
Novel Approaches to the Fast Simulation of the ATLAS Calorimeter
and
Performance Studies of Track-Assisted Reclustered Jets for Searches for
Resonant $X \rightarrow SH \rightarrow b\bar{b}WW^*$ Production with the ATLAS Detector

Dissertation

zur Erlangung des mathematisch-naturwissenschaftlichen Doktorgrades
„Doctor rerum naturalium“
der Georg-August-Universität Göttingen

im Promotionsprogramm Physik
der Georg-August University School of Science (GAUSS)

vorgelegt von

Joshua Falco Beirer

aus Schaffhausen

Göttingen, 2023

Betreuungsausschuss

Prof. Dr. Stanley Lai
Prof. Dr. Ariane Frey
Dr. Michael Dührssen-Debling

Mitglieder der Prüfungskommission:

Referent: Prof. Dr. Stanley Lai
II. Physikalisches Institut, Georg-August-Universität Göttingen

Koreferent: Dr. Michael Dührssen-Debling
Europäische Organisation für Kernforschung (CERN)

Weitere Mitglieder der Prüfungskommission:

PD Dr. Jörn Große-Knetter
II. Physikalisches Institut, Georg-August-Universität Göttingen

Prof. Dr. Hans Hofsäss
II. Physikalisches Institut, Georg-August-Universität Göttingen

Prof. Dr. Ariane Frey
II. Physikalisches Institut, Georg-August-Universität Göttingen

Prof. Dr. Steffen Schumann
Institut für Theoretische Physik, Georg-August-Universität Göttingen

Tag der mündlichen Prüfung: 12. Mai 2023

Referenz: II. Physik-UniGö-Diss-2023/01

Abstract

One of the most essential tools for all measurements and searches at the LHC experiments is the Monte Carlo (MC) simulation of proton-proton collisions. The generation of MC events, in particular the simulation of the detector response, is a very CPU-intensive process. Presently, the limited availability of MC statistics ranks among the most significant sources of systematic uncertainties in numerous ATLAS physics analyses. The primary bottleneck of the detector simulation is the detailed simulation of electromagnetic and hadronic showers in the ATLAS calorimeter system using GEANT4. To increase the number of available MC events, ATLAS has successfully employed a fast calorimeter simulation (FASTCALOSIM) during Run 1 and Run 2 of the LHC that reduces the simulation time by almost an order of magnitude. FASTCALOSIM parametrizes the energy response of particles in the calorimeter cells, taking into account the lateral shower profile and the correlation between the energy depositions in the various layers of the calorimeter.

This thesis presents a significantly improved version of FASTCALOSIM, which makes use of machine learning techniques such as principal component analysis. The new fast calorimeter package is named FASTCALOSIMV2 and is embedded in the state-of-the-art fast simulation software suite ATLFAST3, which was used to simulate about 7 billion events during the Run 2 MC reprocessing campaign in ATLAS. New developments for fast calorimeter simulation in ATLAS for Run 3 and the high-luminosity era of the LHC are also discussed. Among other aspects, new models for a precise data-driven fast simulation of electromagnetic showers are presented, as well as first efforts towards a major structural refactoring of the ATLAS simulation infrastructure, which is anticipated to greatly streamline the overall ATLAS simulation workflow in the coming years.

The second part of this work presents performance studies of Track-Assisted Reclustered (TAR) jets employed in searches for resonant Higgs boson pair production with the ATLAS detector. Many beyond the Standard Model extensions predict an additional massive scalar boson X that subsequently decays into two SM-like Higgs bosons ($X \rightarrow HH$) or into another spin-0 particle S in conjunction with a SM-like Higgs boson ($X \rightarrow SH$). The products of the subsequent scalar decays are heavily boosted and therefore cannot be resolved individually. Instead, the decay products are reconstructed as single large jets and information about their substructure is used to identify their origin. Traditional jet reconstruction algorithms rely solely on topological calorimeter clusters, which limits the resolution of jet substructure variables. TAR jets effectively overcome this limitation by exploiting angular information from the Inner Detector, and allow for a flexible choice of the reconstruction algorithm and jet size. This thesis presents performance studies targeting the $b\bar{b}WW^*$ decay mode that aim to find the optimal TAR jet configuration for two HH searches in the boosted 0ℓ and 1ℓ final states, and a SH search in the split-boosted 0ℓ final state.

This work has been sponsored by the Wolfgang Gentner Programme of the German Federal Ministry of Education and Research (grant no. 13E18CHA)

*Für meine wundervollen Eltern,
Heike und Michael*

Contents

I	Introduction	1
1	The Standard Model and Beyond	5
1.1	Particle Content	7
1.2	Principles and Symmetries	9
1.3	Gauge Interactions	12
1.3.1	Quantum Electrodynamics	13
1.3.2	Weak Interaction	13
1.3.3	Electroweak Unification	14
1.3.4	Quantum Chromodynamics	16
1.4	Electroweak Symmetry Breaking	17
1.5	Higgs Boson Production and Decay	21
1.6	Shortcomings of the Standard Model	23
1.7	Higgs Boson Pair Production	26
1.8	Scalar Sector Extensions	28
1.8.1	Two-Higgs-Doublet Model	29
1.8.2	Two-Real-Scalar-Singlet Model	32
2	The ATLAS Experiment at the LHC	35
2.1	The Large Hadron Collider	35
2.2	Proton-Proton Collisions	38
2.3	The ATLAS Experiment	40
2.3.1	Coordinate System	41
2.3.2	Inner Detector	42
2.3.3	Calorimeter System	44
2.3.4	Muon Spectrometer	48
2.3.5	Trigger System	50
3	Physics of Particle Shower Development	51
3.1	Particle Interactions with Matter	52

3.2	Development of Particle Showers	56
3.2.1	Electromagnetic Showers	56
3.2.2	Hadronic Showers	59
II	Fast Calorimeter Simulation in ATLAS	63
4	The AtlFast3 Simulator	67
4.1	Parametrization Input	68
4.2	Longitudinal Shower Development	74
4.3	Average Lateral Shower Shape	81
4.4	Simulation of Calorimeter Hits	84
4.5	Energy and Shape Interpolation	88
4.6	Computing Performance	89
5	Energy Response Corrections	93
5.1	Energy Response of Hadrons	94
5.2	Reconstructed Energy Response	102
6	Shower Centre Determination	107
6.1	ATLAS Track Extrapolation	108
6.2	Shower Centre Computation	110
6.2.1	Cylinder Intersections	111
6.2.2	Closest Point of Approach	113
6.2.3	Algorithmic Description and Choice of Cylinders	114
6.3	Validation and Limitations	119
7	Physics Performance of AtlFast3	125
7.1	Validation data sets	126
7.2	Performance	127
7.2.1	Electrons and Photons	127
7.2.2	Jets	129
7.2.3	Hadronic Tau Lepton Decays	132
7.2.4	Muons	135
8	Tuning of Electromagnetic Shower Properties to Data	137
8.1	Electromagnetic Shower Shapes	138
8.2	Selection of Photon Candidates	140
8.3	Tuning Model	143
8.3.1	Two-dimensional $\eta - \phi$ scaling model for middle layer shower shapes	145
8.3.2	Cubic η scaling model for strip layer shower shapes	148
8.3.3	Results	151

8.4	Extension of Model Coverage	152
8.4.1	Optimal Model Parameters	155
8.4.2	Model Parameter Interpolation	160
8.4.3	$H \rightarrow \gamma\gamma$ Model Validation	161
8.5	Model Application on Electron-induced Showers	163
8.6	Conclusion	169
9	FastCaloSim as a Geant4 Fast Simulation Model	175
9.1	Complexity Development	176
9.2	Fast Simulation in Geant4	180
9.3	Interfacing FastCaloSimV2 with Geant4	183
9.4	Choice of Trigger Volumes	185
9.5	Computational Performance	194
9.6	Outlook	196
III	Search for Resonant Higgs Boson Pair Production	199
10	Jets in Boosted Topologies	203
10.1	Jet Reconstruction	204
10.2	Jet Substructure	206
10.3	Track-Assisted Reclustered Jets	207
11	Track-Assisted Reclustered Jet Optimization	211
11.1	Methodology	211
11.2	Event Selection	213
11.3	Jet Tagging	215
11.3.1	$H \rightarrow WW^*$ tagging	215
11.3.2	$S \rightarrow W_{\text{had}}$ tagging	216
11.4	Jet Performance	218
11.4.1	Signal Event Efficiency	218
11.4.2	Top Quark Pair Event Rejection	221
11.4.3	Efficiency Ratios	222
11.5	Summary and Recommendations	224
11.6	Impact on Exclusion Limits	225
IV	Conclusion	227
A	Energy Response Corrections for Run 3	235
B	FastCaloSim in Geant4 Hit Validation	243
C	TAR Jet Tagging Performance	251

Introduction

Throughout history, curiosity has been a cornerstone of human progress, driving us to continually push the boundaries of our understanding and venture into uncharted territories of knowledge. The field of particle physics has played a crucial role in exploring some of the most fundamental aspects of our universe, revealing its underlying structure and the laws that govern its existence. Today, our best understanding of the fundamental constituents of the universe and their interactions is encoded in the Standard Model (SM) of particle physics, a theory that has been immensely successful in predicting the existence of many elementary particles and a wide range of physical phenomena. The Higgs boson, first observed in 2012 by the ATLAS and CMS collaborations at CERN, was the last particle to be discovered, completing the particle spectrum predicted by the SM.

Despite its enormous success, the SM leaves many questions unanswered and cannot account for all experimental observations. Among its most notable shortcomings are the exclusion of gravity and the inability to provide a satisfactory explanation for the compelling astrophysical and cosmological evidence supporting the existence of non-baryonic dark matter. Over the past few decades, numerous theories, such as supersymmetry, have been proposed that attempt to extend the SM to address some of its limitations. However, extensive searches for experimental signatures predicted by theories beyond the Standard Model (BSM) have failed to provide any direct evidence for the existence of new particles or forces.

To continue probing the SM at unprecedented energies, extensive research programmes at collider experiments will remain indispensable. In the coming years, the Large Hadron Collider (LHC) at CERN will undergo massive upgrades and deliver proton-proton collisions at record energies and instantaneous luminosities. To maintain or improve current levels of physics performance with the significantly increased data processing rates expected in the near future, novel developments and enhancements in the computing and software infrastructure of the experiments will be essential.

A crucial tool for all measurements and searches at the LHC experiments is the Monte Carlo (MC) simulation of proton-proton collisions. Generating MC events, in particular simulating

the detector response, is a very CPU-intensive process. Presently, the limited availability of MC statistics ranks among the most significant sources of systematic uncertainties in numerous ATLAS physics analyses. The primary bottleneck of the detector simulation is the detailed simulation of electromagnetic and hadronic showers in the ATLAS calorimeter system using GEANT4. To increase the amount of available MC events, ATLAS has successfully employed a fast calorimeter simulation (FASTCALOSIM) during Run 1 and Run 2 of the LHC that reduces the simulation time by almost an order of magnitude. FASTCALOSIM parametrizes the energy response of particles in the calorimeter cells, taking into account the lateral shower profile and the correlation between the energy depositions in the various layers of the calorimeter.

The main scope of this thesis is to present a number of novel developments in the fast simulation of the ATLAS calorimeter. In particular, a significantly improved version of FASTCALOSIM, which makes use of machine learning techniques such as principal component analysis and neural networks, is presented. The new fast calorimeter package is named FASTCALOSIMV2 and is embedded in the state-of-the-art fast simulation software suite ATLFAST3, which was used to simulate about 7 billion events during the Run 2 MC reprocessing campaign. This thesis also discusses new approaches for a precise data-driven simulation of electromagnetic showers and large structural refactorings in the ATLAS simulation infrastructure, which will become increasingly important in Run 3 and beyond.

In addition, this work also presents performance studies of Track-Assisted Reclustered (TAR) jets that have been conducted in the context of searches for resonant Higgs boson pair production with the ATLAS detector. Many BSM extensions predict an additional massive scalar boson X that subsequently decays into two SM-like Higgs bosons ($X \rightarrow HH$) or into another spin-0 particle S in conjunction with a SM-like Higgs boson ($X \rightarrow SH$). The products of the subsequent scalar decays are heavily boosted and therefore cannot be resolved individually. Instead, the decay products are reconstructed as single large jets and information about their substructure is used to identify their origin. Traditional jet reconstruction algorithms rely solely on topological calorimeter clusters, which limits the resolution of jet substructure variables. TAR jets effectively overcome this limitation by exploiting angular information from the Inner Detector, and allow for a flexible choice of the reconstruction algorithm and jet size. This thesis presents performance studies targeting the $b\bar{b}WW^*$ decay mode that aim to find the optimal TAR jet configuration for two HH searches in the boosted 0ℓ and 1ℓ final states, and a SH search in the split-boosted 0ℓ final state.

This thesis is divided into four parts:

Part I gives an introduction to the concepts that are relevant to this work. A brief overview of the Standard Model, its limitations and possible BSM extensions that predict resonant HH and SH production is given. This part also includes a description of the main components of the ATLAS detector and an introduction to the physics of particle shower development, which is particularly relevant to Part II of this work.

Part II focuses on fast simulation in ATLAS. It introduces the novel state-of-the-art fast

simulation tool ATLFast3 with particular emphasis on the simulation of the calorimeter with FASTCALOSIMV2. Among other aspects, this part discusses new algorithmic approaches for determining the centre of simulated showers within ATLFast3 and introduces newly developed corrections to improve the simulation of the energy response with FASTCALOSIMV2. This part also describes the development and optimization of a data-driven model incorporated in FASTCALOSIMV2, which substantially improves the simulation of electromagnetic shower properties in ATLAS and surpasses the performance of the detailed GEANT4 simulation. Lastly, this part showcases a first implementation of FASTCALOSIMV2 as a GEANT4 fast simulation model, marking the first step towards a significant refactoring of the simulation infrastructure anticipated to greatly streamline the overall ATLAS simulation workflow.

Part III presents an overview of jet reconstruction in boosted topologies with particular emphasis on TAR jets. This is followed by performance studies designed to find optimal TAR jet configurations for two searches for resonant $X \rightarrow HH \rightarrow b\bar{b}WW^*$ production in the boosted 0ℓ and 1ℓ final states and a search for $X \rightarrow SH \rightarrow b\bar{b}WW^*$ production in the split-boosted 0ℓ final state.

Part IV briefly summarizes the main results obtained in this thesis and presents an outlook for future developments in fast simulation and searches for resonant Higgs boson pair production in ATLAS.

The Standard Model and Beyond

DEVELOPED over more than a century of scientific research and experimentation, the Standard Model (SM) of particle physics is the most accurate theoretical framework for describing the fundamental building blocks of matter and their interactions that we have at our disposal today. Despite not including gravity, the SM encompasses all other fundamental forces of nature and its predictions have been extensively tested and confirmed through experiments across a wide range of energy scales.

The first discovery of the electron by J.J. Thomson in 1897 [1] marked the beginning of a series of groundbreaking discoveries and theoretical discussions that paved the way for the development of the SM as we know it today. In 1905, Einstein proposed the quantization of electromagnetic (EM) radiation in the form of photons as a result of the photoelectric effect [2]. Based on Einstein's interpretations, the field of Quantum Mechanics (QM), a term first coined by M. Born, W. Heisenberg, and W. Pauli [3], began to emerge in the 1920s. Around the same time, experiments conducted by E. Rutherford [4] led to the discovery of the proton and the atomic substructure.

As the foundations of QM were still being developed in 1926, the first Quantum Field Theory (QFT) emerged naturally from the study of EM interactions. Influenced by classical field theory, M. Born, W. Heisenberg, and P. Jordan developed a first QFT of the free EM field with its description as a set of quantum harmonic oscillators [5]. In 1928, Paul Dirac united the concepts of quantum mechanics with special relativity, leading to the prediction of antimatter [6], which was experimentally confirmed shortly thereafter [7]. The final breakthrough was accomplished with Fermi's description of the β decay [8] in 1933. For the first time, it was possible to describe the creation and annihilation of fermions as well as particle decays in a self-contained quantum field theory, which would later become the groundwork for today's theory of weak interactions.

Despite the early success of QFT, the initial enthusiasm was quickly halted by increasing theoretical hurdles arising from divergences of integrals found in perturbative corrections, leading to nonsensical results in the prediction of the most basic physical quantities. It was not until 1949, more than 15 years following the publication of Fermi’s theory, that the method of renormalization, a systematic procedure for handling problematic divergences, was developed for electromagnetic fields [9–22]. With the development of a systematic renormalization technique that enabled the calculation of theoretical predictions at any order of perturbation theory, Quantum Electrodynamics (QED) was established.

Driven by the invention of the bubble and spark chambers, and the availability of the first particle colliders, the time span between the 1950s and 1960s was rich in new discoveries of composite particles. A timeline of major particle discoveries illustrating the vast amount of particles detected during this period is given in Figure 1.1. Based on the schematic classification




Figure 1.1: Timeline of major particle discoveries from the discovery of the electron in 1897 until the discovery of the Higgs boson in 2012. The time span between the 50s and the 60s marked a period of many new particle discoveries, resulting in what is colloquially known as the *particle zoo*. Modified from Ref. [23].

of the newly discovered particles, Gell-Mann postulated the existence of three internal flavours known as *quarks* [24]. Each discovered particle was suggested to consist of either two or three quarks or anti-quarks that formed a *meson* or *baryon*, respectively. While in the years to come quarks were merely treated as mathematical necessities of flavour patterns, their existence as physical objects was experimentally verified in 1968 [25, 26].

Originally proposed by Yang and Mills in the 1950s [27], a non-abelian gauge theory known

as Quantum Chromodynamics (QCD) was developed to explain the strong interaction that binds quarks in compound states, which gained widespread acceptance after 't Hooft proved its renormalizability in 1973 [28, 29].

In 1961, a substantial advancement towards a unified description of nature was made by S. Glashow as he devised a unification of the electromagnetic and weak interactions into a single electroweak theory [30]. Only three years later, in 1964, the generation of gauge boson masses via electroweak symmetry breaking was formalized within the Brout-Englert-Higgs mechanism [31–33], which was shortly after incorporated into Glashow's electroweak theory [34, 35], shaping the SM as we know it today.

With the observation of the Z and W bosons [36, 37], the electroweak theory was confirmed in 1983. Finally, with the discovery of the Higgs boson [38, 39] in 2012 by the ATLAS and CMS collaborations at the Large Hadron Collider at CERN, the full range of elementary particles predicted by the SM was observed for the first time.

Despite the enormous success of the SM in the description of matter and its interactions, various questions, such as the nonrenormalizability of a QFT of gravity [40], remain open to the present day. While high energy physics continues to probe the most fundamental structures in our universe, multiple experiments have reported tensions to the prediction of the SM in recent years. Most notably, in 2021, a measurement of the muon anomalous magnetic dipole moment [41] revealed significant tensions to the prediction of the SM and a high-precision measurement of the W boson mass [42] showed similarly striking deviations in 2022.

This chapter offers a concise overview of the SM, and its limitations, as well as possible extensions. The particle content and interactions of the SM are described qualitatively in Section 1.1. A more rigorous description of the mathematical structure of the electroweak and strong gauge interactions follows in Section 1.2 and Section 1.3. The mechanism of electroweak symmetry breaking is of particular interest to this thesis and is discussed in detail in Section 1.4, while the main modes of Higgs boson production and decay in the SM are outlined in Section 1.5. The most relevant limitations and possible extensions of the SM are discussed in Section 1.6. Lastly, the chapter concludes by exploring Higgs boson pair production in the SM and beyond, and briefly discusses two interesting extensions of the scalar sector of the SM in Section 1.8, which give rise to resonant Higgs boson pair production and are of particular interest to this work.

1.1 Particle Content

The Standard Model describes all currently known fundamental particles and their interactions. Particles with half-integer spin are referred to as *fermions* and are classified into *leptons* and *quarks*. Each fermion comes in three generations with identical quantum numbers but distinct masses. The first, second and third generation of charged leptons consist of electrons, muons and τ -leptons, respectively. Each of them carries a negative electromagnetic charge

Generation	Leptons				Quarks			
	Particle	Charge [e]	Mass [GeV]	Particle	Charge [e]	Mass [GeV]		
First	electron neutrino	ν_e	0	$< 2 \times 10^{-9}$	up	u	$+\frac{2}{3}$	0.002
	electron	e^-	-1	0.0005	down	d	$-\frac{1}{3}$	0.005
Second	muon neutrino	ν_μ	0	$< 1.9 \times 10^{-4}$	charm	c	$+\frac{2}{3}$	1.27
	muon	μ^-	-1	0.1057	strange	s	$-\frac{1}{3}$	0.096
Third	tauon neutrino	ν_τ	0	$< 1.8 \times 10^{-2}$	top	t	$+\frac{2}{3}$	173.21
	tauon	τ^-	-1	1.777	bottom	b	$-\frac{1}{3}$	4.66

Table 1.1: Summary of the properties of the fermions in the Standard Model. The values are taken from Ref. [43].

and is associated to a neutral neutrino partner. In the case of quarks, the first, second and third generations consist of the *up-type* quarks u , c and t and the *down-type* quarks d , s and b . Up-type quarks possess a fractional electromagnetic charge of $+\frac{2}{3}$, whereas down-type quarks possess a fractional charge of $-\frac{1}{3}$. Finally, every fermionic particle in the SM has a corresponding antiparticle with identical properties, but opposite charge. The mass of fermions increases across generations. As a result, fermions of the third and second generations are short-lived and tend to decay into those of the first generation. A list of the fermions in the SM and a summary of their properties is shown in Table 1.1.

Interactions in the SM are mediated via vector bosons, that is, spin-1 particles, which act as force carriers of the electromagnetic, weak and strong interaction. With the exception of neutrinos, all fermions carry electromagnetic charge and are hence subject to electromagnetic interactions, which are mediated by massless photons.

Electroweak interactions are mediated by neutral Z and charged W^\pm bosons and affect all fermions. Within each generation of quarks and leptons, *left-handed* fermions, that is, fermions with negative chirality, are placed in isospin doublets. Up-type quarks are assigned a positive weak isospin of $I_3 = +\frac{1}{2}$, whereas down-type quarks have weak isospin of $-\frac{1}{2}$. Under the influence of the electroweak interaction, up-type quarks can be converted into down-type quarks and vice-versa through the exchange of a W^\pm boson. Similarly, electrons, muons and τ -leptons are placed into isospin doublets and have $I_3 = -\frac{1}{2}$, while their corresponding neutrinos have positive weak isospin. The corresponding anti-fermions have reversed chirality and opposite I_3 signs. Right-handed fermions and left-handed anti-fermions do not carry weak isospin and as such do not couple to W^\pm , but can still interact weakly via the neutral current that is mediated

by the Z boson.

Quarks are the only fermions that can interact via the strong force, which is mediated by gluons. Each type of quark takes one of three internal degrees of freedom, labelled as red, green or blue, a property known as *colour charge*. Gluons themselves come in an eightfold copy of colour charge combinations and can only couple to particles carrying colour charge. Similar to the electromagnetic charge and the weak isospin, antiquarks carry one of the anticolor states: anti-red, anti-green or anti-blue.

Finally, the Higgs boson is the only scalar particle predicted by the Standard Model and couples to all massive particles in the SM. It plays a crucial role in the generation of masses of the Z and W^\pm bosons as well as of fermions via the mechanism of electroweak symmetry breaking (see Section 1.4). A summary of all properties of the bosons in the Standard Model and their associated forces is shown in Table 1.2.

Force	Strength	Boson		J^P	Charge [e]	Mass [GeV]
Strong	1	Gluon ($\times 8$)	g	1^-	0	0
EM	$\sim 10^{-3}$	Photon	γ	1^-	0	0
Weak	$\sim 10^{-8}$	W	W^\pm	1^-	± 1	80.4
		Z	Z	1^-	0	91.2
		Higgs	H	0^+	0	125.1

Table 1.2: Summary of the properties of the bosons in the Standard Model. J^P denotes the spin J and the parity P of the particle. The strengths of the forces associated with each vector boson are given relative to the strength of the strong force and are approximate indicative values for two elementary particles at a distance of 1 fm. In comparison, gravity has a relative strength of 10^{-37} and is therefore multiple orders of magnitude weaker than the weak force. Values taken from Ref. [43].

1.2 Principles and Symmetries

As a quantum field theory, the Standard Model describes all fundamental objects as quantized fields $\Phi_i(x_\mu)$, which are defined at any arbitrary spacetime point x_μ . In particular, the fields described by the SM are the fermionic spin- $\frac{1}{2}$ fields ψ , the spin-1 electroweak boson fields W_μ^a ($a = 1, 2, 3$) and B , the spin-1 gluon fields G_μ^a ($a = 1, \dots, 8$) and the spin-0 Higgs field ϕ . The observable particles described in Section 1.1 emerge as excitations of these underlying fields.

In analogy to the Lagrangian formalism in classical field theory, the field equations that govern the evolution of quantized fields can be obtained by minimizing the Lagrangian density

$$\mathcal{L} = T - V, \quad (1.1)$$

where T is the kinetic and V the potential energy of the system. The Lagrangian can be typically decoupled into a component describing the dynamics of the free fields and another that characterizes the interactions between them, such that

$$\mathcal{L} = \mathcal{L}_{\text{free}} + \mathcal{L}_{\text{int}}. \quad (1.2)$$

Free Field Equations

The equations governing the dynamics of fields depend strongly on the spin of the field. A non-interacting spin- $\frac{1}{2}$ spinor field ψ describing a free fermion of mass m can be described with the Lagrangian

$$\mathcal{L}_D = i\bar{\psi}\gamma^\mu\partial_\mu\psi - m\bar{\psi}\psi, \quad (1.3)$$

where γ^μ are known as gamma matrices and are defined by the anticommutation relation $\{\gamma^\mu, \gamma^\nu\} = 2\eta^{\mu\nu}\mathbb{I}_4$. The corresponding equation of motion, known as Dirac equation [6], can be obtained with the Euler-Lagrange formalism and can be written as

$$i\gamma^\mu\partial_\mu\psi - m\psi = 0. \quad (1.4)$$

The Dirac equation and its adjoint version fully determine the dynamic evolution of free fermion fields ψ and their adjoint counterparts $\bar{\psi}$, which are associated to the corresponding anti-fermions.

Free spin-0 scalar particles are described as excitations of scalar fields for which the Lagrangian can be identified as

$$\mathcal{L}_S = \frac{1}{2}(\partial_\mu\phi)(\partial^\mu\phi) - \frac{1}{2}m^2\phi^2. \quad (1.5)$$

The minimization of the Lagrangian yields the Klein-Gordon equation [44, 45] and reads

$$\partial_\mu\partial^\mu\phi + m^2\phi = 0, \quad (1.6)$$

which is satisfied by all scalar fields ϕ associated with scalar particles of mass m .

Similarly, for a spin-1 field A_μ , the Lagrangian is given by

$$\mathcal{L}_{\text{Proca}} = -\frac{1}{4}F^{\mu\nu}F_{\mu\nu} + \frac{1}{2}m^2A^\nu A_\nu, \quad (1.7)$$

where m is the mass of the associated spin-1 boson and $F^{\mu\nu} = \partial^\mu A^\nu - \partial^\nu A^\mu$ is the electromagnetic field tensor. The equation of motion resulting from the minimization of the Lagrangian is

$$\partial_\mu F^{\mu\nu} + m^2A^\nu = 0, \quad (1.8)$$

and is known as the Proca equation [46].

Interaction Terms

For non-interacting fermionic and bosonic fields, that is, for terms $\subset \mathcal{L}_{\text{free}}$, analytical expressions that solve the equations of motion can be found trivially. The addition of terms $\subset \mathcal{L}_{\text{int}}$ describing the coupling between fields, such as interactions between photons and fermions render the situation considerably more intricate.

For all practical purposes, the probability of a general interaction of two incoming and n outgoing particles $1 + 2 \rightarrow 3 + 4 + \dots + n$ can be computed with Fermi's golden rule in terms of the cross section

$$\sigma = \frac{S}{4\sqrt{(p_1 p_2)^2 - (m_1 m_2)^2}} \int |\mathcal{M}_{fi}|^2 (2\pi)^4 \delta(p_1 + p_2 - p_3 \cdots - p_n) \times \prod_{j=3}^n 2\pi \delta(p_j^2 - m_j^2) H(p_j^0) \frac{d^4 p_j}{(2\pi)^4}, \quad (1.9)$$

where S is a constant statistical factor to avoid double-counting identical particles, p_i the four-momenta and m_i the masses of the incoming and outgoing particles. The term $\delta(p_1 + p_2 - p_3 \cdots - p_n)$ ensures energy-momentum conservation, while $\delta(p_j^2 - m_j^2)$ guarantees that all outgoing particles lie on their respective mass shell. The Heaviside step function $H(p_j^0)$ ensures that outgoing particles always have positive energies.

The *matrix element*, often referred to as S -matrix, reads

$$\mathcal{M}_{fi} = \lim_{t_2 \rightarrow +\infty} \lim_{t_1 \rightarrow -\infty} \langle \Phi_f | U(t_2, t_1) | \Phi_i \rangle, \quad (1.10)$$

and encodes the coupling of the initial state $|\Phi_i\rangle$ to the set of final states $|\Phi_f\rangle$. Here, U is known as the time-evolution operator that acts on the initial state and is only dependent on the interaction terms of the Lagrangian density. Exact computations of the matrix element are generally not possible. Instead, the S -matrix can be evaluated by perturbative expansion of the time-evolution operator, resulting in what is known as Dyson series [47]. Each term in the series can be represented by a sum of Feynman diagrams [48], a graphical representation of the underlying physical process. Feynman diagrams correspond to uniquely defined mathematical expressions that immensely simplify the calculation of physical observables.

The perturbative expansion of the matrix element into an infinite series inevitably leads to divergent terms. Systematic regularization and renormalization procedures [49] are required to avoid infinite nonsensical quantities and obtain meaningful physical predictions from the SM. Regularization introduces a cut-off energy scale Λ beyond which the perturbative methods

become incorrect and render the SM invalid. The value of the cut-off scale is not predicted by the SM, but must be below the energy scale at which effects of gravity become relevant, that is, $\Lambda \lesssim \Lambda_P = 1.22 \cdot 10^{19}$ GeV, where Λ_P is known as the Planck scale [50]. The divergent integrals can then be absorbed into measurable quantities by renormalizing the coupling constants and masses that appear as part of the perturbative expansion. The renormalization procedure introduces a dependence on the energy scale of the couplings, which has important physical consequences, and is especially relevant for the strong interaction (see Section 1.3.4 for more details).

Gauge Symmetries

A cornerstone concept in the mathematical formulation of the Standard Model is the inherent inclusion of physical conservation laws in terms of symmetries. Symmetries are incorporated by requiring invariance of the Lagrangian density under certain transformations belonging to Lie groups. To be precise, the Lagrangian must remain unchanged under

$$\Phi \rightarrow \Phi' = g \cdot \Phi \quad g \in G \quad (1.11)$$

transformations of the fields, where g is an arbitrary element of the Lie group G . Any group element can be represented as $g = e^{i\theta_a T^a}$, where T^a are known as the generators of the group G and θ_a are phase parameters of the group. If the parameters θ_a are space-time dependent $\theta_a(x^\mu)$, the Lagrangian exhibits a *local gauge invariance*, while a constant phase results in a *global gauge invariance*.

According to Noether's theorem [51], the exhibition of such continuous symmetries automatically implies the existence of corresponding quantities that are conserved in time. For instance, the SM Lagrangian exhibits invariance under global transformations of the Poincaré group $\mathbf{R}^{1,3} \times \text{SL}(2, \mathbf{C})$ [52], which encompasses the full symmetry of special relativity and implies the conservation of energy and momentum.

1.3 Gauge Interactions

As a gauge theory, the SM describes electroweak and strong interactions with fields that preserve local gauge invariance under

$$\underbrace{\text{SU}(3)_c}_{\text{strong}} \times \underbrace{\text{SU}(2)_L \times \text{U}(1)_Y}_{\text{electroweak}}$$

symmetry transformations. The Lagrangian's presented in Equations (1.3), (1.5) and (1.7) for free fields do not exhibit local gauge invariance, but can be suitably modified to conform

with this principle. The resulting Lagrangian densities and their physical implications for the electroweak and strong interactions are briefly described in the following.

1.3.1 Quantum Electrodynamics

Quantum Electrodynamics (QED) is an abelian gauge theory that imposes local gauge invariance under U(1) symmetry transformations. The Lagrangian density describing QED must entail terms that characterize the propagation of free spin- $\frac{1}{2}$ fermions and massless spin-1 photons, as well as their interactions. Such a Lagrangian can be constructed by combining the Dirac Lagrangian from Equation (1.3) and the massless Proca Lagrangian from Equation (1.7):

$$\mathcal{L}_{\text{EM}}^{\text{free}} = \mathcal{L}_D + \mathcal{L}_{\text{Proca}}^{m=0}. \quad (1.12)$$

The resulting free Lagrangian does not satisfy U(1) local gauge invariance. However, invariance can be easily established by replacing the derivative ∂_μ with the *covariant* derivative \mathcal{D}_μ , that is,

$$\partial_\mu \rightarrow \mathcal{D}_\mu = \partial_\mu + ig_e A_\mu \quad (1.13)$$

The replacement results in the QED Lagrangian

$$\mathcal{L}_{\text{EM}} = \mathcal{L}_{\text{EM}}^{\text{free}} + g_e \bar{\psi} \gamma^\mu \psi A_\mu, \quad (1.14)$$

where the covariant derivative introduces a term describing the interaction between photon fields A_μ and fermionic fields ψ , and g_e is a constant that can be identified as the coupling strength between the spin- $\frac{1}{2}$ and massless spin-1 fields.

1.3.2 Weak Interaction

Weak interactions in the Standard Model are described by requiring local gauge invariance under $SU(2)_L$ transformations,

$$\varphi(x) \rightarrow \varphi'(x) = e^{i\frac{g_W}{2}\theta(x)\cdot\sigma} \varphi(x), \quad (1.15)$$

where the 2×2 complex Pauli matrices σ are a representation of the generators of the $SU(2)_L$ group and g_W the weak coupling constant.

Local gauge invariance can only be satisfied by the introduction of three gauge fields W_μ^a ($a = 1, 2, 3$), where the physical W^\pm bosons can be associated with the linear combinations

$$W_\mu^\pm = \frac{1}{\sqrt{2}} \left(W_\mu^1 \mp iW_\mu^2 \right) \quad (1.16)$$

of the underlying fields. The wave functions from Equation (1.15) take the form of isospin doublets

$$\begin{pmatrix} u \\ d' \end{pmatrix}_L \quad \begin{pmatrix} \nu_\ell \\ \ell \end{pmatrix}_L,$$

which group together fermion flavours differing by unit charge that can participate with interactions mediated by W^\pm bosons. For quarks, this corresponds to the up-type quarks with weak isospin $I_3 = +\frac{1}{2}$ and the down-type quarks with weak isospin $I_3 = -\frac{1}{2}$. Similarly, leptons ℓ have negative half-integer weak isospin, while the associated neutrino isospin is flipped in sign.

Only left-handed (LH) chiral fermion states and right-handed (RH) chiral anti-fermion states interact via the charged weak interaction. As such, and as indicated by the index L , only LH particles and RH anti-particles are placed in isospin doublets. In turn, RH chiral particle and LH chiral anti-particle states are placed in isospin singlets with weak isospin $I_3 = 0$ and consequently remain unaffected by local $SU(2)_L$ gauge transformations.

The weak isospin doublets are constructed by taking into account an important mismatch between the fermion states that participate in the weak interaction, known as *weak eigenstates*, and the observable *mass eigenstates*. The mismatch can be parametrized with the CKM matrix [53] in the case of quarks and with the PMNS matrix [54] in the case of leptons, and has important physical consequences. Among other aspects, the PMNS matrix explains the observed spontaneous oscillation between neutrino flavours [55], while both, CKM and PMNS matrices allow for direct sources of $C\mathcal{P}$ violation [56].

In analogy to Equation (1.14) from QED, the requirement of local gauge invariance gives rise to an additional interaction term between the three gauge fields W_μ^a ($a = 1, 2, 3$) and the LH isospin doublets φ_L that reads

$$ig_W \frac{1}{2} \sigma_k \gamma^\mu \mathbf{W}_\mu^k \varphi_L \quad (1.17)$$

and induces three weak currents in the SM. The charged current (CC) Lagrangian that describes the interaction between W^\pm bosons and fermions then takes the form

$$\mathcal{L}_{CC} = \frac{g_W}{\sqrt{2}} \left\{ W_\mu^\dagger \left[\bar{u} \gamma^\mu \hat{P}_L d + \bar{\nu}_\ell \gamma^\mu \hat{P}_L \ell \right] + \text{h.c.} \right\}, \quad (1.18)$$

where $\hat{P}_L = \frac{1}{2} (1 - \gamma_5)$ is the left-handed projection operator and h.c. denotes the corresponding Hermitian conjugate.

1.3.3 Electroweak Unification

In addition to charged current interactions mediated by the W^\pm bosons, the $SU(2)_L$ gauge invariance of the weak interaction induces a neutral current associated to the third gauge field

W_μ^3 . Albeit with differing strengths, the physical Z boson in the SM is observed to couple to RH particles and LH anti-particles, and hence cannot be directly associated with the W_μ^3 field. Instead, the situation is resolved with a unified description of electromagnetic and weak forces. More precisely, the $U(1)$ gauge symmetry from QED that is associated with the electric charge is replaced with a local $U(1)_Y$ gauge symmetry that generates a new electroweak gauge field B_μ , which couples to the *hypercharge*

$$Y = 2 \cdot (Q - I_3), \quad (1.19)$$

and unifies the fundamental properties of electromagnetic charge and weak isospin.

The observable photon and Z boson fields are then the result of a linear *mixing* of the underlying electroweak fields B_μ and W_μ^3 , and can be parametrized according to

$$\begin{pmatrix} A \\ Z \end{pmatrix} = \begin{pmatrix} \cos \theta_W & \sin \theta_W \\ -\sin \theta_W & \cos \theta_W \end{pmatrix} \begin{pmatrix} B \\ W^3 \end{pmatrix}, \quad (1.20)$$

where θ_W is a free parameter of the SM and known as the weak mixing or Weinberg angle.

The underlying Lagrangian density describing electroweak interactions must be invariant under $SU(2)_L \times U(1)_Y$ local gauge transformations which can be accomplished with the covariant derivative

$$\partial_\mu \rightarrow \mathcal{D}_\mu = \partial_\mu - i \frac{g_Z}{2} Y B_\mu - i \frac{g_W}{2} \sigma_j W_\mu^j, \quad (1.21)$$

where g_Z is the coupling constant corresponding to the weak hypercharge and is directly related to g_W through the weak mixing angle $g_W = g_Z \cos \theta_W$. The resulting local gauge invariant neutral current (NC) Lagrangian describing the interaction between Z bosons and fermions can be expressed as

$$\mathcal{L}_{\text{NC}} = g_Z \bar{f} \gamma_\mu \frac{1}{2} (c_V^f - c_A^f \gamma_5) f Z^\mu, \quad (1.22)$$

where c_V and c_A are known as vector and axial-vector couplings of the fermion f and are directly related to the weak isospin and electromagnetic charge of the particle according to

$$\begin{aligned} c_V^f &= I_3^f - 2 Q^f \sin^2 \theta_w \\ c_A^f &= I_3^f. \end{aligned} \quad (1.23)$$

As experimentally observed, the different vector and axial-vector couplings imply different coupling strengths to LH and RH chiral fermion states.

As a non-abelian gauge theory, the generators of the $SU(2)_L$ group are non-commutative and lead to interdependent transformation properties of the underlying gauge fields that spoil the necessary local gauge invariance. However, invariance can be recovered by imposing the correct

transformation behaviour of the gauge fields under local $SU(2)_L$ gauge transformation

$$W_k^\mu \rightarrow W_k'^\mu = W_k^\mu - \partial^\mu \theta_k - g_W \epsilon_{ijk} \theta_i W_j^\mu, \quad (1.24)$$

where ϵ_{ijk} is the totally antisymmetric Levi–Civita tensor. The required transformation properties give rise to triple and quartic gauge boson self-couplings of the electroweak mass eigenstates.

It is to be emphasized that the electroweak Lagrangian densities considered thus far do not encompass mass terms such as those appearing in the free field equations in Section 1.2. The addition of terms accounting for fermion and gauge boson masses unavoidably breaks the necessary gauge invariance. As the electroweak gauge bosons W^\pm and Z as well as fermions are experimentally observed to be massive, an alternative mechanism of mass generation is needed and provided by the Brout-Englert-Higgs mechanism (see Section 1.4).

1.3.4 Quantum Chromodynamics

Quantum Chromodynamics (QCD) is a non-abelian gauge theory that describes the strong interaction between quarks and gluons by imposing local gauge invariance under $SU(3)_c$ symmetry transformations. The subscript refers to the associated *colour charge* of the strong interaction, where each quark carries one of the discrete quantum states: red, green or blue.

In analogy to Equation (1.14) from QED, the QCD Lagrangian density describing the free field propagation of the gluon fields G_μ^a ($a = 1, \dots, 8$), the quark field q_i of colour i ($i = 1, 2, 3$) and mass m_q , and the interactions thereof reads

$$\mathcal{L}_{\text{QCD}} = \bar{q}_i \left\{ i\gamma^\mu \left(D_\mu \right)_{ij} - m_q \delta_{ij} \right\} q_j - \frac{1}{4} G_{\mu\nu}^a G_a^{\mu\nu}, \quad (1.25)$$

where

$$G_{\mu\nu}^a = \partial_\mu G_\nu^a - \partial_\nu G_\mu^a + g_s f^{abc} G_\mu^b G_\nu^c \quad (1.26)$$

is the gluon field strength tensor and the derivative from the free spin-1 field equation is replaced with the covariant derivative

$$\left(D_\mu \right)_{ij} = \partial_\mu \delta_{ij} - ig_s (T_a)_{ij} G_\mu^a \quad (1.27)$$

to ensure invariance of the Lagrangian under local $SU(3)_c$ gauge transformations. The covariant derivative couples the gluon fields with the quark field via the strong coupling constant g_s , where T_a ($a = 1, \dots, 8$) are the generators of the $SU(3)_c$ group that correspond to one of the eight massless physical gluon states, each carrying a distinct colour charge combination. An explicit representation of the generators is given through the 3×3 Gell-Mann matrices λ_a [57], where $T_a = \frac{1}{2}\lambda_a$. The *structure constants* f^{abc} are an expression of the non-abelian nature of

the $SU(3)_c$ group and directly related to the commutation relations of the group generators. In analogy to Equation (1.24) from the theory of weak interactions, the gluon fields must transform according to

$$G_\mu^a \rightarrow G_\mu^{a'} = G_\mu^a - \partial_\mu \theta_a - g_s f_{abc} \theta_a G_\mu^b \quad (1.28)$$

in order to preserve local gauge invariance of the underlying Lagrangian density. As such, the non-vanishing structure constants give rise to triple and quartic gluon self-interactions.

Within the SM, the strong interaction is unique in multiple aspects. In contrast to the electroweak force, the strong coupling constant of QCD is large and takes values in the order of unity at low-energy scales, thus precluding a perturbative expansion as discussed in Section 1.2. However, at large momentum transfers of $O(100 \text{ GeV})$, the renormalized strong coupling constant becomes sufficiently small such that a perturbative treatment is feasible. This property of QCD is referred to as *asymptotic freedom* and allows for first-principle computation of observables involving the strong interaction at high-energies, but generally requires phenomenological models at low momentum transfers.

While the existence of quarks has been experimentally verified on many occasions, fractionally charged particles have never been observed directly. A free propagation of quarks is thought to be impossible as a result of gluon-gluon self-interactions that give rise to a phenomenon known as *colour confinement* that only permits the existence of quarks in colourless singlet states. As a low-energy phenomenon in the non-perturbative regime, colour confinement is a phenomenological concept that is currently not accessible with analytical computations. Qualitatively, two quarks are confined in a colourless state and bound by the exchange of virtual gluons that form a *flux tube*. By adding energy to the bound state, the quarks can be separated from each other, effectively leading to an elongation of the flux tube. However, the energy stored in the gluon field is proportional to the distance between the quarks, such that, at a certain distance threshold, the creation of a new $q\bar{q}$ pair is energetically more favourable than a further expansion of the tube. The process of forming *hadrons*, that is, colourless bound states of quarks, is known as *hadronization* and plays an important role in proton-proton collisions.

1.4 Electroweak Symmetry Breaking

The QCD Lagrangian from Equation (1.25) respects invariance under local $SU(3)_c$ gauge transformations and incorporates mass terms that allow for an explanation of quark masses, while no terms are required to be added to account for mediating gauge bosons, as gluons are observed to be massless. In contrast, the necessary addition of mass terms of the massive Z and W bosons in the electroweak sector spoils the invariance under $SU(2)_L \times U(1)_Y$ transformations. Moreover, the addition of fermionic mass terms (including mass terms for quarks) would imply a mixture of left-handed $SU(2)$ doublets and right-handed singlets that would similarly break local gauge invariance. As a consequence, the joint electroweak and QCD Lagrangian, which must be invariant under $SU(3)_c \times SU(2)_L \times U(1)_Y$ gauge transformation, is unable

to describe the origin of fermionic and gauge boson masses. First developed in the 1960s, the Brout-Englert-Higgs (BEH) mechanism [31–33] provides a natural explanation of the acquisition of particle masses in the SM through electroweak symmetry breaking (EWSB) that is described in the following.

The underlying core concept of EWSB is the observation that the SM Lagrangian must be invariant under local gauge transformations, but specific states are not bound by the same restrictions. Symmetries imposed by gauge groups can therefore be *spontaneously broken* if a system falls into such a state. EWSB in the SM is realized by the addition of a charged ϕ^+ and neutral ϕ^0 complex scalar field placed in a SU(2) isospin doublet

$$\Phi = \begin{pmatrix} \phi^+ \\ \phi^0 \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} \phi_1 + i\phi_2 \\ \phi_3 + i\phi_4 \end{pmatrix}, \quad (1.29)$$

where the doublet field Φ takes the hypercharge $Y_\Phi = 1$ and is known as the *Higgs field*. The addition of the Higgs field leads to an additional component of the SM Lagrangian density and reads

$$\mathcal{L}_H = (D_\mu \Phi)^\dagger (D^\mu \Phi) - V(\Phi), \quad (1.30)$$

where the covariant D_μ ensures $SU(2)_L \times U(1)_Y$ local gauge invariance and takes the form

$$D_\mu = \partial_\mu - ig_W \frac{\sigma^a}{2} W_\mu^a - ig' \frac{Y_\Phi}{2} B_\mu. \quad (1.31)$$

The term $V(\Phi)$ is known as the Higgs potential, which in its most general form can be expressed as

$$V(\Phi) = \mu^2 \Phi^\dagger \Phi + \lambda (\Phi^\dagger \Phi)^2, \quad (1.32)$$

where $\lambda > 0$ and μ^2 are free parameters of the potential. The restriction of λ to positive values ensures a lower bound of the potential, which is necessary to guarantee a stable vacuum state, whereas the sign of μ^2 defines the form of the potential. In the case that $\mu^2 > 0$, the potential exhibits a stable minimum at $\phi_i = 0$, such that the $SU(2)_L \times U(1)_Y$ symmetry remains unbroken. If instead $\mu^2 < 0$, the potential leads to an infinite set of non-zero stable minima that satisfy

$$\sum_i \phi_i^2 = \frac{-\mu^2}{\lambda} \equiv v^2, \quad (1.33)$$

where v is referred to as vacuum expectation value (VEV) of the Higgs field. In this case, the doublet acquires one of the infinite degenerate states and the $SU(2)_L \times U(1)_Y$ symmetry is spontaneously broken. By insertion of the ground state

$$\Phi_0 = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v \end{pmatrix} \quad (1.34)$$

into Equation (1.30), it can be easily shown that the Lagrangian density of the Higgs doublet is invariant under local $U(1)_{\text{EM}}$ gauge transformation. As such, the $SU(2)_L \times U(1)_Y$ symmetry is spontaneously broken into the electromagnetic subgroup $U(1)_{\text{EM}}$

$$SU(2)_L \times U(1)_Y \xrightarrow{\text{EWSB}} U(1)_{\text{EM}}, \quad (1.35)$$

which, by construction, remains a symmetry of the vacuum state. A sketch of a simplified one-dimensional complex Higgs potential illustrating the form of the Higgs potential and its implication for EWSB is shown in Figure 1.2.

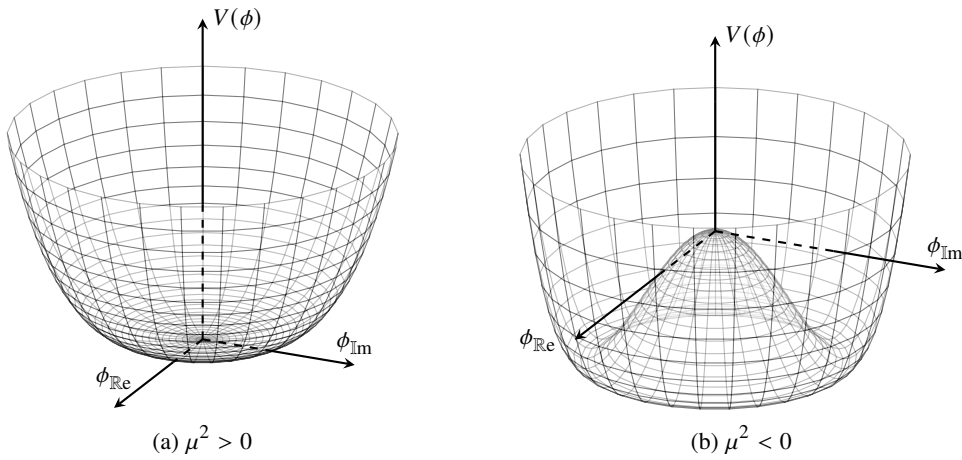


Figure 1.2: Sketch of a simplified Higgs potential $V(\phi)$ for a complex scalar field ϕ in the case where (a) a single stable vacuum state can be identified and in the case that (b) the choice of one of the infinite set of vacuum states leads to electroweak symmetry breaking.

As predicted by Goldstone's theorem [58], electroweak symmetry breaking implies the existence of one additional scalar particle for each generator of the broken symmetry, that is, one massive and three massless scalar bosons. As a result of the local $SU(2)_L$ gauge invariance, the massless Goldstone bosons can be absorbed into a *unitary gauge*, where the Higgs doublet Φ is written as an expansion of the complex fields around the vacuum expectation value

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v + H \end{pmatrix}, \quad (1.36)$$

where H is the elementary scalar field associated with the Higgs boson of the SM. With the insertion of the Higgs doublet in unitary gauge into the Lagrangian from Equation (1.30), and

the identification of the gauge fields as linear combinations of the underlying fields

$$W_\mu^\pm = \frac{W_\mu^1 \mp iW_\mu^2}{\sqrt{2}} \quad Z_\mu = \frac{g_W W_\mu^3 - g' B_\mu}{\sqrt{g_W^2 + g'^2}} \quad A_\mu = \frac{g' W_\mu^3 + g_W B_\mu}{\sqrt{g_W^2 + g'^2}}, \quad (1.37)$$

the mass terms of the gauge fields emerge naturally after electroweak symmetry breaking. While photons remain massless, the gauge bosons acquire the masses

$$m_H = \sqrt{2\lambda v^2} \quad m_W = \frac{1}{2}g_W v \quad m_Z = \frac{1}{2}v\sqrt{g_W^2 + g'^2}. \quad (1.38)$$

The resulting Lagrangian additionally gives rise to triple VVH and quartic $VVHH$ couplings (where V is either W^\pm or Z) to the massive gauge bosons, where the coupling is proportional to the masses of the electroweak mediators. Moreover, the Lagrangian includes terms that allow for triple and quartic self-couplings that are directly related to the Higgs boson mass and the VEV

$$\lambda_{HHH} = \frac{m_H^2}{2v} \quad \lambda_{HHHH} = \frac{m_H^2}{8v^2}. \quad (1.39)$$

The BEH mechanism is fully constrained with the two parameters μ^2 and λ that define the shape of the Higgs potential and the two coupling constants g_W and g' , which have all been experimentally measured. In particular, the VEV of the Higgs field is determined [43] to be

$$v = \frac{2m_W}{g_W} \approx 246.22 \text{ GeV}. \quad (1.40)$$

As opposed to the mass generation of gauge bosons, no mass terms for fermions arise naturally from the BEH mechanism. The addition of explicit mass terms to the Lagrangian would spoil the $SU(2)_L \times U(1)_Y$ gauge invariance as a result of the differing transformation behaviour of left-handed doublets and right-handed singlets. However, the Lagrangian can be expanded with Yukawa terms of the form

$$\mathcal{L}_{\text{Yukawa}} = -g_f \left(\bar{\psi}_L \Phi \psi_R + \bar{\psi}_R \Phi^\dagger \psi_L \right), \quad (1.41)$$

where g_f is known as the Yukawa coupling. After EWSB, the Lagrangian includes mass terms that explain the origin of fermionic masses and take the form

$$\mathcal{L}_{\text{Yukawa}} = -m_f \bar{\psi} \psi - \frac{m_f}{v} \bar{\psi} \psi H, \quad (1.42)$$

where the right-hand term additionally gives rise to a direct coupling between the Higgs boson

and the leptons and quarks in the SM that is proportional to the fermion mass

$$g_f = \frac{m_f}{v} \sqrt{2}. \quad (1.43)$$

1.5 Higgs Boson Production and Decay

The phenomenological predictions of the SM Higgs boson production and decay rates are driven by the strength of the gauge boson and Yukawa coupling constants to the Higgs boson. While in principle the Higgs boson can directly interact with all SM fermions and massive gauge bosons, it preferably couples to heavy particles.

The Higgs boson production mechanisms and rates vary depending on the available initial-state particles and centre-of-mass energy. The four leading production mechanisms for a pp collider operating at $\sqrt{s} = 13$ TeV, such as the LHC, are depicted as Feynman diagrams in Figure 1.3. The leading mechanism of Higgs boson production in this case is gluon-gluon

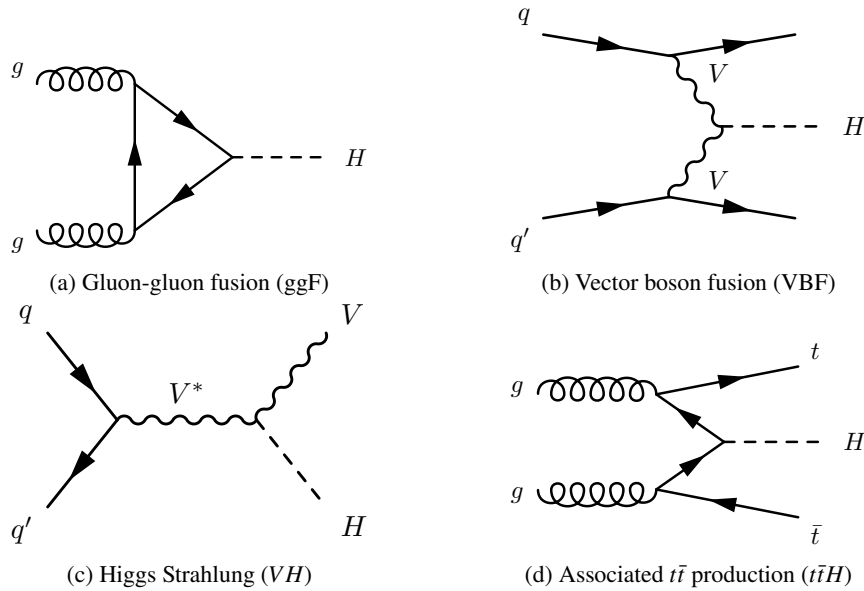


Figure 1.3: Feynman diagrams for the production of a SM Higgs boson via gluon-gluon fusion, vector boson fusion and the two leading associated production modes VH and $t\bar{t}H$.

fusion (ggF), where two initial-state gluons form a Higgs boson via a top-quark loop, which is favourable as a result of the large mass of the top quark. Contributions from loops of other quark flavours are also possible, but suppressed proportionally to m_q^2 . The predicted cross section is known at N^3 LO QCD and NLO EW precision and is $\sigma_{\text{ggF}}^{\text{N}^3\text{LO}} = 48.6 \pm 2.4$ [43] at $\sqrt{s} = 13$ TeV. The second-leading production mechanism, which is suppressed by a full order

of magnitude with respect to gluon-gluon fusion, is vector boson fusion (VBF), where two quarks scatter via the exchange of one of the weak mediators W^\pm or Z that subsequently form a Higgs boson. The predicted cross section is known at NNLO QCD and NLO EW precision and is $\sigma_{\text{VBF}}^{\text{NNLO}} = 3.78 \pm 0.08$ [43] at $\sqrt{s} = 13$ TeV. The next most relevant Higgs boson production mechanism is through the associated production of W^\pm and Z bosons, also referred to as Higgs Strahlung (VH). In this case, two quarks fuse to a virtual weak boson that then decays via simultaneous radiation of a Higgs boson. Finally, the fourth-leading production mechanism is through the associated production of a $t\bar{t}$ pair ($t\bar{t}H$). Two initial-state gluons each decay into a $t\bar{t}$ pair, which is followed by the combination of a top quark and anti-top quark pair to form the SM Higgs boson. Associated $t\bar{t}$ production is of particular relevance as it allows to gain a direct probe of the top-Higgs Yukawa coupling. The cross sections of VH and $t\bar{t}H$ are also known at NNLO QCD and NLO EW precision and are provided in Table 1.3, along with the values for less relevant production mechanisms such as for the associated production with a $b\bar{b}$ pair and a single top quark.

Production Mode	σ @ 13 TeV [pb]	Decay channel	$\mathcal{B}r$ [%]
ggF ($gg \rightarrow H$)	48.6 ± 2.4	$H \rightarrow b\bar{b}$	58.2 ± 0.7
VBF ($qq' \rightarrow Hqq'$)	3.78 ± 0.08	$H \rightarrow WW^*$	21.4 ± 0.3
WH ($qq' \rightarrow WH$)	1.373 ± 0.028	$H \rightarrow gg$	8.53 ± 0.85
ZH ($qq/gg \rightarrow ZH$)	0.88 ± 0.04	$H \rightarrow \tau\tau$	6.27 ± 0.10
$t\bar{t}H$ ($qq/gg \rightarrow t\bar{t}H$)	0.51 ± 0.05	$H \rightarrow c\bar{c}$	2.89 ± 0.24
$b\bar{b}H$ ($qq/gg \rightarrow b\bar{b}H$)	0.49 ± 0.12	$H \rightarrow ZZ^*$	2.62 ± 0.04
tH ($qq/gg \rightarrow tH$)	0.09 ± 0.01	$H \rightarrow \gamma\gamma$	0.227 ± 0.005
		$H \rightarrow Z\gamma$	0.153 ± 0.009
Total	55.1 ± 3.3	$H \rightarrow \mu\mu$	0.022 ± 0.001

Table 1.3: Predicted SM Higgs boson cross sections for gluon-gluon fusion (ggF), vector boson fusion (VBF) and the leading five associated production modes in pp collisions at $\sqrt{s} = 13$ TeV as well as the predicted Higgs boson branching ratios for the various decay channels. The values are taken from Ref. [43, 59, 60].

The Higgs boson is an unstable particle with an extremely short lifetime of $O(10^{-22})$ s and subsequently decays to fermions or gauge bosons. The leading Feynman diagrams for the decay of the Higgs boson are shown in Figure 1.4. Despite that the top quark exhibits the strongest coupling with the Higgs boson, a direct $H \rightarrow t\bar{t}$ decay is kinematically not allowed as $m_H < 2m_{t\bar{t}}$, which would spoil energy-momentum conservation. Instead, the largest branching ratio is the decay into a $b\bar{b}$ pair with a predicted branching ratio of $\mathcal{B}r(H \rightarrow b\bar{b}) = 58.2 \pm 0.7\%$ [43], followed by the decay into a W^\pm pair with $\mathcal{B}r(H \rightarrow WW^*) = 21.4 \pm 0.3\%$ [43]. The bottom quark is the heaviest fermion for which a decay into a $f\bar{f}$ pair is kinematically allowed. The weak W^\pm and Z gauge bosons are significantly heavier, but direct decays are kinematically




Figure 1.4: Leading Feynman diagrams for the direct decay of the SM Higgs boson into (a) fermions, (b) massive bosons as well as for (c) the indirect decay via a top loop to massless gluons and photons.

only permitted if one of the bosons is *virtual* or *off-shell*, that is, temporarily does not satisfy the energy–momentum relation. Off-shell decays are quantum mechanically suppressed, but still significant as a result of the high masses of the weak gauge bosons. As the Higgs boson exclusively couples to massive particles, direct decays into photons and gluons are not possible. Instead, the decay to massless gauge bosons can occur through intermediate top quark loops, leading to the third-leading branching ratio $\mathcal{B}r(H \rightarrow gg) = 8.53 \pm 0.85\%$ [43]. The predicted branching ratios for further decay modes of the SM Higgs boson are provided in Table 1.3.

In 2012, the ATLAS and CMS collaborations announced the discovery [38, 39] of a scalar particle consistent with the SM Higgs boson. In the years following the discovery, the properties of the newly discovered particle were thoroughly scrutinized, including the confirmation of its scalar nature and even positive parity [61] as well as through measurements of the Higgs bosons production and decay rates [62], among other aspects. Recent measurements of the production cross sections and decay rates of the Higgs boson at $\sqrt{s} = 13$ TeV is shown in Figure 1.5. All measured quantities are compatible with the predictions of the SM. The current most precise world average measurement of the Higgs boson mass is $m_H = 125.25 \pm 0.17$ GeV [43].

1.6 Shortcomings of the Standard Model

Despite the enormous success of the SM in precisely describing the fundamental particles of nature and their interactions, many questions remain unanswered. This section aims to provide a brief summary of the shortcomings of the SM and possible enhancements to it.

The limitations of the Standard Model can be broadly categorized into two types: fundamental and aesthetic shortcomings. The first category encompasses inherent limitations of the SM in explaining some physical observations. For instance, the SM fails to unite the concepts of QFT with general relativity [64] in order to incorporate gravitational interactions. Many attempts have been made to quantize gravitational interactions, but the resulting theories of quantum gravity are perturbatively non-renormalizable [65–68] and fail to make meaningful physical




Figure 1.5: ATLAS measurement of the production cross sections and decay rates of the Higgs boson at $\sqrt{s} = 13 \text{ TeV}$. The cross sections are measured relative to the leading gluon-gluon fusion production mode, whereas the branching ratios are measured relative to the decay into a ZZ^* pair. All measured quantities are normalized to the values predicted by the Standard Model. [63]

predictions. Compared to the weak force, gravity is many orders of magnitude weaker (see Table 1.2). As such, quantum effects of gravity are only expected to become relevant close to the Planck scale Λ_{P} [50], an energy regime that is and will remain far beyond the reach of existing and future accelerator technologies.

Many lines of evidence point to the existence of a form of matter for which the SM cannot provide an adequate explanation. Among other aspects, astrophysical observations of galaxy rotation curves [69], the cosmic microwave background [70] and gravitational lensing [71], indicate a significant excess of neutral, non-baryonic matter over visible matter. This form of matter is referred to as *dark matter* (DM) and has not yet been directly observed in experiments. The absence of a direct observation suggests that DM interacts very weakly (or not at all) with baryonic matter in the SM. Good candidate particles to explain DM are weakly interacting massive particles (WIMPs), for which the SM does not provide a suitable candidate. In the Lambda-CDM model [72–74] of cosmology and according to measurements from the Planck spacecraft [70], the universe consists of about 27% dark matter and only 5% baryonic matter. The remaining 68% is expected to be *dark energy*, a constant energy density of the vacuum that is uniformly distributed throughout the universe and is thought to explain its accelerated expansion [75, 76], but cannot be accounted for in the SM.

Shortly after the Big Bang, the universe is thought to have been in a state of equilibrium, with equal amounts of matter and antimatter. To explain the overwhelming excess of baryonic over non-baryonic matter observed today, a process called baryogenesis [77] is thought to have caused the imbalance between matter and antimatter in the early universe. According to the Sakharov conditions [78], baryogenesis implies interactions out of thermal equilibrium, baryon number violation, as well as C and CP violating processes. Although the Sakharov conditions can be satisfied by the SM, the possible CP violation in the electroweak quark and neutrino sector of the SM is not sufficient to explain the overwhelming excess of matter.

The problems described in the first category are fundamental in the sense that they inevitably require modifications to the SM in order to accurately describe all physical phenomena. Therefore, the SM is expected to be a low-energy approximation or *effective field theory* (EFT) of a more fundamental underlying theory. The second category of shortcomings in the SM encompasses mostly issues of aesthetic nature. These are not inherently problematic, but rather arise from a sense of theoretical intuition and experience.

Similar to the coupling constants, particle masses in the SM undergo mass renormalization. In the process, the Higgs boson mass $m_H = m_0 + \delta m_H$, receives massive radiative one-loop corrections of $\mathcal{O}(\delta m_H) = \Lambda_p$, which is about 16 orders of magnitude larger than the experimentally observed value [79]. The *bare* mass m_0 of the Higgs boson can be arbitrarily adjusted to exactly match the observed Higgs boson mass. However, such extreme *fine-tuning* over many orders of magnitude may seem unnatural and is known as the *hierarchy problem*.

The unification of the electromagnetic and weak interactions into a single electroweak force at high energies is an experimentally verified fact. A unification of all three fundamental forces described by the SM at even higher energies therefore seems like a natural extension. The SM is inherently unable to describe a unification of all three coupling constants, but many Grand Unified Theories (GUTs) predict such a unified description around the GUT scale $\Lambda_{\text{GUT}} = \mathcal{O}(10^{16} \text{ GeV})$. The simplest GUT model, referred to as the Georgi–Glashow model [80], combines the SM symmetry into a simple $SU(5)$ gauge group that spontaneously breaks into the $SU(3)_c \times SU(2)_L \times U(1)_Y$ subgroup of the SM at energies below Λ_{GUT} . While such theories are aesthetically appealing, a direct probe of a unification of the electroweak and strong force will remain experimentally out of reach, as the assumed GUT scale is only three orders of magnitude smaller than the Planck scale. In addition, many GUTs such as the Georgi–Glashow model are experimentally disfavoured because they predict the decay of protons, which has not yet been observed.

In the last decades, a large number of models have been proposed that extend the SM in one way or another and aim to solve the problems mentioned above. Supersymmetry (SUSY) [81] is one of the most popular and theoretically appealing beyond the Standard Model (BSM) extensions. At the very least, SUSY doubles the particle content by introducing fermionic and bosonic *superpartners*, which differ by half a unit of spin with respect to the associated SM particles and are known as *sparticles*. The introduction of sparticles allows for an exact cancellation of quadratic divergences in the loop corrections to the Higgs boson mass and

provides an elegant solution to the hierarchy problem. SUSY also permits additional sources of CP violation which may explain the imbalance between baryonic and non-baryonic matter and allows for an effective unification of the electroweak and strong force around the Planck scale. Gravity can be incorporated as a gauge theory of local supersymmetry, which is referred to as supergravity (SUGRA) [82] and provides a natural candidate for a potential spin-2 graviton. If SUSY is incorporated in nature, it must be spontaneously broken. Otherwise, sparticles are predicted to have identical masses with respect to their superpartners and would have already been observed at the LHC.

A self-consistent SUSY theory that is free of gauge anomalies, requires at least two Higgs doublets to explain the generation of masses for up- and down-type quarks. The SUSY theory with exactly two Higgs doublets and the minimum number of additional required particles is known as Minimal Supersymmetric Standard Model (MSSM) [83]. In the MSSM, an additional \mathbb{Z}_2 symmetry is imposed on the supersymmetric fields, such that R -parity

$$P_R = (-1)^{3(B-L)+2s} \quad (1.44)$$

is conserved. In this case, B is the baryon number, L is the lepton number and s is the spin of the particle. All SM particles have a positive integer parity $P_R = +1$, while their superpartners have $P_R = -1$. The conservation of R -parity implies the stability of the proton, which is necessary to explain the non-observation of the proton decay. Furthermore, the \mathbb{Z}_2 symmetry also implies the stability of the lightest supersymmetric particle (LSP), naturally providing a suitable WIMP candidate for DM [84]. In addition to the 26 free parameters of the SM, the MSSM adds over 100 additional unpredicted parameters, which renders phenomenological studies very complex. The MSSM is therefore typically constrained to satisfy well-motivated conditions such as first and second generation universality and the lack of flavour changing neutral currents (FCNC). This constrained model is known as pMSSM [85] and has a reduced phase space of 19 parameters.

1.7 Higgs Boson Pair Production

In the SM, the simultaneous production of two Higgs bosons is permitted via the trilinear Higgs boson and top-Yukawa coupling for which the main production modes at the LHC are shown in Figure 1.6. The first Feynman diagram is the analogue counterpart to gluon-gluon fusion in single Higgs boson production that is discussed in Section 1.5. In this case, however, the produced Higgs boson is off-shell and immediately decays into two real SM Higgs bosons. As the trilinear self-coupling is fundamentally related to the VEV in Equation (1.39), its measurement constitutes a direct probe of the shape of the Higgs potential and EWSB. The second most prominent production mode at the LHC is the next-to-leading order production via a direct coupling to the top quark, which permits experimental access to the top-Yukawa coupling of the Higgs boson. In the computation of the total Higgs boson pair production rate




Figure 1.6: Leading Feynman diagrams in a pp collider at $\sqrt{s} = 13$ TeV for non-resonant Higgs pair production as predicted by the Standard Model via (a) Higgs boson self-coupling or (b) the top-Yukawa coupling. The circles indicate the self-coupling strength κ_λ of the trilinear Higgs boson coupling and the top-Yukawa coupling κ_t .

predicted by the SM, both Feynman diagrams interfere destructively such that the resulting cross section of $\sigma^{\text{NNLO}}(pp \rightarrow HH) = 31.1_{-7.2}^{+2.1}$ fb for $\sqrt{s} = 13$ TeV [86–90] is extremely low. With respect to the total single Higgs production cross section listed in Table 1.3, Higgs boson pair production is suppressed by three orders of magnitude and is therefore an extremely rare process.

Many BSM models allow for an enhanced cross sections of non-resonant production by modifying the self-coupling strength κ_λ of the SM Higgs boson or the top-Yukawa coupling κ_t , effectively altering the interference behaviour between both production modes [91–93]. More relevant to this thesis, multiple SM extensions predict new heavy spin-0 or spin-2 resonances that can decay into two SM Higgs bosons and are expected to exhibit a resonant peak structure in the m_{HH} mass spectrum. Examples for BSM theories that predict such heavy resonances are the MSSM [83], composite Higgs models [94, 95] and twin Higgs models [96]. One more exotic model that, among other aspects, explains the weakness of the gravitational force with warped extra dimensions is referred to as Randall–Sundrum model [97, 98] and also predicts heavy candidate resonances in the form of spin-0 radions and spin-2 gravitons. Some models also introduce a second scalar particle S of differing mass, such that in addition to HH production processes of the form $pp \rightarrow X \rightarrow HH$, a decay into the scalar particle S and a SM Higgs boson $pp \rightarrow X \rightarrow SH$ is also possible, and in the following is referred to as SH production. The Feynman diagrams for both, resonant HH and SH production, are shown in Figure 1.7. Most of the aforementioned models that predict new heavy resonances can be classified into a set of more general BSM extensions, depending on the type of modifications that are imposed to the scalar sector of the SM. In Section 1.8, two of these Higgs sector extensions that are most relevant for this thesis are briefly discussed.

Finally, as the decay of the Higgs bosons in HH production is a statistically independent process, the decay rates can be trivially calculated from the single Higgs boson branching ratios and are given in Table 1.4. In the case of SH production, the branching ratios are naturally model-dependent.




Figure 1.7: Feynman diagrams of BSM resonant Higgs boson pair production. A new heavy resonant spin-0 or spin-2 scalar X decays into (a) two SM Higgs bosons H or (b) into a second new scalar resonance S and the SM Higgs boson.

$\mathcal{B}r$ [%]	$b\bar{b}$	$W W^*$	$\tau\tau$	$Z Z^*$	$\gamma\gamma$
$b\bar{b}$	34				
$W W^*$	25	4.6			
$\tau\tau$	7.3	2.7	0.39		
$Z Z^*$	3.1	1.1	0.33	0.07	
$\gamma\gamma$	0.2	0.1	0.02	0.01	< 0.001

Table 1.4: Predicted branching ratios of relevant Higgs boson pair decays resulting from the multiplicative combination of the single Higgs boson decay rates from Table 1.4.

1.8 Scalar Sector Extensions

In many BSM models, the scalar sector of the SM is extended to include additional degrees of freedom that address some of the shortcomings of the SM described in Section 1.6. The simplest possible extension, known as the RxSM model [99], adds a real scalar singlet to the SM Higgs potential and is the minimal model that can provide a suitable DM candidate. Its complex counterpart, the CxSM model [100], adds a complex singlet instead, and is the simplest extension that provides scenarios for the production of Higgs boson pairs with different masses. The RxSM and CxSM models are often used to define benchmark scenarios for Higgs-to-Higgs decay searches [101]. As briefly mentioned in Section 1.6, the MSSM extends the scalar sector of the SM with one real scalar doublet and is therefore an example of a Two-Higgs-doublet model (2HDM) [102]. The 2HDM also provides suitable DM candidates and is the simplest extension of the scalar sector that can provide new sources of \mathcal{CP} violation. In analogy to the CxSM model, the 2HDM can be extended by an additional complex phase, a model known as C2HDM [103], that permits an explicit \mathcal{CP} violation in the Higgs sector without the addition of new fermions. Some models also extend the scalar sector with more than one singlet or doublet. For instance, the N2HDM [104] extends the 2HDM by an additional real singlet which allows for the possibility of Higgs-to-Higgs decays with three different Higgs bosons and provides a benchmark for the Next-to-Minimal Supersymmetric Standard Model (NMSSM) [105],

which solves the μ -problem [106] of the MSSM. Finally, the Two-Real-Scalar-Singlet Model (TRSM) [107] enhances the scalar sector by two real singlets, leading to predicted experimental signatures of three- and four-Higgs final states. A summary of these extensions and their predicted particle content in the Higgs sector is given in Table 1.5.

		RxSM	CxSM	2HDM	C2HDM	N2HDM	TRSM
Higgs Sector Extension	Real Singlet	✓				✓	✓ (2)
	Real Doublet			✓		✓	
	Complex Singlet		✓				
	Complex Doublet				✓		
Particle Content	H^0						
	\mathcal{CP} -even	2	3	2	-	3	3
	\mathcal{CP} -odd	-	-	1	-	1	-
	\mathcal{CP} -mixed	-	-	-	3	-	-
	H^\pm	-	-	1	1	1	-
Ref.		[99]	[100]	[102]	[103]	[104]	[107]

Table 1.5: Summary of the most relevant BSM models that extend the scalar sector of the SM by a singlet or doublet. The values indicate the number of particles predicted by each extension in the Higgs sector.

All models described result in heavier neutral or charged copies of the SM Higgs boson with \mathcal{CP} properties depending on the type of scalar extension. The two most interesting models for the search for resonant pair production in the $b\bar{b}V V^*$ decay channel that is presented in Part III of this thesis, are the 2HDM and the TRSM. The 2HDM model is mainly motivated by the MSSM and predicts a \mathcal{CP} -even copy of the SM Higgs boson, which acts as a possible candidate for resonant HH production. The TRSM predicts two \mathcal{CP} -even copies of the SM Higgs boson and as such, additionally allows for resonant SH production. Resonant SH production is also allowed in other models such as the CxSM and the N2HDM, but the TRSM is particularly appealing due to the fact that it predicts a significant enhancement in the $SH \rightarrow b\bar{b}V V^*$ branching ratio relative to other decay channels [107]. The most relevant features of the 2HDM and TRSM are briefly discussed in Section 1.8.1 and Section 1.8.2, respectively.

1.8.1 Two-Higgs-Doublet Model

In analogy to the SU(2) isospin doublet introduced in the BEH mechanism in Equation (1.29), the 2HDM model constructs the scalar sector with two real scalar doublets

$$\Phi_a = \begin{pmatrix} \phi_a^+ \\ \phi_a^0 \end{pmatrix} \quad a \in [1, 2], \quad (1.45)$$

where the $SU(3)_c \times SU(2)_L \times U(1)_Y$ structure of the SM is preserved and both doublets take the hypercharge $Y = +1$. In the most general form, the vacuum structure of 2HDM models is very complex and the scalar potential contains a total of 14 free parameters. The large parameter space makes any phenomenological studies impractical, such that several simplifying assumptions are typically made. CP conservation in the Higgs sector is usually imposed and allows a clear distinction between scalar and pseudoscalar fields. Moreover, if one assumes that CP is not spontaneously broken, and that discrete symmetries eliminate all odd quartic terms in either of the doublets, the most general form of the potential [102] reads

$$\begin{aligned}
 V = & m_{11}^2 \Phi_1^\dagger \Phi_1 + m_{22}^2 \Phi_2^\dagger \Phi_2 - m_{12}^2 \left(\Phi_1^\dagger \Phi_2 + \Phi_2^\dagger \Phi_1 \right) + \frac{\lambda_1}{2} \left(\Phi_1^\dagger \Phi_1 \right)^2 + \frac{\lambda_2}{2} \left(\Phi_2^\dagger \Phi_2 \right)^2 \\
 & + \lambda_3 \Phi_1^\dagger \Phi_1 \Phi_2^\dagger \Phi_2 + \lambda_4 \Phi_1^\dagger \Phi_2 \Phi_2^\dagger \Phi_1 + \frac{\lambda_5}{2} \left[\left(\Phi_1^\dagger \Phi_2 \right)^2 + \left(\Phi_2^\dagger \Phi_1 \right)^2 \right],
 \end{aligned} \tag{1.46}$$

where λ_i and m_{ij} are the reduced set of parameters. After EWSB, and assuming CP conservation in both VEVs of the scalar doublets, the doublets acquire two of the infinite degenerate ground states

$$\langle \Phi_a \rangle_0 = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ v_a \end{pmatrix}, \tag{1.47}$$

where v_1 and v_2 are the two vacuum expectation values of the doublets.

The expansion of the fields around their minima leads to eight fields that can be expressed as

$$\Phi_a = \begin{pmatrix} \phi_a^+ \\ (v_a + \rho_a + i\eta_a) / \sqrt{2} \end{pmatrix}, \tag{1.48}$$

from which three correspond to massless Goldstone bosons which generate the masses of the W^\pm and Z bosons through the same mechanism as in the SM. The other five fields correspond to two charged H^\pm , a light neutral CP -even scalar H and a heavy copy X , as well as a neutral CP -odd pseudoscalar A , as already indicated in Table 1.5. The light CP -even scalar H is assumed to be an SM-like Higgs boson and its heavy counterpart X gives rise to resonant Higgs boson pair production. The fields corresponding to the physical light and heavy CP -even scalar particles arise through mixtures of the two ρ_i fields as

$$H = \rho_1 \sin \alpha - \rho_2 \cos \alpha \quad \text{and} \quad X = -\rho_1 \cos \alpha - \rho_2 \sin \alpha, \tag{1.49}$$

where α is a mixing angle that is dependent on the couplings in the potential. The CP -odd pseudoscalar A and the charged Higgs bosons H^\pm can be expressed as linear combinations of ϕ_a^\pm and η_i according to

$$H^\pm = \frac{1}{v} [v_2 \phi_1^\pm - v_1 \phi_2^\pm] \quad \text{and} \quad A = \frac{1}{v} [v_2 \eta_1 - v_1 \eta_2], \tag{1.50}$$

where $v^2 = v_1^2 + v_2^2$ is the quadratic sum of the VEVs of the two isospin doublets. Together with the mixing parameter α , the rotation angle that diagonalizes the mass-squared matrices of the charged and pseudoscalar Higgs bosons $\beta = \tan^{-1} v_2/v_1$ is the most important parameter of the 2HDM and fully determines the couplings of the Higgs fields with the fermions and vector bosons.

Without further constraints, the general potential from Equation (1.46) allows for tree-level flavour-changing neutral currents (FCNCs). FCNCs are predicted by the SM beyond leading order, but are heavily suppressed by the GIM [108] mechanism. Searches for FCNCs have set strict experimental limits, such that their existence at tree-level is experimentally disfavoured [109]. It is possible to avoid FCNCs in the 2HDM model altogether by imposing additional constraints on the coupling of fermions and quarks to the Higgs doublets. More specifically, to ensure the absence of FCNCs, all fermions with identical quantum numbers must couple to the same Higgs doublet, which can be enforced by imposing additional discrete \mathbb{Z}_2 symmetries. In this case, a total of four possible models remain that are summarised in Table 1.6 and are classified according to the type of fermion couplings they exhibit. By

Type	u_R	d_R	ℓ_R	Imposed Discrete \mathbb{Z}_2 Symmetry
Type I	Φ_2	Φ_2	Φ_2	$\mathbb{Z}_2 : \Phi_1 \rightarrow -\Phi_1$
Type II	Φ_2	Φ_1	Φ_1	$\mathbb{Z}_2 : \Phi_1 \rightarrow -\Phi_1, d_R \rightarrow -d_R$
Lepton-specific	Φ_2	Φ_2	Φ_1	$\mathbb{Z}_2 : \Phi_1 \rightarrow -\Phi_1, \ell_R \rightarrow -\ell_R$
Flipped	Φ_2	Φ_1	Φ_2	$\mathbb{Z}_2 : \Phi_1 \rightarrow -\Phi_1, d_R \rightarrow -d_R, \ell_R \rightarrow -\ell_R$

Table 1.6: Summary of the different types of 2HDMs with discrete \mathbb{Z}_2 symmetries that ensure natural flavour conservation, and the resulting couplings of up-type quarks u , down-type quarks d and leptons ℓ to the two Higgs doublets. The couplings are taken from Ref. [102].

convention, up-type quarks exclusively couple to the second Higgs doublet Φ_2 in all cases. For Type I HDMs, down-type quarks also couple to the second doublet, whereas for type II HDMs, down-type quarks couple to the first doublet. The other two models are of lepton-specific type, where all quarks couple to the second doublet, while leptons couple to Φ_1 , and the flipped type, in which only the down-type quarks couple to the first doublet and all other particles to Φ_2 .

The Yukawa interactions of the 2HDM can be expressed [110] in terms of mass eigenstates of the Higgs bosons as

$$\begin{aligned}
 \mathcal{L}_{\text{Yukawa}}^{2\text{HDM}} = & - \sum_{f=u,d,\ell} \frac{m_f}{v} \left(\xi_H^f \bar{f} f H + \xi_X^f \bar{f} f X - i \xi_A^f \bar{f} \gamma_5 f A \right) \\
 & - \left\{ \frac{\sqrt{2} V_{ud}}{v} \bar{u} \left(m_u \xi_A^u P_L + m_d \xi_A^d P_R \right) d H^+ + \frac{\sqrt{2} m_\ell \xi_A^\ell}{v} \bar{\nu}_L \ell_R H^+ + \text{h.c.} \right\}, \quad (1.51)
 \end{aligned}$$

where the sum runs over all leptons as well as up- and down-type quarks, and ξ_H , ξ_X and ξ_A are the couplings of the fermions to the respective neutral Higgs bosons. As mentioned before, the couplings are solely dependent on the two angles α and β , and are given in Table 1.7 for the various types of natural flavour conserving 2HDMs. All 2HDM models absent of FCNCs

Fermion	Boson	Type I	Type II	Lepton-specific	Flipped
u	H	$\cos \alpha / \sin \beta$	$\cos \alpha / \sin \beta$	$\cos \alpha / \sin \beta$	$\cos \alpha / \sin \beta$
	X	$\sin \alpha / \sin \beta$	$\sin \alpha / \sin \beta$	$\sin \alpha / \sin \beta$	$\sin \alpha / \sin \beta$
	A	$\cot \beta$	$\cot \beta$	$\cot \beta$	$\cot \beta$
d	H	$\cos \alpha / \sin \beta$	$-\sin \alpha / \cos \beta$	$\cos \alpha / \sin \beta$	$-\sin \alpha / \cos \beta$
	X	$\sin \alpha / \sin \beta$	$\cos \alpha / \cos \beta$	$\sin \alpha / \sin \beta$	$\cos \alpha / \cos \beta$
	A	$-\cot \beta$	$\tan \beta$	$-\cot \beta$	$\tan \beta$
ℓ	H	$\cos \alpha / \sin \beta$	$-\sin \alpha / \cos \beta$	$-\sin \alpha / \cos \beta$	$\cos \alpha / \sin \beta$
	X	$\sin \alpha / \sin \beta$	$\cos \alpha / \cos \beta$	$\cos \alpha / \cos \beta$	$\sin \alpha / \sin \beta$
	A	$-\cot \beta$	$\tan \beta$	$\tan \beta$	$-\cot \beta$

Table 1.7: Yukawa couplings ξ for up-type quarks u , down-type quarks d , and charged leptons ℓ , with the neutral Higgs bosons H , X and A in the four different models. The couplings are taken from Ref. [102].

forbid couplings between the pseudoscalar A and vector bosons. The couplings between the $C\mathcal{P}$ -even scalar particles are identical in all models and are suppressed by $\sin(\beta - \alpha)$ for the light Higgs boson H and by $\cos(\beta - \alpha)$ for its heavy counterpart X , with respect to the SM coupling to either W^\pm or Z boson.

1.8.2 Two-Real-Scalar-Singlet Model

In general, the SM Higgs potential $V_{\text{SM}}(\Phi)$ can be extended with additional N scalar fields, such that the resulting BSM potential is given by

$$V(\phi_i, \Phi) = V_{\text{singlets}}(\phi_i, \Phi) + V_{\text{SM}}(\Phi), \quad (1.52)$$

where the most general form of the singlets potential V_{singlets} [107] is

$$\begin{aligned} V_{\text{singlets}}(\phi_i, \Phi) = & a_i \phi_i + m_{ij} \phi_i \phi_j + T_{ijk} \phi_i \phi_j \phi_k + \lambda_{ijkl} \phi_i \phi_j \phi_k \phi_l \\ & + T_{iHH} \phi_i (\Phi^\dagger \Phi) + \lambda_{ijHH} \phi_i \phi_j (\Phi^\dagger \Phi), \end{aligned} \quad (1.53)$$

with real coefficient tensors. In the following, the scalar fields are assumed to be real, but the potential can be expanded into the same form for complex singlets.

In the Two-Real-Scalar-Singlet (TRSM) model [107], two such singlets are added ($N = 2$) that are referred to as X and S . In order to reduce the parameter space, two additional discrete \mathbb{Z}_2 symmetries

$$\begin{aligned}\mathbb{Z}_2^S &: S \rightarrow -S, X \rightarrow X, \text{SM} \rightarrow \text{SM} \\ \mathbb{Z}_2^X &: X \rightarrow -X, S \rightarrow S, \text{SM} \rightarrow \text{SM}\end{aligned}\quad (1.54)$$

are typically imposed. The most general renormalizable potential [107] that is invariant under $\mathbb{Z}_2^S \otimes \mathbb{Z}_2^X$ transformations is given by

$$\begin{aligned}V = V_{\text{SM}}(\Phi) &+ \mu_S^2 S^2 + \lambda_S S^4 + \mu_X^2 X^2 + \lambda_X X^4 \\ &+ \lambda_{\Phi S} \Phi^\dagger \Phi S^2 + \lambda_{\Phi X} \Phi^\dagger \Phi X^2 + \lambda_{SX} S^2 X^2,\end{aligned}\quad (1.55)$$

where, in addition to the λ and μ^2 parameters of the SM Higgs potential, the seven coefficients fully determine the shape of the scalar TRSM potential.

