

Search for Axions at the LHC

Dissertation

zur Erlangung des Grades
„Doktor der Naturwissenschaften“

JOHANNES GUTENBERG
UNIVERSITÄT MAINZ



Fachbereich Physik, Mathematik und Informatik

Friedemann Neuhaus

geboren in Marburg
Mainz, den 9.7.2022



1. Berichterstatter

Prof. Dr. Matthias Schott

2. Berichterstatter

[REDACTED]

Datum der mündlichen Prüfung 06.07.2022

Versicherung

Ich, Friedemann Neuhaus, versichere, dass ich die Arbeit selbstständig verfasst und keine anderen als die angegebenen Quellen und Hilfsmittel benutzt sowie Zitate kenntlich gemacht habe.

Mainz, den 9.7.2022

Unterschrift

Abstract

The Standard Model of Particle Physics (SM) is able to describe almost all measurements in particle physics with extraordinary precision. However, few phenomena within the greater scope of particle physics lack an explanation by the SM. Although there have been many searches for all kinds of new particles at the Large Hadron Collider (LHC), to the surprise of many physicists, no beyond Standard Model particle has been discovered to date. Most measurements focus on the search for heavy, centrally produced particles from interactions of two protons. As new particles do not have to be heavy but could be very light, these experiments cover only a fraction of the yet unexplored parameter space. In this thesis, two different approaches are presented to search for a group of light particles summarized under the term Axion-Like-Particles (ALPs).

The first approach utilizes the LHC as a photon collider to probe for the resonant production of an Axion-Like-Particle (ALP) in photon-photon scattering within the ATLAS detector. This measurement not only results in the first direct observation of the light-by-light scattering process, but also to the most stringent limits on the ALP to photon coupling in a mass range of $m_a = 6 - 100 \text{ GeV}$.

The second search is based on a new experiment at the LHC called the FASER experiment. The experiment was installed in March 2021. Data taking will start once the LHC resumes operations in 2022. In this thesis, the contributions to the design and commissioning of the FASER detector are presented. This includes the development of a dedicated calibration system for the calorimeter and commissioning of various sub-systems. In addition, a possible design for a future detector upgrade is evaluated. Its realization will improve the sensitivity for ALP to di-photon decay signatures.

Kurzzusammenfassung

Das Standardmodell der Teilchenphysik (SM) kann nahezu alle Messungen in der Teilchenphysik exakt berechnen und vorhersagen. Es gibt allerdings noch immer Beobachtungen, die sich nicht mit dem Modell beschreiben lassen und darauf hindeuten, dass das SM noch nicht vollständig ist. Zur Überraschung vieler Physiker konnten auch nach großen Bemühungen und der Untersuchung verschiedenster Modelle am Large Hadron Collider (LHC) bisher keine über das SM hinausgehenden Teilchen nachgewiesen werden. Die meisten Messungen konzentrieren sich dabei jedoch auf die Suche nach schweren, zentral produzierten Teilchen, die bei der Wechselwirkungen von zwei Protonen entstehen. In dieser Arbeit werden zwei verschiedene Ansätze für die Suche nach leichten Teilchen vorgestellt, die unter dem Begriff Axion-Like-Particles (ALPs) zusammengefasst werden.

Der erste Ansatz nutzt dazu den LHC als Photonenbeschleuniger und sucht nach der resonanten Produktion eines Axion-Like-Particle (ALP) in der Photon-Photon-Streuung mit Hilfe des ATLAS-Experiments. Diese Messung führt nicht nur zur ersten direkten Beobachtung der Streuung von Licht an Licht, sondern auch zu den strengsten Grenzwerten für die Kopplung von ALP mit Photonen in einem Massenbereich von $m_a = 6 - 100 \text{ GeV}$.

Die zweite Suche basiert auf einem neuen Experiment am LHC, dem FASER-Experiment. Das Experiment wurde im März 2021 installiert, allerdings beginnt die Datennahme erst, wenn der LHC im Jahr 2022 den Betrieb wieder aufnimmt. In dieser Arbeit werden die Beiträge zum Design und zur Inbetriebnahme des FASER-Detektors vorgestellt. Dazu gehören die Entwicklung eines speziellen Kalibrierungssystems für das Kalorimeter und die Inbetriebnahme verschiedener Subsysteme. Darüber hinaus wird ein Design für ein zukünftiges Detektor-Upgrade vorgestellt, welches die Messempfindlichkeit für Zerfälle von ALP in zwei Photonen verbessert.

