctuations. The hypothesis \mathcal{H} is said to be rejected if the observed p -value is below a certain threshold [147].

In the particle physics community, the p -value is commonly converted into an equivalent significance Z , defined as the number of standard deviations σ that a Gaussian-distributed random variable is found above its mean. The precise formula relating the p -value and Z

¹ The definitions of what are the null and alternative hypotheses can however vary depending on the approach used in the statistical analysis.




Figure 6.1: Pre-fit BDT distributions of observed data, expected background, and the non-resonant VBF $\kappa_{2V} = 0$, SM VBF, and SM ggF signals. The distributions of the SM VBF and SM ggF signals were scaled for visualisation purposes. All sources of uncertainties have been included. No data has been observed in the rightmost signal-like bin.

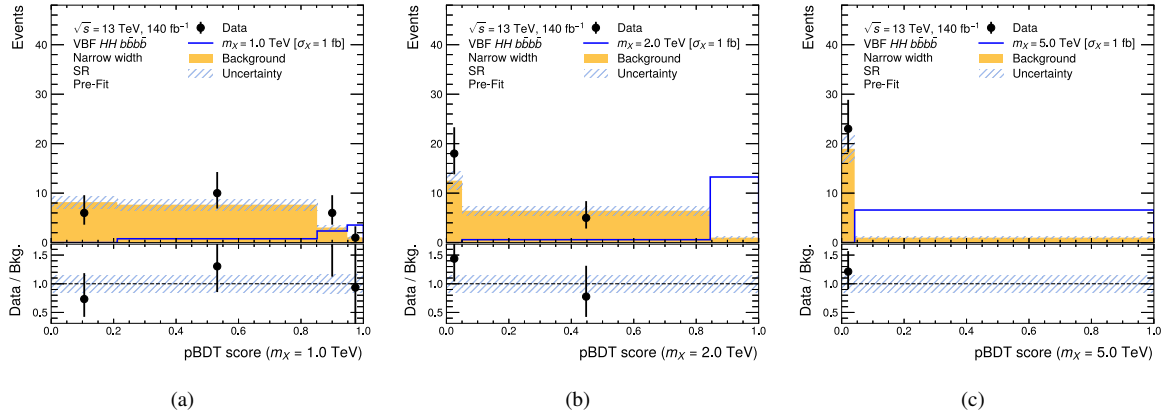


Figure 6.2: Pre-fit pBDT distributions of observed data, expected background, and signal under the narrow-width assumption evaluated at three different values of m_X . The total signal cross-section has been normalised to 1 fb. All sources of uncertainties have been included. No data has been observed in the rightmost signal-like bin in all scenarios where $m_X \geq 1.6$ TeV.

reads

$$Z = \Phi^{-1}(1 - p), \quad (6.1)$$

where Φ^{-1} is the inverse function of the cumulative Gaussian distribution. An illustrative representation of the equation relating Z and the p -value is shown in Figure 6.3(a). The common practice in particle physics searches is to claim the discovery of a new process if the deviation from the null hypothesis has a significance of at least 5σ , or, equivalently, a p -value smaller than 2.87×10^{-7} . In many situations, however, the signal process under study may just be too small to provide such a 5σ deviation in the observed data. The common approach in these cases is to test the agreement of data with the alternative hypothesis that a signal of a given size exists with a *signal strength* μ — a multiplicative factor to mimic an enhanced predicted signal cross-section. An upper limit on the signal strength is then defined as the value of μ above which the p -value drops below 0.05 ($Z = 1.64$).

The likelihood

A central object widely used in particle physics analyses is the *likelihood*, which represents the probability function that describes how well data agree with a hypothesis \mathcal{H} . The simplest scenario is a standard counting experiment that essentially only compares the expected and observed event yields without any consideration on differential distributions; the probability of observing n events given an expected number of events λ is described by the Poisson distribution

$$\text{Pois}(n|\lambda) = \frac{\lambda^n}{n!} e^{-\lambda}. \quad (6.2)$$

The expected number of events is typically written as the sum of signal (s) and background (b) via

$$\lambda(\mu, \theta) = \mu \cdot s(\theta) + b(\theta), \quad (6.3)$$

where the parameter of interest μ in the analysis is the aforementioned signal strength and θ is an ensemble of N_{NP} *nuisance parameters* to account for the statistical and systematic uncertainties that affect both the signal and background yields.

In a more realistic analysis, some variable (e.g. BDT output score) is chosen to construct a histogram with N_{bins} bins for data, signal, and background. Each bin of the histogram is then treated as an independent counting experiment with n_i observed events. The likelihood function $L(\mu, \theta)$ is given by the product of the Poisson probabilities multiplied by constraint functions $f_j(\theta_j)$, one for each nuisance parameter θ_j [147]:

$$L(\mu, \theta) = \prod_{i=1}^{N_{\text{bins}}} \frac{(\mu s_i(\theta) + b_i(\theta))^{n_i}}{n_i!} e^{-(\mu s_i(\theta) + b_i(\theta))} \prod_{j=1}^{N_{\text{NP}}} f_j(\theta_j). \quad (6.4)$$

The functions $f_j(\theta_j)$, which can be e.g. Gaussian or log-normal constraints, are fundamental to control the variation of the nuisance parameters when maximising the likelihood function




Figure 6.3: Illustration of (a) the relation between the p -value and the significance Z in the Gaussian distribution, and (b) the relation between the p -value and the observed test statistic $t_{\mu,obs}$.

to adjust the signal and background to best fit the observed data.

The test statistics

To test a hypothesized value of μ , the common approach in LHC analyses is to use the profile likelihood ratio

$$\Lambda(\mu) = \frac{L(\mu, \hat{\hat{\theta}})}{L(\hat{\mu}, \hat{\theta})}. \quad (6.5)$$

In the definition above, $\hat{\hat{\theta}}$ is the conditional maximum-likelihood estimator of the nuisance parameters and denotes the values of θ that maximise the likelihood function at a fixed value of μ , while $\hat{\mu}$ and $\hat{\theta}$ are the unconstrained fitted values that maximise the likelihood function. By construction, $0 \leq \Lambda(\mu) \leq 1$, with smaller values of $\Lambda(\mu)$ indicating a larger disagreement between data and the signal-plus-background hypothesis at a given μ , $\mathcal{H}(\mu s + b)$. Within this definition for the profile likelihood ratio, the likelihood at a chosen μ is compared to the best-fit result.

For convenience, the log of the profile likelihood ratio defines the test statistic

$$t_\mu \equiv -2 \ln \Lambda(\mu). \quad (6.6)$$

It is possible to directly measure the discrepancy between the observed data and the hypothesis $\mathcal{H}(\mu s + b)$ through the calculation of the p -value

$$p_\mu = \int_{t_{\mu,obs}}^{\infty} f(t_\mu|\mu) dt_\mu, \quad (6.7)$$

where $t_{\mu,\text{obs}}$ is the test statistic obtained for observed data and $f(t_\mu|\mu)$ is the probability distribution function of the test statistic t_μ . An illustration of how the p -value is calculated from the test statistic t_μ is shown in Figure 6.3(b).

The test statistic as defined in Equation 6.6 a priori allows positive and negative values of μ . A best-fit $\hat{\mu} < 0$ may occur for instance if the inclusion of a new signal sample reduces the total expected yield due to interference effects with the background processes or if there is some under fluctuation of data compared to the expectation, but may also have an unphysical interpretation when arising from fit instabilities because of low statistics. If by construction the signal sample should only increase the total event yield, the constraint $\mu \geq 0$ is taken into account using the alternative test statistic \tilde{t}_μ defined as

$$\tilde{t}_\mu = \begin{cases} -2 \ln \frac{L(\mu, \hat{\theta})}{L(\hat{\mu}, \hat{\theta})} & \text{if } \hat{\mu} \geq 0 \\ -2 \ln \frac{L(\mu, \hat{\theta})}{L(0, \hat{\theta}(0))} & \text{if } \hat{\mu} < 0 \end{cases}, \quad (6.8)$$

which makes the best level of agreement between the hypothesis and the data occur at $\mu = 0$ if $\hat{\mu} < 0$. The p -value is calculated analogously to Equation 6.7.

For the purposes of setting upper limits on the signal strength μ , the test statistic has to be further modified in order to avoid rejecting the hypothesis $\mathcal{H}(\mu s + b)$ in the case of upward fluctuations of data ($\hat{\mu} > \mu$). Essentially, the test statistic should still indicate good agreement between data and $\mathcal{H}(\mu s + b)$ for all $\mu < \hat{\mu}$, so by construction the alternative test statistic is set to zero — which is the value obtained at $\mu = \hat{\mu}$ — for all μ smaller than $\hat{\mu}$. The test statistic used for upper limits setting is thus the one-sided log-likelihood ratio

$$\tilde{q}_\mu = \begin{cases} \tilde{t}_\mu & \text{if } \hat{\mu} \leq \mu \\ 0 & \text{if } \hat{\mu} > \mu \end{cases} = \begin{cases} -2 \ln \frac{L(\mu, \hat{\theta})}{L(\hat{\mu}, \hat{\theta})} & \text{if } 0 \leq \hat{\mu} \leq \mu \\ -2 \ln \frac{L(\mu, \hat{\theta})}{L(0, \hat{\theta}(0))} & \text{if } \hat{\mu} < 0 \\ 0 & \text{if } \hat{\mu} > \mu \end{cases}. \quad (6.9)$$

The corresponding p -value is calculated via

$$p_\mu = \int_{\tilde{q}_{\mu,\text{obs}}}^{\infty} f(\tilde{q}_\mu|\mu) d\tilde{q}_\mu, \quad (6.10)$$

now written in terms of the probability distribution function of \tilde{q}_μ .

Limit setting

A natural way of defining an exclusion limit on the hypothesis $\mathcal{H}(\mu s + b)$ is via a method referred to as the CL_{s+b} method. The p -value $p_\mu^{(\text{CL}_{s+b})}$ is calculated directly using Equation 6.10 assuming $\mathcal{H}(\mu s + b)$. If it is found to lie below a predefined threshold α , the hypothesis

is said to be excluded. The common choice is setting $\alpha = 0.05$, in which the hypothesis is excluded at 95% confidence level (CL).

There are however some drawbacks in the CL_{s+b} method when the number of expected background events is significantly greater than the predicted signal yield and/or when the number of observed events is smaller than the expected background yield. In these scenarios, the CL_{s+b} method may lead to a naive exclusion of the signal plus background hypothesis even if there is little or no sensitivity in the search. In order to find to what degree the alternative hypothesis $\mathcal{H}(\mu + b)$ can be distinguished from the background-only null hypothesis $\mathcal{H}(\mu = 0)$, one can calculate the probability that, assuming $\mathcal{H}(\mu = 0)$ is true, a measurement would yield a value \tilde{q}_μ that indicates better agreement between data and the alternative hypothesis. This probability is given by the p -value

$$p_\mu^{(\text{CL}_b)} = 1 - \int_{\tilde{q}_{\mu,obs}}^{\infty} f(\tilde{q}_\mu|0) d\tilde{q}_\mu = \int_{-\infty}^{\tilde{q}_{\mu,obs}} f(\tilde{q}_\mu|0) d\tilde{q}_\mu. \quad (6.11)$$

In the CL_s method [148], conventionally used in many LHC analyses, the exclusion criterion is based on the ratio of p -values

$$p_\mu^{(\text{CL}_s)} = \frac{p_\mu^{(\text{CL}_{s+b})}}{1 - p_\mu^{(\text{CL}_b)}}, \quad (6.12)$$

where the penalty term $(1 - p_\mu^{(\text{CL}_b)})^{-1}$ is large if the alternative and null hypotheses are almost indistinguishable. Notice that, by construction, $p_\mu^{(\text{CL}_s)} \geq p_\mu^{(\text{CL}_{s+b})}$, with $p_\mu^{(\text{CL}_s)} \approx p_\mu^{(\text{CL}_{s+b})}$ only when the alternative and null hypotheses are clearly different from each other. The CL_s method thus commonly yields more conservative upper limits compared to the CL_{s+b} method.

For future reference, an observed limit is defined as the limit obtained when using real observed data. Expected limit refers to the limits when using an *Asimov* dataset that is identical to the background expectation to represent data. For a precise limit, a fine grid of μ -values has to be scanned around the region where the p -value crosses the threshold α .

A word about the test statistic distributions

A precise knowledge of the sampling distributions $f(\tilde{q}_\mu|\mu)$ and $f(\tilde{q}_\mu|0)$ is crucial for reliable calculations of $p_\mu^{(\text{CL}_{s+b})}$ and $p_\mu^{(\text{CL}_s)}$. A possible way of estimating them is numerically via MC generation of a large number of pseudo-experiments, commonly referred to as toys. In this approach, each pseudo-experiment is conducted by creating pseudo-data from variations of the observed data using a Poisson distribution with $\mu s + b$ expected events in each bin of the histogram. The observed test statistic $\tilde{q}_{\mu,obs}$ is calculated for each set of pseudo-data and, for a sufficiently large number of pseudo-experiments, a one-dimensional histogram in \tilde{q}_μ can be constructed to represent the corresponding probability density functions. Given that

the sampling distributions should be known for each value of μ and signal hypothesis tested in the analysis, their numerical estimations are computationally very intensive.

However, thanks to the results obtained by A. Wald [149] and S. Wilks [150], analytic approximations of the sampling distributions can be derived if the number of observed events is sufficiently large. It has been found out that the numerical estimations and the analytic formula already produce similar results for sample sizes of $O(10)$ events. The asymptotic formulas for the probability density functions of several test statistics are available in Ref. [147].

A word about the significance

Having the analytic approximations of the samples distributions also allows writing the significance at which the background-only hypothesis (here identified as the null hypothesis) is rejected as [147]

$$Z_0 = \sqrt{q_0}, \quad (6.13)$$

where q_0 stands for

$$q_0 = \begin{cases} -2 \ln \frac{L(0, \hat{\theta})}{L(\hat{\mu}, \hat{\theta})} & \text{if } \hat{\mu} \geq 0 \\ 0 & \text{if } \hat{\mu} < 0 \end{cases}. \quad (6.14)$$

Assuming the simple scenario where the probability function of data agreeing with an alternative hypothesis $\mathcal{H}(\mu s + b)$ is described by one single Poisson distribution and in the absence of nuisance parameters, i.e.,

$$L(\mu) = \frac{(\mu s + b)^n}{n!} e^{-\mu s + b}, \quad (6.15)$$

the best fit value of the signal strength reads $\hat{\mu} = (n - b)/s$. Hence, the significance can be evaluated for positive $\hat{\mu}$ as

$$Z_0 = \sqrt{2 \left[n \ln \left(\frac{n}{b} \right) + b - n \right]}. \quad (6.16)$$

Considering that data has been observed as $n = s + b$, one obtains

$$Z_0 = \sqrt{2(s + b) \ln \left(1 + \frac{s}{b} \right) - 2s}, \quad (6.17)$$

which in the limit of $s \ll b$ becomes

$$Z_0 = \frac{s}{\sqrt{b}} \left[1 + O\left(\frac{s}{b}\right) \right]. \quad (6.18)$$

This asymptotic approximation for the calculation of the significance, although not used for final results presented in the following sections as it is a simplified approach, is widely used

to help designing the optimal strategy for the analysis, as has been discussed in Chapter 5.

A word about the nuisance parameters

The fitted nuisance parameters carry relevant information that can indicate potential issues in the fit or in the estimation of some systematic uncertainties.

The pull of a nuisance parameter θ_j is defined as

$$x_j = \frac{\hat{\theta}_j - \theta_j^0}{\Delta\theta_j}, \quad (6.19)$$

where $\hat{\theta}_j$ is the fitted value, θ_j^0 is its initial value and $\Delta\theta_j$ its initial uncertainty. These latter two values depend on the constraint functions used in Equation 6.4; for the typical Gaussian constraint, both values are zero, which means $x_j = \hat{\theta}_j$. A very large pull x_j often indicates that the data does not agree well with the background expectation, and therefore the fit adjusts the background by varying it within the particular uncertainty parametrised by θ_j , so that the post-fit background better matches the data. This might be the case if there is either a severe mismodelling of the background expectation or in the presence of signal. Moreover, a small pull error given by the standard deviation of the fitted $\hat{\theta}_j$ can indicate that the dataset to which to fit is being done may contain information to reduce to size of the corresponding systematic uncertainty compared to what has been obtained/modelled. In other words, this might suggest that the corresponding systematic uncertainty is too conservative.

Finally, regardless whether the nuisance parameters were treated as correlated or uncorrelated, the post-fit can introduce correlations among them due e.g. to similar shape variations in the distributions.

6.3 Non-resonant interpretation

In this analysis, a binned maximum likelihood fit to the BDT distribution in the analysis SR is performed including all systematic uncertainties parametrised as nuisance parameters. The fit is performed assuming the background-only hypothesis and the resulting BDT response distributions are shown in Figure 6.4. No clear qualitative deviation of observed data from the expected background has been detected. In view of the extremely low SM signal yields compared to the background expectation (recall Table 6.1), there is no sensitivity to detect SM Higgs-boson pairs production. Under the hypothesis that there is no VBF HH production, the observed (expected) upper limit on the SM VBF signal strength is found to be $\mu_{\text{SM,VBF}} < 551(672)$ at 95% CL. The expected 1σ (2σ) interval of $\mu_{\text{SM,VBF}}$ is $[484, 1075]$ ($[361, 1829]$). To calculate this limit, the SM ggF HH production is included as an extra background. None of the nuisance parameters related to the systematic uncertainties have a large pull or are largely constrained, as seen in Figure 6.5.




Figure 6.4: Post-fit BDT output score distributions in data, the expected background, SM VBF and SM ggF HH production, and the BSM scenario $\kappa_{2V} = 0$. All sources of uncertainties have been included. The distributions of the SM VBF and SM ggF signals were scaled for visualisation purposes. The observed data agrees well with the total background expectation [75].

The low sensitivity to the SM VBF signal with a search relying on the boosted topology only was already expected (recall Section 2.3.3). However, the signal cross-section is significantly enhanced for BSM values of κ_{2V} due to the reduction of interference effects, and therefore it is possible to define an interval in κ_{2V} in which the signal-plus-background hypothesis is not excluded by the observed data at a certain confidence level.

There are two methods to obtain confidence intervals on the parameter κ_{2V} . The first, and nominal, method to constrain κ_{2V} relies on the calculation of the observed and expected test statistic $t_\mu = -2 \ln \Lambda$, defined in Equation 6.6, as a function of κ_{2V} , while keeping the other Higgs-boson couplings fixed to their SM values. A value of κ_{2V} is excluded at 95% CL (68% CL) if the calculated test statistic is larger than $t_\mu = 3.84$ ($t_\mu = 1.0$)². The BDT response distribution and total cross-section for a signal sample with generic values of κ_{2V} , κ_λ , and κ_V are obtained using the technique described in Section 3.3.2. The observed and expected t_μ curves as a function of κ_{2V} are shown in Figure 6.6(a). The interval of κ_{2V} is estimated to be $[0.52, 1.52]$ ($[0.37, 1.71]$ expected) at 95% CL. Moreover, the scenario $\kappa_{2V} = 0$ is excluded with 3.4σ (2.9σ expected) of significance. The observed results are more stringent than the expected results because of the already mentioned deficit of observed data in the signal-like bin of the BDT distribution.

² These numbers are obtained numerically based on Equation 6.7 and depend on the number of degrees of freedom of the statistical analysis.