After EWSB, the fields expanded around the minimum can be written as

$$\Phi = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ \phi_H + v \end{pmatrix}, \quad S = \frac{\phi_S + v_S}{\sqrt{2}}, \quad X = \frac{\phi_X + v_X}{\sqrt{2}}\quad (1.56)$$

in unitary gauge, where v corresponds to the SM VEV of the Higgs potential and v_S and v_X are the vacuum expectation values of the scalar singlets S and X , respectively. If one of the VEVs vanishes, the associated field does not mix with the other two fields and consequently does not acquire any of the SM couplings. It is further stabilized by its \mathbb{Z}_2 symmetry, making it a viable DM candidate. More interesting for this thesis is the case where $v_S, v_X \neq 0$, such that the \mathbb{Z}_2 symmetries are spontaneously broken, and the three fields ϕ_H, ϕ_S and ϕ_X mix to form three observable mass eigenstates h_i that are related as

$$\begin{pmatrix} h_1 \\ h_2 \\ h_3 \end{pmatrix} = R \begin{pmatrix} \phi_H \\ \phi_S \\ \phi_X \end{pmatrix},\quad (1.57)$$

with the mixing matrix R . By definition, the Higgs bosons are ordered according to their masses $m_1 \leq m_2 \leq m_3$, where the lightest Higgs boson h_1 can be associated to the SM Higgs boson, and the heavier twins, h_2 and h_3 , give rise to the possibility of resonant SH production. The mixing matrix can be parametrized according to

$$R = \begin{pmatrix} c_1 c_2 & -s_1 c_2 & -s_2 \\ s_1 c_3 - c_1 s_2 s_3 & c_1 c_3 + s_1 s_2 s_3 & -c_2 s_3 \\ c_1 s_2 c_3 + s_1 s_3 & c_1 s_3 - s_1 s_2 c_3 & c_2 c_3 \end{pmatrix},\quad (1.58)$$

where the coefficients s_i and c_i are only dependent on the mixing angles of the scalar fields and

can be expressed as

$$\begin{aligned} s_1 &= \sin \theta_{HS}, & s_2 &= \sin \theta_{HX}, & s_3 &= \sin \theta_{SX} \\ c_1 &= \cos \theta_{HS}, & c_2 &= \cos \theta_{HX}, & c_3 &= \cos \theta_{SX} \end{aligned} \quad (1.59)$$

As pure gauge singlets, the singlets introduced in the potential from Equation (1.55) do not induce any further gauge interactions. Furthermore, they cannot interact directly with SM fermions, as any attempt to add coupling terms would either spoil the local gauge invariance or lead to a non-renormalizable theory. However, the observable mass eigenstates h_i inherit some of the Yukawa couplings of the Higgs doublet through the mixing. Moreover, the mixing gives rise to trilinear couplings and self-interactions between the three fields of the form λ_{abc} , λ_{abb} and λ_{aaa} with coupling strengths

$$\lambda_{abc} = \sum_i \frac{R_{ai}R_{bi}R_{ci}}{v_i} \sum_i m_i^2, \quad (1.60)$$

that are only dependent on the Higgs boson masses, the mixing matrix and the VEVs of the fields.

The ATLAS Experiment at the LHC

THE ATLAS [111] experiment at the European Organization for Nuclear Research (CERN) is one of the four main detectors at the Large Hadron Collider (LHC) [112]. High-density beams of protons or heavy ions are produced at unprecedented energies by a series of accelerators and brought to collision at four interaction points along the beam pipe. Specially designed detectors, including the ATLAS experiment, are placed around each interaction point to precisely measure the properties of the resulting collision products.

This chapter gives a brief description of the accelerator chain and the experimental setup of the ATLAS experiment at the LHC.

2.1 The Large Hadron Collider

The LHC is the largest and highest-energy particle accelerator ever built by humankind. Located beneath the French-Swiss border near Geneva at an average depth of 100 m, it is a circular hadron collider with a circumference of 26.7 km. Bunches of protons or heavy ions are injected into two adjacent parallel beam pipes that travel in opposite directions and are brought to collision at the interaction points of the ATLAS [111], CMS [113], ALICE [114] and LHCb [115] experiments. The ATLAS and CMS detectors are multipurpose experiments designed to precisely measure a wide range of physics processes for precision measurements as well as for searches for beyond the Standard Model particles. ALICE is a detector dedicated to heavy-ion physics, aiming to study the physics of strongly interacting matter at extreme energy densities, a state known as quark-gluon plasma. Finally, LHCb specializes in b -physics and is primarily designed to investigate $C\mathcal{P}$ violation in the b -sector with the aim of understanding the matter-antimatter asymmetry of the universe.

For proton-proton collisions, the LHC is designed to produce collisions with a centre-of-mass energy of up to $\sqrt{s} = 14$ TeV. To achieve these unprecedented energies, protons are not injected directly into the LHC ring, but are accelerated successively through various pre-accelerators. First, negatively charged hydrogen ions are injected into the linear accelerator 4 (LINAC4) and accelerated to 160 MeV. The ions are then stripped of their two electrons and the resulting protons are subsequently injected into the Proton Synchrotron Booster (PSB), where they reach energies of up to 2 GeV. The protons are then further accelerated to 26 GeV in the Proton Synchrotron (PS) and to 450 GeV in the Super Proton Synchrotron (SPS), before being injected into the LHC, where protons can be accelerated to the TeV energy regime. A sketch of the CERN accelerator complex as of 2022 is shown in Figure 2.1.

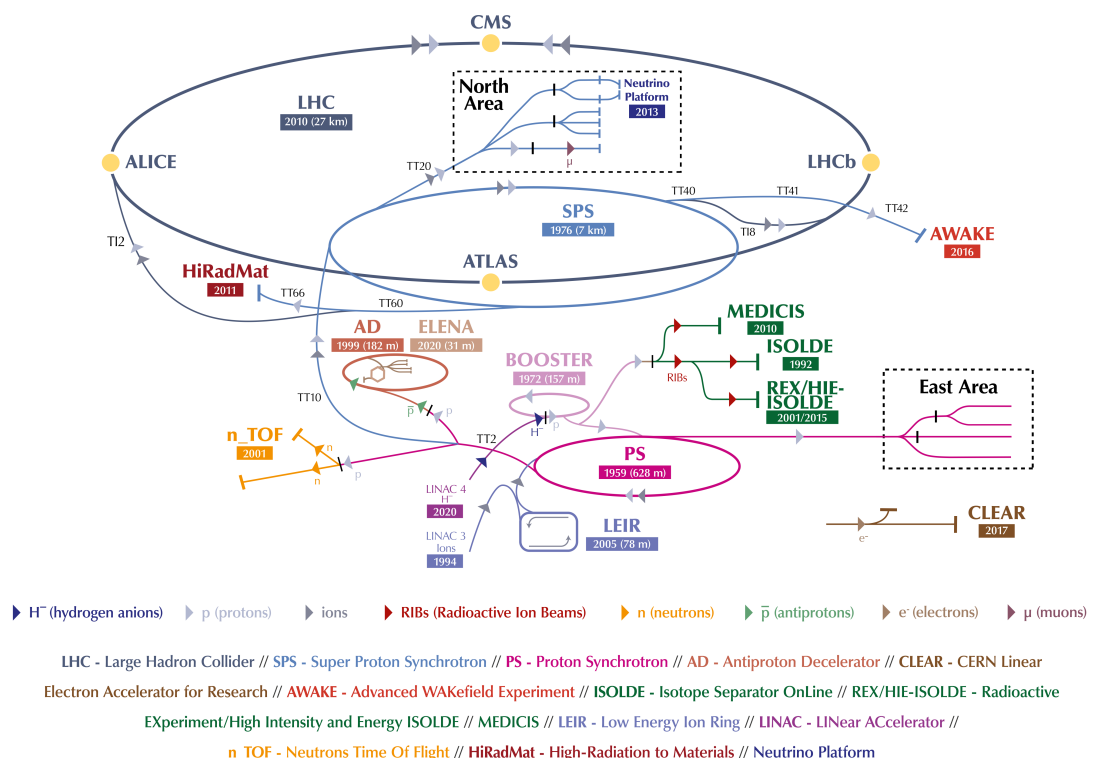


Figure 2.1: Sketch of the CERN accelerator complex as of 2022. The light (dark) grey arrows indicate the typical flight direction of protons (heavy ions) through the complex. The years indicate the start dates of the experiments and accelerators shown. [116]

Protons are injected into the LHC storage ring as separated high-density particle bunches. The LHC is designed to hold up to 2835 bunches per beam, where each bunch carries approximately 10^{11} particles with a bunch separation of 25 ns. The spacing is necessary to accommodate for the dead time of the detector systems and allows proton-proton collisions at a rate of up to 40 MHz. The bunches are held in their circular trajectory by 1232 superconducting Nb-Ti dipole

magnets, each 15 metres long, that produce large magnetic fields of up to 8.3 T. The magnets are operated at temperatures below 2 K and cooled with superfluid helium. Proton beams, which have a natural tendency to disperse due their positive charge, are focused by quadrupole magnets that are symmetrically arranged around the beam pipe. The dipole magnets are also equipped with higher-order magnets to correct for magnetic field imperfections and to enhance the stability of the beam.

Apart from the centre-of-mass energy of an accelerator, the frequency of proton-proton collisions is a key benchmark value to assess the performance of an accelerator and is directly related to the instantaneous luminosity

$$\mathcal{L} = f \frac{N_1 N_2 n_b}{4\pi\sigma_x\sigma_y}, \quad (2.1)$$

where $N_{1,2}$ are the number of particles in the two adjacent beams, n_b is the total number of bunches in the storage ring, f is the frequency at which bunches are brought to collision, and $\sigma_{x,y}$ are the beam dimensions in the horizontal and vertical direction. The luminosity is proportional to the interaction rate

$$\frac{dN}{dt} = \mathcal{L} \cdot \sigma, \quad (2.2)$$

where σ is the cross section of an arbitrary process. The integrated luminosity over time is typically used as a measure to characterize the amount of data collected by an experiment over a given period of time.

The LHC is designed to provide a peak instantaneous luminosity of up to $\mathcal{L} = 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. In Run 1 of the LHC, the machine was operated at $\sim 75\%$ of the design luminosity and delivered a total integrated luminosity of 5.46 fb^{-1} at $\sqrt{s} = 7 \text{ TeV}$ and 22.8 fb^{-1} at $\sqrt{s} = 8 \text{ TeV}$. The machine subsequently underwent a period of maintenance and upgrades known as long shutdown 1 (LS1) and restarted its operation in 2015 at a centre-of-mass energy of $\sqrt{s} = 13 \text{ TeV}$. During Run 2, the machine reached its design luminosity and delivered a total integrated luminosity of 156 fb^{-1} . The second run concluded in late 2018, marking the start of long shutdown 2 (LS2), which was used to perform critical updates to the detectors. The shutdown ended with the start of Run 3 in 2022, during which collision energies of $\sqrt{s} = 13.6 \text{ TeV}$ were achieved for the first time and peak luminosities of up to $2 \cdot 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$ are expected. In 2026, the machine will enter a third three-year shutdown period and will be upgraded as part of the High-Luminosity Large Hadron Collider (HL-LHC) [117] project. At its expected start in early 2029, the HL-LHC is expected to provide a 5- to 7.5-fold increase in peak luminosity compared to the design value. The HL-LHC is expected to deliver up to $3 - 4 \text{ ab}^{-1}$ in integrated luminosity after about 12 years of operation. A timeline of the operational history of the LHC from the start of Run 1 to the projected operation of the HL-LHC until 2040 is shown in Figure 2.2.

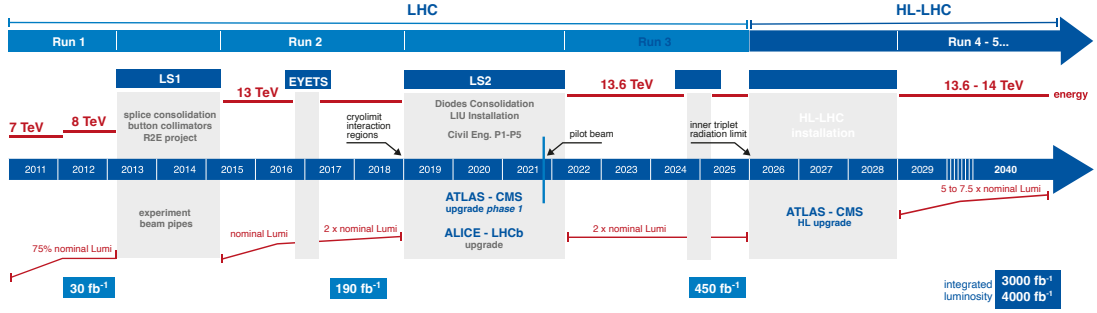


Figure 2.2: Operational history of the LHC from the start of Run 1 to the beginning of Run 3 in 2022 and projected timeline of the LHC and the high luminosity upgrade HL-LHC until 2040. The shown luminosity numbers refer to the total cumulated integrated luminosities over time. [118]

2.2 Proton-Proton Collisions

The ATLAS experiment is specialized in measuring the products of high-energetic proton-proton (pp) collisions. In comparison to lepton colliders, pp colliders produce very rich final states with a multitude of particles resulting from various underlying processes. Figure 2.3 shows an illustration of a typical proton-proton collision with a top quark pair in the final state.

Protons are composite particles made up of three valence quarks uud . In high-energy proton-proton collisions, interactions occur between the constituent partons. This is because, in the relativistic limit, the valence quarks are asymptotically free and behave essentially as free particles. The total proton momentum is distributed among its constituents, but the momentum of individual quarks is unknown and cannot be calculated analytically. However, experimentally determined Parton Density Functions (PDFs) $f_i(x, Q^2)$ can be used to encode the probability to find a parton of species i with a proton momentum fraction x at a fixed momentum transfer Q^2 at which the proton is probed. The cross section of an arbitrary process $pp \rightarrow X$ can then be factorized [120] into a cross section at parton level $\hat{\sigma}_{q_1 q_2 \rightarrow X}$ and PDFs f_{q_i} that account for the inner structure of the proton according to

$$\sigma_{pp \rightarrow X} = \int dx_1 dx_2 \sum_{q_1, q_2} f_{q_1}(x_1, \mu_F^2) f_{q_2}(x_2, \mu_F^2) \hat{\sigma}_{q_1 q_2 \rightarrow X}(\hat{s}, \mu_F^2), \quad (2.3)$$

where the sum runs over all possible parton combinations in the proton. The factorization scale μ_F^2 is the momentum scale at which perturbative calculations of the partonic cross sections become impossible and are taken into account by the density functions. The PDFs are typically determined using deep inelastic scattering data from $e^\pm p$, pp or $p\bar{p}$ collisions at a fixed energy scale. The PDFs at arbitrary Q^2 can then be determined with the DGLAP evolution equations [121–123].

The hard scattering process of the proton constituents may produce short-lived resonances,




Figure 2.3: Illustration of a typical $pp \rightarrow t\bar{t}$ event. The partons of two incoming protons scatter to form a few outgoing particles that produce parton showers. The partons within the shower hadronize to form colourless states, which eventually decay into stable hadrons. ISR (FSR) refers to initial (final) state radiation. [119]

such as electroweak gauge bosons or top quarks, as shown in Figure 2.3. Quarks and gluons created in the hard scatter process immediately produce parton showers that form colourless hadrons, a process known as *hadronization*. The process is non-perturbative and described by phenomenological models such as the Lund string model [124]. Some of the produced hadrons will be unstable and decay further into secondary hadrons. The collections of stable nearby hadrons resulting from these hadronization processes form cone-like structures known as *jets*.

In addition to the hard-scatter interaction, other processes play an important role in pp collisions. For example, the partonic scattering initiators can produce Initial State Radiation (ISR) in the form of gluons, which will start additional hadronization processes and reduce the hard-scatter momentum transfer. Similarly, Final-State Radiation (FSR) can be produced from the hard scattering itself or from any subsequent resonance decay. Scattering processes can also occur from partons of the incident beams, which are known as Multiple Parton Interactions

(MPIs). In addition to the hard scattering, multiple softer hadron-hadron collisions can occur within a very short time. In experiments, these collisions are recorded in conjunction with the hard scattering and are known as *pileup*, which includes additional pp collisions from adjacent bunch crossings. Proton remnants from hard collisions and the combination of all detector activity which does not stem from the hard scattering is known as the underlying event (UE).

2.3 The ATLAS Experiment

The ATLAS experiment is a multipurpose detector at the LHC that is 44 m long, 25 m in diameter and weighs about 7000 tonnes. It consists of a series of subdetectors that form concentric layers around the interaction point. With additional endcap layers, the detector covers almost the entire 4π solid angle, allowing a precise measurement of most physics processes. Figure 2.4 shows a schematic illustration of the ATLAS detector and its various subsystems, including new systems installed during LS2. The main components of the detector




Figure 2.4: Schematic illustration of the ATLAS detector and its various subsystems, including new systems installed during LS2. [125]

are the Inner Detector (ID), the calorimeter system and the Muon Spectrometer (MS), and are explained in more detail in the following.

2.3.1 Coordinate System

To describe the properties of particles in the detector, the ATLAS experiment uses a right-handed cylindrical reference frame. The z -axis is aligned with the beam line and the $x - y$ plane defines the plane transverse to beam. The x -axis points towards the centre of the LHC ring, while the y -axis points towards the surface of the earth. The angle $\phi \in [-\pi, \pi]$ is defined as the azimuth around the beam line. Instead of using the polar angle $\theta \in [0, \pi]$ along the longitudinal direction, the pseudorapidity

$$\eta = -\ln \left[\tan \left(\frac{\theta}{2} \right) \right] \quad (2.4)$$

is typically used. In the relativistic limit, pseudorapidity corresponds approximately to the rapidity

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right). \quad (2.5)$$

A sketch of the coordinate system used in ATLAS is shown in Figure 2.5.




Figure 2.5: Sketch of the coordinate system employed in ATLAS. Particles travel in opposite directions through the beam pipe (illustrated as grey tube) and collide at the interaction point (IP). Adapted from Ref. [126].

Differences in rapidity are invariant under Lorentz boosts along the longitudinal axis. Similarly, in the relativistic limit, differences in pseudorapidity are approximately invariant under Lorentz boosts in the direction of the beam line.

The angular separation between two objects in the detector is typically described with

$$\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}, \quad (2.6)$$

which, in the relativistic limit, is invariant under Lorentz boosts along the z -direction.

2.3.2 Inner Detector

The Inner Detector (ID) [127, 128] is the innermost part of the ATLAS experiment and is dedicated to measure and reconstruct the tracks of charged particles. It is surrounded by a solenoid magnet which generates a highly homogeneous magnetic field of 2 T parallel to the beam line. The magnetic field causes the charged particles to bend, which allows to measure their direction, momentum, and charge from the observed curvature of their trajectory. The ID has a diameter of 2.1 m and a length of 6.2 m and consists of three main parts: the Pixel Detector, the Semiconductor Tracker (SCT), and the Transition Radiation Tracker (TRT). A sketch of the Inner Detector in the central region of the detector is shown in Figure 2.6. The overall resolution of the ID degrades as the transverse momentum of the incoming particle increases, due to the smaller track curvatures, and is $\sigma_{p_T}/p_T = 0.05\% \cdot p_T \oplus 1\%$.




Figure 2.6: Sketch of the Inner Detector in the barrel region. The values indicate the radial distances of the layers to the interaction point. [129]

Pixel Detector

The innermost subsystem of the ID is the Pixel detector, which consists of four finely segmented concentric layers of two-dimensional arrays of silicon cells in the central barrel region and three endcap layers in the form of disks. It starts at a radial distance of 3.3 cm from the beam axis and extends to a longitudinal length of 6.2 m and a radius of 1.2 m. The Pixel detector consists of about 92 million pixels with a pixel size of $50\ \mu\text{m} \times 400\ \mu\text{m}$ for the three outermost layers. The innermost insertable *B*-layer (IBL) [130] in the barrel region has smaller pixels of size $50\ \mu\text{m} \times 250\ \mu\text{m}$ and was inserted during LS1 to provide a significantly enhanced spatial resolution of tracks.

As a charged particle travels through the sensitive material, it ionizes the atoms in the material, creating electron-hole pairs. A voltage bias is applied across the silicon sensors that creates an electric field which causes both charges to separate and drift towards the electrodes. The electric current generated by the drifting charges is detected and amplified by the readout electronics creating a *hit* in the respective pixel. Series of hits in the layers are then used to reconstruct particle track in the detector. A more detailed description of track reconstruction in ATLAS is provided in Section 6.1.

Semiconductor Tracker

The measurement principle of the Semiconductor Tracker (SCT) is essentially identical to that of the Pixel detector. However, instead of pixel arrays, it uses narrow silicon strips to detect charged particles. It consists of four barrel layers with pairs of strip sensors that are rotated by 40 mrad with respect to each other. Since the strip sensors can only perform measurements in one direction, the information from both strip module pairs is typically combined. A total of nine endcap discs are placed in the longitudinal direction. In total, the SCT encompasses over 4000 modules with approximately 6 million readout chips. In the barrel region, the strips are placed every $80\ \mu\text{m}$ and have a length of about 120 mm. The intrinsic spatial resolution in the $r - \phi$ direction is approximately $17\ \mu\text{m}$, while the resolution in the $r - z$ direction is about $580\ \mu\text{m}$ in the barrel (end-caps).

Transition Radiation Tracker

The Transition Radiation Tracker (TRT) is the outermost part of the ID. It consists of about 3×10^5 gas-filled straw tubes of 4 mm in diameter which each have a $30\ \mu\text{m}$ gold-plated tungsten wire in their centre. A total of 50 000 tubes, each 144 cm long, are arranged parallel to the beam line in the barrel, and 250 000 tubes, each 39 cm long, are arranged perpendicular to the beam line in the endcap region.

The TRT essentially acts as a drift chamber. Particles passing through the detector ionize the Xe/Ar gas mixture in the drift tubes, producing freely charged particles. A voltage is applied

between the tubes and the tungsten wire, which act as the electrodes of the TRT. The charged particles then drift towards the electrodes, where the signal is collected and amplified.

The region between the tubes is filled with polymer fibers and foils that increase the transition radiation (TR) of particles traversing the straw tubes. As the production of TR is sensitive to the Lorentz factor γ , its measurement allows to infer information on the particles mass and provides important aid in the identification of charged particles. As a result, the TRT is especially useful for discriminating between electrons and charged hadrons. The timing information of the drift tube readout allows to measure the track position in the $r - \phi$ direction with a resolution of $130 \mu\text{m}$.

During LS3, the ID is expected to be replaced with an all-silicon Inner Tracker (ITk) [131, 132]. The new tracker will significantly increase the size of the current pixel system and feature a much more segmented strip detector. The upgrade is essential to cope with the much harsher pileup conditions at the HL-LHC, while maintaining the performance of the ID from previous runs.

2.3.3 Calorimeter System

Calorimeters are destructive devices used to provide energy measurements and directional information for incident particles, whose measurement principle is based on the complete stopping of an incident particle within a block of matter. An incident electromagnetically or hadronically interacting particle will initiate a *shower cascade* that ideally deposits all its energy within the measuring device. The deposited energy induces a signal in the detector which is proportional to the energy of the incident particle. As the physics of shower cascades is of particular relevance to this thesis, a detailed description is provided in Chapter 3.

The ATLAS experiment uses *sampling calorimeters* with alternating layers of active and absorbing material. The absorber consists of a high-density material, chosen to maximize the interaction probability of electromagnetic or hadronic interacting particles, to ensure that a shower cascade is rapidly initiated, and all shower remnants are stopped within the calorimeter system. The non-absorbing layers, on the other hand, consist of active material in which detectable electronic signals are created that are used to measure the energy of the incident particle. Compared to *homogeneous calorimeters*, sampling calorimeters have the advantage that they can be easily segmented laterally and longitudinally, providing improved position measurement and particle identification capabilities. In addition, the alternating structure allows for significantly more compact devices. As hadronic showers penetrate deep into the material, sampling calorimeters are typically the only option for hadronic calorimeters.

The calorimeter system in ATLAS uses electromagnetic and hadronic calorimeters that surround the Inner Detector and cover at total pseudorapidity range $|\eta| < 4.9$. The innermost part consists of the electromagnetic calorimeter (ECAL) and is followed by the hadronic calorimeter (HCAL), which both cover the region $|\eta| < 3.2$. The acceptance is extended to

$|\eta| < 4.9$ with a dedicated forward calorimeter that provides electromagnetic and hadronic calorimetry. A cut-way sketch of the ATLAS calorimeter system is provided in Figure 2.7.




Figure 2.7: Cut-away sketch of the ATLAS calorimeter system. [111]

The ECAL is divided into barrel (EMB) and endcap (EMEC) modules which consist of high-granularity lead/liquid-argon (LAr) calorimeters. The lead layers in the calorimeter vary in thickness from 1.1 to 2.2 mm, depending on the specific detector region, and are separated by 4 mm active liquid argon-filled gaps. The thickness of the EM calorimeter is typically measured in terms of radiation lengths X_0 , and is $\geq 22X_0$ for all regions in the detector.

Traditional sampling calorimeters are positioned perpendicular to the direction of the incident particles. In ATLAS, the LAr calorimeter has a characteristic accordion geometry that allows the readout electrodes to be placed close to the active detector material, resulting in much faster signal collection. This fast response is essential for coping with the dense collision environment at the LHC. A sketch of the ATLAS LAr calorimeter with its distinct accordion structure around $\eta \approx 0$ is shown in Figure 2.8.

In the barrel region, the LAr calorimeter covers the pseudorapidity range $|\eta| < 1.475$ and consists of three segmented compartments. The first layer is finely segmented in the η direction with 4 mm strips and allows a good separation between prompt photons and $\pi^0 \rightarrow \gamma\gamma$ decays. The second compartment is segmented in the η and ϕ direction with towers of size $4 \times 4 \text{ cm}^2$. A




Figure 2.8: Schematic depiction of the ATLAS LAr calorimeter granularity around $\eta \approx 0$. [133]

typical EM shower will deposit most of its energy in this second compartment. The third layer is used to aid the energy measurement of very high energetic EM showers and has a factor of 2 coarser granularity compared to the second layer. In addition to these three layers, a thin 11 mm presampler is employed in front of the ECAL that provides a first sampling of the showers in front of the accordion LAr EM calorimeter.

The EMEC covers the range $1.5 < |\eta| < 3.2$ and has a similar layout as the EMB, but not all layers cover the full pseudorapidity range. The exact η coverage of the EM modules as well as of the hadronic and forward calorimeters is provided in Table 2.1.

The HCAL is positioned in the outer part of the calorimeter system and consists of the Tile in the barrel region and the hadronic endcap calorimeter (HEC) in the endcap region. The tile calorimeter consists of sampling layers in the central region and extended barrels that cover a total pseudorapidity region $|\eta| < 1.7$. Between the barrel and the extended barrel, gap and crack scintillators are disposed to help collect energy otherwise lost in these transition regions. The total thickness of the tile, in units of radiation lengths λ , is 7.4λ . It uses iron as the absorber material and scintillating tiles as active layers and has a $\Delta\eta \times \Delta\phi$ granularity of 0.1×0.1 .

The hadronic endcap calorimeter (HEC) covers the region $1.5 < |\eta| < 3.2$ and uses copper and LAr as passive and active medium with a granularity of $\Delta\eta \times \Delta\phi = 0.1 \times 0.1$ and 0.2×0.2 ,

2.3 The ATLAS Experiment

Calorimeter	Layers	Module Name	η -coverage	Sampling Layer
EM calorimeters	4	EM Barrel (EMB)	$ \eta < 1.5$	PreSamplerB, EMB1, EMB2, EMB3
	4	EM Endcap (EMEC)	$1.5 < \eta < 1.8$	PreSamplerE
			$1.5 < \eta < 3.2$	EME1, EME2
			$1.5 < \eta < 2.5$	EME3
Hadronic calorimeters	4	Hadronic Endcap (HEC)	$1.5 < \eta < 3.2$	HEC0, HEC1, HEC2, HEC3
	3	Tile Barrel (TileBar)	$ \eta < 1.0$	TileBar0, TileBar1, TileBar2
	3	Tile Extended Barrel (TileExt)	$0.8 < \eta < 1.7$	TileExt0, TileExt1, TileExt2
	3	Tile Gap (TileGap)	$1.0 < \eta < 1.6$	TileGap1, TileGap2, TileGap3
Forward calorimeter	3	FCal	$3.1 < \eta < 4.9$	FCal0, FCal1, FCal2
Transition regions		between barrel and endcap	$ \eta \approx 1.45$	
		between outer and inner wheel of endcap	$ \eta = 2.5$	
		between endcap and FCal	$ \eta \approx 3.2$	

Table 2.1: Breakdown of the various calorimeter systems in the ATLAS detector, including the various module and sampling layer names and the respective η coverages. The transition regions between the different modules are also shown. [134]

depending on the pseudorapidity of the respective region. The total thickness of the HCAL is about 11λ .

Finally, the forward calorimeter (FCal) completes the almost full solid angle coverage of the detector with a pseudorapidity range $3.1 < |\eta| < 4.9$. It uses LAr as active material, and copper and tungsten as absorber material. Copper is found in the innermost layer of the FCal and is designed for electromagnetic calorimetry, while tungsten is placed in the two outer layers and is optimized for hadronic measurements. The EM FCal layer has a depth of $28X_0$, and all three FCal modules have a combined depth of 10λ .

The ECAL, HCAL and FCal almost completely cover the full solid 4π angle of the detector. However, the transition region between the barrel and end-cap cryostats at $1.37 < |\eta| < 1.52$ suffers from a significant loss in energy resolution and particles in this region are typically excluded in most use cases. Similar transition regions are also found between the outer and inner wheel of the endcap at $|\eta| = 2.5$ and between the endcap and the FCal at $|\eta| \approx 3.2$. A visualization of the full ATLAS calorimeter geometry with all its 24 sampling layers is shown in Figure 2.9.

The intrinsic energy resolution of the calorimeter system depends on the specific calorimeter technology and can be parametrized as

$$\frac{\Delta E}{E} = \frac{a}{\sqrt{E}} \oplus \frac{b}{E} \oplus c, \quad (2.7)$$

where the first term $\sim E^{-1/2}$ is the stochastic term and results from intrinsic statistical shower and sampling fluctuations. The second term $\sim E^{-1}$ is the noise term and absorbs effects from electronic and instrumental noise of the readout chain. Finally, the constant term takes into

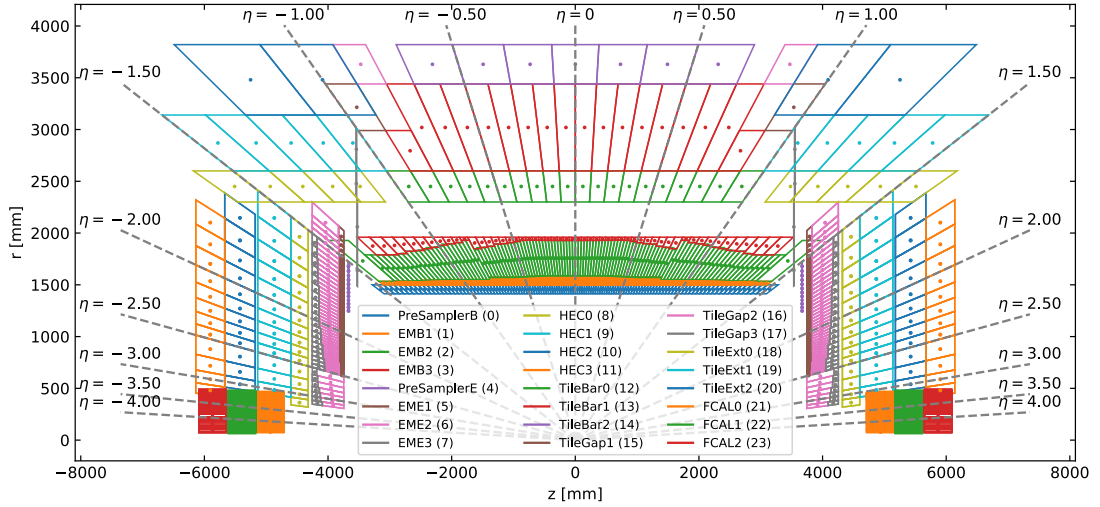


Figure 2.9: Visualization of the ATLAS calorimeter geometry. The Liquid Argon, the Tile and the Forward calorimeter subsystems are shown. The values in brackets indicate the number used internally to identify each layer. [135]

account any further contributions that are not dependent on the energy of the incident particle. Among other aspects, these include any lateral and longitudinal energy leakage or leakages due to cracks and dead regions in the detector. At the high energies of the LHC, the noise term typically does not play a role.

The ECAL has a design resolution of $\sigma_E/E = 10\%/\sqrt{E} \oplus 0.7\%$, whereas the hadronic calorimeter in the barrel and end-cap offer $\sigma_E/E = 50\%/\sqrt{E} \oplus 3\%$ and $\sigma_E/E = 100\%/\sqrt{E} \oplus 10\%$ in the forward direction. The actual values are typically determined from test beam measurements or simulations and can exhibit important dependencies on the pseudorapidity region of the detector. This is particularly the case for hadron-initiated showers, where the intrinsic fluctuations are large.

2.3.4 Muon Spectrometer

At the LHC, muons are typically produced with energies ranging between a few 100 MeV and a few 100 GeV. In this range, muons, which are about 200 times more massive than electrons, ionize matter very weakly and therefore deposit very little energy in the detector. As a result, and due to their long mean lifetime of $c\tau \approx 660$ m, muons pass through the detector with minimal material interactions and energy loss.

ATLAS uses a dedicated Muon Spectrometer (MS) [136] for the detection and identification of muons, located in the outermost part of the detector. A sketch of the Muon Spectrometer and its subsystems is shown in Figure 2.10. The MS is immersed in a dedicated radial magnetic field




Figure 2.10: Illustration of the ATLAS Muon Spectrometer. [111]

that reaches field strengths of up to 4 T. It is generated by a large toroidal magnet in the barrel region and two additional toroids in the end-caps, each consisting of 8 superconducting coils. The MS covers an acceptance of $|\eta| < 2.7$ and extends to the full 11 m radius of the detector. It uses 1171 chambers with a total of 354 240 Monitored Drift Tubes (MDTs) to measure the curvature of muon tracks. The 0.85-6.5 m long MDTs are 3 cm in diameter and have a spatial resolution of about $35 \mu\text{m}$ in the longitudinal direction. Cathode Strip Chambers (CSCs) are placed in the very forward direction of the MS, providing higher granularity and a resolution of $40 \mu\text{m} \times 5 \text{ mm}$ in $z \times r$. The long response time of the MDTs makes their use in the ATLAS trigger system (see Section 2.3.5) impractical. Therefore, Resistive Plate Chambers (RPCs) and Thin Gap Chambers (TGCs), which have faster response times, are additionally placed in the barrel and endcap regions, respectively, and cover a pseudorapidity region $|\eta| < 2.4$.

During LS2, the small wheel of the MS was replaced with the New Small Wheel (NSW) [137] that uses novel small-strip Thin Gap Chambers (sTGC) and Micromegas detectors, providing a $50 \mu\text{m}$ position resolution and a 1 mrad angular resolution. The NSW is required to maintain the same level of efficiency and momentum resolution in the harsher environment of the HL-LHC, and is shown as part of the sketch of the ATLAS detector in Figure 2.4. The overall relative transverse momentum resolution of the MS is p_T dependent and is 10% for particles with $p_T = 1 \text{ TeV}$. A table summarizing the resolution of all discussed detector subsystems is shown in Table 2.2.

Detector component	Resolution	η coverage	
		Measurement	Trigger
Tracking	$\sigma_{p_T}/p_T = 0.05\% \cdot p_T \oplus 1\%$	$ \eta < 2.5$	
TRT		$ \eta < 2.0$	
ECAL	$\sigma_E/E = 10\%/\sqrt{E} \oplus 0.7\%$	$ \eta < 3.2$	$ \eta < 2.5$
HCAL			
barrel and end-cap	$\sigma_E/E = 50\%/\sqrt{E} \oplus 3\%$	$ \eta < 3.2$	$ \eta < 3.2$
forward	$\sigma_E/E = 100\%/\sqrt{E} \oplus 10\%$	$3.1 < \eta < 4.9$	$3.1 < \eta < 4.9$
Muon spectrometer	$\sigma_{p_T}/p_T = 10\%$ at $p_T = 1$ TeV	$ \eta < 2.7$	$ \eta < 2.4$

Table 2.2: Design values for the resolution and η coverage of the main subsystems of the ATLAS detector. The energy and transverse momenta are in units of GeV. The values are taken from Ref. [111].

2.3.5 Trigger System

As mentioned in Section 2.1, the LHC features bunch crossings every 25 ns, giving a nominal bunch crossing rate of 40 MHz. Between 2015 and 2018, the average number of interactions per crossing was at 33.7, which equates to a total of about 1.7 billion proton-proton collisions per second. Limitations in data processing speed and storage capacity make it impossible to record the immense amount of data produced by the experiment. In addition, most of the collisions result in well-known Standard Model processes that are of little interest. ATLAS therefore uses a multi-level trigger system designed to quickly discard uninteresting events and record only those with interesting properties.

During Run 2, ATLAS used a two-level trigger system consisting of a hardware-based level-1 (L1) trigger and a software-based High Level Trigger (HLT) [138]. The L1 trigger uses coarse information from the calorimeter system and the MS to quickly identify events with interesting features such as high- p_T leptons and jets. For interesting events, the L1 trigger identifies Regions of Interest (ROIs) in the detector, which are further analysed by the HLT in the second step of the trigger system. The L1 trigger operates with a latency of $2.5 \mu\text{s}$ and reduces the event rate from 40 MHz to 100 kHz.

The High Level Trigger (HLT) is purely software-based and analyses the events accepted by the L1 trigger on a CPU farm of approximately 40 thousand processing units. The HLT reconstructs physics objects such as photons, electrons and jets similar to the full offline reconstruction of recorded events. However, to keep the trigger latency to a minimum, only the ROIs identified by the L1 trigger are used in the reconstruction. The reconstruction allows imposing specific requirements on the presence of physics objects and their kinematic properties. The HLT operates at a latency in the order of 100 ms and reduces the event rate from 100 kHz to about 1 kHz. At this rate, the output stream amounts to about 1 Gb/s and events can be permanently recorded for offline analysis.

Physics of Particle Shower Development

THE development and usage of calorimeters in high energy physics experiments has a long-standing history. Today, calorimeters have become indispensable instruments for measuring the energy of particles as well as for distinguishing between different particle types.

Fundamentally, the energy measurement of an incident particle relies on the deposition of its energy in a block of matter. As a particle enters the volume, it recurrently loses some of its energy as a consequence of a series of interactions with the material. The aggregation of resulting particles from these interactions is known as a *shower cascade*. Ideally, the particle and its remnants are fully halted within the block, such that no energy escapes from the measuring device. The deposited energy induces a signal in the detector, either in the form of electric charge or in the form of light, which is proportional to the energy of the incident particle. While the underlying measurement principle is simple, calorimeters are complex devices as a result of the large variety of interactions that a particle can undergo when traversing matter. It is therefore important to understand the processes that electromagnetic (EM) and hadronic interacting particles undergo when travelling through matter.

In this chapter, the types of interactions that neutral and charged particles undergo at different energy scales is briefly outlined. Moreover, the development of electromagnetic and hadronic showers and their respective properties are described.

3.1 Particle Interactions with Matter

Photons

Photons travelling through a block of matter can undergo a variety of interactions. Whether a photon is more likely to experience one or another process, strongly depends on the initial energy of the particle as well as on the material properties of the medium. In general, the interaction probability in a homogenous medium scales with the atomic number Z of the constituent atoms. The most dominant processes in the according energy range are:

1. **Photoelectric effect:** the predominant interaction for low energy photons in the eV to MeV range with energies larger than the electron binding energy. Electrons in the shell of an atom absorb the energy of the incoming photon and are eventually ejected. The resulting vacancy can be filled with electrons from outer atomic shells. The energy released in the process can in turn be emitted in the form of X -ray photons or Auger electrons. The cross section approximately scales with $\sigma \sim Z^{4-5}$.
2. **Rayleigh scattering:** low energy photons in the eV to MeV range may also undergo elastic scattering with atomic electrons. No energy transfer occurs in the process and atoms persist in their ground state. The cross section typically scales with $\sigma \sim Z^2$.
3. **Compton scattering:** the predominant interaction for mid-ranged energy photons in the MeV range. Photons undergo inelastic scattering with an electron, such that part of the incoming photons' energy is transferred to the recoiling electron. The material sensitivity of the cross-section is significantly lower compared to Rayleigh scattering and the photoelectric effect and only scales approximately linearly with Z .
4. **Electron pair production:** The dominant process for high energetic photons in the range between a few MeV and higher. In the presence of an electromagnetic field, a photon can convert into an electron positron pair ($\gamma \rightarrow e^+e^-$) in the case that $E_\gamma > 2m_e$ such that energy-momentum conservation is satisfied. Similar to Rayleigh scattering, the cross section scales with approximately $\sigma \sim Z^2$.
5. **Photonuclear reactions:** Very high energetic photons can also undergo photonuclear reactions in which high energy gamma rays are absorbed in the nucleus. Consequently, the excited nucleus emits a neutron, proton or alpha particle. While the cross section for nuclear processes is very low in general, specific photon energies can cause numerous de-excitation events, most notably at the giant dipole resonance [139].

The partial photon interaction cross section for all the mentioned processes for lead ($Z = 82$) spanning several orders of magnitude between 10 eV and 100 GeV are shown in Figure 3.1(a). The circles indicate the experimentally measured values. As indicated, the photoelectric effect clearly dominates for photons with $10 \text{ eV} < E_\gamma < 1 \text{ MeV}$. For slightly higher photon energies and in a relatively short energy window, Compton scattering is dominant, while electron-pair production becomes dominant at approximately 9 MeV. Most of the electron-positron pairs

produced originate from interactions in the electric field of the nucleus. The probability of pair production in the field of the atomic electrons is approximately an order of magnitude lower. The giant dipole resonance is the sole relevant contributor in terms of nuclear effects and is clearly visible around a photon energy of $E_\gamma \approx 12$ MeV.

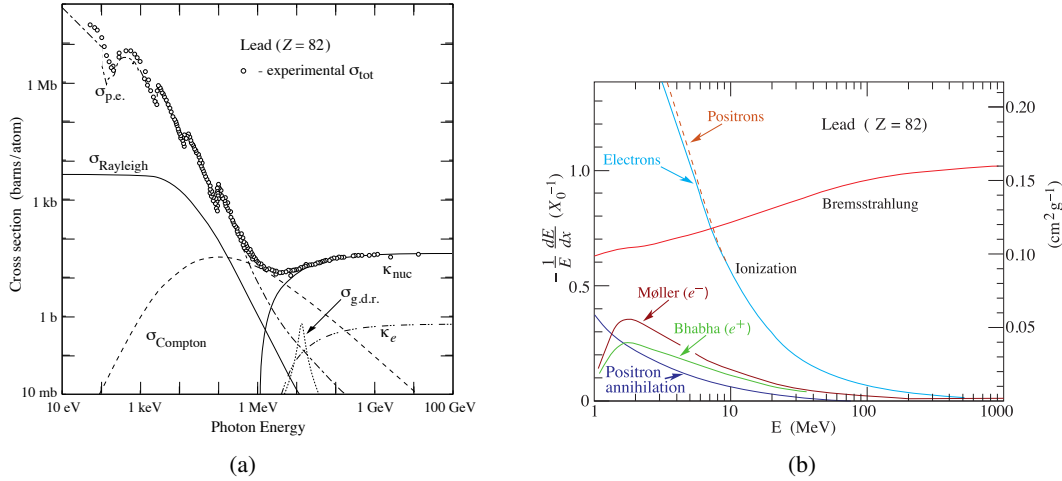


Figure 3.1: (a) Photon interaction cross section σ for a variety of processes as a function of the photon energy in lead. The circles indicate the experimentally measured total cross section. The partial cross sections correspond to the photoelectric effect ($\sigma_{\text{p.e.}}$), Rayleigh scattering (σ_{Rayleigh}), Compton scattering (σ_{Compton}), pair production in the field of the nucleus (κ_{nuc}), pair production in the field of the electrons (κ_e), and the giant dipole resonance ($\sigma_{\text{g.d.r.}}$), a photonuclear process. (b) Fractional energy loss through various processes of electrons and positrons as a function of the initial energy in lead. [43]

Charged Particles

In general, charged particles lose their energy through Coulomb interactions with the atomic electrons, through interactions with the nuclei or through radiative Bremsstrahlung. For electrons and positrons up to energies in the order of MeV, muons of a few hundred GeV and even higher energetic hadrons, the energy loss through interactions with shell electrons is dominant. The atomic energy loss includes ionization processes, thermal excitation of the atoms and the scattering with shell electrons. At higher energies, radiative losses start to become dominant. Highly energetic charged particles interact with the electromagnetic field of the atoms and are deflected. As a consequence, the charged particle is decelerated and emits a stream of photons, known as Bremsstrahlung. The energy loss through Bremsstrahlung is approximately proportional to Z^2 and inversely proportional to the quadratic rest mass of the incoming particle, and thus muons lose ~ 200 times less energy as a result of Bremsstrahlung compared to electrons of identical energy.

While significantly less likely, Coulomb interactions with the nuclei also occur and dominate the deflection of charged particles in matter. In dense media, charged particles undergo millions of small displacements due to multiple scattering events. On an event level, the outcome of such a series of deflections can be well described on a statistical basis using Molière theory [140]. Even though nuclear interactions play a large role in the macrophysical development of a multiple scattering event, the resulting energy loss is generally insignificant. As an example, Figure 3.1(b) shows the fractional energy loss per unit path length $-dE/dx$ for various processes of electrons and positrons in lead. One can identify two distinct regions: in the low energy regime up to 10 MeV, the energy loss through collisions with atoms, in particular via ionization processes, dominates. Other processes such as electron-electron (*Møller*) and electron-positron (*Bhabha*) scattering also play a non-negligible role in this region, although not to the same extent as Bremsstrahlung. For $E_{e^\pm} \gtrsim 10$ MeV, energy losses through ionization decrease sharply and Bremsstrahlung becomes the dominant source of energy loss. From approximately 100 MeV onwards, almost all energy loss is due to radiative effects, while no further energy dependence in the fractional energy loss for electrons and positrons is observed for $E_{e^\pm} \gtrsim 1$ GeV. At energies below 10 MeV minor differences between electrons and positrons emerge in the ionization scale, whereas for higher energy scales any discrepancies in the energy loss between the two particle types vanish.

The mean rate of energy loss by heavy charged moderately relativistic particles can be described by the *Bethe* equation [43] and is also known as *mass stopping power*:

$$\left\langle -\frac{dE}{dx} \right\rangle = Kz^2 \frac{Z}{A} \frac{1}{\beta^2} \left[\frac{1}{2} \ln \frac{2m_e c^2 \beta^2 \gamma^2 W_{\max}}{I^2} - \beta^2 - \frac{\delta(\beta\gamma)}{2} \right], \quad (3.1)$$

where z is the electric charge number of the incoming particle, K a material-independent proportionality factor, A the atomic mass of the absorber, I the mean excitation energy. W_{\max} is the maximum possible energy transfer to an electron in a single collision and $\delta(\beta\gamma)$ is a density effect correction [141] to ionization energy loss, effectively taking into account the polarization of the material. The parametrization is valid in the ionization dominated region $0.1 < \beta\gamma < 1000$ with errors in the order of magnitude of a few percent with respect to the measured data.

As an example, Figure 3.2 shows the mass stopping power of a muon entering a copper volume for $0.001 < \beta\gamma < 10^5$. The red line shows the validity region of the *Bethe* equation. As the latter is derived from a first order Born approximation, the model is not valid for particles with $\beta\gamma < 0.1$ and higher order corrections need to be taken into account to calculate accurate predictions below that threshold. In the very low energy regime, e.g. for hadrons of less than several hundred eV, nuclear losses dominate the total energy loss [142], while in the keV region, the mass stopping power is proportional to β and is accurately described by the *Lindhard-Scharff* model [143]. In the range between $0.01 < \beta\gamma < 0.05$, no suitable theory exists, such that this region is usually described using phenomenological fits first performed by

Andersen and Ziegler [144]. Particles with momentum at the minimum ionization scale are known as *minimum ionizing particles*. It should be emphasized that for all minimum ionizing particles $\beta\gamma \approx 3$ holds, and the mass stopping power is $1\text{-}2 \text{ MeV g}^{-1} \text{ cm}^{-2}$ in most materials. For larger values of $\beta\gamma$, radiative losses gradually increase. At the critical energy ϵ_c , the energy lost as a result of ionization equals the energy lost due to Bremsstrahlung. After exceeding the critical energy, radiative losses become dominant at TeV momentum scales.

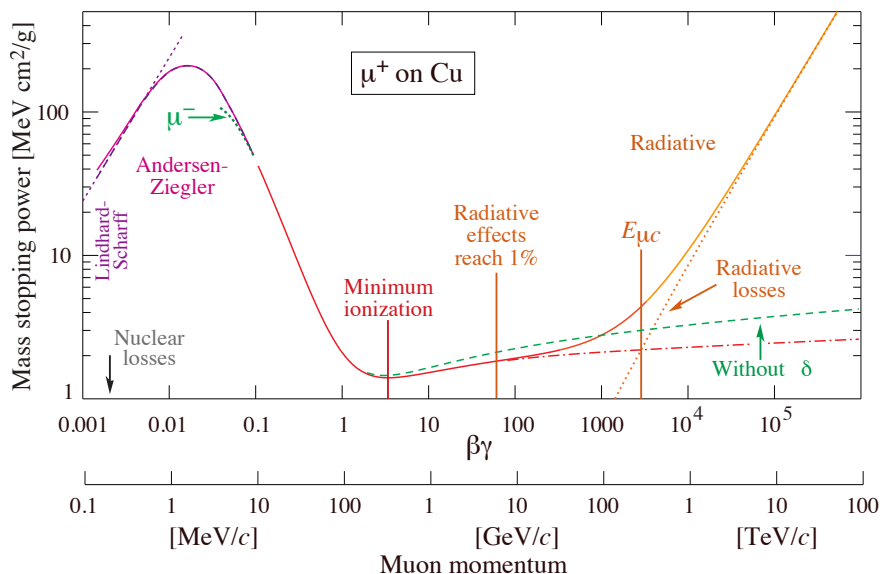


Figure 3.2: Mass stopping power of a muon entering a copper volume over several magnitudes of muon momentum. [43]

Hadrons

In addition to the general statements made in the case of charged particles, hadrons can also undergo hadron-nucleus interactions via the strong force. Most notably, hadrons with energies in the order of a few MeV and higher can interact strongly with the nucleons of the atom and produce new hadrons as a result. The production of hadrons can consequently lead to intranuclear cascades [145] in which high energetic hadrons can escape the nucleus altogether. After this initial fast process with a timescale of $\mathcal{O}(10^{-22} \text{ s})$, the nucleus is generally left in an excited state. As a result, the atom undergoes a process of atomic relaxation, known as *evaporation*, in which free nucleons and photons are emitted isotropically. This second stage of the *spallation* process is much slower, with $\mathcal{O}(10^{-16} \text{ s})$. In the case that the energy of the incident hadron is high enough, the nuclei can split in its entirety into two or more smaller nuclei, a process known as *nuclear fission*.

Other noteworthy processes include the decay of neutral hadrons. In shower cascades, pions are produced extensively. While charged pions can interact electromagnetically, their neutral counterpart decays almost instantly ($\tau \approx 8.52 \cdot 10^{-17}$ s) into a pair of photons with a branching ratio of $\Gamma(\pi^0 \rightarrow \gamma\gamma) = 98.823\%$ [43]. Neutrons, on the other hand, can also decay ($n^0 \rightarrow p^+ + e^- + \bar{\nu}_e$), but have a significantly larger lifetime ($\tau \approx 878.4$ s), such that most slow neutrons are captured by the nuclei. The resulting atom is typically in an excited state and quickly decays to the ground state by the emittance of highly energetic gamma rays. In fact, the only way for neutrons to interact with material is through nuclear reactions.

3.2 Development of Particle Showers

While the physics of individual particles traversing matter is well understood, the properties of shower cascades are described in a statistical manner for all practical purposes. The longitudinal and lateral energy profiles for EM showers, consisting of photons, electrons and positrons, can be precisely described in terms of simple parametrizations. The development of hadronic showers, in turn, is much more complex as a result of the vast variety of possible interactions involving electromagnetic, strong and nuclear interactions, such that analytical treatments are not available in most cases. In this section, the mechanisms of EM and hadronic shower development as well as the most important shower properties are described.

3.2.1 Electromagnetic Showers

When a high energy photon or electron enters an absorber volume, Bremsstrahlung and pair production generates additional photons and e^\pm pairs, respectively. These secondaries will only carry a fraction of the incident particle. Through the same mechanism, the resulting secondary particles will create further particles, which in turn create more particles, effectively forming the electromagnetic shower cascade. The number of particles in the shower increases until the energy of the secondary particles falls below a critical energy ϵ_c at which the energy loss occurs mainly due to dissipating processes such as ionization and excitation. As a result, very few particles are created and the shower eventually comes to a halt.

Most EM shower properties can be described using the *radiation length* X_0 , which is defined as the average characteristic length that an electron travels such that it loses $1 - 1/e \approx 63.3\%$ of its original energy. The average energy of electrons with initial energy E_0 after travelling a distance x is then given by

$$\langle E(x) \rangle = E_0 \cdot e^{-\frac{x}{X_0}}. \quad (3.2)$$

The radiation length is a material-dependent constant and can be approximately described [146]

by

$$X_0 \approx \frac{716.4 \text{ g cm}^{-2} A}{Z(Z+1) \ln \frac{287}{\sqrt{Z}}}, \quad (3.3)$$

where values for materials typically used in high-energy experiments range from 0.6 cm for lead to 13.84 cm for iron and 14 cm for liquid argon [43].

While electrons entering a material immediately start dissipating their energy via Bremsstrahlung, high energy photons, on average, cover the *mean free path* $\lambda = \frac{9}{7} \cdot X_0$ before decaying into an electron-positron pair. As a result, showers initiated by photons start slightly later and deposit their energy deeper in the material. Moreover, due to the statistical nature of the mean free path, fluctuations in the energy deposited in a given block of matter are larger for photon-induced showers. Figure 3.3 shows an illustration of the development of a photon-induced electromagnetic shower cascade.




Figure 3.3: Sketch of the development of an electromagnetic shower cascade originating from a photon entering an absorber volume. Adapted from Ref. [147].

The longitudinal energy profile can be well described using a gamma distribution [148]:

$$\frac{dE}{dt} = E_0 b \frac{(bt)^{a-1} e^{-bt}}{\Gamma(a)}, \quad (3.4)$$

where $t = x/X_0$ is the penetration depth measured in multiples of radiation length, E_0 the initial energy of the shower inducing particle and a and b are free, particle dependent parameters. Figure 3.4(a) shows the longitudinal shower profile of photons between 1 GeV and 1 TeV in lead as simulated using GEANT4 [149], a toolkit for the precise simulation of the passage of particles through matter. Fits to the simulation from a nonlinear regression using Equation (3.4) are illustrated as dashed lines and are shown to effectively describe the underlying functional form.

The shower maximum exhibits a logarithmic dependence on the initial energy of the particle:

$$t_{\max} \simeq \ln \frac{E_0}{\epsilon_c} + t_0, \quad (3.5)$$

where $t_0 = 1/2$ for photons and $t_0 = -1/2$ for electrons. As a consequence of this small energy dependence, calorimeters can be built as very compact devices such that even for particles in the TeV range, most of the energy can be absorbed for penetration depths in the order of $20X_0$. Figure 3.4(b) shows a comparison of the fractional energy deposit as a function of the penetration depth for photons and electrons with an energy of 10 GeV. As mentioned earlier, photon showers are initiated slightly later such that their energy is deposited deeper in the material and the shower maximum is shifted to larger values.

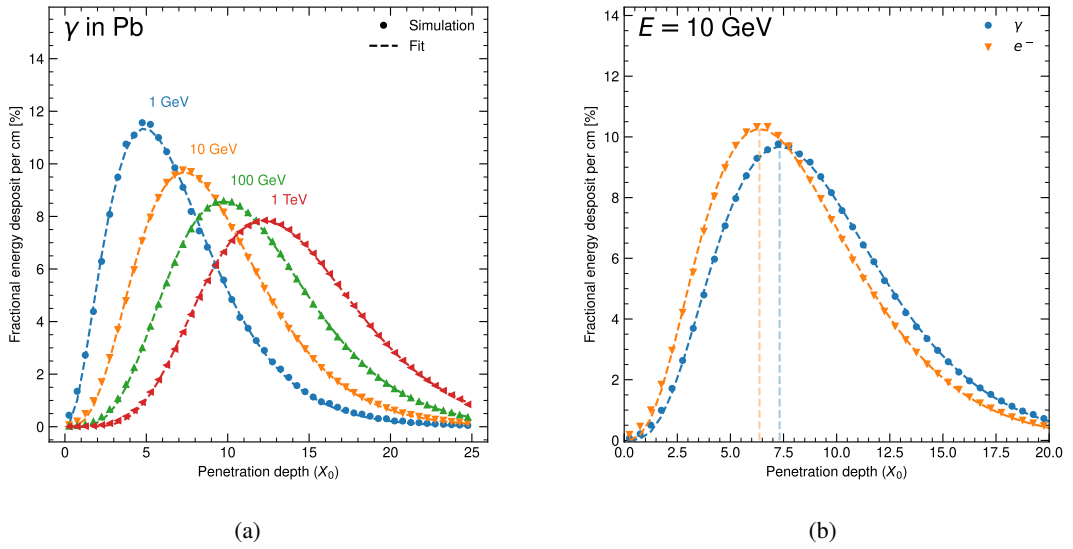


Figure 3.4: (a) GEANT4 simulation of the fractional energy deposit per cm as a function of the penetration depth in multiples of X_0 in lead for photons with energies between 1 GeV and 1 TeV. The dashed lines indicate fits of the simulations with gamma distributions. (b) Comparison of the fractional energy deposit as a function of the penetration depth in lead between photons and electrons with an energy of 10 GeV. The dashed lines indicate the penetration depth at the shower maximum.

While the energy loss in the longitudinal direction is mainly dominated by Bremsstrahlung and dissipative effects such as ionization and excitation, the transversal shower development is driven by multiple scattering processes of low energy electrons and positrons in the EM field of the atoms. A measure of the transverse spread of the EM shower is the *Molière radius* R_M , which is defined as the radius of a hypothetical cylinder which, on average, contains 90% of the shower energy. The Molière radius is approximately independent of the inducing particle type

and can be solely expressed as material-dependent quantity [146]:

$$R_M \simeq 21 \text{ MeV} \frac{X_0}{\epsilon_c [\text{MeV}]}, \quad (3.6)$$

where values for typical calorimeter systems are in the order of a few cm. The Molière radius for liquid argon is about 9 cm for iron 1.7 cm and for lead 1.6 cm [43].

3.2.2 Hadronic Showers

While the development of hadronic showers is conceptually similar to the creation of EM cascades, hadron showers are significantly more complex due the diverse particle content and much larger number of interactions that strongly interacting particles can undergo. If a high energy hadron enters an absorber volume, it deposits some of its original energy via ionization of the material. After travelling some distance, the hadron undergoes a nuclear reaction with the material. As a result, new high energy secondary hadrons, most of which are pions, emerge. The impacted atoms remain in an excited state and emit γ -rays in consequence of the evaporation process. Secondary hadrons then repeat this process, effectively forming the hadronic cascade.

The average distance a hadron travels through a slab of material before undergoing a nuclear interaction is known as the interaction length λ_{int} , and, similar to X_0 for EM showers, is a characteristic quantity governing the evolution of hadronic showers. The characteristic length of EM showers is generally much shorter ($\lambda_{\text{int}} \gg X_0$), such that hadronic showers tend to penetrate much deeper into the material and form sub-showers across the material, which make them much wider compared to their electromagnetic counterpart. Lighter hadrons such as mesons have a lower probability of interacting with the atomic nucleus, such that the interaction length is generally larger compared to baryons. Strong interactions are charge independent, such that, on average, 2/3 of the created pions are π^\pm , while 1/3 consist of neutral pions. The subsequent $\pi \rightarrow \gamma\gamma$ decay initiates an electromagnetic component to the hadronic shower cascade and irreversibly transfers energy from hadronic to electromagnetic interactions.

An illustration of the development of a neutron-initiated hadronic shower cascade is shown in Figure 3.5. Approximately 1/3 of the hadrons initial energy will be carried away by the EM component of the hadronic cascade after the first interaction. As $\lambda_{\text{int}} \gg X_0$, the energy of the EM component is likely to be entirely deposited in the absorber volume before secondary hadrons undergo further nuclear reactions. As a result, the initial stage of the hadronic shower development is fully dominated by the EM component and the penetration depth at which most of the energy is deposited approximately independent of the hadron's incident energy. This effect is shown in Figure 3.6, which shows a GEANT4 simulation of the longitudinal shower profile of pion-induced hadronic showers in a lead volume with incident energies between 10 GeV and 200 GeV. It must be emphasized, however, that the longitudinal energy profiles




Figure 3.5: Sketch of the development of a hadronic shower cascade originating from a neutron entering an absorber volume. The shower develops an electromagnetic and a hadronic component. The circles indicate hard fragments in the cascade. Adapted from Ref. [147].




Figure 3.6: Geant4 simulation of the longitudinal shower profile of pions entering a lead absorber volume with incident energies between 10 GeV and 200 GeV.

shown here are averaged over many pion-induced shower cascades. Longitudinal shower profiles of individual showers can differ substantially from the average energy depositions. In fact, the physics behind hadronic shower cascades is fully driven by event-to-event fluctuations in the energy profile. While π^0 mesons are typically created after the first nuclear reaction, individual showers may only create neutral pions after several interaction lengths. Consequently, the EM component of such showers will develop in completely different regions of the material, creating showers that look very different compared to the average.

It is also noteworthy that soft neutrons and photons emitted from nuclear reactions can carry a significant energy fraction of the shower initiating particle, but remain mostly undetected in calorimeter devices due to their delayed emittance. This *invisible energy* created in hadronic showers make the development of precise hadronic calorimeters a very challenging endeavour.

Fast Calorimeter Simulation in ATLAS

ONE of the most important backbones of the ATLAS physics programme is the extensive software and computing infrastructure that is required for running the experiment as well as collecting, processing, storing and analysing data effectively. Hence, it is vital for the collaboration to monitor the development of its current and future computing resources in order to allow for a smooth functioning of the experiment.




Figure II.1: Projected evolution of CPU usage from 2020 to 2036, under a conservative and aggressive R&D scenario. The black lines indicate the evolution under the assumption of a total 10% and 20% computing capacity increase. The vertical shaded areas show projected data taking periods. [150]

With the start of Run 3, the LHC operates with a record number of collisions per bunch

crossing of $\langle\mu\rangle \sim 55$ which is expected to be increased to up to 140 for the start of the HL-LHC in the late 2020s and up to 200 in Run 5. With significantly higher luminosity in Run 3 and up to 7 to 10 times increased data rates at the HL-LHC, simulated data with even higher statistics compared to the collected data must be produced. In particular, SM background processes will need to be modelled with unprecedented precision and more exotic phase spaces for Beyond the Standard Model searches need to be populated.

However, the production of MC events is a very CPU intensive task and the limited amount of MC events is already now one of the largest sources of systematic uncertainties in many ATLAS analyses. While the rising CPU requirements can be partially alleviated with increasing computing capacities, the collaboration is limited by a yearly flat budget with a 10-20% annual capacity increase. As a consequence, it is crucial for ATLAS to find alternative means of saving computing resources without sacrificing physics performance. Figure II.1 shows the projected annual CPU consumption from 2020 to 2036, under a conservative and aggressive R&D scenario in ATLAS. The conservative scenario assumes the completion of projects that can be performed with a sustained level of person-power, whereas the aggressive scenario includes projects which are at high risk of not converging without additional collaborators and a higher time fraction spent on software and computing activities from existing contributors.




Figure II.2: Distribution of ATLAS CPU usage by activity in 2018. Adapted from Ref. [151].

Figure II.2 shows the fractional CPU usage in 2018 split by the various computing activities in ATLAS. With 38% of the total CPU time, the simulation of MC events is the most intensive task by a wide margin, followed by the reconstruction step with 18% of the total usage. In contrast, event generation makes up a relatively small fraction of approximately 11%. For the

detector simulation, approximately 80% of the CPU time is typically spent on the simulation of electromagnetic and hadronic showers in the calorimeter. In particular, the `GEANT4` simulation of the complex accordion structure in ATLAS has a significant impact on the simulation speed.

In order to mitigate the CPU load from detector simulation, ATLAS has successfully employed a fast simulation tool known as `ATLFASTII` (AF2) [152], which is based on a parametrized simulation of the calorimeter known as `FASTCALOSIM` [153] and `GEANT4` for the simulation of the Inner Detector. For the simulation of complex physics processes, AF2 is about 6 to 7 times faster compared to `GEANT4` and has been used to produce approximately 32 out of the 52 billion events ($\sim 60\%$) that have been used to analyse Run 2 data. Consequently, AF2 had a significant impact in the reduction of overall CPU usage. While AF2 generally offers good physics performance, the simulation is known to have non-negligible limitations, in particular in the simulation of jets reconstructed with large-radius clustering algorithms and the modelling of their substructure variables. For this reason, many analyses were unable to use fast simulation in ATLAS until now and improvements in physics performance of fast simulation remains a major goal of the collaboration. In fact, one of the major factors in the reduction of CPU usage between the conservative and aggressive R&D scenario shown in Figure II.1, is the assumption that the fraction of full simulation used by ATLAS can be reduced to 25 and 10% for Run 3 and Run 4, respectively. Figure II.3 illustrates the projected breakdown of computing usage for the two scenarios in Run 4. In the conservative scenario, full simulation accounts for up to 24% of




Figure II.3: Projection for Run 4 of the breakdown of compute usage for the conservative (b) and aggressive (a) R&D scenario. [150]

the total CPU consumption, whereas in the aggressive scenario, the usage of `GEANT4` can be restricted to account for approximately 11%.

In order to reach these targets, ATLAS released a successor of AF2 known as `ATLFAST3` [134] with significantly improves physics accuracy, especially in the simulation of jet substructure. `ATLFAST3` is based on an updated parametrized simulation (`FASTCALOSIMV2`) and a machine-

learning based tool using Generative Adversarial Networks (GANs) [154] for the simulation of hadron showers known as FastCaloGAN [155]. ATLFast3 has already been used extensively to re-simulate 7 billion Run 2 events and is the default simulation tool for Run 3.

It is noteworthy that while the simulation of the calorimeter represents the largest single share of CPU usage and hence gives the greatest potential for saving computing resources, further tools to speed up the MC production of physics events are already being employed or are in active development. With ATLFast3 in place, the overall simulation time is dominated by the simulation of particle tracks in the Inner Detector. ATLAS is therefore also developing a fast simulation of charged particle propagation that uses a simplified geometry and parametrized material effects known as Fast Track Simulation (FATRAS) [156]. As track reconstruction scales combinatorially with increasing number of charged particles, it is expected to become the most computationally intensive component of the production chain at the HL-LHC and beyond (see Figure II.3). ATLAS therefore is also developing fast reconstruction techniques such as truth-seeded track reconstruction [157] as well as parametrized methods for fast digitization. Furthermore, instead of simulating pile-up by sampling individual inelastic interactions, ATLAS will overlay pre-sampled minimum-bias-events to the hard-scatter interaction, a method known as RDO-overlay [158]. The combination of all these tools aiming to take shortcuts in the pile-up and detector simulation, as well as reconstruction and digitization is what is known as the Fast Chain [159]. The ultimate goal of Fast Chain is to provide a single pipeline without intermediate steps that enables the fast simulation, digitization and reconstruction of complex physics events in the order of a few seconds, while keeping any degradation in physics performance to a minimum.

In this part, ATLFast3, the state-of-the-art fast simulation tool employed for Run 3, is described in detail in Chapter 4, with a particular focus on FASTCALOSIMV2. In Chapter 5, the derivation of energy response corrections aimed at improving the total energy response simulation of EM and hadronic showers with FASTCALOSIMV2 is presented. Furthermore, the development of a new shower centre position extrapolation engine, required for the correct assignment of energies within the calorimeter geometry, is described in Chapter 6. In Chapter 7, the physics performance of the ATLFast3 version employed in the re-processing of Run 2 samples is presented, with a particular emphasis on reconstructed objects typically used in ATLAS analyses. In Chapter 8, the development of two distinct models for tuning electromagnetic shower shapes to data are described. Finally, a proof-of-concept integration of FASTCALOSIMV2 as a GEANT4 fast simulation model is presented in Chapter 9.

The AtFast3 Simulator

ATlFast3 [134] is the next generation of fast simulation which is to be employed for the official MC production campaign during Run 3 and beyond. It replaces the costly simulation of EM and hadronic showers using GEANT4 [149] with a parametrized detector simulation (FASTCALOSIMV2). Instead of simulating every physics process and material interaction of primary and secondary particles produced in the particle showers, energy deposits are sampled from pre-defined parametrizations that depend on the particle type, incident energy and direction in the detector.

Hadrons in the intermediate energy range are simulated with a set of Generative Adversarial Networks [154], specifically Wasserstein GANs with gradient penalty [160, 161], which are embedded in a separate simulation package known as FASTCALOGAN [155]. FASTCALOGAN has been shown to perform especially well in simulating correlated fluctuations between the longitudinal and lateral energy distributions in hadronic showers and is thus used in phase space regions in which it can efficiently aid the baseline simulation with FASTCALOSIMV2.

Very low energetic pions with $E_{\text{kin}} < 200$ MeV and other hadrons with $E_{\text{kin}} < 400$ MeV are simulated with GEANT4, as their energy response cannot be precisely parametrized using AF3. As low energetic hadrons are very fast to simulate, this choice does not negatively alter the overall speed of the simulation.

The Inner Detector is fully simulated using GEANT4. FATRAS [156], which is under active development, offers an alternative approach for fast track simulation in the Inner Detector, but does not yet reach the required precision to be used for the production of MC samples which are to be used for general physic analyses.

As muons penetrate most material without undergoing interactions and do not create shower cascades, their simulation is fast and can be performed using GEANT4 in the calorimeter and the Muon Spectrometer, without affecting general computing performance.

Secondary particles created in hadronic showers may escape the calorimeter system and penetrate the Muon Spectrometer. These *muon punch-through* particles create hits in the MS and therefore need to be modelled accurately. As any information on the trajectory and particle properties of the muon punch-through particles is lost through the parametrization procedure, the punch-through probability of particles is modelled dependent on the momentum and direction of the incoming particle. Punch-through particles are then sampled accordingly and passed to be simulated with GEANT4. An illustration giving an overview of ATLFast3 with its various simulators is given in Figure 4.1.

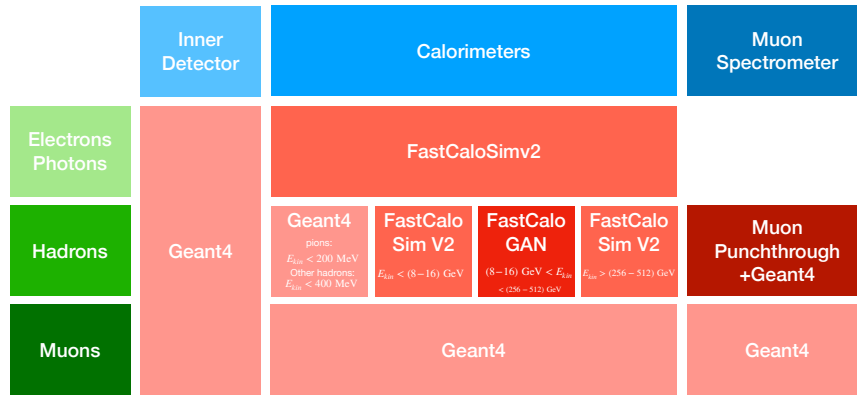


Figure 4.1: Illustration of the ATLFast3 configuration showing the various simulators that are being used depending on particle type, detector component and kinematic range. [134].

ATLFast3 is currently fully integrated in the Integrated Simulation Framework (ISF) [162] within the ATLAS software framework Athena [163]. However, efforts have started to abandon ISF altogether and replace its functionality by integrating fast simulation tools such as FASTCALOSIMV2 directly into GEANT4 (see Chapter 9).

In this chapter, the parametrized fast calorimeter simulation employed in ATLAS, FASTCALOSIMV2, is discussed in detail. The parametrization and simulation of the lateral and longitudinal shower profiles of electromagnetic and hadronic shower is described and some benchmarks assessing the computing performance of ATLFast3 and FASTCALOSIMV2 are presented. The overall physics performance of ATLFast3 is presented in Chapter 7.

4.1 Parametrization Input

In order to precisely simulate the calorimeter response to complex physics processes, the energy response as well as the longitudinal and lateral development of showers induced by particles reaching the calorimeters needs to be modelled precisely. Electromagnetic and hadronic shower

development will, however, not only vary by type and energy of the shower inducing particle, but also strongly depend on the position at which the shower is initiated.

The fundamental strategy behind FASTCALOSIMV2 is to exploit information from precisely pre-simulated EM and hadronic showers initiated by single particles which is used to derive a parametrization of the calorimeter response. At the simulation stage, the appropriate parametrization is chosen depending on the type of the particle, its kinematic properties and incoming direction to simulate its response, effectively avoiding the computationally expensive simulation and propagation of each primary and secondary particle in the material.

Only relatively stable particles reach the calorimeter surface and consequently only these will deposit energy in the calorimeter system and are required to be parametrized. Other particles that have very short lifetimes and decay before reaching the calorimeter system, such as τ -leptons ($c\tau \approx 87 \mu\text{m}$) and neutral pions ($c\tau \approx 25 \text{ nm}$), can be indirectly parametrized through the showers induced by their respective decay products. Three particle types are used to parametrize the calorimeter response to hadronic and electromagnetic showers:

1. **Electromagnetic showers** are parametrized using single photons, electrons and positrons. As a result of the slightly later onset of EM showers induced by photons, their energy is deposited deeper in the calorimeter such that a separate parametrization is needed. Differences in showers induced by electrons and positrons are negligible and as such, the parametrization is based on an equal charge combination of e^\pm .
2. **Hadronic showers** are parametrized with a combination of positively and negatively charged pions. As for electrons and pions, differences in the shower development between π^+ and π^- are negligible, and thus the parametrization is derived from an equal charge combination. Other stable hadrons may also induce hadronic showers in the calorimeter, but at significantly lower rates for most processes of interest. The shower development of other stable hadrons that induce hadronic showers in the calorimeter is regarded as sufficiently close enough to be described using the pion parametrization. Remaining differences in the total energy response are corrected using ad-hoc correction factors (see Chapter 5).

The parametrization of the calorimeter response is performed by generating single particles at a virtual boundary defined using a set of three cylinders, which approximately matches the physical boundary between the Inner Detector and the calorimeter system. The radius and half-lengths of the corresponding cylinders used to define the boundary are given in Table 4.1. The cylinder C_1 with radius $R = 1148 \text{ mm}$ and half-length $Z = 3550 \text{ mm}$ covers the boundary between ID and EM barrel and endcap calorimeter. In addition to partially covering the boundary between ID and endcap, the cylinder C_2 with $R = 120 \text{ mm}$ and $Z = 4587 \text{ mm}$ also covers the boundary between ID and FCal. Finally, the cylinder C_3 with $R = 41 \text{ mm}$ and $Z = 6783 \text{ mm}$ defines the beam pipe and covers the remaining boundary between the beam pipe and the FCal. A three-dimensional representation of the boundary as defined within FASTCALOSIMV2 is shown in Figure 4.2.

	C_1	C_2	C_3
Radius R	1148 mm	120 mm	41 mm
Half-length Z	3550 mm	4587 mm	6783 mm

Table 4.1: Geometric values of the three cylinders C_1 , C_2 and C_3 that define the boundary between ID and calorimeter used for the parametrization of FASTCALOSIMV2.




Figure 4.2: Three-dimensional representation of the boundary between Inner Detector and calorimeter system which is used to generate particles that are used for the parametrization derivation of FASTCALOSIMV2.

For the parametrization, single photons, electrons and pions are generated in a total of 100 distinct $|\eta|$ regions ranging from 0 to 5 with a step size of 0.05. The parametrization extends to slightly larger values than the calorimeter acceptance (see Table 2.1), in order to include particles that only deposit a fraction of their energy in the calorimeter system. Within the pseudorapidity bins, the particles are sampled from a uniform η distribution, covering both negative and positive values. Given the approximate symmetry of the detector in the azimuthal angle, the parametrization is not dependent on the ϕ angle of the incoming particle such that particles are assigned uniformly distributed azimuthal angles. The energy dependence in the shower development is taken into account by dividing the parametrization into a total of 17 logarithmically spaced bins of total truth momentum ranging from 64 MeV to 4 TeV. The particles are generated at the virtual boundary between ID and calorimeter in order to avoid any interference from material interactions in the Inner Detector. The momentum of the generated particles is fixed to one of the 17 parametrization points and the direction is chosen to be consistent with the particle production at the interaction point.

An illustration showing the spatial position and direction of particles generated at the boundary between ID and calorimeter system for a reduced number of pseudorapidity values is shown in Figure 4.3. The number of single particle events generated and used for the




Figure 4.3: Spatial position and direction of particles generated on the boundary between the ID and calorimeter system (blue) to derive the FASTCALOSIMV2 parametrization. The black dots show the positions on the boundary and the orange arrows their respective momentum directions that are used to generate the particles. The dotted pink lines indicate integer values of η for reference. For illustration purposes, the step size between neighbouring η values is increased by a factor of 5 in comparison to the true granularity used in the parametrization.

parametrization depends on the momentum of the shower inducing particles. For low to intermediate momenta between 64 MeV and 524 GeV, a total of 10 000 events are simulated. For larger momenta, the number of events is gradually reduced as a result of the increasing computational expense of shower simulation with GEANT4. A summary table showing the number of events simulated depending on the truth momentum of the generated particle is shown in Table 4.2. In the following, the particles from a given energy parametrization point are often referred to be $E = X$ GeV particles. In this case, X refers to the total (truncated) truth momentum of the shower inducing particle from which the parametrization is derived from. A 65 GeV photon, for instance, corresponds to the energy parametrization point for which particles with a truth momentum of 65536 MeV on the logarithmically spaced energy scale were generated. As massless particles, this corresponds to the incident energy of photons, but does not equal to the energies of pions and electrons.

The showers induced by the generated particles on the boundary between ID and calorimeter system are simulated using the GEANT4 toolkit in order to create a set of reference data sets on

Momentum	64 MeV-262 GeV	524 GeV	1 TeV	2 TeV	4 TeV
Events	10 000	5000	3000	2000	1000

Table 4.2: Number of single particle events generated depending on the truth momentum of the shower-inducing particles.

which the parametrization is based on. For the Run 2 parametrization, showers were simulated using `GEANT4` 10.1.3 [149] with the `FTFP_BERT_ATL` physics list [164] that uses the Bertini intra-nuclear cascade model [165] below 9 GeV and transitions to the Fritiof model [166, 167] with a pre-compound model for 12 GeV and higher. The default `GEANT4` physics list is employed for the simulation of EM showers. While the standard `GEANT4` configuration employed by ATLAS during Run 2 uses a library of Frozen Showers [152] as a fast simulation technique to simulate EM showers in the FCal, the option is deactivated for the simulation of the input samples. Furthermore, the `GEANT4` simulation is run in a special configuration with significantly smaller `GEANT4` steps, such that details of the spatial position of deposited energies can be stored and used for the derivation of the lateral energy parametrization (see Section 4.3). As the interaction region is not point-like as a result of the non-negligible spread of the LHC beam bunches, a smearing of the vertex position is typically applied in the simulation of physics samples. For the purpose of generating the samples used for the parametrization, no vertex smearing is applied. Moreover, energy deposits are digitized without the inclusion of electronic noise and cross-talk between calorimeter cells in the readout electronics.

As an example, Figure 4.4 shows the total deposited energy of the single photons simulated with `GEANT4` in the $p - |\eta|$ phase space. Individual parametrization points are marked with white crosses. As expected, the total deposited energy is directly proportional to the incident photon momentum and remains constant for most of the detector regions. In the transition regions between barrel and endcap ($|\eta| \approx 1.45$) as well as between outer and inner wheel of endcap ($|\eta| = 2.5$) and between endcap and FCal ($|\eta| \approx 3.2$) not all energy from the incident photon is deposited in the active material, such that the total deposited energy is smaller in these regions. Similarly, the total deposited energy decreases when nearing very forward regions that are outside the FCal acceptance ($|\eta| \approx 4.9$).

Finally, the total energies deposited by photons in the EM barrel region $0.20 < |\eta| < 0.25$ and with varying momentum from 64 MeV (leftmost distribution) to 4 TeV (rightmost distribution) corresponding to the 17 parametrization points is shown in Figure 4.5(a). The widths of the distributions become negligible in comparison to the shown energy scale for large photon momenta, such that the distributions are seen as individual lines. Figure 4.5(b) shows the mean total energy deposited by photons with momenta corresponding to the parametrization points as function of $|\eta|$. The transition region and out-of-acceptance effects in the average deposited energy are clearly visible.




Figure 4.4: Total energy deposited in the ATLAS calorimeter from a photon as simulated using GEANT4 depending on incident momentum and incoming direction. The white crosses indicate the points in the momentum- $|\eta|$ phase space that are used to derive the parametrization used in FASTCALOSIMV2.




Figure 4.5: (a) Total energy deposited by photons with varying momentum scales ranging from 64 MeV (leftmost) to 4 TeV (rightmost) and incident direction range of $0.20 < |\eta| < 0.25$, simulated using GEANT4. For large photon momenta, the widths of the distributions become negligible in comparison to the shown energy scale, such that the distributions are visible as individual lines. (b) Mean total energy deposited by photons with varying momentum scales ranging from 64 MeV (bottom) to 4 TeV (top) as a function of incident direction, simulated using GEANT4.

4.2 Longitudinal Shower Development

As shown in Table 2.1 in Chapter 2, the calorimeter system in ATLAS consists of a total of 21 distinct layers from the EM and hadronic calorimeter system in addition to three FCal layers in the forward region. A particle inducing a shower in the calorimeter system will therefore deposit its energy across various layers. The amount of energy deposited in each layer will strongly depend on the initial energy of the shower initiating particle and its incoming direction. In order to simulate the overall detector response, the energy deposit in each of the layers needs to be determined. However, the amount of energy deposited in each layer is highly correlated between layers, such that it is not possible to easily parametrize the response in each layer independently. In order to decorrelate the energies deposited in the layers, a Principal Component Analysis (PCA) [168] is performed for each particle within each η region and truth momentum point used in the parametrization. For the energy parametrization, only layers with energies larger than 0.1% of the total deposited shower energy are taken into account and are hereinafter referred to as relevant layers. In order to derive the energy parametrization, the following steps are performed:

1. For each parametrization point and particle type, the energy fractions in the relevant layers computed from the GEANT4 reference samples are transformed into Cumulative

Distribution Functions (CDFs) by the integration over the probability densities of the energy fractions. Example energy fractions for the PreSamplerB and EM Barrel layers as well as the resulting CDF is shown in Figure 4.6 for photons with energies between 32 and 131 GeV in the region $0.20 < |\eta| < 0.25$. In this case, most of the energy is deposited in the EMB2 and EMB1 layers, whereas only little energy is found in the PreSamplerB and EMB3 layer. Higher energy photons that penetrate deeper in the material, deposit larger energy fractions in the outer layers of the calorimeter system, whereas less energetic photons deposit more energy in the layers closer to the calorimeter boundary.

2. The energy fractions in each layer are transformed using the inverse error function $\text{erf}^{-1}(x)$, specifically using the transformation $\epsilon \mapsto \pi/2 \cdot \text{erf}^{-1}(y)$, where $y = 2 \cdot \text{CDF}(\epsilon) - 1$ and ϵ is the respective energy fraction, effectively transforming the energies into Gaussian distributions which are centred at zero. The resulting transformed energy distributions and their correlations for the first four calorimeter layers are shown in Figure 4.7 for 32 GeV, 65 GeV and 131 GeV photons in the region $0.20 < |\eta| < 0.25$. As expected, the transformed energies exhibit strong correlations. For instance, the energies deposited in the layers EMB1 and EMB2 are almost fully anti-correlated with a Pearson correlation coefficient of $\rho = -0.99$ in the case of 65 GeV photons. The black arrows show the orthogonal eigenvectors of the covariance matrix which point towards the maximum variance of the energies and are of relevance for the Principal Component Analysis in the next step.
3. The Gaussian transformed energy deposits computed in the previous step are used as input to the Principal Component Analysis. Traditionally used for the purpose of data reduction, PCA is an orthogonal linear transformation which *rotates* the energies into a new set of coordinates, in this case linear combinations of the transformed energies, in a way that the variance of the data is highest for the first coordinate which is generally known as the leading principal component. Subsequent components have gradually decreasing variance, that is, encode less information of the data set. PCA effectively converts the transformed energies in a set of approximately uncorrelated variables. The principal components are eigenvectors of the covariance matrix and the corresponding eigenvalues encode the amount of variance along the eigenvector direction. Figure 4.8 shows the distributions and correlations of the principal components of the transformed energy distributions for 32 GeV, 65 GeV and 131 GeV photons in the region $0.20 < |\eta| < 0.25$. The distributions are largely uncorrelated and the eigenvectors of the PCA transformed data approximately point towards the directions of the principal component axes. The leading PCA component, denoted as PCA0, encodes the largest variance ($\approx 82\%$) of the data, which becomes especially apparent in comparison with PCA3, the component with the least information content ($\approx 0.6\%$ of the total variance). In the transformed coordinate system, the data points are stretched along a large range of PCA0 values, but have very little spread in PCA3. The leading PCA component (PCA0) is used to divide

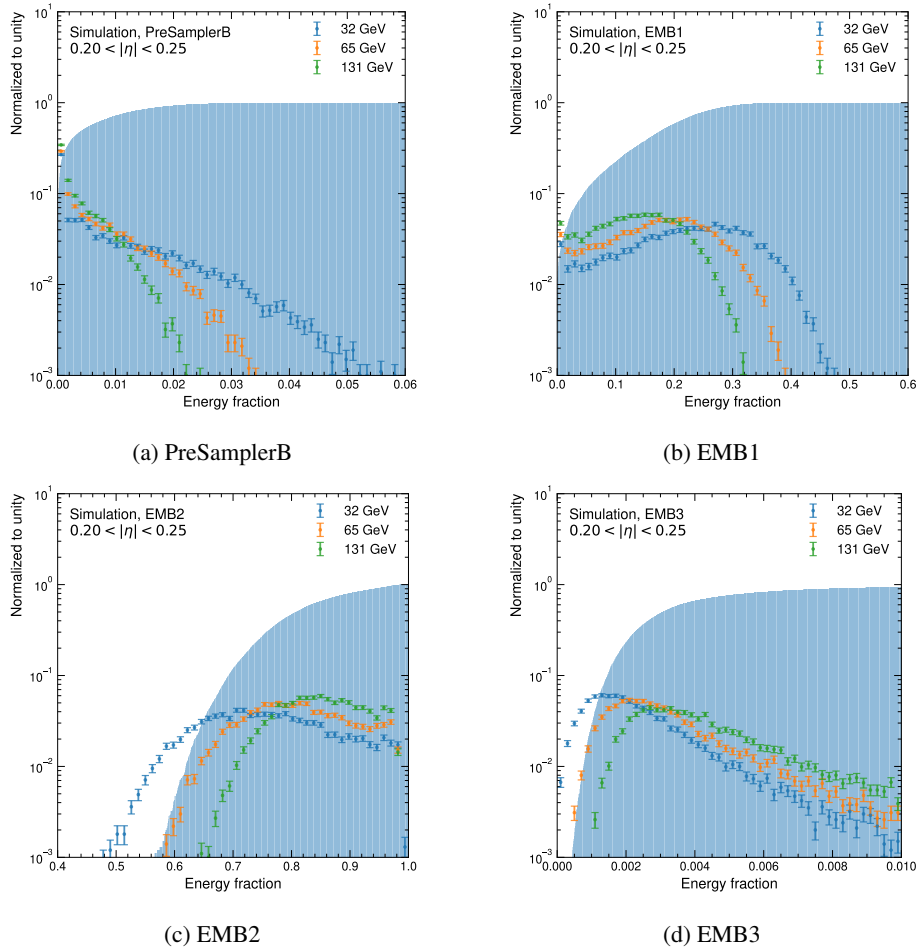


Figure 4.6: Fraction of energy deposited in the first four layers of the calorimeter for EM showers initiated by 32 GeV (blue), 65 GeV (orange) and 131 GeV photons (green) in the region $0.20 < |\eta| < 0.25$. The shaded area shows the cumulative distribution, which are used as input for the Gaussian transform, in the case of 65 GeV photons. Note the very different energy fraction axis ranges for the various calorimeter layers.