Contents

1. Introduction	1
2. Theoretical background	3
2.1. The Standard Model	3
2.1.1. Elementary particles	3
2.1.2. Mathematical formulation of interactions	4
2.2. Physics beyond the Standard Model	9
2.3. The axion	10
2.4. Axion-Like Particles	11
3. Machine learning	13
3.1. Feedforward neural networks	13
3.1.1. The perceptron	14
3.1.2. Multilayer perceptron	15
3.2. Layer types	16
3.2.1. Dense	16
3.2.2. Dropout	16
3.2.3. Convolutional layers	17
3.2.4. Pooling layers	17
3.2.5. Flatten layers	18
3.2.6. Concatenate layers	18
3.3. Activation functions	18
3.4. Training of neural networks	19
3.4.1. Model architecture	19
3.4.2. Training data	20
3.4.3. Performance metric	21
3.4.4. Optimizer	21
3.4.5. Training procedure	21
3.5. Variable Importance	22
4. CERN & the LHC	25
4.1. The accelerator complex	25
4.2. Large Hadron Collider	25
4.3. Heavy ions and the LHC as a photon accelerator	27
4.4. Measurement of the luminosity	27
I. Light-by-light scattering	29
5. The ATLAS experiment	31
5.1. Inner detector and magnetic field	32
5.2. Calorimeters	34

5.3.	Muon spectrometer	36
5.4.	Luminosity detectors	37
5.5.	Reconstructing objects	37
5.5.1.	Charged particle tracks	37
5.5.2.	Electromagnetic particles	38
5.5.3.	Jet reconstruction	38
5.6.	Trigger system	39
5.7.	Event simulation	39
5.7.1.	Event generation	40
5.7.2.	Detector simulation	40
5.7.3.	Digitization	41
5.7.4.	Reconstruction	41
6.	Measurement of light-by-light scattering and ALP limits	43
6.1.	Measurement procedure	44
6.2.	Measured and simulated samples	45
6.3.	Event selection	46
6.3.1.	Trigger	46
6.3.2.	Object selection	48
6.4.	Object and particle reconstruction	49
6.4.1.	Track reconstruction	49
6.4.2.	Photon reconstruction	49
6.5.	Photon identification	50
6.5.1.	Shower shape variables	53
6.5.2.	Photon and background samples	53
6.5.3.	Custom cut-based identification	54
6.5.4.	Neural network photon identification	55
6.5.5.	Measurement of the photon identification efficiency	59
6.6.	Background estimation	62
6.6.1.	Electron-positron background	62
6.6.2.	Central exclusive production	64
6.6.3.	Fake photon background	66
6.7.	Kinematic distributions	68
6.8.	Systematic and statistical uncertainties	70
6.8.1.	Systematic uncertainties on the signal	70
6.8.2.	Systematic uncertainties on the background	72
6.8.3.	Statistical uncertainties	73
6.8.4.	Luminosity uncertainty	73
6.9.	Integrated fiducial cross-section	74
6.10.	Differential fiducial cross-sections	75
6.10.1.	Binning of distributions	75
6.10.2.	Bin-by-bin unfolding	75
6.10.3.	Methodology of Bayesian unfolding	77
6.10.4.	MC prior	78
6.10.5.	Fiducial and efficiency corrections and response matrices	78
6.10.6.	Closure studies	79
6.10.7.	Comparison of the two unfolding methods	83
6.10.8.	Results	85
6.11.	Limits on the production of axion-like-particles	86

6.12. Conclusion	93
II. FASER	95
7. Introduction	97
7.1. ALP production	98
7.2. ALP decay	99
7.3. Outline	100
8. Detector	101
8.1. Overview of the detector components	101
8.2. Upper support frame	102
8.3. Scintillator stations	102
8.3.1. Veto station	103
8.3.2. Timing station	104
8.3.3. Preshower station	105
8.4. Electromagnetic calorimeter	105
8.5. Tracker	107
8.6. Magnet system	107
8.7. Readout and reconstruction	108
9. Tracker hardware and tracker commissioning	109
9.1. SCT module	109
9.2. Module quality assurance procedure	111
9.3. Results from the module quality assurance	114
9.4. Plane and station commissioning	115
9.4.1. Commissioning setups	116
9.4.2. Testing procedure	117
9.4.3. Results	119
10. Calorimeter	121
10.1. Initial module testing	121
10.2. Photomultiplier and power supply	122
10.3. Optical filter	124
10.4. Mechanical assembly	125
11. Calorimeter calibration system	127
11.1. Requirements	127
11.2. LED driver circuit	128
11.2.1. The Kapustinsky circuit	128
11.2.2. Transistor + thyristor	129
11.2.3. First prototypes	131
11.3. Calibration board	133
11.3.1. Tuning the driver behaviour and LED selection	134
11.4. Fibre optics and mechanics	135
11.4.1. Calorimeter fibres	135
11.4.2. Scintillator fibres	137
11.5. Characterization and commissioning of calibration board	137
11.5.1. Measurement setup	137

