

Paula Martínez Suárez

Search for a spin-zero particle in Higgs boson decays or top-associated production using the ATLAS detector

Ph. D. Thesis

Institut de Física d'Altes Energies (IFAE)

Departament de Física

Facultat de Ciències

Universitat Autònoma de Barcelona (UAB)

Thesis supervisor: Imma Riu Dachs

Thesis tutor: José María Crespo Vicente

11/03/2025

Abstract

This thesis contains two searches for light pseudoscalar particles using 140 fb^{-1} proton-proton (pp) collision data at a centre-of-mass energy of 13 TeV collected by the ATLAS detector at the Large Hadron Collider during Run 2. Both searches focus on multi- b final states, using novel techniques for b -quark identification. Only events with two leptons are considered, in order to reduce contributions from difficult-to-model hadronic background processes. The first search explores light pseudoscalar production in Higgs boson decays, $H \rightarrow aa$, with an a -boson mass between 12 and 60 GeV. Each pseudoscalar decays to two b -quarks, $a \rightarrow b\bar{b}$, which translates into a final state with four b -quarks in total. The Higgs boson is produced in association with a Z boson that decays to either electrons or muons. No significant excess of events above the expected Standard Model background is observed. Upper limits at 95% confidence level are set for the branching ratio of the Higgs boson decaying to two light pseudoscalars, $\text{BR}(H \rightarrow aa \rightarrow 4b)$, between 3% and 25%. The second search targets light pseudoscalar production in association with top quarks, $t\bar{t}a$. This production mechanism allows to explore the pseudoscalar mass range from 12 to 100 GeV. The $t\bar{t}$ pair is required to decay leptonically to electrons, muons or both. The light pseudoscalar is studied via its decay to b -quarks. As no significant excess is observed, upper limits for the production cross section of a light pseudoscalar in association with top quarks, $\sigma(pp \rightarrow t\bar{t}a) \times \text{BR}(a \rightarrow b\bar{b})$, are presented. They range between 0.1 and 0.9 pb.

Agradecimientos

Esta tesis doctoral es el resultado de años de trabajo en colaboración con muchas personas e instituciones sin las cuales no habría sido capaz de llevar a cabo este proyecto.

En primer lugar, gracias al equipo de ATLAS del Institut de Física d'Altes Energies de Barcelona, en particular a Imma Riu y Aurelio Juste, ya que ellos son quienes me dieron la oportunidad de embarcarme en esta nueva etapa de mi vida, y ellos son quienes me han guiado durante los últimos cuatro años. Gracias al equipo de DESY y al de la Universidad de Massachusetts Amherst, con quienes tuve la oportunidad de trabajar en el análisis de $H \rightarrow aa \rightarrow 4b$. Gracias también al equipo de la Universidad Johannes Gutenberg en Mainz, con quienes desarrollamos el trigger de ATLAS para la Run 3.

En segundo lugar, gracias a mis compañeros de trabajo, en especial a Adrián Salvador, Javier Jiménez, Yuan-Tang Chou y Judit Höfer, de los cuales he aprendido casi todo lo que sé hoy en día, y a quienes les deseo lo mejor en sus futuros proyectos.

En tercer lugar, gracias a mis profesores de la asignatura de Modelo Estándar en el máster de Física Teórica del Instituto de Física Teórica de Madrid, María José Herrero y Carlos Pena, por sus excelentes clases en las cuales me he inspirado para redactar el correspondiente capítulo de esta tesis. Gracias también al equipo de CMS de la Universidad de Oviedo, por inspirarme durante la carrera y descubrirme una pasión por la Física de Partículas.

Finalmente, gracias a mis amigos y a mi familia, por el apoyo incondicional que me han dado y, en ocasiones, por los sacrificios que han tenido que hacer para que un día como hoy pueda alcanzar una de las metas más importantes de mi vida.

Contents

Introduction	1
I Theoretical framework	
1 The Standard Model of Particle Physics	5
1.1 The building blocks of the Standard Model	5
1.2 Mathematical formulation	6
1.2.1 Quantum electrodynamics	6
1.2.2 Electroweak theory	7
1.2.3 Quantum chromodynamics	10
1.3 The Brout-Englert-Higgs mechanism	12
1.3.1 Spontaneous symmetry breaking and mass generation . . .	12
1.3.2 Fermion mass generation	15
1.4 Feynman rules of the Standard Model	16
1.4.1 Electroweak interactions	16
1.4.2 Strong interactions	19
1.4.3 Higgs interactions	19
2 Standard Model measurements	21
2.1 Experimental successes of the Standard Model	21
2.1.1 Particles predicted by the Standard Model	21
2.1.2 Cross section measurements at the LHC	26
2.2 Shortcomings	27
2.2.1 Theoretical problems	27
2.2.2 Experimental problems	28
3 Physics beyond the Standard Model	33
3.1 Axion-like particles	33
3.1.1 What is an axion?	33
3.1.2 Generalisation of the axion theory	33
3.2 Two Higgs doublet model	34
3.2.1 General formulation	34
3.2.2 Phenomenology of the two Higgs doublet model	36

3.2.3	Two Higgs doublet model + pseudoscalar	38
3.3	Simplified models for LHC Physics	40
II	Experimental setup	
4	The ATLAS experiment	45
4.1	The Large Hadron Collider	45
4.1.1	LHC layout	45
4.1.2	LHC performance in Run 2	48
4.2	The ATLAS Experiment	49
4.2.1	Coordinate system	50
4.2.2	Inner Detector	51
4.2.3	Calorimeters	53
4.2.4	Muon spectrometer	55
4.3	Trigger and data acquisition	57
4.3.1	Level-1 Trigger	58
4.3.2	High-Level Trigger	62
5	Physics simulation of proton-proton collisions	63
5.1	Event simulation	63
5.1.1	Hard scattering	64
5.1.2	Parton shower	67
5.1.3	Hadronisation and decay	69
5.1.4	Underlying events	71
5.2	Detector simulation	71
6	Object reconstruction	73
6.1	Tracks and vertices	73
6.2	Leptons	74
6.2.1	Electrons	74
6.2.2	Muons	74
6.3	Jets	75
6.3.1	General aspects of jet reconstruction	75
6.3.2	Small- R jets	76
6.3.3	Large- R jets	76
6.4	Flavour tagging	78
6.4.1	b -jets	79
6.4.2	B -jets	81
6.4.3	Soft secondary vertices	82

6.5	Overlap removal	83
6.6	μ -in-jet p_T correction	83
6.7	Missing transverse energy	84
 III Analysis tools		
7	Machine learning	87
7.1	Neural networks	87
7.1.1	General definition	87
7.1.2	Training methods	88
7.1.3	Performance	90
7.1.4	Parametrisation	91
7.2	Boosted decision trees	93
8	Statistical methods	95
8.1	Hypothesis testing	95
8.2	Statistical tests	96
8.2.1	The binned profile likelihood method	96
8.2.2	p -value and confidence intervals	98
8.3	Software implementation	99
8.3.1	The TRExFITTER package	99
8.3.2	Pruning, smoothing and symmetrisation	100
 IV Search for decays of the Higgs boson into pseudoscalar particles decaying to four bottom quarks using proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector		
9	Analysis overview	103
9.1	Introduction	103
9.2	Data and simulated samples	106
9.2.1	Data samples	106
9.2.2	Signal Monte Carlo samples	107
9.2.3	Background Monte Carlo samples	107
9.3	Event selection	110
9.3.1	Trigger	110
9.3.2	Preselection	111
9.4	Background modelling	112

9.4.1	$t\bar{t}$ +jets reweighting	113
9.4.2	Z +jets reweighting	122
9.5	Analysis strategy	128
9.5.1	Event reconstruction for $H \rightarrow aa \rightarrow 4b$	128
9.5.2	BDT for signal versus background discrimination	132
9.5.3	Signal and control regions	136
9.6	Systematic uncertainties	140
9.6.1	Experimental uncertainties	140
9.6.2	Modelling uncertainties	143
10	Results and conclusions	147
10.1	Fit setup	147
10.1.1	General considerations	147
10.1.2	Normalisation factors for Z +jets and $t\bar{t}$ +jets	147
10.1.3	Binning, smoothing and pruning	148
10.2	Fit results	148
10.2.1	B-only fit to the Asimov dataset	149
10.2.2	S+B fit to data	155
10.2.3	Expected and observed limits	161
10.3	Conclusions	163
V	Search for a new pseudoscalar decaying into a pair of bottom and anti-bottom quarks in top-associated production using proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector	
11	Analysis overview	167
11.1	Introduction	167
11.2	Data and simulated samples	170
11.2.1	Data samples	170
11.2.2	Signal Monte Carlo samples	170
11.2.3	Background Monte Carlo samples	171
11.3	Event selection	173
11.3.1	Trigger	173
11.3.2	Preselection	173
11.4	Background modelling	174
11.4.1	$t\bar{t}$ +jets reweighting	175
11.5	Analysis strategy	182
11.5.1	BDT for event reconstruction	182
11.5.2	NN for signal versus background discrimination	188

11.6	Systematic uncertainties	197
11.6.1	Experimental uncertainties	198
11.6.2	Modelling uncertainties	200
12	Results and conclusions	207
12.1	Fit setup	207
12.1.1	General considerations	207
12.1.2	Signal and control regions	207
12.1.3	Normalisation factors for $t\bar{t}$ +jets	208
12.1.4	Binning, smoothing and pruning	210
12.2	Fit results	211
12.2.1	B-only fit to the Asimov dataset	211
12.2.2	S+B fit to data	220
12.2.3	Expected and observed limits	227
12.3	Conclusions	228
	Summary	229
	Appendices	231
A	Phase-I upgrades for the Level-1 trigger	231
A.1	Level-1 calorimeter trigger	232
A.2	Level-1 muon trigger	233
A.3	Level-1 topological trigger	234
B	Blinded fits to data	239
B.1	$H \rightarrow aa \rightarrow 4b$ analysis	239
B.2	$t\bar{t}a, a \rightarrow b\bar{b}$ analysis	243
C	Injection tests for the $t\bar{t}a, a \rightarrow b\bar{b}$ analysis	249
	Bibliography	255
	List of Figures	271
	List of Tables	281

Introduction

The discovery of the Higgs boson in 2012 by the ATLAS and CMS experiments at CERN [1, 2] was one of the most important achievements in the field of Particle Physics in the last decades, constituting the last missing piece of the Standard Model. Since then, an extensive programme has been carried out by both experiments in order to characterise the properties of this particle, which are free parameters of the theory and therefore need to be determined experimentally. The Higgs boson is not only relevant for Standard Model studies, but also for searches of new physics beyond the Standard Model. The upper limit on the branching ratio of the Higgs boson into undetected particles is constrained by global analyses of Higgs boson properties to approximately 12% [3, 4]. Higgs boson decays are particularly sensitive to new physics due to the small total decay width, of about 4 MeV. Even very small couplings to new particles can give sizable branching ratios compatible with available measurements.

This work focuses on extensions of the Standard Model that include new light pseudoscalars, sometimes referred to as axion-like particles, or simply a -bosons. They appear in theories with extended Higgs sectors, such as the next-to-minimal supersymmetric Standard Model [5–9], dark matter models [10–14], models with a first-order electroweak phase transition [15, 16] and theories of neutral naturalness [17, 18]. In general, models featuring axion-like particles have the flexibility to accommodate new physics in the electroweak scale, which is the energy scale accessible at the Large Hadron Collider. Two searches for axion-like particles are described in this thesis. One of them investigates the Standard Model Higgs boson decay into a pair of light pseudoscalars, $H \rightarrow aa$, and the other, direct production of a light pseudoscalar in association with top quarks, $t\bar{t}a$. In many existing theories, the a -boson inherits Yukawa couplings to fermions, which are directly proportional to the fermion mass. Under that assumption, the $a \rightarrow b\bar{b}$ decay provides the largest branching ratio, reason why it is used in both analyses.

This thesis is organised as follows. Part I describes the theoretical framework and is divided in three chapters. Chapter 1 introduces the Standard Model of Particle Physics and its mathematical formulation, Chapter 2 describes its experimental successes and shortcomings and Chapter 3 proposes some well-known theoretical extensions to the Standard Model that are explored at the Large Hadron Collider. Part II focuses on the experimental setup that is used to obtain the data to perform the searches. Chapter 4 provides an overview of the Large Hadron Collider and the ATLAS Experiment. Chapters 5 and 6 describe the Monte Carlo simulation of proton-proton collisions and the object reconstruction algorithms, respectively. Part III introduces the analysis tools used in the $H \rightarrow aa$ and $t\bar{t}a$, $a \rightarrow b\bar{b}$ searches, including the machine learning framework in Chapter 7 and statistical methods in Chapter 8. Parts IV and V contain a detailed description of the $H \rightarrow aa \rightarrow 4b$ and $t\bar{t}a$, $a \rightarrow b\bar{b}$ searches, respectively. They are structured in two chapters each, one providing the analysis overview and the

other providing the final results and conclusions. A final summary is included at the end. Additionally, Appendix A offers an overview of the Run 3 updates to the ATLAS Level-1 trigger system, and Appendices B and C describe additional tests performed during the development of the $H \rightarrow aa \rightarrow 4b$ and the $t\bar{t}a, a \rightarrow b\bar{b}$ searches.

Part I

Theoretical framework

Chapter 1

The Standard Model of Particle Physics

The Standard Model of Particle Physics (SM) is a quantum field theory that describes Physics at the subatomic level. It was conceived during the second half of the 20th century, taking its final form during the 1970s. Since then, it has proven to be a solid theory of Particle Physics, correctly predicting, for example, the existence of the $W^\pm$ and Z bosons, carriers of the weak force, the top quark, the only quark that does not form hadrons, or the Higgs boson, a fundamental piece in electroweak theory. The following sections provide an overview of the Standard Model, describing its building blocks and mathematical formulation.

1.1 The building blocks of the Standard Model

The Standard Model is a quantum field theory, meaning that particles are excitations of quantum fields defined in all points of spacetime. It is built from the $SU(3)_C \times SU(2)_L \times U(1)_Y$ gauge symmetry group, where $SU(3)_C$ is the symmetry governing the strong interactions and $SU(2)_L \times U(1)_Y$ corresponds to the electroweak sector. There are three generations of fermions, which are the constituents of matter, and have spin 1/2. These fermions can be either quarks or leptons, depending on whether they interact via the strong force or not. Interactions are governed by four spin-1 gauge bosons resulting from the corresponding gauge symmetries. The photon (γ), Z and $W^\pm$ bosons are responsible for the electroweak force, while gluons (g) are the carriers of the strong force. The Higgs boson is the only spin-0 particle, and it interacts with all the massive particles in the model. Figure 1.1 summarises the particle content of the SM and its fundamental interactions.




Figure 1.1: (a) Particle content of the Standard Model and (b) fundamental interactions between Standard Model particles.

1.2 Mathematical formulation

Quantum field theory (QFT) is a theoretical framework that combines quantum mechanics and special relativity. In the case of the SM, it provides a description of Particle Physics in the four-dimensional Minkowski spacetime, which requires the SM lagrangian to be local and invariant under Poincaré transformations. QFT is usually studied as a perturbation theory in which particle interactions are represented as small perturbations with respect to the free theory. This is done by expanding the interaction potential as a power series of the coupling constant. The SM is also required to be renormalisable, meaning that it can not contain terms of dimension larger than four in the fields and their derivatives. Renormalisation ensures that SM observables are finite after taking into account physical effects due to higher orders in perturbation theory, which lead to scale-dependent couplings, also known as running couplings.

1.2.1 Quantum electrodynamics

The mathematical formulation of the SM begins with the description of quarks and leptons as free particles. Free fermion fields propagate according to the Dirac lagrangian:

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

where ψ refers to the fermion field and $\bar{\psi} = \psi^\dagger \gamma^0$ is its Dirac conjugate. γ^μ with $\mu = 1, 2, 3, 4$ are the four Dirac matrices and m is the mass of the fermion. This lagrangian is invariant under charge conjugation (C), parity (P) and time reversal (T) transformations, but when exploring more complex scenarios, CPT is the only combination of C , P , and T that is observed to be an exact symmetry of nature at the fundamental level.

The Dirac lagrangian has a $U(1)_{\text{EM}}$ global symmetry, meaning that it is invariant under phase transformations such as $\psi \rightarrow e^{ieQ\theta}\psi$, where Q is the electrical charge operator, measured in units of the electron charge e , and θ is a global free parameter. According to Nöther's theorem, this symmetry can be associated to one conserved quantity, the electrical charge, which is the eigenvalue of the operator Q . This rule also applies to charge conjugated particles, which have the same eigenvalue under Q with opposite sign. They are commonly known as antiparticles. Local interactions of charged fermion fields are introduced by promoting the $U(1)_{\text{EM}}$ global symmetry to a $U(1)_{\text{EM}}$ local (gauge) symmetry $\psi \rightarrow e^{ieQ\theta(x)}\psi$. This ensures that the physical laws governing electromagnetic interactions remain invariant under arbitrary local phase transformations. In order to make the lagrangian gauge-invariant, one gauge field must be introduced for each generator of the symmetry group. In the case of $U(1)_{\text{EM}}$, there is only one generator, and therefore, one gauge field. This gauge field, denoted as A_μ , is a vector field, corresponding to a spin-1 particle, which transforms under the local $U(1)_{\text{EM}}$ group in a way that ensures that the overall physical laws remain invariant:

$$\begin{aligned} \psi(x) &\rightarrow e^{ieQ\theta(x)}\psi(x), \\ A_\mu(x) &\rightarrow A_\mu(x) - \partial_\mu\theta(x). \end{aligned} \quad (1.2)$$

This results in the following lagrangian:

$$\mathcal{L} = \bar{\psi}(i\mathcal{D} - m)\psi, \text{ with } \mathcal{D} = \gamma^\mu D_\mu = \gamma^\mu(\partial_\mu + ieQA_\mu). \quad (1.3)$$

D_μ is known as the covariant derivative. The dynamics of this new spin-1 particle, the photon, can be described using a gauge-invariant field strength tensor:

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu. \quad (1.4)$$

The electromagnetic field tensor, $F_{\mu\nu}$ is nothing else than the QFT formulation of Maxwell's equations of electromagnetism. Introducing this kinetic term results in the lagrangian for Quantum Electrodynamics (QED):

$$\mathcal{L}_{\text{QED}} = \bar{\psi}(i\mathcal{D} - m)\psi - \frac{1}{4}F_{\mu\nu}F^{\mu\nu} \quad (1.5)$$


QED has one coupling constant, e , which can also be written as $\alpha_{\text{EM}} = e^2/4\pi$. This corresponds to the so-called bare coupling, which does not account for quantum fluctuations, which contribute to the interaction strength at different energy scales. After including these effects through renormalisation, one obtains the true coupling, which is not constant anymore:

$$\alpha_{\text{EM}}(Q^2) = \frac{\alpha_{\text{EM}}(\mu^2)}{1 - \frac{\alpha_{\text{EM}}(\mu^2)}{3\pi} \log\left(\frac{Q^2}{\mu^2}\right)}, \quad (1.6)$$

where Q^2 refers to the energy scale of the interaction and μ^2 is the reference momentum scale from which $\alpha_{\text{EM}}(Q^2)$ is calculated. The running coupling constant $\alpha_{\text{EM}}(Q^2)$ describes how the effective charge depends on the separation of the two charged particles. In the case of QED, it decreases at lower energy (larger distance). This is a feature of abelian QFTs.

1.2.2 Electroweak theory

At the time of its conception in 1927, QED was very successful in describing photon-mediated interactions between electrically charged particles at the quantum relativistic level, but it was not able to explain experimental phenomena such as the neutron beta decay ($n \rightarrow p + e + \bar{\nu}_e$), the muon decay ($\mu \rightarrow e + \bar{\nu}_e + \nu_\mu$) or the electron-neutrino scattering ($e + \nu \rightarrow e + \nu$). In 1932, Enrico Fermi postulated his theory of beta decay as a four-point interaction with a coupling G_F [19]:



$$= G_F(\bar{\nu}_e \gamma^\mu p)(\bar{\nu}_e \gamma_\mu e), \quad (1.7)$$


with $G_F \approx 1.166 \times 10^{-5} \text{GeV}^{-2}$. Later, in 1956, an experiment by Chien-Shiung Wu [20] established that conservation of parity was violated by the weak interaction in beta

1. The Standard Model of Particle Physics

decay. This meant that the interaction was different depending on whether a fermion was left-handed (spin aligned with the direction of motion) or right-handed (spin anti-aligned with the direction of motion). For a Dirac fermion, the two chiral projections of a massless particle can be defined using the γ^5 matrix as:

$$\psi = \psi_L + \psi_R \text{ where } \begin{cases} \psi_L = P_L \psi = \frac{1}{2}(1 - \gamma^5)\psi & \text{(left-handed)} \\ \psi_R = P_R \psi = \frac{1}{2}(1 + \gamma^5)\psi & \text{(right-handed)} \end{cases}. \quad (1.8)$$

This was introduced in the originally proposed Fermi theory as follows:



$$= G_F [\bar{n} \gamma^\mu (1 - \gamma^5) p] [\bar{v}_e \gamma_\mu (1 - \gamma^5) e], \quad (1.9)$$

implying that weak interactions are only sensitive to the left-handed projection. All this motivated the development of the electroweak sector of the Standard Model, governed by the $SU(2)_L \times U(1)_Y$ gauge symmetry, years later.

The $SU(2)_L$ symmetry group implies that the electroweak force is chiral. Left-handed spinors transform as doublets of $SU(2)_L$, while right handed spinors remain invariant, and are represented as scalars:

$$\begin{aligned} \text{Left-handed fermions: } \ell_L &= \begin{pmatrix} \nu_{eL} \\ e_L \end{pmatrix}, \begin{pmatrix} \nu_{\mu L} \\ \mu_L \end{pmatrix}, \begin{pmatrix} \nu_{\tau L} \\ \tau_L \end{pmatrix} \\ q_L &= \begin{pmatrix} u_L \\ d_L \end{pmatrix}, \begin{pmatrix} c_L \\ s_L \end{pmatrix}, \begin{pmatrix} t_L \\ b_L \end{pmatrix} \end{aligned} \quad (1.10)$$

$$\begin{aligned} \text{Right-handed fermions: } \ell_R &= e_R, \mu_R, \tau_R, \nu_{eR}, \nu_{\mu R}, \nu_{\tau R} \\ q_R &= u_R, c_R, t_R, d_R, s_R, b_R \end{aligned}$$

Therefore, under a $SU(2)_L$ rotation:

$$\begin{aligned} \psi_L &\rightarrow e^{i\mathbf{T} \cdot \boldsymbol{\theta}} \psi_L, \text{ with } \mathbf{T} = \frac{\boldsymbol{\sigma}}{2}, \\ \psi_R &\rightarrow \psi_R, \end{aligned} \quad (1.11)$$

where $\mathbf{T}$ is the weak isospin vector, which contains the three generators of the $SU(2)_L$ symmetry group, proportional to the Pauli matrices, $\boldsymbol{\sigma}$. According to Nöther's theorem, if $SU(2)_L$ is a global symmetry, the eigenvalues corresponding to the three components of $\mathbf{T}$ are conserved quantities. However, the $SU(2)_L$ group is non-abelian and the three generators do not commute:

$$[T_i, T_j] = i\epsilon_{ijk} T_k, \quad (1.12)$$

where ϵ_{ijk} is the Levi-Civita symbol. At the quantum level, this implies that the three separate eigenvalues can not be measured simultaneously. Instead, the quantum numbers of choice are the total weak isospin T and the projection along the z -axis, T_3 :

$$[T^2, T_3] = 0, \text{ with } T^2 = T_1^2 + T_2^2 + T_3^2. \quad (1.13)$$

The $U(1)_Y$ symmetry group arises from the non-commutation of the generators of $SU(2)_L$ and $U(1)_{EM}$, which can be seen in the fact that particles belonging to the same $SU(2)_L$ doublet can have different values of the electric charge. Instead, the theory is completed by introducing a new quantum operator, called weak hypercharge, $Y_W/2$, such that:

$$[T_i, Y_W] = 0 \text{ for all } i = 1, 2, 3. \quad (1.14)$$

Then, electric charge is defined as:

$$Q = T_3 + \frac{Y_W}{2}. \quad (1.15)$$

Following a similar procedure to QED, the $SU(2)_L \times U(1)_Y$ symmetry group can be gauged by defining a new covariant derivative:

$$\mathcal{L} = i\bar{\psi}\not{D}\psi, \text{ with } \not{D} = \gamma^\mu \left(\partial_\mu - ig\mathbf{T} \cdot \mathbf{W}_\mu - ig' \frac{Y_W}{2} B_\mu \right). \quad (1.16)$$

In this equation, g and g' are the coupling constants of the $SU(2)_L$ and $U(1)_Y$ groups, respectively, and $\mathbf{W}_\mu$ and B_μ are the boson fields that mediate the interactions associated to each of the two symmetry groups. It is important to note that mass terms $m\bar{\psi}\psi$ do not respect the $SU(2)_L$ symmetry, reason why they must be removed from the lagrangian. Kinetic terms for the $\mathbf{W}_\mu$ and B_μ bosons can be written using a field strength tensor, similarly to the photon field, with the only difference that the $SU(2)_L$ is non-abelian, providing a self-interacting term for the $\mathbf{W}$ boson:

$$\begin{aligned} W_{\mu\nu}^i &= \partial_\mu W_\nu^i - \partial_\nu W_\mu^i + g\epsilon^{ijk} W_\mu^j W_\nu^k, \text{ with } i, j, k = 1, 2, 3, \\ B_{\mu\nu} &= \partial_\mu B_\nu - \partial_\nu B_\mu. \end{aligned} \quad (1.17)$$

In summary, the electroweak lagrangian can be written as follows:

$$\mathcal{L}_{EW} = i\bar{\psi}\not{D}\psi - \frac{1}{4}\mathbf{W}_{\mu\nu} \cdot \mathbf{W}^{\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu}, \quad (1.18)$$

where ψ represents the different types of fermions, whose quantum numbers are shown in Tables 1.1. The electroweak lagrangian has an additional $U(1)$ accidental global symmetry under transformations of the type $\psi \rightarrow e^{iL\theta}\psi$. Here L is the lepton number operator, whose eigenvalues are 1 for leptons, -1 for antileptons and 0 for quarks and bosons.

Electroweak theory has two coupling constants: $\alpha_W = g^2/4\pi$, associated to the $SU(L)_L$ group and $\alpha_Y = g'^2/4\pi$, associated to the $U(1)_Y$ group. After renormalisation, they acquire a scale dependency:

$$\begin{aligned} \alpha_W(Q^2) &= \frac{\alpha_W(\mu^2)}{1 + \frac{41\alpha_Y(\mu^2)}{20\pi} \log\left(\frac{Q^2}{\mu^2}\right)}, \\ \alpha_Y(Q^2) &= \frac{\alpha_Y(\mu^2)}{1 - \frac{\alpha_W(\mu^2)}{2\pi} \log\left(\frac{Q^2}{\mu^2}\right)}. \end{aligned} \quad (1.19)$$

1. The Standard Model of Particle Physics

Lepton	Q	T	T_3	Y_W
ℓ_L	-1	1/2	-1/2	-1
$\nu_{\ell L}$	0	1/2	1/2	-1
ℓ_R	-1	0	0	-2
$\nu_{\ell R}$	0	0	0	0

Quark	Q	T	T_3	Y_W
u_L	2/3	1/2	1/2	1/3
d_L	-1/3	1/2	-1/2	1/3
u_R	2/3	0	0	4/3
d_R	-1/3	0	0	-2/3

Table 1.1: Electroweak charges associated to each fermion of the SM. For simplicity, $\ell = e, \mu, \tau$, $u = u, c, t$ and $d = d, s, b$. Antiparticles have the opposite electrical charge, which leads to the sign of T_3 and Y_W being flipped as well. Because right-handed neutrinos are not charged under any symmetry, they are usually excluded from the model.

The coupling strength $\alpha_W(Q^2)$ decreases with energy. This implies that particles that have low energy or are very separated interact stronger. At very low energies, when the coupling strength becomes of $O(1)$, perturbative theory is not valid anymore and particles are said to be confined. This is not relevant in the case of weak interactions because the confinement scale is well below the energy scale of the $SU(2)_L \times U(1)_Y \rightarrow U(1)_{EM}$ spontaneous symmetry breaking that generates mass terms for the $W^\pm$ and Z bosons (Section 1.3). The coupling strength $\alpha_Y(Q^2)$ behaves similarly to the QED coupling strength, $\alpha_{EM}(Q^2)$.

1.2.3 Quantum chromodynamics

The third and last gauge symmetry group of the SM is $SU(3)_C$, with C referring to the colour charge. This is the group that defines the theory of strong interactions, also known as Quantum Chromodynamics (QCD). Leptons do not have colour charge, therefore they do not participate in the strong interaction. Quarks, on the other hand, have three possible colour charges, usually denoted as R , G and B ¹. They can be represented in the colour basis as:

$$\psi^R = \begin{pmatrix} 1 \\ 0 \\ 0 \end{pmatrix}, \quad \psi^G = \begin{pmatrix} 0 \\ 1 \\ 0 \end{pmatrix}, \quad \psi^B = \begin{pmatrix} 0 \\ 0 \\ 1 \end{pmatrix}. \quad (1.20)$$

Under a colour rotation, they transform as follows:

$$\psi^A \rightarrow \Omega^{AB} \psi^B, \quad \text{where } \Omega^{AB} = e^{iT^{AB} \cdot \theta} \quad \text{with } T^{AB} = \frac{\lambda^{AB}}{2}, \quad (1.21)$$

where the colour indices $A, B = R, G, B$ are made explicit. The vector operator $\mathbf{T}$ is used to denote the 8 generators of $SU(3)_C$, which are proportional to the Gell-Mann matrices λ . With the $SU(3)_C$ global symmetry, there is an associated conservation law for the eigenvalues of the generators of the group. Similarly to $SU(2)_L$, $SU(3)_C$ is non-abelian, which means that not all of the generators commute, implying that they can not be measured simultaneously:

$$[T_a^{AB}, T_b^{AB}] = if_{abc} T_c^{AB}, \quad (1.22)$$

¹ Antiquarks are charged with the corresponding anticolours, usually represented as $\bar{R}$, $\bar{G}$ and $\bar{B}$.

where a, b, c run over the indices of the Gell-Mann matrices and f_{abc} are the structure constants analogous to the Levi-Civita symbol ϵ_{jkl} of $SU(2)_L$. Despite the non-abelian nature of the $SU(3)_C$ symmetry group, it is possible to find a set of conserved quantum numbers for quarks that can be measured simultaneously, corresponding to the T_3 and T_8 components.

Promoting the $SU(3)_C$ symmetry to a gauge symmetry yields the following lagrangian for coloured fermions:

$$\mathcal{L}_{\text{QCD}} = i\bar{\psi}^A \not{D}^{AB} \psi^B, \text{ with } \not{D}^{AB} = \gamma^\mu \left(\partial_\mu \delta^{AB} - ig_s \mathbf{T}^{AB} \cdot \mathbf{G}_\mu \right). \quad (1.23)$$

In this equation, $\mathbf{G}_\mu$ is the gluon field and g_s is the strong coupling constant. The field strength tensor governing the kinematics of the gluon field can be written as:

$$G_{\mu\nu}^a = \partial_\mu G_\nu^a - \partial_\nu G_\mu^a + g_s f^{abc} G_\mu^b G_\nu^c, \text{ with } a, b, c = 1, \dots, 8, \quad (1.24)$$

leading to the following QCD lagrangian:

$$\mathcal{L}_{\text{QCD}} = i\bar{\psi} \not{D} \psi - \frac{1}{4} G_{\mu\nu}^a G_a^{\mu\nu}. \quad (1.25)$$

Explicit colour indices are not shown for simplicity. Fermion mass terms are not included in order to preserve the $SU(2)_L$ symmetry.

For historical reasons, conserved quantities in QCD are usually defined from its global symmetries and not from its gauge colour symmetry. First, QCD has an accidental $U(1)$ global symmetry under transformations of the type $\psi \rightarrow e^{iB\theta} \psi$, where B is the baryon number operator. This global symmetry leads to the conservation of the baryon number, which is 1 for baryons (made of three quarks, qqq), -1 for antibaryons (made of three antiquarks, $\bar{q}\bar{q}\bar{q}$) and 0 for mesons (made of a $q\bar{q}$ pair). Additionally, QCD has a flavour symmetry $SU(N_F)$ that arises from the fact that the strong interaction is unable to distinguish quark flavours if they are massless. In reality, the flavour symmetry is broken by c -, b - and t -quarks, but it can be taken as an approximate symmetry for u , d and s , whose masses are closer to 0 and way below the QCD scale $\Lambda_{\text{QCD}} \sim 1$ GeV. This approximate symmetry is the origin of light mesons such as pions or kaons, which are, in fact, nothing else than pseudo-Nambu-Goldstone bosons from the approximate $SU(3)$ QCD flavour symmetry. Quantum numbers associated to the $SU(3)$ flavour symmetry are isospin (I_3), hypercharge (Y) and strangeness (S). However, they are only approximately conserved by strong interactions, and not conserved at all by weak interactions. They relate to each other, as well as to the baryon number and electric charge, as follows:

$$Y = B + S = 2(Q - I_3). \quad (1.26)$$

Table 1.2 summarises the quantum numbers associated to the baryon number symmetry and the QCD flavour symmetry.

The scale-dependent coupling strength for QCD can be defined as:

$$\alpha_s(Q^2) = \frac{\alpha_s(\mu^2)}{1 + \frac{\alpha_s(\mu^2)}{2\pi} \left(11 - \frac{2}{3}N_f \right) \log \left(\frac{Q^2}{\mu^2} \right)}, \quad (1.27)$$

1. The Standard Model of Particle Physics

Quarks	Q	B	I_3	Y	S
u	2/3	1/3	1/2	1/3	0
d	-1/3	1/3	-1/2	1/3	0
s	-1/3	1/3	0	-2/3	-1
c	2/3	1/3			
b	-1/3	1/3			
t	-1/3	1/3			

Table 1.2: QCD charges associated to light and heavy quarks. For antiparticles, all charges have the same value and opposite sign. Quantum numbers I_3 , Y and S are not shown for heavy quarks as the approximate QCD flavour symmetry applies to light quarks only.

where N_f is the number of quark flavours at the scale Q^2 . Similarly to $\alpha_w(Q^2)$, the strong coupling $\alpha_s(Q^2)$ also decreases with energy. For QCD, the confinement scale is set at $\Lambda_{\text{QCD}} \sim 1 \text{ GeV}$, that is, around the proton mass, such that $\alpha_s(\Lambda_{\text{QCD}}) \sim 1$. At this point, QCD becomes non-perturbative and the large coupling strength forces all coloured particles to combine in order to form colour-neutral states (hadrons).

1.3 The Brout-Englert-Higgs mechanism

As it was originally formulated, the Standard Model had one big flaw: there was no way of writing mass terms for the fermions and gauge bosons without breaking the electroweak symmetry. The Brout-Englert-Higgs mechanism was introduced in 1964 as a solution to this problem inspired by the Bardeen-Cooper-Schrieffer model of superconductivity [21, 22] and the Nambu-Goldstone theorem [23–25]: if a quantum field theory has a global continuous symmetry that is not a symmetry of the vacuum, there is a massless spin-0 boson associated to each generator of the symmetry group that does not annihilate the vacuum. When the broken symmetry is a gauge symmetry, the Nambu-Goldstone bosons no longer remain massless particles. Instead, they are absorbed by the gauge bosons associated with the broken symmetry, which gives them a mass term.

1.3.1 Spontaneous symmetry breaking and mass generation

Experimentally, it was known that the weak force was short-ranged, but the electromagnetic force was not. This suggested that the mediators of the weak force had to be massive gauge bosons, while the mediator of the electromagnetic force had to be massless.

The way of achieving this is to break the $SU(2)_L \times U(1)_Y$ symmetry down to $U(1)_{\text{EM}}$ at a certain energy scale by introducing a new field with non-vanishing quantum numbers with respect to the symmetries that need to be broken:

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix}, \quad (1.28)$$

where Φ is a complex scalar doublet of $SU(2)_L$ with $T = 1/2$ and $Y = 1$, and four degrees of freedom. The kinetic term for a complex scalar field is written using the covariant derivative as:

$$\mathcal{L} = (D_\mu \Phi)^\dagger (D^\mu \Phi), \text{ with } D_\mu = \left(\partial_\mu - ig\mathbf{T} \cdot \mathbf{W}_\mu - ig' \frac{Y_W}{2} B_\mu \right). \quad (1.29)$$

By adding the so-called Higgs potential $V(\Phi)$, one obtains the lagrangian for spontaneous symmetry breaking (SSB) of the electroweak symmetry:

$$\begin{aligned} \mathcal{L}_{\text{SSB}} &= (D_\mu \Phi)^\dagger (D^\mu \Phi) - V(\Phi) \\ &= (D_\mu \Phi)^\dagger (D^\mu \Phi) - [\mu^2 (\Phi^\dagger \Phi) + \lambda (\Phi^\dagger \Phi)^2], \end{aligned} \quad (1.30)$$

with $\lambda > 0$ in order to have any stable minima. Figure 1.2 shows the shape of $V(\Phi)$ for two different situations: if $\mu^2 \geq 0$, the minimum occurs at $\langle \Phi \rangle = 0$. However, if $\mu^2 < 0$, then $\langle \Phi \rangle = 0$ is a local maximum. In that case, the ground state is degenerate over a circle with $|\Phi|^2 = v^2$, where $v = -\mu^2/\lambda$ is the vacuum expectation value of Φ . Since $V(\Phi)$ does not depend on $\theta = \text{Arg}(\Phi)$, one can choose, without loss of generality, the ground state to be:

$$\langle \Phi \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix}. \quad (1.31)$$

Due to the conservation of charge imposed by the $U(1)_{\text{EM}}$ symmetry, only the neutral component of the Φ field, ϕ^0 , acquires a vacuum expectation value.




Figure 1.2: Higgs potential $V(\Phi)$ for $\mu^2 \geq 0$ (left) and $\mu^2 < 0$ (right).

Small fluctuations around the vacuum are parametrised as:

$$\Phi(x) = \frac{1}{\sqrt{2}} e^{i\mathbf{T} \cdot \boldsymbol{\xi}(x)/v} \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix} \approx \frac{1}{\sqrt{2}} \begin{pmatrix} \frac{1}{2}(\xi_2 + i\xi_3) \\ v + H - \frac{1}{2}i\xi_3 \end{pmatrix}, \quad (1.32)$$

where $\boldsymbol{\xi}$ are the three Nambu-Goldstone bosons associated to each of the three broken generators of $SU(2)_L$ and $H(x)$ is the remaining degree of freedom. Given that the SM lagrangian is invariant under local $SU(2)_L$ transformations, the exponential term in Equation 1.32 can be removed by applying a $SU(2)_L$ rotation. This transformation

1. The Standard Model of Particle Physics

provides the unitary gauge representation, in which the degrees of freedom corresponding to the Nambu-Goldstone bosons vanish:

$$\Phi(x) = \frac{1}{\sqrt{2}} e^{i\mathbf{T} \cdot \boldsymbol{\xi}(x)/v} \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix} \xrightarrow[\text{gauge}]{\text{unitary}} \Phi(x) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + H(x) \end{pmatrix}. \quad (1.33)$$

Replacing this expression in the derivatives from Equation 1.30 yields:

$$\begin{aligned} D_\mu \Phi &= \frac{1}{\sqrt{2}} \left[\partial_\mu \begin{pmatrix} 0 \\ v + H \end{pmatrix} - ig\mathbf{T} \cdot \mathbf{W}_\mu \begin{pmatrix} 0 \\ v + H \end{pmatrix} - ig' \frac{Y}{2} B_\mu \begin{pmatrix} 0 \\ v + H \end{pmatrix} \right] \\ &= \frac{1}{\sqrt{2}} \left[\partial_\mu \begin{pmatrix} 0 \\ H \end{pmatrix} - i \frac{g}{2} \begin{pmatrix} W_\mu^3 & W_\mu^1 - iW_\mu^2 \\ W_\mu^1 + iW_\mu^2 & -W_\mu^3 \end{pmatrix} \begin{pmatrix} 0 \\ v + H \end{pmatrix} - i \frac{g'}{2} B_\mu \begin{pmatrix} 0 \\ v + H \end{pmatrix} \right] \\ &= \frac{1}{\sqrt{2}} \left[\partial_\mu \begin{pmatrix} 0 \\ H \end{pmatrix} - i \frac{g}{2} \begin{pmatrix} W_\mu^1 - iW_\mu^2 \\ -W_\mu^3 \end{pmatrix} (v + H) - i \frac{g'}{2} \begin{pmatrix} 0 \\ B_\mu \end{pmatrix} (v + H) \right], \end{aligned} \quad (1.34)$$

which results in the following kinetic term:

$$\begin{aligned} (D_\mu \Phi)^\dagger (D^\mu \Phi) &= \frac{1}{2} (\partial_\mu H) (\partial^\mu H) + \frac{g^2}{8} (W_\mu^1 + iW_\mu^2) (W_\mu^1 - iW_\mu^2) (v + H)^2 \\ &\quad + \frac{1}{2} \left[\frac{g^2}{4} W_\mu^3 W_\mu^3 - \frac{gg'}{2} B_\mu W_\mu^3 + \frac{g'^2}{4} B_\mu B_\mu \right] (v + H)^2 \\ &= \frac{1}{2} (\partial_\mu H) (\partial^\mu H) + \frac{g^2}{8} (W_\mu^1 + iW_\mu^2) (W_\mu^1 - iW_\mu^2) (v + H)^2 \\ &\quad + \frac{1}{8} (gW_\mu^3 - g'B_\mu) (gW_\mu^3 + g'B_\mu) (v + H)^2. \end{aligned} \quad (1.35)$$

This string can be further simplified by switching to the mass eigenstates of the vector bosons:

$$\begin{aligned} W_\mu^\pm &= \frac{1}{\sqrt{2}} (W_\mu^1 \mp iW_\mu^2), \\ A_\mu &= B_\mu \cos \theta_W + W_\mu^3 \sin \theta_W, \\ Z_\mu &= -B_\mu \sin \theta_W + W_\mu^3 \cos \theta_W, \end{aligned} \quad (1.36)$$

where θ_W is the weak mixing angle or Weinberg angle, defined as:

$$g \sin \theta_W = g' \cos \theta_W = e, \quad (1.37)$$

resulting in:

$$\begin{aligned} (D_\mu \Phi)^\dagger (D^\mu \Phi) &= \frac{1}{2} (\partial_\mu H) (\partial^\mu H) + \left(\frac{g^2}{4} W_\mu^- W_\mu^+ + \frac{1}{2} \frac{(g^2 + g'^2)}{4} Z_\mu Z^\mu \right) (v + H)^2 \\ &= \frac{1}{2} (\partial_\mu H) (\partial^\mu H) \\ &\quad + \frac{g^2 v^2}{4} \left(W_\mu^- W_\mu^+ + \frac{1}{v} W_\mu^- W_\mu^+ H + \frac{1}{v^2} W_\mu^- W_\mu^+ H^2 \right) \\ &\quad + \frac{1}{2} \frac{(g^2 + g'^2) v^2}{4} \left(Z_\mu Z^\mu + \frac{1}{v} Z_\mu Z^\mu H + \frac{1}{v^2} Z_\mu Z^\mu H^2 \right). \end{aligned} \quad (1.38)$$

This expression provides mass terms for the W and Z gauge bosons, as well as a description of their interaction with the Higgs boson H . Defining m_W and m_Z as:

$$\begin{aligned} m_W &= gv/2, \\ m_Z &= \sqrt{g^2 + g'^2}v/2, \end{aligned} \quad (1.39)$$

leads to the following result:

$$\begin{aligned} (D_\mu \Phi)^\dagger (D^\mu \Phi) &= \frac{1}{2}(\partial_\mu H)(\partial^\mu H) \\ &+ m_W^2 \left(W_\mu^- W_+^\mu + \frac{1}{v} W_\mu^- W_+^\mu H + \frac{1}{v^2} W_\mu^- W_+^\mu H^2 \right) \\ &+ \frac{1}{2} m_Z^2 \left(Z_\mu Z^\mu + \frac{1}{v} Z_\mu Z^\mu H + \frac{1}{v^2} Z_\mu Z^\mu H^2 \right). \end{aligned} \quad (1.40)$$

This result can be related with the Fermi theory that inspired the formulation of the electroweak interaction by rewriting the Fermi constant as:

$$\frac{G_F}{\sqrt{2}} = \frac{g^2}{8m_W^2} = \frac{1}{\sqrt{2}v^2}. \quad (1.41)$$

The potential term can be developed in a similar manner to obtain the self-interaction terms for the Higgs boson:

$$\begin{aligned} -V(\Phi) &= -\frac{\mu^2}{2}(v+H)^2 - \frac{\lambda}{4}(v+H)^4 \\ &= -\lambda v^2 H^2 - \lambda v H^3 - \frac{\lambda}{4} H^4 \\ &= -\frac{1}{2} m_H^2 H^2 - \lambda v H^3 - \frac{\lambda}{4} H^4, \end{aligned} \quad (1.42)$$

where $m_H = \sqrt{2\lambda}v = \sqrt{-2\mu^2}$ is the mass of the Higgs boson and λv , $\lambda/4$ are the trilinear and quartic couplings of the three-point and four-point self-interaction, respectively.

1.3.2 Fermion mass generation

Gauge invariant fermion mass terms can be written using the Φ field as follows:

$$\begin{aligned} \mathcal{L}_{\text{Yukawa}} &= - \left[\lambda_e \bar{\ell}_L \Phi e_R + \lambda_\nu \bar{\ell}_L \Phi \nu_R + \lambda_u \bar{q}_L \tilde{\Phi} u_R + \lambda_d \bar{q}_L \Phi d_R + \text{h.c.} \right] \\ &= -\frac{\lambda_e}{\sqrt{2}}(v+H)(\bar{e}_L e_R + \bar{e}_R e_L) \\ &\quad -\frac{\lambda_u}{\sqrt{2}}(v+H)(\bar{u}_L u_R + \bar{u}_R u_L) - \frac{\lambda_d}{\sqrt{2}}(v+H)(\bar{d}_L d_R + \bar{d}_R d_L), \end{aligned} \quad (1.43)$$

where $e = e, \mu, \tau$, $\nu = \nu_e, \nu_\mu, \nu_\tau$, $u = u, c, t$ and $d = d, s, b$. $\tilde{\Phi} = i\sigma_2 \Phi^*$ such that $Y_W \tilde{\Phi} = -1$. The term that includes right-handed neutrinos is shown in grey. This type of neutrinos

1. The Standard Model of Particle Physics


does not interact with any particle of the SM, and it is usually removed from the lagrangian. With a bit of math, this lagrangian can be further simplified into:

$$\mathcal{L}_{\text{Yukawa}} = - \left(m_\psi \bar{\psi} \psi + \frac{m_\psi}{v} \bar{\psi} \psi H \right). \quad (1.44)$$

this expression contains two terms: a mass term proportional to $\bar{\psi} \psi$, where $m_\psi = \lambda_\psi v / \sqrt{2}$ is the mass of the fermion, and an interaction term with the Higgs boson, where the coupling strength is proportional to m_ψ . The free parameters λ_ψ are known as the Yukawa couplings of the Higgs boson to the SM fermions.

1.4 Feynman rules of the Standard Model

This section summarises all the Standard Model propagators and vertices at tree level² via their Feynman diagram representation, using the mass eigenstates for the gauge bosons. Feynman diagrams are always read from left to right. Fermion propagators are represented with a straight line and an arrow that shows the direction of the spin. Electroweak gauge boson propagators (γ , $W^\pm$, Z) are represented with a wavy line. Gluon propagators are represented with a curly line. Finally, the Higgs boson propagator (and any other spin-0 particle) is represented with a dashed line:

	fermion propagator
	anti-fermion propagator
	$\gamma/W^\pm/Z$ propagator
	gluon propagator
	Higgs propagator

1.4.1 Electroweak interactions

The electroweak lagrangian can be divided in two components. Interactions between fermions and gauge bosons are described by the following term in the EW lagrangian:


$$\begin{aligned} \mathcal{L}_{\text{EW}} &\ni i \bar{\psi} \gamma^\mu \left(-ig \mathbf{T} \cdot \mathbf{W}_\mu - ig' \frac{Y_W}{2} B_\mu \right) \psi \\ &= i \bar{\psi} \gamma^\mu \left[-i \frac{g}{2} \begin{pmatrix} W_\mu^3 & W_\mu^1 - i W_\mu^2 \\ W_\mu^1 + i W_\mu^2 & -W_\mu^3 \end{pmatrix} - ig' \frac{Y_W}{2} B_\mu \right] \psi. \end{aligned} \quad (1.45)$$

Interactions via a charged $W^\pm$ boson are governed by the W_μ^1 and W_μ^2 components associated to $\text{SU}(2)_L$. They only apply to left-handed fermions:

$$\begin{aligned} \mathcal{L}_{\text{EW}} &\ni \frac{g}{2} \bar{\psi}_L \gamma^\mu \begin{pmatrix} 0 & W_\mu^1 - i W_\mu^2 \\ W_\mu^1 + i W_\mu^2 & 0 \end{pmatrix} \psi_L \\ &= \frac{g}{\sqrt{2}} (\bar{\nu}_{\ell L} \ \bar{\ell}_L) \gamma^\mu \begin{pmatrix} 0 & W_\mu^+ \\ W_\mu^- & 0 \end{pmatrix} \begin{pmatrix} \nu_{\ell L} \\ \ell_L \end{pmatrix}, \end{aligned} \quad (1.46)$$

²Additional terms (ghosts) need to be included to perform loop calculations.

where $\ell = e, \mu, \tau, d, s, b$, and $\nu_\ell = \nu_e, \nu_\mu, \nu_\tau, u, c, t$ for simplicity. The sum over quark colours is not shown explicitly. This expression yields the following interaction vertices:



Tree-level transitions from lepton to quark (or viceversa) via the weak interaction are forbidden by the conservation of baryon and lepton numbers, but flavour-changing charged currents occur in the Standard Model. They are modulated by the Cabibbo–Kobayashi–Maskawa (CKM) matrix for quarks [26, 27]:

$$V_{\text{CKM}} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} = \begin{pmatrix} \text{large} & \text{small} & \text{very small} \\ \text{small} & \text{large} & \text{very small} \\ \text{very small} & \text{very small} & \text{large} \end{pmatrix}. \quad (1.47)$$

This mixing matrix appears due to a mismatch between gauge symmetry eigenstates and mass eigenstates. Transitions of the type $W^\pm \rightarrow ij$ are proportional to $|V_{ij}|^2$. Diagonal transitions are favoured by the CKM matrix probability, as it is shown in the graphical representation of the magnitude of its components.

Interactions via the neutral γ/Z bosons are a combination of W_μ^3 and B_μ and include both left- and right-handed fermions:

$$\begin{aligned} \mathcal{L}_{\text{EW}} &\ni \bar{\psi}_L \gamma^\mu \left[\frac{g}{2} \begin{pmatrix} W_\mu^3 & 0 \\ 0 & -W_\mu^3 \end{pmatrix} + g' \frac{Y_W}{2} B_\mu \right] \psi_L + \bar{\psi}_R \gamma^\mu \left(g' \frac{Y_W}{2} B_\mu \right) \psi_R \\ &= g(\bar{\nu}_{\ell L} \ \bar{\ell}_L) \gamma^\mu T^3 W_\mu^3 \begin{pmatrix} \nu_{\ell L} \\ \ell_L \end{pmatrix} + \frac{g'}{2} (\bar{\nu}_{\ell L} \ \bar{\ell}_L) \gamma^\mu Y_W B_\mu \begin{pmatrix} \nu_{\ell L} \\ \ell_L \end{pmatrix} \\ &\quad + \frac{g'}{2} \bar{\ell}_R \gamma^\mu Y_W B_\mu \ell_R + \frac{g'}{2} \bar{\nu}_{\ell R} \gamma^\mu Y_W B_\mu \nu_{\ell R}. \end{aligned} \quad (1.48)$$

Replacing W_μ^3, B_μ by A_μ, Z_μ yields:


$$\begin{aligned} \mathcal{L}_{\text{EW}} &\ni g \cos \theta_W (\bar{\nu}_{\ell L} \ \bar{\ell}_L) \gamma^\mu T^3 Z_\mu \begin{pmatrix} \nu_{\ell L} \\ \ell_L \end{pmatrix} + e (\bar{\nu}_{\ell L} \ \bar{\ell}_L) \gamma^\mu T^3 A_\mu \begin{pmatrix} \nu_{\ell L} \\ \ell_L \end{pmatrix} \\ &\quad - \frac{g' \sin \theta_W}{2} (\bar{\nu}_{\ell L} \ \bar{\ell}_L) \gamma^\mu Y_W Z_\mu \begin{pmatrix} \nu_{\ell L} \\ \ell_L \end{pmatrix} + \frac{e}{2} (\bar{\nu}_{\ell L} \ \bar{\ell}_L) \gamma^\mu Y_W A_\mu \begin{pmatrix} \nu_{\ell L} \\ \ell_L \end{pmatrix} \\ &\quad - \frac{g' \sin \theta_W}{2} \bar{\ell}_R \gamma^\mu Y_W Z_\mu \ell_R + \frac{e}{2} \bar{\ell}_R \gamma^\mu Y_W A_\mu \ell_R \\ &\quad - \frac{g' \sin \theta_W}{2} \bar{\nu}_{\ell R} \gamma^\mu Y_W Z_\mu \nu_{\ell R} + \frac{e}{2} \bar{\nu}_{\ell R} \gamma^\mu Y_W A_\mu \nu_{\ell R}. \end{aligned} \quad (1.49)$$

1. The Standard Model of Particle Physics

Using that $Q = T_3 + Y_W/2$ and $T_3 = 0$ for right-handed particles:

$$\begin{aligned} \mathcal{L}_{EW} \ni & \frac{g}{\cos \theta_W} (\bar{\nu}_{\ell L} \bar{\ell}_L) \gamma^\mu (T^3 - Q \sin^2 \theta_W) Z_\mu \begin{pmatrix} \nu_{\ell L} \\ \ell_L \end{pmatrix} \\ & - \frac{g'}{\sin \theta_W} (\bar{\nu}_{\ell R} \gamma^\mu (Q \sin^2 \theta_W) Z_\mu \nu_{\ell R} + \bar{\ell}_R \gamma^\mu Q Z_\mu \ell_R) \\ & + eQ (\bar{\nu}_{\ell L} \bar{\ell}_L) \gamma^\mu A_\mu \begin{pmatrix} \nu_{\ell L} \\ \ell_L \end{pmatrix} + eQ (\bar{\ell}_R \gamma^\mu A_\mu \ell_R + \bar{\nu}_{\ell R} \gamma^\mu A_\mu \nu_{\ell R}). \end{aligned} \quad (1.50)$$


In this expression, it is easy to see that the Z boson is chiral, because it couples differently to left- and right-handed fermions, while the photon is not. The corresponding interaction vertices are:



The only interactions remaining are interactions between the electroweak gauge bosons. They can be derived from the non-abelian term in the field strength of the $\mathbf{W}$ boson. Three-point interactions are obtained from:

$$\begin{aligned} \mathcal{L}_{EW} \ni & -\frac{g}{4} \left[(\partial_\mu W_\nu^i - \partial_\nu W_\mu^i) (\epsilon_{ijk} W_\mu^j W_\nu^k) + (\epsilon^{ijk} W_j^\mu W_k^\nu) (\partial_\mu W_\nu^i - \partial_\nu W_\mu^i) \right] \\ & = -\frac{g}{2} (\partial_\mu W_\nu^i - \partial_\nu W_\mu^i) (\epsilon_{ijk} W_\mu^j W_\nu^k). \end{aligned} \quad (1.51)$$


Introducing the mass eigenstates generates terms proportional to $(W^- W^+ \gamma)$ and $(W^- W^+ Z)$:



Four-point interactions can be derived from the fourth-order term:

$$\begin{aligned} \mathcal{L}_{EW} \ni & -\frac{g^2}{4} \epsilon^{ijk} W_\mu^j W_\nu^k \epsilon_{imn} W_m^\mu W_n^\nu \\ & = \frac{g^2}{4} (W_\mu^i W_\nu^j W_i^\mu W_j^\nu - W_\mu^j W_\nu^i W_i^\mu W_j^\nu). \end{aligned} \quad (1.52)$$

Expanding this expression in terms of the mass eigenstates yields terms proportional to $(W^- W^+ W^- W^+)$, $(W^- W^+ ZZ)$, $(W^- W^+ AA)$ and $(W^- W^+ ZA)$. They can be represented as:




1.4.2 Strong interactions

Strong interactions between quarks and gluons are given by the QCD covariant derivative:

$$\mathcal{L}_{\text{QCD}} \ni \bar{\psi}^A \left(g_s \mathbf{T}^{AB} \cdot \mathbf{G}_\mu \right) \psi^B. \quad (1.53)$$


The strong force does not distinguish between left- or right-handed quarks, and it is also independent of the quark flavour. Because of that, the expression can be written in terms of $\psi = u, d, c, s, t, b$. Upper case indices refer to the sum over colour, and the dot product is computed over components in the basis of the Gell-Mann matrices. The interaction vertex is proportional to $\mathbf{T}^{AB}$, meaning that the gluon interacts with quarks by exchanging colour charge. There are eight gluons in QCD, one for each matrix, with different colour charges.



Similar to the electroweak lagrangian, QCD also contains a self-interaction term for the gluon which includes three- and four-point interactions:

$$\mathcal{L}_{\text{QCD}} \ni -\frac{g_s}{2} (\partial_\mu G_\nu^a - \partial_\nu G_\mu^a) (f_{abc} G_b^\mu G_c^\nu) - \frac{g_s^2}{4} f_{abc} G_\mu^b G_\nu^c f_{amn} G_m^\mu G_n^\nu. \quad (1.54)$$

They can be represented as:




1.4.3 Higgs interactions

The Higgs boson interacts with all massive particles in the SM, including itself. Interactions with gauge bosons are given by the derivative terms in the SSB lagrangian:

$$\mathcal{L}_{\text{SSB}} \ni m_W^2 \left(\frac{1}{v} W_\mu^- W_+^\mu H + \frac{1}{v^2} W_\mu^- W_+^\mu H^2 \right) + \frac{1}{2} m_Z^2 \left(\frac{1}{v} Z_\mu Z^\mu H + \frac{1}{v^2} Z_\mu Z^\mu H^2 \right). \quad (1.55)$$

They correspond to:




1. The Standard Model of Particle Physics

Higgs self-interactions are obtained from the Higgs potential $V(\Phi)$:

$$\mathcal{L}_{\text{SSB}} \ni -\lambda v H^3 - \frac{\lambda}{4} H^4, \quad (1.56)$$


and can be represented as:



Finally, interactions of the Higgs boson with fermions are included in the SM via the Yukawa term:

$$\mathcal{L}_{\text{Yukawa}} \ni -\frac{m_\psi}{v} \bar{\psi}\psi H. \quad (1.57)$$

They are proportional to the mass of the fermion, and they do not modify its flavour.



Chapter 2

Standard Model measurements

Since its conception, the Standard Model has undergone extensive experimental testing, and its predictions have been remarkably accurate across a wide range of phenomena in Particle Physics. However, it is not a complete theory; it does not account for gravity, dark matter, dark energy, or neutrino masses. This section explores some of the experimental successes of the SM throughout its history, as well as other phenomena that hint the existence of Physics beyond the Standard Model (BSM).

2.1 Experimental successes of the Standard Model

2.1.1 Particles predicted by the Standard Model

Quantum chromodynamics

Before the 1960s, protons, neutrons and light mesons were considered to be elementary particles. This notion was changed when Gell-Mann and Zweig proposed the quark model [28, 29]. Shown in Figure 2.1, it was a classification scheme for hadrons in terms of their valence quarks, closely related to the SU(3) flavour symmetry of QCD. The quark model explained the masses of the already known mesons and baryons based on the existence of three constituent quarks: u , d and s . In 1968, the first high energy electron-proton scattering experiments at the Stanford Linear Accelerator Center (SLAC) provided evidence for the existence of point-like particles inside protons and neutrons, corresponding to the u - and d -quarks [30]. The s -quark was indirectly validated by the SLAC experiment, since it provided an explanation for the existence of pions (π) and kaons (K).




Figure 2.1: In the initially proposed quark model, u , d and s quarks (blue) formed a triplet of the SU(3) flavour symmetry, which could be combined to create mesons and baryons. This figure contains two examples: the meson octet for pseudoscalar mesons (green) and the baryon octet for spin-1/2 baryons (red). They are represented in the $I_3 - Y$ plane of the QCD flavour symmetry. The line $Q = Y/2 + I_3$ is also included.

2. Standard Model measurements

The existence of the c -quark was first suggested by Bjorken and Glashow in 1964 [31], and later predicted by the Glashow–Iliopoulos–Maiani (GIM) mechanism [32], which required the existence of a fourth quark flavour to explain the suppression of flavour-changing neutral currents (FCNCs) in neutral kaon decays using the Cabibbo mixing angle [26]. This is represented in Figure 2.2. Charm quarks were produced simultaneously in 1974 by SLAC and the Brookhaven National Laboratory in the form of the $J/\psi(c\bar{c})$ meson [33, 34].




Figure 2.2: Example Feynman diagrams contributing to neutral kaon decay. This decay is suppressed at loop level due to the opposite sign from the u - and c -mediated diagrams with $m_u \neq m_c$. θ_C is the Cabibbo angle.

The number of quark flavours grew to the current six in 1973, when Kobayashi and Maskawa noted that the experimental observation of indirect CP violation in kaon decays (Figure 2.3) could be explained should there be another pair of quarks [27]. This is because CP violation is related to complex coupling constants. In the four quark model, the mixing matrix is a 2×2 unitary matrix, which can always be reduced into a real matrix plus a global phase:

$$\begin{pmatrix} V_{ud} & V_{us} \\ V_{cd} & V_{cs} \end{pmatrix} = \begin{pmatrix} \cos \theta_C & \sin \theta_C \\ -\sin \theta_C & \cos \theta_C \end{pmatrix} e^{i\delta} \quad (2.1)$$

where θ_C is the Cabibbo angle. A six-quarks model allows for a CP-violating phase δ_{13} :

$$\begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} = \begin{pmatrix} c_{12}c_{13} & s_{12}c_{13} & s_{13}e^{-i\delta_{13}} \\ -s_{12}c_{23} - c_{12}s_{23}s_{13}e^{i\delta_{13}} & c_{12}c_{23} - s_{12}s_{23}s_{13}e^{i\delta_{13}} & s_{23}c_{13} \\ s_{12}s_{23} - c_{12}c_{23}s_{13}e^{i\delta_{13}} & -c_{12}s_{23} - s_{12}c_{23}s_{13}e^{i\delta_{13}} & c_{23}c_{13} \end{pmatrix} \quad (2.2)$$

where $c_{ij} = \cos \theta_{ij}$ and $s_{ij} = \sin \theta_{ij}$. θ_{12} is the Cabibbo angle.