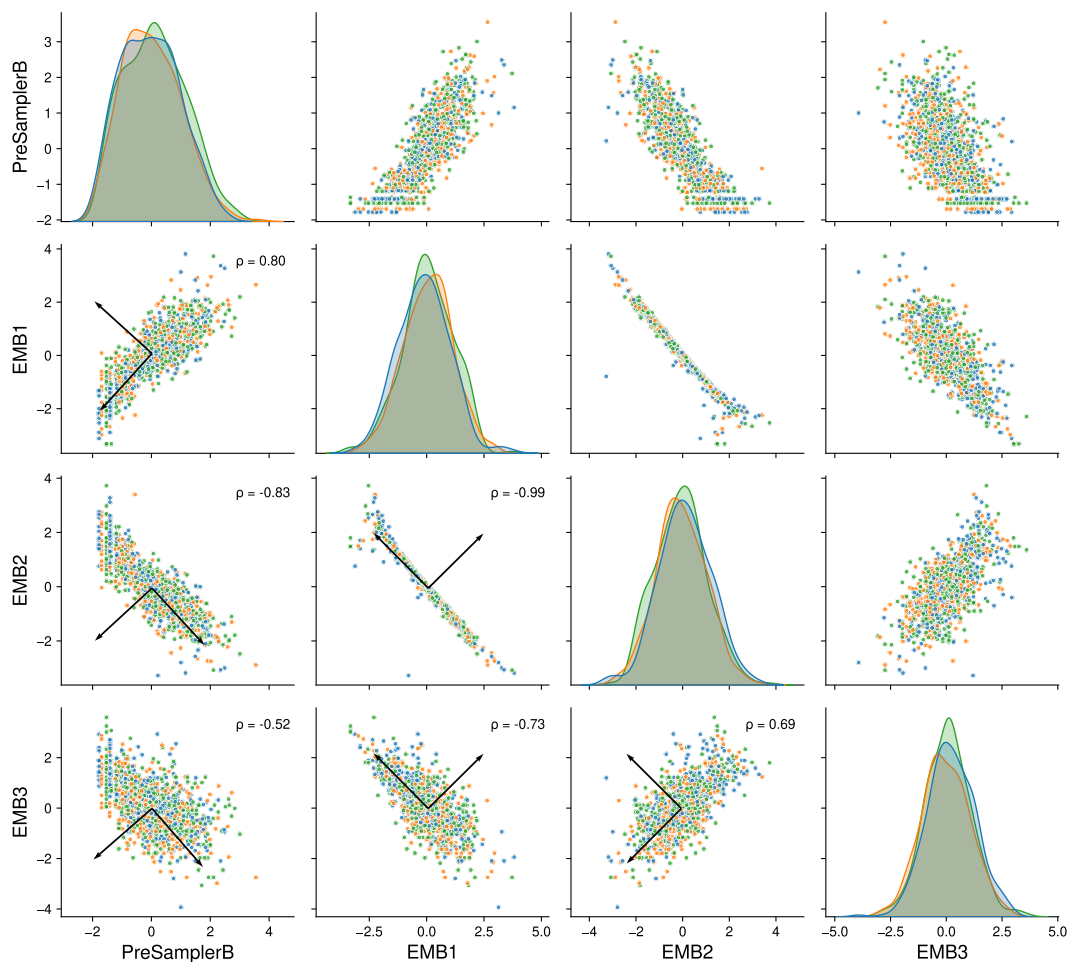


Figure 4.7: Distributions and correlations of the transformed energies in the first four layers of the calorimeter before performing the PCA rotation for 32 GeV (blue), 65 GeV (orange) and 131 GeV photons (green) in the region $0.20 < |\eta| < 0.25$. The black arrows indicate the orthogonal eigenvectors of the covariance matrix that point towards the directions of maximum variance. For each pairwise combination of layers, the Pearson correlation coefficient is given. For illustration purposes, the eigenvectors and correlation coefficients are only shown in the case of 65 GeV photons.

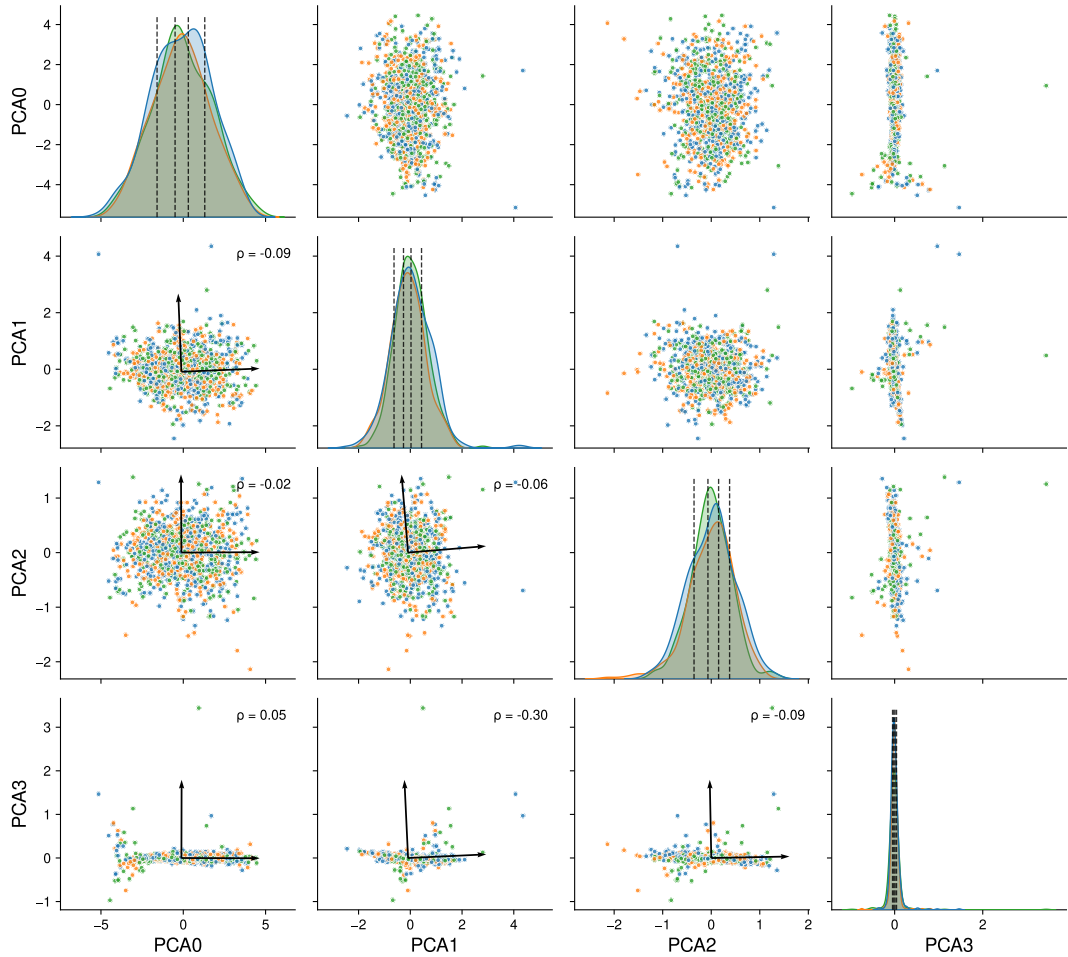


Figure 4.8: Distributions and correlations of the principal components after performing the first PCA rotation for 32 GeV (blue), 65 GeV (orange) and 131 GeV photons (green) in the region $0.20 < |\eta| < 0.25$. The energies are projected along the directions of maximum variance, resulting in a set of uncorrelated variables. The dashed vertical lines indicate the PCA bin borders constructed such that each bin approximately contains the same number of shower events. For each pairwise combination of layers, the Pearson correlation coefficient is given. For illustration purposes, the eigenvectors and correlation coefficients are only shown in the case of 65 GeV photons.

the GEANT4 dataset into five equal sized bins, that is, each event contains approximately the same number ($\approx 20\%$) of shower events. The resulting bin borders for all PCA components are also shown in Figure 4.8, but only the leading component is used in most of the cases. Within each PCA bin, the electromagnetic and hadronic showers exhibit similar features. More specifically, the showers are approximately classified by shower depth, defined as

$$d_{\text{shower}} = \frac{\sum_i^{\text{layers}} E_i \cdot d_i}{\sum_i^{\text{layers}} E_i}, \quad (4.1)$$

where d_i is the radial central position in case of barrel layers, z_i the z position of the endcap and forward layers centre and E_i the energy deposited in layer i . Figure 4.9 shows the shower depth of electromagnetic and hadronic showers induced by 65 GeV photons, electrons and pions in the region $0.20 < |\eta| < 0.25$ split into the five PCA bins which the showers are assigned to. In all cases, a fair separation in shower depths is observed for the various PCA bins. The figures also illustrate the very distinct features of electromagnetic and hadronic showers. While the depth of showers induced by photons and electrons is concentrated along a narrow distance between 1.66 m and 1.78 m, pion showers penetrate the calorimeter much deeper and the range of shower depths expands to almost 2 m. It is noteworthy that while in most cases a total of 5 PCA bins are used from the leading PCA component, this is not always the case. The exact number of bins used is based on a χ^2 test used to optimize the modelling of the energy fractions. Especially in the transition regions of the calorimeter system, the sub-leading PCA component is used additionally to define two-dimensional PCA bins.

4. After the classification of the showers in one of the PCA bins, a second Principal Component Analysis is performed within each of the bins. Identical to the first PCA transform, the energy fractions of showers in each first-PCA bin as well as the total energy are converted into Gaussian distributions which serve as inputs to the second PCA. The aim of the second PCA is to produce fully uncorrelated Gaussian distribution, effectively removing any remaining correlations between the first-PCA components. At this point, the mean and the standard deviation of the Gaussianized energy distributions of the second PCA, the correlation matrices, the PCA bin probabilities as well as the inverse cumulative distributions are stored and later used in the simulation.

During the simulation process, the described steps are performed backwards. Once a particle reaches the boundary between Inner Detector and calorimeter system, the appropriate parametrization is loaded from memory. According to the PCA bin probabilities, a random PCA bin is generated. As such, this step makes a random choice on how deep the energy of the shower will be deposited in the calorimeter. Then, uncorrelated random numbers are generated from the Gaussian distributions obtained from the second PCA and mapped into correlated Gaussian numbers using the inverse correlation matrix. With the error function and inverse CDFs computed in the first PCA, these correlated numbers are finally used to obtain the total energy

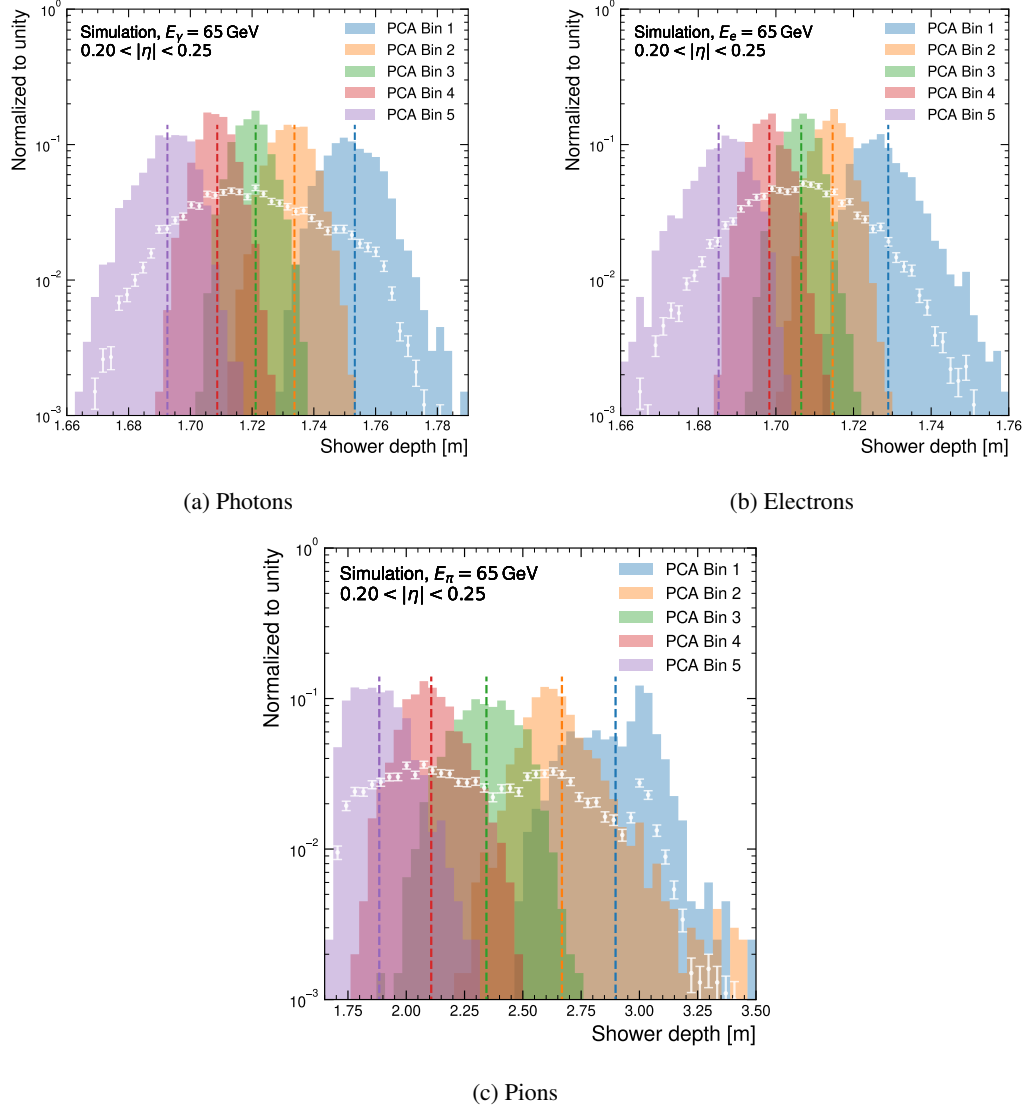


Figure 4.9: Shower depth of electromagnetic and hadronic showers induced by 65 GeV photons, electrons and pions in the region $0.20 < |\eta| < 0.25$ classified into the five PCA bins after the first PCA rotation. The vertical dashed lines indicate the mean of the respective distributions. The white overlaid data points show the shower depths of all showers combined and can be considered a linear superposition of the showers in the individual PCA bins.

as well as the energy fractions deposited in each layer. This method ensures that the correlations between energy deposits in the various layers remain intact after simulation. Figure 4.10 shows the total energy energy deposited and the energy fractions in the first four layers of the calorimeter of 65 GeV photons for photons simulated with GEANT4 and FASTCALOSIMV2. In all cases, the distributions are compatible within uncertainties.

4.3 Average Lateral Shower Shape

In the previous section, the simulation of the total energy deposited by an incident particle and the distribution across the layers of the calorimeter is described. While this defines the longitudinal development of the shower, the shape within each layer, that is, the lateral shower shape is also required to be modelled precisely. The accurate modelling of the lateral shape of showers is especially important for the identification of particles and hence is crucial to many ATLAS analyses (see Chapter 8). Fundamentally, the lateral shower shape is defined by the distribution of spatial energy deposits in each layer and the energy that is assigned to these hits. The task of the simulation hence consists of sampling hit positions from pre-defined parametrizations and using adequate models to assign energy fractions obtained from the simulation of the longitudinal shower development to the hits in the respective layers.

In order to define an average lateral shower shape within each layer, the spatial energy deposits from GEANT4 are grouped into volumes, hereinafter called *voxels*, and are designed to be significantly smaller compared to the individual calorimeter cell dimensions. These voxels define three dimensional bins which are used for the parametrization of the average lateral shower shape in each of the layers. The lateral coordinates of the GEANT4 energy deposits in the relevant layers are described relative to the extrapolated shower centre position of the particle in that layer and are denoted $\Delta\eta = \eta^{\text{hit}} - \eta_c^{\text{extr}}$ and $\Delta\phi = \phi^{\text{hit}} - \phi_c^{\text{extr}}$. The extrapolation is performed by transporting the particle across the calorimeter, based on its initial momentum when entering the calorimeter system and taking into account the magnetic field in the detector (see Chapter 6 for an in-depth explanation). The relative coordinates $\Delta\eta$ and $\Delta\phi$ are hereinafter often referred to as distance from the shower centre in the direction of the respective coordinate. For a more convenient description, the relative hit coordinates are transformed into units of millimeters:

$$\Delta\eta^{\text{mm}} = \Delta\eta \times \eta_{\text{Jacobi}} \times \sqrt{r_{\text{cell}}^2 + z_{\text{cell}}^2} \quad (4.2)$$

$$\Delta\phi^{\text{mm}} = \Delta\phi \cdot r_{\text{cell}}, \quad (4.3)$$

where r_{cell} and z_{cell} are the spatial coordinates of the individual GEANT4 hits in the calorimeter

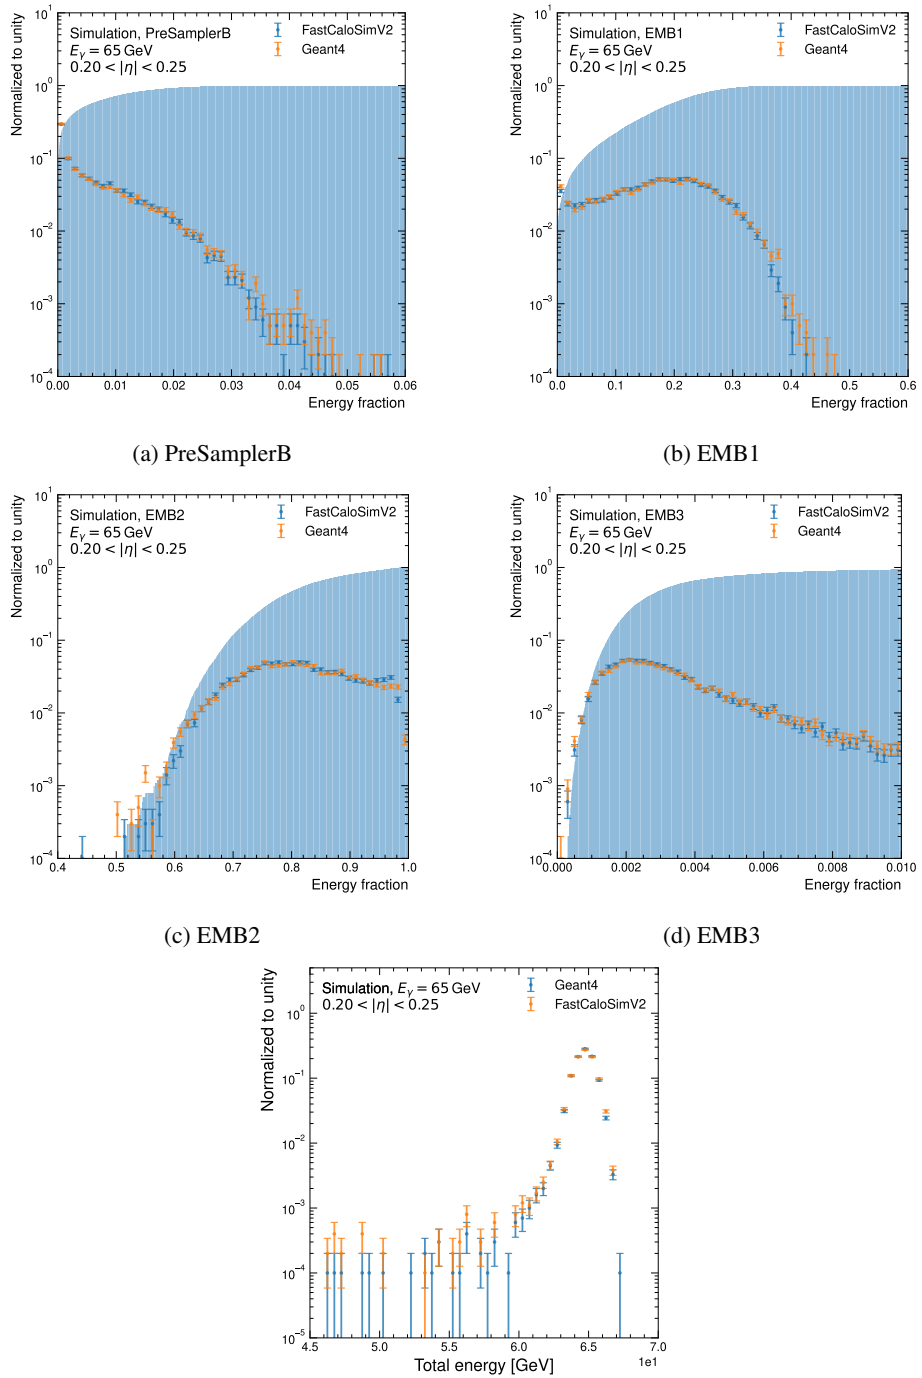


Figure 4.10: Fraction of energy deposited in the first four layers of the calorimeter and total energy for EM showers initiated by 65 GeV photons in the region $0.20 < |\eta| < 0.25$ as simulated by GEANT4 and FASTCALOSIMV2. The shaded area shows the cumulative distribution.

and

$$\eta_{\text{Jacobi}} = \left| \frac{2 \cdot \exp(-\eta_{cs}^{\text{extr}})}{1 + \exp(-2\eta_c^{\text{extr}})} \right|. \quad (4.4)$$

The $\eta - \phi$ coordinates are then transformed into polar coordinates describing the radial distance from the shower centre ΔR^{mm} (in millimeters) and the polar angle α in the $\Delta\eta - \Delta\phi$ plane:

$$\Delta R^{\text{mm}} = \sqrt{(\Delta\eta^{\text{mm}})^2 + (\Delta\phi^{\text{mm}})^2} \quad (4.5)$$

$$\alpha = \arctan2(\Delta\phi^{\text{mm}}, \Delta\eta^{\text{mm}}) \quad (4.6)$$

The lateral parametrization is then derived from the GEANT4 samples in eight distinct bins in the polar angle. In the radial direction, a binning of 1 mm is used for the highly granular EMB1 and EME1 layers and 5 mm for all others. The number of bins is tuned to cover more than two Molière radii, that is, containing approximately 99.5% of the shower's energy deposition. The parametrization is performed separately for each particle type, η bin and discrete value of truth momentum as well as PCA bin the shower is assigned to. In order to create a Probability Density Function (PDF) of the average lateral shower shape, the energy deposits are normalized to the total energy deposited in each parametrization slice. Figure 4.11 shows an example PDF of the average lateral shower shape for EM and hadronic showers initiated by 265 GeV photons and pions in the region $0.55 < |\eta| < 0.60$ for the EMB2 and TileBar2 layers, respectively. The densities here are shown for showers from all PCA bins, while the simulation takes into account differences in the lateral shower shape between the various PCA bins. The shown PDFs




Figure 4.11: Lateral shower energy density in the $\Delta\eta - \Delta\phi$ plane for photons (a) and pions (b) with a momentum of 265 GeV and $0.55 < |\eta| < 0.60$ in the EM Barrel 2 and the Tile Barrel 2, respectively, for all PCA bins combined. The radial (circular) lines indicate the boundaries between the angular (radial) bins used in the shape parametrization. [134]

demonstrate the fact that the energy in EM showers is largely distributed around the core of the shower, while hadronic showers are, on average, much wider compared to EM showers and can deposit their energies very far away from the shower centre. In order to reduce the memory footprint of the parametrization, the ϕ -symmetry of the detector is exploited and only $|\phi|$ values are used and mirrored along the ϕ axis to extend the parametrization to negative ϕ values.

In principle, the parametrization of the average shower shape projects the hits within a layer (and hence within voxels) into the 2D $\Delta\eta - \Delta\phi$ space, such that information on the longitudinal position of hits within each layer is lost. In order to retain the information, the average longitudinal position in the GEANT4 samples for each of the parametrizations is stored for every layer and PCA bin, and is applied as a correction to the simulated hit positions (see next Section 4.4). A more detailed description on how the longitudinal position within the individual layers is taken into account is provided in Chapter 6.

4.4 Simulation of Calorimeter Hits

As a rough approximation, the 2D lateral shape histograms can be directly used for simulating the lateral shape of the EM and hadronic showers, which in fact, is the approach taken by the predecessor of FASTCALOSIMV2. While this approach works reasonably well for the simulation of showers induced by photons and electrons, it fails to accurately simulate the complex structure of hadronic showers as a result of their intrinsically large fluctuations.

Hit Distribution

Instead of directly using the average lateral shower shape, FASTCALOSIMV2 draws a random sample of hits based on the lateral shower shape PDFs described in Section 4.3 which are consequently corrected for their longitudinal position in the individual layers. In the interest of computing performance, the resulting hits are mapped onto a simplified geometry of the calorimeter, effectively avoiding the direct placement of energy deposits on the complex liquid-argon accordion structure. In the simplified geometry, each cell is described as a cuboid in η , ϕ and r for barrel layers, η , ϕ and z for endcap layers and in cartesian coordinates for forward layers. The assignment of hits on cuboids of the simplified geometry introduces a small bias in the energy distributions with respect to GEANT4, which is corrected by the addition of a small, random displacement in the ϕ position of the hit prior to mapping it onto the simplified geometry. In all cases, the hit energies are normalized such that the addition of energy deposits exactly sum up to the total energy deposited in the respective layer.

In order to reproduce the fluctuations in the lateral shower shape of EM and hadronic showers, the number of hits that are drawn from the average shape PDFs is tuned to reproduce the intrinsically expected resolution of the respective calorimeter technology. In order to recreate

the resolution σ_E/E in a given layer, the number of hits N_{hits} drawn is determined from sampling from a Poisson distribution with an expectation value

$$\langle N_{\text{hits}} \rangle^{-1} = \left(\frac{\sigma_E}{E} \right)^2. \quad (4.7)$$

Figure 4.12 shows the probability of drawing N hits for photons and electrons with varying incident energies for layers in the LAr and Tile calorimeters. The resolutions to compute the distributions are taken from Ref. [134]. With increasing incident particle energy, the number of generated hits increases. As the LAr EM layers have significantly better energy resolution than the Tile layers, more hits are typically created for LAr layers.

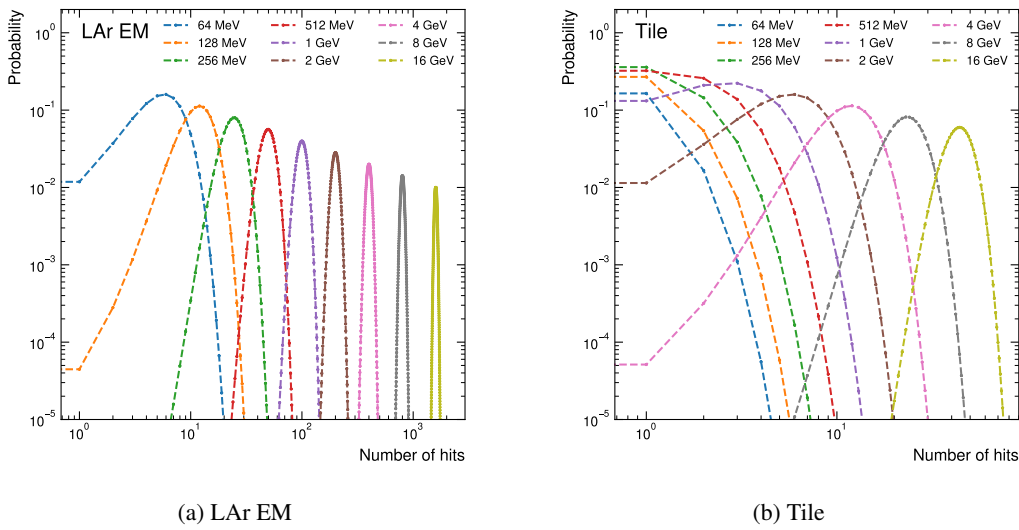


Figure 4.12: Poisson probability distributions used to simulate the number of hits distributed across the layers in the case of photons and electrons with varying incident energies between 64 MeV and 16 GeV for the LAr EM (a) and Tile (b) layers. The means of the distributions correspond to the inverse of the squared resolution of the respective detector technology.

In principle, the same strategy is applied to EM and hadronic showers. For hadron-initiated shower, however, the relative resolutions in individual layers are significantly larger and exhibit a non-negligible η dependence, which has significant implications for the modelling of hadronic showers. These differences in η are taken into account by extracting the constant and stochastic terms by computing the sampling fraction $f(E)$ from special charged pion simulations in each η slice and fitting their relative resolution σ_E/E to Equation (2.7), with exclusion of the noise term. The fitted values of the stochastic term are between 30-40% in the EM layers of the calorimeter, 50-60% for the Tile calorimeter, 60-80% for the HEC layers and 80-100% for the FCal. The constant term, on the other hand, ranges between 1-10%.

Hit Energy assignment

Once the hits are sampled from the lateral shape parametrization described in Section 4.3 and mapped onto the simplified detector geometry, the total energy simulated within each detector layer is distributed among the sampled hits.

For photon and electron induced showers, the energies are distributed equally among the hits:

$$E_{\text{hit}} = \frac{E_i}{N_{\text{hits}}}, \quad (4.8)$$

where E_{hit} is the energy of the individual hits, E_i the total energy deposited in a layer i and N_{hits} the Poisson distributed number of hits sampled within a layer. As photon and electrons will deposit most of their energy in the LAr EM layers and the energy resolution is dominated by the stochastic term of 10.1% (see Section 2.3.3), hits produced by EM showers will typically be assigned a value of

$$\langle E_{\text{hit}} \rangle = \frac{E_i}{\langle N_{\text{hits}} \rangle} \approx 10 \text{ MeV} \quad (4.9)$$

for the hit energy.

To some extent, the uniform distributions of energies among the hits, hereinafter called equal hit model, is an arbitrary choice which is observed to precisely reproduce reconstructed quantities related to the simulation of EM showers. For the simulation of charged pions, however, the equal hit model is shown to produce significant differences in the formation of calorimeter clusters and hence in the reconstruction of hadronic objects. The underlying reason becomes apparent when analyzing the dependence between the voxel energy and hit energy energy ratio $E_{\text{voxel}}/E_{\text{hit}}$ and the radial distance ΔR^{mm} from the shower centre as shown in Figure 4.13 for 65 GeV charged pions in the EMB2 layer in the range $0.20 < |\eta| < 0.25$. Figure 4.13(a) shows the distribution in the GEANT4 simulation: for distances far away from the shower centre, only few voxels contain energies of the order of E_{hit} and most contain only a small fraction. In the case of the equal hit model shown in Figure 4.13(b), the voxels contain, by definition, multiples of E_{hit} . As a result, FASTCALOSIMV2 will create a significant amount of voxels containing energies of the order of E_{hit} far away from the shower center. Due to the large stochastic terms for hadrons, the resulting hit energies can range between 100 and 300 MeV, which in many cases is sufficient to seed the creation of a new cluster. These differences in the number of clusters created can lead to significant mismodelling in reconstructed hadronic objects, particularly in the description of jet substructure.

In order to improve the modelling of the lateral shape in the simulation of hadrons, the number of hits sampled in each ΔR^{mm} bin is modified to better reproduce the RMS of the $E_{\text{voxel}}/E_{\text{hit}}$ distributions. The energy that is assigned to the hits is scaled with ΔR^{mm} dependent weights ω :

$$\widetilde{E}_{\text{hit}} = \frac{E_i}{N_{\text{hits}}} \cdot \omega. \quad (4.10)$$

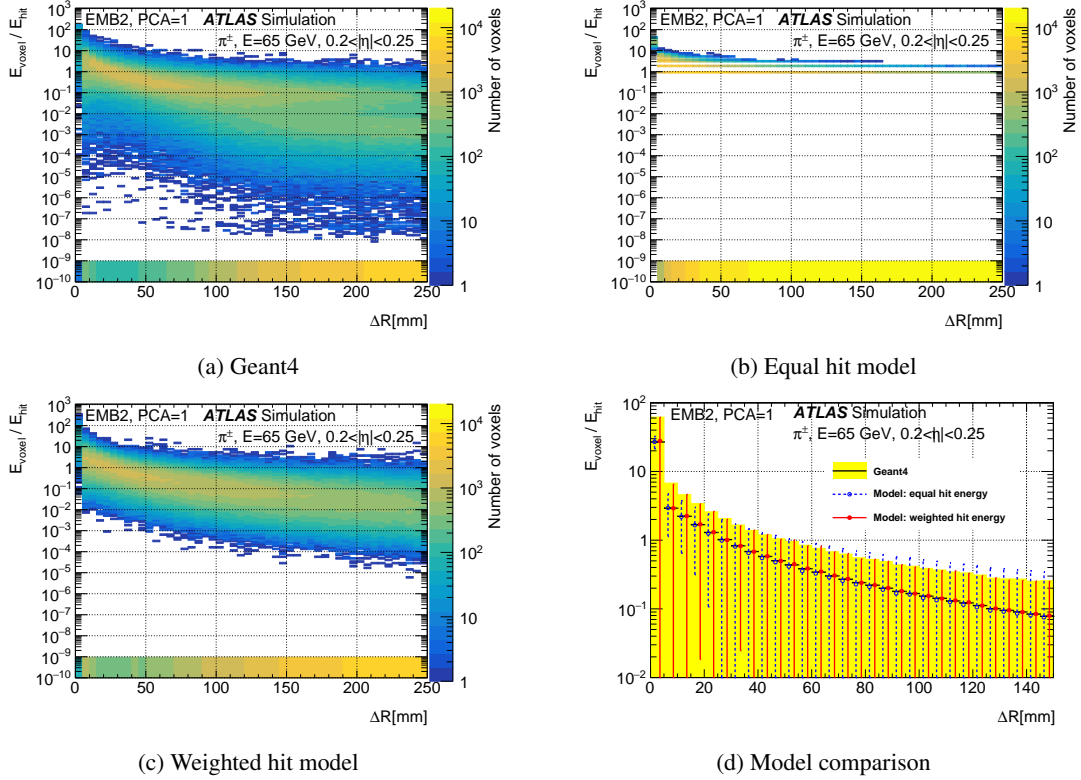


Figure 4.13: Two-dimensional distribution of voxels in dependence of the ratio between voxel energy and hit energy, and the radial distance ΔR^{mm} from the shower centre as simulated with 65 GeV charged pions in EMB2 in the range $0.20 < |\eta| < 0.25$ in the first bin of the leading PCA (PCA=1). Entries with zero voxel energy are shown in the underflow bin. The distributions are shown for GEANT4 (a), for a model using equal distributed hit energies (b) and in a model using weighted hit energies (c). (d) Comparison of the mean (central value) and the RMS (error bars) between the GEANT4 simulation and the equal and weighted hit models. The yellow band indicates the 1σ uncertainty for GEANT4. [134]

A gaussian smearing of the hit energy is additionally applied in some cases in order to improve the modelling of the lateral shape further. A more in-depth description on the derivation of the weights is provided in Ref. [134]. Figure 4.13(c) shows the voxel distribution for the weighted hit model, which shows significantly better agreement with GEANT4. The weighted hit model allows the creation of voxels with energies well below the GEANT4 hit energy and effectively prevents the creation of high energy and potentially cluster seeding hits in regions far away from the shower center. Figure 4.13(d) shows the comparison of the mean and RMS of the $E_{\text{voxel}}/E_{\text{hit}}$ distributions in each of the ΔR^{mm} bins for the equal and weighted hit energy model. The means of the distributions are reproduced correctly in both cases, but the RMS is significantly better modelled in the model with weighted hits.

4.5 Energy and Shape Interpolation

The energy parametrization of FASTCALOSIMV2 is based on discrete, logarithmically spaced truth momentum points. Particles entering the calorimeter volume, however, can have very different energies that do not necessarily exactly match one of the energy points used to derive the parametrization.

In order to interpolate the longitudinal and lateral shower shape properties between the logarithmically spaced energy grid, a parametrization is randomly selected based on the energy difference between the true kinetic energy of the particle and the kinetic energy of the neighbouring parametrization points. In that regard, the probability of choosing a parametrization linearly decreases in $\log(E_{\text{kin}})$ with larger differences in kinetic energy to the parametrization point.

Without additional corrections, this procedure would introduce a bias in the simulated energy of individual showers, as the energy depositions in the calorimeter produced by particles of neighbouring parametrizations can vary significantly. To take this into account, the energy response

$$\mathcal{R}(E_{\text{kin}}^{\text{truth}}) = \frac{\langle E_{\text{dep}} \rangle}{E_{\text{kin}}^{\text{truth}}}, \quad (4.11)$$

of the GEANT4 simulated showers is interpolated using a piecewise cubic spline. Here, $E_{\text{kin}}^{\text{truth}}$ is the true kinetic energy of the incident particles and E_{dep} the total deposited energy in the calorimeter. The interpolations are performed for each particle type and η slice independently and are used to re-scale the total energy response from the parametrization points that is later distributed across the layers. Figure 4.14 shows two example spline interpolations for photons in the Barrel and Endcap regions. The energy response in the barrel decreases for very energetic photons as a result of leakage into the Tile layers. The energy response beyond the energy parametrization range, that is, for very low and high energetic particles, is computed using a linear extrapolation using the two adjacent parametrization points.

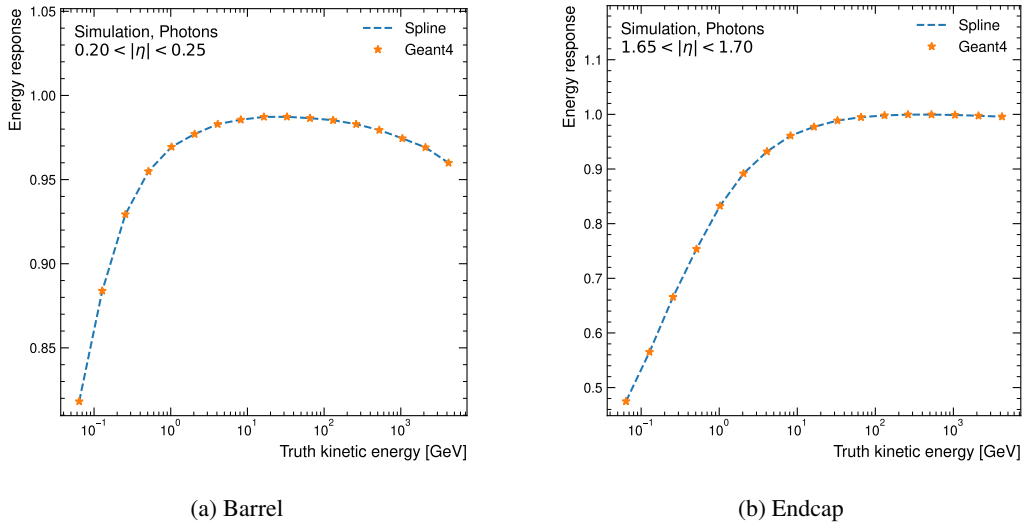


Figure 4.14: Energy response of the FASTCALOSIMV2 parametrization points and cubic spline interpolation used to re-scale the total energy response from the parametrization points for photons in the region (a) $0.20 < |\eta| < 0.25$ and (b) $1.65 < |\eta| < 1.70$.

4.6 Computing Performance

In realistic conditions and with respect to GEANT4, ATLFAST3 achieves a speed-up factor of 5 to 6. Figure 4.15 shows the simulation time of 250 $t\bar{t}$ events, a complex physics process that is commonly used as benchmark scenario, for a GEANT4 and ATLFAST3 simulation using Run 2 data taking conditions. The alternative simulator AF3F, which uses FATRAS for fast track simulation in the Inner Detector is also shown for reference. The benchmarks are performed on a 64-core AMD EPYCTM 7702 CPU at 2 GHz base clock. On average, a $t\bar{t}$ event is simulated in 113.8 ± 2.1 seconds with GEANT4, while ATLFAST3 only takes 19.1 ± 0.5 s, giving approximately a factor of 6 increase in simulation speed. The mean simulation time for ATLFAST3 when using FATRAS as fast track simulation (AF3F) is another order of magnitude smaller with a mean simulation speed of 0.921 ± 0.024 s per event, demonstrating that the simulation time of ATLFAST3 is fully dominated by the simulation of particle tracks in the ID with GEANT4.

If one only considers the simulation of the calorimeter, the difference in simulation speed between ATLFAST3 and GEANT4 is consequently significantly larger. Figure 4.16(a) shows the simulation time of EM showers in the calorimeters induced by single photons with varying energies that have been sampled from a combination of Fermi-Dirac distributions. Note that the axes showing the simulation time is discontinued, and both axes span very different orders of magnitude. On average, GEANT4 takes 122 ± 6 s to simulate these inclusive photon events, while ATLFAST3 just takes 0.04738 ± 0.00023 s amounting to an overall speed up in the order of $O(2500)$. However, these very large differences in simulation times are mostly driven from the




Figure 4.15: Simulation time per event of 250 $t\bar{t}$ events with GEANT4, ATLFAST3 and ATLFAST3F, which uses FATRAS for the simulation of the Inner Detector using Run 2 data taking conditions. The grey dashed lines indicate the means of the distributions. Benchmarks are performed on a 64-core AMD EPYC™ 7702 CPU at 2 GHz base clock.

simulation of very energetic photons. As Figure 4.16(b) demonstrates, the simulation time of GEANT4 scales exponentially with the particle's incident energy. This is not surprising, as the number of particles in the electromagnetic shower, whose passage through matter GEANT4 has to individually simulate, scales accordingly. In contrast, the simulation time of FASTCALOSIMV2, as a parametrized simulation, is independent of the particle's energy and as such, gains in the simulation speed are especially apparent in the high energy regime. For a 65 GeV photon, for instance, the simulation with FASTCALOSIMV2 is faster by a factor of approximately 1800, while for slightly higher energetic photons with 130 GeV, the simulation is faster by more than a factor of 3000. Note, however, that these numbers are averaged over all directions in the calorimeter and that simulation times might vary considerably depending on the incoming direction of the particle.

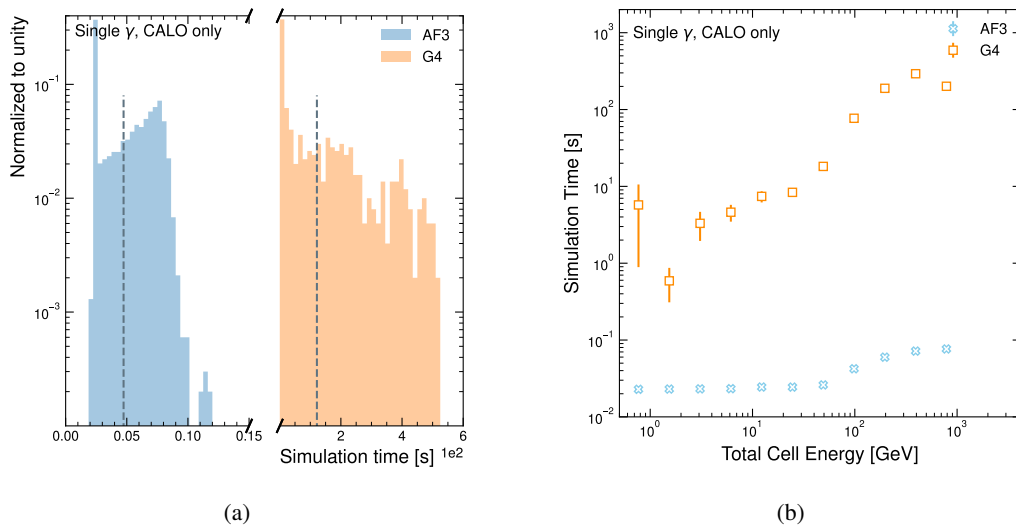


Figure 4.16: (a) Simulation time of EM showers induced by single photons as simulated with ATLFast3 and GEANT4. The grey dashed lines indicate the means of the distributions. Note that the two axis are discontinued and span over different orders of magnitude. (b) Simulation time of EM showers induced by single photons as simulated with ATLFast3 and GEANT4 in dependence of the total cell energy deposited in the calorimeter, which is proportional to the energy of the incident photon. In both cases, single photons are generated at the calorimeter surface with energies sampled from a Fermi-Dirac distribution and simulated using Run 2 data taking conditions. Benchmarks are performed on a 64-core AMD EPYCTM 7702 CPU at 2 GHz base clock.

Energy Response Corrections

WITH substantially increasing luminosity in the coming years, precision measurements will significantly gain in importance in the near future. In particular, the further understanding of electroweak symmetry breaking through precise measurements of the Higgs bosons properties remains a major goal of the ATLAS physics programme and may open the window to new physics. Furthermore, recent measurements of the W boson mass with the CDF II detector have revealed significant tensions to the Standard Model predictions [42], providing additional motivation for precise mass measurements for indirect BSM searches.

In order to enable these highly precise measurements, a very accurate simulation of electromagnetic and hadronic showers is of utmost importance. While the simulation of the total energy response of particle showers with ATLFast3 is very accurate, small differences with respect to GEANT4 are observed, which can affect the modelling of high-level observables and may limit the usage of ATLFast3 for precision measurements.

Differences in the simulated total energy response between FASTCALOSIMV2 and GEANT4 can originate from two distinct effects. On one hand, imperfections in the energy parametrization procedure can lead to small discrepancies in the total energy response, which can be further amplified by the subsequent reconstruction and digitization of the simulated energy deposits. On the other hand, the simulation of all hadronic showers, regardless of the type of the shower-initiating particle, is performed using the parametrization derived from pions. This approximation results in inherent differences in the energy response simulation for hadrons other than pions.

In this chapter, correction methods which are developed to account for these differences and improve the total energy response simulation of EM and hadronic showers are presented. In Section 5.1 the correction of the energy response of neutrons, protons, kaons and their corresponding antiparticles is discussed. Finally, corrections to account for imperfections in

the energy parametrization and discrepancies induced by the reconstruction of cell energies are presented in Section 5.2.

5.1 Energy Response of Hadrons

The energy response of hadrons simulated by FASTCALOSIMV2 is based on the charged pion parametrization. In other words, hadrons other than pions that reach the boundary between Inner Detector and calorimeter system are intercepted by FASTCALOSIMV2 and simulated as an equal-charge mixture of pions. In principle, separate parametrizations for each respective hadron type could be used during simulation. However, as the full parametrization is required to be loaded into memory, the addition of new parametrizations would result in a drastically higher memory usage to an extent that FASTCALOSIMV2 could not be employed on many machines.

For many physics searches and measurements, this approximation is sufficient for two reasons: the properties of charged pion-initiated showers are reasonably similar to showers initiated by other hadrons, and as such the energy response does not differ significantly. Secondly, pions are the predominant hadrons created through hadronization in a typical jet which will consequently need to be simulated. Hence, a mismodelling in the energy response simulation of other hadrons will have a significantly smaller impact on the simulation performance of reconstructed objects. Figure 5.1 shows the averaged fraction of predominant particles that reach the boundary between Inner Detector and calorimeter system in a typical $t\bar{t}$ event. Neutral and charged pions are created in approximately equal fractions and constitute $\sim 60\%$ of the total hadrons that are simulated with FASTCALOSIMV2. This is followed by neutrons and protons with $\sim 15\%$ each. The charged K^\pm , the long-lived neutral K_L as well as antiprotons and antineutrons each only constitute less than 2% of the total hadrons that reach the calorimeter system.

For the purpose of precision measurements, the correction of the total energy response of hadrons other than pions, especially of the predominant protons and neutrons, becomes increasingly more important. Figure 5.2 shows the total energy response of protons, neutrons, kaons and their corresponding antiparticles as simulated with GEANT4 for particles with incident truth momentum of 65 GeV in the $0.20 < |\eta| < 0.25$ region. As the particles which are used for the parametrization are generated in truth momentum instead of kinetic energy, intrinsic differences in the total energy response between pions and other hadrons are expected as a result of their differing rest masses within one energy slice. More precisely, the expected difference in the energy response as a result of the differing rest masses in a given $p - |\eta|$ slice is given by the ratio of the particle's true kinetic energies:

$$\frac{E_{\text{kin,true}}^h}{E_{\text{kin,true}}^\pi} = \frac{m_h \cdot (m_h - 1) + p^2}{m_\pi \cdot (m_\pi - 1) + p^2}, \quad (5.1)$$

where h indicates hadrons other than pions. For instance, for the 65 GeV particles shown,




Figure 5.1: Averaged fraction of predominant particles that reach the boundary between Inner Detector and calorimeter system in a typical $t\bar{t}$ event normalized to the most prevalent hadrons (blue) and to the total number of particles reaching the boundary, including photons and electrons (orange).

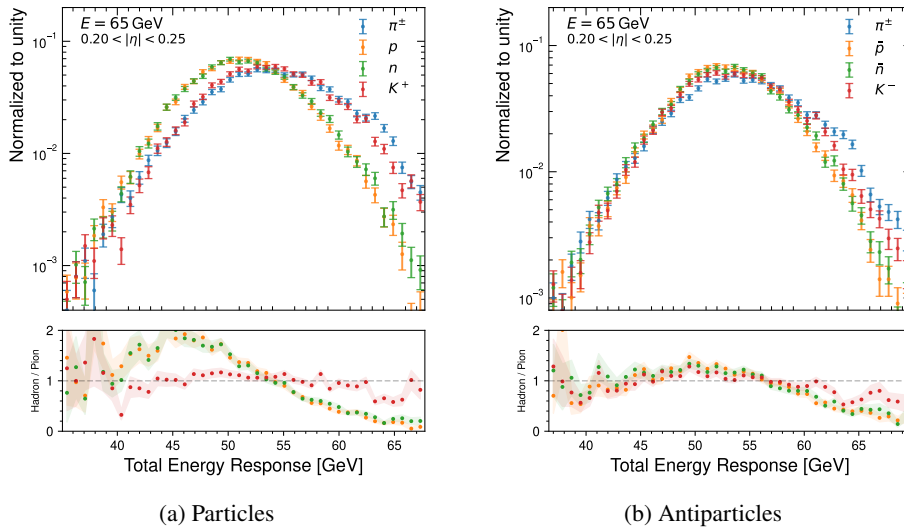


Figure 5.2: Total energy response of (a) protons, neutrons, kaons and (b) their corresponding antiparticles as simulated with GEANT4 for particles generated at the boundary between Inner Detector and calorimeter system with incident truth momentum of 65 GeV in the $0.20 < |\eta| < 0.25$ region. The lower panel shows the energy response ratio with respect to the pion response.

a shift of approximately 790 MeV is expected for neutrons and protons, and a shift of 350 MeV for kaons, due to their different rest masses. For antineutrons and antiprotons, the shift with respect to the pion response expected from their rest masses is as for neutrons and pions, respectively. However, antiprotons and antineutrons that reach the boundary between Inner Detector and calorimeter are annihilated, which will result in an additional energy of $2 \cdot m_{p/n}$ deposited in the calorimeter. During simulation with FASTCALOSIMV2, both effects are taken into account, that is, other hadrons are simulated using the pion parametrization, but the total energy response is chosen according to the true kinetic energy of the incident particle from the energy interpolation discussed in Section 4.5. To account for the additional energy deposited by the antiprotons and antineutrons as a result of the annihilation process, their kinetic energy used to evaluate the total energy response is calculated by adding $2 \cdot m_p$ or $2 \cdot m_n$, respectively.

FASTCALOSIMV2 can easily take into account overall variations in the total energy response originating from kinematic differences. However, more subtle differences between the energy response of pion and other hadron initiating showers resulting from slightly different hadronic interactions, which may depend on the incident particle energy and direction, cannot be taken into account from first principle computations. Figure 5.3 shows the hadron-to-pion energy response in terms of the true kinetic energy ratio $E_{\text{kin,true}}^h/E_{\text{kin,true}}^\pi$ and the observed mean total energy response ratio $\langle E_h^{G4} \rangle / \langle E_\pi^{G4} \rangle$ in dependence of the truth momentum of the incident particle for protons, neutrons, kaons and their corresponding antiparticles in the $0.20 < |\eta| < 0.25$ region. For large truth momenta, the differences in the rest masses become

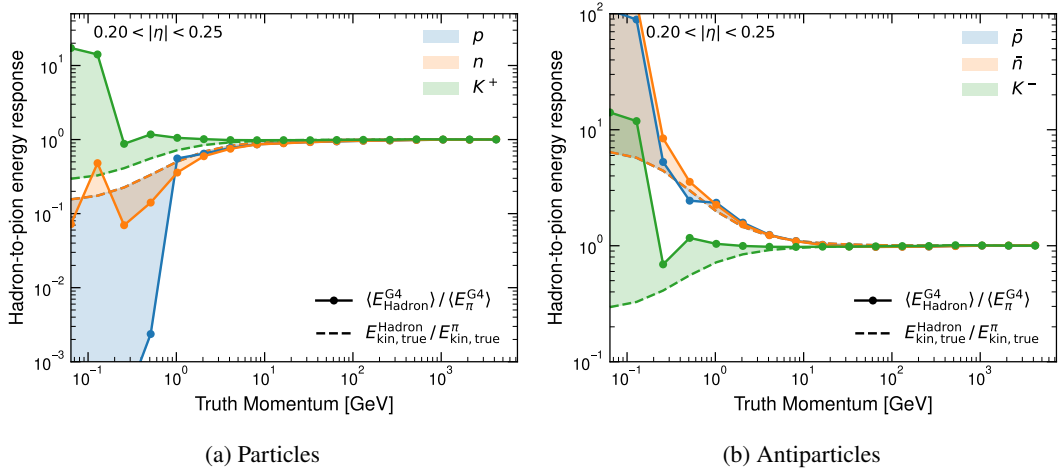


Figure 5.3: True kinetic energy ratio between hadrons and pions (dashed line) and observed energy response ratio for single particles generated at the boundary between Inner Detector and calorimeter system, as simulated with GEANT4 in the $0.20 < |\eta| < 0.25$ region (solid line) in dependence of the truth momentum of the incident particle. The ratios are shown for (a) protons, neutrons, kaons and (b) their corresponding antiparticles. The shaded areas show the discrepancies between both ratios and indicate variations in the energy response which the nominal FASTCALOSIMV2 simulation cannot encode.

negligible and the ratio of kinetic energies converges to unity in all cases. Similarly, the ratio between the observed mean energy response of pions and other hadrons can be approximated by one for large incident particle momentum, that is, differences in the energy response from varying hadronic interactions are more predominant in the low energy region. The shaded areas show the difference between the two ratios and indicate variations in the energy response that FASTCALOSIMV2 cannot take into account without additional corrections.

In order to correct for these variations, $p - |\eta|$ dependent correction factors are derived for the full parametrization grid that is used for FASTCALOSIMV2. The correction factors are derived from GEANT4 simulations of single hadrons produced at the boundary between Inner Detector and calorimeter surface. In order to match the pion simulation used to generate the parametrization, no vertex smearing is applied, and energy deposits are digitized without the inclusion of electronic noise and cross-talk between calorimeter cells in the readout electronics. The minimum number of events generated corresponds to the number of events used for the parametrization of hadronic showers and is shown in Table 4.2. As FASTCALOSIMV2 intrinsically takes into account differences in the total energy response resulting from kinematic considerations, the mean energy response ratio $\langle E_h^{G4} \rangle / \langle E_\pi^{G4} \rangle$ is scaled with the inverse kinetic energy ratio. Thus, the mean energy response of hadrons is scaled according to $\langle \widetilde{E}_h \rangle = \lambda \cdot \langle E_h \rangle$, where

$$\lambda = \frac{\langle E_h^{G4} \rangle}{\langle E_\pi^{G4} \rangle} \times \frac{E_{\text{kin,true}}^\pi}{E_{\text{kin,true}}^h} \quad (5.2)$$

is the corresponding correction factor that is derived for each $p - |\eta|$ point. Figure 5.4(a) shows the ratio between true kinetic energy of pions and other hadrons and Figure 5.4(b) the observed energy response ratio in the $0.20 < |\eta| < 0.25$ region for protons, neutrons, kaons and their antiparticles in dependence of the truth momentum of the incident particle. The product between both quantities yields the final correction factors that are used to scale the total energy response of hadrons other than pions in order to remove the shift in the total energy response observed between the FASTCALOSIMV2 and GEANT4 simulation of hadrons.

The resulting correction factors in dependence of the true kinetic energy of the incident particle for three $|\eta|$ bins corresponding to the barrel, endcap and FCal regions are shown in Figure 5.5. The corrections are piecewise linearly interpolated between the parametrization points. An interpolation using cubic splines as employed for the general interpolation of the energy response was investigated, but was found to generate undesired oscillations in the low energy regime, such that this option is disfavoured. In general, the corrections are very similar for particles and antiparticles. For kaons with kinetic energies below 10-20 GeV, the correction factors are in most cases significantly larger than one, while for higher energetic particles the corrections generally fall slightly below unity and remain mostly constant in a 2 – 3% window. The corrections for protons and neutrons are fairly similar and are typically in a 5 – 10% correction range for particles with kinetic energies around 10 GeV and show an approximately linear increase until reaching unity for particles in the TeV range.

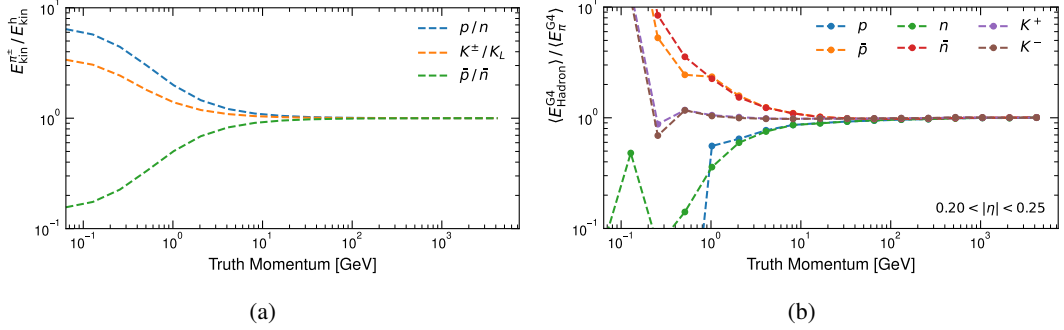


Figure 5.4: (a) The ratio between true kinetic energy of pions and other hadrons as a function of true momentum of protons, kaons, neutrons and the corresponding antiparticles. Hadrons with nearly identical mass are grouped together, as their ratio is approximately equal. (b) Observed energy response ratio between hadrons and pions generated at the boundary between Inner Detector and calorimeter region, and simulated with GEANT4 in the $0.20 < |\eta| < 0.25$ region. The product between both quantities constitutes the final correction factors.

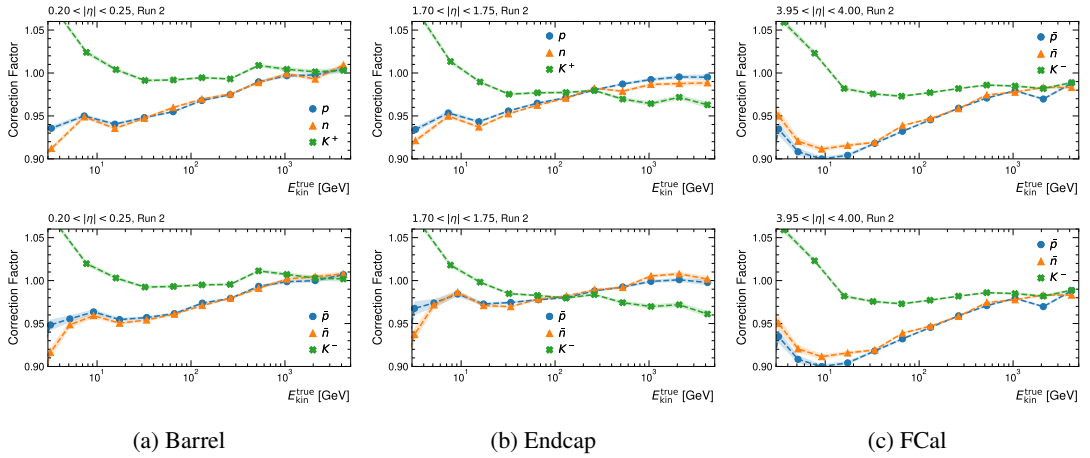


Figure 5.5: Example energy response correction factors as a function of true kinetic energy of the incident hadron for protons, neutrons and kaons (top) and their corresponding antiparticles (bottom) in the (a) Barrel, (b) Endcap and (c) FCal regions. The correction factors are piecewise linearly interpolated, and the bands show the associated uncertainties to the correction factors.

The full set of colour encoded energy response correction factors in the two-dimensional $p - |\eta|$ parametrization space is shown in Figure 5.6. White colours indicate correction factors close to unity, while blue indicates values for which the corrected energy response is shifted towards lower values ($\lambda < 1$) and red values for which the corrected energy response is shifted towards larger values ($\lambda > 1$). Note that colours at the very edge of the displayed colour bar do not necessarily indicate the corresponding correction factor, but values outside the displayed range.

It becomes apparent, that the energy response of kaons and antikaons for all $|\eta|$ ranges is corrected towards larger energy responses ($\lambda > 1.1$) up to truth momenta between 16 and 32 GeV. For higher energies, the energy response is corrected towards the opposite direction. In the case of protons, the correction factors are generally below unity, except for the very high $|\eta|$ region, that is, approximating the out-of-acceptance regimen with $|\eta| > 4.5$, where the total energy response is scaled upwards for medium to high-energy particles. For antiprotons, neutrons and antineutrons, similar effects are visible with the difference that the $\lambda > 1$ corrections close to the out-of-acceptance region extend to the low energy regime. Furthermore, for low energetic particles, except for large parts of the FCal region, the correction factors are well below unity.

In order to validate the set of corrections, the corrected simulation is compared with the nominal FASTCALOSIMV2 and GEANT4 simulation for the specific parametrization points. As an example, Figure 5.7 shows the total energy response of the three simulators for 65 GeV hadrons in the $0.20 < |\eta| < 0.25$ region for protons, neutrons, kaons and their antiparticles and denotes the shift of the mean total energy response observed with respect to the GEANT4 simulation before and after the application of the correction. As expected, the energy shift is compatible with zero after the application of the correction. For protons, for instance, a shift of $|\Delta E| \sim 2.36$ GeV is observed in the total energy response with respect to GEANT4, which is fully removed after application of the correction. For all hadrons, the corrected shift is approximately 2 GeV with exception of kaons, where the corrected shift is significantly smaller and ranges between 0.4 and 0.2 GeV. This is expected, as the correction factors for kaons are significantly closer to unity in the shown energy regime. With regard to the RMS of the distributions, minor differences remain. However, these higher order differences play a minor role and require significantly more sophisticated methods to correct.

The samples employed to derive the hadron corrections presented in this chapter were produced using the same configuration as for the GEANT4 reference samples used for the parametrization of FASTCALOSIMV2, as specified in Chapter 4. Among other aspects, the GEANT4 version 10.1.3 [149] and the FTFP_BERT_ATL [164] physics list was used. This was the default configuration for generating all full simulation samples in ATLAS during Run 2. For Run 3, ATLAS will employ GEANT4 10.6, which introduces significant changes in the simulation of hadronic physics. Small changes in the correction factors are expected, such that the correction factors have been re-derived for the Run 3 configuration of FASTCALOSIMV2. In addition, correction factors have also been derived for long-lived kaons K_L , which are among the most predominant hadrons that reach the calorimeter surface (see Figure 5.1). The

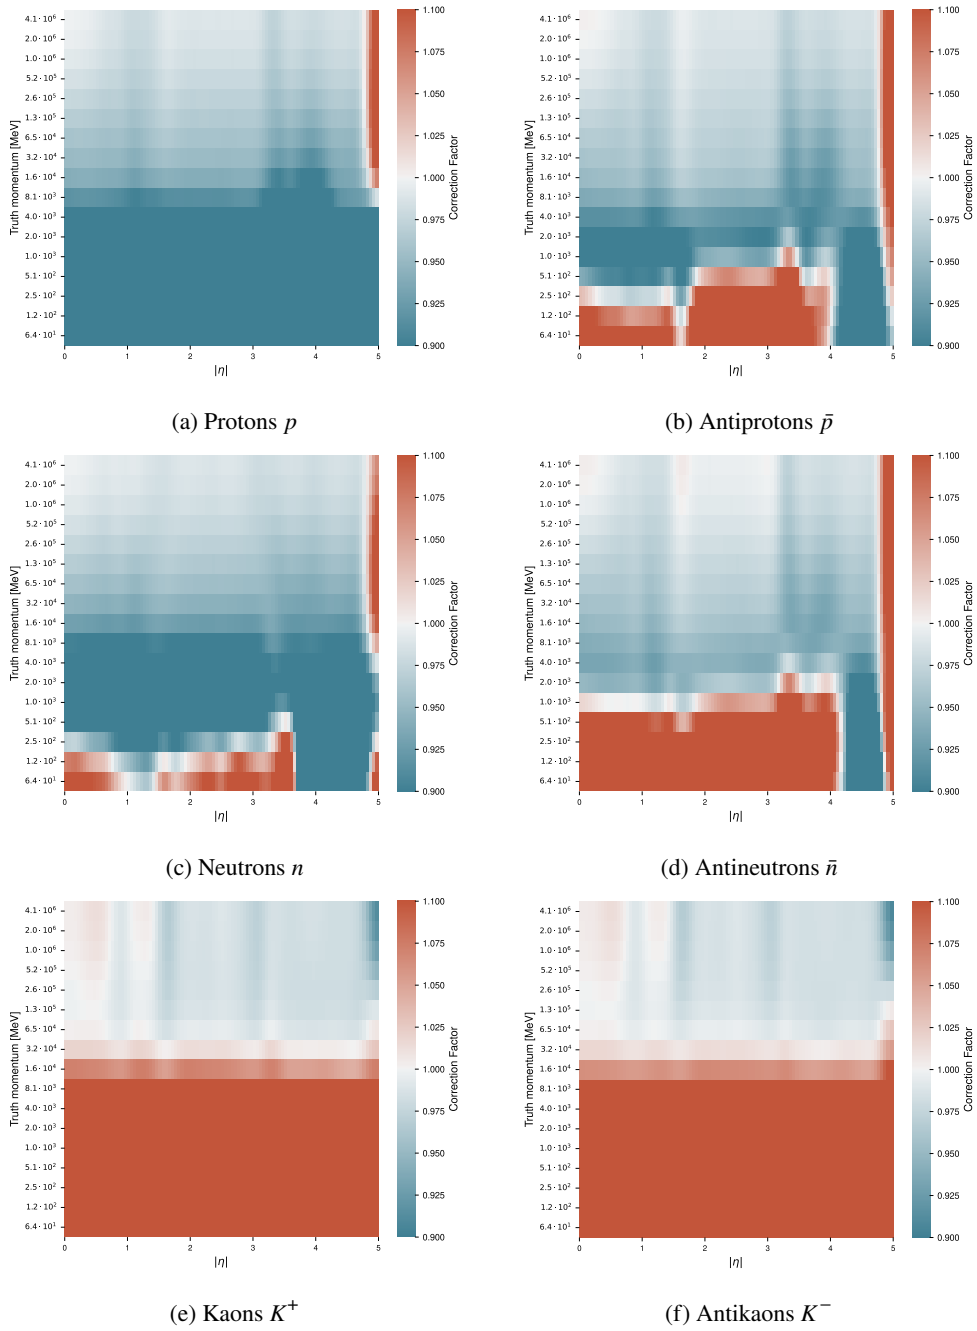


Figure 5.6: Energy response correction factors for proton, neutrons and kaons for the full FASTCALOSIMV2 parametrization grid, that is, in dependence of $|\eta|$ and the truth momentum of the incident particle. A Gaussian filter is applied to smooth the transition between bordering bins. The colour range of the correction factors is chosen such that most of the fluctuations of the correction factors can be resolved. Note that colours at the very edge of the displayed colour bar do not necessarily indicate the corresponding correction factor, but values outside the displayed range.

5.1 Energy Response of Hadrons

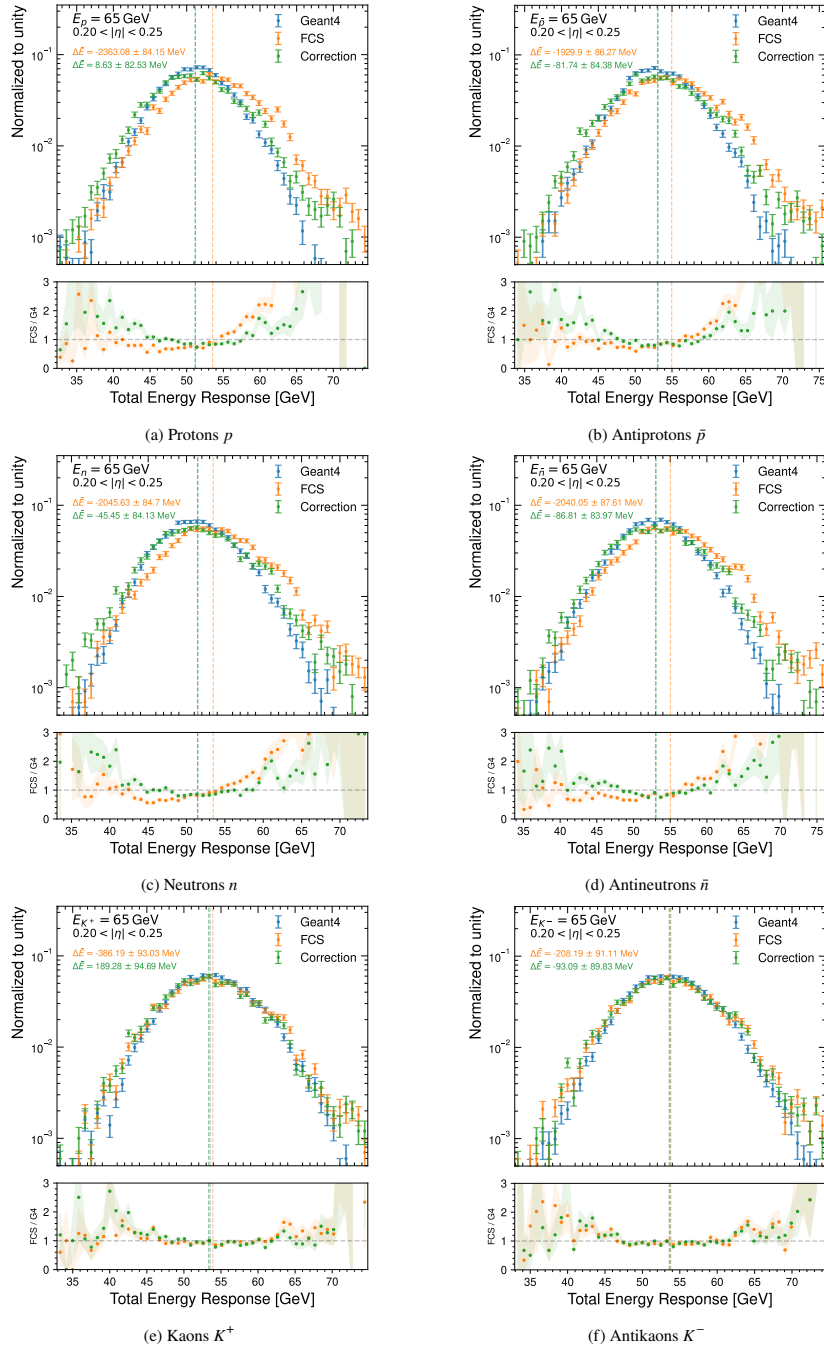


Figure 5.7: Total energy response of 65 GeV protons, neutrons, kaons and their corresponding antiparticles, as simulated with `GEANT4` as well as with the nominal and corrected `FASTCALOSIMV2` parametrization in the $0.20 < |\eta| < 0.25$ region. The dashed vertical lines indicate the mean of the distributions and the lower pad shows the ratio between the `FASTCALOSIMV2` and `GEANT4` simulation. The observed energy shifts with respect to `GEANT4` are annotated for the nominal (orange) and corrected (green) `FASTCALOSIMV2` parametrization.

corresponding correction factors for Run 3 are given in Appendix A for all detector regions.

Given that the energy response of hadrons in a specific $p - |\eta|$ slice simulated by FASTCALOSIMV2 corresponds to the pion response scaled to the kinetic energy of the respective hadron, the corrected mean energy response can be expressed as:

$$\begin{aligned}
 \overline{\langle E_h \rangle} &= \lambda \cdot \langle E_h \rangle \\
 &= \lambda \cdot \langle E_\pi \rangle \cdot \frac{E_{\text{kin,true}}^h}{E_{\text{kin,true}}^\pi} \\
 &= \underbrace{\frac{\langle E_\pi \rangle}{\langle E_\pi^{\text{G4}} \rangle}}_{:= \xi_{\text{off}}} \cdot \langle E_h^{\text{G4}} \rangle = \xi_{\text{off}} \cdot \langle E_h^{\text{G4}} \rangle
 \end{aligned} \tag{5.3}$$

Therefore, the shift in the mean hadron energy response between the GEANT4 and the corrected FASTCALOSIMV2 simulation is given by

$$\begin{aligned}
 \Delta E &= \langle E_h^{\text{G4}} \rangle - \overline{\langle E_h \rangle} \\
 &= \langle E_h^{\text{G4}} \rangle \cdot (1 - \xi_{\text{off}}) \xrightarrow{\xi_{\text{off}} \rightarrow 1} 0
 \end{aligned} \tag{5.4}$$

Hence, by construction, the correction of the total energy response for hadrons other than pions can only be fully effective if $\xi_{\text{off}} = 1$, that is, no differences in the mean energy response between the FASTCALOSIMV2 and GEANT4 simulation of the pion energy response exists. In order to make sure that this is indeed the case, the pion response is corrected using dedicated scaling factors which are described in Section 5.2. These correction factors are applied to the total energy response simulations shown in Figure 5.7, and thus guarantee that the energy shifts with respect to GEANT4 of hadrons other than pions are fully removed.

5.2 Reconstructed Energy Response

Small discrepancies in the simulation of the energy response are also observed for particles, which are directly used to parametrize the simulation of EM and hadronic showers in FASTCALOSIMV2. These differences can originate from an imperfect parametrization and from discrepancies introduced during the reconstruction and digitization of the deposited energies. In order to correct for these differences, a similar procedure as for the simulation of hadrons other than pions is employed. The mean energy response of electrons, photons and pions is scaled using the ratio between the mean total energy response as simulated with GEANT4 and the uncorrected

FASTCALOSIMV2 total energy, that is, the corrected mean energy response is given by

$$\overline{\langle E_p \rangle} = \frac{\langle E_p^{G4} \rangle}{\langle E_p \rangle} \cdot \langle E_p \rangle \quad p \in \{e, \gamma, \pi^\pm\} \quad (5.5)$$

The correction factors are calculated for all $p - |\eta|$ parametrization points and after the full simulation and reconstruction chain in order to match GEANT4 at reconstruction level. The FASTCALOSIMV2 response is computed by simulating corresponding single particle events with the same number of events generated as for the samples employed for the parametrization. No vertex smearing is applied, and energy deposits are digitized without the inclusion of electronic noise and cross-talk between calorimeter cells in the readout electronics.

Figure 5.8 shows example energy response correction factors for photons, electrons and pions as a function of the true kinetic energy of the incident particle in the barrel and endcap regions. As for the correction of hadrons other than pions, the correction factors are piecewise linearly interpolated between energy parametrization points. In most cases, the derived correction

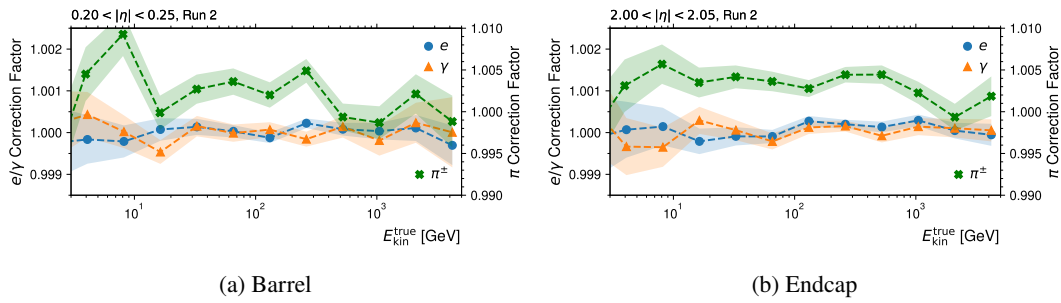


Figure 5.8: Example energy response correction factors for photons, electrons and pions as a function of true kinetic energy of the incident particle in the (a) barrel and the (b) endcap. The correction factors are piecewise linearly interpolated, and the coloured bands indicate the associated statistical uncertainties. Note the two different axes for the e/γ and π^\pm correction factors.

factors for photons and electrons are compatible with unity within statistical uncertainties. In the case of pions, the correction factors are generally larger than one and fluctuate within a 5% window.

The full set of colour encoded correction factors for the $p - |\eta|$ parametrization grid are shown in Figure 5.9. In the case of pions, the corrections factors are generally above unity for particles with truth momentum above a few GeV. Below this energy threshold, the correction factors tend to fall below one. Similarly, the correction factors for electrons fall mostly below unity for low energy particles, except for a few well-isolated $|\eta|$ regions. For photons, the picture is reversed and most correction factors in the low energy regime fall above one. In both cases and as a result of their correction factors being significant closer to one, the overall fluctuations around unity are larger for EM showers compared to hadronic showers in the case

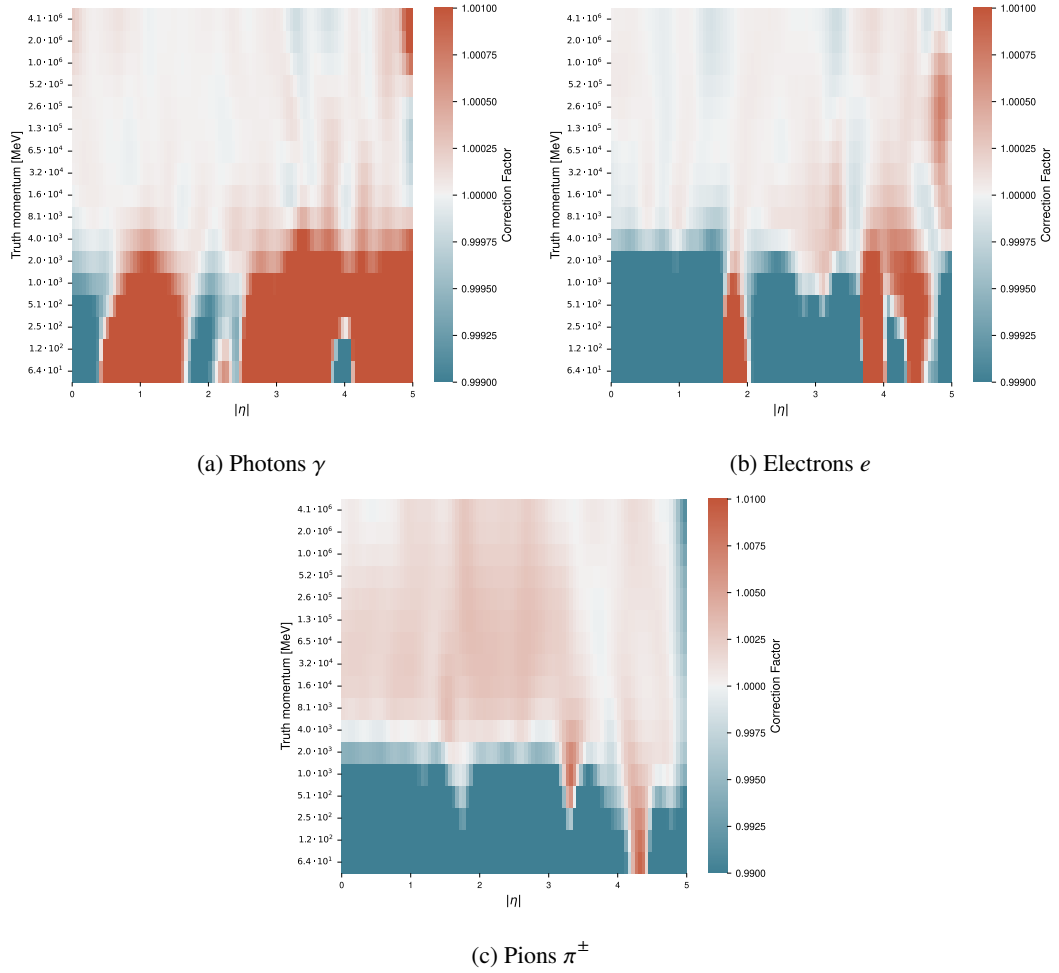


Figure 5.9: Energy response correction factors for (a) photons, (b) electrons and (c) pions for the full FASTCALOSIMV2 parametrization grid, that is, in dependence of $|\eta|$ and the truth momentum of the incident particle. A Gaussian filter is applied to smooth the transition between bordering bins. The colour range of the correction factors is chosen such that most of the fluctuations of the correction factors can be resolved. Note that colours at the very edge of the displayed colour bar do not necessarily indicate the corresponding correction factor, but values outside the displayed range.

of high energetic incident particles.

As the correction factors are generally very small and often compatible with unity, especially for electrons and photons, the correction is only applied in the case that it is statistically significant. More specifically, the correction is applied only if $\sigma_\lambda < |1 - \lambda|$, where λ denotes the correction factor and σ_λ the associated statistical uncertainty.

The effects of the reconstructed energy corrections are typically not visible in corresponding energy distributions. The correction factors are generally derived by using a random seed in the FASTCALOSIMV2 simulation in order to avoid any potential bias. However, for validation purposes, the corrections are derived with a fixed random seed. The resulting corrections are then applied to the simulation, using the identical seed and the correction factors are re-derived in order to make sure that all differences in the energy response are effectively removed. Figure 5.10 shows example correction factors that are derived before and after the application of the correction for photons in the $0.20 < |\eta| < 0.25$ region. As expected, the correction factors derived after application of the correction are all at unity.




Figure 5.10: Energy response correction factor in dependence of true kinetic energy of photons in the $0.20 < |\eta| < 0.25$ region before (blue) and after (orange) applying the energy correction. In both cases, the simulation is performed with the same random seed in order to avoid any fluctuations in the simulated energy response. The shaded bands indicate the size of the associated uncertainties.

While the derived corrections of the reconstructed energy response are generally small, especially for electrons and photons in the middle to high energy range, the compatibility of higher level observables to GEANT4, such as the mean of the simulated invariant mass of $Z \rightarrow ee$ and $H \rightarrow \gamma\gamma$ events is found to be improved with respect to the predecessor of FASTCALOSIMV2 (see Chapter 7).

Shower Centre Determination

THE lateral energy parametrization in `ATLFAST3` describes the energy density and is parametrized depending on the position of the shower centre in each layer (see Section 4.3). During simulation, the distribution of energy deposits within each layer will therefore strongly depend on the centre position of the shower.

In order to determine the shower centre, the current implementation of `ATLFAST3` relies on tools employed for the reconstruction of particle tracks in the ATLAS detector. More specifically, particles that reach the calorimeter envelope are transported through the calorimeter based on their incident momentum, while taking into account the magnetic field which they are subject to. The transportation yields a series of intersections with active calorimeter surfaces that the particle is predicted to create within the calorimeter system, if it would not initiate a showering process.

The hits created by the ATLAS track reconstruction tools within the calorimeter system and along the particle trajectory are not sufficient to fully determine the positions to which energy deposits simulated by `ATLFAST3` need to be assigned. On one hand, the particle transport from the calorimeter envelope does not necessarily lead to direct hits on the boundary between calorimeter system and ID, as used to generate the single particles that are employed to derive the parametrization (see Figure 4.3). As such, the particle position is required to be extrapolated to the parametrization boundary based on the hits generated by the tracking algorithm. The result of the extrapolation will hence determine the choice of the parametrization for all incident particles and is therefore a crucial component of `ATLFAST3`. Moreover, shower-initiating particles can deposit energies in calorimeter layers through which no direct transport of the particle track is performed. In order to ensure that energies can be assigned correctly to these calorimeter regions, sensible shower centres need to be determined for all layers of the calorimeter based on the information available from the track transport.

In the predecessor of ATLFast3, a basic linear extrapolation algorithm was employed to determine the shower centres which, however, is observed to fail to assign the correct shower centre position in a significant fraction of cases, especially for very low energetic tracks which exhibit significant curvatures when being transported through the calorimeter system.

In this chapter, a fully redesigned algorithm is presented that is used to reliably determine the shower centres in each calorimeter layer, even in the case of low energetic particles. In Section 6.1, the transport of track parameters across the calorimeter as used in the ATLAS track reconstruction is briefly outlined. The novel extrapolation algorithm is described in Section 6.2 and an overview of the strategies employed to validate the new shower centre computation is given in Section 6.3.

6.1 ATLAS Track Extrapolation

Track reconstruction is a crucial component of the ATLAS reconstruction chain and aims to estimate the trajectories of charged particles throughout the detector based on a set of measurements. The resulting tracks are, in many cases, the input to subsequent reconstruction algorithms that are used to build higher-level physics objects such as jets and τ -leptons.

The reconstruction of tracks in ATLAS is essentially composed of three main steps. First, the raw measurements generated by the pixel and SCT detectors are clustered by a connected component analysis [169], which groups relevant pixels and strips within a sensor. These clusters are then used to obtain three-dimensional measurements, in the following referred to as hits, which correspond to the points where a particle traverses the active material of the ID. A seeding algorithm [170] then attempts to find triplets of hits that are likely to belong to the same track. The triplets are iteratively filtered by implying certain conditions on the momentum and impact parameters of the seed triplets which are estimated under the assumption of a perfect helical track in a homogenous magnetic field. Finally, track candidates are reconstructed using a combinatorial Kalman filter [171]. Within the Kalman formalism, track parameters are estimated at each intersection of the track with active layers of the detector, and rely on the measurement as well as the prediction of the track parameters at each active surface. Figure 6.1 illustrates a typical extrapolation within a Kalman filter step. At each active surface, the track state is predicted by a dedicated extrapolation algorithm and consequently filtered by incorporating the measurement associated to the step. Under the assumption of Gaussian distributed uncertainties of the track parameters entering the algorithm, the Kalman filter can be shown to yield the optimal estimation of the development of the true track parameters across the detector [173].

The *extrapolation* or *transportation* of the track parameters through the detector is one of the main ingredients to the track reconstruction in ATLAS. The travel of a charged particle through the detector is mainly driven by the presence of the magnetic field as well as by material interactions and energy losses due to Bremsstrahlung or ionization.