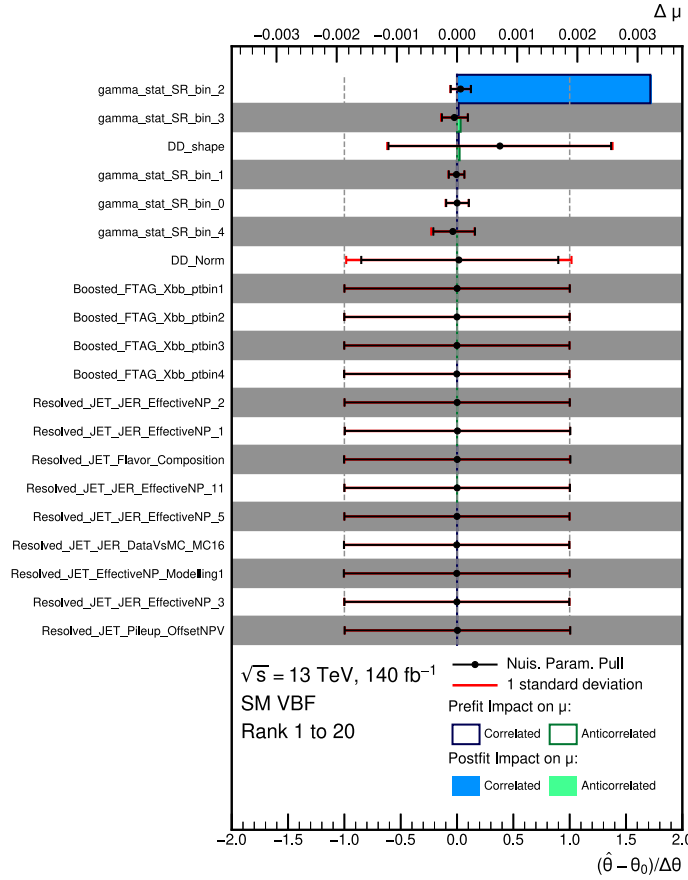


Figure 6.5: Pulls (lower x-axis) of the nuisance parameters that parametrise the systematic uncertainties or the uncorrelated statistical uncertainties at each distribution bin (prefix `gamma_`) and their impact $\Delta\mu$ on the signal strength (upper x-axis) for the VBF SM signal. No fitted nuisance parameter related to a systematic uncertainty has either a large pull or is over-constrained. Only the top 20 most impactful uncertainties are included in the plot. The names of the nuisance parameters are described in Appendix B.

Another approach to estimate the κ_{2V} interval is to directly determine the upper limits at 95% CL on the VBF $pp \rightarrow HH$ cross-section $\sigma_{\text{VBF},HH}$ as a function of κ_{2V} and compare it with the theoretically predicted cross-section. The regions in κ_{2V} where the theory predicted cross-section is higher than the observed (expected) cross-section upper limit are said to be excluded at 95% CL. The cross-section upper limits as a function of κ_{2V} are presented in Figure 6.6(b). In this method, the extracted interval of κ_{2V} at 95% CL is $[0.46, 1.57]$ ($[0.33, 1.70]$ expected).

The interval obtained for κ_{2V} is wider in the cross-section upper limit method when compared to the likelihood approach. This is mainly due to the fact that, while the likelihood approach is based on the CL_{s+b} method, the cross-section approach uses the CL_s method, which by construction produces more conservative results. The hypotheses tested in each of these approaches are also distinct. While in the likelihood method the presence of signal




Figure 6.6: Derivation of the confidence interval on κ_{2V} using (a) observed and expected values of the test statistic $-2 \ln \Lambda$ as a function of κ_{2V} in the boosted analysis. The dashed horizontal lines indicate the thresholds above which a hypothesis is excluded at 68% and 95% CL. With the alternative method (b) using the 95% CL upper limit on the VBF $pp \rightarrow HH$ cross-section as a function of κ_{2V} . The solid and dashed lines indicate respectively the observed and expected limits. The cyan and yellow bands represent the 1σ and 2σ intervals of the expected limit. The star marker indicates the SM prediction. In both approaches the κ_λ and κ_V coupling strengths are fixed to their SM values.

in data is assumed, this is not the case in the cross-section upper limit approach. Finally, in the cross-section upper limit approach no uncertainties on the calculated cross-section value are included as the cross-section itself is a fit parameter, given its relation to the signal strength. These theory uncertainties are, however, relatively small and do not give meaningful contributions to the final results. As a general ATLAS recommendation for Higgs-boson analyses, the final results for coupling estimations/constraints should be based on the likelihood scan method.

Combination with previous resolved analysis

Even more stringent constraints on κ_{2V} can be derived through a combination with the previous non-resonant ggF+VBF resolved analysis [73], where each of the b -jets from the $H \rightarrow b\bar{b}$ decays were reconstructed as individual small- R jets. The resolved analysis provides in particular better sensitivity to the SM VBF and ggF signals, and to variations of the trilinear coupling strength κ_λ . This is because, on contrary to what is seen for κ_{2V} variations in Figure 2.10(a), variations of κ_λ tend to produce a softer m_{HH} spectrum. As a consequence, the lower energetic Higgs-boson decays are hardly captured by a single large- R jet, hence making a resolved-based strategy better suited for the event reconstruction. Full orthogonality between the boosted and resolved analyses is ensured by removing from the boosted SR all events that are also inside the resolved SR (recall Section 5.1.5). All experimental systematic uncertainties that are present in both analyses are fully correlated in the combined fits.




Figure 6.7: Observed (a) and expected (b) distributions of the likelihood $-2 \ln \Delta$ as a function of κ_{2V} in the boosted (blue dashed), resolved (green dotted), and combined (black solid) analyses. The dashed horizontal lines indicate the thresholds above which a hypothesis is excluded at 68% and 95% CL. Both κ_λ and κ_V are fixed to their SM values [75].

With a dedicated analysis targeting the ggF HH production mode as well now included in the combined channel, the observed (expected) upper limit on the signal strength of the SM ggF+VBF Higgs-boson pair production $\mu_{\text{SM}, \text{ggF}+\text{VBF}}$ is estimated to be 5.3 (8.1) at 95% CL. The expected 1σ (2σ) interval of $\mu_{\text{SM}, \text{ggF}+\text{VBF}}$ is [5.8, 12.2] ([4.3, 19.1]). There is therefore still not enough sensitivity to detect SM Higgs-boson pair production.

The combined boosted and resolved analyses yield a κ_{2V} interval [0.55, 1.49] ([0.37, 1.67] expected) at 95% CL using the likelihood scan method. In the combined analysis, the $\kappa_{2V} = 0$ scenario is also excluded with an observed (expected) significance of 3.8σ (3.3σ), a slight improvement compared to the boosted-only analysis. The observed and expected $-2 \ln \Delta$ distributions as a function of κ_{2V} in the boosted, resolved, and combined channels are shown in Figure 6.7. In the cross-section upper limit approach, the interval of κ_{2V} is estimated to be [0.49, 1.54] ([0.38, 1.66] expected) at 95% CL. The 95% CL cross-section $\sigma_{\text{VBF}, HH}$ upper limits as a function of κ_{2V} in the boosted, resolved, and combined channels are presented in Figure 6.8. This figure shows that the resolved analysis has much higher sensitivity than the boosted analysis for $\kappa_{2V} \approx 1.0$, but has significantly less sensitivity for non-SM values of κ_{2V} . The combined result is thus heavily dominated by the new boosted search of Higgs-boson pairs.

In addition to estimating κ_{2V} , the combined channel also provides enough sensitivity to set constraints on the trilinear Higgs-boson self-coupling. The coupling strength κ_λ interval is estimated to be $[-3, 11]$ ($[-5, 11]$ expected) at 95% CL in the likelihood method, with the resolved channel dominating the sensitivity of the final result, as shown in the observed and expected $-2 \ln \Delta$ distributions as a function of κ_λ in Figure 6.9. The result for κ_λ at 95% CL is almost identical when using the upper limit cross-section approach; the corresponding exclusion range is found to be $[-4, 11]$ ($[-5, 11]$ expected). The upper limits at 95% CL




Figure 6.8: Upper limit at 95% CL on the VBF $pp \rightarrow HH$ cross-section as a function of κ_{2V} . The solid and dashed black lines indicate respectively the observed and expected combined limits. The cyan and yellow bands represent the 1σ and 2σ intervals of the expected limit. The blue dash-dotted and the green dotted curves represent the expected upper limits in the boosted-only and resolved-only analyses, respectively. The star marker indicates the SM prediction. Both κ_λ and κ_V are fixed to their SM values [75].




Figure 6.9: Observed (a) and expected (b) distributions of the likelihood $-2 \ln \Lambda$ as a function of κ_λ in the boosted, resolved, and combined analyses. The dashed horizontal lines indicate the thresholds above which a hypothesis is excluded at 68% and 95% CL. Both κ_{2V} and κ_V are fixed to their SM values [75].

on the cross-section $\sigma_{\text{ggF+VBF},HH}$ as a function of κ_λ in the three channels are presented in Figure 6.10.

Finally, the increased sensitivity to κ_λ variations from the inclusion of the resolved analysis makes possible to study the more general scenario where both κ_{2V} and κ_λ coupling strengths are varied, while κ_V and all the other Higgs-boson couplings are kept fixed to their SM values. The observed and expected two-dimensional likelihood contours at 68% and 95% CL in the $\kappa_{2V} - \kappa_\lambda$ plane in the boosted, resolved, and combined analyses are presented in Figure 6.11. The κ_{2V} sensitivity is mainly driven by the boosted analysis, while the κ_λ sensitivity is




Figure 6.10: Upper limits at 95% CL on the ggF+VBF $pp \rightarrow HH$ cross-section as a function of κ_λ . The solid and dashed black lines indicate respectively the observed and expected combined limits. The cyan and yellow bands represent the 1σ and 2σ intervals of the expected limit. The blue dash-dotted and the green dotted curves represent the expected upper limits in the boosted-only and resolved-only analyses, respectively. The star marker indicates the SM prediction. Both κ_{2V} and κ_V are fixed to their SM values [75].




Figure 6.11: Observed (a) and expected (b) likelihood contours at 68% and 95% CL in the $\kappa_{2V} - \kappa_\lambda$ plane in the boosted (blue curves), resolved (red curves), and combined (black curves) analyses. The solid (dashed) curves indicate the limits at 95% (68%) CL. The cross-marker indicates the observed best-fit values of κ_{2V} and κ_λ , while the star marker indicates the SM prediction. All other Higgs-boson couplings are fixed to their SM values. The results for $\kappa_\lambda > 15$ are not plotted for better readability [75].

dominated by the resolved analysis. The observed best-fit values of κ_{2V} and κ_λ are compatible with the SM prediction within 68% CL.

6.4 Resonant interpretation

In the resonant analysis, the results for each mass hypothesis m_X are obtained using as fit observable the corresponding BDT response distributions from the pBDT model with the parametric mass input set to m_X . The results presented below are based on the model-independent narrow-width assumption, where the heavy spin-0 particle width Γ_X lies well below the detector resolution. An alternative broad-width interpretation based on the composite Higgs model is presented at the end of this section. In this resonant search, the non-resonant SM HH samples should be interpreted as extra background components; they are however not included in the analysis in view of their negligible yields compared to the main, data-driven background expectation (recall Table 6.1).

The post-fit distributions for the mass point hypotheses 1.0, 2.0, and 5.0 TeV are shown in Figure 6.12. No clear qualitative deviation of the observed data from the expected background has been detected. For masses above 1.6 TeV, it is necessary to enforce the positive signal strength constraint $\mu_{\text{VBF},X \rightarrow HH} > 0$ in order to ensure fit stability, as will be discussed later, otherwise the fit favours unphysical large negative values of $\mu_{\text{VBF},X \rightarrow HH}$ until the probability distribution functions become negative. The fitted values $\hat{\mu}_{\text{VBF},X \rightarrow HH}$ for each mass point and their upper errors divided into statistical and systematic sources are shown in Table 6.2. The lower errors are not shown because of the aforementioned positive constraint on the signal strength. In all cases the statistical error dominates. In Figure 6.13, the pulls of the nuisance parameters and their impact on the signal strength are shown for the mass hypotheses $m_X = 1.0$ TeV and $m_X = 1.5$ TeV. In the $m_X = 1.0$ TeV scenario, no nuisance parameters with large constraints or pulls have been found. The dominant systematic uncertainties are coming from the data-driven background modelling and from the double b -tagger. In the $m_X = 1.5$ TeV scenario, the impacts of the systematic uncertainties on the signal strength are essentially one-sided and have almost no effect due to $\hat{\mu}_{\text{VBF},X \rightarrow HH} \approx 0$.

The main reason of the fit instabilities for $m_X \geq 1.6$ TeV is the already mentioned lack of observed data in the last pBDT distribution bin for these masses. The stability issue can be better understood when analysing the likelihood $-2 \ln \Lambda$ curves as a function of the parameter of interest while fixing the nuisance parameters to their best-fit values, as shown in Figure 6.14. While for $m_X = 1.0$ TeV both the observed and expected likelihood curves exhibit a parabola-like shape with a well-defined minimum, there is no local minimum for all $m_X \geq 1.6$ TeV hypotheses.

Upper limits at 95% CL on the VBF resonant production of Higgs-boson pairs cross-section $\sigma_{\text{VBF}}(X \rightarrow HH)$ are set using the CL_s based on the \tilde{q}_μ test statistic, as described in Section 6.2. In view of the fit instabilities due to low statistics, it is necessary to verify if the upper limits derived using analytic approximations are comparable to the results obtained from the generation of pseudo-experiments (toys). In Table 6.3, a comparison between the expected and observed limits on the cross-section when using the asymptotic formula and toys is presented. For all mass points, the differences in the upper limits are below 10%

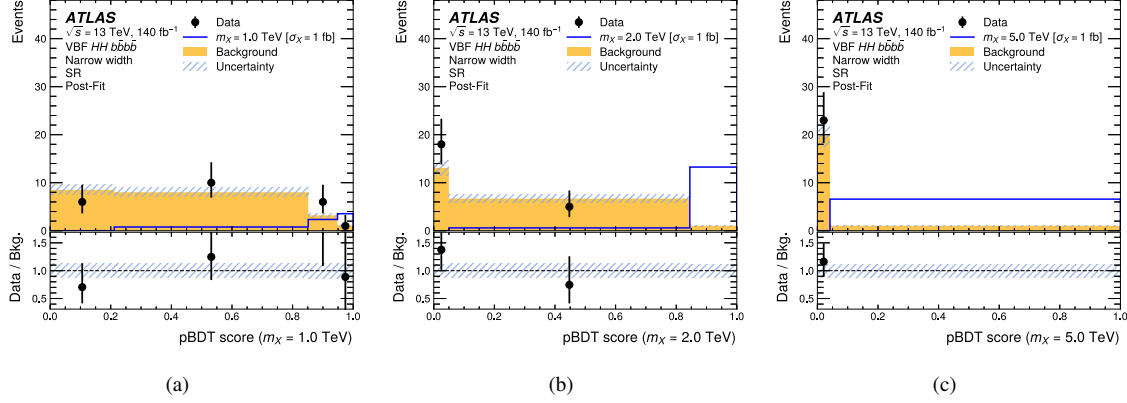


Figure 6.12: Post-fit pBDT distributions of observed data, expected background, and signal evaluated for three representative values of m_X . All sources of uncertainties have been included. The narrow width scenario is considered and the signal cross-section is normalised to 1 fb. The distributions observed in data agree well with the total background expectation [75].

Table 6.2: Fitted signal strength for each mass hypothesis in the narrow-width assumption. The total upper error of the fit parameter is divided into statistical and systematic components. The constraint $\mu_{\text{VBF}, X \rightarrow HH} > 0$ has been applied in the fit. The signal cross-section was normalised to 1 fb.

m_X (TeV)	$\hat{\mu}_{\text{VBF}, X \rightarrow HH}$	Upper error
1.0	0.231	$0.480_{\text{stat}} + 0.317_{\text{syst}} = 0.575_{\text{tot}}$
1.1	0.076	$0.284_{\text{stat}} + 0.168_{\text{syst}} = 0.330_{\text{tot}}$
1.2	0.072	$0.176_{\text{stat}} + 0.104_{\text{syst}} = 0.204_{\text{tot}}$
1.3	0.001	$0.130_{\text{stat}} + 0.048_{\text{syst}} = 0.139_{\text{tot}}$
1.4	0.001	$0.115_{\text{stat}} + 0.038_{\text{syst}} = 0.121_{\text{tot}}$
1.5	0.000	$0.112_{\text{stat}} + 0.048_{\text{syst}} = 0.122_{\text{tot}}$
1.6	0.000	$0.043_{\text{stat}} + 0.015_{\text{syst}} = 0.045_{\text{tot}}$
1.8	0.000	$0.039_{\text{stat}} + 0.015_{\text{syst}} = 0.041_{\text{tot}}$
2.0	0.000	$0.037_{\text{stat}} + 0.015_{\text{syst}} = 0.040_{\text{tot}}$
2.5	0.000	$0.038_{\text{stat}} + 0.017_{\text{syst}} = 0.042_{\text{tot}}$
3.0	0.000	$0.042_{\text{stat}} + 0.021_{\text{syst}} = 0.047_{\text{tot}}$
4.0	0.000	$0.054_{\text{stat}} + 0.029_{\text{syst}} = 0.061_{\text{tot}}$
5.0	0.000	$0.076_{\text{stat}} + 0.042_{\text{syst}} = 0.087_{\text{tot}}$

and therefore the asymptotic formula is used for the final results.³ The upper limits on $\sigma_{\text{VBF}}(X \rightarrow HH)$ at 95% CL as a function of m_X in the narrow-width scenario are presented in Figure 6.15. The difference between the observed and expected limits becomes relatively

³ In LHC analyses the typical threshold to allow the use of the asymptotic formula for the upper limits is 20%. Having the final results based on the asymptotic formula facilitates potential future combinations with other analyses.




Figure 6.13: Pulls (lower x-axis) of the nuisance parameters that parametrise the systematic uncertainties or the uncorrelated statistical uncertainties at each distribution bin (prefix `gamma_`) and their impact on the signal strength (upper x-axis) for (a) $m_X = 1.0$ TeV and (b) $m_X = 1.5$ TeV. No fitted nuisance parameter related to a systematic uncertainty has either a large pull or is over-constrained. Only the top 20 most impactful uncertainties are included in the plots. The names of the nuisance parameters are described in Appendix B.

constant for $m_X \geq 1.6$ TeV, again because of the deficit of observed data in the most sensitive pBDT bin. Since the narrow-width assumption is model-independent, no theory curve is depicted in the upper limit plot.

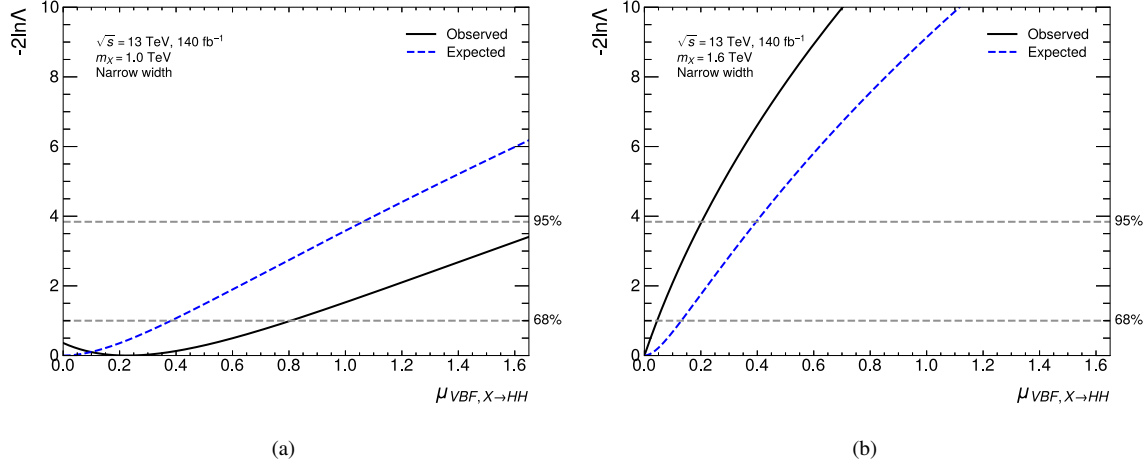


Figure 6.14: Observed and expected likelihood $-2 \ln \Lambda$ scan as a function of the signal strength in the scenarios (a) $m_X = 1.0$ TeV and (b) $m_X = 1.6$ TeV under the narrow-width assumption. Only for low masses the likelihood scan exhibits a parabola-like shape with well-defined minimum.