Figure 2.3: Feynman diagrams contributing to neutral kaon oscillation $K_0 \leftrightarrow \bar{K}_0$. CP violation appears because this process does not occur with the same probability in both directions.

In 1977, the bottom quark was observed at Fermilab in the form of the so-called $\Upsilon(b\bar{b})$ meson [35]. The t -quark, with a mass much larger than the rest of its partners, was

discovered in the Tevatron proton-antiproton collider (Fermilab) in 1995 [36, 37]. Direct CP violation in neutral kaon oscillations was later observed by Fermilab and CERN in 1999 [38, 39].

The gluon was discovered in 1979 at the electron-positron collider PETRA from DESY through the observation of three-jet events corresponding to $ee \rightarrow q\bar{q}g$ [40–42]. One year later, the newly discovered gluon was confirmed to be a spin-1 particle, as predicted by QCD.

Electroweak theory

Charged currents were already known previous to the formulation of the electroweak theory, thanks to multiple studies in nuclear physics, including phenomena such as neutron beta decay. They were initially described as a four-point interaction by Fermi, and later incorporated in the Standard Model as the two spin-1 particles emerging from the $SU(2)_L$ gauge symmetry, the $W^\pm$ bosons. However, the EW theory also predicted the existence of a neutral current, associated with the $U(1)_Y$ gauge symmetry group, which had never been observed before. These neutral currents were measured for the first time in 1974 by the Gargamelle experiment at CERN, in hadron-neutrino scattering of the type $\nu + n \rightarrow \nu + \text{hadrons}$ [43].

The next particle to be discovered was the τ lepton, in 1975. At that time, a third lepton generation was not strongly motivated by the theory, but it was found anyway in SLAC via anomalous lepton production in electron-positron collisions, $ee \rightarrow e\mu$, which in reality corresponded to the process $ee \rightarrow \tau\tau$ with $\tau \rightarrow e/\mu + \nu$ [44].

The discovery of the $W^\pm$ and Z bosons had to wait for the construction of a particle accelerator powerful enough to produce them. That was the Super Proton Synchrotron (SPS) at CERN, operating at a centre-of-mass energy of 540 GeV for proton-antiproton collisions. In 1983, $W \rightarrow e\nu$ decays were observed [45], and later the same year, $Z \rightarrow ee/\mu\mu$ [46]. In 1989, the Large Electron-Positron Collider (LEP), predecessor of the current LHC, was conceived to precisely measure the properties of the electroweak gauge bosons. One of the highlights of the LEP Physics programme was the confirmation of the number of neutrino families in the SM, which took place during 1989 by performing precision measurements of the Z boson properties [47]. Figure 2.4 shows the Z lineshape measurement that lead to the determination of the number of neutrino families to be 2.9840 ± 0.0082 .

The existence of three charged leptons (e, μ, τ), plus the number of neutrino families determined by LEP motivated the construction of the DONUT experiment at Fermilab to find a third neutrino, ν_τ , given that only two neutrinos had been observed so far at that point in history. The τ neutrino was discovered in the year 2000, by directing a ν_τ beam into an iron plate and observing the τ leptons produced due to the interaction between the neutrino and the nuclei [48].

The Higgs boson

The Brout-Englert-Higgs mechanism for mass generation via spontaneous symmetry breaking was formulated in 1964, but it was not confirmed experimentally until several decades later. The discovery of the Higgs boson in 2012 by the ATLAS and CMS

2. Standard Model measurements




Figure 2.4: Measurement of the $ee \rightarrow \text{hadronic}$ production cross section around the Z resonance. The curves indicate the predicted cross section for two, three and four neutrino species with SM couplings and negligible mass [47].

collaborations at CERN [1, 2] at a mass of 125 GeV marked a major milestone in Particle Physics. The initial measurements during the LHC Run 1 included the gluon-gluon fusion (ggF), vector-boson fusion (VBF), associated production with a vector boson (WH/ZH) and associated production with a $t\bar{t}$ pair ($t\bar{t}H$) Higgs production mechanisms, whose Feynman diagrams are shown in Figure 2.5. The $H \rightarrow \gamma\gamma/WW^*/ZZ^*$ decay modes were used, which offered a clean signature with either photons or leptons in the final state. Predicted SM Higgs cross sections and branching fractions are shown in Figure 2.6.




Figure 2.5: Feynman diagrams for the main Higgs production mechanisms in the LHC. VH can refer to both WH and ZH .

The outstanding performance of the LHC Run 2 made it possible for the ATLAS and CMS experiments to independently and unambiguously establish the couplings of the Higgs boson to the charged fermions of the third generation (t -quark, b -quark and τ lepton) and gauge bosons (γ , W , Z) [3, 4]. In all observed production and decay modes measured so far, the rates and differential measurements have been found to be




Figure 2.6: (a) SM Higgs boson production cross sections as a function of the centre-of-mass energy, $\sqrt{s}$, for proton-proton collisions and (b) branching ratios for the main decays of the SM Higgs boson around $m_H = 125$ GeV. Data from [49].

consistent, within experimental and theoretical uncertainties, with the SM predictions. And, in high mass resolution decay channels, such as $H \rightarrow ZZ^* \rightarrow 4\ell$ or $H \rightarrow \gamma\gamma$, the mass of the Higgs boson has been measured at an unprecedented precision level [50, 51]. Figure 2.7 shows the most up-to-date ATLAS results using Run 1 and Run 2 data for the Higgs couplings and mass.




Figure 2.7: (a) Reduced Higgs boson coupling strength modifiers and their uncertainties from Run 2 [3] and (b) summary of Higgs mass measurements from the $H \rightarrow ZZ^* \rightarrow 4\ell$ and $H \rightarrow \gamma\gamma$ decay channels and their combination [50].

2. Standard Model measurements

However, within the current precision, a more complex Higgs sector with additional states, although significantly constrained, is not ruled out. This is why both the present and future strategy for CERN includes improving the precision on the Higgs boson decay width and branching fractions, as well as measuring the Higgs triple and quartic self-couplings, only accessible via HH and HHH production.

2.1.2 Cross section measurements at the LHC

There has been a very extensive programme during LHC Run 1 and Run 2 to measure as many SM processes as possible, while looking for possible deviations from SM predictions and, at the same time, improving the description of well-known processes that will later become the backgrounds of more intricated searches. Figure 2.8 shows an overview of the most up-to-date cross section measurements, spanning more than ten orders of magnitude. So far, no significant differences with respect to the SM have been observed.




Figure 2.8: Overview of cross section measurements of selected SM processes compared to the corresponding theoretical expectations. The figure includes results from Run 1, Run 2 and the early stages of Run 3, using the ATLAS detector [52].

2.2 Shortcomings

The Standard Model is a remarkably successful theory for describing Particle Physics, but it has some theoretical limitations and open questions that hint that there might be a more fundamental theory underneath.

2.2.1 Theoretical problems

Theoretical problems arise in the SM for two different reasons. The first one is that it fails to describe all fundamental forces in the universe, as it does not include gravity. Gravity is described at the classical level by the theory of General Relativity as the curvature of the four-dimensional spacetime and its relation with the energy-momentum tensor. Attempts to quantise gravity lead to mathematical inconsistencies, which is why it is not currently possible to interpret it in terms of Particle Physics. A theory of Quantum Gravity is fundamental to understand extreme scenarios such as the very early stages of the Big Bang or the proximity of the singularity of a black hole, which are characterised by the Planck energy scale $\Lambda_{\text{Planck}} \sim 10^{19}$ GeV.

The second reason is that some features of the Standard Model are *ad hoc* additions that are not problems *per se*, but reflect a lack of understanding of the underlying physics. Examples of these are:

- **The Higgs hierarchy problem.** The mass of the Higgs boson is a free parameter of the SM that has been measured to be around 125 GeV. This mass is not protected by any symmetry, which means that no symmetry is restored when it is set to 0. Therefore, if the SM is taken as a low-energy effective field theory valid up to a cutoff scale Λ , then the bare mass of the Higgs boson can receive additive corrections proportional to Λ^2 during renormalisation. Should there not be new physics up to the Planck scale, where quantum gravity effects become non-negligible, then the mass of the Higgs boson would receive corrections of $O(\Lambda_{\text{Planck}}^2)$. Therefore, reproducing the 125 GeV measured mass would require a very large fine-tuning in the bare mass:

$$(m_H^{\text{exp}})^2 = (m_H^{\text{bare}})^2 + O(\Lambda_{\text{Planck}}^2). \quad (2.3)$$

- **The SM flavour structure.** All particle masses are free parameters in the SM. This also raises the question of whether there is a theory that predicts the observed mass relations and why some particles are so heavy while others are so light. From experimental measurements, it is observed that $m_{\text{1st family}} \ll m_{\text{2nd family}} \ll m_{\text{3rd family}}$ and $|V_{ij}^{\text{CKM}}(i=j)| \gg |V_{ij}^{\text{CKM}}(i \neq j)|$.
- **The strong CP problem.** Symmetry-wise, QCD allows for a CP-violating term:

$$\mathcal{L}_{\text{CP}} = \theta \frac{g_s^2}{32\pi^2} G_{\mu\nu} \tilde{G}^{\mu\nu}, \quad (2.4)$$

where θ is a constant parameter representing the strength of CP-violation. Experimental measurements show that this term is strongly suppressed, with

2. Standard Model measurements

$\theta \lesssim 10^{-10}$. While this value could be accidental, there could also be additional symmetries that explain why θ is so small.

2.2.2 Experimental problems

Although the Standard Model has successfully predicted most of the new phenomena observed in collider physics during the XX and XXI centuries, there are several experimental observations that can not be fully explained by the SM alone, suggesting that it is an incomplete theory of Particle Physics. The following list contains some examples:

- **Dark matter.** Dark matter is a hypothetical form of matter that may only interact with ordinary matter via the gravitational force. Experimental evidences include the anomalous galaxy rotation curves and galaxy cluster interactions, as well as the mechanism for large-scale structure formation in the universe (Figures 2.9 and 2.10). From this type of observations, it is calculated that dark matter could make up about 27% of the universe's mass-energy content.
- **Dark energy.** Cosmological observations of distant supernovae suggest that the universe is expanding at an accelerating rate. Dark energy is a hypothetical form of energy that drives this expansion, that would constitute around 68% of the universe's mass-energy content.
- **Matter-antimatter asymmetry.** The universe is made almost entirely of matter, with very little antimatter content, even though the Big Bang should have produced equal amounts of matter and antimatter. The SM includes a source of CP violation in weak interactions, but this source is too small to account for the observed matter-antimatter asymmetry. This implies that there might be additional CP-violating processes beyond the SM.
- **Neutrino oscillations.** Neutrinos are massless particles in the SM, and do not acquire a mass term after SSB. Nevertheless, the mass difference between different neutrino flavours, Δm_{ij} , has been measured to be different than 0, meaning that at least two neutrinos are massive. The two possible hypotheses for neutrino mass ordering are shown in Figure 2.11. The mass of each individual neutrino flavour remains unknown. The fact that neutrinos are massive and their masses are very small induces another phenomenon called neutrino oscillation, which refers to the non-conservation of lepton flavour, and arises from the mixing between the flavour and mass eigenstates of neutrinos, dictated by the Pontecorvo–Maki–Nakagawa–Sakata (PMNS) matrix [53, 54]:

$$U_{\text{PMNS}} = \begin{pmatrix} U_{e1} & U_{e2} & U_{e3} \\ U_{\mu1} & U_{\mu2} & U_{\mu3} \\ U_{\tau1} & U_{\tau2} & U_{\tau3} \end{pmatrix} = \begin{pmatrix} \text{green} & \text{grey} & \text{red} \\ \text{grey} & \text{green} & \text{grey} \\ \text{red} & \text{grey} & \text{green} \end{pmatrix}, \quad (2.5)$$

where e, μ, τ are the flavour eigenstates and 1, 2, 3 are the mass eigenstates.

- **Muon ($g-2$) anomaly.** The muon ($g-2$) anomaly refers to an observed discrepancy between the experimentally measured value and the theoretical prediction for the magnetic moment of the muon. The most precise measurement of the muon ($g-2$) was carried out by the Muon $g-2$ collaboration at Fermilab in 2023 [55], showing a 5.2σ deviation from the 2020 calculations of the SM prediction. While this type of anomaly could be explained by some BSM scenarios, there is also a lot of effort involved in improving both the experimental measurements as well as the theory calculations from lattice QCD. The current status is shown in Figure 2.12.




Figure 2.9: Superposition of 21 galaxy rotation curves. The plot shows the speed of the objects in the galactic disk with respect to the distance to the galactic nucleus. This form of the rotation curves suggests that the total mass of the galaxy is not condensed around the centre, where ordinary matter lies, but also at large radius in the form of dark matter [56].

2. Standard Model measurements




Figure 2.10: (a) Image of the bullet cluster, comprised of two colliding galaxy clusters. The contour lines represent the mass distribution obtained via gravitational lensing and the gray dots correspond to the X-ray image from ordinary matter [57]. Dark matter has a much lower interaction cross section, and therefore suffers less deceleration during the collision. (b) Comparison of the galaxy distribution obtained from the Sloan Digital Sky Survey (blue) and theoretical mock catalogs from the Millenium Simulation. The pattern of galaxy clustering through gravity can help determining the properties of hypothetical dark matter candidates [58].




Figure 2.11: Scheme of the two distinct neutrino mass hierarchies, from measurements of solar (*sol*) and atmospheric (*atm*) neutrinos. The colour code indicates the fraction of each neutrino flavour (e, μ, τ) present in each of the mass eigenstates (1, 2, 3) [59].




Figure 2.12: Summary of theoretical calculations and experimental measurements of the muon anomaly $a_\mu = (g - 2)/2$ [60].

Chapter 3

Physics beyond the Standard Model

While the Standard Model has been able to provide a highly successful description of Particle Physics, it fails to account for key phenomena such as those described in Section 2.2. Physics beyond the Standard Model include extensions of the already existing theory, for example through additional symmetries, as well as fundamentally different models, such as String Theory. The question of which theory is the right one, or at least the best one, can only be settled via experimental observation, and is one of the main research topics in the LHC. The work presented in this thesis aims to probe SM extensions containing axion-like particles.

3.1 Axion-like particles

3.1.1 What is an axion?

Axions are hypothetical elementary particles proposed as a solution to the strong CP problem in QCD, which arises from the fact that the CP violating term that is allowed in the QCD lagrangian:

$$\mathcal{L}_{\text{CP}} = \theta \frac{g_s^2}{32\pi^2} G_{\mu\nu} \tilde{G}^{\mu\nu}, \quad (3.1)$$

is strongly suppressed, with $\theta \lesssim 10^{-10}$ according to experimental observations. The reason why this is considered a problem is due to the high level of fine-tuning necessary for the SM to agree with the experiment, which would not exist if there was an additional symmetry. In 1977, the Peccei-Quinn mechanism [61, 62] was postulated. This model introduces a $U(1)_{\text{PQ}}$ global symmetry, which is spontaneously broken at a certain energy scale f_a , which dynamically drives θ to 0, solving the strong CP problem. The pseudo-Nambu-Goldstone boson that originates from this broken symmetry is the axion, usually denoted as a , whose mass can be defined as:

$$m_a \sim \left(\frac{10^9 \text{ GeV}}{f_a} \right) \text{ meV}. \quad (3.2)$$

Axions have been studied in a variety of astrophysical, cosmical and laboratory experiments, leading to $f_a \gg \Lambda_{\text{EW}} \sim \mathcal{O}(100 \text{ GeV})$. For values of f_a much larger than the EW symmetry breaking scale Λ_{EW} , the axion becomes inaccessible to collider experiments due to its very low mass.

3.1.2 Generalisation of the axion theory

Many extensions of the SM feature one or several spontaneously broken global $U(1)$ symmetries. Similar to the Peccei-Quinn mechanism, these broken symmetries produce pseudo-Nambu-Goldstone bosons below the symmetry breaking energy scale, which

can naturally be light with respect to the EW energy scale. Axion-like particles, also known as ALPs, is the name given to these "axions" that are not necessarily related to the strong CP problem, thus allowing more flexibility in their theoretical and experimental properties, and making them more accessible for searches in collider experiments. They appear in theories with extended Higgs sectors, such as the next-to-minimal supersymmetric SM (NMSSM) [5–9], dark matter models [10–14], models with a first-order EW phase transition [15, 16] and theories of neutral naturalness [17, 18]. When studying ALPs at the LHC, it is common to use a general theoretical framework at the LHC energy scale that can be later interpreted in different UV theories¹. Examples of these are the ALP effective field theory (EFT) [63], the two Higgs doublet model (2HDM) [64] and simplified models where a spin-0 singlet is coupled to the SM [65].

The work in this thesis is based on simplified ALP models derived from the 2HDM+ a . Section 3.2 introduces the general 2HDM potential and the 2HDM+ a extension, and shows the assumptions that lead to a simplified lagrangian for describing ALP production in association with a $t\bar{t}$ pair and decaying to two b -quarks, corresponding to the $t\bar{t}a$, $a \rightarrow b\bar{b}$ analysis described in Part V. Section 3.3 on the other hand, shows how a simplified model can also be used to probe ALP production at the LHC, in this case via Higgs boson decays, corresponding to the $H \rightarrow aa \rightarrow 4b$ analysis described in Part IV.

3.2 Two Higgs doublet model

3.2.1 General formulation

The SM Higgs sector is made up of a single $SU(2)_L$ doublet Φ with weak hypercharge $Y_W = 1$. Adding another doublet to this minimal scenario is one of the simplest extensions. The two Higgs doublet model (2HDM) expands the SM by introducing an additional complex scalar doublet, Φ_2 , alongside the SM scalar doublet Φ_1 . This extension provides an extended scalar sector, containing five different spin-0 particles after electroweak symmetry breaking: two neutral CP-even (h and H), one neutral CP-odd (A) and two charged CP-even ($H^\pm$). The 2HDM scalar potential for two Higgs doublets is given by [64]:

$$\begin{aligned} V_{2\text{HDM}}(\Phi_1, \Phi_2) = & \mu_1^2 \Phi_1^\dagger \Phi_1 + \mu_2^2 \Phi_2^\dagger \Phi_2 - \mu_3^2 (\Phi_1^\dagger \Phi_2 + \text{h.c.}) \\ & + \frac{1}{2} \lambda_1 (\Phi_1^\dagger \Phi_1)^2 + \frac{1}{2} \lambda_2 (\Phi_2^\dagger \Phi_2)^2 + \lambda_3 (\Phi_1^\dagger \Phi_1) (\Phi_2^\dagger \Phi_2) \\ & + \lambda_4 (\Phi_1^\dagger \Phi_2) (\Phi_2^\dagger \Phi_1) + \frac{1}{2} \lambda_5 \left[(\Phi_1^\dagger \Phi_2)^2 + \text{h.c.} \right], \end{aligned} \quad (3.3)$$

where all the parameters are real. In this formulation, a series of assumptions are implicit. First, CP is conserved in the Higgs sector and not spontaneously broken. In addition, the hypercharge of the two doublets is $Y_W = 1$. This condition is imposed so

¹UV theory refers to a more general quantum field theory that is well-defined at arbitrarily high energies.

that the model is consistent with the experimental observation:

$$\rho \equiv \frac{m_W^2}{m_Z^2 \cos \theta_W} = \frac{\sum_{i=1}^N \left[T_{3,i}(T_{3,i} + 1) - \frac{1}{4} Y_{W,i}^2 \right] v_i}{\sum_{i=1}^N \frac{1}{2} Y_{W,i}^2 v_i} \approx 1, \quad (3.4)$$

where N is the number of complex scalar doublets. If the SM Higgs hypothesis is assumed, then $\rho = 1$. Furthermore, a $\mathbb{Z}_2$ symmetry is imposed to suppress FCNC terms, which are strongly suppressed in the SM and highly constrained by the experiment as well. This symmetry is allowed to be softly broken by the term $-\mu_3^2(\Phi_1^\dagger \Phi_2 + \Phi_2^\dagger \Phi_1)$, which provides the mixing between the two doublets and can give rise to FCNCs in loop diagrams.

The two Higgs doublets can be parametrised in the same way as the SM Higgs boson, with one charged component and one neutral component:

$$\Phi_i = \begin{pmatrix} \phi_i^+ \\ \phi_i^0 \end{pmatrix} \text{ with } i = 1, 2. \quad (3.5)$$

Similarly, they also develop a vacuum expectation value:

$$\langle \Phi_i \rangle = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v_i \end{pmatrix} \text{ with } i = 1, 2. \quad (3.6)$$

A general way of parametrisng the oscillations around the vacuum is:

$$\Phi_i = \begin{pmatrix} \phi_i^+ \\ (v_i + \rho_i + i\eta_i)/\sqrt{2} \end{pmatrix} \quad (3.7)$$

which corresponds to eight degrees of freedom. The mass term for the charged scalar is given by:

$$\mathcal{L}_{\phi^\pm} \ni [\mu_3^2 - (\lambda_4 + \lambda_5)v_1v_2](\phi_1^-, \phi_2^-) \begin{pmatrix} v_2/v_1 & -1 \\ -1 & v_1/v_2 \end{pmatrix} \begin{pmatrix} \phi_1^+ \\ \phi_2^+ \end{pmatrix}. \quad (3.8)$$

This matrix has one zero eigenvalue corresponding to the charged Nambu-Goldstone boson that is absorbed by the $W^\pm$. The non-zero eigenvalue is the mass of the charged Higgs:

$$m_{H^\pm}^2 = [\mu_3^2 - (\lambda_4 + \lambda_5)v_1v_2] \frac{v_1^2 + v_2^2}{v_1v_2}. \quad (3.9)$$

The mass terms for the pseudoscalars are derived from the neutral imaginary component of the doublets:

$$\mathcal{L}_\eta \ni [\mu_3^2 - 2\lambda_5v_1v_2](\eta_1, \eta_2) \begin{pmatrix} v_2/v_1 & -1 \\ -1 & v_1/v_2 \end{pmatrix} \begin{pmatrix} \eta_1 \\ \eta_2 \end{pmatrix}. \quad (3.10)$$

This matrix yields one neutral Nambu-Goldstone boson that is absorbed by the Z and one massive pseudoscalar particle, A :

$$m_A^2 = [\mu_3^2 - 2\lambda_5v_1v_2] \frac{v_1^2 + v_2^2}{v_1v_2}. \quad (3.11)$$

3. Physics beyond the Standard Model

The two remaining degrees of freedom correspond to the two neutral CP-even states:

$$\mathcal{L}_\rho \ni -(\rho_1, \rho_2) \begin{pmatrix} \mu_3^2(v_2/v_1) + \lambda_1 v_1^2 & -\mu_3^2 + (\lambda_3 + \lambda_4 + \lambda_5)v_1 v_2 \\ -\mu_3^2 + (\lambda_3 + \lambda_4 + \lambda_5)v_1 v_2 & \mu_3^2(v_1/v_2) + \lambda_2 v_2^2 \end{pmatrix} \begin{pmatrix} \rho_1 \\ \rho_2 \end{pmatrix}. \quad (3.12)$$

It is common to use the mixing angle α to define the mass eigenstates h and H as:

$$\begin{pmatrix} H \\ h \end{pmatrix} = - \begin{pmatrix} \cos \alpha & \sin \alpha \\ -\sin \alpha & \cos \alpha \end{pmatrix} \begin{pmatrix} \rho_1 \\ \rho_2 \end{pmatrix}, \quad (3.13)$$

which characterises the rotation that diagonalises the mass-squared matrix. The other mixing angle that is used to define the model is β :

$$\tan \beta \equiv \frac{v_2}{v_1}, \quad (3.14)$$

which is the angle that diagonalises the mass-squared matrices of the charged and pseudoscalar sectors:

$$\begin{pmatrix} G^+ \\ H^+ \end{pmatrix} = - \begin{pmatrix} \cos \beta & \sin \beta \\ -\sin \beta & \cos \beta \end{pmatrix} \begin{pmatrix} \phi_1^+ \\ \phi_2^+ \end{pmatrix} \quad \text{and} \quad \begin{pmatrix} G^0 \\ A \end{pmatrix} = - \begin{pmatrix} \cos \beta & \sin \beta \\ -\sin \beta & \cos \beta \end{pmatrix} \begin{pmatrix} \eta_1 \\ \eta_2 \end{pmatrix}, \quad (3.15)$$

with $G^\pm$ and G^0 referring to the charged and neutral NG bosons, respectively. Additionally, if one redefines the Φ_1, Φ_2 doublets as follows:

$$\begin{pmatrix} H_1 \\ H_2 \end{pmatrix} = \begin{pmatrix} \cos \beta & \sin \beta \\ -\sin \beta & \cos \beta \end{pmatrix} \begin{pmatrix} \Phi_1 \\ \Phi_2 \end{pmatrix}, \quad (3.16)$$

then H_1 has a vacuum expectation value of $v/\sqrt{2}$, with $v^2 \equiv v_1^2 + v_2^2 \approx (246 \text{ GeV})^2$, and H_2 has a vacuum expectation value of 0.

3.2.2 Phenomenology of the two Higgs doublet model

In a 2HDM, contributions from the two scalar doublets to the fermion mass terms can produce FCNCs at tree-level, which are heavily constrained by the experiment. One possible solution is to have fermions with the same quantum numbers couple to the same Higgs multiplet, an assumption that cancels first-order FCNCs. This is the reason why the $\mathbb{Z}_2$ symmetry is usually introduced. Depending on the choice of $\mathbb{Z}_2$ symmetry, four variants of 2HDM can be defined: type I, type II, type III (also known as type Y or flipped) and type IV (also known as type X or lepton-specific). They are summarised in Table 3.1.

The α and β parameters determine the interactions of the various Higgs fields with vector bosons and fermions, being crucial to discuss the phenomenology of the model. Table 3.2 shows the coupling strength modifiers (ξ) to the SM-like Yukawa couplings of the neutral h, H and A to fermions. They appear in the 2HDM Yukawa lagrangian as:

$$\begin{aligned} \mathcal{L}_{\text{2HDM}}^{\text{Yukawa}} \ni & - \sum_f \frac{m_f}{v} (\xi_{h,f} \bar{f} f h + \xi_{H,f} \bar{f} f H - i \xi_{A,f} \bar{f} \gamma^5 f A) \\ & - \frac{1}{v/\sqrt{2}} [V_{ud} \bar{u} (m_u \xi_{A,u} P_L + m_d \xi_{A,d} P_R) d H^+ + m_\ell \xi_{A,\ell} \bar{\nu}_L \ell_R H^+ + \text{h.c.}], \end{aligned} \quad (3.17)$$

Type	Φ_1	Φ_2	ℓ_L	ℓ_R	q_L	u_R	d_R
I	+	−	+	−	+	−	−
II	+	−	+	+	+	−	+
III	+	−	+	−	+	−	+
IV	+	−	+	+	+	−	−

Table 3.1: $\mathbb{Z}_2$ charges in the four 2HDM types. The notation $\ell = e, \mu, \tau$, $u = u, c, t$ and $d = d, s, b$ is used. By convention, u_R always couples to Φ_2 .

and they vary depending on the 2HDM type. The couplings to vector bosons are the same in all models. The coupling of h to either WW or ZZ is the same as the SM coupling times $\sin(\beta - \alpha)$, and the coupling of H is the same as the SM times $\cos(\alpha - \beta)$. The coupling of the A boson is zero.

Coupling	Type I	Type II	Type III	Type IV
$\xi_{h,\ell}$	$\cos \alpha / \sin \beta$	$-\sin \alpha / \cos \beta$	$-\sin \alpha / \cos \beta$	$\cos \alpha / \sin \beta$
$\xi_{h,u}$	$\cos \alpha / \sin \beta$	$\cos \alpha / \sin \beta$	$\cos \alpha / \sin \beta$	$\cos \alpha / \sin \beta$
$\xi_{h,d}$	$\cos \alpha / \sin \beta$	$-\sin \alpha / \cos \beta$	$\cos \alpha / \sin \beta$	$-\sin \alpha / \cos \beta$
$\xi_{H,\ell}$	$\sin \alpha / \sin \beta$	$\cos \alpha / \cos \beta$	$\cos \alpha / \cos \beta$	$\sin \alpha / \sin \beta$
$\xi_{H,u}$	$\sin \alpha / \sin \beta$	$\sin \alpha / \sin \beta$	$\sin \alpha / \sin \beta$	$\sin \alpha / \sin \beta$
$\xi_{H,d}$	$\sin \alpha / \sin \beta$	$\cos \alpha / \cos \beta$	$\sin \alpha / \sin \beta$	$\cos \alpha / \cos \beta$
$\xi_{A,\ell}$	$-\cot \beta$	$\tan \beta$	$\tan \beta$	$-\cot \beta$
$\xi_{A,u}$	$\cot \beta$	$\cot \beta$	$\cot \beta$	$\cot \beta$
$\xi_{A,d}$	$-\cot \beta$	$\tan \beta$	$-\cot \beta$	$\tan \beta$

Table 3.2: Yukawa coupling strength for the neutral h , H and A in the four 2HDM models. The notation $\ell = e, \mu, \tau$, $u = u, c, t$ and $d = d, s, b$ is used.

One region of interest is the decoupling limit $\alpha \rightarrow \beta - \pi/2$, such that $\sin(\beta - \alpha) = 1$. This is the region of parameter space in which the h -boson has SM-like EW couplings and can be identified with the 125 GeV SM Higgs boson. Under this assumption, the H -boson does not interact with the W or Z at tree-level, given that $\cos(\alpha - \beta) = 0$. The resulting effective theory is then similar to the SM Higgs sector, with small corrections to the various couplings due to the heavy sector. In this case, the couplings to fermions shown in Table 3.2 are simplified by using the following equivalence:

$$\alpha \rightarrow \beta - \pi/2 \Rightarrow \begin{cases} \sin \alpha = -\cos \beta \\ \cos \alpha = \sin \beta \end{cases} \quad (3.18)$$

The summary of the Yukawa coupling strengths in the decoupling limit are shown in Table 3.3.

3. Physics beyond the Standard Model

Coupling	Type I	Type II	Type III	Type IV
$\xi_{h,\ell}$	1	1	1	1
$\xi_{h,u}$	1	1	1	1
$\xi_{h,d}$	1	1	1	1
$\xi_{H,\ell}$	$-\cot \beta$	$\tan \beta$	$\tan \beta$	$-\cot \beta$
$\xi_{H,u}$	$-\cot \beta$	$-\cot \beta$	$-\cot \beta$	$-\cot \beta$
$\xi_{H,d}$	$-\cot \beta$	$\tan \beta$	$-\cot \beta$	$\tan \beta$
$\xi_{A,\ell}$	$-\cot \beta$	$\tan \beta$	$\tan \beta$	$-\cot \beta$
$\xi_{A,u}$	$\cot \beta$	$\cot \beta$	$\cot \beta$	$\cot \beta$
$\xi_{A,d}$	$-\cot \beta$	$\tan \beta$	$-\cot \beta$	$\tan \beta$

Table 3.3: Yukawa coupling strength for the neutral h , H and A in the four 2HDM models at the decoupling limit $\alpha \rightarrow \beta - \pi/2$. The notation $\ell = e, \mu, \tau$, $u = u, c, t$ and $d = d, s, b$ is used.

3.2.3 Two Higgs doublet model + pseudoscalar

One variation to the 2HDM is the 2HDM+ a [66], where a is a pseudoscalar singlet. This theory is usually preferred in dark matter (DM) models, where the a boson acts as a BSM pseudoscalar mediator between the SM and some hypothetical fermionic DM candidate and mixes with the pseudoscalar partner of the SM Higgs boson.

A simple formulation of the 2HDM+ a is derived starting from the tree-level potential of the 2HDM:

$$\begin{aligned}
 V_{2\text{HDM}}(\Phi_1, \Phi_2) = & \mu_1^2 \Phi_1^\dagger \Phi_1 + \mu_2^2 \Phi_2^\dagger \Phi_2 - \mu_3^2 (\Phi_1^\dagger \Phi_2 + \text{h.c.}) \\
 & + \frac{1}{2} \lambda_1 (\Phi_1^\dagger \Phi_1)^2 + \frac{1}{2} \lambda_2 (\Phi_2^\dagger \Phi_2)^2 + \lambda_3 (\Phi_1^\dagger \Phi_1) (\Phi_2^\dagger \Phi_2) \\
 & + \lambda_4 (\Phi_1^\dagger \Phi_2) (\Phi_2^\dagger \Phi_1) + \frac{1}{2} \lambda_5 \left[(\Phi_1^\dagger \Phi_2)^2 + \text{h.c.} \right],
 \end{aligned} \tag{3.19}$$

and adding a CP-odd mediator, P , that mixes with the CP-odd state A from Equation 3.19. This can be achieved by considering the following interaction terms:

$$V_P = \frac{1}{2} m_P^2 P^2 + P(ib_P \Phi_1^\dagger \Phi_2 + \text{h.c.}) + P^2(\lambda_{P,1} \Phi_1^\dagger \Phi_1 + \lambda_{P,2} \Phi_2^\dagger \Phi_2), \tag{3.20}$$

where m_P and b_P are parameters with dimensions of mass. b_P is assumed to be real so that CP is conserved. As a consequence, P does not develop a vacuum expectation value.

The interactions in $V_{2\text{HDM}}$ mix the neutral CP-even eigenstates, which can be related to the mass eigenstates h and H by the rotation angle denoted as α . The portal coupling b_P in V_P mixes the two neutral CP-odd eigenstates with a mixing angle θ , leading to the mass eigenstates A and a . The scalar spectrum also contains two charged mass eigenstates $H^\pm$ of identical mass, related to the mixing angle β . Diagonalising the mass-squared matrices of the scalar states leads to relations between the fundamental parameters in $V_{2\text{HDM}}$ and V_P . In the broken EW phase, the physics of the proposed 2HDM+ a is fully captured by the mixing angles α , β and θ , the EW vacuum expectation

value v , the quartic couplings λ_3 , $\lambda_{P,1}$ and $\lambda_{P,2}$ and the masses m_h , m_H , m_A , m_a and $m_{H^\pm}$. In order to avoid FCNCs at tree-level, different types of 2HDM+ a can be defined, the same way it is done in the previously described 2HDM.

In this model, a hypothetical fermionic DM particle χ can be coupled to the SM via the pseudoscalar mediator P as follows:

$$\mathcal{L}_\chi = -y_\chi P \bar{\chi} \gamma^5 \chi, \quad (3.21)$$

which introduces two additional parameters to the analysis, the DM coupling to P , y_χ , and its mass, m_χ .

In the decoupling limit, where $\alpha \rightarrow \beta - \pi/2$ and the h -boson has SM-like couplings to the EW bosons, the most relevant interactions with fermions are given by the following expression:

$$\begin{aligned} \mathcal{L}_{2\text{HDM}+a}^{\text{Yukawa}} \ni & -\frac{m_t}{v} \bar{t} \left[h + \xi_t^M H + \xi_t^M i\gamma^5 (A \cos \theta - a \sin \theta) \right] t \\ & -\frac{m_b}{v} \bar{b} \left[h + \xi_b^M H - \xi_b^M i\gamma^5 (A \cos \theta - a \sin \theta) \right] b \\ & -\frac{m_t}{v} V_{tb} \xi_t^M H^+ \bar{t}_R b_L - \frac{m_b}{v} V_{tb} \xi_b^M H^+ \bar{t}_L b_R + \text{h.c.} \\ & -y_\chi (A \sin \theta + a \cos \theta) \bar{\chi} (i\gamma^5) \chi, \end{aligned} \quad (3.22)$$

where m_f/v are the SM Yukawa couplings and V_{ij} are the elements of the CKM matrix. The coupling modifiers ξ_f^M encode the information about the 2HDM type, with $M = \text{I, II, III or IV}$. The $t\bar{t}a$, $a \rightarrow b\bar{b}$ search described in this thesis is based on the type II 2HDM+ a , and therefore:

$$\xi_t^{\text{II}} = -\cot \beta, \quad \xi_b^{\text{II}} = \tan \beta. \quad (3.23)$$

The model is further simplified by assuming that the H , A and $H^\pm$ bosons, as well as the DM candidate χ , are too massive to be observed in the LHC². Table 3.4 summarises the values of the other free parameters of the model that are used in the $t\bar{t}a$, $a \rightarrow b\bar{b}$ MC simulation.

Parameter	Value
$\tan \beta$	1
$\sin \theta$	0.7
λ_3	3
$\lambda_{P,1}$	3
$\lambda_{P,2}$	3
y_χ	1

Table 3.4: 2HDM+ a parameters used in the $t\bar{t}a$, $a \rightarrow b\bar{b}$ analysis.

This results in a very simplified Yukawa lagrangian for the $t\bar{t}a$, $a \rightarrow b\bar{b}$ model [67]:

$$\mathcal{L}_{t\bar{t}a, a \rightarrow b\bar{b}}^{\text{Yukawa}} = -\frac{g_t y_t}{\sqrt{2}} a \bar{t} (i\gamma^5) t - \frac{g_b y_b}{\sqrt{2}} a \bar{b} (i\gamma^5) b, \quad (3.24)$$

²In MC simulation, their mass is set to 10^5 GeV.

3. Physics beyond the Standard Model

where $y_i/\sqrt{2} = m_i/v$ is the SM-like Yukawa coupling of particle i to the a -boson and g_i is the coupling modifier. According to the values in Table 3.4, $g_t = g_b \simeq 0.5$. Only the $a \rightarrow b\bar{b}$ decay is considered in the ATLAS search, since it has the largest branching ratio according to this particular model. A comparison of various branching ratios for the dominant $a \rightarrow XX$ decays as a function of the a -boson mass is shown in Figure 3.1.




Figure 3.1: $\text{BR}(a \rightarrow XX)$ with $X = b, c, \tau, g$ for $\tan \beta = 1$, $\cos \theta = 0.7$ and $m_\chi \gg m_a$ in the type II 2HDM+ a as a function of the a -boson mass. The couplings of the a -boson to WW and ZZ are zero due to CP conservation. The mass of the light CP-even state h , equivalent to the SM Higgs boson, is marked in grey.

3.3 Simplified models for LHC Physics

Simplified models [68] are formulations of a hypothetical BSM theory involving only a few particles and interactions. Many simplified models result from taking low-energy limits of more general new physics scenarios where all but a few particles are integrated out. They are usually parametrised by observables from collider experiments, such as particle masses or branching ratios, which are related to the couplings between different objects. Equation 3.24 is an example of a simplified model derived from the 2HDM+ a , which is often used in DM studies, and currently used in ALP searches in ATLAS. The primary applications of this type of models are identifying the boundaries of experimental sensitivity, characterising new Physics signals and derive limits on more general models.

For the $H \rightarrow aa \rightarrow 4b$ search described in this thesis, only three parameters are needed in order to build a reference model: the mass of the new particle, denoted as m_a , and its couplings to the SM Higgs boson, λ_a , and the b -quark, y_b . The associated BSM lagrangian reads as follows:

$$\mathcal{L}_{\text{BSM}} \ni \frac{1}{2}(\partial_\mu a)(\partial^\mu a) - \frac{1}{2}m_a^2 a^2 - \frac{1}{2}\lambda_a a^2 H - y_b a \bar{b}(i\gamma^5)b, \quad (3.25)$$

where a is a pseudoscalar singlet. In general, it is easier to consider each simplified model with branching ratios set to 1. Then, models with multiple decay modes can be studied by taking linear combinations of results from single decay modes. For the MC simulation of the $H \rightarrow aa \rightarrow 4b$ process, a $\text{BR}(a \rightarrow b\bar{b}) = 1$ is assumed.

Part II

Experimental setup

Chapter 4

The ATLAS experiment

Situated on the Franco-Swiss border near Geneva, the European Organization for Nuclear Research (CERN), stands as a global epicenter for particle physics exploration. Since its inception, CERN has been at the forefront of unraveling the mysteries of the universe, driven by a collective ambition to understand the fundamental building blocks of matter and the forces that govern them. Founded in 1952, CERN has been home to many groundbreaking experiments, which have lead to exciting scientific discoveries, such as the first observation of the W and Z bosons (1983), the determination of the number of neutrino families (1989) and the more recent discovery of the SM Higgs boson (2012) at the Large Hadron Collider (LHC), which is, to this day, the world's largest and most powerful particle accelerator. The ATLAS detector (A Toroidal LHC ApparatuS) is one of the four main experiments located at the LHC, designed to record the high-energy proton-proton collisions, which take place at a rate of over a billion interactions per second. The work in this thesis is based on proton-proton collisions collected by this experiment, at a center of mass energy of 13 TeV, during the Run 2 period (2015-2018).

4.1 The Large Hadron Collider

The LHC is a circular particle accelerator with a circumference of 27 km, situated an average of 100 m below ground. It comprises two separate beam pipes, where particles are accelerated in opposite directions up to very high energies, using a combination of magnetic fields and radiofrequency cavities, before they collide in each of the four interaction points. The LHC primarily collides bunches of protons, but it can also accelerate beams of heavy-ions (such as lead), which are used to study quark-gluon plasma.

4.1.1 LHC layout

Before entering the LHC ring, proton bunches are extracted and accelerated gradually using a series of machines, shown in Figure 4.1. The first step to obtain proton bunches is to extract them from a hydrogen source using strong electrical fields, which break down the atoms into its constituent protons and electrons. Protons are then injected into the Linear Accelerator 2 (LINAC 2), where they reach an energy of 50 MeV. Next, they enter the first circular accelerator, the Proton Synchrotron Booster, where they are accelerated to 1.4 GeV and prepared to enter the Proton Synchrotron (PS), which pushes the beam up to 26 GeV. The next step in the acceleration chain is the Super Proton Synchrotron (SPS), where protons reach 450 GeV. Finally, they are transferred to the LHC beam pipes, where half of the bunches circulate clockwise and the other half, anti-clockwise. Here they are given the final boost where they reach their maximum

4. The ATLAS experiment

energy, 6.5 TeV. This is equivalent to a centre-of-mass energy of 13 TeV in the final proton-proton collision¹.




Figure 4.1: Layout of the CERN accelerator complex [69].

A second injection chain exists for heavy ions. In this case, ions are removed from the atoms inside a plasma and injected into the Linear Accelerator 3 (LINAC 3), where they are accelerated up to 4.2 MeV. They continue their path through the Low Energy Ion Ring (LEIR), where they are further accelerated to 72 MeV, compatible with injection into the PS. Once they enter the PS, the acceleration chain is the same as for the protons.

Colliding two beams of equally charged particles requires opposite magnet dipole fields in both beams. This is why within the LHC tunnel there are two beam pipes, with separate magnet fields and vacuum chambers and common sections only at the intersection points where the experimental detectors are located. Radiofrequency (RF) cavities situated along the beam pipes provide additional energy boosts to accelerate the particles to their desired collision energy. There are a total of 16 RF cavities in the LHC, 8 per beam. Each cavity is powered by a 2 MV electric field, which oscillates at a frequency of 400 MHz, synchronised with the circulating proton bunches. Charged particles passing through the cavity feel the overall force and direction of the resulting electromagnetic field, which transfers energy to push them forwards along the accelerator. Maximum energy is reached in around 15 minutes, the bunches having passed the cavities around 1 million times. These cavities are built from copper coated

¹In Run 3 (2022 – now), the centre-of-mass energy has been increased to 13.6 TeV.

in superconducting niobium and operate at 4.5 K. Figure 4.2 shows a schematic view of a RF cavity, and how it operates.




Figure 4.2: Diagram showing the basic functioning of the LHC RF cavities.

The LHC operates using superconducting magnets, which guide the particle beams along their circular trajectory. These magnets use superconducting coils made from a niobium-titanium alloy, which allows them to operate at very high magnetic fields with minimal electrical resistance. They are chilled to temperatures close to absolute zero (1.9 K) using liquid helium to achieve superconductivity, minimizing energy loss. There are two types of magnets in the LHC, displayed in Figure 4.3. Dipole magnets are used to bend the paths of charged particles so that they follow a circular trajectory. Quadrupole magnets are used for focusing the particle beams. Unlike dipole magnets, quadrupole magnets act like lenses, converging or diverging the particles to keep them tightly packed.




Figure 4.3: Diagram showing the cross section of a LHC dipole (left) and quadrupole (right) magnet. Adapted from [70, 71].

4. The ATLAS experiment

4.1.2 LHC performance in Run 2

LHC Run 2 corresponds to the data taking period from 2015 to 2018, during which about 10^{16} proton-proton collisions (events) were delivered. Due to the statistical nature of LHC data analyses, this number of collisions is key for making precise measurements of known particles and their interactions, as well as for searching for new particles and phenomena that may occur at very low probabilities. The number of events for a given process can be calculated as follows:

$$N_{\text{events}} = \sigma L, \quad (4.1)$$

where σ is the corresponding cross section and L is the integrated luminosity, calculated as the integral over time of the instantaneous luminosity. The cross section for a given process depends on the physics of such process, and typically increases with the centre-of-mass energy ($\sqrt{s}$), as it can be seen in Figure 4.4.




Figure 4.4: Cross section versus centre-of-mass energy for different processes in the LHC. The dashed vertical lines indicate the 7 and 8 TeV energies used in LHC Run 1, the 13 TeV energy used in LHC Run 2, and the 100 TeV goal for future colliders. Adapted from [72].

The integrated luminosity, L , depends on the accelerator, and is calculated as:

$$L = \int dt \mathcal{L}, \text{ with } \mathcal{L} = \frac{N_b N_1 N_2 f}{4\pi \sigma_x \sigma_y}. \quad (4.2)$$

L represents the total luminosity over a period of time, that is, the number of proton-proton collisions per unit area. $\mathcal{L}$ denotes the instantaneous luminosity, and it is

computed from the number of colliding bunches (N_b), the number of protons per bunch (N_1 and N_2), the frequency of the accelerator (f) and the transverse beam sizes in the x and y directions (σ_x, σ_y). Each bunch crossing can produce numerous simultaneous interactions, leading to a complex environment where signals from different collisions overlap. The average number of collisions per bunch crossing is known as pileup, and denoted as $\langle\mu\rangle$. Figure 4.5 (a) shows the delivered luminosity by LHC during Run 2, which reached 140 fb^{-1} [73]. The pileup distribution in the ATLAS detector is shown in Figure 4.5 (b), averaging 33.7 interactions per bunch crossing for the whole Run 2 data-taking period.




Figure 4.5: (a) Cumulative integrated luminosity at a centre-of-mass energy of 13 TeV versus time during Run 2. (b) Luminosity-weighted distribution of the mean number of interactions per crossing for the 2015–2018 proton-proton collision data [74].

4.2 The ATLAS Experiment

ATLAS (A Toroidal LHC Apparatus) is one of the two general-purpose detectors at the LHC, and is installed in a cavern 100 m below ground near the main CERN site. It is the largest detector ever constructed for a particle collider, with 46 metres long and 25 metres in diameter. The data taken by the ATLAS Experiment is used to investigate a wide range of SM and BSM physics by measuring the proton-proton collisions taking place at the centre of the detector. Several different detecting subsystems arranged in layers around the interaction point record the paths, momentum, and energy of the particles, allowing them to be individually identified. In addition, a large magnet system allows to determine their electrical charge. Not all interactions that take place in the detector are relevant for the physics programme, which is why the trigger and data acquisition systems select only the most interesting collision events to study. Figure 4.6 shows the different components of the ATLAS detector, which are described in more detail throughout this section.

4. The ATLAS experiment




Figure 4.6: The ATLAS detector at CERN [75].

4.2.1 Coordinate system

All ATLAS analyses use a common coordinate system that is Lorentz invariant. The origin is set at the interaction point (IP), in the centre of the detector, and the z -axis runs along the beamline anti-clockwise. The part of the detector that corresponds to positive z values is called the A-side, while the part that corresponds to negative values is known as the C-side. The x -axis points from the IP to the center of the LHC and the y -axis points upwards. These two axes constitute the transverse plane. The projection of the momentum vector over the transverse plane, also called transverse momentum (p_T), is often used in physics analyses:

$$p_T = \sqrt{p_x^2 + p_y^2}, \quad (4.3)$$

where $p_{x,y}$ are the x, y components of the momentum $\vec{p}$ of a certain particle. Given that proton-proton collisions occur along the z -axis, the total momentum in the transverse plane is 0. This allows to define the so-called missing transverse energy (E_T^{miss}), which accounts for the energy of "invisible" particles, that is, those that do not interact with the detector (e.g. neutrinos):

$$E_T^{\text{miss}} = - \sum p_T^{\text{visible particles}}. \quad (4.4)$$

In addition, spherical coordinates are usually employed to determine the position of particles. The polar angle θ is the angle between the z -axis and the position of the particle. It is usually replaced by the pseudorapidity (η) which is defined as follows:

$$\eta = -\ln \left[\tan \left(\frac{\theta}{2} \right) \right]. \quad (4.5)$$

In the limit where the speed of the particle is close to the speed of light, or equivalently, when the mass of the particle is much smaller than its energy, the pseudorapidity becomes a good approximation of rapidity (y):

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right) \xrightarrow{m \ll E} \frac{1}{2} \ln \left(\frac{|\mathbf{p}| + p_z}{|\mathbf{p}| - p_z} \right) = \frac{1}{2} \ln \left(\frac{1 + \cos \theta}{1 - \cos \theta} \right) = -\ln \left[\tan \left(\frac{\theta}{2} \right) \right]. \quad (4.6)$$

This quantity is very useful in collider physics, given that differences in rapidity between particles are Lorentz invariant under boosts along the longitudinal axis. The azimuthal angle is denoted as ϕ and it represents the angle between the x -axis and the $x - y$ projection of a point. It ranges from $-\pi$ to π . Using these two angles, it is possible to define the angular distance between two objects in the $\eta - \phi$ plane (ΔR):

$$\Delta R = \sqrt{(\Delta \eta)^2 + (\Delta \phi)^2}, \quad (4.7)$$

where $\Delta \eta$ and $\Delta \phi$ represent the difference in η and ϕ between two objects.

A schematic representation of the ATLAS coordinate system is shown in Figure 4.7.




Figure 4.7: Coordinate system used in the ATLAS Experiment. Adapted from [76].

4.2.2 Inner Detector

The inner detector (ID) [77–79] is the closest sub-detector to the IP, located only a few centimetres away from the beampipe. The ID measures the direction, momentum, and charge of electrically-charged particles produced in proton-proton collisions with very high precision. It has three main components: the pixel detector, the semiconductor tracker (SCT), and the transition radiation tracker (TRT). Together, they provide a coverage within $|\eta| < 2.5$ for any charged particle with $p_T > 0.5$ GeV. Figure 4.8 shows a schematic view of the different subsystems in the ATLAS ID.

Pixel detector

The pixel detector [80, 81] is designed to provide a very high-granularity, high-precision set of measurements as close to the interaction point as possible. It consists of four

4. The ATLAS experiment




Figure 4.8: Schematic view of a quarter section of the ATLAS ID [80].

barrels with a radius of approximately 33.5, 50.5, 88.5 and 122.5 mm each, and a set of three disks on each side, at a distance of 495, 580 and 650 mm from the IP along the z -axis, respectively. There are over 2000 sensor modules, each of them containing 16 units with 2880 pixels, read out by a common electronic chip. The pixels are made of silicon, a semiconductor material that produces a measurable electric current when a charged particle passes through it. Each of them has an area of $50 \times 250 \mu\text{m}^2$ in the innermost barrel layer (also known as Insertable B-layer or IBL) and $50 \times 400 \mu\text{m}^2$ in the rest of the detector. The pixel provides tracking information with a resolution of $8 \times 40 \mu\text{m}^2$ in the IBL and $10 \times 115 \mu\text{m}^2$ in the remaining layers, in $r\phi \times z$.

Semiconductor tracker

The SCT [82] is a precise silicon microstrip detector, which extends the tracking volume to radial distances of $299 < r < 560$ mm. It consists of four layers of silicon strip sensors in the barrel and nine disks in each of the endcaps. In total, it contains around 2112 barrel sensor modules and 2×988 endcap sensor modules, each of them consisting of four silicon sensors. Each sensor has 768 silicon strips with a length of 6 cm and a pitch of $80 \mu\text{m}$. They are arranged in pairs to form a 12 cm strip. A second pair of identical sensors is glued back-to-back with the first pair, rotated by 40 mrad , enabling z -position measurement. Overall, the SCT can measure charged particle tracks with a resolution of $17 \times 580 \mu\text{m}^2$ in $r\phi \times z$.

Transition radiation tracker

The TRT [83] is the outermost part of the ID. It consists of 370000 thin-walled straws, also known as drift tubes. The barrel section of the TRT covers $560 < r < 1080$ mm

and $|z| < 712$ mm and has the straws aligned parallel to the direction of the beam axis. The two endcap sections cover $617 < r < 1106$ mm and $827 < |z| < 2744$ mm and have the straws arranged radially in wheels. Each straw is about 4 mm in diameter and up to 1.44 m long, and is filled with a gas mixture of xenon (70%), carbon dioxide (27%) and oxygen (3%) at room temperature. Inside, there is a gold-plated tungsten wire with a difference in potential with respect to the surface of the tube of 1.5 kV. When a charged particle traverses a straw it ionises the gas inside and the resulting electrons drift towards the central wire, thus generating an electrical signal. The TRT provides tracking information only in the $r\phi$ direction, as the tubes are parallel to the beam line. The resolution in the barrel and endcap is 120 and 130 μm , respectively.

Magnetic field

The ID is immersed in a 2 T magnetic field parallel to the beamline, generated by a central solenoid, which extends over a length of 5.3 m with a diameter of 2.5 m. By measuring the curvature of the charged particle tracks caused by this field, their momentum and sign of their electrical charge can be determined.

4.2.3 Calorimeters

The next layer of the ATLAS detector is the calorimeter system. The calorimeters provide precise measurements of the total energy carried by both charged and neutral particles generated in the proton-proton collisions. There are two types of calorimeters in the ATLAS detector, electromagnetic and hadronic. Electromagnetic calorimeters measure the energy of electrons and photons as they produce an electromagnetic particle cascade in the detector material. Hadronic calorimeters sample the energy of hadrons as they interact with matter via the strong force, producing a hadronic particle cascade. The ATLAS calorimeter system is divided in two sub-detectors: the liquid argon (LAr) [84, 85] calorimeter and the tile calorimeter (TileCal) [86–89]. Both of them are sampling calorimeters, which consist of alternating layers of an absorber, a dense material used to degrade the energy of the incident particle, and an active medium that provides the detectable signal. Figure 4.9 shows the layout of the ATLAS calorimeters in the $r - z$ plane.

Liquid argon calorimeter

The ATLAS LAr calorimeter is a sampling calorimeter consisting of electromagnetic and hadronic sections with cryogenically-cooled (-184°C) liquid argon as the active medium. The LAr barrel and the LAr EM endcap calorimeters use lead plates as the absorber material, with electrodes arranged in an accordion-like structure. They cover the pseudorapidity range $|\eta| < 1.475$ and $1.375 < |\eta| < 3.2$, respectively. When an incoming electron or photon hits the lead absorber, an electromagnetic shower is produced in the liquid argon and the resulting current is collected by a set of electrodes. The high density of lead enhances the probability of interactions like bremsstrahlung and pair production, while ensuring that electromagnetic showers develop over a short distance. This leads to a compact shower shape, which is easier to contain and measure accurately within a limited volume. The forward calorimeters are built in copper and tungsten, with rod electrode structures, and cover the range $3.1 < |\eta| < 4.9$.

4. The ATLAS experiment




Figure 4.9: Schematic view of a quarter section of the ATLAS calorimeters. Adapted from [90].

The LAr hadronic endcap consists of parallel copper plate electrodes, extending in the pseudorapidity region $1.5 < |\eta| < 3.2$. Hadrons interact with the nuclei of copper atoms, initiating complex cascades of secondary particles. Copper is a more suitable material for hadronic calorimetry because it is less dense than lead, providing better separation of the individual hadronic interactions within the shower.

The time resolution of the LAr calorimeter can be parametrised as a function of the energy of the particle, E :

$$\sigma_t = \frac{p_0}{E} + p_1. \quad (4.8)$$

For $E \lesssim 30$ GeV, $p_0 \approx 2.0$ ns GeV and $p_1 \approx 0.2$ ns. For $E \gtrsim 30$ GeV, $p_0 \approx 3.6$ ns GeV and $p_1 \approx 0.2$ ns. This amounts for a time resolution of $O(0.1$ ns). Similarly, the energy resolution can be parametrised as:

$$\frac{\sigma_E}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c. \quad (4.9)$$

For the LAr EM calorimeters, $a \sim 10\%$ is the stochastic term, $b \sim 0.4$ GeV is the noise term and $c \sim 0.7\%$ is a constant term that accounts for systematic uncertainties such as calibration errors, non-uniformity in the detector response, and other non-statistical effects. In the endcap hadronic calorimeters, the b/E term becomes negligible, and the energy resolution is determined by $a \sim 60\%$ and $c \sim 2\%$.

Tile calorimeter

The ATLAS TileCal consists of tiles of plastic scintillator as active material and low-carbon steel absorber plates. It surrounds the LAr calorimeter, providing coverage in a

radius from 2.28 m to 4.23 m. The calorimeter is divided into three segments along the z -axis, having one central barrel section that extends up to $|\eta| = 1.0$ and two extended barrel sections on either side of the central barrel that cover $0.8 < |\eta| < 1.7$. Each TileCal barrel consists of 64 modules in the angular direction ϕ allowing full azimuth coverage. Each module is composed of alternating layers of steel and scintillator tiles with a thickness of 14 mm and 3 mm, respectively. The scintillator tiles are arranged in a staggered pattern to optimise coverage and minimise gaps. Light produced in the scintillators by passing particles is collected by wavelength-shifting fibers and directed to photomultiplier tubes located at the outer edges of the modules. Each module contains several hundred tiles, with a total of about 9852 readout channels in the entire TileCal system.

Similarly to the LAr calorimeter, the time and energy resolution of the ATLAS TileCal can also be parametrised as a function of the energy:

$$\sigma_t = \frac{p_0}{\sqrt{E}} + \frac{p_1}{E}. \quad (4.10)$$

For $E \lesssim 30$ GeV, $p_0 \approx 1.5$ ns GeV^{1/2} and $p_1 \approx 0.4$ ns GeV. For $E \gtrsim 30$ GeV, $p_0 \approx 1.7$ ns GeV^{1/2} and $p_1 \approx 17$ ns. The time resolution is therefore $\mathcal{O}(0.1$ ns). In the same manner, the energy resolution is parametrised as:

$$\frac{\sigma_E}{E} = \frac{a}{\sqrt{E}} \oplus c, \quad (4.11)$$

where $a \sim 50\%$ and $c \sim 2\%$.

4.2.4 Muon spectrometer

The muon spectrometer (MS) sits on the outer layer of the ATLAS experiment. It consists of monitored drift tubes (MDTs) and Cathode Strip Chambers (CSCs) for precision tracking and Resistive Plate Chambers (RPCs) and Thin Gap Chambers (TGCs) for triggering in the barrel and endcap, respectively. The barrel chambers form three cylinders at radii of about 5, 7.5, and 10 m. They cover the pseudorapidity range $|\eta| < 1$. The endcap chambers cover the range $1 < |\eta| < 2.7$ and are arranged in four disks at distances of approximately 7, 10, 14, and 21 m from the interaction point, concentric with the beam axis. Figure 4.10 shows a schematic view of the ATLAS MS.

Muon tracking

The precise measurement of muon tracks is done in the $r - z$ plane, parallel to the magnetic field's bending direction. The z -coordinate is measured in the barrel region, while the r -coordinate is measured in the transition and endcap regions. This is primarily achieved using MDT chambers, which cover most of the spectrometer's solid angle. MDT chambers offer a single-wire resolution of about 80 μm when operated at a high gas pressure of 3 bar, and they ensure reliable operation due to the mechanical isolation of neighbouring wires. In the first station of the endcap region and for pseudorapidities $|\eta| > 2$, CSCs are used. CSCs provide finer granularity to handle the high rate and background conditions in this area.

4. The ATLAS experiment




Figure 4.10: Schematic view of a quarter section of the ATLAS muon spectrometer [91].

Muon triggering

RPCs are used primarily in the barrel region of the MS, covering the pseudorapidity range $|\eta| < 1.05$. They consist of two parallel resistive plates, made of bakelite, separated by an insulating gas (tetrafluoroethane, $C_2H_2F_4$). The signal is read out via capacitive coupling by metal strips on both sides of the detector. When a muon passes through the gas, it ionises the gas molecules. The resulting ions are amplified by a uniform electric field of 4.5 kV/mm, creating an avalanche that induces a signal on the readout strips.

TGCs are utilised in the endcap regions of the MS, covering the pseudorapidity range $1.05 < |\eta| < 2.7$. They consist of an array of parallel wires, which act as electrodes, placed within a gas filled chamber (55% CO_2 and 45% n -pentane, $n-C_5H_{12}$). Each TGC has a cathode-cathode distance (gas gap) of 2.8 mm, a wire pitch of 1.8 mm, and a wire diameter of 50 μm . The operating voltage is 3.1 kV. When a muon crosses a TGC, it ionises the gas within the chamber, and the resulting electrons are collected by the closely spaced wires. Signals from the anode wires, arranged parallel to the MDT wires, provide the trigger information together with readout strips arranged orthogonal to the wires.

Both these subsystems are required to have a time resolution below the LHC bunch spacing of 25 ns and a granularity of approximately 1 cm. Additionally, measurement of the second coordinate in a direction orthogonal to the one measured in the precision chambers must have a resolution between 5 and 10 mm. In total, the RPCs and the TGCs cover an area of 3650 m^2 in the barrel and 2900 m^2 in the endcap region. The total number of readout channels is about 350000 for the barrel and 440000 in the endcaps.

Magnetic field

The magnet system consists of 3 sets of air-core superconducting toroids, one for the barrel and two for the endcaps, each with 8 coils. It is designed to produce a large-volume magnetic field covering the pseudorapidity range $0 \leq |\eta| < 2.7$. The barrel toroid coils are each 25×7 m and the endcap coils are 9×4 m. The performance in terms of bending power is characterised by the bending power $\int B dl$, where B is the field component normal to the muon direction and the integral is computed along an infinite-momentum muon trajectory, between the innermost and outermost muon-chamber planes. The barrel toroid provides 1.5 to 5.5 Tm of bending power in the pseudorapidity range $0 \leq |\eta| < 1.4$, while the endcap toroids provide from 1 to 7.5 Tm in the region $1.6 < |\eta| < 2.7$. The bending power is lower in the transition regions where the two magnets overlap ($1.4 < |\eta| < 1.6$). Figure 4.11 shows a schematic view of the ATLAS magnet system, as well as the bending power of the spectrometer magnetic field with respect to the η coordinate, for different values of the azimuthal angle ϕ .