Figure 6.1: Simplified illustration of a typical extrapolation within a Kalman filter step. The track representation on the first module is propagated onto the next measurement surface, which results in the track production on the second module. The traversing of the material layer between the two modules causes an increase of the track direction uncertainties and thus, by correlation, an increased uncertainty of the predicted track parameters. [172]

In a homogenous magnetic field and in the absence of material interactions, particles follow helical tracks which, in principle, can be determined using purely analytical computations. However, while the solenoid in ATLAS creates a constant magnetic field near the centre of the detector, it exhibits significant inhomogeneities at the boundaries. Hence, analytical computations are not possible and numerical integration techniques are used instead. For this purpose, ATLAS employs the Runge-Kutta-Nyström method [174–176] to integrate the equations of motion and iteratively approximate the analytical solution.

While travelling throughout the detector material, the path of the particle can deviate as a result of multiple scatter interactions. However, as an independent statistical process, multiple scatterings generally lead to very minimal total mean deflections from the unperturbed particle trajectory, which tends to be negligible for many scattering events. Hence, the effects as a result of multiple scattering processes can be typically absorbed with an increase of uncertainties associated to the track parameters [177] as shown in Figure 6.1.

Moreover, the trajectory of the particle can deviate due to radiative processes. Ionization

losses are typically taken into account by altering the trajectory energy at discrete surfaces along the track, while losses as a result of Bremsstrahlung are required to be incorporated in a continuous manner throughout the extrapolation process.

The ATLAS tracking tools implement two distinct methods to account for both types of energy losses. On one hand, material interactions are approximated by averaging the real material on discrete surfaces. Once a track is transported through the surface, the track parameters and uncertainties are altered according to the averaged expected material effect. On the other hand, continuous energy losses in dense materials are also taken into account [178, 179] and allow for a precise description of track deviations resulting from Bremsstrahlung losses.

6.2 Shower Centre Computation

In order to assign the total energy simulated with ATLFast3 to sensible positions within each calorimeter layer, the track extrapolation described in Section 6.1 is employed to transport tracks throughout the calorimeter system. The transport of particles from a $t\bar{t}$ event is shown in Figure 6.2. Particles that reach the calorimeter envelope are passed to the transport engine,




Figure 6.2: Transport of particles from a $t\bar{t}$ event from the calorimeter envelope (black solid lines) throughout the calorimeter system using the ATLAS tracking tools (black dashed line) resulting in a series of intersections with active calorimeter layers (black dots). The coloured dots show the entrance of each of the calorimeter layers in the ATLAS detector. Note that the shown calorimeter envelope does not exactly correspond to the boundary used to derive the ATLFast3 parametrization. The particle that is not compatible with a production at the origin is the result of a backscatter.

which yields a series of intersections with active calorimeter layers. The coloured dots show only the entrance of the different calorimeter layers and therefore do not exactly coincide with the active layer intersections computed from the track transportation algorithm.

In order to determine sensible positions for each calorimeter layer based on the information provided by the track transport, the layers of the calorimeters can be approximately described as cylinders. The shower centre for a layer can then be determined by computing the intersection between the polygonal chain spanned by the intersections of the extrapolated track path with active layers. In the case that no intersections are found, the position on the layer can be determined as the Closest Point of Approach (CPA) of the polygonal chain to the respective cylinder. Hence, the main geometric task of the ATLF_{FAST3} extrapolation algorithm consists in finding intersections between (infinite) lines and cylinders as well as the closest point of approach of a line segment to an arbitrary cylinder. In Section 6.2.1 and Section 6.2.2 and overview of the mathematical background used to compute the cylinder-line intersections and Closest Points of Approach is given, while Section 6.2.3 provides a detailed algorithmic description of the shower centre determination employed in ATLF_{FAST3}.

6.2.1 Cylinder Intersections

Given a cylinder with radius R and half-length Z , a three-dimensional line given in the coordinate system, where the z -axis corresponds to the symmetry axis of the cylinder, can intersect with

- the two circular faces of the cylinder $C_F = \{x^2 + y^2 = R^2 \mid x, y \in \mathbb{R} \wedge z = \pm Z\}$ or
- with the curved surface $C_C = \{x^2 + y^2 = R^2 \mid x, y \in \mathbb{R} \wedge z \in [-Z, +Z]\}$,

where the cylinder is assumed to be centred around $z = 0$.

The intersection of a given line, parametrized as $\mathbf{l}(\lambda) = \mathbf{m} + \lambda \cdot \mathbf{d}$ with an arbitrary plane defined by all points \mathbf{x} that satisfy

$$\mathbf{n} \cdot (\mathbf{x} - \mathbf{p}) = 0, \quad (6.1)$$

is found for

$$\lambda = \frac{\mathbf{n} \cdot (\mathbf{p} - \mathbf{m})}{\mathbf{n} \cdot \mathbf{d}}, \quad (6.2)$$

where \mathbf{n} is the normal vector orthogonal to the plane and \mathbf{p} and arbitrary point that lies within the plane. Hence, the intersections of a line with the two circular faces C_F of an arbitrary cylinder can be determined by computing the intersections with the (infinite) planes placed at $z = \pm Z$ with the additional constraint that $x^2 + y^2 \leq R^2$.

To compute the intersections with the curved surface of the cylinder, the line spanned by the two points $\mathbf{p}_1 = (p_{1x}, p_{1y}, p_{1z})$ and $\mathbf{p}_2 = (p_{2x}, p_{2y}, p_{2z})$ is projected onto the $x - y$ plane,

where it can be described as $y = m \cdot x + d$ with

$$m = \frac{p_{2y} - p_{1y}}{p_{2x} - p_{1x}} \quad (6.3)$$

and

$$d = \frac{p_{2x}p_{1y} - p_{1x}p_{2y}}{p_{2x} - p_{1x}}. \quad (6.4)$$

The problem is reduced to the computation of the two-dimensional intersection between the projected line and the circle, which can be easily determined by solving the quadratic equation $x^2 + y^2 = R^2$. The corresponding z component can be computed by insertion of the solution into the original line equation.

A three-dimensional illustration of results from the intersection finding algorithm employed in the ATLFast3 shower centre determination algorithm is shown in Figure 6.3 and illustrates intersections found between randomly generated lines and the circular faces and curved surfaces of the cylinders.

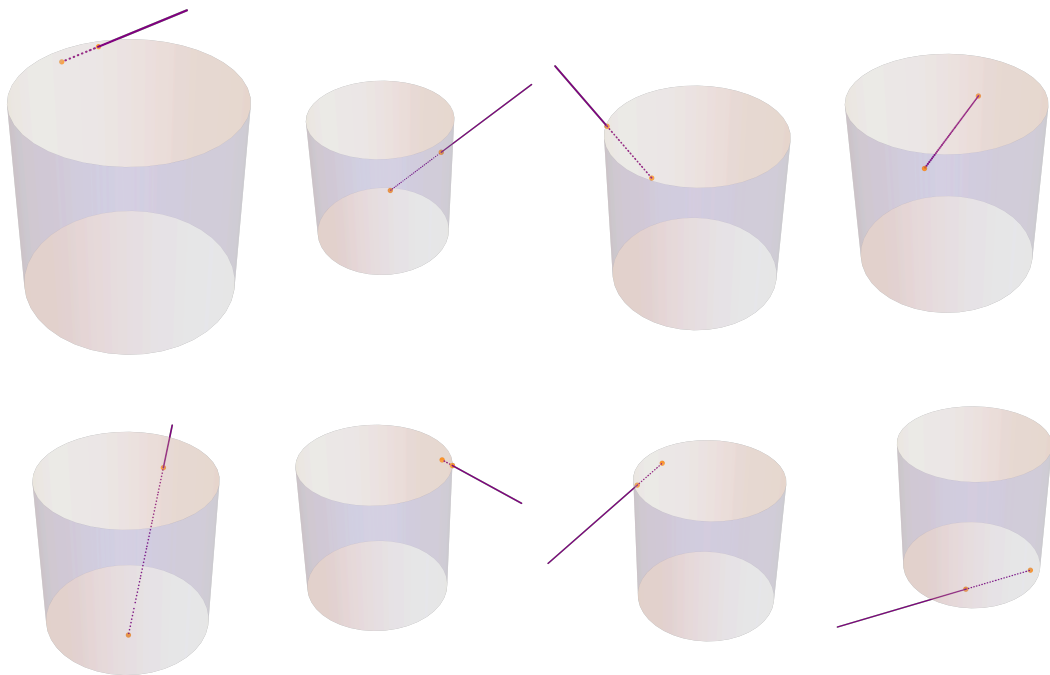


Figure 6.3: Three-dimensional illustration of results from the intersection finding algorithm. An (infinite) line defined by two points is given as input (purple line) and the resulting cylinder intersections (orange points) are computed. The dashed part of the line segment indicates parts of the line that are enclosed within the cylinder.

6.2.2 Closest Point of Approach

In the case that no cylinder intersections of the polygonal chain are found with the cylinder, the shower centre is determined by computing the CPA of the chain to the cylinder. To compute the point of the closest approach of a line segment defined by two endpoints $\mathbf{p}_{A/B}$ on an arbitrary cylinder, various cases are considered.

In the case that the two points that define the line segment are positioned within the Z bounds of the cylinder, that is, $|p_{i,z}| < Z$, the CPA between the cylinder and the line corresponds to the closest approach between the line segment and the z -axis of the cylinder and as such, the problem is reduced to the computation of the closest approach between two lines. Given the two lines

$$\begin{aligned}\mathbf{l}_a(\lambda) &= \mathbf{m}_a + \lambda \cdot \mathbf{d}_a \\ \mathbf{l}_b(\xi) &= \mathbf{m}_b + \xi \cdot \mathbf{d}_b,\end{aligned}\tag{6.5}$$

the vector between any points on the two lines is described by

$$\Delta \mathbf{l}(\lambda, \xi) = \mathbf{l}_b(\xi) - \mathbf{l}_a(\lambda) = \Delta \mathbf{m} + \xi \cdot \mathbf{d}_b - \lambda \cdot \mathbf{d}_a,\tag{6.6}$$

where $\Delta \mathbf{m} = \mathbf{m}_b - \mathbf{m}_a$. The vector between the two closest points $\mathbf{k}(\lambda_0, \xi_0)$ is perpendicular to the direction vectors \mathbf{d}_a and \mathbf{d}_b . As a result, the two linear equations

$$\begin{aligned}\mathbf{k}(\lambda_0, \xi_0) \cdot \mathbf{d}_a &= \Delta \mathbf{m} \cdot \mathbf{d}_a + \xi_0 \cdot \mathbf{d}_a \cdot \mathbf{d}_b - \lambda_0 = 0 \\ \mathbf{k}(\lambda_0, \xi_0) \cdot \mathbf{d}_b &= \Delta \mathbf{m} \cdot \mathbf{d}_b + \xi_0 - \lambda_0 \cdot \mathbf{d}_a \cdot \mathbf{d}_b = 0\end{aligned}\tag{6.7}$$

need to be satisfied. The system of equations can be easily solved and yields:

$$\lambda_0 = \frac{(\Delta \mathbf{m} \cdot \mathbf{d}_a) - (\Delta \mathbf{m} \cdot \mathbf{d}_b) (\mathbf{d}_a \cdot \mathbf{d}_b)}{1 - (\mathbf{d}_a \cdot \mathbf{d}_b)^2}\tag{6.8}$$

and

$$\xi_0 = -\frac{(\Delta \mathbf{m} \cdot \mathbf{d}_b) - (\Delta \mathbf{m} \cdot \mathbf{d}_a) (\mathbf{d}_a \cdot \mathbf{d}_b)}{1 - (\mathbf{d}_a \cdot \mathbf{d}_b)^2}.\tag{6.9}$$

In the case that one of the points which defines the line segment lies outside the Z bounds of the cylinder, that is, $|p_{i,z}| > Z$, analytical solutions are more complex such that the CPA is computed using an iterative approach. A notable exception occurs when the line segment is approximately parallel to the circular faces of the cylinder. In such cases, the problem is reduced to the computation of two-dimensional circle-line intersections and is performed as described in the case of the intersection finding procedure. For all other cases, the line segment is first divided into equidistant parts. Then, the CPA finding algorithm iteratively evaluates the distance between the line at its current position and the projected position on the cylinder. The iterations are performed from both sides of the segment. The algorithm terminates its execution




Figure 6.4: Sketch of the iterative Closest Point of Approach (CPA) finding procedure. A segment defined by two bounds $B_{1/2}$ is divided into equidistant parts of size S . The distance between the current position on the line segment and the projected position on the cylinder is iteratively computed at each step. The algorithm starts the computation from both bounds and moves towards the centre of the segment. The computation stops if the computed distance to the cylinder increases from one iteration to another. The CPA is then taken as the centre between the current and previous iteration. The precision of the CPA finding procedure is fully determined by the choice of the step size.

once the distance to the cylinder from the current evaluation increases with respect to the prior step. In this case, the true CPA must lay between both iterations and is hence taken as average position between both bounds. A sketch of the iterative CPA finding procedure is shown in Figure 6.4. The precision of the computed CPA is fully determined by the step size chosen in the division of the line segment and is set to 0.01 mm. It is noteworthy that the iterative CPA computation is typically only called in rare edge cases and is not observed to negatively affect the computational performance in any way.

A three-dimensional representation of results from the CPA finding algorithm is shown in Figure 6.5. A non-intersecting line segment defined by two endpoints is given as input and the resulting CPA is computed on the line segment and then projected onto the cylinder surface.

6.2.3 Algorithmic Description and Choice of Cylinders

The extrapolation algorithm employed in ATLFast3 is wrapped within a single `extrapolate` method that consists of consecutive calls of three main methods and receives the state of the truth particle which is to be subsequently simulated with `FASTCALOSIMV2` or `FASTCALOGAN` with the aid of the resulting extrapolation state. A simplified control-flow graph of the ATLFast3 extrapolation engine with a representation of the most important methods is shown in Figure 6.6. The following methods are executed in sequence:

- **caloHits**: based on the state of the particle received at the boundary between ID and calorimeter envelope, the particle is transported throughout the calorimeter system using the ATLAS tracking tools as described in Section 6.1. The output is a set of intersections with active calorimeter layers (*hits*) ordered according to the time the intersection is recorded, that is, according to the path the particle is transported through the detector.
- **extrapolateToID**: based on the $R - Z$ value pairs shown in Table 4.1, the position of the particle is extrapolated onto the boundary between ID and calorimeter system. Three

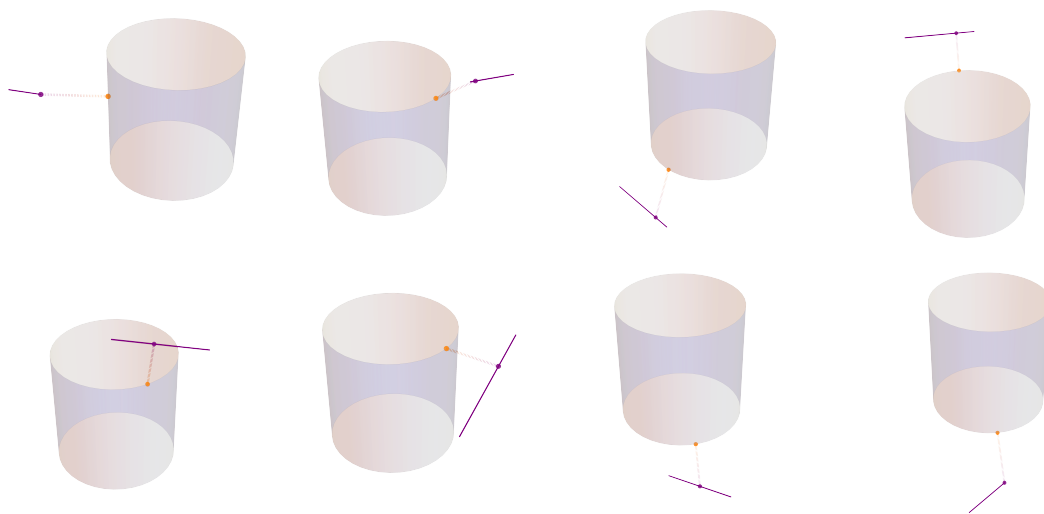


Figure 6.5: Three-dimensional representation of results from the Closest Point of Approach (CPA) finding algorithm. A non-intersecting line segment defined by two endpoints is given as input (purple line) and the resulting CPA is computed on the line segment (purple point) and then projected onto the cylinder surface (orange point). The dashed line indicates the path with the closest distance between the cylinder and the line segment.

cylinders are constructed with the corresponding radii and half-lengths and the position is extrapolated onto each of the cylinders based on the vector of hits obtained from the `caloHits` method.

More precisely, the vector of hits is passed to the `extrapolateToCylinder` method. In the case that the vector contains one hit, the projected hit position on the cylinder is taken as the extrapolated position. If the vector contains more than one hit, possible intersections of the polygonal chain spanned by the set of hits with the cylinders are computed within `extrapolateWithIntersection`. The method performs a pairwise loop over the hit vector and returns an intersection in the case that one of the hits is already positioned on the cylinder surface or a hit segment spanned by two hits crosses the cylinder surface. If there is no direct intersection between the line segments spanned by each of the hit pairs, the last two hit positions are used to linearly extrapolate to the cylinder surface by computing potential forward intersections, that is, intersections in the direction of the last recorded hit position.

Backward intersections, that is intersections in the direction opposite to the particles' flight path are only allowed under certain circumstances in the case of the first hit pair. This is required as a result of a small mismatch between the definition of the boundary between ID and calorimeter system as defined within ISF, that is employed by `FASTCALOSIMV2`, and as defined within `ATHENA`, which is used by the ATLAS tracking

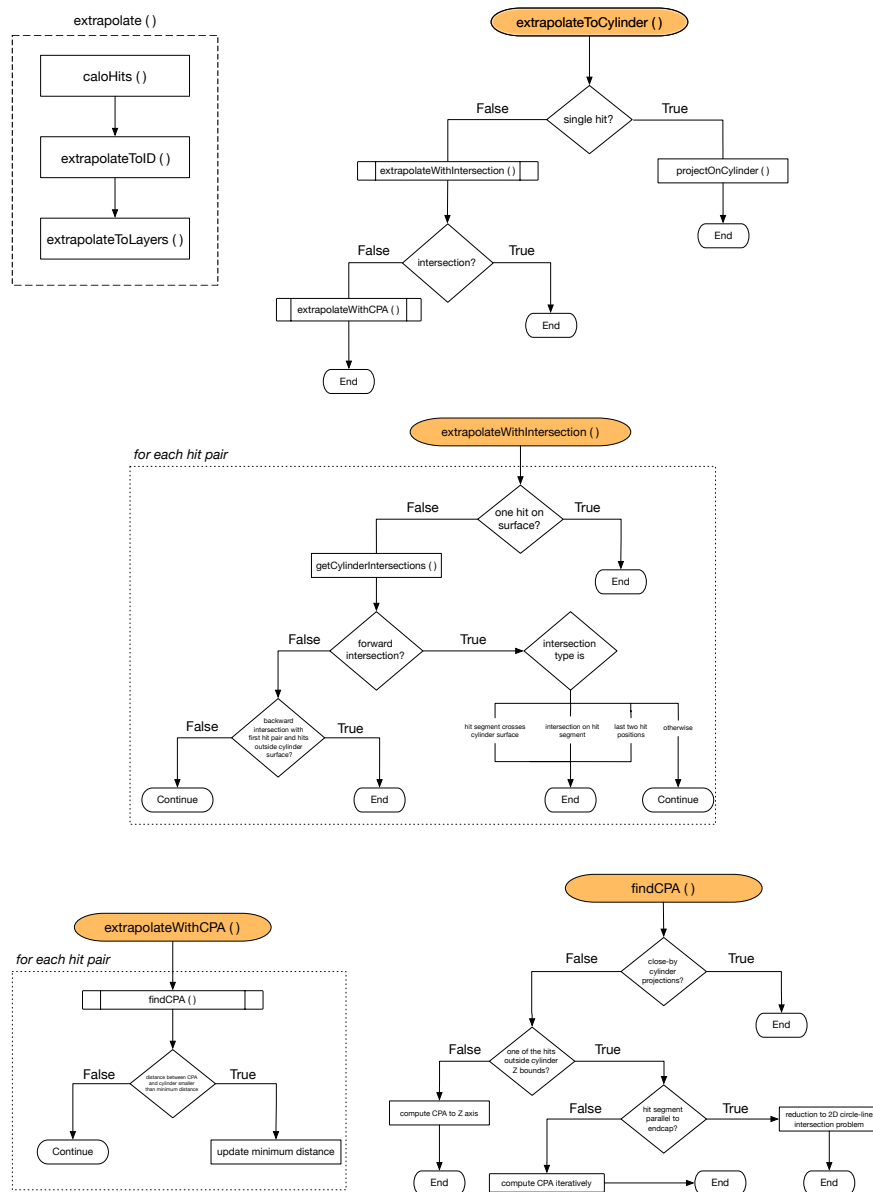


Figure 6.6: Simplified control-flow graph of the ATLFast3 extrapolation engine with a representation of the most important methods. The dashed box gives a high-level-overview of the programme flow starting from the track transport with `caloHits` to the extrapolation of the track parameters to the boundary between ID and calorimeter system with `extrapolateToID` and to all layers of the calorimeter with `extrapolateToLayers`. Code blocks within the dotted regions are executed for each pair of hits which is returned from the track transport.

tools (see Chapter 9 for further details).

In the case that no intersections are found, the hit vector is passed to the `extrapolateWithCPA` method, which finds the CPA on the cylinders for each hit segment of the polygonal chain, and returns the CPA with the minimum distance to the cylinders.

Finally, from the extrapolated points on the three cylinders, the position that is most compatible with lying on the actual boundary between ID and calorimeter system is selected and used to choose the appropriate parametrization during the simulation with `FASTCALOSIMV2` or `FASTCALOGAN`.

- **extrapolateToLayers:** suitable positions are computed for the entrance, middle and exit positions of all calorimeter layers by using the same `extrapolateToCylinder` method employed for the extrapolation of the hit vector to the boundary between ID and calorimeter system. For barrel layers, the radius of the cylinder onto which the extrapolation is performed is determined from the layer radius at the η position of the particle at the calorimeter surface. In the case of the EMB0 - EMB3 layers, the half-lengths of the cylinders are chosen as the z positions that correspond to the entrance of the EME1 layer. For all other barrel layers, the half-length is chosen according to their actual physical measures.

Similarly, in the case of endcap layers, the cylinder half-length is chosen according the real calorimeter layer measurements, depending on the η position of the particle on the calorimeter surface. Moreover, the corresponding cylinder radius is determined from the maximum η extension of the layer. More specifically, the maximum azimuthal angle is calculated as

$$\theta_{\max} = 2 \cdot \arctan(e^{-\eta_{\max}}), \quad (6.10)$$

such that the maximum cylinder radius is given by

$$R_{\max}^{\text{cyl}} = z^{\text{cyl}} \cdot \sqrt{\sec^2 \theta_{\max} - 1}. \quad (6.11)$$

Finally, the extrapolated positions are scaled to fit the radius (half-length Z) of the cylinder in the case of barrel (endcap) layers. The applied scaling to the extrapolated r and z values are only non-unitary if the extrapolation on a barrel layer yields a position on one of the two circular faces of the cylinder, or if the extrapolation on an endcap layer yields a position on the curved surface of the cylinder. In both cases, the extrapolated ϕ and η positions remain intact. This method allows to extrapolate the set of hits to cylinders that approximately match the physical calorimeter layers, while ensuring that the determined positions lie on sensible positions with regard to the orientation of the calorimeter layers. The extrapolated positions on the various calorimeter layers based on the collection of hits shown in Figure 6.2 are shown in the $R - Z$ plane in Figure 6.7.




Figure 6.7: Extrapolated positions for the various calorimeter layers from a $t\bar{t}$ event in the $R - Z$ plane. The extrapolation is based on the identical set of hits shown in Figure 6.2. As a reference, the grey dashed lines show the particle transport path through the calorimeter system, that is, the connection between the set of hits generated by the ATLAS tracking tools. The grey solid line indicates the boundary between ID and calorimeter system as defined within ISF.

The final shower centre position in each calorimeter layer is calculated by weighting the extrapolated positions from the entrance and exit surface of the layer according to

$$x_c = (1 - \omega_{\text{extr}}) \cdot x_{\text{entrance}} + \omega_{\text{extr}} \cdot x_{\text{exit}}, \quad (6.12)$$

where ω_{extr} is the extrapolation weight, and x_{entrance} (x_{exit}) are the extrapolated positions at the entrance and exit of the layers, respectively. In principle, the shower centre can be computed with respect to the longitudinal centre of the calorimeter layer, that is, $\omega_{\text{extr}} = 0.5$. However, for a more accurate description of the shower position, the weight is chosen as the energy weighted average longitudinal depth in each layer, and computed independently in each PCA bin of the parametrization. In the future, it is planned to employ a neural-network-based prediction of the extrapolation weights, which can further help to improve the description of the shower position. A simplified illustration showing the computed shower centres based on the extrapolated particle positions on the entrance and exit of a calorimeter layer for various extrapolation weights is shown in Figure 6.8.




Figure 6.8: Simplified illustration showing the computed shower centres from the extrapolated particle positions on the entrance and exit of the calorimeter layer for various extrapolation weights and incoming particle directions compatible with neutral particles produced at the origin. For charged particles that are bent in the magnetic field, the extrapolations do not necessarily need to lie on a straight line.

6.3 Validation and Limitations

The performance of the new shower centre determination algorithm in `ATLFAST3` has been extensively validated. First, the main methods `extrapolateWithIntersection` and `extrapolateWithCPA` that calculate the intersections between lines and cylinders as well as the closest point of approach between cylinders and line segments are validated with unit tests designed to cover a broad range of possible scenarios which are expected to arise in real extrapolation processes. Similar to the illustrations shown in Figure 6.3 and Figure 6.5, the results of these unit tests are graphically validated.

Moreover, the results of real extrapolation processes to the various calorimeter layers are graphically visualized and compared against the results of the extrapolation algorithm employed in `ATLFASTII`. In the high energy regime, where the curvature of the particles in the toroidal magnetic field is negligible or sufficiently small, the results of the updated extrapolation algorithm typically match with the computations used in the predecessor of `ATLFAST3`. In contrast, for low energy particles, the updated extrapolator is observed to perform significantly better. More specifically, for particles up to a few GeV, the predecessor algorithm often fails to assign suitable positions on the calorimeter layers or entirely fails to find the shower centre position. In these cases, the total energy simulated cannot be assigned to the calorimeter cells in the layer, which will result in missing energy in the reconstruction of physics objects. An illustration of a real extrapolation process to the `PreSamplerB` layer of a 100 GeV e^+ and a 1 GeV proton is shown in Figure 6.9. The results of the extrapolation are shown for the old (yellow dot) and updated algorithm (cyan dot). In the high-energy case, the positions calculated on the calorimeter layer match exactly. In the low energy case, where the curvature of the particle in

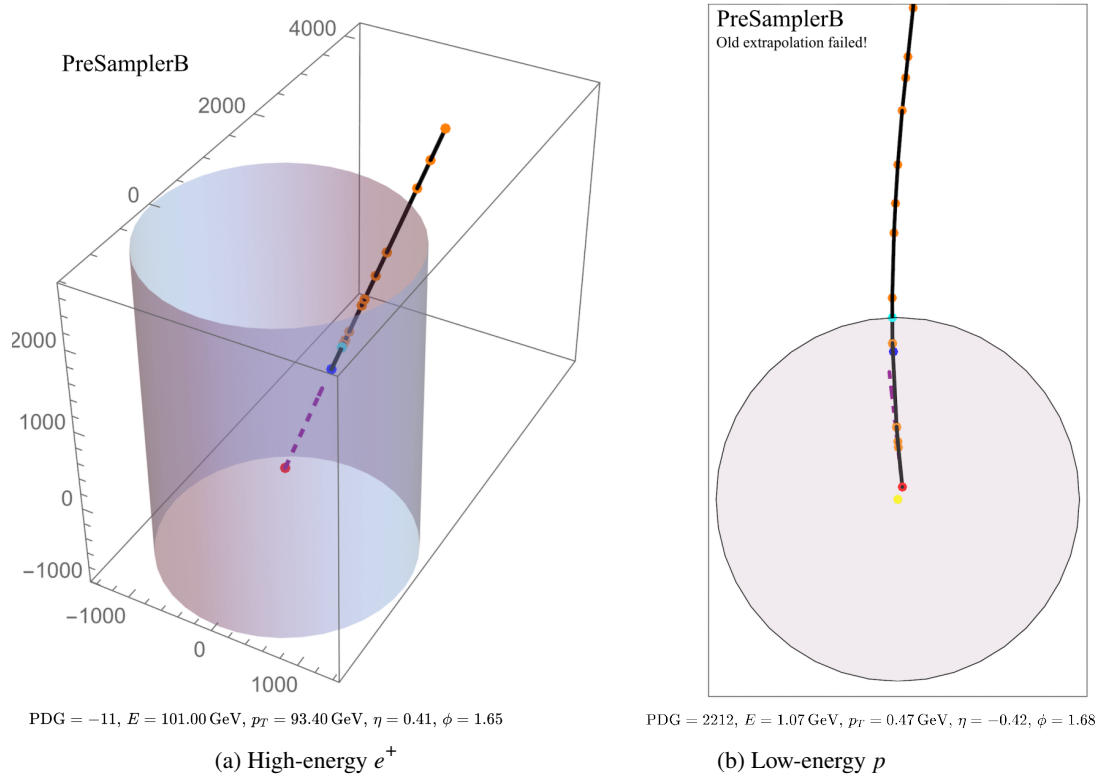


Figure 6.9: Representation of an extrapolation process for a (a) high-energy positron and a (b) low-energy proton to the PreSamplerB layer. The red dot indicates the production vertex, while the dashed purple line shows the momentum direction at the vertex of the particle. The blue and cyan dots show the extrapolated position on the boundary between ID and calorimeter system and the PreSamplerB layer, respectively. The orange dots are the positions from the track transport upon which the extrapolation is based on. The yellow dot illustrates the position determined by the old extrapolation algorithm. It fully overlaps with the cyan dot in (a) and is assigned a dummy value in (b) as a result of a failure to find a suitable position on the layer. The new algorithm finds a sensible position on the calorimeter layer.

the electromagnetic field is clearly visible, ATLFastII fails to find a suitable intersection, while the new extrapolation engine successfully determines the correct position on the calorimeter layer.

To determine the fraction of failed extrapolations resulting from the shower centre finding algorithm employed in ATLFastII and the updated algorithm used within ATLFast3, extrapolations are performed from particles originating from a total of 1000 $t\bar{t}$ events. Particles that are unlikely to reach the boundary between Inner Detector and calorimeter system and are thus not simulated with FASTCALOSIMV2 or FASTCALOGAN are not considered:

- Very forward particles with $|\eta| > 5$, which are unlikely to be secondary particles and are produced close to the intersection point, with vertex positions $|V_x| < 10$ mm, $|V_y| < 10$ mm and $|V_z| < 200$ mm, are rejected
- Interactions very close to the beam pipe, that is, particles with production vertex $\sqrt{V_x^2 + V_y^2} < 50$ mm which subsequently travel in the very forward direction $|\eta| > 5$ are rejected
- Interactions outside the detector range, that is, particles that travel along the beam pipe and undergo interactions with the beam pipe and have $V_z > 7000$ mm are rejected.

The resulting fraction of failed extrapolation with the old and new shower centre determination algorithm for 1000 $t\bar{t}$ events is shown in dependence of η and p_T in Figure 6.10. The majority




Figure 6.10: Fraction of failed extrapolations with the old (blue) and updated (orange) shower centre determination algorithm employed in ATLFast3 as simulated with 1000 $t\bar{t}$ events in dependence of true particle η and transverse momentum. For particles with transverse momenta in the order of a few GeV, the number of failed extrapolations is negligible in both cases.

of extrapolation failures in the decommissioned algorithm is observed in the central region

$|\eta| < 1$ and for very low energy particles, particularly in a $0.2 < p_T < 0.4$ GeV window. In the central region, the fraction of failed extrapolations is as high as 1.6% and reaches up to 3.5% in certain p_T bins. The overall fraction of failed extrapolations is at approximately 0.6%. In the re-designed extrapolation algorithm, the number of extrapolation failures is greatly reduced. While a slight accumulation of failures in the central region is still observed, the failure fraction is well below 0.1% for $|\eta| < 1$ and below 0.05% for more forward regions. It is noteworthy that particles with $E_{\text{kin}} < 10$ MeV, pions with $E_{\text{kin}} < 200$ MeV and other hadrons with $E_{\text{kin}} < 400$ MeV are currently simulated with GEANT4 (see Chapter 4) and as such, very low energy particles are not required to be extrapolated with ATLFast3. However, the choice of thresholds may change in future releases and therefore correct results are also important in the very low energy regime.

The remaining extrapolation failures that are seen with the re-designed shower determination algorithm are graphically analysed and observed to stem from fundamental limitations of the linear extrapolation model. Very low energy particles can create helix-like trajectories in the toroidal magnetic field. As a result of the strong curvature, the linear extrapolation based on the last two hit positions will in many cases lead to an unsuitable prediction of the particles' intersection with the calorimeter layer, which can subsequently lead to extrapolation failures. Figure 6.11 shows a three-dimensional illustration of two extrapolation processes of very low-energy pions to the PreSamplerE layer. In Figure 6.11(a), a clear helical structure is visible from the connection of the sets of hits generated by the track transport, while in Figure 6.11(b), the helical path of the particle can be assumed from the last active layer intersection, effectively creating a *kink* in the polygon spanned by the hits. In both cases, the extrapolation to the PreSamplerE layer is purely based on the last two hits, which does not necessarily correspond to the true intersection of the particle with the respective calorimeter layer. It may be possible to develop more involved models that have the ability to detect helical paths or kinks in the track transport, such that more suitable positions can be assigned to the layers in these cases. However, the fraction of extrapolation failures overall and in the low p_T region is small. Moreover, energy depositions that would have originated from low energy spiralling particles causing ill-defined extrapolation calls are unlikely to have any significant impact on reconstructed physics objects.

Finally, as a last method of validation, the shower centre determination algorithm was extensively validated as part of the overall ATLFast3 performance validation of reconstructed physics objects that is described in Chapter 7.




Figure 6.11: Three-dimensional illustration of two extrapolation processes of very low-energy pions in the order of 10-50 MeV which exhibit a (a) helical structure and a (b) kink. In these cases, the algorithm is not guaranteed to provide reasonable results, which can lead to extrapolation failures. The blue and cyan dots show the extrapolated position on the boundary between ID and calorimeter system and the PreSamplerE layer, respectively. The orange dots are the positions from the track transport upon which the extrapolation is based on. The yellow dot indicates the position as computed with the decommissioned extrapolator.

Physics Performance of AtIFAST3

THE paramount objective of any fast simulation tool in ATLAS is to reproduce the detailed GEANT4-based simulation with utmost precision while maintaining minimal CPU resource utilization. In practice, an exact agreement between fast and full simulation is not possible for all regions of phase space and observables of interest. Consequently, fast simulation tools generally sacrifice some of the full simulation accuracy in exchange for a considerable reduction in simulation time. At the same time, however, an exact agreement at the simulation level is in many cases not necessary. Additional soft pp collisions from pileup and detector effects such as noise from the readout electronics, which are independent of the simulation, may smear out residual differences, leaving the underlying simulation sufficiently accurate for most physics analyses. Therefore, fast simulation merely aims to provide a level of accuracy such that no sizeable differences from the GEANT4 simulation can be resolved by the algorithms used to reconstruct physics objects such as electrons, photons, τ -leptons and jets, allowing most analyses to entirely replace the use of full simulation with ATLFAST3.

As part of the re-processing of Run 2 physics samples, ATLFAST3 was used to simulate the detector response of approximately 7 billion events from a wide variety of physics samples. Prior to the production campaign, ATLFAST3 was thoroughly validated by comparing the modelling of reconstructed observables typically used in physics analyses, and was shown to outperform its predecessor AFII in almost all aspects.

This chapter discusses the main improvements of ATLFAST3 with respect to its predecessor, and presents comparisons of reconstructed quantities with the detailed GEANT4 simulation. A more comprehensive overview of the physics performance of ATLFAST3 is described in Ref. [134], on which this chapter is loosely based. It is worth noting that the FASTCALOSIMV2 energy response corrections (see Chapter 5) and the novel extrapolation engine employed in ATLFAST3 (see Chapter 6) are used in the ATLFAST3 release for which the following benchmark results have been obtained.

7.1 Validation data sets

To validate the physics performance of ATLFast3, a wide range of different processes are simulated and compared to GEANT4. The simulated samples are chosen to reflect the broad physics programme of the ATLAS Collaboration and to probe the simulation of all relevant reconstructed physics objects. Table 7.1 summarizes the full set of samples used for the validation and the event generator settings used in each case.

Process	Generator	ME Order	PDF	Parton Shower	Tune
$t\bar{t}$	POWHEG BOX r2330.3	NLO	CT10	PYTHIA 6.427	P2012
$Z^*/\gamma^* \rightarrow \tau\tau$	POWHEG BOX r2856	NLO	CTEQ6L1	PYTHIA 8.186	AZNLO
$Z \rightarrow \mu^+\mu^-$	POWHEG BOX r2856	NLO	CTEQ6L1	PYTHIA 8.186	AZNLO
$Z \rightarrow ee$	POWHEG BOX r2856	NLO	CTEQ6L1	PYTHIA 8.186	AZNLO
$H \rightarrow \gamma\gamma$	POWHEG v2 NNLOPS	NNLO	PDF4LHC15	PYTHIA 8.230	AZNLO
Dijet	PYTHIA 8.186	NLO	NNPDF23	PYTHIA 8.186	A14
$W'(13 \text{ TeV})$	PYTHIA 8.235	LO	NNPDF23LO	PYTHIA 8.235	A14
$\rightarrow WZ \rightarrow 4q$					
$Z'(4 \text{ TeV})$	PYTHIA 8.235	LO	NNPDF23LO	PYTHIA 8.235	A14
$\rightarrow t\bar{t}$					

Table 7.1: Validation datasets employed to validate the physics performance of ATLFast3 and the Monte Carlo generator settings used in each case. ME refers to the matrix element order used in each case. Adapted from Ref. [134]

The simulation of electrons, muons, and b - and light quark initiated jets can be probed with semi-leptonic $t\bar{t}$ decays, where one of the top quark creates a lepton $\ell \in [e, \mu]$ in the final state $t \rightarrow bW(W \rightarrow \ell\nu_\ell)$ and the other top quark decays hadronically $t \rightarrow bW(W \rightarrow q\bar{q})$. For this purpose, semi-leptonic $t\bar{t}$ decays are simulated with POWHEG BOX r2330.3 [180], and interfaced with PYTHIA 6.427 [181] for the parton shower and hadronization modelling with the CT10 [182] set of parton distribution functions (PDFs) and the P2012 set of tuned parameters [183].

To validate the simulation of electrons, muons and τ -leptons, the decay of Z bosons into a pair of leptons is simulated using POWHEG BOX r2856 at NLO in QCD using the CTEQ6L1 [184] PDF set. The events were interfaced with PYTHIA 8.186 [185] for the parton shower and hadronization modelling using the AZNLO tune [186].

Similarly, the simulation of photons is probed in $H \rightarrow \gamma\gamma$ decays. The considered Higgs boson production mode is gluon-gluon fusion only and was simulated using POWHEG v2 NNLOPS [187–191] at NNLO accuracy with the PDF4LHC15 PDF set [192] and the AZNLO [186] tune. The simulation was interfaced with PYTHIA 8.230 [193] for the parton shower and non-perturbative hadronization effects.

To study the general performance of jet simulation in a broad p_T range, multijet events were generated with PYTHIA 8.186 [185] interfaced with EvtGen 1.20 [194] for simulating the decay of heavy-flavour mesons. The NNPDF23 [195] PDF set and the A14 tune [196] were used.

Finally, to validate the simulation performance of boosted jets, particularly their substructure, hypothetical spin-1 W' and Z' resonances are generated using PYTHIA 8.235 [193] at LO with the NNPDF23LO [197] PDF set and the A14 tune [196]. A detailed description of boosted jet substructure is provided in Part III of this thesis. The W' resonance was generated with a mass of 13 TeV and decayed into a WZ pair that subsequently decays into a fully hadronic final state. The Z' resonance is simulated in the $t\bar{t}$ decay mode and assumed to have a mass of 4 TeV. In both cases, the differential cross-section is reweighted to yield a flat jet p_T distribution to populate kinematic regions with low statistics.

For all considered data sets, pile-up is modelled by overlaying signals from simulated inelastic pp events with the hard-scattering event [158]. An average number of pile-up interactions per pp bunch crossing of 38 with a standard deviation of 12 was used, resembling the conditions during the Run 2 data taking period. The pile-up events were generated with PYTHIA 8.186 [185] and use the NNPDF23 [195] PDF set and the A3 tune [198]. The detector response of the pile-up events was simulated with GEANT4 in all cases.

7.2 Performance

This section presents comparisons between GEANT4, ATLFast3 and its predecessor AFII for a variety of reconstructed objects typically used in physics analyses, such as photons, electrons, jets, hadronically decaying τ -leptons and muons. The modelling of most properties of reconstructed objects in ATLFast3 does not show any sizeable differences with respect to GEANT4. Therefore, this section focuses on the most important features with respect to its predecessor. In particular, this section does not present comparisons for variables that depend primarily on the Inner Detector, as no inherent differences are expected compared to GEANT4 in this case (see Chapter 4). A more detailed overview of the physics performance of ATLFast3 can be found in Ref. [134].

7.2.1 Electrons and Photons

Electron and photon candidates in ATLAS are reconstructed using topological clusters [199] of energy deposits [200]. Electrons are objects consisting of clusters of energy deposits in the calorimeter, which are matched to a track in the tracking system. Converted photons that interact with the material in the ID and undergo $\gamma \rightarrow e^+e^-$ transitions are clusters matched to a conversion vertex, whereas unconverted photons are clusters without a matched track in the ID. A more detailed discussion on track reconstruction in ATLAS is provided in Chapter 6.

Photons and electrons are identified primarily based on electromagnetic shower shape properties. Three identification categories are typically defined for photons [133] and electrons [201] depending on the stringency of the imposed identification criteria: *loose*, *medium* and *tight* (see Chapter 8 for more details). The correct modelling of identification efficiencies in simulation is crucial to ensure a correct prediction of physics processes with photons and electrons in the final state. Figure 7.1 shows the tight identification efficiencies for photons from $H \rightarrow \gamma\gamma$ decays and single electrons as simulated using GEANT4, ATLFast3 and AFII. In both cases, the

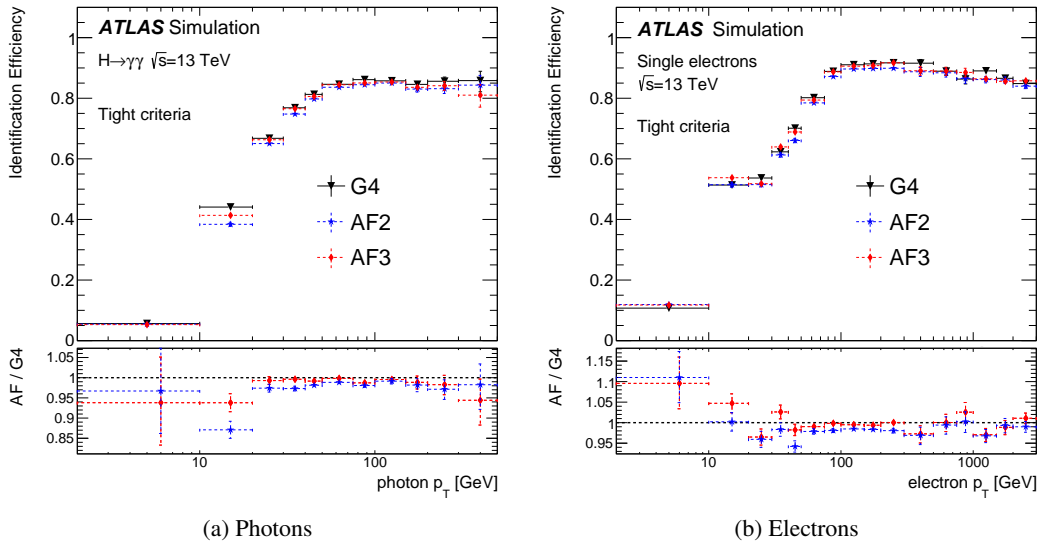


Figure 7.1: Tight identification efficiencies for (a) photons from $H \rightarrow \gamma\gamma$ decays and (b) single electrons as a function of p_T in the pseudorapidity region $|\eta| < 2.5$. The lower pad shows the ratio between the fast and full simulation. [134]

modelling is significantly improved compared to AFII. For electrons in a p_T range between 30 and 300 GeV, AFII agrees with the full simulation within a 5% margin, while ATLFast3 achieves agreement within 2% over most of the kinematic range. Similarly, except for the very low p_T region, the identification of photons in ATLFast3 is correctly modelled to within a few percent.

To further probe the simulation performance of photons and electrons in higher-level variables, the reconstructed invariant di-photon mass $m_{\gamma\gamma}$ from simulated $H \rightarrow \gamma\gamma$ events and the reconstructed dielectron mass m_{ee} from $Z \rightarrow ee$ events is shown in Figure 7.2. Diphoton events are selected by requiring two photons with $p_T > 0.35m_{\gamma\gamma}$ and $p_T > 0.25m_{\gamma\gamma}$, selections that are typically imposed in $H \rightarrow \gamma\gamma$ analyses to avoid a bias in the Higgs boson line shape. Dielectron events are selected by requiring two electrons with $p_T > 25$ GeV each. In both cases, the transition region between the barrel and endcap of the calorimeter is excluded and the pseudorapidity range is limited to $|\eta| < 2.47$. Excellent modelling of the reconstructed invariant

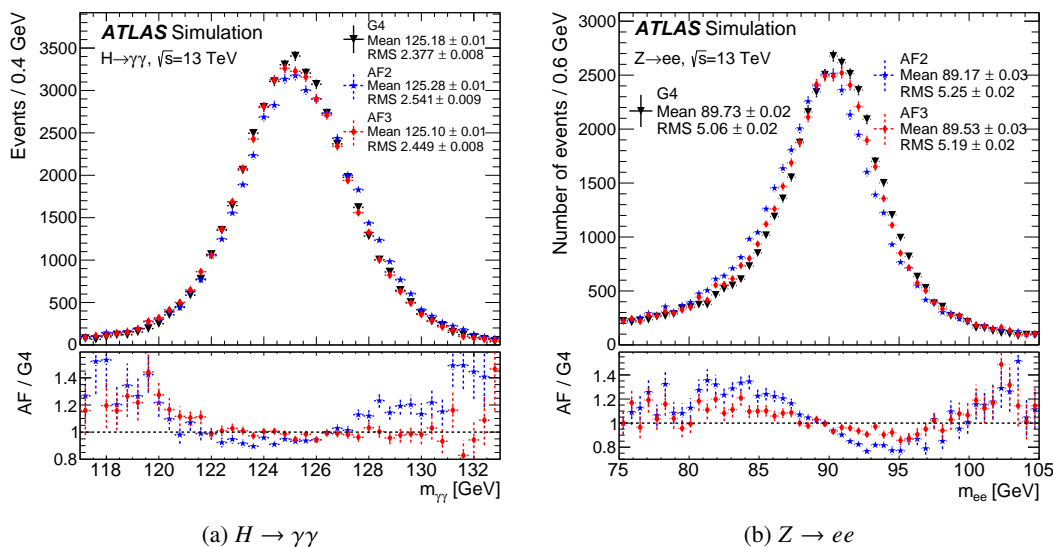


Figure 7.2: (a) Reconstructed di-photon invariant mass from simulated $H \rightarrow \gamma\gamma$ decays and (b) reconstructed dielectron invariant mass from simulated $Z \rightarrow ee$ decays. Photons are required to fulfil $p_T > 0.35m_{\gamma\gamma}$ and $p_T > 0.25m_{\gamma\gamma}$, while both electrons are required to have $p_T > 25$ GeV. [134]

masses is required for many precision measurements in ATLAS. With respect to ATLFAST3, the modelling of $m_{\gamma\gamma}$ and m_{ee} is improved significantly as a result of the energy response corrections discussed in Chapter 5 and a dedicated correction of the energy resolution [134]. The mean of the diphoton invariant mass is well reproduced with both fast simulation tools, with a small improvement in the relative difference to GEANT4 from 0.08% in AFII to 0.06% in ATLFAST3. The relative difference for the standard deviation (RMS) improves more significantly from 6.6% to 3.0%. For reconstructed $Z \rightarrow ee$ events, the relative difference for the mean of the dielectron invariant mass distribution is improved from 0.6% in AFII to 0.2% in ATLFAST3, while the mean is improved from about 3.7% to 2.5%.

7.2.2 Jets

In ATLAS, jets are reconstructed using a variety of reconstruction algorithms and use information from the calorimeter system and the ID. A more detailed description of jet reconstruction is provided in Part III of this thesis.

The overall ATLFAST3 jet performance is assessed using top jets reconstructed from a $t\bar{t}$ sample using EMPFlow [202] jets, which are constructed with the anti- k_t [203] clustering algorithm with a size parameter $R = 0.4$ and use topological clusters [199] at the EM scale and tracks from the ID as input.

The kinematic properties of these small- R jets are very well modelled by ATLFast3. Figure 7.3 shows the transverse momentum of p_T -leading EMPFlow jets and the pseudorapidity of p_T -sub-leading EMPFlow jets reconstructed in a $t\bar{t}$ sample simulated with GEANT4, ATLFast3 and AFII. Both ATLFast3 and AFII are consistent with the GEANT4 simulation at the percent level,

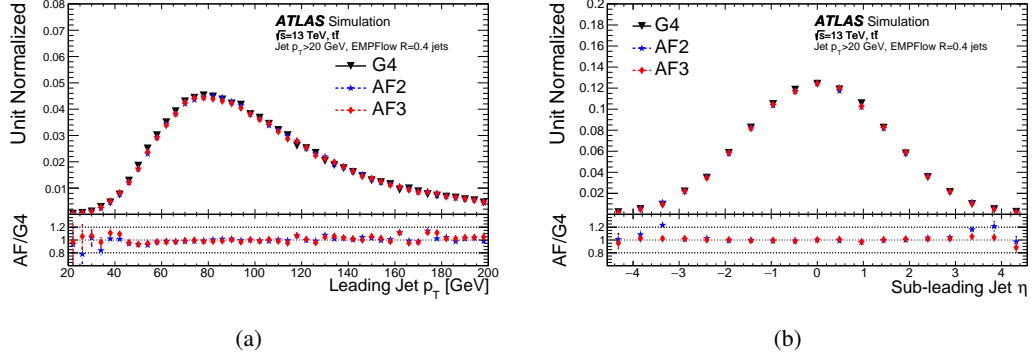


Figure 7.3: (a) Transverse momentum of the p_T -leading jet and (b) pseudorapidity of the p_T -sub-leading jet reconstructed in a $t\bar{t}$ sample. Jets are reconstructed as EMPFlow $R = 0.4$ jets and required to have $p_T > 20$ GeV. [134]

but ATLFast3 shows improved modelling in the very forward region of the detector as a result of a dedicated parametrization for the FCal.

The decay products of massive particles that receive a sufficiently large Lorentz boost are typically reconstructed using large- R jets. The internal substructure of these jets is then exploited to identify their origin and discriminate between signal and background jets. A more detailed discussion on boosted jets and their substructure is provided in Part III of this thesis. To investigate the performance of ATLFast3 on large- R jets, LCTopo [199] and UFO [204] jets, both reconstructed with the anti- k_r [203] clustering algorithm with a size parameter $R = 1.0$, are used. LCTopo jets use locally calibrated topological clusters as input and were the default choice for reconstructing large- R jets in ATLAS during Run 2. UFO jets are constructed from a novel type of jet input known as Unified Flow Objects (UFOs) [204] that combine information from the calorimeter and the ID, which has been shown to provide an improved jet mass resolution and to be beneficial for background jet rejection.

A precise simulation of the detailed structure within high- p_T jets plays a crucial role in their correct identification. Figure 7.4 shows the number of jet constituents of p_T -leading jets in a $Z'(4 \text{ TeV}) \rightarrow t\bar{t}$ sample as simulated with GEANT4, ATLFast3 and AFII for EMPFlow $R = 0.4$ and UFO $R = 1.0$ jets. In AFII, the number of jet constituents is significantly underestimated, particularly for small- R jets. ATLFast3 accurately reproduces the number of jet constituents observed for UFO $R = 1.0$ jets in the GEANT4 simulation. Similarly, the number of constituents is well simulated for EMPFlow $R = 0.4$ jets with many constituents. Some mismodelling is observed for jets with fewer than ~ 15 constituents, but the difference with respect to GEANT4

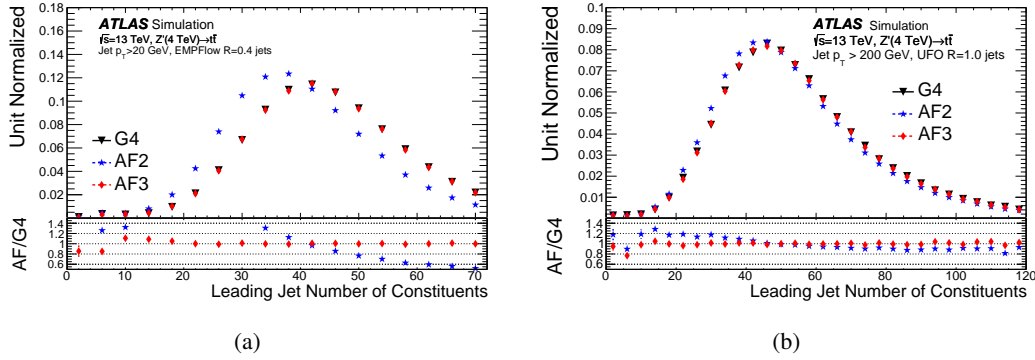


Figure 7.4: Number of jet constituents of p_T -leading jets in a $Z'(4 \text{ TeV}) \rightarrow t\bar{t}$ sample reconstructed as (a) EMPFlow $R = 0.4$ and (b) UFO $R = 1.0$ jet. Small- R (large- R) jets are required to have $p_T > 20$ (200) GeV. [134]

is reduced from about 20% in AFII to 10% in ATLFast3.

Figure 7.5 shows the simulated energy correlation function D_2 [205] and the N -subjettiness ratio τ_{32} [206], substructure variables that are typically used to identify boosted two-body and three-body decays, respectively (see Section 10.2 for the exact definitions). The D_2 distribution

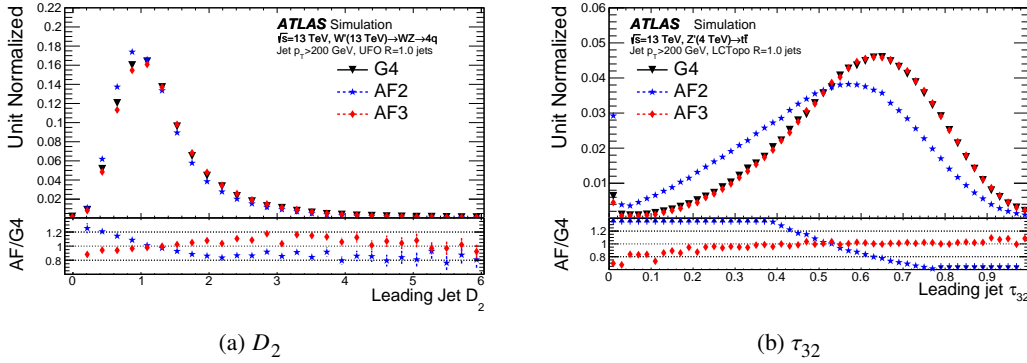


Figure 7.5: Substructure observables (a) D_2 for p_T -leading UFO $R = 1.0$ jets from simulated resonant $W'(13 \text{ TeV}) \rightarrow WZ \rightarrow q\bar{q}$ decays and (b) τ_{32} for p_T -leading LCTopo $R = 1.0$ jets from simulated resonant $Z'(4 \text{ TeV}) \rightarrow t\bar{t}$ decays. The reconstructed jets are required to have $p_T > 200$ GeV. [134]

is shown for p_T -leading UFO $R = 1.0$ jets from resonant $W'(13 \text{ TeV}) \rightarrow WZ \rightarrow q\bar{q}$ decays and the τ_{32} observable is shown for p_T -leading LCTopo $R = 1.0$ jets from resonant $Z'(4 \text{ TeV}) \rightarrow t\bar{t}$ decays. In the case of D_2 , a significant slope is observed in the AFII/GEANT4 ratio, while ATLFast3 improves the modelling significantly, particularly at lower values of D_2 . The modelling of τ_{32} is very poor with AFII and significantly improved with ATLFast3, reaching agreement with GEANT4 within 20%. Note that improvements in jet substructure are generally expected to be more pronounced for LCTopo jets and less so for UFO jets, as the latter use

additional information from the Inner Detector.

The improvements in the simulation of the internal structure of high- p_T jets with ATLFast3 are primarily due to the weighted hit model discussed in Section 4.4, which has been shown to significantly improve the modelling of cluster formation in the calorimeter, and the usage of FastCaloGAN for medium-energy hadrons. The precise simulation of calorimeter clusters also results in a significantly better modelled jet mass. Figure 7.6 shows the reconstructed mass of p_T -leading UFO $R = 1.0$ jets from $Z'(4 \text{ TeV}) \rightarrow t\bar{t}$ events. AFII significantly underestimates




Figure 7.6: Mass of the p_T -leading UFO $R = 1.0$ jet reconstructed from $Z'(4 \text{ TeV}) \rightarrow t\bar{t}$ events. The reconstructed jets are required to have $p_T > 200 \text{ GeV}$. [134]

the jet mass and the width of the two peaks at the W boson and top quark masses. ATLFast3, on the other hand, simulates the jet mass with much better accuracy and agrees with GEANT4 within a maximum deviation of 20% in the fast over full simulation ratio.

The presented improvements in the accuracy of the simulation of the complex substructure of boosted jets and their mass are crucial to ensure that most physics analyses searching for heavy resonances from BSM theories will be able to fully replace the GEANT4 simulation with ATLFast3, so that the usage targets for fast simulation imposed by the collaboration can be met.

7.2.3 Hadronic Tau Lepton Decays

The reconstruction of hadronically decaying τ -leptons in ATLAS is seeded by the presence of a jet and uses the subsequent τ decay products of one or three charged hadrons and additional neutral particles for the reconstruction [207–209]. The reconstructed hadrons are used to classify the τ decay modes, which are typically labelled $YpXn$, where Y refers to the number of

reconstructed charged particles and X to the additional number of neutral particles produced in the decay.

Figure 7.7 shows the hadronic τ -lepton decay modes reconstructed in a $Z^*/\gamma^* \rightarrow \tau\tau$ Drell-Yan sample with an off-shell mass of 2-2.25 TeV for τ candidates spatially matched to generator level τ -leptons (*true* τ -leptons) and for candidates incorrectly identified as τ -leptons (*fake* τ -leptons). With the exception of the 3pXn category in the case of true τ -leptons, the




Figure 7.7: Hadronic τ -lepton decay modes reconstructed in a $Z^*/\gamma^* \rightarrow \tau\tau$ Drell-Yan sample with an off-shell mass of 2-2.25 TeV for (a) true and (b) fake τ -leptons. Decays with one (three) charged particle tracks are denoted by 1p (3p), while $X \in [1, 2, 3]$ refers to the number of neutral particles. [134]

modelling of the decay modes is noticeably improved with ATLFAST3, and agrees to better than 10% (5%) for true (fake) τ -leptons. The improvement is particularly pronounced for categories with many neutral particles, due to the improved simulation of lateral correlations of calorimeter clusters in ATLFAST3.

Similar to the simulation of jet substructure in boosted decays, the precise modelling of the inner structure of τ -jets is crucial in obtaining an accurate simulation of τ -leptons. Figure 7.8 shows the simulated number of clusters in 1-prong hadronic τ -lepton candidates for true and fake τ -leptons reconstructed in a $Z^*/\gamma^* \rightarrow \tau\tau$ Drell-Yan sample with an off-shell mass of 2-2.25 TeV. In both cases, the number of clusters is significantly underestimated by AFII, but is nearly perfectly modelled in the ATLFAST3 simulation.

Finally, Figure 7.9 shows the reconstructed visible di-tau invariant mass from hadronically decaying τ -leptons in the same $Z^*/\gamma^* \rightarrow \tau\tau$ sample. Both simulators model the invariant di-tau mass reasonably well and are compatible with the GEANT4 simulation within 10% for

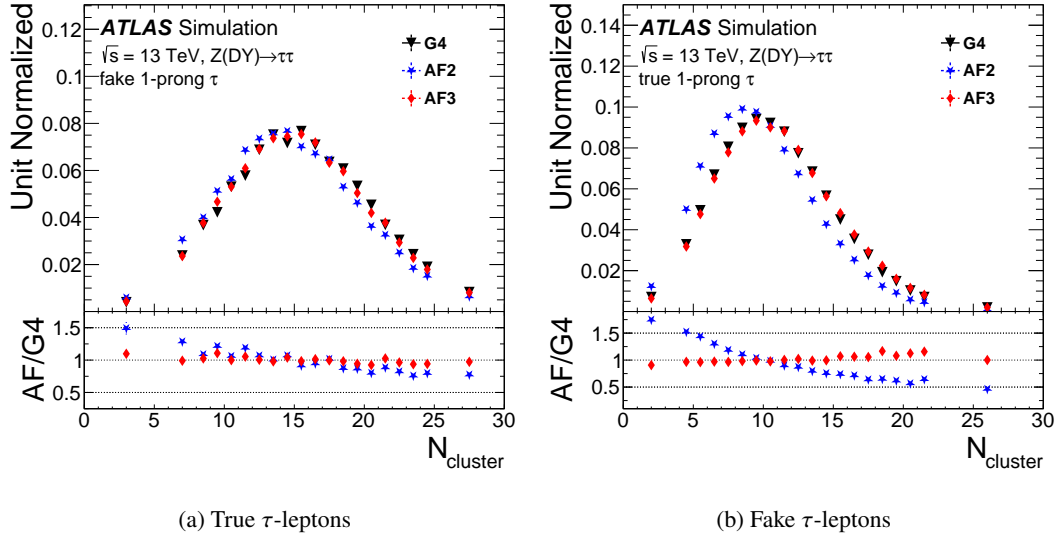


Figure 7.8: Number of clusters in 1-prong hadronic τ -lepton candidates for (a) true and (b) fake τ -leptons reconstructed in a $Z^*/\gamma^* \rightarrow \tau\tau$ Drell-Yan sample with an off-shell mass of 2-2.25 TeV. [134]




Figure 7.9: Visible invariant di-tau mass reconstructed from hadronically decaying τ -leptons in a $Z^*/\gamma^* \rightarrow \tau\tau$ Drell-Yan sample with an off-shell mass of 2-2.25 TeV. [134]

most of the $m_{\tau\tau}$ spectrum. AFII slightly underestimates the width of the distribution, while the opposite is true for ATLFast3.

7.2.4 Muons

Muons in ATLAS are reconstructed from tracks in the Muon Spectrometer and signals from the Inner Detector [210], and are simulated exclusively with GEANT4 within ATLFast3 (see Chapter 4). Therefore, no differences in the simulation of prompt muon observables are expected with respect to the full simulation. However, secondary particles produced in hadronic showers may escape the calorimeter system and enter the Muon Spectrometer. These muon punch-through (MPT) particles create hits in the MS and are required to be modelled for an accurate simulation. While AFII does not take MPT particles into account, ATLFast3 parametrizes the punch-through probability of particles depending on their incident momentum and direction, generates secondary particles accordingly and passes these to GEANT4 for simulation.

Figure 7.10 shows the number of muon segments created from EMPFlow $R = 0.4$ jets reconstructed in $Z'(4\text{ TeV}) \rightarrow t\bar{t}$ events. Both fast simulators reproduce the distribution for




Figure 7.10: Number of muon segments created from EMPFlow $R = 0.4$ jets reconstructed in $Z'(4\text{ TeV}) \rightarrow t\bar{t}$ events. Jets are required to have $p_T > 20\text{ GeV}$. [134]

events with up to four muon segments. As expected, AFII significantly underestimates the number of muon segments for larger values. The current punch-through simulation of ATLFast3 effectively recovers the segments, but tends to produce slightly too many. Future updates are expected to further improve the simulation of MPT particles.

Tuning of Electromagnetic Shower Properties to Data

FINAL states with photons and electrons are an important component of the ATLAS physics programme. Phenomena with prompt photons, that is, photons that originate from the hard scatter interaction and are not the result of subsequent hadron decays, play a significant role for probing certain aspects of QCD, such as the inner structure of protons [211]. Moreover, photon and electron final states have been proven fundamental for the discovery of the Higgs boson [38, 39] and are crucial for precision measurements and BSM searches.

In hadronic collisions, such as at the LHC, the identification of prompt leptons is a challenging task. Most of the reconstructed photon and electron candidates originate from decays of neutral hadrons that are created within jets as well as hadrons that deposit a significant fraction of their energy in the electromagnetic calorimeter and mimic the signature of photons and electrons. Additionally, reconstructed electrons are required to be separated from converted photons that decay into e^+e^- pairs before reaching the calorimeter system.

For the identification of photons and electrons, ATLAS heavily relies on calorimetric variables that describe the lateral and longitudinal properties of electromagnetic showers and provide good discriminatory power against fake signatures from QCD jets and decays of neutral hadrons in jets. As a result, a precise simulation of the electromagnetic shower development is crucial to ensure the correct identification of photon and electron candidates. However, the GEANT4 simulation of the ATLAS detector is known to imperfectly reproduce the properties of electromagnetic showers observed in data. In fact, the identification efficiency of photons, which is important to predict physics processes, is only modelled with $\mathcal{O}(2 - 5\%)$ accuracy and is mostly limited by the imperfect description of EM shower variables [133].

In general, longitudinal shower shapes appear better modelled in comparison to the transverse shower development. While these differences mainly exhibit in shifts in the central values of the

shower shape distributions, their origin remains unknown. Traditionally, these discrepancies are corrected using *ad hoc* methods through the parametrization and application of shifts to the corresponding distributions after simulation, a procedure known as *fudging* [133].

In this chapter, it is shown that FASTCALOSIMV2 can be modified in a way that most of the photon shower shapes observed in data are accurately reproduced by the simulation. Two separate models are developed, one targeting shower shape variables specific to the EM middle layer of the calorimeter and another which targets shower properties as measured in the finely segmented strip layer of the calorimeter system. In Section 8.1, the targeted shower shapes and their usage in the identification algorithms employed by ATLAS are discussed. In Section 8.2, the selection of isolated photon candidates that are used for the tuning procedure is presented. In Section 8.3, the development of the two underlying models are described for photons with $p = 65$ GeV and $0.20 < |\eta| < 0.25$, corresponding to a single parametrization point in FASTCALOSIMV2. The extension of the tune to the full detector region is discussed in Section 8.4. Finally, the application of the tune to electron candidates is discussed in Section 8.5.

8.1 Electromagnetic Shower Shapes

In order to identify prompt photons, ATLAS employs a set of rectangular cuts on calorimetric variables that characterize the lateral and longitudinal EM shower development and help to distinguish energy deposits resulting from jets and decays of neutral hadrons [133]. Typically, two distinct working points are employed for the identification of photon candidates: a *loose* identification only uses information based on the distribution of energy deposits within the middle layer of the calorimeter and the leakage of energy into the hadronic calorimeter, while the *tight* identification additionally uses the EM shower properties as measured in the finely segmented strip layer of the calorimeter. Similarly, for electrons, a likelihood discriminant is built based on the same shower shape variables, which is designed to reject light-flavour jets, photon conversions and non-prompt electrons [201]. For the development of the models to tune electromagnetic shower shape properties to data, the most relevant variables in both the middle and strip layers of the calorimeter are considered. The definitions and a sketch of the shower shape variables targeted in the development of the data tuning models are shown in Figure 8.1.

For the EM middle layer, the variables $w_{\eta 2}$, R_{η} and R_{ϕ} are targeted in the tune. $w_{\eta 2}$ is a measure of the proper width of the showers. R_{η} is defined as an energy ratio and gives a measure of the showers' width in η direction. As hadronic showers tend to be broader, energy deposits originating from fake jets or from the decay of natural hadrons have lower R_{η} and higher $w_{\eta 2}$ values. Similarly, R_{ϕ} describes the shower width in the ϕ direction and is particularly useful to discriminate between converted and unconverted photons, as the magnetic field separates the resulting e^+e^- pair from the conversion in ϕ direction. Figure 8.2 shows the comparison of the middle layer shower shape variables as simulated with GEANT4 and FASTCALOSIMV2 as well as

Category	Description	Name
EM middle layer	Ratio of the sum of the energies of the cells contained in a 3×7 rectangle (measured in cell units) to the sum of the cell energies in a 7×7 rectangle, both centred around the most energetic cell	R_η
	Lateral shower width, $\sqrt{\sum E_i \eta_i^2 / \sum E_i - (\sum E_i \eta_i / \sum E_i)^2}$, where E_i is the energy and η_i is the pseudorapidity of cell i and the sum is calculated within a window of 3×5 cells	$w_{\eta 2}$
	Ratio of the sum of the energies of the cells contained in a 3×3 rectangle (measured in cell units) to the sum of the cell energies in a 3×7 rectangle, both centred around the most energetic cell	R_ϕ
EM strip layer	Lateral shower width, $\sqrt{\sum E_i (i - i_{\max}^2) / \sum E_i}$, where i runs over all cells in a window of 3 cells around the highest-energy cell	$w_{\eta 1} / w_{s3}$
	Total lateral shower width, $\sqrt{\sum E_i (i - i_{\max})^2 / \sum E_i}$, where i runs over all cells in a window of $\Delta\eta \approx 0.0625$ and i_{\max} is the index of the highest-energy cell	w_{stot}
	Energy outside the core of the three central strips but within seven strips divided by energy within the three central strips	F_{side}
	Ratio of the energy difference between the maximum energy deposit and the energy deposit in the secondary maximum in the cluster to the sum of these energies	E_{ratio}

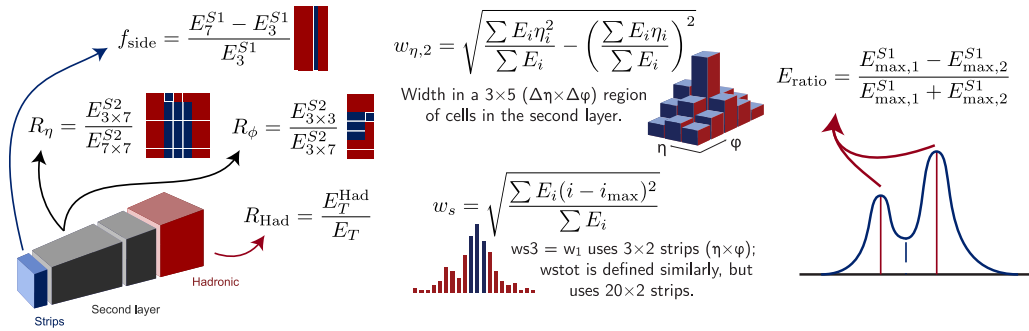


Figure 8.1: Pictorial representations and definitions of the various variables describing the shape of showers in the calorimeter that are employed in the identification of electromagnetic showers and targeted in the development of the data tune. Modified from Ref. [212, 213]

from photons collected in data for 65 GeV photons in the $0.20 < |\eta| < 0.25$ region. The exact selection used to collect the photons in data is discussed in Section 8.2.

Based on the energies deposited in the strip layer of the calorimeter system, the variables $w_{\eta 1}$, F_{side} , E_{ratio} and w_{stot} are targeted in the tune. F_{side} measures the fraction of energy deposited away from the shower centre. Therefore, showers initiated by prompt photons tend to have lower F_{side} values compared to background processes. Both $w_{\eta 1}$ and w_{stot} measure the width of the shower, where $w_{\eta 1}$ measures the width based on the central strip and its neighbours and w_{stot} employs a much larger number of strips, typically in the order of 20 strips. Neutral mesons that decay into photon pairs before reaching the calorimeter system, e.g. $\pi^0 \rightarrow \gamma\gamma$ and $\rho \rightarrow \gamma\gamma$, leave a typical double-peak structure in the calorimeter. In order to distinguish the signature from prompt photons, the variable E_{ratio} gives a measure on how evenly the energy is distributed between the double peak structure. As such, typical values for E_{ratio} for prompt photons are very close to one, while photons originating from neutral meson decays have significantly lower values. Figure 8.3 shows the comparison of the strip layer shower shape variables as simulated with GEANT4 and FASTCALOSIMV2 as well as from photons collected in data for 65 GeV photons in the $0.20 < |\eta| < 0.25$ region.

In the case of both, strip and middle layer shower shapes, the FASTCALOSIMV2 simulation precisely reproduces the properties of the photon showers as simulated with GEANT4, with only minor differences in the tails of some shower shapes. As expected, in all cases with exception to E_{ratio} , a clear shift in the simulated distributions with respect to data is observed.

8.2 Selection of Photon Candidates

In order to develop a tuning model that can intrinsically simulate the shower shapes of photons observed in data, a clean sample of photons is required to be collected which is used as a reference for the model. A high purity fraction of prompt over non-prompt photons is needed to avoid any significant bias in the tuning procedure.

The tuning model described in the following is derived from a pp collision dataset of 58.5 pb^{-1} collected in 2018 during Run 2 with the ATLAS detector at $\sqrt{s} = 13 \text{ TeV}$. The mean number of interactions per bunch crossing was $\langle \mu \rangle = 36.1$. A Good Run List (GRL) selection is applied to only select events recorded from good luminosity blocks, in which a proper functioning of the detector can be ensured. Photons are collected by an OR of all available single photon triggers from the 2018 trigger menu [214], most of which are expected to stem from γ +jets processes.

For the nominal and tuned version of the simulation, single photons are simulated using FASTCALOSIMV2. The photons are generated at the origin of the detector and pile-up is modelled by overlaying signals from simulated inelastic pp events with the hard-scattering event [158], resembling the conditions during the Run 2 data taking period. An average number of pile-up interactions per pp bunch crossing of 38 was used, reasonably similar to the

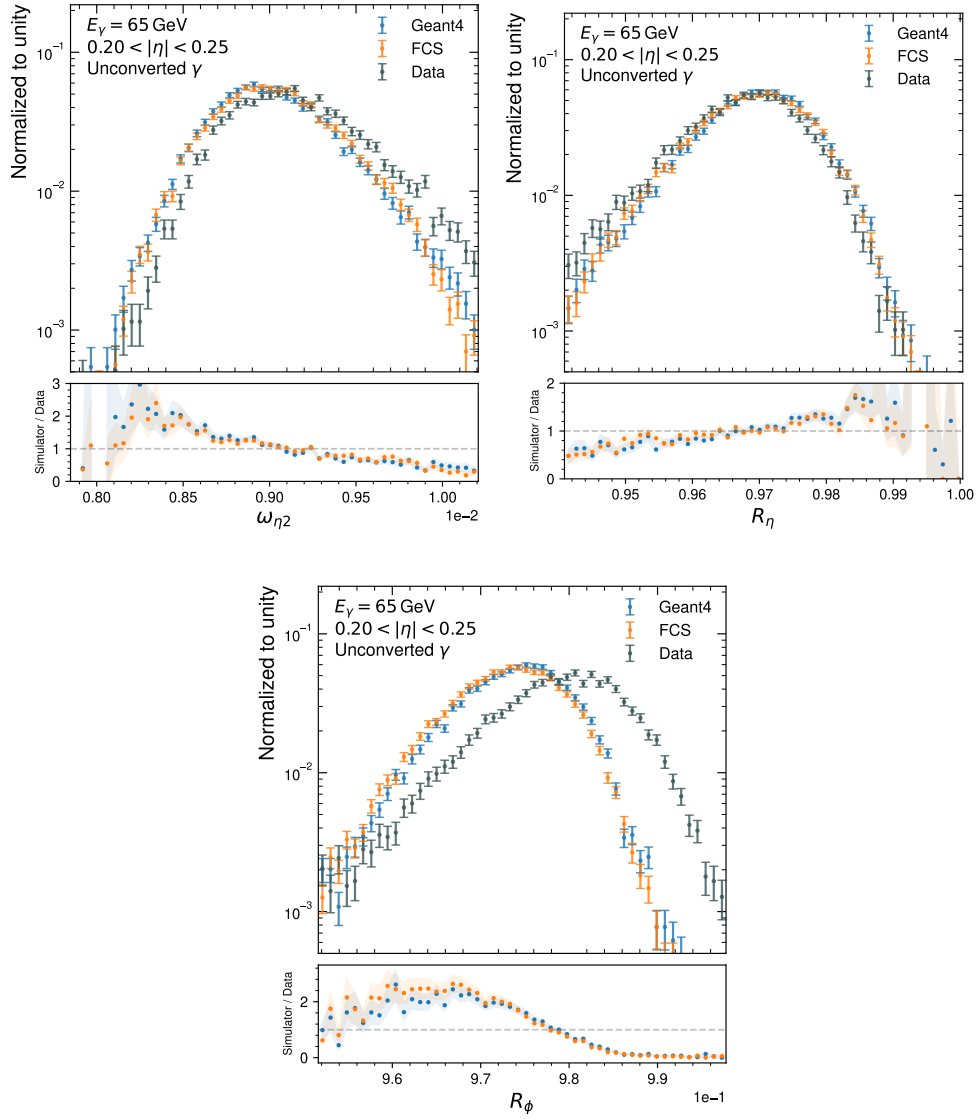


Figure 8.2: Middle layer shower shapes as simulated with GEANT4 and FASTCALOSIMV2 in comparison with the shapes observed in data for unconverted 65 GeV photons in the $0.20 < |\eta| < 0.25$ region. The lower panel shows the ratio between the simulation and the data. Single particles are simulated at the origin and pile-up interactions are overlaid to mimic real data taking conditions.

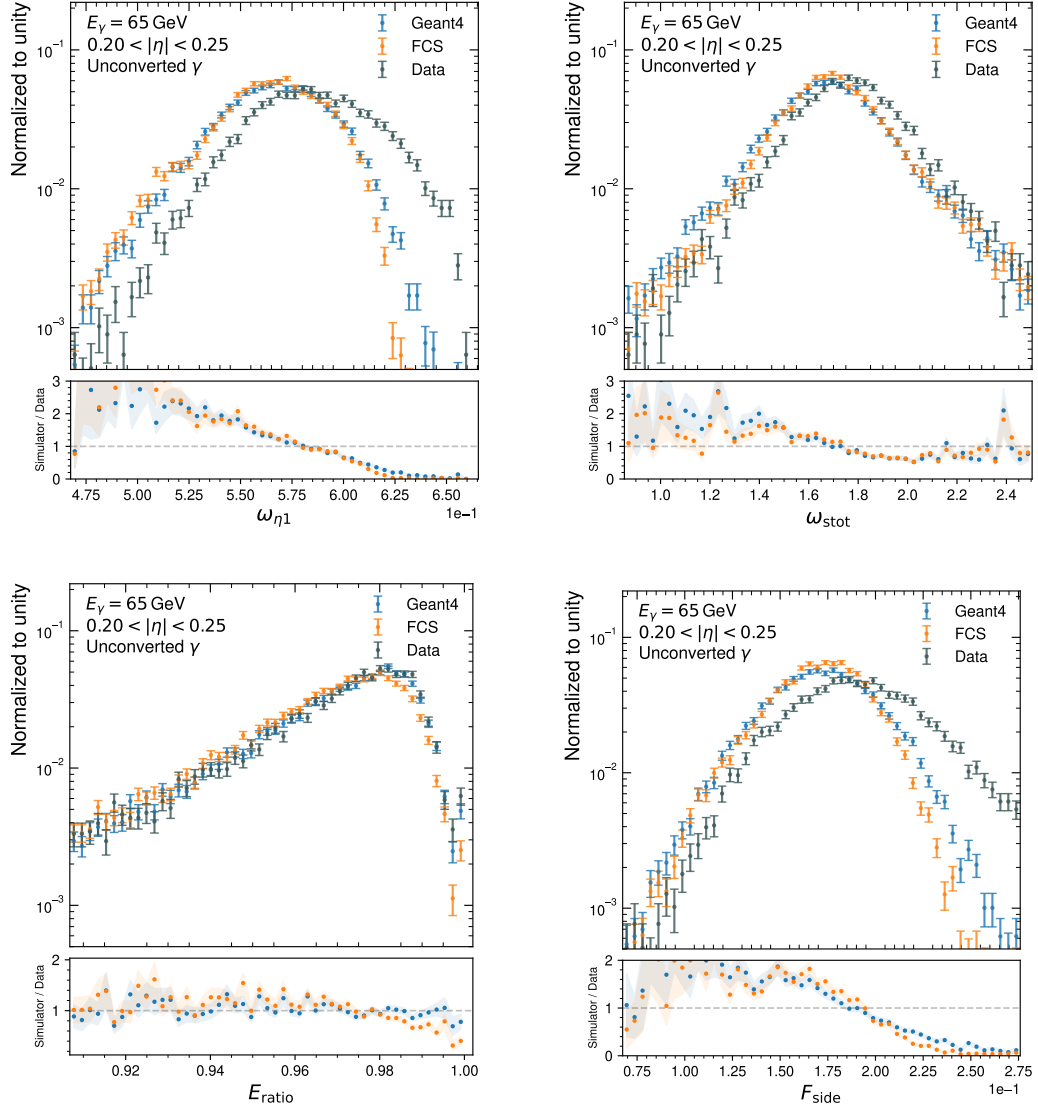


Figure 8.3: Strip layer shower shapes as simulated with GEANT4 and FASTCALOSIMV2 in comparison with the shapes observed in data for converted 65 GeV photons in the $0.20 < |\eta| < 0.25$ region. The lower panel shows the ratio between the simulation and the data. Single particles are simulated at the origin and pile-up interactions are overlaid to mimic real data taking conditions.

pile-up distribution recorded by the ATLAS experiment during 2018. The presence of pile-up produces low- E_T activity in the detector, which results in additional energy deposits in the EM calorimeter. Hence, with a larger number of overlaid pp events, the photon shower shapes are smeared out, while the means of the shower shape distributions remain fairly intact. The pile-up samples are generated using PYTHIA 8.186 using the NNPDF23 set of PDFs [195] and the A3 tune [198].

To avoid electrons resulting from photon conversions from interfering in the tuning procedure, only photons classified as unconverted by the reconstruction algorithm [200] are selected. The purity of the sample is further increased by requiring the photon candidates to pass the `Tight` identification criteria and the `FixedCutLoose` [200] isolation working point.

The purity of the photon sample collected in data is observed to increase with larger photon momenta and is expected to be very high ($> 90\%$) for photons with energies above 100 GeV [133]. Any remaining falsely classified non-prompt photons may slightly broaden the photon shower shape distributions, but are not expected to have a significant impact on their mean.

8.3 Tuning Model

The development of an appropriate tuning model that achieves the reproduction of the shower shapes observed in data is a non-trivial task for various reasons. First, the origin of the observed discrepancies remains unknown until this point and as such, no *a priori* knowledge is available which may aid the development of an appropriate model. Secondly, the majority of the shower shapes of interest are highly correlated. As an example, Figure 8.4 shows two-dimensional distributions of the middle layer shower shapes R_η , R_ϕ and $w_{\eta 2}$ for 65 GeV photons in the $0.20 < |\eta| < 0.25$ region, all showing large correlations. Hence, any alterations in the simulation procedure which improves the modelling of a subset of the shower shapes, may negatively impact the simulation of other variables. Lastly, the tune is required to perform well across a large region of phase space, that is, for a large range of photon momenta and for directions that are relevant for most physics analyses. In this section, the development of two independent models for the middle and strip layers of the calorimeter is presented, based on a single parametrization point used in FASTCALOSIMV2, that is, for photons with $p = 65$ GeV in the $0.20 < |\eta| < 0.25$. For this purpose, the photon candidates selected in the simulation and data are required to have $60 \text{ GeV} < E_T < 70 \text{ GeV}$. The models are then extended to account for the full phase space of interest in Section 8.4.

On a fundamental level, the reconstructed shower shapes are dependent on the geometrical distribution of the hits in the various layers of the calorimeter. In the case of photons and electrons, the assigned hit energy in FASTCALOSIMV2 is independent of the position of the hit (see Section 4.4), such that the hit positions in the $\eta - \phi$ plane constitute the only degrees of freedom that will impact the simulated properties of the electromagnetic shower




Figure 8.4: Correlations between the middle layer shower shapes R_η , R_ϕ and $w_{\eta 2}$ for 65 GeV photons in the $0.20 < |\eta| < 0.25$ region. Single photons are generated at the origin and simulated with GEANT4. No pile-up interactions are overlaid that may smear out the correlations between the shower shapes. In all cases, 50×50 bins are used create a two-dimensional colour map, where bin borders are smoothed using a Gaussian interpolation. The Pearson correlation factor ρ is displayed in all cases.

development. By design, the geometric distribution of the hit positions can be arbitrarily altered in FASTCALOSIMV2, while a similar procedure is not possible in GEANT4. As a result, a powerful procedure to tune the fast simulation to data, is to transform the hit positions within FASTCALOSIMV2 in a way that the shower shapes observed in data are reproduced.

In order to accomplish the latter, the azimuthal angle ϕ of the hit position is transformed into a new angle ϕ' according to


$$\phi \rightarrow \phi' = \phi_c + \phi_s \cdot \Delta\phi, \quad (8.1)$$

where ϕ_c is the azimuthal angle of the shower centre, $\Delta\phi$ the projected distance to the shower centre in ϕ direction and ϕ_s an arbitrary model parameter. The computation of the shower centres is based on the trajectory of the incident particle and extrapolated to the various layers of the calorimeter (see Chapter 6). The transformed distance to the shower centre in ϕ is then given by $\Delta\phi' = \phi_s \cdot \Delta\phi$, such that ϕ_s is a measure of the scale of transformation. Analogously, the identical transformation can be performed on the η of the hit positions. While positive values of ϕ_s or η_s will shift hits away from the shower centre, negative values will cause an inwards movement of the hits.

8.3.1 Two-dimensional $\eta - \phi$ scaling model for middle layer shower shapes


As the middle layer of the calorimeter is finely grained in both η and ϕ directions, shower shapes are expected to be sensitive to shifts around the shower centre in both directions. The most simple transformation in this case, is to apply constant scaling factors ϕ_s and η_s . Figure 8.5 shows the hit distributions in the $\Delta\eta - \Delta\phi$ space after application of a constant scaling in η and ϕ direction as well as for a radial scaling, in which case the identical scaling is applied in both directions. According to expectations, the hits are shifted further away from the shower centre in η direction when setting a large value for η_s . In the opposite case of setting a large value for ϕ_s , the hit density is stretched in ϕ direction. Lastly, setting both parameters to equally large values leads to hits being moved radially outwards from the shower centre such that the hit density decreases within the shower core and increases towards the outer regions.