11.5.2. Linearity of the bias voltage	138
11.5.3. Pulse amplitude scans	138
11.5.4. Frequency scans	139
11.5.5. Uniformity of the fibre bundles	140
11.5.6. Long-term stability	141
11.6. Test of the calibration system in a test beam	142
11.6.1. Calibration procedure	143
11.6.2. Analysis of the test beam	144
11.7. Conclusion	146
12. Preshower upgrade	147
12.1. Prototype design	148
12.2. Electron test beam	152
12.2.1. Test beam setup	152
12.2.2. Analysis procedure	152
12.2.3. Resolution and alignment	155
12.2.4. Results	157
12.3. Detector simulation	158
12.3.1. Material interaction	158
12.3.2. Hit digitization	160
12.3.3. Tuning of the simulation	161
12.3.4. Simulated samples	161
12.4. Neural network	164
12.4.1. Training setup	164
12.4.2. Selection of training data	165
12.4.3. Neural network architecture	165
12.4.4. Performance evaluation	167
12.4.5. Layer importance	169
12.4.6. Physics reach	169
12.5. Conclusion and further prospects	172
13. Summary	173
 III. Appendix	 175
A. Light-by-light scattering	179
B. FASER	185
B.1. Calibration system measurements	185
B.2. Lab control	186
Bibliography	191
Acronyms	203
List of Figures	205
List of Tables	212

1. Introduction

To date the Standard Model of Particle Physics (SM) is the best theory available describing fundamental particles and their interactions. However, from various observations it becomes obvious that the SM in its current formulation is not complete. This motivates a range of extensions which aim to build upon the SM. Most of those extensions are based on the existence of one or multiple new particles, of which have not yet been observed. This thesis focusses on a set of models, which all postulate a particle with similar properties which are summarized under the term Axion-Like-Particles (ALPs). All of these models share the fact that the ALP is a scalar or pseudo-scalar Boson with relatively low mass and couples to two photons.

The search for ALPs is conducted using the A Toroidal LHC ApparatuS (ATLAS) detector at the Large Hadron Collider (LHC) in photon-photon scattering events. The initial state photons stem from the strong magnetic field produced by the charged beam particles moving at relativistic velocities. The measurements are performed using runs with Lead ions rather than protons as the photon flux increases with Z^4 .

However, the established LHC experiments are not the only option to search for new physics. In the recent years, more small-scale experiments are proposed and built with the goal to quickly and relatively cheaply probe for a limited set of new models [1]. One of those new experiments is the FASER experiment. It is a detector located approximately 480 m away from the ATLAS Interaction Point (IP) on the line of sight of the collision axis. It is searching for light, weakly interacting particles like dark photons or with direct relevance to this thesis ALPs.

The thesis opens in chapter 2 with an introduction into the theoretical background of the relevant topics focussing on the SM and ALPs. As both analysis strategies involve the use of artificial neural networks, chapter 3 explains the basics of machine learning. The general introduction is concluded with a description of the European Organization for Nuclear Research (CERN) research laboratory, with attention on the overall accelerator complex and the LHC.

The main part of the thesis is divided into two parts. The first part deals with the measurement of light-by-light scattering and the search for ALPs. This part is opened with a description of the ATLAS experiment in chapter 5. It is followed by chapter 6 focussing on the measurement of light-by-light scattering.

The second part shifts the focus away from ATLAS towards the FASER experiment. It opens with an introduction of the experiment and the physics motivation in chapter 7. Chapter 8 gives a description of the FASER detector, including a summary of the contributions to the detector support structure. It is followed by the commissioning of the tracker and calorimeter in chapters 9 and 10. Chapter 11 outlines the development and testing of an optical system for calibrating the calorimeter energy response. Finally, a possible upgrade for the FASER detector is presented aiming to improve the detection capabilities for ALPs in chapter 12.

The closing chapter 13 gives a summary of the most important aspects presented in this thesis and comments on possible future studies based on the results presented.

Parts of this thesis have been published in [2], [3] and [4] where I was among the lead authors.

2. Theoretical background

2.1 The Standard Model

The SM is a quantum field theory which describes the elementary particles and the fundamental forces except for the gravitational force. It can be used to predict the interactions between the elementary particles which form all matter. While all predictions to this day meet the measured values, there are phenomena which cannot be explained by the SM as it stands right now. One large problem is that the SM cannot explain the matter-antimatter-asymmetry in the universe. Another open question is how dark matter can fit into the model. Therefore, a multitude of extensions to the SM are postulated which can explain one or more of those open questions. So far, there is no clear evidence for any of such models being correct, but it is rather an open field of research.

In this chapter, the SM as it is formulated as of today is explained along with the elementary particles known so far. This section is based on [5]. It is followed by possible extensions with the focus on Axion-Like-Particle (ALP) models.