Figure 4.11: Schematic view of the ATLAS magnets. The barrel toroid is shown in red and the endcap toroids are shown in green. The blue cylinder is the central solenoid, responsible for the magnetic field in the ID.

4.3 Trigger and data acquisition

The ATLAS trigger and data acquisition (TDAQ) system is responsible for online processing, selecting and storing events of interest for offline analysis. The TDAQ system processes up to 40 million collision events per second using two layers of trigger systems. The Level-1 (L1) trigger uses custom hardware to reduce the event rate to around 100 kHz, which is the maximum detector readout rate. The high-level trigger (HLT) then further processes these events using software tools, reducing the rate to approximately 1 kHz, which corresponds to a data storage capability of a few hundred MB/s. This selection process ensures that the most interesting collision data is captured while maintaining manageable data volumes. After the event is accepted by the HLT, it is transferred to local storage and exported to a Tier-0 facility at the CERN

4. The ATLAS experiment

computing centre for offline reconstruction, and then moved to permanent storage. Figure 4.12 shows a schematic view of the ATLAS TDAQ system.




Figure 4.12: Schematic of the ATLAS TDAQ system in Run 2 [92].

4.3.1 Level-1 Trigger

The L1 trigger is a hardware-based system. It makes an initial selection based on reduced-granularity information from the calorimeters and the muon detectors, with a latency of $\sim 2.5 \mu\text{s}$. It consists of four main subsystems: the Level-1 Calorimeter (L1Calo) trigger, the Level-1 Muon (L1Muon) trigger, the Level-1 topological (L1Topo) trigger and the central trigger processor (CTP).

Level-1 calorimeter trigger

The L1Calo trigger [93, 94] processes information from the calorimeters in real time and sends trigger signals to the L1Topo processor and the CTP using field programmable gate arrays (FPGAs). Figure 4.13 gives an overview of the L1Calo subsystem from Run 2. It receives signals from sets of multiple cells in the electromagnetic and hadronic calorimeters, called trigger towers, with a $\Delta\eta \times \Delta\phi$ resolution of 0.1×0.1 and 0.2×0.2 , respectively. These signals are sampled in the pre-processor modules at 80 MHz and the resulting digital data is sent to the cluster processor (CP) and jet energy processor (JEP). The CP covers the region $|\eta| < 2.5$ and contains 56 cluster processor modules.

Each of them processes an area of calorimeter data to identify and count energy deposits indicative of isolated e/γ and τ particles. The JEP covers the region $|\eta| < 3.2$ and contains 32 jet/energy modules. Each of them identifies and counts energetic jet candidates and computes transverse energy quantities for the area of calorimeter examined, including the total transverse energy ($\sum E_T$), the missing transverse energy (E_T^{miss}) and the transverse energy significance (XS). Additional signals from the forward calorimeters ($|\eta| > 3.2$) are used to calculate them. The CP and JEP do not actually make a decision on whether to keep the event (Level-1 Accept). The common merger modules (CMX) send the L1Calo object counts to the CTP for this purpose. In addition to real-time trigger processing, the CMX FPGAs provide region of interest (RoI) and readout data to the HLT and DAQ systems, which is transmitted via the existing readout driver (ROD) modules.

The L1Calo subsystem also generates trigger objects (TOBs), which are bit arrays encoding the position (η , ϕ), p_T , and quality flags (e.g. isolation) from the physics objects that are reconstructed in the electromagnetic and hadronic calorimeters. TOBs for the whole η , ϕ range processed by L1Calo are transmitted optically by the CMX to L1Topo, which also receives data from L1Muon.




Figure 4.13: Schematic of the ATLAS L1Calo system in Run 2. ECAL and HCAL refer to the EM and hadronic calorimeters, respectively [94].

Level-1 muon trigger

The L1Muon trigger [91, 93, 94] uses information from the RPC, the TGC and the TileCal. It selects muon candidates, assigns them to the correct LHC bunch crossing and classifies them into one of six p_T threshold classes. Following a Level-1 Accept decision, the p_T thresholds and the corresponding detector regions (RoIs), are sent to the HLT for further consideration. The typical dimensions of the RoIs in $\Delta\eta \times \Delta\phi$ are 0.1×0.1 in the RPC and 0.03×0.03 in the TGC. The L1 trigger decision in the barrel region ($|\eta| < 1.05$) is based on the coincidence of hits from three (two) concentric RPC stations for the three high- p_T (low- p_T) thresholds. In the endcap region ($1.05 < |\eta| < 2.4$), the trigger decision is based on the coincidence of hits in the TGC stations of the middle layer, called the big wheel. The main source of trigger background is low-momentum charged particles emerging from the endcap toroid magnets and beam

4. The ATLAS experiment

shielding, which are suppressed by requiring a coincidence between the big wheel and the TGC forward inner (TGC-FI) chambers. A coincidence between the TGC chambers and the TileCal assists in the rejection of fake muon triggers in the region $1.05 < |\eta| < 1.3$. The muon trigger detectors are grouped into sectors, 64 for the barrel, 96 for the endcap and 48 for the forward region (208 sectors in total). Each sector can provide information on up to two muon candidates (priority is given to the highest- p_T candidates). The information from all the sectors is combined in the muon-to-CTP interface (MUCTPI), which counts the number of muon candidates for each of the six p_T thresholds, and sends the multiplicity information to the CTP. The MUCTPI [95] is responsible for detecting cases where muons traverse more than one sector due to chamber overlaps, making sure that they are counted only once. Overlaps within sectors are handled by the logic specific to the barrel and endcap subsystems. The MUCTPI also interacts with the L1Topo system, sending (η, ϕ) information with a coarser granularity ($\Delta\eta \times \Delta\phi \approx 0.3 \times 0.1$) via the MUCTPI-to-Topo interface. A diagram of the L1Muon trigger data processing is shown in Figure 4.14.




Figure 4.14: Schematic of the ATLAS L1Muon system in Run 2.

Level-1 topological trigger

The L1Topo trigger [94, 96] was introduced in Run 2 in order to maintain the 100 kHz L1 output rate after the luminosity increase with respect to Run 1 while keeping the most interesting physics events. It consists of two modules, each of them containing two processor FPGAs to run the algorithms and one controller FPGA responsible for the readout and communication to the external trigger systems. In total, the four L1Topo processor FPGAs that can run up to 128 algorithms. This allows to perform real-time event selection based on geometric and kinematic relationships between different trigger objects, i.e. electrons/photons, jets, taus and muons, as well as event-level

quantities such as E_T^{miss} . The data format used in L1Topo is called trigger object (TOB). TOBs are bit arrays encoding the energy, position and further qualifying information (e.g. isolation) for each object. The maximum numbers of different TOBs that can be received by the L1Topo system in one event are: 120 EM (e/γ) TOBs, 64 jet TOBs, 120 tau TOBs, 32 muon TOBs and one E_T^{miss} TOB. The total latency of the L1Topo system is ~ 200 ns, 75 of which are dedicated to the execution of the topological algorithms. Here, the first 50 ns are used to filter the input TOBs in order to reduce the number of possible combinations when considering relations among objects. This is achieved by creating two types of filtered TOB lists, by either sorting the TOBs by p_T and taking the first six leading objects or selecting the first ten objects above a certain p_T threshold as ordered in the input source channel and received in the FPGA. The remaining 25 ns in the algorithm execution are used to run the decision algorithms on the desired TOB lists. This process is summarised in Figure 4.15. The list of L1Topo decision algorithms implemented in Run 2 is summarised in Table 4.1.




Figure 4.15: Schematic view of the L1Topo data flow [96].

Central trigger processor

The CTP [94, 97] is responsible for processing trigger signals from sub-detectors and making high-speed decisions to determine whether collision events should be kept for further analysis. It receives input signals from L1Calo, L1Topo, the MUCTPI and various forward detectors and aligns them in time, ensuring that each trigger decision is associated with the correct bunch crossing. The CTP generates a Level-1 Accept (L1A) signal from the trigger inputs according to the Level-1 trigger menu. This menu contains up to 256 trigger items, each of which is a combination of one or more conditions on trigger inputs. The L1A is the logical OR of all trigger items, that is, the event is kept for further processing if it fires one or more triggers from the menu. The CTP is also used for prescaling and dead time control. Prescaling reduces the rate of specific trigger items by accepting only a fraction of the corresponding events, while dead time introduces a recovery period between consecutive trigger acceptances to

4. The ATLAS experiment

L1Topo algorithm	Definition
Pseudorapidity distance	$\Delta\eta_{\min} \leq \Delta\eta = \eta_1 - \eta_2 \leq \Delta\eta_{\max}$
Azimuthal distance	$\Delta\phi_{\min} \leq \Delta\phi = \phi_1 - \phi_2 \leq \Delta\phi_{\max}$
Box cut	$\Delta\eta_{\min} \leq \Delta\eta \leq \Delta\eta_{\max}$ and $\Delta\phi_{\min} \leq \Delta\phi \leq \Delta\phi_{\max}$
Window cut	$\eta_{\min} < \eta < \eta_{\max}$ and $\phi_{\min} < \phi < \phi_{\max}$
Angular distance	$\Delta R_{\min}^2 \geq \Delta R^2 = \Delta\eta^2 + \Delta\phi^2 \leq \Delta R_{\max}^2$
Disambiguation	$\eta_1 \neq \eta_2$ or $\phi_1 \neq \phi_2$ or $\Delta R > \Delta R_{\min}$
Ratio	$f(\text{TOB}_1) \geq \alpha f(\text{TOB}_2)$ with $\eta_1 = \eta_2$, $\phi_1 = \phi_2$ and $\alpha = \text{constant}$
Invariant mass	$m_{\text{inv},\min}^2 \leq m_{\text{inv}}^2 = 2E_T^1 E_T^2 (\cosh\Delta\eta - \cos\Delta\phi) \leq m_{\text{inv},\max}^2$
Transverse mass	$m_{T,\min}^2 \leq m_T^2 = 2E_T^1 E_T^{\text{miss}} (1 - \cos\Delta\phi) \leq m_{T,\max}^2$
Event hardness	$H_{T,\min} < H_T = \sum p_T^{\text{jets}}$
Simple cone	$E_{T,\min} < E_T^{\text{cone}} = \sum_{\Delta R < 1.0} E_T^{\text{jets}}$
Late muon	Finds the highest- p_T μ in the next BC and combines it with the input lists associated with the current BC.

Table 4.1: List of topological algorithms implemented in Run 2. Trigonometric functions are calculated using look-up tables.

prevent a system overload. Additionally, the timing, trigger, and control (TTC) system ensures the synchronisation of signals throughout the ATLAS detector.

4.3.2 High-Level Trigger

The HLT [94, 98, 99] is a software-based subsystem capable of reducing the event rate from 100 kHz (L1 output rate) to approximately 1 kHz within a latency of about 200 ms. It runs on a computer farm consisting of 40000 CPU cores. Each of these cores runs a processing unit (PU), which runs a hypothesis algorithm to decide whether the trigger condition is satisfied or not based on the reconstructed event. The HLT operates by analysing data from specific regions of the detector, known as Regions of Interest (RoIs), which are identified by the L1 trigger and used to run a partial event reconstruction with full detector granularity. Events that are accepted after partial reconstruction are then reconstructed in the event builder (EB) using the full event data stored in the readout buffers. Similarly to the L1 trigger, a trigger menu is used to define the set of conditions that are required to keep the event. Once an event is accepted by the HLT, the sub-farm output (SFO) sends the data to permanent storage for offline reconstruction and exports the data to the Tier-0 facility [100] at CERN's computing centre. The SFO system provides up to 48 hours of temporary storage for accepted events. Background jobs copy event files to permanent storage locations and delete them from the local disk only when they are safely on tape.

Chapter 5

Physics simulation of proton-proton collisions

Event simulation is a fundamental tool for analysing and interpreting data at the LHC. It allows to model and predict the outcomes of particle collisions, facilitating the understanding of complex processes and the comparison between theoretical predictions and experimental results. The simulation of proton-proton collisions involves physics processes at very different energy scales, from the high energy deep-inelastic scattering between the partons in the protons, to the low energy final state hadrons. While a fully analytic calculation of these events is not possible, they can be factorised into different energy regimes. The simulation of the hard interaction can be computed up to a fixed order in perturbation theory, while the description of the softer scales can be done with phenomenological models. In order to obtain a simulated physics model that is directly comparable with measured data, interactions of the final-state particles with the detector are also simulated.

5.1 Event simulation

A proton-proton collision at the LHC can be broken down into several stages, each representing different aspects of the physics involved. Figure 5.1 shows a schematic view of the process.




Figure 5.1: Schematic view of a LHC proton-proton collision. Adapted from [101].

The hard scattering is the main interaction, where two partons (quarks or gluons) from the colliding protons interact, typically involving a high momentum transfer. These interactions are described by perturbative QCD, which allows the use of Feynman diagrams to calculate the cross sections of the different processes. After the hard

5. Physics simulation of proton-proton collisions

scattering, the high energy quarks and gluons produced radiate additional QCD particles as they move away from each other, due to the asymptotic freedom of the strong force. This process is called parton shower, and it can be seen as a cascade of emissions where quarks and gluons radiate more quarks and gluons, which continue until the energy of the particles is too low for further radiation. During both the hard scattering and parton shower stages, charged particles can also emit photons via QED bremsstrahlung. When the energy of the quarks and gluons reaches the non-perturbative QCD scale, $\Lambda_{\text{QCD}} \sim 1 \text{ GeV}$, they hadronise, which means that they combine to form colour-neutral hadrons, such as pions, kaons, protons, and neutrons. The fact that this process is non-perturbative means that it has to be described using phenomenological models. Once the first hadrons are produced, many of them are unstable and decay into lighter, more stable particles. These are usually the particles that interact with the detector. The initial partons can radiate gluons and photons before the hard scattering, producing what is called initial state radiation (ISR). Once the partons have undergone the primary interaction, they continue to radiate gluons and photons, generating final state radiation (FSR) as well. In addition, multiple proton-proton collisions can occur in the same bunch crossing, overlaying additional interactions on top of the primary collision. These additional collisions are known as pileup. Overall, a typical LHC event spans a wide range of energy scales and needs to be modelled using a combination of perturbative and non-perturbative models [102].

5.1.1 Hard scattering

The cross section for a hard scattering process initiated by two protons with four-momenta P_1 and P_2 , $\sigma_{pp \rightarrow X}$, can be computed using the collinear factorisation [103] equation:

$$\sigma_{pp \rightarrow X} = \sum_{i,j} \int_0^1 dx_1 dx_2 f_i(x_1, \mu_F^2) f_j(x_2, \mu_F^2) \hat{\sigma}_{ij \rightarrow X} \left(x_1 P_1, x_2 P_2, \frac{Q^2}{\mu_F^2}, \frac{Q^2}{\mu_R^2} \right), \quad (5.1)$$

where the sum runs over the parton types that can initiate the process ($g, u, \bar{u}, d, \bar{d}, s, \bar{s}, \dots$). In this equation, $f_i(x_j, \mu_F^2)$ is the parton distribution function (PDFs) of the proton. It describes the probability to find a parton of species i with momentum fraction x_j when a proton is probed at an energy scale μ_F . The factorisation scale μ_F is an arbitrary parameter that represents the boundary between high- and low-energy physics. A parton emitted with a four-momentum $q^2 < \mu_F^2$ is considered part of the hadron structure and absorbed into the PDF. A parton emitted with a four-momentum $q^2 > \mu_F^2$ is included in the parton-level cross section. μ_F^2 is typically chosen to be the same value as the hard scale Q^2 , which is the energy scale of the parton-parton interaction. Figure 5.2 offers a visual representation of Equation 5.1.

PDFs play a central role in event generators, for the simulation of hard processes, parton showers and multiple parton interactions. Therefore, the choice of PDF set influences both cross sections and event shapes. These distributions can not be predicted from first principles, since they depend on the non-perturbative physics of the proton wave function. They are primarily determined from data in global fits and extrapolated to




Figure 5.2: Feynman diagram representing the various parameters used in Equation 5.1.

higher energy scales using the Dokshitzer-Gribov-Lipatov-Altarelli-Parisi (DGLAP) equations [104–106], which describe how PDFs evolve with the energy scale μ_F in the perturbative regime, $\alpha_s \ll 1$ [107]:

$$\frac{d}{d \ln \mu_F} \begin{pmatrix} q(x, \mu_F^2) \\ g(x, \mu_F^2) \end{pmatrix} = \frac{\alpha_s(\mu_F)}{2\pi} \begin{pmatrix} P_{q \leftarrow q}(z) & P_{q \leftarrow g}(z) \\ P_{g \leftarrow q}(z) & P_{g \leftarrow g}(z) \end{pmatrix} \otimes \begin{pmatrix} q(x, \mu_F^2) \\ g(x, \mu_F^2) \end{pmatrix}, \quad (5.2)$$

where $q(x, \mu_F^2)$ and $g(x, \mu_F^2)$ are PDFs of the quarks and the gluon, respectively, and $P_{j \leftarrow i}$ are the splitting functions that describe the evolution from parton i to parton j due to initial state radiation as a function of the momentum fraction z (Figure 5.3). Dedicated collaborations such as NNPDF, CTEQ and MSTW provide PDFs for physics analyses [108]. Figure 5.4 shows the NNPDF3.1 PDF set for the different proton partons and two different values of the factorisation scale.




Figure 5.3: Graphic representation of the DGLAP splitting functions.

The PDFs are combined with the parton-level cross section for the production of the final state X through the initial partons i and j , denoted as $\hat{\sigma}_{ij \rightarrow X}$. This cross section is computed at a fixed order in perturbation theory, which makes it dependent on the factorisation scale μ_F and the renormalisation scale μ_R . Both μ_F and μ_R are non-observable quantities, and dependence on them decreases at higher order in perturbation theory. They are usually chosen to have the same value, $\mu_F^2 = \mu_R^2 = Q^2$. The computation of the parton-level cross section for a general process $ij \rightarrow X$ is related to the matrix element amplitude as [110]:

$$\hat{\sigma}_{ij \rightarrow X} \propto \sum_{k=0}^{\infty} \int d\Phi_{X+k} \left| \sum_{l=0}^{\infty} \mathcal{M}_{X+k}^{(l)} \right|^2. \quad (5.3)$$

The sum over k represents a sum over additional real-emission corrections, also called legs, and the sum over l runs over additional virtual corrections, that is, loops. By

5. Physics simulation of proton-proton collisions




Figure 5.4: Global fit by the NNPDF Collaboration. PDFs are shown at factorisation scales of 10 GeV^2 (left) and 10^4 GeV^2 (right). PDFs evaluated at values of μ_F are connected by the DGLAP evolution equations [109].

assigning different values to k and l , one can obtain different levels of precision in the resulting cross section¹:

- $k = 0, l = 0 \Rightarrow \text{LO inclusive } ij \rightarrow X \text{ production.}$
- $k = N, l = 0 \Rightarrow \text{LO } ij \rightarrow X + N \text{ jets production.}$
- $k + l \leq p \Rightarrow \text{N}^p\text{LO inclusive } ij \rightarrow X \text{ production.}$
It includes N^{p-1}LO for $ij \rightarrow X + 1 \text{ jet}$, N^{p-2}LO for $ij \rightarrow X + 2 \text{ jets}$, etc.

In Equation 5.3, $d\Phi_{X+k}$ is the differential phase space element over the final state particles and $\mathcal{M}_{X+k}^{(l)}$ can be computed from the corresponding Feynman diagram with k legs and l loops. As long as one is dealing with infrared (IR) safe observables², the LO prediction only needs to include tree-level Feynman diagrams avoiding kinematic regions that contain soft and/or collinear emissions. The only difficulty then is to integrate over a multidimensional phase space, a computation that can easily lead to convergence problems for a large number of final-state particles. In these cases, the fastest converging numerical integration algorithm is Monte Carlo integration [111], which has a purely stochastic error $\propto O(1/\mathcal{N})$, with $\mathcal{N}$ being the number of generated points. This is an advantage over any other algorithms, where the error scales with powers of dimension. Therefore, virtually all numerical cross section

¹LO = leading order, NLO = next-to-leading order, N^pLO = p -times next-to-leading order.

²IR safe observables are quantities that are not affected by adding any number of infinitely soft particles or splitting an existing particle up into two comoving particles with arbitrary momentum fractions z and $1 - z$, respectively. If both conditions are satisfied, non-perturbative effects from soft and collinear contributions are negligible.

calculations are based on Monte Carlo techniques in one form or another, reason why event simulation is commonly referred to as Monte Carlo simulation. At NLO and beyond, that is, when $l > 0$, the virtual amplitudes from the extra loops are divergent for any point in phase space. However, the KLN theorem [112, 113] establishes that the IR divergences originating from virtual amplitudes³ cancel with those coming from soft and/or collinear real emission, order by order, making the matrix element computation finite. A numerical implementation of this method presents some difficulty, since they are related to phase spaces of different dimensionality. Equation 5.4 illustrates a common procedure for dealing with this calculation:

$$\begin{aligned}\sigma^{\text{NLO}} &= \int_N d\sigma^{\text{LO}} + \int_N d\sigma^{\mathcal{V}} + \int_{N+1} d\sigma^{\mathcal{R}} \\ &= \int_N d\sigma^{\text{LO}} + \int_N \left[d\sigma^{\mathcal{V}} + \int_1 d\sigma^{\mathcal{A}} \right] + \int_{N+1} \left[d\sigma^{\mathcal{R}} - d\sigma^{\mathcal{A}} \right],\end{aligned}\tag{5.4}$$

where the NLO cross section for a process with N particles in the final state is the combination of the LO cross section plus the 1-loop virtual contribution ($\mathcal{V}$) and the additional emission of 1 jet ($\mathcal{R}$). The general idea of the subtraction method is to use the following identity [114]:

$$d\sigma^{\text{NLO}} = d\sigma^{\text{LO}} + (d\sigma^{\mathcal{V}} + d\sigma^{\mathcal{A}}) + (d\sigma^{\mathcal{R}} - d\sigma^{\mathcal{A}}),\tag{5.5}$$

where the counterterm $d\sigma^{\mathcal{A}}$ must fulfil a series of conditions:

- It has to be soft and collinear, so that it does not modify an IR safe observable.
- It has to cancel the singular behaviour of $d\sigma^{\mathcal{R}}$.
- It has to be integrable analytically in $4 - 2\epsilon$ dimensions so that the result can be easily added to the loop diagram and cancel its IR divergences, which manifest as double poles of the type $1/\epsilon^2$ with $\epsilon \rightarrow 0^+$.

5.1.2 Parton shower

Fixed-order calculations can accurately describe the hard scattering process of an LHC event, but fail to reproduce the effects of soft and collinear emissions by the initial- and final-state particles. The effect of all higher orders can be simulated through a parton shower algorithm [115], which is an iterative algorithm based on the successive random generation of gluon emissions and gluon splittings ($g \rightarrow qq$), from the high-energy scales associated with the hard process to the low-energy scales where QCD enters the non-perturbative regime. At this point, partons are bound by confinement into colour-neutral hadrons and can not be separately resolved.

The parton shower algorithm is developed based on the general formula for the emission probability in the collinear approximation. For any matrix element calculation

³UV divergences occur when $p \rightarrow \infty$, and are regularised, usually by dimensional regularisation, before renormalisation. They appear as poles of the type $1/\epsilon$ with $\epsilon \rightarrow 0^+$, and can be cancelled by using the adequate counterterms.

5. Physics simulation of proton-proton collisions

that produces partons of any flavour i , the cross section for the additional emission of a parton j with momentum fraction z is given by [102]:

$$d\sigma \approx \sigma_{\text{ME}} \sum_i \frac{\alpha_s}{2\pi} \frac{d\theta^2}{\theta^2} dz P_{j \leftarrow i}(z), \quad (5.6)$$

where $\theta \ll 1$ is the splitting angle between the two partons, i and j and $P_{j \leftarrow i}(z)$ are the DGLAP splitting functions shown in Figure 5.3. This is the building block of the parton shower algorithm. Since this is a completely general expression for any hard process to be accompanied by a collinear splitting, it can be iterated, using it on the hard process to generate one collinear splitting and then treating the final state of that splitting as a new hard process, generating an even more collinear splitting from it, and so on. It only remains to discuss what happens when $\theta \rightarrow 0$. The point here is that any physical measurement can not distinguish an exactly collinear pair of partons from a single parton with the same total momentum and other quantum numbers, and therefore this divergence is associated with a process that can not be detected. A resolution criterion is introduced, which is equivalent to say that only distributions of resolvable partons are taken into account. This can be done in many different ways, as any variable proportional to θ^2 yields a mathematically equivalent expression:

$$t^2 \propto \theta^2 \Rightarrow \frac{dt^2}{t^2} = \frac{d\theta^2}{\theta^2}. \quad (5.7)$$

Using this general formulation, one can write the total probability for a parton i of emitting a parton j between t and $t + dt$ in the soft and collinear limit as:

$$d\mathcal{P}_i(t) = \frac{\alpha_s}{2\pi} \frac{dt^2}{t^2} \int_{t_0^2/t^2}^{1-t_0^2/t^2} dz P_{j \leftarrow i}(z). \quad (5.8)$$

In order to construct the probability distribution of the first branching, i.e. the one that yields the largest t^2 , one needs to calculate the probability of no emission above a certain t^2 , $\overline{\mathcal{P}}_i(t_{\text{max}}^2, t^2)$, for $t^2 \leq t_{\text{max}}^2$. This is done by using the multiplication rule of probability:

$$\begin{aligned} \overline{\mathcal{P}}_i(t_{\text{max}}^2, t^2) &= \lim_{N \rightarrow \infty} \prod_{i=0}^{N-1} \overline{\mathcal{P}}_i(t_{i+1}^2, t_i^2) \\ &= \lim_{N \rightarrow \infty} \prod_{i=0}^{N-1} [1 - \mathcal{P}_i(t_{i+1}^2, t_i^2)] \\ &= \exp \left[- \lim_{N \rightarrow \infty} \mathcal{P}_i(t_{i+1}^2, t_i^2) \right] \\ &= \exp \left[- \int_{t^2}^{t_{\text{max}}^2} d\mathcal{P}_i(\tilde{t}) \right], \end{aligned} \quad (5.9)$$

where $\overline{\mathcal{P}}_i + \mathcal{P}_i = 1$ is assumed. Substituting $d\mathcal{P}_i$ by its definition according to Equation 5.8, one can obtain the final form of the probability of no emission of a

parton j above t^2 from a parton i , also known as de Sudakov form factor, and usually denoted as $\Lambda_i(t_{\max}^2, t^2)$:

$$\Lambda_i(t_{\max}^2, t^2) = \exp \left[\int_{t^2}^{t_{\max}^2} \frac{\alpha_s}{2\pi} \frac{d\tilde{t}^2}{\tilde{t}^2} \int_{t_0^2/\tilde{t}^2}^{1-t_0^2/\tilde{t}^2} dz P_{j \leftarrow i}(z) \right]. \quad (5.10)$$

The Monte Carlo implementation of this equation is then quite straightforward: a random number $\rho \in [0, 1]$ is chosen, and the equation $\Lambda_i(t_{\max}^2, t^2) = \rho$ is solved for t^2 . If $t_0^2 < t^2 \leq t_{\max}^2$, a resolvable branching is generated at scale t^2 , and $t_{\max}^2 = t^2$ is used as the scale for the next branching. Otherwise, the evolution terminates. It is important to note that the strong coupling α_s increases as the parton shower evolves and the transverse momentum of the emitted partons becomes smaller. This also implies that the cutoff value t_0 must be well above the non-perturbative limit, $\Lambda_{\text{QCD}} \sim 1 \text{ GeV}$, since the parton shower algorithm is derived from QCD perturbation theory. The ordering variable t can vary depending on which parton shower algorithm is used. Some of them use what is known as p_T -ordering, meaning that harder QCD emissions are produced first, while others use angular ordering instead.

There are several different strategies for combining matrix element calculations and parton showers such that any double-counting or gap in the parameter space is avoided. Two widely-used methods are the Catani-Krauss-Kuhn-Webber (CKKW) algorithm [116] and the M. L. Mangano (MLM) algorithm [117]. The CKKW algorithm works by generating partonic events using matrix elements for hard emissions, then reclustering the final state using the k_T algorithm [118] to form a jet history. At each emission vertex, Sudakov form factors are applied to account for the probability of no additional radiation between scales, ensuring soft emissions are handled by the parton shower. The event is reweighted and passed to the parton shower, which generates softer emissions below a specific resolution scale. The MLM algorithm, on the other hand, relies on generating the parton level event using matrix elements for hard emissions up to $n = 0, 1, \dots, N$ jets and the corresponding parton showers, then reclustering all partons resulting from the shower evolution using a cone jet algorithm with radius R_{cluster} and associating each parton from the matrix element calculation to one and only one of the reconstructed jets. Starting from the highest- p_T parton, a jet is matched when $\Delta R(\text{jet}, \text{parton}) < R_{\text{cluster}}$. If $n < N$, the event is accepted when all the partons are matched to a jet and there are no extra jets. If $n = N$, the event is accepted if $N_{\text{jets}} > N$ as well, provided the non-matched jets are softer than all the matched ones. After matching, the exclusive samples with $n = 0, 1, \dots, N - 1$ jets are combined with the inclusive $n = N$ sample.

In addition to QCD parton showers, electromagnetic radiation, such as photon emission from quarks, is included using a similar showering algorithm based on QED.

5.1.3 Hadronisation and decay

Hadronisation occurs when the partons originating from the proton-proton collision enter the non-perturbative regime, corresponding to energies below $\Lambda_{\text{QCD}} < 1 \text{ GeV}$. At such low energies, the QCD coupling α_s becomes very large, causing what is called

5. Physics simulation of proton-proton collisions

QCD confinement, meaning that all coloured particles combine to form colour-neutral states, also known as hadrons. Because hadronisation is a non-perturbative process, phenomenological models are used to describe the transition from partons to hadrons in event simulations. The two most commonly used models for hadronisation in Monte Carlo event generators are the string model and the cluster model. A schematic approach to both procedures is depicted in Figure 5.5.




Figure 5.5: Schematic representations of (a) the string hadronisation model and (b) the cluster hadronisation model [119].

The string model [120] is based on the observation that at large distances the potential energy of colour sources increases linearly with their separation, similar to a string. As these particles move apart, the string stretches and its potential energy grows at the expense of the partons' kinetic energy. When the potential energy becomes of the order of hadron masses, it becomes energetically favourable for the string to break at some point along its length, leading to the creation of a new quark-antiquark pair. The two string segments then begin to stretch and break again, and so on until all the energy has been converted into quark-antiquark pairs connected by short string segments, which can be identified with hadrons. The pattern of the hadronisation depends on the colour structure of the system. Each parton in the system has a unique colour partner, connected to it by a string segment that stretches and breaks as described above in its rest frame. Gluons produce a kink on the string, which becomes sharper at higher momentum values. These factors determine the spatial distribution for hadron production.

The cluster model [121] is based on the idea of preconfinement, where the colour-singlet nature of the system is restored before hadronisation occurs. Preconfinement dictates that at evolution scales much smaller than the hard subprocess scale, $q \ll Q$, the partons in a shower are clustered in colourless groups with an invariant mass distribution that is independent of the nature and scale of the hard subprocess, depending only on q and Λ_{QCD} . It is then natural to identify these clusters at the hadronisation scale Q_0 as proto-hadrons that decay into the observed final-state

hadrons. In practice, gluons, which are colour-octets, are split into quark-antiquark pairs before clustering. Comparing this method with the string model, one can see that here the string is always broken at a gluon, rather than just having a kink.

Many of the primary hadrons generated by these models are unstable and decay further at various time scales. Hadron decays are typically implemented as separate algorithms, which include information about the hadron lifetime and branching ratios. These parameters are determined experimentally and collected in the Review of Particle Physics [122].

5.1.4 Underlying events

In a proton-proton collision, many different subprocesses contribute to the total observed activity. The underlying event (UE) represents the additional activity that is not directly associated with the primary interaction. It can receive contributions from initial- and final-state radiation, from the QCD evolution of colour connections between the hard scattering and the beam-proton remnants, and from additional hard scatters in the same proton-proton collision, also known as multiple parton interactions. Because of the low energy scale of these processes, phenomenological models have to be used where the parameters are tuned based on experimental data [123]. It is impossible to uniquely separate the UE from the hard scattering process on an event-by-event basis, but observables can be defined which are particularly sensitive to the properties of the underlying event. Minimum-bias events are used for these studies.

5.2 Detector simulation

The final output of a Monte Carlo generator is a list of four-vectors of all stable particles produced in the event, after hadronisation and decay of the intermediate unstable particles. This output is used in order to study the physics processes at the particle level but, in order to compare it with the experimentally recorded data, the simulation needs to include detector effects. In the ATLAS detector, a detailed simulation is implemented to obtain an event format that is identical to that of the true detector [124]. The simulation programme is integrated into the ATLAS software framework, Athena [125], and uses the GEANT4 simulation toolkit [126].

Figure 5.6 shows all the steps involved in the simulation of the ATLAS detector. The Monte Carlo generator produces events in HepMC format [127]. These events can be filtered at generation time so that only events with the desired properties are kept. The entire connected tree of the resulting HepMC event record is stored as the MC Truth, which is the input to the simulation algorithm. Particles are then propagated through the ATLAS detector simulation using GEANT4. Particles resulting from decays within the detector or interaction with its materials are added to the MC Truth and their energy depositions in the detector, also known as hits, are stored in a hits file.

The hits are then digitised, that is, they are converted into digits that serve as input to the RODs from the detector electronics. Then, the RODs are emulated and the output is stored in a Raw Data Object (RDO), which is later sent to the reconstruction software.

5. Physics simulation of proton-proton collisions

The output from the ATLAS detector itself is converted from bytestream format to RDO format, so that it can be reconstructed in a similar manner. The hits are also mapped to the MC Truth particles that deposited the energy, producing Simulated Data Objects (SDOs), which are fed to the reconstruction algorithm so that they can be used for validation.




Figure 5.6: Schematic representation of the ATLAS simulation software, from event generation (top left) to event reconstruction (top right). Algorithms are placed in rectangular boxes and persistent data objects are placed in rounded boxes. The optional pileup portion of the chain, used only when events are overlaid, is shown in blue. Adapted from [124].

Large computing resources are required to accurately model the complex detector geometry and physics descriptions in the standard ATLAS detector simulation. This has led to the development of several varieties of fast simulation. In this context, the standard GEANT4 simulation is usually referred to as full simulation (FS). Almost 80% of the full simulation time with GEANT4 is spent simulating the progression of particle showers traversing the calorimeter. The ATLAS fast simulation II (AFII), used in this thesis, uses the FastCaloSim algorithm [128] to speed up this part of the simulation by replacing low-energy electromagnetic particles in the calorimeter with presimulated showers stored in memory as libraries. This simulation is optimal for studies that require very high statistics but do not need the level of precision from the FS.

Chapter 6

Object reconstruction

The ATLAS detector records data from proton-proton collisions by measuring signals in its various subdetectors, for example, tracks in the inner detector, energy deposits in the calorimeters or hits in the muon chambers. Object reconstruction is the method of converting raw data from the detector into meaningful physical objects that can be analysed, such as electrons, muons, jets or missing transverse energy (E_T^{miss}). This section describes the reconstruction algorithms for all the objects used in this thesis.

6.1 Tracks and vertices

The ATLAS ID provides position measurements for charged particles in the pseudorapidity range $|\eta| < 2.5$ by combining information from its three subdetectors (pixel, SCT and TRT) and the inner magnetic field. Tracks [129] are reconstructed from clusters of energy deposits in the pixel and SCT detectors. Seeds are generated from triplets of clusters using a loose compatibility criterion to ensure high reconstruction efficiency. A Kalman filter algorithm [130] is then used to build track candidates from the chosen seeds by incorporating additional space-points from the remaining layers of the pixel and SCT detectors that are compatible with the preliminary trajectory. The filter creates multiple track candidates per seed if more than one compatible space-point extension exists on the same layer. In order to solve the ambiguity, tracks are required to fulfill a set of quality criteria based on p_T , $|\eta|$ and their signature pattern on the pixel and SCT detectors. Then, a high-resolution fit is performed using all the available information. The TRT, which is the most external layer of the ATLAS ID, provides a complementary measurement to the silicon-based pixel and SCT detectors, improving the p_T resolution of the tracks [131].

Tracks are used to reconstruct vertices based on a pattern recognition algorithm [132]. The primary vertex, which denotes the origin point of the hard-scattering interaction, is of particular importance. The input to the vertex reconstruction is a collection of reconstructed tracks that must pass a set of quality requirements based on their p_T and their pixel and SCT hits pattern. A seed position for the vertex is selected, and then the tracks and the seed are used to fit the best vertex position. This fit follows an iterative procedure where less compatible tracks are down-weighted in each iteration and the vertex position is recomputed. After the vertex position is determined, tracks that are incompatible with the vertex are removed from it and allowed to be used in the determination of another vertex. Vertices are required to have at least two associated tracks. The procedure is repeated until no unassociated tracks are left in the event or no additional vertex is found in the remaining set of tracks.

6.2 Leptons

6.2.1 Electrons

Electron candidates are reconstructed from energy deposits (clusters) in the electromagnetic calorimeter associated to reconstructed tracks in the ID [133]. The reconstruction starts with the identification of topological clusters, that is, energy deposits in topologically connected calorimeter cells which have a high signal to noise ratio. Next, tracks are reconstructed from hits in the silicon detectors using pattern recognition algorithms and then matched to topological clusters. Following track matching, the reconstruction algorithm generates superclusters, which are dynamic, variable-size clusters that change in size as needed to recover energy from bremsstrahlung photons. They consist of one topological cluster (seed) and its satellite clusters. The seed is required to have $E_T \geq 1$ GeV and be matched to a track with at least four hits in the silicon tracking detectors. After the electron superclusters are built, an energy calibration and position correction is applied to them, and tracks are matched to electron superclusters in the same way as they were matched to the original clusters.

Electron candidates are required to satisfy $p_T > 10$ GeV and $|\eta| < 2.47$, excluding the calorimeter transition region $1.37 < |\eta| < 1.52$. They must also satisfy a tight likelihood-based identification criterion and a multivariate isolation criteria based on the transverse momentum of calorimeter cluster and tracks around the electron, as well as the properties of highly displaced tracks in the same region. They are further required to have $|z_0 \sin \theta| < 0.5$ mm and $|d_0/\sigma(d_0)| < 5$, where z_0 is the longitudinal impact parameter relative to the primary vertex and d_0 , with uncertainty $\sigma(d_0)$, is the transverse impact parameter relative to the beam line.

6.2.2 Muons

Muon candidates are reconstructed from track segments in the various layers of the MS and matched with tracks from the ID [134]. The reconstruction of tracks in the MS starts with the identification of short local track segments reconstructed from hits in an individual MS station. Segments in the different stations are combined into preliminary track candidates using a parabolic trajectory that constitutes a first-order approximation to the muon bending in the magnetic field. Information from precision measurements in the bending plane is combined with measurements of the second coordinate from the trigger detectors to create three-dimensional track candidates. Finally, a global χ^2 fit of the muon trajectory through the magnetic field is performed, taking into account the effects of possible interactions in the detector material as well as the effects of possible misalignments between the different detector chambers. Combined muons are identified by matching MS tracks to ID tracks and performing a combined track fit based on the ID and MS hits, taking into account the energy loss in the calorimeters. Based on the particle trajectory from the combined fit, the muon spectrometer hits associated with the track may again be updated and the track fit repeated.

The final muon candidates are required to satisfy $p_T > 10$ GeV and $|\eta| < 2.5$. They must also pass the Medium quality requirements and analogous multivariate isolation criteria as described for electrons. Muons are also required to satisfy $|z_0 \sin \theta| < 0.5$ mm and $|d_0/\sigma(d_0)| < 4$.

6.3 Jets

Quarks and gluons produced in a proton-proton collision can not exist as free particles due to colour confinement. Instead, they form jets, which are the collimated streams of particles that arise from the hadronisation of high-energy partons. Jet reconstruction aims to obtain a physics object whose characteristics are as close as possible to those of the initial parton.

6.3.1 General aspects of jet reconstruction

Jet reconstruction starts in the hadronic calorimeter with the so-called topological clusters, which are formed by applying a clustering algorithm to the calorimeter cells in order to identify and group together nearby energy deposits likely to originate from the same hadronic shower. The measurements from the topological clusters are improved by making more complete use of the information from both the tracking and calorimeter systems. The particle flow (PFlow) algorithm [135], introduced in Run 2, combines measurements from both the tracker and the calorimeter to form the input signals for jet reconstruction. First, well-measured tracks are selected following a set of quality criteria. The algorithm then attempts to match each track to a single topological cluster in the calorimeter. The expected energy in the calorimeter, deposited by the particle that also created the track, is computed based on the topological cluster position and the track momentum. It is relatively common for a single particle to deposit energy in multiple topological clusters, so for each track-cluster system, the algorithm evaluates the probability that the particle energy was deposited in more than one topological cluster and decides if it is necessary to add more topological clusters to the track-cluster system to recover the full shower energy. The expected energy deposited in the calorimeter by the particle that produced the track is subtracted cell by cell from the set of matched topological clusters. If the remaining energy in the system is consistent with the expected shower fluctuations of a single particle's signal, the topological cluster remnants are removed. This procedure is applied to tracks sorted in descending p_T -order, firstly to the cases where only a single topological cluster is matched to the track, and then to the other selected tracks.

Improved calorimeter signals using the PFlow algorithm are then passed as inputs to the jet clustering algorithms. Sequential clustering algorithms [136] are a common choice for jet reconstruction. They assume that particles within jets have small differences in transverse momentum, grouping particles based on momentum space. All sequential clustering algorithms have a similar method. First, the distance between particles i and j is computed as:

$$d_{ij} = \min \left(p_{T,i}^{2p}, p_{T,j}^{2p} \right) \frac{\Delta_{ij}^2}{R}, \quad (6.1)$$

6. Object reconstruction

where $\Delta_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$ represents the distance between particles i and j in the $y - \phi$ plane, R is the radius parameter which determines the final size of the jet and p is a parameter whose value depends on the algorithm of choice. The second distance variable is d_{iB} , which is the momentum space distance between the beam axis and particle i :

$$d_{iB} = p_{T,i}^{2p}. \quad (6.2)$$

Sequential clustering algorithms work by first finding the minimum distance in the entire set $\{d_{ij}, d_{iB}\}$. If d_{ij} is the minimum, then particles i and j are combined into one using summation of four-vectors, after which i and j are removed from the list of particles. If d_{iB} is the minimum, then i is labelled a final jet and removed from the list of particles. This process is repeated until either all particles are part of a jet with the distance between the jet axes $R_{ij} > R$ (inclusive clustering) or until a desired amount of jets is found (exclusive clustering).

For $p = 1$ one obtains the k_T algorithm [137]. The case of $p = 0$ corresponds to the inclusive Cambridge/Aachen algorithm [138, 139]. In ATLAS, the usual choice for jet clustering is the anti- k_T algorithm [140], which has $p = -1$. Using a negative p value favours the clustering of hard particles first, given that the distance d_{ij} will be smaller when the p_T of the particles involved is larger. The behaviour of different jet algorithms is illustrated in Figure 6.1. The anti- k_T algorithm is preferred for reconstructing jets because it clusters high- p_T particles first, producing jets with well-defined, circular shapes of radius R that are less sensitive to soft radiation or pileup. The k_T and Cambridge/Aachen algorithms, on the other hand, cluster soft and/or collinear particles first, making them better suited for resolving jet substructure, as they naturally break down jets into smaller components, revealing details of its constituents, if any.

6.3.2 Small- R jets

Small- R jets are reconstructed by clustering PFlow objects using the anti- k_T algorithm with a radius parameter $R = 0.4$ and a four-momentum recombination scheme, i.e. the kinematic variables of the jet are given by direct addition of the four-momenta of its individual massless constituents. The energy of the jet is corrected to the particle level by the application of a jet energy scale calibration derived from 13 TeV data and simulation [141]. Baseline jets are required to have $p_T > 20$ GeV and $|\eta| < 2.5$. They are also required to pass a Tight pileup rejection based on the Jet Vertex Tagger (JVT) [142] score, w_{JVT} . For jets with $p_T \in [20, 60]$ GeV and $|\eta| < 2.4$, $w_{JVT} > 0.5$ is required. Due to the low energy nature of some of the processes described in this work, the corrections on the jet energy centrally derived by the ATLAS Collaboration were extended to cover the low jet p_T range between 15 and 20 GeV. Extra studies were performed for the JVT and b -tagging efficiency measurements down to this transverse momentum range.

6.3.3 Large- R jets

Large- R jets are reclustered with the anti- k_T algorithm by using a larger radius parameter of $R = 0.8$ to form the Ak8 reclustered jets. Setting the distance parameter




Figure 6.1: Jet clustering comparison performed on the same data with the same input radius, $R = 1$. Results using (a) the k_T algorithm, (b) the Cambridge/Aachen algorithm and (c) the anti- k_T algorithm are shown. Different colours are used to represent different jets and their areas in the $y - \phi$ plane. Adapted from [140].

to $R = 0.8$ when reconstructing the reclustered jets gives a larger radius for track association, thus potentially allowing more tracks from the targeted double b -hadron decays to be associated to a Ak_8 reclustered jet. The tracks in and around the small- R jet are selected with a Loose track selection [143] and they are associated to a given reclustered jet through ghost association [144–146]. In this method, tracks are treated as infinitesimally soft particles, by setting their p_T to a value close to zero. These tracks are then added to the list of inputs for jet finding. The low p_T scale means that the tracks do not influence the reconstruction of the Ak_4 PFlow jet. However, after jet finding, it is possible to identify which tracks are clustered into which subjects. In the following, the constituent Ak_4 PFlow jets will be denoted as isolated jets and the collection of ghost-associated tracks will be referred to as the Ak_8 track jet.

Additional substructure information is obtained by reconstructing the two $Exk_T^{(2)}$ track subjects inside the Ak_8 track jet. These track subjects are derived using the ghost-associated tracks to each Ak_8 jet as inputs to the exclusive- k_T (Exk_T) method [147]. The selected tracks for a given jet are clustered using the k_T algorithm with a distance parameter of $R = 0.8$. The clustering stops when only two (2) track clusters are left. These clusters are used as the $Exk_T^{(2)}$ track subjects associated with a given jet. For signal events, each $Exk_T^{(2)}$ track subject is expected to originate from the decay of one

b -hadron. The associated Ak8 track jet is required to satisfy $|\eta| < 2.0$ to account for the extended radius and the acceptance of the ID. Furthermore, each $\text{Ex}k_{\text{T}}^{(2)}$ track subjet is required to satisfy $p_{\text{T}} > 5$ GeV, obtained from the sum of the tracks' four-momenta. The four-momentum of the Ak8 track jet is defined as the sum of the four-momenta of its $\text{Ex}k_{\text{T}}^{(2)}$ track subjets. The resolution of the invariant mass of the Ak8 jet is further improved by using a deep neural network (DNN) estimate trained using the seed Ak4 PFlow jet, the tracks of the Ak8 jet and non-isolated leptons inside the jet as inputs.

Secondary vertices (SVs) inside the Ak8 jets are reconstructed as well to help the identification of large- R jets containing a pair of b -quarks. SVs in jets have been explored by many flavour tagging algorithms in ATLAS. For the $H \rightarrow aa \rightarrow 4b$ analysis, an algorithm to reconstruct multiple decay vertices was developed by combining the track-cluster-based low- p_{T} vertex tagger (TC-LVT) [148] and the multiple secondary vertex finder algorithms (MSVF) [149]. The TC-LVT algorithm has been developed for soft b -hadron tagging and optimised to reconstruct low- p_{T} b -hadron decays. In this thesis, the clustering algorithm from TC-LVT is used to identify the collection of tracks that may have at least one displaced secondary vertex, and the MSVF algorithm is used afterwards to identify multiple SVs in the track cluster.

6.4 Flavour tagging

The identification of jets resulting from the fragmentation of b -quarks, usually referred to as flavour tagging, is of uttermost importance for analyses with a large number of b -quarks in the final state, such as the ones described in this work. As shown in Figure 6.2, jets originating from a b -quark have a different substructure, which allows to separate them from other types of jets. B -hadrons have a lifetime of ~ 1.5 ps, which is enough for them to travel a few millimetres away from the primary vertex before decaying, producing what is known as a secondary vertex. Often, the b -hadron will also decay semileptonically, providing an additional soft lepton that can be used for jet identification as well.




Figure 6.2: Schematic representation of a b -jet.

6.4.1 *b*-jets

Jets that originate from the decay of one *b*-hadron are identified (*b*-tagged) using the DL1r tagger [150, 151], which is a deep neural network that combines the output quantities of the low-level algorithms such as IP2D, IP3D, SV1 and JetFitter, plus the jet RNNIP output probabilities. In addition, kinematic properties of the jets, namely p_T and $|\eta|$, are included in the training in order to take advantage of the correlations with the other input variables. The IP2D and IP3D algorithms are based on the track impact parameter (IP). The IP2D tagger makes use of the signed transverse impact parameter significance of tracks to construct a discriminating variable, whereas IP3D uses both the signed transverse and signed longitudinal impact parameter significances in a two-dimensional template to account for their correlation. Due to computational limitations, these algorithms do not take into consideration possible correlations among tracks within a jet. This is why the RNNIP was later introduced. RNNIP exploits a recurrent neural network (RNN) to learn track impact-parameter correlations in order to further improve the jet flavour discrimination. The secondary-vertex-tagging algorithm, SV1, reconstructs one displaced secondary vertex in a jet by analysing all tracks associated with it. The topological multi-vertex finding algorithm, JetFitter, exploits the topological structure of weak *b*- and *c*-hadron decays inside the jet and tries to reconstruct the full *b*-hadron decay chain using track information as well.

The DL1r NN has a multidimensional output corresponding to the probabilities for a jet to be a *b*-jet, a *c*-jet or a light-flavour jet. The DL1r tagging discriminant is defined as:

$$D_{\text{DL1r}} = \ln \left(\frac{p_b}{f_c p_c + (1 - f_c) p_l} \right), \quad (6.3)$$

where p_b , p_c and p_l refer to the *b*-, *c*- and light jet probabilities and f_c denotes the effective *c*-jet fraction in the background hypothesis. This value is chosen a posteriori in order to optimise the performance of the algorithm at physics analysis level. Jets with a DL1r score above a certain threshold are then defined as *b*-tagged jets. Figure 6.3 shows the distribution of the DL1r *b*-jet discriminant for light-, *c*- and *b*-jets.




Figure 6.3: Normalised distributions of the DL1r *b*-jet discriminant corresponding to light-, *c*- and *b*-jets for a $t\bar{t}$ test sample using $f_c = 0.018$ [150].

6. Object reconstruction

The 60%, 70%, 77%, and 85% DL1r efficiency working points (WPs) are common in ATLAS physics analyses. In this thesis, pseudo-continuous (PC) b -tagging is used, meaning that five tag-weight bins are defined which correspond to b -jet tagging efficiency intervals of 100–85%, 85–77%, 77–70%, 70–60% and 60–0%. Most of the light flavour jets fall within the 100–85% bin, while the 60–0% bin is enriched in b -jets. All five working points are calibrated and can be used simultaneously. Each WP is assigned an integer value, called b -tagging score, which is used to define relevant variables in the analysis. Table 6.1 summarises the equivalences between b -tagging efficiency bins, WP notation and integer notation.

DL1r ϵ_b	DL1r WP	b -tagging score
100-85%	None	1
85-77%	85	2
77-70%	77	3
70-60%	70	4
60-0%	60	5

Table 6.1: Correspondence between DL1r b -tagging efficiency (ϵ_b) and the associated working points (WP) and b -tagging scores.

The performance of the b -tagging algorithm is characterised by the probability or efficiency of correctly tagging a signal jet, and the probability of mistakenly identifying a background jet, referred to as the mistag rate. Figure 6.4 shows the b -tagging efficiency (ϵ_b) and the associated light- and c -jet rejection rates ($1/\epsilon_{\text{light}}$ and $1/\epsilon_c$, respectively). The mistag rates for light- and c -jets are included in Figure 6.5. As expected, larger efficiency translates into higher acceptance and, consequently, higher mistag rates.




Figure 6.4: DL1r light- and c -flavour rejection with respect to the b -jet efficiency [152].




Figure 6.5: Mistag rates for (a) light- and (b) c-jets for the 85%, 77%, 70% and 60% DL1r b -tagging WPs as a function of the jet p_T [151].

6.4.2 B -jets

The DeXTer algorithm [153] is a double b -tagger based on a deep sets neural network architecture, designed to identify pairs of b -jets that are too close to be resolved and tagged individually using DL1r. DeXTer is designed to perform flavour tagging of reconstructed large- R jets in two transverse momentum ranges: a low p_T range between 20 and 200 GeV and a high p_T one, above 200 GeV. Throughout this thesis, double b -tagging is referred to as B -tagging, and merged double b -jets are denoted as B -jets. The NN architecture uses track information as input, together with additional data from reconstructed secondary vertices and kinematical variables from the Ak4 PFlow jet that seeds the Ak8 candidate. The architecture comprises two separate feed-forward neural networks which serve as the feature extractor for the tracks in the $R = 0.8$ track jet and for the secondary vertices. An additional global feed-forward neural network combines the output of the feature-extraction neural networks with the PFlow jet kinematics to learn correlations and predict the probability for each flavour label.

The DeXTer algorithm is trained to classify reconstructed jets into three different categories based on the input features: a B -jet, a b -jet or a light jet. They are combined in a B -tagging discriminant, defined as:

$$D_{\text{DeXTer}} = \ln \left(\frac{p_B}{(1 - f_b)p_l + f_b p_b} \right), \quad (6.4)$$

where p_B , p_b and p_l refer to the B -, b - and light jet probabilities and f_b is the effective b -jet fraction, which can be optimised for background rejection. The separation between the three categories is shown in Figure 6.6 (a).

Two working points are defined: the 0-40% tagging interval (Tight) and the 40-60% tagging interval (Loose). Similarly to the DL1r tagger, both working points are calibrated and can be used simultaneously. One desirable feature of the tagger is to be independent of the parent particle mass. In order to minimise this correlation,

and obtain a consistent WP definition, the B -enriched sample used to train DeXter contains $H \rightarrow aa \rightarrow 4b$ and $t\bar{t}a, a \rightarrow b\bar{b}$ samples with different values of the a -boson mass. The results are shown in Figure 6.6 (b). The B -labeled jet efficiency on the two reference working points differs by approximately 2-3% across the examined values of m_a . The B -jet efficiency used to define the WPs is obtained from the training with mixed samples.




Figure 6.6: (a) DeXter discriminant distribution evaluated using B -labeled jets from $H \rightarrow aa \rightarrow b\bar{b}b\bar{b}$ and $t\bar{t}a, a \rightarrow b\bar{b}$ samples and b - and light-labeled jets from $t\bar{t}$ samples with $f_b = 0.4$. (b) Performance for different a -boson masses in $H \rightarrow aa \rightarrow 4b$ events. Images from [153].

6.4.3 Soft secondary vertices

The use of calorimeter jets as a starting point for b -tagging sets constraints on the energy of reconstructed b -hadrons. As mentioned previously, standard b -tagging techniques in ATLAS limit their realm of applicability to calorimeter jets with $p_T > 20$ GeV. However, lower- p_T b -hadrons that are not energetic enough to produce a jet can still be tagged based on the presence of a secondary vertex produced by the b -hadron decay.

Soft secondary vertices, here denoted as v , are built from a set of seed tracks with $p_T > 1.5$ GeV and a large transverse impact parameter significance, $d_0/\sigma(d_0) > 0.5$. Track clusters are built around these tracks, and a vertex finding algorithm is run to produce candidate vertices. The four-momentum of each vertex is defined as the sum of the track momenta. The set of soft secondary vertices used in this thesis is selected by requiring $p_T > 3$ GeV and $m_v > 600$ MeV.

6.5 Overlap removal

During object reconstruction, it is possible to identify one detector signal as two different physics objects. Overlap removal ensures that reconstructed objects are correctly categorised based on their dominant characteristics.

First, electrons sharing a track with a muon are removed. To prevent double-counting of electron energy deposits as jets, the closest jet within $\Delta R = 0.2$ of a selected electron is removed. If the nearest jet surviving that selection is within $\Delta R = 0.4$ of the electron, the electron is discarded. To reduce the background from muons from heavy flavour decays inside jets, muon candidates are required to be separated by $\Delta R > 0.4$ from the nearest jet, removing the muon if the jet has at least three associated tracks, and removing the jet otherwise. This avoids an inefficiency for high-energy muons undergoing significant energy loss in the calorimeter. In this work, the removed soft muons are re-added to the jet four-momentum for b -tagged jets, as described in the next section.

Jet overlap removal is done as follows. First, every Ak4 jet is tested to see if it is eligible for being DeXTer tagged. This requires the Ak4 jet to have $p_T > 20$ GeV and be isolated. These jets are then tested by the DeXTer tagger: if a jet passes the Loose working point selection, it is promoted to a B -tagged Ak8 jet and removed from the Ak4 jet collection. If the jet fails the DeXTer Loose working point, it is not promoted to an Ak8 jet and it is kept as an Ak4 jet. Therefore, two jet collections are built: the Ak8 jets, which are all successfully DeXTer tagged, and the Ak4 jets, which are either ineligible for DeXTer or have failed the DeXTer tagging requirements. For the Ak8 jets, the overlap removal is repeated to remove any overlap of the leptons within the $0.4 < R < 0.8$ cone ring.

6.6 μ -in-jet p_T correction

Decays of b -hadrons have a higher likelihood of containing additional leptons in the decay process than light hadrons; around 10% of all b -hadron decays produce a soft muon inside of the resulting jet. The standard procedure, however, removes these muons during the overlap removal. Since the analyses described in this thesis rely on the invariant mass of the b -tagged jets, they are sensitive to this effect. Soft muons reconstructed inside a DL1r b -jet are added to the four-momentum of the respective b -jet. These soft muons are required to have $p_T > 4$ GeV and $|\eta| < 2.5$ and fulfil the Medium soft muon quality requirement. In the case of a DeXTer B -jet, the soft muons are first matched to the $\text{Ex}k_T^{(2)}$ track subsets if $\Delta R(\text{Ex}k_T^{(2)}, \mu) < 0.3$. From the muons that fulfil this requirement, at least the leading two are included for each $\text{Ex}k_T^{(2)}$ track subset, and any muon is only matched once to the closest subset. Then, the four-momentum of the matched soft-muons is added to that of the $\text{Ex}k_T^{(2)}$ track subset.

6.7 Missing transverse energy

Missing transverse momentum, often referred to as missing transverse energy or E_T^{miss} , is defined as the negative sum of the transverse momenta of all reconstructed particles in a collision event. This quantity represents the energy imbalance in the plane perpendicular to the beam axis (the transverse plane) due to invisible particles that do not interact with the detector, such as neutrinos or potential candidates for BSM particles.

E_T^{miss} includes contributions from both the hard event signals, which consist of fully reconstructed and calibrated particles and jets, and the soft event signals, which come from reconstructed charged-particle tracks associated with the primary vertex but not linked to any hard objects. The missing transverse energy is computed as [154]:

$$E_T^{\text{miss}} = \sqrt{(E_x^{\text{miss}})^2 + (E_y^{\text{miss}})^2}, \quad (6.5)$$

where the x, y components are given by:

$$E_{x,y}^{\text{miss}} = - \sum_{i \in \text{hard}} p_{x,y}^i - \sum_{i \in \text{soft}} p_{x,y}^i. \quad (6.6)$$

In the calculation of E_T^{miss} , the contributing objects need to be reconstructed from mutually exclusive detector signals to avoid double-counting. The most commonly used order for the E_T^{miss} reconstruction sequence for the hard contribution starts with electrons, followed by photons, then hadronically decaying τ -leptons, and finally jets. Muons are mainly reconstructed from ID and MS tracks alone, leading to little or no signal overlap with the other reconstructed particles in the calorimeter.

Part III

Analysis tools

Chapter 7

Machine learning

Machine learning (ML) is a field of artificial intelligence that enables computers to learn patterns from data in order to make decisions or predictions. ML is widely used in the ATLAS experiment in particle identification and reconstruction, event triggering and data analysis. This section focuses on the latter, highlighting the ML techniques that are used in this thesis.

7.1 Neural networks

7.1.1 General definition

A neural network (NN) [155] is a model that can recognise patterns and make predictions based on learned information. The basic unit of a neural network is the neuron, which receives input, processes it and produces an output. Neurons are organised in layers with different functions each. There is an input layer that receives the initial data, $\mathbf{x}$. This layer has as many neurons as features, which is the name used to refer to the different inputs. Examples of features in an ATLAS search would be the particle p_T , η or the number of jets in the event. Then, there are one or more hidden layers, where data is transformed. A neural network is called a deep neural network (DNN) when it has a large number of hidden layers. The depth and size of the hidden layers is a free parameter of the NN design, but needs to be optimised, since a too simple or a too complex architecture can lead to undesired results. Finally, there is an output layer that provides the final prediction. It can have one neuron (binary choice classification) or more (multiple choice classification). Figure 7.1 shows the general structure of a NN.

Each layer of the NN is given by a linear combination of the type:

$$\text{Linear}(\mathbf{x}) = \mathbf{x}A^T + \mathbf{b}, \quad (7.1)$$

where $\mathbf{x}$ is the data vector, A^T is the weight matrix, which represents the strength or importance of the connections and $\mathbf{b}$ is the bias, which helps to adjust the output value. These parameters are learned during training, allowing the network to focus on relevant features. An activation function $g(x)$ is introduced to break the linearity of the NN by transforming or deactivating output from a node, allowing it to capture more complex patterns:

$$\text{Layer}(\mathbf{x}) = g[\text{Linear}(\mathbf{x})] = g(\mathbf{x}A^T + \mathbf{b}). \quad (7.2)$$

Commonly used activation functions are shown in Table 7.1.

The choice of input variables or features for the NN plays an important role in its performance. It is common in particle physics to divide these inputs into two

7. Machine learning




Figure 7.1: Schematic representation of the basic constituents of a NN. The input layer receives the initial information $\mathbf{x}$, which is transformed in the hidden layers by a combination of linear and non-linear functions, resulting in the predicted output $\hat{y}$.

ReLU	Leaky ReLU	Sigmoid	SoftMax
$g(x) = \max(0, x)$	$g(x) = \begin{cases} x & \text{if } x > 0 \\ \alpha x & \text{if } x \leq 0 \end{cases}$	$g(x) = \frac{1}{1 + e^{-x}}$	$g(\mathbf{x}) = \frac{e^{x_i}}{\sum_{j=1}^N e^{x_j}}$
Adds non-linearity to the input and hidden layers.	Adds non-linearity to the input and hidden layers avoiding "dead neurons" ($= 0$).	Normalises the NN score in a binary classification problem.	Normalises the NN score in a multiple classification problem.