In order to find suitable model parameters η_s and ϕ_s , a two-dimensional phase space scan in the region $\eta_s \times \phi_s = [0.9, 1.1] \times [0.8, 1.1]$ with step size $\Delta\phi_s = 0.1$ and $\Delta\eta_s = 0.01$ is performed using a standalone FASTCALOSIMV2 shower simulation that does not rely on the full simulation and reconstruction chain employed in ATLAS. In each point of the $\eta_s - \phi_s$ phase space, the mean of the shower shapes are evaluated, and intermediate values are interpolated using natural neighbour interpolation [215]. Employing latter interpolation, contour values corresponding to the desired means observed in data are extracted for each of the shower shapes. The two-dimensional phase space scan in the $\eta - \phi$ scaling model is shown in Figure 8.6 for w_{η^2} , R_η and R_ϕ . As expected, w_{η^2} and R_η show a strong sensitivity for a scaling in η and a low sensitivity in ϕ scaling, while the opposite is the case for R_ϕ . The extracted target contour



(a) η scaling ($\eta_s = 10$, $\phi_s = 1$)

(b) ϕ scaling ($\eta_s = 1$, $\phi_s = 10$)



(c) Radial scaling ($\eta_s = \phi_s = 20$)

Figure 8.5: Hit distributions in the $\Delta\eta - \Delta\phi$ space after application of a (a) constant scaling in η and (b) ϕ direction as well as for a (c) radial scaling. The colour scheme encodes the hit density, while the white dots show the positions of the individual hits. [213]




Figure 8.6: Two-dimensional phase space scan in the linear scaling model for the middle layer shower shapes (a) $w_{\eta 2}$, (b) R_{η} and (c) R_{ϕ} . The white crosses show the points in the scaling space at which the mean values of the shower shapes are evaluated. Natural neighbour interpolation is employed to interpolate values in between in order to extract contour lines, which correspond to the target shower shape means. The shaded regions show the statistical uncertainties of the unscaled means of the corresponding shower shapes. The red circles indicate the intersections between the target contour lines. [213]

lines for w_{η^2} and R_{η} show an almost parallel behaviour, but are close enough to each other, indicating that the same set of parameters is suitable to tune both variables. In both cases, a clear intersection with the contour line extracted for R_{ϕ} is visible, i.e. all target contour lines cross at one single point within statistical uncertainties. A two-dimensional constant scaling model is therefore expected to be able to fully describe the differences observed in the central values of the respective shower shapes between simulation and data. The optimal parameters chosen to validate the model (see Section 8.3.3) correspond to the intersection of the target contours given by w_{η^2} and R_{ϕ} and are $\phi_s = 0.861$ and $\eta_s = 1.028$.

8.3.2 Cubic η scaling model for strip layer shower shapes

As the strip layer of the ATLAS calorimeter consists of finely segmented strips, the shower shapes are expected to be very sensitive to hit displacements in η direction, while very little sensitivity is expected for displacements in ϕ . As a result, the ϕ_s scaling applied for the development of the data tune targeting middle layer variables described in Section 8.3.1 becomes inapplicable. However, a constant η_s factor is insufficient to correctly model all the target shower shapes. Under the assumption that there exists a functional form $\eta_s(\Delta\eta)$ that can describe the differences of electromagnetic shower properties in simulation compared to data, one can expand η_s as series in the distance to the shower centre in η direction:

$$\eta_s = 1 + a_0 + a_1 \cdot \frac{\Delta\eta}{\bar{\Delta\eta}} + a_2 \cdot \left(\frac{\Delta\eta}{\bar{\Delta\eta}}\right)^2 + \dots + a_n \cdot \left(\frac{\Delta\eta}{\bar{\Delta\eta}}\right)^n, \quad (8.2)$$

where a_n are arbitrary model parameters and $\bar{\Delta\eta} \approx 0.0039$ is the mean distance of the hits to the shower centre in η direction. For $n = 1$, η_s is an affine function of $\Delta\eta$ and is described by two free parameters a_0 and a_1 , such that a similar two-dimensional phase space scan as described in Section 8.3.1 can be performed. While such a linear model might describe some targeted variables correctly, it is deemed to insufficiently model the differences observed in data for all targeted shower shapes.

The transition from a two-dimensional to higher dimensional models is however non-trivial, as graphical approaches to study the model phase space become impractical and are significantly more expensive to evaluate in regard to the required CPU consumption. Models with $n > 1$ may therefore offer improved performance, but require more involved optimization algorithms such as derivative-free black box optimization [216] methods to analyse.

The time required to evaluate the shower shape means for a single point in the model parameter space with the standalone FASTCALOSIMV2 simulation is in the order of a few minutes for approximately 1000 events. Without parallelization, the total evaluation time of the model scales with the number of free parameters and the granularity of the scanned phase space. Under the realistic assumption of an average three- to five-minute evaluation time of a single set of model parameters, the time required to evaluate the two-dimensional $\eta - \phi$ scaling

model described in Section 8.3.1 using the shown grid division is already in the order of 1-2 days. With an increasing number of model parameters, the evaluation time scales accordingly, such that an effective model evaluation is not possible anymore. Furthermore, there is no a priori reason to believe that the same set of parameters performs equally well across all photon energies and directions, and as such, the model parameters are required to be optimized for each relevant parametrization point employed in FASTCALOSIMV2 (see Section 8.4).

In order to allow for a large-scale optimization procedure, a dedicated and robust parallelization infrastructure is developed to ensure the optimal use of resources. ORION [217], a framework for large-scale asynchronous distributed black-box optimization, is employed in order to evaluate these higher dimensional models. A simplified sketch of the optimization set-up is shown in Figure 8.7. A user launches the optimization procedure from a host machine by defining the

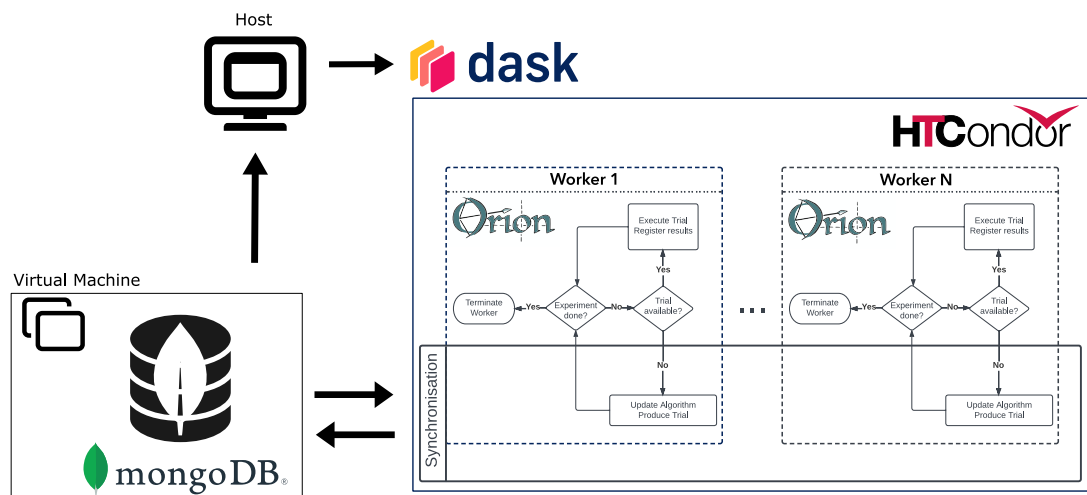


Figure 8.7: Simplified sketch of the technical set-up of the optimization procedure used to evaluate the tuning model. An arbitrary number of ORION workers, each running on separate machines as part of a computer cluster, evaluate model parameters in parallel. Trials are suggested by ORION based on their common shared history stored in a database. Results can be retrieved by the host at anytime during the execution via access to the database.

properties of the experiments, such as the desired search space and optimization algorithm, that ought to be run as well as the number of workers that should be deployed in parallel. With DASK [218] as an executor backend, one machine per ORION worker is spawned from a computer cluster managed by HTCONDOR [219]. ORION then suggests trials, that is, a specific set of model parameters, to the individual workers, which subsequently evaluate the models based on the suggested trials. Once a worker has concluded the model evaluation for a suggested trial, the result is stored in a common database. The database is hosted on a virtual machine to which all individual workers have access. Hence, while all workers run on individual machines, the results of all prior trials are known to ORION, such that the decision of what trials come next

is based on the results of the previously evaluated model parameters. In principle, the set-up allows for arbitrary horizontal scaling and is only limited by the number of available machines in the respective computer cluster and the computational resources on the virtual machine that hosts the database. It is noteworthy that the time required to evaluate the two-dimensional $\eta - \phi$ scaling model described in Section 8.3.1 for a single FASTCALOSIMV2 parametrization point with this set-up is effectively reduced from 1-2 days to less than 30 minutes.

With ORION as a black-box optimizer, a random search in higher dimensional strip layer phase spaces is performed with the objective of minimizing

$$\chi^2 = \frac{1}{N} \sum_{i \in S} \left(\frac{\bar{X}_i - \bar{X}'_i}{\sigma_i} \right)^2, \quad (8.3)$$

where the sum runs over the target shower shapes $S = \{w_{\eta 1}, w_{\text{stot}}, E_{\text{ratio}}, F_{\text{side}}\}$, \bar{X}_i is the mean resulting from the black-box evaluation given a set of arbitrary parameters a_n , \bar{X}'_i is the target mean extracted from data and σ_i the standard deviation of the respective shower shape in the nominal simulation. The normalization factor $N = 4$ is chosen as the number of targeted variables, such that on average the black-box evaluated mean and the target mean have a distance of 1σ to each other for a value of $\chi^2 = 1$. The baseline error with the nominal simulation is at $\chi^2 \approx 0.165$. While several variations in the model dimension have been studied, a χ^2 value in the sub-percent region is reached with four free parameters ($n = 3$).

More sophisticated optimization algorithms other than a random search were studied, but were not found to provide significant advantages. An initial phase space scan in a relatively large search region of $a_0 \in [-3, 3]$ and $a_k \in [-1, 1]$ with $k \in \{1, 2, 3\}$ is performed. Figure 8.8 shows the development of the χ^2 error in dependence of the trials ordered by the time they are suggested by ORION. The red line indicates the error of the current best trial at a given point in time. While most of the evaluations in the phase space yield large errors in the range between 1 and 30, the minimal error drops significantly after 293 trials to $\chi^2 \approx 0.0023$. The corresponding optimal parameters are given by $a_n = \{0.2486, 0.2944, -0.3486, 0.1633\}$.

These results show that with higher dimensional η_s scaling models, an effective tuning of electromagnetic shower properties to data is feasible. In fact, it has been shown that a reduction in the size of the search space around regions of interest can decrease the model error even further. The resulting functional dependence of the scaling factor η_s on the distance to the shower centre as well as the transformed distance $\Delta\eta^*$ as a function of the original distance projected onto the η direction is shown in Figure 8.9. Hits in the shower core ($\Delta\eta = 0$) are moved outwards with a constant scaling $\eta_s = 1 + a_0 \approx 1.2486$. The scaling first increases moderately until $\Delta\eta \approx 0.008$, at which point the scaling factor begins to increase strongly. The scaling factor diverges for $\Delta\eta \rightarrow \infty$, which however, has no physical implications as the majority of hits do not receive such large scaling factors as shown by the shaded red hit distribution in Figure 8.9.




Figure 8.8: Model error χ^2 as a function of trials ordered by suggested time for the data tune targeting the strip layer variables $w_{\eta 1}$, w_{stot} , E_{ratio} and F_{side} . The phase space is scanned with a random search algorithm in the region $a_0 \in [-3, 3]$ and $a_k \in [-1, 1]$ with $k \in \{1, 2, 3\}$. The blue dots show the individual trials, while the red line indicates the error of the current best trial. The orange box shows the optimal set of parameters found after 350 trials. [213]

8.3.3 Results

The phase space analysis performed in Section 8.3 suggests that both models should be able to reproduce the means of the shower shapes observed in data correctly. However, the analysis was performed using the FASTCALOSIMV2 standalone implementation, and it is crucial to ensure that the model accomplishes similar performance when propagating through the full simulation and reconstruction chain employed by ATLAS. While not expected, it is also necessary to make certain that the model does not significantly distort the shape of the distributions beyond a simple shift in the central values.

Figure 8.10 shows the comparison between the data, the nominal and the tuned simulation for the shower shape variables $w_{\eta 2}$, R_η and R_ϕ . In the case of $w_{\eta 2}$, the mean of the nominal simulation is shifted by 1.26% in comparison to data. The two-dimensional $\eta - \phi$ scaling model entirely eliminates this shift and is furthermore compatible with the shape observed in data within statistical uncertainties. Similarly, a small shift seen for R_η is reduced from 0.16% to 0.06%, without impacting the shape of the distribution. A large shift to data is seen for the nominal simulation in the case of R_ϕ , which is effectively reduced from 0.55% to 0.15%. The difference in the RMS to data is slightly increased by 6% from 20% to 26% and the tails of the distribution show minor mismodelling.




Figure 8.9: (a) Scaling factor η_s as a function of the distance to the shower centre $\Delta\eta$ in η direction and (b) new distance to the shower centre $\Delta\eta^*$ as a function of the original distance to the shower centre $\Delta\eta$. The blue dashed line indicates the functional form of the nominal simulation, while the orange line shows the dependence for the tuned simulation. The shaded red area displays the normalized shape of the hit distribution in $\Delta\eta$. For illustration purposes, the distribution is reflected across the $\Delta\eta$ axis in case (b). The vertical dashed lines indicate the positions of the strip edges, assuming that the shower centre is situated in the centre of a strip. [213]

Figure 8.11 shows the validation of the cubic η scaling model for the strip layer variables $w_{\eta 1}$, w_{stot} , E_{ratio} and F_{side} . The largest shift between the data and the nominal simulation is observed for F_{side} and is effectively reduced from 10.4% to 0.4%, followed by a shift of 3.8% in the central value of w_{stot} , which is decreased to 0.8%. Similarly, the difference in $w_{\eta 1}$ is reduced from 3.5% to 0.4%. In all three cases, the shape of the shower shape variables generated using the tuned simulation shows excellent agreement with data, with very small differences in the high-end tail of the $w_{\eta 1}$ distribution. In the case of E_{ratio} , a small degradation in the relative mean difference to data is observed from 0.25% to 0.4%, but still remaining well under the percent level. A summary of the relative mean and RMS differences to data for the nominal and the tuned simulation for all targeted shower shape variables is shown in Table 8.1.

8.4 Extension of Model Coverage

The two different models developed for the strip layer in Section 8.3.2 and the middle layer in Section 8.3.1 are found to work well in describing most of the shower shapes within a single

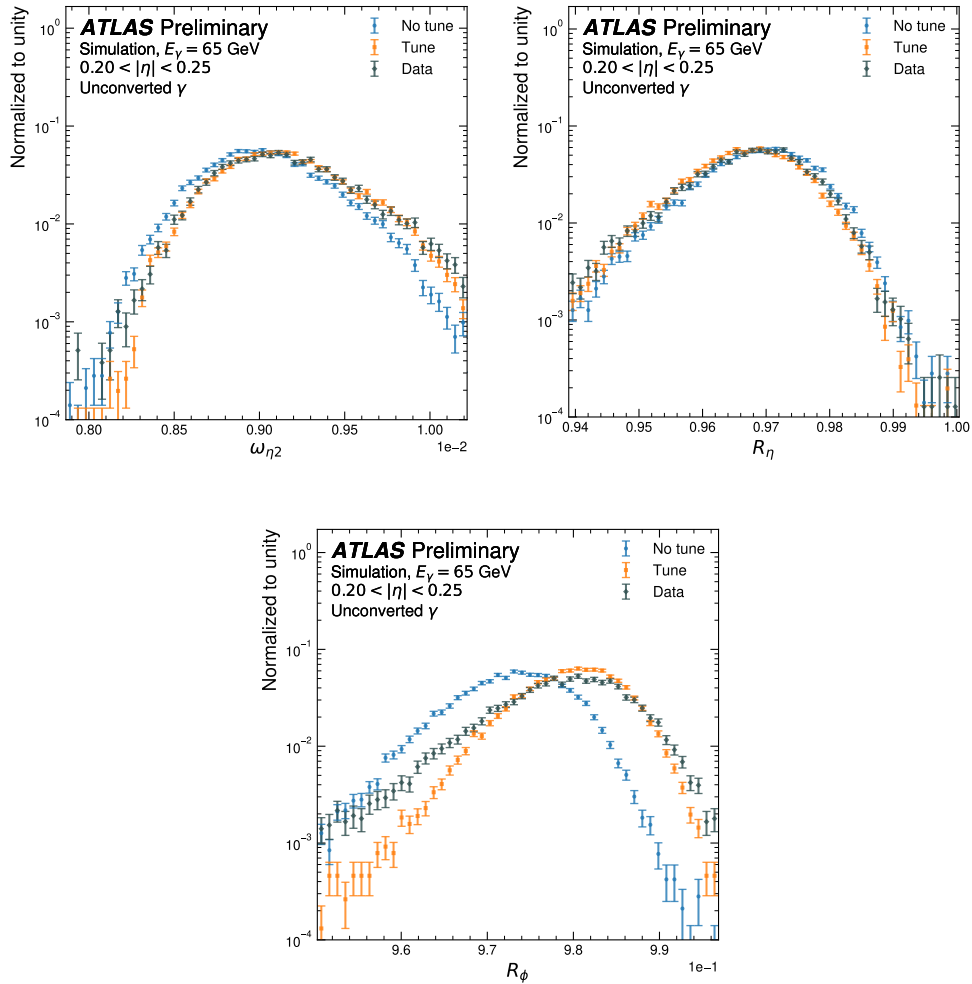


Figure 8.10: Validation of the linear hit scaling model for the most important middle layer shower shape variables. The distributions for the nominal simulation are shown in blue, while the tuned simulation is shown in orange and the data in green. The uncertainties on the data include statistical sources only and no background subtraction is performed. [213]

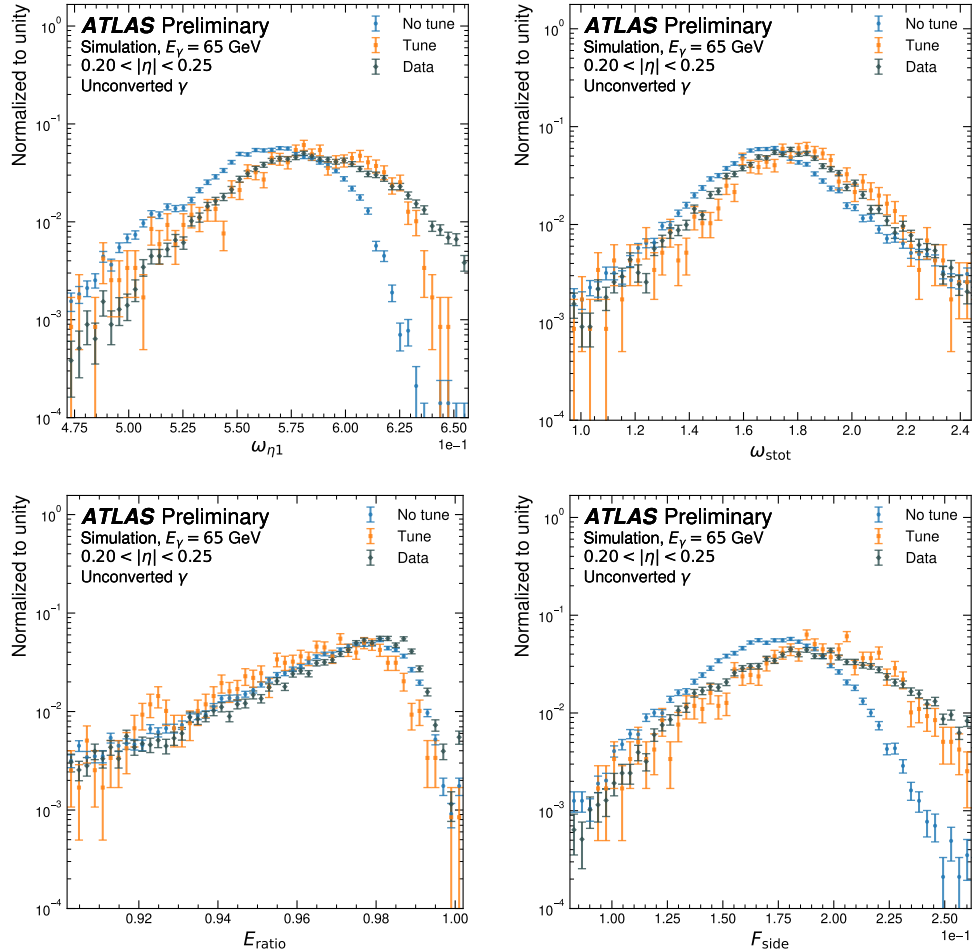


Figure 8.11: Validation of the cubic hit scaling model for the most important strip layer shower shape variables. The distributions for the nominal simulation are shown in blue, while the tuned simulation is shown in orange and the data in green. The uncertainties on the data include statistical sources only and no background subtraction is performed. [213]

Variable	RMD (%)		RRD (%)	
	No tune	Tune	No tune	Tune
$w_{\eta 2}$	1.26 ± 0.06	0.01 ± 0.06	6 ± 1	4 ± 1
R_{η}	0.16 ± 0.01	0.06 ± 0.01	5 ± 1	5 ± 1
R_{ϕ}	0.55 ± 0.01	0.15 ± 0.01	20 ± 1	26 ± 1
$w_{\eta 1}$	3.54 ± 0.07	0.4 ± 0.2	11 ± 1	6 ± 3
w_{stot}	3.8 ± 0.2	0.8 ± 0.5	2 ± 1	4 ± 3
E_{ratio}	0.24 ± 0.04	0.50 ± 0.08	6 ± 1	14 ± 3
F_{side}	10.4 ± 0.2	0.4 ± 0.5	22 ± 1	9 ± 3

Table 8.1: Absolute relative mean difference (RMD) and absolute relative RMS difference (RRD) to data for the nominal and the tuned simulation for the various targeted shower shape variables. [213]

FASTCALOSIMV2 parametrization point, that is, for 65 GeV photons in the $0.20 < |\eta| < 0.25$ region. In this section, the extension of the model to other detector regions and photon energies is discussed.

8.4.1 Optimal Model Parameters

For practical reasons, the optimal model parameters for the extension of the tuning models to the full FASTCALOSIMV2 parametrization grid are determined using the distributed asynchronous optimization procedure developed for the strip model in the case of both strip and middle layer variables. The inputs used for minimizing the χ^2 error as defined in Equation (8.3) are:

1. The baseline means and standard deviations of the shower shapes of the nominal simulation. The values are computed from single particle simulations within the standalone FASTCALOSIMV2 framework, that is, without the inclusion of noise and the addition of pile-up interactions. The respective means and RMS values as a function of incident direction for photons of energies between 2 GeV and 4 TeV is shown in Figure 8.12 and Figure 8.13, respectively. As expected, larger incident photon energies cause narrower showers in both, η and ϕ directions. The dependence of the shower shapes on the incident photon direction is the result of the variation in the calorimeter geometry. For instance, the shower shape means and standard deviations show a clear discontinuity around $|\eta| \sim 0.8$ as a result of a change in the longitudinal depth in the barrel in this region.
2. Scaling factors that encode by what amount the shower shape mean from the nominal simulation is required to be scaled to fit the mean shower shapes observed in data. The scaling factors are extracted from comparing the shower shape means as simulated




Figure 8.12: Mean of the strip and middle layer shower shapes of photons with varying energies between 2 GeV and 4 TeV in dependence of the incident direction as computed using the standalone FASTCALOSIMV2 simulation. As such, no additional pp interactions are overlaid and no noise is simulated. The hatched area indicates the transition regions, in which reconstructed photons are typically excluded.




Figure 8.13: Standard deviations of the strip and middle layer shower shapes of photons with varying energies between 2 GeV and 4 TeV in dependence of the incident direction as computed using the standalone FASTCALOSIMV2 simulation. As such, no additional pp interactions are overlaid and no noise is simulated. The hatched area indicates the transition regions, in which reconstructed photons are typically excluded.

using FASTCALOSIMV2 with those observed from photon candidates collected in data. In this case, the FASTCALOSIMV2 simulation is propagated through the full ATLAS reconstruction and digitization chain, additional pp interactions are overlaid and contributions from noise terms are simulated in order to best mimic data taking conditions. The mean scaling factors are computed for each FASTCALOSIMV2 parametrization point in the case that a sufficient amount of photon candidates is available in the respective region of phase space. The energy region used to select photons is chosen as the $\pm 10\%$ window around the respective energy parametrization point. As a consequence, the model parameters are derived for all parametrization points between 32 GeV and 1 TeV in the logarithmically spaced energy grid employed by FASTCALOSIMV2 and only in the $|\eta|$ regions in which enough photons are available in data and the simulation. The resulting mean scaling factors are shown in Figure 8.14 and employed in the FASTCALOSIMV2 standalone optimization procedure to calculate the target means of the variables.

In all cases, the computed mean scaling factors remain relatively constant until $|\eta| \sim 0.5$ and approximately correspond to the values shown in Table 8.1. For more forward regions, the scaling factor for $w_{\eta 1}$ increases from $\sim 4\%$ to $\sim 8\%$ and remains approximately constant. In both central and forward directions, no significant energy dependence is observed. For E_{ratio} , the scaling factors are generally close to zero, but can increase to up to $1 - 2\%$ for large photon energies. The mean scaling factors for w_{stot} are within a $0 - 5\%$ range in the central detector region and a $5 - 20\%$ range in the endcaps. Low energetic photons show a significantly stronger deviation from the means observed in data in this case. A similar behaviour is observed for F_{side} , for which the overall deviation to data is found largest. In the central region, the scaling factors lie in a $0 - 10\%$ window, whereas the scaling factors strongly increase in the forward region and fluctuate between $10 - 40\%$. Similar to $w_{\eta 1}$, no significant energy dependence is found for $w_{\eta 2}$. In this case, the scaling factor ranges between $1 - 4\%$ in the central region and $4 - 7\%$ in the endcap. For R_{η} , the scaling factors are in the range between -0.25% and -1.25% with the same tendency of larger corrections in the endcap region. Lastly, R_{ϕ} has a constant scaling of $\sim 0.5\%$ until $|\eta| \sim 0.5$ which drops up to -1% for photons in the low energy regime and fluctuates close to zero in the endcap region. Note that scaling factors are only computed in the case that at least 500 photon candidates pass the selection described in Section 8.2. As such, the model parameters are only optimized for the points shown in Figure 8.14. For all other cases, the model parameters are interpolated as described in Section 8.4.2.

In order to effectively optimize the model parameters, suitable search spaces need to be chosen. If the search space of choice is too small, one risks excluding important minima in the black-box optimization procedure. On the other hand, if the search space is too large, a random search algorithm is likely unable to find the (global) minimum within the specified space after a reasonable amount of trials. Taking into consideration the very different scalings which are necessary in the central and forward regions of the detector, a different choice of search spaces in the barrel and endcap regions is shown to be beneficial. The search spaces are iteratively optimized by taking into account the partial dependencies [220] of the model parameters, that

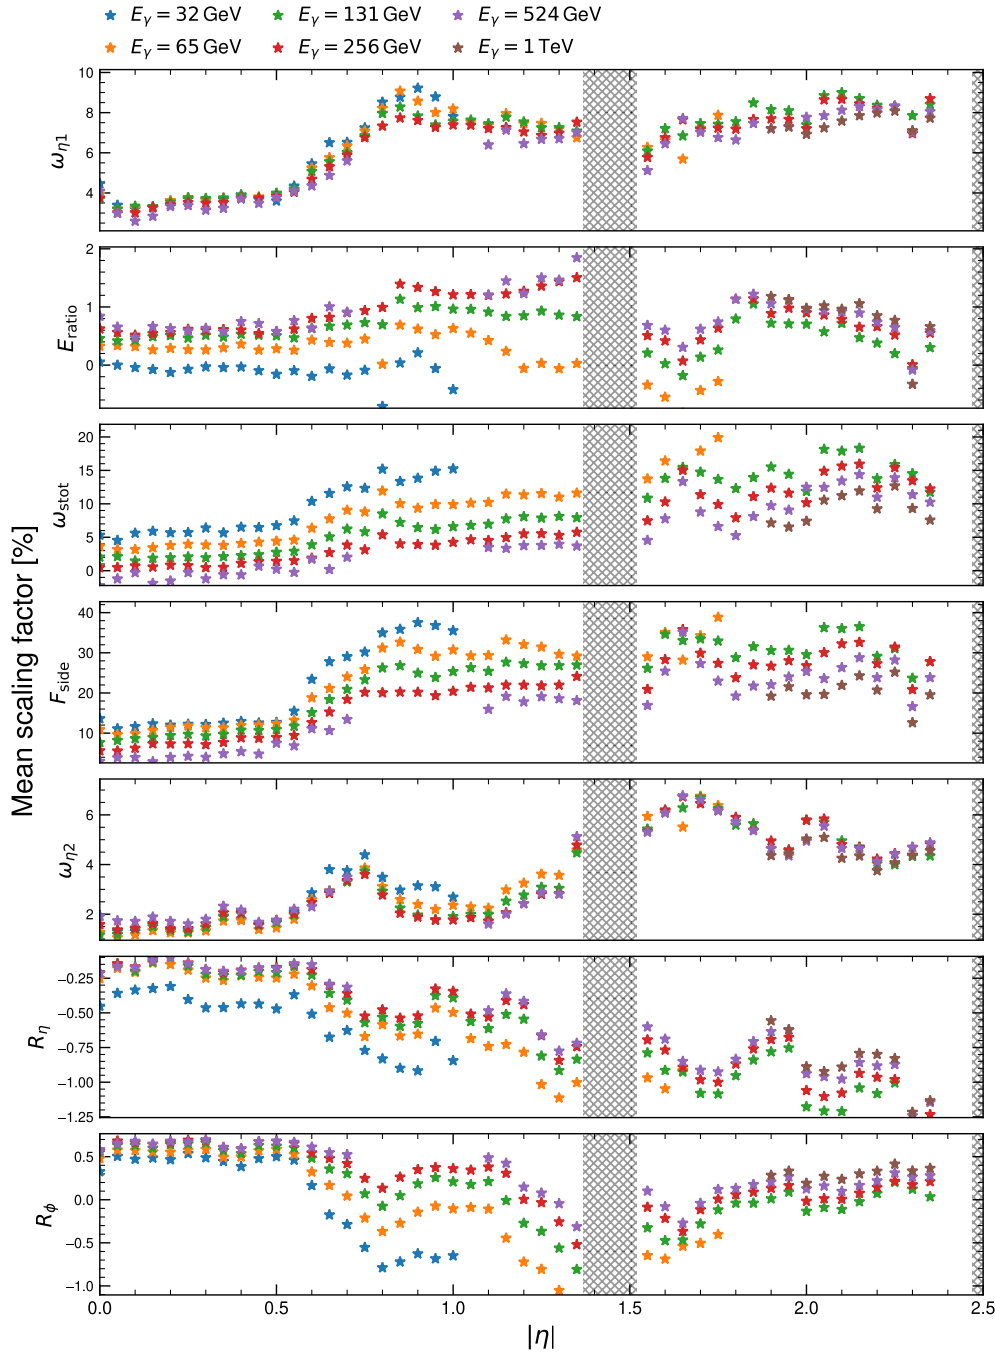


Figure 8.14: Mean shower shape scaling factor for photons with energies ranging from 32 GeV and 1 TeV simulated with the nominal FASTCALOSIMV2 simulation and photons collected in data in dependence of the incident photon direction. Single photons are simulated including noise at the origin and additional pp interactions are overlaid to mimic data taking conditions. The hatched area indicates the transition regions, in which reconstructed photons are typically excluded. Scaling factors are not computed in case that insufficient reconstructed photons candidates are available in either simulation or data. 159

is, the marginal effect that each parameter has on the χ^2 error. The optimization is performed for the $0.20 < |\eta| < 0.25$ and $2.00 < |\eta| < 2.05$ regions, which serves as proxy for the barrel and endcap region, respectively. The resulting optimized search spaces are shown in Table 8.2 for the strip and middle layer model. While the search spaces are optimized for the specific

		Strip layer		Middle layer	
		Barrel	Endcap	Barrel	Endcap
a_0	[0.10, 0.35]	[0.70, 0.90]	η_s	[0.90, 1.10]	[1.10, 1.25]
a_1	[0.20, 0.50]	[-0.10, 0.10]	ϕ_s	[0.70, 0.90]	[0.98, 1.25]
a_2	[-0.60, -0.30]	[-0.40, -0.20]			
a_3	[0.20, 0.40]	[0.20, 0.40]			

Table 8.2: Uniform search spaces for the strip and middle layer tuning models optimized for the barrel and endcap regions and used to find optimal model parameters for the full detector coverage.

barrel and endcap regions, the optimization with each search space is performed across the full $|\eta|$ range, where the model parameters with minimum χ^2 are chosen.

The optimized model parameters for the middle and strip layer of the calorimeter are shown in Figure 8.15 for 131 GeV photons. The lower panel shows the ratio between the baseline and optimized χ^2 error. Up to $|\eta| \sim 0.7$, the η_s and ϕ_s scalings remain approximately constant and are very similar to the two-dimensional proof-of-concept optimization described in Section 8.3.1. For larger $|\eta|$, the optimal parameters are found within the search space optimized for the endcap region and increase accordingly. In most cases, the χ^2 is significantly improved by two to three orders of magnitude.

A similar picture is found in the case of the strip layer model. The model parameters a_k remain approximately constant until $|\eta| \sim 0.7$ at which a sudden increase from $a_0 \sim 0.25$ to $a_0 \sim 0.75$ and decrease from $a_1 \sim 0.35$ to $a_1 \sim 0.05$ is observed. As before, the rapid change is the result of the optimal parameters being chosen from the region optimized for the endcap. In most cases, the χ^2 is significantly improved by 1-2 orders of magnitude. The least improvement is observed very close to the transition region between barrel and endcap.

8.4.2 Model Parameter Interpolation

In general, the energy dependence of the optimized parameters is modest, and significantly more pronounced in the middle layer model. Figure 8.16 shows the optimized model parameters as a function of photon energy for the middle and strip layer model in the $0.20 < |\eta| < 0.25$ region. While the energy dependence is very much dependent on the direction of the incident photon, the η_s and ϕ_s scaling factors are observed to decrease in many cases with higher

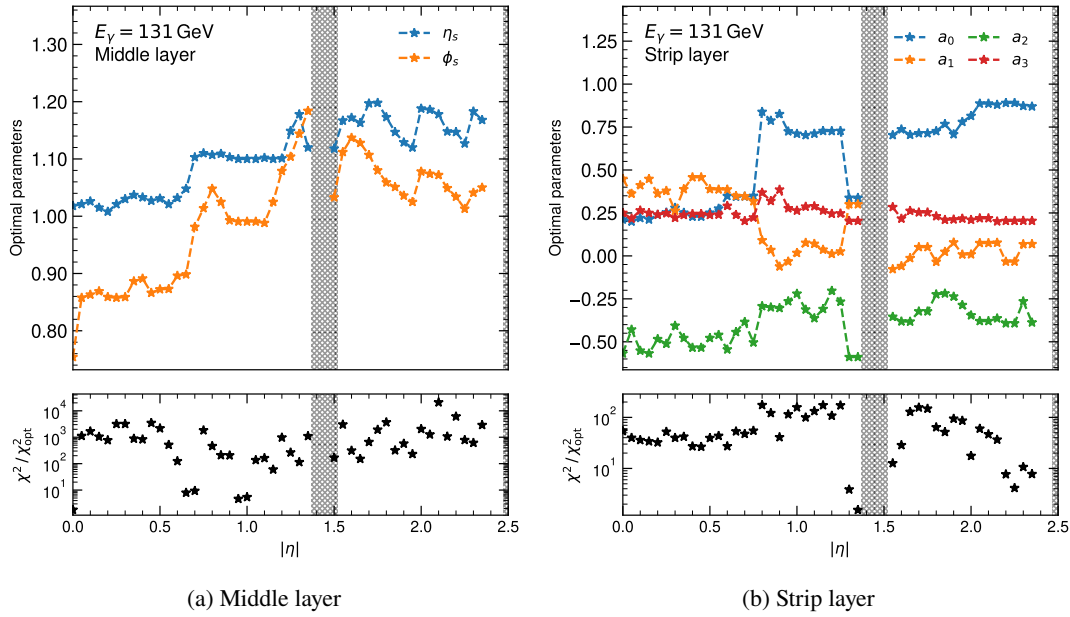


Figure 8.15: Optimal model parameters in dependence of $|\eta|$ for the (a) middle and (b) strip layer of the calorimeter for 131 GeV photons. The lower panel shows the ratio between the baseline error and the evaluated error from the optimized model parameters.

photon energies. In contrast, the strip layer model parameters do not exhibit a distinct energy dependence and remain largely constant for most detector regions.

Similar to the interpolation of the energy parametrization discussed in Section 4.5, the set of model parameters is chosen depending on the direction of the incident photon. The model parameters are then piecewise linearly interpolated for photon energies between the energy points for which the parameters are optimized. As a conservative approach, no extrapolation of the model parameters beyond the optimized energy range is employed. Instead, the model parameters for photons with very low and very high energies are fixed to the values at the edges of the available energy range (see Figure 8.16).

8.4.3 $H \rightarrow \gamma\gamma$ Model Validation

In order to validate the model beyond the means computed from the standalone FASTCALOSIMV2 simulation, approximately 40 million $H \rightarrow \gamma\gamma$ events are simulated with the tuned FASTCALOSIMV2 simulation. The process is simulated with the identical Run 2 data taking conditions as the single photon samples that are used to compute the mean scaling factors with respect to data, and the same selection criteria described in Section 8.2 is employed to select photon candidates.

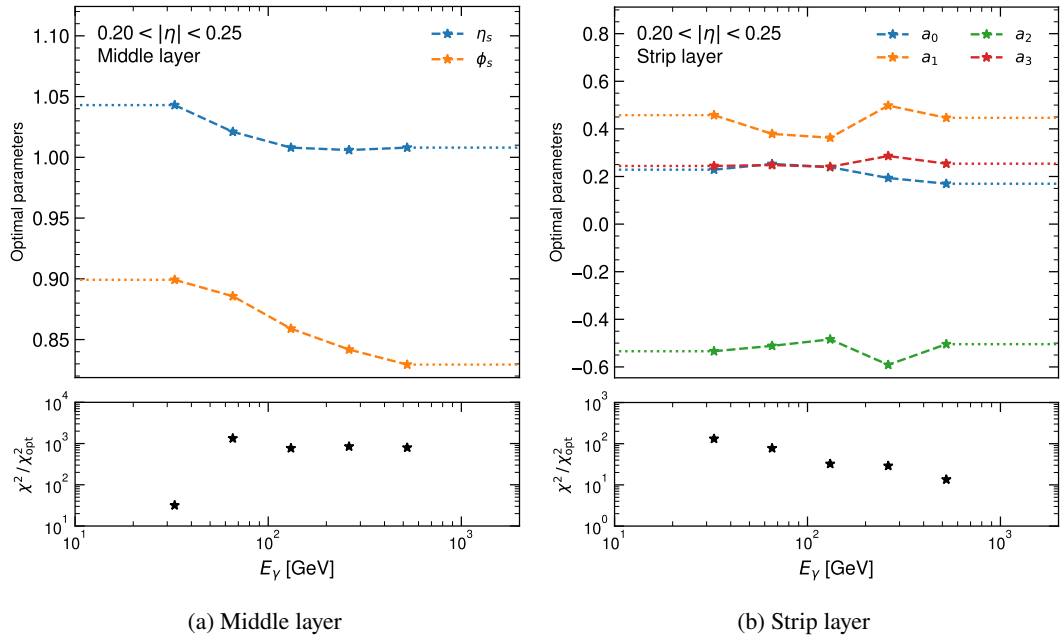


Figure 8.16: Optimized model parameters for the (a) middle and (b) layer model as a function of the incident photon energy in the $0.20 < |\eta| < 0.25$ region. The dashed lines indicate values that are piecewise linearly interpolated between optimized energy points. The dotted lines show model parameters used outside the optimized energy range. The lower panel shows the ratio between the baseline error and the evaluated error from the optimized model parameters.

The shower shapes resulting from photons selected from the tuned simulation of $H \rightarrow \gamma\gamma$ events are compared to those from the nominal simulation of the reference single photon sample. The ratio of the shower shape means of both simulations with respect to data for the various targeted variables as a function of incident photon directions is shown in Figure 8.17 for 131 GeV and 262 GeV unconverted photons. As expected, the tuned simulation performs similarly well for both photon energies and across the detector regions. The means of the shower widths R_η and R_ϕ in the middle layer can be impressively reproduced across the full $|\eta|$ range. Similarly, the proper mean widths $w_{\eta 1}$ and $w_{\eta 2}$ of the strip and middle layer, fluctuate close to unity in a $1 - 2\%$ window. The mean of w_{stot} is generally well reproduced across most regions, but a degradation of the tune is observed close to the transition region, which is the result of the reduced χ^2 ratio already observed in Figure 8.15 for the strip layer model. With an improvement of up to 20%, F_{side} exhibits the largest improvement in the mean shower shape ratio. In this case, the tuned simulation performs excellently in the central region and slightly deteriorates for larger $|\eta|$. As already seen in the development of the strip layer model for a single parametrization point (see Section 8.3.3), a consistent deterioration in the modelling of E_{ratio} is observed. The worsening is moderate in the central region and increases incrementally




Figure 8.17: Ratio of the shower shape mean observed for the nominal (blue) and tuned (orange) simulation with respect to data in dependence of the incident photon direction for (a) 131 GeV and (b) 262 GeV unconverted photons. The grey dashed line shows the optimal ratio at unity. The hatched area indicates the transition regions, in which reconstructed photons are typically excluded.

in the forward barrel and endcap region. However, the maximum difference of the mean to data in the tuned simulation does not exceed more than approximately 2 %.

As an example, Figure 8.18 and Figure 8.19 show the comparison between the nominal and tuned simulation for 131 GeV photons in the barrel region ($0.20 < |\eta| < 0.25$) for the strip and middle layer shower shapes, respectively. The corresponding distributions for the endcap region ($2.00 < |\eta| < 2.05$) are shown for the strip and middle layer shower shapes in Figure 8.20 and Figure 8.21, respectively.

In all cases, except of E_{ratio} , the shower shapes appear significantly better modelled in the tuned simulation of FASTCALOSIMV2 and only small differences remain in the tails of some distributions.

8.5 Model Application on Electron-induced Showers

While the models developed for the strip and middle layer of the calorimeter are optimized for photons only, they can be equally applied to the simulation of electron-induced showers within

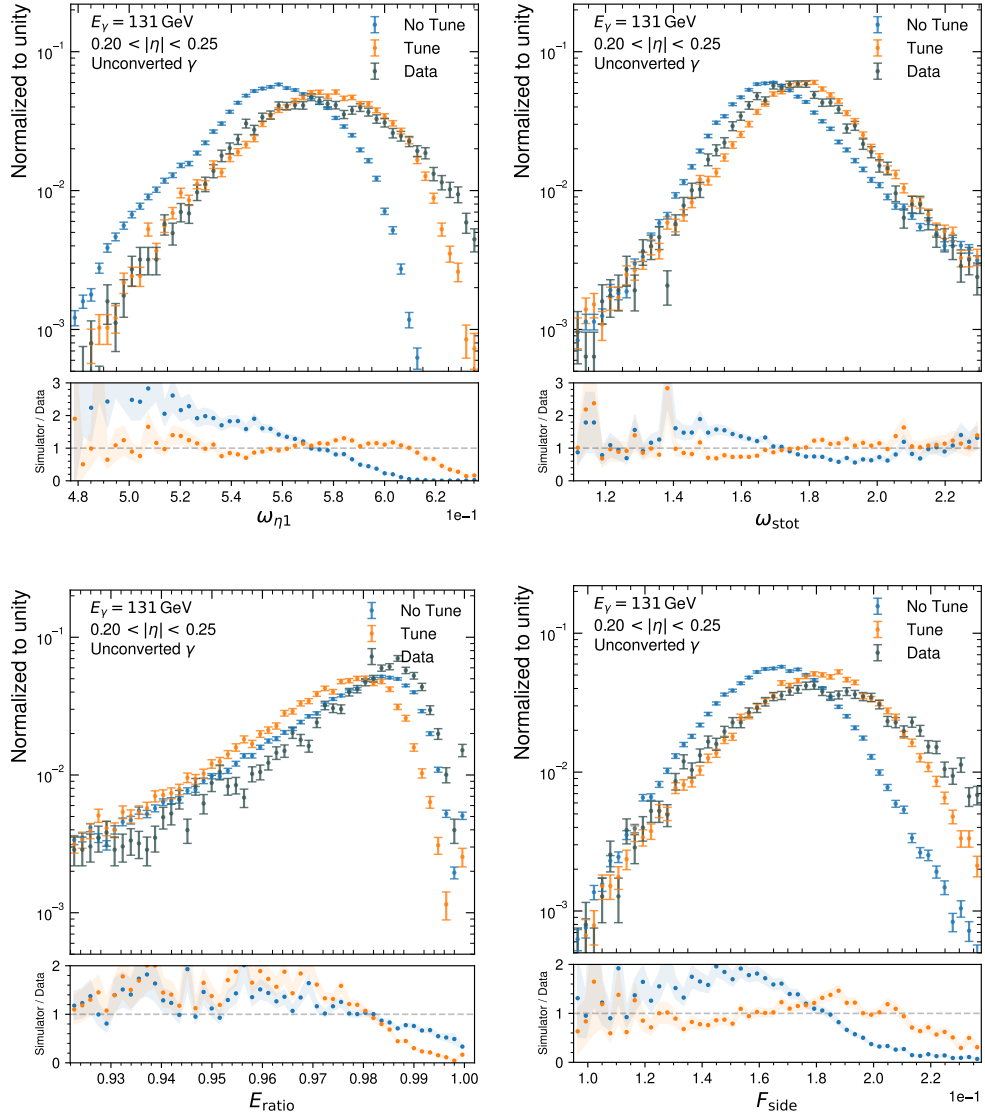


Figure 8.18: Strip layer shower shapes as simulated with the nominal and tuned FASTCALOSIMV2 simulation, and as observed in data for 131 GeV unconverted photons in the barrel region of the calorimeter. The lower panel shows the ratio between the simulation and the data.

8.5 Model Application on Electron-induced Showers

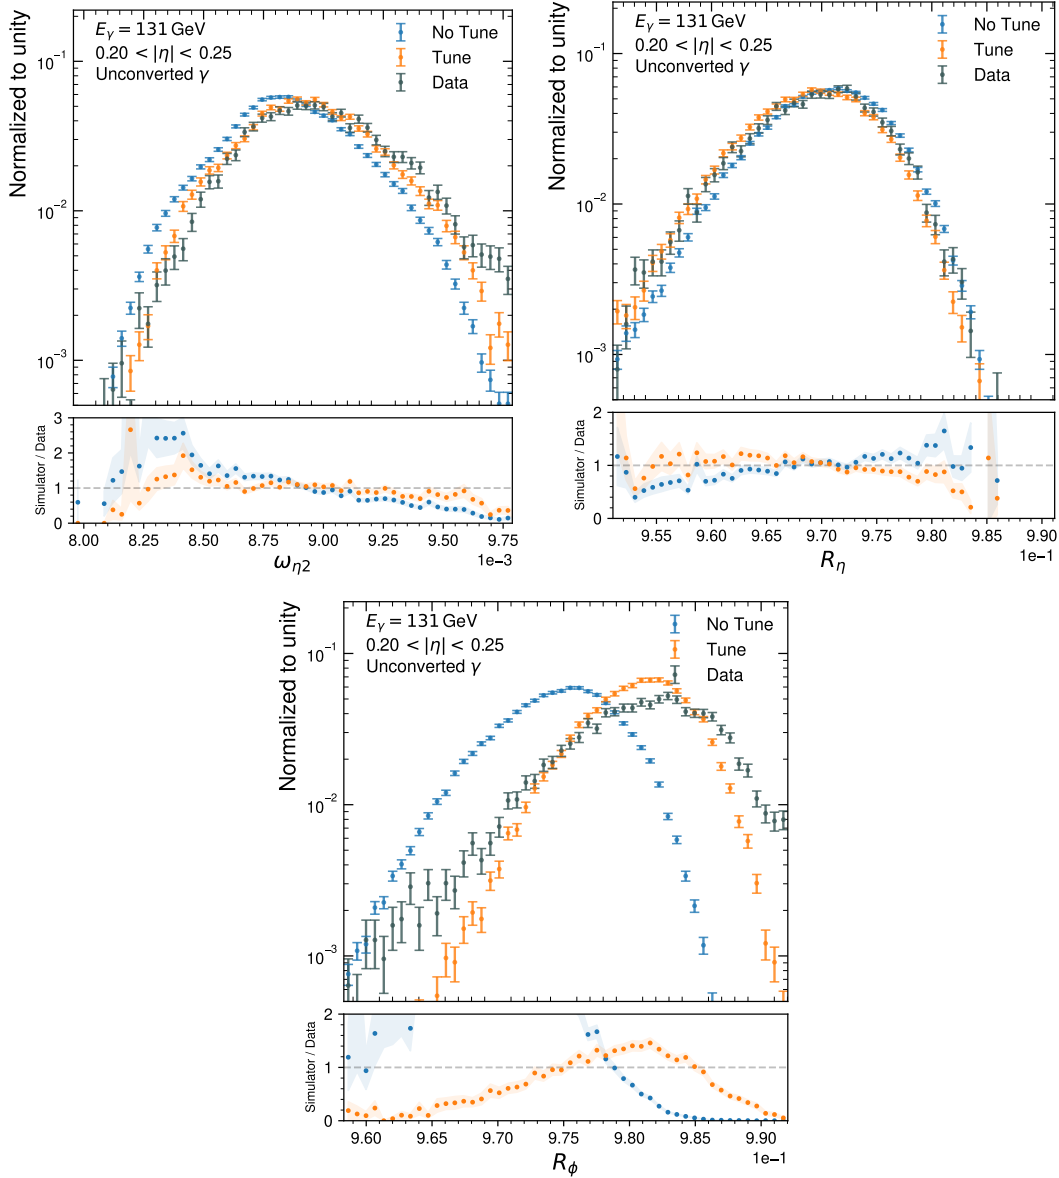


Figure 8.19: Middle layer shower shapes as simulated with the nominal and tuned FASTCALOSIMV2 simulation, and as observed in data for 131 GeV unconverted photons in the barrel region of the calorimeter. The lower panel shows the ratio between the simulation and the data.

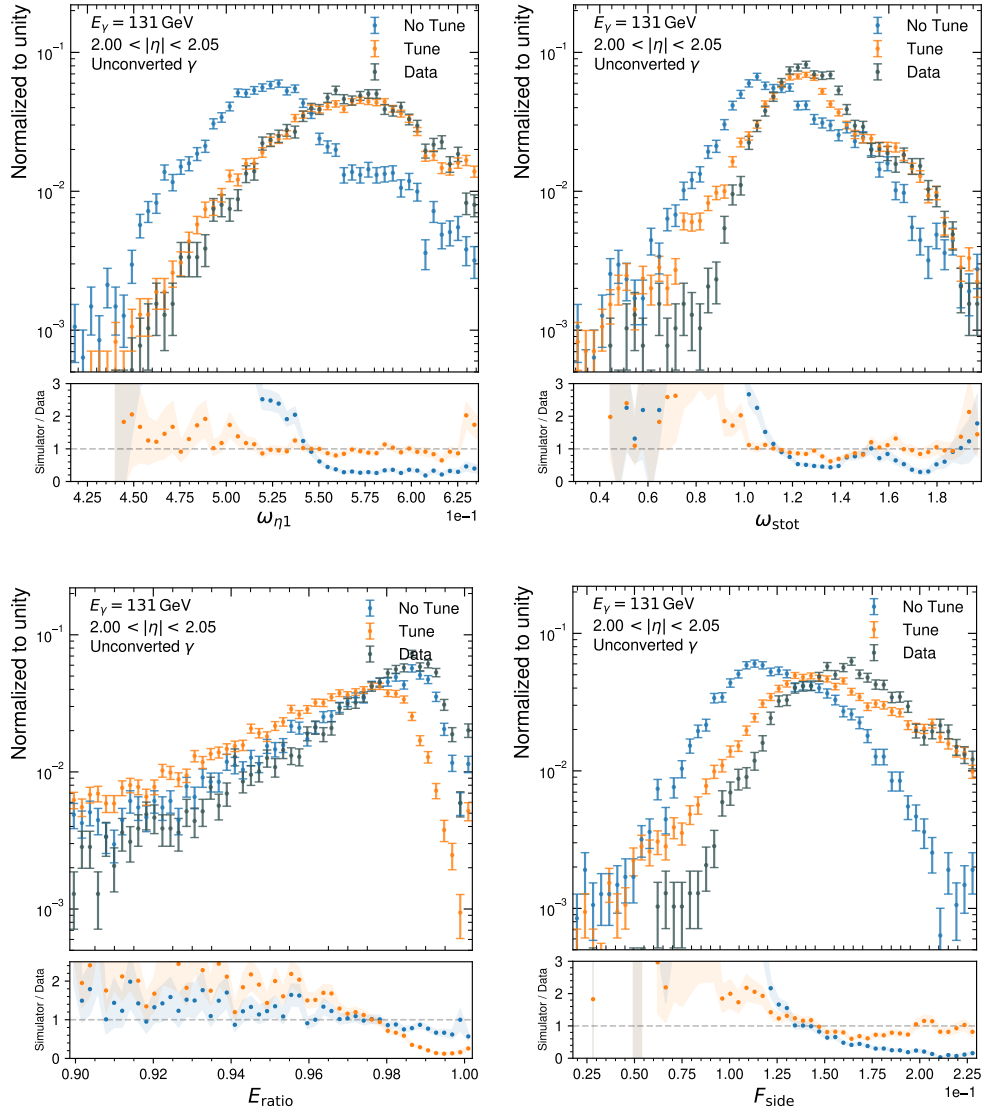


Figure 8.20: Strip layer shower shapes as simulated with the nominal and tuned FASTCALOSIMV2 simulation, and as observed in data for 131 GeV unconverted photons in the endcap region of the calorimeter. The lower panel shows the ratio between the simulation and the data.

8.5 Model Application on Electron-induced Showers

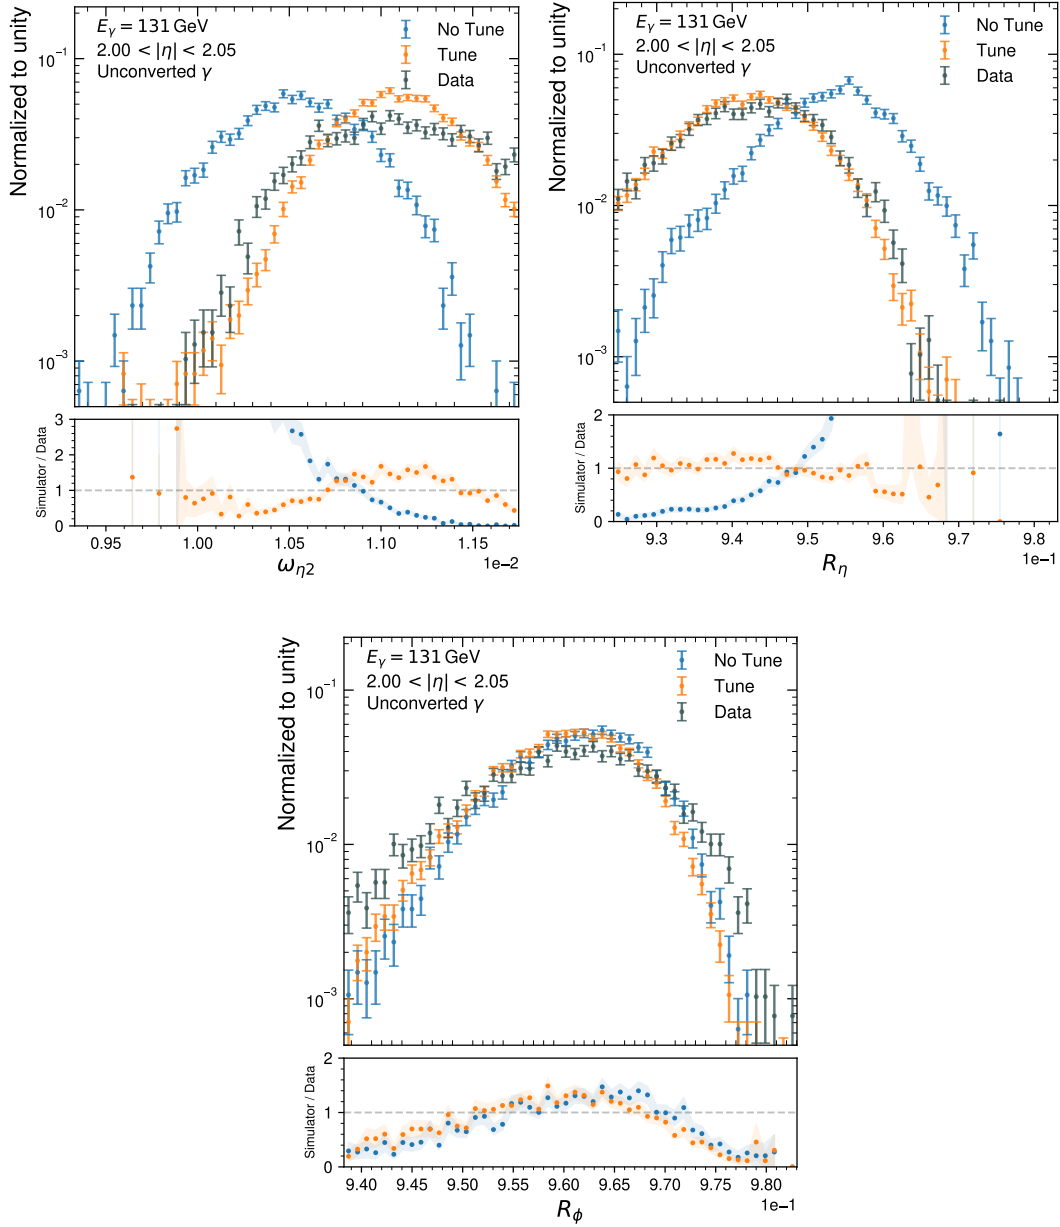


Figure 8.21: Middle layer shower shapes as simulated with the nominal and tuned FASTCALOSIMV2 simulation, and as observed in data for 131 GeV unconverted photons in the endcap region of the calorimeter. The lower panel shows the ratio between the simulation and the data.

FASTCALOSIMV2. As the electron- and photon-induced shower development is fairly similar, the tune is also expected to be beneficial in this case.

In order to study the effect of the application of the model to electron-induced showers, the shower shapes of converted photons are compared, while all other selections described in Section 8.2 remain unchanged. The comparison of converted photons is an indirect probe of the application of the tuning model to electron showers, as the photon decays to an e^+e^- pair prior to reaching the boundary between the Inner Detector and the calorimeter. As such, the energy response induced by the e^+e^- pair is simulated with the FASTCALOSIMV2 electron parametrization to which the tuning model is applied as well. Note, however, that the purity of converted photons is shown to be significantly lower (between 60-80% for $E_\gamma \gtrsim 100$ GeV) compared to unconverted photons, such that some bias in the shower shapes resulting from background contaminations is to be expected.




Figure 8.22: Ratio of the shower shape mean observed for the nominal (blue) and tuned (orange) simulation with respect to data in dependence of the incident photon direction for (a) 131 GeV and (b) 262 GeV converted photons. The grey dashed line shows the optimal ratio at unity. The hatched area indicates the transition regions, in which reconstructed photons are typically excluded.

Figure 8.22 shows the ratio of the shower shape means of the nominal and tuned simulation with respect to data for the various targeted variables as a function of incident photon directions for 131 GeV converted photons. As for unconverted photons, the energy dependence of the model performance is modest. In general, the tuned simulation shows a fair improvement in the

modelling of most of the shower shape means, but to a significantly lesser extent compared to photon-induced showers. In the case of $w_{\eta 2}$ and w_{stot} , the middle and strip layer models overestimate the necessary scaling of the means in the central region, such that absolute difference observed in data is similar or worse than the nominal simulation. However, in both cases, the means are significantly better simulated in the endcap region when applying both tuning models to electron-initiated showers. For F_{side} and R_{η} , the tune is shown to be beneficial across the full $|\eta|$ range, while E_{ratio} exhibits a similar degradation as observed for the simulation of photon showers. Lastly, the middle layer model exhibits a fair improvement in the modelling of R_{ϕ} , while no significant differences are observed in the endcap region with respect to the nominal simulation.

As an example, Figure 8.23 and Figure 8.24 show the comparison between the nominal and tuned simulation for 131 GeV photons in the barrel region ($0.20 < |\eta| < 0.25$) for the strip and middle layer shower shapes, respectively. The corresponding distributions for the endcap region ($2.00 < |\eta| < 2.05$) are shown for the strip and middle layer shower shapes in Figure 8.25 and Figure 8.26, respectively. In almost all cases, the shower shapes in the tuned simulation appear significantly better modelled. In the case of the barrel, the means of the widths in the first layer $w_{\eta 1}$ and w_{stot} exhibit a shift in their means that is slightly too large, while excellent modelling is found in the endcap region for both variables. R_{ϕ} , on the other hand, is very well reproduced in the central region, while no significant improvement is observed in the endcap.

8.6 Conclusion

Differences between the simulation of electromagnetic shower properties and data are a long-standing problem in ATLAS and could only be tackled using *ad hoc* methods such as the application of *fudge factors* until now. This chapter presented two different models embedded in FASTCALOSIMV2 that are able to precisely reproduce most of the photon shower shapes measured in the first and second layer of the calorimeter for a large range of incident photon energies and directions. The data tune is directly integrated in the simulation toolkit and does not employ any *ad hoc* post-simulation corrections.

Both models are shown to effectively reproduce the widths of the photon-induced showers in the strip and middle layer as well as the fraction of energy deposited far away from the shower centre. This, however, has been shown to cause a small, but consistent degradation in the modelling of E_{ratio} , that is, the energy distribution between the two largest energy deposits, of up to a few percent. Attempts to improve its overall modelling by weighting the importance of E_{ratio} by large factors in the optimization metric given in Equation (8.3) have not shown any significant improvements. Similarly, the addition of a Gaussian smearing term to the η_s scaling from Equation (8.2) in the strip layer model is not found to be beneficial. While the extension of the η_s series to higher order terms ($n > 3$) might result in minor benefits at the expense of a larger parameter space, both observations seem indicative of an intrinsic model limitation.

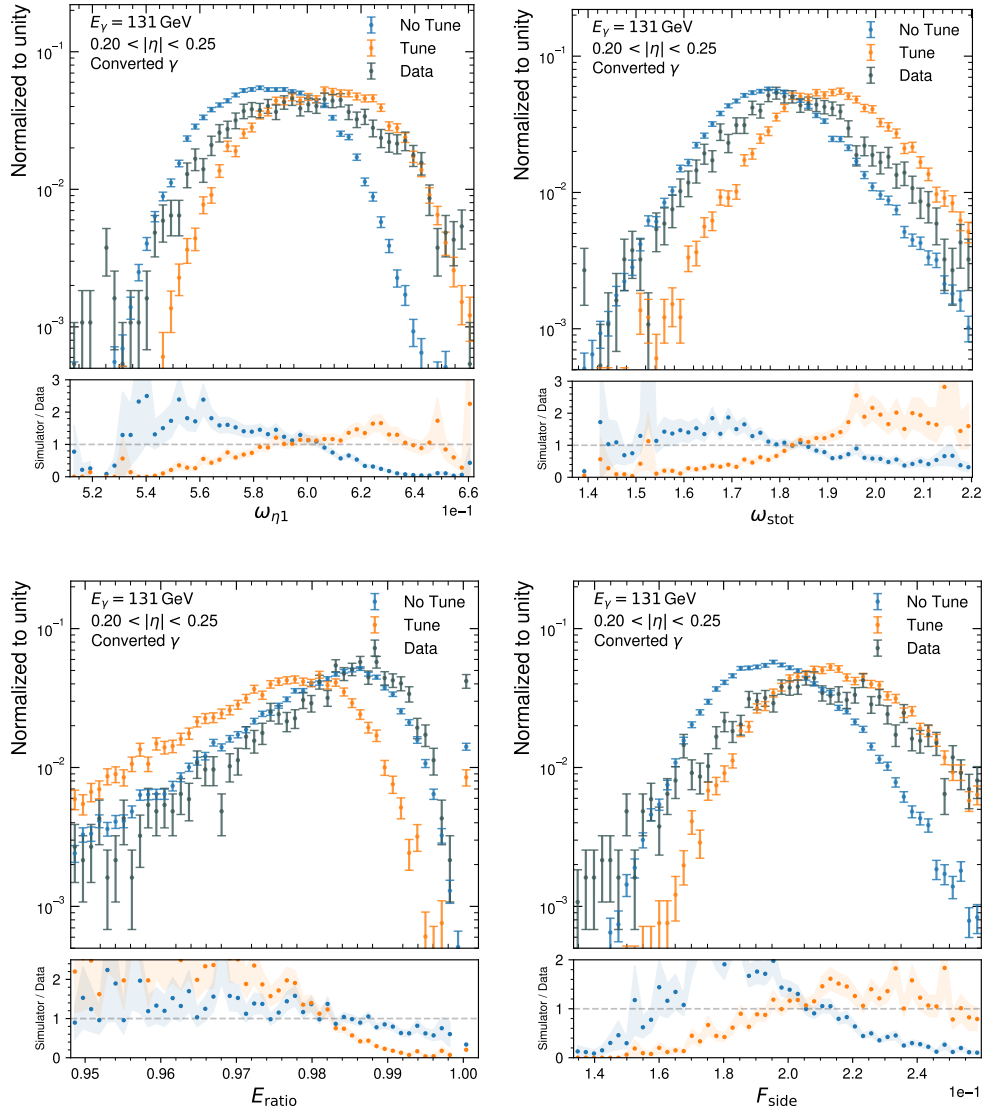


Figure 8.23: Strip layer shower shapes as simulated with the nominal and tuned FASTCALOSIMV2 simulation, and as observed in data for 131 GeV converted photons in the barrel region of the calorimeter. The lower panel shows the ratio between the simulation and the data.

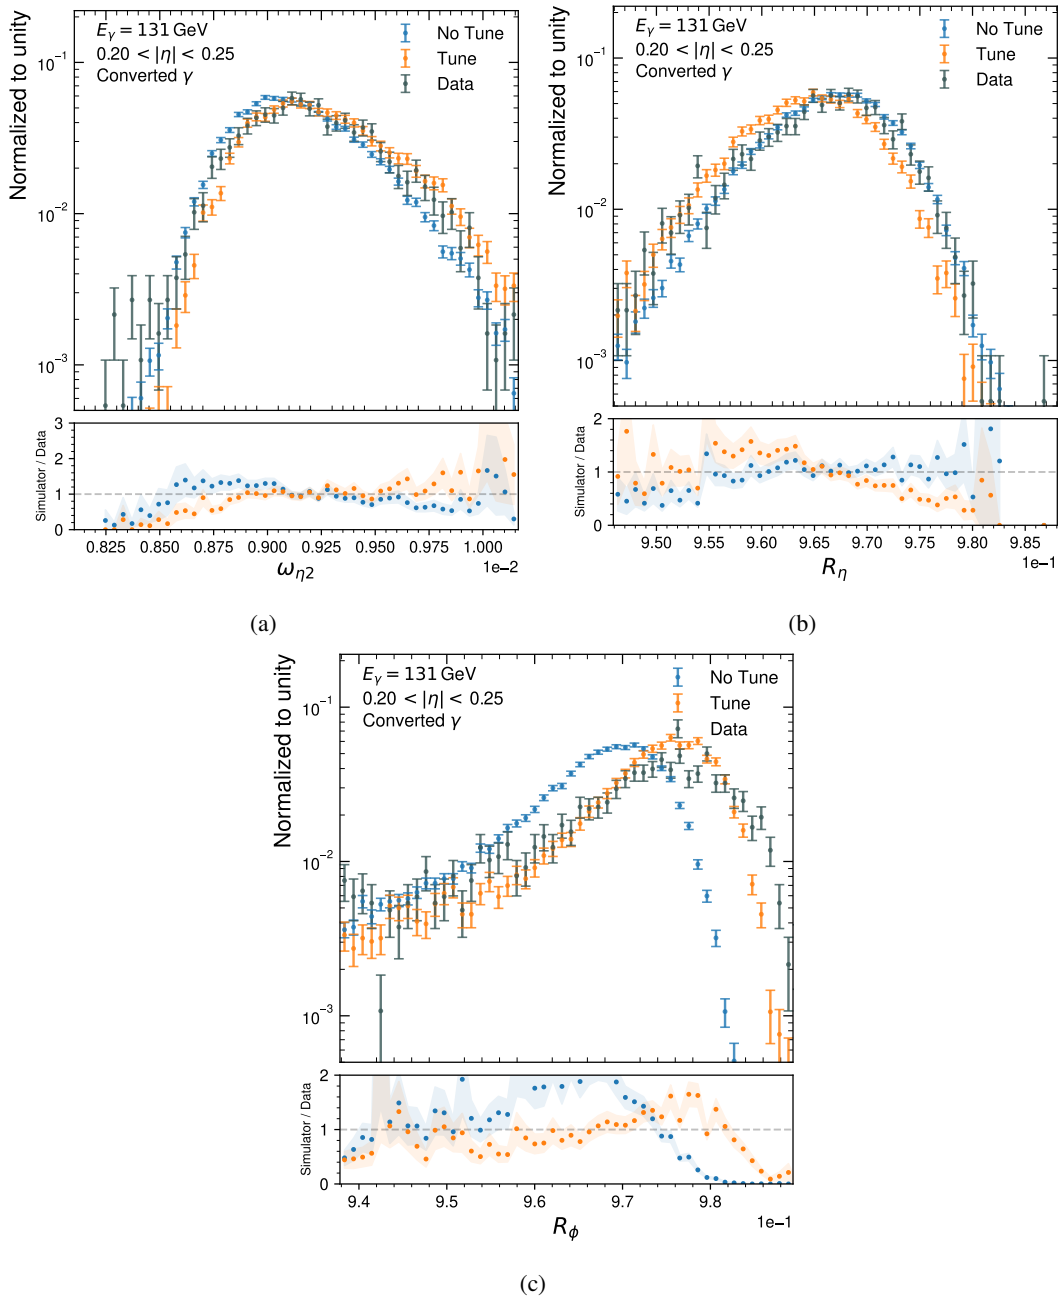


Figure 8.24: Middle layer shower shapes as simulated with the nominal and tuned FASTCALOSIMV2 simulation, and as observed in data for 131 GeV converted photons in the barrel region of the calorimeter. The lower panel shows the ratio between the simulation and the data.

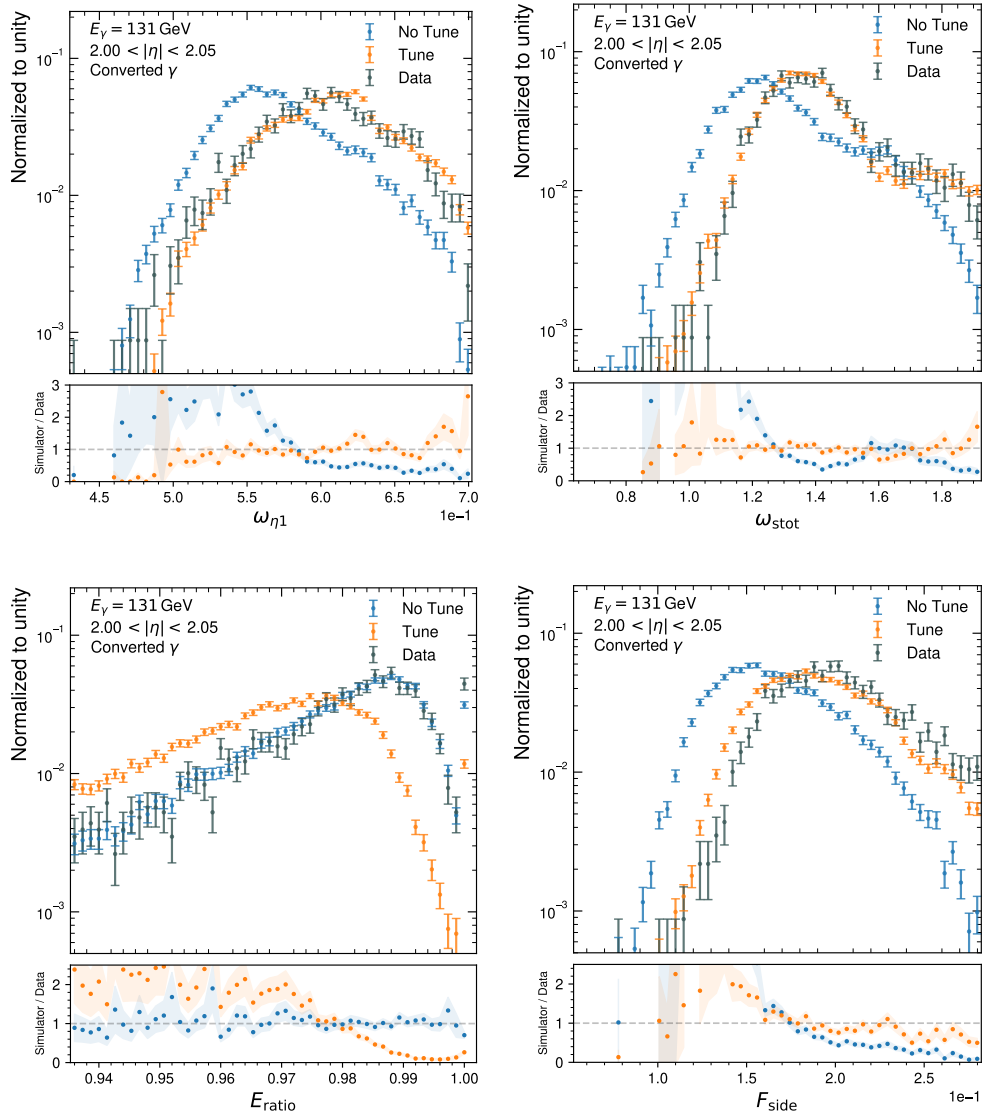


Figure 8.25: Strip layer shower shapes as simulated with the nominal and tuned FASTCALOSIMV2 simulation, and as observed in data for 131 GeV converted photons in the endcap region of the calorimeter. The lower panel shows the ratio between the simulation and the data.

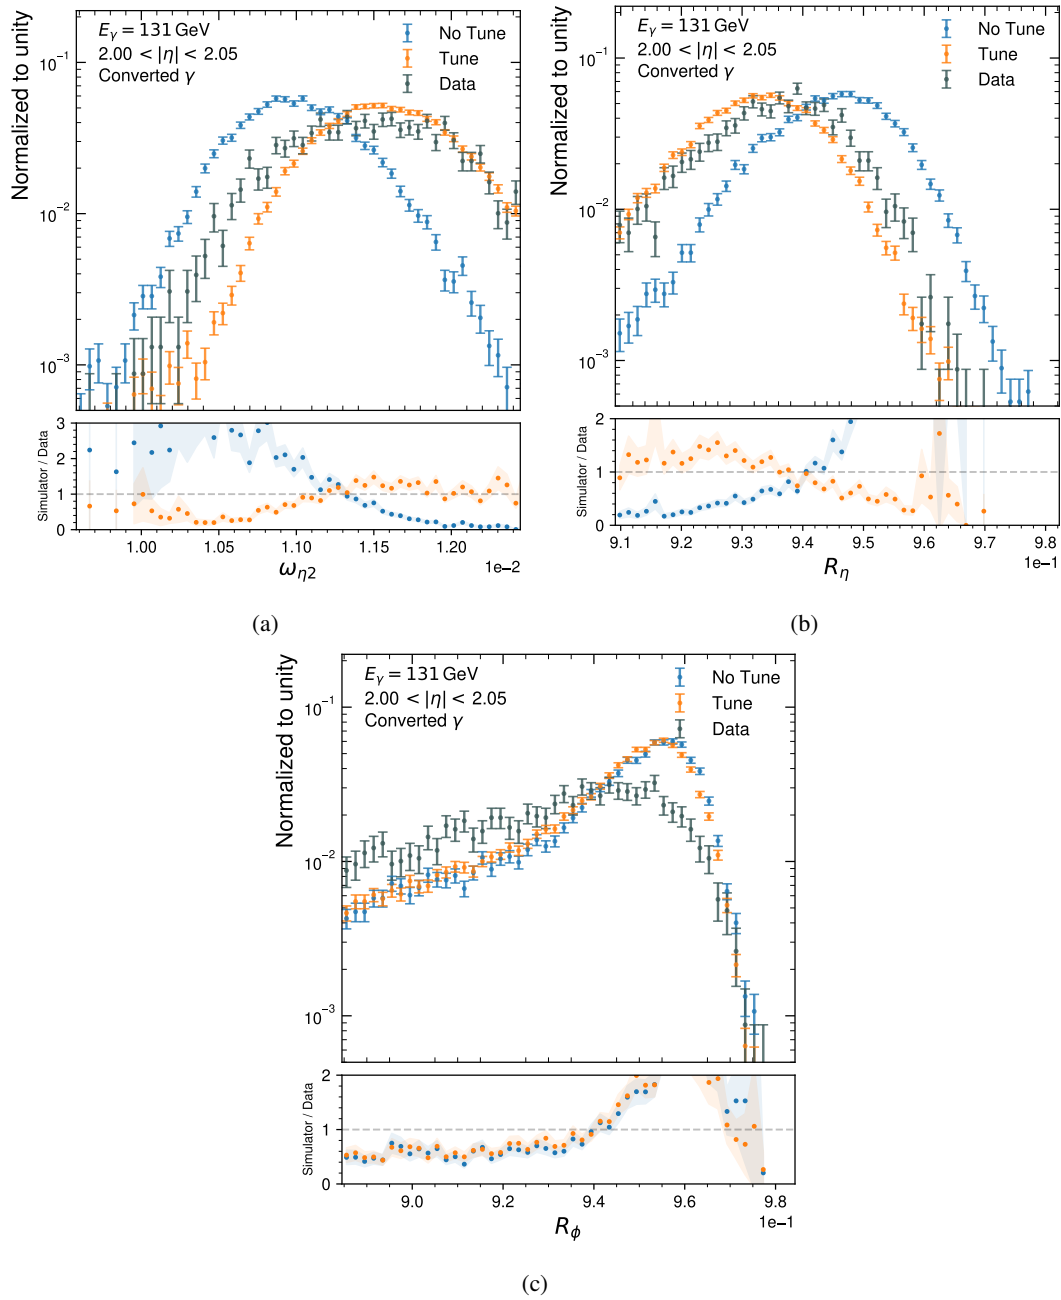


Figure 8.26: Middle layer shower shapes as simulated with the nominal and tuned FASTCALOSIMV2 simulation, and as observed in data for 131 GeV converted photons in the endcap region of the calorimeter. The lower panel shows the ratio between the simulation and the data.

While not optimal, the application of both models to the simulation of electron-induced showers is shown to be beneficial for the simulation of most of the variables, especially in the endcap region. In the future, the model parameters for electron showers could be re-optimized based on isolated electrons collected from data in $Z \rightarrow ee$ and $J/\psi \rightarrow ee$ events with dielectron triggers. However, the bias on the shower shapes as a result of misidentified prompt electrons remains to be understood, and a dedicated background estimation might be necessary.

The model parameters presented in this chapter are optimized for a total of six parametrization points employed by FASTCALOSIMV2, covering the photon energy range between 32 GeV to 1 TeV. While the parameters generally do not exhibit a strong energy dependence, the simulation of showers initiated by lower energy particles might benefit from a dedicated optimization, especially for middle layer shower shapes. The inclusive photon sample collected by single photon triggers in data is limited by the prescaling of single photon triggers in the low energy regime and by event yields at high energies. However, photons originating from radiative $Z \rightarrow \ell^+ \ell^- \gamma$ decays collected using unprescaled single-lepton and dilepton triggers with low momentum thresholds could be used to extend the optimization of the models to photons with energies as low as 10 GeV. A small background contamination is expected from $Z \rightarrow \ell\ell$ decays, but the purity of such a sample is very high ($\gtrsim 95\%$) [133]. Photons from radiative Z boson decays up to approximately 100 GeV can be used, beyond which event yields are insufficient, such that the high photon purity can also prove beneficial in this energy regime.

The model errors for specific incident photon directions can be potentially improved further with a specific optimization of the search space in the regions of interest. This is especially the case for neighbouring $|\eta|$ regions, where model parameters are chosen from the barrel- and endcap-optimized search spaces and show significant discontinuities. In some cases, such as for $|\eta| \approx 1$ in the middle layer model and close to the transition region between barrel and endcap in the strip layer model, a drop to values below 10 in the χ^2 ratio is observed (see Figure 8.15). The degradation of the model performance in these regions can be potentially recovered with a dedicated search space optimization.

Finally, the derived models and scalings required to match the photon shower shapes observed in data might provide useful insights and possible explanations on the mismodelling observed in the GEANT4 implementation in ATLAS. For instance, in the central region, a modest scaling in the order of $\eta_s \sim 1 - 3\%$ is needed for the middle layer, while a large widening of the showers by $\eta_s (\Delta\eta = 0) \sim 25\%$ is required in the strip layer. Furthermore, the optimal scalings increase significantly in the forward direction, starting from approximately $|\eta| \approx 0.7$ in both the strip and middle layer shower shapes. For the strip layer, the η_s scaling factor in the shower centre is approximately 75 – 85 % in the forward region, while the η_s increases to up to 20 % in the middle layer. The very different widening of the showers that is required within the two calorimeter layers and the increase of scaling factors at around $|\eta| \approx 0.7$ in both models appears very indicative of an underlying systematic issue in the GEANT4 simulation as currently implemented in the ATLAS simulation framework.

FastCaloSim as a Geant4 Fast Simulation Model

Together with `GEANT4`, `FASTCLOGAN`, and other fast simulation tools, `FASTCALOSIMV2` is currently embedded in the Integrated Simulation Framework (ISF) [162]. The framework is a powerful tool that allows combining various full and fast simulation approaches, in a way that different simulators can be called depending on the type of incoming particle and the detector region they enter. At the same time, ISF implements complex particle routing algorithms that take into account user-defined simulation selection rules, which may or may not change for each simulated hard scattering event, and become increasingly complex as new simulators and their associated detector regions are added.

On a technical level, a wrapper Athena service known as `FastCaloSimV2Svc` is employed that receives particles from the `ISF SimulationKernel` in case that they are directed to be simulated using the parametrized calorimeter simulation. Within the service, the `FASTCALOSIMV2` parametrization is called for each incoming particle which then returns a dedicated simulation state that holds all the registered energy deposits in the calorimeter system. The energy deposits in the individual cells are stored in a reconstruction level data format known as the `CaloCellContainer`, which remains in transient memory and is only recorded at the end of each Athena event. After each event, the `CaloCellContainer` is converted into the `HITS` data format, which corresponds to the format used by `GEANT4` and ensures compatibility to the subsequent digitization and reconstruction steps.

Over the past few years, ISF has accumulated a significant amount of technical complexity, making it increasingly difficult to maintain. At the same time, recent advances in `GEANT4` may allow to fully replace ISF and entirely outsource its functionality as a particle stack dispatcher to the `GEANT4` toolkit. This approach would greatly simplify the simulation workflow in ATLAS,

eliminate the heavy maintenance overhead generated by ISF, and ensure a good maintainability of the simulation infrastructure for Run 4 and beyond.

Furthermore, any performance overhead introduced through the particle routing mechanisms implemented in ISF could potentially be avoided, as particles are handled within a single, highly optimized simulator and do not need to be routed through complex algorithms. In fact, ATLAS has observed a $\sim 10\%$ simulation time overhead introduced by the ISF implementation of GEANT4 compared to the pure, full GEANT4 simulation.

Finally, the implementation of existing simulators with a direct interface to the GEANT4 toolkit is a first step towards enabling the usage of these tools by actors outside the ATLAS collaboration. This is particularly relevant for FASTCALOSIMV2, which is not intrinsically experiment dependent and may be of interest to other experimental endeavours.

In this chapter, the need for an alternative implementation of the simulation infrastructure is further motivated in Section 9.1 and a first implementation of FASTCALOSIMV2 as a GEANT4 fast simulation engine is presented. The general concepts of GEANT4 fast simulation models are introduced in Section 9.2 and the new implementation of FASTCALOSIMV2 in GEANT4 is described in Section 9.3 and Section 9.4. Performance studies of the new implementation are presented in Section 9.5.

9.1 Complexity Development

The Integrated Simulation Framework is the main component of the overall simulation infrastructure that is used by ATLAS and consists of approximately 43 thousand lines of code (LoC). As a critical component to the success of the ATLAS physics programme, it is important to ensure its maintainability so that new developments can be effectively integrated without a significant technical overhead and without a continuous increase in complexity.

Figure 9.1 shows a hierarchical visualization of the complete simulation infrastructure, including an expanded view into the code structure of ISF. Each circle represents a folder structure in the codebase, with the radius scaled to the lines of code in the respective component. The coloured circles represent individual files and the colours encode the overall health of the code. Green represents healthy and easily maintainable code, yellow represents parts of the code with increased maintenance needs, and red represents unhealthy, high-risk code with significant problems. The metric from Ref. [221] takes into account a broad variety of factors such as nested complexity and developer congestion.

It becomes apparent that ISF is the largest and most complex component in the simulation infrastructure, followed by a package for GEANT4-related utilities (`Geant4Utilities`) and the pure, ISF-independent implementation of GEANT4 in ATLAS (`G4Atlas`). Within ISF, the FASTCALOSIMV2 implementation makes up the largest part of the codebase, followed by the ISF core implementation and the GEANT4 and FATRAS implementations within

the Integrated Simulation Framework. Generally, the codebase dedicated for the core ISF implementation, its configuration and simulator selection, represents a significant part of the overall ISF implementation, including all the different simulators used, and thus represents a major technical burden for the overall maintenance of the simulation infrastructure.

In addition, the overall maintainability is affected by the general code health, encoded in the colours of the individual files and reveals that a relevant fraction of the codebase within the individual implementation of the simulators as well as in the pure ISF related parts of the codebase is in need of increased maintenance. A direct integration of the simulators within `GEANT4` would eliminate the ISF core implementation and configuration, including all dedicated particle routing algorithms, and potentially simplify the configuration of the individual simulators.

A metric commonly used to measure code complexity is related to the number of linearly independent paths through the source code of a program, known as *cyclomatic complexity* [222], and is useful for assessing the maintainability of a given codebase. Figure 9.2(a) shows the cyclomatic complexity of methods in Athena, the reconstruction codebase, the complete simulation infrastructure and the Integrated Simulation Framework. Cyclomatic complexity is typically divided into four different categories: methods with cyclomatic complexities between 1 and 10 are considered simple and well-maintainable, cyclomatic complexities between 11 and 22 indicate more complex methods with moderate maintainability risk, a score between 21 and 50 indicates very complex, high-risk code with significant problems. Cyclomatic complexities above 50 usually indicate completely untestable code with an extremely high maintenance burden.

For all the components examined, the vast majority of methods (> 90%) fall into the first category indicating good maintainability. Most of the technical complexity is caused by a relatively small number of very complex methods. For ISF, a total of 82 methods (~ 3.24%) fall into the second category with cyclomatic complexity values between 10 and 20, requiring increased maintenance. A total of 27 methods (~ 1.07%) have cyclomatic complexities between 20 and 50 and are therefore extremely hard to maintain. Finally, a total of 7 methods, some of which could be removed entirely in an ISF-independent implementation, have even higher complexity values and should be considered untestable. The relative distributions in each category are reasonably similar for Athena as well as for the reconstruction and simulation codebases and are shown in Table 9.1.