2.1.1 Elementary particles

The SM includes only a limited number of elementary particles and forces. Elementary particles are the smallest building blocks and cannot be broken down any further. All other particles, such as protons or neutrons, are combinations of these elementary particles. The particles are grouped into multiple categories based on their fundamental properties. The main separation is done with regard of the spin and splits the particles into fermions with a spin of $1/2$ and bosons with a spin of 1 . The fermions thereby are the particles which are combined to form larger particles making up the visible matter in the universe. They interact with each other by exchanging bosons, which are also called force carriers. The SM includes all but one of the four fundamental forces with only the gravity missing in the model. However, the interaction strength of the gravitational force is much weaker than the one of the other three fundamental forces and has negligible influence on interactions at the scale of elementary particles. In Figure 2.1 an overview of all particles and force carriers is given.

Fermions

The fermions can be further grouped into six quarks and six leptons. In addition, each fermion has a corresponding anti-particle, which features the same properties as the particle, except for all charges being inverted. The fermions exist in three generations with each generation being made up from two quarks and two leptons as shown in Figure 2.1. The composition and properties of the generations are the same except for the mass of the particles which for each generation is larger with respect to the corresponding particle in the lower generation. Only the

fermions in the first generation are stable while the ones in the second and third generation have limited lifetimes and decay into particles from the lower generations. All fermions carry the charge of the weak interaction called weak isospin and therefore, couple together by means of the weak interaction, which is also responsible for the decay into other generations.

Quarks carry, in addition, a colour charge and an electric charge. As quarks and gluons are the only fundamental particles having colour charge, they are the only particles interacting via the strong interaction by the exchange of gluons. Regarding the electric charge there are three up-type quarks with a charge of $2/3$ being the up, charm and top quark and three down-type quarks with a charge of $-1/3$ being the down, strange and bottom quark.

The leptons are composed out of three pairs of one charged particle called electron, muon and tau and one corresponding neutral particle called neutrino (ν_e, ν_μ, ν_τ). The electron, muon and tau carry an electric charge of -1 and therefore, interact both electromagnetically and weakly. As stated, neutrinos are neutral and therefore, carry no electric charge. Hence, the only interaction neutrinos take part in is via the weak interaction. One speciality of neutrinos is their extremely small mass. It is so small that in the SM they are considered massless. As of today the masses of the neutrinos are still not exactly known, but only upper limits are found experimentally.

Bosons

In addition to the fermions, the SM includes five bosons mediating the fundamental forces. Gauge bosons have a spin of 1 and each boson is responsible for a corresponding force. The gluon (g) mediates colour charges and hence, couples only to particles carrying colour charge themselves. As gluons carry a colour charge as well, they couple not only to quarks but also among themselves. The photon (γ) does not carry any charge or mass and mediates the electromagnetic force. Hence, it only couples to charged particles, i.e. quarks, charged leptons and the $W^\pm$ boson.

The two bosons associated with the weak interaction are called the $W^\pm$ - and Z-boson. As the name $W^\pm$ suggests there are two “versions” of the W-boson where one carries positive electric charge and one carries negative electric charge. The W-boson always changes the flavour. The Z boson on the other hand is not charged. Both the W- and Z-boson couple to all fermions while the $W^\pm$ can only couple together different flavours which differ by one electric charge due to charge conservation. For leptons the W-boson only couples together the charged lepton with its corresponding neutrino: (e, ν_e) , (μ, ν_μ) and (τ, ν_τ) . Opposite to the gluon and photon the $W^\pm$ - and Z-boson are massive which explains the limited range of the interaction. They have a mass of $m_{W^\pm} = 80.379 \text{ GeV}$ and $m_Z = 91.188 \text{ GeV}$ respectively [7].

The last boson in the SM is the Higgs boson. It is a massive particle with a mass of $m_H = 125.10 \text{ GeV}$ [7] and has a spin of 0. Its existence is a consequence of the Higgs-field which gives massive fundamental particles their mass. Being measured for the first time in 2012 by both the ATLAS and the Compact Muon Solenoid (CMS) experiment [8, 9] it was the last missing particle in the current formulation of the SM to be discovered.

2.1.2

Mathematical formulation of interactions

In the previous section, the elementary particles in the SM were introduced. In this section, the interactions between them and their mathematical representation is explained in more detail. In the SM, each fundamental force is described by a Quantum Field Theory (QFT). The




Figure 2.1: Diagram showing the elementary particles within the Standard Model of particle physics. The particles are grouped by particle type. The grey boxes indicate to which force the fermions are coupling. The graphic is based on [6].