Table 6.3: Upper limit of $\sigma_{\text{VBF}}(X \rightarrow HH)$ [fb] at 95% CL calculated using the asymptotic approximation and toys for each mass hypothesis. A total of 100 000 pseudo-experiments were generated for each m_X point. The narrow-width samples are used here.

m_X (TeV)	Asymp. (exp)	Toys (exp)	Difference (exp,%)	Asymp. (obs)	Toys (obs)	Difference (obs,%)
1.0	3.148	3.235	2.763	4.634	4.902	5.784
1.1	2.01	2.051	2.055	2.543	2.7	6.142
1.2	1.501	1.535	2.271	1.715	1.841	7.342
1.3	1.395	1.406	0.781	1.279	1.339	4.702
1.4	1.29	1.298	0.573	1.165	1.237	6.194
1.5	1.241	1.28	3.155	1.152	1.229	6.691
1.6	1.173	1.171	-0.167	0.752	0.746	-0.831
1.8	1.158	1.154	-0.267	0.731	0.719	-1.583
2.0	1.161	1.144	-1.499	0.729	0.716	-1.704
3.0	1.49	1.424	-4.461	0.922	0.896	-2.826
4.0	2.079	1.984	-4.551	1.267	1.251	-1.258
5.0	3.008	2.825	-6.063	1.863	1.77	-4.975




Figure 6.15: Upper limit at 95% CL on the VBF $pp \rightarrow X \rightarrow HH$ cross-section as a function of m_X in the narrow-width scenario. The solid and dashed lines indicate, respectively, the observed and expected limits. The cyan and yellow bands represent the 1σ and 2σ intervals of the expected limit. A linear interpolation was used to derive the upper limits for intermediate mass values based on the values obtained for available signal samples (black dots) [75].




Figure 6.16: Upper limit at 95% CL on the VBF $pp \rightarrow X \rightarrow HH$ cross-section as a function of m_X in the broad-width scenario using the Composite Higgs model. The solid and dashed lines indicate, respectively, the observed and expected limits. The cyan and yellow bands represent the 1σ and 2σ intervals of the expected limit. A linear interpolation was used to derive the upper limits for intermediate mass values based on the values obtained for available signal samples (black dots). The theory predicted cross-section is represented by the red curve [75].

Broad-width interpretation

For the alternative broad-width interpretation the composite Higgs model is used and the width Γ_X of the heavy spin-0 particle is set to 20% of m_X , as discussed in Section 3.3.3. In this interpretation, mass values between 1.2 and 2.0 TeV are considered. Similarly to what has been seen in the narrow-width scenario, in the broad-width assumption the upper limits based on the asymptotic formula agree well with the ones based on pseudo-experiments and hence are safe to use for the final results. The upper limits on the cross-section $\sigma_{\text{VBF}}(X \rightarrow HH)$ at 95% CL as a function of m_X in the broad-width scenario are presented in Figure 6.16. Since the broad-width interpretation is model-dependent, a comparison between the observed limits and the theory prediction (red-curve) provides meaningful information. The observed limits on the cross-section are considerably larger than the cross-section predicted within the chosen benchmark model. Hence, there is not enough sensitivity to exclude mass regions within the chosen composite Higgs model benchmark across the tested mass range. It is also important to point out that the upper limits in the broad-width interpretation are slightly weaker compared to the ones in the narrow-width assumption. The main reason is the degradation of the signal kinematic distributions, which make them more similar to the background distributions.

6.5 Discussion

The search targeting VBF HH production in the Lorentz-boosted kinematic regime with four bottom-quarks in the final state provide constraints on κ_{2V} and on the two-dimensional $\kappa_{2V} \times \kappa_\lambda$ plane that show a huge improvement compared to a previous ATLAS search that targets VBF production of $HH \rightarrow b\bar{b}b\bar{b}$ in the resolved topology. Compared to the resolved-only analysis, the κ_{2V} limits in the combined boosted+resolved channel are 56% more stringent, constraining κ_{2V} to $[0.55, 1.49]$ ($[0.37, 1.77]$ expected) at 95% CL. It is now also possible in an ATLAS search to exclude the $\kappa_{2V} = 0$ scenario with an observed (expected) significance of 3.8σ (3.3σ).

An analogous search for Higgs-boson pairs in the boosted regime conducted by the CMS collaboration [151], however, yields more stringent constraints. The CMS VBF search of boosted Higgs-boson pairs decaying into $b\bar{b}b\bar{b}$ constrains κ_{2V} in the interval $[0.62, 1.41]$ ($[0.66, 1.37]$ expected) at 95% CL and additionally excludes $\kappa_{2V} = 0$ with an impressive 6.3σ of observed significance. The larger sensitivity of the CMS analysis is believed to be mostly due to a more complex Graph Neural Network (GNN) based algorithm [152] to select large- R jets resulting from $H \rightarrow b\bar{b}$ decays with great signal purity. As a comparison, for the 60% signal efficiency WP, the CMS tagger has a rate of 0.3% efficiency in keeping multijet background events; for the ATLAS tagger used in the analysis presented in this thesis, this number is 0.92% — about three times larger.

The results of the non-resonant analysis presented in this thesis were also recently combined [29] with other ATLAS HH searches targeting the $b\bar{b}\tau^+\tau^-$ [96], $b\bar{b}\gamma\gamma$ [97], multilepton [98], and $b\bar{b}\ell\ell + E_T^{\text{miss}}$ [99] final states. All these analyses used the full Run-2 dataset collected with the ATLAS detector. In Figure 6.17, the likelihood scan for the individual analyses and their statistical combination as a function of κ_{2V} while keeping all other Higgs-boson couplings fixed to their SM value are summarised. The combined observed (expected) κ_{2V} interval is found to be $[0.57, 1.48]$ ($[0.41, 1.65]$) at 95% CL. It is striking how the $b\bar{b}b\bar{b}$ presented in this thesis drives almost the entire combined result for κ_{2V} .

The analysis presented in this thesis is also the first to include a resonant interpretation for masses above 1.0 TeV in a search for VBF HH production. No evidence of a signal has been found and upper limits on the total cross-section $\sigma_{\text{VBF}}(X \rightarrow HH)$ have been set. In the narrow-width scenario, the observed (expected) limits range from 4.6 fb (3.1 fb) for $m_X = 1.0$ TeV to 1.9 fb (3.0 fb) for $m_X = 5.0$ TeV, with the most stringent limit of 0.7 fb (1.2 fb) obtained for $m_X = 1.8$ TeV. In the broad-width hypothesis, the observed (expected) limits decrease from 2.5 fb (2.1 fb) for $m_X = 1.2$ TeV to 0.8 fb (1.3 fb) for $m_X = 2.0$ TeV.

The sensitivity in both the non-resonant and resonant analyses is heavily dependent on the performance of the double b -tagging algorithm used to identify $H \rightarrow b\bar{b}$ decays. Taggers with better performance would allow the analysis to utilise WPs that provide much higher background rejection while still maintaining great signal efficiency. Moreover, it is also possible to reduce the impact of the tagger systematic uncertainties by employing better calibration strategies.

The more powerful double b -tagging algorithm used by the CMS collaboration is one of




Figure 6.17: Observed and expected distributions of the likelihood $-2 \ln \Lambda$ as a function of κ_{2V} for the combined (black) HH analyses. The expected results for the individual $b\bar{b}b\bar{b}$ (blue), $b\bar{b}\tau^+\tau^-$ (green), $b\bar{b}\gamma\gamma$ (purple), multilepton (cyan), and $b\bar{b}\ell\ell + E_T^{\text{miss}}$ (brown) analyses are also shown [29].

the key reasons on why the CMS search has obtained more stringent results than the ATLAS analysis presented in this thesis. Future ATLAS searches of Higgs-boson pairs decaying into $b\bar{b}b\bar{b}$ in the boosted topology will benefit by a new transformer-based double b -tagging algorithm being developed by the ATLAS Flavour Tagging Group [153]. In light of the statistical limitations in Higgs-boson pair production searches, the sensitivity of the current analysis should further increase with the inclusion of the Run-3 dataset collected by the ATLAS detector. The next chapter of this thesis is dedicated to an on-going ATLAS search for $HH \rightarrow b\bar{b}b\bar{b}$ in the boosted topology using the new tagger and combining the full Run-2 dataset with Run-3 data.

Prospects of VBF $HH \rightarrow b\bar{b}b\bar{b}$ searches in the HL-LHC

The High-Luminosity LHC (HL-LHC) should be operational in 2030 and is expected to last about 10 years, providing a total integrated luminosity of 3000 fb^{-1} at $\sqrt{s} = 14 \text{ TeV}$ center-of-mass energy. Many upgrades to the ATLAS detector are being prepared, such as the development of a new trigger system [154], the construction of a new Inner Tracker (also known as the ITk detector) [155], and a new High-Granularity Timing Detector [156] for precision timing identification and luminosity measurement.

The significant increase in luminosity is expected to enhance the sensitivity in HH searches. In order to estimate how much the analysis presented in this thesis would benefit from the

HL-LHC data, an extrapolation procedure is performed [157].

First, a common scale factor to account for the increase in luminosity is applied to signal and background samples. It is defined as the ratio $L_{\text{new}}/L_{\text{Run-2}}$ between the targeted integrated luminosity (e.g $L_{\text{new}} = 3000 \text{ fb}^{-1}$) and the Run-2 integrated luminosity $L_{\text{Run-2}} = 140 \text{ fb}^{-1}$. In sequence, to consider the impact of the increase of \sqrt{s} to 14 TeV in the cross-sections, a scale factor of 1.19 is applied to the VBF HH signals and a scale factor of 1.18 is applied to the data-driven background estimation [157]. Finally, in view of the impact of the double b -tagging algorithm in the analysis sensitivity, it is interesting to check scenarios with different tagger efficiencies. This is effectively taken into account assuming scale factor of $(x/0.6)^2$ applied exclusively to signal events, where 0.6 comes from the 60% tagger WP used in the analysis and x is a new (hypothetical) tagger efficiency under the same background rejection.

In Figure 6.18, the expected κ_{2V} interval at 95% CL is shown as a function of the double b -tagger efficiency for different integrated luminosity scenarios. The increase in luminosity from 1000 fb^{-1} to 2000 fb^{-1} has a greater impact in the κ_{2V} interval size compared to the change from 2000 fb^{-1} to 3000 fb^{-1} . This reflects the fact that, although the higher luminosity increases the number of produced HH events, it also increases the background yield. In this context, more performant double b -tagging algorithms become the key elements in $HH \rightarrow b\bar{b}b\bar{b}$ searches.




Figure 6.18: Extrapolation of the expected κ_{2V} interval at 95% CL in the boosted VBF $HH \rightarrow b\bar{b}b\bar{b}$ search for the HL-LHC as a function of the double b -tagger signal efficiency and for different integrated luminosities. The yellow vertical line represents the tagger WP in the analysis presented in this thesis. The green vertical line is the equivalent tagger WP under the same background rejection for the new double b -tagging algorithm discussed in Section 7.1 [157].

Boosted VBF $HH \rightarrow b\bar{b}b\bar{b}$ search with Run-3 data

The start of the Run-3 programme of the LHC in 2022 at the unprecedented $\sqrt{s} = 13.6$ TeV center-of-mass energy, increased instantaneous luminosity, and the vastly improved range of reconstruction and identification algorithms developed over the long LHC shutdown, along with hardware updates, brought a huge excitement to explore a vast range of physics processes. While the increased center-of-mass energy opens the possibility to explore phase-space regions never reached before, the increase in luminosity is of great importance to improve measurements explored in the earlier Runs of the LHC and to increase sensitivity in searches for extremely rare processes.

In the context of the VBF Higgs-boson pair production in the Lorentz-boosted regime, the large- R jets can now be reconstructed with improved resolution using a different set of input constituents based on particle flow and calorimeter clusters objects, and a new, much more performant double b -tagger algorithm is available. In addition to the increased statistics when combining the Run-2 and Run-3 datasets, the updated techniques are expected to enhance the sensitivity to κ_{2V} .

In fact, a new ATLAS $HH \rightarrow b\bar{b}b\bar{b}$ analysis has been launched using Run-3 data collected in 2022 and 2023 combined to the full Run-2 dataset. The total integrated luminosity amounts to 192 fb^{-1} . The new analysis targets both the non-resonant ggF and VBF production modes, exploring the resolved and boosted topologies in the ggF production mode and only the Lorentz-boosted regime in the VBF scenario. No resonant interpretation is included in this analysis. The aspects of this new on-going VBF $HH \rightarrow b\bar{b}b\bar{b}$ search, some preliminary expected limits, and comparisons to the main analysis presented in this thesis are shown in this chapter.

To avoid confusions among the main analysis presented in this thesis and the one going to be discussed in this chapter, from now on the main analysis will be referred to as “Run-2 analysis”, while the one presented in this chapter will be called “Run-3 analysis” or, in some

cases for convenience, simply “new analysis”. Furthermore, in view of the fact that the Run-3 analysis strategy is largely based on that of the Run-2 analysis and it is still an on-going work, only the most relevant information and differences with respect to the Run-2 analysis are shown here in detail.

7.1 The next-generation of the double b -tagger

A new double b -tagger algorithm has been developed by the ATLAS collaboration for Run-3 analyses to identify the Higgs-boson decay into a pair of bottom quarks reconstructed as a single large- R jet. The new algorithm, GN2X, benefits from a transformer-based network architecture [158] that allows feeding input of variable size into the model [153], like the track vectors within a jet.

The GN2X model takes as input the p_T , mass, and pseudorapidity of the large- R jet, as well as 20 variables related to each track associated to the jet, like impact parameters and number of hits in the Pixel Detector, with a maximum of 100 tracks per jet.¹ The large- R jet and individual track inputs are first combined into a jet-track sequence vector and passed to linear neural network layers to project the input representation into a larger embedding dimension. These representations are then fed into a transformer encoder corresponding to several blocks of multi-head self-attention layers.

In a nutshell, an input X to a head self-attention layer is first projected into three matrices $Q = XW^Q$, $K = XW^K$, and $V = XW^V$, where W^Q , W^K , and W^V are weight matrices determined during model training. For each head in the multi-head self-attention layer, the Attention function is computed as

$$\text{Attention}(Q, K, V) = \text{softmax} \left(\frac{QK^T}{\sqrt{d}} \right) V, \quad (7.1)$$

where d is the embedded input dimension and the softmax function ensures the attention scores will follow a valid probability density function. The multi-head self-attention layer combines the information from all its heads and outputs an array with the same dimensions of the input X . When passing X through the encoder blocks, all jet-track sequences influence their new representations in the self-attention layers; the output thus encodes an information concerning how relevant are the different jet-tracks and their variables, and how they are correlated.

Finally, the output representations of all jet-track sequences are combined to form a global representation of the large- R jet that is used to perform a classification. The GN2X model provides probabilities that the large- R jet arise from a $H \rightarrow b\bar{b}$ decay (p_{Hbb}), a $H \rightarrow c\bar{c}$ decay (p_{Hcc}), a fully hadronic top-quark decay (p_{top}), or as from a QCD multijets process

¹ In case a jet contains more than 100 tracks, the tracks are sorted by decreasing transverse impact parameter significance d_0/σ_{d_0} . This ensures the tracks associated to a displaced vertex are more likely to be picked.

7.1 The next-generation of the double b -tagger




Figure 7.1: Rejection rates for the (a) multijet and (b) hadronic top-decay backgrounds as a function of the large- R jet p_T when considering the 50% signal efficiency WP. The new GN2X model (blue) exhibits a much better performance compared to the previous tagger (pink) and the alternative approach tagging the sub-jets inside the large- R jet (green). The ratio pad is with respect to the previous D_{Xbb} tagger performance [153].

$(p_{\text{multijet}})^2$.² In analogy to the previous tagger used in the Run-2 analysis, all these probabilities are combined into a single discriminant as

$$D_{\text{GN2X}} = \ln \frac{p_{\text{Hbb}}}{f_{\text{Hcc}} \cdot p_{\text{Hcc}} + f_{\text{top}} \cdot p_{\text{top}} + (1 - f_{\text{Hcc}} - f_{\text{top}}) \cdot p_{\text{multijet}}}, \quad (7.2)$$

where the free parameters f_{Hcc} and f_{top} were assigned the values of $f_{\text{Hcc}} = 0.02$ and $f_{\text{top}} = 0.25$ to maximise the background rejection to a given signal efficiency [153]. Different tagger WPs are defined by placing different selection requirements on D_{GN2X} .

In Figure 7.1, the rejection rates for the multijets and hadronic top-decay backgrounds are shown as a function of the large- R jet p_T when using a 50% signal efficiency WP to detect $H \rightarrow b\bar{b}$ events. The new GN2X model outperforms significantly the previous D_{Xbb} tagger (described in Section 4.2.5) and the baseline approach where one tags directly the individual b -jets inside the large- R jet reconstructed with the variables-radius jet algorithm (also already described in Section 4.2.5).

At the time of the writing of this thesis, the calibration of the GN2X tagger has not yet been completed. Therefore, in the following scale factors of 1.0 are applied in simulated events.

² The previous tagger used in the Run-2 analysis provided no probability of the jet originating from a $H \rightarrow c\bar{c}$ decay.

7.2 UFO jets

In Run-3 analyses, large- R jets are commonly reconstructed using Unified Flow Objects (UFO) [159], a special type of particle-flow algorithm that combines the advantages of particle-flow objects (recall Section 4.2.1) and Track-CaloClusters (TCC) objects. TCC objects, like particle-flow objects, combine both tracking and calorimeter information, but differs from the particle-flow technique with regard to the energy subtraction algorithm, to the track to calorimeter-cluster matching algorithm, and is optimised for the high- p_T kinematic regime [160]. The TCC jets standalone, however, were found to yield poorer tagging performance compared to the baseline calorimeter jets used in the Run-2 analyses, and are generally more sensitive to pile-up effects.

Nevertheless, when combining particle-flow and TCC objects and techniques in a unified version for the jet reconstruction, an optimal overall tagging performance across the full kinematic p_T range is achieved. The UFO algorithm uses primarily the standard particle-flow algorithm but further apply a cluster splitting technique of the TCC algorithm if the particle-flow object was either identified as neutral or was found to be in a detector region with high track density.

7.3 Event selection and analysis regions

In the new VBF $HH \rightarrow b\bar{b}b\bar{b}$ search, the procedure to reconstruct the Higgs-boson candidates and the VBF jets is similar to that used in the Run-2 analysis. The Higgs-boson candidates are reconstructed as the two leading- p_T UFO large- R jets and the VBF jets are the two leading- p_T small- R jets with no angular overlap to any of the Higgs-boson candidates. Only events passing at least one of the lowest unscaled high- p_T jet triggers are processed; the online jet- p_T trigger thresholds are the same as to that in the Run-2 analysis.