Another metric of interest is the complexity evolution of the codebase over time. Figure 9.2(b) shows the indent-based complexity evolution [223] from late 2014 to late 2022 for the ATLAS reconstruction as well as for the entire simulation infrastructure and the ISF codebase. Indent-based complexity uses only indentation as a lightweight and language independent measure of complexity and is found to correlate with more traditional measures of complexity. In all cases, the complexity is found to increase linearly over time as shown by the regression analysis performed. The reconstruction component features the steepest slope, followed by the Simulation and ISF components. However, the relative increase in complexity is greatest for the

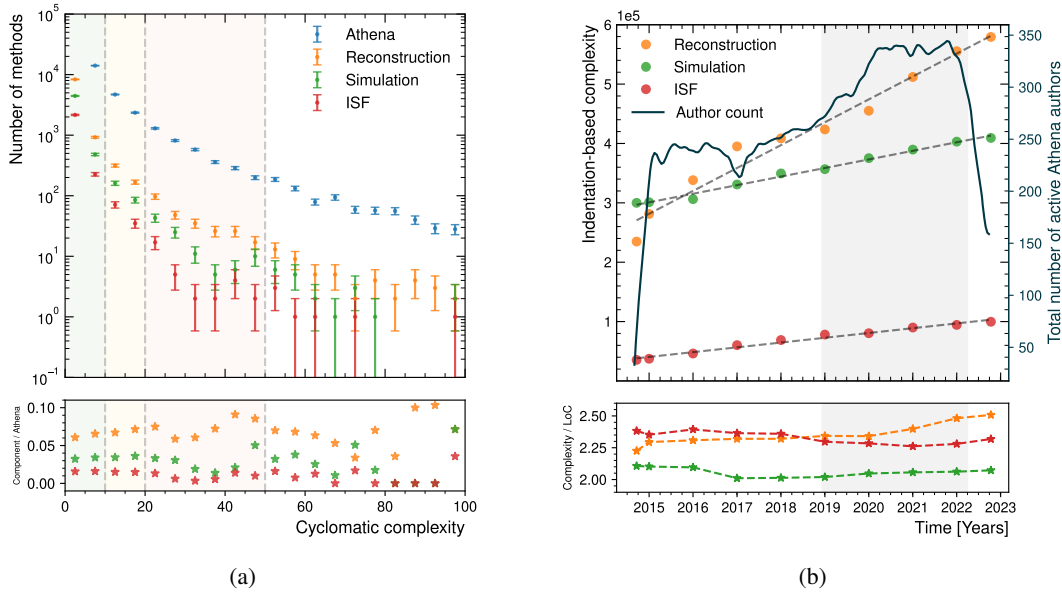


Figure 9.2: (a) Cyclomatic complexity of methods in Athena, the reconstruction codebase, the complete simulation infrastructure and the Integrated Simulation Framework. The dashed vertical lines show thresholds typically used to categorize cyclomatic complexity, with lower values indicating healthy, maintainable code and higher values indicating complex, very difficult or impossible to maintain code. The bottom panel shows the ratio of the complexity of each subsystem to the total complexity of Athena. (b) Evolution of the indentation-based complexity of the Integrated Simulation Framework, as well as of the entire ATLAS simulation and reconstruction infrastructure, from the end of 2014 to late 2022. The dashed grey lines show linear fits performed to the complexity evolution. The blue solid line shows the evolution of active Athena contributors over the same period. The grey vertical band shows the time span of LS2, a period of increased authorship. The bottom panel shows the ratio of complexity to Lines of Code (LoC) to highlight non-linearities in complexity growth relative to the overall growth of the codebase.

Complexity	Athena	Reconstruction	Simulation	ISF
0 < CC < 10	92.91	92.12	93.00	94.27
10 < CC < 20	3.56	4.01	3.68	3.24
20 < CC < 50	1.98	2.22	1.59	1.07
CC > 50	0.57	0.57	0.43	0.28

Table 9.1: Relative distribution in percent of methods in the various categories of cyclomatic complexity (CC) for Athena and the different components of interest.

ISF: its complexity almost tripled from the end of 2014 to late 2022 by about 181%, compared to an increase of 147% for the reconstruction codebase and 36.4% for the entire simulation infrastructure. For reference, the overall relative complexity of Athena has increased by only 79.5% over the same period. The ratio of complexity to lines of code has remained roughly constant for ISF and the overall simulation infrastructure, but has increased in recent years for the reconstruction component. Non-linearities in the growth of complexity compared to the overall growth of the codebase are in many cases an indication of deteriorating code quality and thus have a negative impact on its maintainability. Figure 9.2(b) also shows the evolution of active Athena authors over time. An active author is defined as any unique individual who has contributed to the Athena repository in a three-month sliding window at each point in time. During Run 2, the number of active authors has remained relatively stable at around 250. During LS2 and in preparation for Run 3, ATLAS has experienced a period of increased authorship, with a peak of around 350 authors, which has fallen dramatically since the start of Run 3.

In summary, ISF is a sophisticated simulation framework that allows a flexible combination of different simulators in different detector regions. At the same time, ISF is inherently complex and has experienced a disproportionate growth in relative complexity compared to other components and Athena as a whole. Historically, the complexity of the components has grown linearly over time, while the total number of Athena authors has remained fairly constant at around 250, with a temporary increase during LS2. Undoubtedly, a codebase that continues to grow in complexity becomes increasingly difficult to maintain, especially if the number of contributors remains roughly constant. Reversing this trend is a non-trivial undertaking, but may be achievable through major structural refactorings. As the main component of the ATLAS simulation infrastructure, the retirement of ISF and the integration of its functionality directly into GEANT4 is expected to have a significant impact on the overall complexity and maintainability of the simulation infrastructure of the ATLAS collaboration, and to greatly simplify the integration of future fast simulation developments.

9.2 Fast Simulation in Geant4

The main purpose of GEANT4 is to simulate the detector response of incident particles with the highest possible accuracy. However, as a simulation toolkit, GEANT4 also offers the possibility to implement arbitrary fast simulation models and therefore has the potential to replace the functionality of ISF.

To implement parametrizations, GEANT4 allows to shortcut its full simulation by attaching fast simulation models (or any other arbitrary model) to specific volumes or sets of volumes. A particle entering such a volume is intercepted by the attached fast simulation model and treated according to the implemented parametrization. Within the respective model, the properties of the particle can be arbitrarily manipulated. In fact, no GEANT4 transport takes place within

the specific region, such that the model is required to describe exactly how the particle is to be manipulated, e.g. where the particle travels to, what energy is deposited within the sensitive detectors, how the momentum of the particle develops and what secondary particles are produced. Figure 9.3 shows a sketch showing the principle of the full GEANT4 simulation and a possible configuration using the GEANT4 fast simulation framework, for a charged particle entering the detector system. In the full simulation, the detector response is simulated at each

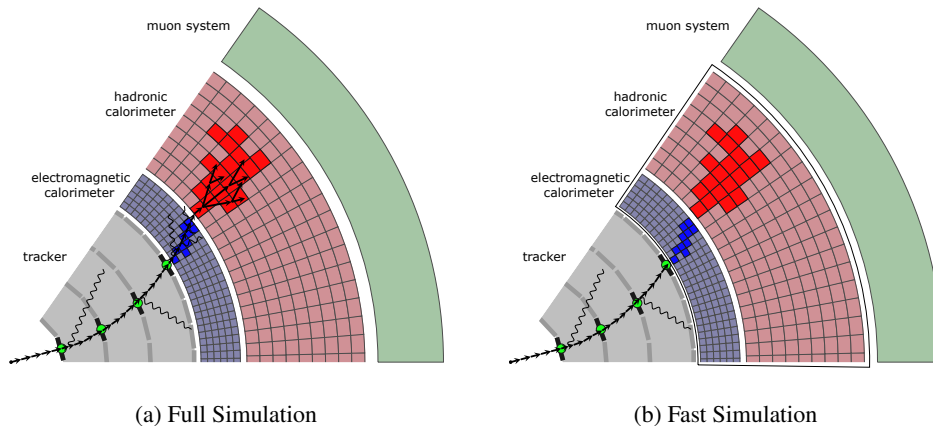


Figure 9.3: (a) Sketch showing the principle of the full GEANT4 simulation. A charged particle enters the detector and is tracked with small steps along its trajectory in the detector. The resulting energy deposits are registered in the sensitive detectors. (b) Illustration showing a possible fast simulation configuration within GEANT4. A charged particle is fully simulated in the tracker. Once it reaches the envelope around the calorimeter system, the particle is intercepted and an arbitrary parametrization can be invoked. Figures modified from Ref. [224].

small GEANT4 step and propagated through the tracker to the calorimeter system, where the resulting energy deposits are registered in the sensitive detectors. In the example of a possible fast simulation configuration, the calorimeter region, which includes the EM and hadronic calorimeter system, is attached to a fast simulation model. The particle is simulated with GEANT4 in the tracker, but is intercepted at the envelope around the calorimeter system and then passed to the attached parametrization, which generates the respective energy deposits independently of GEANT4. The attachment of fast simulation models is by no means limited to the calorimeter system and can be configured flexibly.

On a more technical level, fast simulation models are attached to `G4Region` objects, which can be identified as the envelopes enclosing a physical volume or set of volumes. The attached parametrization is active for all daughter `G4LogicalVolume`'s held by the respective `G4Region`. Concrete fast simulation models are bound to the envelope by a `G4FastSimulationManager` instance, which can hold any number of fast simulation models. The interface between the tracking and the fast simulation is provided by a `G4FastSimulationManagerProcess` which invokes the parametrization models when certain trigger conditions are met (see Section 9.3)

and is part of the process list associated to the particle that ought to be parametrized. When a `G4Track` is transported inside the envelope, the `G4FastSimulationManagerProcess` looks for the existence of a `G4FastSimulationManager`. If it exists, the fast simulation models held by the `G4FastSimulationManager` check if the appropriate trigger conditions are met to invoke the parametrizations. If this is the case, the model is applied at the current position of the `G4Track`. In the contrary case, the tracking proceeds with a nominal `G4Step` within the envelope. Figure 9.4 shows a simplified illustration of the described fast simulation mechanism embedded within `GEANT4`.




Figure 9.4: Simplified illustration of the `GEANT4` fast simulation mechanism. Components embedded in the blue dashed squares are specific to the fast simulation infrastructure, while the red squares indicate generic simulation components. Detector illustration modified from Ref. [224].

In some cases, the definition of suitable envelopes based on the existing tracking geometry may not be trivial or even possible. For example, one may want to group certain detector components together to define an envelope. However, the respective envelope may overlap with the existing tracking geometry (*geometry clashes*), which may lead to undefined behaviour during the simulation process. To avoid this problem, `GEANT4` provides the concept of *parallel* or *ghost* geometries [225]. These ghost geometries provide independent navigation managed by the `G4FastSimulationManagerProcess` and can therefore be placed at arbitrary positions within the existing tracking geometry. In addition, ghost volumes can be sensitive to the particle type, such that these volumes are invisible to particles that should not to be parametrized in a certain region. However, the additional navigation can cause a significant overhead in simulation

time. For many applications, where only a few ghost volumes are used, the additional geometry calculations may be tolerable. The employment of parallel geometries in ATLAS, however, would require the use of large sets of ghost volumes, which would introduce a performance overhead that runs counter to the very principle of fast simulations. The use of ghost volumes in ATLAS is therefore ruled out for the purpose of fast simulation, and other means of defining the appropriate trigger regions must be found. The choice of these regions for the FASTCALOSIMV2 interface to GEANT4 is discussed in Section 9.4.

9.3 Interfacing FastCaloSimV2 with Geant4

The concrete fast simulation model that is added to the `G4FastSimulationManager` and attached to a `G4Region` is implemented by inheriting from the `G4VFastSimulationModel` class. The abstract class defines a set of three pure virtual methods that need to be implemented to provide information about what static and dynamic conditions an incident particle must satisfy to trigger the parametrization, and what exactly should happen in place of the detailed simulation. Listing 1 shows the basic structure of the abstract `G4VFastSimulationModel` class, which contains the three pure virtual methods `IsApplicable`, `ModelTrigger` and `DoIt` used to define the fast simulation model.

```
class G4VFastSimulationModel
{
    public:
        G4VFastSimulationModel(const G4String& aName);

        // Is the model applicable for the incident particle type?
        virtual G4bool IsApplicable(const G4ParticleDefinition&) = 0;
        // Model triggered for kinematic properties of particle?
        virtual G4bool ModelTrigger(const G4FastTrack&) = 0;
        // Invoke the parametrisation of choice
        virtual void DoIt(const G4FastTrack&, G4FastStep&) = 0;
}
```

Listing 1: Simplified structure of the abstract `G4VFastSimulationModel` class with the three main pure virtual methods `IsApplicable`, `ModelTrigger` and `DoIt`, which is used to interface the FASTCALOSIMV2 parametrization with the GEANT4 toolkit.

The implementation of the three methods used to interface FASTCALOSIMV2 with GEANT4 is as follows:

1. `IsApplicable` checks that the static conditions of the fast simulation model are satisfied. `G4ParticleDefinition` provides access to all relevant intrinsic particle properties

such as mass, charge, spin, etc. In the `FASTCALOSIMV2` implementation, the method returns `true` if the incident particle is a photon, electron, positron or hadron. Particles that do not deposit energy in the calorimeter system, such as muons and neutrinos, are considered not applicable.

2. `ModelTrigger` is only called if the parametrization is applicable and checks whether the dynamic conditions of the model are fulfilled. `G4FastTrack` provides access to the current `G4Track` and thus to all relevant kinematic particle properties such as position, momentum and energy. In the `FASTCALOSIMV2` implementation, all particles are required to have a kinetic energy of at least 50 keV. This selection is implemented in order to match the ISF behaviour, where particles simulated with `GEANT4` that have energies below this threshold are not returned to the `ISF ParticleBroker`. In addition, pions are required to have $E_{\text{kin}} > 200$ GeV and other hadrons $E_{\text{kin}} > 400$ GeV. This is to match the kinetic energy thresholds of low-energy hadrons implemented in `ATLFAST3` and motivated in Chapter 4. Finally, an additional boundary check is performed to ensure that the model is only triggered when the incident particle is positioned on the boundary between the ID and the calorimeter system, where an invocation of `FASTCALOSIMV2` is well-defined. A detailed discussion of why this is presently necessary can be found in Section 9.4.
3. `DoIt` implements the fast simulation model and is only invoked when the model is triggered. The `G4FastTrack` provides all input information of the incident particle required to apply the model. The final state of the particle after the `G4Step` must be returned by the `G4FastStep` reference. For the `FASTCALOSIMV2` implementation, the entire simulation of the calorimeter response is handled by `FASTCALOSIMV2`, so that instead of providing information to `G4FastStep`, the particle is killed at the very end, such that no further `G4Steps` for the primary are performed within the envelope. Incoming particles with $E_{\text{kin}} < 10$ MeV are killed immediately and not simulated with `FASTCALOSIMV2`, in order to mimic the current behaviour within ISF. For all other particles, their properties, including their position and momentum from the `G4FastTrack`, are passed to `FASTCALOSIMV2` specific data types. The particle is then transported through the calorimeter system using `FASTCALOSIMV2`'s own extrapolation engine and finally the detector response is simulated using the `FastCaloSimV2Svc` wrapper Athena service. The currently used extrapolation engine (see Chapter 6 for more details) causes a heavy coupling to the Athena tracking tools. In the future, the transport through the calorimeter system is expected to be replaced by a pure `GEANT4`-based approach. Similar to ISF, the resulting energy deposits are recorded in a `CaloCellContainer` which remains in transient storage until the end of each Athena event, when it is finally converted to the `GEANT4`-specific HITS data format to allow for compatibility in the later digitization and reconstruction stages. More specifically, the conversion is performed by writing out hits to the corresponding sensitive detectors, taking into account their respective sampling fractions. In the multithreaded version of Athena, `G4VFastSimulationModel`

instances are not necessarily unique to Athena threads, such that the cell energies cannot be recorded directly into the `CaloCellContainer` from within the `DoIt` method. To ensure thread-safety, the registration of energy deposits in the `CaloCellContainer` is outsourced to a new `G4VSensitiveDetector` called `CaloCellContainerSD`, which is guaranteed to be unique to an Athena thread.

9.4 Choice of Trigger Volumes

Within ISF, particles are intercepted and passed to `FASTCALOSIMV2` at the boundary between Inner Detector and the calorimeter system that is used to derive the parametrization (see Figure 4.5), hereafter referred to as the ISF or parametrization boundary. The particle is then transported through the detector based on its momentum at the position it is handed over to `FASTCALOSIMV2`. The parametrization is chosen depending on the extrapolated position of the particle on the boundary between the ID and the calorimeter system (see Chapter 6). Therefore, to ensure a correct choice of the parametrization and thus a correct simulation, the incident particles must be intercepted at the ISF boundary or at a reasonably close distance to the ISF boundary. Furthermore, if a particle is handed over only after it has traversed some material in the calorimeter system, superfluous `GEANT4` hits are created, which are expensive to simulate and affect the overall computational performance. Therefore, the fast simulation model described in Section 9.3 is ideally attached to an envelope that exactly follows the ISF boundary. However, such an envelope does not exist in the Run 3 geometry of the ATLAS detector, so different approaches have to be investigated.

The purpose of this section is to motivate the choice of fast simulation trigger volumes using the ATLAS Run 3 geometry with the goal of reproducing the simulation of the ISF configuration as closely as possible. While this does not necessarily represent the final implementation that may be used for the production of physics samples in the future, a reproduction of the ISF-integrated simulation ensures an understanding of the effects that may alter the simulation in undesirable ways in the novel `GEANT4` implementation.

The ATLAS Run 3 geometry implements a calorimeter envelope referred to as `CALO: :CALO`, which encloses the entire ATLAS calorimeter system and approximately follows the ISF boundary. A three-dimensional illustration of the envelope is shown in Figure 9.5. The red point cloud indicates the position of the ISF boundary where particles are ideally intercepted by `FASTCALOSIMV2`. In principle, `CALO: :CALO` can be used directly to attach the fast simulation model. However, since the envelope does not exactly follow the ISF boundary, some features may lead to different simulation results compared to the nominal `FASTCALOSIMV2` implementation:

1. The `CALO: :CALO` envelope as defined in the Run 3 geometry encloses the Minimum Bias Trigger Scintillators (MBTS) [226] of the ATLAS detector. The MBTS, consisting of a set of 16 plastic scintillators, is mounted on the inner face plate of the EM calorimeter endcaps and is employed as trigger for selecting minimum bias events. As such, the

MBTS is not used for energy measurements and is therefore not taken into account in the definition of the ISF boundary, i.e. particles in FASTCALOSIMV2 are parametrized only after passing through the MBTS. In addition, radiation shielding components, in particular the Moderator Shielding (JM) disk [111], which protects the Inner Detector from neutrons back-splash from the calorimeter, are also placed inside the ISF boundary. If the CALO::CALO envelope is used as a direct fast simulation trigger, incoming particles in the endcap region would be intercepted slightly earlier. In the ISF implementation, a particle passing through the MBTS and the JM Disk is simulated with GEANT4. Material interactions may produce very low-energy secondary particles, many of which will fall below the ISF-imposed $E_{\text{kin}} < 50$ keV threshold required to invoke fast simulation tools, and consequently will produce GEANT4 hits in the calorimeter system. Therefore, using CALO::CALO as a fast simulation trigger will result in an artificial speed-up of the simulation caused by the early interception of particles by FASTCALOSIMV2 (see Section 9.5) and a lower number of GEANT4 hits generated in the endcap region of the calorimeter.

2. The CALO::CALO envelope does not fully enclose the ISF boundary in the FCal region around the beam pipe, creating a gap referred to as the *no trigger region*. Particles travelling in the very forward direction in a window of approximately $4.18 < |\eta| < 4.28$ will travel slightly further than the ISF boundary before being intercepted by FASTCALOSIMV2, such that the simulation within the gap volume is performed using GEANT4. However, no material is associated with the gap volume, so no material interactions are simulated, and no additional hits are generated. Furthermore, the proprietary FASTCALOSIMV2 extrapolation algorithm (see Chapter 6) should to ensure that the correct parametrization is chosen even in these corner cases. All in all, the gap volume should not alter the simulation significantly.
3. Particles that punch-through the calorimeter system may generate secondary particles that may backscatter into the calorimeter. In this case, the particles would be intercepted by the fast simulation model at the outer boundary of the CALO::CALO envelope, leading to ill-defined FASTCALOSIMV2 calls. In general, this is expected to reduce the number of GEANT4 hits compared to the ISF implementation.

To assess the simulation similarity between the nominal FASTCALOSIMV2 implementation and its GEANT4 counterpart, the properties of the simulated hits are compared. Figure 9.6(a) shows the simulated hit energies from single photon events with varying incident photon energies and directions for the full GEANT4 simulation and FASTCALOSIMV2 as implemented in ISF and GEANT4. The GEANT4 simulation generates many low energy hits and the energy distribution falls quickly and smoothly at higher energies. In the case of FASTCALOSIMV2, the hit energies show a periodic peak structure as a result of the discrete energies assigned to the hits (see Section 4.4) and extend over much larger ranges. The energy distributions for the ISF and GEANT4 implementations of FASTCALOSIMV2 are fairly similar, but show some non-negligible




Figure 9.5: Three-dimensional representation of the CALO : CALO envelope from the ATLAS Run 3 geometry (grey) and the boundary between calorimeter system and Inner Detector as defined within ISF (red). The upper panel shows a zoom into the endcap region: the MBTS and the moderator shielding protrude into the ISF boundary. The bottom panel shows a zoom into the FCAL region: the envelope does not completely enclose the boundary defined within ISF, that is, there exists a *gap* between the envelope and the ISF boundary referred to as *no trigger region*.

differences of up to 4%.

Since very low energy particles are not passed to FASTCALOSIMV2, a small number of GEANT4 hits are simulated when using FASTCALOSIMV2 as a simulator. To better study the effects resulting from the two different implementations, it is useful to split the simulation into hits directly generated by FASTCALOSIMV2 and hits generated by GEANT4. Figure 9.6(b) shows the hit energy distributions of FASTCALOSIMV2 hits (top) and GEANT4 hits (bottom) as generated by the GEANT4 implementation of FASTCALOSIMV2 and the nominal FASTCALOSIMV2 implementation in ISF. The number of FASTCALOSIMV2 hits is approximately compatible in both implementations, while the number of GEANT4 hits is systematically lower in the case of the GEANT4 implementation. The overall ratio of the number of hits generated by GEANT4 to

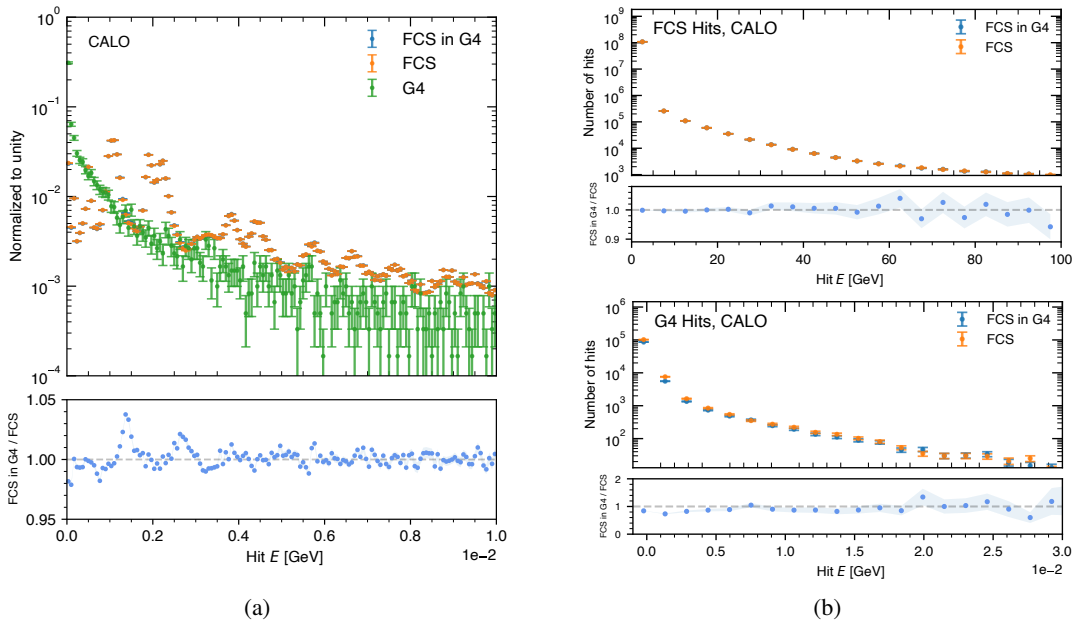


Figure 9.6: (a) Simulated hit energies from single photon events for the GEANT4 implementation of FASTCALOSIMV2 (FCS in G4), the nominal FASTCALOSIMV2 implementation in ISF (FCS) and for the full GEANT4 simulation. Note that energies created by FASTCALOSIMV2 hits extend to significantly larger energies than the shown range. (b) Simulated hit energies for the GEANT4 and ISF implementation of FASTCALOSIMV2 split into energies deposited by FASTCALOSIMV2 (top) and GEANT4 (bottom). In many cases the dots overlap entirely, but differences become apparent in the lower pad, which shows the ratio between the distributions in the ISF and GEANT4 implementation of FASTCALOSIMV2. In both cases, the fast simulation model is triggered for particles that reach the CALO : : CALO envelope.

the number of hits generated by FASTCALOSIMV2 is 0.105% for the nominal implementation and about 0.088% when using CALO : : CALO as the triggering volume, effectively indicating a small mismatch in both implementations. Since the ISF boundary does not exactly match the CALO envelope, these small discrepancies are to be expected.

In order to guarantee the reproduction of validated physics results and to align both configurations exactly, the following three approaches were considered:

1. The introduction of a very thin vacuum volume that exactly follows the ISF boundary in the ATLAS geometry. The additional thin volume would force the `GEANT4` simulation to stop exactly at the ISF boundary, where the particle can be intercepted by `FASTCALOSIMV2`. Although much less likely, backscattered particles from the Muon Spectrometer may still reach the ID-Calo boundary, and as such an additional check may be necessary to ensure that the particle originates from the origin. However, the introduction of such a volume is not possible without creating geometry clashes, which can lead to undefined simulation behaviour, and is therefore disfavoured. Alternatively, only the gap along the beam pipe can be closed with additional vacuum volumes and without introduction new geometry clashes. Since the small gaps are not expected to alter the simulation significantly, a change in the ATLAS geometry and the subsequent maintenance burden is considered unnecessary.
2. Use the concept of `GEANT4` ghost geometries to define an envelope that groups together the calorimeter system according to the ISF boundary. Physics-wise, the correct simulation would be inherently ensured, similar to the first approach. However, as discussed in Section 9.2, the resulting performance overhead due to the independent navigation is expected to defeat the very principle of a fast simulation, such that this option is also disfavoured.
3. Use the `CALO::CALO` envelope or a combination of existing volumes implemented within the Run 3 geometry with additional `ModelTrigger` requirements that ensure that the fast simulation is triggered only when the particle is positioned at the ISF boundary and its momentum direction points from the ID outwards to the calorimeter system. Ideally, one chooses the minimum set of volumes required to cover the ISF boundary in order to minimize unnecessary `ModelTrigger` calls.

The first two options are excluded for the aforementioned reasons, so that the choice of trigger volumes is constrained by the existing Run 3 geometry. As a first approach, the two positional and directional requirements in `ModelTrigger` are introduced for `CALO::CALO` as a fast simulation trigger, which should ensure that `FASTCALOSIMV2` is only called once a particle has crossed the MBTS and avoid any ill-defined fast simulation calls resulting from backscatter events. However, within the Run 3 geometry, the `CALO::CALO` envelope is not implemented as the root volume of all enclosed sub-volumes that may directly border the ISF boundary, such that particles cannot be intercepted by `FASTCALOSIMV2`. For example, if a particle passes through the endcap region, the fast simulation model will not be triggered within the MBTS disk as a result of failing the positional check in `ModelTrigger`. The next sub-volume that the particle traverses is referred to as `LAr::Endcap::Cryostat::Cylinder::Mixed` and exactly bounds the ISF boundary. However, this volume is not a daughter volume of `CALO::CALO`, such that the fast simulation model is not active in that region. As a result, the particle will continue to be simulated with `GEANT4`, generating redundant and computationally expensive `GEANT4` hits.

These problems can potentially be avoided by using an explicit list of volumes to trigger the fast simulation. As a minimum requirement, the volumes must cover approximately the entire ISF boundary. Figure 9.7 shows a three-dimensional representation of the volumes which are enclosed by the `CALO::CALO` envelope in the barrel region. The ISF boundary is completely covered by the `LAr::Barrel::Cryostat::MotherVolume`, except for a very small gap in the transition region to the endcap. In principle, one can use the underlying daughter volumes, such as `LAr::Barrel::Cryostat::InnerWall`, but additional gaps would be created along the ISF boundary, such that this option is disfavoured. Similarly, the same volume decomposition is shown in the endcap region in Figure 9.8. In this case, `LAr::Endcap::Cryostat::MotherVolume` also covers the same part of the ISF boundary as the `CALO::CALO` envelope. However, the decomposition into its daughter volumes `LAr::Endcap::Cryostat::Cylinder::Mixed` with the addition of the `ModeratorTube` provides the same ISF boundary coverage and a natural way to exclude the MBTS, i.e. without additional `ModelTrigger` calls, such that this option is preferred.

Based on these considerations, three different choices of fast simulation trigger volumes are studied and the resulting simulation at hit level is compared to the nominal implementation of `FASTCALOSIMV2` in ISF:

1. `CALO`: Usage of the calorimeter envelope as sole fast simulation trigger, which is regarded as the baseline scenario.
2. Explicit: An explicit list of volumes consisting of
 - a) `LAr::Barrel::Cryostat::MotherVolume`
 - b) `LAr::Endcap::Cryostat::Cylinder::Mixed`
 - c) `ModeratorTube`,

that provides approximately the same ISF boundary coverage as `CALO`, but naturally excludes the MBTS.

3. `CALO & Explicit`: The use of `CALO::CALO` and the explicit set of volumes as fast simulation trigger, motivated by the fact that not all volumes enclosed in the calorimeter envelope are implemented as daughter volumes. With the inclusion of a positional and directional check, this option should inherently result in the best alignment with the ISF configuration, albeit with the largest number of superfluous `ModelTrigger` calls.

All three configurations are examined with and without an additional positional and directional check in `ModelTrigger`. Table 9.2 shows the ratio between the number of hits generated by `GEANT4` and the number of hits generated by `FASTCALOSIMV2` for the different configurations tested and is a good indication of the compatibility of the configuration with the ISF implementation. As noted above, the hit ratio for the baseline `CALO` option without boundary check is lower than in the nominal implementation in ISF. This is the result of the lower number of `GEANT4` hits generated through the early interception of particles entering the MBTS and the interception of backscattered particles by `FASTCALOSIMV2`. The introduction

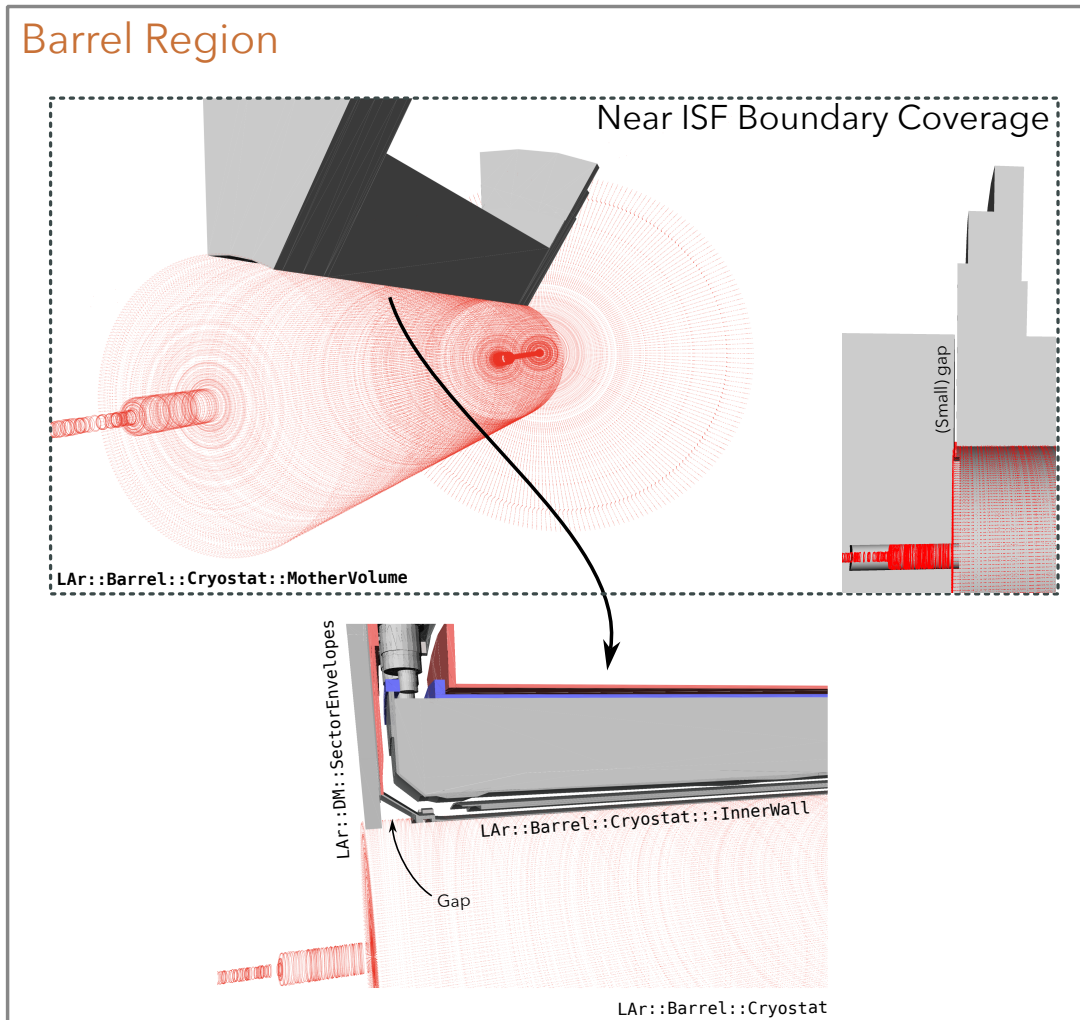


Figure 9.7: Three-dimensional representation motivating the choice of trigger volumes in the barrel region. The barrel mother volume provides near ISF boundary coverage, except for a small gap in the transition region between the barrel and endcap volumes. The explicit usage of daughter volumes of the barrel as fast simulation trigger would create additional gaps along the barrel.

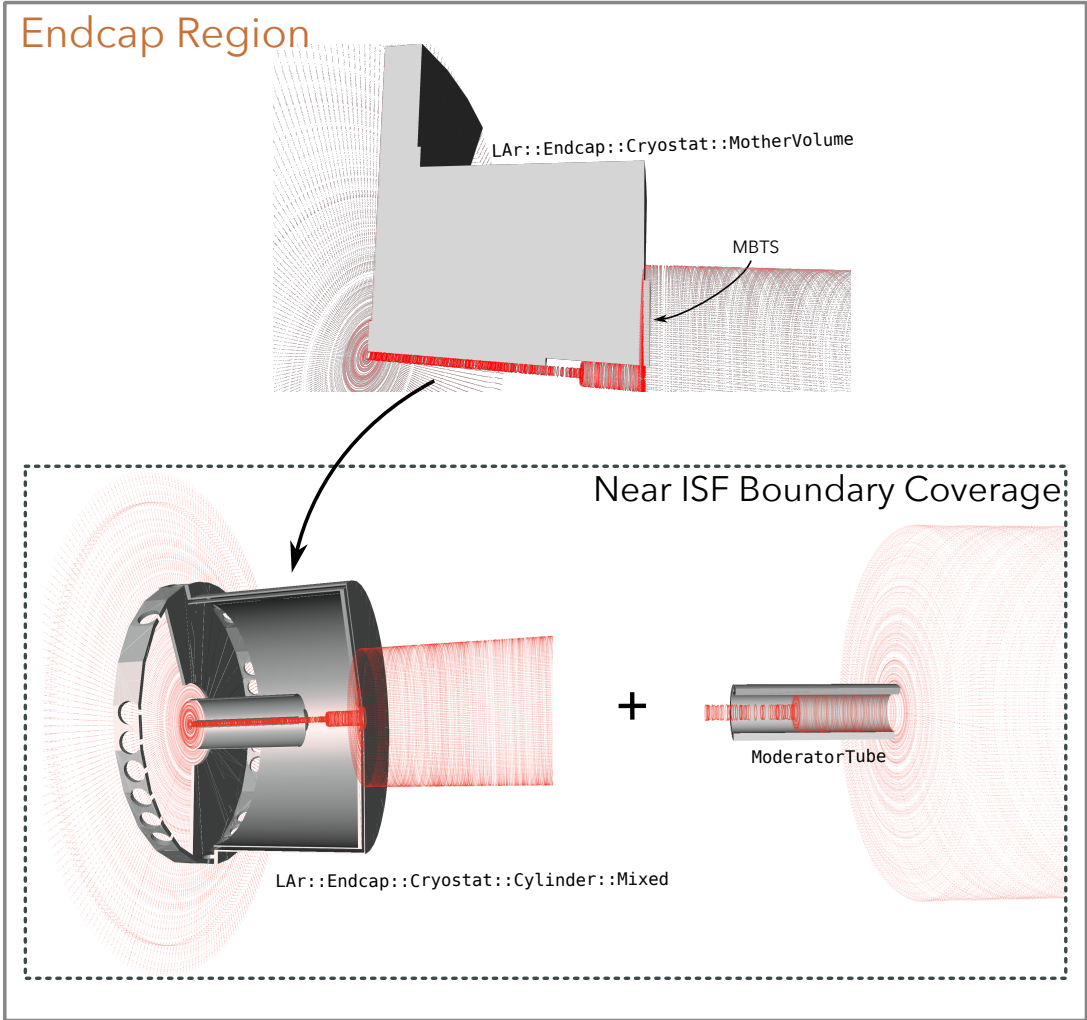


Figure 9.8: Three-dimensional representation motivating the choice of trigger volumes in the endcap region. The endcap mother volume provides near ISF boundary coverage, except for a gap along the beam pipe, but includes the MBTS that protrudes into the ISF boundary. The MBTS can be excluded by using the endcap daughter volumes and the moderator tube as trigger envelopes, while keeping the same ISF boundary coverage.

$N_{\text{hits}}^{\text{G4}} / N_{\text{hits}}^{\text{FCS}}$ [%]	CALO	Explicit	Explicit & CALO	Nominal (ISF)
w/o boundary check	0.088	0.100	0.088	0.105
with boundary check	0.121	0.110	0.105	

Table 9.2: The ratio between number of hits generated by GEANT4 and the number of hits generated by FASTCALOSIMV2 for the different tested configurations of the FASTCALOSIMV2 implementation in GEANT4. The configuration with the CALO: :CALO envelope and the explicit volume set (in bold) leads to the same hit fraction as in the nominal implementation of FASTCALOSIMV2 in ISF.

of the boundary check increases the ratio to 0.121%, exceeding the ratio of the nominal implementation. This is the result of particles failing the positional and directional check and encountering volumes along the ISF boundary for which the fast simulation model is not active, thus generating additional GEANT4 hits. As expected, the best hit ratio is observed for the CALO & Explicit configuration, which exactly reproduces the ratio observed for ISF. In fact, it is shown that this configuration precisely reproduces the simulation of ISF at hit level. As is shown for the baseline CALO option in Figure 9.6(a), Figure 9.9 shows the simulated hit energies for the CALO & Explicit configuration for the energy range in which the largest fraction of hits is found. The simulation at hit level is shown to be fully compatible with the ISF implementation within statistical fluctuations. Furthermore, the simulated hit energy and the distributions of the cell positions of the simulated hits, split into FASTCALOSIMV2 and GEANT4 hits, are shown in Figure 9.10. For FASTCALOSIMV2 hits, the hit energies and positions are very well reproduced, and only very minor differences are observed in the ratio between both implementations. For GEANT4 hits, differences are slightly more apparent. However, the ratio between both implementations fluctuates around unity, such that these minor differences are likely purely driven by statistical fluctuations. The corresponding distributions for all other studied choices of trigger volumes are shown in Appendix B. The additional ModelTrigger calls that are required with the CALO & Explicit configuration is not found to create any significant performance overhead, such that this option is chosen for the Run 3 proof-of-concept implementation.

It should be emphasized that the described complications in the choice of trigger volumes are purely a result of the mismatch between the definition of the ID-Calo boundary that is used to derive the FASTCALOSIMV2 parametrization and the calorimeter envelope as defined in the ATLAS geometry. For Run 4 and beyond, the FASTCALOSIMV2 parametrization can be easily altered to align with the GEANT4 geometry, effectively avoiding any ambiguity in the choice of trigger volumes.




Figure 9.9: Simulated hit energies from single photon events for the `GEANT4` implementation of `FASTCALOSIMV2` (FCS in G4), the nominal `FASTCALOSIMV2` implementation in ISF (FCS) and for the full `GEANT4` simulation. Note that energies created by `FASTCALOSIMV2` hits extend to significantly larger energies than the shown range. The `CALO::CALO` envelope and the set of explicit volumes is employed as fast simulation trigger with an additional positional and directional check in `ModelTrigger`. In almost all cases, the dots from the distributions of both implementations entirely overlap. The lower pad shows the ratio between the distributions in the ISF and `GEANT4` implementation.

9.5 Computational Performance

Apart from aligning the novel `GEANT4` implementation of `FASTCALOSIMV2` with the ISF counterpart, the resulting computational performance is of interest. In terms of memory consumption, it is shown that both implementations perform similarly, since the amount of memory required is dominated by the loading of the `FASTCALOSIMV2` parametrization file, which is kept in memory in both cases.

In terms of CPU usage, Figure 9.11 shows the relative simulation speed-up as a function of the hit energy weighted η average in comparison to the nominal ISF simulation for the baseline configuration with `CALO::CALO` as the trigger volume and for the configuration that best mimics the ISF implementation, i.e., using the set of explicit volumes and imposing an additional boundary check. The simulation speed-up for all investigated trigger volume configurations is shown in Figure B.6 in the appendix. In the baseline scenario, an overall speed-up of about 2.9% is achieved, with improvements of up to 9% for very forward particles. However, the simulation speed-up is of artificial nature, as it is caused by the lower number of `GEANT4` hits being simulated in this configuration compared to the ISF implementation. Especially in the

9.5 Computational Performance

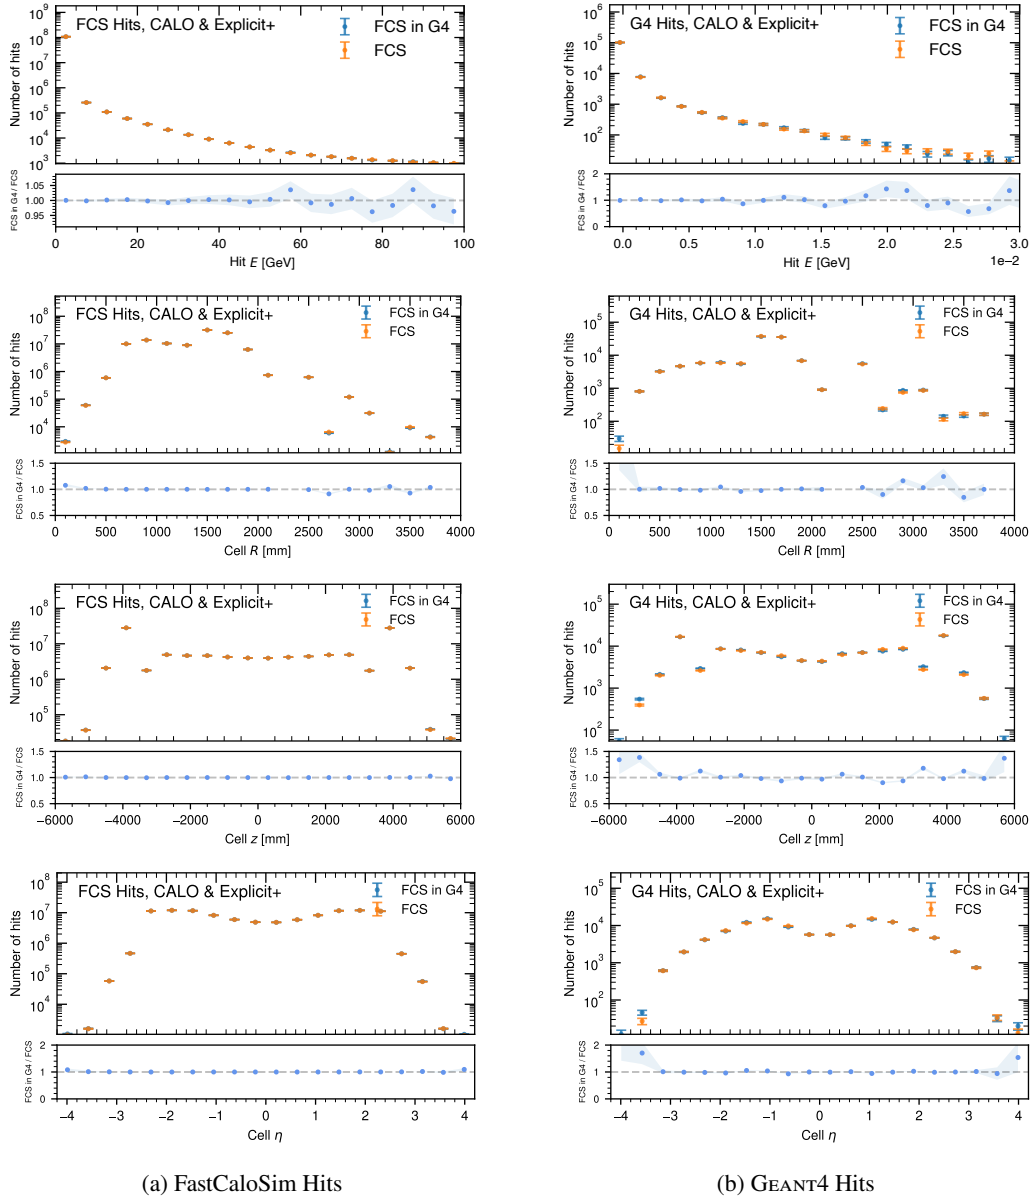


Figure 9.10: Simulated hit energy and distributions of the cell R , η and z position of the simulated hits for the GEANT4 and ISF implementation of FASTCALOSIMV2, split into the simulation of FASTCALOSIMV2 hits (left) and GEANT4 hits (right). The Run 3 calorimeter envelope and a complementary list of explicit volumes is used as fast simulation trigger with a dedicated boundary check in ModelTrigger.




Figure 9.11: Simulation speed-up in dependence of the hit energy weighted η average for the configuration with CALO::CALO as trigger volume without additional boundary checks (a) and speed-up for the configuration with CALO::CALO and the set of explicit volumes as trigger volumes, with inclusion of the additional positional and directional boundary checks (b). Benchmarks are performed on a 64-core AMD EPYC™ 7702 CPU at 2 GHz base clock.

forward region, where particles are intercepted by FASTCALOSIMV2 upon reaching the MBTS, fewer GEANT4 hits are simulated, resulting in the observed performance improvements.

For the chosen configuration, which reproduces the same GEANT4 to FASTCALOSIMV2 hit ratio, the simulation speed remains approximately the same, with an overall negligible speed-up of 0.13%. Thus, the approximately 10% overhead introduced by the ISF implementation of GEANT4 over the pure, full GEANT4 simulation, is not recovered. Presumably, the overhead of the ISF implementation is dominated by the GEANT4 simulation of the calorimeter, and since the calorimeter simulation is replaced by FASTCALOSIMV2 in both cases, the overhead is not observed.

9.6 Outlook

This chapter motivated the need for an alternative implementation of fast simulation models in ATLAS. A proof-of-concept implementation of FASTCALOSIMV2 within GEANT4 was presented and has been shown to match the simulation of the ISF implementation with the choice of appropriate trigger volumes.

For future FASTCALOSIMV2 releases in Run 4 and beyond, the boundary used to derive the parametrization can be easily adapted to the ATLAS geometry, so that the choice of appropriate

trigger volumes should become trivial. As has been shown in Section 9.5, the inclusion of the MBTS simulation within the `FASTCALOSIMV2` parametrization could have noticeable performance benefits for the simulation of particles with $|\eta| > 2.0$.

In the near future, and using the presented proof-of-concept implementation of `FASTCALOSIMV2`, it is planned to directly interface other simulators such as `FASTCALOGAN` and `FATRAS` with `GEANT4`, with the vision of entirely replacing `ISF` as particle dispatcher and dramatically simplifying the simulation workflow.

Search for Resonant Higgs Boson Pair Production

IN recent years, searches for resonant and non-resonant Higgs boson pair production have been conducted by the ATLAS Collaboration in a multitude of final states, including the $4b$ [227, 228], the $b\bar{b}\tau\tau$ [229] and the $bbWW^*$ [230] final states. Figure III.1 shows the combined upper limits at 95% CLs [231] on the cross-section of resonant Higgs boson pair production for a spin-0 heavy scalar particle and a spin-2 KK graviton. In most mass regions, $4b$ is the

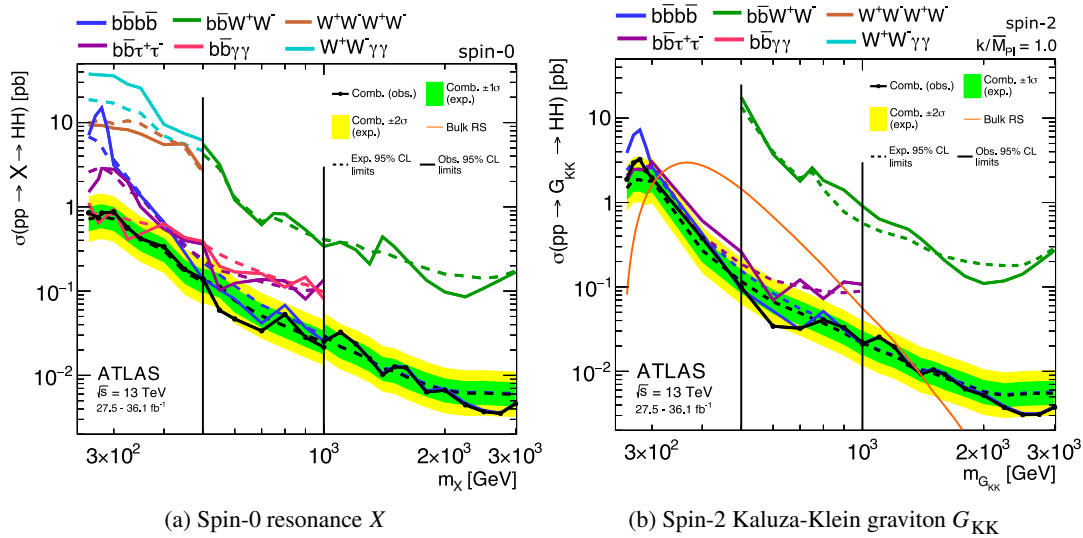


Figure III.1: Upper limits at 95% CL on the cross-section of resonant Higgs boson pair production for (a) a spin-0 heavy scalar particle and (b) a spin-2 KK graviton with $k/\bar{M}_{\text{PL}} = 1$. The observed (expected) limits are shown as solid (dashed) lines. The $\pm 1\sigma$ and $\pm 2\sigma$ bands are only shown for the expected limits of the combination. The vertical black lines in each panel indicate mass intervals where different final states are combined. Adapted from Ref. [232].

most sensitive final state, followed by $b\bar{b}\tau\tau$. Previous searches in the $b\bar{b}WW^*$ final state do not contribute significantly to the combined upper limits. Nevertheless, this channel remains promising for exploring resonant Higgs boson pair production, as it offers the largest di-Higgs branching ratio after the $4b$ final state (see Table 1.4), and is the main focus of this work.

The $b\bar{b}WW^*$ final state is further split according to the subsequent decay of the W bosons:

- 0-lepton (“0 ℓ ”) final state - both W bosons decay hadronically resulting in the largest HH branching ratio $\mathcal{B}r(HH \rightarrow b\bar{b}WW^* \rightarrow b\bar{b}4q) \approx 11.4\%$
- 1-lepton (“1 ℓ ”) final state - one W boson decays leptonically, while the other decays hadronically, resulting in a HH branching ratio $\mathcal{B}r(HH \rightarrow b\bar{b}WW^* \rightarrow b\bar{b}\ell\nu_\ell qq') \approx 7.19\%$
- 2-lepton (“2 ℓ ”) final state - both W bosons decay leptonically resulting in the smallest HH branching ratio $\mathcal{B}r(HH \rightarrow b\bar{b}WW^* \rightarrow b\bar{b}\ell\nu_\ell\ell'\bar{\nu}_{\ell'}) \approx 1.14\%$

The 0 ℓ final state is advantageous due to the fact that it has the highest branching ratio, while the presence of a lepton in the 1 ℓ final state is particularly useful to reject multijet background. The 2 ℓ final state is not considered as a result of the very small branching ratio.

While the search for HH production is limited to the boosted and resolved topologies, the extension to SH production opens the window to many more interesting topologies. The search for resonant SH production is of particular interest in the $b\bar{b}WW^*$ decay channel, as it is the final state with the largest predicted branching ratio for a number of promising benchmark scenarios (see Section 1.8.2).

Figure III.2 shows sketches of the various topologies of resonant $X \rightarrow SH \rightarrow b\bar{b}WW^*$ production in the case that both W bosons decay hadronically. The topologies depend on the values of the resonance mass m_X and the additional scalar particle mass m_S . A categorization of the topologies in the two-dimensional $m_S - m_X$ plane, based on the angular separation of the Higgs and W boson decay products is shown in Figure III.3 for the 0 ℓ and 1 ℓ final states. The categorization is only slightly different for the two final states. By convention, the masses are ordered as $m_X \geq m_S > m_H$, where the resonance mass is restricted to $m_X > m_S + m_H$ for SH production and $m_X > 2m_H$ for HH production. In the following, the Higgs boson is assumed to decay exclusively into a pair of b quarks, while the boson S is assumed to have $m_S > 2m_W$ and to decay into a pair of W bosons. For $m_X, m_S \lesssim 1$ TeV (Cat. 1), all decay products can be resolved as single isolated jets.

The most relevant topologies for this work are the boosted (Cat. 5) and split-boosted (Cat. 4) topologies. In the first case, the scalar particle masses satisfy $m_X \gtrsim 1$ TeV and $m_X \gg m_H, m_S$, such that the Higgs boson and the S boson are produced at large transverse momenta and their decay products appear close in the laboratory frame, and can be reconstructed as a single large jet. The boosted 1 ℓ final state is particularly difficult to reconstruct because of the expected overlap between the lepton and the hadronic decay products of the jet originating from the WW^* decay. In the split-boosted case, the energy available to produce the decay products is particularly asymmetric and the masses of the scalar particles fulfil $m_X \gtrsim 1$ TeV and

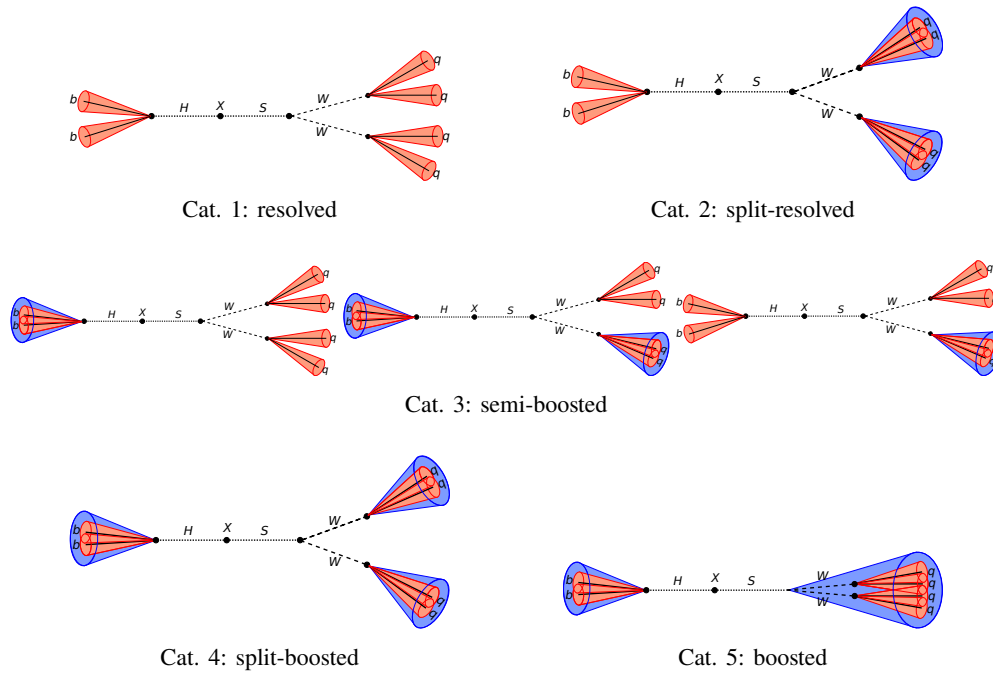


Figure III.2: Sketch of the different topologies of resonant $X \rightarrow SH \rightarrow b\bar{b}WW^*$ production in the 0ℓ final state. The topology depends on the two resonance masses m_X and m_S and is classified according to the angular separation of the decay products. [233]




Figure III.3: Topology categorization of resonant $X \rightarrow SH \rightarrow b\bar{b}WW^*$ production in the (a) 0ℓ and (b) 1ℓ final state in the $m_X - m_S$ plane. The labelling is based on the angular distance between the Higgs boson and W bosons decay products. [233]

$m_S \gtrsim 0.3m_X$. When the mass of the S boson is sufficiently close to the value of m_X , the angular separation between the W boson pair is large enough to resolve their decay products using isolated jets. For smaller values of m_S , the angular separation between the W bosons decreases and their decay products begin to overlap. In all cases, the masses of X and S are sufficiently large such that both W bosons remain boosted, and their decay products are reconstructed as single jets.

At the present day, five different decay channels which offer the most promising sensitivities to Higgs boson pair production in the $b\bar{b}WW^*$ final state, are under investigation: the resolved, the split-boosted and the boosted topologies in the 1ℓ final state and the split-boosted and boosted topologies in the 0ℓ final state.

One of the most challenging tasks in the boosted and split-boosted topologies is the identification of Higgs jets and their distinction from top and QCD jets from background processes. Jet substructure variables are typically used to aid the identification of boosted jets, but the finite size of topological clusters in the calorimeter limits their resolution. Track-assisted Reclustered (TAR) jets [234] jets effectively overcome this limitation by exploiting angular information from the Inner Detector, and allow for a flexible choice of the reconstruction algorithm and jet size.

This part is structured as follows: In Chapter 10, a brief introduction to jet reconstruction in boosted topologies is given, with a particular focus on TAR jets, and is largely taken from Ref. [235]. Chapter 11 presents TAR jet optimization studies in the boosted 0ℓ [236] and 1ℓ [233] final states, as well as in the split-boosted 0ℓ [237] final state. The goal of the studies is to find a TAR jet configuration that provides improved sensitivity in all channels with respect to the default jet reconstruction algorithm and parameters typically used in ATLAS. Note that this thesis does not aim to discuss the analyses in full detail. A more comprehensive discussion is found in Refs. [233, 236, 237].

Jets in Boosted Topologies

THE high energies available at the LHC allow for the production of jets with transverse momenta of up to a few TeV, resulting from the decay of boosted particles. Since many BSM models predict new heavy resonances, the study of these high- p_T jets opens an important window for the search for new physics.

In the decay of strongly boosted objects, the decay products are collimated in the particles' flight direction, so that their angular separation in the laboratory frame is typically very small. Figure 10.1 illustrates the difference between the experimental signature of a boosted decay and that of a particle at rest. When a particle decays at rest, the decay products are well-separated in the detector and can be reconstructed as individual jets. In the boosted case, however, the decay products of the individual partons are in many cases not resolvable and form a single *fat* or *large- R* jet. More precisely, for a two-body decay, the angular separation of the decay products [238] is approximated by

$$\Delta R \approx \frac{2m}{p_T}, \quad (10.1)$$




Figure 10.1: Sketch of the production of two jets from the decay of a particle X , when the particle is at rest (left) and when the particle is boosted (right). The fraction of the particle momentum taken by one of the jets is denoted by z . Modified from Ref. [238]

where m and p_T are the mass and transverse momentum of the decaying particle, respectively. The decay products begin to overlap at angular distances $\Delta R \lesssim 2R$, where R is the minimum jet radius required to contain all the decay products and thus sets an upper limit on the jet size needed to resolve the decay products of both jet-initiating particles. For a Higgs boson with a transverse momentum of 500 GeV, the angular separation between the decay products is $\Delta R \approx 0.5$, so that the decay products begin to overlap at a radius of $R \gtrsim 0.25$ and cannot be resolved with the conventional calorimeter jets used in the LHC experiments, which typically cover angular distances between 0.4 and 0.6. Instead, the products of boosted decays are typically reconstructed in a single large- R jet, where the default jet radius used in ATLAS is $R = 1.0$ (see Section 10.1 for more details) and allows for the capture of all decay products for particles with $p_T \gtrsim 2m$.

In the search for new physics, the hadronic branching ratios of boosted Higgs bosons, W/Z bosons and top quark decays are typically larger compared to their leptonic counterparts, making the exploitation of boosted hadronic decays especially important in the search for BSM particles. At the same time, the distinction between jets from boosted decays and background jets that originate from high- p_T light quarks and gluons (*QCD jets*) is a major experimental challenge at hadron colliders. The need to use large- R jets in the reconstruction of boosted decays further complicates the situation, as these are more susceptible to underlying event and pile-up contributions.

In order to effectively identify the origin of boosted jets, one typically exploits the heavy flavour content, the mass, and the internal structure of the large- R jet with the help of substructure observables specifically designed to discriminate between signal jets (typically from boosted electroweak bosons, Higgs bosons and top quark decays) and background jets (typically from QCD jets and top quark initiated jets from $t\bar{t}$ decays). However, the finite size of topological clusters [199] in the calorimeter limits the resolution of substructure variables. To recover the resolution of substructure variables for high p_T jets, Track-Assisted Reclustered (TAR) jets [234], which aid the reconstruction of boosted jets by using angular information from the Inner Detector, can be used.

This chapter focuses on the description of typical jet reconstruction algorithms used in ATLAS and the definition of jet substructure variables used to identify boosted jets. In addition, the chapter provides a brief introduction to TAR jets, which are the focus of the performance studies presented in Chapter 11.

10.1 Jet Reconstruction

Partons are not directly observable, but fragment and form a collimated spray of hadrons in the detector. While all hadrons deposit energy in the calorimeter, charged hadrons leave additional tracks in the ID. This information is then used to group together stable hadrons, ideally reconstructing all hadrons originating from the original parton into a single object known

as a jet. This procedure allows to infer the properties from the original parton, thus providing the link between the most fundamental interactions and the observable particles in the detector. While jet reconstruction algorithms are well-motivated *ad-hoc* methods, they do not follow an underlying physical principle. However, there is a broad consensus on the requirements that these algorithms must meet. For example, jet reconstruction algorithms should be infrared (IR) and collinear safe (IRC safe) [239]. In IR safe algorithms, the addition of soft particles does not significantly distort the output of the algorithm, such that sensitivity to the underlying event and pile-up is greatly reduced. In collinear safe algorithms, the invariance of the algorithm's output after the splitting of a high- p_T particle into two collinear particles is ensured. Currently, IRC safe sequential jet clustering algorithms are the most widely used. The most common one is known as the generalized- k_t algorithm and follows the following procedure:

1. Based on given input, create a list of proto-jets and compute the distance defined by a certain metric d_{ij} between every pair of proto-jets i and j .
2. Compute the metric $d_{i,B} = p_{T,i}^{2p}$ between each proto-jet i and the beam B .
3. Find the minimum distance.
4. If the minimum distance is between the proto-jets, merge both objects to form a new proto-jet and remove the individual proto-jets from the list. If the minimum distance is between the proto-jet and the beam, remove the proto-jet from the list and call it a jet.
5. Repeat the procedure until no proto-jets are left.

The metric d_{ij} is given by

$$d_{ij} = \min(p_{T,i}^{2p}, p_{T,j}^{2p}) \frac{\Delta R_{ij}^2}{R^2} \quad \text{with} \quad p = \begin{cases} +1, & k_t \\ 0, & \text{Cambridge/Aachen} \\ -1, & \text{anti-}k_t \end{cases}, \quad (10.2)$$

where p defines the type of the clustering algorithm and R is known as the size parameter of the jet.¹ A value of $p = +1$ corresponds to the k_t (KT) algorithm [240], which begins by clustering low- p_T proto-jets first. In contrast, the anti- k_t (AKT) algorithm [203], defined by $p = -1$, starts by clustering high- p_T proto-jets first. In contrast, the Cambridge/Aachen (CA) algorithm [241] offers a p_T independent jet clustering. Note that while the jets produced by the AKT algorithm have regular and conical shapes, the same event reconstructed with the CA algorithm produces more irregular jet shapes. The size parameter R can be tuned for optimal performance. While a large jet radius is advantageous to ensure full containment of the hadrons resulting from the original parton, a smaller radius helps to reduce contributions from pile-up and the underlying event, thus avoiding the overestimation of the jet energy and mass.

¹ In this case, the angular separation ΔR_{ij} is defined by $\Delta R_{ij}^2 = (y_i - y_j)^2 + (\phi_i - \phi_j)^2$ with the rapidity y as opposed to the more common definition, which uses the pseudorapidity η .

In recent years, extensions to the above jet clustering algorithms have been developed, particularly in the context of boosted analysis. The variable- R (VR) [242] modification introduces a jet p_T dependence in the size parameter by replacing it with an effective size parameter

$$R \rightarrow R_{\text{eff}}(p_T) = \frac{\rho}{p_T}, \quad (10.3)$$

where ρ is a tunable parameter known as the *energy scale* and determines the absolute scaling. The inverse proportionality of the effective jet size is particularly useful in the reconstruction of boosted topologies, as the size can be dynamically adapted to the shrinking cone that is required to collect all the decay products of high- p_T jets. Minimum and maximum cut-off values R_{min} and R_{max} are defined for R_{eff} . These prevent jets from shrinking below detector resolution at high p_T and jets becoming too large at low p_T . All three parameters ρ , R_{min} and R_{max} can be tuned to give the best performance for the specific application. In the following, the VR extensions to the AKT and CA algorithms are denoted AKT-VR and CA-VR, while the standard fixed- R algorithms are denoted AKT-FR and CA-FR, respectively.

A number of different jet reconstruction algorithms are used in ATLAS. During Run 2, LCTopo jets [203] reconstructed using the AKT-FR algorithm with a size parameter $R = 1.0$, were the default choice for reconstructing large- R jets in ATLAS, and are hereafter simply referred to as large- R jets. LCTopo jets use noise suppressed and locally calibrated topological clusters [199] from the calorimeter as input to the jet reconstruction. For the identification of b -hadrons, ATLAS also uses tracks from the Inner Detector to reconstruct small radius subjets within large- R jets, which are referred to as *track jets* [243]. For boosted topologies, track jets reconstructed with the AKT-VR [244] algorithm are typically used as they have been shown to significantly improve the performance for high- p_T jets. A more detailed overview on the jet collections used in ATLAS is provided in Chapter 7 in the context of the ATLFast3 physics performance validation.

10.2 Jet Substructure

Although it is not possible to resolve the decay products individually in strongly boosted topologies, information about the original jet-initiating partons can be obtained using jet substructure variables. A widely used variable is known as N -subjettiness τ_N [206]. First, one reclusters the jet constituents of the candidate jet into exactly N KT subjets. Using these subjets, τ_N is defined as

$$\tau_N = \frac{\sum_{i \in J} p_{T,i} \min(\Delta R_{k,i} \mid k = 1, \dots, N)}{\sum_{i \in J} p_{T,i} R}, \quad (10.4)$$

where the sum is over the constituent particles of a given jet, and $\Delta R_{k,i}$ is the angular separation between a candidate subjet k and the constituent particle i . N -subjettiness is a measure of how well a given jet is described by being composed of at least N subjets. In practice, the ratio $\tau_{NM} = \tau_N/\tau_{M-1}$ is used. For a particular jet containing N -subjets (“ N -prong jet”), τ_N is expected to be small, while τ_{M-1} is expected to be large such that $\tau_{NM} \ll 1$. The N -subjettiness ratios have been shown to be effective in identifying top [245], W [246] and $H \rightarrow WW^*$ [247] jets. While the subjet axis for calculating the N -subjettiness variables is usually defined by the vector sum of all its constituents, the axis can also be defined by its hardest constituent, in which case it is known as the Winner Takes All (WTA) axis [248].

A further set of interesting variables are known as energy correlation functions (ECFs) [205] and are defined as

$$\text{ECF}(N, \beta) = \sum_{i_1 < i_2 < \dots < i_N \in J} \left(\prod_{a=1}^N p_{T_{i_a}} \right) \left(\prod_{b=1}^{N-1} \prod_{c=b+1}^N \Delta R_{i_b i_c} \right)^\beta, \quad (10.5)$$

where the sum goes over all constituents i of a jet J and β is a positive tunable parameter. In analogy to the ratio τ_{NM} , one defines the potentially discriminating dimensionless ratios

$$r_N^{(\beta)} = \frac{\text{ECF}(N+1, \beta)}{\text{ECF}(N, \beta)}, \quad (10.6)$$

as well as

$$C_N^{(\beta)} = \frac{r_N^{(\beta)}}{r_{N-1}^{(\beta)}} = \frac{\text{ECF}_{N+1} \times \text{ECF}_{N-1}}{\text{ECF}_N^2} \quad (10.7)$$

and

$$D_N^{(\beta)} = \frac{\text{ECF}_{N+1} \times \text{ECF}_{N-1} \times \text{ECF}_1^N}{\text{ECF}_N^3}. \quad (10.8)$$

The variables C_2^β and D_2^β are a measure for the transverse energy distribution within a jet, and have been shown to effectively identify 2-prong decays, such as jets originating from W -boson decays [246].

10.3 Track-Assisted Reclustered Jets

Due to the finite angular resolution of the ATLAS calorimeter, the reconstruction of jet substructure variables from large- R jets becomes increasingly difficult in boosted topologies where the jets appear very close in the detector. To help overcome this limitation, it is possible to assist the jet reconstruction with track information from the ID, which offers excellent angular resolution. As a standard procedure, ATLAS already uses the track-assisted (TA) jet mass for

standard large- R jets. Conventionally, the mass of a jet J reconstructed in the calorimeter is defined as

$$m^{\text{calo}} = \sqrt{\left(\sum_{i \in J} E_i\right)^2 - \left(\sum_{i \in J} \vec{p}_i\right)^2}, \quad (10.9)$$

where the sum goes over all cluster constituents i of the jet J . In order to include track information, the particle tracks in the ID are associated with the calorimeter clusters by scaling the momentum of the tracks to negligible values and adding them to the list of inputs for the corresponding clustering algorithm, a procedure known as ghost-association [249]. The scaling ensures that only angular information is added, and the jet reconstruction is not distorted. The TA jet mass is then defined as [250]


$$m^{\text{TA}} = \frac{p_T^{\text{calo}}}{p_T^{\text{track}}} \cdot m^{\text{track}}, \quad (10.10)$$

where m_{track} is the mass and p_T^{track} the transverse momentum of the associated tracks. The reweighting $p_T^{\text{calo}}/p_T^{\text{track}}$ is required to ensure the inclusion of the energy of neutral hadrons to which the ID is not sensitive.

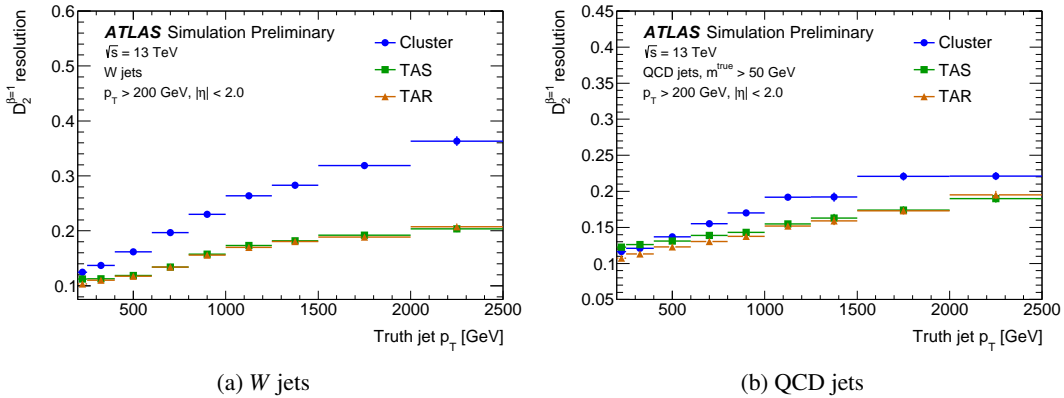
While this approach improves the resolution of jet mass measurements and could in principle be extended to other jet substructure variables, it does not account for local differences in the energy deposited by neutral and charged hadrons within the jet. A novel jet collection known as track-assisted reclustered (TAR) jets [234] aims to circumvent this problem by applying local corrections to the p_T of nearby tracks. TAR jets are built from tracks and calibrated AKT-FR jets with a size parameter of $R = 0.2$ reconstructed from energy clusters in the calorimeters. The AKT-FR jets are reclustered into large- R jets and then trimmed [251] with the parameter $f_{\text{cut}} = 0.05$ such that soft components are removed. The tracks are assigned to the remaining AKT-FR $R = 0.2$ through *ghost-association*. If tracks are not associated to any of the jets, they are matched to the closest $R = 0.2$ jet in case the jet axis lies within $\Delta R < 0.3$ of the track. If a jet axis does not lie within this value, the non-associated track is discarded. After the matching procedure, the p_T of the associated tracks is rescaled according to

$$p_T^{\text{track, new}} = p_T^{\text{track, old}} \times \frac{p_{T,j}^{\text{subject}}}{\sum_{i \in j} p_{T,i}^{\text{track, old}}}, \quad (10.11)$$

where the sum runs over all tracks i matched to the subject j . Similar to the track-assisted mass, the rescaling aims to account for the energy deposition of neutral hadronic components. A large- R jet is then formed from the jets with the rescaled associated tracks and is hereafter referred to as TAR jet. The jet substructure variables including the jet mass, m^{TAR} , are then computed using the rescaled tracks of the TAR jets. A schematic illustration of this procedure is shown in Figure 10.2. In contrast to the calculation of the track-assisted mass m^{TA} , the


 Figure 10.2: Schematic illustration of the track-assisted jet reclustering algorithm in the $\eta - \phi$ plane. [235]

procedure has the advantage that the missing energy from neutral hadrons is corrected on a subject-by-subject rather than on a jet-by-jet basis. Thus, local effects are taken into account and the resolution of the jet substructure variables can be significantly improved. Figure 10.3 shows the resolution of $D_2^{\beta=1}$ for W and QCD jets as a function of truth jet p_T for TAR jets and for jets reconstructed using the calorimeter-only based standard reconstruction algorithm. A significant improvement in the resolution is observed for TAR jets, with increasing relative gain at higher jet p_T . In addition, TAR jets also offer the advantage that calibrations and uncertainties can be propagated from the constituent $R = 0.2$ jets and tracks, so that no additional custom calibration procedures are required. As a result, TAR jets are not limited to AKT-FR reclustering with a fixed size parameter $R = 1.0$, but provide the flexibility to use whichever jet definition that provides the best performance for the specific use case.


 Figure 10.3: Resolution of $D_2^{\beta=1}$ for (a) W and (b) QCD jets as a function of truth jet p_T . The TAR jet reconstruction algorithm is compared to the standard algorithm labelled as “Cluster” [234].

Track-Assisted Reclustered Jet Optimization

WHILE most analyses in ATLAS targeting boosted topologies rely on large- R jets reconstructed using standard algorithms and parameters for which calibrations are readily available, the use of TAR jets allows the free choice of almost arbitrary jet definitions. This additional freedom allows for the optimization of jet definitions based on the specific target signature and phase space of an analysis, with the ultimate goal of maximizing the overall sensitivity of the search.

This chapter presents performance studies targeting the $b\bar{b}WW^*$ decay mode, aimed at finding the optimal TAR jet configuration for Run 2 HH searches in the boosted 0ℓ and 1ℓ final states, and a SH search in the split-boosted 0ℓ final state.

11.1 Methodology

In this work, FR TAR jets with fixed radius and VR TAR jets with fixed energy scale are considered, reconstructed using either the AKT or the CA algorithm. For FR TAR jets, the radius is varied between $R_{\min} = 0.5$ and $R_{\max} = 1.5$ with a step size of 0.25, while for VR TAR jets the energy scale is varied between $\rho_{\min} = 150$ GeV and $\rho_{\max} = 350$ GeV with a step size of 50 GeV. For convenience, the results of the performance studies for FR and VR TAR jets are often presented on the same axis. Note, however, that there is no one-to-one correspondence between the considered jet radii of FR TAR jets and the energy scales of VR TAR jets. For example, a jet with $p_T = 250$ GeV will have $R = 1.0$ at $\rho = 250$ GeV, but will have a radius of 0.6 at $\rho = 150$ GeV, which does not correspond to the smallest radius considered for FR TAR jets.

To evaluate the performance of TAR jets, two distinct approaches can be taken:

1. *jet-by-jet*: jet taggers are defined based on a selection that exploits the heavy flavour content, the mass and the substructure of the TAR jet. The performance of these taggers can be studied by comparing their ability to reject background jets, such as top and QCD jets.
2. *event-by-event*: instead of studying the performance of a specific tagger, the signal event efficiencies and background event rejections are compared for the TAR jet definitions of interest.

Both approaches are valid, but they examine very different aspects of TAR jet performance. In the first case, the performance results are specific to the tagger studied, but independent of specific analyses. Thus, jet-by-jet performance studies allow very general statements to be made about a tagger's ability to identify the origin of jets. In contrast, event-by-event studies take into account any further kinematic selection that an analysis may impose. Therefore, the results in this case are selection specific and allow less general statements about tagging performance, but may be more useful for the individual analysis.

In this work, both approaches have been investigated in the context of analyses searching for resonant Higgs boson pair production in the boosted 0ℓ [233] and 1ℓ [236] final states and in the 0ℓ split boosted [237] final state. For the jet-by-jet studies, the TAR jet performance was examined in terms of the ability to reject top and QCD jets, and is based on taggers that approximately correspond to the selections employed by the analyses. For the boosted 0ℓ channel, a tagger is defined based on a two-sided mass cut and a selection on C_4 to identify four-prong $H \rightarrow W_{\text{had}}W_{\text{had}}$ jets. For the purpose of both boosted channels, the tagging performance was also investigated for $H \rightarrow b\bar{b}$ jets, in which case the tagger is based solely on a two-sided mass cut. Finally, for the split-boosted topology, the ability to tag isolated $S \rightarrow W_{\text{had}}$ jets using a one-sided selection on D_2 was studied as well. In all cases, the taggers are constructed to be p_T -dependent in a way that the tagging efficiency of signal jets is independent of jet p_T . This has the advantage that the taggers are independent of the considered mass point and the background jet rejection can be easily compared between different jet definitions. The performance differences between the investigated TAR jet definitions are found to be very small in the majority of cases. Therefore, the results of the jet-by-jet studies are only taken account complementarily in the decision of a jet configuration and are presented in Appendix C.

The selection of the optimal TAR jet algorithm and size parameter is therefore driven by event-by-event performance studies of the TAR jets. For this purpose, the event efficiencies are computed for the different TAR jet configurations, taking into account the preliminary event selections of the most sensitive signal regions of the different analyses. For the boosted topologies, a subset of three mass resonances $m_X = 1, 2$ and 3 TeV is considered, while in the split-boosted case the mass of the X resonance is fixed at 3 TeV and m_S is varied between $0.75, 1.5$ and 2.5 TeV.

At the time of conducting these studies, a QCD multijet background estimate was not available, so that only $t\bar{t}$ events are considered as background processes. For the 1ℓ boosted

final state, $t\bar{t}$ is indeed the leading background process, such that a restriction to $t\bar{t}$ events is expected to be a good approximation. In contrast, for the 0ℓ final states, multijet processes are expected to be the major contributors to the overall background estimate. In this case, the event-by-event studies may be less meaningful for the overall sensitivity of the analysis and the QCD jet tagging studies presented in Appendix C may be taken into account in a complementary way. In any case, the overall sensitivity difference between the default large- R jet size parameter and the optimized TAR jet definition with inclusion of the full background estimate is compared in Section 11.6 and shows a significant improvement in all cases.

11.2 Event Selection

In the boosted 0ℓ analysis, at least two TAR jets with $p_T > 300$ GeV and $|\eta| < 2.0$ are required, where the p_T -(sub)-leading TAR jet must fulfil $p_T > 450(300)$ GeV and $m_{\text{TAR}} > 90$ GeV. If a third TAR jet is reconstructed in the event, it must have $m_{\text{TAR}} < 50$ GeV. Events with lepton candidates are vetoed. In addition, exactly one double b -tagged TAR jet with the 77% DL1r working point is required, which is chosen as the reconstructed $H \rightarrow b\bar{b}$ candidate. The DL1r algorithm is based on a deep-learning neural network that exploits the distinctive features of b -hadrons and is an improved version of the tagger described in Ref. [252]. The remaining p_T -leading TAR jet is defined as the $H \rightarrow W_{\text{had}}W_{\text{had}}^*$ candidate. In the boosted 0ℓ analysis, $W_{\text{had}}W_{\text{had}}$ jets originating from the decay of a Higgs boson are identified using a dedicated neural network known as four prong tagger (FPT) [236]. For the purpose of this study, a p_T -dependent tagger consisting of a two-sided mass window cut and an additional cut on C_4 is used as a proxy for the FPT, designed to select hadronically decaying $H \rightarrow WW^*$ candidates with an 80% signal efficiency over the entire jet p_T spectrum for all jet definitions. The choice of the selection is motivated by the fact, that the TAR jet mass and the C_4 substructure variable have been shown to be particularly important in the decisions taken by the FPT in discriminating against multijet events. The proxy taggers are defined individually for each of the jet collections studied in order to allow for a fair comparison of TAR jet performance. The use of proxy taggers effectively avoids the need to retrain the FPT on each jet collection considered, which would be a much more a laborious task. The definition of the p_T dependent $H \rightarrow W_{\text{had}}^*W_{\text{had}}$ proxy taggers is briefly described in Section 11.3.1 and a summary table of all applied selection criteria is provided in Table 11.1.

In the boosted 1ℓ analysis, at least two TAR jets with $p_T > 100$ GeV and $|\eta| < 2.0$ are required, where exactly one TAR jet must be double b -tagged using the 77% DL1r working point. The presence of exactly one reconstructed muon in the event is required, where the angular distance between the muon and the three p_T -leading TAR jets must be $\min \Delta R(\mu, \text{jet}) < 1$. The W_{had} candidate is chosen as the TAR jet with the minimum angular distance $\Delta R(\mu, \text{jet})$ from the muon. The $H \rightarrow b\bar{b}$ candidate is then chosen as the p_T -leading TAR jet that has been b -tagged (excluding of the W_{had} TAR jet candidate). Finally, the invariant mass of the $H \rightarrow b\bar{b}$

Topology		Boosted		Split-Boosted
Channel		0-lepton	1-lepton	0-lepton
	Count	≥ 2	≥ 2	≥ 3
	$ \eta $	< 2.0	< 2.0	< 2.0
TAR Jet	p_T [GeV]	> 450 (> 300)	> 100	> 200
	m_{TAR} [GeV]	> 90 (> 90) (< 50)	\times	> 50
	2 b -tag	77% DL1r	77% DL1r	\times
Lepton		veto	$= 1\mu$	veto
$\min \Delta R(\ell, J)$		\times	< 1	\times
Candidate Selections		4-prong tagger (80% m_{TAR} + 80% C_4)	$m_{b\bar{b}} > 90$ GeV $m_{b\bar{b}} < 140$ GeV	$m_W > 50$ GeV $m_W < 105$ GeV D_2 tagger (50%)
Ref.		[236]	[233]	[237]

Table 11.1: Overview of the selections applied in the split-boosted and boosted 0ℓ and 1ℓ channels for the TAR jet optimization studies. The selections approximately correspond to the most sensitive signal region of the individual analyses. The numbers in brackets correspond to the values for the p_T -sub- and p_T -sub-sub-leading TAR jets, in the case that different requirements are placed depending on the transverse jet momentum.

candidate is constrained to $90 \text{ GeV} < m_{H \rightarrow b\bar{b}} < 140 \text{ GeV}$, but no additional mass cut is applied to the W_{had} candidate. A summary of the selection criteria applied in the boosted 1ℓ final state is shown in Table 11.1.

In the case of the split-boosted 0ℓ analysis, at least three TAR jets with $m_{\text{TAR}} > 50 \text{ GeV}$, $p_T > 200 \text{ GeV}$ and $|\eta| < 2.0$ are required. Events with the presence of a reconstructed lepton are vetoed. The $H \rightarrow b\bar{b}$ candidate is selected as follows: if no double b -tagged (77% DL1r) TAR jets are present in the event, the jet with the single largest track jet b -tagging score is selected. Otherwise, if at least one b -tagged track jet is present, the TAR jet with the highest number of b -tagged track jets is selected. If there are multiple TAR jets with the same number of b -tagged track jets, the jet with the highest TAR jet mass is selected. The remaining p_T -leading TAR jets are then classified as $S \rightarrow W_{1,2}$ candidates. The $S \rightarrow W_{\text{had}}$ candidates must pass a $50 \text{ GeV} < m_{\text{TAR}} < 105 \text{ GeV}$ mass window cut as well as a p_T dependent one-sided cut on $D_2^{\beta=1.5}$. Similar to the FPT proxy taggers defined for the 0ℓ boosted final state, the tagger definitions are re-derived for all considered jet configurations in order to provide a 50% flat signal efficiency across the entire jet p_T spectrum in all cases. The procedure used to derive the p_T -dependent $S \rightarrow W_{\text{had}}$ taggers for all jet definitions is described in Section 11.3.2 and a summary of all applied selection criteria is provided in Table 11.1.

11.3 Jet Tagging

In this section, the development of two taggers for the identification of $H \rightarrow W_{\text{had}}W_{\text{had}}$ jets in the boosted 0ℓ final state and for the identification of $S \rightarrow W_{\text{had}}$ jets in the split-boosted 0ℓ final state, is described. The taggers are optimized for each considered TAR jet definition, and are designed to allow for a fair comparison between the performance of all considered jet definitions. The conceptual derivation for the taggers is described, as they are used to identify $H \rightarrow W_{\text{had}}W_{\text{had}}$ and $S \rightarrow W_{\text{had}}$ jets in the event-by-event studies presented in Section 11.4. The performance of both taggers in rejecting top and QCD jets is not presented in the main part of this work, but is provided in Appendix C, which also includes TAR jet performance studies for a tagger designed to identify boosted $H \rightarrow b\bar{b}$ jets.

The underlying strategy to derive the taggers is essentially identical in both cases. In each jet p_T bin, one- or two-sided cuts on discriminating variables against QCD and top jets are derived that are designed to provide a fixed, pre-defined signal jet efficiency, hereafter referred to as *working point*. The cut values are then fitted with a logarithmic regression of the form

$$f(p_T) = p_0 + p_1 \cdot \log p_T + p_2 \cdot \log^2 p_T, \quad (11.1)$$

where p_i are arbitrary floating parameters. The aim is to find a functional form that allows to achieve a flat efficiency in jet p_T after application of the tagger. The fits are repeated for each considered jet definition and the binning in transverse jet momentum is manually optimized to ensure a stable fit and an identical jet efficiency across all TAR jet definitions.

11.3.1 $H \rightarrow WW^*$ tagging

As mentioned in Section 11.2, a dedicated neural network known as the four-prong tagger (FPT) is used in the boosted 0ℓ channel to identify $W_{\text{had}}W_{\text{had}}$ jets originating from the decay of a Higgs boson, by exploiting various jet substructure variables (see Ref. [236] for a detailed description). As a proxy for the FPT, a p_T -dependent mass window cut is first defined for each jet collection. The mass window is selected by iteratively finding the smallest window around the median that contains 80% of the signal jets. After applying the TAR jet mass selection, an additional one-sided upper cut is applied on C_4 , also using an 80% working point. Thus, the total signal efficiency of the tagger is approximately 60%.