symmetries of the SM are described by the unitary product group of

$$SU(3)_C \times SU(2)_L \times U(1)_Y. \quad (2.1)$$

The indices C and Y indicate the conservation quantities of the symmetry groups while L is indicating that it only applies to left-handed particles. The fundamental particles are not modelled by discrete points but rather are described by excitations in the corresponding quantum field. These quantum fields can interact and thereby implement the interaction between particles. Hereby fermions are described as Dirac-spinors $\Psi(x)$ while bosons are represented by vector fields $A_\mu(x)$. The Lagrangian density $\mathcal{L}$ determines the dynamics of the quantum state and fundamental fields. Using the *principle of least action* the Euler-Lagrange equations describing the motion of particles can be obtained from the Lagrangian. The Lagrangian itself holds the symmetries given in Equation 2.1 and hence, is invariant under some types of field variations. Following the Noether-theorem each of the symmetries gives rise to conserved quantities respectively the various charges in the SM. Hereby, the $SU(3)_C$ group represents the strong interaction as described by Quantum Chromodynamics (QCD). The accompanying conserved quantity is the colour charge C. It is invariant under $SU(2)_L \times U(1)_Y$ which is the unification of the electromagnetic and weak force. The gauge fields give rise to the gauge bosons mediating the force in interactions. Gauge bosons are required to be massless, however, the W and Z bosons are observed to be massive. In the SM, they are considered generated bosons which receive their mass by the Higgs mechanism which implies the electroweak symmetry breaking.

Electroweak interactions

While originally being formulated independently, the SM includes a unification of the electromagnetic and weak interaction into the electroweak interaction. It is an $SU(2)_L \times U(1)_Y$ group. One feature of the weak force is that it allows parity violation in the charged weak force and hence, the coupling depends on the handedness of particles. In particular the charged weak force is only coupling to left-handed fermions and right-handed anti-fermions. To differentiate the handedness, also called chirality, the spinors are extended by the so-called chirality operators:

$$\Psi_L = \frac{1}{2}(1 - \gamma_5)\Psi \quad \Psi'_L = \frac{1}{2}(1 + \gamma_5)\Psi \quad (2.2)$$

$$\Psi_R = \frac{1}{2}(1 + \gamma_5)\Psi \quad \Psi'_R = \frac{1}{2}(1 - \gamma_5)\Psi, \quad (2.3)$$

where Ψ is the Dirac spinor of a particle and Ψ' the spinor of an anti-particle. It is then $\Psi = \Psi_L + \Psi_R$.

In the $SU(2)$ the left-handed fermions are expressed as doublets of the neutrino and the corresponding charged lepton, e.g.

$$\varphi_L = \begin{pmatrix} \nu_e \\ e \end{pmatrix}. \quad (2.4)$$

For left-handed quarks, the weak isospin doublets are constructed from the weak eigenstates which accounts for the mixing in the Cabibbo-Kobayashi-Maskawa (CKM) matrix. The quark doublets are therefore written as:

$$\varphi_L = \begin{pmatrix} u \\ d' \end{pmatrix}, \quad (2.5)$$

where d' is the flavour mixed quark as described by:

$$\begin{pmatrix} d' \\ s' \\ b' \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} d \\ s \\ b \end{pmatrix}. \quad (2.6)$$

The left- and right-handed fermion fields are differentiated by their weak isospin I_W . For the doublet (ν_e, e^-) the total weak isospin is $1/2$ with the third component being $I_W^{(3)}(\nu_e) = +\frac{1}{2}$ and $I_W^{(3)}(e^-) = -\frac{1}{2}$. As the W bosons carry a weak isospin of 1, they couple to these weak isospin doublets. To express that right-handed particles and left-handed anti-particles do not interact with the W they are placed in weak isospin singlets with a weak isospin of $I_W = I_W^3 = 0$.

The gauge bosons of the electroweak group are the W_i with $i = 1, 2, 3$ and the B . The photon γ and the Z -boson are expressed by

$$\begin{pmatrix} \gamma \\ Z \end{pmatrix} = \begin{pmatrix} \cos \theta_W & \sin \theta_W \\ -\sin \theta_W & \cos \theta_W \end{pmatrix} \begin{pmatrix} B \\ W_3 \end{pmatrix} \quad (2.7)$$

with θ_W being the weak mixing angle which is also called Weinberg angle. The W_1 and W_2 on the other hand combine to produce the charged massive $W^\pm$ -bosons:

$$W^\pm = \frac{1}{\sqrt{2}}(W_1 \mp W_2). \quad (2.8)$$

The weak mixing angle also gives a relationship between the masses of the W - and Z boson:

$$m_Z = \frac{m_W}{\cos \theta_W}. \quad (2.9)$$

Strong interaction

The strong interaction is formulated by Quantum Chromodynamics (QCD) and describes the interaction between particles which carry colour charges which are only quarks and gluons. It is based on an $SU(3)$ gauge symmetry resulting in eight massless gauge bosons called gluons. Each gluon carries a combination of a colour and anti-colour charge, which with three colours explains the necessity for introducing eight gluons. Due to the gluons only carrying two colour charges, each gluon has a net colour charge which in turn enables the gluons to interact with each other.

This self-interaction of the quarks is believed to give rise to the observation that quarks or gluons are never found as free particles, but always in a coupled state of two or more quarks. When trying to separate two quarks from each other, the self-interaction of the gluons exchanged between the quarks leads to the colour field being squeezed in a thin tube between the quarks. At large distances the energy density in this field is constant and therefore an infinite amount of energy would be required to separate two free quarks to infinity. As a consequence, coloured objects always form colourless objects with no colour field in between. This phenomenon is called colour confinement or just confinement.