Table 7.1: Common activation functions used in NN training and possible use cases.

categories: the low-level features and the high-level features. Low-level features are direct measurements or outputs from the object reconstruction that provide basic information about the collision event. In the ATLAS experiment, these can be the energy or momentum of individual particles, their angular position and their track and hit information from the ID and the calorimeter. Low-level variables are typically used in DNNs, which are able to infer patterns and correlations due to their complex internal structure. High-level features are processed kinematic variables created by the user based on their knowledge in order to capture key event properties. These features summarise and interpret the raw data to make important aspects more accessible for machine learning models. They include, for example, invariant masses, E_T^{miss} , angular separations between objects or flavour-tagging scores. In general, they do not require complex NN architectures to achieve a good performance, making the training process faster.

7.1.2 Training methods

Training is the process of exposing a NN to a set of data repeatedly, so that the parameters of the model can be optimised. It can be supervised, where the model

learns from labeled data with known outcomes, or unsupervised, where it finds patterns or groupings in unlabeled data. This section focuses in supervised neural networks, which are the ones used in this thesis.

Training a NN consists in finding the optimal values of the parameters of the model (the weights and the biases) using a given dataset. The input dataset is usually divided in a training dataset, shown to the NN, and a validation dataset, used to test the predictions made by the NN. At the beginning of the training, the weights and biases are initialised with a random value. Then, data is passed through each layer of the network, activating the neurons. This step is known as forward propagation. Here, each neuron calculates a weighted sum of its inputs and applies an activation function. The result is passed to the next layer until the output layer is reached, generating the final prediction. After the forward pass, the NN predictions are compared to the true data labels in the validation dataset using a loss function, which measures the difference between the predictions and the actual values. A common choice of loss function is the minimum squared error (MSE):

$$L_{\text{MSE}}(\hat{\mathbf{y}}, \mathbf{y}) = \frac{1}{N} \sum_{i=1}^N (y_i - \hat{y}_i)^2, \quad (7.3)$$

where $\hat{\mathbf{y}}$ is the predicted output using the training dataset and $\mathbf{y}$ is the test output using the validation dataset, both of them with dimension N . When the output of the model is to be interpreted as a probability between 0 and 1, the binary cross entropy (BCE) loss function is used instead:

$$L_{\text{BCE}}(\hat{\mathbf{y}}, \mathbf{y}) = -\frac{1}{N} \sum_{i=1}^N [y_i \ln(\hat{y}_i) + (1 - y_i) \ln(1 - \hat{y}_i)]. \quad (7.4)$$

The loss function of choice then needs to be minimised. This is done with a method called backpropagation [156], which consists of calculating the gradients of the loss function with respect to each of the parameters of the model, layer by layer in reverse order, by using the chain rule for derivatives. This is possible because the NN output is a set of nested functions, exactly one per layer:

$$\begin{aligned} \hat{\mathbf{y}} &= g_n [\mathbf{x}_n \mathbf{A}_n^T + \mathbf{b}_n] \\ &= g_n [g_{n-1} (\mathbf{x}_{n-1} \mathbf{A}_{n-1}^T + \mathbf{b}_{n-1}) \mathbf{A}_n^T + \mathbf{b}_n] \\ &= g_n [g_{n-1} (g_{n-2} (\dots) \mathbf{A}_{n-1}^T + \mathbf{b}_{n-1}) \mathbf{A}_n^T + \mathbf{b}_n]. \end{aligned} \quad (7.5)$$

Once gradients are computed, an optimisation algorithm is used to adjust the weights and biases of the NN. Two examples of optimisation algorithms are the stochastic gradient descent (SGD) and the adaptative moment estimation (Adam) [157]. The learning rate η is introduced to modulate the size of the steps during optimisation. For example, in the SGD:

$$\theta \rightarrow \theta - \eta \nabla_{\theta} L(\hat{\mathbf{y}}, \mathbf{y}), \quad (7.6)$$

where θ represents a parameter in the model and $L(\hat{\mathbf{y}}, \mathbf{y})$ is the loss function. The value of η needs to be adjusted for each individual model. A too large learning rate would

7. Machine learning

lack precision, while a too small one could lead to convergence problems. Once the parameters are updated, the training dataset is processed through the NN again, so that the model can be further optimised. Each pass over the full dataset constitutes an epoch. In general, the training is stopped when the loss function does not improve by a given amount over a certain number of epochs, also called patience. This is done to avoid overfitting, which refers to the scenario where the NN starts to memorise the training dataset and its predictions are not general anymore.

Preparing the datasets correctly is of uttermost importance to guarantee an optimal training. One easy thing to do that improves convergence is to normalise the inputs such that they have mean $\langle \mathbf{z} \rangle = 0$ and standard deviation $\sigma = 1$. This is achieved by using the following redefinition:

$$z'_i = \frac{z_i - \langle \mathbf{z} \rangle}{\sigma} = \frac{z_i - \langle \mathbf{z} \rangle}{\sqrt{\frac{1}{N} \sum_{i=1}^N (z_i - \langle \mathbf{z} \rangle)^2}}, \quad (7.7)$$

where $\mathbf{z}$ is the vector of values from each feature, with size N . Normalising the inputs can accelerate convergence and improve model stability. Another common problem is having an imbalanced dataset. This occurs, for example, when the Monte Carlo simulations of background processes have many more events than their signal counterparts. Reweighting each sample so that they have a similar number of events is important because it prevents biases towards dominant classes and improves the training performance overall.

7.1.3 Performance

The general performance of a NN can be monitored using the loss function to check if the model is undertrained or overtrained, as illustrated in Figure 7.2. If both training and validation losses remain high, the model might be underfitted, meaning that it is too simple to capture the underlying patterns in data and needs more training. If the validation loss starts to increase while the training loss decreases, this may indicate overfitting, which happens when the NN memorises the training data rather than learning its general patterns. The choice of training and validation datasets can vary from one training to another. In general, the data used to train a model should not be evaluated by the same model to avoid introducing biases. One way to guarantee this is to use cross-validation, which consists in dividing the dataset in two parts, A and B , training first with one and then with the other, providing two NN models. This way, the model trained with dataset A is used to evaluate the dataset B , and viceversa. A more complex version of this method is the k -fold cross-validation, where the dataset is randomly divided in k parts (typically 5) of similar size, such that each subset $k = i$ is evaluated with the model trained in the remaining $k - 1$ subsets with $k \neq i$.

The receiver operating characteristic (ROC) curve is another performance evaluation metric for binary classification models. It is a representation of the true positive rate




Figure 7.2: Schematic representation of the loss function with respect to the number of epochs for both the training and validation datasets. Comparing the training and validation loss functions allows to monitor a possible underfitting or overfitting of the model.

(TPR) versus the false positive rate (FPR), both ranging from 0 to 1:

$$\begin{aligned} \text{TPR} &= \frac{\text{True positives}}{\text{True positives} + \text{False negatives}}, \\ \text{FPR} &= \frac{\text{False positives}}{\text{False positives} + \text{True negatives}}. \end{aligned} \quad (7.8)$$

In the context of a signal versus background classification problem, the true positives represent signal events correctly classified as signal and false negatives correspond to signal events incorrectly classified as background. Similarly, false positives represent background events incorrectly classified as signal and true negatives refer to background events correctly classified as background. A high TPR means that the NN is effective at identifying signal events, while a high FPR indicates that the NN is good at rejecting background events. The area under the ROC curve (AUC) is a metric that quantifies the model's overall ability to distinguish between the two different classes (e.g. signal and background). It ranges from 0 to 1. An AUC = 1 means that the NN can separate the two classes perfectly, an AUC = 0.5 is equivalent to random guessing, that is, the NN is not capable to distinguish the two different classes and an AUC < 0.5 means that the NN is performing worse than random guessing, implying that there are errors in the classification procedure. Figure 7.3 shows the AUC for three cases with good, random and underperforming NN results.

7.1.4 Parametrisation

A parametrised NN [158] is a neural network that includes external parameters in its architecture, allowing it to adapt its output based on them. A parametrised network can replace a set of individual networks trained for specific cases, as well as smoothly interpolate to scenarios that have not been used in the training. Parametrised NNs are useful in BSM searches where the mass of a particle from the signal hypothesis is unknown. In this case, a parametrised NN with the parameter being the mass of the BSM particle is proven to be more efficient, in both computation and performance, than using one individual NN for each mass hypothesis. Parametrisation is used in

7. Machine learning




Figure 7.3: Visual representation of different NN performances with their corresponding NN scores and ROC curves.

both the analyses presented in this thesis. In parametrised NNs, each signal event is provided an additional label corresponding to its generated mass, and each background event is provided a randomly distributed value from the set of generated masses. This label is used as one more feature in the training, therefore the evaluation depends on the input value for the mass and the final NN score is a function of this parameter. A schematic representation of a parametrised NN is shown in Figure 7.4.




Figure 7.4: Schematic representation of the basic constituents of a parametrised NN.

7.2 Boosted decision trees

A boosted decision tree (BDT) is a machine learning algorithm widely used in data analysis for classifying events or making predictions. It is created from multiple decision trees, which are flowchart-like structures where each internal node or splitting represents a test on an attribute, each branch represents the outcome of the test (yes or no) and each leaf node represents a class label, that is, the final decision taken after computing all attributes. Figure 7.5 shows a schematic view of a decision tree.




Figure 7.5: Schematic representation of a decision tree. An input event is evaluated in each node, and classified accordingly until a final decision (leaf) is reached.

Single decision trees can be unstable and prone to overfitting, as they often memorise details specific to the training data. Boosted decision trees combine the predictive power of multiple decision trees to produce a more accurate classification or regression. This combination, known as boosting, consists on training multiple models sequentially, with each new model focusing on the errors of the previous one, turning a set of weak learners (models only slightly better than random guessing) into a strong learner by combining them. The final BDT score, $\hat{y}(\mathbf{x})$, is defined as a linear combination of $\hat{y}_i(\mathbf{x})$, which are the predictions from each of the N decision trees, weighted by a factor α_i :

$$\hat{y}(\mathbf{x}) = \sum_{i=1}^N \alpha_i \hat{y}_i(\mathbf{x}). \quad (7.9)$$

Different boosting techniques exist, such as AdaBoost (adaptive boosting) [159] or XGBoost (extreme gradient boosting) [160], both of them used in the analyses presented in this thesis. In AdaBoost, the weights of training samples are iteratively adjusted based on whether they are correctly or incorrectly classified by the previous decision tree. Samples that are misclassified are given higher weights, while correctly classified samples receive lower weights in the next iteration. This process forces each subsequent decision tree to focus more on the difficult cases, thereby refining the overall model accuracy. XGBoost, on the other hand, improves the prediction by minimising a loss function via the gradient descent method. A regularisation term penalises very complex trees to avoid overfitting.

7. Machine learning

The performance of a BDT can be enhanced by using similar techniques to those mentioned in the previous section. Choosing the optimal architecture (number of decision trees, tree depth, learning rate) is fundamental. Cross-validation techniques are used as well, and overfitting can be prevented by monitoring the BDT loss function. The AUC metric is also used to evaluate the efficiency of the BDT prediction.

Chapter 8

Statistical methods

This chapter describes a set of statistical methods that are used to test the compatibility between data and Monte Carlo simulation, and to make statements about the presence or absence of a particular physics process. In a typical ATLAS search for BSM physics, rare events are often hidden within extensive backgrounds of common processes. Monte Carlo simulations are used to model expected particle interactions and to estimate backgrounds, allowing for its comparison with experimental data. The frequentist approach is often employed to establish confidence levels for new physics signals.

8.1 Hypothesis testing

In frequentist statistics, the probability $\mathcal{P}$ of an outcome A is defined as:

$$\mathcal{P}(A) = \lim_{N \rightarrow \infty} \frac{\text{Times that the outcome is } A}{N}. \quad (8.1)$$

The interpretation of $\mathcal{P}(A)$ as a probability is a natural choice, given that ATLAS studies are based on proton-proton collisions that are repeated over time with different outcomes, resulting in a very large dataset, denoted as $\mathbf{x}$. Frequentist statistics quantify what to expect, under the assumption of certain probabilities, about hypothetical repeated observations. The preferred hypotheses are those that predict a high probability for data to be like the observed data. The conditional probability of an outcome A given a condition B is defined as:

$$P(A|B) = \frac{P(A \cap B)}{P(B)}. \quad (8.2)$$

If A and B are independent, then $P(A \cap B) = P(A)P(B)$ and $P(A|B) = P(A)$.

The observed dataset, $\mathbf{x}$, can be represented in many different ways, one of them being a binned distribution where the bin contents are given by $\mathbf{n}$. In order to understand the underlying physics that can lead to the observable $\mathbf{n}$, data is compared with Monte Carlo simulations of signal and background processes, which depend on a series of parameters of interest (POIs) and nuisance parameters (NPs). In the searches performed in this thesis, the POI corresponds to the signal strength, denoted as μ , a multiplicative factor to the cross section of the signal process, and the NPs, denoted as $\boldsymbol{\theta}$, which include systematic uncertainties, background modelling parameters and statistical uncertainties. The general procedure used to search for a new phenomenon is to use a frequentist statistical test. For the purpose of discovering a new signal process, one defines the null or background-only hypothesis, H_0 , as to describe only known processes corresponding to the SM. This is to be tested against the alternative H_1 , which includes the background plus some sort of BSM signal. When setting limits,

the model with signal plus background plays the role of H_0 , and is tested against the background-only hypothesis.

8.2 Statistical tests

8.2.1 The binned profile likelihood method

A widely used procedure to establish discovery (or exclusion) in particle physics is based on a frequentist significance test using a binned likelihood ratio as a test statistic [161]. The binned likelihood function $\mathcal{L}(\mu, \boldsymbol{\theta})$ measures how well a statistical model describes observed data by calculating the probability of observing that data under different values of the parameters in the model. It is defined as:

$$\mathcal{L}(\mu, \boldsymbol{\theta}) = \mathcal{P}(\mathbf{n} \mid \mu, \boldsymbol{\theta}) = \prod_{i \in \text{bins}} \mathcal{P}(n_i \mid \mu S_i(\boldsymbol{\theta}) + B_i(\boldsymbol{\theta})), \quad (8.3)$$

where S_i and B_i are the signal and background predictions for the bin i , respectively, and n_i is the bin content. Assuming that the data follows a Poisson distribution, the probability for the observed data to be produced by the model in each bin is:

$$\mathcal{P}(n_i \mid \mu S_i(\boldsymbol{\theta}) + B_i(\boldsymbol{\theta})) = \frac{[\mu S_i(\boldsymbol{\theta}) + B_i(\boldsymbol{\theta})]^{n_i}}{n_i!} e^{-[\mu S_i(\boldsymbol{\theta}) + B_i(\boldsymbol{\theta})]}. \quad (8.4)$$

Systematic uncertainties arise from imperfect knowledge of the parameters of the model, and decrease the sensitivity of the analysis to the POI. However, systematic uncertainties can be constrained from auxiliary measurements. These constraints are implemented in the form of a penalty term $\rho(\theta_j)$ for each component of the set of NPs, θ_j . This penalty term typically follows a gaussian distribution:

$$\rho(\theta) = \frac{1}{\sqrt{2\pi}\Delta\theta_j} e^{-\frac{(\theta_j - \hat{\theta}_j)^2}{2(\Delta\theta_j)^2}}, \quad (8.5)$$

which encodes the information about the best estimate, $\hat{\theta}_j$ and width, $\Delta\theta_j$. Typically, they are redefined such that $\hat{\theta}_j = 0$ and $\Delta\theta_j = 1$, so that the post-fit parameters can be easily compared with their pre-fit values. A fitted value displaced from 0 indicates that the observed data is able to pull the NPs from the Monte Carlo in order to reach a better agreement. Pulled NPs in the fit can be problematic if there are strong correlations between parameters, causing instability and preventing the fit from finding a positive signal strength if there are signal events. Fitted errors smaller than 1 indicate that the initial uncertainty from the penalty term was too large and the observed data has enough statistical power to reduce the allowed range for the systematic variation. It is said in this case that the systematic is constrained. Constraints can translate into a reduction of the total systematic uncertainty in the final results, but they can also lead to biased or underestimated errors if not treated correctly.

Additionally, there are other NPs that are not associated to systematic uncertainties and are treated differently. Normalisation factors (NFs) are NPs that are included in

the likelihood as free parameters, with no penalty term. NFs are multiplicative factors that are used to correct the cross sections of certain processes, and they are treated in the same manner as the POI. There are also statistical uncertainties which arise from the limited number of Monte Carlo events used to build S_i and B_i in each bin. They are modulated using a Poisson distribution.

Adding the penalty term to the likelihood results in the following expression:

$$\mathcal{L}(\mu, \boldsymbol{\theta}) = \mathcal{P}(\mathbf{n} \mid \mu, \boldsymbol{\theta}) = \prod_{i \in \text{bins}} \mathcal{P}(n_i \mid \mu S_i(\boldsymbol{\theta}) + B_i(\boldsymbol{\theta})) \prod_{j \in \text{syst}} \rho(\theta_j). \quad (8.6)$$

The way to find the optimal parameters for the model is to maximise the profile likelihood ratio, defined as:

$$\lambda(\mu) = \frac{\mathcal{L}(\mu, \hat{\boldsymbol{\theta}}(\mu))}{\mathcal{L}(\hat{\mu}, \hat{\boldsymbol{\theta}})}, \quad (8.7)$$

where $\mathcal{L}(\mu, \hat{\boldsymbol{\theta}}(\mu))$ is the (conditional) maximum likelihood for a given μ and $\mathcal{L}(\hat{\mu}, \hat{\boldsymbol{\theta}})$ is the (unconditional) maximum likelihood across all parameters. $\lambda(\mu)$ can range from 0 to 1. The closer to 1, the better the agreement between the observed data and the proposed μ value. It is convenient to use the test statistic t_μ :

$$t_\mu = -2 \ln \lambda(\mu), \quad (8.8)$$

such that higher values of t_μ correspond to a larger incompatibility between μ and the observed data. If the search is targetting a BSM signal additional to the SM background, it is safe to assume that $\mu \geq 0$, with $\mu = 0$ corresponding to the background-only hypothesis. Therefore, if $\hat{\mu} < 0$ is found, it is established that the most compatible signal strength is $\mu = 0$. This leads to an alternative test statistic, $\tilde{t}_\mu$:

$$\tilde{t}_\mu = -2 \ln \tilde{\lambda}(\mu), \text{ with } \tilde{\lambda}(\mu) = \begin{cases} \frac{\mathcal{L}(\mu, \hat{\boldsymbol{\theta}}(\mu))}{\mathcal{L}(0, \hat{\boldsymbol{\theta}}(0))} & \text{if } \hat{\mu} < 0 \\ \frac{\mathcal{L}(\mu, \hat{\boldsymbol{\theta}}(\mu))}{\mathcal{L}(\hat{\mu}, \hat{\boldsymbol{\theta}})} & \text{if } \hat{\mu} \geq 0 \end{cases}. \quad (8.9)$$

The background-only hypothesis can be tested by setting $\mu = 0$:

$$\tilde{t}_0 = -2 \ln \tilde{\lambda}(0) = \begin{cases} 0 & \text{if } \hat{\mu} < 0 \\ -2 \ln \frac{\mathcal{L}(0, \hat{\boldsymbol{\theta}}(0))}{\mathcal{L}(\hat{\mu}, \hat{\boldsymbol{\theta}})} & \text{if } \hat{\mu} \geq 0 \end{cases}. \quad (8.10)$$

This expression allows to reject the background-only hypothesis only if $\hat{\mu} > 0$. A value of $\hat{\mu}$ smaller than 0 might also constitute an evidence against the background-only hypothesis, but it does not translate into an excess caused by additional signal events.

8. Statistical methods

In order to establish an upper limit on the signal strength, the following expression is used:

$$\tilde{t}_\mu = \begin{cases} -2 \ln \tilde{\lambda}(\mu) & \text{if } \hat{\mu} \leq \mu \\ 0 & \text{if } \hat{\mu} > \mu \end{cases} = \begin{cases} -2 \ln \frac{\mathcal{L}(\mu, \hat{\boldsymbol{\theta}}(\mu))}{\mathcal{L}(0, \hat{\boldsymbol{\theta}}(0))} & \text{if } \hat{\mu} \leq 0 \\ -2 \ln \frac{\mathcal{L}(\mu, \hat{\boldsymbol{\theta}}(\mu))}{\mathcal{L}(\hat{\mu}, \hat{\boldsymbol{\theta}})} & \text{if } 0 \leq \hat{\mu} \leq \mu \\ 0 & \text{if } \hat{\mu} > \mu \end{cases} \quad (8.11)$$

When the best-fit value $\hat{\mu}$ is negative, μ is tested against the background-only hypothesis to see if a positive μ improves the modelling of the observed data. When the best-fit value is positive and $\mu > \hat{\mu}$, μ is tested against $\hat{\mu}$ as a plausible upper bound. When the hypothesised μ value is below the best-fit value $\hat{\mu}$, μ is not a valid upper bound for the signal strength anymore, and $\tilde{t}_\mu$ is set to 0.

8.2.2 p -value and confidence intervals

The outcome of a search is summarised by quantising the level of agreement of the observed data with a given hypothesis H_μ by computing a p -value, denoted as p_μ . This number represents the probability, under the assumption of H_μ , of finding data of equal or greater incompatibility with the predictions of H_μ . It is defined as:

$$p_\mu = \int_{t_{\mu, \text{obs}}}^{\infty} f(t_\mu | H_\mu) dt_\mu, \quad (8.12)$$

where $t_{\mu, \text{obs}}$ is the value of the test statistic t_μ observed from the data and $f(t_\mu | H_\mu)$ denotes the probability distribution function of $\tilde{t}_\mu$ under the assumption of hypothesis H_μ ¹. The gaussian significance Z is often used instead of the p -value. It is defined as the number of standard deviations (σ) that a Gaussian variable would fluctuate in one direction to give the same p -value:

$$Z = \Phi^{-1}(1 - p_\mu), \quad (8.13)$$

where Φ refers to the cumulative Gaussian distribution.

In order to exclude a model with a given μ , p_μ is calculated and compared with a reference threshold α . If $p_\mu \leq \alpha$, then μ is said to be excluded at a confidence level (CL) of $1 - \alpha$. The value of α in this scenario is usually set to 0.05, providing a 95% CL exclusion limit. This corresponds to a gaussian significance $Z = 1.64\sigma$. When searching for a positive signal, the limits are more stringent. In order to claim evidence, the background-only hypothesis ($\mu = 0$) must be rejected with $p_0 \leq 1.3 \times 10^{-3}$, equivalent to $Z \geq 3\sigma$. In order to claim a discovery, the background-only hypothesis must be rejected with $p_0 \leq 2.9 \times 10^{-7}$, equivalent to $Z \geq 5\sigma$. Figure 8.1 shows a visual representation of different p -values and the corresponding gaussian significance.

The CL_s method [162] presents an alternative to the CL exclusion limits that is commonly used in experiments like ATLAS. This method is especially useful when

¹This expression holds when using $\tilde{t}_\mu$ instead of t_μ .




Figure 8.1: Thresholds for exclusion ($Z = 1.64\sigma$), evidence ($Z = 3\sigma$) and discovery ($Z = 5\sigma$) using a gaussian distribution with mean = 0 and standard deviation = 1. The associated p -value corresponds to the area under the curve at the tail of the distribution.

potential signals are weak or obscured by background noise, as it helps avoid excluding hypotheses that could be plausible under the background-only assumption. Mathematically, it is defined as:

$$\text{CL}_s = \frac{p_\mu}{1 - p_0}, \quad (8.14)$$

where the p -values p_μ and p_0 quantify the compatibility between the observed data and the H_μ and H_0 hypotheses, respectively. Dividing by $1 - p_0$ means that a signal hypothesis is not excluded unless the confidence in the background-only hypothesis is low. This avoids cases where fluctuations in the background data might lead to falsely ruling out the signal. Similarly to the standard CL exclusion limit, a signal with a certain value of μ is said to be excluded at a 95% CL if $\text{CL}_s \leq 0.05$.

8.3 Software implementation

8.3.1 The TRExFITTER package

TRExFITTER [163] is a ROOT-based [164] framework used for binned profile likelihood fits in particle physics analyses in ATLAS. It uses the HistFactory [165] and RooFit [166] libraries to generate a binned likelihood model based on template histograms of signal and background processes. Systematic uncertainties are included as shape and normalisation variations of these histograms, weighted by the corresponding penalty terms. All the information necessary for the binned likelihood fit is structured in a XML file, which is fed to the RooStats [167] library to compute and minimise likelihood ratios and estimate confidence intervals. When allowed by the model, asymptotic approximations are used to speed up the computation. These are based on Wilks' theorem [168], which states that when the sample size is very

large, the test statistic $-2 \ln \lambda(0)$ for the background-only hypothesis behaves like a $\chi^2(N)$ distribution, where N is the number of POIs.

8.3.2 Pruning, smoothing and symmetrisation

Modern ATLAS analyses often involve hundreds of systematic uncertainties that can make the statistical model overly complex. Systematic uncertainties are removed (pruned) in order to reduce the CPU time for the computation of the fit, as well as to improve convergence. This is done by considering the normalisation and shape uncertainties relative to each process in each region. In the analyses presented in this thesis, the pruning threshold is set to 1%, meaning that if the effect of a nuisance parameter is smaller than 1% (separately for shape and normalisation), then it is not used in the fit.

Smoothing refers to a technique for reducing statistical noise in the shape of systematic uncertainty distributions, particularly those coming from shape systematics, which affect the bin-by-bin yields in histograms used for the fit. Smoothing is essential when systematic variations exhibit large statistical fluctuations, as this noise can lead to unstable or biased fit results. Several algorithms are available for smoothing [169]. In this thesis, two of them are used: MAXVARIATION and COMMONSMOOTHTOOLPARABOLIC. The MAXVARIATION algorithm is preferred for systematics in samples with very low statistics, for example, instrumental systematics that apply to all background processes. It consists of merging close bins until the statistical uncertainty of the resulting combined bin is lower than a certain threshold X . If the number of slope changes in the binned distribution of choice, N_{slopes} , is larger than two, X is halved and the merging is repeated. Then, the TH1 : : Smooth option from the ROOT package algorithm is applied, to avoid artificially flat uncertainties due to rebinning. The COMMONSMOOTHTOOLPARABOLIC method uses a threshold of 5% statistical uncertainty for rebinning. Once the statistical uncertainty is reduced, pairs of neighbouring bins with the lowest χ^2 are combined, until $N_{\text{slopes}} < 3$. Then, each bin, except the first and last, is smoothed using the formula:

$$n'_i = \frac{n_{i-1} + 2n_i + n_{i+1}}{4}, \quad (8.15)$$

where n_i is the content of bin i .

In TREXFITTER, systematics are often symmetrised to simplify the handling of uncertainties and to ensure that the fitting process remains stable and interpretable. Symmetrizing systematics means that the variations due to a particular systematic uncertainty are treated as having equal magnitude in both positive and negative directions, even if the actual impact is slightly asymmetric. When comparing two different MC setups (e.g. ME computation or PS algorithm), symmetrisation reflects the fact that the actual uncertainty is not very well known, and the purpose is only to obtain a rough, conservative estimation of the effect.

Part IV

Search for decays of the Higgs boson into pseudoscalar particles decaying to four bottom quarks using proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector

Chapter 9

Analysis overview

Since the discovery of the Higgs boson with a mass near 125 GeV by the ATLAS and CMS collaborations [1, 2], an extensive measurement program has been undertaken to characterise its properties. So far, the branching ratio of the Higgs boson into undetected BSM particles has been constrained by global analyses down to 12%, approximately [3, 4]. Higgs boson decays are particularly sensitive to new physics due to the particle's small total width of 4 MeV. Even minimal couplings to new particles can result in sizeable branching ratios, remaining consistent with existing SM measurements [170].

9.1 Introduction

This search focuses on extensions of the SM that include additional pseudoscalar particles, denoted as a from now on. This new particle is studied via exotic decays of the SM Higgs boson, $H \rightarrow aa$. In scenarios where the a -boson mixes with the Higgs boson and inherits its Yukawa couplings to fermions, decays of the a -boson into heavy fermions such as b -quarks are favoured, and the process $H \rightarrow aa \rightarrow 4b$ is expected to have a sizeable branching ratio in the mass range $2m_b < m_a < m_H/2$. This analysis uses a simplified model of SM+ a described by the following BSM lagrangian:

$$\mathcal{L}_{\text{BSM}} \ni \frac{1}{2}(\partial_\mu a)(\partial^\mu a) - \frac{1}{2}m_a^2 a^2 - \frac{1}{2}\lambda_a a^2 H - y_b a \bar{b}(i\gamma^5)b \quad (9.1)$$

where λ_a and y_b are the $H \rightarrow aa$ coupling and the $a \rightarrow b\bar{b}$ coupling, respectively. They are treated as free parameters of the theory. This signature is studied via associated Higgs production with a Z boson that decays to two leptons, either ee or $\mu\mu$. These two leptons provide a clear trigger signal for the detector and helps to reduce the all-hadronic QCD background. Figure 9.1 shows the tree-level and 1-loop diagrams considered in this analysis. The 1-loop diagrams are expected to contribute to around 10% of the total signal cross section.




Figure 9.1: Feynman diagrams for $H \rightarrow aa \rightarrow 4b$ production in association with a Z boson decaying to two leptons at tree-level and 1-loop.

Figures 9.2 to 9.4 summarise the kinematic characteristics of the $ZH, H \rightarrow aa \rightarrow 4b$ process. ZH production is independent of the mass of the a -boson, and lepton kinematics from $Z \rightarrow \ell\bar{\ell}$ decay are expected to be the same for any value of m_a . Figure 9.2 shows the Higgs boson and Z boson p_T distributions, which are identical for different values of m_a . The $H \rightarrow aa$ and $a \rightarrow b\bar{b}$ decays, on the other hand, are sensitive to the a -boson mass. Higgs boson decays into lighter pseudoscalars are more likely to carry a high momentum, leading to boosted $a \rightarrow b\bar{b}$ decays in which the two b -quarks are close together in the ΔR plane. On the contrary, Higgs boson decays to heavier pseudoscalars tend to carry less energy and produce well separated or resolved $a \rightarrow b\bar{b}$ decays. Figures 9.3 and 9.4 compare the p_T of the a -bosons from $H \rightarrow aa$ decay for different mass hypotheses and the angular separation between the resulting b -quarks, respectively.




Figure 9.2: (a) Higgs and (b) Z -boson p_T from $ZH, H \rightarrow aa \rightarrow 4b$ production using different mass hypotheses. The kinematics of the ZH system do not depend on the value of m_a .

The signal process studied in this analysis is characterised by a large jet multiplicity in the final state, which leads to the decay products having low p_T in general. In order to recover potential signal events, this analysis extends the lower-bound for jet p_T calibration from the standard 20 GeV down to 15 GeV. In addition, several novel reconstruction techniques are used to identify the b -hadron decay products for different mass hypotheses with different topologies. For low m_a , the a -boson has a large Lorentz boost and its decay products are collimated. Consequently, the $a \rightarrow b\bar{b}$ final state is reconstructed as a single jet that contains the hadronisation products of the two b -quarks. Due to the relatively low mass of the Higgs boson, the merged $a \rightarrow b\bar{b}$ jet will have low p_T . This analysis uses the DeXTer $b\bar{b}$ -tagger to identify low-mass, merged, $b\bar{b}$ -jets, referred to as B -jets from now on. For high m_a , the b -quarks from $a \rightarrow b\bar{b}$ decay tend to be well separated, and the reconstructed jets capture the hadronisation of a single b -quark, denoted as b -jet. They are identified using the DL1r b -tagger. In some cases, particularly at high m_a , the p_T of the decay products may be too low to produce a reconstructed jet. In these cases, some of the signal efficiency can be recovered by identifying the soft secondary vertex (v) that corresponds to the




Figure 9.3: (a) Leading and (b) sub-leading a -boson p_T from $H \rightarrow aa$ decay for different mass hypotheses. These distributions show how decay products with larger m_a tend to carry less momentum.




Figure 9.4: (a) ΔR separation between the two b -quarks from the (a) leading and (b) sub-leading $a \rightarrow b\bar{b}$ decay for different mass hypotheses. For lower values of m_a , the system is boosted, while for higher values of m_a , the system is resolved.

low energy b -quark, using the TC-LVT algorithm.

This analysis is a search for the exotic Higgs boson decay $H \rightarrow aa \rightarrow 4b$ over the mass range $12 \leq m_a \leq 60$ GeV using the full Run 2 dataset of proton-proton collisions at $\sqrt{s} = 13$ TeV recorded with the ATLAS detector, corresponding to a total integrated luminosity of 140 fb^{-1} . Similar searches in the $4b$ decay channel have been performed by the ATLAS Collaboration [171–173] and the CMS Collaboration [174]. This analysis aims to improve the sensitivity of previous ATLAS results by using a larger integrated luminosity and dedicated heavy flavor reconstruction algorithms. It is also complementary to other searches for $H \rightarrow aa$ decays performed by the ATLAS

9. Analysis overview

and CMS collaborations using Run 1 and Run 2 data in several final states, including 4μ [175, 176], $2\mu 2\tau$ [177–180], 4τ [178, 180, 181], $2b2\mu$ [178, 182–185], $2b2\tau$ [185–187], 4γ [188–191], and $2\gamma 2j$ [192].

9.2 Data and simulated samples

This section describes the data and Monte Carlo samples used for this analysis. In both cases, the datasets are divided in three subsets, corresponding to the three data-taking periods of Run 2, which are 2015+2016, 2017 and 2018. The signal and background samples are simulated using different matrix element generators, interfaced with various parton shower algorithms. A reweighting method is applied to the simulated samples such that they match the pileup conditions in data. For the detector simulation two different approaches are used: the full ATLAS detector simulation (FS), based on the standard GEANT4 simulation, and the fast detector simulation (AFII) which uses a parametrisation of the calorimeter response. Both MC and data events are processed with the same reconstruction and analysis software.

9.2.1 Data samples

This analysis uses proton–proton collision data collected from 2015 to 2018 by the ATLAS detector at $\sqrt{s} = 13$ TeV. Selected events are recorded using unprescaled triggers, as detailed in Section 9.3. Standard data taking quality requirements are applied, such as stable LHC beams and fully operational ATLAS detector conditions. These requirements define the Good Run Lists (GRLs) of events that can be used for physics analyses. The GRLs used for this analysis for each year of data taking are listed in Table 9.1, together with the corresponding integrated luminosity. The total integrated luminosity for the full Run 2 dataset is 140.1 fb^{-1} , usually written as 140 fb^{-1} . The uncertainty is 0.83% [193], obtained using the LUCID-2 detector [194] for the primary luminosity measurements.

Year	$\int \mathcal{L} dt [\text{pb}^{-1}]$	GRL
2015	3244.54	20170619/data15_13TeV.periodAllYear_DetStatus-v89-pro21-02_Unknown_PHYS_StandardGRL_All_Good_25ns.xml
2016	33402.2	20170605/data16_13TeV.periodAllYear_DetStatus-v89-pro21-01_DQDefects-00-02-04_PHYS_StandardGRL_All_Good_25ns.xml
2017	44630.6	20180619/data17_13TeV.periodAllYear_DetStatus-v99-pro22-01_Unknown_PHYS_StandardGRL_All_Good_25ns_TriggerNo17e33prim.xml
2018	58791.6	20181105/data18_13TeV.periodAllYear_DetStatus-v102-pro22-04_Unknown_PHYS_StandardGRL_All_Good_25ns_TriggerNo17e33prim.xml

Table 9.1: GRLs used for each of the four years of data-taking, together with the corresponding integrated luminosity for unprescaled triggers, computed with the 0flLumi-13TeV-011 luminosity tag [195]. All the listed GRLs are available in [196].

9.2.2 Signal Monte Carlo samples

The signal samples consist of ZH production at tree-level ($q\bar{q} \rightarrow ZH$) and 1-loop ($gg \rightarrow ZH$), generated using the POWHEGBox-v2 MC generator [197] at NLO accuracy in QCD with the NNPDF3.0 NLO PDF set [198]. A Higgs boson mass of 125 GeV is assumed. The events are interfaced with the PYTHIA 8.244 generator that models the decay of the Higgs boson into a pair of light pseudoscalars with $m_a \in [12, 16, 20, 25, 30, 40, 50, 60]$ GeV. They are decayed into a pair of b -quarks with $\text{BR}(H \rightarrow aa \rightarrow 4b) = 1$. Arbitrary values of the branching ratio can be obtained a posteriori by applying a constant factor to the cross section. Events are showered and hadronised with the PYTHIA 8.244 PS algorithm [199] using the AZNLO tune to model the underlying event. Events from minimum-bias interactions are simulated with the PYTHIA 8.244 generator with A3 tune. They are overlaid on the simulated signal events according to the luminosity profile of the recorded data. Finally, the generated samples are processed through a simulation of the detector geometry and response using AFII. All samples are processed through the same reconstruction software as the data. Simulated events are corrected so that the object identification efficiencies, energy scales and energy resolutions match those determined from data control samples. Samples are normalised using dedicated higher-order cross section predictions calculated at NNLO in QCD with NLO electroweak corrections.

9.2.3 Background Monte Carlo samples

$t\bar{t}$ + jets

The production of $t\bar{t}$ + jets events is modelled using the POWHEGBox-v2 generator at NLO in QCD with the five flavour scheme (5FS) NNPDF3.0 NLO PDF set. The renormalisation and factorisation scales are set to the default scale $\mu_R = \mu_F = \sqrt{m_t^2 + p_{T,t}^2}$. It is combined with PYTHIA 8.230 for the parton shower and hadronisation, using the A14 tuned parameters and the NNPDF2.3 LO PDF set. The $t\bar{t}$ + jets sample is normalised to the cross section prediction at NNLO in QCD including the resummation of NNLL soft-gluon terms calculated using TOP++ 2.0 [200]. In order to enhance the statistics in the phase space relevant to this analysis, dedicated filtered samples are produced, requiring additional c - or b -hadrons apart from those originating from top quark decays. They are categorised as follows:

- **BBFilt**: at least two additional b -hadrons with $p_T > 5$ GeV matched to truth jets.
- **BFiltBBVeto**: one additional b -hadron with $p_T > 5$ GeV matched to a truth jet.
- **CFiltBVeto**: at least one additional c -hadron with $p_T > 5$ GeV matched to truth jets and no additional b -hadrons.

Simulated events in the inclusive sample are replaced with simulated events from the corresponding filtered samples to avoid duplication and to enhance MC statistics.

9. Analysis overview

The $t\bar{t}$ + jets MC samples are used to model $t\bar{t}$ + light and $t\bar{t}$ + $\geq 1c$, that is, top quark pair production where the additional jets that arise from QCD radiation originate from gluons, light quarks (u, d, s) or c -hadrons.

$t\bar{t}$ + $b\bar{b}$

The production of $t\bar{t}$ + $\geq 1b$ events is modelled using a dedicated MC simulation for top quark pair production in association with a b -quark pair, $t\bar{t}$ + $b\bar{b}$. The MC predictions are calculated using the POWHEGBOX-Res framework [201] at NLO with massive b -quarks, using the four flavour scheme (4FS) NNPDF3.0 NNLO PDF set [202]. The renormalisation scale is set to the geometric average of the of t - and b -quark transverse energies, $\mu_R = \sqrt[4]{\prod_i E_{T,i}}$ with $i = t, \bar{t}, b, \bar{b}$. The factorisation scale is related to the sum of the transverse mass of the outgoing partons in the matrix element calculation and is defined as $\mu_F = \frac{1}{2}(\sum_i E_{T,i})$, with $i = t, \bar{t}, b, \bar{b}, j$. The PYTHIA 8.230 parameters for parton shower and hadronisation modelling are set to the A14 tune and the NNPDF2.3 LO PDF set.

Single-top

Single-top production includes the t -channel, the s -channel, and tW production. It is modelled using POWHEGBOX-v2 at NLO in QCD. For t -channel production, events are generated with the 4FS NNPDF3.0 NLO PDF set, and the renormalisation and factorisation scales are set to $\mu_R = \mu_F = 4\sqrt{m_b^2 + p_{T,b}^2}$ [203]. For the s -channel and tW production, events are generated with the 5FS NNPDF3.0 NLO PDF set, and the renormalisation and factorisation scales are set to the default value $\mu_R = \mu_F = m_t$. For tW production, the diagram removal (DR) scheme [204] is applied to handle the interference with $t\bar{t}$ production at NLO. Events are showered with PYTHIA 8.230. All single-top samples are normalised to the cross section prediction calculated at NLO+NNLL in QCD [205].

$t\bar{t}H$, $t\bar{t}Z$ and $t\bar{t}W$

The production of $t\bar{t}H$ events is modelled using the POWHEGBOX-v2 generator at NLO in QCD with the NNPDF3.0 NLO PDF set. The resulting events are interfaced with PYTHIA 8.230 for the PS using the A14 tune and the NNPDF2.3 LO PDF set. The decays of bottom and charm hadrons are performed by EVTGEN 1.6.0 [206]. The cross section is calculated at NLO in the QCD and EW couplings [49].

$t\bar{t}Z$ and $t\bar{t}W$ events are produced using the MADGRAPH5_aMC@NLO v2.3.3 generator [207] at NLO with the NNPDF3.0 NLO PDF set. Events are interfaced with PYTHIA 8.210 using the A14 tune and the and the NNPDF2.3 LO PDF set. The decays of bottom and charm hadrons are simulated using the EVTGEN 1.2.0 program. Both cross sections are calculated at NLO in the QCD and EW couplings [49].

Rare processes

Rare processes refer to single-top production in association with vector bosons such as tZq and tWZ , which have a very low cross section and do not contribute significantly to the total background. The production of tZq events is performed using MADGRAPH5_aMC@NLO v2.3.3 at LO in QCD in the 4FS with the CTEQ6L1 PDF

set [208]. Events are showered with PYTHIA 8.212. The production of tWZ events is modelled using the MADGRAPH5_aMC@NLO v2.2.3 generator at NLO with the NNPDF3.0 NLO PDF set, interfaced with Pythia 8.212. The DR scheme is employed to handle the interference between tWZ and $t\bar{t}Z$ at NLO.

$Z/W + \text{jets}$

The production of $Z/W + \text{jets}$ is simulated with the SHERPA 2.2.11 generator [209] using NLO matrix elements for up to two partons, and LO matrix elements for up to four partons, calculated with the COMIX [210] and OPENLOOPS [211] libraries. The MC sampling is biased towards $\max(p_{T,Z/W}, H_T)^2$ in order to enhance statistics in the tail of the p_T distribution. They are matched with the SHERPA PS [212] using the MEPS@NLO prescription [213] with the set of tuned parameters developed by the SHERPA authors. The NNPDF3.1 NNLO set of PDFs is used and the cross section is normalised to the NNLO prediction [214].

Diboson

Diboson final states include processes such as ZZ , ZW and WW . They are simulated with the SHERPA 2.2.1 or 2.2.2 generator depending on whether the bosons decay to hadrons or leptons and including off-shell effects and Higgs boson contributions when appropriate. Fully leptonic and semileptonic final states are generated using matrix elements at NLO accuracy in QCD for up to one additional parton emission and at LO accuracy for up to three additional parton emissions. Samples for the loop-induced processes $gg \rightarrow ZZ/ZW/WW$ are generated using LO-accurate matrix elements for up to one additional parton emission for both the fully leptonic and semileptonic final states. The matrix element calculations are matched and merged with the Sherpa parton shower based on the Catani–Seymour dipole factorisation [212] using the MEPS@NLO prescription. Virtual QCD corrections are provided by the OPENLOOPS library. The NNPDF3.0 NNLO set of PDFs is used, along with a dedicated set of tuned PS parameters developed by the SHERPA authors.

Summary

The full list of signal and background Monte Carlo samples is summarised in Table 9.2.

9. Analysis overview

Process	ME generator	PDF set	PS generator	Norm.	Sim.
$H \rightarrow aa \rightarrow 4b$ Tree-level	POWHEGBox-v2	NNPDF3.0 NLO	PYTHIA 8.244	NNLO _{QCD} + NLO _{EW}	AFII
$H \rightarrow aa \rightarrow 4b$ 1-loop	POWHEGBox-v2	NNPDF3.0 NLO	PYTHIA 8.244	NNLO _{QCD} + NLO _{EW}	AFII
$t\bar{t}$ + jets	POWHEGBox-v2	NNPDF3.0 NLO	PYTHIA 8.230	(NLO+NNLL) _{QCD}	FS
$t\bar{t} + b\bar{b}$	POWHEGBox-Res	NNPDF3.0 NNLO	PYTHIA 8.230	–	FS
Single-top	POWHEGBox-v2	NNPDF3.0 NLO	PYTHIA 8.230	(NLO+NNLL) _{QCD}	FS
$t\bar{t}H$	POWHEGBox-v2	NNPDF3.0 NLO	PYTHIA 8.230	NLO _{QCD} +EW	FS
$t\bar{t}Z$	MG5_aMC@NLO v2.3.3	NNPDF3.0 NLO	PYTHIA 8.210	NLO _{QCD} +EW	FS
$t\bar{t}W$	MG5_aMC@NLO v2.3.3	NNPDF3.0 NLO	PYTHIA 8.210	NLO _{QCD} +EW	FS
Rare processes	MG5_aMC@NLO v2.3.3	CTEQ6L1 NNPDF3.0 NLO	PYTHIA 8.212	–	FS
Z/W + jets	SHERPA 2.2.11	NNPDF3.1 NNLO	SHERPA	NNLO _{QCD}	FS
Diboson	SHERPA 2.2.1 SHERPA 2.2.2	NNPDF3.1 NNLO	SHERPA	–	FS

Table 9.2: Summary of signal and background samples used in the $H \rightarrow aa \rightarrow 4b$ analysis, including the ME generator, the PDF set, the PS generator, the cross section normalisation and the type of detector simulation. Note that MADGRAPH is abbreviated to MG.

9.3 Event selection

9.3.1 Trigger

This search is based on proton-proton collision data collected by the ATLAS experiment between 2015 and 2018, corresponding to the full Run 2 data taking period. Only events recorded with a single-electron or single-muon trigger under stable beam conditions and for which all detector subsystems were operational are considered. Single-lepton triggers with different p_T thresholds are combined in a logical OR in order to increase the overall efficiency. Table 9.3 lists all the trigger chains used during the four years of data-taking.

Years	Single electron triggers	Single muon triggers
2015	HLT_e24_lhmedium_L1EM20VH HLT_e60_lhmedium HLT_e120_lhloose	HLT_mu20_iloose_L1MU15 HLT_mu50
2016-2018	HLT_e26_lhtight_nod0_ivarloose HLT_e60_lhmedium_nod0 HLT_e140_lhloose_nod0	HLT_mu26_ivarmedium HLT_mu50

Table 9.3: Single-electron and single-muon trigger chains, depending on the year of data-taking.

The triggers with low p_T thresholds include isolation requirements on the candidate lepton, resulting in inefficiencies at high p_T that are recovered by the triggers with higher p_T thresholds.

9.3.2 Preselection

Events are required to have exactly two leptons (ee , $\mu\mu$ or $e\mu$) with opposite charge, satisfying the criteria defined in Section 6.2. Since single-lepton triggers are used, at least one of the two reconstructed leptons is required to have $p_T > 27$ GeV and match a lepton with the same flavour reconstructed by the trigger algorithm within $\Delta R < 0.15$. In the ee and $\mu\mu$ channels, the dilepton invariant mass must be above 50 GeV. This cut is used to reduce the contribution from $t\bar{t}$ + jets. Further suppression of the background is achieved by applying a cut based on a combination of b -objects:

$$2N_B + N_b + N_v \geq 3, \quad (9.2)$$

where N_B refers to the number of B -jets tagged using the DeXTer 60% WP, N_b refers to the number of b -jets tagged with the DL1r 85% WP and N_v corresponds to the number of soft secondary vertices tagged using TC-LVT tagger. Events with exactly three b -objects are used to derive corrections to the background modelling in simulation, as will be described in Section 9.4. Events with at least four b -objects are used to define the signal regions (SRs), while events with three b -objects are used to define the control regions (CRs). This is described in Section 9.5. Tables 9.4 and 9.5 show the event yields in the preselection region for background and signal samples, respectively.

Sample	Events	% over total
$t\bar{t}$ +light	210531	54%
$t\bar{t}$ + $\geq 1c$	33652	9%
$t\bar{t}$ + $\geq 1b$	31476	8%
tW	9574	2%
$t\bar{t}H$	927	< 1%
$t\bar{t}Z$	1370	< 1%
$t\bar{t}W$	708	< 1%
tq, tZ, tWZ	336	< 1%
Z +jets	101803	26%
W +jets	73	< 1%
WW, ZZ, WZ	2005	1%
Total	392455	100%

Table 9.4: Background composition in the preselection region, corresponding to 140 fb^{-1} .

m_a [GeV]	Events
12	1950
16	1750
20	1573
25	1416
30	1338
40	1313
50	1376
60	1439

Table 9.5: Number of signal events in the preselection region, corresponding to 140 fb^{-1} .

9.4 Background modelling

The main backgrounds for this analysis are top quark pair production and Z boson production in association with extra jets ($t\bar{t}$ +jets and Z+jets). They are particularly relevant when they have additional contributions from QCD radiation causing a large b -quark multiplicity in the final state, similar to the signal hypothesis. Example Feynman diagrams of $t\bar{t}$ +jets and Z+jets with additional b -quarks are shown in Figures 9.5 and 9.6.




Figure 9.5: Dominant Feynman diagram for $t\bar{t}$ +jets production at tree level with four b -quarks in the final state.




Figure 9.6: Examples of Feynman diagrams for Z+jets production at tree level with two, three and four b -quarks in the final state.

Previous ATLAS studies [215, 216] have shown that the modelling of the $t\bar{t}$ +jets and Z+jets backgrounds in Monte Carlo simulations is not fully accurate, especially when

there is heavy-flavour (HF) radiation, that is, additional b - and c -jets produced in association to the $t\bar{t}$ pair or the Z boson. To address this, a series of data-driven corrections is applied to improve the agreement between the simulated and observed data. This section outlines the reweighting procedure for these two backgrounds, which are corrected independently using dedicated regions where the signal contribution is negligible. The methodology for reweighting $t\bar{t}$ +jets and Z +jets is detailed in Sections 9.4.1 and 9.4.2, respectively. The associated uncertainties are discussed in Section 9.6. These corrections provide first-order improvements to the modelling, while residual discrepancies are further mitigated by the NPs in the final fit model. Given that the $t\bar{t}$ +jets reweighting region has very little contamination from other backgrounds, $t\bar{t}$ +jets corrections are calculated first and then applied when computing the Z +jets reweighting.

9.4.1 $t\bar{t}$ +jets reweighting

Reweighting region

The control region used for $t\bar{t}$ +jets reweighting requires two leptons with opposite charges and different flavours (one electron and one muon), in order to suppress the Z +jets background. The invariant mass of the dilepton system must be within the Z -mass peak, $|m_{\ell\ell} - m_Z| < 20$ GeV, which is where signal events are expected to be more abundant. Events must contain at least two b -jets, tagged using the DL1r 85% WP, given that the main t -quark decay is $t \rightarrow W\bar{b}$. This region contains a fraction of $t\bar{t}$ +jets above 95%, as shown in Table 9.6.

Sample	Events	% over total
$t\bar{t}$ +light	137325	86%
$t\bar{t}+\geq 1c$	10879	7%
$t\bar{t}+\geq 1b$	5846	4%
tW	5327	3%
$t\bar{t}H$	171	< 1%
$t\bar{t}Z$	193	< 1%
$t\bar{t}W$	183	< 1%
tq, tZ, tWZ	22	< 1%
Z +jets	139	< 1%
W +jets	42	< 1%
WW, ZZ, WZ	155	< 1%
Total	160282	100%

Table 9.6: Background composition in the $t\bar{t}$ +jets reweighting region.

Heavy-flavour reweighting

The fraction of $t\bar{t}$ +HF is typically underestimated in MC simulation. In order to have a more accurate flavour composition, a first reweighting is done based on the truth categorisation of the $t\bar{t}$ +jets sample, which is derived from the flavour of the extra jets of the event that do not originate from $t\bar{t}$ decay. These extra jets are reconstructed

9. Analysis overview

from stable truth particles using the anti- k_T algorithm with a radius parameter $R = 0.4$, and they are required to have $p_T > 15$ GeV and $|\eta| < 2.5$. Events are labelled as $t\bar{t}+\geq 1b$ if at least one extra jet is matched within $\Delta R < 0.4$ to a b -hadron with $p_T > 5$ GeV not originating from the $t\bar{t}$ decay based on MC truth information. Similarly, if at least one particle jet is matched to a c -hadron, which is not a decay product of a b -hadron or a W boson, with $p_T > 5$ GeV, the event is labelled as $t\bar{t}+\geq 1c$. These categories also include events where the additional jets are coming from multiparton interaction (MPI) or final state radiation (FSR). The remaining events, including those with no additional jets matched to a b -hadron or c -hadron, are labelled as $t\bar{t}$ +light. Since the tW process interferes with $t\bar{t}$ +light at NLO, it is included in the $t\bar{t}$ +light category for reweighting purposes.

This procedure aims to correct the underestimation of the $t\bar{t}$ +HF production rate in the MC prediction compared to data. Scale factors for the three different components are derived using a likelihood fit over a distribution that provides a separation between the different categories. The variable of choice is the sum of the pseudo-continuous (PC) b -tagging score of all the AK4 jets from each event (sumPCBTag)¹. This PC b -tagging score is an integer number that is assigned to each jet and represents the tightest DL1r WP that each jet fulfills, according to the map shown in Table 9.7.

b -tagged?	DL1r WP	b -tagging score
No	None	1
Yes	85%	2
Yes	77%	3
Yes	70%	4
Yes	60%	5

Table 9.7: Correspondence between the DL1r WP and the b -tagging score for each jet. Low values of the DL1r WP correspond to tight b -tagging WPs, and viceversa.

Systematic uncertainties from DL1r tagging are taken into account in this fit. Figure 9.7 shows the sumPCBTag distribution before and after the fit. $t\bar{t}$ +light is the dominant category at low values of sumPCBTag, while $t\bar{t}+\geq 1b$ is the dominant category at high values. The resulting normalisation factors are shown in Table 9.8.

Category	Norm. factor
$t\bar{t}$ +light, tW	0.90 ± 0.2
$t\bar{t}+\geq 1c$	1.79 ± 0.12
$t\bar{t}+\geq 1b$	1.20 ± 0.10

Table 9.8: Normalisation factors for the three $t\bar{t}$ +jets categories resulting from the likelihood fit performed using the sumPCBTag distribution.

¹Only the 85% WP is calibrated for jets with $p_T < 20$ GeV. However, it was checked that the contribution from those jets to the reweighting procedure is negligible.




Figure 9.7: Data versus MC comparison of the sum of the PC b -tagging score of all the jets per event (a) before and (b) after applying the $t\bar{t}$ +jets normalisation correction. The error band includes the statistical error and the DL1r uncertainties.

The post-fit distribution shows a better agreement between data and MC. The resulting normalisation factors are larger than 1 for $t\bar{t}+\geq 1c$ and $t\bar{t}+\geq 1b$, and smaller than 1 for $t\bar{t}$ +light. These factors are consistent with an underestimation of the $t\bar{t}$ +HF rate and provide an approximate correction to facilitate the convergence of the final fit.

Kinematic reweighting

After applying the normalisation factors for the different $t\bar{t}$ +jets categories, there are residual shape mismodellings affecting distributions such as the jet multiplicity, N_{jets} and the transverse momentum of both jets and leptons that need to be corrected as well. In order to reduce the data to MC discrepancies as much as possible, a reweighting procedure in two steps is proposed, correcting first the mismodelling with respect to N_{jets} and second, the mismodelling observed in the p_T distributions. In both cases, the correction factors are calculated assuming that all the mismodelling comes from the background that is being reweighted. They are computed over the data and MC binned distributions as:

$$R(x) = \frac{\text{Data}(x) - \text{MC}_{\text{other}}(x)}{\text{MC}_{t\bar{t}+\text{jets}, tW}(x)}, \quad (9.3)$$

where x is the variable to be corrected. In the kinematic reweighting, the $t\bar{t}$ +jets and tW backgrounds are treated as a single sample and receive the same correction factors. No systematic uncertainties are considered.

The first correction is applied per-bin to the N_{jets} distribution, using the ratio defined in Equation 9.3 is used. It is calculated in the range between 2 and 8, with the last bin including $N_{\text{jets}} \geq 8$. Figure 9.8 shows the data versus MC comparison over the number of jets per event in the $t\bar{t}$ +jets reweighting region, before and after the correction. Note

9. Analysis overview

that the HF correction is already applied. The values of the coefficients are shown in Table 9.9.




Figure 9.8: Data versus MC comparison of the N_{jets} distribution (a) before and (b) after applying the N_{jets} correction to the $t\bar{t}$ +jets and tW backgrounds. Systematic uncertainties are not included.

N_{jets}	Weight
2	1.043 ± 0.005
3	1.006 ± 0.005
4	1.001 ± 0.005
5	1.023 ± 0.008
6	1.072 ± 0.011
7	1.155 ± 0.018
≥ 8	1.23 ± 0.03

Table 9.9: Correction factors for each bin of the $t\bar{t}$ +jets and tW N_{jets} distribution.

The second correction is meant to fix the p_T mismodelling observed in both jets and leptons. For $t\bar{t}$ +jets, the variable of choice is the event hardness or H_T , which is defined as the scalar sum of the p_T of all the jets and leptons in the event. As shown in Figure 9.9 (a), this variable depends strongly on N_{jets} . Using the event hardness to calculate a p_T correction would heavily affect the previous reweighting step, unless one reweighting function was derived for each value of the jet multiplicity. In order to avoid that, a new variable called H_T^{red} is defined:

$$H_T^{\text{red}}(n) = H_T(n) - (n - n_{\text{min}})\Delta H_T(n), \quad (9.4)$$

where n is the number of jets in the event, n_{min} is the minimum number of jets per event allowed by the preselection (two, in this particular case) and $\Delta H_T(n)$ is an approximate

value that represents the offset in H_T caused by the addition of each extra jet to the event, that is:

$$\Delta H_T(n) = \left\langle \frac{\langle H_T(n) \rangle - \langle H_T(n_{\min}) \rangle}{n - n_{\min}} \right\rangle_{n=3, \dots, \geq 8}, \quad (9.5)$$

where $H_T(n)$ is the H_T distribution for the subset of events with n jets. The H_T^{red} distribution is illustrated in Figure 9.9 (b) for events with different number of jets.




Figure 9.9: Data-only (a) H_T and (b) H_T^{red} normalised distributions as a function of the number of jets in the $t\bar{t}$ +jets reweighting region. The histograms show how each additional jet in the event causes an offset in H_T , which is mitigated in H_T^{red} .

For this correction, Equation 9.3 is applied over the binned H_T^{red} distribution and the results are fitted using a continuous hyperbolic function of the type:

$$f(x) = c_0 + \frac{c_1}{(x + c_2)^{c_3}}, \quad (9.6)$$

as shown in Figure 9.10. The parameters of the fitted hyperbola are included in Table 9.10. Values of $H_T^{\text{red}} \leq 50$ GeV are reweighted using the value of the function at 50 GeV, in order to avoid very large values when $H_T^{\text{red}} \rightarrow 0$. Figure 9.11 shows the N_{jets} distribution before and after applying the correction. As expected, no significant changes are observed. Figures 9.12 to 9.14 show several p_T distributions, both for jets and leptons, before and after applying the correction. A general improvement in the shape mismodelling is achieved.

9. Analysis overview




Figure 9.10: $t\bar{t}$ +jets correction factor as a function of H_T^{red} .

Parameter	Value
c_0	1.283 ± 0.003
c_1	3.515 ± 0.004
c_2	35 ± 7
c_3	0.0773 ± 0.0003

Table 9.10: Fitted parameters to the hyperbolic function for the H_T correction. The table shows the numerical values from all reweighting parameters, with their corresponding uncorrelated uncertainties.




Figure 9.11: N_{jets} distribution (a) before and (b) after applying the H_T correction using H_T^{red} . Systematic uncertainties are not included.




Figure 9.12: Leading (1) and sub-leading (2) jet p_T distributions (a) before and (b) after applying the H_T correction using H_T^{red} . Systematic uncertainties are not included.

9. Analysis overview




Figure 9.13: Jet 3 and jet 4 p_T distributions (a) before and (b) after applying the H_T correction using H_T^{red} . Jets are ordered by p_T . Systematic uncertainties are not included.




Figure 9.14: Leading (1) and sub-leading (2) lepton p_T distributions (a) before and (b) after applying the H_T correction using H_T^{red} . Systematic uncertainties are not included.

9.4.2 Z+jets reweighting

Reweighting region

The Z+jets reweighting region contains two leptons with opposite charges and the same flavour (either ee or $\mu\mu$). The invariant mass of the system is required to be within the Z-mass peak, $|m_{\ell\ell} - m_Z| < 20$ GeV. In addition, events must have at least two jets and satisfy the condition $2N_B + N_b + N_v = 3$. This selection is very close to the signal regions; consequently, both the composition and mismodelling are expected to be similar. This region has a contribution from $t\bar{t}$ +jets, which is reduced by applying an extra cut in the missing transverse energy, $E_T^{\text{miss}} < 60$ GeV. The background composition in this region is summarised in Table 9.11.

Sample	Events	% over total
$t\bar{t}$ +light	12149	12%
$t\bar{t}+\geq 1c$	1509	2%
$t\bar{t}+\geq 1b$	1060	1%
tW	508	1%
tq, tZ, tWZ	159	< 1%
$t\bar{t}H$	25	< 1%
$t\bar{t}Z$	335	< 1%
$t\bar{t}W$	17	< 1%
Z+jets	81337	83%
W+jets	2	< 1%
WW, ZZ, WZ	1389	1%
Total	98490	100%

Table 9.11: Background composition in the Z+jets reweighting region.

HF reweighting

A HF correction for the Z+jets background is not considered since both the reweighting region and the signal regions are strongly dominated by the $Z+\geq 1b$ category, as shown in Table 9.12. The N_{jets} correction in the next section provides an implicit normalisation correction for Z+jets that is only possible under the assumption that $Z+\text{jets} \approx Z+\geq 1b$. This is not the case for $t\bar{t}$ +jets. Table 9.13 shows how all of the $t\bar{t}$ +jets components appear in non-negligible fractions for different b -object multiplicities.

Region	Category	% over total Z+jets
3 b -objects	Z+light	8%
	Z+ $\geq 1c$	11%
	Z+ $\geq 1b$	81%
4 b -objects	Z+light	5%
	Z+ $\geq 1c$	8%
	Z+ $\geq 1b$	87%
≥ 5 b -objects	Z+light	3%
	Z+ $\geq 1c$	3%
	Z+ $\geq 1b$	94%

Table 9.12: Z+jets composition in the Z+jets reweighting region, as well as additional regions with higher b -object multiplicities, where most of the signal events are expected to be. The fraction of Z+ $\geq 1b$ is above 80% in all of them.