The TAR jets used to derive the tagger are selected by requiring the presence of at least two TAR jets in an event with $p_T > 250$ GeV and $|\eta| < 2.0$. Track jets associated with the TAR jets must satisfy $p_T > 15$ GeV, $|\eta| < 2.5$ and consist of at least two constituents. Truth $H \rightarrow W_{\text{had}}W_{\text{had}}$ TAR jets are built from truth small- R jets using the AKT algorithm with a size parameter of $R = 1.6$ and must be associated with exactly two truth W_{had} -bosons. Signal candidate jets are then selected as the TAR jets closest to the truth $H \rightarrow W_{\text{had}}W_{\text{had}}$ TAR jet.

Figure 11.1 shows example fits for AKT-FR $R = 1.0$ TAR jets used for the definition of the p_T dependent taggers, consisting of a two-sided TAR jet mass cut, followed by a one-sided upper cut on C_4 for identifying $H \rightarrow W_{\text{had}}W_{\text{had}}$ jets. The fit fairly describes the functional form

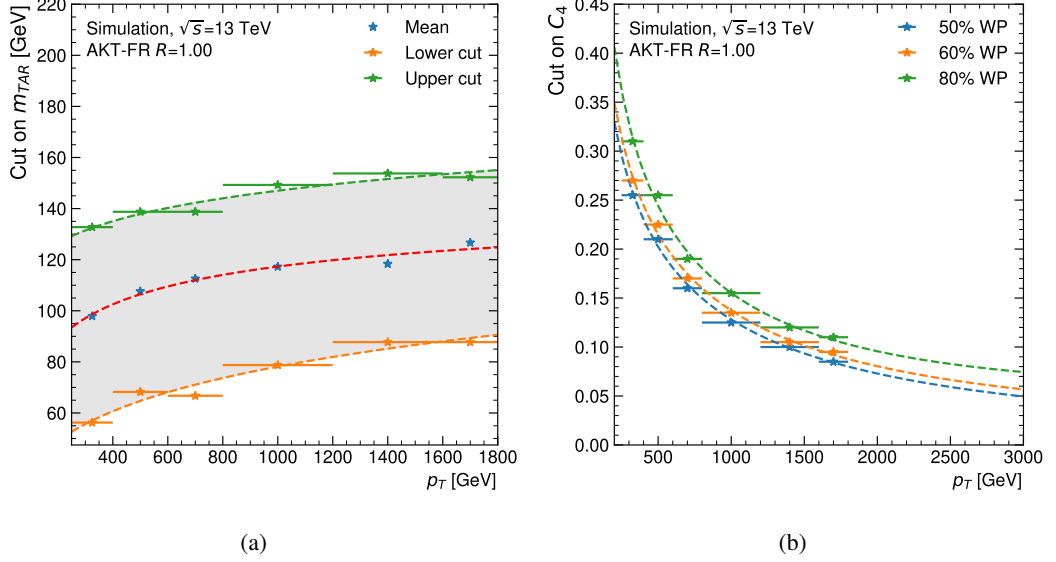


Figure 11.1: (a) Logarithmic regression of lower and upper cuts used to define a p_T dependent mass window cut around the median resulting in an 80% signal jet efficiency for AKT-FR $R=1.0$ TAR jets. The blue stars indicate the means in each p_T bin and the red dashed curve the best fit of the mean, which are both shown for reference. The shaded area indicates the selected region of phase space. (b) Logarithmic regression of a p_T dependent upper cut on C_4 resulting in a 50%, 60% and 80% signal jet efficiency for AKT-FR $R=1.0$ TAR jets. The fits on C_4 are derived after application of the 80% mass window cut.

of both cuts and leads to a flat efficiency in jet p_T for all TAR jet definitions within 5-10% (see Figure C.6 in Appendix C).

The best fit values for the parameters p_i used in the definition of the p_T dependent $H \rightarrow W_{\text{had}}W_{\text{had}}$ tagger are shown in Figure 11.2 for the TAR jet mass window selection and the C_4 cut as a function of the jet size for AKT-FR TAR jets. Note that the functional form of the fit changes for TAR jet sizes $R \leq 1$ in the case of the jet mass window cut, and $R \leq 0.65$ for the upper C_4 cut. The change is driven by effects in the low jet p_T region up to 500 GeV, where not all $H \rightarrow W_{\text{had}}W_{\text{had}}$ decay products can be collected with such small jet sizes.

11.3.2 $S \rightarrow W_{\text{had}}$ tagging

The tagger defined for to study the selection of $S \rightarrow W_{\text{had}}$ (W_{had}) candidate jets in the 0ℓ split-boosted final state consists of a p_T dependent one-sided upper cut on $D_2^{\beta=1.5}$ defined to

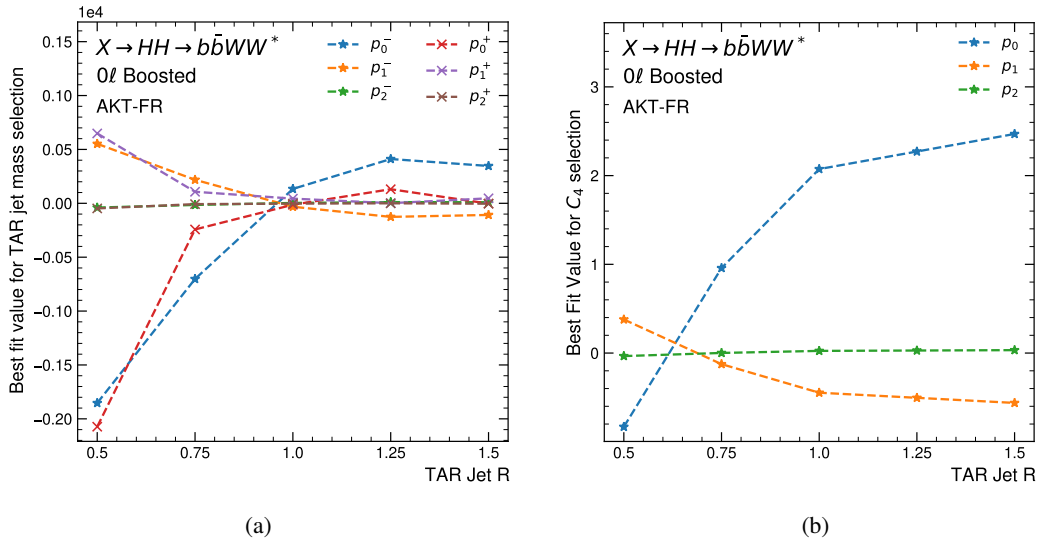


Figure 11.2: Best fit values for the parameters p_i used in the definition of the p_T dependent $H \rightarrow W_{\text{had}}W_{\text{had}}$ tagger as a function of the jet size for AKT-FR TAR jets for (a) the TAR jet mass window and (b) the substructure variable C_4 .

provide a 50% jet efficiency throughout over the entire p_T spectrum.

TAR jets are selected by requiring the presence of at least three TAR jets in an event with $p_T > 200$ GeV, $|\eta| < 2.0$ and $50 \text{ GeV} < m_{\text{TAR}} < 105$ GeV. Track jets associated with TAR jets must satisfy $p_T > 15$ GeV, $|\eta| < 2.5$ and consist of at least two constituents. Truth $S \rightarrow W_{\text{had}}$ TAR jets are built from truth small- R jets using the AKT algorithm with a size parameter of $R = 1.6$ and must be associated with exactly one truth W_{had} -boson. Signal candidate jets are then selected as the TAR jets closest to the isolated truth $S \rightarrow W_{\text{had}}$ TAR jet. A $S \rightarrow W_{\text{had}}$ TAR jet is defined as being isolated if the angular distance to all other non truth-matched W_{had} jets satisfies $\Delta R > 2.0$.

A typical fit used for the definition of the taggers is shown in Figure 11.3(a) for AKT-FR $R = 1.0$ TAR jets. For all jet definitions, the resulting fits describe the functional form of the cuts on $D_2^{\beta=1.5}$ in the different p_T bins well and result in a flat 50% jet efficiency in p_T within a margin of $\sim 5\%$ (see Figure C.11 in Appendix C).

The best fit values for the parameters p_i used in the definition of the p_T dependent $S \rightarrow W_{\text{had}}$ tagger are shown in Figure 11.3(b) as a function of the jet size for AKT-FR TAR jets. The apparent large difference between the fit parameters in the case of FR $R = 0.75$ jets is driven by small differences in the case of low p_T jets with transverse momenta below 500 GeV, but the functional form of the upper cuts remains well described for all jet definitions.




Figure 11.3: (a) Logarithmic regression of a p_T dependent upper cut on $D_2^{\beta=1.5}$ resulting in a 50%, 60% and 80% signal jet efficiency for AKT-FR $R=1.0$ TAR jets and (b) best fit values of the parameters p_i in dependence of the TAR jet size for AKT-FR jets for the 50% working point.

11.4 Jet Performance

As mentioned in Chapter 10, the TAR jet performance in this work is studied based on the signal event efficiency

$$\epsilon_s = \frac{N_{\text{events}}^{\text{pass}}}{N_{\text{events}}^{\text{total}}}, \quad (11.2)$$

where $N_{\text{events}}^{\text{total}}$ is the total number of events and $N_{\text{events}}^{\text{pass}}$ is the number of events passing the selections described in Section 11.2. The background efficiency ϵ_b is defined analogously, while ϵ_b^{-1} is referred to as background event rejection.

11.4.1 Signal Event Efficiency

Figure 11.4 shows the signal event efficiencies for all considered TAR jet definitions in the boosted 0ℓ and 1ℓ final state for the resonance masses $m_X = 1, 2$ and 3 TeV. The efficiencies are interpolated as a function of the TAR jet size for FR jets and in dependence of the energy scale for VR jets using cubic splines. Note that in the 1ℓ final state, the lower limit of TAR jet parameters is extended with a size parameter of $R = 0.25$ for FR TAR jets and an energy scale of $\rho = 100$ GeV for VR TAR jets. These values are below the thresholds used for standard small- R jets ($R = 0.4$) employed in ATLAS and are shown only for reference.

In the 1ℓ boosted final state, the signal efficiency for FR jets tends to increase until it

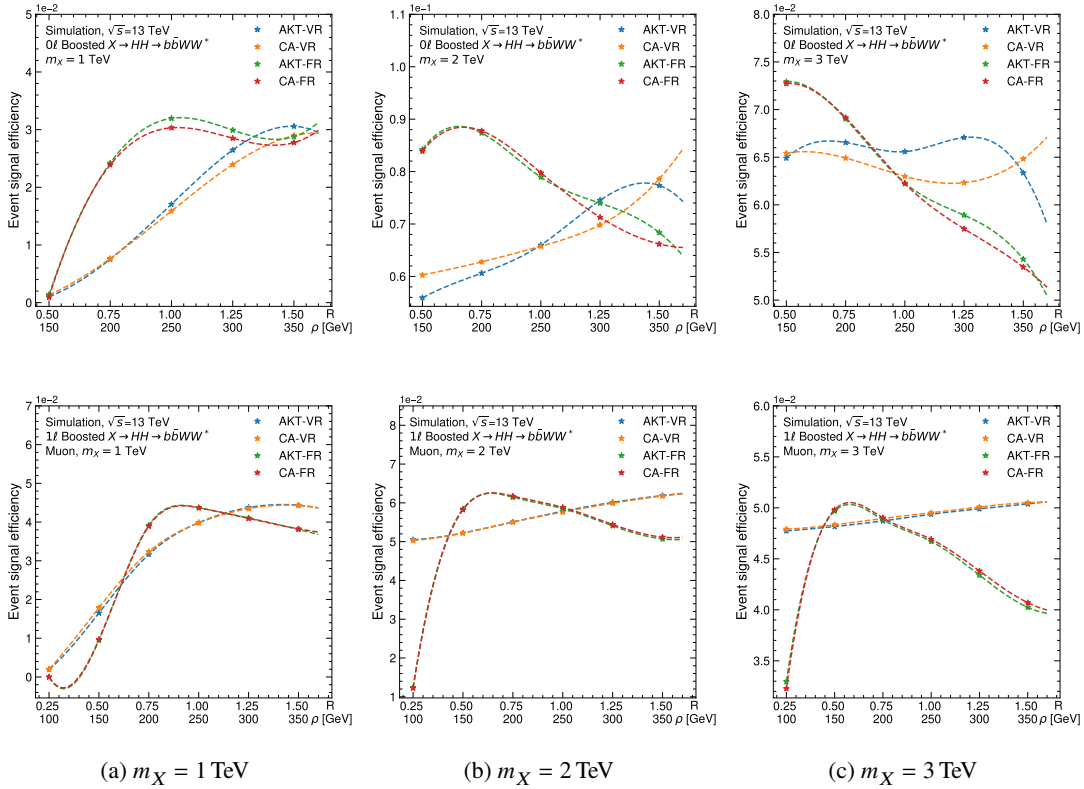


Figure 11.4: Signal event efficiencies for (a) $m_X = 1 \text{ TeV}$, (b) $m_X = 2 \text{ TeV}$ and (c) $m_X = 3 \text{ TeV}$ in the boosted 0ℓ (top row) and 1ℓ channel (bottom row) for the various TAR jet configurations. The efficiencies are interpolated using cubic splines as a function of the size parameter R in the case of fixed radius jets, and in dependence of the energy scale ρ in the case of variable radius jets. Note that the lower limit of TAR jet parameters is extended with a size parameter of $R = 0.25$ for FR TAR jets and an energy scale of $\rho = 100 \text{ GeV}$ for VR TAR jets in the boosted 1ℓ final state.

reaches a maximum, which is shifted to smaller jet sizes in the case of heavier resonances. For $m_X = 1$ TeV the maximum is at $R \sim 0.85$, for 2 TeV at $R \sim 0.6$, and for 3 TeV at $R \sim 0.55$. For jet sizes above this threshold, the efficiency drops significantly, with the slope being largest for the heaviest resonance considered. This effect is explained by the fact that for sufficiently small jet sizes, the TAR jets fail to collect all the decay products of the $H \rightarrow b\bar{b}$ decay, so that the efficiency drops as a consequence of the fixed $H \rightarrow b\bar{b}$ mass window cut imposed in this final state. On the other hand, for larger jet sizes, the probability of reconstructing the $H \rightarrow b\bar{b}$ and the $H \rightarrow W_{\text{had}}W_{\text{lep}}$ decay products in a single TAR jet increases, and consequently reduces the event signal efficiency. The optimum jet size with maximum efficiency shifts to smaller radii for heavier resonance masses, as the decay products become more boosted and as a result, smaller TAR jet sizes are sufficient to collect all decay products, and the probability of collecting all energy deposits in a single jet is reduced. The dependence of the effective jet size on the transverse jet momentum in the case of VR TAR jets allows to recover the loss of efficiency observed for FR jets for very low and large energy scales. In this case, the efficiency generally increases with larger values of ρ , with the slope being the largest for $m_X = 1$ TeV. In most cases, a tiny improvement is observed for CA jets with respect to AKT jets, but this can be considered negligible for all practical purposes.

In the 0ℓ boosted final state, the situation is more intricate due to the more complex event selection, but very similar effects can be observed. With respect to the 1ℓ boosted final state, the efficiencies with VR TAR jets are significantly worse compared to FR TAR jets, while the overall trend remains similar. In this case, the difference between the AKT and CA algorithms is much more pronounced. A slight improvement is observed when using the AKT algorithm for FR TAR jets, especially for large jet sizes, while for VR TAR jets the improvement is much more pronounced.

Figure 11.5 shows the signal event efficiencies for all considered TAR jet definitions in the split-boosted 0ℓ final state for the resonance masses $m_S = 0.75, 1.5$ and 2.5 TeV, where m_X is fixed to 3 TeV. For $m_S \leq 1.5$ TeV, the signal efficiency for FR TAR jets is largest for small jet sizes, peaking at $R \sim 0.75$ for $m_S = 750$ GeV and $R \sim 0.60$ for $m_S = 1.5$ TeV. For larger jet sizes, the event efficiency drops substantially from $\sim 9\%$ to $\sim 1\%$ for $m_S = 750$ GeV and from $\sim 11\%$ to $\sim 6\%$ for $m_S = 1.5$ TeV with respect to the smallest and largest jet sizes. Similar to the boosted final states, the effect is explained by the fact that the probability of reconstructing only two TAR jets in an event increases with larger jet sizes, such that the signal efficiency decreases accordingly. This feature is not observed for VR TAR jets, where the signal efficiency remains approximately constant over all energy scales considered, and shows efficiencies similar to small radius FR TAR jets. For $m_S = 2.5$ TeV, the value of m_S is sufficiently close to $m_X = 3$ TeV that the Higgs boson is no longer boosted. As a result, the angular distance between the b quarks (or W -bosons) decreases and the probability of reconstructing a B -hadron as part of a W boson jet increases. For FR TAR jets, this leads to an optimal jet size at $R \sim 0.8$, where the efficiency decreases for smaller and larger jet sizes. For VR TAR jets, on the other hand, the efficiency increases continuously from 5% for $\rho = 150$ GeV to 10% for $\rho = 350$ GeV. In the




Figure 11.5: Signal event efficiencies for (a) $m_S = 0.75$ TeV, (b) $m_S = 1.5$ TeV and (c) $m_S = 2.5$ TeV in the 0ℓ split-boosted channel for the various TAR jet configurations and resonant Higgs boson masses. The efficiencies are interpolated using cubic splines as a function of the size parameter R in the case of fixed radius jets, and in dependence of the energy scale ρ in the case of variable radius jets.

case of FR TAR jets, the CA reclustering algorithm features significant efficiency gains for most jet definitions, while no significant differences between AKT and CA jets are observed for VR TAR jets.

11.4.2 Top Quark Pair Event Rejection

In addition to the signal efficiency, a study of the background rejection is of course equally important to understand how the choice of the TAR jet definition may affect the overall sensitivity of the analyses. For this purpose, the $t\bar{t}$ event rejections are shown in Figure 11.6 for all studied TAR jet definitions for the 0ℓ and 1ℓ boosted, and 0ℓ split-boosted final states.

In the 0ℓ boosted final state, the largest $t\bar{t}$ rejection is achieved with VR jets for the full range of studied parameters, while the trend is reversed in the case of the 1ℓ boosted final state, where FR jets provide a better background rejection. In both cases, the rejection remains approximately constant for $R \geq 1$ ($\rho \geq 250$ GeV), but increases strongly for smaller jet sizes and energy scales. Top jets from $t\bar{t}$ events are typically reconstructed in a single large- R jet, where the B -hadron and the W boson from the $t \rightarrow bW$ decay are entirely contained within the jet. If one aims to discriminate against $H \rightarrow b\bar{b}$ jets, the TAR jet is typically required to contain two double b -tagged track jets, as is the case in both boosted final states. As a result, for large jet sizes, the TAR jet will only be misidentified if one of the subjets is mistagged by the b -jet identification algorithm. In the case of a sufficiently small jet size, not all $t \rightarrow bW$ decay products are collected, and the probability of collecting the single B -hadron decreases accordingly. Therefore, a misidentification of the TAR jets requires a mistag of both subjets, which effectively increases the rejection of top jets. Note that while this is generally true, the

increased $t\bar{t}$ rejection for small jet sizes and energy scales in the boosted 0ℓ final state is largely the result of the strict $p_T > 300$ GeV requirement on the jets. No significant differences are observed in the $t\bar{t}$ event rejection between the AKT and CA reclustering algorithms in the boosted final states.




Figure 11.6: Top-antitop ($t\bar{t}$) event rejections in the (a) 0ℓ boosted, the (b) 1ℓ boosted and the (c) 0ℓ split-boosted channel for the various TAR jet configurations. The rejections are interpolated using cubic splines as a function of the size parameter R in the case of fixed radius jets, and in dependence of the energy scale ρ in the case of variable radius jets. Note that the lower limit of TAR jet parameters is extended with a size parameter of $R = 0.25$ for FR TAR jets and an energy scale of $\rho = 100$ GeV for VR TAR jets in the boosted 1ℓ final state.

In the 0ℓ split-boosted final state, the $t\bar{t}$ rejection with FR TAR jets outperforms the rejection with VR TAR jets over the entire parameter range and up to a full order of magnitude for the smallest jets with $R = 0.5$ and $\rho = 150$ GeV. Similar to the boosted 0ℓ final state, the rejection increases for smaller jet sizes and energy scales as a result of the strict p_T and m_{TAR} requirements which are placed on the three jets whose presence is required in each event. The $t\bar{t}$ rejection reaches a minimum at $\rho = 200$ GeV for VR TAR jets and $R = 1.0$ for FR TAR jets. For larger jet sizes and energy scales, the rejection increases slightly as a consequence of the imposed $S \rightarrow W_{\text{had}}$ mass window cut. While the CA jets slightly outperform the AKT jets for small jet sizes and energy scales, the opposite is true for very large jets. In all cases, the observed differences between AKT and CA jets are negligible for practical purposes.

11.4.3 Efficiency Ratios

To facilitate the selection of an appropriate TAR jet definition, the quantity $\log(\epsilon_s/\epsilon_b)$ is calculated in each case and interpolated using cubic splines as a function of the TAR jet size and energy scale.

Figure 11.7 and Figure 11.8 show the interpolated logarithmic efficiency ratios for all

considered resonance masses in the boosted and split-boosted final states, respectively. In the

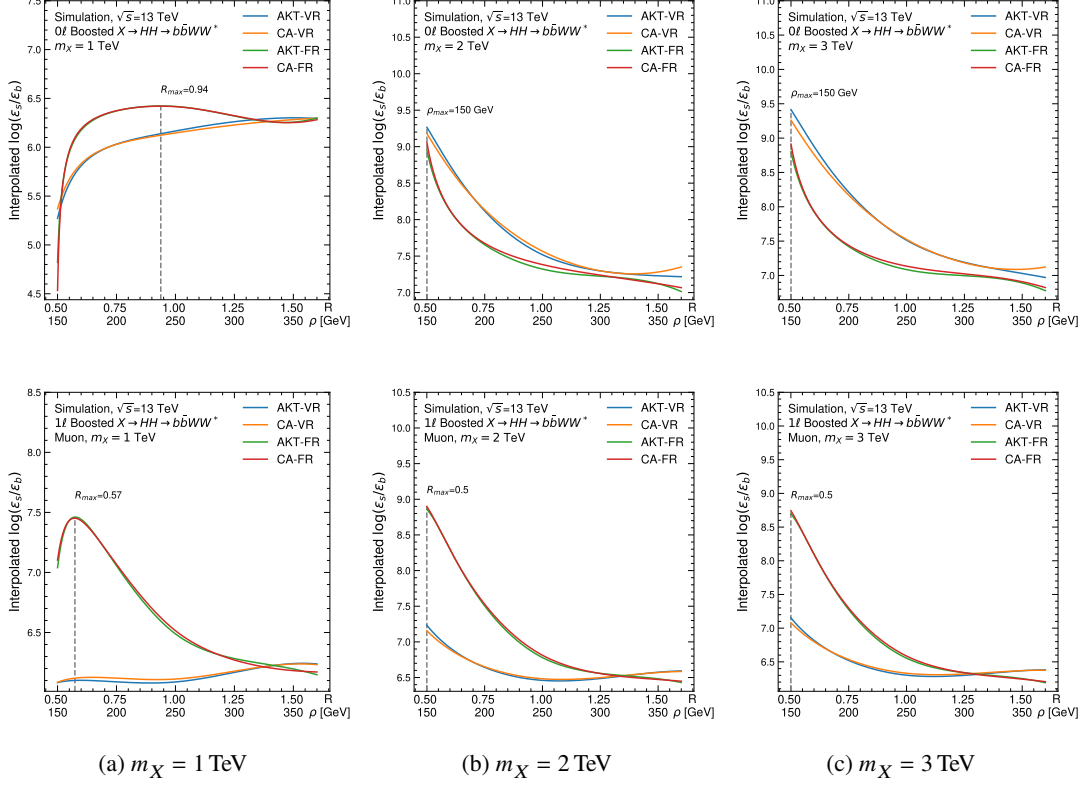


Figure 11.7: Interpolated logarithmic efficiency ratio $\log \epsilon_s/\epsilon_b$ for (a) $m_X = 1$ TeV, (b) $m_X = 2$ TeV and (c) $m_X = 3$ TeV in the boosted 0ℓ (top row) and in the boosted 1ℓ final state (bottom row) for the various TAR jet definitions as a function of the jet size and energy scale for FR and VR TAR jets, respectively.

case of the 1ℓ boosted final state, FR jets clearly have a better efficiency ratio with compared to VR jets, with the differences being largest for small jet sizes and energy scales. For VR TAR jets, the ratio remains approximately flat over the range of energy scales studied, while for FR jets it is optimal at $R \sim 0.57$ in the case of $m_X = 1$ TeV. For $m_X \geq 2$ TeV, the (local) maximum is at $R = 0.5$, which corresponds to the lower limit of the studied jet sizes.

For the 0ℓ boosted final state, a similar increase in the efficiency ratio is observed for smaller jet sizes and energy scales for $m_X \geq 2$ TeV, but VR jets perform fairly better over the studied parameter range. For $m_X = 1$ TeV, on the other hand, FR jets have a moderately better efficiency ratio, and for both, FR (VR) TAR jets, the ratio drops sharply for $R \leq 0.75$ ($\rho \leq 200$ GeV).

Finally, in the split-boosted final state, FR TAR jets generally provide a better efficiency ratio for small jet sizes and energy scales, while VR TAR jets are slightly better for $R \geq 1.25$ ($\rho \geq 300$ GeV) in most cases. For all studied mass points, the (local) maximum of the efficiency ratio is




Figure 11.8: Interpolated logarithmic efficiency ratio $\log \epsilon_s/\epsilon_b$ for (a) $m_S = 0.75$ TeV, (b) $m_S = 1.5$ TeV and (c) $m_S = 2.5$ TeV in the boosted 0ℓ (top row) and in the boosted 1ℓ final state (bottom row) for the various TAR jet definitions as a function of the jet size and energy scale for FR and VR TAR jets, respectively.

at $R = 0.5$, which corresponds to the lower limit of the studied jet sizes.

11.5 Summary and Recommendations

The final decision on the choice of the TAR jet definition is based on the following considerations:

1. *Signal efficiency*: considering all resonance masses studied, FR TAR jets with $R \sim 0.75$ offer the best compromise in terms of event signal efficiency for the 0ℓ boosted final state. For the 1ℓ boosted final state, FR TAR jets with $R \sim 0.75$ also yield close to optimal efficiency, while VR TAR jets with large energy scales are also a viable option. In the split-boosted case, VR TAR jets generally offer better efficiencies, but the gain with respect to FR TAR jets with $R \sim 0.75$ is fairly limited.
2. *$t\bar{t}$ event rejection*: in all cases, small jet sizes and energy scales are beneficial for $t\bar{t}$ background rejection. In the 0ℓ boosted final state, VR TAR jets offer better rejections, while FR TAR jets feature better rejections in the 1ℓ boosted and split-boosted final states. These results are particularly relevant in the 1ℓ boosted case, where $t\bar{t}$ is the leading background.
3. *$\log(\epsilon_s/\epsilon_b)$ ratio*: the ratio clearly favours FR TAR jets in the 1ℓ boosted and split-boosted final states, but shows a small benefit for VR TAR jets for $m_X \leq 2$ TeV in the 0ℓ boosted final state. Moreover, the smallest studied TAR jet size $R = 0.5$ is preferred in all cases, except for $m_X = 1$ TeV due to the steep drop in efficiency for smaller radii in the boosted final states.

In principle, the TAR jet definitions can be chosen independently in each final state. However, to allow a later combination of results between the 0ℓ boosted and split-boosted final states, both analyses must use the same jet definition. The 1ℓ final state does not have such a restriction, but is also chosen to use the same jet definition for the sake of harmonization. Taking this and the above findings into account, AKT-FR $R = 0.75$ jets are chosen for all final states. This choice allows a near-optimal signal efficiency in all cases, taking advantage of the increased background rejection for smaller jet sizes, and avoiding the steep drop in efficiency for even smaller jet radii in the case of lighter resonance masses. Since in general no significant differences between AKT and CA jets are observed, the AKT reclustering algorithm is chosen, which is the standard algorithm used in ATLAS.

11.6 Impact on Exclusion Limits

As a final validation of the choice of the TAR jet definition, upper limits on the resonant Higgs boson pair production cross section are computed with AKT-FR $R = 0.75$ jets, and compared with the upper limits observed with the standard large- R jet size of $R = 1.0$ used in ATLAS. The results are obtained by taking into account the full set-up of the respective 0ℓ boosted [236], 0ℓ split-boosted [237] and 1ℓ boosted [233] analyses. In particular, this includes a preliminary estimate of the multijet background, which is particularly relevant in the 0ℓ final states. Note that at the time of conducting these studies, the analyses for all final states had not been completed. As such, the shown upper limits are subject to change and are intended only to illustrate the relative impact that optimizing the TAR jet definition can have on the final exclusion limits.

Figure 11.9 shows the expected 95% CLs upper limits on resonant Higgs boson production $\sigma(pp \rightarrow X \rightarrow HH)$ for $R = 1.0$ and $R = 0.75$ FR TAR jets for both boosted final states. For the 1ℓ boosted final state, the expected upper limits are slightly improved with $R = 0.75$ TAR jets for all values of m_X . For the 0ℓ boosted final state, the limits tend to improve especially in the highly boosted region, while a slight degradation is observed in the low mass region. However, the FPT employed for the identification of boosted $H \rightarrow W_{\text{had}}W_{\text{had}}$ jets used to obtain these results is trained on standard FR $R = 1.0$ TAR jets. Retraining the FPT on the recommended smaller jet size is expected to improve tagging performance, so significantly improved exclusion limits are expected for $R = 0.75$ TAR jets over the entire mass range.

Figure 11.10 shows the expected 95% CLs upper limits on $\sigma(pp \rightarrow X \rightarrow SH) \times \mathcal{B}r(SH \rightarrow b\bar{b}WW^*)$ as a function of m_S for the 0ℓ split-boosted final state for standard FR $R = 1.0$ and optimized $R = 0.75$ TAR jets in the case of $m_X = 2$ TeV and $m_X = 3$ TeV. For all studied masses, a significant improvement in the upper limits is observed, especially for low values of m_S . The largest improvement is observed for $m_S = 400$ GeV ($m_X = 2$ TeV) and $m_S = 500$ GeV ($m_X = 3$ TeV), where the limits with FR $R = 0.75$ are improved by almost an order of magnitude with respect to the standard $R = 1.0$ jets.




Figure 11.9: Expected 95% CLs upper limits on $\sigma(pp \rightarrow X \rightarrow HH)$ as a function of m_X for the (a) 0ℓ and (b) 1ℓ final state for standard FR $R = 1.0$ and optimized $R = 0.75$ TAR jets. Note that the NN-based four-prong tagger employed in the 0ℓ boosted case is optimized for $R = 1.0$ FR TAR jets in both cases. The shown uncertainty bands correspond to standard FR $R = 1.0$ TAR jets.




Figure 11.10: Expected 95% CLs upper limits on $\sigma(pp \rightarrow X \rightarrow SH) \times \mathcal{B}r(SH \rightarrow b\bar{b}WW^*)$ as a function of m_S for the 0ℓ split-booster final state for standard FR $R = 1.0$ and optimized $R = 0.75$ TAR jets in the case of (a) $m_X = 2 \text{ TeV}$ and (b) $m_X = 3 \text{ TeV}$. The shown uncertainty bands correspond to standard FR $R = 1.0$ TAR jets.

Conclusion

This thesis introduced the state-of-the-art fast simulation tool `ATLFAST3`, which has been employed by the ATLAS Collaboration for simulating about 7 billion events for the Run 2 MC reprocessing campaign and is expected to be used for Run 3 and beyond. The computationally demanding simulation of electromagnetic and hadronic showers using `GEANT4` is replaced with a parametrized simulation of the calorimeter response with `FASTCALOSIMV2`, which is the central focus of this work.

Many significant improvements in the new fast calorimeter simulation used in ATLAS were discussed in detail. Among various aspects, a novel procedure to correct for differences in the simulated total energy response between `FASTCALOSIMV2` and `GEANT4` was presented. For the simulation of photons, electrons and pions, differences in the simulated energy response resulting from imperfections in the energy parametrization and the subsequent reconstruction and digitization algorithms are corrected. In addition, a correction model has been developed for the charged pion parametrization that allows to accurately simulate the energy response of protons, neutrons, kaons and their respective antiparticles. Both corrections have been rigorously validated and have been shown to precisely reproduce the detailed `GEANT4` simulation.

This work also provided insights into the prospective developments of fast calorimeter simulation in ATLAS for Run 3 and the high-luminosity era of the LHC. It was shown that the `FASTCALOSIMV2` simulation of electromagnetic shower properties can be modified to accurately describe the shower shapes observed in data. This was achieved by implementing two separate mathematical models, one for the strip layer and another for the EM middle layer of the ATLAS calorimeter. A distributed large-scale optimization procedure was developed to extract optimal model parameters for all parametrized energies and relevant detector regions. The optimized tuning model has been extensively validated on photon- and electron-induced showers from unconverted and converted photons reconstructed in $H \rightarrow \gamma\gamma$ events, and has been shown to reproduce the mean of shower shape variables typically used to identify photon and electron candidates with an accuracy of less than a few percent in most cases. This is a vast improvement over the detailed `GEANT4` simulation in ATLAS, which has been known to inaccurately describe

electromagnetic shower shapes since the start of Run 1 of the LHC. In the near future, the models are expected to be further developed with the aim of completely eliminating the need for *ad-hoc* post-simulation corrections of shower shapes, which are presently necessary as a result of inaccurately simulated electromagnetic showers and require significant resources to develop.

This thesis also presented a first ISF-independent implementation of FASTCALOSIMV2 as a fast simulation model in the GEANT4 simulation toolkit. The novel implementation has been shown to accurately reproduce the current simulation and marks the first step towards a substantial refactoring of the simulation infrastructure, which is anticipated to greatly streamline the overall ATLAS simulation workflow. Moreover, the disentanglement of FASTCALOSIMV2 from ISF constitutes a first step towards broadening its exclusive use in the ATLAS experiment to other current or future high energy physics experiments as a versatile framework for fast calorimeter simulation. In the near future, ATLAS intends to integrate all available fast simulation tools such as FASTCALOGAN and FATRAS into GEANT4, effectively minimizing the reliance on any external tools. This updated configuration is expected to offer a robust, easily maintainable, and unified simulation infrastructure for the high-luminosity phase of the LHC.

Finally, this work presented performance studies of Track-Assisted Reclustered jets that have been conducted in the context of two searches for $X \rightarrow HH \rightarrow b\bar{b}WW^*$ production in the boosted 0ℓ and 1ℓ final states and a search for $X \rightarrow SH \rightarrow b\bar{b}WW^*$ production in the split-boosted 0ℓ final state. The use of information from the ID allows for a significantly improved resolution of jet substructure variables, enhancing the tagging capabilities of TAR jets over conventional jets whose reconstruction relies solely on topological clusters in the calorimeter. With respect to alternative jet reconstruction algorithms that also use information from the ID, such as UFO jets, TAR jets offer enhanced flexibility thanks to its bottom-up propagation of uncertainties. The exact reconstruction algorithm and jet size can be arbitrarily adapted to the specific analysis without the need of developing custom calibration procedures. With respect to the default $R = 1.0$ jet size used to reconstruct boosted topologies in ATLAS, important improvements in the sensitivity of the searches could be achieved with AKT-FR $R = 0.75$ jets in all cases considered. The smaller jet size is found to provide improved background jet rejection, while providing near-optimal signal efficiencies. The overall improvements in the expected 95% CLs upper limits range from a few percent to up to a full order of magnitude for low values of m_S in the split-boosted topology.

At the time of writing, Run 3 of the LHC is actively underway, delivering proton-proton collision data at a record energy of 13.6 TeV. The LHC will undergo major upgrades during LS3 and is scheduled to start in the HL-LHC configuration as early as 2029, delivering a 5- to 7.5-fold increase in instantaneous luminosity with respect to the LHC's design value. The recent advancements in fast simulation in ATLAS will allow most measurements and searches to use ATLFAST3 to simulate the detector response of physics processes, almost entirely eliminating the need for GEANT4. As a result, it will be possible to simulate many more MC events while adhering to computational and budgetary constraints, effectively reducing the

systematic uncertainties of many physics analyses and ultimately increasing the discovery potential of BSM searches and the accuracy of precision measurements.

By 2035, the LHC is anticipated to accumulate up to 3 ab^{-1} of integrated luminosity with which sensitivity to the Higgs boson self-coupling in non-resonant Higgs boson pair production is expected to be gained. With more and more data delivered by the LHC, searches for resonant Higgs boson pair production will also become increasingly relevant. These searches will benefit in particular from the greatly improved simulation of boosted jet substructure with ATLFast3, which will help in the rejection of the vast background expected in the dense and challenging environment of the high-luminosity era of the LHC.

Author's contribution

With around 6000 members, the ATLAS experiment is one of the largest collaborative efforts ever attempted in science. The results obtained from performance work and the analysis of physics data are typically a result of a combined effort of people from various research groups located around the world. The purpose of this statement is to provide a brief description of the author's personal contributions.

For the past three years, the author has been the lead developer and maintainer of FASTCALOSIMV2, and contributed significantly to the success of the Run 2 MC reprocessing campaign during which about 7 billion ATLFAST3 events were simulated. The most significant contributions were:

- Energy Response Corrections (Chapter 5): The author was fully responsible for the development and derivation of the energy response corrections for the Run 2 and Run 3 production of FASTCALOSIMV2. This includes the development of both correction models and the private simulation of a vast number of GEANT4 samples used to derive the corrections for protons, neutrons, kaons, and their antiparticles, as well as the simulation of FASTCALOSIMV2 events for the corrections of the energy response for electrons, photons and pions. The produced GEANT4 hadron samples are presently also being used for the training of Generative Adversarial Networks in FASTCALOGAN.
- Shower Centre Determination (Chapter 6): The author was fully responsible for the redesign of the extrapolation engine used in ATLFAST3 to determine the positions of shower centres within the ATLAS calorimeter, including all algorithmic developments and efforts to validate the novel approach.
- Physics Validation of ATLFAST3 (Chapter 7): The validation of the overall physics performance was a joint effort of members of the FASTCALOSIMV2 team and experts from the various performance groups in ATLAS.
- Tuning of Electromagnetic Shower Properties to Data (Chapter 8): The author was fully responsible for the development, optimization and validation of the data-driven models for electromagnetic shower simulation, including the development of a distributed large-scale black-box infrastructure. Relevant shower shape variables targeted in the data tune were identified together with e/γ experts in ATLAS. The data set used for the tuning procedure was also provided by the e/γ group. The implementation of the shower shape variables within FASTCALOSIMV2 is documented in Ref. [253] and was carried out with the assistance of a student supervised by the author. A report documenting the development of the models and their optimization for a single parametrization point

used in FASTCALOSIMV2 was prepared by the author and published by the ATLAS Collaboration in Ref. [213].

- FASTCALOSIMV2 as a GEANT4 Fast Simulation Model (Chapter 9): The author was fully responsible for the development of an ISF-independent implementation of FASTCALOSIMV2 as a GEANT4 fast simulation model, including the elaboration of a strategy to ensure unaltered physics results and the extensive validation performed of the new configuration.

A research paper presenting the ATLFast3 simulator, including the energy response corrections and the redesigned shower centre determination in the calorimeter described in Chapter 5 and Chapter 6 of this thesis, has been published by the ATLAS Collaboration in Ref. [134].

In addition, the author also made significant contributions to searches for resonant Higgs boson pair production (resonant HH and SH searches) in the $b\bar{b}WW^*$ final state:

- Track-Assisted Reclustered Jet Optimization (Chapter 11): The author was fully responsible for the performance studies on TAR jets conducted in the context of two searches for HH production in the boosted 0ℓ and 1ℓ final states and a search for SH production in the split-boosted 0ℓ final state. The studies are based on methods which were initially developed and published by the author in Ref. [235]. The respective analysis groups provided the preliminary selections used to perform the studies as well as the final 95 % CLs upper limits obtained with the optimized TAR jet configurations.
- Statistical analysis in the boosted 1ℓ state (ongoing): Based upon early initial studies published in Ref. [254], the author was the main responsible for the statistical analysis in the search for $X \rightarrow HH \rightarrow b\bar{b}WW^*$ production in the boosted 1ℓ final state, including the integration of systematic uncertainties as nuisance parameters. These studies remain to be finalized and are therefore not presented in this work.

Acknowledgements

Working at the world's largest particle physics laboratory over the past few years has taught me many lessons. One of them is that collaboration is the essence of any great scientific endeavour. This thesis would not have been possible without the support of many people to whom I would like to express my deepest gratitude.

My greatest thanks go to my supervisors, Stanley Lai and Michael Dürrssen-Debling. Stan, when I started my undergraduate studies in Göttingen, I did not expect to graduate with a PhD from the same university. Your guidance and support played a big part in keeping me attached to my *alma mater*. Thank you for not letting me go. Michael, I cannot stress enough how grateful I am to have had such a knowledgeable advisor throughout my time as a PhD student. You always found time for me and my questions, no matter how much else was going on, and your advice was always very valuable to me. Thank you both for your unwavering support and for making this journey as smooth as I could have ever hoped for.

I would like to especially thank everyone that indirectly contributed to this work:

- *John Derek Chapman* for your support in fighting all the little quirks of the Athena software and *Marilena Bandieramonte* for your support with GEOMODEL, which helped to develop the first implementation of FASTCALOSIMV2 as a GEANT4 fast simulation model.
- *Zachary Marshall and Andrea Dell'Acqua* for your support and knowledge on the GEANT4 side, which helped to identify the optimal strategy to ensure unaltered physics results in the ISF-independent implementation of FASTCALOSIMV2.
- *Andreas Salzburger* for his insights into the ATLAS tracking software, which helped in the development of the ATLFAST3 shower centre determination algorithm.
- *Julien Maurer* for your expertise in the e/γ reconstruction and identification software, which was crucial to the development of the data-driven model of fast EM shower simulation.

- *Xavier Bouthillier* and *Ben Jones* for your help in setting up the necessary infrastructure for a reliable large-scale distributed black-box optimization with Orion, which was essential for the optimization of the data-driven model of fast EM shower simulation.

I would like to extend my gratitude to all members of the FASTCALOSIM and the *bbWW** analysis group. It has been a truly enjoyable experience collaborating with each of you throughout these years.

My thanks also go to all the colleagues I have met over the years who have made my work more enjoyable in so many ways. Kira, I will remain forever indebted to you for your endless support. My first encounter with high energy physics was under your supervision, and your guidance played a significant role in fostering my interest in the field. A big thank you to all my colleagues at CERN: Corentin Allaire, Robert Langenberg, Martin Errenst, Andreas Stefl, Bastian Schlag, Stephen Swatman, Paul Gessinger, Joana Niermann, Hans Joos and Alexander Pflieger. Thank you for the many well-deserved coffee breaks and the interesting discussions during lunch. Thanks also to all my former colleagues from Göttingen, including Lino Gerlach, Petar Bokan, Serhat Ördek, Andrés Melo and Naman Kumar Bhalla.

I am deeply grateful to the German Federal Ministry of Education and Research for awarding me the Wolfgang Gentner Fellowship, which has provided me with the opportunity to conduct my entire PhD research at CERN.

Abschließend möchte ich meiner Familie, meinen Freunden und all denjenigen meine Dankbarkeit ausdrücken, die mein Leben außerhalb der Teilchenphysik bereichert haben. Mein größter Dank gebührt meinen wundervollen Eltern, Heike und Michael. Ihr habt maßgeblich dazu beigetragen, dass ich zu der Person geworden bin, die ich heute bin, und mich bei all meinen Entscheidungen uneingeschränkt unterstützt. Dafür werde ich euch für immer dankbar sein.

Köszönöm, Gabi, hogy támogatál a doktori képzés alatt, és még akkor is, amikor elvesztem a részecskefizikában, te mindig megtaláltál! Szeretlek.

APPENDIX A

Energy Response Corrections for Run 3

Appendix A Energy Response Corrections for Run 3

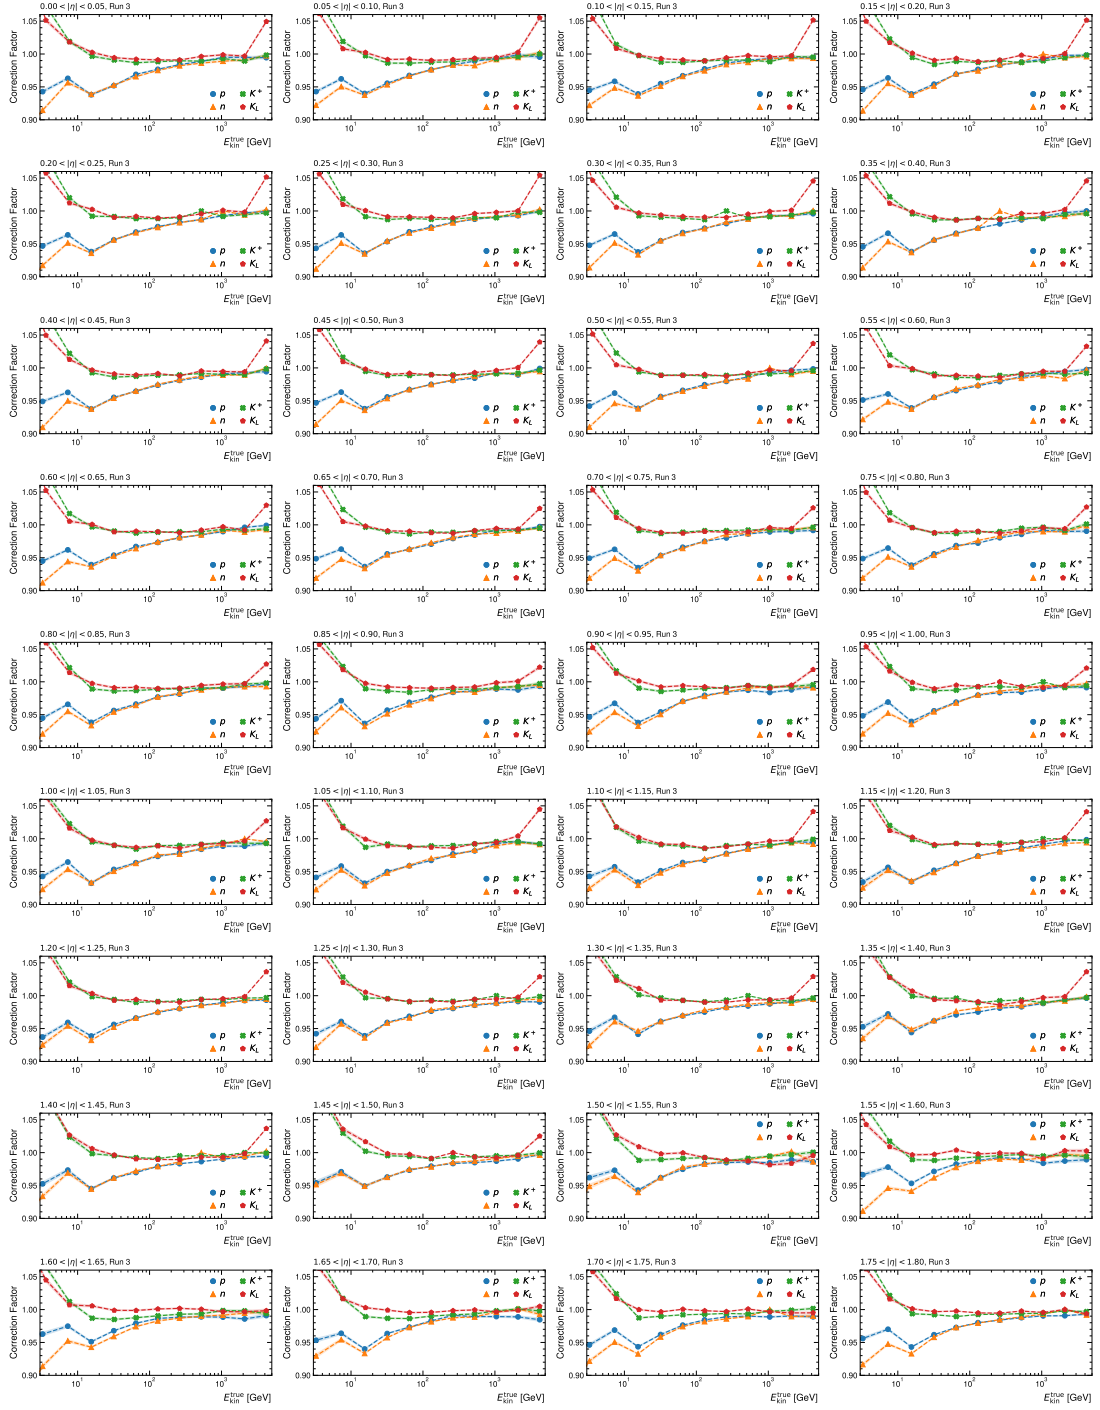


Figure A.1: Energy response correction factors as a function of true kinetic energy of the incident hadron for protons, neutrons, positively charged kaons, and long-lived kaons as derived for the Run 3 production of FASTCALOSIMV2 in the region of $0 < |\eta| < 1.80$. The shaded bands indicate the statistical uncertainties.

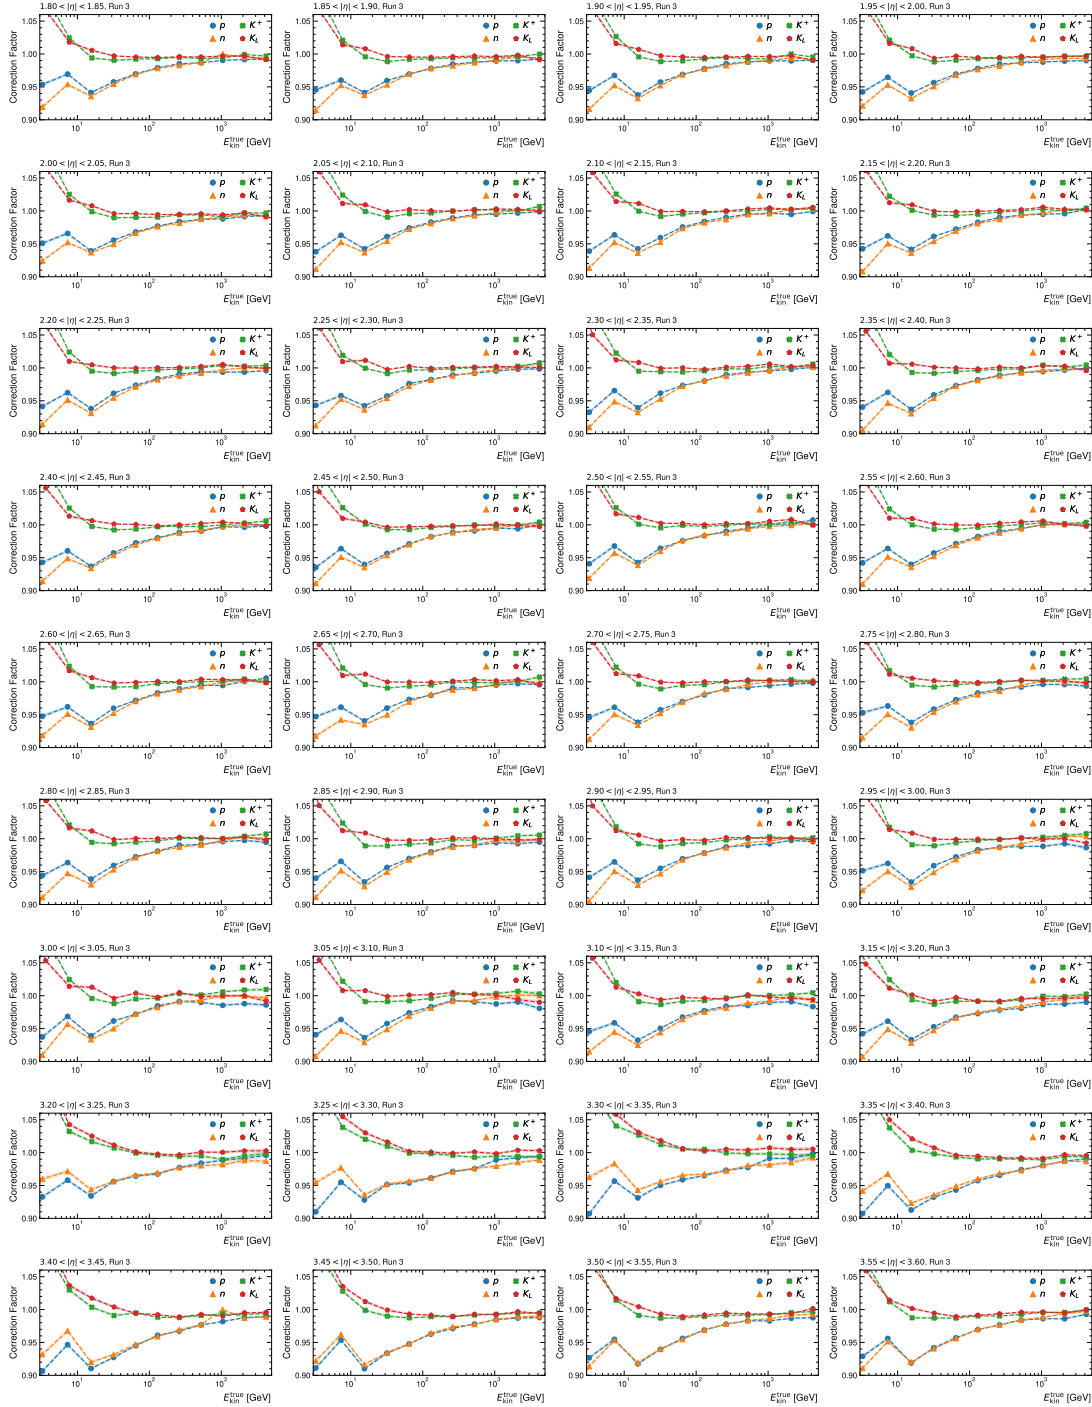


Figure A.2: Energy response correction factors as a function of true kinetic energy of the incident hadron for protons, neutrons, positively charged kaons, and long-lived kaons as derived for the Run 3 production of FASTCALOSIMV2 in the region of $1.80 < |\eta| < 3.60$. The shaded bands indicate the statistical uncertainties.

Appendix A Energy Response Corrections for Run 3

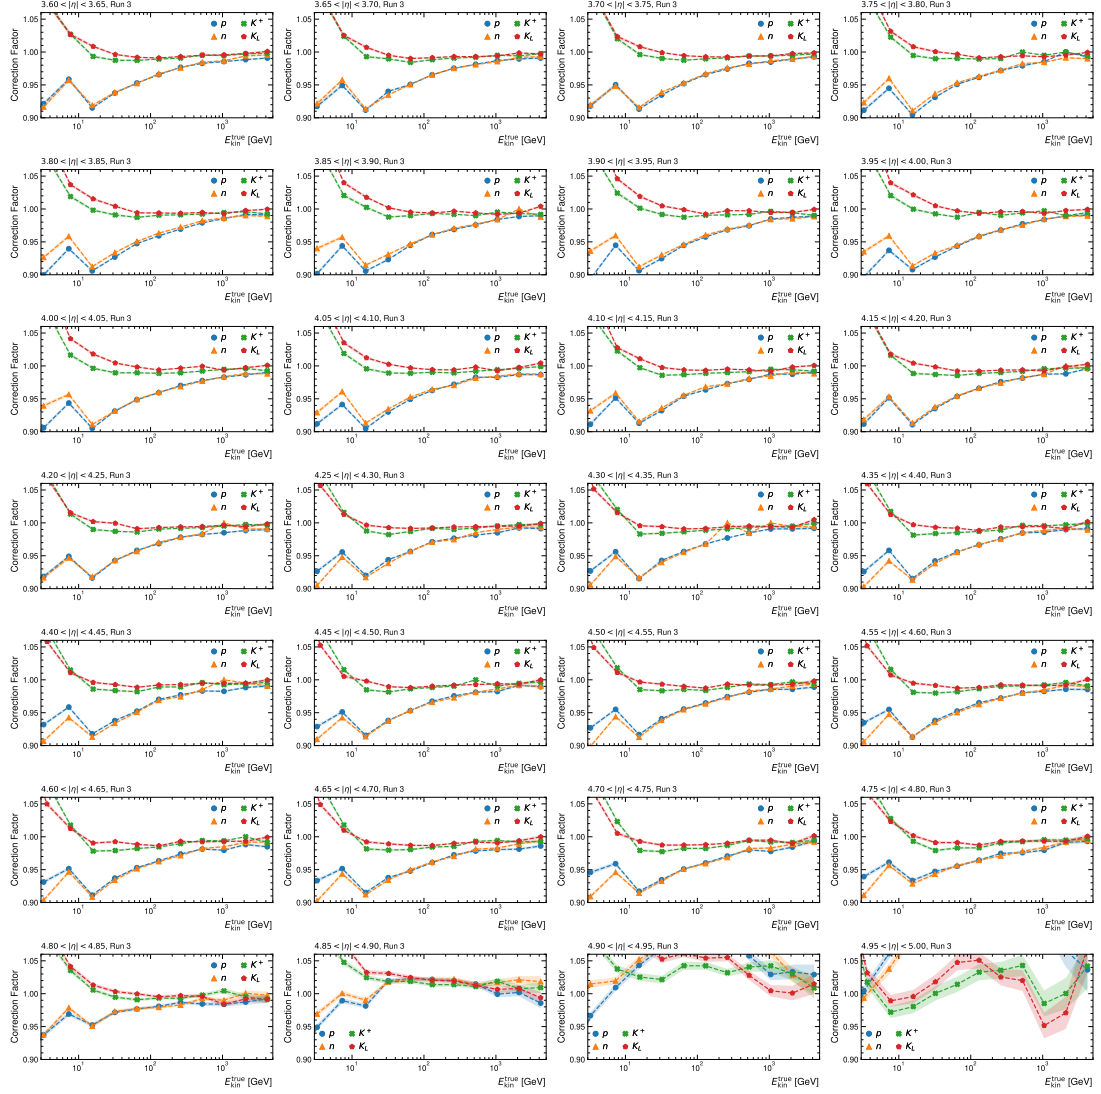


Figure A.3: Energy response correction factors as a function of true kinetic energy of the incident hadron for protons, neutrons, positively charged kaons, and long-lived kaons as derived for the Run 3 production of FASTCALOSIMV2 in the region of $3.60 < |\eta| < 5.0$. The shaded bands indicate the statistical uncertainties.

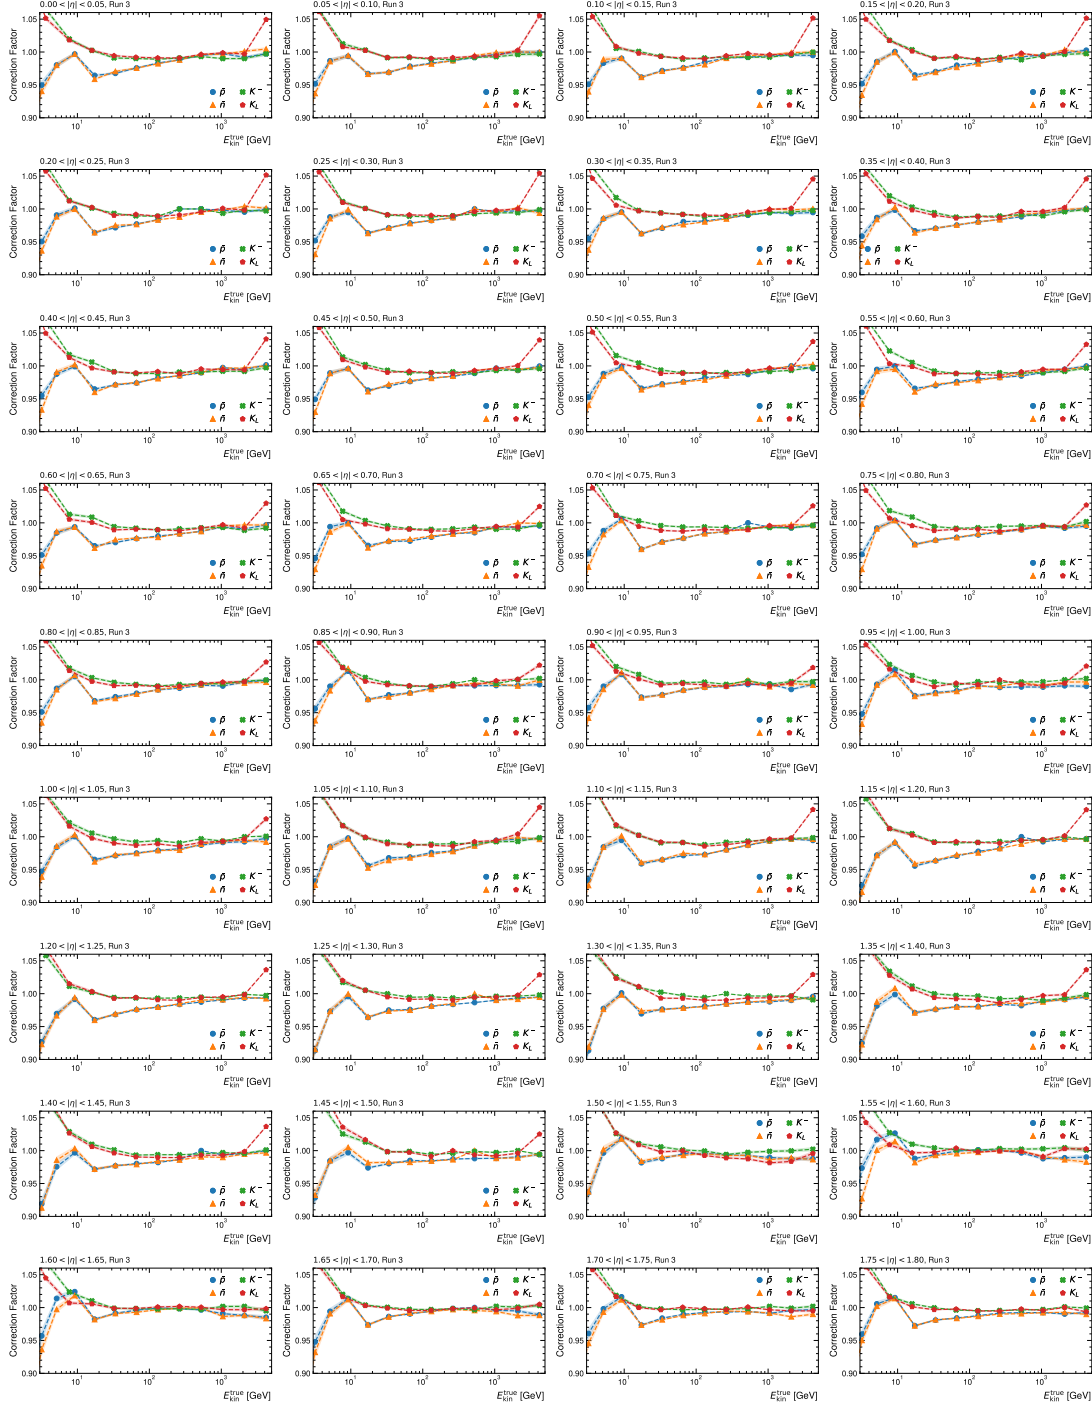


Figure A.4: Energy response correction factors as a function of true kinetic energy of the incident hadron for antiprotons, antineutrons, negatively charged kaons, and long-lived kaons as derived for the Run 3 production of FASTCALOSIMV2 in the region of $0 < |\eta| < 1.80$. The shaded bands indicate the statistical uncertainties.

Appendix A Energy Response Corrections for Run 3

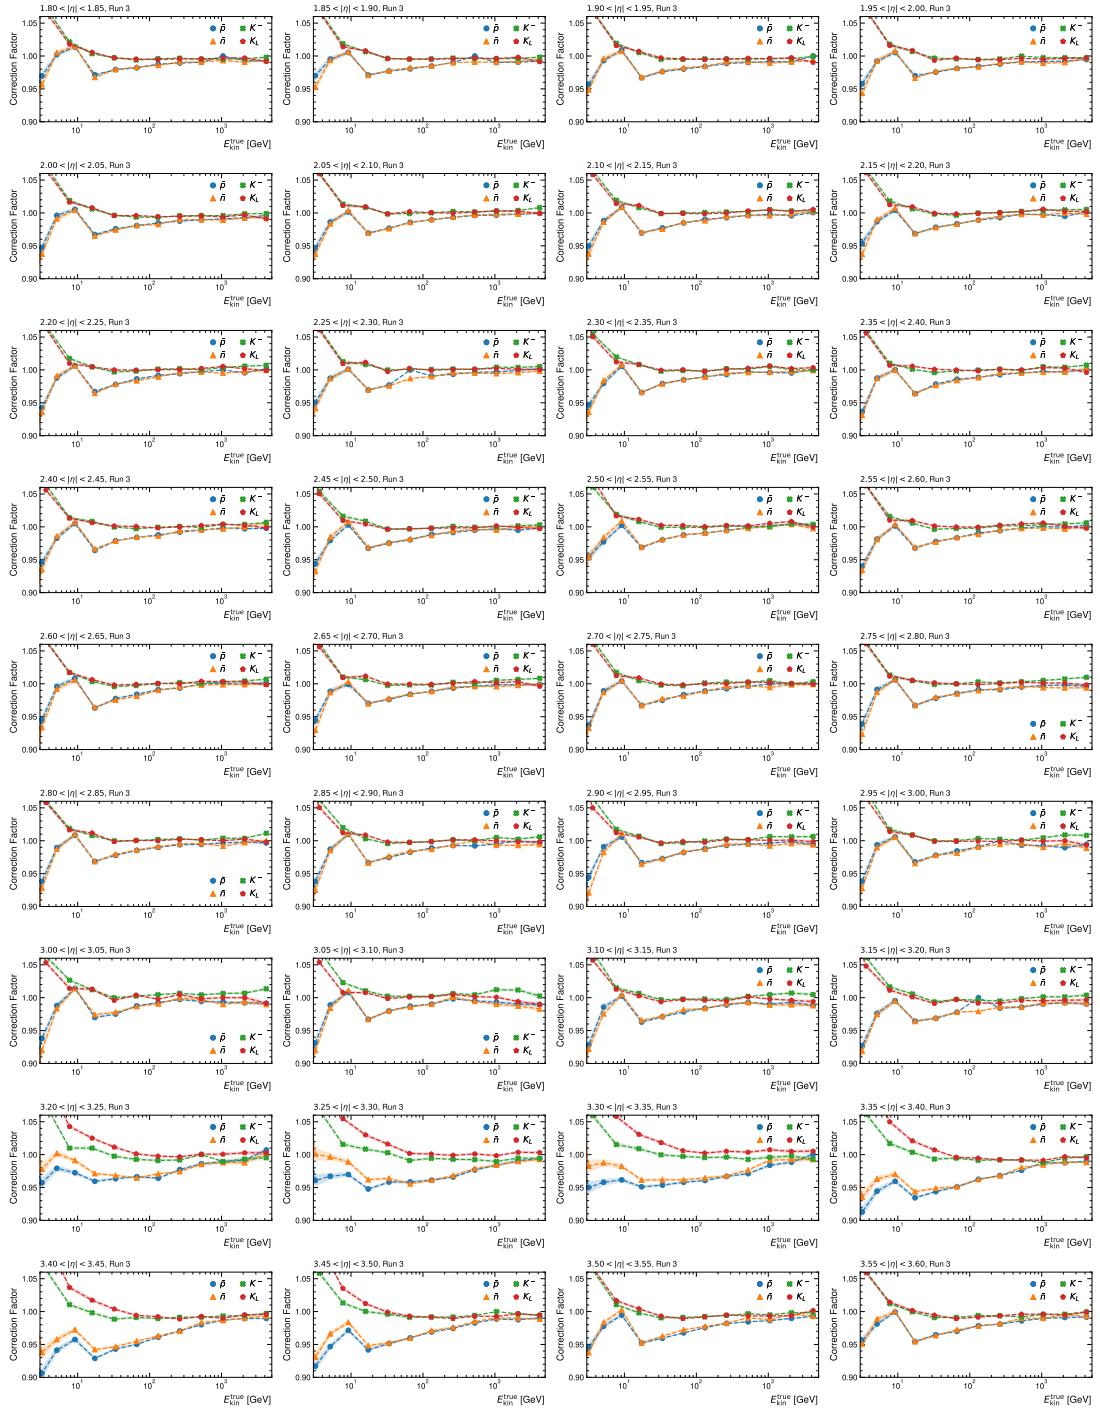


Figure A.5: Energy response correction factors as a function of true kinetic energy of the incident hadron for antiprotons, antineutrons, negatively charged kaons, and long-lived kaons as derived for the Run 3 production of FASTCALOSIMV2 in the region of $1.80 < |\eta| < 3.60$. The shaded bands indicate the statistical uncertainties.

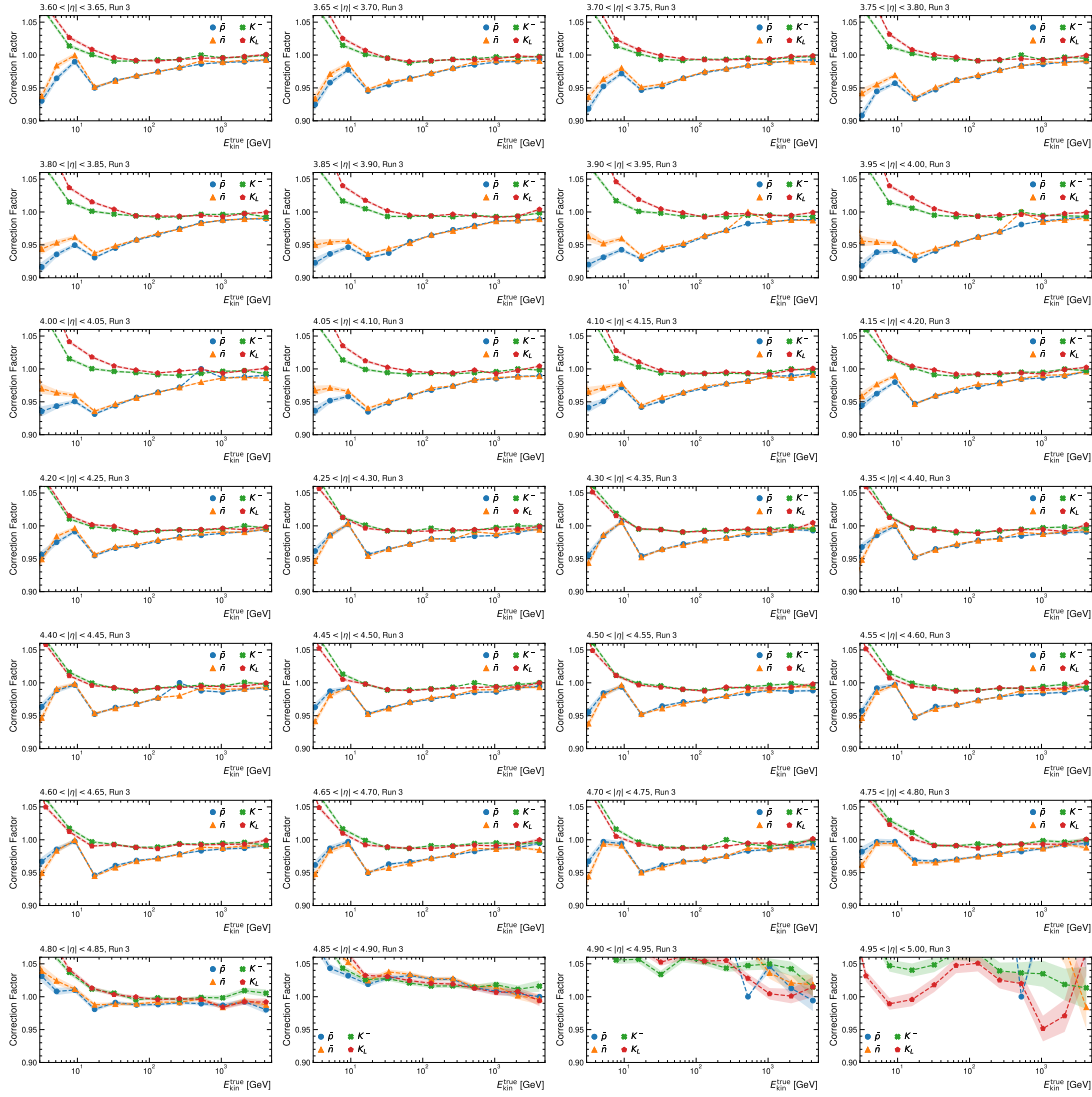


Figure A.6: Energy response correction factors as a function of true kinetic energy of the incident hadron for antiprotons, antineutrons, negatively charged kaons, and long-lived kaons as derived for the Run 3 production of FASTCALOSIMV2 in the region of $3.60 < |\eta| < 5.0$. The shaded bands indicate the statistical uncertainties.

APPENDIX **B**

FastCaloSim in Geant4 Hit Validation

Appendix B FastCaloSim in Geant4 Hit Validation

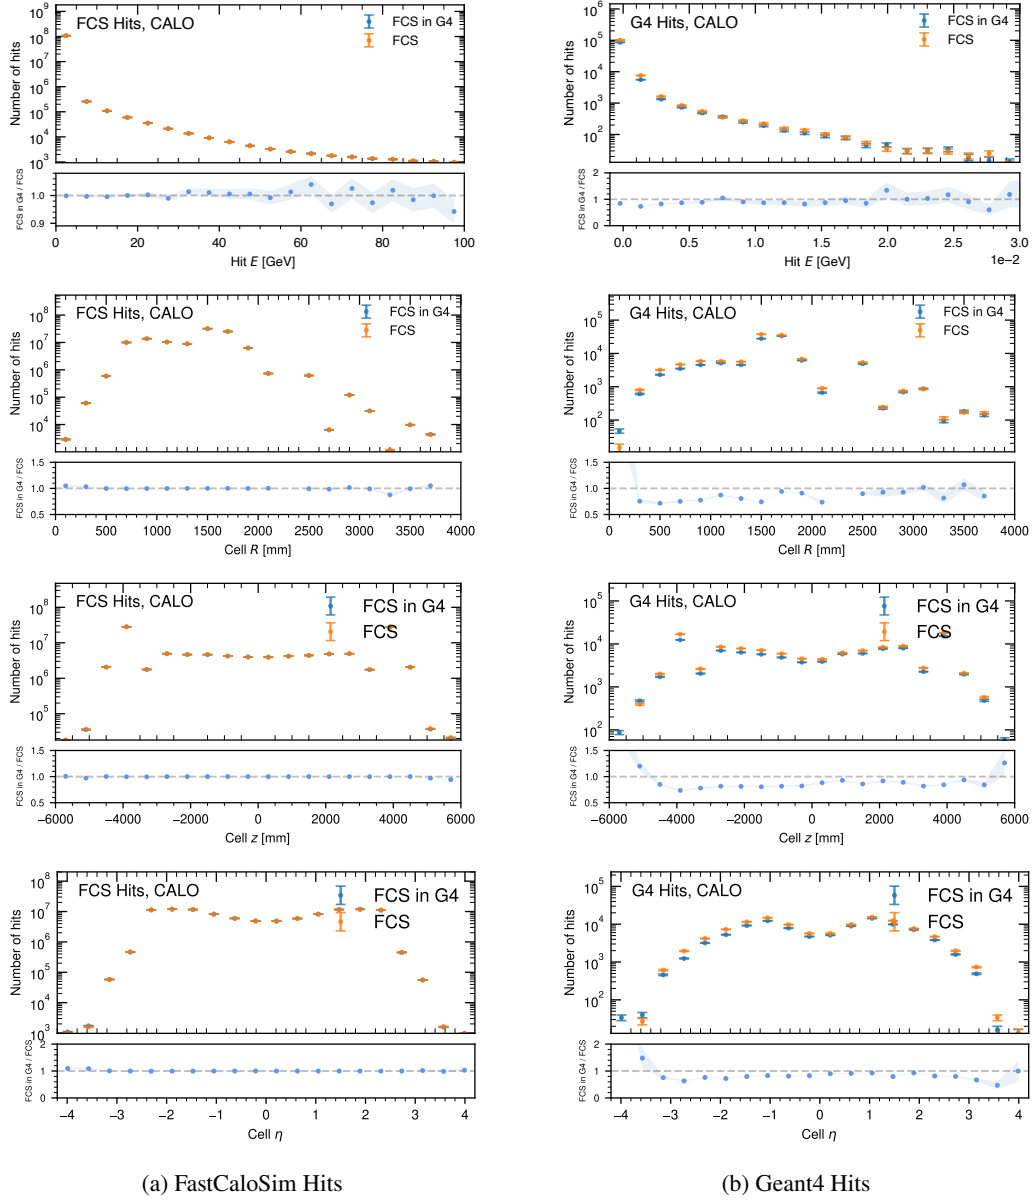










Figure B.1: Simulated hit energy and distributions of the cell R , η and z position of the simulated hits for the GEANT4 and ISF implementation of FASTCALOSIMV2, split into the simulation of FASTCALOSIMV2 hits (left) and GEANT4 hits (right). The CALO::CALO envelope is used as trigger volume without additional boundary checks.



(a) FastCaloSim Hits

(b) Geant4 Hits

Figure B.2: Simulated hit energy and distributions of the cell R , η and z position of the simulated hits for the GEANT4 and ISF implementation of FASTCALOSIMV2, split into the simulation of FASTCALOSIMV2 hits (left) and GEANT4 hits (right). The CALO+ :CALO envelope is used as trigger volume with an additional positional and directional boundary check.

Appendix B FastCaloSim in Geant4 Hit Validation

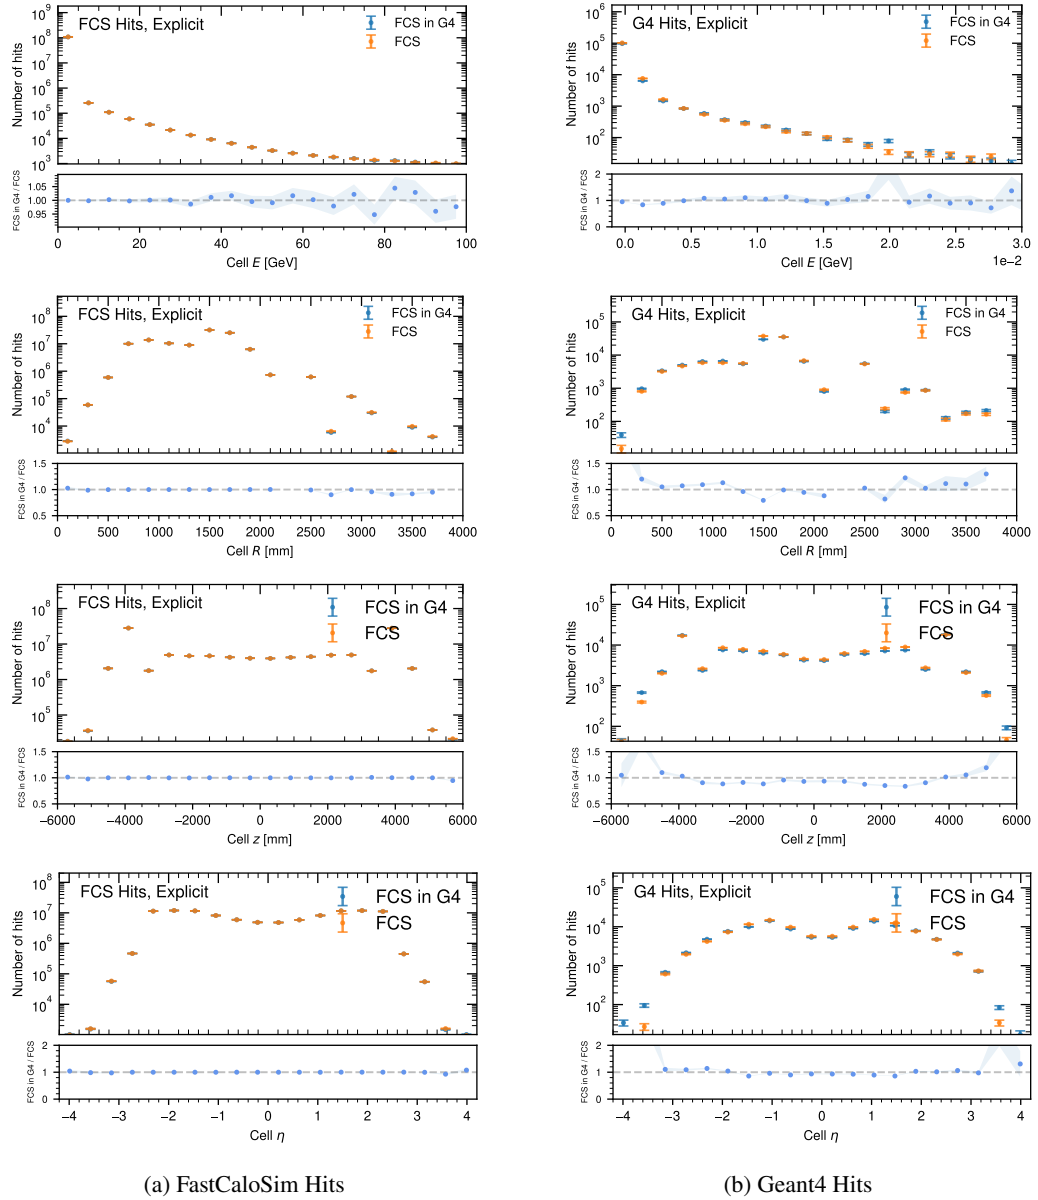










Figure B.3: Simulated hit energy and distributions of the cell R , η and z position of the simulated hits for the GEANT4 and ISF implementation of FASTCALOSIMV2, split into the simulation of FASTCALOSIMV2 hits (left) and GEANT4 hits (right). The set of explicit volumes is used as trigger volumes without additional boundary checks.



(a) FastCaloSim Hits

(b) Geant4 Hits

Figure B.4: Simulated hit energy and distributions of the cell R , η and z position of the simulated hits for the GEANT4 and ISF implementation of FASTCALOSIMV2, split into the simulation of FASTCALOSIMV2 hits (left) and GEANT4 hits (right). The set of explicit volumes is used as trigger volumes with an additional positional and directional boundary check.

Appendix B FastCaloSim in Geant4 Hit Validation

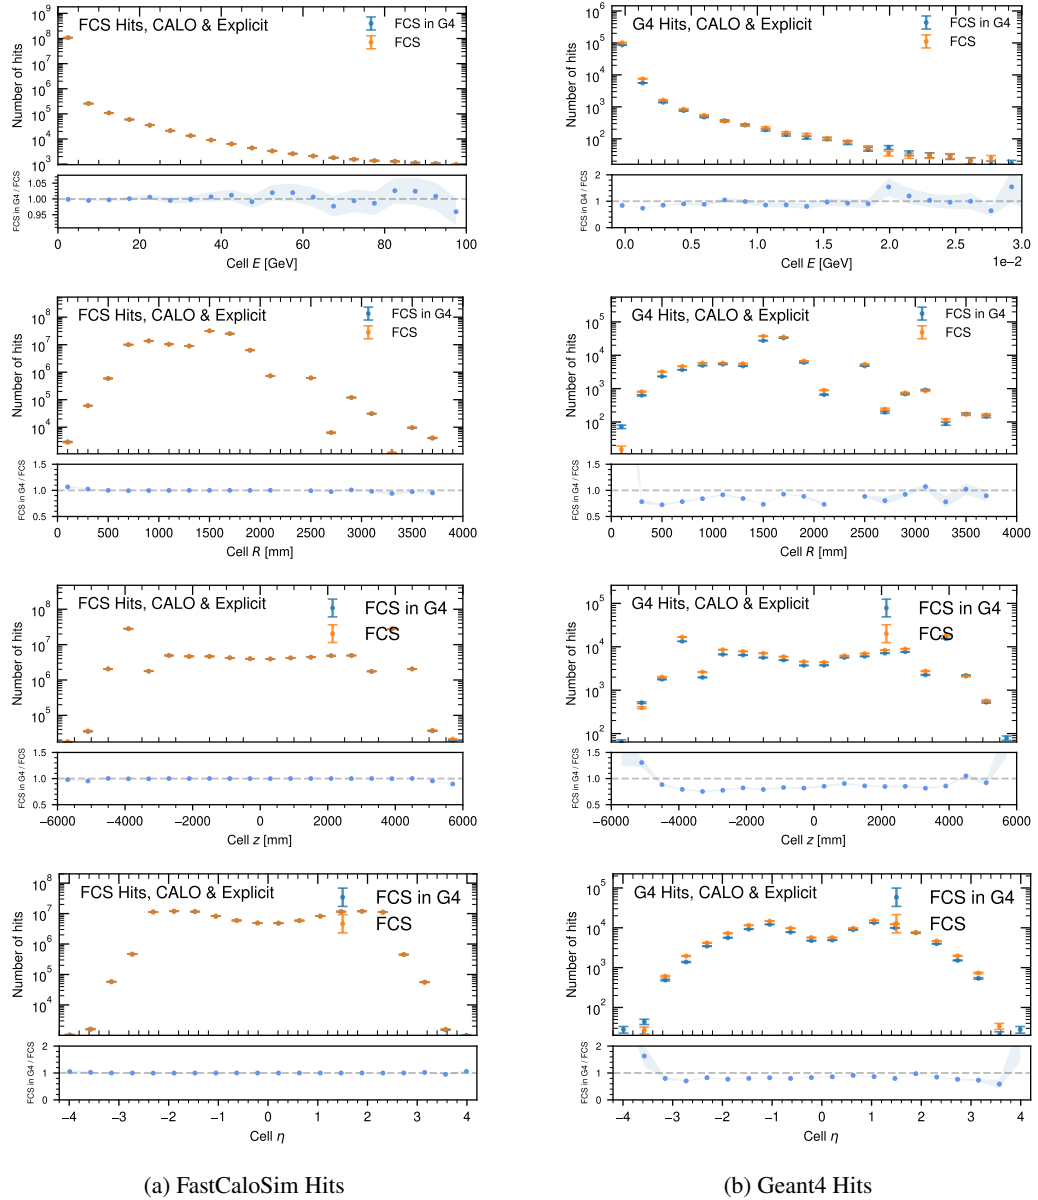


Figure B.5: Simulated hit energy and distributions of the cell R , η and z position of the simulated hits for the GEANT4 and ISF implementation of FASTCALOSIMV2, split into the simulation of FASTCALOSIMV2 hits (left) and GEANT4 hits (right). The Run 3 calorimeter envelope and a complementary list of explicit volumes is used as fast simulation trigger without additional boundary checks.