Depending on the number of quarks in a particle it is called either a meson for particles consisting out of a quark-anti-quark pair ($q\bar{q}$) or baryon for a combination of three quarks (qqq).

Since gluons which mediate the strong force carry colour charge as well, they cannot travel

macroscopic distances and consequently, the strong force only has a short effective range of $10 \cdot 10^{-15}$ m which is less than the radius of a proton.

Another process related to confinement is the so-called hadronization. While the SM allows decays leading to free quarks or gluons these free colour charged particles instantly form hadrons by combining with other quarks e.g. created due to vacuum fluctuations.

Feynman diagrams

A common way of describing the particle interactions is using Feynman diagrams. Those diagrams were introduced by Richard Feynman back in 1949 [10]. The diagrams can be translated into the required math to calculate the corresponding interaction cross section by following a few simple rules. The diagrams have a time and a spatial axis which are rarely shown in the diagram. The most common convention is to have the time on the horizontal and the spatial dimension on the vertical. An example for a Feynman diagram showing the electron-positron annihilation and pair-production is given in Figure 2.2. As can be seen there, particles are travelling forward in time as denoted by the small arrow and anti-particles travel backward in time. In addition, some basic vertices are shown exemplary in Figure 2.3. For more details on Feynman diagrams and rules refer to [5].




Figure 2.2: Feynman diagram for the electron-positron annihilation and subsequent pair-production. The time axis is oriented horizontally. The initial state is on the left and the final state on the right.




Figure 2.3: Diagrams showing the vertices described by the Standard Model of particle physics. The graphic is based on [5].

2.2 Physics beyond the Standard Model

While the Standard Model can explain most of the results observed in particle physics there are some phenomena and experimental results which cannot be explained by it. One obvious phenomenon is the gravitational force which currently is not included in the SM at all. To include gravity one would need to describe gravity by a QFT which so far was not achieved. The addition of gravity to the SM, thus forming a theoretical model being able to describe everything, is also called Grand Unified Theory (GUT). While different models were postulated, so far no strong evidence has been found for any of the models.

Another observation not explainable with the SM to day is the asymmetry between matter and anti-matter. If, as expected, in the early beginning of the universe matter and anti-matter existed in equal parts, the two should have annihilated nearly completely. In the SM this holds true even when taking into account *CP* violation. However, the observable universe today consists mostly of matter.

Observations in cosmology also gave rise to another open question. It was observed already back in 1933 by F. Zwicky that the velocity dispersion in the Coma galaxy cluster is so high that with the mass density observed the system would be unstable and hence, some kind of invisible matter must exist in the galaxy cluster [11]. Even though other researchers found similar behaviour it was not before the 1970s when more and more examples for this phenomenon were observed that this missing matter was considered problematic [12]. In addition to the hints for the existence of dark matter from observing the rotational velocity of galaxies similar effects are seen in gravitational lensing where the lensing effect is stronger than expected from the visible matter. With those two observations alone, one could also come to the conclusion that dark matter may not be the only explanation. Instead, the problem might arise from errors in general relativity and that general relativity requires modifications to explain the phenomenon. However, the observation of the so-called Bullet Cluster [13] which was created in a collision between two galaxies is favouring the existence of dark matter. This is due to the reconstructed centre of mass being offset from the centre of mass of the observed baryonic mass by 8σ [14]. While this can be easily explained by the vast majority of dark matter models it cannot be explained by adjustments to the general relativity. Today the amount of regular matter is expected to account only for 5 % of the energy in the universe. Another 27 % is accounted for by dark matter. The remaining 68 % are made up by what is called dark energy [15]. Dark energy is needed to account for the accelerated expansion of the universe.

This, however, does not yet explain what dark matter is actually made out of. The fact that dark matter is invisible implies that it is not interacting via the strong or electromagnetic force, with only the weak interaction remaining in the SM. Neutrinos are already ruled out as candidates for dark matter by an experiment showing that the measured number of neutrinos is much lower than what is required to explain dark matter. Henceforth, it must be made up from a yet unknown kind of particle outside the SM.

An unpleasant aspect of the SM is found in the strong interaction and is called the strong *CP* problem. Naively, one would expect an additional parameter θ in the QCD-Lagrangian giving rise to *CP* violation in the strong force [16, 17]:

$$\mathcal{L}_\theta = \theta \frac{g_3^2}{32\pi^2} G_{\mu\nu}^a G_a^{\mu\nu}, \quad (2.10)$$

where θ is an angular parameter with $\theta \in [0, 2\pi]$. However, in case there were *CP* violation in QCD, the electric dipole moment of the neutron would be in the magnitude of $10 \cdot 10^{-18}$ e m. Ex-

perimentally though an upper limit of $3.0 \cdot 10^{-26}$ e cm is found [18] and no hints on CP violation were found in other measurements. While θ in principle could take any value within $[0, 2\pi]$ the measured upper limit constrains the value of θ to $\theta < 10^{-9}$. While this is not a problem in itself, it would be favourable to have an explanation to why this parameter must be so small.