As the kinematics of UFO and calorimeter jets are not identical, the offline jet- p_T value above which the trigger efficiency is found to be maximal is expected to change, especially in view that the relevant trigger chains are based on calorimeter jets. In fact, the trigger turn-on region moved from 450 GeV (recall Section 5.1.1) to around 520 GeV now that offline UFO jets are used. Hence, the leading Higgs-boson candidate must now pass the tighter requirement $p_T(H_1) > 520$ GeV, while the p_T requirement on the sub-leading Higgs boson is kept at 250 GeV. Both large- R jets must also have a mass larger than 50 GeV in order to lie within the region in which the jet calibration is valid.

Each Higgs-boson candidate is tagged using the new double b -tagger algorithm described in Section 7.1 and employing the 0.46% QCD multijets efficiency WP, which roughly corresponds to a 56% signal efficiency. Again, events in which two (one) Higgs-boson candidates pass the tagger WP are referred to as 2Pass (1Pass) events.

The choice of the tagger WP was obtained while simultaneously optimising the selection requirements on the VBF system. With a much more powerful tagger algorithm at hand, it is important to have a compromise between a high performance in detecting the signal

of interest, while still keeping a reasonable amount of background events to avoid running into fit instabilities in the statistical analysis due to very low background rates. In this Run-3 analysis, tighter WPs require loosening the requirements applied on the VBF system to slightly increase the rate of background events. Therefore, for each WP a scan over the values of the invariant mass of the VBF system and the pseudorapidity separation between the VBF jets, m_{jj} and $|\Delta\eta(j_1, j_2)|$, was performed and, for each scenario, the expected κ_{2V} interval at 95% CL was calculated, as well as the expected $\kappa_{2V} = 0$ exclusion significance, the overall number of expected background events, and the average systematic uncertainties arising from the background modelling (see Section 7.4). In Figure 7.2, the quantities are shown for the different choices of m_{jj} and $|\Delta\eta(j_1, j_2)|$ and for the nominal analysis WP. These numbers were obtained when already making use of a multivariate algorithm to improve the discrimination power between signal and background events, as described in Section 7.5. The preferred m_{jj} requirement is clearly seen for the quantities shown in Figures 7.2(a) and 7.2(b) (notice the colour scales are inverted). Therefore, the final requirement placed on the forward jets reads

- $m_{jj} > 400 \text{ GeV}$,
- $|\Delta\eta(j_1, j_2)| > 2.0$.

The analysis SR is kept the same as in the Run-2 analysis (see Equation 5.1), as it was found to bound well the region where most signal events lie in this Run-3 analysis.³ This can be seen more clearly in the mass planes of Figure 7.3 for signal events in the 2Pass region under the VBF $\kappa_{2V} = 1$ and $\kappa_{2V} = 0$ hypotheses, as well as data events in the 1Pass and 2Pass regions. Since this is an on-going analysis, the 2Pass SR is still blinded for data.

The VR and CR were slightly modified to a quadrant based version as it was found to present more robust estimations for the background modelling systematic uncertainties, with better consistency across different tagger WP choices. It additionally aligns well with the strategy being employed in the ggF analysis the team is also studying. The outer boundary of the contour is given in Equation 5.6, while the quadrants are split according to the mass centers of the SR formula. More precisely,

- VR: $(m_{H_1} > 124 \text{ GeV} \ \& \ m_{H_1} < 117 \text{ GeV}) \ || \ (m_{H_1} < 124 \text{ GeV} \ \& \ m_{H_1} > 117 \text{ GeV})$,
- CR: $(m_{H_1} > 124 \text{ GeV} \ \& \ m_{H_1} > 117 \text{ GeV}) \ || \ (m_{H_1} < 124 \text{ GeV} \ \& \ m_{H_1} < 117 \text{ GeV})$.

³ A re-optimisation of the SR contour was also considered, but the choice made for the Run-2 analysis was still found to rank amongst the best choices.




Figure 7.2: (a) Value of the expected κ_{2V} interval size at 95% CL, (b) the expected $\kappa_{2V} = 0$ exclusion significance, and (c) the overall background yield for different selection requirements on m_{jj} and $|\Delta\eta(j_1, j_2)|$. The nominal tagger WP corresponding to 0.46% of QCD multijets efficiency is used here. The background is estimated as described in Section 7.4.

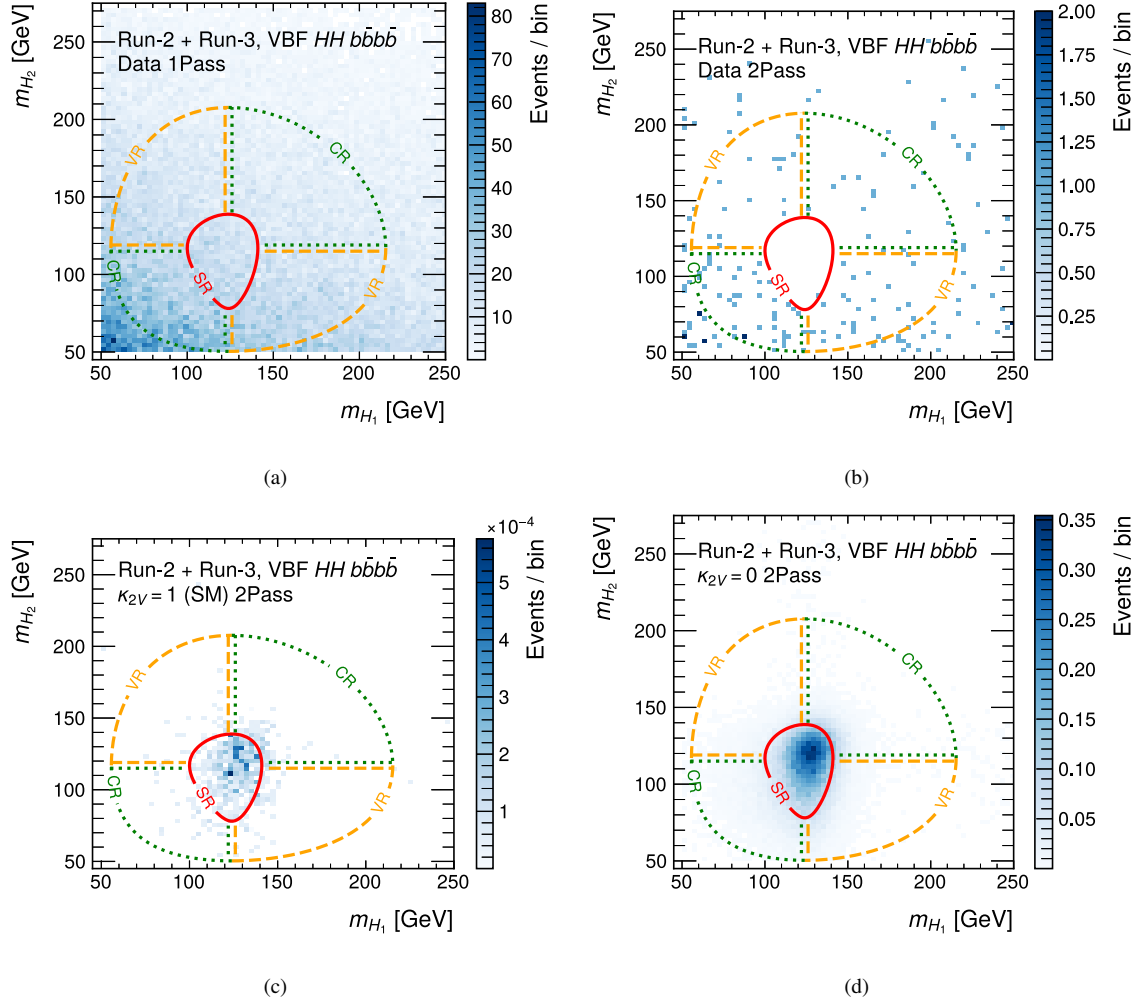


Figure 7.3: Reconstructed m_{H_1} - m_{H_2} plane for the (a) the 1Pass and (b) the 2Pass regions in data. The same mass plane for the 2Pass region for two representative signal hypotheses: (c) SM VBF $\kappa_{2V} = 1$ and (d) BSM VBF $\kappa_{2V} = 0$. The SR is described by the continuous red line. The VR and CR are split in quadrants and delimited by the continuous dashed yellow and dotted green lines, respectively. 2Pass data events are still blinded inside the SR.




Figure 7.4: Distributions of the DNN output score of data in the VR 2Pass and VR 1Pass regions. The shape systematic uncertainties are estimated from the ratio pads.

7.4 Background estimation

The background modelling in the SR 2Pass region is performed exactly as in the Run-2 analysis, using a single normalisation factor w to re-scale 1Pass events to the 2Pass region. The machine-learning based reweighting described in Appendix A was also implemented but again found to provide negligible impact. The value of w was derived in the analysis CR and found to be

$$w = 0.00334 \pm 0.00044, \quad (7.3)$$

where the uncertainty is of systematic nature and was estimated comparing with the result obtained in the VR.

The overall shape systematic uncertainty associated to the background modelling is derived by comparing the VR 2Pass and VR 1Pass data distributions (unity normalised) in the variable used as observable for the statistical inference, as shown in Figure 7.4. In this Run-3 analysis, a DNN is used instead of a BDT to provide further discrimination between signal and background events, as discussed in the next section.

7.5 Multivariate analysis

A dense neural network (DNN) is used to increase the discrimination power between signal and background events. The BDT approach used in the Run-2 analysis was also considered but it was found to provide slightly poorer κ_{2V} intervals in the statistical analysis. Hence, the DNN is the preferred choice in this Run-3 analysis.

Table 7.1: Kinematic variables used as input features for the DNN model.

Physics objects	DNN input variables
Higgs-boson Candidate ($H_i, i = 1, 2$)	$\log(p_T^{H_i}), \eta_{H_i}$
Di-Higgs System (HH)	$\log(p_T^{HH}), \eta_{HH}, \log(m_{HH}), \Delta\eta(H_1, H_2) , \Delta\Phi(H_1, H_2) $
VBF Jets ($j_i, i = 1, 2$)	$\log(p_T^{j_i}), \eta_{j_i}, \log(E_{j_i})$
VBF system ($j j$)	$\log(m_{jj})$

In a nutshell, the DNN consists of a one-dimensional input layer, some hidden layers, and an output layer. Each *node* in the input layer represents an input feature. The set of inputs is then passed through the multiple hidden layers, each containing a user-specified number of nodes that apply a weighted sum of the inputs followed by an activation function to introduce some non-linearity. These weights are adjusted during model training using the backpropagation technique [161], with the goal to minimise a loss function that quantifies the difference between the model’s predictions and the true values. The output layer produces the predictions, with the number of nodes depending on the task.

In this Run-3 analysis, the kinematic variables used as input features are all available in Table 7.1. The log function is applied to variables with a large range of values for better model performance. The model architecture was optimised as three hidden layers containing 50 nodes each. The swish activation function⁴ is applied to each node. The output layer consists of one single node using the sigmoid activation function, ensuring the model output can be interpreted as a probability quantifying how signal-like an event is. The training was performed using the binary cross-entropy loss function and the Adam optimiser for the backpropagation. The learning rate was kept fixed to 0.003.

Similarly to what has been seen in the Run-2 analysis, the kinematic distributions for signals with $\kappa_{2V} \neq 1$ are all very similar. This makes it possible to use just one single proxy signal sample when training the DNN without risking a worse performance of the network in detecting other signal hypotheses. Based on the latest κ_{2V} constraints derived in the ATLAS and CMS Run-2 searches, in this Run-3 analysis the VBF $\kappa_{2V} = 0.5$ hypothesis is chosen as the proxy signal sample.⁵ Again, the background used during training is the data-driven background model added to the SM VBF and SM ggF samples in order to enhance model performance in detecting BSM signals.

To ensure that any potential over-fit bias is not propagated into the analysis results, all events were first classified according to their “event number” parity. Two sets of models were trained. Models trained using events with even (odd) event number are used to make inferences on events with odd (even) event number. Therefore, each trained DNN model

⁴ The swish activation function is a smooth, non-monotonic function that combines elements of the well-known ReLU and sigmoid activation functions. It is defined as $f(x) = x \cdot \text{sigmoid}(x)$.

⁵ The same studies were also performed using $\kappa_{2V} = 0$ as the proxy sample. The differences were found to be negligible.




Figure 7.5: ROC curves and AUC evaluated for several samples and the background using the DNN classifier. The ratio pad compares the performances of the samples with respect to the proxy VBF $\kappa_{2V} = 0.5$ sample used in the model training. The dashed gray line indicates the hypothetical ROC curve if the model made random predictions only.

makes predictions only on events on which it was not trained.

The performance of the DNN in detecting signal events can be visualised in the ROC curve of Figure 7.5 for several non-SM samples, as well as for the SM VBF and SM ggF samples. The overall accuracy is higher compared to what has been seen in the Run-2 analysis with the BDT. However, this should not be interpreted as the DNN outperforming the previous BDT. Rather, this is a reflection of the already better signal over background separation in the Run-3 analysis when applying the baseline event selection, mostly thanks to the more performant double b -tagger algorithm.

Again, the model output distributions are used as observables in the statistical inference. The optimal binning was derived using the Transformation D method of the ATLAS fitting package `TRExFitter` [140] and fixed as $[0, 0.12, 0.31, 0.584, 0.856, 1.0]$. The DNN output distributions for several samples and the data-driven background modelling are presented in Figure 7.6, showing great separation power between background and the signal hypotheses.

7.6 Preliminary expected results

At the time of the writing of this thesis, data events in the analysis SR are still blinded. Therefore, it is only possible to evaluate the expected constraints on κ_{2V} while assuming the observed data is represented by the data-driven background model. At this stage, only a few systematic uncertainties are considered, as discussed in the following.

When including only the background modelling systematic uncertainties, the new analysis combining the full Run-2 dataset and partial Run-3 data expects to constrain $\kappa_{2V} \in [0.61, 1.40]$ at 95% CL using the likelihood approach and expects to exclude the $\kappa_{2V} < 0$ hypothesis with




Figure 7.6: Distributions of the DNN output score of several samples and the data-driven background model using the optimal binning. The SM samples were scaled for visualisation purposes. The uncertainties are of statistical nature only.

6.55σ of significance. In comparison to the main Run-2 analysis presented in this thesis but using the same setup for the uncertainties included, the expected improvement on the size of the κ_{2V} interval is of about 33% and about 56% on the $\kappa_{2V} < 0$ exclusion significance. If instead using only the full Run-2 dataset but with the tools and techniques of the new analysis, the improvements would be of 25% and 33%, respectively. This shows that the new tools for object reconstruction and identification, mainly the better GN2X double b -tagger, are the leading components to improve the sensitivity in the $\text{VBF } HH \rightarrow b\bar{b}b\bar{b}$ search — even more important than the increase of integrated luminosity with the addition of partial Run-3 data. The likelihood scan for these different scenarios as a function of κ_{2V} can be seen in Figure 7.7.

The Run-2 analysis was heavily impacted by the inclusion of the systematic uncertainties related to the double b -tagger. Since the calibration of the new GN2X tagger is still on-going work, it is imperative to estimate how these uncertainties might affect the final results in this analysis. Hence, for each bin in jet p_T used in the calibration,⁶ a dummy systematic uncertainty of 30% on the scale-factors is applied — a number slightly smaller compared to the average of 40% seen in the calibration of the previous double b -tagger. The central values of the scale-factors, however, are for now kept fixed at 1.0. These systematic uncertainties are included in the analysis as described in Section 5.4.3. The likelihood scan as a function of κ_{2V} using the full Run-2 dataset combined with the partial Run-3 data, with and without the dummy systematic uncertainties of the GN2X double b -tagger, is shown in Figure 7.8. The expected κ_{2V} interval is found to be $[0.55, 1.47]$ at 95% CL and the $\kappa_{2V} < 0$ hypothesis is expected to be excluded with 3.98σ of significance, diminishing significantly the analysis sensitivity.

⁶ Here, the bins in jet p_T for which the calibration is obtained is $[250, 520, 600, 1000]$ GeV, which very likely should be the setup for the final calibration.




Figure 7.7: Expected likelihood $-2 \ln \Lambda$ distribution as a function of κ_{2V} while fixing the other Higgs-boson couplings to their SM prediction. The curves in red and cyan are from the new analysis based on the GN2X tagger, but with different choices for the dataset. The curve in black represents the main Run-2 analysis of this thesis based on the previous tagger, but when including only statistical and the background modelling systematic uncertainties.




Figure 7.8: Expected likelihood $-2 \ln \Lambda$ distribution in the new analysis as a function of κ_{2V} while fixing the other Higgs-boson couplings to their SM prediction. The red curve corresponds to a scenario where only statistical errors and systematic uncertainties from the background modelling are included, while the black curve is obtained when additionally including the impact from the dummy GN2X double b -tagger uncertainties.

Transformers for triple-Higgs searches

In addition to the trilinear Higgs-boson coupling κ_λ and the coupling between two Higgs bosons and two vector bosons κ_{2V} discussed in the previous chapters, the Higgs potential (see Equation 2.36) in the SM also induces a direct interaction between four Higgs bosons governed by the quartic coupling, here identified in terms of its coupling modifier as κ_4 . This interaction is primarily accessible via the extremely rare process of triple Higgs-boson production, which is also sensitive to the κ_λ coupling.

The HHH process has been targeted directly for the very first time in a recent ATLAS analysis [30] using the full Run-2 dataset. The search targeted final states with six b -jets, corresponding to the triple-Higgs decay mode with the largest branching ratio. Other channels, although with lower branching ratios, would benefit from the cleaner backgrounds [162]. But because of the severe statistical limitation, searches targeting the other final states would rely even more on data from the Run-3 and future Runs of the LHC.

In general, any HHH search faces the difficult task of isolating a small signal from a huge background. Hence, in this chapter a proof-of-principle study using a transformer-based network architecture to enhance the discrimination power between signal and background events in HHH searches is presented. Due to the sophisticated transformer architecture that makes it possible to use low-level variables like jet four-vectors as input features, it is expected that the transformer performance should supersede simpler network architectures like DNNs.

This work is done upon the published ATLAS $HHH \rightarrow b\bar{b}b\bar{b}b\bar{b}$ search¹ with the goal to implement the techniques studied here in a novel ATLAS $HHH \rightarrow b\bar{b}b\bar{b}\tau^+\tau^-$ analysis that has recently started.

A brief overview of the HHH production in the SM is given in Section 8.1 and the relevant details of the ATLAS $HHH \rightarrow b\bar{b}b\bar{b}b\bar{b}$ search are shown in Section 8.2. The studies involving transformers in the context of HHH searches are presented in Section 8.3.

¹ The author of this thesis thanks the $HHH \rightarrow b\bar{b}b\bar{b}b\bar{b}$ analysis team for providing all the relevant files containing data and simulated reconstructed events.




Figure 8.1: Leading order Feynman diagrams representing the ggF process to produce three Higgs bosons in the SM. The κ -modifiers at each Higgs vertex are included in the diagrams to mark the couplings for the different vertices.

8.1 HHH production in the SM

The main mode of triple Higgs-boson production HHH in the SM is via quark-loop induced ggF , as depicted in the leading order Feynman diagrams of Figure 8.1. Diagram a) provides unique access to the quartic coupling between four Higgs bosons, κ_4 , while diagrams b) and c) depend on the trilinear Higgs-boson self-coupling κ_λ .

The total ggF cross-section in the SM at next-to-next-to-leading order at $\sqrt{s} = 13$ TeV has been calculated to be

$$\sigma_{ggF}(pp \rightarrow HHH) = 0.079^{+0.012}_{-0.013}(\text{scale} + m_t) \text{ fb}, \quad (8.1)$$

based on the finite top-quark mass approximation and a Higgs-boson mass of 125 GeV. This value was calculated by extrapolating the cross-sections at higher \sqrt{s} -values provided in Ref. [163]. Similarly to the HH production, the cross-section is, however, significantly enhanced for non-SM values of κ_λ and κ_4 due to the reduction of destructive interference effects among the leading order diagrams.