Region	Category	% over total $t\bar{t}$ +jets
3 b -objects	$t\bar{t}$ +light	70%
	$t\bar{t}$ + $\geq 1c$	20%
	$t\bar{t}$ + $\geq 1b$	10%
4 b -objects	$t\bar{t}$ +light	40%
	$t\bar{t}$ + $\geq 1c$	26%
	$t\bar{t}$ + $\geq 1b$	34%
≥ 5 b -objects	$t\bar{t}$ +light	16%
	$t\bar{t}$ + $\geq 1c$	22%
	$t\bar{t}$ + $\geq 1b$	62%

Table 9.13: $t\bar{t}$ +jets composition in the Z+jets reweighting region, as well as additional regions with higher b -object multiplicities, where most of the signal events are expected to be. The $t\bar{t}$ +jets composition varies strongly from one region to another.

Kinematic reweighting

The N_{jets} distribution for the Z+jets sample is corrected following the same procedure as for $t\bar{t}$ +jets. Previously derived $t\bar{t}$ +jets corrections are applied to the Z+jets control region, so that the remaining mismodelling can be attributed to Z+jets. Correction factors are computed as:

$$R(x) = \frac{\text{Data}(x) - \text{MC}_{\text{other}}(x)}{\text{MC}_{Z+\text{jets}}(x)} \quad (9.7)$$

over the binned N_{jets} distribution in the same range as for the $t\bar{t}$ +jets correction. Figure 9.15 shows the data versus MC comparison over the N_{jets} distribution in the Z+jets control region before and after applying the correction. The values of each coefficient are shown in Table 9.14.

9. Analysis overview




Figure 9.15: Data versus MC comparison of the N_{jets} distribution (a) before and (b) after applying the N_{jets} correction to the Z+jets background. Systematic uncertainties are not included.


N_{jets}	Weight
2	1.083 ± 0.010
3	1.137 ± 0.010
4	1.233 ± 0.013
5	1.304 ± 0.018
6	1.31 ± 0.03
7	1.36 ± 0.04
≥ 8	1.42 ± 0.05

Table 9.14: Correction factors for each bin of the Z+jets N_{jets} distribution.


The p_T correction for Z+jets is carried out using the transverse momentum of the Z boson instead of H_T for two reasons: it is more precisely measured, given that it is calculated from leptons only, and it is not strongly correlated with jet multiplicity. As a result, no additional variables need to be defined. Similarly to the $t\bar{t}$ +jets H_T reweighting, this correction is calculated by computing $R(x)$ over the Z p_T distribution and fitting it to a continuous function. In this case, the expression of choice is a parabolic function with an exponential tail:

$$f(x) = c_1 + c_1(x + c_2)^2 e^{-c_3(x - c_4)}. \quad (9.8)$$

The result of this fit is shown in Figure 9.16. The values of the fit parameters are summarised in Table 9.15. Additionally, the N_{jets} distribution before and after the Z p_T correction is shown in Figure 9.17 with no significant differences. The effect on other p_T distributions can be seen in Figures 9.18 and 9.19. Note that the p_T distributions for jets 3 and 4 are not shown because of the high overlap with the signal regions.

Figure 9.16: Z+jets correction factor as a function of $Z p_T$.

Parameter	Value
c_0	0.972 ± 0.007
c_1	0.00003 ± 0.00003
c_2	-43 ± 5
c_3	0.0104 ± 0.0019
c_4	30 ± 100

Table 9.15: Values of the parameters from the Z+jets $Z p_T$ fit.Figure 9.17: N_{jets} distribution (a) before and (b) after applying the $Z p_T$ correction. Systematic uncertainties are not included.

9. Analysis overview




Figure 9.18: Leading (1) and sub-leading (2) jet p_T distributions (a) before and (b) after applying the $Z p_T$ correction. Systematic uncertainties are not included.




Figure 9.19: Leading (1) and sub-leading (2) lepton p_T distributions (a) before and (b) after applying the $Z p_T$ correction. Systematic uncertainties are not included.

9.5 Analysis strategy

This section describes the various machine learning techniques employed for reconstructing the $H \rightarrow aa \rightarrow 4b$ final state and create a set of signal and control regions to be used in the final fit. This is done in two steps. First, a mass parametrised NN is used to identify the optimal pairing hypothesis for the final-state b -objects. This allows to define variables such as invariant masses and angular distances based on the best pairing. Then, these variables are used as inputs to a BDT for signal versus background discrimination. The resulting BDT score distribution is the variable used in the fit to calculate the expected and observed limits.

9.5.1 Event reconstruction for $H \rightarrow aa \rightarrow 4b$

The four b -hadrons from $a \rightarrow b\bar{b}$ decay can be reconstructed as different b -objects, which include B -jets, b -jets and soft secondary vertices. In order to classify the different final states, a hypothesis testing NN is trained to identify the best combination under the signal hypothesis. The $4b$ final state can be reconstructed as multiple combinations of b -objects with varying numbers of input candidates, as shown in Figure 9.20. The possible categories include two B -jets (2B), one B -jet and two b -jets (1B2b), one B -jet, one b -jet and one soft secondary vertex (1B1b1v), four b -jets (4b) and three b -jets and one soft secondary vertex (3b1v). B -jets are tagged using the DeXTer 60% WP and b -jets, using the DL1r 85% WP. This causes an issue with conventional NNs, which require fixed-size inputs. However, it is covered by using DeepSets [217], which is an advanced NN architecture capable of handling input vectors with variable length. Using this method, the different b -object reconstruction information is mapped into a common latent space vector that can handle the various topologies simultaneously.

The basic structure of the NN used for event reconstruction is shown in Figure 9.21. It can be divided into three parts: the object encoder, the a -boson reconstruction and the Higgs reconstruction. All input features to the object encoders and the Higgs reconstruction NN are summarised in Table 9.16. The input candidates are encoded into common-size latent space vectors with dimension N , which are later propagated to a permutation invariant sum pooling layer. The aggregated information goes through a parametrised shared weight a -boson NN to learn the latent space representation of the a -boson. The a -boson NN is parametrised by an extra parameter, m_a , representing the a -boson mass of the tested hypothesis. The last sum pooling layer enforces the permutation invariant between the two a -bosons. The summed a -boson vector is concatenated with extra information from Z kinematics as input to the final Higgs NN. A sigmoid function is used in the output layer to predict the score of a given pairing hypothesis presented to the hypothesis testing NN. A correct pairing hypothesis will have a score of 1, whereas an incorrect pairing hypothesis based on the truth record will have a score of 0. For each event, the hypothesis with the highest score is the predicted pairing hypothesis according to the NN. The predicted pairing hypothesis allows to identify the optimal pairing strategy in order to compute relevant kinematic variables. When applied to the background, the pairing hypothesis score provides an estimate of how confident one can be that a given event looks like a $H \rightarrow aa \rightarrow 4b$




Figure 9.20: Illustration of all object combinations considered for the hypothesis testing NN in the $4b$ final state. The B -jet can either represent a boosted $a \rightarrow b\bar{b}$ decay reconstructed as one single Ak8 jet ($\textcircled{B}\square$) or a merged jet containing one b from each a -boson ($\textcircled{B}\square$).

event based on the input information. This can be exploited to improve signal versus background discrimination and to define signal and control regions. These two topics are covered in Sections 9.5.2 and 9.5.3.

The NN is trained using signal and background simulated events, which contain information about the true pairing that can be used to evaluate the training performance. Only the correct pairing from a signal process is labelled as 1. Incorrect pairings from the signal and any pairings from the background are labelled as 0.

9. Analysis overview




Figure 9.21: Diagram of the NN used for event reconstruction. All filled blocks are dense multilayer perceptions (MLPs) that share common weights between MLPs of the same colour. This particular example shows the case of a hypothesis with one Ak8 jet candidate, two Ak4 jet candidates and one soft secondary vertex candidate.

Object	Feature	Description
Ak8 jets	$\log m$	DNN-corrected track jet mass
	$\log p_T$	Transverse momentum
	η	Pseudorapidity
	ϕ	Azimuthal angle
	isDeXTer60WP	True if the jet is DeXTer tagged with 60% WP
	isDeXTer40WP	True if the jet is DeXTer tagged with 40% WP
Ak4 jets	$\log m$	Invariant mass
	$\log p_T$	Transverse momentum
	η	Pseudorapidity
	ϕ	Azimuthal angle
	DL1r b -tagging score	PC DL1r tagging score
Soft v	$\log m$	Track mass
	$\log p_T$	Transverse momentum
	η	Pseudorapidity
	ϕ	Azimuthal angle
	L_{3D}	Decay length relative to the primary vertex
	$S_{L_{3D}}$	Decay length significance
Z boson candidate (ee or $\mu\mu$)	p_T	Transverse momentum
	η	Pseudorapidity
	ϕ	Azimuthal angle
	m	Invariant mass

Table 9.16: List of features used as inputs to the hypothesis testing NN.

Examples of the pairing hypothesis score distributions are shown in Figures 9.22 to 9.24 for the different categories contemplated in the 4b final state: 2B, 1B2b, 1B1b1v, 3B1v and 4b. Very low values of m_a , such as 12 GeV, are more likely to fall in the 2B category. Masses around 20 GeV can populate regions such as 1B2b and 1B1b1v. For higher values of m_a , the 4b and 3b1v regions are the most sensitive. With the exception of the 2B region, it is observed that a large fraction of background events have a NN score below 0.05.




Figure 9.22: Highest hypothesis score distribution in the 2B region for $m_a = 12$ GeV.




Figure 9.23: Highest hypothesis score distribution in the (a) 1B2b and (b) 1B1b1v regions for $m_a = 20$ GeV.




Figure 9.24: Highest hypothesis score distribution in the (a) 4b and (b) 3b1v regions for $m_a = 40$ GeV.

9.5.2 BDT for signal versus background discrimination

The hypothesis testing NN selects the best combination of four b -objects that reconstructs the decay chain $H \rightarrow aa \rightarrow 4b$. Since the main sources of background do not have a similar cascade decay, the score of the best hypothesis provides a strong signal-discrimination. However, it does not fully exploit the kinematics of the dominant background sources, which can provide additional separation between signal and background. To optimise the result, a new multivariate observable is obtained by training a BDT to discriminate between signal and background events, using reconstructed variables from the NN as inputs.

The objects selected by the event reconstruction as coming from the $H \rightarrow aa \rightarrow 4b$ decay are used to build a representation of the four final-state particles in the decay of the two a -bosons. In the case of objects reconstructed from the hadronisation of a single b -quark, like a DL1r tagged Ak4 jet or a TC-LVT tagged soft secondary vertex, the association is trivial. In the case of DeXTer tagged Ak8 jets, the four-momentum of the PF1ow jet is decomposed in two components using projections onto the $\text{Ex}k_T^{(2)}$ track subjets. The four b 's are represented as b_{11} , b_{12} , b_{21} and b_{22} , where the first index distinguishes the two a -bosons and the second differentiates between the two b -quarks from each of them. Then, discriminating variables depending on the mass and spin of the $H \rightarrow aa \rightarrow 4b$ decay are defined. For example, the reduced mass of the a -bosons:

$$m_{a_1}^{\text{red}} = \sqrt{(p_{b_{11}} + p_{b_{12}})^2 - m_a}, \quad m_{a_2}^{\text{red}} = \sqrt{(p_{b_{21}} + p_{b_{22}})^2 - m_a}, \quad (9.9)$$

where m_a refers to the true mass hypothesis. Similarly, one can compute the reduced

mass of the Higgs boson:

$$m_H^{\text{red}} = \sqrt{(p_{b_{11}} + p_{b_{12}} + p_{b_{21}} + p_{b_{22}})^2 - m_H^2 - m_{a_1}^{\text{red}} - m_{a_2}^{\text{red}}}, \quad (9.10)$$

where $m_H = 125$ GeV. Other interesting variables that provide additional discrimination are $\cos \theta^*$, which is the polar angle of the leading a -boson in the Higgs boson rest frame, and the Collins-Soper angle, $\cos \theta_{\text{CS}}$, defined as:

$$\cos \theta_{\text{CS}} = \frac{2(\ell_1^+ \ell_2^- - \ell_1^- \ell_2^+)}{m_{\ell_1 \ell_2} \sqrt{m_{\ell_1 \ell_2}^2 + p_{T, \ell_1 \ell_2}^2}}, \quad (9.11)$$

with $\ell_i^\pm = (E_i \pm p_{z,i})/\sqrt{2}$, where E_i and $p_{z,i}$ are the energy and momentum along the z -direction of lepton i . These two angles exploit the angular properties associated to the spin-0 and spin-1 nature of the Higgs and Z bosons, respectively.

Table 9.17 summarises all the input features used in the BDT training. Examples of some of these distributions are shown in Figure 9.25.

Feature	Description
$a_1 p_T$	Transverse momentum of the leading a
$a_2 p_T$	Transverse momentum of the sub-leading a
$H p_T$	Transverse momentum of the Higgs boson
$m_{a_1}^{\text{red}}$	Reduced mass of the leading a
$m_{a_2}^{\text{red}}$	Reduced mass of the sub-leading a
m_H^{red}	Reduced mass of the Higgs boson
$\cos \theta^*$	Cosine of the a -bosons' polar angle in the Higgs boson rest frame
$\cos \theta_{\text{CS}}$	Cosine of the Collins-Soper angle
$Z p_T$	Transverse momentum of the Z boson candidate
$m_{\ell\ell}$	Invariant mass of the Z boson candidate
$ \Delta\eta_{Z,H} $	Pseudorapidity distance between the Z boson candidate and the Higgs boson
$ \Delta\phi_{Z,H} $	Azimuthal angle between the Z boson candidate and the Higgs boson
NN score	Hypothesis testing NN highest hypothesis score

Table 9.17: List of features used as inputs to the hypothesis testing NN.

The training of the BDT is performed using the XGBoost package. The depth of each tree is chosen depending on the acceptance of each mass reconstruction category. Cases with large acceptance use trees with four splittings, while cases with small acceptance have two splittings. This choice is made to avoid overtraining, which is assessed by dividing the full dataset into a training dataset (60% of events) and a validation dataset (40% of events). Signs of overtraining are tested by comparing the performance on the training and validation datasets. One BDT is trained for each NN reconstruction hypothesis and a -boson mass, using signal and background simulated events. The final BDT score distribution is arranged in three bins, with low, intermediate and high signal content. They are denoted as Loose, Medium and Tight, respectively. Figure 9.26 shows examples of the BDT score distribution for a subset of m_a values for different reconstruction and m_a hypotheses.

9. Analysis overview




Figure 9.25: Examples of BDT input distributions for different reconstruction and a -boson mass hypotheses. From top to bottom: leading $a p_T$ and E_T^{miss} for the $m_a = 12$ GeV hypothesis in the 2B region, $\cos \theta^*$ and reduced Higgs boson mass for the $m_a = 20$ GeV hypothesis in the 1B2b region, reduced leading and sub-leading a -boson mass for the $m_a = 40$ GeV hypothesis in the 4b region.




Figure 9.26: Examples of BDT score distributions for different reconstruction and a -boson mass hypotheses. From top to bottom, $m_a = 12$ GeV hypothesis in the 2B region, $m_a = 20$ GeV hypothesis in the 1B2b and 1B1b1v regions and $m_a = 40$ GeV hypothesis in the 4b and 3b1v regions.

9.5.3 Signal and control regions

Signal regions

Signal regions targeting the $H \rightarrow aa \rightarrow 4b$ final state are required to fulfill the preselection requirements in Section 9.3. The signal is produced in association with a Z boson, reason why the two leptons are required to have the same flavour (SF), either ee or $\mu\mu$, and opposite charge. Additionally, the invariant mass of the dilepton system is restricted to be within the range $71 \leq m_{\ell\ell} \leq 111$ GeV, denoted as the loose Z -mass window. Events in the signal region are required to have a score from the hypothesis testing NN above 0.05, which discards a large fraction of background events. The only exception is the 2B region, where the NN score is set to be > 0.5 to enhance statistics in the complementary control regions, defined in the next paragraph. There are five signal regions in total, one for each NN pairing hypothesis. Table 9.18 summarises the five signal regions and the selection they correspond to. Figure 9.27 shows the background composition of the five signal regions in the $m_a = 12, 20$ and 40 GeV mass hypotheses.

Region	SF	Z-mass	NN score
2B	Yes	Loose	> 0.5
1B2b	Yes	Loose	> 0.05
1B1b1v	Yes	Loose	> 0.05
4b	Yes	Loose	> 0.05
3b1v	Yes	Loose	> 0.05

Table 9.18: Signal regions for the $H \rightarrow aa \rightarrow 4b$ analysis.

Z+jets control regions

The Z +jets control regions are selected to be enriched in Z -boson events while staying close to the signal regions. They fulfill the preselection requirements from Section 9.3. Additionally, the two leptons in the event are required to have the same flavour (ee or $\mu\mu$) and opposite charge, and their invariant mass must be in the range $81 \leq m_{\ell\ell} \leq 101$ GeV. This is referred to as the tight Z -mass window. In order to increase the purity of the control regions, a missing transverse momentum cut $E_T^{\text{miss}} < 60$ GeV is included. This requirement reduces the contribution from $t\bar{t}$ +jets. The hypothesis testing NN score selection is the opposite as in the signal regions, that is ≤ 0.5 in the 2B region and ≤ 0.05 in the rest of control regions. Note that this guarantees orthogonality between SRs and CRs associated to the same mass hypothesis, but there can be overlap among SRs or CRs for different values of the a -boson mass.

On top of this, the Z +jets control regions are split according to the number of B -jets in each event. Two types of B -jets are contemplated: tight (T) B -jets (DeXTer tagged at 40% WP) and loose (L) B -jets (DeXTer tagged at 60% WP but not at 40% WP). Table 9.19 contains the list of control regions used for the Z +jets background. Three columns with the number of b -jets, soft secondary vertices and non-tagged jets are included in order to clarify which regions are inclusive in these b -objects and which ones are not. For b -jets, only the DL1r 85% WP is used, given that it is the only WP calibrated down

to $p_T = 15$ GeV. Figure 9.28 shows the background composition of the six Z +jets CRs in the $m_a = 12, 20$ and 40 GeV mass hypotheses.

Region	SF	Z-mass	NN score	N_B^{tight}	N_B^{loose}	N_b	N_ν	N_{jets}
SFonZ.1i.2i(T+L).0i	Yes	Tight	≤ 0.5	≥ 1	≥ 2 (T+L)	≥ 0	≥ 0	≥ 0
SFonZ.0.2i.0i	Yes	Tight	≤ 0.5	$=0$	≥ 2	≥ 0	≥ 0	≥ 0
SFonZ.1.0.2i	Yes	Tight	≤ 0.05	$=1$	$=0$	≥ 2	≥ 0	≥ 0
SFonZ.0.1.2i	Yes	Tight	≤ 0.05	$=0$	$=1$	≥ 2	≥ 0	≥ 0
SFonZ.0.0.4i	Yes	Tight	≤ 0.05	$=0$	$=0$	≥ 4	≥ 0	≥ 0
SFonZ.0.0.3.1ivorj*	Yes	Tight	≤ 0.05	$=0$	$=0$	$=3$	$\geq 1^*$	$\geq 1^*$

Table 9.19: Control regions for the Z +jets background. Each control region exists for every mass hypothesis, as they all depend on the score of the hypothesis testing NN. SFonZ indicates that the two leptons have the same flavour and are within the Z -mass window. The naming scheme can be read as (number of B tight).(number of B loose).(number of b). The letter i represents that the selection is inclusive. *"vorj" means that $N_{\text{jets}} \geq 1$ is required only if $N_\nu = 0$.

$t\bar{t}$ +jets control regions

The $t\bar{t}$ +jets control regions also fulfill the preselection requirements from Section 9.3. These regions benefit from the fact that the $t\bar{t}$ system can decay to ee , $\mu\mu$ and $e\mu$. In order to suppress both the signal and the Z +jets background, leptons are required to have different flavour (DF), that is, $e\mu$, and opposite charge. No additional cuts on the invariant mass or the hypothesis testing NN are applied. The list of $t\bar{t}$ +jets control regions is shown in Table 9.20. Figure 9.29 shows the background composition of the six $t\bar{t}$ +jets CRs. They are common for all a -boson mass hypotheses.

Region	SF	Z-mass	NN score	N_B^{tight}	N_B^{loose}	N_b	N_ν	N_{jets}
DF.1i.2i(T+L).0i	No	–	–	≥ 1	≥ 2 (T+L)	≥ 0	≥ 0	≥ 0
DF.0.2i.0i	No	–	–	$=0$	≥ 2	≥ 0	≥ 0	≥ 0
DF.1.0.2i	No	–	–	$=1$	$=0$	≥ 2	≥ 0	≥ 0
DF.0.1.2i	No	–	–	$=0$	$=1$	≥ 2	≥ 0	≥ 0
DF.0.0.4i	No	–	–	$=0$	$=0$	≥ 4	≥ 0	≥ 0
DF.0.0.3.1ivorj*	No	–	–	$=0$	$=0$	$=3$	$\geq 1^*$	$\geq 1^*$

Table 9.20: Control regions for the $t\bar{t}$ +jets background. DF indicates that the two leptons have different flavour. The naming scheme can be read as (number of B tight).(number of B loose).(number of b). The letter i represents that the selection is inclusive. *"vorj" means that $N_{\text{jets}} \geq 1$ is required only if $N_\nu = 0$.

Ph. D. Thesis

P. Martinez

Signal regions





Figure 9.27: Background composition in the five signal regions under the $m_a = 12, 20$ and 40 GeV mass hypotheses. All signal regions are dominated by the Z +jets background due to the SF and Z -mass cuts. In general, the $t\bar{t}$ +light background contributes significantly in the boosted regime (2B, 1B2b and 1B1b1v), while $t\bar{t}$ +HF is more prominent in the resolved regime (4b, 3b1v).

Ph. D. Thesis

P. Martinez


Control regions
SFonZ

Figure 9.28: Background composition in the six Z+jets control regions under the $m_a = 12, 20$ and 40 GeV mass hypotheses. All Z+jets control regions are dominated by the Z+jets background due to the SF and Z-mass cuts. The E_T^{miss} veto contributes to suppress the $t\bar{t}$ +jets background as well.

Ph. D. Thesis

P. Martinez

Control regions
DF

■ Z + jets
■ tt + light
■ tt + $\geq 1c$
■ tt + $\geq 1b$
■ tW
■ Other




Figure 9.29: Background composition in the six $t\bar{t}$ +jets control regions, common to all mass hypotheses. All $t\bar{t}$ +jets control regions are dominated by the $t\bar{t}$ +jets background due to the DF cut. The $t\bar{t}$ +light category constitutes most of the 2B control regions (1i.2i(T+L).0i and 0.2i.0i), while others have a larger fraction of $t\bar{t}$ +HF in comparison.

9.6 Systematic uncertainties

Various sources of systematic uncertainties are considered in this analysis. Each systematic uncertainty is introduced as a nuisance parameter in the statistical analysis described in Chapter 8. Section 9.6.1 describes all experimental uncertainties, related to the luminosity, the pileup and the reconstruction and identification of jets and leptons. These uncertainties are applied to all MC samples equally. Signal and background modelling uncertainties are detailed in Section 9.6.2. They are associated to the theoretical predictions for each process, and can vary depending on the sample. Systematic uncertainties are corrected according to the data-driven reweighting described in Section 9.4.

9.6.1 Experimental uncertainties

Luminosity and pileup

The uncertainty on the integrated luminosity for the full Run 2 dataset is 0.83% [73], obtained using the LUCID-2 detector [194].

A variation in the pileup reweighting of simulated events is included to cover the uncertainty in the ratio of the predicted and measured inelastic cross sections in the fiducial volume defined by $M_X > 13$ GeV, where M_X is the mass of the hadronic system [218]. In practice, the nominal scale factor applied on the data pileup distribution when performing the pileup reweighting is changed into 1.0/0.99 or 1.0/1.07 to derive the up and down systematic uncertainty, instead of its nominal value of 1.0/1.03.

Leptons

Lepton systematic uncertainties are related to the trigger, reconstruction, identification and isolation, as well as the lepton energy or momentum scale and resolution.

The reconstruction, identification, and isolation efficiency of electrons and muons, as well as the efficiency of the trigger used to record the events, differ slightly between data and simulation, and is corrected by dedicated scale factors. Efficiency scale factors are measured using tag-and-probe techniques on $Z \rightarrow \ell\ell$ data and simulated samples [219, 220], and are applied to the simulation to correct for differences. The effect of these scale factors and their uncertainties are propagated as corrections to the MC event weight. In total, four independent components are considered for electrons and eight for muons.

Additional sources of uncertainty originate from the corrections applied to adjust the lepton momentum scale and resolution in the simulation to match those in data, measured using reconstructed distributions of the $Z \rightarrow \ell\ell$ and $J/\psi \rightarrow \ell\ell$ invariant masses, as well as the E/p ratio measured in $W \rightarrow e\nu$ events where E and p are the electron energy and momentum measured by the calorimeter and the tracker, respectively [133, 220]. To evaluate the effect of momentum scale uncertainties, the event selection is redone with the lepton energy or momentum varied by $\pm 1\sigma$. The event selection is redone for the momentum resolution uncertainties by smearing the lepton energy or momentum. In total, two independent components are considered for electrons, and four for muons.

Jets

Uncertainties associated with jets arise from the efficiency of pileup rejection by the Jet Vertex Tagger (JVT), from the Jet Energy Scale (JES) and Resolution (JER), and from the different flavour tagging algorithms, which include DL1r, DeXTer and TC-LVT in this analysis.

Scale factors are applied to correct for discrepancies between data and MC for JVT efficiencies. These scale factors are estimated using $Z \rightarrow \mu\mu$ events with tag-and-probe techniques [221]. The effect of these scale factors, as well as of their uncertainties, are propagated as corrections to the MC event weight.

The jet energy scale and its uncertainty are derived by combining information from test-beam data, LHC collision data and simulation [222]. The uncertainties from these measurements are factorised into eight independent sources. The jet energy resolution was measured in Run 2 data and MC simulation as a function of the jet p_T and rapidity using di-jet events with a similar method as described in Ref. [223]. The combined uncertainty is propagated by smearing the jet p_T in the MC simulation, yielding another eight independent sources.

To correct DL1r flavour tagging efficiencies in simulated samples to match efficiencies in data, scale factors are derived by the flavour-tagging working group. They are calculated as a function of p_T for b -jets, c -jets and light jets separately in dedicated calibration analyses. For b -jet efficiencies, $t\bar{t}$ -jets events in the dilepton topology are used, exploiting the very pure sample of b -jets arising from the decay of the top

quarks [224]. For c -jet mistag rates, $t\bar{t}$ +jets events in the single-lepton topology are used, exploiting the c -jets from the hadronically decaying W boson [225]. The so-called negative-tag method is used in Z +jets events [226] for the light-jets mistag rates. In the three calibration analyses, a large number of uncertainty components are considered, and a principal component analysis is performed, yielding 45, 20, and 20 eigen-variations for b -, c - and light-jets, respectively, which are taken as uncorrelated sources of uncertainty. These eigen-variations correspond to the number of p_T bins (9, 4, and 4, respectively), multiplied by the number of DL1r bins (5). Dedicated calibrations were performed to extend jets down to $p_T = 15$ GeV for b - and light-jets. The uncertainty is 11.5% for the b -jet scale factor and 21.9% for the light-jet scale factor. For c -jets, an inflated uncertainty from the closest (lowest) p_T bin of the standard c -jet calibration is considered, which is 11%.

The use of DeXTer introduces additional scale factors to correct the efficiency difference between simulated samples and data. The scale factors for the B -tagger are provided by the DeXTer development team in collaboration with the ATLAS flavour-tagging working group [153]. They are derived for B - and b -jets, separately for each p_T bin and for each B -tagging WP. The calibration measurements with data are performed using both $t\bar{t}$ +jets and Z +jets events simultaneously in order to measure the B -jet tagging and b -jet mistagging efficiencies.

To correct for a potential mismodelling in the reconstruction efficiency and fake rate of the TC-LVT tagger, a calibration is carried out [227]. An efficiency scale factor and a fake rate scale factor are derived simultaneously in dileptonic $t\bar{t}$ +jets events, separated into regions with different b -jet multiplicities, different number of soft secondary vertices and different levels of pileup. Three independent uncertainty components are derived, one for the efficiency scale factor and two mistag rate scale factors, for low and high pileup conditions.

Tracks

Systematic uncertainties related to the track selection efficiency are determined by changing the amount of tracker material and the physical models in the GEANT4 simulation. Systematic variations on the number of fake tracks are applied based on the recommendation for the Loose track selection working point estimated and provided by the ATLAS tracking combined performance group [228, 229]. Dedicated systematic uncertainties are considered for the track parameters, including the transverse and longitudinal impact parameters and the track sagitta. The event selection is redone for each systematics variation.

To correct the mismodelling in the Ak8 track jet mass, additional mass scale corrections are estimated from data. The Ak8 track-jet mass scale is varied by $\pm 5\%$ and compared to the nominal results with a mass scale of 1.01.

Missing transverse energy

All previously described uncertainties on energy scales or resolutions of the reconstructed objects (hard components) are propagated to the missing transverse momentum. Additional uncertainties in the scale and resolution of the soft term are

considered, to account for the disagreement between data and MC for the p_T balance between the hard and soft components, for a total of three independent sources: an offset along the hard component p_T axis and the smearing resolution along and perpendicular to this axis [154, 230].

Summary

Table 9.21 summarises all the experimental systematic uncertainties, including the type (normalisation, shape or both) and the number of components. They are applied equally to all MC samples.

Systematic uncertainty	Type	Components
General		
Luminosity	N	1
Pileup	SN	1
Physics objects		
Electrons	SN	6
Muons	SN	12
Jets	SN	32+1
Tracks	SN	13+1
E_T^{miss}	SN	3
Flavour tagging		
DL1r b	SN	45+1
DL1r c	SN	20+1
DL1r light	SN	20+1
DeXTer	SN	12
TC-LVT	SN	3

Table 9.21: List of experimental systematic uncertainties included in the analysis. "N" means that the uncertainty is taken as normalisation-only for all processes and channels affected, while "SN" means that the uncertainty is taken on both the shapes and the normalisation. Some of the systematic uncertainties are split into several components for a more accurate treatment. Additional systematic uncertainties due to the low- p_T jets and the Ak8 track mass scale, particular to this analysis, are shown separately with a "+" sign.

9.6.2 Modelling uncertainties

General modelling uncertainties

General modelling uncertainties are those associated to the missing higher orders in the perturbative expansion of the partonic cross section and the uncertainties in the proton PDFs. These uncertainties are applied to the signal and to the two main backgrounds, $t\bar{t}$ +jets and Z +jets. They are assumed to be negligible for the rest of the samples.

Variations in the renormalisation and factorisation scales are used to estimate the uncertainty due to missing higher order corrections. The samples used include on-the-fly variations corresponding to seven different configurations of these two parameters:

9. Analysis overview

$\{\mu_R, \mu_F\} \times \{0.5, 0.5\}, \{0.5, 1\}, \{1, 0.5\}, \{1, 1\}, \{2, 1\}, \{1, 2\}, \{2, 2\}$. The uncertainties are combined by taking an envelope of all the variations in each bin of the fitted distribution:

$$\max[O(\mu_{R,i}, \mu_{F,i}) - O(\mu_{R,0}, \mu_{F,0})], \quad (9.12)$$

where O is the value of the observable in that bin, i refers to the i th scale variation and 0 refers to the central value.

The uncertainty due to ISR is estimated by simultaneously changing μ_R and μ_F in the ME and μ_R^{ISR} in the PS, while the uncertainty due to FSR is estimated by changing μ_R^{FSR} in the PS. For the ISR, the amount of radiation is increased (decreased) by scaling μ_R and μ_F by a factor of 0.5 (2) and by varying the renormalisation scale for QCD emission in the ISR by a factor of 0.549 (1.960), corresponding to the Var3cUp (Var3cDown) variation from the A14 tune which sets α_s^{ISR} to 0.140 (0.115) instead of the nominal 0.127. For the FSR, the amount of radiation is increased (decreased) by varying the renormalisation scale for QCD emission in the FSR by a factor of 0.5 (2), corresponding to $\alpha_s^{\text{FSR}} = 0.1423$ (0.1147) instead of the nominal 0.127. This uncertainty is not available for the Z +jets sample. Additional Z +HF modelling uncertainties are derived by comparing the nominal sample with alternative ones with varied settings for the overlap between ME and PS emissions and for the resummation scale [231]. The ME matching uncertainty (CKKW) is estimated by varying the scale used to calculate the overlap between jets from the ME and the PS. The nominal value for this parameter is 20 GeV. The up variation corresponds to a value of 30 GeV, while the down variation corresponds to a value to 15 GeV. The resummation scale (QSF) uncertainty is estimated by varying the scale used for the resummation of soft gluon emission by a factor of 2 and by a factor of 0.5 with respect to the nominal scale. The total uncertainty is derived from the sum in quadrature of the two uncertainties.

The PDF uncertainties follow the PDF4LHC recommendations [108] and are available as weights in the samples. The α_s uncertainty is derived using the same PDF set evaluated with two different α_s values. The two uncertainties from the PDF and α_s are added in quadrature.

Signal modelling uncertainties

The nominal $H \rightarrow aa \rightarrow 4b$ samples generated using POWHEGBox-v2 + PYTHIA 8 are compared to POWHEGBox-v2 + HERWIG 7 to extract a parton shower modelling uncertainty.

Background modelling uncertainties

The impact of the matrix element corrections applied to the top decay in POWHEGBox and the NLO matching uncertainty in $t\bar{t}$ +jets are evaluated by comparing the nominal sample to an alternative sample of events generated with MADGRAPH5_aMC@NLO v2.6.0 interfaced with PYTHIA 8.230. For parton shower and hadronisation uncertainties, the nominal POWHEGBox + PYTHIA 8 sample is compared to the POWHEGBox + HERWIG 7 alternative. Due to the unavailability of $t\bar{t}+b\bar{b}$ 4FS alternative samples at the time of the analysis, 5FS alternative samples are used to build the ME and PS uncertainties for $t\bar{t}$ +light, $t\bar{t}+\geq 1c$ and $t\bar{t}+\geq 1b$.

The nominal Z +jets sample is compared to an alternative sample simulated with MADGRAPH5_aMC@NLO v2.6.5 with NLO matrix elements for up to three partons, using the NNPDF3.0 NLO set of PDFs. Events are interfaced with PYTHIA 8.240 for the modelling of the parton shower, hadronisation, and underlying event, in the A14 tune. The different jet multiplicities are merged using the FxFx prescription [232].

Reweightings uncertainties

In order to propagate the errors coming from the data-driven reweighting (RW) procedure, the $t\bar{t}$ +jets and Z +jets RW functions are assumed not to be correlated. This can be done because the $t\bar{t}$ +jets RW function is calculated in a region where Z +jets events are largely suppressed. For the $t\bar{t}$ +jets background, the first step is to perform a HF normalisation. This is done using TREXFITTER, which provides three normalisation factors with their corresponding uncertainties and correlation matrix. Because the correlations are not negligible, three variations (up and down) are calculated by diagonalising the covariance matrix. The other two steps apply to both $t\bar{t}$ +jets and Z +jets. For the N_{jets} correction, because it is calculated per-bin, each point in $R(x)$ (Equation 9.3) is uncorrelated. The error per bin is simply the propagated statistical error. For the H_T (or $Z p_T$) correction, there is an additional step, which is a fit to a continuous function. In this case, the errors are derived from the fit results, diagonalising the covariance matrix to obtain as many variations as free parameters in the fit. These variations are applied by implementing an alternative reweighting function for each of them. The alternative reweighting functions are derived by replacing the nominal value by a variation, and propagating it through the remaining RW steps.

Summary

Table 9.22 includes the weight-based modelling uncertainties. Table 9.23 shows the alternative MC samples used to derive the MC generator and parton shower uncertainties.

9. Analysis overview

Systematic uncertainty	Type	Components
Signal		
μ_R, μ_F	SN	3
ISR, FSR	SN	2
PDF+ α_s	SN	32
$t\bar{t}$+jets		
μ_R, μ_F	SN	3
ISR, FSR	SN	2
PDF+ α_s	SN	31
Reweighting	SN	16
tW		
Reweighting	SN	16
Z+jets		
μ_R, μ_F	SN	3
CKKW+QSF	N	1
PDF+ α_s	SN	101
Reweighting	SN	13

Table 9.22: List of modelling systematic uncertainties included in the analysis. "N" means that the uncertainty is taken as normalisation-only for all processes and channels affected, while "SN" means that the uncertainty is taken on both the shapes and the normalisation. Some of the systematic uncertainties are split into several components for a more accurate treatment. Systematics derived from alternative samples are not included.

Process	ME generator	PDF set	PS generator	Sim.
$H \rightarrow aa \rightarrow 4b$	POWHEGBox-v2	NNPDF3.0 NLO	PYTHIA 8.244	AFII
	POWHEGBox-v2	NNPDF3.0 NLO	HERWIG 7.0.4	AFII
$t\bar{t}$ + jets	POWHEGBox-v2	NNPDF3.0 NLO	PYTHIA 8.230	AFII
	MG5_aMC@NLO v2.6.0	NNPDF3.0 NLO	PYTHIA 8.230	AFII
	POWHEGBox-v2	NNPDF3.0 NLO	HERWIG 7.1.3	AFII
$t\bar{t} + b\bar{b}$	POWHEGBox-v2	NNPDF3.0 NLO	PYTHIA 8.230	AFII
	MG5_aMC@NLO v2.6.0	NNPDF3.0 NLO	PYTHIA 8.230	AFII
	POWHEGBox-v2	NNPDF3.0 NLO	HERWIG 7.1.3	AFII
Z + jets	SHERPA 2.2.11	NNPDF3.1 NNLO	SHERPA	AFII
	MG5_aMC@NLO v2.6.5	NNPDF3.0 NLO	PYTHIA 8.240	AFII

Table 9.23: Summary of alternative MC samples used in the $H \rightarrow aa \rightarrow 4b$ analysis. Note that MADGRAPH is abbreviated to MG. All alternative samples are simulated using AFII. They are compared to the AFII version of the nominal FS MC, highlighted in grey.

Chapter 10

Results and conclusions

10.1 Fit setup

10.1.1 General considerations

Different distributions in each of the regions considered (control and signal regions) are combined to test for the presence of a signal. The statistical analysis is based on a binned likelihood function $\mathcal{L}(\mu, \theta)$ constructed as a product of Poisson probability terms over all bins, as described in Section 8.2. This function depends on the signal-strength parameter μ , a multiplicative factor to the theoretical signal production cross section for ZH production, and θ , a set of nuisance parameters that encode the effect of statistical or systematic uncertainties on the signal and background expectation values. Therefore, the total number of expected events in a given bin depends on μ and θ . The nuisance parameters θ allow variations of the expectations for signal and background according to the corresponding systematic uncertainties. Their fitted values correspond to the deviations from the nominal expectations that globally provide the best fit to the data. This procedure reduces the impact of systematic uncertainties on the search sensitivity by taking advantage of the highly populated background-dominated regions included in the likelihood fit. It requires a good understanding of the systematic variations affecting the shapes of the discriminant distributions, for which detailed validation studies have been performed using the simulation.

10.1.2 Normalisation factors for Z +jets and $t\bar{t}$ +jets

The normalisations of the $t\bar{t}$ +jets and Z +jets backgrounds are determined from the fit simultaneously with μ , and are constrained by the systematic uncertainties and the data. Statistical uncertainties in each bin are taken into account by dedicated parameters in the fit. Each of these two background normalisations is separated in three categories according to the flavour of the additional parton radiation different from those used in the previous chapter, in order to account for possible differences in their normalisation. They are defined as follows:

- $Z/t\bar{t} + \geq 1B$. Events with at least one B , defined as an Ak4 jet with $p_T > 15$ GeV and $|\eta| < 2.5$ matched to at least two b -hadrons that not originating from the $Z/t\bar{t}$ system.
- $Z/t\bar{t} + \geq 1b$. Events with no B 's and at least one b , defined as an Ak4 jet with $p_T > 15$ GeV and $|\eta| < 2.5$ matched to one b -hadron not originating from the $Z/t\bar{t}$ system.
- $Z/t\bar{t} + \text{light}/c$. Events with no B 's and no b 's, which can contain light and/or c radiation, defined as Ak4 jets with $p_T > 15$ GeV and $|\eta| < 2.5$ matched to c - or

10. Results and conclusions

light-hadrons not originating from the $Z/t\bar{t}$ system.

For the Z +jets background, only the $Z+\geq 1B$ and $Z+\geq 1b$ components are normalised in the fit, by one common free-floating factor that is constrained using the Z +jets CRs. Due to the lack of available statistics, the fit does not have the ability to constrain additional Z +HF categories separately. This is accounted for by the Z +HF modeling uncertainties instead. For the $t\bar{t}$ +jets background, the overall normalisation and the $t\bar{t}+\geq 1B$ and $t\bar{t}+\geq 1b$ normalisations are left free to float separately. They are constrained by the dedicated $t\bar{t}$ +jets CRs. In total, there are five free-floating parameters in the fit: the signal strength μ , the $Z+\geq 1B/b$ normalisation factor and the $t\bar{t}+\geq 1B$, $t\bar{t}+\geq 1b$ and $t\bar{t}$ +jets normalisation factors.

10.1.3 Binning, smoothing and pruning

The fit uses five signal regions, six control regions for Z +jets and six control regions for $t\bar{t}$ +jets. Signal regions are separated in three bins, depending on their BDT score (loose, medium or tight). Each control region enters the fit as one single bin.

Systematic uncertainties are derived from MC binned distributions for each sample in each signal and control region. This method is susceptible to statistical fluctuations that can arise when the number of events is very small. If these fluctuations become too large, they can lead to a bad template for the uncertainty. In order to mitigate these effects and avoid problems with the convergence of the minimisation, smoothing algorithms are applied to the histograms prior to the fit. Smoothing is applied to modelling uncertainties of the main backgrounds. For smaller backgrounds, for systematic uncertainties with weights close to unity, or when using the same events but with different calibrations, no smoothing is applied.

Systematic uncertainties are pruned in order to reduce the CPU time used for fitting and also to help with the convergence of the fit. This is done by considering bin-by-bin variations of each input histogram, relative to each sample in each region. The pruning threshold is set to 1%, meaning that if the effect of a nuisance parameter is smaller than that, then it does not contribute to the fit. Each nuisance parameter is separated into a shape component and a normalisation component, and the pruning criterion applies for the two of them independently.

10.2 Fit results

This Section summarises the results of the analysis. Expected results based on the Asimov dataset are presented in Section 10.2.1. Observed results from the fits to data are shown in Section 10.2.2. Only the fits under the 25 GeV mass hypothesis are shown. Expected and observed upper limits on the branching ratio of $H \rightarrow aa \rightarrow 4b$, together with the list of most relevant systematic uncertainties, are included in Section 10.2.3. Blinded fits to data, used for tests, are presented in Appendix B.1.


10.2.1 B-only fit to the Asimov dataset

This section shows the results from the fits to the Asimov dataset. In this test, the background-only (B-only) hypothesis is assumed ($\mu = 0$), where the Asimov dataset becomes identical to the background MC. The Asimov dataset is built such that the event count in each bin is set to the expected event yield for the chosen model parameters. The NPs are fitted in both the signal and control regions. This setup provides an expected exclusion limit, which illustrates the experimental sensitivity of the analysis under the assumption that no signal exists in the observed data, and it serves as a cross-check for the fit model.


Figure 10.1 shows the signal and background composition in each region, before and after performing the Asimov fit with $\mu = 0$. Figure 10.2 contains the normalisation factors for the different components of Z +jets and $t\bar{t}$ +jets and the signal strength, with their corresponding uncertainties. The largest uncertainty appears in the $t\bar{t} \geq 1B$ normalisation factor, which corresponds to the $t\bar{t}$ +jets category with less statistics. Figure 10.3 shows the correlations between NPs. Due to the choice of a single-bin CRs, some level of correlation is expected between the Z +jets and $t\bar{t}$ +jets NFs and their modelling uncertainties, as well as between the NFs themselves. Figure 10.4 shows the constraints on the relevant nuisance parameters, that is, those that have not been pruned. Some systematic uncertainties such as DeXTer B -tagging and low- p_T DL1r b -tagging are expected to be constrained, as the pre-fit uncertainties are propagated from the calibration in order to avoid profiling the same systematics twice in the fit. The Ak8 track mass scale uncertainty is also expected to be constrained. Some of the modelling uncertainties have a manifest degree of constraint as well. These are, for example, $Z \geq 1B/b$ scale variations, the $Z \geq 1B/b$ MADGRAPH5_aMC@NLO alternative sample or the $t\bar{t}$ +light HERWIG 7 alternative sample. These constraints are understood by studying the behaviour of the systematic uncertainty in each region individually. It is observed that regions that are dominated by $Z \geq 1B/b$ or $t\bar{t}$ +light with enough statistics prevent these NPs from varying within the full $\pm 1\sigma$ range, given their large impact on the cross section. Figures 10.5 to 10.7 illustrate the previously described constraints.

The ranked impact from the dominant systematics to the signal strength is shown in Figure 10.8. The most significant systematic uncertainties originate from Z +jets and $t\bar{t}$ +jets modelling, DeXTer B -tagging and DL1r b -tagging. The signal strength is also sensitive to the $Z \geq 1B/b$ and $t\bar{t}$ +jets normalisation factors.

10. Results and conclusions



(a)



(b)

Figure 10.1: Predicted yields for the signal and control regions in the B-only fit to the Asimov dataset with $m_a = 25$ GeV (a) before and (b) after the fit. The expected signal yield for the 25 GeV mass hypothesis assuming the SM production cross section for ZH and $\text{BR}(H \rightarrow aa \rightarrow 4b) = 1$ is overlayed in the pre-fit plot.

Ph. D. Thesis P. Martinez




Figure 10.2: Normalisation factors for the Z +jets and $t\bar{t}$ +jets samples in fit to the Asimov dataset with $m_a = 25$ GeV.

Ph. D. Thesis P. Martinez




Figure 10.3: Correlation matrix for the relevant NPs in the B-only fit to the Asimov dataset with $m_a = 25$ GeV.

10. Results and conclusions




Figure 10.4: Fitted nuisance parameters from the B-only fit to the Asimov dataset with $\sqrt{s} = 25$ GeV.




Figure 10.5: Up (red) and down (blue) 1σ variation of the third eigenvalue of the DeXTer systematic uncertainty in the 2B SR, corresponding to the (a) $Z + \geq 1b$ MC sample and the (b) 25 GeV signal MC sample.




Figure 10.6: Up (red) and down (blue) 1σ variation of the MADGRAPH5_aMC@NLO ME uncertainty in the (a) $Z + \text{jets}$ 4b CR and the (b) 4b SR, corresponding to the $Z + \geq 1b$ MC sample.

10. Results and conclusions




Figure 10.7: Up (red) and down (blue) 1σ variation of the HERWIG 7 PS uncertainty for $t\bar{t}$ +light in the (a) the $t\bar{t}$ +jets 4b CR and (b) the 4b SR.




Figure 10.8: Ranked systematic impacts to the signal strength from the B-only fit to the Asimov dataset with $m_a = 25$ GeV.


10.2.2 S+B fit to data

This section shows the results from the fit to data, which is used to derive the final limits. In this fit, the signal plus background (S+B) hypothesis is assumed, and the signal strength is left as a free-floating parameter.


Figure 10.9 summarises the data to MC comparison in both the SRs and CRs, before and after the fit. A good post-fit agreement between data and simulation can be observed in in all regions. Figure 10.10 shows the normalisation factors for the different components of Z +jets and $t\bar{t}$ +jets and their corresponding uncertainties. They are consistent with the values expected from the Asimov fit and the blinded fit to data. Figure 10.11 shows the correlations between NPs, which are also consistent with the previous tests. Figure 10.12 shows the pulls and constraints for all relevant NPs. There are some small constraints in the jet reconstruction systematics, which do not have a significant impact in the signal strength. Constraints in the flavour tagging and Ak8 track mass scale uncertainties are expected, as it was argued in the previous sections, due to the way that they are propagated from the calibration results. Constraints in the modelling systematics from the Z +jets and $t\bar{t}$ +jets samples are expected as well, due to the low statistical uncertainties in some of the fitted bins. Even though some NPs are pulled, they are all within $\pm 1\sigma$.

The fitted signal strength resulting from the S+B fit to data for $m_a = 25$ GeV is $\hat{\mu} = 0.11^{+0.05}_{-0.04}$. The ranked impact from the dominant systematics to the signal strength is shown in Figure 10.13. The leading systematic uncertainty is the PS modelling of the signal MC sample. Other relevant systematics are DeXTer B -tagging, DL1r b -tagging and modelling systematic uncertainties. Table 10.1 summarises the impact on $\hat{\mu}$ from the various sources of statistical and systematic uncertainty, grouped by their type. Finally, Tables 10.2 to 10.4 show the post-fit event yields in each of the signal and control regions for the 25 GeV a -boson mass hypothesis. For signal and background MC samples, the full systematic uncertainty is included.

10. Results and conclusions




(a)



(b)

Figure 10.9: Predicted yields for the signal and control regions in the S+B fit to data with $m_a = 25$ GeV (a) before and (b) after the fit. The fitted signal yield for the 25 GeV mass hypothesis is stacked in the post-fit plot.



10. Results and conclusions




Figure 10.12: Fitted nuisance parameters from the S+B fit to data with $m_a = 25$ GeV.




Figure 10.13: Ranked impact of the systematic uncertainties to the signal strength from the S+B fit to data with $m_a = 25 \text{ GeV}$.

Category	Impact on $\hat{\mu}$
Total stat.	33.1%
Total syst.	22.5%
General	2.4%
Lepton reconstruction	0.3%
Jet reconstruction	7.8%
Track reconstruction	1.5%
E_T^{miss}	0.8%
Flavour tagging	11.6%
Signal modelling	8.2%
Z+jets modelling	5.4%
$t\bar{t}$ +jets modelling	4.8%
Reweighting	0.7%
MC stats.	13.0%
Signal MC stats.	1.7%
Norm. factors	3.5%

Table 10.1: Relative impacts over the fitted signal strength $\hat{\mu}$ obtained from the S+B fit to data, grouped by the type of uncertainty. The impact of the statistical uncertainty (without MC stats.) is also included for comparison.

10. Results and conclusions

Sample	SR 2B	SR 1B2b	SR 1B1b1v	SR 4b	SR 3b1v
Signal 25 GeV	6 ± 2	15 ± 6	3 ± 1	9 ± 4	3 ± 1
$Z \geq 1B$	34 ± 11	110 ± 40	50 ± 19	40 ± 18	21 ± 13
$Z \geq 1b$	109 ± 13	560 ± 40	340 ± 30	540 ± 50	360 ± 40
$Z \geq 1c$	1.0 ± 0.9	32 ± 10	20 ± 5	60 ± 30	53 ± 16
$Z + \text{light}$	0.36 ± 0.14	5 ± 4	10 ± 5	4 ± 4	10 ± 20
$t\bar{t} \geq 1B$	10 ± 9	30 ± 20	4 ± 2	20 ± 20	3 ± 3
$t\bar{t} \geq 1b$	5 ± 2	50 ± 20	11 ± 5	110 ± 40	34 ± 14
$t\bar{t} \geq 1c$	5 ± 2	70 ± 20	17 ± 6	80 ± 40	30 ± 13
$t\bar{t} + \text{light}$	25 ± 5	80 ± 20	57 ± 8	25 ± 10	23 ± 6
Other	9 ± 2	32 ± 4	12 ± 2	53 ± 5	19 ± 2
Total MC	203 ± 13	1000 ± 30	530 ± 20	950 ± 30	560 ± 20
Data	202	1001	534	956	550

Table 10.2: Post-fit signal and background yields in the five signal regions for the $m_a = 25$ GeV hypothesis. The total MC and data yields are included in the last two rows.

Sample	SFonZ 1i.2i(T+L).0i	SFonZ 0.2i.0i	SFonZ 1.0.2i	SFonZ 0.1.2i	SFonZ 0.0.4i	SFonZ 0.0.3.1ivorj
Signal 25 GeV	1.5 ± 0.7	1.0 ± 0.6	2.5 ± 1.1	3.8 ± 1.7	1.0 ± 0.5	1.5 ± 0.6
$Z \geq 1B$	90 ± 40	70 ± 40	160 ± 70	200 ± 90	40 ± 20	210 ± 80
$Z \geq 1b$	640 ± 50	690 ± 40	660 ± 70	1880 ± 110	620 ± 60	3900 ± 400
$Z \geq 1c$	5 ± 5	15 ± 5	13 ± 10	130 ± 30	110 ± 40	1600 ± 400
$Z + \text{light}$	0	5 ± 6	17 ± 14	40 ± 30	16 ± 15	500 ± 200
$t\bar{t} \geq 1B$	9 ± 6	6 ± 4	32 ± 19	40 ± 20	12 ± 9	29 ± 16
$t\bar{t} \geq 1b$	9 ± 5	7 ± 2	40 ± 20	90 ± 40	80 ± 40	140 ± 60
$t\bar{t} \geq 1c$	9 ± 2	13 ± 5	45 ± 11	120 ± 40	80 ± 40	350 ± 130
$t\bar{t} + \text{light}$	84 ± 11	110 ± 15	38 ± 9	120 ± 20	23 ± 8	320 ± 80
Other	18 ± 3	20 ± 3	32 ± 3	87 ± 8	60 ± 6	211 ± 16
Total MC	860 ± 30	930 ± 30	1040 ± 30	2720 ± 50	1040 ± 30	7330 ± 100
Data	850	942	1050	2703	1048	7330

Table 10.3: Post-fit signal and background yields in the six Z +jets control regions (SFonZ) for the $m_a = 25$ GeV hypothesis. The total MC and data yields are included in the last two rows.

Sample	DF 1i.2i(T+L).0i	DF 0.2i.0i	DF 1.0.2i	DF 0.1.2i	DF 0.0.4i	DF 0.0.3.1ivorj
Signal 25 GeV	0	0	0	0	0	0
$Z \geq 1B$	1.0 ± 0.7	0.04 ± 0.11	0.4 ± 0.4	0.8 ± 0.5	0.6 ± 0.3	2.5 ± 1.6
$Z \geq 1b$	2.3 ± 0.9	3.2 ± 1.7	2.2 ± 1.1	8.9 ± 3.7	3.5 ± 1.3	43 ± 12
$Z \geq 1c$	0	0	0	0.6 ± 0.4	0	22 ± 12
$Z + \text{light}$	0	0	0	0.1 ± 0.2	0	1.5 ± 1.0
$t\bar{t} \geq 1B$	160 ± 100	90 ± 50	600 ± 300	700 ± 400	300 ± 200	1300 ± 700
$t\bar{t} \geq 1b$	140 ± 60	120 ± 50	600 ± 300	1400 ± 600	1800 ± 700	5000 ± 2000
$t\bar{t} \geq 1c$	190 ± 60	210 ± 60	800 ± 200	2200 ± 600	1600 ± 600	14000 ± 4000
$t\bar{t} + \text{light}$	1200 ± 90	1300 ± 70	720 ± 140	1900 ± 400	450 ± 160	12000 ± 3000
Other	50 ± 7	54 ± 8	86 ± 8	198 ± 18	211 ± 14	1098 ± 70
Total MC	1750 ± 40	1780 ± 40	2870 ± 70	6440 ± 120	4460 ± 140	33400 ± 700
Data	1762	1765	2858	6452	14396	36495

Table 10.4: Post-fit signal and background yields in the six $t\bar{t}$ +jets control regions (DF) for the $m_a = 25$ GeV hypothesis. The total MC and data yields are included in the last two rows.

10.2.3 Expected and observed limits

The expected and observed upper limits on $\text{BR}(H \rightarrow aa \rightarrow 4b)$ resulting from the S+B fit to data are shown in Figure 10.14 (a) as a function of the a -boson mass, which ranges from 12 to 60 GeV. No significant excess over the background expectation is observed for any signal mass hypothesis. This result is compared to the previous Run 2 analysis using a partial Run 2 dataset with 36 fb^{-1} of proton-proton collisions at $\sqrt{s} = 13 \text{ TeV}$, as described in Refs. [172, 173]. The relative contribution from each signal region to the fit can be studied by looking at the expected limits from the Asimov fit using one signal region at a time. Figure 10.14 (b) shows the expected limits calculated using each SR individually, in combination with the Z +jets and $t\bar{t}$ +jets control regions. The 2B SR is the most sensitive region at very low masses, such as 12 and 16 GeV, and plays a key role in improving sensitivity in the boosted regime with respect to previous analyses. The 1B2b and 1B1b1v regions correspond to the mixed regime, where one a -boson is boosted and the other is resolved. They are particularly relevant around 20 GeV. The 4b and 3b1v SRs, corresponding to the fully resolved regime, show the best performance for values of the a -boson mass above 30 GeV.




Figure 10.14: (a) Expected and observed 95% CL upper limits of $\text{BR}(H \rightarrow aa \rightarrow 4b)$ as a function of m_a . The red and blue lines correspond to the two previous analyses using 36 fb^{-1} of proton-proton collisions at $\sqrt{s} = 13 \text{ TeV}$. (b) Expected 95% CL upper limits of $\text{BR}(H \rightarrow aa \rightarrow 4b)$ as a function of m_a for each SR. The dashed line corresponds to the combined expected limit from (a). All limits are calculated for the available signal mass points and smoothed using ROOT for easier visualisation.

The upper limits for $\text{BR}(H \rightarrow aa \rightarrow 4b)$ are compared to other ATLAS searches targeting exotic $H \rightarrow aa$ decays with b -quarks, in particular, the $2b2\tau$ [186] and $2b2\mu$ [182] final states. Figure 10.15 (a) compares the expected and observed limits at 95% confidence level, all of them obtained using the full Run 2 dataset. This plot shows two analyses that use the DeXTer B -tagger in the low- m_a regime ($4b$ and $2b2\tau$) and one that does not ($2b2\mu$). It is observed that the DeXTer tagger for boosted $a \rightarrow b\bar{b}$ decays allows to extend the range of a -boson mass hypotheses down to 12 GeV, and prevents the drop in sensitivity that can be observed below $m_a \sim 20$ GeV in the $2b2\mu$ search. Figure 10.15 (b), on the other hand, compares the observed limits from ATLAS with the observed limits from equivalent CMS analyses [174, 185]. Similarly, the CMS searches show less sensitivity at low values of m_a , due to the lack of a dedicated tagger for low- p_T collimated decays. The observed upper limits are similar in the $2b2\mu$ final state, where DeXTer is not used.




Figure 10.15: (a) Expected and observed 95% CL upper limits of $\text{BR}(H \rightarrow aa \rightarrow XYY)$ as a function of m_a using the full Run 2 dataset recorded by the ATLAS experiment. (b) Observed 95% CL upper limits of $\text{BR}(H \rightarrow aa \rightarrow XYY)$ as a function of m_a using the full Run 2 dataset recorded by the ATLAS and CMS experiments. Only analyses with at least two b -quarks in the final state are included.

10.3 Conclusions

A search for exotic decays of the Higgs boson into a pair of scalar particles, $H \rightarrow aa$, where each a -boson decays into two b -quarks has been performed. This analysis is part of a more general study including $H \rightarrow a_1 a_2$ decays, where a_1 and a_2 are two a -bosons with different mass, also decaying to b -quarks. The analysis is performed using 140 fb^{-1} of proton-proton collision data at a centre-of-mass energy of 13 TeV recorded by the ATLAS detector at the LHC between 2015 and 2018. The search makes use of heavy-flavour-tagging techniques to target collimated $a \rightarrow b\bar{b}$ decays, as well as soft secondary vertices outside of jets to improve the sensitivity at large values of m_a . No significant excess above the SM background expectation is observed, and upper limits at 95% confidence level are set on $\text{BR}(H \rightarrow aa \rightarrow 4b)$ between 5% and 25% for a -boson mass values in the range $12 \text{ GeV} \leq m_a \leq 60 \text{ GeV}$. These results contribute to the broad program of searches for $H \rightarrow aa$ decays and can be used to set constraints on a variety of BSM scenarios featuring exotic Higgs decays.

Part V

Search for a new pseudoscalar decaying into a pair of bottom and anti-bottom quarks in top-associated production using proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector

Chapter 11

Analysis overview

ALPs are present in several extensions of the SM where they mix with fields in an extended Higgs sector, inheriting the Yukawa couplings to fermions. Under this hypothesis, a new pseudoscalar particle would couple strongly to the top quark. ALP searches in association with a $t\bar{t}$ pair complement the previously described $H \rightarrow aa$ studies and allow to probe a -boson candidates with $m_a > m_H/2$.

11.1 Introduction

This search is motivated by SM extensions including a light pseudoscalar particle that mixes with the Higgs sector. The analysis is based on the 2HDM+ a at the decoupling limit, where the light CP-even state can be identified with the SM Higgs boson, here denoted as H . The $t\bar{t}a$, $a \rightarrow b\bar{b}$ process can be described in terms of a simplified model, where couplings of the a -boson with the bottom and top quarks appear in the SM lagrangian as:

$$\mathcal{L}_{\text{BSM}} \ni -\frac{g_t y_t}{\sqrt{2}} a \bar{t}(i\gamma^5)t - \frac{g_b y_b}{\sqrt{2}} a \bar{b}(i\gamma^5)b, \quad (11.1)$$

where g_t , g_b are multiplicative factors that parametrise the deviation from the SM Yukawa couplings, $y_t = m_t/v$ and $y_b = m_b/v$ for the top and bottom quark, respectively. This process is studied for dileptonic $t\bar{t}$ decays of the type ee , $\mu\mu$ or $e\mu$. These two leptons provide the trigger signal and contribute to suppress the all-hadronic QCD background. Figure 11.1 shows two example Feynman diagrams for $t\bar{t}$ -associated production, $t\bar{t}a$, and tW -associated production, $tW a$. The sub-leading $tW a$ process contributes to around 10% of the total cross section¹. A comparison between the $t\bar{t}a$ and $tW a$ cross sections from the MC samples used in the analysis can be found in Figure 11.2. They are both linearly dependent on the a -boson mass.




Figure 11.1: Example Feynman diagrams for a -boson production in association with $t\bar{t}$ and with tW in the dilepton channel with $a \rightarrow b\bar{b}$.

¹Throughout this thesis, the $t\bar{t}a + tW a$ signal will be referred to as $t\bar{t}a$ for simplicity.




Figure 11.2: Cross sections for the $t\bar{t}a$ and tWa signal processes assuming $g_t = 0.5$ and $\text{BR}(a \rightarrow b\bar{b}) = 1$.