A vast amount of possible extensions to the SM exist tackling one or multiple of the open questions listed above. As this thesis is focussed on the search for ALPs, the following section will present the properties of ALPs as well as how it could solve the strong CP problem.

2.3 The axion

As mentioned before CP is not broken in QCD without a conclusive argument for that behaviour. In 1977 Peccei and Quinn proposed a theory which extends the QCD Lagrangian with an additional global chiral $U(1)$ symmetry, now often denoted as $U(1)_{PQ}$. This dynamic field hereby cancels out the CP violating term instead of requiring θ to be a constant close to 0 [19]. The QCD Lagrangian would then be:

$$\mathcal{L} = \mathcal{L}_{SM} + \theta \frac{g_3^2}{32\pi^2} G_{\mu\nu}^a \tilde{G}_a^{\mu\nu} - \frac{1}{2} \partial_\mu a \partial^\mu a + \mathcal{L}_{int} + \zeta \frac{a}{f_a} \frac{g_3^2}{32\pi^2} G_{\mu\nu}^a \tilde{G}_a^{\mu\nu}, \quad (2.11)$$

with a being the axion field, $\mathcal{L}_{int}$ being the interaction term and f_a the order parameter of the symmetrie breaking. The last term represents an effective potential for the axion field with a minimum at $\langle a \rangle = -f_a / \zeta \theta$. At this minimum it becomes obvious that the term is cancelling the θ -term and hence, provides a dynamical solution to the CP problem. This global $U(1)$ symmetry would be spontaneously broken by the vacuum expectation value obtained by this scalar field. As every broken symmetry gives rise to a massless Goldstone-boson, this introduces a new particle called the axion.

While in the original model, f_a was postulated to be coincident with the electroweak breaking, so $f_a = v_f$. This though is long ruled out by measurements of the Branching Ratio (BR) of $K^+ \rightarrow \pi^+ + a$ for which an upper bound is found to be magnitudes smaller than the prediction from the axion theory [20, 21]. However, invisible axion models are still viable which introduce scalar fields which carry PQ charge but are $SU(2) \times U(1)$ singlets. Therefore, the vacuum expectation value can be much higher than for the weak interaction $f_a \gg v_f$. Two such models, both named after the original authors, are the KSVZ- [22, 23] and the DFSZ-model [24, 25] which both allow $f_a \gg v_f$ and therefore, arise a very light axion. The axion mass m_a in those models can be expressed in terms of f_a ¹:

$$m_a = \frac{m_\pi f_\pi}{f_a} \frac{\sqrt{z}}{z+1} \approx 6.3 \text{ eV} \left(\frac{10^6 \text{ GeV}}{f_a} \right), \quad (2.12)$$

with m_π and f_π being the pion mass and the pion decay constant and z the ratio of the light quark masses $z = m_u / m_d$.

All models have in common that the axion is coupling to two photons with [27]:

$$\mathcal{L}_{a\gamma\gamma} = \frac{g_{a\gamma\gamma}}{4} a F_{\mu\nu} \tilde{F}^{\mu\nu}, \quad (2.13)$$

¹Following from [26] after applying some relations

where $g_{a\gamma\gamma}$ is the coupling strength and $F^{\mu\nu}$ the electromagnetic field strength tensor. The axion-photon coupling $g_{a\gamma\gamma}$ is inversely proportional to the symmetry breaking scale f_a and as $f_a \gg v_F$ this results in a very weakly coupled axion. This $a \rightarrow \gamma\gamma$ channel is the predominantly used channel to search for axions and is also used in both experiments studied in the scope of this thesis.

2.4 Axion-Like Particles

Many extensions of the SM rely on the introduction of additional global symmetries which are spontaneously broken similar to the theory of the axion. These symmetries in turn predict light particles with weak couplings such as the axion and hence, are commonly summarized under the term Axion-Like-Particles (ALPs).

In addition to being a solution for the strong CP problem, the ALP is also a candidate for the cold dark matter of the universe as it fulfils both requirements for cold dark matter: The universe could be full of invisible ALPs with a population sufficiently large to constitute the expected dark matter energy density and they are effectively collisionless [28, 29].