8.2 The ATLAS $HHH \rightarrow b\bar{b}b\bar{b}b\bar{b}$ search

The baseline strategy employed in the ATLAS $HHH \rightarrow b\bar{b}b\bar{b}b\bar{b}$ search [30] to select the events of interest is the following: all jets are reconstructed as small- R jets and identified whether originating or not from b -quarks using the DL1r single b -tagging algorithm introduced in Section 4.2.5. Events with exactly six b -tagged jets ($6b$) are assigned to the analysis SR, while $4b$ and $5b$ events are assigned to the analysis CRs and are mainly used for the background modelling in the SR. The background is heavily dominated by QCD multijets

with some real and some mis-identified b -jets.

To reconstruct the Higgs-boson candidates, the six leading- p_T b -tagged jets are selected. If an event has fewer than six b -tagged jets, the extra non- b -tagged jets with highest p_T are picked. The jet pairings to reconstruct the three $H \rightarrow b\bar{b}$ decays are determined by minimising the sum

$$|m_{H_1} - 120 \text{ GeV}| + |m_{H_2} - 115 \text{ GeV}| + |m_{H_3} - 110 \text{ GeV}| \quad (8.2)$$

for all possible permutations of jet pairs while imposing the p_T -ordering $p_T(H_1) > p_T(H_2) > p_T(H_3)$. In the equation above, the mass centres are chosen based on the peaks of the m_{H_i} distributions in simulated signal events.

The $6b$ background is primarily estimated using $5b$ data events assuming the kinematic properties of the background processes in the $5b$ and $6b$ regions are similar. Some correction factors to account for modelling imperfections when transitioning from the $5b$ to the $6b$ region are derived comparing the $4b$ and $5b$ data distributions in the final discriminant. More precisely, the bin-by-bin number of background events $\text{bkg}_{6b}(i)$ in a generic distribution is modelled as

$$\text{bkg}_{6b}(i) = \text{data}_{5b}(i) \times \frac{R_{5b/4b}(i)}{R_{5b/4b}} \times R_{6b/5b}, \quad (8.3)$$

where $\text{data}_{5b}(i)$ is the number of $5b$ data events in bin i , $R_{5b/4b}$ stands for the ratio between the total $5b$ and $4b$ data yields, $R_{5b/4b}(i)$ is the same ratio but only considering events inside of bin i , and $R_{6b/5b}$ is a global normalising factor to correct $5b$ data to the expected $6b$ yield. To evaluate $R_{6b/5b}$, a DNN has first been trained using the SM signal sample and $5b$ data,² which works as the first approximation for the background expectation in terms of kinematics. The region in the DNN distribution where the presence of signal is found to be minimal was then used to access $6b$ data events so that $R_{6b/5b}$ could be determined without having the danger of including signal contamination into the background model. The same DNN was used later as a fit observable in the statistical data analysis.

In this analysis, the very first two-dimensional exclusion limit curves in the $\kappa_\lambda \times \kappa_4$ plane were calculated, as depicted in Figure 8.2. The gray dashed curve in the figure was calculated in Ref. [162] and delimits the region inside which the Lagrangian is unitary. The current sensitivity using only the Run-2 dataset collected with the ATLAS detector is still not enough to exclude regions at 95% CL in the $\kappa_\lambda \times \kappa_4$ plane that are in principle allowed by unitarity.

The DNN is one of the key elements of this ATLAS analysis, since it plays an important role for the background modelling and its performance in discriminating signal from background events directly influence the analysis sensitivity in constraining the Higgs-boson self-couplings. The baseline DNN architecture consists of three hidden layers with 24 nodes each. The output layer is a single node with the softmax function applied. A total of 10 reconstructed

² The original analysis actually trained many DNNs because several benchmark scenarios including the presence of BSM particles were also studied. Here for simplicity only the case for non-resonant HHH production with κ -variations is discussed.




Figure 8.2: Observed and expected likelihood contours at 68% and 95% CL in the κ_λ - κ_4 plane. The unitary bound delimited by the gray dashed curve was calculated in Ref. [162]. The red star indicates the SM prediction. The notation κ_3 was used to represent κ_λ in the plot [30].

high-level variables are used as input features and their definitions are provided in Table 8.1.

In the next section, it is discussed how much performance improvement is expected when using a transformer-based architecture.


8.3 The transformer-based architecture

In the new transformer-based architecture, the main set of inputs consists of jet sequences — each containing low-level jet properties like p_T , η , etc. No b -tagging information is used due to its high-correlation with the background modelling strategy. All the jet information is passed to a transformer encoder with the goal of finding correlations among the jets composing the event, while the high-level variables of Table 8.1 are passed to a feed-forward network formed by three linear layers with 24 nodes each. The transformer encoder consists of three blocks of a multi-head self-attention layer with four heads followed by a normalisation layer and a linear layer of 32 nodes (recall Section 7.1 for attention definition). In the end, the multidimensional outputs of the transformer encoder and the feed-forward network are combined and passed to a single-node layer with softmax activation function to perform the event classification. The optimised model was found by scanning several architecture parameters and comparing the model performances.

A performance comparison between the baseline DNN and the transformer to identify SM signal and background events can be visualised in the model output distributions of Figure 8.3. The presence of more signal (background) events in the right-most (left-most) bin shows that the transformer-based architecture slightly outperforms the baseline DNN model.

Table 8.1: High-level kinematic variables used as input features for the DNN and transformer-based classifiers.

Variable	Definition
ΔR_{H_1}	Angular separation between the reconstructed jets composing H_1 .
ΔR_{H_2}	Angular separation between the reconstructed jets composing H_2 .
ΔR_{H_3}	Angular separation between the reconstructed jets composing H_3 .
m_{H_1}	Mass of the reconstructed Higgs-boson H_1 .
m_H -radius	Euclidean distance between the reconstructed Higgs-boson masses and the centres of Equation 8.2 in the $(m_{H_1}, m_{H_2}, m_{H_3})$ volume.
$\text{RMS}(m_{jj})$	Root-mean-squared (RMS) of the invariant mass of all possible jet pairings that can reconstruct a Higgs-boson candidate.
$\text{RMS}(\Delta R_{jj})$	RMS of the angular distance of all possible jet pairings that can reconstruct a Higgs-boson candidate.
$\text{RMS}(\eta)$	RMS of the pseudorapidity of the reconstructed Higgs-boson candidates.
Transverse $6j$ -Sphericity	p_T -isotropy within the x - y plane of the six jets selected to reconstruct the Higgs-boson candidates.
$6j$ -Aplanarity	p_T -fraction of the six jets selected to reconstruct the Higgs-boson candidates lying outside the plane formed by the two p_T -leading jets.

Figure 8.3: Background (cyan) and SM HHH signal (red) model output distributions when using the baseline DNN (solid lines) and the transformer-based architectures (dashed lines).

To check how the new transformer model impacts the analysis sensitivity, in Figure 8.4 it is shown the expected $-2 \ln \Lambda$ scan as a function of κ_4 , while fixing $\kappa_\lambda = 1$, using the model distributions as fit observables for the statistical inference. No systematic uncertainties are included here. To model the signal distributions for arbitrary κ_4 , a reweighing technique using nine basis samples is used — similarly to what has been done in the $HH \rightarrow b\bar{b}b\bar{b}$ analysis.

The gains from the transformer-based network on the κ_4 limits are found to be small, having only about 4% of improvement on the size of the κ_4 interval at 95% CL when keeping




Figure 8.4: Expected likelihood $-2 \ln \Lambda$ as a function of κ_4 while fixing the other Higgs-boson couplings to their SM prediction. The red (cyan) curve is the result obtained using the transformer (baseline DNN) output as a fit observable. Only statistical uncertainties are included in the fit. The dashed gray line delimits the region where a hypothesis is excluded at 95% CL.

the other Higgs-boson couplings fixed to their SM prediction. This is primarily due to fact that the search is severely statistics limited, which mitigates the impact from a better signal over background discriminant when the background yield is significantly larger. Besides, the $b\bar{b}b\bar{b}b\bar{b}$ final state offers a priori only jet properties as relevant information to be passed to the transformer encoder.

In a more complex final state like $b\bar{b}b\bar{b}\tau^+\tau^-$, the τ -decays reconstruction requires accessing a more extensive set of physics objects, such as muons, electrons, and E_T^{miss} . Added to the standard jets, a much richer list of inputs would be available to be passed to the transformer encoder with the goal to find nuances hardly captured by a standard DNN. Therefore, it is expected that in the new ATLAS $HHH \rightarrow b\bar{b}b\bar{b}\tau^+\tau^-$ search the use of a transformer network architecture should be even more relevant.

Conclusions

In this thesis, a search for the production of a pair of Higgs bosons via VBF decaying into four bottom quarks was presented. Both non-resonant and resonant production modes were considered. The former are predicted in the SM as well as in BSM scenarios and provide unique access to the coupling between two vector bosons and a pair of Higgs bosons, κ_{2V} . The resonant HH production from the decay of a heavy scalar particle originating from the vector-boson scattering is a BSM-scenario-only. The main results presented in this thesis rely on the full Run-2 dataset with a total integrated luminosity of 140 fb^{-1} collected with the ATLAS detector in pp collisions at $\sqrt{s} = 13 \text{ TeV}$.

In view of the low expected sensitivity in observing the SM signal due to the small cross-section in combination with the high rate of QCD multijets background events in the SR, the non-resonant analysis focused on anomalous values of κ_{2V} , with the goal of setting exclusion limits as a function of κ_{2V} at 95% CL and excluding hypotheses with $\kappa_{2V} < 0$. The signal cross-section for $\kappa_{2V} \neq 1$ is significantly enhanced due to the reduction of destructive interference effects between the leading order Feynman diagrams of the process and shift the energy spectrum of the produced Higgs-boson pairs towards higher values. As a consequence of the latter, the pairs of bottom quarks originating from the Higgs-boson decay become more collimated and tend to be found in the central region of the detector. This is also the case in the resonant production of Higgs-boson pairs when the intermediate particle is sufficiently heavy.

The analysis strategy was therefore designed to target boosted Higgs-boson decay topologies, where the pair of bottom quarks from each Higgs-boson decay is reconstructed as a single large- R jet. To identify whether the large- R originated from a $H \rightarrow b\bar{b}$ decay or other background, a specialised machine-learning based double b -tagger algorithm was used. To account for the VBF signature of the HH production mode, the presence of two additional small- R jets with no overlap to any of the Higgs-bosons candidates was required. In order to enhance the overall analysis sensitivity, BDTs were used to discriminate between signal and background events. One single BDT was used in the non-resonant analysis to differentiate

BSM signals with $\kappa_{2V} \neq 1$ and background, while in the resonant analysis a parametric BDT trained with the true mass of the heavy scalar particle as an additional parameter was implemented. The BDT outputs were used as fit observables in the statistical analysis.

In the non-resonant search, the boosted analysis provides an observed (expected) κ_{2V} constrain in $[0.52, 1.52]$ ($[0.37, 1.71]$) at 95% CL. The $\kappa_{2V} < 0$ hypothesis is excluded with an observed (expected) significance of 3.4σ (2.9σ). To further increase the analysis sensitivity, the boosted search was combined with a previous search targeting the resolved decay topology, where each of the four b -quarks were reconstructed individually as a single small- R jet. The combination specifically improves the sensitivity in the $\kappa_{2V} \approx 1.0$ region, where the resolved reconstruction yields a higher signal efficiency due to the presence of Higgs bosons with smaller Lorentz boost. The combination provides an observed (expected) κ_{2V} constraint of $[0.55, 1.49]$ ($[0.37, 1.67]$) at 95% CL and increases the observed (expected) $\kappa_{2V} < 0$ exclusion significance to 3.8σ (3.3σ). All these results were obtained assuming all other Higgs-boson couplings were fixed to their SM expectation. Two-dimensional exclusion regions were also derived as a function of κ_{2V} and the trilinear Higgs-boson self-coupling κ_λ (see Figure 6.11). The VBF boosted search provides the stronger constraint on κ_{2V} compared to the resolved analysis, but it is much less sensitive to κ_λ . Hence, the κ_λ constraint comes almost exclusively from the resolved analysis, which also includes the ggF production mode with higher cross-section.

The search results were also interpreted in the context of a spin-0 resonance mediating the VBF Higgs-boson pair production. Both signal hypotheses with narrow- and broad-resonances widths were taken into account. In the model-independent, narrow-width approximation, mass values ranging from 1.0 TeV to 5.0 TeV were considered. In the model-dependent, broad-width scenario with $\Gamma_X/m_X = 0.2$ and using a composite Higgs benchmark model, the search was limited to mass values ranging between 1.2 TeV and 2.0 TeV. No deviations from the SM background were observed and upper limits on the resonant production cross-section were set at 95% CL. The results are given in Figures 6.15 and 6.16.

Even though the low data statistics in the signal region is a limiting factor for VBF $HH \rightarrow b\bar{b}b\bar{b}$ searches, the key factor limiting the sensitivity of the search has been found to be the efficiency of the double b -tagger algorithm in detecting signal events while rejecting background events at high-rates. The systematic uncertainties arising from the tagger scale-factors to be applied in simulated events also reduces significantly the analysis sensitivity. A much more powerful tagging algorithm based on transformer networks is now being employed in an on-going $HH \rightarrow b\bar{b}b\bar{b}$ search using both the full Run-2 dataset as well as Run-3 data, as shown in Chapter 7. The results presented in this thesis demonstrate that the tagger performance is the key element for future, more precise determinations of κ_{2V} and κ_λ .

Finally, in Chapter 8 a study into the use of transformers to improve the discrimination power between signal and background events in searches involving the production of three Higgs bosons was presented. Although the HHH production cross-section in the SM is very small, it provides unique access to the quartic coupling between four Higgs bosons. Additionally, many BSM models involving the presence of new heavy particles can be explored in such searches. The techniques discussed in this chapter can be used in a new

ATLAS $HHH \rightarrow b\bar{b}b\bar{b}\tau^+\tau^-$ search using the Run-2 and Run-3 datasets of the ATLAS detector that has recently started.

In conclusion, this thesis contains work on searches for both HH and HHH production in final states with at least four bottom quarks. The results form a strong basis for the upcoming and exciting set of searches for multi-Higgs production based on a novel set of analysis tools and using the new Run-3 data.

Appendix

Data-driven background modelling via reweighting

This appendix describes a machine-learning based method to derive the reweighting function introduced in Section 5.2 to model the expected background of the analysis presented in this thesis.

In order to learn the reweighting

$$w(\mathbf{x}) = \frac{p_{2\text{Pass}}(\mathbf{x})}{p_{1\text{Pass}}(\mathbf{x})} \quad (\text{A.1})$$

between the data 2Pass and 1Pass probability density distributions as a function of a set \mathbf{x} of kinematic variables, an interesting trick is to train a neural network using 2Pass and 1Pass data events with the goal to minimise the cost function [164]

$$\mathcal{L}[Q] = \mathbb{E}_{\mathbf{x} \sim p_{2\text{Pass}}} [e^{Q(\mathbf{x})/2}] + \mathbb{E}_{\mathbf{x} \sim p_{1\text{Pass}}} [e^{-Q(\mathbf{x})/2}], \quad (\text{A.2})$$

where $Q(\mathbf{x})$ is the one-dimensional neural network output based on the input \mathbf{x} . This cost function is minimised at $Q^*(\mathbf{x}) = \log w^*(\mathbf{x})$.

A neural network has been trained using the input variables displayed in Table A.1. The model architecture consists of four hidden layers with 10 nodes each. The first three layers use the ReLU activation function while the last one uses the sigmoid activation function as it was found to produce better training stabilisation. The output layer with only one single node has no activation function and it provides directly the logarithm of the reweighting function.

Because of the extreme low statistics available, the neural network training was performed in an extended region defined as the whole mass-plane excluding the signal region and the low-mass region $m_{H_1} + m_{H_2} < 130$ GeV enriched by multijet QCD background events. The training was performed 10 times independently and the final output is the average of all models output.

The impact from the reweighting is found to be negligible, as shown in Figure A.1 for a few kinematic distributions. Therefore, the best strategy is to perform the simple re-scaling between 1Pass and 2Pass data events described in greater detail in Section 5.2.

Table A.1: Kinematic variables used as input for the neural-network based reweighting.

Physics objects	Input variables
Higgs-boson Candidate ($H_i, i = 1, 2$)	$\log(p_T^{H_i}), \eta_{H_i}$
Di-Higgs System (HH)	$\log(p_T^{HH}), \eta_{HH}, \log(m_{HH})$
VBF Jets ($j_i, i = 1, 2$)	$\log(p_T^{j_i}), \eta_{j_i}, \log(E_{j_i})$

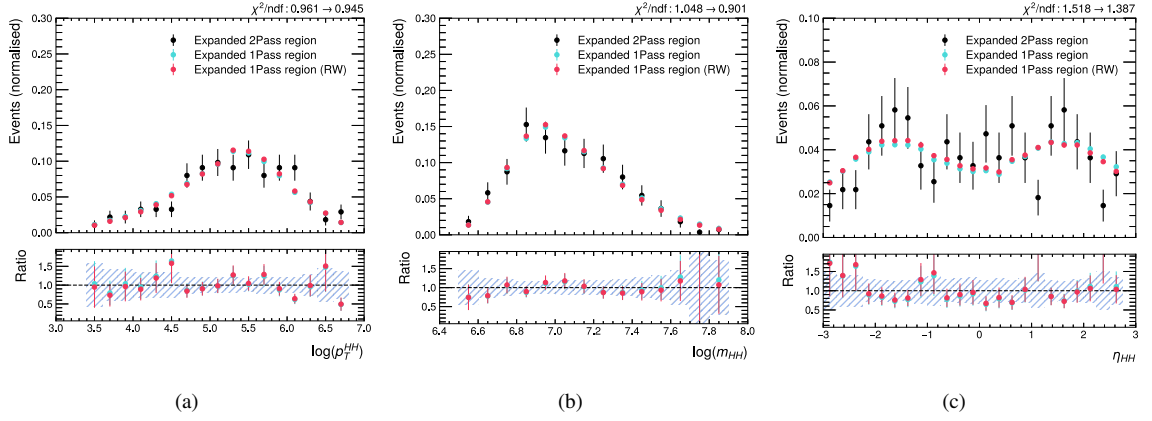


Figure A.1: Comparison of 2Pass, 1Pass, and reweighted 1Pass data distributions in an extended region of the mass-plane of the reconstructed Higgs bosons.

APPENDIX B

Naming conventions for the nuisance parameters

This appendix is dedicated to describe the naming conventions for the nuisance parameters used in Figures 6.5 and 6.13. Their definitions are shown in Table B.1. It only includes the relevant names appearing in the plots.

Table B.1: Name conventions for the nuisance parameters.

Name	Associated uncertainty
$\text{gamma_stat_SR_bin_}i$	Statistical uncertainty in the i -th bin of the distribution.
DD_shape	Shape uncertainties of the data-driven background model.
DD_norm	Normalisation uncertainty of the data-driven background model.
AltPS_*	Parton shower uncertainties
BR_bb	Uncertainty on the $H \rightarrow b\bar{b}$ branching ratio.
Theory_*	Scale variations (QCD) and PDF+ α_s (PDFaS) uncertainties.
Boosted_FTAG_Xbb_ptbini	Double b -tagger uncertainty associated to the i -th calibration bin.
Boosted_JET_*	Systematic uncertainty components associated to the large- R jet calibration.
Resolved_JET_*	Systematic uncertainty components associated to the small- R jet calibration.