The kinematics of the $t\bar{t}$ system are not very sensitive to the a -boson mass, as it can be seen in Figure 11.3, showing the top quark p_T and the angular distance between the two leptons from the $t\bar{t}$ system. However, differences can be observed in the kinematics of the $a \rightarrow b\bar{b}$ decay, shown in Figures 11.4 and 11.5. Higher values of m_a lead to more energetic a -bosons with more energetic decay products. Nevertheless, boosted $a \rightarrow b\bar{b}$ decays are more prominent at low m_a , especially below 30 GeV, where the angular separation between b -quarks, ΔR_{bb} , becomes smaller than 0.4, which is the standard ATLAS jet cone width.




Figure 11.3: (a) Top quark p_T and (b) angular separation between the two leptons from the $t\bar{t}$ system under different a -boson mass hypotheses. The kinematics of the $t\bar{t}$ system are not sensitive to the value of m_a .




Figure 11.4: (a) a -boson p_T and (b) angular separation between the two b -quarks from the $a \rightarrow b\bar{b}$ under different a -boson mass hypotheses. These two distributions show how higher masses tend to carry more energy, but their decay products have a bigger angular separation.




Figure 11.5: (a) Leading and (b) sub-leading b -quark p_T from $t\bar{t}a$, $a \rightarrow b\bar{b}$ production using different mass hypotheses. For higher masses, the a -boson is more energetic, and tends to produce particles with larger p_T .

The signal process studied in this analysis is characterised by a large jet multiplicity in the final state. However, preliminary tests showed that it does not require the extended low- p_T jet calibration that is used in the $H \rightarrow aa \rightarrow 4b$ analysis or the soft secondary vertex tagging for very low energy b -objects. It benefits, though, from the DeXTer B -tagger, particularly at $m_a \lesssim 30$ GeV. Because jets are required to have $p_T \geq 20$ GeV, the standard ATLAS jet calibration is used, allowing to include different DL1r b -tagging WPs in the analysis.

This analysis is a search for top-associated production of a light pseudoscalar particle decaying to b -quarks, $t\bar{t}a$, $a \rightarrow b\bar{b}$, over the mass range $12 \leq m_a \leq 100$ GeV using the

full Run 2 dataset of proton-proton collisions at $\sqrt{s} = 13$ TeV recorded with the ATLAS detector, corresponding to a total integrated luminosity of 140 fb^{-1} . This is the first search of $t\bar{t}a$, $a \rightarrow b\bar{b}$ performed at the LHC, and it is complementary to other ATLAS studies of $t\bar{t}a$ production currently ongoing. Already published results from ATLAS and CMS include the $a \rightarrow \mu\mu$ [233, 234] and $a \rightarrow ee$ [234] decay channels.

11.2 Data and simulated samples

This section describes the data and Monte Carlo samples used for this analysis. The datasets are divided in three subsets, corresponding to the three data-taking periods of Run 2, which are 2015+2016, 2017 and 2018. The signal and background samples are simulated using different matrix element generators, interfaced with various parton shower algorithms. A reweighting method is applied to the simulated samples such that they match the pileup conditions in data. For the detector simulation two different approaches are used: the full ATLAS detector simulation (FS), based on the standard GEANT4 simulation, and the fast detector simulation (AFII) which uses a parametrisation of the calorimeter response. Both MC and data events are processed with the same reconstruction and analysis software.

11.2.1 Data samples

This analysis uses the same data samples as the $H \rightarrow aa \rightarrow 4b$ analysis. Selected events are recorded using unprescaled triggers, as detailed in Section 11.3. Standard data taking quality requirements are applied, such as stable LHC beams and fully operational ATLAS detector conditions. These requirements define the GRLs, which are the same as those in Section 9.2. The total integrated luminosity for the full Run 2 dataset is 140.1 fb^{-1} , usually written as 140 fb^{-1} , with an uncertainty of 0.83%.

11.2.2 Signal Monte Carlo samples

The signal samples consist mainly of $t\bar{t}$ -associated production of a light pseudoscalar particle, $t\bar{t}a$, with a mass smaller than the Higgs boson mass, which is assumed to be 125 GeV. The associated production of a light pseudoscalar with a single top quark and a W boson, $tW a$, is also included, although its contribution is subdominant. The mass hypotheses considered in the MC simulation include $m_a \in [12, 16, 20, 25, 30, 40, 50, 60, 80, 100] \text{ GeV}$. The analysis targets the dileptonic decay channel of the $t\bar{t}$ pair, together with the a -boson decaying into a $b\bar{b}$ pair, which would be the most favoured decay channel in the case of Yukawa-like couplings. The $t\bar{t}a$ signal samples are simulated with the MADGRAPH5_aMC@NLO v2.3.3 generator at NLO with the NNPDF3.0 NLO PDF set. The $tW a$ signal samples are simulated with the MADGRAPH5_aMC@NLO v2.3.3 generator at LO, also with the NNPDF3.0 NLO PDF set. In both cases, events are interfaced with PYTHIA 8.230 using the A14 tune. Decays of bottom and charm hadrons are performed by EVTGEN 1.7.0. The generated samples are processed through a simulation of the detector geometry and response using AFII. All samples are processed through the same reconstruction software as the data. Simulated

events are corrected so that the object identification efficiencies, energy scales and energy resolutions match those determined from data control samples.

11.2.3 Background Monte Carlo samples

The final states of the $H \rightarrow aa \rightarrow 4b$ and the $t\bar{t}a, a \rightarrow b\bar{b}$ processes are very similar, therefore the background composition is expected to be roughly the same in both analyses. Because the $t\bar{t}a, a \rightarrow b\bar{b}$ analysis was developed later in time, some of the MC background samples have been updated with respect to the ones used in the $H \rightarrow aa \rightarrow 4b$ analysis. This section summarises the full list of background MC used in this search, with special emphasis on the newly introduced samples.

$t\bar{t} + \text{jets}$

The $t\bar{t} + \text{light}$ and $t\bar{t} + \geq 1c$ categories from the $t\bar{t} + \text{jets}$ sample are modelled using the POWHEGBox-v2 generator at NLO in QCD with the 5FS NNPDF3.0 NLO PDF set, combined with PYTHIA 8.230 for the parton shower and hadronisation. The $t\bar{t} + \text{jets}$ sample is normalised to the cross section prediction at NNLO in QCD including the resummation of NNLL soft-gluon terms calculated using TOP++ 2.0. In order to enhance the statistics of the HF additional radiation, the enriched BBFilt, BFiltBBVeto and CFiltBVeto samples are included.

$t\bar{t} + b\bar{b}$

The production of $t\bar{t} + \geq 1b$ events is modelled using a dedicated MC simulation for top quark pair production in association with a b -quark pair, $t\bar{t} + b\bar{b}$. A new version with respect to the $H \rightarrow aa \rightarrow 4b$ analysis is used. The sample is produced using the POWHEGBox-Res framework at NLO with massive b -quarks and the 4FS NNPDF3.1 NNLO PDF set. The PowhegBox internal parameter h_{damp} is set to $H_T/2$, where H_T corresponds to the scalar sum of all transverse momenta over leptons and jets. The renormalisation scale is set to half the geometric average of the transverse energy of t - and b -quarks, $\mu_R = \frac{1}{2} \sqrt[4]{\prod_i E_{T,i}}$ with $i = t, \bar{t}, b, \bar{b}$. The factorisation scale is related to the sum of the transverse mass of the outgoing partons in the matrix element calculation and is defined as $\mu_F = \frac{1}{2} (\sum_i E_{T,i})$, with $i = t, \bar{t}, b, \bar{b}, j$. The POWHEGBox internal parameter h_{bzd} , which controls the fraction of events produced in singular and finite regions, is set to 5 instead of the previous value of 2. This choice of scales follows the latest recommendations from the LHC Higgs Working Group, described in Ref. [202]. The POWHEGBox ME calculation is matched to PYTHIA 8.244 with the A14 tuned parameters for the parton shower and hadronisation.

Single top

Single-top production includes the t -channel, the s -channel, and tW production. It is modelled using POWHEGBox-v2 at NLO in QCD with the NNPDF3.0 NLO PDF set, and showered with PYTHIA 8.230. All single-top samples are normalised to the cross section prediction calculated at NLO+NNLL in QCD.

$t\bar{t}H, t\bar{t}Z$ and $t\bar{t}W$

The production of $t\bar{t}H$ events is modelled using the PowhegBox-v2 generator at NLO

in QCD with the NNPDF3.0 NLO PDF set, interfaced with PYTHIA 8.230 for the PS. The cross section is calculated at NLO in the QCD and EW couplings.

$t\bar{t}Z$ and $t\bar{t}W$ events are produced using the MADGRAPH5_aMC@NLO v2.3.3 generator at NLO with the NNPDF3.0 NLO PDF set. Events are interfaced with Pythia 8.210 using the A14 tune and the NNPDF2.3 LO PDF set, interfaced with PYTHIA 8.210. Both cross sections are calculated at NLO in the QCD and EW couplings.

Rare processes

The production of tZq events is performed using MADGRAPH5_aMC@NLO v2.3.3 at LO in QCD in the 4FS with the CTEQ6L1 PDF set. Events are showered with PYTHIA 8.212. The production of tWZ events is modelled using the MADGRAPH5_aMC@NLO v2.2.3 generator at NLO with the NNPDF3.0 NLO PDF set, interfaced with PYTHIA 8.212.

$Z/W + \text{jets}$

The production of $Z/W + \text{jets}$ is simulated with the SHERPA 2.2.11 generator using NLO matrix elements for up to two partons, and LO matrix elements for up to four partons, calculated with the COMIX and OPENLOOPS libraries. They are matched with the SHERPA PS using the MEPS@NLO prescription with the set of tuned parameters developed by the Sherpa authors. The NNPDF3.1 NNLO set of PDFs is used and the cross section is normalised to the NNLO prediction. Dedicated low-mass $Z + \text{jets}$ MC samples with $10 < m_{\ell\ell} < 40$ GeV are included for accurate background modelling at low energies. These are not used in the $H \rightarrow aa \rightarrow 4b$ analysis due to the $m_{\ell\ell} > 50$ GeV cut that is applied in the preselection.

Diboson

Diboson final states are simulated with the SHERPA 2.2.1 or 2.2.2 generator depending on whether the bosons decay to hadrons or leptons. Fully leptonic and semileptonic final states are generated using matrix elements at NLO accuracy in QCD for up to one additional parton emission and at LO accuracy for up to three additional parton emissions. Samples for the loop-induced processes $gg \rightarrow ZZ/ZW/WW$ are generated using LO-accurate matrix elements for up to one additional parton emission for both the cases of fully leptonic and semileptonic final states. The ME calculations are matched and merged with the SHERPA parton shower based on the Catani–Seymour dipole factorisation using the MEPS@NLO prescription. Virtual QCD corrections are provided by the OPENLOOPS library. The NNPDF3.0 NNLO set of PDFs is used, along with a dedicated set of tuned PS parameters developed by the SHERPA authors.

Summary

The full list of signal and background Monte Carlo samples is summarised in Table 11.1.

Process	ME generator	PDF set	PS generator	Norm.	Sim.
$t\bar{t}a, a \rightarrow b\bar{b}$	MG_aMC@NLO v2.3.3	NNPDF3.0 NLO	PYTHIA 8.230	–	AFII
$tW a, a \rightarrow b\bar{b}$	MG_aMC@NLO v2.3.3	NNPDF3.0 NLO	PYTHIA 8.230	–	AFII
$t\bar{t} + \text{jets}$	POWHEGBOX-v2	NNPDF3.0 NLO	PYTHIA 8.230	$(\text{NLO}+\text{NNLL})_{\text{QCD}}$	FS
$t\bar{t} + b\bar{b}$	POWHEGBOX-Res	NNPDF3.1 NNLO	PYTHIA 8.244	–	FS
Single-top	POWHEGBOX-v2	NNPDF3.0 NLO	PYTHIA 8.230	$(\text{NLO}+\text{NNLL})_{\text{QCD}}$	FS
$t\bar{t}H$	POWHEGBOX-v2	NNPDF3.0 NLO	PYTHIA 8.230	$\text{NLO}_{\text{QCD}+\text{EW}}$	FS
$t\bar{t}Z$	MG5_aMC@NLO v2.3.3	NNPDF3.0 NLO	PYTHIA 8.210	$\text{NLO}_{\text{QCD}+\text{EW}}$	FS
$t\bar{t}W$	MG5_aMC@NLO v2.3.3	NNPDF3.0 NLO	PYTHIA 8.210	$\text{NLO}_{\text{QCD}+\text{EW}}$	FS
Rare processes	MG5_aMC@NLO v2.3.3	CTEQ6L1 NNPDF3.0 NLO	PYTHIA 8.212	–	FS
$Z/W + \text{jets}$	SHERPA 2.2.11	NNPDF3.1 NNLO	SHERPA	NNLO_{QCD}	FS
Diboson	SHERPA 2.2.1 SHERPA 2.2.2	NNPDF3.1 NNLO	SHERPA	–	FS

Table 11.1: Summary of signal and background samples used in the $t\bar{t}a, a \rightarrow b\bar{b}$ analysis, including the ME generator, the PDF set, the PS generator, the cross section normalisation and the type of detector simulation. Note that MADGRAPH is abbreviated to MG.

11.3 Event selection

11.3.1 Trigger

Only events recorded with a single-electron or single-muon trigger under stable beam conditions and for which all detector subsystems were operational are considered. Single-lepton triggers with different p_T thresholds are combined in a logical OR in order to increase the overall efficiency. The trigger chains used during the four years of data-taking are the same as for the $H \rightarrow aa \rightarrow 4b$ analysis. They are listed in Section 9.3.

The triggers with the lower p_T threshold include isolation requirements on the candidate lepton, resulting in inefficiencies at high p_T that are recovered by the triggers with higher p_T threshold.

11.3.2 Preselection

Events are required to have exactly two leptons ($ee, \mu\mu$ or $e\mu$) with opposite charge, satisfying the criteria defined in Section 6.2. Since single-lepton triggers are used, at least one of the two reconstructed leptons is required to have $p_T > 27$ GeV and match a lepton with the same flavour reconstructed by the trigger algorithm within $\Delta R < 0.15$. In the ee and $\mu\mu$ channels, the dilepton invariant mass must be above 15 GeV. This cut

is used to remove low-mass dilepton resonances such as those originating from J/ψ and Υ mesons. The dilepton invariant mass is also required to be outside of the Z -mass window, such that $|m_{\ell\ell} - m_Z| > 8$ GeV, in order to suppress the Z +jets background. Additionally, events are required to have at least three jets, either Ak4 or Ak8, and at least one b -tagged jet with DL1r 85% WP. This requirement helps to increase signal purity and reduce the size of the event files. Tables 11.2 and 11.3 show the event yields in the preselection region for background samples and signal samples, respectively.

Sample	Events	% over total
$t\bar{t}$ +light	1025033	66%
$t\bar{t}+\geq 1c$	102006	7%
$t\bar{t}+\geq 1b$	46669	3%
tW	52304	3%
$t\bar{t}H$	1621	<1%
$t\bar{t}Z$	2518	<1%
$t\bar{t}W$	2061	<1%
tq, tZ, tWZ	557	<1%
Z +jets	308425	20%
W +jets	1354	<1%
WW, ZZ, WZ	10055	1%
Total	1552603	100%

Table 11.2: Background composition in the preselection region, corresponding to 140 fb^{-1} .

m_a [GeV]	Events
12	662
16	642
20	631
30	595
40	572
60	500
80	426
100	364

Table 11.3: Number of signal events in the preselection region, corresponding to 140 fb^{-1} .

11.4 Background modelling

As shown in Table 11.2, the $t\bar{t}$ +jets background is the main background process in the search for a $t\bar{t}a, a \rightarrow b\bar{b}$ signal. Possible mismodellings of the $t\bar{t}$ +jets MC simulation are studied in a data-driven way, by comparing it to data in a control region with negligible signal content. This section describes the set of corrections that are applied to the $t\bar{t}$ +jets background in order to improve the modelling of the MC simulation. These corrections also apply to the tW sample, which constitutes a non-negligible contribution to the total background and interferes at NLO with $t\bar{t}$ +jets.

11.4.1 $t\bar{t}$ +jets reweighting

Reweighting region

The region of choice for calculating the corrections satisfies the preselection cuts described in Section 11.3, in addition to the requirement of at least two b -tagged jets using the DL1r 70% WP, and no B -tagged jets using the DeXTer 40% WP. Additionally, only the $e\mu$ region is considered, in order to suppress the Z +jets contribution. The background composition in this region is shown in Table 11.4. Signal contamination is smaller than 1% for all mass hypotheses. The $t\bar{t}$ +jets and tW backgrounds constitute almost 100% of the events in this reweighting region.

Sample	Events	% over total
$t\bar{t}$ +light	197633	82%
$t\bar{t}+\geq 1c$	20326	8%
$t\bar{t}+\geq 1b$	14223	6%
tW	6706	3%
$t\bar{t}H$	443	<1%
$t\bar{t}Z$	427	<1%
$t\bar{t}W$	375	<1%
tq, tZ, tWZ	30	<1%
Z +jets	229	<1%
W +jets	14	<1%
WW, ZZ, WZ	69	<1%
Total	240475	100%

Table 11.4: Background composition of the $t\bar{t}$ +jets control region. The $t\bar{t}$ +jets and tW backgrounds constitute around 100% of the events.

Heavy-flavour reweighting

It is known from previous analyses that the rate of $t\bar{t}$ +HF events is underestimated in MC simulation. Due to the high b -object multiplicity of the $t\bar{t}a, a \rightarrow b\bar{b}$ signal, the $t\bar{t}$ +HF categories represent a large fraction of the $t\bar{t}$ +jets background in the signal regions, and therefore they need to be corrected. In order to have a more accurate flavour composition, a reweighting procedure is applied based on the sumPCBTag distribution. This variable is defined as the sum of the pseudo-continuous b -tagging score of all the jets in the event, according to the values in Table 11.5. Figure 11.6 (a) shows the sumPCBTag distribution in the $t\bar{t}$ +jets reweighting region. It can be seen that $t\bar{t}$ +light is dominant at low values of sumPCBTag, while $t\bar{t}$ +HF populates the tail of the distribution. The disagreement between data and MC is more pronounced in the tail of the distribution, dominated by $t\bar{t}$ +HF. The correction procedure consists in deriving three normalisation factors, one for each component, using a likelihood fit of the sumPCBTag MC distribution with respect to data. In this fit, DL1r b -tagging systematic uncertainties are included. Figure 11.6 (b) shows the improved agreement after the fit. The three normalisation values with their corresponding uncertainties are shown in Table 11.6. The correction factor corresponding to the $t\bar{t}+\geq 1b$ category

11. Analysis overview

is smaller than in the previously described $H \rightarrow aa \rightarrow b\bar{b}$ analysis, thanks to the use of the new version of the $t\bar{t}+b\bar{b}$ MC sample.

b -tagged?	DL1r WP	b -tagging score
No	None	1
Yes	85%	2
Yes	77%	3
Yes	70%	4
Yes	60%	5

Table 11.5: correspondence between the DL1r WP and the b -tagging score for each jet. Low values of the DL1r WP correspond to tight b -tagging WPs, and viceversa.

HF category	Norm. factor
$t\bar{t}+\text{light}, tW$	0.91 ± 0.03
$t\bar{t}+\geq 1c$	1.58 ± 0.14
$t\bar{t}+\geq 1b$	1.13 ± 0.07

Table 11.6: Normalisation factors for the three $t\bar{t}$ +jets categories resulting from the likelihood fit performed using the sumPCBTag distribution.




Figure 11.6: Data versus MC comparison of the sum of the PC b -tagging score of all the jets per event (a) before and (b) after applying the $t\bar{t}$ +jets normalisation correction. The error band includes the statistical error and the DL1r uncertainties.

As a consequence of this correction, not only the sumPCBTag distribution is corrected, but also the N_{jets} distribution, which is known to be mismodelled in the simulation as well. Figure 11.7 shows the corresponding histogram, before and after applying the normalisation factors from Table 11.6.




Figure 11.7: Data versus MC comparison of the N_{jets} distribution (a) before and (b) after applying the tt +jets normalisation correction. Systematic uncertainties are not included.

Kinematic reweighting

After applying the $t\bar{t}$ +HF correction, the shape of the p_T distributions of the jets and leptons is mismodelled with respect to data for the $t\bar{t}$ +jets 5FS MC sample, that is, the $t\bar{t}$ +light and $t\bar{t}+\geq 1c$ components. The aim of the so-called kinematic reweighting is to improve the general modelling of the p_T distributions by correcting the binned distributions of relevant kinematic variables using the following reweighting factor:

$$R(x) = \frac{\text{Data}(x) - \text{MC}_{\text{other}}(x)}{\text{MC}_{t\bar{t}(5\text{FS}), tW}(x)}, \quad (11.2)$$

where x is the variable to be corrected. The mismodelling described in this section is assumed to be independent of the flavour of the extra radiation, and it is applied equally to the $t\bar{t}$ +light, $t\bar{t}+\geq 1c$ and tW backgrounds.

The variable of choice for the p_T correction is the event hardness or H_T , which is defined as the scalar sum of the p_T of all the jets and leptons in the event. This variable can not be reweighted directly, due to its strong correlation with the number of jets, shown in Figure 11.8 (a). In order to reduce this dependency, a new variable is defined:

$$H_T^{\text{red}}(n) = H_T(n) - (n - n_{\min})\Delta H_T(n), \quad (11.3)$$

where n is the number of jets, $n_{\min}$ is the minimum number of jets per event allowed by the preselection (three, in this particular case) and $\Delta H_T(n)$ is the average offset in H_T caused by the addition of each extra jet to the event:

$$\Delta H_T(n) = \left\langle \frac{\text{Mo}[H_T(n)] - \text{Mo}[H_T(n_{\min})]}{n - n_{\min}} \right\rangle_{n=4,5,\dots,\geq 10}, \quad (11.4)$$

where $H_T(n)$ is the H_T distribution for the subset of events with n jets and "Mo" is the mode of the distribution. Note that the number 3 that appears in both equations is the minimum number of jets required by the preselection. The H_T^{red} distribution is illustrated in Figure 11.8 (b).




Figure 11.8: Data-only (a) H_T and (b) H_T^{red} normalised distributions as a function of N_{jets} in the $t\bar{t}$ +jets reweighting region. The histograms show how each additional jet in the event causes an offset in H_T , while this effect is mitigated in H_T^{red} .

For this correction, equation 11.2 is applied over a binned H_T^{red} distribution, and the results are fitted using a continuous hyperbolic function of the type:

$$f(x) = c_0 + \frac{c_1}{x^{c_2}}, \quad (11.5)$$

as shown in Figure 11.9. Values of $H_T^{\text{red}} \leq 50$ GeV are reweighted using the value of the function at 50 GeV, in order to avoid very large values when $H_T^{\text{red}} \rightarrow 0$. The results of the hyperbolic fit are included in Table 11.7. Figure 11.10 shows the N_{jets} distribution before and after applying the correction. No significant changes are observed. Figures 11.11 to 11.13 show several p_T distributions for jets and leptons, including E_T^{miss} , before and after applying the correction. There is a general improvement in the shape mismodelling in all cases.

Parameter	Value
c_0	0.69 ± 0.06
c_1	7 ± 3
c_2	0.55 ± 0.11

Table 11.7: Fitted parameters to the hyperbolic function for the H_T correction. The table shows the numerical values from all reweighting parameters, with their corresponding uncorrelated uncertainties.




Figure 11.9: $t\bar{t}$ +jets and tW correction function as a function of H_T^{red} .




Figure 11.10: Data versus MC comparison of the N_{jets} distribution (a) before and (b) after correcting the H_T distribution using H_T^{red} . Systematic uncertainties are not included.




Figure 11.11: Leading (1) and sub-leading (2) jet p_T distributions (a) before and (b) after applying the H_T correction using H_T^{red} . Systematic uncertainties are not included.




Figure 11.12: Leading (1) and sub-leading (2) lepton p_T distributions (a) before and (b) after applying the H_T correction using H_T^{red} . Systematic uncertainties are not included.




Figure 11.13: Missing transverse energy distribution (a) before and (b) after applying the H_T correction using H_T^{red} . Systematic uncertainties are not included.

11.5 Analysis strategy

This section describes the various machine learning techniques employed to identify the $t\bar{t}a$, $a \rightarrow b\bar{b}$ final state and to create a set of signal and control regions to be used in the final fit. This is done in two steps. First, a set of BDTs is used to find the best jet pairing from the $t \rightarrow j\ell$ and $a \rightarrow jj$ decays. This allows to define variables such as invariant masses, angular distances or jet ID². Then, a parametrised NN is trained for signal versus background discrimination in four separate regions targetting different topologies of the final state. The resulting NN score distribution is the variable that will be used in the fit to calculate the expected and observed limits.

11.5.1 BDT for event reconstruction

Reconstruction BDT for $t \rightarrow j\ell$ decay

A BDT is trained to identify the jets originating from the decay of the top and anti-top quarks and pair them with their corresponding leptons. The design is based on the DNN for event reconstruction used in the template method measurement of the top quark mass in the dilepton decay channel [235], plus various improvements in terms of flexibility and efficiency: it allows non b -tagged jets to be labelled as the top or anti-top b -quarks, it examines up to the 5th leading jet of the event and it targets the reconstruction of the top/anti-top lepton-jet pair independently. No attempt of reconstructing the two neutrinos from the top/anti-top quark decay is performed in this analysis.

²Jet ID is the jet index, ranging from 1 to N_{jets} , from the list of p_T -ordered Ak4 jets.

The jets corresponding to the top/anti-top b -quark in the MC samples are identified using truth information. Given that the fraction of reconstructed leptons not originating from the decay of a top or anti-top quark is below 1%, all leptons are assumed to be correctly reconstructed and they are assigned to the top or anti-top quark based on their electric charge. The BDT is trained with a subset of events from the $t\bar{t}a$ ³ and 5FS $t\bar{t}$ +jets MC samples, including the HF-filtered $t\bar{t}$ +jets samples⁴, as shown in Table 11.8. Note that the events listed in the table are not normalised to the luminosity and cross section of each MC sample; they correspond to the real number of generated MC events. During the training, the correct pair of positive lepton and jet matched to the b -quark from the top quark decay is labelled as signal, while all other positive-lepton/jet pairs are labelled as background. For those events in which the b -quark from the top quark decay is not matched to any reconstructed jet, all possible positive-lepton/jet pairs are used as background. The same procedure is followed with the negative lepton and the b -quark from the anti-top quark.

Sample	Events
$t\bar{t}a$ 12 GeV	30000
$t\bar{t}a$ 16 GeV	30000
$t\bar{t}a$ 20 GeV	30000
$t\bar{t}a$ 30 GeV	30000
$t\bar{t}a$ 40 GeV	30000
$t\bar{t}a$ 60 GeV	30000
$t\bar{t}a$ 80 GeV	30000
$t\bar{t}a$ 100 GeV	30000
$t\bar{t}$ +jets	30000
$t\bar{t}$ +jets BBFilt	30000
$t\bar{t}$ +jets BFiltBBVeto	30000
$t\bar{t}$ +jets CFiltBVeto	30000

Table 11.8: Number of $t\bar{t}a$ and $t\bar{t}$ +jets MC events used in the training of the $t \rightarrow j\ell$ reconstruction BDTs.

The input variables to the BDT are summarised in Table 11.9. They include jet and b -jet multiplicities, as well as low-level variables from the leptons and Ak4 jets in the event. The BDT for top (anti-top) reconstruction uses information from both the test jet-lepton pair, which is the one associated to the top (anti-top) and the auxiliary jet-lepton pair, which is the one coming from the anti-top (top). This procedure improves the performance of the algorithm.

The training of the BDT is performed with the ROOT TMVA package [236] and the AdaBoost algorithm. It consists of 500 trees, with a maximum depth of four layers. Before the training, the dataset is split into two, using one half for training and the

³The $t\bar{t}a$ samples with $m_a = 25$ and 50 GeV and the $tW a$ samples were not used in the training, as it was performed prior to their inclusion in the analysis. No significant impact on the BDT performance was observed.

⁴No significant differences were observed when comparing the BDT outputs using the 5FS $t\bar{t}$ +jets sample and the $t\bar{t}+b\bar{b}$ 4FS sample for the $t\bar{t} \rightarrow \geq 1b$ component.

Object	Variables
Full event	$N_{\text{jets}}, N_{b\text{-jets}}$ (85% WP)
Test/aux. $j\ell$ pair	$m, p_T, \eta, \Delta R$
Test/aux. jet	$p_T, \eta, b\text{-tagging score, jet ID}$
Test/aux. lepton	p_T, η
$t\bar{t}$ pair	$m, p_T, \eta, \Delta R, \Delta\phi$
jj pair	ΔR

Table 11.9: Input variables to the $t \rightarrow j\ell$ reconstruction BDTs. Kinematic variables of the $t\bar{t}$ pair are computed using $j\ell$ and $j\bar{\ell}$.

other half for testing purposes, in order to prevent overtraining. A comparison of the BDT score evaluated using the training dataset and the validation dataset for signal (correct pairings) and background (incorrect pairings) is shown in Figure 11.14.




Figure 11.14: Top BDT response for signal and background, using the training and the test dataset, trained with (a) even and (b) odd events.

Figure 11.15 shows examples of relevant variables extracted from the $t \rightarrow j\bar{\ell}$ BDT, for both the background and the signal, including the $m_a = 12, 30, 60$ and 80 GeV hypotheses. The BDT score in Figure 11.15 (a) shows little separation between signal and background, illustrating that the kinematics of the $t\bar{t}$ system in the $t\bar{t}$ +jets background are not that different from those of the signal. Figures 11.15 (b) to (d) show the jet-lepton p_T and their separation in $\Delta\eta$ and $\Delta\phi$. In general, the jet-lepton system with the highest BDT score tends to have a smaller angular separation and to be more energetic for the signal samples. It is also observed that the behaviour of the $t\bar{t}$ pair from $t\bar{t}a$ events does not vary significantly as a function of the a -boson mass.




Figure 11.15: Normalised MC distributions for (a) the highest hypothesis score from the $t \rightarrow j\bar{\ell}$ BDT, (b) the lepton-jet p_T , (c) the difference in pseudorapidity and (d) the difference in azimuthal angle. Signal distributions for $m_a = 12, 30, 60$ and 80 GeV are overlayed for comparison.

Reconstruction BDT for $a \rightarrow jj$ decay

A second BDT is designed to identify the two jets originating from the a -boson decay. It targets the resolved decay and does not consider the case where the two b -quarks from $a \rightarrow b\bar{b}$ are reconstructed as one Ak8 jet. The $a \rightarrow jj$ reconstruction BDT is very similar to the $t \rightarrow j\bar{\ell}$ reconstruction BDT. Truth information is used to identify the two b -quarks from a -boson decay. For the training, the five leading jets in the event are considered. The correct pair formed by the two jets from the a -boson decay is used as signal, while all other pairs are used as background. The BDT is trained with a

11. Analysis overview

subset of $t\bar{t}a$ ⁵ and $t\bar{t}$ +jets events, as shown in Table 11.10, using the same architecture as the top/anti-top BDT, with the only difference that all jet permutations from the $t\bar{t}$ +jets samples are labelled as background.

Sample	Events
$t\bar{t}a$ 12 GeV	106067
$t\bar{t}a$ 16 GeV	103016
$t\bar{t}a$ 20 GeV	107246
$t\bar{t}a$ 30 GeV	110104
$t\bar{t}a$ 40 GeV	112470
$t\bar{t}a$ 60 GeV	114470
$t\bar{t}a$ 80 GeV	115680
$t\bar{t}a$ 100 GeV	116434
$t\bar{t}$ +jets	120000
$t\bar{t}$ +jets BBFilt	120000
$t\bar{t}$ +jets BFiltBBVeto	120000
$t\bar{t}$ +jets CFiltBVeto	120000

Table 11.10: Number of $t\bar{t}a$ and $t\bar{t}$ +jets MC events used in the training of the $a \rightarrow jj$ reconstruction BDT.

Inputs to the BDT include jet and b -jet multiplicities, sumPCBTag and low-level kinematic variables from the possible jet permutations. The full list of variables is shown in Table 11.11. Figure 11.16 shows the comparison between the BDT score obtained using the training dataset and the test or validation dataset for signal (correct $a \rightarrow jj$ permutations from $t\bar{t}a$ events) and background (incorrect $a \rightarrow jj$ permutations from $t\bar{t}a$ events and $t\bar{t}$ +jets events). There are no signs of overtraining.

Object	Variables
Full event	$N_{\text{jets}}, N_{b\text{-jets}}$ (85% WP), sumPCBTag
jj pair	$m, p_T, \eta, E, \phi, \Delta R$
Test jet 1	$p_T, \eta, b\text{-tagging score, jet ID}$
Test jet 2	$p_T, \eta, b\text{-tagging score, jet ID}$

Table 11.11: Input variables to the $a \rightarrow jj$ reconstruction BDT.

⁵The $t\bar{t}a$ samples with $m_a = 25$ and 50 GeV and the $tW a$ samples were not used in the training, as it was performed prior to their inclusion in the analysis. No significant impact on the BDT performance was observed.




Figure 11.16: $a \rightarrow jj$ BDT response for signal and background, using the training and the test dataset, trained with (a) even and (b) odd events.

Figure 11.17 shows examples of relevant variables extracted from the $a \rightarrow jj$ BDT, for both the background and the signal, including the $m_a = 12, 30, 60$ and 80 GeV hypotheses. The BDT score in Figure 11.17 (a) illustrates how the performance of the $a \rightarrow jj$ reconstruction improves for higher values of the a -boson mass. This is expected, since the $a \rightarrow jj$ decay, with j being an Ak4 jet, reproduces the topology from the resolved regime. Figures 11.17 (b) to (d) show the di-jet angular distance, p_T and invariant mass, respectively. In general, the di-jet system with the highest BDT score tends to have a smaller angular separation and to be more energetic for the signal samples. It also produces a resonance in the invariant mass distribution, with the exception of very low a -boson mass hypotheses such as $m_a = 12$ GeV, where the a -boson is more likely to decay to a boosted $b\bar{b}$ pair reconstructed as an Ak8 jet.




Figure 11.17: Normalised MC distributions for (a) the highest hypothesis score from the $a \rightarrow jj$ BDT, (b) the angular separation of the di-jet pair, (c) the di-jet p_T and (d) the di-jet invariant mass. Signal distributions for $m_a = 12, 30, 60$ and 80 GeV are overlayed for comparison.

11.5.2 NN for signal versus background discrimination

Events from the preselection region are classified into four non-overlapping regions based on the total number of b -objects (B -jets and b -jets), in order to reduce the fraction of background events. There are two "resolved" regions, denoted as 0B4b and 0B3b, and two "boosted" regions, denoted as 1B2b and 1B1b. They are defined in Table 11.12. A mass-parametrised NN is trained in each of these signal-enriched regions, using input features that vary depending on the final state topology.

Region	<i>B</i> -jets (40% WP)	<i>b</i> -jets (70% WP)
0B4b	= 0	≥ 4
0B3b	= 0	= 3
1B2b	= 1	≥ 2
1B1b	= 1	= 1

Table 11.12: Signal regions for signal versus background discrimination. All regions are required to have at least 2 Ak4 jets for the BDT variables to be well-defined.

The training is performed using $t\bar{t}a$ ⁶ events plus all background samples available, as shown in 11.13. Because the number of events is unbalanced, the loss function is weighted differently for signal and background samples, assigning a larger weight to the signal samples, which are underrepresented in the dataset. This prevents the NN from giving more importance to the background samples, thus ignoring key features from the $t\bar{t}a$ MC events.

Sample	0B4b	0B3b	1B2b	1B1b
$t\bar{t}a$ 12 GeV	–	15353	11096	16269
$t\bar{t}a$ 16 GeV	–	13669	11028	15878
$t\bar{t}a$ 20 GeV	1045	15440	10148	14896
$t\bar{t}a$ 30 GeV	3447	20801	7239	9964
$t\bar{t}a$ 40 GeV	5721	25618	5157	6891
$t\bar{t}a$ 60 GeV	8624	30650	3604	3803
$t\bar{t}a$ 80 GeV	10087	33083	3261	2925
$t\bar{t}a$ 100 GeV	10973	33916	3310	2685
$t\bar{t}$ +light	96	34149	2438	352628
$t\bar{t} \geq 1c$	1412	122056	6643	48749
$t\bar{t} \geq 1b$	49703	446433	61772	78655
tW	175	2857	280	3867
Other	195973	715973	71470	141348

Table 11.13: Number of signal and background MC events used in the training of the signal versus background discriminator for each training region. In this table, "Other" includes the $t\bar{t}H$, $t\bar{t}Z$, $t\bar{t}W$, tq , tZ , tWZ , Z/W +jets and diboson background processes. The 12 and 16 GeV $t\bar{t}a$ samples are excluded from the 0B4b training due to low statistics.

The mass-parametrised NN for signal versus background discrimination is trained using the PyTorch module [237], with the hyperparameters from Table 11.14. The NN contains two hidden layers with twice as many nodes as the input layer, connected by ReLu activation functions. No improvement was found from increasing the number of layers or nodes. The final layer is a single node, normalised by a sigmoid function, which provides a NN score ranging from 0 to 1. The list of input features is summarised in Table 11.15. They include some of the BDT outputs plus low- and high-level variables

⁶The $t\bar{t}a$ samples with $m_a = 25$ and 50 GeV and the $tW a$ samples were not used in the training, as it was performed prior to their inclusion in the analysis. No significant impact on the NN performance was observed.

11. Analysis overview

related to lepton and jet kinematics. Figure 11.18 illustrates some of the most important variables for signal versus background discrimination in the 0B4b and 0B3b (resolved) regions, such as the hadronic event hardness (H_T^{jets}) and the average of the angular separation between bb pairs ($\langle |\Delta\eta_{bb}| \rangle$). Figure 11.19 shows the p_T and invariant mass of the B -jet from the 1B2b and 1B1b (boosted) regions.

Hyperparameter	Value
Hidden layers	2
Hidden size	$2 \times N_{\text{features}}$
Activation function	ReLU
Learning rate	10^{-4}
Epochs	500
Patience	4
Dropout	0.3

Table 11.14: List of hyperparameters used in the NN training. The number of epochs corresponds to the maximum number allowed. The training is stopped if the loss does not improve after 4 epochs (patience). The choice of the values is based on the NN performance and the total training time.

Object	Variables
BDT $t \rightarrow j\ell$	Score, $p_T^{\ell\ell}$, $\Delta R_{j\ell}$, $\Delta\eta_{j\ell}$, $\Delta\phi_{j\ell}$, jet ID
BDT $a \rightarrow jj$	Score, p_T^{jj} , η_{jj} , m_{jj} , ΔR_{jj} , $\Delta\eta_{jj}$, $\Delta\phi_{jj}$, jet ID
Small- R jets	p_T , η , b -tagging score p_T^{bb} , m_{bb} , m_{bbb} , m_{bbbb} , ΔR_{bb} , $\Delta\eta_{bb}$, $\Delta\phi_{bb}$, $\Delta\phi_{E_T^{\text{miss}},b}$
Large- R jets	p_T , η , m ΔR_{Bb} , $\Delta\phi_{E_T^{\text{miss}},B}$
Leptons	$\Delta R_{\ell\ell}$, $\Delta\eta_{\ell\ell}$, $\Delta\phi_{\ell\ell}$, $\Delta\phi_{E_T^{\text{miss}},\ell}$ $\Delta R_{\ell\ell,b}$, $\Delta R_{\ell\ell,B}$, $\Delta R_{\ell\ell,b}$
Event	N_{jets} , H_T^{jets} , E_T^{miss}

Table 11.15: NN input variables. For bb variables, both the pair with maximum p_T and minimum ΔR are included. Angular variables with one b and/or one B use the minimum ΔR pair. m_{bbbb} and m_{bbb} correspond to the combination with maximum p_T .




Figure 11.18: MC distributions of (a,c) the hadronic event hardness and (b,d) the average of the pseudorapidity separation for $b\bar{b}$ pairs in the 0B4b and 0B3b regions, respectively. Signal distributions for $m_a = 12, 30, 60$ and 80 GeV are overlayed for comparison. The 12 GeV mass hypothesis is not displayed the 0B4b region due to low statistics.




Figure 11.19: MC distributions of (a,c) the B -jet p_T and (b,d) the B -jet invariant mass in the 1B2b and 1B1b regions, respectively. Signal distributions for $m_a = 12, 30, 60$ and 80 GeV are overlaid for comparison.

The training is performed using the k -fold method with $k = 5$. The training dataset is generated using 80% of the events and the remaining 20% is used as the validation dataset. The partition of events is done based on the event number, in such a way that each block of 20% of events is evaluated with the model that was trained with the remaining 80%. Figures 11.20 to 11.23 show the normalised NN score for signal and background for each signal hypothesis, together with the corresponding ROC curves.

Ph. D. Thesis P. Martinez
NN region 0B4b

— Signal
— Background




Figure 11.20: NN score and ROC curve in the 0B4b region for each mass hypothesis.

Ph. D. Thesis P. Martinez
NN region 0B3b

— Signal
— Background




Figure 11.21: NN score and ROC curve in the 0B3b region for each mass hypothesis.

Ph. D. Thesis P. Martinez
NN region 1B2b

— Signal
— Background




Figure 11.22: NN score and ROC curve in the 0B4b region for each mass hypothesis.

11. Analysis overview

Ph. D. Thesis P. Martinez
NN region 1B1b

— Signal
— Background




Figure 11.23: NN score and ROC curve in the 0B4b region for each mass hypothesis.

The results are consistent with the assumption that high values of m_a tend to fall in the resolved regimes (0B4b and 0B3b), which translates into a better performance of the NN for the mass points in these regions. Similarly, low values of m_a are more likely to fall into the boosted regimes (1B2b and 1B1b), and the NN provides a very good signal versus background separation in these cases. A summary of the results is provided in Figure 11.24, which shows the AUC as a function of the a -boson mass for each training region.




Figure 11.24: AUC versus mass for each signal hypothesis in the NN regions (a) 0B4b, (b) 0B3b, (c) 1B2b and (d) 1B1b.

11.6 Systematic uncertainties

Various sources of systematic uncertainties are considered in this analysis. Each systematic uncertainty is introduced as a nuisance parameter in the statistical analysis described in Chapter 8. Section 11.6.1 describes all experimental uncertainties, related to the luminosity, the pileup and the reconstruction and identification of jets and

leptons. These uncertainties are applied to all MC samples equally. Signal and background modelling uncertainties are detailed in Section 11.6.2. They are associated to the theoretical predictions for each process, and can vary depending on the sample. All systematic uncertainties use the data-driven corrections described in Section 11.4, except for the alternative MC samples, which have dedicated reweighting procedures.

11.6.1 Experimental uncertainties

Luminosity and pileup

The uncertainty on the integrated luminosity for the full Run 2 dataset is 0.83% [73], obtained using the LUCID-2 detector [194].

A variation in the pileup reweighting of simulated events is included to cover the uncertainty in the ratio of the predicted and measured inelastic cross sections in the fiducial volume defined by $M_X > 13$ GeV, where M_X is the mass of the hadronic system [218]. In practice, the nominal scale factor applied to the data pileup distribution when performing the pileup reweighting is changed into 1.0/0.99 or 1.0/1.07 to derive the up and down systematic uncertainty, instead of its nominal value of 1.0/1.03.

Leptons

Lepton systematic uncertainties are related to the trigger, reconstruction, identification and isolation, as well as the lepton energy or momentum scale and resolution.

The reconstruction, identification, and isolation efficiency of electrons and muons, as well as the efficiency of the trigger used to record the events, differ slightly between data and simulation, and is corrected by dedicated scale factors. Efficiency scale factors are measured using tag-and-probe techniques on $Z \rightarrow \ell\ell$ data and simulated samples [219, 220], and are applied to the simulation to correct for differences. The effect of these scale factors and their uncertainties are propagated as corrections to the MC event weight. In total, four independent components are considered for electrons and eight for muons.

Additional sources of uncertainty originate from the corrections applied to adjust the lepton momentum scale and resolution in the simulation to match those in data, measured using reconstructed distributions of the $Z \rightarrow \ell\ell$ and $J/\psi \rightarrow \ell\ell$ invariant masses, as well as the E/p ratio measured in $W \rightarrow e\nu$ events where E and p are the electron energy and momentum measured by the calorimeter and the tracker, respectively [133, 220]. To evaluate the effect of momentum scale uncertainties, the event selection is redone with the lepton energy or momentum varied by $\pm 1\sigma$. The event selection is redone for the momentum resolution uncertainties by smearing the lepton energy or momentum. In total, two independent components are considered for electrons, and four for muons.

Jets

Uncertainties associated with jets arise from the efficiency of pileup rejection by the Jet Vertex Tagger (JVT), from the Jet Energy Scale (JES) and Resolution (JER), and from the different flavour tagging algorithms, which include DL1r and DeXTer.

Scale factors are applied to correct for discrepancies between data and MC for JVT efficiencies. These scale factors are estimated using $Z \rightarrow \mu\mu$ events with tag-and-probe techniques [221]. The effect of these scale factors, as well as of their uncertainties, are propagated as corrections to the MC event weight.

The jet energy scale and its uncertainty are derived by combining information from test-beam data, LHC collision data and simulation [222]. The uncertainties from these measurements are factorised into 8 independent sources. The jet energy resolution was measured in Run 2 data and MC simulation as a function of jet p_T and rapidity using di-jet events with a similar method as in Ref. [223]. The combined uncertainty is propagated by smearing the jet p_T in the MC simulation, yielding 8 independent sources.

To correct the DL1r flavour tagging efficiencies in simulated samples to match efficiencies in data, scale factors are derived by the flavour-tagging working group. They are calculated as a function of p_T for b -jets, c -jets and light jets separately in dedicated calibration analyses. For b -jet efficiencies, $t\bar{t}$ +jets events in the dilepton topology are used, exploiting the very pure sample of b -jets arising from the decay of the top quarks [224]. For c -jet mistag rates, $t\bar{t}$ +jets events in the single-lepton topology are used, exploiting the c -jets from the hadronically decaying W boson [225]. The so-called negative-tag method is used in Z +jets events [226] for the light-jets mistag rates. In the three calibration analyses, a large number of uncertainty components are considered, and a principal component analysis is performed, yielding 45, 20, and 20 eigen-variations for b -, c - and light-jets, respectively, which are taken as uncorrelated sources of uncertainty. These eigen-variations correspond to the number of p_T bins (9, 4, and 4, respectively), multiplied by the number of DL1r bins (5).

The use of DeXTer introduces additional scale factors to correct the efficiency difference between simulated samples and data. The scale factors for the B -tagger are provided by the DeXTer development team in collaboration with the ATLAS flavour-tagging working group [153]. They are derived for B - and b -jets, separately for each p_T bin and for each B -tagging WP. The calibration measurements with data are performed using both $t\bar{t}$ +jets and Z +jets events simultaneously in order to measure the B -jet tagging and b -jet mistagging efficiencies. Additional scale factors were computed a posteriori to extend DeXTer B -tagging to jets with $p_T > 200$ GeV, with their corresponding systematic variations.

Tracks

Systematic uncertainties related to the track selection efficiency are determined by changing the amount of tracker material and the physical models in the GEANT4 simulation. Systematic variations on the number of fake tracks are applied based on the recommendation for the Loose track selection working point estimated and provided by the ATLAS tracking combined performance group. Dedicated systematic uncertainties are considered for the track parameters, including the transverse and longitudinal impact parameters and the track sagitta. The event selection is redone for each systematics variation.

To correct the mismodelling in the Ak8 track jet mass, additional mass scale corrections

are estimated from data. The Ak8 track-jet mass scale is varied by $\pm 5\%$ and compared to the nominal results with a mass scale of 1.01.

Missing transverse energy

All previously described uncertainties on energy scales or resolutions of the reconstructed objects (hard components) are propagated to the missing transverse momentum. Additional uncertainties in the scale and resolution of the soft term are considered, to account for disagreement between Data and MC for the p_T balance between the hard and soft components, for a total of 3 independent sources: an offset along the hard component p_T axis and the smearing resolution along and perpendicular to this axis [154, 230].

Summary

Table 11.16 summarises all the experimental systematic uncertainties, including the type (normalisation, shape or both) and the number of components. They are applied equally to all MC samples.

Systematic uncertainty	Type	Components
General		
Luminosity	N	1
Pileup	SN	1
Physics objects		
Electrons	SN	6
Muons	SN	12
Jets	SN	32
Tracks	SN	13+1
E_T^{miss}	SN	3
Flavour tagging		
DL1r b	SN	45
DL1r c	SN	20
DL1r light	SN	20
DeXTer	SN	12+4

Table 11.16: List of experimental systematic uncertainties included in the analysis. "N" means that the uncertainty is taken as normalisation-only for all processes and channels affected, while "SN" means that the uncertainty is taken on both the shapes and the normalisation. Some of the systematic uncertainties are split into several components for a more accurate treatment. Additional systematic uncertainties due to the Ak8 track mass scale and the high- p_T DeXTer tagging, particular to this analysis, are shown separately with a "+" sign.

11.6.2 Modelling uncertainties

This section details the modelling uncertainties associated to the missing higher orders in the perturbative expansion of the partonic cross section and the uncertainties in the proton PDFs. These uncertainties are applied to the signal and the main backgrounds,

$t\bar{t}$ +jets, tW , $t\bar{t}H$ and $t\bar{t}Z$. They are assumed to be negligible for any other background process.

General modelling uncertainties

Variations in the renormalisation and factorisation scales are used to estimate the uncertainty due to missing higher order corrections. The samples used include on-the-fly variations corresponding to seven different configurations of these two parameters: $\{\mu_R, \mu_F\} \times \{0.5, 0.5\}, \{0.5, 1\}, \{1, 0.5\}, \{1, 1\}, \{2, 1\}, \{1, 2\}, \{2, 2\}$. The uncertainties are combined by taking an envelope of all the variations. The cross section normalisation component of the scale uncertainties is not considered for the $t\bar{t}$ +jets and tW backgrounds, whose normalisation is calculated in the fit. In that case, only the shape uncertainty is used.

The uncertainty due to ISR is estimated by simultaneously changing μ_R and μ_F in the ME and μ_R^{ISR} in the PS, while the uncertainty due to FSR is estimated by changing μ_R^{FSR} in the PS. For the ISR, the amount of radiation is increased (decreased) by scaling μ_R and μ_F by a factor of 0.5 (2) and by varying the renormalisation scale for QCD emission in the ISR by a factor of 0.549 (1.960), corresponding to the Var3cUp (Var3cDown) variation from the A14 tune which sets α_s^{ISR} to 0.140 (0.115) instead of the nominal 0.127. For the FSR, the amount of radiation is increased (decreased) by varying the renormalisation scale for QCD emission in the FSR by a factor of 0.5 (2), corresponding to $\alpha_s^{\text{FSR}} = 0.1423$ (0.1147) instead of the nominal 0.127⁷.

PDF uncertainties follow the PDF4LHC recommendations [108] and are available as weights in the samples. The α_s uncertainty is derived using the same PDF set evaluated with two different α_s values. The two uncertainties from the PDF and α_s are added in quadrature⁸.

Signal modelling uncertainties

The nominal $t\bar{t}a$ samples generated using MADGRAPH5_aMC@NLO + PYTHIA 8 are compared to MADGRAPH5_aMC@NLO + HERWIG 7 to extract a parton shower modelling uncertainty. The comparison is done after normalising both $t\bar{t}a$ samples to the same cross section. There are no alternative samples generated with MADGRAPH5_aMC@NLO + HERWIG 7 for the tW process; nonetheless, the resulting uncertainty is not significant for any of the signal mass hypotheses.

Background modelling uncertainties

Matrix element uncertainties in the $t\bar{t}$ +jets samples are taken into account by comparing the 5FS nominal MC with an alternative pThard=1 sample. For parton shower uncertainties, the nominal POWHEGBOX + PYTHIA 8 sample is compared to the POWHEGBOX + HERWIG 7 sample.

The modelling uncertainties for the 4FS $t\bar{t}+b\bar{b}$ sample follow the recommendations in Ref. [202]. The uncertainty associated to the matching between the ME calculations and the parton shower is taken into account by comparing the nominal MC with an

⁷ISR and FSR uncertainties are not available for the signal and $t\bar{t}Z$.

⁸ α_s variations are not available for tW and $t\bar{t}Z$.

alternative $p_{\text{Thard}}=1$ sample. For PS and hadronisation uncertainties, the nominal POWHEGBOX + PYTHIA 8 sample is compared to the POWHEGBOX + HERWIG 7 sample. The effect of the choice of the h_{bzd} is evaluated by comparing the nominal sample ($h_{\text{bzd}} = 5$) with an alternative sample in which the scale has been set to 2. The uncertainty due to the recoiler-choice of ISR emissions is evaluated by comparing the nominal POWHEGBOX + PYTHIA 8 sample, in which the whole final state acts as recoil of the ISR emission, with an alternative POWHEGBOX + PYTHIA 8 sample in which only one final-state parton takes the recoil of the emission.

For the tW sample, ME uncertainties are evaluated by comparing the nominal POWHEGBOX + PYTHIA 8 sample with an alternative using MADGRAPH5_aMC@NLO + PYTHIA 8. PS uncertainties are treated in the same way as it is done for $t\bar{t}$ +jets. In addition, a third alternative sample is included to account for uncertainties on the interference between $t\bar{t}$ +jets and tW , where the nominal sample generated using diagram removal (DR) is compared to another one generated using diagram subtraction (DS).

The nominal $t\bar{t}H$ samples are compared to MADGRAPH5_aMC@NLO + PYTHIA 8 for the matrix element uncertainties and POWHEGBOX + HERWIG 7 for the parton shower uncertainties. The nominal $t\bar{t}Z$ samples are compared to an alternative sample generated using SHERPA 2.2.0.

Because all the $t\bar{t}$ +jets and tW background normalisations are left as free floating in the fit, differences in total yields are not taken into account when comparing the nominal samples with their corresponding alternatives. In practice, this means that the yields of each component in the alternative samples are normalised to the HF corrected nominal yield. Table 11.17 shows the correction factor applied to each process in each alternative sample. In addition, a H_T shape-correction is calculated for each set of alternative samples (Nominal AFII, $p_{\text{Thard}}=1/\text{MADGRAPH5_aMC@NLO} + \text{PYTHIA 8 AFII}$ and POWHEGBOX + HERWIG 7 AFII), in the same way as it is done for the nominal sample. This last step does not apply to $t\bar{t}+b\bar{b}$ (4FS), since the corresponding nominal sample does not receive any kinematic reweighting. In the case of the tW DS sample, the nominal reweighting is used, given that the contribution of this process in the reweighting region is very small. Figure 11.25 shows the H_T correction for each of the alternative samples, in comparison to the nominal one.

Process	Sample	Norm. factor
$t\bar{t}$ +light	PhPy8 5FS FS	0.91
	PhPy8 5FS AFII	0.91
	PhPy8 pThard=1 5FS AFII	0.91
	PhH7 5FS AFII	0.86
$t\bar{t}+\geq 1c$	PhPy8 5FS FS	1.58
	PhPy8 5FS AFII	1.58
	PhPy8 pThard=1 5FS AFII	1.54
	PhH7 5FS AFII	1.71
$t\bar{t}+\geq 1b$	PhPy8 4FS bzd5 FS	1.13
	PhPy8 4FS bzd5 AFII	1.12
	PhPy8 pThard=1 4FS bzd5 FS	1.14
	PhH7 4FS bzd5 AFII	1.11
	PhPy8 4FS pTdef1 (bzd=2) AFII	1.15
	PhPy8 4FS pTdef1 (bzd=2) dipole recoil AFII	1.15
tW	PhPy8 DR FS	0.91
	PhPy8 DR AFII	0.90
	MG_aMC@NLOPy8 AFII	0.85
	PhH7 AFII	0.99

Table 11.17: Correction factors applied to each of the $t\bar{t}$ +jets and tW samples. The first row in each sample corresponds to the nominal $t\bar{t}$ +HF reweighting, which is shown for comparison. Since all alternative samples are simulated in AFII, an AFII version of the nominal is included to be used as reference. The following abbreviations are used: "Ph" = POWHEGBox, "Py8" = Pythia 8, "H7" = HERWIG 7, "MG" = MADGRAPH. Nominal samples (FS and AFII) are highlighted in grey.




Figure 11.25: $t\bar{t}$ +jets and tW correction as a function of H_T^{red} for the nominal (black line) and alternative samples (grey, red and green lines).

Reweighting uncertainties

The reweighting function used to correct the $t\bar{t}$ +jets and tW backgrounds consists in two steps. First, a HF normalisation correction and second, a p_T shape correction using H_T . As described in Section 11.4, these corrections are obtained by performing two fits, the first one with three free-floating normalisation parameters and the second one with three free-floating parameters corresponding to a hyperbolic function. Two covariance matrices corresponding to the two fits are extracted. In the ideal case where the three parameters in the fits were uncorrelated, the uncertainties on each parameter would correspond to the square root of the diagonal elements of the covariant matrix. Given that the resulting matrices contain non-diagonal elements, the covariance matrices are diagonalised and six variations (three up and three down) are extracted per matrix. The corresponding systematic uncertainties are added in the analysis by implementing an alternative reweighting function for each correlated variation. This function is calculated by replacing the nominal values by each of the variations. In the case of the HF variations, new fits to H_T^{red} are performed after having individually applied each of the HF variations, hence propagating the effect of one step of the reweighting into the next. In conclusion, the analysis includes three systematic uncertainties corresponding to the variations of $t\bar{t}$ +light and tW , $t\bar{t}+\geq 1c$ and $t\bar{t}+\geq 1b$ normalisation factors and three additional ones for the variations of the three parameters of the hyperbolic fit (c_0 , c_1 and c_2).

Summary

Table 11.18 summarises the weight-based modelling uncertainties. Table 11.19 shows the alternative MC samples used to derive the MC generator and parton shower uncertainties.

Systematic uncertainty	Type	Components
Signal		
μ_R, μ_F	SN	3
PDF+ α_s	SN	32
$t\bar{t}$+jets		
μ_R, μ_F	S	3
ISR, FSR	SN	2
PDF+ α_s	SN	31
Reweighting	SN	6
tW		
μ_R, μ_F	S	3
ISR, FSR	SN	2
PDF	SN	30
Reweighting	SN	6
$t\bar{t}H$		
μ_R, μ_F	SN	3
ISR, FSR	SN	2
PDF+ α_s	SN	32
$t\bar{t}Z$		
μ_R, μ_F	SN	3
PDF	SN	100

Table 11.18: List of modelling systematic uncertainties included in the analysis. "N" means that the uncertainty is taken as normalisation-only for all processes and channels affected, while "SN" means that the uncertainty is taken on both the shapes and the normalisation. Some of the systematic uncertainties are split into several components for a more accurate treatment. Systematics derived from alternative samples are not included.

11. Analysis overview

Process	ME generator	PDF set	PS generator	Sim.
$t\bar{t}a, a \rightarrow b\bar{b}$	MG_aMC@NLO v2.3.3	NNPDF3.0 NLO	PYTHIA 8.244	AFII
	MG_aMC@NLO v2.3.3	NNPDF3.0 NLO	HERWIG 7.2.1	AFII
$t\bar{t} + \text{jets}$	POWHEGBox-v2	NNPDF3.0 NLO	PYTHIA 8.230	AFII
	POWHEGBox-v2 pThard=1	NNPDF3.0 NLO	PYTHIA 8.306	AFII
	POWHEGBox-v2	NNPDF3.0 NLO	HERWIG 7.1.3	AFII
$t\bar{t} + b\bar{b}$	POWHEGBox-Res	NNPDF3.1 NNLO	PYTHIA 8.230	AFII
	POWHEGBox-Res pThard=1	NNPDF3.1 NNLO	PYTHIA 8.306	FS
	POWHEGBox-Res	NNPDF3.1 NNLO	HERWIG 7.1.3	AFII
	POWHEGBox-Res bzd=2	NNPDF3.0 NNLO	PYTHIA 8.230	AFII
	POWHEGBox-Res bzd=2	NNPDF3.0 NNLO	PYTHIA 8.230 Dipole recoil	AFII
tW	POWHEGBox-v2	NNPDF3.0 NLO	PYTHIA 8.230	AFII
	POWHEGBox-v2 (DS)	NNPDF3.0 NLO	PYTHIA 8.230	FS
	MG5_aMC@NLO v2.6.2	NNPDF3.0 NLO	PYTHIA 8.230	AFII
	POWHEGBox-v2	NNPDF3.0 NLO	HERWIG 7.0.4	FS
$t\bar{t}H$	POWHEGBox-v2	NNPDF3.0 NLO	PYTHIA 8.230	AFII
	MG5_aMC@NLO v2.6.0	NNPDF3.0 NLO	PYTHIA 8.230	AFII
	POWHEGBox-v2	NNPDF3.0 NLO	HERWIG 7.0.4	AFII
$t\bar{t}Z$	MG5_aMC@NLO v2.3.3	NNPDF3.0 NLO	PYTHIA 8.210	AFII
	SHERPA 2.2.0	NNPDF3.0 NNLO	SHERPA	AFII

Table 11.19: Summary of alternative MC samples used in the $t\bar{t}a, a \rightarrow b\bar{b}$ analysis. Note that MADGRAPH is abbreviated to MG. All alternative samples are simulated using AFII with the exception of the pThard=1 alternative for $t\bar{t}+b\bar{b}$ and the DS alternative for tW . The AFII samples are compared to the AFII version of the nominal FS MC, highlighted in grey.

Chapter 12

Results and conclusions

12.1 Fit setup

12.1.1 General considerations

In order to test the presence of a $t\bar{t}a$ signal, a binned maximum-likelihood fit to the data is performed simultaneously in all signal regions, for each mass hypothesis. The procedure used to quantify the level of agreement with the B-only or S+B hypothesis and to determine exclusion limits is based on the profile likelihood ratio test and the CL_s method. The parameter of interest is the signal strength, μ , a multiplicative factor to the cross section of the signal process. For simplicity, it is normalised to a reference value of 32.058 fb, corresponding to the observed $t\bar{t}H$, $H \rightarrow b\bar{b}$ dileptonic cross section from Ref. [238]. This cross section is expected to be larger than the MC prediction for $t\bar{t}a$, $a \rightarrow b\bar{b}$. Therefore, it provides a conservative reference for blinding during the design of the analysis strategy.

12.1.2 Signal and control regions

Signal regions

The signal regions for the $t\bar{t}a$, $a \rightarrow b\bar{b}$ analysis are based on the four NN regions used to train the signal versus background discriminator. They all fulfil the preselection requirements from Section 11.3, and are further classified according to their b -object multiplicity. There are two regions targetting the resolved $a \rightarrow b\bar{b}$ decay, denoted as 0B4b and 0B3b, and two regions targetting the boosted $a \rightarrow b\bar{b}$ decay, denoted as 1B2b and 1B1b+1bL. The 1B1b+1bL signal region is derived from the 1B1b NN region, with at least one additional loose b -jet (+1bL) tagged with the DL1r 85% WP. This condition reduces the fraction of $t\bar{t}$ +light events, which leads to a better convergence of the fit and less constrained NPs. The fitted distribution in these signal regions is the NN score from the signal versus background discriminator, which ranges from 0 to 1.