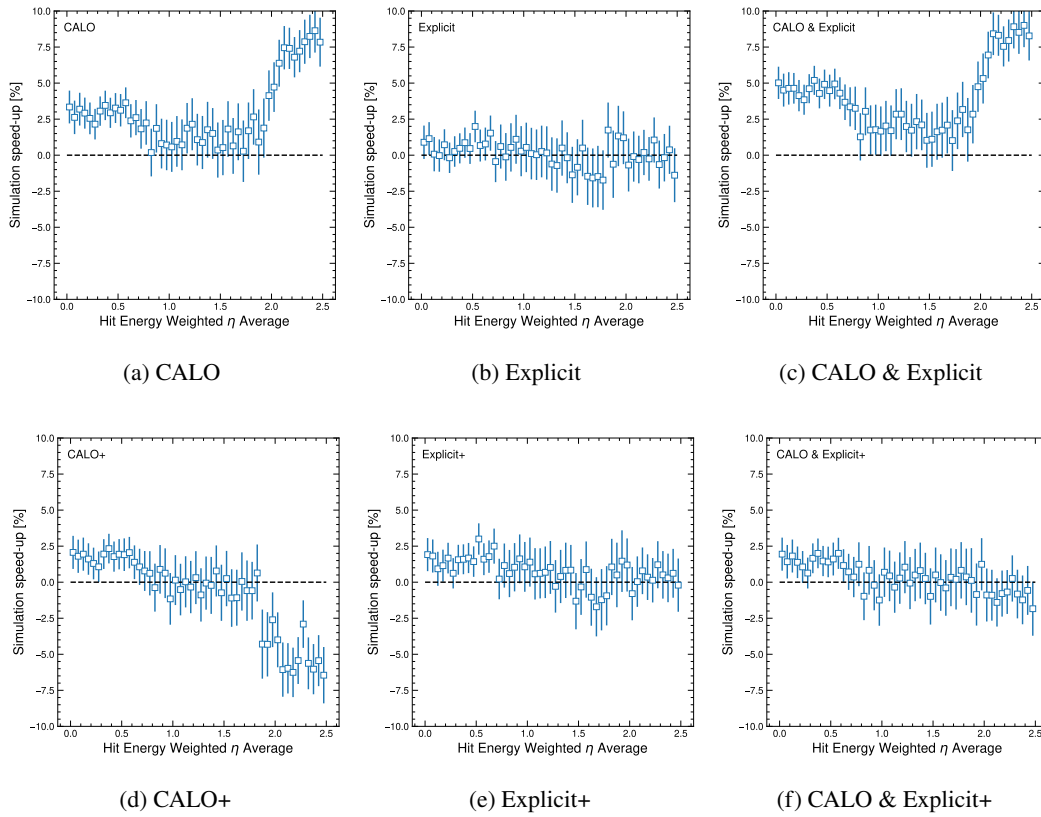


Figure B.6: Simulation speed-up in dependence of the hit energy weighted η average for the various tested fast simulation trigger volume configurations. The "+" denotes cases in which an additional positional and directional boundary check is imposed. Benchmarks are performed on a 64-core AMD EPYC™ 7702 CPU at 2 GHz base clock.

TAR Jet Tagging Performance

In this section, three simple cut-based taggers to identify $H \rightarrow b\bar{b}$, $H \rightarrow WW^*$ and isolated $S \rightarrow W_{\text{had}}$ jets are developed and the tagging performance is evaluated in terms of QCD and top jet rejection for all studied TAR jet configurations.

C.1 $H \rightarrow b\bar{b}$

The tagger defined for studying the selection of the $H \rightarrow b\bar{b}$ candidate jets consists of a p_T dependent two-sided mass cut defined to provide an 80% signal jet efficiency throughout the full p_T spectrum. TAR jets are selected by requiring the presence of at least two double b -tagged (77% DLR1) TAR jets in an event with $p_T > 250$ GeV and $|\eta| < 2.0$. Track jets associated to the TAR jets are required to fulfil $p_T > 15$ GeV, $|\eta| < 2.5$ and consist of at least two constituents. Truth $H \rightarrow b\bar{b}$ TAR jets are built from truth small- R jets using the AKT algorithm with a size parameter of $R=1.6$ and are required to be associated with exactly two truth B -hadrons. The signal candidate jets are then selected as the TAR jets closest to the truth $H \rightarrow b\bar{b}$ TAR jets. Top and QCD jets are selected as p_T -leading TAR jets from $t\bar{t}$ and QCD di-jet samples. In order to define the tagger, the $m_\chi = 1, 2$ and 3 TeV mass resonances are used, which cover a broad p_T spectrum.

For each jet collection, a p_T -dependent 80% mass window cut is defined. The mass window is selected by iteratively finding the smallest window around the median that contains 80% of the signal jets. The upper and lower limits of the window are used to perform a logarithmic fit using $f(p_T) = p_0 + p_1 \cdot \log p_T + p_2 \cdot \log^2 p_T$ as fit function in order to find a functional form that allows to achieve a flat efficiency in p_T . Exemplary fits for AKT-FR ($R = 1.0$) and AKT-VR ($\rho = 250$ GeV) TAR jets used in the definition of the mass cut are shown in Figure C.1. The orange and green stars show the lower and upper cuts of the TAR jet mass window in

the individual p_T bins, respectively, while the blue stars show the mean of the TAR jet mass distributions in the different p_T bins. The dashed green and orange lines show the respective fits, which enclose the shaded grey area, that represents the allowed phase space of the tagger. The fits describe the mass window cuts fairly well and lead to a flat $H \rightarrow b\bar{b}$ jet efficiency within a 5-10% margin as shown in Figure C.2.




Figure C.1: Exemplary fits for AKT-FR ($R = 1.0$) and AKT-VR ($\rho = 250$ GeV) TAR jets used for the definition of the p_T -dependent two-sided 80% TAR jet mass cuts for identifying $H \rightarrow b\bar{b}$ jets.

Figure C.1 shows the effect of the mass cut on the TAR jet mass reconstructed using AKT-FR jets with $R = 0.50$, $R = 1.0$ and $R = 1.5$ for signal ($m_X = 2$ TeV), top and QCD jets. As expected, the tagger removes the very low and high mass tails from the $H \rightarrow b\bar{b}$ signal TAR jet mass peak, while keeping most of the jets in the core of the distribution. For $R = 1.5$ and $R = 1.0$, the top jet TAR jet mass shows a clear peak at $m_{\text{TAR}} \sim m_{\text{top}}$, while in the $R = 0.5$ case, the top decay products are not fully contained within the jet such that the peak is shifted towards $m_{\text{TAR}} \sim 75$ GeV. After applying the p_T dependent mass cut, most of the high m_{TAR} top jets are effectively removed and only jets in a $40 \text{ GeV} < m_{\text{TAR}} < 140 \text{ GeV}$ TAR jet mass window remain. Previous to the mass cut, the QCD jet mass distribution shows a clear peak around $m_{\text{TAR}} \sim 60$ GeV, which gets increasingly smeared out with larger jet sizes. With the application of the tagger, QCD jets in the core of the distribution as well as in the long high m_{TAR} tail are effectively suppressed.

In order to compare the performance of the various TAR jet configurations, the QCD and top jet rejections for the 80% mass cut are shown in Figure C.4 for AKT-FR and AKT-VR jets for all considered parameters.




Figure C.2: $H \rightarrow b\bar{b}$ signal jet efficiency after application of the 80% mass window cut for AKT-FR and AKT-VR TAR jets.

AKT-FR and AKT-VR TAR jets show a similar top jet rejection in the low- p_T region. However, as the top jet rejection increases with increasing jet p_T , AKT-FR TAR jets significantly outperform the AKT-VR jet collection in the high- p_T region. In both cases, larger size parameters seem to be beneficial for the rejection of top jets. In the case of QCD jets, the rejection remains fairly constant throughout the jet p_T spectrum and no significant differences are observed for the different jet configurations. No significant differences are observed between CA and AKT jet reclustering.

C.2 $H \rightarrow WW^*$

In the boosted 0ℓ channel, a dedicated neural network ("Four-prong tagger") is used to identify $W_{\text{had}}W_{\text{had}}$ jets originating from the decay of a Higgs boson by exploiting various jet substructure variables. In order to study the $H \rightarrow WW^*$ tagging performance of the various TAR jet definitions, a two-sided 80% mass window cut with an additional one-sided 80% cut on C_4 is employed to function as a proxy to the four-prong tagger. TAR jets are selected by requiring the presence of at least two TAR jets in an event with $p_T > 250$ GeV and $|\eta| < 2.0$. Track jets associated to the TAR jets are required to fulfill $p_T > 15$ GeV, $|\eta| < 2.5$ and consist of at least two constituents. Truth $H \rightarrow WW^*$ TAR jets are built from truth small- R jets using the AKT algorithm with a size parameter of $R=1.6$ and are required to be associated with exactly two truth W_{had} -bosons. Signal candidate jets are then selected as the TAR jets closest to

Appendix C TAR Jet Tagging Performance

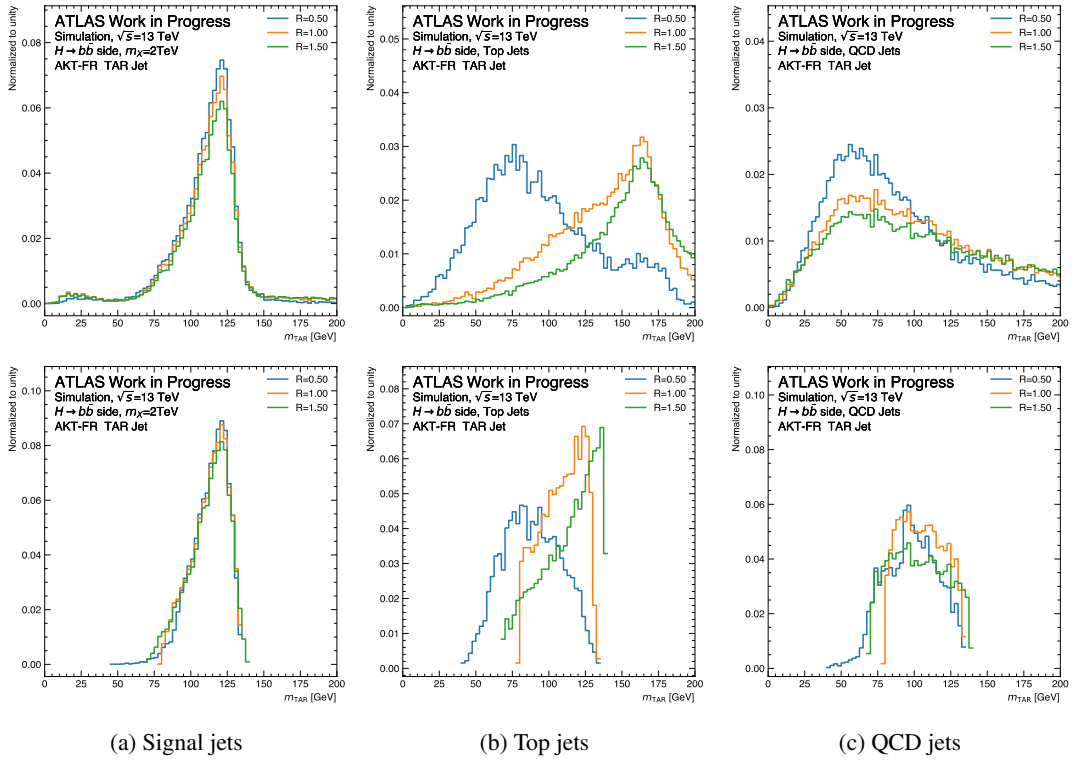
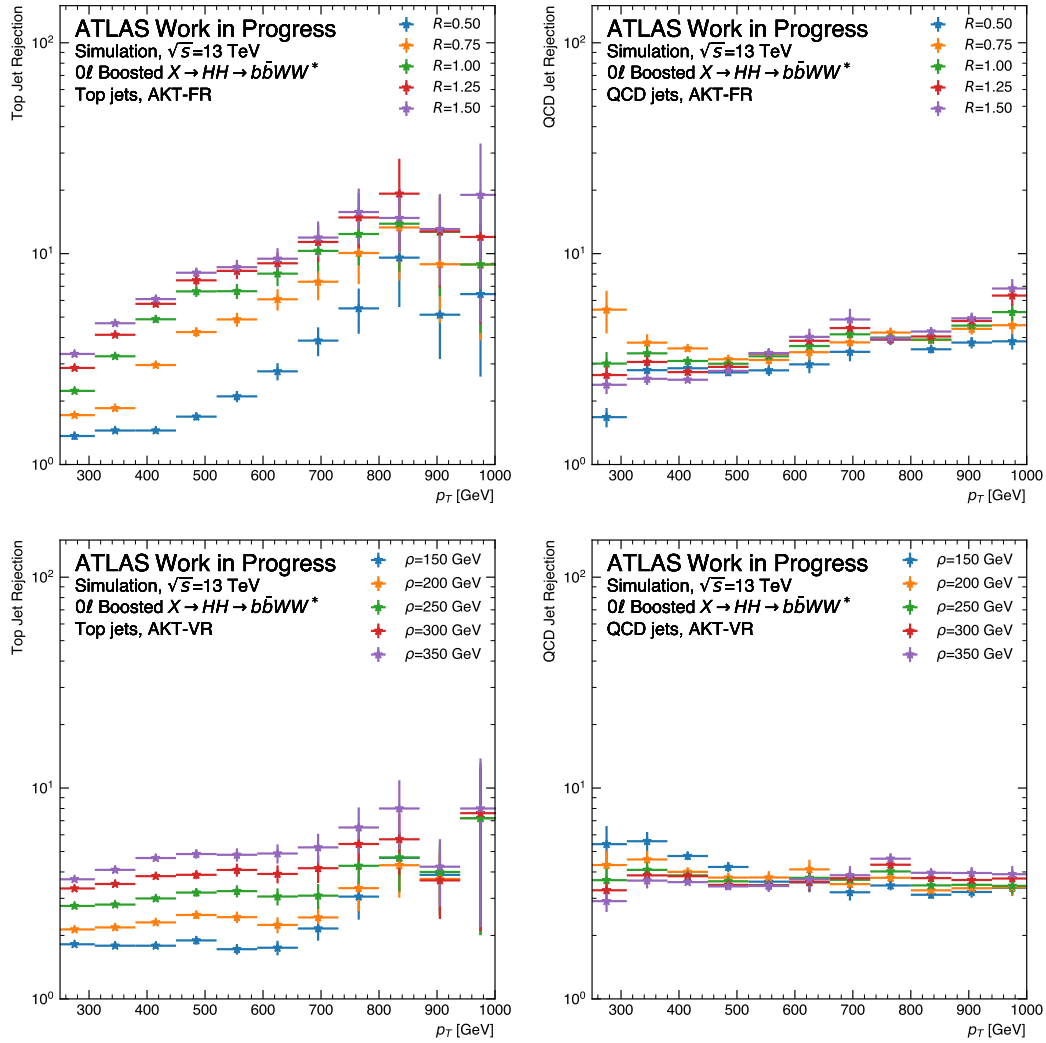


Figure C.3: TAR Jet mass of signal (2TeV), top and QCD jets without (top) and with the 80% mass cut applied (bottom) for AKT-FR jets with various size parameters.

the truth $H \rightarrow WW^*$ TAR jet. Top and QCD jets are selected as p_T -leading TAR jets from $t\bar{t}$ and QCD di-jet samples. In order to define the tagger, the $m_X = 1, 2$ and 3 TeV mass resonances are used, which cover a broad p_T spectrum.

Figure C.5 shows exemplary fits for AKT-FR ($R = 1.0$) TAR jets used for the definition of the p_T dependent taggers, consisting of a two-sided TAR jet mass cut, followed by a one-sided cut on C_4 for identifying $H \rightarrow WW^*$ jets. The fit fairly describes the functional form of both cuts for the studied jet parameters and as expected, yields a flat signal jet efficiency around 0.6 within a 10% (5%) margin for AKT-FR (AKT-VR) TAR jets (see Figure C.6). In the case of AKT-FR jets, the tagger exhibits a small hierarchy between the jet definitions in the high- p_T region, which needs to be taken into account when examining the background rejection of top and QCD jets.

Figure C.7 shows the effect of the tagger on the TAR jet mass distributions for signal, top and QCD jets reconstructed using AKT-FR TAR jets with $R = 0.5$, $R = 1.0$ and $R = 1.50$. As expected, the tagger fully removes jets in high and low mass tails, with the exception of $R = 0.5$ TAR jets, in which case the fit used to define the two-sided mass window cut yields negative



(a) Top Jets

(b) QCD Jets

Figure C.4: Top and QCD jet rejections for AKT-FR (top) and AKT-VR (bottom) TAR jets for the various studied parameters in the boosted 0ℓ channel.

Appendix C TAR Jet Tagging Performance

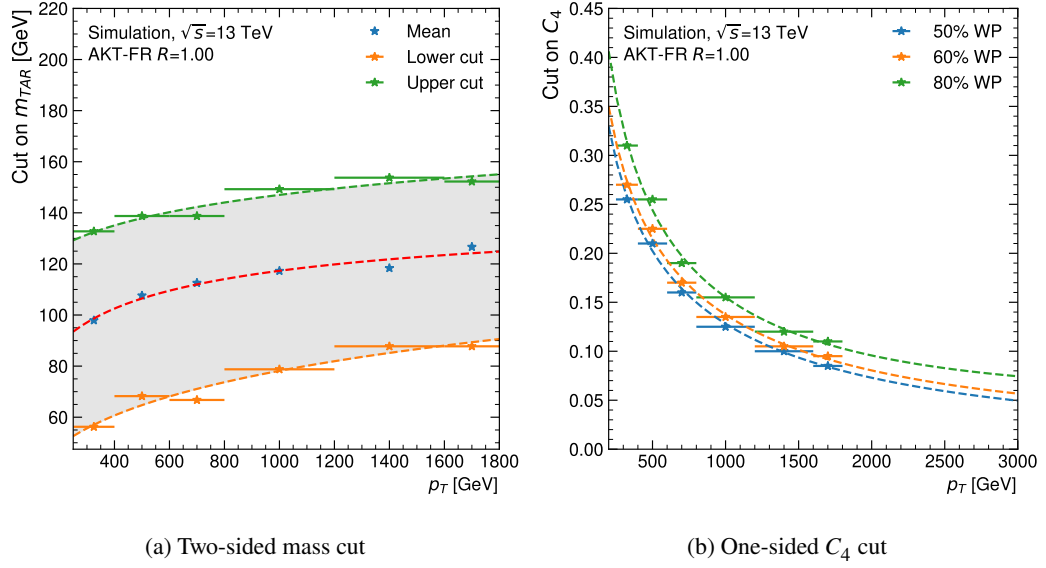


Figure C.5: Exemplary fits for AKT-FR $R = 1.0$ TAR jets used for the definition of the p_T dependent taggers, consisting of a two-sided TAR jet mass cut and a one-sided cut on C_4 for identifying $H \rightarrow WW^*$ jets.




Figure C.6: $H \rightarrow W_{had}W_{had}^*$ signal jet efficiency after application of the 80% mass window and the 80% cut on C_4 for AKT-FR and AKT-VR TAR jets.

lower cuts on the TAR jet mass in the low p_T region, such that no lower cut on the TAR jet mass is applied in these cases. The m_{TAR} distribution for top jets shows the expected peak at the top jet mass as well as a peak at $m_{\text{TAR}} \sim m_W$, resulting from cases in which the jet fails to collect the b -quark and its hadronization products. For $R = 0.5$ jets, the top quark and its decay products are not fully contained, such that the distribution falls quickly for $m_{\text{TAR}} > m_W$. For QCD jets, the m_{TAR} shape shows a peak at 25 GeV with a long-ranged tail towards larger TAR jet masses. Note that the differences between the QCD and top jet distributions in the case of $H \rightarrow b\bar{b}$ tagging (see Figure C.3) are only a result of the additional b -tagging requirement which is applied to the jets. A large fraction of QCD and top jets are effectively removed after applying the $H \rightarrow WW^*$ tagger, especially in the very low and very high m_{TAR} region.

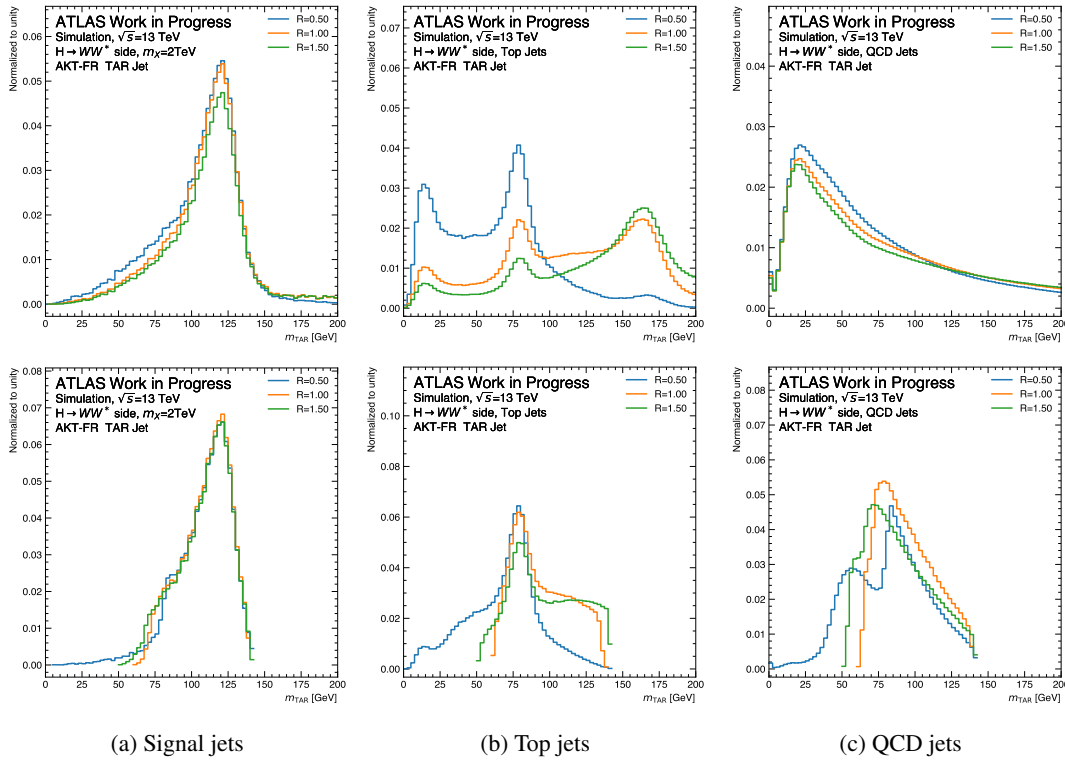
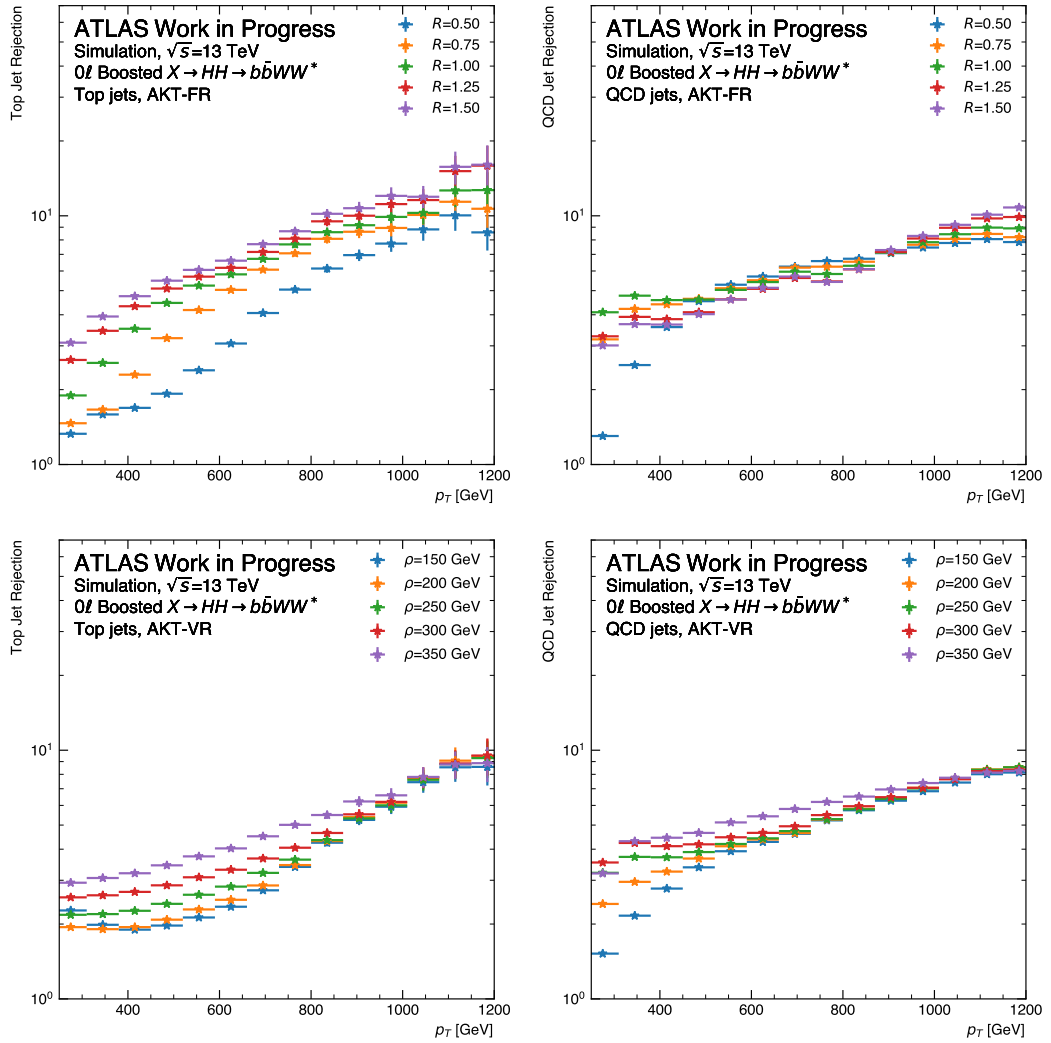


Figure C.7: TAR Jet mass of signal (2TeV), top and QCD jets without (top) and with the $H \rightarrow WW^*$ tagger applied (bottom) for AKT-FR TAR jets with various jet sizes.

Figure C.8 shows the top and QCD jet rejections of the tagger for AKT-FR and AKT-VR TAR jets for the various TAR jet parameters. AKT-FR TAR jets generally offer a better top jet rejection and a very similar QCD jet rejection. In both cases the rejection increases with increasing jet p_T . For top jets, AKT-FR and AKT-VR TAR jets show better rejection for larger jet sizes, while the differences are reduced in the high- p_T region. However, as the efficiencies



(a) Top Jets

(b) QCD Jets

Figure C.8: Top and QCD jet rejections for AKT-FR (top) and AKT-VR (bottom) TAR jets for the various TAR jet parameters in the boosted 0ℓ channel.

of the taggers show a slight hierarchy between the different jet configurations (see Figure C.6) for AKT-FR jets, differences in the true top jet rejection are likely smaller in this case. No significant differences in the QCD rejection between the various size parameters has been observed. Also, no significant differences have been observed between AKT and CA TAR jets.

C.3 Isolated $S \rightarrow W_{\text{had}}$ tagging

The tagger defined for studying the selection of isolated $S \rightarrow W_{\text{had}}$ candidate jets consists of a p_T dependent one-sided cut on $D_2^{\beta=1.5}$ defined to provide a 50% jet efficiency throughout the full p_T spectrum. TAR jets are selected by requiring the presence of at least three TAR jets in an event with $p_T > 200$ GeV, $|\eta| < 2.0$ and $50 \text{ GeV} < m_{\text{TAR}} < 105$ GeV. Track jets associated to the TAR jets are required to fulfill $p_T > 15$ GeV, $|\eta| < 2.5$ and consist of at least two constituents. Truth $S \rightarrow W_{\text{had}}$ TAR jets are built from truth small- R jets using the AKT algorithm with a size parameter of $R=1.6$ and are required to be associated with exactly one truth W_{had} -boson. Signal candidate jets are then selected as the TAR jets closest to the isolated truth $S \rightarrow W_{\text{had}}$ TAR jet. A $S \rightarrow W_{\text{had}}$ TAR jet is defined as being isolated if the angular distance to all other non truth-matched W_{had} jets is $\Delta R > 2.0$. Top and QCD jets are selected as p_T -leading TAR jets from $t\bar{t}$ and QCD di-jet samples. The tagging performance has been studied using three different SH samples with $m_X = 3$ TeV and $m_S \in [0.75 \text{ TeV}, 1.5 \text{ TeV}, 2.5 \text{ TeV}]$.

As the following studies only target isolated $S \rightarrow W_{\text{had}}$ jets, it is important to understand the number of isolated W_{had} candidates expected for the different mass points. Figure C.9 shows the event fraction observed with the specified number of isolated W_{had} candidates after applying the event selection specified in Section 11.2 for the different values of m_S . In the case of $m_S = 0.75$ TeV, up to 90% of the jets are not fully isolated. This fraction significantly increases for $m_S = 1.5$ TeV, where $\approx 80\%$ of the events contain at least one isolated W_{had} candidate. Finally, for $m_S = 2.5$ TeV, this number remains roughly constant, but most events only contain one isolated jet. While small differences are seen when varying the jet size of AKT-FR TAR jets, no significant discrepancies are found between the different energy scales studied for AKT-VR TAR jets.

TAR jets passing the mass cut are used to define a p_T -dependent tagger based on a lower cut on $D_2^{\beta=1.5}$. Exemplary fits for AKT-FR ($R = 1.0$) and AKT-VR ($\rho = 250$ GeV) TAR jets used for the definition of tagger are shown in Figure C.10. The resulting fits describe well the functional form of the cut on $D_2^{\beta=1.5}$ in the different p_T bins. The resulting $S \rightarrow W_{\text{had}}$ signal jet efficiencies after application of the 50% cut on $D_2^{\beta=1.5}$ for AKT-FR and AKT-VR TAR jets are shown in Figure C.11. As expected, the efficiency shows a flat spectrum in p_T around 0.5 within a 5% margin for all jet collections. The effect of the tagger on the $D_2^{(\beta=1.0)}$ distributions of signal, top and QCD jets reconstructed using AKT-FR TAR jets are shown in Figure C.12. The D_2 distribution of signal jets shows a peak around 0.8 with a long tail towards larger

Appendix C TAR Jet Tagging Performance

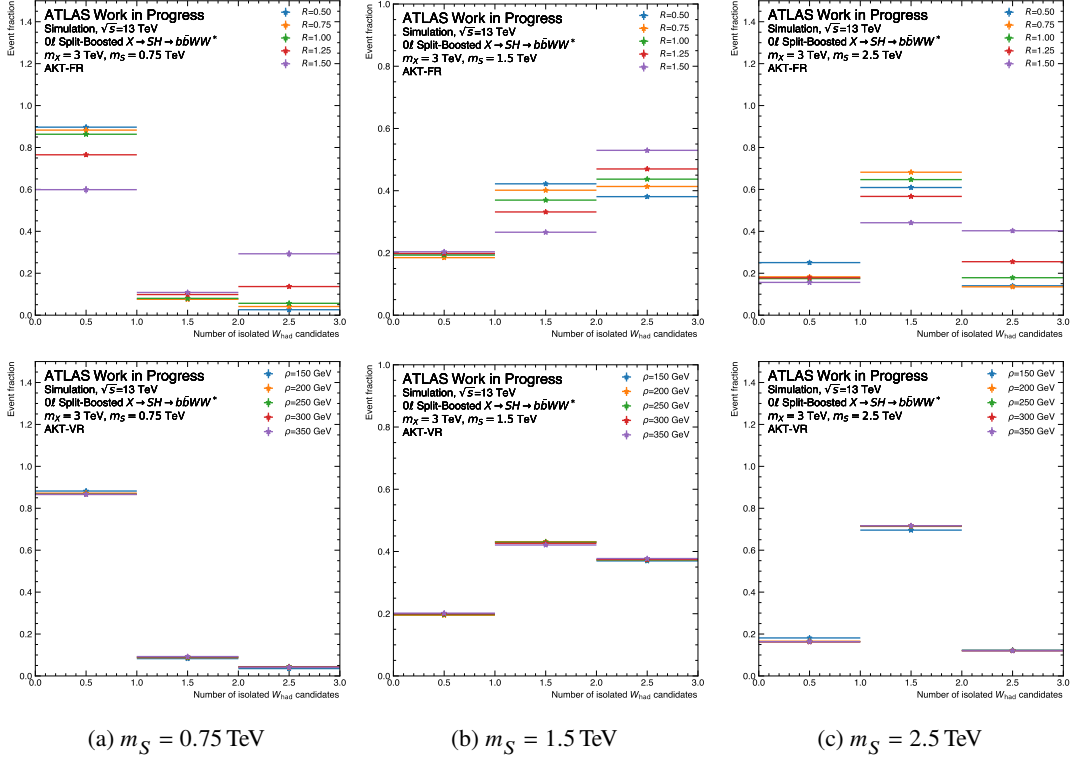


Figure C.9: Number of isolated W_{had} candidates for AKT-FR (top) and AKT-VR (bottom) TAR jets for different values of m_S and the different jet parameters.

values. Similarly, top and QCD jets show a peak at 1.3 and 1.5, respectively, with large tails towards larger values of D_2 . In all cases, the tagger effectively removes jets with large D_2 substructure values. The top and QCD jet rejections for AKT-FR and AKT-VR TAR jets with the various jet parameters are shown in Figure C.13. AKT-FR TAR jets show a slightly better top jet rejection and very similar QCD jet rejection. In all cases, the top and QCD jet rejections increase with larger jet p_T . No significant differences are observed between the various TAR jet size parameters and energy scales.




Figure C.10: Exemplary fits for AKT-FR ($R = 1.0$) and AKT-VR ($\rho = 250$ GeV) TAR jets used for the definition of the p_T -dependent one-sided $D_2^{\beta=1.5}$ cuts for identifying isolated $S \rightarrow W_{\text{had}}$ jets.




Figure C.11: $S \rightarrow W_{\text{had}}$ signal jet efficiency after application of the 50% cut on $D_2^{\beta=1.5}$ for AKT-FR and AKT-VR TAR jets.

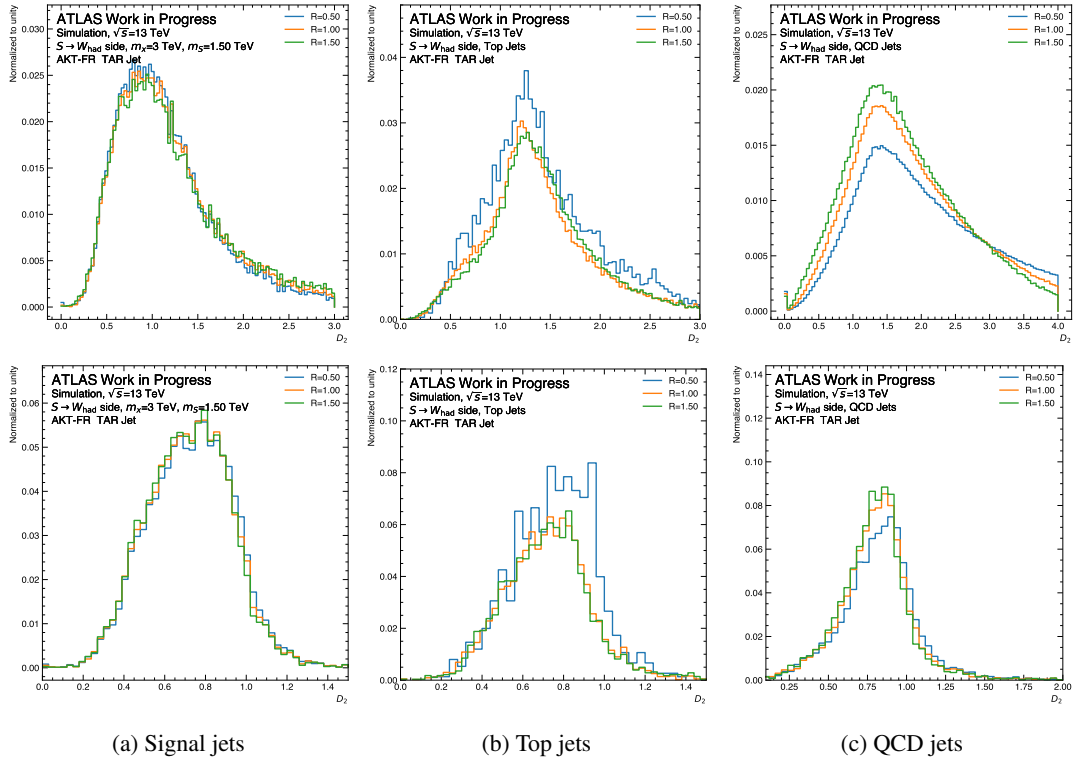
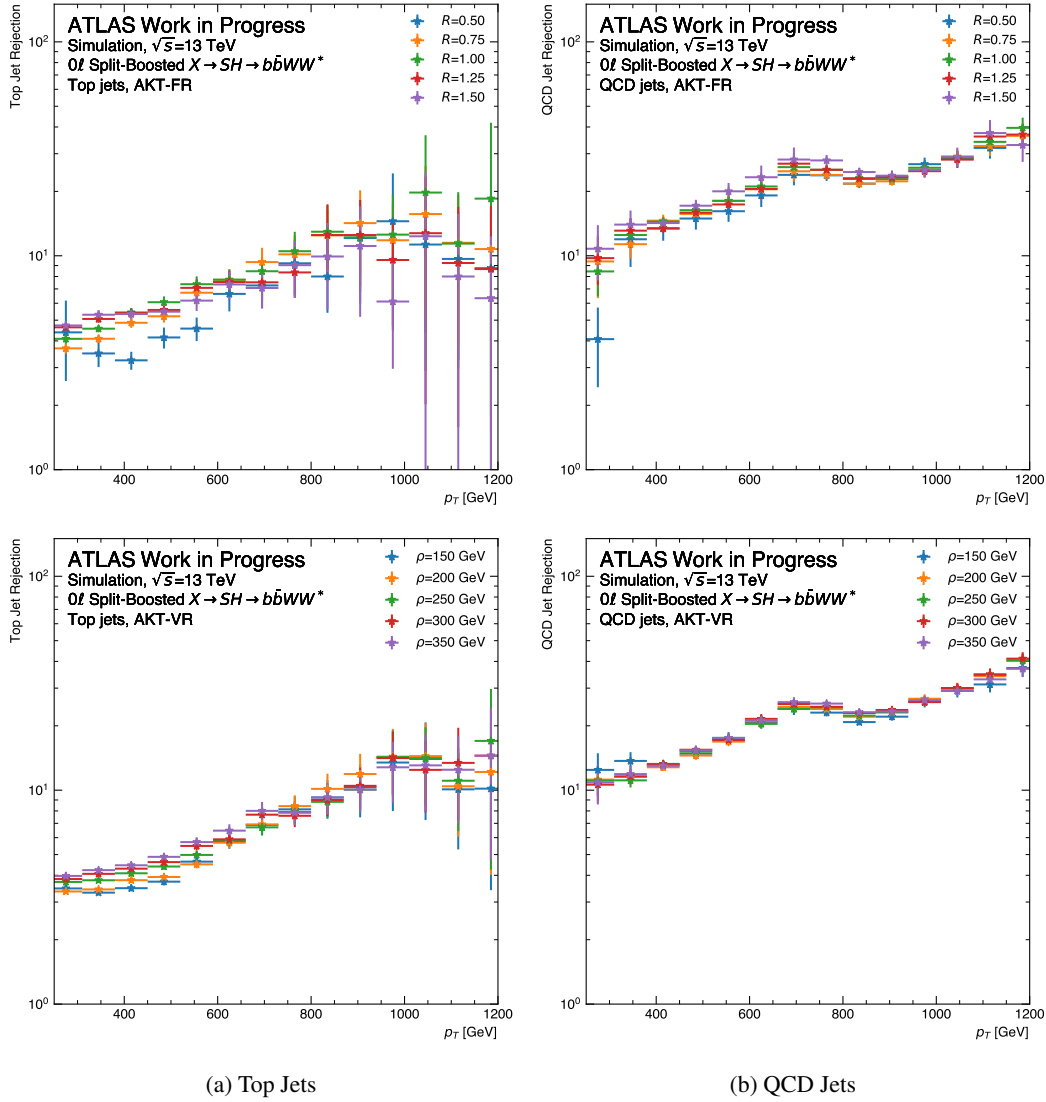


Figure C.12: TAR jet $D_2^{\beta=1}$ distributions of signal ($m_S = 1.5$ TeV), top and QCD jets without (top) and with the $S \rightarrow W_{\text{had}}$ tagger applied (bottom).



(a) Top Jets

(b) QCD Jets

Figure C.13: Top and QCD jet rejections for AKT-FR (top) and AKT-VR (bottom) TAR jets with the various jet configurations in the split-boosted 0ℓ channel.

Bibliography

- [1] J. J. Thomson, *Cathode rays*, *Phil. Mag. Ser. 5* **44** (1897) 293 (cit. on p. 5).
- [2] A. Einstein, *Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt*, *Ann. Phys.* **322** (1905) 132 (cit. on p. 5).
- [3] M. Born and P. Jordan, *Zur Quantenmechanik*, *Z. Phys.* **34** (1925) 858 (cit. on p. 5).
- [4] E. Rutherford,
Collision of α particles with light atoms. IV. An anomalous effect in nitrogen,
Phil. Mag. Ser. 6 **37** (1919) 581 (cit. on p. 5).
- [5] M. Born, W. Heisenberg and P. Jordan, *Zur Quantenmechanik. II.*,
Z. Phys. **35** (1926) 557 (cit. on p. 5).
- [6] P. A. M. Dirac and R. H. Fowler, *The quantum theory of the electron*,
Proc. R. Soc. Lond. A **117** (1928) 610 (cit. on pp. 5, 10).
- [7] C. D. Anderson, *The Positive Electron*, *Phys. Rev.* **43** (6 1933) 491 (cit. on p. 5).
- [8] E. Fermi, *An attempt of a theory of beta radiation. 1.*, *Z. Phys.* **88** (1934) 161
(cit. on p. 5).
- [9] H. A. Bethe, *The Electromagnetic Shift of Energy Levels*, *Phys. Rev.* **72** (4 1947) 339
(cit. on p. 6).
- [10] J. Schwinger, *On Quantum-Electrodynamics and the Magnetic Moment of the Electron*,
Phys. Rev. **73** (4 1948) 416 (cit. on p. 6).
- [11] J. Schwinger, *Quantum Electrodynamics. II. Vacuum Polarization and Self-Energy*,
Phys. Rev. **75** (4 1949) 651 (cit. on p. 6).
- [12] R. P. Feynman, *Space-Time Approach to Non-Relativistic Quantum Mechanics*,
Rev. Mod. Phys. **20** (2 1948) 367 (cit. on p. 6).
- [13] R. P. Feynman, *A Relativistic Cut-Off for Classical Electrodynamics*,
Phys. Rev. **74** (8 1948) 939 (cit. on p. 6).
- [14] R. P. Feynman, *Relativistic Cut-Off for Quantum Electrodynamics*,
Phys. Rev. **74** (10 1948) 1430 (cit. on p. 6).

- [15] S. Tomonaga,
On a Relativistically Invariant Formulation of the Quantum Theory of Wave Fields,
Prog. Theor. Phys. **1** (1946) 27 (cit. on p. 6).
- [16] Z. Koba, T. Tati and S.-i. Tomonaga,
*On a Relativistically Invariant Formulation of the Quantum Theory of Wave Fields. II:
Case of Interacting Electromagnetic and Electron Fields*,
Prog. Theor. Phys. **2** (1947) 101 (cit. on p. 6).
- [17] Z. Koba, T. Tati and S.-i. Tomonaga,
*On a Relativistically Invariant Formulation of the Quantum Theory of Wave Fields. III:
Case of Interacting Electromagnetic and Electron Fields*,
Prog. Theor. Phys. **2** (1947) 198 (cit. on p. 6).
- [18] S. Kanesawa and S. Tomonaga,
*On a Relativistically Invariant Formulation of the Quantum Theory of Wave Fields. V:
Case of Interacting Electromagnetic and Meson Fields*, Prog. Theor. Phys. **3** (1948) 1
(cit. on p. 6).
- [19] S. Kanesawa and S. Tomonaga,
*On a Relativistically Invariant Formulation of the Quantum Theory of Wave Fields V:
Case of Interacting Electromagnetic and Meson Fields*,
Prog. Theor. Phys. **3** (1948) 101 (cit. on p. 6).
- [20] Z. Koba and S. Tomonaga,
*On Radiation Reactions in Collision Processes. I: Application of the “Self-Consistent”
Subtraction Method to the Elastic Scattering of an Electron*,
Prog. Theor. Phys. **3** (1948) 290 (cit. on p. 6).
- [21] S. Tomonaga and J. R. Oppenheimer,
On Infinite Field Reactions in Quantum Field Theory, Phys. Rev. **74** (2 1948) 224
(cit. on p. 6).
- [22] F. J. Dyson, *The Radiation Theories of Tomonaga, Schwinger, and Feynman*,
Phys. Rev. **75** (3 1949) 486 (cit. on p. 6).
- [23] Izaak Neutelings, *History timeline and energy scales*, Licenced under the Creative
Commons Attribution-ShareAlike 4.0 International (CC BY-SA 4.0) licence,
URL: <https://tikz.net/timeline> (cit. on p. 6).
- [24] M. Gell-Mann, *A schematic model of baryons and mesons*, Phys. Lett. **8** (1964) 214
(cit. on p. 6).
- [25] E. D. Bloom et al., *High-Energy Inelastic $e - p$ Scattering at 6° and 10°* ,
Phys. Rev. Lett. **23** (16 1969) 930 (cit. on p. 6).
- [26] M. Breidenbach et al.,
Observed Behavior of Highly Inelastic Electron-Proton Scattering,
Phys. Rev. Lett. **23** (16 1969) 935 (cit. on p. 6).

-
- [27] C. N. Yang and R. L. Mills, *Conservation of Isotopic Spin and Isotopic Gauge Invariance*, Phys. Rev. **96** (1 1954) 191 (cit. on p. 6).
- [28] G. Hooft, *Renormalization of massless Yang-Mills fields*, Nucl. Phys. B **33** (1971) 173 (cit. on p. 7).
- [29] G. Hooft, *Renormalizable Lagrangians for massive Yang-Mills fields*, Nucl. Phys. B **35** (1971) 167 (cit. on p. 7).
- [30] S. L. Glashow, *Partial Symmetries of Weak Interactions*, Nucl. Phys. **22** (1961) (cit. on p. 7).
- [31] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, Phys. Rev. Lett. **13** (9 1964) 321 (cit. on pp. 7, 18).
- [32] G. S. Guralnik, C. R. Hagen and T. W. B. Kibble, *Global Conservation Laws and Massless Particles*, Phys. Rev. Lett. **13** (20 1964) 585 (cit. on pp. 7, 18).
- [33] P. W. Higgs, *Broken Symmetries and the Masses of Gauge Bosons*, Phys. Rev. Lett. **13** (16 1964) 508 (cit. on pp. 7, 18).
- [34] S. Weinberg, *A Model of Leptons*, Phys. Rev. Lett. **19** (21 1967) 1264 (cit. on p. 7).
- [35] A. Salam, *Weak and Electromagnetic Interactions*, Conf. Proc. C **680519** (1968) 367 (cit. on p. 7).
- [36] UA1 Collaboration, *Experimental Observation of Lepton Pairs of Invariant Mass Around $95 \text{ GeV}/c^2$ at the CERN SPS Collider*, Phys. Lett. **126B** (1983) (cit. on p. 7).
- [37] UA1 Collaboration, *Experimental Observation of Isolated Large Transverse Energy Electrons with Associated Missing Energy at $\sqrt{s} = 540 \text{ GeV}$* , Phys. Lett. **122B** (1983) (cit. on p. 7).
- [38] ATLAS Collaboration, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, Phys. Lett. B **716** (2012) (cit. on pp. 7, 23, 137).
- [39] CMS Collaboration, *Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC*, Phys. Lett. B **716** (2012) (cit. on pp. 7, 23, 137).
- [40] S. Hawking and R. Laflamme, *Baby universes and the non-renormalizability of gravity*, Phys. Lett. B **209** (1988) 39 (cit. on p. 7).
- [41] Muon g-2 Collaboration, *Measurement of the Positive Muon Anomalous Magnetic Moment to 0.46 ppm* , Phys. Rev. Lett. **126** (14 2021) (cit. on p. 7).
- [42] CDF Collaboration, *High-precision measurement of the W boson mass with the CDF II detector*, Science **376** (2022) (cit. on pp. 7, 93).

- [43] Particle Data Group, *Review of Particle Physics*, Prog. Theor. Exp. Phys **2022** (2022) (cit. on pp. 8, 9, 20–23, 53–57, 59).
- [44] W. Gordon, *Der Comptoneffekt nach der Schrödingerschen Theorie*, Z. Phys. **40** (1926) 117 (cit. on p. 10).
- [45] O. Klein, *Quantentheorie und fünfdimensionale Relativitätstheorie*, Z. Phys. **37** (1926) 895 (cit. on p. 10).
- [46] A. Proca, *Sur la theorie ondulatoire des electrons positifs et negatifs*, J. Phys. Radium **7** (1936) 347 (cit. on p. 11).
- [47] F. J. Dyson, *The S Matrix in Quantum Electrodynamics*, Phys. Rev. **75** (11 1949) 1736 (cit. on p. 11).
- [48] R. P. Feynman, *Space-Time Approach to Quantum Electrodynamics*, Phys. Rev. **76** (6 1949) 769 (cit. on p. 11).
- [49] G. 't Hooft and M. Veltman, *Regularization and renormalization of gauge fields*, Nucl. Phys. B **44** (1972) 189 (cit. on p. 11).
- [50] M. Planck, *Ueber irreversible Strahlungsvorgänge*, Ann. Phys. **306** (1900) 69 (cit. on pp. 12, 24).
- [51] E. Noether, *Invariante Variationsprobleme*, Nachr. Ges. Wiss. Gottingen, Math. Phys. Kl. **1918** (1918) 235 (cit. on p. 12).
- [52] M. H. Poincaré, *Sur la dynamique de l'électron*, Rend. Circ. Mat. Palermo **21** (1906) 129 (cit. on p. 12).
- [53] M. Kobayashi and T. Maskawa, *CP-Violation in the Renormalizable Theory of Weak Interaction*, Prog. Theor. Phys. **49** (1973) 652 (cit. on p. 14).
- [54] Z. Maki, M. Nakagawa and S. Sakata, *Remarks on the Unified Model of Elementary Particles*, Prog. Theor. Phys. **28** (1962) 870 (cit. on p. 14).
- [55] Super-Kamiokande Collaboration, *Evidence for Oscillation of Atmospheric Neutrinos*, Phys. Rev. Lett. **81** (8 1998) 1562 (cit. on p. 14).
- [56] J. H. Christenson, J. W. Cronin, V. L. Fitch and R. Turlay, *Evidence for the 2π Decay of the K_2^0 Meson*, Phys. Rev. Lett. **13** (4 1964) 138 (cit. on p. 14).
- [57] M. Gell-Mann, *Symmetries of Baryons and Mesons*, Phys. Rev. **125** (3 1962) 1067 (cit. on p. 16).
- [58] J. Goldstone, *Field theories with Superconductor solutions*, Il Nuovo Cimento **19** (1961) 154 (cit. on p. 19).

-
- [59] ATLAS Collaboration, *Measurements of the Higgs boson production and decay rates and coupling strengths using pp collision data at $\sqrt{s} = 7$ and 8 TeV in the ATLAS experiment*, Eur. Phys. J. C **76** (2016) 6 (cit. on p. 22).
- [60] ATLAS Collaboration, *Higgs boson production cross-section measurements and their EFT interpretation in the 4ℓ decay channel at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Eur. Phys. J. C **80** (2020) 957 (cit. on p. 22).
- [61] A. Collaboration, *Evidence for the spin-0 nature of the Higgs boson using ATLAS data*, Phys. Lett. B **726** (2013) 120 (cit. on p. 23).
- [62] ATLAS and CMS Collaboration, *Measurements of the Higgs boson production and decay rates and constraints on its couplings from a combined ATLAS and CMS analysis of the LHC pp collision data at $\sqrt{s} = 7$ and 8 TeV*, JHEP **08** (2016) 045 (cit. on p. 23).
- [63] ATLAS Collaboration, *Combined measurements of Higgs boson production and decay using up to 80 fb^{-1} of proton-proton collision data at $\sqrt{s} = 13$ TeV collected with the ATLAS experiment*, Phys. Rev. D **101** (2020) 012002 (cit. on p. 24).
- [64] A. Einstein, *Die Grundlage der allgemeinen Relativitätstheorie*, Ann. Phys. **354** (1916) 769 (cit. on p. 23).
- [65] K. G. Wilson and M. E. Fisher, *Critical Exponents in 3.99 Dimensions*, Phys. Rev. Lett. **28** (4 1972) 240 (cit. on p. 23).
- [66] F. J. Wegner and A. Houghton, *Renormalization Group Equation for Critical Phenomena*, Phys. Rev. A **8** (1 1973) 401 (cit. on p. 23).
- [67] K. G. Wilson and J. Kogut, *The renormalization group and the ϵ expansion*, Phys. Rep. **12** (1974) 75 (cit. on p. 23).
- [68] K. G. Wilson, *The renormalization group and critical phenomena*, Rev. Mod. Phys. **55** (3 1983) 583 (cit. on p. 23).
- [69] E. Corbelli and P. Salucci, *The extended rotation curve and the dark matter halo of M33*, Mon. Not. R. Astron. Soc. **311** (2000) 441 (cit. on p. 24).
- [70] *Planck 2018 results - I. Overview and the cosmological legacy of Planck*, Astron. Astrophys. **641** (2020) (cit. on p. 24).
- [71] A. N. Taylor, S. Dye, T. J. Broadhurst, N. Benítez and E. van Kampen, *Gravitational Lens Magnification and the Mass of Abell 1689*, Astrophys. J. **501** (1998) 539 (cit. on p. 24).

- [72] P. J. E. Peebles and B. Ratra, *The cosmological constant and dark energy*, Rev. Mod. Phys. **75** (2 2003) 559 (cit. on p. 24).
- [73] P. J. E. Peebles, *Tests of cosmological models constrained by inflation*, Astrophys. J. **284** (1984) 439 (cit. on p. 24).
- [74] S. M. Carroll, *The Cosmological Constant*, Living Rev. Relativ. **4** (2001) 1 (cit. on p. 24).
- [75] A. G. Riess et al., *Observational Evidence from Supernovae for an Accelerating Universe and a Cosmological Constant*, Astrophys. J. **116** (1998) 1009 (cit. on p. 24).
- [76] S. Perlmutter et al., *Measurements of Ω and Λ from 42 High-Redshift Supernovae*, Astrophys. J. **517** (1999) 565 (cit. on p. 24).
- [77] J. D. Barrow and M. S. Turner, *Baryosynthesis and the origin of galaxies*, Nature **291** (1981) 469 (cit. on p. 25).
- [78] A. D. Sakharov, *Violation of CP invariance, C asymmetry, and baryon asymmetry of the universe*, Phys. Uspekhi **34** (1991) 392 (cit. on p. 25).
- [79] P. Grangé, J. Mathiot, B. Mutet and E. Werner, *Aspects of fine-tuning of the Higgs mass within finite field theories*, Phys. Rev. D **88** (12 2013) 125015 (cit. on p. 25).
- [80] H. Georgi and S. L. Glashow, *Unity of All Elementary-Particle Forces*, Phys. Rev. Lett. **32** (8 1974) 438 (cit. on p. 25).
- [81] S. P. Martin, “A Supersymmetry Primer”, *Perspectives on Supersymmetry*, World Scientific, 2010 1 (cit. on p. 25).
- [82] P. van Nieuwenhuizen, *Supergravity*, Phys. Rep. **68** (1981) 189 (cit. on p. 26).
- [83] C. Csaki, *The Minimal Supersymmetric Standard Model*, Mod. Phys. Lett. A **11** (1996) 599 (cit. on pp. 26, 27).
- [84] G. Jungman, M. Kamionkowski and K. Griest, *Supersymmetric dark matter*, Phys. Rep. **267** (1996) 195 (cit. on p. 26).
- [85] LHC New Physics Working Group, *Simplified models for LHC new physics searches*, J. Phys. G **39** (2012) 105005 (cit. on p. 26).
- [86] J. Baglio et al., *The measurement of the Higgs self-coupling at the LHC: theoretical status*, JHEP **2013** (2013) 151 (cit. on p. 27).
- [87] R. Frederix et al., *Higgs pair production at the LHC with NLO and parton-shower effects*, Phys. Lett B **732** (2014) 142 (cit. on p. 27).

-
- [88] L. S. Ling et al., *NNLO QCD corrections to Higgs pair production via vector boson fusion at hadron colliders*, Phys. Rev. D **89** (7 2014) (cit. on p. 27).
- [89] F. A. Dreyer and A. Karlberg, *Fully differential vector-boson fusion Higgs pair production at next-to-next-to-leading order*, Phys. Rev. D **99** (7 2019) 074028 (cit. on p. 27).
- [90] F. A. Dreyer and A. Karlberg, *Vector-boson fusion Higgs pair production at N³LO*, Phys. Rev. D **98** (11 2018) 114016 (cit. on p. 27).
- [91] A. Azatov, R. Contino, G. Panico and M. Son, *Effective field theory analysis of double Higgs boson production via gluon fusion*, Phys. Rev. D **92** (3 2015) 035001 (cit. on p. 27).
- [92] P. Huang, A. Joglekar, M. Li and C. E. M. Wagner, *Corrections to di-Higgs boson production with light stops and modified Higgs couplings*, Phys. Rev. D **97** (7 2018) 075001 (cit. on p. 27).
- [93] G. Branco et al., *Theory and phenomenology of two-Higgs-doublet models*, Phys. Rep. **516** (2012) 1 (cit. on p. 27).
- [94] J. Mrazek et al., *The other natural two Higgs doublet model*, Nucl. Phys. B **853** (2011) 1 (cit. on p. 27).
- [95] R. Gröber and M. Mühlleitner, *Composite Higgs boson pair production at the LHC*, JHEP **2011** (2011) 20 (cit. on p. 27).
- [96] Z. Chacko, Y. Nomura, M. Papucci and G. Perez, *Natural little hierarchy from a partially goldstone twin Higgs*, JHEP **2006** (2006) 126 (cit. on p. 27).
- [97] L. Randall and R. Sundrum, *Large Mass Hierarchy from a Small Extra Dimension*, Phys. Rev. Lett. **83** (17 1999) 3370 (cit. on p. 27).
- [98] K. Cheung, *Phenomenology of the radion in the Randall-Sundrum scenario*, Phys. Rev. D **63** (5 2001) 056007 (cit. on p. 27).
- [99] D. O’Connell, M. J. Ramsey-Musolf and M. B. Wise, *Minimal extension of the standard model scalar sector*, Phys. Rev. D **75** (2007) 037701 (cit. on pp. 28, 29).
- [100] V. Barger, P. Langacker, M. McCaskey, M. Ramsey-Musolf and G. Shaughnessy, *Complex singlet extension of the standard model*, Phys. Rev. D **79** (1 2009) 015018 (cit. on pp. 28, 29).
- [101] T. Robens and T. Stefaniak, *LHC benchmark scenarios for the real Higgs singlet extension of the standard model*, Eur. Phys. J. C **76** (2016) 268 (cit. on p. 28).
- [102] G. Branco et al., *Theory and phenomenology of two-Higgs-doublet models*, Phys. Rep. **516** (2012) 1 (cit. on pp. 28–32).

- [103] D. Fontes et al., *The C2HDM revisited*, JHEP **02** (2018) 073 (cit. on pp. 28, 29).
- [104] C.-Y. Chen, M. Freid and M. Sher, *Next-to-minimal two Higgs doublet model*, Phys. Rev. D **89** (7 2014) 075009 (cit. on pp. 28, 29).
- [105] U. Ellwanger, C. Hugonie and A. M. Teixeira, *The Next-to-Minimal Supersymmetric Standard Model*, Phys. Rep. **496** (2010) 1 (cit. on p. 28).
- [106] K. J. Bae, H. Baer, V. Barger and D. Sengupta, *Revisiting the SUSY μ problem and its solutions in the LHC era*, Phys. Rev. D **99** (11 2019) 115027 (cit. on p. 29).
- [107] T. Robens, T. Stefaniak and J. Wittbrodt, *Two-real-scalar-singlet extension of the SM: LHC phenomenology and benchmark scenarios*, Eur. Phys. J. C **80** (2020) 151 (cit. on pp. 29, 32, 33).
- [108] S. L. Glashow, J. Iliopoulos and L. Maiani, *Weak Interactions with Lepton-Hadron Symmetry*, Phys. Rev. D **2** (7 1970) 1285 (cit. on p. 31).
- [109] M. Raidal et al., *Flavor physics of leptons and dipole moments*, Eur. Phys. J. C **57** (2008) 13 (cit. on p. 31).
- [110] M. Aoki, S. Kanemura, K. Tsumura and K. Yagyu, *Models of Yukawa interaction in the two Higgs doublet model, and their collider phenomenology*, Phys. Rev. D **80** (1 2009) 015017 (cit. on p. 31).
- [111] ATLAS Collaboration, *The ATLAS Experiment at the CERN Large Hadron Collider*, JINST **3** (2008) (cit. on pp. 35, 45, 49, 50, 186).
- [112] L. Evans and P. Bryant, *LHC Machine*, JINST **3** (2008) (cit. on p. 35).
- [113] CMS Collaboration, *The CMS Experiment at the CERN LHC*, JINST **3** (2008) S08004 (cit. on p. 35).
- [114] ALICE Collaboration, *The ALICE experiment at the CERN LHC*, JINST **3** (2008) S08002 (cit. on p. 35).
- [115] LHCb Collaboration, *The LHCb Detector at the LHC*, JINST **3** (2008) S08005 (cit. on p. 35).
- [116] E. Lopienska, *Complexe des accélérateurs du CERN en janvier 2022*, CERN-GRAPHICS-2022-001 (2022) (cit. on p. 36).
- [117] I. Zurbano Fernandez et al., *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*, CERN-2020-010 (2020) (cit. on p. 37).
- [118] I. Zurbano Fernandez et al., *The High Luminosity LHC Project*, 2022, URL: <https://hilumilhc.web.cern.ch/content/hl-lhc-project> (cit. on p. 38).

-
- [119] C. Bierlich et al., *A comprehensive guide to the physics and usage of PYTHIA 8.3*, SciPost Phys. Codebases 8 (2022) (cit. on p. 39).
- [120] J. C. Collins, D. E. Soper and G. F. Sterman, *Factorization of Hard Processes in QCD*, Adv. Ser. Direct. High Energy Phys. **5** (1989) 1 (cit. on p. 38).
- [121] G. Altarelli and G. Parisi, *Asymptotic Freedom in Parton Language*, Nucl. Phys. B **126** (1977) 298 (cit. on p. 38).
- [122] Y. L. Dokshitzer, *Calculation of the Structure Functions for Deep Inelastic Scattering and e^+e^- Annihilation by Perturbation Theory in Quantum Chromodynamics.*, Sov. Phys. JETP **46** (1977) 641 (cit. on p. 38).
- [123] V. Gribov and L. Lipatov, *Deep inelastic electron scattering in perturbation theory*, Physics Letters B **37** (1971) 78 (cit. on p. 38).
- [124] B. Andersson, G. Gustafson, G. Ingelman and T. Sjostrand, *Parton Fragmentation and String Dynamics*, Phys. Rept. **97** (1983) 31 (cit. on p. 39).
- [125] ATLAS Collaboration, *ATLAS experiment schematic illustration*, ATLAS-PHOTO-2022-055 (2022) (cit. on p. 40).
- [126] Izaak Neutelings, *CMS coordinate system*, Licenced under the Creative Commons Attribution-ShareAlike 4.0 International (CC BY-SA 4.0) licence, URL: https://tikz.net/axis3d_cms/ (cit. on p. 41).
- [127] ATLAS Collaboration, *ATLAS inner detector: Technical Design Report 1*, ATLAS-TDR-4 (1997) (cit. on p. 42).
- [128] ATLAS Collaboration, *ATLAS inner detector: Technical Design Report 2*, ATLAS-TDR-5 (1997) (cit. on p. 42).
- [129] ATLAS Collaboration, *Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2*, Eur. Phys. J. C **77** (2017) 673 (cit. on p. 42).
- [130] ATLAS Collaboration, *ATLAS Insertable B-Layer Technical Design Report*, ATLAS-TDR-19 (2010) (cit. on p. 43).
- [131] ATLAS Collaboration, *Technical Design Report for the ATLAS Inner Tracker Strip Detector*, ATLAS-TDR-025 (2017) (cit. on p. 44).
- [132] ATLAS Collaboration, *Technical Design Report for the ATLAS Inner Tracker Pixel Detector*, ATLAS-TDR-030 (2017) (cit. on p. 44).
- [133] ATLAS Collaboration, *Measurement of the photon identification efficiencies with the ATLAS detector using LHC Run 2 data collected in 2015 and 2016*, Eur. Phys. J. C **79** (2019) 205 (cit. on pp. 46, 128, 137, 138, 143, 174).

- [134] ATLAS Collaboration, *AtlFast3: the next generation of fast simulation in ATLAS*, *Comput. Softw. Big Sci.* **6** (2022) 7 (cit. on pp. 47, 65, 67, 68, 83, 85, 87, 88, 125–135, 232).
- [135] P. Gessinger et al., *The Acts project: track reconstruction software for HL-LHC and beyond*, *EPJ Web Conf.* **245** (2020) 10003 (cit. on p. 48).
- [136] ATLAS Collaboration, *ATLAS muon spectrometer: Technical Design Report*, ATLAS-TDR-10 (1997) (cit. on p. 48).
- [137] ATLAS Collaboration, *New Small Wheel Technical Design Report*, ATLAS-TDR-020 (2013) (cit. on p. 49).
- [138] ATLAS Collaboration, *Operation of the ATLAS trigger system in Run 2*, *JINST* **15** (2020) P10004 (cit. on p. 50).
- [139] G. C. Baldwin and G. S. Klaiber, *Photo-fission in heavy elements*, *Phys. Rev.* **71** (1947) 3 (cit. on p. 52).
- [140] G. Moliere, *Theorie der Streuung schneller geladener Teilchen I. Einzelstreuung am abgeschirmten Coulomb-Feld*, *Z. Naturforsch. A* **2** (1947) 133 (cit. on p. 54).
- [141] R. M. Sternheimer, *The density effect for the ionization loss in various materials*, *Phys. Rev.* **88** (1952) 851 (cit. on p. 54).
- [142] J. P. Biersack and J. F. Ziegler, *The Stopping and Range of Ions in Solids*, Springer, 1982 (cit. on p. 54).
- [143] J. Lindhard and M. Scharff, *Energy dissipation by ions in the keV region*, *Phys. Rev.* **124** (1961) 128 (cit. on p. 54).
- [144] H. H. Andersen and J. F. Ziegler, *Hydrogen Stopping powers and ranges in all elements*, (1977) (cit. on p. 55).
- [145] H. W. Bertini, *Low-energy intranuclear cascade calculation*, *Phys. Rev.* **131** (1963) 1801 (cit. on p. 55).
- [146] C. W. Fabjan and F. Gianotti, *Calorimetry for particle physics*, *Rev. Mod. Phys.* **75** (4 2003) 1243 (cit. on pp. 56, 59).
- [147] T. S. Virdee, “Experimental Techniques”, In: *1998 European School on High-Energy Physics* (St. Andrews, Scotland), 1998 (cit. on pp. 57, 60).
- [148] E. Longo and I. Sestili, *Monte Carlo calculation of photon-initiated electromagnetic showers in lead glass*, *Nucl. Instrum. Methods* **128** (1975) (cit. on p. 57).
- [149] Geant4 Collaboration, *GEANT4—a simulation toolkit*, *Nucl. Instrum. Meth. A* **506** (2003) (cit. on pp. 57, 67, 72, 99).

-
- [150] ATLAS Collaboration, *ATLAS Software and Computing HL-LHC Roadmap (2022)*, (), URL: <https://cds.cern.ch/record/2802918> (cit. on pp. 63, 65).
- [151] ATLAS Collaboration, *ATLAS HL-LHC Computing Conceptual Design Report*, (2020), URL: <http://cds.cern.ch/record/2729668> (cit. on p. 64).
- [152] ATLAS Collaboration, *The ATLAS Simulation Infrastructure*, Eur. Phys. J. C **70** (2010) (cit. on pp. 65, 72).
- [153] ATLAS Collaboration, *The simulation principle and performance of the ATLAS fast calorimeter simulation FastCaloSim*, ATL-PHYS-PUB-2010-013 (2010) (cit. on p. 65).
- [154] I. Goodfellow et al., “Generative Adversarial Nets”, In: *Adv. Neural Inf. Process. Syst.* (Montreal, Canada), 2014 (cit. on pp. 66, 67).
- [155] ATLAS Collaboration, *Fast simulation of the ATLAS calorimeter system with Generative Adversarial Networks*, ATL-SOFT-PUB-2020-006 (2020) (cit. on pp. 66, 67).
- [156] ATLAS Collaboration, *FATRAS: The ATLAS fast track simulation project*, J. Phys. Conf. Ser. **331** (2011) 032046 (cit. on pp. 66, 67).
- [157] ATLAS Collaboration, *Fast Track Reconstruction for HL-LHC*, ATL-PHYS-PUB-2019-041 (2019) (cit. on p. 66).
- [158] ATLAS Collaboration, *Emulating the impact of additional proton–proton interactions in the ATLAS simulation by presampling sets of inelastic Monte Carlo events*, Comput. Softw. Big Sci. **6** (2022) 3 (cit. on pp. 66, 127, 140).
- [159] ATLAS Collaboration, *The Fast Simulation Chain in the ATLAS experiment*, "ATL-SOFT-PROC-2021-017 (2021) (cit. on p. 66).
- [160] M. Arjovsky, S. Chintala and L. Bottou, “Wasserstein Generative Adversarial Networks”, In: *Proceedings of the 34th International Conference on Machine Learning* (Sydney, Australia), JMLR, 2017 (cit. on p. 67).
- [161] I. Gulrajani, F. Ahmed, M. Arjovsky, V. Dumoulin and A. C. Courville, “Improved training of wasserstein gans”, In: *Adv. Neural Inf. Process. Syst.* (Long Beach, California), 2017 (cit. on p. 67).
- [162] Ritsch, E., *ATLAS detector simulation in the integrated simulation framework applied to the w boson mass measurement*, Innsbruck, Univ., Diss. (2014) (cit. on pp. 68, 175).
- [163] ATLAS Collaboration, *Athena*, URL: <https://doi.org/10.5281/zenodo.4772550> (cit. on p. 68).
- [164] Geant4 Collaboration, *Recent developments in Geant4*, Nucl. Instrum. Methods Phys. Res. A **835** (2016) (cit. on pp. 72, 99).

- [165] B. Andersson, G. Gustafson and B. Nilsson-Almqvist, *A model for low- p_T hadronic reactions with generalizations to hadron-nucleus and nucleus-nucleus collisions*, Nucl. Phys. B **281** (1987) (cit. on p. 72).
- [166] B. Andersson, A. Tai and B.-H. Sa, *Final state interactions in the (nuclear) FRITIOF string interaction scenario*, Z. Phys. C **70** (1996) (cit. on p. 72).
- [167] B. Ganhuayag and V. Uzhinsky, *Modified FRITIOF code: Negative charged particle production in high energy nucleus nucleus interactions*, Czech. J. Phys. **47** (1997) (cit. on p. 72).
- [168] I. Jolliffe, *Principal Component Analysis*, *International Encyclopedia of Statistical Science*, 2011 (cit. on p. 74).
- [169] A. Rosenfeld and J. L. Pfaltz, *Sequential Operations in Digital Picture Processing*, J. ACM **13** (1966) 471 (cit. on p. 108).
- [170] M. Aaboud et al., *Performance of the ATLAS Track Reconstruction Algorithms in Dense Environments in LHC Run 2*, Eur. Phys. J. C **77** (2017) 673 (cit. on p. 108).
- [171] R. E. Kalman et al., *A new approach to linear filtering and prediction problems*, J. Basic Eng. **82** (1960) 35 (cit. on p. 108).
- [172] A. Salzburger, *The ATLAS Track Extrapolation Package*, ATL-SOFT-PUB-2007-005 (2007) (cit. on p. 109).
- [173] J. Humpherys, P. Redd and J. West, *A fresh look at the Kalman filter*, SIAM Rev. **54** (2012) 801 (cit. on p. 108).
- [174] C. Runge, *Über die numerische Auflösung von Differentialgleichungen*, Math. Ann. **46** (1895) 167 (cit. on p. 109).
- [175] W. Kutta, *Beitrag zur näherungsweise integration totaler differentialgleichungen*, Z. Math. Phys. **46** (1901) 435 (cit. on p. 109).
- [176] E. J. Nyström, *Über die praktische Auflösung von Integralgleichungen mit Anwendungen auf Randwertaufgaben*, Acta Math. **54** (1930) 185 (cit. on p. 109).
- [177] A. Strandlie and R. Frühwirth, *Track and vertex reconstruction: From classical to adaptive methods*, Rev. Mod. Phys. **82** (2010) 1419 (cit. on p. 109).
- [178] E. Lund, L. Bugge, I. Gavrilenko and A. Strandlie, *Transport of covariance matrices in the inhomogeneous magnetic field of the ATLAS experiment by the application of a semi-analytical method*, JINST **4** (2009) P04016 (cit. on p. 110).

-
- [179] E. Lund et al.,
Treatment of energy loss and multiple scattering in the context of track parameter and covariance matrix propagation in continuous material in the ATLAS experiment,
ATL-SOFT-PUB-2008-003 (2008) (cit. on p. 110).
- [180] T. Ježo, J. M. Lindert, N. Moretti and S. Pozzorini,
New NLOPS predictions for $t\bar{t} + b$ -jet production at the LHC,
Eur. Phys. J. C **78** (2018) 502 (cit. on p. 126).
- [181] T. Sjostrand, S. Mrenna and P. Z. Skands, *PYTHIA 6.4 Physics and Manual*,
JHEP **05** (2006) 026 (cit. on p. 126).
- [182] J. Gao et al., *CT10 next-to-next-to-leading order global analysis of QCD*,
Phys. Rev. D **89** (3 2014) 033009 (cit. on p. 126).
- [183] P. Z. Skands, *Tuning Monte Carlo generators: The Perugia tunes*,
Phys. Rev. D **82** (7 2010) 074018 (cit. on p. 126).
- [184] W. K. Tung,
New generation of parton distributions with uncertainties from global QCD analysis,
Acta Phys. Polon. B **33** (2002) 2933, ed. by J. Tran Thanh Van (cit. on p. 126).
- [185] T. Sjostrand, S. Mrenna and P. Z. Skands, *A Brief Introduction to PYTHIA 8.1*,
Comput. Phys. Commun. **178** (2008) 852 (cit. on pp. 126, 127).
- [186] ATLAS Collaboration, *Measurement of the Z/γ^* boson transverse momentum distribution in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*,
JHEP **09** (2014) 145 (cit. on p. 126).
- [187] P. Nason,
A New method for combining NLO QCD with shower Monte Carlo algorithms,
JHEP **11** (2004) 040 (cit. on p. 126).
- [188] S. Frixione, P. Nason and C. Oleari, *Matching NLO QCD computations with Parton Shower simulations: the POWHEG method*, JHEP **11** (2007) 070 (cit. on p. 126).
- [189] S. Alioli, P. Nason, C. Oleari and E. Re, *A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX*,
JHEP **06** (2010) 043 (cit. on p. 126).
- [190] K. Hamilton, P. Nason, E. Re and G. Zanderighi,
NNLOPS simulation of Higgs boson production, JHEP **10** (2013) 222 (cit. on p. 126).
- [191] K. Hamilton, P. Nason and G. Zanderighi,
Finite quark-mass effects in the NNLOPS POWHEG+MiNLO Higgs generator,
JHEP **05** (2015) 140 (cit. on p. 126).
- [192] J. Butterworth et al., *PDF4LHC recommendations for LHC Run II*,
J. Phys. G **43** (2016) 023001 (cit. on p. 126).

- [193] T. Sjöstrand et al., *An introduction to PYTHIA 8.2*, Comput. Phys. Commun. **191** (2015) 159 (cit. on pp. 126, 127).
- [194] D. J. Lange, *The EvtGen particle decay simulation package*, Nucl. Instrum. Meth. A **462** (2001) 152, ed. by S. Erhan, P. Schlein and Y. Rozen (cit. on p. 127).
- [195] R. D. Ball et al., *Parton distributions with LHC data*, Nucl. Phys. B **867** (2013) 244 (cit. on pp. 127, 143).
- [196] ATLAS Collaboration, *ATLAS Pythia 8 tunes to 7 TeV data*, ATL-PHYS-PUB-2014-021 (2014) (cit. on p. 127).
- [197] S. Carrazza, S. Forte and J. Rojo, “Parton Distributions and Event Generators”, In: *43rd International Symposium on Multiparticle Dynamics*, 2013 89 (cit. on p. 127).
- [198] ATLAS Collaboration, *The Pythia 8 A3 tune description of ATLAS minimum bias and inelastic measurements incorporating the Donnachie–Landshoff diffractive model*, ATL-PHYS-PUB-2016-017 (2016) (cit. on pp. 127, 143).
- [199] ATLAS Collaboration, *Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1*, Eur. Phys. J. C **77** (2017) 490 (cit. on pp. 127, 129, 130, 204, 206).
- [200] ATLAS Collaboration, *Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton–proton collision data*, JINST **14** (2019) P12006 (cit. on pp. 127, 143).
- [201] ATLAS Collaboration, *Electron reconstruction and identification in the ATLAS experiment using the 2015 and 2016 LHC proton–proton collision data at $\sqrt{s} = 13$ TeV*, Eur. Phys. J. C **79** (2019) 639 (cit. on pp. 128, 138).
- [202] ATLAS Collaboration, *Jet reconstruction and performance using particle flow with the ATLAS Detector*, Eur. Phys. J. C **77** (2017) 466 (cit. on p. 129).
- [203] M. Cacciari, G. P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, JHEP **04** (2008) 063 (cit. on pp. 129, 130, 205, 206).
- [204] ATLAS Collaboration, *Optimisation of large-radius jet reconstruction for the ATLAS detector in 13 TeV proton–proton collisions*, Eur. Phys. J. C **81** (2021) 334 (cit. on p. 130).
- [205] A. J. Larkoski, G. P. Salam and J. Thaler, *Energy Correlation Functions for Jet Substructure*, JHEP **06** (2013) 108 (cit. on pp. 131, 207).
- [206] J. Thaler and K. Van Tilburg, *Identifying Boosted Objects with N -subjettiness*, JHEP **03** (2011) 015 (cit. on pp. 131, 206).

-
- [207] ATLAS Collaboration, *Identification and energy calibration of hadronically decaying tau leptons with the ATLAS experiment in pp collisions at $\sqrt{s}=8$ TeV*, Eur. Phys. J. C **75** (2015) 303 (cit. on p. 132).
- [208] ATLAS Collaboration, *Reconstruction, Energy Calibration, and Identification of Hadronically Decaying Tau Leptons in the ATLAS Experiment for Run-2 of the LHC*, ATL-PHYS-PUB-2015-045(2015) (cit. on p. 132).
- [209] ATLAS Collaboration, *Reconstruction of hadronic decay products of tau leptons with the ATLAS experiment*, Eur. Phys. J. C **76** (2016) 295 (cit. on p. 132).
- [210] ATLAS Collaboration, *Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13$ TeV*, Eur. Phys. J. C **81** (2021) 578 (cit. on p. 135).
- [211] ATLAS Collaboration, *Determination of the parton distribution functions of the proton using diverse ATLAS data from pp collisions at $\sqrt{s} = 7, 8$ and 13 TeV*, Eur. Phys. J. C **82** (2022) (cit. on p. 137).
- [212] J. Saxon, *Discovery of the Higgs Boson, Measurements of its Production, and a Search for Higgs Boson Pair Production*, CERN-THESIS-2014-084 (2014) (cit. on p. 139).
- [213] ATLAS Collaboration, *Towards tuning electromagnetic shower properties to data with AtlFast3*, ATL-SOFT-PUB-2021-004 (2021) (cit. on pp. 139, 146, 147, 151–155, 232).
- [214] ATLAS Collaboration, *Performance of electron and photon triggers in ATLAS during LHC Run 2*, Eur. Phys. J. C **80** (2020) 47 (cit. on p. 140).
- [215] S. Park, L. Linsen, O. Kreylos, J. Owens and B. Hamann, *Discrete Sibson interpolation*, IEEE Trans. Vis. Comput. Graph. **12** (2006) 243 (cit. on p. 145).
- [216] J. Larson, M. Menickelly and S. M. Wild, *Derivative-free optimization methods*, Acta Numer. **28** (2019) 287 (cit. on p. 148).
- [217] X. Bouthillier et al., *Epistimio/orion: Asynchronous Distributed Hyperparameter Optimization*, version v0.1.17, 2021 (cit. on p. 149).
- [218] M. Rocklin, “Dask: Parallel computation with blocked algorithms and task scheduling”, In: *Proceedings of the 14th python in science conference*, vol. 130, SciPy Austin, TX, 2015 136 (cit. on p. 149).

- [219] D. Thain, T. Tannenbaum and M. Livny, *Distributed computing in practice: the Condor experience*, *Concurr. Comput. Pract. Exp.* **17** (2005) 323 (cit. on p. 149).
- [220] J. H. Friedman, *Greedy function approximation: a gradient boosting machine*, *Ann. Statist.* (2001) 1189 (cit. on p. 158).
- [221] A. Tornhill, “Assessing technical debt in automated tests with CodeScene”, In: *2018 IEEE International Conference on Software Testing, Verification and Validation Workshops (ICSTW)*, 2018 122 (cit. on p. 176).
- [222] T. J. McCabe, *A complexity measure*, *IEEE Trans. Softw. Eng.* (1976) (cit. on p. 178).
- [223] A. Hindle, M. W. Godfrey and R. C. Holt, “Reading beside the lines: Indentation as a proxy for complexity metric”, In: *16th IEEE International Conference on Program Comprehension*, 2008 133 (cit. on p. 178).
- [224] A. Zaborowska, *Geant4 fast and full simulation for Future Circular Collider studies*, *J. Phys. Conf. Ser.* **898** (2017) (cit. on pp. 181, 182).
- [225] J. Apostolakis et al., “Parallel geometries in Geant4: Foundation and recent enhancements”, In: *2008 IEEE Nuclear Science Symposium Conference Record*, 2008 883 (cit. on p. 182).
- [226] ATLAS Collaboration, *Minimum Bias Trigger Scintillators in ATLAS Run II*, *JINST* **9** (2014) (cit. on p. 185).
- [227] ATLAS Collaboration, *Search for Higgs boson pair production in the $b\bar{b}b\bar{b}$ final state from pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Eur. Phys. J. C* **75** (2015) 412 (cit. on p. 199).
- [228] ATLAS Collaboration, *Search for resonant pair production of Higgs bosons in the $b\bar{b}b\bar{b}$ final state using pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. D* **105** (2022) 092002 (cit. on p. 199).
- [229] ATLAS Collaboration, *Search for resonant and non-resonant Higgs boson pair production in the $b\bar{b}\tau^+\tau^-$ decay channel in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *Phys. Rev. Lett.* **121** (2018) 191801 (cit. on p. 199).
- [230] ATLAS Collaboration, *Search for Higgs boson pair production in the $b\bar{b}WW^*$ decay mode at $\sqrt{s} = 13$ TeV with the ATLAS detector*, *JHEP* **04** (2019) 092 (cit. on p. 199).
- [231] A. L. Read, *Presentation of search results: the CLs technique*, *J. Phys. G* **28** (2002) 2693 (cit. on p. 199).

-
- [232] ATLAS Collaboration, *Combination of searches for Higgs boson pairs in pp collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector*, Phys. Lett. B **800** (2020) 135103 (cit. on p. 199).
- [233] K. Abeling, *Search for resonant Higgs boson pair production in the $b\bar{b}WW^*$ decay channel in the boosted 1-lepton final state using the full Run 2 ATLAS dataset*, PhD thesis: Goettingen U. (2022) (cit. on pp. 201, 202, 212, 214, 225).
- [234] ATLAS Collaboration, *Track assisted techniques for jet substructure*, ATL-PHYS-PUB-2018-012 (2018) (cit. on pp. 202, 204, 208, 209).
- [235] Beirer, J.F., *Performance studies of advanced $H \rightarrow b\bar{b}$ tagging methods for searches for resonant Higgs boson pair production in the $b\bar{b}WW^*$ decay channel with the ATLAS detector*, II.Physik-UniGö-MSc-2019/01 (2019) (cit. on pp. 202, 209, 232).
- [236] B. C. Forland, *Inaugural ATLAS searches for resonant Di-Higgs and SH Signals in the boosted, fully-hadronic $b\bar{b}VV^*$ final state at ATLAS using $\sqrt{s} = 13$ TeV data and novel machine learning techniques*, PhD thesis: Indiana U., Bloomington (2021) (cit. on pp. 202, 212–215, 225).
- [237] ATLAS Collaboration, *Search for resonant boosted HH and SH production in the bbVV decay channel with 0 or 1 lepton in the final state using the full Run 2 ATLAS data*, ATL-COM-PHYS-2022-020 (2022) (cit. on pp. 202, 212, 214, 225).
- [238] S. Schätzel, *Boosted top quarks and jet structure*, Eur. Phys. J. C **75** (2015) 415 (cit. on p. 203).
- [239] G. P. Salam, *Towards Jetography*, Eur. Phys. J. C **67** (2010) 637 (cit. on p. 205).
- [240] S. Catani, Y. L. Dokshitzer, M. H. Seymour and B. R. Webber, *Longitudinally invariant K_t clustering algorithms for hadron hadron collisions*, Nucl. Phys. B **406** (1993) 187 (cit. on p. 205).
- [241] Y. L. Dokshitzer, G. D. Leder, S. Moretti and B. R. Webber, *Better jet clustering algorithms*, JHEP **08** (1997) 001 (cit. on p. 205).
- [242] D. Krohn, J. Thaler and L.-T. Wang, *Jets with Variable R*, JHEP **06** (2009) 059 (cit. on p. 206).
- [243] ATLAS Collaboration, *Boosted Higgs ($\rightarrow b\bar{b}$) Boson Identification with the ATLAS Detector at $\sqrt{s} = 13$ TeV*, ATLAS-CONF-2016-039 (2016) (cit. on p. 206).
- [244] ATLAS Collaboration, *Variable Radius, Exclusive- k_T , and Center-of-Mass Subject Reconstruction for Higgs ($\rightarrow b\bar{b}$) Tagging in ATLAS*, ATL-PHYS-PUB-2017-010 (2017) (cit. on p. 206).

- [245] ATLAS Collaboration, *Identification of high transverse momentum top quarks in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, JHEP **06** (2016) 093 (cit. on p. 207).
- [246] ATLAS Collaboration, *Identification of boosted, hadronically decaying W bosons and comparisons with ATLAS data taken at $\sqrt{s} = 8$ TeV*, Eur. Phys. J. C **76** (2016) 154 (cit. on p. 207).
- [247] CMS Collaboration, *Search for a massive resonance decaying into a Higgs boson and a W or Z boson in hadronic final states in proton-proton collisions at $\sqrt{s} = 8$ TeV*, JHEP **02** (2016) 145 (cit. on p. 207).
- [248] A. J. Larkoski, D. Neill and J. Thaler, *Jet Shapes with the Broadening Axis*, JHEP **04** (2014) 017 (cit. on p. 207).
- [249] M. Cacciari and G. P. Salam, *Pileup subtraction using jet areas*, Phys. Lett. B **659** (2008) 119 (cit. on p. 208).
- [250] ATLAS Collaboration, *Jet mass reconstruction with the ATLAS Detector in early Run 2 data*, ATLAS-CONF-2016-035 (2016) (cit. on p. 208).
- [251] D. Krohn, J. Thaler and L.-T. Wang, *Jet Trimming*, JHEP **02** (2010) 084 (cit. on p. 208).
- [252] ATLAS Collaboration, *Optimisation and performance studies of the ATLAS b-tagging algorithms for the 2017-18 LHC run*, ATL-PHYS-PUB-2017-013 (2017) (cit. on p. 213).
- [253] Neubert, C., *Investigation of shower shape variables in the ATLAS Fast Calorimeter Simulation*, II.Physik-UniGö-BSc-2021/03 (2021) (cit. on p. 231).
- [254] Bhalla, N.K., *Searches for resonant Higgs boson pair production in the $bbWW^*$ decay channel for the boosted single-lepton final state at ATLAS*, II.Physik-UniGö-MSc-2021/03 (2021) (cit. on p. 232).