Bibliography

- [1] S. L. Glashow, *Partial Symmetries of Weak Interactions*, **Nucl. Phys.** **22** (1961) 579 (cit. on pp. 1, 5).
- [2] S. Weinberg, *A Model of Leptons*, **Phys. Rev. Lett.** **19** (1967) 1264 (cit. on pp. 1, 5).
- [3] S. L. Glashow, J. Iliopoulos and L. Maiani, *Weak Interactions with Lepton-Hadron Symmetry*, **Phys. Rev.** **D2** (1970) 1285 (cit. on pp. 1, 5, 6).
- [4] R. P. Feynman, *Space - time approach to quantum electrodynamics*, **Phys. Rev.** **76** (1949) 769 (cit. on p. 1).
- [5] J. S. Schwinger, *Quantum electrodynamics. I A covariant formulation*, **Phys. Rev.** **74** (1948) 1439, ed. by K. A. Milton (cit. on p. 1).
- [6] S. Tomonaga, *On a relativistically invariant formulation of the quantum theory of wave fields*, **Prog. Theor. Phys.** **1** (1946) 27 (cit. on p. 1).
- [7] 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).
- [8] A. Salam, *Weak and Electromagnetic Interactions*, **Conf. Proc. C** **680519** (1968) 367 (cit. on pp. 1, 6).
- [9] D. J. Gross and F. Wilczek, *Ultraviolet Behavior of Nonabelian Gauge Theories*, **Phys. Rev. Lett.** **30** (1973) 1343 (cit. on pp. 1, 7).
- [10] H. D. Politzer, *Reliable Perturbative Results for Strong Interactions?*, **Phys. Rev. Lett.** **30** (1973) 1346 (cit. on pp. 1, 7).
- [11] G. Arnison et al., *Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540$ GeV*, **Phys. Lett. B** **122** (1983) 103 (cit. on p. 1).
- [12] M. Banner et al., *Observation of Single Isolated Electrons of High Transverse Momentum in Events with Missing Transverse Energy at the CERN anti-p p Collider*, **Phys. Lett. B** **122** (1983) 476 (cit. on p. 1).

- [13] F. Abe et al., *Observation of top quark production in $\bar{p}p$ collisions*, *Phys. Rev. Lett.* **74** (1995) 2626, arXiv: [hep-ex/9503002](#) (cit. on p. 1).
- [14] S. Abachi et al., *Observation of the top quark*, *Phys. Rev. Lett.* **74** (1995) 2632, arXiv: [hep-ex/9503003](#) (cit. on p. 1).
- [15] T. Aoyama, T. Kinoshita and M. Nio, *Theory of the Anomalous Magnetic Moment of the Electron*, *Atoms* **7** (2019) 28 (cit. on p. 1).
- [16] 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.* **B716** (2012) 1, arXiv: [1207.7214 \[hep-ex\]](#) (cit. on pp. 1, 14).
- [17] CMS Collaboration, *Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC*, *Phys. Lett.* **B716** (2012) 30, arXiv: [1207.7235 \[hep-ex\]](#) (cit. on pp. 1, 14).
- [18] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, *Phys. Rev. Lett.* **13** (1964) 508 (cit. on pp. 1, 11).
- [19] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, *Phys. Rev. Lett.* **13** (1964) 321 (cit. on pp. 1, 11).
- [20] G. S. Guralnik, C. R. Hagen and T. W. B. Kibble, *Global Conservation Laws and Massless Particles*, *Phys. Rev. Lett.* **13** (1964) 585 (cit. on pp. 1, 11).
- [21] M. C. Gonzalez-Garcia and M. Maltoni, *Phenomenology with Massive Neutrinos*, *Phys. Rept.* **460** (2008) 1, arXiv: [0704.1800 \[hep-ph\]](#) (cit. on p. 1).
- [22] P. A. R. Ade et al., *Planck 2015 results. XIII. Cosmological parameters*, *Astron. Astrophys.* **594** (2016) A13, arXiv: [1502.01589 \[astro-ph.CO\]](#) (cit. on pp. 1, 17).
- [23] G. Arcadi, A. Djouadi and M. Raidal, *Dark Matter through the Higgs portal*, *Phys. Rept.* **842** (2020) 1, arXiv: [1903.03616 \[hep-ph\]](#) (cit. on p. 2).
- [24] H. E. Haber and G. L. Kane, *The Search for Supersymmetry: Probing Physics Beyond the Standard Model*, *Phys. Rept.* **117** (1985) 75 (cit. on pp. 2, 18).
- [25] J. E. Kim, *Light Pseudoscalars, Particle Physics and Cosmology*, *Phys. Rept.* **150** (1987) 1 (cit. on pp. 2, 18).
- [26] G. C. Branco et al., *Theory and phenomenology of two-Higgs-doublet models*, *Phys. Rept.* **516** (2012) 1, arXiv: [1106.0034 \[hep-ph\]](#) (cit. on pp. 2, 18).
- [27] ATLAS Collaboration, *A detailed map of Higgs boson interactions by the ATLAS experiment ten years after the discovery*, *Nature* **607** (2022) 52, [Erratum: *Nature* 612, E24 (2022)], arXiv: [2207.00092 \[hep-ex\]](#) (cit. on pp. 2, 14, 15).

-
- [28] 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. 2, 14–16).
- [29] 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. 2, 3, 15, 16, 118, 119).
- [30] ATLAS Collaboration, *A search for triple Higgs boson production in the $6b$ final state using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, (2024), arXiv: 2411.02040 [hep-ex] (cit. on pp. 2, 3, 81, 133, 134, 136).
- [31] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, *JINST* **3** (2008) S08003 (cit. on pp. 2, 30, 34, 37–39, 41).
- [32] S. Navas et al., *Review of particle physics*, *Phys. Rev. D* **110** (2024) 030001 (cit. on pp. 6, 13, 28).
- [33] M. Gell-Mann, *Quarks*, *Acta Phys. Austriaca Suppl.* **9** (1972) 733 (cit. on p. 6).
- [34] S. Weinberg, *New approach to the renormalization group*, *Phys. Rev. D* **8** (1973) 3497 (cit. on p. 6).
- [35] A. Pich, “Quantum chromodynamics”, *High energy physics. Proceedings, 2nd European School, Sorrento, Italy, August 29-September 1, 1994. Vol. 1, 2*, 1995 157, arXiv: hep-ph/9505231 [hep-ph] (cit. on p. 6).
- [36] V. N. Gribov, *The Theory of quark confinement*, *Eur. Phys. J. C* **10** (1999) 91, ed. by J. Nyiri, arXiv: hep-ph/9902279 (cit. on p. 7).
- [37] C. S. Wu, E. Ambler, R. W. Hayward, D. D. Hoppes and R. P. Hudson, *Experimental Test of Parity Conservation in β Decay*, *Phys. Rev.* **105** (1957) 1413 (cit. on p. 8).
- [38] F. J. Hasert et al., *Search for Elastic ν_μ Electron Scattering*, *Phys. Lett. B* **46** (1973) 121 (cit. on p. 8).
- [39] G. Radell and R. Beyer, *Neutrino electron scattering*, *Mod. Phys. Lett. A* **8** (1993) 1067 (cit. on p. 8).
- [40] A. Pich, “The Standard model of electroweak interactions”, *The Standard model and beyond. Proceedings, 22nd International Meeting on Fundamental Physics, Jaca, Spain, February 7-11, 1994*, 1994 1, arXiv: hep-ph/9412274 [hep-ph] (cit. on p. 8).
- [41] R. Hofstadter, *Electron scattering and nuclear structure*, *Rev. Mod. Phys.* **28** (1956) 214 (cit. on p. 10).
- [42] Y. Nambu, *Quasiparticles and Gauge Invariance in the Theory of Superconductivity*, *Phys. Rev.* **117** (1960) 648, ed. by J. C. Taylor (cit. on p. 12).

- [43] J. Goldstone, *Field Theories with Superconductor Solutions*, *Nuovo Cim.* **19** (1961) 154 (cit. on p. 12).
- [44] J. Goldstone, A. Salam and S. Weinberg, *Broken Symmetries*, *Phys. Rev.* **127** (1962) 965 (cit. on p. 12).
- [45] P. Vilain et al., *Precision measurement of electroweak parameters from the scattering of muon-neutrinos on electrons*, *Phys. Lett. B* **335** (1994) 246 (cit. on p. 13).
- [46] D. M. Webber et al., *Measurement of the Positive Muon Lifetime and Determination of the Fermi Constant to Part-per-Million Precision*, *Phys. Rev. Lett.* **106** (2011) 041803, arXiv: [1010.0991 \[hep-ex\]](#) (cit. on p. 13).
- [47] A. H. G. Peter, *Dark Matter: A Brief Review*, (2012), arXiv: [1201.3942 \[astro-ph.CO\]](#) (cit. on p. 17).
- [48] R. Massey, T. Kitching and J. Richard, *The dark matter of gravitational lensing*, *Rept. Prog. Phys.* **73** (2010) 086901, arXiv: [1001.1739 \[astro-ph.CO\]](#) (cit. on p. 17).
- [49] K. Kajantie, M. Laine, K. Rummukainen and M. E. Shaposhnikov, *Is there a hot electroweak phase transition at $m_H \gtrsim m_W$?*, *Phys. Rev. Lett.* **77** (1996) 2887, arXiv: [hep-ph/9605288](#) (cit. on p. 17).
- [50] A. D. Sakharov, *Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe*, *Pisma Zh. Eksp. Teor. Fiz.* **5** (1967) 32 (cit. on p. 17).
- [51] J. H. Christenson, J. W. Cronin, V. L. Fitch and R. Turlay, *Evidence for the 2π Decay of the K_2^0 Meson*, *Phys. Rev. Lett.* **13** (1964) 138 (cit. on p. 17).
- [52] T. D. Lee, *A Theory of Spontaneous T Violation*, *Phys. Rev. D* **8** (1973) 1226, ed. by G. Feinberg (cit. on p. 18).
- [53] A. Pilaftsis and C. E. M. Wagner, *Higgs bosons in the minimal supersymmetric standard model with explicit CP violation*, *Nucl. Phys. B* **553** (1999) 3, arXiv: [hep-ph/9902371](#) (cit. on p. 18).
- [54] J. Braathen and S. Kanemura, *Leading two-loop corrections to the Higgs boson self-couplings in models with extended scalar sectors*, *Eur. Phys. J. C* **80** (2020) 227, arXiv: [1911.11507 \[hep-ph\]](#) (cit. on p. 19).
- [55] J. F. Gunion and H. E. Haber, *The CP conserving two Higgs doublet model: The Approach to the decoupling limit*, *Phys. Rev. D* **67** (2003) 075019, arXiv: [hep-ph/0207010](#) (cit. on p. 19).
- [56] S. Kanemura, S. Kiyoura, Y. Okada, E. Senaha and C. P. Yuan, *New physics effect on the Higgs selfcoupling*, *Phys. Lett. B* **558** (2003) 157, arXiv: [hep-ph/0211308](#) (cit. on p. 19).

-
- [57] R. A. Diaz and R. Martinez, *The Custodial symmetry*, Rev. Mex. Fis. **47** (2001) 489, arXiv: [hep-ph/0302058](#) (cit. on p. 20).
- [58] H. Bahl, J. Braathen and G. Weiglein, *New Constraints on Extended Higgs Sectors from the Trilinear Higgs Coupling*, Phys. Rev. Lett. **129** (2022) 231802, arXiv: [2202.03453 \[hep-ph\]](#) (cit. on p. 20).
- [59] K. Agashe, R. Contino and A. Pomarol, *The Minimal composite Higgs model*, Nucl. Phys. B **719** (2005) 165, arXiv: [hep-ph/0412089](#) (cit. on p. 21).
- [60] F. Bishara, R. Contino and J. Rojo, *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. 21, 25).
- [61] R. Contino, D. Marzocca, D. Pappadopulo and R. Rattazzi, *On the effect of resonances in composite Higgs phenomenology*, JHEP **10** (2011) 081, arXiv: [1109.1570 \[hep-ph\]](#) (cit. on pp. 21, 22).
- [62] J. Gao, L. Harland-Lang and J. Rojo, *The Structure of the Proton in the LHC Precision Era*, Phys. Rept. **742** (2018) 1, arXiv: [1709.04922 \[hep-ph\]](#) (cit. on pp. 22, 24).
- [63] M. Jamin, *QCD and renormalization group methods*, <http://www.maria-laach.tp.nt.uni-siegen.de/downloads/files/2006/Jamin-2006.pdf> (cit. on p. 23).
- [64] G. Altarelli and G. Parisi, *Asymptotic Freedom in Parton Language*, Nucl. Phys. B **126** (1977) 298 (cit. on p. 23).
- [65] V. N. Gribov and L. N. Lipatov, *Deep inelastic $e p$ scattering in perturbation theory*, Sov. J. Nucl. Phys. **15** (1972) 438 (cit. on p. 23).
- [66] Y. L. Dokshitzer, *Calculation of the Structure Functions for Deep Inelastic Scattering and e^+e^- Annihilation by Perturbation Theory in Quantum Chromodynamics.*, Sov. Phys. JETP **46** (1977) 641 (cit. on p. 23).
- [67] L. N. Lipatov, *The parton model and perturbation theory*, Yad. Fiz. **20** (1974) 181 (cit. on p. 23).
- [68] R. D. Ball et al., *Parton distributions from high-precision collider data*, Eur. Phys. J. C **77** (2017) 663, arXiv: [1706.00428 \[hep-ph\]](#) (cit. on p. 23).
- [69] A. D. Martin, W. J. Stirling, R. S. Thorne and G. Watt, *Parton distributions for the LHC*, Eur. Phys. J. C **63** (2009) 189, arXiv: [0901.0002 \[hep-ph\]](#) (cit. on p. 23).
- [70] J. Gao et al., *CT10 next-to-next-to-leading order global analysis of QCD*, Phys. Rev. D **89** (2014) 033009, arXiv: [1302.6246 \[hep-ph\]](#) (cit. on p. 23).

- [71] J. Butterworth et al., *PDF4LHC recommendations for LHC Run II*, *J. Phys. G* **43** (2016) 023001, arXiv: 1510.03865 [hep-ph] (cit. on pp. 23, 24, 44).
- [72] 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. 24).
- [73] 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. 24, 45, 71, 76, 81, 108).
- [74] CMS Collaboration, *Search for Higgs Boson Pair Production in the Four b Quark Final State in Proton-Proton Collisions at $\sqrt{s}=13$ TeV*, *Phys. Rev. Lett.* **129** (2022) 081802, arXiv: 2202.09617 [hep-ex] (cit. on p. 24).
- [75] 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. 25, 27, 77, 79, 80, 93, 106, 109–111, 113, 116, 117).
- [76] F. A. Dreyer and A. Karlberg, *Vector-Boson Fusion Higgs Pair Production at N^3 LO*, *Phys. Rev. D* **98** (2018) 114016, arXiv: 1811.07906 [hep-ph] (cit. on p. 24).
- [77] ATLAS Collaboration, *Measurements of the Higgs boson production and decay rates and coupling strengths using pp collision data at $\sqrt{s} = 7$ and 8 TeV in the ATLAS experiment*, *Eur. Phys. J. C* **76** (2016) 6, arXiv: 1507.04548 [hep-ex] (cit. on p. 26).
- [78] V. Khachatryan et al., *Observation of the Diphoton Decay of the Higgs Boson and Measurement of Its Properties*, *Eur. Phys. J. C* **74** (2014) 3076, arXiv: 1407.0558 [hep-ex] (cit. on p. 26).
- [79] M. E. Peskin and D. V. Schroeder, *An Introduction to quantum field theory*, Reading, USA: Addison-Wesley, 1995, ISBN: 978-0-201-50397-5 (cit. on p. 27).
- [80] *LHC Machine*, *JINST* **3** (2008) S08001, ed. by L. Evans and P. Bryant (cit. on p. 29).
- [81] CMS Collaboration, *The CMS Experiment at the CERN LHC*, *JINST* **3** (2008) S08004 (cit. on p. 30).
- [82] LHCb Collaboration, *The LHCb Detector at the LHC*, *JINST* **3** (2008) S08005 (cit. on p. 30).
- [83] ALICE Collaboration, *The ALICE experiment at the CERN LHC*, *JINST* **3** (2008) S08002 (cit. on p. 30).
- [84] *Views of CERN images gallery*, <https://home.cern/resources/image/cern/views-cern-images-gallery> (cit. on p. 30).

-
- [85] ATLAS Collaboration, *Public ATLAS Luminosity Results for Run-2 of the LHC*, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2> (cit. on p. 32).
- [86] ATLAS Collaboration, *Public ATLAS Luminosity Results for Run-3 of the LHC*, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun3> (cit. on p. 32).
- [87] J. Pequeno and P. Schaffner, “How ATLAS detects particles: diagram of particle paths in the detector”, 2013, URL: <https://cds.cern.ch/record/1505342> (cit. on p. 34).
- [88] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, *Comput. Phys. Commun.* **191** (2015) 159, arXiv: 1410.3012 [hep-ph] (cit. on p. 43).
- [89] M. Bahr et al., *Herwig++ Physics and Manual*, *Eur. Phys. J. C* **58** (2008) 639, arXiv: 0803.0883 [hep-ph] (cit. on p. 43).
- [90] 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 p. 43).
- [91] S. Alioli, P. Nason, C. Oleari and E. Re, *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 p. 43).
- [92] ATLAS Collaboration, *The ATLAS Simulation Infrastructure*, *Eur. Phys. J. C* **70** (2010) 823, arXiv: 1005.4568 [physics.ins-det] (cit. on p. 43).
- [93] S. Agostinelli et al., *GEANT4—a simulation toolkit*, *Nucl. Instrum. Meth. A* **506** (2003) 250 (cit. on p. 43).
- [94] R. D. Ball et al., *Parton distributions for the LHC Run II*, *JHEP* **04** (2015) 040, arXiv: 1410.8849 [hep-ph] (cit. on p. 44).
- [95] D. J. Lange, *The EvtGen particle decay simulation package*, *Nucl. Instrum. Meth. A* **462** (2001) 152, ed. by S. Erhan, P. Schlein and Y. Rozen (cit. on p. 44).
- [96] 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. 45, 118).
- [97] 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. 45, 118).

- [98] 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 pp. 45, 118).
- [99] 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 pp. 45, 118).
- [100] ATLAS Collaboration, *ATLAS Tracking Software Tutorial*, <https://atlassoftwaredocs.web.cern.ch/internal-links/tracking-tutorial/index.html> (cit. on pp. 50, 52).
- [101] 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. 50, 51).
- [102] H. Gray, *Track reconstruction with the ATLAS experiment*, https://indico.cern.ch/event/504284/contributions/2023875/attachments/1240146/1823137/HGray_Zurich_Tracking.pdf (cit. on p. 51).
- [103] ATLAS Collaboration, *A neural network clustering algorithm for the ATLAS silicon pixel detector*, *JINST* **9** (2014) P09009, arXiv: 1406.7690 [hep-ex] (cit. on p. 52).
- [104] E. E. Khoda, *ATLAS pixel cluster splitting using Mixture Density Networks*, tech. rep., CERN, 2019, URL: <https://cds.cern.ch/record/2687968> (cit. on p. 52).
- [105] 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. 52).
- [106] ATLAS Collaboration, *Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC*, *Eur. Phys. J. C* **77** (2017) 332, arXiv: 1611.10235 [physics.ins-det] (cit. on p. 54).
- [107] ATLAS Collaboration, *Performance of vertex reconstruction algorithms for detection of new long-lived particle decays within the ATLAS inner detector*, tech. rep., CERN, 2019, URL: <https://cds.cern.ch/record/2669425> (cit. on p. 55).
- [108] S. D. Ellis and D. E. Soper, *Successive combination jet algorithm for hadron collisions*, *Phys. Rev. D* **48** (1993) 3160, arXiv: hep-ph/9305266 (cit. on pp. 55, 56).