Control regions

A $t\bar{t}$ +light and $t\bar{t}+\geq 1c$ control region is defined in order to improve the performance of the fit. This is necessary because the signal regions do not provide enough constraining power on their own to determine the $t\bar{t}$ +light and $t\bar{t}+\geq 1c$ normalisation factors. This control region, denoted as 0B2b+1bL, is required to have no B -jets tagged with DeXTer 40% WP, exactly two b -jets tagged with DL1r 70% WP and at least one additional b -jet tagged with DL1r 85% WP, which reduces the fraction of $t\bar{t}$ +light in favour of the $t\bar{t}+\geq 1c$ category. The region is binned as a function of the sumPCBTag variable. The central portion of the distribution, $12 \leq \text{sumPCBTag} \leq 15$, is used due to its higher $t\bar{t}+\geq 1c$ content. Figure 12.1 shows the comparison between NFs with and without the $t\bar{t}$ +light/ $\geq 1c$ control region. The use of the CR improves the determination of the

$t\bar{t}$ +light and $t\bar{t} \geq 1c$ normalisation factors.




Figure 12.1: Normalisation factors for the $t\bar{t}$ +jets components (a) without and (b) with the 0B2b+1bL CR. The test is performed using a B-only fit to the bins where the signal to background ratio is inferior to 1%.

Summary

A description of each signal and control region is included in Table 12.1. Their background composition is shown in Figure 12.2. As expected, regions with higher b -jet multiplicity, such as 0B4b, 0B3b and 1B2b are dominated by $t\bar{t} \geq 1b$ and, to lesser extent, $t\bar{t} \geq 1c$. The 1B1b+1bL signal region and the 0B2b+1bL control region receive a larger contribution from $t\bar{t}$ +light.

Region	B -jets (40% WP)	b -jets (70% WP)	b -jets (85% WP)
SR 0B4b	= 0	≥ 4	≥ 4
SR 0B3b	= 0	= 3	≥ 3
SR 1B2b	= 1	≥ 2	≥ 2
SR 1B1b+1bL	= 1	= 1	≥ 2
CR 0B2b+1bL	= 0	= 2	≥ 3

Table 12.1: Signal and control regions used in the fit. All regions are required to have at least two Ak4 jets for the BDT variables to be well-defined. Note that b -jets tagged with the DL1r 70% WP are also tagged with the 85% WP.

12.1.3 Normalisation factors for $t\bar{t}$ +jets

In addition to the signal strength, the test statistics includes three additional free parameters to adjust the normalisation of the three main background components: $t\bar{t}$ +light and tW , $t\bar{t} \geq 1c$ and $t\bar{t} \geq 1b$. Different fit models in which the $t\bar{t} \geq 1b$ normalisation was further split into sub-components were studied, but they were discarded as they suffered from larger uncertainties and statistical fluctuations.

Ph. D. Thesis

P. Martinez




Figure 12.2: Background composition in the signal and control regions of the fit.

Studies on $t\bar{t} + \geq 1b$ normalisation

In some ATLAS analyses, the $t\bar{t} + \geq 1b$ category is divided into further sub-categories with different b and B multiplicities in order to provide a more accurate treatment for their normalisation factors and systematic uncertainties. For this search, three options were compared, as described in Table 12.2.

Test 1	$t\bar{t} + \geq 1b$	Events containing at least one additional b or B
Test 2	$t\bar{t} + \geq 1B$	Events containing at least one additional B
	$t\bar{t} + \geq 1b$	Events containing at least one additional b and no B 's
Test 3	$t\bar{t} + \geq 1B$	Events containing at least one additional B
	$t\bar{t} + 1b$	Events containing exactly one additional b and no B 's
	$t\bar{t} + \geq 2b$	Events containing at least two additional b 's and no B 's

Table 12.2: Test categorisations of the $t\bar{t} + \geq 1b$ sample. In this context, " b " refers to an Ak4 jet with $p_T > 15$ GeV and $|\eta| < 2.5$ matched to one b -hadron not originating from the $t\bar{t}$ system and " B " refers to an Ak4 jet with $p_T > 15$ GeV and $|\eta| < 2.5$ matched to at least two b -hadrons not originating from the $t\bar{t}$ system.

The three configurations were tested in a B -only fit including only the bins where the signal to background ratio is less than 1%, as shown in Figure 12.3. From these results, it was concluded that the regions used in the fit were not capable of constraining the various normalisation factors for the $t\bar{t} + \geq 1b$ sub-categories. As a consequence, only the global $t\bar{t} + \geq 1b$ category was kept.




Figure 12.3: Normalisation factors for the $t\bar{t}$ +jets components with different categorisation schemes for $t\bar{t}+\geq 1b$ as a function of the a -boson mass. These results show how further decomposition of the $t\bar{t}+\geq 1b$ category can lead to fit instability.

12.1.4 Binning, smoothing and pruning

The binning of the NN distributions is optimised for every signal and region in such a way that the signal to background ratio increases from left (low NN score) to right (high NN score). The method is based on the AutoBin method available in TREXFITTER, which determines the binning in an iterative approach based on the following variable:

$$Z = Z_S \frac{n_S}{N_S} + Z_B \frac{n_B}{N_B}, \quad (12.1)$$

where n_S (n_B) represents the signal (background) yields in a single bin and N_S (N_B) refers to the total signal (background) yields of the distribution. The two input parameters, namely Z_S and Z_B , typically sum up to the total number of bins of the rebinning process. The algorithm starts from the lowest boundary of the variable and the upper edge of a bin is defined when the bin satisfies $Z \geq 1$. The values of choice are $Z_S = Z_B = 7$, since they provide the best estimation of the expected limit. $Z_S = Z_B = 2$ is used for $m_a = 12, 16$ GeV in the 0B4b signal region due to the low statistics.

Systematic uncertainties are derived from MC binned distributions for each sample in each signal and control region. This method is susceptible to statistical fluctuations that can arise when the number of events in a bin is very small. If these fluctuations become too large, they can lead to a bad template for the uncertainty. In order to mitigate these effects and avoid problems with the convergence of the minimisation, smoothing algorithms are applied to the histograms prior to the fit. Smoothing is applied to modelling uncertainties of the main backgrounds. For smaller backgrounds, for systematic uncertainties with weights close to unity, or when using the same events but with different calibrations, no smoothing is applied.

Systematic uncertainties are pruned in order to reduce the CPU time used for fitting and also to help with the convergence of the fit. This is done by considering bin-by-bin variations of each input histogram, relative to each sample in each region. The pruning threshold is set to 1%, meaning that if the effect of a nuisance parameter is smaller than that, then it does not contribute to the fit. Each nuisance parameter is separated into a shape component and a normalisation component, and the pruning criterion applies to each of them independently.

12.2 Fit results

This section summarises the results of the analysis. Expected results based on the fits to the Asimov dataset are presented in Section 12.2.1. Observed results from the fits to data are shown in Section 12.2.2. Only the fits under the 30 GeV mass hypothesis are shown. Expected and observed upper limits on $t\bar{t}a$, $a \rightarrow b\bar{b}$ production, together with the list of most relevant systematic uncertainties, are included in Section 12.2.3. Blinded fits to data, used for tests, are presented in Appendix B.2.

12.2.1 B-only fit to the Asimov dataset


This section shows the results from the fit to the Asimov dataset. The Asimov dataset is built such that the event count in each bin is set to the expected event yield for the chosen model parameters. In this test, the B-only hypothesis is assumed ($\mu = 0$), where the Asimov dataset becomes identical to the background MC. The NPs are fitted in both the signal and control regions, and the expected upper limit to the $t\bar{t}a$, $a \rightarrow b\bar{b}$ production cross section is extracted.

Figures 12.4 and 12.5 show the NN score distribution for the background processes in the signal and control regions, before and after performing the Asimov fit with $\mu = 0$, as well as the total error bands. Figure 12.6 contains the normalisation factors for the different components of $t\bar{t}$ +jets and the signal strength, with their corresponding uncertainties. The uncertainties of the three $t\bar{t}$ +jets normalisation factors range between 15% and 30%. Thanks to the 0B4b SR, the $t\bar{t}+\geq 1b$ normalisation is the most precise. Figure 12.7 shows the correlations between NPs. No strong correlations between systematic uncertainties are observed. The only exceptions are some of the DeXTer eigenvalues, whose behaviour is very similar to the $H \rightarrow aa \rightarrow 4b$ results, and the NFs for $t\bar{t}$ +jets, which have some correlations with other NPs that affect the overall normalisation. Figures 12.8 and 12.9 show the constraints on the relevant nuisance

parameters. Some systematic uncertainties such as DeXTer B -tagging are expected to be constrained, as the pre-fit uncertainties are propagated from the calibration in order to avoid profiling the same systematics twice in the fit. The Ak8 track mass scale uncertainty is also expected to be constrained. Some of the modelling uncertainties have a manifest degree of constraint as well. These are, for example, the $t\bar{t}$ +light HERWIG 7 alternative sample or the $t\bar{t}+\geq 1b$ pThard=1 alternative samples. These constraints are understood by studying the behaviour of the systematic uncertainty in each region individually. It is observed that regions that have enough statistics prevent these NPs from varying within the full $\pm 1\sigma$ range, given their large impact on the cross section. Figures 12.10 to 12.12 illustrate the previously described constraints.

The ranked impact from the dominant systematics to the signal strength is shown in Figure 12.13. The most significant systematic uncertainties originate from $t\bar{t}+\geq 1b$ modelling and DeXTer B -tagging. The signal strength is also sensitive to the $t\bar{t}$ +jets normalisation factors.

Additional studies of the fit response under the S+B hypothesis ($\mu > 0$) for various a -boson masses, based on the Asimov dataset as well, are included in Appendix C.

Figure 12.4: Pre-fit distributions from the B-only fit to the Asimov dataset with $m_a = 30$ GeV.

12. Results and conclusions




Figure 12.5: Post-fit distributions from the B-only fit to the Asimov dataset with $m_a = 30$ GeV.




Figure 12.6: Normalisation factors for the $t\bar{t}$ +jets samples in the fit to the Asimov dataset with $m_a = 30$ GeV.




Figure 12.7: Correlation matrix for the relevant NPs in the B-only fit to the Asimov dataset with $m_a = 30$ GeV.

12. Results and conclusions




Figure 12.8: Fitted nuisance parameters from the B-only fit to the Asimov dataset with $m_a = 30$ GeV (I).




Figure 12.9: Fitted nuisance parameters from the B-only fit to the Asimov dataset with $m_a = 30$ GeV (II).

12. Results and conclusions




Figure 12.10: Up (red) and down (blue) 1σ variation of the third eigenvalue of the DeXTer systematic uncertainty with respect to the NN score distribution in the (a) 1B2b SR and (b) 1B1b+1bL SR, corresponding to the $t\bar{t}+1b$ MC sample.




Figure 12.11: Up (red) and down (blue) 1σ variation of the HERWIG 7 PS uncertainty with respect to the sumPCBTag distribution in the 0B2b+1bL CR, corresponding to the $t\bar{t}+\text{light}$ MC sample.




Figure 12.12: Up (red) and down (blue) 1σ variation of the $p_{T\text{Hard}}=1$ uncertainty with respect to the NN score distribution for the (a) $t\bar{t}+1b$ and the (b) $t\bar{t}+2b$ MC samples in the 0B4b SR.




Figure 12.13: Ranked systematic impacts to the signal strength from the B-only fit to the Asimov dataset with $m_a = 30$ GeV.


12.2.2 S+B fit to data

This section shows the results from the fit to data that is used to derive the final limits. In this fit, the signal plus background hypothesis is assumed, and the signal strength is left as a free-floating parameter.

Figures 12.14 and 12.14 summarise the data to MC comparison in both the SRs and CRs, before and after the fit. There is a good post-fit agreement between data and simulation in all regions. Figure 12.16 contains the normalisation factors for the different components of $t\bar{t}$ +jets, with their corresponding uncertainties. They are consistent with the values expected from the Asimov fit. Figure 12.17 shows the correlations between NPs, which are also consistent with the previous tests. Figures 12.18 and 12.19 show the pulls and constraints for all relevant NPs. There are some small constraints in the jet reconstruction systematics, which do not have a significant impact in the signal strength. Constraints in the flavour tagging and Ak8 track mass scale uncertainties are expected, as argued in the previous sections. Constraints in the modelling systematics from the $t\bar{t}$ +jets samples are expected as well, due to the low statistical uncertainties in some of the fitted bins. Even though some NPs are pulled, they are all within $\pm 1\sigma$.

The fitted signal strength resulting from the S+B fit to data for $m_a = 30$ GeV is $\hat{\mu} = 1.42^{+0.86}_{-0.67}$. The ranked impact from the dominant systematics to the signal strength is shown in Figure 12.20. The leading systematic uncertainty is the 0th eigenvalue¹ of the high- p_T ($p_T > 200$ GeV) DeXTer B -tagging uncertainty. Other relevant systematics are $t\bar{t}+\geq 1b$ modelling and signal modelling. Table 12.3 summarises the impact on $\hat{\mu}$ from the various sources of statistical and systematic uncertainty, grouped by their type. Finally, Table 12.4 shows the post-fit event yields in each of the signal and control regions for the 30 GeV a -boson mass hypothesis. For signal and background MC samples, the full systematic uncertainty is included.

¹Eigenvalues are indexed starting from zero.

Figure 12.14: Pre-fit plots from the S+B fit to data with $m_a = 30$ GeV.

12. Results and conclusions





Figure 12.15: Post-fit plots from the S+B fit to data with $m_a = 30$ GeV.



12. Results and conclusions





Figure 12.18: Fitted nuisance parameters from the S+B fit to data with $m_a = 30$ GeV (I).

Figure 12.19: Fitted nuisance parameters from the S+B fit to data with $m_a = 30$ GeV (II).

12. Results and conclusions




Figure 12.20: Ranked impact of the systematic uncertainties to the signal strength from the S+B fit to data with $m_a = 30$ GeV.

Category	Impact on $\hat{\mu}$
Total stat.	30.0%
Total syst.	49.0%
General	0.8%
Lepton reconstruction	0.1%
Jet reconstruction	10.7%
Track reconstruction	3.2%
E_T^{miss}	0.6%
DeXTer B -tagging	39.2%
DL1r b -tagging	7.8%
Signal modelling	16.4%
$t\bar{t}$ +light modelling	4.3%
$t\bar{t} + \geq 1c$ modelling	4.0%
$t\bar{t} + \geq 1b$ modelling	27.7%
tW modelling	1.6%
$t\bar{t}H$ modelling	0.6%
$t\bar{t}Z$ modelling	0.5%
Reweighting	<0.1%
MC stats.	6.6%
Signal MC stats.	6.5%
Norm. factors	14.8%

Table 12.3: Relative impacts over the fitted signal strength $\hat{\mu}$ obtained from the S+B fit to data. The statistical uncertainty (without MC stats.) is included for comparison.

Sample	SR 0B4b	SR 0B3b	SR 1B2b	SR 1B1b+1bL	CR 0B2b+1bL
Signal 30 GeV	28 ± 14	180 ± 100	70 ± 30	35 ± 16	120 ± 60
$t\bar{t}$ +light	4 ± 2	1400 ± 300	130 ± 20	1110 ± 180	17000 ± 3000
$t\bar{t} \geq 1c$	60 ± 20	4700 ± 1200	380 ± 140	700 ± 200	12000 ± 3000
$t\bar{t}+1b$	70 ± 7	4800 ± 500	640 ± 50	350 ± 50	3200 ± 600
$t\bar{t}+2b$	870 ± 40	3700 ± 400	460 ± 40	160 ± 20	1700 ± 300
$t\bar{t} \geq 3b$	85 ± 10	130 ± 20	32 ± 7	6.5 ± 1.0	40 ± 8
$t\bar{t} \geq 1B$	67 ± 11	1130 ± 140	630 ± 70	300 ± 50	650 ± 150
tW	22 ± 12	360 ± 140	50 ± 20	64 ± 17	800 ± 200
$t\bar{t}H$	62 ± 9	220 ± 20	31 ± 4	14 ± 2	136 ± 13
$t\bar{t}Z$	27 ± 6	120 ± 20	15 ± 3	11 ± 2	130 ± 30
$t\bar{t}W$	1.39 ± 0.12	32 ± 2	2.8 ± 0.2	6.6 ± 0.5	92 ± 5
Other	13 ± 2	360 ± 40	44 ± 4	71 ± 10	970 ± 120
Total MC	1310 ± 40	17240 ± 130	2470 ± 50	2870 ± 50	36350 ± 190
Data	1301	17242	2479	2866	36350

Table 12.4: Post-fit signal and background yields in the five signal regions for the $m_a = 30$ GeV hypothesis. The total MC and data yields are included in the last two rows.

12.2.3 Expected and observed limits

The expected and observed upper limits on $\sigma(t\bar{t}a) \times \text{BR}(a \rightarrow b\bar{b})$ resulting from the S+B fit are shown in Figure 12.21 (a) as a function of the a -boson mass, which ranges from 12 to 100 GeV. No significant excess over the background expectation is observed for any signal mass hypothesis. The contribution from each signal region to the fit can be understood by looking at the expected limits from the Asimov fit, as shown in Figure 12.21 (b). This plot shows the expected limits calculated using each SR individually, in combination with the 0B2b+1bL control region. The 1B2b and 1B1b+1bL SRs are the most sensitive at very low masses, such as 12 and 16 GeV, and plays a key role in improving sensitivity in the boosted regime. Between 20 and 40 GeV, the $a \rightarrow b\bar{b}$ decays can be either boosted or resolved, meaning that all regions contribute to this topology. The 0B4b and 0B3b SRs show the best performance for $m_a \gtrsim 30$ GeV, where the $t\bar{t}a$ signal enters the fully resolved regime.




Figure 12.21: (a) Expected and observed 95% CL upper limits of $\sigma(t\bar{t}a) \times \text{BR}(a \rightarrow b\bar{b})$ as a function of m_a . The red lines correspond to different values of the coupling to the top quark. (b) Expected 95% CL upper limits of $\sigma(t\bar{t}a) \times \text{BR}(a \rightarrow b\bar{b})$ as a function of m_a for each SR. The dashed line corresponds to the combined expected limit from (a). All limits are calculated for the available signal mass points and smoothed using ROOT for easier visualisation.

12.3 Conclusions

A search for a light pseudoscalar particle a produced in association with a pair of top-antitop quarks in the dilepton decay channel has been performed using the full Run 2 dataset. The analysis targets the $a \rightarrow b\bar{b}$ decay channel, which is expected to have the largest branching ratio for Yukawa-like couplings of the a -boson to fermions. The search covers the mass range between 12 and 100 GeV, including two kinematic regimes: the boosted regime, where the decay products from the a -boson are collimated into one large B -jet and the resolved regime, where the two b -quarks from a -decay yield two separate b -jets. No significant excess above the SM background expectation is observed. Upper limits to the $t\bar{t}a$, $a \rightarrow b\bar{b}$ production cross section range between 0.1 and 0.9 pb, excluding couplings to the top quark larger than the SM top Yukawa coupling ($g_t > 1$) in the 12 to 100 GeV a -boson mass range, and $g_t > 0.5$ in the 50 to 80 GeV a -boson mass range.

Summary

This dissertation describes two searches for light spin-0 particles (a -bosons) using 140 fb^{-1} of proton-proton collisions at 13 TeV recorded with the ATLAS experiment during the Run 2 data taking period (2015–2018) of the Large Hadron Collider. Both of them focus on multi- b final states, which are favoured under the assumption of Yukawa-like couplings of the a -boson. They benefit from having two leptons in the final state, which provide a clear signal for event triggering and help to reduce difficult-to-model hadronic backgrounds.

Both the $H \rightarrow aa \rightarrow 4b$ and the $t\bar{t}a, a \rightarrow b\bar{b}$ searches have been developed with a similar set of analysis tools, using novel flavour-tagging algorithms for different kinematic regimes and machine learning techniques for event reconstruction and signal versus background discrimination. No clear evidence of signal has been found in these studies, and the corresponding upper limits have been set. The $H \rightarrow aa \rightarrow 4b$ search presented in this thesis represents a significant improvement in sensitivity with respect to the previous analyses using 36 fb^{-1} of Run 2 data, and is complementary to the rest of ATLAS $H \rightarrow aa$ searches in other final states. This analysis also shows an outstanding performance when compared to its CMS counterpart, thanks to the use of dedicated algorithms for the different topologies of the $4b$ final state. The main limitation to the $H \rightarrow aa \rightarrow 4b$ search comes from the statistical component of the total uncertainty. This means that increased luminosity in Run 3 will improve the sensitivity of the analysis. The $t\bar{t}a, a \rightarrow b\bar{b}$ search is the first of its kind at the LHC. This analysis is part of an ongoing ATLAS programme of $t\bar{t}a$ searches in various final states, including $a \rightarrow \mu\mu$ and $a \rightarrow \tau\tau$, none of them studied prior to Run 2. Concerning the precision of the analysis, the $t\bar{t}a, a \rightarrow b\bar{b}$ search is more susceptible to non-statistical effects, such as jet reconstruction uncertainties, flavour tagging and signal and $t\bar{t}$ +jets modelling.

Paper publications for the two analyses presented in this thesis are under review by the ATLAS Collaboration. They are foreseen to be published in a scientific journal in 2025. Possible options include *Physical Review D* (PRD) and the *European Physical Journal C* (EPJC).

In order to qualify for ATLAS authorship, a typical qualification period of one year is required, where the individual needs to contribute to the experiment. This often involves participation in operational tasks, such as data collection, software development or detector maintenance. The Level-1 trigger upgrades associated to this qualification task have been presented in Appendix A. An overview of the Run 3 upgrades to the trigger and data acquisition systems and the trigger performance during the commissioning period in 2022 has been published in the *Journal of Instrumentation* (JINST) [239].

Appendices

A Phase-I upgrades for the Level-1 trigger

The BSM physics searches presented in this thesis are performed using the ATLAS Run 2 dataset from years 2015 to 2018, and the description of the ATLAS detector and trigger system in Chapter 4 corresponds to this period in time. Nevertheless, upgrades to the Level-1 trigger system have been carried out during Long Shutdown 2 (LS2), in the years 2018 to 2022, in preparation for the next period of LHC data taking, known as Run 3 (2022 – now). This Appendix describes the Phase-I upgrades for Run 3 to the ATLAS Level-1 trigger in general, with emphasis on the Level-1 topological trigger [94, 239, 240]. Sections A.1 and A.2 summarise the Phase-I updates to the L1Calo and L1Muons subsystems. Updates to the L1Topo trigger are detailed in Section A.3. Figure A.22 shows a schematic view of the ATLAS TDAQ system for Run 3, including the Run 2 (legacy) subsystems, used for commissioning at the start of Run 3. This is to be compared with Figure 4.13 from Section 4.3, depicting the previous Run 2 setup.




Figure A.22: Schematics of the ATLAS TDAQ system in Run 3.

A.1 Level-1 calorimeter trigger

During LS2, the electronics from the LAr calorimeter have been updated so that the L1Calo trigger can receive information with finer granularity in E_T , η and ϕ . The Run 2 trigger towers, with dimension 0.1×0.1 in $\Delta\eta \times \Delta\phi$, have been replaced by SuperCells, which are sets of four or eight calorimeter cells. Each trigger tower contains ten SuperCells, as shown in Figure A.23. They can provide a minimum granularity of 0.025×0.1 in $\Delta\eta \times \Delta\phi$. In addition, the Run 2 CP and JEP modules have been replaced by the electromagnetic, jet and global feature extractors (abbreviated as eFEX, jFEX and gFEX, respectively), plus improved algorithms for particle reconstruction and identification.

The eFEX module is used to identify e/γ clusters and hadronically decaying τ -leptons in the $|\eta| < 2.5$ range. It receives the full granularity from the SuperCells, that is and an E_T resolution of 0.1 GeV. It also has access to relevant shower shape parameters for particle identification, such as the shower width in the first and second layers of the calorimeter and the energy ratio between the hadronic and electromagnetic calorimeters.

The jFEX module identifies jets, hadronically decaying τ -leptons, E_T^{miss} and $\sum E_T$ in the range $|\eta| < 4.9$. It receives information with variable granularity in $\Delta\eta \times \Delta\phi$, as shown in Figure A.24 (a). The E_T resolution from the hadronic calorimeter is 0.2 GeV, except for E_T^{miss} and $\sum E_T$, which have a precision of 0.1 GeV. Improvements with respect to Run 2 include better reconstruction of nearby jets in events with high jet multiplicity, large- R jets, hadronically-decaying τ leptons in the $|\eta| < 2.5$ region and electromagnetic objects in the forward region.

The gFEX module is used to reconstruct large- R jets, E_T^{miss} and $\sum E_T$. It has a coarser granularity than jFEX, 0.2×0.2 in $\Delta\eta \times \Delta\phi$, so that the data from the entire calorimeter can be processed by a single module, facilitating the identification of boosted objects and global observables.




Figure A.23: Trigger tower composed of ten SuperCells after the Phase-I upgrade for the LAr calorimeter [94].




Figure A.24: Granularity of the inputs to (a) jFEX and (b) gFEX in the $\Delta\eta \times \Delta\phi$ plane [240].

A.2 Level-1 muon trigger

The L1Muon trigger has undergone several updates during LS2, in both the barrel and endcap regions. In the barrel, additional RPC modules have been deployed. New Barrel Inner Small (BIS 7/8) chambers cover the transition region $1.0 < |\eta| < 1.3$ between the barrel and the endcap, providing additional muon track information in the inner layers of the MS. Additionally, in the regions where the inner layer of the muon system provides incomplete coverage due to the presence of the toroid magnets, a coincidence between the endcap inner (EI) TGC chambers and the TileCal is required in order to improve the rejection of fake muons. In the endcap, the existing TGC inner stations have been replaced by the New Small Wheel (NSW) for larger η coverage and improved resolution. The upgrades to the L1Muon trigger are illustrated in Figure A.25.




Figure A.25: Phase-I upgrades for the L1Muon trigger [241].

The MUCTPI has also been redesigned for Run 3, being able to provide full granularity information to L1Topo at the BC rate. The inputs from the MUCTPI to L1Topo include the position and p_T threshold passed by the muon candidates from the RPC and TGC, along with track quality flags, geometrical flags, and in the case of the endcap, the electric charge of the muon candidate.

A.3 Level-1 topological trigger

The L1Topo processor receives inputs from L1Calo and L1Muon in TOB format, including kinematic information (E_T , η , ϕ) and particle properties. Phase-I upgrades include new TOB formats with improved granularity and additional object information, as well as new FPGA modules, designed to deal with the new inputs.

The upgrades in L1Calo and L1Muon have translated into better precision in the determination of E_T , η and ϕ . Tables A.5 and A.6 offer a comparison between the L1Topo TOB input granularities in Run 2 and Run 3. In the case of muons, the Phase-I MUCTPI provides 15 p_T thresholds in the TGC region (compared to 6 in Run 2), thanks to the additional bandwidth. In addition, the FEX modules and the MUCTPI are now able to send quality information to the L1Topo processor, which can be used to refine the L1Topo algorithms and improve their efficiency. The following list describes the additional properties that have been added to the new Run 3 TOBs:

eTau TOB (eFEX)

- R_{had} . It is the hadronic E_T fraction, that is, the ratio between the transverse energy deposited in the hadronic calorimeter (E_T^{had}) and the total transverse energy deposited in both the EM and hadronic calorimeters ($E_T^{\text{EM}} + E_T^{\text{had}}$):

$$R_{\text{had}} = E_T^{\text{had}} / (E_T^{\text{EM}} + E_T^{\text{had}}). \quad (\text{A.1})$$

- R_{core} . It is the ratio between the transverse energy deposited in the core of the EM cluster, which is the region closer to the seed, and the total transverse energy deposited in the cluster:

$$R_{\text{core}} = E_T^{\text{core}} / E_T^{\text{cluster}}. \quad (\text{A.2})$$

jTau TOB (jFEX)

- E_T^{iso} . It corresponds to the energy within the $0.2 < R < 0.4$ ring around the $R < 0.2$ cone used to reconstruct the τ candidate.

Muon TOB (MUCTPI)

- **bw2or3** (TGC). It indicates whether there is a 2-station or a 3-station coincidence in the TGC big wheel (BW).
- **innerCoin** (TGC). It indicates whether a coincidence with inner detectors, such as the NSW, is satisfied.

- **goodMF** (TGC). It provides information on whether the quality of the magnetic field in the given RoI is good.
- **Muon charge** (TGC). It indicates the electrical charge of the muon.
- **isMoreCand** (RPC). Each RPC pad, consisting of four RoIs, can only send one muon candidate to L1Topo. The **isMoreCand** flag is set to 1 if there is more than one muon candidate in different RoIs within the same pad. This feature is used to improve the efficiency of close-by di- μ triggers at L1.

Run 2 TOBs				
Name	Description	ΔE_T [GeV]	$\Delta\eta$	$\Delta\phi$
EM	e/γ from CP	0.5	0.1	0.1
Tau	τ from CP			
Jet	Jet from JEP	1	0.2	0.2
MET	Missing transverse energy from JEP	1	–	0.1
Muon	Muon from MUCTPI	RPC p_T thresholds {4, 6, 10, 11, 20, 21}	0.2-0.4	0.1
		TGC p_T thresholds {4, 6, 10, 15, 20}		

Table A.5: L1Topo TOBs available in Run 2 and their granularity in E_T , η and ϕ for $|\eta| < 2.5$.

Run 3 TOBs				
Name	Description	ΔE_T [GeV]	$\Delta\eta$	$\Delta\phi$
eEM	e/γ from eFEX	0.1	0.025	0.1
eTau	τ from eFEX			
jEM	e/γ from jFEX	0.2	0.1	0.1
jTau	τ from jFEX			
jJet	Jet from jFEX			
jLJet	Large- R jet from jFEX			
jXE	Missing transverse energy from jFEX	0.2	–	0.05
jTE	Total transverse energy from jFEX			
gJet	Jet from gFEX	0.2	0.2	0.2
gLJet	Large- R jet from gFEX			
gXE	Missing transverse energy from gFEX	0.2	–	0.2
gTE	Total transverse energy from gFEX			
Muon	Muon from MUCTPI	RPC p_T thresholds {4, 6, 10, 11, 20, 21}	0.025	0.05
		TGC p_T thresholds {3, 4, 5, 6, 7, 8, 9, 10, 11 12, 13, 14, 15, 18, 20}		

Table A.6: L1Topo TOBs available in Run 3 and their granularity in E_T , η and ϕ for $|\eta| < 2.5$.

The L1Topo system in Run 3 consists of three modules, denoted as TOPO1, TOPO2 and TOPO3, with two FPGAs each. TOPO1 module is used to run selections based on multiplicities of L1Calo objects. These algorithms count the number of objects passing a given E_T threshold or located within a given region in η , and the resulting output bits are transmitted to the CTP to perform the trigger decision. The TOPO2 and TOPO3 modules are used for the topological selections as in the Run 2 L1Topo system. The number of TOBs that can be processed has also increased with respect to Run 2. Tables A.7 and A.8 compare the maximum number of inputs per TOB type to the L1Topo trigger in Run 2 and Run 3, respectively.

Run 2 max. TOBs per event	
EM	120
Tau	120
Jet	64
MET	1
Muon	32

Table A.7: Maximum number of TOBs received by the Run 2 L1Topo system.

Run 3 max. TOBs per event	
eEM	144
eTau	144
jEM	5
jTau	6
jJet	168
jLJet	24
jXE	7
jTE	7
gJet	6
gLJet	3
gXE	3
gTE	1
Muon	32

Table A.8: Maximum number of TOBs received by the Run 3 L1Topo system.

Improving the precision of the L1Topo inputs has also allowed to create new L1Topo decision algorithms for Run 3. For example, the jEM TOB from jFEX has been used to design a L1 trigger for $Z \rightarrow ee$ processes where one of the electrons falls in the forward calorimeter region, where EM calorimetry is not available. The measurement of the $Z \rightarrow ee$ process is of interest to determine the Z -boson forward-backward asymmetry, which is closely related to the weak mixing angle, θ_w . Upgrades in the L1Muon trigger and the MUCTPI have been used to design new L1Topo triggers for the ATLAS B -physics programme, targetting two- and three-muon signatures from B -hadron decays. This sort of physics processes typically occur at low energies, and

therefore can benefit greatly from topological algorithms based on invariant masses, angular distances, muon charge and muon quality flags.

One last highlight of the Phase-I L1Topo processor is the successful implementation of the combined τ (cTau) TOB. This TOB is designed to reduce the L1 rate of τ leptons at low energy ($E_T < 25$ GeV) by applying a tighter selection. It consists of an eTau TOB with an isolation requirement computed using information from both eFEX and jFEX. The latest implementation in L1Topo is the following (Figure A.26):

$$\text{Isolation} = \frac{E_T^{\text{iso}}(\text{jTau}) + \alpha [E_T(\text{jTau}) - E_T(\text{eTau})]}{E_T(\text{eTau})} < r, \quad (\text{A.3})$$

where α and r are fixed values.




Figure A.26: Visual representation of the trigger towers used to determine $E_T(\text{jTau})$, $E_T^{\text{iso}}(\text{jTau})$ and $E_T(\text{eTau})$. Each square corresponds to one trigger tower, some of them containing several SuperCells. HAD and EM refer to the hadronic and EM calorimeters, respectively.


B Blinded fits to data

A blinded fit to data refers to a fit where certain aspects of the data, such as its behaviour in the signal regions, are deliberately hidden (blinded) until the analysis method is fully developed and validated. This ensures that the final results are not unintentionally biased by any prior knowledge. It is usually computed over the control regions or, equivalently, over the bins of the fitted distributions where the signal to background ratio is expected to be low. A common threshold for this ratio is 5%. However, a more conservative approach with a smaller signal to background ratio can be used as well. Since this type of fit is performed in background-only regions, it also serves as a cross-check that the predicted background from the MC simulation is in agreement with the observed data.


B.1 $H \rightarrow aa \rightarrow 4b$ analysis

This section shows the results of the blinded fit to data from the $H \rightarrow aa \rightarrow 4b$ search for the 25 GeV a -boson mass hypothesis. It is performed using the Z +jets and $t\bar{t}$ +jets control regions only, where the signal content is negligible. The B-only hypothesis ($\mu = 0$) is assumed.

Figure B.27 shows the data to MC comparison per control region, before and after performing the blinded fit. A good post-fit agreement is observed between data and MC. Figure B.28 contains the normalisation factors for the different components of Z +jets and $t\bar{t}$ +jets obtained from the fit, which are all compatible with unity. Figure B.29 shows the correlations between NPs. Figure B.30 shows the relevant nuisance parameters. There is a similar pattern to the one observed in the Asimov fit. In both the correlation matrix and the NP plot, the most significant correlations and constraints originate from flavour-tagging and background modelling. Because the blinded fit is performed with real data, NP pulls are observed in the NP plot. The strongest pulls appear in some of the Z +jets and $t\bar{t}$ +jets modelling uncertainties derived from alternative MC samples, however they are well within the $\pm 1\sigma$ range.



(a)



(b)

Figure B.27: Predicted yields for the signal and control regions in the blinded fit to data with $m_a = 25$ GeV (a) before and (b) after the fit. Data points are drawn only in the signal-depleted control regions that are used in the fit. The expected signal yield for the 25 GeV mass hypothesis assuming the SM production cross section for ZH and $\text{BR}(H \rightarrow aa \rightarrow 4b) = 1$ is overlayed in the pre-fit plot.




Figure B.28: Normalisation factors for the Z +jets and $t\bar{t}$ +jets samples in the blinded fit to data with $m_a = 25$ GeV.




Figure B.29: Correlation matrix for the relevant NPs in the blinded fit to data with $m_a = 25$ GeV.




Figure B.30: Fitted nuisance parameters from the blinded fit to data with $m_a = 25$ GeV.

B.2 $t\bar{t}a, a \rightarrow b\bar{b}$ analysis

This section shows the results of the blinded fit to data from the $t\bar{t}a, a \rightarrow b\bar{b}$ search for the 30 GeV a -boson mass hypothesis. It is performed using the 0B2b+1bL control region and the bins in the signal regions where the signal to background ratio is below 3% (SR control bins). The background-only hypothesis ($\mu = 0$) is assumed.

Figures B.31 and B.32 summarise the data to MC comparison per control region, before and after performing the blinded fit to data. Figure B.33 contains the normalisation factors for the different components of $t\bar{t}$ +jets, and Figure B.34 shows the correlations between NPs. Figures B.35 and B.36 show the relevant nuisance parameters. The output from the blinded fit shows a good post-fit agreement between data and MC in the SR control bins and the 0B2b+1bL control region. The normalisation factors for the $t\bar{t}$ +jets background are within the expected range and compatible with unity. Results also show a similar pattern in both the correlation matrix and the NP plot, where the most significant correlations and constraints originate from flavour-tagging and background modelling. Because the blinded fit is performed with real data, NP pulls are observed in the NP plot. The strongest pulls appear in the DeXTer uncertainties and in some of the $t\bar{t}$ +jets alternative samples, but they are well within the $\pm 1\sigma$ range.





Figure B.31: Pre-fit plots from the blinded fit to data with $m_a = 30$ GeV.


 Figure B.32: Post-fit plots from the blinded fit to data with $m_a = 30$ GeV.




Figure B.33: Normalisation factors for the $t\bar{t}$ +jets samples in the blinded fit to data with $m_a = 30$ GeV.





Figure B.34: Correlation matrix for the relevant NPs in the blinded fit to data with $m_a = 30$ GeV.

Figure B.35: Fitted nuisance parameters from the blinded fit to data with $m_a = 30$ GeV (I).




Figure B.36: Fitted nuisance parameters from the blinded fit to data with $m_a = 30$ GeV (II).

C Injection tests for the $t\bar{t}a$, $a \rightarrow b\bar{b}$ analysis

An injection test is a S+B Asimov fit used to validate the statistical analysis framework and to study the sensitivity of a search in the presence of a BSM signal. It involves injecting a signal ($\mu > 0$) into the Asimov dataset and evaluating the analysis' ability to detect and characterise that signal. The results of the injection tests are compared with the B-only Asimov fit ($\mu = 0$), in order to ensure that the nuisance parameters are well-behaved in the S+B hypothesis, and to evaluate the fit response to different hypothetical signal cross sections. Several injection tests are performed for the a -boson mass hypotheses of 12, 30 and 60 GeV. For each mass, signal strengths of $\mu = 0.5$, 1 and 2 are tested. All signal samples are normalised to a cross section of 32.058 fb, to which the signal strength is a multiplicative factor.

Figure C.37 shows the expected limits from the various fits to the Asimov dataset with $\mu = 0$ (B-only hypothesis) and $\mu > 0$ (S+B hypothesis). They provide an estimation of the fit response under different signal strength hypotheses. These results show that the analysis is more sensitive to high values of the a -boson mass, such as $m_a = 60$ GeV. For low a -boson masses, the search is limited by the DeXTer B -tagging systematic uncertainties in addition to the general $t\bar{t}$ +jets modelling uncertainties. Therefore, a low-mass signal will have less significance than a high-mass signal at the same cross section. Figure C.38 shows the fitted signal strength, $\hat{\mu}$, with respect to the a -boson mass for each injection test. A value of $\hat{\mu}$ close to the injected signal strength, μ , indicates that the fit is able to identify the presence of a BSM signal in the provided dataset. In general, a positive response is observed for the injected a -boson mass hypothesis, but also for the adjacent mass points, due to the limited mass resolution in the $a \rightarrow b\bar{b}$ decay channel. The nuisance parameters with the largest pulls for each injection test are represented in Figure C.39. The comparison shows a similar degree of constraint for different values of the injected signal strength. The pulls in all nuisance parameters are negligible. This check allows to ensure that the fit is not trying to replicate the injected BSM signal by pulling the systematics instead of adjusting the fitted signal strength.




Figure C.37: Expected limits from the (a) $m_a = 12$ GeV, (b) $m_a = 30$ GeV and (c) $m_a = 60$ GeV injection tests. They are computed for $\mu = 0$ (dashed black line), $\mu = 0.5$ (solid red line), $\mu = 1$ (solid purple line) and $\mu = 2$ (solid blue line).





Figure C.38: Fitted signal strength with respect to the a -boson mass from the (a) $m_a = 12$ GeV, (b) $m_a = 30$ GeV and (c) $m_a = 60$ GeV injection tests. The different colours correspond to $\mu = 0.5$ (red), $\mu = 1$ (purple) and $\mu = 2$ (blue).

Ph. D. Thesis P. Martinez
Injection test 12 GeV


— $\mu = 0.5$
— $\mu = 1$
— $\mu = 2$



(a)

Ph. D. Thesis P. Martinez
Injection test 30 GeV

— $\mu = 0.5$
— $\mu = 1$
— $\mu = 2$



(b)




Figure C.39: Relevant nuisance parameters from the (a) $m_a = 12$ GeV, (b) $m_a = 30$ GeV and (c) $m_a = 60$ GeV injection tests. The different colours correspond to $\mu = 0.5$ (red), $\mu = 1$ (black) and $\mu = 2$ (blue).

Bibliography

- [1] ATLAS Collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, Phys. Lett. B vol. 716 (2012), arXiv: 1207.7214.
- [2] CMS Collaboration, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, Phys. Lett. B vol. 716 (2012), arXiv: 1207.7235.
- [3] ATLAS Collaboration, *A detailed map of Higgs boson interactions by the ATLAS experiment ten years after the discovery*, Nature vol. 607, no. 7917 (2022), arXiv: 2207.00092.
- [4] CMS Collaboration, *A portrait of the Higgs boson by the CMS experiment ten years after the discovery*, Nature vol. 607, no. 7917 (2022), arXiv: 2207.00043.
- [5] B. A. Dobrescu and K. T. Matchev, *Light axion within the next-to-minimal supersymmetric standard model*, JHEP vol. 09 (2000), arXiv: hep-ph/0008192.
- [6] U. Ellwanger et al., *Towards a no lose theorem for NMSSM Higgs discovery at the LHC* (2003), arXiv: hep-ph/0305109.
- [7] R. Dermisek and J. F. Gunion, *Escaping the large fine tuning and little hierarchy problems in the next to minimal supersymmetric model and $h \rightarrow aa$ decays*, Phys. Rev. Lett. vol. 95 (2005), arXiv: hep-ph/0502105.
- [8] S. Chang et al., *Nonstandard Higgs Boson Decays*, Ann. Rev. Nucl. Part. Sci. vol. 58 (2008), arXiv: 0801.4554.
- [9] D. E. Morrissey and A. Pierce, *Modified Higgs Boson Phenomenology from Gauge or Gaugino Mediation in the NMSSM*, Phys. Rev. D vol. 78 (2008), arXiv: 0807.2259.
- [10] V. Silveira and A. Zee, *Scalar Phantoms*, Phys. Lett. B vol. 161 (1985).
- [11] M. Pospelov, A. Ritz and M. B. Voloshin, *Secluded WIMP Dark Matter*, Phys. Lett. B vol. 662 (2008), arXiv: 0711.4866.
- [12] P. Draper et al., *Dark Light Higgs*, Phys. Rev. Lett. vol. 106 (2011), arXiv: 1009.3963.
- [13] S. Ipek, D. McKeen and A. E. Nelson, *A Renormalizable Model for the Galactic Center Gamma Ray Excess from Dark Matter Annihilation*, Phys. Rev. D vol. 90, no. 5 (2014), arXiv: 1404.3716.
- [14] A. Martin, J. Shelton and J. Unwin, *Fitting the Galactic Center Gamma-Ray Excess with Cascade Annihilations*, Phys. Rev. D vol. 90, no. 10 (2014), arXiv: 1405.0272.
- [15] J. Kozaczuk, M. J. Ramsey-Musolf and J. Shelton, *Exotic Higgs boson decays and the electroweak phase transition*, Phys. Rev. D vol. 101, no. 11 (2022), arXiv: 1911.10210.

- [16] M. Carena, Z. Liu and Y. Wang, *Electroweak phase transition with spontaneous Z_2 -breaking*, JHEP vol. 08 (2020), arXiv: 1911.10206.
- [17] N. Craig et al., *Naturalness in the Dark at the LHC*, JHEP vol. 07 (2015), arXiv: 1501.05310.
- [18] D. Curtin and C. B. Verhaaren, *Discovering Uncolored Naturalness in Exotic Higgs Decays*, JHEP vol. 12 (2015), arXiv: 1506.06141.
- [19] E. Fermi, *Tentativo di una Teoria Dei Raggi β* , Nuovo Cim. vol. 11 (1934).
- [20] C. S. Wu et al., *Experimental Test of Parity Conservation in Beta Decay*, Phys. Rev. vol. 105 4 (1957).
- [21] J. Bardeen, L. N. Cooper and J. R. Schrieffer, *Microscopic Theory of Superconductivity*, Phys. Rev. vol. 106 1 (1957).
- [22] J. Bardeen, L. N. Cooper and J. R. Schrieffer, *Theory of Superconductivity*, Phys. Rev. vol. 108 5 (1957).
- [23] Y. Nambu, *Quasi-Particles and Gauge Invariance in the Theory of Superconductivity*, Phys. Rev. vol. 117 3 (1960).
- [24] J. Goldstone, *Field theories with «Superconductor» solutions*, Nuovo Cim. vol. 19 (1961).
- [25] J. Goldstone, A. Salam and S. and Weinberg, *Broken Symmetries*, Phys. Rev. vol. 127 3 (1962).
- [26] N. Cabibbo, *Unitary Symmetry and Leptonic Decays*, Phys. Rev. Lett. vol. 10 (1963).
- [27] M. Kobayashi and T. Maskawa, *CP Violation in the Renormalizable Theory of Weak Interaction*, Prog. Theor. Phys. vol. 49 (1973).
- [28] M. Gell-Mann, *A Schematic Model of Baryons and Mesons*, Phys. Lett. vol. 8 (1964).
- [29] G. Zweig, *An $SU(3)$ model for strong interaction symmetry and its breaking. Version 1* (1964). CERN-TH-401.
- [30] E. D. Bloom et al., *High-Energy Inelastic $e - p$ Scattering at 6° and 10°* , Phys. Rev. Lett. vol. 23 (1969).
- [31] J. D. Bjorken and S. L. Glashow, *Elementary Particles and $SU(4)$* , Phys. Lett. vol. 11 (1964).
- [32] S. L. Glashow, J. Iliopoulos and L. Maiani, *Weak Interactions with Lepton-Hadron Symmetry*, Phys. Rev. D vol. 2 (1970).
- [33] E598 Collaboration, *Experimental Observation of a Heavy Particle J*, Phys. Rev. Lett. vol. 33 (1974).
- [34] SLAC-SP-017 Collaboration, *Discovery of a Narrow Resonance in e^+e^- Annihilation*, Phys. Rev. Lett. vol. 33 (1974).
- [35] E288 Collaboration, *Observation of a Dimuon Resonance at 9.5 GeV in 400 GeV Proton-Nucleus Collisions*, Phys. Rev. Lett. vol. 39 (1977).

- [36] CDF Collaboration, *Observation of top quark production in $\bar{p}p$ collisions*, Phys. Rev. Lett. vol. 74 (1995), arXiv: hep-ex/9503002.
- [37] D0 Collaboration, *Observation of the top quark*, Phys. Rev. Lett. vol. 74 (1995), arXiv: hep-ex/9503003.
- [38] KTeV Collaboration, *Observation of direct CP violation in $K_{S,L} \rightarrow \pi\pi$ decays*, Phys. Rev. Lett. vol. 83 (1999), arXiv: hep-ex/9905060.
- [39] NA48 Collaboration, *A New measurement of direct CP violation in two pion decays of the neutral kaon*, Phys. Lett. B vol. 465 (1999), arXiv: hep-ex/9909022.
- [40] TASSO Collaboration, *Evidence for Planar Events in e^+e^- Annihilation at High-Energies*, Phys. Lett. B vol. 86 (1979).
- [41] MARK J Collaboration, *Discovery of Three Jet Events and a Test of Quantum Chromodynamics at PETRA Energies*, Phys. Rev. Lett. vol. 43 (1979).
- [42] PLUTO Collaboration, *Evidence for Gluon Bremsstrahlung in e^+e^- Annihilations at High-Energies*, Phys. Lett. B vol. 86 (1979).
- [43] Gargamelle Neutrino Collaboration, *Observation of Neutrino Like Interactions without Muon or Electron in the Gargamelle Neutrino Experiment*, Nucl. Phys. B vol. 73 (1974).
- [44] SLAC-SP-017 Collaboration, *Evidence for Anomalous Lepton Production in e^+e^- Annihilation*, Phys. Rev. Lett. vol. 35 (1975).
- [45] UA1 Collaboration, *Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540$ GeV*, Phys. Lett. B vol. 122 (1983).
- [46] UA2 Collaboration, *Evidence for $Z^0 \rightarrow e^+e^-$ at the CERN $\bar{p}p$ Collider*, Phys. Lett. B vol. 129 (1983).
- [47] ALEPH, DELPHI, L3, OPAL and SLD Collaborations, LEP Electroweak Working Group, SLD Electroweak and Heavy Flavour Groups, *Precision electroweak measurements on the Z resonance*, Phys. Rept. vol. 427 (2006), arXiv: hep-ex/0509008.
- [48] DONUT Collaboration, *Observation of tau neutrino interactions*, Phys. Lett. B vol. 504 (2001), arXiv: hep-ex/0012035.
- [49] LHC Higgs Cross Section Working Group, *Handbook of LHC Higgs Cross Sections: 4. Deciphering the Nature of the Higgs Sector*, vol. 2/2017 (2016), arXiv: 1610.07922.
- [50] ATLAS Collaboration, *Combined Measurement of the Higgs Boson Mass from the $H \rightarrow \gamma\gamma$ and $H \rightarrow ZZ^* \rightarrow 4\ell$ Decay Channels with the ATLAS Detector Using $\sqrt{s}=7, 8$, and 13 TeV pp Collision Data*, Phys. Rev. Lett. vol. 131, no. 25 (2023), arXiv: 2308.04775.
- [51] CMS Collaboration, *A measurement of the Higgs boson mass in the diphoton decay channel* (2019). CMS-PAS-HIG-19-004.

- [52] ATLAS Collaboration, *Standard Model Summary Plots June 2024* (2024). ATL-PHYS-PUB-2024-011.
- [53] B. Pontecorvo, *Inverse beta processes and nonconservation of lepton charge*, Zh. Eksp. Teor. Fiz. vol. 34 (1957).
- [54] Z. Maki, M. Nakagawa and S. Sakata, *Remarks on the unified model of elementary particles*, Prog. Theor. Phys. vol. 28 (1962).
- [55] Muon g-2 Collaboration, *Measurement of the Positive Muon Anomalous Magnetic Moment to 0.20 ppm*, Phys. Rev. Lett. vol. 131, no. 16 (2023), arXiv: 2308.06230.
- [56] V. C. Rubin, N. Thonnard and W. K. Ford Jr., *Rotational properties of 21 SC galaxies with a large range of luminosities and radii, from NGC 4605 /R = 4kpc/ to UGC 2885 /R = 122 kpc/*, Astrophys. J. vol. 238 (1980).
- [57] M. Markevitch et al., *Direct constraints on the dark matter self-interaction cross-section from the merging galaxy cluster 1E0657-56*, Astrophys. J. vol. 606 (2004), arXiv: astro-ph/0309303.
- [58] J. Bland-Hawthorn and P. J. E. Peebles, *Near-Field Cosmology*, Science vol. 313, no. 5785 (2006).
- [59] KM3Net Collaboration, *Letter of intent for KM3NeT 2.0*, J. Phys. G vol. 43, no. 8 (2016), arXiv: 1601.07459.
- [60] A. Boccaletti et al., *High precision calculation of the hadronic vacuum polarisation contribution to the muon anomaly* (2024), arXiv: 2407.10913.
- [61] Peccei, R. D. and Quinn, H. R., *CP Conservation in the Presence of Pseudoparticles*, Phys. Rev. Lett. vol. 38 25 (1977).
- [62] Peccei, R. D. and Quinn, H. R., *Constraints imposed by CP conservation in the presence of pseudoparticles*, Phys. Rev. D vol. 16 6 (1977).
- [63] M. Bauer, M. Neubert and A. Thamm, *Collider Probes of Axion-Like Particles*, JHEP vol. 12 (2017), arXiv: 1708.00443.
- [64] G. C. Branco et al., *Theory and phenomenology of two-Higgs-doublet models*, Phys. Rept. vol. 516 (2012), arXiv: 1106.0034.
- [65] M. R. Buckley, D. Feld and D. Goncalves, *Scalar Simplified Models for Dark Matter*, Phys. Rev. D vol. 91 (2015), arXiv: 1410.6497.
- [66] M. Bauer, U. Haisch and F. Kahlhoefer, *Simplified dark matter models with two Higgs doublets: I. Pseudoscalar mediators*, JHEP vol. 05 (2017), arXiv: 1701.07427.
- [67] M. Casolino et al., *Probing a light CP-odd scalar in di-top-associated production at the LHC*, Eur. Phys. J. C vol. 75 (2015), arXiv: 1507.07004.
- [68] LHC New Physics Working Group, *Simplified Models for LHC New Physics Searches*, J. Phys. G vol. 39 (2012), arXiv: 1105.2838.
- [69] E. A. Mobs, *The CERN accelerator complex*, <https://cds.cern.ch/record/2225847>. Accessed: 18/03/2024.

-
- [70] AC Team, *Diagram of an LHC dipole magnet*, <https://cds.cern.ch/record/40524>. Accessed: 28/06/2024.
- [71] J. L. Caron, *LHC quadrupole cross section*, <https://cds.cern.ch/record/841485>. Accessed: 01/07/2024.
- [72] F. Sauerburger, *LHC cross section plot*, <https://lhc-xsecs.org/%5C#/>. Accessed: 28/06/2024.
- [73] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s}=13$ TeV using the ATLAS detector at the LHC*, Eur. Phys. J. C vol. 83, no. 10 (2023), arXiv: 2212.09379.
- [74] ATLAS Collaboration, *Public ATLAS Luminosity Results for Run-2 of the LHC*, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/LuminosityPublicResultsRun2>. Accessed: 28/06/2024.
- [75] J. Pequeno, *Computer generated image of the whole ATLAS detector*, <https://cds.cern.ch/record/1095924>. Accessed: 08/07/2024.
- [76] I. Neutelings (2017), https://tikz.net/axis3d_cms/. Accessed: 05/01/2025.
- [77] ATLAS Collaboration, *ATLAS inner detector: Technical design report. Vol. 1* (1997). CERN-LHCC-97-16.
- [78] ATLAS Collaboration, *ATLAS inner detector: Technical design report. Vol. 2* (1997). CERN-LHCC-97-17.
- [79] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, JINST vol. 3 (2008).
- [80] The ATLAS IBL Collaboration, *Production and Integration of the ATLAS Insertable B-Layer*, JINST vol. 13, no. 05 (2018), arXiv: 1803.00844.
- [81] B. Mandelli, *The Pixel Detector of the ATLAS Experiment for the Run 2 at the Large Hadron Collider*, Nucl. Part. Phys. Proc. vol. 273-275 (2016).
- [82] ATLAS Collaboration, *Operation and performance of the ATLAS semiconductor tracker in LHC Run 2*, JINST vol. 17, no. 01 (2022), arXiv: 2109.02591.
- [83] V. A. Mitsou, *The ATLAS transition radiation tracker*, 8th International Conference on Advanced Technology and Particle Physics (ICATPP 2003): Astroparticle, Particle, Space Physics, Detectors and Medical Physics Applications (2003), arXiv: hep-ex/0311058.
- [84] ATLAS Collaboration, *ATLAS liquid argon calorimeter: Technical design report* (1996). CERN-LHCC-96-41.
- [85] M. Spalla, *ATLAS LAr Calorimeter performance in LHC Run 2*, JINST vol. 15, no. 04 (2020).
- [86] ATLAS Collaboration, *ATLAS tile calorimeter: Technical design report* (1996). CERN-LHCC-96-42.
- [87] ATLAS Collaboration, *Operation and performance of the ATLAS tile calorimeter in LHC Run 2* (2024), arXiv: 2401.16034.

- [88] R. Leitner, P. Tas and V. V. Shmakova, *Time resolution of the ATLAS Tile calorimeter and its performance for a measurement of heavy stable particles* (2007). ATL-TILECAL-PUB-2007-002, ATL-COM-TILECAL-2007-008.
- [89] L. Heelan, *Performance of the ATLAS Tile Calorimeter*, J. Phys. Conf. Ser. vol. 623, no. 1 (2015).
- [90] W. Lampl, *Optimizing the energy measurement of the ATLAS electromagnetic calorimeter*, Ph.D. Thesis, Vienna Tech. University (2006). CERN-THESIS-2006-002.
- [91] ATLAS Collaboration, *Performance of the ATLAS muon triggers in Run 2*, JINST vol. 15, no. 09 (2020), arXiv: 2004.13447.
- [92] ATLAS Collaboration, *Public ATLAS DAQ plots*, <https://twiki.cern.ch/twiki/bin/view/AtlasPublic/ApprovedPlotsDAQ>. Accessed: 06/08/2024.
- [93] ATLAS Collaboration, *ATLAS level-1 trigger: Technical Design Report* (1998). ATLAS-TDR-12, CERN-LHCC-98-014, CERN-LHCC-98-14, ATLAS-TDR-12.
- [94] ATLAS Collaboration, *Technical Design Report for the Phase-I Upgrade of the ATLAS TDAQ System* (2013). CERN-LHCC-2013-018, ATLAS-TDR-023.
- [95] S. Artz et al., *The ATLAS Level-1 Muon Topological Trigger Information for Run 2 of the LHC*, JINST vol. 10, no. 02 (2015).
- [96] ATLAS Collaboration, *Performance of the ATLAS Level-1 topological trigger in Run 2*, Eur. Phys. J. C vol. 82, no. 1 (2022), arXiv: 2105.01416.
- [97] S. Ask et al., *The ATLAS central level-1 trigger logic and TTC system*, JINST vol. 3 (2008).
- [98] ATLAS Collaboration, *ATLAS high-level trigger, data acquisition and controls: Technical design report* (2003). CERN-LHCC-2003-022, ATLAS-TRD-016.
- [99] ATLAS Collaboration, *Operation of the ATLAS trigger system in Run 2*, JINST vol. 15, no. 10 (2020), arXiv: 2007.12539.
- [100] M. Elsing et al., *The ATLAS Tier-0: Overview and operational experience*, J. Phys. Conf. Ser. vol. 219 (2010).
- [101] I. Georgescu, *Simulating parton showers*, Nat. Rev. Phys. vol. 3 (2021).
- [102] A. Buckley et al., *General-purpose event generators for LHC physics*, Phys. Rept. vol. 504 (2011), arXiv: 1101.2599.
- [103] R. K. Ellis, W. J. Stirling and B. R. Webber, *QCD and collider physics*, Cambridge monographs on particle physics, nuclear physics, and cosmology, Cambridge: Cambridge University Press, 2003.
- [104] 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 vol. 46 (1977).
- [105] V. N. Gribov and L. N. Lipatov, *Deep inelastic $e p$ scattering in perturbation theory*, Sov. J. Nucl. Phys. vol. 15 (1972).

- [106] G. Altarelli and G. Parisi, *Asymptotic Freedom in Parton Language*, Nucl. Phys. B vol. 126 (1977).
- [107] G. P. Salam, *Elements of QCD for hadron colliders*, CERN Yellow Report vol. CERN-2010-002 (2010).
- [108] J. Butterworth et al., *PDF4LHC recommendations for LHC Run II*, J. Phys. G vol. 43 (2016), arXiv: 1510.03865.
- [109] L. del Debbio, *Parton distributions in the LHC era*, EPJ Web Conf. vol. 175 (2018).
- [110] P.Z. Skands, *QCD for Collider Physics*, 2010 European School of High Energy Physics (2011), arXiv: 1104.2863.
- [111] F. James, *Monte Carlo Theory and Practice*, Rept. Prog. Phys. vol. 43 (1980).
- [112] T. Kinoshita, *Mass singularities of Feynman amplitudes*, J. Math. Phys. vol. 3 (1962).
- [113] T. D. Lee and M. Nauenberg, *Degenerate Systems and Mass Singularities*, Phys. Rev. vol. 133 (1964), ed. by Feinberg, G.
- [114] S. Catani and M. H. Seymour, *A General algorithm for calculating jet cross-sections in NLO QCD*, Nucl. Phys. B vol. 485 (1997), arXiv: hep-ph/9605323.
- [115] G. C. Fox and S. Wolfram, *A Model for Parton Showers in QCD*, Nucl. Phys. B vol. 168 (1980).
- [116] S. Catani et al., *QCD matrix elements + parton showers*, JHEP vol. 11 (2001), arXiv: hep-ph/0109231.
- [117] M. L. Mangano et al., *Matching matrix elements and shower evolution for top-quark production in hadronic collisions*, JHEP vol. 01 (2007), arXiv: hep-ph/0611129.
- [118] S. Catani et al., *New clustering algorithm for multi - jet cross-sections in e^+e^- annihilation*, Phys. Lett. B vol. 269 (1991).
- [119] B. R. Webber, Scholarpedia, *Parton shower Monte Carlo event generators*, http://www.scholarpedia.org/article/Parton_shower_Monte_Carlo_event_generators. Accessed: 26/09/2024.
- [120] B. Andersson et al., *Parton Fragmentation and String Dynamics*, Phys. Rept. vol. 97 (1983).
- [121] B. R. Webber, *A QCD Model for Jet Fragmentation Including Soft Gluon Interference*, Nucl. Phys. B vol. 238 (1984).
- [122] Particle Data Group, *Review of particle physics*, Phys. Rev. D vol. 110, no. 3 (2024).
- [123] ATLAS Collaboration, *Measurement of charged-particle distributions sensitive to the underlying event in $\sqrt{s} = 13$ TeV proton-proton collisions with the ATLAS detector at the LHC*, JHEP vol. 03 (2017), arXiv: 1701.05390.
- [124] ATLAS Collaboration, *The ATLAS Simulation Infrastructure*, Eur. Phys. J. C vol. 70 (2010), arXiv: 1005.4568.