-
- [109] Y. L. Dokshitzer, G. D. Leder, S. Moretti and B. R. Webber, *Better jet clustering algorithms*, *JHEP* **08** (1997) 001, arXiv: [hep-ph/9707323](#) (cit. on pp. 55, 56).
- [110] G. P. Salam and G. Soyez, *A Practical Seedless Infrared-Safe Cone jet algorithm*, *JHEP* **05** (2007) 086, arXiv: [0704.0292 \[hep-ph\]](#) (cit. on p. 55).
- [111] M. Cacciari, G. P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063, arXiv: [0802.1189 \[hep-ph\]](#) (cit. on pp. 55, 56).
- [112] ATLAS Collaboration, *Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run I*, *Eur. Phys. J. C* **77** (2017) 490, arXiv: [1603.02934 \[hep-ex\]](#) (cit. on p. 57).
- [113] 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. 57).
- [114] ATLAS Collaboration, *Performance of pile-up mitigation techniques for jets in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector*, *Eur. Phys. J. C* **76** (2016) 581, arXiv: [1510.03823 \[hep-ex\]](#) (cit. on pp. 58, 60).
- [115] 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. 59).
- [116] ATLAS Collaboration, *Performance of jet substructure techniques for large- R jets in proton-proton collisions at $\sqrt{s} = 7$ TeV using the ATLAS detector*, *JHEP* **09** (2013) 076, arXiv: [1306.4945 \[hep-ex\]](#) (cit. on p. 60).
- [117] ATLAS Collaboration, *Boosted Object Tagging with Variable- R Jets in the ATLAS Detector*, tech. rep., CERN, 2016, URL: <https://cds.cern.ch/record/2199360> (cit. on p. 60).
- [118] ATLAS Collaboration, *Jet mass reconstruction with the ATLAS Detector in early Run 2 data*, tech. rep., CERN, 2016, URL: <https://cds.cern.ch/record/2200211> (cit. on p. 61).
- [119] 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. 61–63).
- [120] ATLAS Collaboration, *Monte Carlo Calibration and Combination of In-situ Measurements of Jet Energy Scale, Jet Energy Resolution and Jet Mass in ATLAS*, tech. rep., CERN, 2015, URL: <https://cds.cern.ch/record/2044941> (cit. on pp. 61, 62).
- [121] ATLAS Collaboration, *In situ calibration of large-radius jet energy and mass in 13 TeV proton–proton collisions with the ATLAS detector*, *Eur. Phys. J. C* **79** (2019) 135, arXiv: [1807.09477 \[hep-ex\]](#) (cit. on pp. 62, 63, 91).

- [122] ATLAS Collaboration, *Configuration and performance of the ATLAS b -jet triggers in Run 2*, *Eur. Phys. J. C* **81** (2021) 1087, arXiv: 2106.03584 [hep-ex] (cit. on p. 64).
- [123] 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 pp. 64, 65).
- [124] 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 p. 64).
- [125] D. Krohn, J. Thaler and L.-T. Wang, *Jets with Variable R* , *JHEP* **06** (2009) 059, arXiv: 0903.0392 [hep-ph] (cit. on p. 64).
- [126] ATLAS Collaboration, *Variable Radius, Exclusive- k_T , and Center-of-Mass Subject Reconstruction for Higgs($\rightarrow b\bar{b}$) Tagging in ATLAS*, tech. rep., CERN, 2017, URL: <https://cds.cern.ch/record/2268678> (cit. on p. 65).
- [127] ATLAS Collaboration, *Identification of Boosted Higgs Bosons Decaying Into $b\bar{b}$ With Neural Networks and Variable Radius Subjects in ATLAS*, tech. rep., CERN, 2020, URL: <https://cds.cern.ch/record/2724739> (cit. on p. 66).
- [128] ATLAS Collaboration, *Efficiency corrections for a tagger for boosted $H \rightarrow b\bar{b}$ decays in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, tech. rep., CERN, 2021, URL: <https://cds.cern.ch/record/2777811> (cit. on p. 67).
- [129] R. Frühwirth, *A Gaussian-mixture approximation of the Bethe–Heitler model of electron energy loss by bremsstrahlung*, *Computer Physics Communications* **154** (2003) 131, ISSN: 0010-4655, URL: <https://www.sciencedirect.com/science/article/pii/S0010465503002923> (cit. on p. 68).
- [130] 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 p. 68).
- [131] 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. 68).
- [132] ATLAS Collaboration, *Performance of missing transverse momentum reconstruction with the ATLAS detector using proton-proton collisions at $\sqrt{s} = 13$ TeV*, *Eur. Phys. J. C* **78** (2018) 903, arXiv: 1802.08168 [hep-ex] (cit. on p. 69).

-
- [133] ATLAS Collaboration, *Selection of jets produced in 13 TeV proton-proton collisions with the ATLAS detector*, tech. rep., CERN, 2015, URL: <https://cds.cern.ch/record/2037702> (cit. on p. 74).
- [134] B. Abbott et al., *Extraction of the width of the W boson from measurements of $\sigma(p\bar{p} \rightarrow W + X) \times B(W \rightarrow e\nu)$ and $\sigma(p\bar{p} \rightarrow Z + X) \times B(Z \rightarrow ee)$ and their ratio*, *Phys. Rev. D* **61** (2000) 072001, arXiv: [hep-ex/9906025](https://arxiv.org/abs/hep-ex/9906025) (cit. on p. 81).
- [135] ATLAS Collaboration, *Search for heavy neutral Higgs bosons decaying into a top quark pair in 140 fb⁻¹ of proton-proton collision data at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **08** (2024) 013, arXiv: [2404.18986](https://arxiv.org/abs/2404.18986) [[hep-ex](#)] (cit. on p. 81).
- [136] J. H. Friedman, *Greedy function approximation: A gradient boosting machine.*, *Annals Statist.* **29** (2001) 1189 (cit. on pp. 82, 83).
- [137] J. R. Quinlan, *Induction of decision trees*, *Machine learning* **1** (1986) 81 (cit. on p. 82).
- [138] T. Chen and C. Guestrin, *XGBoost: A Scalable Tree Boosting System*, (2016), arXiv: [1603.02754](https://arxiv.org/abs/1603.02754) [[cs.LG](#)] (cit. on p. 83).
- [139] P. Baldi, K. Cranmer, T. Faucett, P. Sadowski and D. Whiteson, *Parameterized neural networks for high-energy physics*, *Eur. Phys. J. C* **76** (2016) 235, arXiv: [1601.07913](https://arxiv.org/abs/1601.07913) [[hep-ex](#)] (cit. on p. 85).
- [140] *TRExFitter Documentation*, URL: <https://trexfitter-docs.web.cern.ch/trexfitter-docs/> (cit. on pp. 88, 130).
- [141] ATLAS Collaboration, *ATLAS data quality operations and performance for 2015–2018 data-taking*, *JINST* **15** (2020) P04003, arXiv: [1911.04632](https://arxiv.org/abs/1911.04632) [[physics.ins-det](#)] (cit. on p. 91).
- [142] 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](https://arxiv.org/abs/2212.09379) [[hep-ex](#)] (cit. on p. 91).
- [143] G. Avoni et al., *The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS*, *JINST* **13** (2018) P07017 (cit. on p. 91).
- [144] ATLAS Collaboration, *JESJERRecommendationsR21*, <https://twiki.cern.ch/twiki/bin/view/AtlasProtected/JetUncertaintiesRel21Summer2018SmallR?rev=11> (cit. on p. 91).
- [145] *SM Higgs Branching Ratios*, URL: <https://twiki.cern.ch/twiki/bin/view/LHCPhysics/CERNYellowReportPageBR?rev=22> (cit. on p. 95).

- [146] B. Mellado Garcia, P. Musella, M. Grazzini and R. Harlander, *CERN Report 4: Part I Standard Model Predictions*, (2016), URL: <https://cds.cern.ch/record/2150771> (cit. on p. 95).
- [147] G. Cowan, K. Cranmer, E. Gross and O. Vitells, *Asymptotic formulae for likelihood-based tests of new physics*, *Eur. Phys. J. C* **71** (2011) 1554, [Erratum: *Eur.Phys.J.C* 73, 2501 (2013)], arXiv: [1007.1727](https://arxiv.org/abs/1007.1727) [[physics.data-an](#)] (cit. on pp. 98, 100, 104).
- [148] 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 p. 103).
- [149] A. Wald, *Tests of Statistical Hypotheses Concerning Several Parameters When the Number of Observations is Large*, *Transactions of the American Mathematical Society*, Vol. 54, No. 3 (Nov., 1943), pp. 426-482 () (cit. on p. 104).
- [150] S. S. Wilks, *The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses*, *Annals Math. Statist.* **9** (1938) 60 (cit. on p. 104).
- [151] CMS Collaboration, *Search for Nonresonant Pair Production of Highly Energetic Higgs Bosons Decaying to Bottom Quarks*, *Phys. Rev. Lett.* **131** (2023) 041803, arXiv: [2205.06667](https://arxiv.org/abs/2205.06667) [[hep-ex](#)] (cit. on p. 118).
- [152] H. Qu and L. Gouskos, *ParticleNet: Jet Tagging via Particle Clouds*, *Phys. Rev. D* **101** (2020) 056019, arXiv: [1902.08570](https://arxiv.org/abs/1902.08570) [[hep-ph](#)] (cit. on p. 118).
- [153] ATLAS Collaboration, *Transformer Neural Networks for Identifying Boosted Higgs Bosons decaying into $b\bar{b}$ and $c\bar{c}$ in ATLAS*, tech. rep., CERN, 2023, URL: <https://cds.cern.ch/record/2866601> (cit. on pp. 119, 122, 123).
- [154] ATLAS Collaboration, *Technical Design Report for the Phase-II Upgrade of the ATLAS Trigger and Data Acquisition System - Event Filter Tracking Amendment*, tech. rep., CERN, 2022, URL: <http://cds.cern.ch/record/2802799> (cit. on p. 119).
- [155] ATLAS Collaboration, *Technical Design Report for the ATLAS Inner Tracker Pixel Detector*, tech. rep., CERN, 2017, URL: <https://cds.cern.ch/record/2285585> (cit. on p. 119).
- [156] ATLAS Collaboration, *Technical Design Report: A High-Granularity Timing Detector for the ATLAS Phase-II Upgrade*, tech. rep., CERN, 2020, URL: <https://cds.cern.ch/record/2719855> (cit. on p. 119).
- [157] ATLAS Collaboration, *HL-LHC prospects for the search of boosted Higgs boson pair production via vector-boson fusion in the $b\bar{b}b\bar{b}$ final state at the ATLAS experiment*, tech. rep., CERN, 2025, URL: <https://cds.cern.ch/record/2925813> (cit. on p. 120).
- [158] A. Vaswani et al., *Attention is all you need*, *Advances in Neural Information Processing Systems* (2017) (cit. on p. 122).

-
- [159] ATLAS Collaboration, *Optimisation of large-radius jet reconstruction for the ATLAS detector in 13 TeV proton–proton collisions*, *Eur. Phys. J. C* **81** (2021) 334, arXiv: [2009.04986 \[hep-ex\]](#) (cit. on p. [124](#)).
- [160] ATLAS Collaboration, *Improving jet substructure performance in ATLAS using Track-CaloClusters*, tech. rep., CERN, 2017, URL: <https://cds.cern.ch/record/2275636> (cit. on p. [124](#)).
- [161] D. E. Rumelhart, G. E. Hinton and R. J. Williams, *Learning representations by back-propagating errors*, *Nature* **323** (1986) 533 (cit. on p. [129](#)).
- [162] P. Stylianou and G. Weiglein, *Constraints on the trilinear and quartic Higgs couplings from triple Higgs production at the LHC and beyond*, *Eur. Phys. J. C* **84** (2024) 366, arXiv: [2312.04646 \[hep-ph\]](#) (cit. on pp. [133](#), [135](#), [136](#)).
- [163] D. de Florian, I. Fabre and J. Mazzitelli, *Triple Higgs production at hadron colliders at NNLO in QCD*, *JHEP* **03** (2020) 155, arXiv: [1912.02760 \[hep-ph\]](#) (cit. on p. [134](#)).
- [164] G. V. Moustakides and K. Basioti, *Training neural networks for likelihood/density ratio estimation*, (2019), arXiv: [1911.00405](#) (cit. on p. [145](#)).

List of Figures

2.1	Potential $V(\phi)$ of a complex scalar field $\phi(x)$ as given in Equation 2.25 for $\lambda > 0$ and $\mu^2 < 0$. Excitations of the field $\phi(x)$ from the ground state are parametrised in terms of the real fields $\varphi_1(x)$ and $\varphi_2(x)$ according to Equation 2.28.	12
2.2	Measurements of the reduced couplings κ of the Higgs boson to fermions and vector boson performed by the ATLAS collaboration [27]. Two scenarios have been considered to cope with the low sensitivity to the charm coupling: in the first, the charm coupling modifier is set to be equal to the top coupling modifier (coloured circle markers); in the second, the charm coupling modifier is left as a free floating parameter in the fits (grey cross markers). Notice the y-axis is different for the vector bosons and is expressed in terms of $\sqrt{\kappa}$ instead of κ	15
2.3	Upper limits at 95% CL on the HH production cross-section as a function of κ_λ obtained by the ATLAS collaboration [29]. The combined result (black) is the combination from several analyses. Regions where the predicted cross-section is higher than the observed cross-section are excluded at 95% CL.	16
2.4	Upper limits at 95% CL on the HH production cross-section as a function of κ_{2V} obtained by the CMS collaboration [28]. Regions where the predicted cross-section is higher than the observed cross-section are excluded at 95% CL.	16
2.5	Trilinear Higgs-boson self-coupling modifier κ_λ as a function of the pseudo-scalar Higgs mass m_A at one-loop (dashed curve) and two-loop (solid curve) orders [58]. The mass of the charged Higgs boson is set to be equal to m_A , while the heavy Higgs-boson mass $m_{\tilde{H}}$ and the scale M are fixed to 600 GeV. Further improvements on the experimental side to constrain the values of κ_λ will decrease the range of allowed values of m_A that can be explored. . . .	20
2.6	Schematic representation of the cross-section factorisation mechanism. A parton a with momentum fraction x_1 coming from a proton with momentum p_1 interacts with a parton b with momentum fraction x_2 coming from the second proton with momentum p_2 . The partonic cross-section $\hat{\sigma}$ depends on the process under study.	23

2.7	PDF4LHC PDF sets [71] as a function of x , the incoming proton momentum fraction, for different choices of the factorisation scale [62].	24
2.8	Feynman diagrams at leading order representing the gluon-gluon fusion process to produce a pair of Higgs bosons. The κ -modifiers at each Higgs vertex are included in the diagrams to mark the couplings for the different vertices [75].	25
2.9	Feynman diagrams at leading order representing the vector-boson fusion process to produce a pair of Higgs bosons. The κ -modifiers at each Higgs vertex are included in the diagrams to mark the couplings for the different vertices [75].	25
2.10	Reconstructed invariant mass distributions of the HH system. In (a) the pair of Higgs bosons are produced via non-resonant VBF under different assumptions for κ_{2V} . In (b) the BSM resonant VBF production mode for resonances with different masses m_X . All histograms are normalised to unity.	26
2.11	Leading order Feynman diagram for the resonant production of a pair of Higgs bosons via vector-boson fusion via an intermediate heavy X [75]. . .	27
2.12	Branching ratios of a pair of Higgs bosons with $m_H = 125$ GeV. The dominant decay mode is $HH \rightarrow b\bar{b}b\bar{b}$ and corresponds to almost 34% of the total HH decays. Numbers based on Ref. [32].	28
3.1	Schematic illustration of the LHC complex [84].	30
3.2	(a) Peak instantaneous luminosity delivered to the ATLAS detector during stable beams condition as a function of time in 2018. (b) Integrated luminosity as a function of the time during the Run-2 delivered by the LHC, recorded by ATLAS, and certified to be of good quality data [85].	32
3.3	Luminosity-weighted distribution of the average number of interactions per bunch crossing $\langle\mu\rangle$ for different years during the Run-2 data-taking [85]. . .	32
3.4	(a) The ATLAS detector [31]. (b) Cross-sectional view of the ATLAS barrel to exemplify the interaction between the different particle types with the detector elements [87].	34
3.5	Illustration of the ATLAS Inner Detector. The Pixel Detector is the closest to the beam axis. In sequence, there are the Semiconductor Tracker and the Transition Radiation Tracker. The IBL is not displayed in the figure as it was included in the Pixel Detector only after the end of the Run-1 [31].	37
3.6	Illustration of the ATLAS calorimeter system [31]. The Electromagnetic Calorimeter surrounds the Pixel Detector and is overlaid by the Hadronic Calorimeter. The calorimeter layout and technologies are different in the barrel and end-cap regions. The main components are pointed in the figure.	38
3.7	Schematic representation of the electromagnetic calorimeter. The thickness of each layer is given in terms of radiation length X_0 (see text) [31].	39

3.8	Schematic representation of the ATLAS Muon Spectrometer [31]. The Monitored Drift Tubes and Cathode Strip chambers are dedicated for precise muon trajectory reconstruction, while the Resistive Plate Chambers and Thin Gap Chambers are used primarily to provide fast response to the ATLAS trigger system.	41
3.9	Comparison of the m_{HH} distribution for the VBF $\kappa_{2V} = 0$ signal sample if using directly simulated events or the reweighting technique of Equation 3.12.	46
4.1	Schematic perigee representation of the track parameters [100]. Here, the perigee representation describes the particle trajectory at its point of closest approach to the reference point.	50
4.2	Illustration of the ATLAS track reconstruction steps. Each yellow dot represents a hit in the Inner Detector. The solid curves represent selected triplets of hits to form the initial track seeds. The dashed curve represents the adding of extra hits to the particle trajectory compatible with the original track seed. All particle trajectories passing a stringent set of requirements are then extended to the TRT to improve the resolution of the track momenta [102].	51
4.3	(a) Several clusters, each activated by a single ionising particle. (b) Cluster activated by several ionising particles. Different colors refer to the energy deposits of different particles [100].	52
4.4	(a) Average number of pixel clusters used to reconstruct a track that were identified as merged by the neural network as a function of $\Delta R(\text{trk}, \text{jet})$, for two different intervals of jet p_T . (b) Normalised average r_{trk} as a function of jet p_T . Both plots use data collected in 2018.	54
4.5	Illustrative examples of the shape of the reconstructed jets using (a) the k_t algorithm, (b) the Cambridge/Aachen algorithm, and (c) the anti- k_t algorithm for the same set of input particles. In all cases the radius parameter was set to 1.0 for simplicity. Different colors correspond to different jets [111].	56
4.6	Comparison of hard-scatter and pile-up jets in the range $20 \text{ GeV} < p_T < 30 \text{ GeV}$ in the variables (a) corrJVF, (b) R_{p_T} , and (c) JVT [115].	59
4.7	Schematic representation of how the angular separation between the decay products of the process $A \rightarrow B + C$ becomes smaller for increasing transverse momentum of the mother particle A	60
4.8	Illustration of an interaction producing one b -jet and two light-flavour jets. The b -hadron travels a distance L_{xy} before decaying and producing displaced tracks originated from a secondary vertex. Compared to prompt tracks originating from the primary vertex, the displaced tracks have a large transverse impact parameter d_0 [122].	64
4.9	Background rejection in simulated $t\bar{t}$ events as a function of the b -jet efficiency for the different single b -tagging algorithms. The two lower ratio-pads is with respect to the MV2c10 algorithm. In the labels, f_c is a parameter used to construct the tagger discriminants [123].	65

4.10	Multijet (a) and top (b) jet rejection as a function of large- R jet p_T for the 60% efficiency $H \rightarrow b\bar{b}$ WP. The double b -tagger performance is shown as blue dots and supersedes the alternative methods [127].	66
4.11	Signal efficiency scale factors for the double b -tagger at 60% WP in four large- R jet p_T bins. The black error bars indicate the statistical component of the scale factor uncertainties and the green bands indicate the combined statistical plus systematic uncertainties. The calibration was performed using $Z(\rightarrow b\bar{b}) + \text{jets}$ and $Z(\rightarrow b\bar{b}) + \gamma$ events. This figure is adapted from Ref. [128] based on the b -tagger version used in the analysis presented here.	67
5.1	Schematic representation of the VBF production of pairs of boosted Higgs-bosons decaying into four b -quarks. Each large- R jet represents a Higgs boson decaying into a pair of b -quarks passing a double b -tagging algorithm criteria. The two extra small- R jets account for the VBF signature of the production mode.	72
5.2	Trigger efficiency as a function of the leading large- R jet p_T for the trigger chains of (a) the combined 2015+2016, (b) the 2017, and (c) the 2018 datasets. The vertical dashed line indicates the plateau region at $p_T = 450$ GeV above which all trigger chains present a stable efficiency. The non-resonant VBF $\kappa_{2V} = 0$ signal has been used in this study.	73
5.3	Reconstructed masses of the leading (solid line) and the sub-leading (dashed line) Higgs-bosons candidates for (a) the non-resonant VBF $\kappa_{2V} = 0$ and (b) resonant narrow width $m_X = 2.0$ TeV signals.	75
5.4	Distributions for (a) the invariant mass and (b) the pseudorapidity separation of the system of VBF jets for the MC simulated background, the non-resonant VBF $\kappa_{2V} = 0$ signal, and the resonant narrow-width $m_X = 2.0$ TeV signal. The vertical gray dashed lines indicate the thresholds where a selection requirement based on the variables is applied to select events with signal-like characteristics.	76
5.5	Reconstructed m_{H_1} - m_{H_2} plane for the (a) 1PASS and (b) 2PASS selections of the analysis, shown for the data events. The mass planes for the 2PASS selection of the analysis are also shown for the (c) VBF SM $\kappa_{2V} = 1$, (d) VBF $\kappa_{2V} = 0$, and (e) $m_X = 1.0$ TeV spin-0 narrow-width resonance HH samples. The SR is described by the continuous red line and captures most of the signal events, as seen in c), d), and e). The VR is defined as the region between the continuous red line and the dashed yellow line. The CR is delimited between the dashed yellow line and the dotted green line [75].	77
5.6	Cumulative acceptance times efficiency as a function of κ_{2V} for each event selection step in the VBF non-resonant analysis. “Upstream selection” refers to selections applied when reconstructing objects and cleaning events [75].	79

5.7	Cumulative acceptance times efficiency in the resonant search as a function of m_X for each event selection step for the (a) narrow-width and (b) broad-width hypotheses. “Upstream selection” refers to selections applied when reconstructing objects and cleaning events [75].	80
5.8	Data distributions in the CR 2Pass and CR 1Pass regions of (a) the invariant mass of the HH system, (b) the reconstructed mass of the leading Higgs boson, and (c) the reconstructed p_T of the leading Higgs boson. All uncertainties are of statistical nature.	82
5.9	Schematic representation of a decision tree. It starts with the root containing all samples meant to be classified. A series of decision requirements are applied to split the samples across different internal nodes.	83
5.10	Distributions of (a) the invariant mass of the HH system and transverse momentum of (b) the leading and (c) the sub-leading Higgs-bosons candidates, for data-driven background and non-resonant signal events. The SM $\kappa_{2V} = 1$ distributions are shown in black and differ notably from the $\kappa_{2V} \neq 1$ distributions.	85
5.11	ROC curves and AUC evaluated for several non-SM VBF HH signal events and the background components using the BDT classifier. The ratio plot compares the performances of the signal samples with respect to the proxy $\kappa_{2V} = 0$ signal. The dashed gray line indicates the hypothetical ROC curve if the model makes random predictions only.	86
5.12	Distributions of (a) the invariant mass of the HH system and transverse momentum of (b) the leading and (c) the sub-leading Higgs-bosons candidates, for the data-driven background and resonant signal events under the narrow-width hypothesis. The kinematic distributions of signal events depend heavily on the value of m_X	86
5.13	ROC curves and AUC evaluated for the background and several (a) narrow-width and (b) broad-width signals using the pBDT classifier. The dashed gray line indicates the hypothetical ROC curve if the model makes random predictions only.	87
5.14	Invariant mass m_{HH} distributions for the resonant $m_X = 2.0$ TeV signal under the narrow-width (solid line) and broad-width (dashed line) hypotheses.	88
5.15	Impact of the systematic uncertainties of the double b -tagger scale factors per calibration p_T -bin on the VBF non-resonant $\kappa_{2V} = 0$ signal as a function of the BDT score.	92
5.16	BDT distributions in data in the VR 2Pass and VR 1Pass regions using (a) the optimal binning for the non-resonant analysis and (b) the two-bin smoothed version of the binning. The shape systematic uncertainties are estimated from the ratio pads [75].	93

5.17	Estimated parton shower normalisation uncertainties in the narrow-width resonant analysis. The blue points indicate estimations from the direct comparison between nominal and alternative parton shower samples. The red dashed line is obtained via linear interpolation. The same uncertainties are used for broad-width samples.	96
6.1	Pre-fit BDT distributions of observed data, expected background, and the non-resonant VBF $\kappa_{2V} = 0$, SM VBF, and SM ggF signals. The distributions of the SM VBF and SM ggF signals were scaled for visualisation purposes. All sources of uncertainties have been included. No data has been observed in the rightmost signal-like bin.	99
6.2	Pre-fit pBDT distributions of observed data, expected background, and signal under the narrow-width assumption evaluated at three different values of m_χ . The total signal cross-section has been normalised to 1 fb. All sources of uncertainties have been included. No data has been observed in the rightmost signal-like bin in all scenarios where $m_\chi \geq 1.6$ TeV.	99
6.3	Illustration of (a) the relation between the p -value and the significance Z in the Gaussian distribution, and (b) the relation between the p -value and the observed test statistic $t_{\mu,\text{obs}}$	101
6.4	Post-fit BDT output score distributions in data, the expected background, SM VBF and SM ggF HH production, and the BSM scenario $\kappa_{2V} = 0$. All sources of uncertainties have been included. The distributions of the SM VBF and SM ggF signals were scaled for visualisation purposes. The observed data agrees well with the total background expectation [75]. . . .	106
6.5	Pulls (lower x-axis) of the nuisance parameters that parametrise the systematic uncertainties or the uncorrelated statistical uncertainties at each distribution bin (prefix <code>gamma_</code>) and their impact $\Delta\mu$ on the signal strength (upper x-axis) for the VBF SM signal. No fitted nuisance parameter related to a systematic uncertainty has either a large pull or is over-constrained. Only the top 20 most impactful uncertainties are included in the plot. The names of the nuisance parameters are described in Appendix B.	107
6.6	Derivation of the confidence interval on κ_{2V} using (a) observed and expected values of the test statistic $-2 \ln \Lambda$ as a function of κ_{2V} in the boosted analysis. The dashed horizontal lines indicate the thresholds above which a hypothesis is excluded at 68% and 95% CL. With the alternative method (b) using the 95% CL upper limit on the VBF $pp \rightarrow HH$ cross-section as a function of κ_{2V} . The solid and dashed lines indicate respectively the observed and expected limits. The cyan and yellow bands represent the 1σ and 2σ intervals of the expected limit. The star marker indicates the SM prediction. In both approaches the κ_λ and κ_V coupling strengths are fixed to their SM values. .	108

6.7	Observed (a) and expected (b) distributions of the likelihood $-2 \ln \Lambda$ as a function of κ_{2V} in the boosted (blue dashed), resolved (green dotted), and combined (black solid) analyses. The dashed horizontal lines indicate the thresholds above which a hypothesis is excluded at 68% and 95% CL. Both κ_λ and κ_V are fixed to their SM values [75].	109
6.8	Upper limit at 95% CL on the VBF $pp \rightarrow HH$ cross-section as a function of κ_{2V} . The solid and dashed black lines indicate respectively the observed and expected combined limits. The cyan and yellow bands represent the 1σ and 2σ intervals of the expected limit. The blue dash-dotted and the green dotted curves represent the expected upper limits in the boosted-only and resolved-only analyses, respectively. The star marker indicates the SM prediction. Both κ_λ and κ_V are fixed to their SM values [75].	110
6.9	Observed (a) and expected (b) distributions of the likelihood $-2 \ln \Lambda$ as a function of κ_λ in the boosted, resolved, and combined analyses. The dashed horizontal lines indicate the thresholds above which a hypothesis is excluded at 68% and 95% CL. Both κ_{2V} and κ_V are fixed to their SM values [75].	110
6.10	Upper limits at 95% CL on the ggF+VBF $pp \rightarrow HH$ cross-section as a function of κ_λ . The solid and dashed black lines indicate respectively the observed and expected combined limits. The cyan and yellow bands represent the 1σ and 2σ intervals of the expected limit. The blue dash-dotted and the green dotted curves represent the expected upper limits in the boosted-only and resolved-only analyses, respectively. The star marker indicates the SM prediction. Both κ_{2V} and κ_V are fixed to their SM values [75].	111
6.11	Observed (a) and expected (b) likelihood contours at 68% and 95% CL in the $\kappa_{2V} - \kappa_\lambda$ plane in the boosted (blue curves), resolved (red curves), and combined (black curves) analyses. The solid (dashed) curves indicate the limits at 95% (68%) CL. The cross-marker indicates the observed best-fit values of κ_{2V} and κ_λ , while the star marker indicates the SM prediction. All other Higgs-boson couplings are fixed to their SM values. The results for $\kappa_\lambda > 15$ are not plotted for better readability [75].	111
6.12	Post-fit pBDT distributions of observed data, expected background, and signal evaluated for three representative values of m_X . All sources of uncertainties have been included. The narrow width scenario is considered and the signal cross-section is normalised to 1 fb. The distributions observed in data agree well with the total background expectation [75].	113
6.13	Pulls (lower x-axis) of the nuisance parameters that parametrise the systematics uncertainties or the uncorrelated statistical uncertainties at each distribution bin (prefix gamma_) and their impact on the signal strength (upper x-axis) for (a) $m_X = 1.0$ TeV and (b) $m_X = 1.5$ TeV. No fitted nuisance parameter related to a systematic uncertainty has either a large pull or is over-constrained. Only the top 20 most impactful uncertainties are included in the plots. The names of the nuisance parameters are described in Appendix B.	114

6.14	Observed and expected likelihood $-2 \ln \Lambda$ scan as a function of the signal strength in the scenarios (a) $m_X = 1.0$ TeV and (b) $m_X = 1.6$ TeV under the narrow-width assumption. Only for low masses the likelihood scan exhibits a parabola-like shape with well-defined minimum.	115
6.15	Upper limit at 95% CL on the VBF $pp \rightarrow X \rightarrow HH$ cross-section as a function of m_X in the narrow-width scenario. The solid and dashed lines indicate, respectively, the observed and expected limits. The cyan and yellow bands represent the 1σ and 2σ intervals of the expected limit. A linear interpolation was used to derive the upper limits for intermediate mass values based on the values obtained for available signal samples (black dots) [75].	116
6.16	Upper limit at 95% CL on the VBF $pp \rightarrow X \rightarrow HH$ cross-section as a function of m_X in the broad-width scenario using the Composite Higgs model. The solid and dashed lines indicate, respectively, the observed and expected limits. The cyan and yellow bands represent the 1σ and 2σ intervals of the expected limit. A linear interpolation was used to derive the upper limits for intermediate mass values based on the values obtained for available signal samples (black dots). The theory predicted cross-section is represented by the red curve [75].	117
6.17	Observed and expected distributions of the likelihood $-2 \ln \Lambda$ as a function of κ_{2V} for the combined (black) HH analyses. The expected results for the individual $b\bar{b}b\bar{b}$ (blue), $b\bar{b}\tau^+\tau^-$ (green), $b\bar{b}\gamma\gamma$ (purple), multilepton (cyan), and $b\bar{b}\ell\ell + E_T^{\text{miss}}$ (brown) analyses are also shown [29].	119
6.18	Extrapolation of the expected κ_{2V} interval at 95% CL in the boosted VBF $HH \rightarrow b\bar{b}b\bar{b}$ search for the HL-LHC as a function of the double b -tagger signal efficiency and for different integrated luminosities. The yellow vertical line represents the tagger WP in the analysis presented in this thesis. The green vertical line is the equivalent tagger WP under the same background rejection for the new double b -tagging algorithm discussed in Section 7.1 [157].	120
7.1	Rejection rates for the (a) multijet and (b) hadronic top-decay backgrounds as a function of the large- R jet p_T when considering the 50% signal efficiency WP. The new GN2X model (blue) exhibits a much better performance compared to the previous tagger (pink) and the alternative approach tagging the sub-jets inside the large- R jet (green). The ratio pad is with respect to the previous D_{Xbb} tagger performance [153].	123
7.2	(a) Value of the expected κ_{2V} interval size at 95% CL, (b) the expected $\kappa_{2V} = 0$ exclusion significance, and (c) the overall background yield for different selection requirements on m_{jj} and $ \Delta\eta(j_1, j_2) $. The nominal tagger WP corresponding to 0.46% of QCD multijets efficiency is used here. The background is estimated as described in Section 7.4.	126

7.3	Reconstructed m_{H_1} - m_{H_2} plane for the (a) the 1Pass and (b) the 2Pass regions in data. The same mass plane for the 2Pass region for two representative signal hypotheses: (c) SM VBF $\kappa_{2V} = 1$ and (d) BSM VBF $\kappa_{2V} = 0$. The SR is described by the continuous red line. The VR and CR are split in quadrants and delimited by the continuous dashed yellow and dotted green lines, respectively. 2Pass data events are still blinded inside the SR.	127
7.4	Distributions of the DNN output score of data in the VR 2Pass and VR 1Pass regions. The shape systematic uncertainties are estimated from the ratio pads.	128
7.5	ROC curves and AUC evaluated for several samples and the background using the DNN classifier. The ratio pad compares the performances of the samples with respect to the proxy VBF $\kappa_{2V} = 0.5$ sample used in the model training. The dashed gray line indicates the hypothetical ROC curve if the model made random predictions only.	130
7.6	Distributions of the DNN output score of several samples and the data-driven background model using the optimal binning. The SM samples were scaled for visualisation purposes. The uncertainties are of statistical nature only. .	131
7.7	Expected likelihood $-2 \ln \Lambda$ distribution as a function of κ_{2V} while fixing the other Higgs-boson couplings to their SM prediction. The curves in red and cyan are from the new analysis based on the GN2X tagger, but with different choices for the dataset. The curve in black represents the main Run-2 analysis of this thesis based on the previous tagger, but when including only statistical and the background modelling systematic uncertainties.	132
7.8	Expected likelihood $-2 \ln \Lambda$ distribution in the new analysis as a function of κ_{2V} while fixing the other Higgs-boson couplings to their SM prediction. The red curve corresponds to a scenario where only statistical errors and systematic uncertainties from the background modelling are included, while the black curve is obtained when additionally including the impact from the dummy GN2X double b -tagger uncertainties.	132
8.1	Leading order Feynman diagrams representing the ggF process to produce three Higgs bosons in the SM. The κ -modifiers at each Higgs vertex are included in the diagrams to mark the couplings for the different vertices. . .	134
8.2	Observed and expected likelihood contours at 68% and 95% CL in the κ_λ - κ_4 plane. The unitary bound delimited by the gray dashed curve was calculated in Ref. [162]. The red star indicates the SM prediction. The notation κ_3 was used to represent κ_λ in the plot [30].	136
8.3	Background (cyan) and SM HHH signal (red) model output distributions when using the baseline DNN (solid lines) and the transformer-based architectures (dashed lines).	137

8.4	Expected likelihood $-2 \ln \Lambda$ as a function of κ_4 while fixing the other Higgs-boson couplings to their SM prediction. The red (cyan) curve is the result obtained using the transformer (baseline DNN) output as a fit observable. Only statistical uncertainties are included in the fit. The dashed gray line delimits the region where a hypothesis is excluded at 95% CL.	138
A.1	Comparison of 2Pass, 1Pass, and reweighted 1Pass data distributions in an extended region of the mass-plane of the reconstructed Higgs bosons. . . .	146

List of Tables

2.1	Properties of the elementary particles described by the Standard Model. The second, third, and fourth columns (in blue) show the properties of the fermions. The upper half shows the electron (e), muon (μ), tau (τ), together with their corresponding neutrinos ν . The lower half describes the quarks up (u), down (d), strange (s), charm (c), bottom (b), and top (t). The last column, in red, shows the properties, respectively, of the bosons W^\pm , Z , gluon (g), and photon (γ). The very last row in the table (in orange) contains the properties of the Higgs boson (H). Uncertainties and additional properties can be found in Ref. [32].	6
5.1	Event yields after each selection step for data and several representative signal hypotheses: non-resonant SM ggF, non-resonant SM VBF, non-resonant VBF $\kappa_{2V} = 0$, and resonant narrow width $m_X = 1.0$ TeV and $m_X = 5.0$ TeV signal samples. All signal yields except “Raw events” are MC generator-weighted and normalised to 140 fb^{-1} . Resonant samples were additionally normalised using $\sigma(pp \rightarrow X \rightarrow HH) = 1 \text{ fb}$. “Upstream selection” refers to selections applied when reconstructing objects and cleaning events. No “Veto resolved selection” is applied to resonant signals. The yields for the broad-width signal samples are similar to those of the narrow-width signal samples listed here [75].	80
5.2	Kinematic variables used as input features for the BDTs models. The same set of variables is used in the non-resonant and resonant searches.	84
5.3	Bins edges of the non-resonant BDT model and of resonant pBDTs for each signal hypothesis. The optimal binnings were obtained using the “transformation D” algorithm of the ATLAS fitting package TRExFitter. If the lowest bin during the optimisation had a width lower than 0.01, it has been merged to the following one.	89
5.4	Systematic uncertainties considered in the analysis and the effects on the fitting variable (normalization or shape effects).	90

6.1	Total yields of observed data, expected background, and SM VBF and ggF signals, in the non-resonant and resonant searches. The total uncertainty of the expected background in the resonant analysis depends on the mass hypothesis.	98
6.2	Fitted signal strength for each mass hypothesis in the narrow-width assumption. The total upper error of the fit parameter is divided into statistical and systematic components. The constraint $\mu_{\text{VBF}, X \rightarrow HH} > 0$ has been applied in the fit. The signal cross-section was normalised to 1 fb.	113
6.3	Upper limit of $\sigma_{\text{VBF}}(X \rightarrow HH)$ [fb] at 95% CL calculated using the asymptotic approximation and toys for each mass hypothesis. A total of 100 000 pseudo-experiments were generated for each m_X point. The narrow-width samples are used here.	115
7.1	Kinematic variables used as input features for the DNN model.	129
8.1	High-level kinematic variables used as input features for the DNN and transformer-based classifiers.	137
A.1	Kinematic variables used as input for the neural-network based reweighting.	146
B.1	Name conventions for the nuisance parameters.	147