- [125] ATLAS Collaboration, *ATLAS computing: Technical design report* (2005). CERN-LHCC-2005-022, ATLAS-TRD-017.
- [126] GEANT4 Collaboration, *GEANT4 – a simulation toolkit*, Nucl. Instrum. Meth. A vol. 506 (2003).
- [127] M. Dobbs and J. B. Hansen, *The HepMC C++ Monte Carlo event record for High Energy Physics*, Comput. Phys. Commun. vol. 134 (2001).
- [128] T. Yamanaka, *The ATLAS calorimeter simulation FastCaloSim*, J. Phys. Conf. Ser. vol. 331 (2011).
- [129] ATLAS Collaboration, *Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2*, Eur. Phys. J. C vol. 77, no. 10 (2017), arXiv: 1704.07983.
- [130] R. Fruhwirth, *Application of Kalman filtering to track and vertex fitting*, Nucl. Instrum. Meth. A vol. 262 (1987).
- [131] A. Vogel, *ATLAS Transition Radiation Tracker (TRT): Straw tube gaseous detectors at high rates*, Nucl. Instrum. Meth. A vol. 732 (2013).
- [132] F. Meloni, *Primary vertex reconstruction with the ATLAS detector*, JINST vol. 11, no. 12 (2016).
- [133] ATLAS Collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton-proton collision data*, JINST vol. 14, no. 12 (2019), arXiv: 1908.00005.
- [134] 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 vol. 81, no. 7 (2021), arXiv: 2012.00578.
- [135] ATLAS Collaboration, *Jet reconstruction and performance using particle flow with the ATLAS Detector*, Eur. Phys. J. C vol. 77, no. 7 (2017), arXiv: 1703.10485.
- [136] R. Atkin, *Review of jet reconstruction algorithms*, J. Phys. Conf. Ser. vol. 645, no. 1 (2015).
- [137] S. D. Ellis and D. E. Soper, *Successive combination jet algorithm for hadron collisions*, Phys. Rev. D vol. 48 (1993), arXiv: hep-ph/9305266.
- [138] Y. L. Dokshitzer et al., *Better jet clustering algorithms*, JHEP vol. 08 (1997), arXiv: hep-ph/9707323.
- [139] M. Wobisch and T. Wengler, *Hadronization corrections to jet cross-sections in deep inelastic scattering* (1998), arXiv: hep-ph/9907280.
- [140] M. Cacciari, G. P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, JHEP vol. 04 (2008), arXiv: 0802.1189.
- [141] 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 vol. 81, no. 8 (2021), arXiv: 2007.02645.
- [142] ATLAS Collaboration, *Identification and rejection of pile-up jets at high pseudorapidity with the ATLAS detector*, Eur. Phys. J. C vol. 77, no. 9 (2017), arXiv: 1705.02211.

- [143] ATLAS Collaboration, *Deep Sets based Neural Networks for Impact Parameter Flavour Tagging in ATLAS* (2020). ATL-PHYS-PUB-2020-014.
- [144] M. Cacciari and G. P. Salam, *Pileup subtraction using jet areas*, Phys. Lett. B vol. 659 (2008), arXiv: 0707.1378.
- [145] M. Cacciari, G. P. Salam and G. Soyez, *The Catchment Area of Jets*, JHEP vol. 04 (2008), arXiv: 0802.1188.
- [146] 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 vol. 09 (2013), arXiv: 1306.4945.
- [147] ATLAS Collaboration, *Variable Radius, Exclusive- k_T , and Center-of-Mass Subject Reconstruction for Higgs($\rightarrow b\bar{b}$) Tagging in ATLAS* (2017). ATL-PHYS-PUB-2017-010.
- [148] ATLAS Collaboration, *Soft b -hadron tagging for compressed SUSY scenarios* (2019). ATLAS-CONF-2019-027.
- [149] V. Kostyukhin, *Secondary vertex based b -tagging* (2003). ATL-PHYS-2003-033, ATL-COM-PHYS-2003-032, CERN-ATL-PHYS-2003-033.
- [150] ATLAS Collaboration, *ATLAS flavour-tagging algorithms for the LHC Run 2 pp collision dataset*, Eur. Phys. J. C vol. 83, no. 7 (2023), arXiv: 2211.16345.
- [151] ATLAS Collaboration, *Monte Carlo to Monte Carlo scale factors for flavour tagging efficiency calibration* (2020). ATL-PHYS-PUB-2020-009.
- [152] ATLAS Collaboration, *Expected performance of the 2019 ATLAS b -taggers*, <https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/FTAG-2019-005/>. Accessed: 10/08/2024.
- [153] ATLAS Collaboration, *DeXTer: Deep Sets based Neural Networks for Low- p_T $X \rightarrow b\bar{b}$ Identification in ATLAS* (2022). ATL-PHYS-PUB-2022-042.
- [154] 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 vol. 78, no. 11 (2018), arXiv: 1802.08168.
- [155] W. S. McCulloch and W. Pitts, *A logical calculus of the ideas immanent in nervous activity*, Bull. Math. Biol. vol. 5, no. 4 (1943).
- [156] D. E. Rumelhart, G. E. Hinton and R. J. Williams, *Learning representations by back-propagating errors*, Nature vol. 323, no. 6088 (1986).
- [157] D. P. Kingma and J. Ba, *Adam: A Method for Stochastic Optimization* (2014), arXiv: 1412.6980.
- [158] P. Baldi et al., *Parameterized neural networks for high-energy physics*, Eur. Phys. J. C vol. 76, no. 5 (2016), arXiv: 1601.07913.
- [159] Y. Freund and R. E. Schapire, *A Decision-Theoretic Generalization of On-Line Learning and an Application to Boosting*, J. Comput. Syst. Sci. vol. 55, no. 1 (1997).
- [160] T. Chen and C. Guestrin, *XGBoost: A Scalable Tree Boosting System* (2016), arXiv: 1603.02754.

- [161] G. Cowan et al., *Asymptotic formulae for likelihood-based tests of new physics*, Eur. Phys. J. C vol. 71 (2011), arXiv: 1007.1727.
- [162] T. Junk, *Confidence level computation for combining searches with small statistics*, Nucl. Instrum. Meth. A vol. 434 (1999), arXiv: hep-ex/9902006.
- [163] ATLAS Collaboration, *TRExFitter repository*, <https://gitlab.cern.ch/TRExStats/TRExFitter>. Accessed: 05/11/2024.
- [164] R. Brun and F. Rademakers, *ROOT: An object oriented data analysis framework*, Nucl. Instrum. Meth. A vol. 389 (1997).
- [165] K. Cranmer et al., *HistFactory: A tool for creating statistical models for use with RooFit and RooStats* (2012). CERN-OPEN-2012-016.
- [166] W. Verkerke and D. P. Kirkby, *The RooFit toolkit for data modeling*, eConf vol. C0303241 (2003), arXiv: physics/0306116.
- [167] L. Moneta et al., *The RooStats Project*, PoS vol. ACAT2010 (2010), arXiv: 1009.1003.
- [168] S. S. Wilks, *The Large-Sample Distribution of the Likelihood Ratio for Testing Composite Hypotheses*, Annals Math. Statist. vol. 9, no. 1 (1938).
- [169] ATLAS Physics Exotics Group, *CommonSystSmoothingTool documentation*, https://gitlab.cern.ch/atlas-phys/exot/CommonSystSmoothingTool/-/blob/master/doc/algorithm_description.md. Accessed: 05/11/2024.
- [170] D. Curtin et al., *Exotic decays of the 125 GeV Higgs boson*, Phys. Rev. D vol. 90, no. 7 (2014), arXiv: 1312.4992.
- [171] ATLAS Collaboration, *Search for the Higgs boson produced in association with a W boson and decaying to four b-quarks via two spin-zero particles in pp collisions at 13 TeV with the ATLAS detector*, Eur. Phys. J. C vol. 76, no. 11 (2016), arXiv: 1606.08391.
- [172] ATLAS Collaboration, *Search for the Higgs boson produced in association with a vector boson and decaying into two spin-zero particles in the $H \rightarrow aa \rightarrow 4b$ channel in pp collisions at $\sqrt{s}=13$ TeV with the ATLAS detector*, JHEP vol. 10 (2018), arXiv: 1806.07355.
- [173] ATLAS Collaboration, *Search for Higgs boson decays into two new low-mass spin-0 particles in the $4b$ channel with the ATLAS detector using pp collisions at $\sqrt{s} = 13$ TeV*, Phys. Rev. D vol. 102, no. 11 (2020), arXiv: 2005.12236.
- [174] CMS Collaboration, *Search for the decay of the Higgs boson to a pair of light pseudoscalar bosons in the final state with four bottom quarks in proton-proton collisions at $\sqrt{s} = 13$ TeV*, JHEP vol. 06 (2024), arXiv: 2403.10341.
- [175] ATLAS Collaboration, *Search for Higgs boson decays to beyond-the-Standard-Model light bosons in four-lepton events with the ATLAS detector at $\sqrt{s}=13$ TeV*, JHEP vol. 06 (2018), arXiv: 1802.03388.
- [176] ATLAS Collaboration, *Search for Higgs bosons decaying into new spin-0 or spin-1 particles in four-lepton final states with the ATLAS detector with 139 fb^{-1} of pp collision data at $\sqrt{s}=13$ TeV*, JHEP vol. 03 (2022), arXiv: 2110.13673.

- [177] ATLAS Collaboration, *Search for Higgs bosons decaying to aa in the $\mu\mu\tau\tau$ final state in pp collisions at $\sqrt{s}=8$ TeV with the ATLAS experiment*, Phys. Rev. D vol. 92, no. 5 (2015), arXiv: 1505.01609.
- [178] CMS Collaboration, *Search for light bosons in decays of the 125 GeV Higgs boson in proton-proton collisions at $\sqrt{s} = 8$ TeV*, JHEP vol. 10 (2017), arXiv: 1701.02032.
- [179] CMS Collaboration, *Search for an exotic decay of the Higgs boson to a pair of light pseudoscalars in the final state of two muons and two τ leptons in proton-proton collisions at $\sqrt{s} = 13$ TeV*, JHEP vol. 11 (2018), arXiv: 1805.04865.
- [180] CMS Collaboration, *Search for light pseudoscalar boson pairs produced from decays of the 125 GeV Higgs boson in final states with two muons and two nearby tracks in pp collisions at $\sqrt{s}=13$ TeV*, Phys. Lett. B vol. 800 (2020), arXiv: 1907.07235.
- [181] CMS Collaboration, *Search for a very light NMSSM Higgs boson produced in decays of the 125 GeV scalar boson and decaying into τ leptons in pp collisions at $\sqrt{s}=8$ TeV*, JHEP vol. 01 (2016), arXiv: 1510.06534.
- [182] ATLAS Collaboration, *Search for Higgs boson decays into a pair of pseudoscalar particles in the $b\bar{b}\mu\mu$ final state with the ATLAS detector in pp collisions at $\sqrt{s}=13$ TeV*, Phys. Rev. D vol. 105, no. 1 (2022), arXiv: 2110.00313.
- [183] ATLAS Collaboration, *Search for Higgs boson decays into a pair of light bosons in the $b\bar{b}\mu\mu$ final state in pp collision at $\sqrt{s}=13$ TeV with the ATLAS detector*, Phys. Lett. B vol. 790 (2019), arXiv: 1807.00539.
- [184] CMS Collaboration, *Search for an exotic decay of the Higgs boson to a pair of light pseudoscalars in the final state with two muons and two b quarks in pp collisions at 13 TeV*, Phys. Lett. B vol. 795 (2019), arXiv: 1812.06359.
- [185] CMS Collaboration, *Search for exotic decays of the Higgs boson to a pair of pseudoscalars in the $\mu\mu b\bar{b}$ and $\tau\tau b\bar{b}$ final states*, Eur. Phys. J. C vol. 84, no. 5 (2024), arXiv: 2402.13358.
- [186] ATLAS Collaboration, *Search for decays of the Higgs boson into a pair of pseudoscalar particles decaying into $b\bar{b}\tau^+\tau^-$ using pp collisions at $\sqrt{s}=13$ TeV with the ATLAS detector*, Phys. Rev. D vol. 110, no. 5 (2024), arXiv: 2407.01335.
- [187] CMS Collaboration, *Search for an exotic decay of the Higgs boson to a pair of light pseudoscalars in the final state with two b quarks and two τ leptons in proton-proton collisions at $\sqrt{s}=13$ TeV*, Phys. Lett. B vol. 785 (2018), arXiv: 1805.10191.
- [188] ATLAS Collaboration, *Search for new phenomena in events with at least three photons collected in pp collisions at $\sqrt{s}=8$ TeV with the ATLAS detector*, Eur. Phys. J. C vol. 76, no. 4 (2016), arXiv: 1509.05051.
- [189] ATLAS Collaboration, *Search for short- and long-lived axion-like particles in $H \rightarrow aa \rightarrow 4\gamma$ decays with the ATLAS experiment at the LHC*, Eur. Phys. J. C vol. 84, no. 7 (2024), arXiv: 2312.03306.

- [190] CMS Collaboration, *Search for the exotic decay of the Higgs boson into two light pseudoscalars with four photons in the final state in proton-proton collisions at $\sqrt{s} = 13$ TeV*, JHEP vol. 07 (2023), arXiv: 2208.01469.
- [191] CMS Collaboration, *Search for exotic Higgs boson decays $H \rightarrow \mathcal{A}\mathcal{A} \rightarrow 4\gamma$ with events containing two merged diphotons in proton-proton collisions at $\sqrt{s}=13$ TeV*, Phys. Rev. Lett. vol. 131 (2023), arXiv: 2209.06197.
- [192] ATLAS Collaboration, *Search for Higgs boson decays into pairs of light (pseudo)scalar particles in the $\gamma\gamma jj$ final state in pp collisions at $\sqrt{s}=13$ TeV with the ATLAS detector*, Phys. Lett. B vol. 782 (2018), arXiv: 1803.11145.
- [193] ATLAS Collaboration, *Luminosity determination in pp collisions at $\sqrt{s}=13$ TeV using the ATLAS detector at the LHC* (2022), arXiv: 2212.09379.
- [194] G. Avoni et al., *The new LUCID-2 detector for luminosity measurement and monitoring in ATLAS*, JINST vol. 13, no. 07 (2018).
- [195] ATLAS Data Preparation, *Luminosity For Physics*, <https://twiki.cern.ch/twiki/bin/viewauth/Atlas/LuminosityForPhysics>. Accessed: 20/11/2024.
- [196] ATLAS Data preparation, *Web directory for GRLs in atlas-groupdata area*, <https://atlas-groupdata.web.cern.ch/atlas-groupdata/GoodRunsLists/>. Accessed: 20/11/2024.
- [197] S. Alioli et al., *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*, JHEP vol. 06 (2010), arXiv: 1002.2581.
- [198] NNPDF Collaboration, *Parton distributions for the LHC Run II*, JHEP vol. 04 (2015), arXiv: 1410.8849.
- [199] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, Comput. Phys. Commun. vol. 191 (2015), arXiv: 1410.3012.
- [200] M. Czakon and A. Mitov, *Top++: A Program for the Calculation of the Top-Pair Cross-Section at Hadron Colliders*, Comput. Phys. Commun. vol. 185 (2014), arXiv: 1112.5675.
- [201] T. Ježo and P. Nason, *On the Treatment of Resonances in Next-to-Leading Order Calculations Matched to a Parton Shower*, JHEP vol. 12 (2015), arXiv: 1509.09071.
- [202] ATLAS Collaboration, *Studies of Monte Carlo predictions for the $t\bar{t}b\bar{b}$ process* (2022). ATL-PHYS-PUB-2022-006.
- [203] R. Frederix, E. Re and P. Torrielli, *Single-top t -channel hadroproduction in the four-flavour scheme with POWHEG and aMC@NLO*, JHEP vol. 09 (2012), arXiv: 1207.5391.
- [204] ATLAS Collaboration, *Studies on top-quark Monte Carlo modelling for Top2016* (2016). ATL-PHYS-PUB-2016-020.
- [205] N. Kidonakis, *Top Quark Production* (2014), arXiv: 1311.0283.

-
- [206] D. J. Lange, *The EvtGen particle decay simulation package*, Nucl. Instrum. Meth. A vol. 462 (2001).
- [207] 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 vol. 07 (2014), arXiv: 1405.0301.
- [208] J. Pumplin et al., *New generation of parton distributions with uncertainties from global QCD analysis*, JHEP vol. 07 (2002), arXiv: hep-ph/0201195.
- [209] SHERPA Collaboration, *Event Generation with Sherpa 2.2*, SciPost Phys. vol. 7, no. 3 (2019), arXiv: 1905.09127.
- [210] T. Gleisberg and S. Hoeche, *Comix, a new matrix element generator*, JHEP vol. 12 (2008), arXiv: 0808.3674.
- [211] F. Bucci et al., *OpenLoops 2*, Eur. Phys. J. C vol. 79, no. 10 (2019), arXiv: 1907.13071.
- [212] S. Schumann and F. Krauss, *A Parton shower algorithm based on Catani-Seymour dipole factorisation*, JHEP vol. 03 (2008), arXiv: 0709.1027.
- [213] S. Hoeche et al., *QCD matrix elements + parton showers: The NLO case*, JHEP vol. 04 (2013), arXiv: 1207.5030.
- [214] C. Anastasiou et al., *High precision QCD at hadron colliders: Electroweak gauge boson rapidity distributions at NNLO*, Phys. Rev. D vol. 69 (2004), arXiv: hep-ph/0312266.
- [215] ATLAS Collaboration, *Measurements of inclusive and differential fiducial cross-sections of $t\bar{t}$ production with additional heavy-flavour jets in proton-proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, JHEP vol. 04 (2019), arXiv: 1811.12113.
- [216] ATLAS Collaboration, *Measurements of the production cross-section for a Z boson in association with b-jets in proton-proton collisions at $\sqrt{s}=13$ TeV with the ATLAS detector*, JHEP vol. 07 (2020), arXiv: 2003.11960.
- [217] M. Zaheer et al., *Deep Sets* (2018), <https://arxiv.org/abs/1703.06114>, arXiv: 1703.06114.
- [218] ATLAS Collaboration, *Measurement of the Inelastic Proton-Proton Cross Section at $\sqrt{s}=13$ TeV with the ATLAS Detector at the LHC*, Phys. Rev. Lett. vol. 117, no. 18 (2016), arXiv: 1606.02625.
- [219] ATLAS Collaboration, *Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton-proton collision data at $\sqrt{s}=13$ TeV*, Eur. Phys. J. C vol. 79, no. 8 (2019), arXiv: 1902.04655.
- [220] ATLAS Collaboration, *Muon reconstruction performance of the ATLAS detector in proton-proton collision data at $\sqrt{s}=13$ TeV*, Eur. Phys. J. C vol. 76, no. 5 (2016), arXiv: 1603.05598.
- [221] 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 vol. 76, no. 11 (2016), arXiv: 1510.03823.

- [222] ATLAS Collaboration, *Jet energy scale measurements and their systematic uncertainties in proton-proton collisions at $\sqrt{s}=13$ TeV with the ATLAS detector*, Phys. Rev. D vol. 96, no. 7 (2017), arXiv: 1703.09665.
- [223] ATLAS Collaboration, *Determination of jet calibration and energy resolution in proton-proton collisions at $\sqrt{s}=8$ TeV using the ATLAS detector*, Eur. Phys. J. C vol. 80, no. 12 (2020), arXiv: 1910.04482.
- [224] 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 vol. 79, no. 11 (2019), arXiv: 1907.05120.
- [225] ATLAS Collaboration, *Measurement of the c -jet mistagging efficiency in $t\bar{t}$ events using pp collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS detector*, Eur. Phys. J. C vol. 82, no. 1 (2022), arXiv: 2109.10627.
- [226] ATLAS Collaboration, *Calibration of light-flavour jet b -tagging rates on ATLAS proton-proton collision data at $\sqrt{s}=13$ TeV* (2018). ATLAS-CONF-2018-006.
- [227] ATLAS Collaboration, *Calibration of a soft secondary vertex tagger using proton-proton collisions at $\sqrt{s}=13$ TeV with the ATLAS detector*, Phys. Rev. D vol. 110, no. 3 (2024), arXiv: 2405.03253.
- [228] ATLAS Combined Performance, *Tracking CP Recommendations for Early 2018 Analyses*, <https://twiki.cern.ch/twiki/bin/viewauth/AtlasProtected/TrackingCPRecsEarly2018>. Accessed: 19/01/2025.
- [229] ATLAS Combined Performance, *Final Tracking CP Recommendations for Run 2*, <https://twiki.cern.ch/twiki/bin/view/AtlasProtected/TrackingCPRecsRun2Final>. Accessed: 19/01/2025.
- [230] ATLAS Collaboration, *E_T^{miss} performance in the ATLAS detector using 2015-2016 LHC pp collisions* (2018). ATLAS-CONF-2018-023.
- [231] F. Cascioli et al., *Precise Higgs-background predictions: merging NLO QCD and squared quark-loop corrections to four-lepton + 0,1 jet production*, JHEP vol. 01 (2014), arXiv: 1309.0500.
- [232] R. Frederix and S. Frixione, *Merging meets matching in MC@NLO*, JHEP vol. 12 (2012), arXiv: 1209.6215.
- [233] ATLAS Collaboration, *Search for a new pseudoscalar decaying into a pair of muons in events with a top-quark pair at $\sqrt{s}=13$ TeV with the ATLAS detector*, Phys. Rev. D vol. 108, no. 9 (2023), arXiv: 2304.14247.
- [234] CMS Collaboration, *Search for physics beyond the standard model in multilepton final states in proton-proton collisions at $\sqrt{s}=13$ TeV*, JHEP vol. 03 (2020), arXiv: 1911.04968.
- [235] ATLAS Collaboration, *Measurement of the top-quark mass in $t\bar{t} \rightarrow$ dilepton events with the ATLAS experiment using the template method in 13 TeV pp collision data* (2022). ATLAS-CONF-2022-058.
- [236] A. Hoecker et al., *TMVA - Toolkit for Multivariate Data Analysis* (2009), arXiv: physics/0703039.

-
- [237] A. Paszke et al., *PyTorch: An Imperative Style, High-Performance Deep Learning Library*, Advances in Neural Information Processing Systems 32 (2019), <http://papers.neurips.cc/paper/9015-pytorch-an-imperative-style-high-performance-deep-learning-library.pdf>.
 - [238] ATLAS Collaboration, *Measurement of Higgs boson decay into b -quarks in associated production with a top-quark pair in pp collisions at $\sqrt{s}=13$ TeV with the ATLAS detector*, JHEP vol. 06 (2022), arXiv: 2111.06712.
 - [239] ATLAS Collaboration, *The ATLAS trigger system for LHC Run 3 and trigger performance in 2022*, JINST vol. 19, no. 06 (2024), arXiv: 2401.06630.
 - [240] ATLAS Collaboration, *The ATLAS experiment at the CERN Large Hadron Collider: a description of the detector configuration for Run 3*, JINST vol. 19, no. 05 (2024), arXiv: 2305.16623.
 - [241] ATLAS Collaboration, *The Phase-1 Upgrade of the ATLAS Level-1 Endcap Muon Trigger*, Springer Proc. Phys. vol. 212 (2018), arXiv: 1806.09234.

List of Figures

1.1	(a) Particle content of the Standard Model and (b) fundamental interactions between Standard Model particles.	5
1.2	Higgs potential $V(\Phi)$ for $\mu^2 \geq 0$ (left) and $\mu^2 < 0$ (right).	13
2.1	In the initially proposed quark model, u , d and s quarks (blue) formed a triplet of the $SU(3)$ flavour symmetry, which could be combined to create mesons and baryons. This figure contains two examples: the meson octet for pseudoscalar mesons (green) and the baryon octet for spin-1/2 baryons (red). They are represented in the $I_3 - Y$ plane of the QCD flavour symmetry. The line $Q = Y/2 + I_3$ is also included.	21
2.2	Example Feynman diagrams contributing to neutral kaon decay. This decay is suppressed at loop level due to the opposite sign from the u - and c -mediated diagrams with $m_u \neq m_c$. θ_C is the Cabibbo angle.	22
2.3	Feynman diagrams contributing to neutral kaon oscillation $K_0 \leftrightarrow \bar{K}_0$. CP violation appears because this process does not occur with the same probability in both directions.	22
2.4	Measurement of the $ee \rightarrow$ hadronic production cross section around the Z resonance. The curves indicate the predicted cross section for two, three and four neutrino species with SM couplings and negligible mass [47].	24
2.5	Feynman diagrams for the main Higgs production mechanisms in the LHC. VH can refer to both WH and ZH	24
2.6	(a) SM Higgs boson production cross sections as a function of the centre-of-mass energy, $\sqrt{s}$, for proton-proton collisions and (b) branching ratios for the main decays of the SM Higgs boson around $m_H = 125$ GeV. Data from [49].	25
2.7	(a) Reduced Higgs boson coupling strength modifiers and their uncertainties from Run 2 [3] and (b) summary of Higgs mass measurements from the $H \rightarrow ZZ^* \rightarrow 4\ell$ and $H \rightarrow \gamma\gamma$ decay channels and their combination [50].	25
2.8	Overview of cross section measurements of selected SM processes compared to the corresponding theoretical expectations. The figure includes results from Run 1, Run 2 and the early stages of Run 3, using the ATLAS detector [52]. . . .	26
2.9	Superposition of 21 galaxy rotation curves. The plot shows the speed of the objects in the galactic disk with respect to the distance to the galactic nucleus. This form of the rotation curves suggests that the total mass of the galaxy is not condensed around the centre, where ordinary matter lies, but also at large radius in the form of dark matter [56].	29

2.10	(a) Image of the bullet cluster, comprised of two colliding galaxy clusters. The contour lines represent the mass distribution obtained via gravitational lensing and the gray dots correspond to the X-ray image from ordinary matter [57]. Dark matter has a much lower interaction cross section, and therefore suffers less deceleration during the collision. (b) Comparison of the galaxy distribution obtained from the Sloan Digital Sky Survey (blue) and theoretical mock catalogs from the Millenium Simulation. The pattern of galaxy clustering through gravity can help determining the properties of hypothetical dark matter candidates [58].	30
2.11	Scheme of the two distinct neutrino mass hierarchies, from measurements of solar (<i>sol</i>) and atmospheric (<i>atm</i>) neutrinos. The colour code indicates the fraction of each neutrino flavour (e, μ, τ) present in each of the mass eigenstates (1, 2, 3) [59].	30
2.12	Summary of theoretical calculations and experimental measurements of the muon anomaly $a_\mu = (g - 2)/2$ [60].	31
3.1	$\text{BR}(a \rightarrow XX)$ with $X = b, c, \tau, g$ for $\tan \beta = 1$, $\cos \theta = 0.7$ and $m_\chi \gg m_a$ in the type II 2HDM+ a as a function of the a -boson mass. The couplings of the a -boson to WW and ZZ are zero due to CP conservation. The mass of the light CP-even state h , equivalent to the SM Higgs boson, is marked in grey.	40
4.1	Layout of the CERN accelerator complex [69].	46
4.2	Diagram showing the basic functioning of the LHC RF cavities.	47
4.3	Diagram showing the cross section of a LHC dipole (left) and quadrupole (right) magnet. Adapted from [70, 71].	47
4.4	Cross section versus centre-of-mass energy for different processes in the LHC. The dashed vertical lines indicate the 7 and 8 TeV energies used in LHC Run 1, the 13 TeV energy used in LHC Run 2, and the 100 TeV goal for future colliders. Adapted from [72].	48
4.5	(a) Cumulative integrated luminosity at a centre-of-mass energy of 13 TeV versus time during Run 2. (b) Luminosity-weighted distribution of the mean number of interactions per crossing for the 2015–2018 proton-proton collision data [74]. . .	49
4.6	The ATLAS detector at CERN [75].	50
4.7	Coordinate system used in the ATLAS Experiment. Adapted from [76].	51
4.8	Schematic view of a quarter section of the ATLAS ID [80].	52
4.9	Schematic view of a quarter section of the ATLAS calorimeters. Adapted from [90].	54
4.10	Schematic view of a quarter section of the ATLAS muon spectrometer [91]. . . .	56
4.11	Schematic view of the ATLAS magnets. The barrel toroid is shown in red and the endcap toroids are shown in green. The blue cilinder is the central solenoid, responsible for the magnetic field in the ID.	57
4.12	Schematic of the ATLAS TDAQ system in Run 2 [92].	58
4.13	Schematic of the ATLAS L1Calo system in Run 2. ECAL and HCAL refer to the EM and hadronic calorimeters, respectively [94].	59
4.14	Schematic of the ATLAS L1Muon system in Run 2.	60
4.15	Schematic view of the L1Topo data flow [96].	61
5.1	Schematic view of a LHC proton-proton collision. Adapted from [101].	63
5.2	Feynman diagram representing the various parameters used in Equation 5.1. . .	65

5.3	Graphic representation of the DGLAP splitting functions.	65
5.4	Global fit by the NNPDF Collaboration. PDFs are shown at factorisation scales of 10 GeV^2 (left) and 10^4 GeV^2 (right). PDFs evaluated at values of μ_F are connected by the DGLAP evolution equations [109].	66
5.5	Schematic representations of (a) the string hadronisation model and (b) the cluster hadronisation model [119].	70
5.6	Schematic representation of the ATLAS simulation software, from event generation (top left) to event reconstruction (top right). Algorithms are placed in rectangular boxes and persistent data objects are placed in rounded boxes. The optional pileup portion of the chain, used only when events are overlaid, is shown in blue. Adapted from [124].	72
6.1	Jet clustering comparison performed on the same data with the same input radius, $R = 1$. Results using (a) the k_T algorithm, (b) the Cambridge/Aachen algorithm and (c) the anti- k_T algorithm are shown. Different colours are used to represent different jets and their areas in the $y - \phi$ plane. Adapted from [140].	77
6.2	Schematic representation of a b -jet.	78
6.3	Normalised distributions of the DL1r b -jet discriminant corresponding to light-, c - and b -jets for a $t\bar{t}$ test sample using $f_c = 0.018$ [150].	79
6.4	DL1r light- and c -flavour rejection with respect to the b -jet efficiency [152].	80
6.5	Mistag rates for (a) light- and (b) c -jets for the 85%, 77%, 70% and 60% DL1r b -tagging WPs as a function of the jet p_T [151].	81
6.6	(a) DeXTer discriminant distribution evaluated using B -labeled jets from $H \rightarrow aa \rightarrow b\bar{b}b\bar{b}$ and $t\bar{t}a$, $a \rightarrow b\bar{b}$ samples and b - and light-labeled jets from $t\bar{t}$ samples with $f_b = 0.4$. (b) Performance for different a -boson masses in $H \rightarrow aa \rightarrow 4b$ events. Images from [153].	82
7.1	Schematic representation of the basic constituents of a NN. The input layer receives the initial information $\mathbf{x}$, which is transformed in the hidden layers by a combination of linear and non-linear functions, resulting in the predicted output $\hat{y}$	88
7.2	Schematic representation of the loss function with respect to the number of epochs for both the training and validation datasets. Comparing the training and validation loss functions allows to monitor a possible underfitting or overfitting of the model.	91
7.3	Visual representation of different NN performances with their corresponding NN scores and ROC curves.	92
7.4	Schematic representation of the basic constituents of a parametrised NN.	92
7.5	Schematic representation of a decision tree. An input event is evaluated in each node, and classified accordingly until a final decision (leaf) is reached.	93
8.1	Thresholds for exclusion ($Z = 1.64\sigma$), evidence ($Z = 3\sigma$) and discovery ($Z = 5\sigma$) using a gaussian distribution with mean = 0 and standard deviation = 1. The associated p -value corresponds to the area under the curve at the tail of the distribution.	99

List of Figures

9.1	Feynman diagrams for $H \rightarrow aa \rightarrow 4b$ production in association with a Z boson decaying to two leptons at tree-level and 1-loop.	103
9.2	(a) Higgs and (b) Z -boson p_T from $ZH, H \rightarrow aa \rightarrow 4b$ production using different mass hypotheses. The kinematics of the ZH system do not depend on the value of m_a	104
9.3	(a) Leading and (b) sub-leading a -boson p_T from $H \rightarrow aa$ decay for different mass hypotheses. These distributions show how decay products with larger m_a tend to carry less momentum.	105
9.4	(a) ΔR separation between the two b -quarks from the (a) leading and (b) sub-leading $a \rightarrow b\bar{b}$ decay for different mass hypotheses. For lower values of m_a , the system is boosted, while for higher values of m_a , the system is resolved.	105
9.5	Dominant Feynman diagram for $t\bar{t}$ +jets production at tree level with four b -quarks in the final state.	112
9.6	Examples of Feynman diagrams for Z +jets production at tree level with two, three and four b -quarks in the final state.	112
9.7	Data versus MC comparison of the sum of the PC b -tagging score of all the jets per event (a) before and (b) after applying the $t\bar{t}$ +jets normalisation correction. The error band includes the statistical error and the DL1r uncertainties.	115
9.8	Data versus MC comparison of the N_{jets} distribution (a) before and (b) after applying the N_{jets} correction to the $t\bar{t}$ +jets and tW backgrounds. Systematic uncertainties are not included.	116
9.9	Data-only (a) H_T and (b) H_T^{red} normalised distributions as a function of the number of jets in the $t\bar{t}$ +jets reweighting region. The histograms show how each additional jet in the event causes an offset in H_T , which is mitigated in H_T^{red}	117
9.10	$t\bar{t}$ +jets correction factor as a function of H_T^{red}	118
9.11	N_{jets} distribution (a) before and (b) after applying the H_T correction using H_T^{red} . Systematic uncertainties are not included.	118
9.12	Leading (1) and sub-leading (2) jet p_T distributions (a) before and (b) after applying the H_T correction using H_T^{red} . Systematic uncertainties are not included.	119
9.13	Jet 3 and jet 4 p_T distributions (a) before and (b) after applying the H_T correction using H_T^{red} . Jets are ordered by p_T . Systematic uncertainties are not included.	120
9.14	Leading (1) and sub-leading (2) lepton p_T distributions (a) before and (b) after applying the H_T correction using H_T^{red} . Systematic uncertainties are not included.	121
9.15	Data versus MC comparison of the N_{jets} distribution (a) before and (b) after applying the N_{jets} correction to the Z +jets background. Systematic uncertainties are not included.	124
9.16	Z +jets correction factor as a function of Z p_T	125
9.17	N_{jets} distribution (a) before and (b) after applying the Z p_T correction. Systematic uncertainties are not included.	125
9.18	Leading (1) and sub-leading (2) jet p_T distributions (a) before and (b) after applying the Z p_T correction. Systematic uncertainties are not included.	126
9.19	Leading (1) and sub-leading (2) lepton p_T distributions (a) before and (b) after applying the Z p_T correction. Systematic uncertainties are not included.	127

9.20	Illustration of all object combinations considered for the hypothesis testing NN in the $4b$ final state. The B -jet can either represent a boosted $a \rightarrow b\bar{b}$ decay reconstructed as one single Ak8 jet (B) or a merged jet containing one b from each a -boson (B).	129
9.21	Diagram of the NN used for event reconstruction. All filled blocks are dense multilayer perceptions (MLPs) that share common weights between MLPs of the same colour. This particular example shows the case of a hypothesis with one Ak8 jet candidate, two Ak4 jet candidates and one soft secondary vertex candidate.	130
9.22	Highest hypothesis score distribution in the 2B region for $m_a = 12$ GeV.	131
9.23	Highest hypothesis score distribution in the (a) 1B2b and (b) 1B1b1v regions for $m_a = 20$ GeV.	131
9.24	Highest hypothesis score distribution in the (a) 4b and (b) 3b1v regions for $m_a = 40$ GeV.	132
9.25	Examples of BDT input distributions for different reconstruction and a -boson mass hypotheses. From top to bottom: leading a p_T and E_T^{miss} for the $m_a = 12$ GeV hypothesis in the 2B region, $\cos \theta^*$ and reduced Higgs boson mass for the $m_a = 20$ GeV hypothesis in the 1B2b region, reduced leading and sub-leading a -boson mass for the $m_a = 40$ GeV hypothesis in the 4b region.	134
9.26	Examples of BDT score distributions for different reconstruction and a -boson mass hypotheses. From top to bottom, $m_a = 12$ GeV hypothesis in the 2B region, $m_a = 20$ GeV hypothesis in the 1B2b and 1B1b1v regions and $m_a = 40$ GeV hypothesis in the 4b and 3b1v regions.	135
9.27	Background composition in the five signal regions under the $m_a = 12, 20$ and 40 GeV mass hypotheses. All signal regions are dominated by the Z +jets background due to the SF and Z -mass cuts. In general, the $t\bar{t}$ +light background contributes significantly in the boosted regime (2B, 1B2b and 1B1b1v), while $t\bar{t}$ +HF is more prominent in the resolved regime (4b, 3b1v).	138
9.28	Background composition in the six Z +jets control regions under the $m_a = 12, 20$ and 40 GeV mass hypotheses. All Z +jets control regions are dominated by the Z +jets background due to the SF and Z -mass cuts. The E_T^{miss} veto contributes to suppress the $t\bar{t}$ +jets background as well.	139
9.29	Background composition in the six $t\bar{t}$ +jets control regions, common to all mass hypotheses. All $t\bar{t}$ +jets control regions are dominated by the $t\bar{t}$ +jets background due to the DF cut. The $t\bar{t}$ +light category constitutes most of the 2B control regions (1i.2i(T+L).0i and 0.2i.0i), while others have a larger fraction of $t\bar{t}$ +HF in comparison.	140
10.1	Predicted yields for the signal and control regions in the B-only fit to the Asimov dataset with $m_a = 25$ GeV (a) before and (b) after the fit. The expected signal yield for the 25 GeV mass hypothesis assuming the SM production cross section for ZH and $\text{BR}(H \rightarrow aa \rightarrow 4b) = 1$ is overlayed in the pre-fit plot.	150
10.2	Normalisation factors for the Z +jets and $t\bar{t}$ +jets samples in the fit to the Asimov dataset with $m_a = 25$ GeV.	151
10.3	Correlation matrix for the relevant NPs in the B-only fit to the Asimov dataset with $m_a = 25$ GeV.	151

List of Figures

10.4	Fitted nuisance parameters from the B-only fit to the Asimov dataset with $m_a = 25$ GeV.	152
10.5	Up (red) and down (blue) 1σ variation of the third eigenvalue of the DeXTer systematic uncertainty in the 2B SR, corresponding to the (a) $Z+\geq 1b$ MC sample and the (b) 25 GeV signal MC sample.	153
10.6	Up (red) and down (blue) 1σ variation of the MADGRAPH5_aMC@NLO ME uncertainty in the (a) Z +jets 4b CR and the (b) 4b SR, corresponding to the $Z+\geq 1b$ MC sample.	153
10.7	Up (red) and down (blue) 1σ variation of the HERWIG 7 PS uncertainty for $t\bar{t}$ +light in the (a) the $t\bar{t}$ +jets 4b CR and (b) the 4b SR.	154
10.8	Ranked systematic impacts to the signal strength from the B-only fit to the Asimov dataset with $m_a = 25$ GeV.	154
10.9	Predicted yields for the signal and control regions in the S+B fit to data with $m_a = 25$ GeV (a) before and (b) after the fit. The fitted signal yield for the 25 GeV mass hypothesis is stacked in the post-fit plot.	156
10.10	Normalisation factors for the Z +jets and $t\bar{t}$ +jets samples in the S+B fit to data with $m_a = 25$ GeV.	157
10.11	Correlation matrix for the relevant NPs in the S+B fit to data with $m_a = 25$ GeV.	157
10.12	Fitted nuisance parameters from the S+B fit to data with $m_a = 25$ GeV.	158
10.13	Ranked impact of the systematic uncertainties to the signal strength from the S+B fit to data with $m_a = 25$ GeV.	159
10.14	(a) Expected and observed 95% CL upper limits of $\text{BR}(H \rightarrow aa \rightarrow 4b)$ as a function of m_a . The red and blue lines correspond to the two previous analyses using 36 fb^{-1} of proton-proton collisions at $\sqrt{s} = 13$ TeV. (b) Expected 95% CL upper limits of $\text{BR}(H \rightarrow aa \rightarrow 4b)$ as a function of m_a for each SR. The dashed line corresponds to the combined expected limit from (a). All limits are calculated for the available signal mass points and smoothed using ROOT for easier visualisation.	161
10.15	(a) Expected and observed 95% CL upper limits of $\text{BR}(H \rightarrow aa \rightarrow XXYY)$ as a function of m_a using the full Run 2 dataset recorded by the ATLAS experiment. (b) Observed 95% CL upper limits of $\text{BR}(H \rightarrow aa \rightarrow XXYY)$ as a function of m_a using the full Run 2 dataset recorded by the ATLAS and CMS experiments. Only analyses with at least two b -quarks in the final state are included.	162
11.1	Example Feynman diagrams for a -boson production in association with $t\bar{t}$ and with tW in the dilepton channel with $a \rightarrow b\bar{b}$	167
11.2	Cross sections for the $t\bar{t}a$ and tWa signal processes assuming $g_t = 0.5$ and $\text{BR}(a \rightarrow b\bar{b}) = 1$	168
11.3	(a) Top quark p_T and (b) angular separation between the two leptons from the $t\bar{t}$ system under different a -boson mass hypotheses. The kinematics of the $t\bar{t}$ system are not sensitive to the value of m_a	168
11.4	(a) a -boson p_T and (b) angular separation between the two b -quarks from the $a \rightarrow b\bar{b}$ under different a -boson mass hypotheses. These two distributions show how higher masses tend to carry more energy, but their decay products have a bigger angular separation.	169

11.5	(a) Leading and (b) sub-leading b -quark p_T from $t\bar{t}a$, $a \rightarrow b\bar{b}$ production using different mass hypotheses. For higher masses, the a -boson is more energetic, and tends to produce particles with larger p_T	169
11.6	Data versus MC comparison of the sum of the PC b -tagging score of all the jets per event (a) before and (b) after applying the $t\bar{t}$ +jets normalisation correction. The error band includes the statistical error and the DL1r uncertainties.	176
11.7	Data versus MC comparison of the N_{jets} distribution (a) before and (b) after applying the $t\bar{t}$ +jets normalisation correction. Systematic uncertainties are not included.	177
11.8	Data-only (a) H_T and (b) H_T^{red} normalised distributions as a function of N_{jets} in the $t\bar{t}$ +jets reweighting region. The histograms show how each additional jet in the event causes an offset in H_T , while this effect is mitigated in H_T^{red}	178
11.9	$t\bar{t}$ +jets and tW correction function as a function of H_T^{red}	179
11.10	Data versus MC comparison of the N_{jets} distribution (a) before and (b) after correcting the H_T distribution using H_T^{red} . Systematic uncertainties are not included.	179
11.11	Leading (1) and sub-leading (2) jet p_T distributions (a) before and (b) after applying the H_T correction using H_T^{red} . Systematic uncertainties are not included.	180
11.12	Leading (1) and sub-leading (2) lepton p_T distributions (a) before and (b) after applying the H_T correction using H_T^{red} . Systematic uncertainties are not included.	181
11.13	Missing transverse energy distribution (a) before and (b) after applying the H_T correction using H_T^{red} . Systematic uncertainties are not included.	182
11.14	Top BDT response for signal and background, using the training and the test dataset, trained with (a) even and (b) odd events.	184
11.15	Normalised MC distributions for (a) the highest hypothesis score from the $t \rightarrow j\bar{\ell}$ BDT, (b) the lepton-jet p_T , (c) the difference in pseudorapidity and (d) the difference in azimuthal angle. Signal distributions for $m_a = 12, 30, 60$ and 80 GeV are overlayed for comparison.	185
11.16	$a \rightarrow jj$ BDT response for signal and background, using the training and the test dataset, trained with (a) even and (b) odd events.	187
11.17	Normalised MC distributions for (a) the highest hypothesis score from the $a \rightarrow jj$ BDT, (b) the angular separation of the di-jet pair, (c) the di-jet p_T and (d) the di-jet invariant mass. Signal distributions for $m_a = 12, 30, 60$ and 80 GeV are overlayed for comparison.	188
11.18	MC distributions of (a,c) the hadronic event hardness and (b,d) the average of the pseudorapidity separation for $b\bar{b}$ pairs in the 0B4b and 0B3b regions, respectively. Signal distributions for $m_a = 12, 30, 60$ and 80 GeV are overlayed for comparison. The 12 GeV mass hypothesis is not displayed the 0B4b region due to low statistics.	191
11.19	MC distributions of (a,c) the B -jet p_T and (b,d) the B -jet invariant mass in the 1B2b and 1B1b regions, respectively. Signal distributions for $m_a = 12, 30, 60$ and 80 GeV are overlayed for comparison.	192
11.20	NN score and ROC curve in the 0B4b region for each mass hypothesis.	193
11.21	NN score and ROC curve in the 0B3b region for each mass hypothesis.	194
11.22	NN score and ROC curve in the 0B4b region for each mass hypothesis.	195
11.23	NN score and ROC curve in the 0B4b region for each mass hypothesis.	196

List of Figures

11.24	AUC versus mass for each signal hypothesis in the NN regions (a) 0B4b, (b) 0B3b, (c) 1B2b and (d) 1B1b.	197
11.25	$t\bar{t}$ +jets and tW correction as a function of H_T^{red} for the nominal (black line) and alternative samples (grey, red and green lines).	203
12.1	Normalisation factors for the $t\bar{t}$ +jets components (a) without and (b) with the 0B2b+1bL CR. The test is performed using a B-only fit to the bins where the signal to background ratio is inferior to 1%.	208
12.2	Background composition in the signal and control regions of the fit.	209
12.3	Normalisation factors for the $t\bar{t}$ +jets components with different categorisation schemes for $t\bar{t}+\geq 1b$ as a function of the a -boson mass. These results show how further decomposition of the $t\bar{t}+\geq 1b$ category can lead to fit instability.	210
12.4	Pre-fit distributions from the B-only fit to the Asimov dataset with $m_a = 30$ GeV.	213
12.5	Post-fit distributions from the B-only fit to the Asimov dataset with $m_a = 30$ GeV.	214
12.6	Normalisation factors for the $t\bar{t}$ +jets samples in the fit to the Asimov dataset with $m_a = 30$ GeV.	215
12.7	Correlation matrix for the relevant NPs in the B-only fit to the Asimov dataset with $m_a = 30$ GeV.	215
12.8	Fitted nuisance parameters from the B-only fit to the Asimov dataset with $m_a = 30$ GeV (I).	216
12.9	Fitted nuisance parameters from the B-only fit to the Asimov dataset with $m_a = 30$ GeV (II).	217
12.10	Up (red) and down (blue) 1σ variation of the third eigenvalue of the DeXTer systematic uncertainty with respect to the NN score distribution in the (a) 1B2b SR and (b) 1B1b+1bL SR, corresponding to the $t\bar{t}+1b$ MC sample.	218
12.11	Up (red) and down (blue) 1σ variation of the HERWIG 7 PS uncertainty with respect to the sumPCBTag distribution in the 0B2b+1bL CR, corresponding to the $t\bar{t}$ +light MC sample.	218
12.12	Up (red) and down (blue) 1σ variation of the pThard=1 uncertainty with respect to the NN score distribution for the (a) $t\bar{t}+1b$ and the (b) $t\bar{t}+2b$ MC samples in the 0B4b SR.	219
12.13	Ranked systematic impacts to the signal strength from the B-only fit to the Asimov dataset with $m_a = 30$ GeV.	219
12.14	Pre-fit plots from the S+B fit to data with $m_a = 30$ GeV.	221
12.15	Post-fit plots from the S+B fit to data with $m_a = 30$ GeV.	222
12.16	Normalisation factors for the $t\bar{t}$ +jets samples in the S+B fit to data with $m_a = 30$ GeV.	223
12.17	Correlation matrix for the relevant NPs in the S+B fit to data with $m_a = 30$ GeV.	223
12.18	Fitted nuisance parameters from the S+B fit to data with $m_a = 30$ GeV (I).	224
12.19	Fitted nuisance parameters from the S+B fit to data with $m_a = 30$ GeV (II).	225
12.20	Ranked impact of the systematic uncertainties to the signal strength from the S+B fit to data with $m_a = 30$ GeV.	226

12.21	(a) Expected and observed 95% CL upper limits of $\sigma(t\bar{t}a) \times \text{BR}(a \rightarrow b\bar{b})$ as a function of m_a . The red lines correspond to different values of the coupling to the top quark. (b) Expected 95% CL upper limits of $\sigma(t\bar{t}a) \times \text{BR}(a \rightarrow b\bar{b})$ as a function of m_a for each SR. The dashed line corresponds to the combined expected limit from (a). All limits are calculated for the available signal mass points and smoothed using ROOT for easier visualisation.	228
A.22	Schematics of the ATLAS TDAQ system in Run 3.	231
A.23	Trigger tower composed of ten SuperCells after the Phase-I upgrade for the LAr calorimeter [94].	232
A.24	Granularity of the inputs to (a) jFEX and (b) gFEX in the $\Delta\eta \times \Delta\phi$ plane [240].	233
A.25	Phase-I upgrades for the L1Muon trigger [241].	233
A.26	Visual representation of the trigger towers used to determine $E_T(\text{jTau})$, $E_T^{\text{iso}}(\text{jTau})$ and $E_T(\text{eTau})$. Each square corresponds to one trigger tower, some of them containing several SuperCells. HAD and EM refer to the hadronic and EM calorimeters, respectively.	237
B.27	Predicted yields for the signal and control regions in the blinded fit to data with $m_a = 25$ GeV (a) before and (b) after the fit. Data points are drawn only in the signal-depleted control regions that are used in the fit. The expected signal yield for the 25 GeV mass hypothesis assuming the SM production cross section for ZH and $\text{BR}(H \rightarrow aa \rightarrow 4b) = 1$ is overlayed in the pre-fit plot.	240
B.28	Normalisation factors for the Z +jets and $t\bar{t}$ +jets samples in the blinded fit to data with $m_a = 25$ GeV.	241
B.29	Correlation matrix for the relevant NPs in the blinded fit to data with $m_a = 25$ GeV.	241
B.30	Fitted nuisance parameters from the blinded fit to data with $m_a = 25$ GeV.	242
B.31	Pre-fit plots from the blinded fit to data with $m_a = 30$ GeV.	244
B.32	Post-fit plots from the blinded fit to data with $m_a = 30$ GeV.	245
B.33	Normalisation factors for the $t\bar{t}$ +jets samples in the blinded fit to data with $m_a = 30$ GeV.	246
B.34	Correlation matrix for the relevant NPs in the blinded fit to data with $m_a = 30$ GeV.	246
B.35	Fitted nuisance parameters from the blinded fit to data with $m_a = 30$ GeV (I).	247
B.36	Fitted nuisance parameters from the blinded fit to data with $m_a = 30$ GeV (II).	248
C.37	Expected limits from the (a) $m_a = 12$ GeV, (b) $m_a = 30$ GeV and (c) $m_a = 60$ GeV injection tests. They are computed for $\mu = 0$ (dashed black line), $\mu = 0.5$ (solid red line), $\mu = 1$ (solid purple line) and $\mu = 2$ (solid blue line).	250
C.38	Fitted signal strength with respect to the a -boson mass from the (a) $m_a = 12$ GeV, (b) $m_a = 30$ GeV and (c) $m_a = 60$ GeV injection tests. The different colours correspond to $\mu = 0.5$ (red), $\mu = 1$ (purple) and $\mu = 2$ (blue).	251
C.39	Relevant nuisance parameters from the (a) $m_a = 12$ GeV, (b) $m_a = 30$ GeV and (c) $m_a = 60$ GeV injection tests. The different colours correspond to $\mu = 0.5$ (red), $\mu = 1$ (black) and $\mu = 2$ (blue).	253

List of Tables

1.1	Electroweak charges associated to each fermion of the SM. For simplicity, $\ell = e, \mu, \tau$, $u = u, c, t$ and $d = d, s, b$. Antiparticles have the opposite electrical charge, which leads to the sign of T_3 and Y_W being flipped as well. Because right-handed neutrinos are not charged under any symmetry, they are usually excluded from the model.	10
1.2	QCD charges associated to light and heavy quarks. For antiparticles, all charges have the same value and opposite sign. Quantum numbers I_3 , Y and S are not shown for heavy quarks as the approximate QCD flavour symmetry applies to light quarks only.	12
3.1	$\mathbb{Z}_2$ charges in the four 2HDM types. The notation $\ell = e, \mu, \tau$, $u = u, c, t$ and $d = d, s, b$ is used. By convention, u_R always couples to Φ_2	37
3.2	Yukawa coupling strength for the neutral h , H and A in the four 2HDM models. The notation $\ell = e, \mu, \tau$, $u = u, c, t$ and $d = d, s, b$ is used.	37
3.3	Yukawa coupling strength for the neutral h , H and A in the four 2HDM models at the decoupling limit $\alpha \rightarrow \beta - \pi/2$. The notation $\ell = e, \mu, \tau$, $u = u, c, t$ and $d = d, s, b$ is used.	38
3.4	2HDM+ a parameters used in the $t\bar{t}a$, $a \rightarrow b\bar{b}$ analysis.	39
4.1	List of topological algorithms implemented in Run 2. Trigonometric functions are calculated using look-up tables.	62
6.1	Correspondence between DL1r b -tagging efficiency (ε_b) and the associated working points (WP) and b -tagging scores.	80
7.1	Common activation functions used in NN training and possible use cases. . . .	88
9.1	GRLs used for each of the four years of data-taking, together with the corresponding integrated luminosity for unprescaled triggers, computed with the 0f1Lumi-13TeV-011 luminosity tag [195]. All the listed GRLs are available in [196].	106
9.2	Summary of signal and background samples used in the $H \rightarrow aa \rightarrow 4b$ analysis, including the ME generator, the PDF set, the PS generator, the cross section normalisation and the type of detector simulation. Note that MADGRAPH is abbreviated to MG.	110
9.3	Single-electron and single-muon trigger chains, depending on the year of data-taking.	110
9.4	Background composition in the preselection region, corresponding to 140 fb^{-1} . .	111
9.5	Number of signal events in the preselection region, corresponding to 140 fb^{-1} . .	112
9.6	Background composition in the $t\bar{t}$ +jets reweighting region.	113

List of Tables

9.7	Correspondence between the DL1r WP and the b -tagging score for each jet. Low values of the DL1r WP correspond to tight b -tagging WPs, and viceversa.	114
9.8	Normalisation factors for the three $t\bar{t}$ +jets categories resulting from the likelihood fit performed using the sumPCBTag distribution.	114
9.9	Correction factors for each bin of the $t\bar{t}$ +jets and tW N_{jets} distribution.	116
9.10	Fitted parameters to the hyperbolic function for the H_T correction. The table shows the numerical values from all reweighting parameters, with their corresponding uncorrelated uncertainties.	118
9.11	Background composition in the Z +jets reweighting region.	122
9.12	Z +jets composition in the Z +jets reweighting region, as well as additional regions with higher b -object multiplicities, where most of the signal events are expected to be. The fraction of $Z+\geq 1b$ is above 80% in all of them.	123
9.13	$t\bar{t}$ +jets composition in the Z +jets reweighting region, as well as additional regions with higher b -object multiplicities, where most of the signal events are expected to be. The $t\bar{t}$ +jets composition varies strongly from one region to another.	123
9.14	Correction factors for each bin of the Z +jets N_{jets} distribution.	124
9.15	Values of the parameters from the Z +jets Z p_T fit.	125
9.16	List of features used as inputs to the hypothesis testing NN.	130
9.17	List of features used as inputs to the hypothesis testing NN.	133
9.18	Signal regions for the $H \rightarrow aa \rightarrow 4b$ analysis.	136
9.19	Control regions for the Z +jets background. Each control region exists for every mass hypothesis, as they all depend on the score of the hypothesis testing NN. SFonZ indicates that the two leptons have the same flavour and are within the Z -mass window. The naming scheme can be read as (number of B tight).(number of B loose).(number of b). The letter i represents that the selection is inclusive. *"vorj" means that $N_{\text{jets}} \geq 1$ is required only if $N_v = 0$	137
9.20	Control regions for the $t\bar{t}$ +jets background. DF indicates that the two leptons have different flavour. The naming scheme can be read as (number of B tight).(number of B loose).(number of b). The letter i represents that the selection is inclusive. *"vorj" means that $N_{\text{jets}} \geq 1$ is required only if $N_v = 0$	137
9.21	List of experimental systematic uncertainties included in the analysis. "N" means that the uncertainty is taken as normalisation-only for all processes and channels affected, while "SN" means that the uncertainty is taken on both the shapes and the normalisation. Some of the systematic uncertainties are split into several components for a more accurate treatment. Additional systematic uncertainties due to the low- p_T jets and the Ak8 track mass scale, particular to this analysis, are shown separately with a "+" sign.	143
9.22	List of modelling systematic uncertainties included in the analysis. "N" means that the uncertainty is taken as normalisation-only for all processes and channels affected, while "SN" means that the uncertainty is taken on both the shapes and the normalisation. Some of the systematic uncertainties are split into several components for a more accurate treatment. Systematics derived from alternative samples are not included.	146

9.23	Summary of alternative MC samples used in the $H \rightarrow aa \rightarrow 4b$ analysis. Note that MADGRAPH is abbreviated to MG. All alternative samples are simulated using AFII. They are compared to the AFII version of the nominal FS MC, highlighted in grey.	146
10.1	Relative impacts over the fitted signal strength $\hat{\mu}$ obtained from the S+B fit to data, grouped by the type of uncertainty. The impact of the statistical uncertainty (without MC stats.) is also included for comparison.	159
10.2	Post-fit signal and background yields in the five signal regions for the $m_a = 25$ GeV hypothesis. The total MC and data yields are included in the last two rows.	160
10.3	Post-fit signal and background yields in the six Z +jets control regions (SFonZ) for the $m_a = 25$ GeV hypothesis. The total MC and data yields are included in the last two rows.	160
10.4	Post-fit signal and background yields in the six $t\bar{t}$ +jets control regions (DF) for the $m_a = 25$ GeV hypothesis. The total MC and data yields are included in the last two rows.	160
11.1	Summary of signal and background samples used in the $t\bar{t}a, a \rightarrow b\bar{b}$ analysis, including the ME generator, the PDF set, the PS generator, the cross section normalisation and the type of detector simulation. Note that MADGRAPH is abbreviated to MG.	173
11.2	Background composition in the preselection region, corresponding to 140 fb^{-1}	174
11.3	Number of signal events in the preselection region, corresponding to 140 fb^{-1}	174
11.4	Background composition of the $t\bar{t}$ +jets control region. The $t\bar{t}$ +jets and tW backgrounds constitute around 100% of the events.	175
11.5	correspondence between the DL1r WP and the b -tagging score for each jet. Low values of the DL1r WP correspond to tight b -tagging WPs, and viceversa.	176
11.6	Normalisation factors for the three $t\bar{t}$ +jets categories resulting from the likelihood fit performed using the sumPCBTag distribution.	176
11.7	Fitted parameters to the hyperbolic function for the H_T correction. The table shows the numerical values from all reweighting parameters, with their corresponding uncorrelated uncertainties.	178
11.8	Number of $t\bar{t}a$ and $t\bar{t}$ +jets MC events used in the training of the $t \rightarrow j\ell$ reconstruction BDTs.	183
11.9	Input variables to the $t \rightarrow j\ell$ reconstruction BDTs. Kinematic variables of the $t\bar{t}$ pair are computed using $j\ell$ and $j\bar{\ell}$	184
11.10	Number of $t\bar{t}a$ and $t\bar{t}$ +jets MC events used in the training of the $a \rightarrow jj$ reconstruction BDT.	186
11.11	Input variables to the $a \rightarrow jj$ reconstruction BDT.	186
11.12	Signal regions for signal versus background discrimination. All regions are required to have at least 2 Ak4 jets for the BDT variables to be well-defined.	189
11.13	Number of signal and background MC events used in the training of the signal versus background discriminator for each training region. In this table, "Other" includes the $t\bar{t}H$, $t\bar{t}Z$, $t\bar{t}W$, tq , tZ , tWZ , Z/W +jets and diboson background processes. The 12 and 16 GeV $t\bar{t}a$ samples are excluded from the 0B4b training due to low statistics.	189

List of Tables

11.14	List of hyperparameters used in the NN training. The number of epochs corresponds to the maximum number allowed. The training is stopped if the loss does not improve after 4 epochs (patience). The choice of the values is based on the NN performance and the total training time.	190
11.15	NN input variables. For bb variables, both the pair with maximum p_T and minimum ΔR are included. Angular variables with one b and/or one B use the minimum ΔR pair. m_{bbbb} and m_{bbb} correspond to the combination with maximum p_T	190
11.16	List of experimental systematic uncertainties included in the analysis. "N" means that the uncertainty is taken as normalisation-only for all processes and channels affected, while "SN" means that the uncertainty is taken on both the shapes and the normalisation. Some of the systematic uncertainties are split into several components for a more accurate treatment. Additional systematic uncertainties due to the Ak8 track mass scale and the high- p_T DeXTer tagging, particular to this analysis, are shown separately with a "+" sign.	200
11.17	Correction factors applied to each of the $t\bar{t}$ +jets and tW samples. The first row in each sample corresponds to the nominal $t\bar{t}$ +HF reweighting, which is shown for comparison. Since all alternative samples are simulated in AFII, an AFII version of the nominal is included to be used as reference. The following abbreviations are used: "Ph" = POWHEGBOX, "Py8" = Pythia 8, "H7" = HERWIG 7, "MG" = MADGRAPH. Nominal samples (FS and AFII) are highlighted in grey.	203
11.18	List of modelling systematic uncertainties included in the analysis. "N" means that the uncertainty is taken as normalisation-only for all processes and channels affected, while "SN" means that the uncertainty is taken on both the shapes and the normalisation. Some of the systematic uncertainties are split into several components for a more accurate treatment. Systematics derived from alternative samples are not included.	205
11.19	Summary of alternative MC samples used in the $t\bar{t}a, a \rightarrow b\bar{b}$ analysis. Note that MADGRAPH is abbreviated to MG. All alternative samples are simulated using AFII with the exception of the pThard=1 alternative for $t\bar{t}+b\bar{b}$ and the DS alternative for tW . The AFII samples are compared to the AFII version of the nominal FS MC, highlighted in grey.	206
12.1	Signal and control regions used in the fit. All regions are required to have at least two Ak4 jets for the BDT variables to be well-defined. Note that b -jets tagged with the DL1r 70% WP are also tagged with the 85% WP.	208
12.2	Test categorisations of the $t\bar{t}+\geq 1b$ sample. In this context, " b " refers to an Ak4 jet with $p_T > 15$ GeV and $ \eta < 2.5$ matched to one b -hadron not originating from the $t\bar{t}$ system and " B " refers to an Ak4 jet with $p_T > 15$ GeV and $ \eta < 2.5$ matched to at least two b -hadrons not originating from the $t\bar{t}$ system.	209
12.3	Relative impacts over the fitted signal strength $\hat{\mu}$ obtained from the S+B fit to data. The statistical uncertainty (without MC stats.) is included for comparison.	226
12.4	Post-fit signal and background yields in the five signal regions for the $m_a = 30$ GeV hypothesis. The total MC and data yields are included in the last two rows.	227
A.5	L1Topo TOBs available in Run 2 and their granularity in E_T , η and ϕ for $ \eta < 2.5$	235
A.6	L1Topo TOBs available in Run 3 and their granularity in E_T , η and ϕ for $ \eta < 2.5$	235

A.7	Maximum number of TOBs received by the Run 2 L1Topo system.	236
A.8	Maximum number of TOBs received by the Run 3 L1Topo system.	236