Another model relevant in the scope of this thesis is a low-energy effective theory in which an ALP a couples to vector bosons through dimension-5 interactions. The Lagrangian for this model is given by

$$\mathcal{L}_{\text{int}} = -\frac{1}{4}g_{aBB}aB_{\mu\nu}\tilde{B}^{\mu\nu} - \frac{1}{4}g_{aWW}aW_{\mu\nu}^A\tilde{W}^{A,\mu\nu}, \quad (2.14)$$

where $B_{\mu\nu}$ and $W_{\mu\nu}^A$ are the $U(1)_Y$ and $SU(2)_L$ field strength tensors and g_{aBB} and g_{aWW} the associated coupling constants [30]. For physics generated with a coupling α integrated out at some heavy scale f one expects for the couplings:

$$g_{aBB}, g_{aWW} \sim \frac{\alpha}{2\pi f}. \quad (2.15)$$

This is true for pseudo-Goldstone bosons with non-vanishing axial anomalies [30]. The couplings g_{aBB} and g_{aWW} induce couplings of the ALP to photons and the bosons of the charged and neutral current above the electroweak symmetry breaking with:

$$a \rightarrow \gamma\gamma, \quad a \rightarrow \gamma Z^0, \quad a \rightarrow Z^0 Z^0, \quad a \rightarrow W^+ W^-.$$

In this thesis only couplings to two photons and to the $W^\pm$ bosons are relevant.

3. | Machine learning

Machine learning is a computational model from what is commonly known as Artificial Intelligence. One common kind of machine learning methods are so called Artificial Neural Networks (ANNs) which consist out of a network of single computational units and originally are inspired by the neurons in the human brain. While first computational models for neural networks were published back in 1943 [31] the usage of ANNs has gained popularity in the last two decades which can be explained by increased availability of computing power. Recently the use in particle physics has increased substantially.

The idea of machine learning is that a model is fed with sample data and improved in an iterative process. As a consequence, instead of writing an algorithm by hand a more or less generic model is taken as a starting point and its behaviour is trained with a known data sample. The choice of the model architecture thereby plays an important role in the performance of the trained model. The second important factor is to provide training samples which are ideally free of bias or artefacts which are not present in the real datasets.

Typically, one differentiates into three different approaches of machine learning depending on the feedback used for the training procedure [32]. In the first approach called supervised learning the training data includes the desired output for each sample and both is used in the training procedure. When using unsupervised learning the input datasets do not contain any output labels and it is up to the model to find patterns in the dataset. Finally, in reinforcement learning the model interacts with a dynamic environment and gets a score depending on its actions. Often the chosen actions do not only influence the current step but also influence subsequent scores (e.g. a model learning to play chess against a virtual opponent). Within this thesis only the approach of supervised learning is used as all models are trained with simulated data where a clear label exists for each event.

In supervised learning one commonly distinguished between two kinds of problems: classification and regression. The aim of a classification problem is to assign the correct category to a dataset (e.g. an image) while for regression the dependency between variables have to be determined and through that predict the value of a variable. In this thesis only classification problems are studied and hence, the following explanations are limited to this class of problems.

3.1 Feedforward neural networks

The most basic ANN structure is the so-called feedforward neural network [33]. In this architecture data is flowing only from the input of the network to the outputs with no loops. Consequently, each node in the network only receives inputs from nodes closer to the input than the node itself.

3.1.1 The perceptron

The perceptron is a simple algorithm which forms the basis of all modern ANNs. It was already invented in 1958 [34]. It is a type of linear classifier which is based on a linear function combining weights with the input vector. A diagram for a perceptron is shown in Figure 3.1. It takes an arbitrary number of inputs x_i , applies weights w_i to each input and combines them using the combination function Σ . In general, this combination function just sums all weighted inputs. After adding a bias b , the activation function σ is applied, yielding the perceptrons output y . In the original definition the activation function is the Heaviside step function:

$$H(x) = \begin{cases} 1 & \text{for } x > 0 \\ 0 & \text{for } x \leq 0. \end{cases} \quad (3.1)$$

Mathematically the perceptron can be written as

$$y = \sigma \left(\sum_i (w_i x_i) + b \right). \quad (3.2)$$




Figure 3.1: Graphical representation of a perceptron. The inputs x_i on the left are weighted with the weights w_i and combined by the combination function Σ . After adding the bias b , the activation function σ is applied, yielding the output y .

While the perceptron model might appear relatively simple it can already be used to express some boolean operations like AND, OR and NOT. Implementations of those boolean operators using the original perceptron model with a step function for the activation are shown in Figure 3.2. However, it was later shown that a single perceptron is not able to model an exclusive or (XOR) operation [35].

By today, a more generalized version of the perceptron is used which is, analogous to the human brain, commonly referred to as a neuron. In this more general form, the activation function is not fixed to be a Heaviside function but in principle any function can be used in its place. A selection of common activation functions will be discussed in more detail in section 3.3.




Figure 3.2: Implementation of logic operations using simple perceptrons. The weights for the inputs are denoted at the edges connecting the input to the combination function. The node for the combination and activation functions are merged for a more compact graphic.

3.1.2 Multilayer perceptron

As the complexity of problems that can be described by a single perceptron is quite limited the natural next development was to combine multiple perceptrons [36]. A multilayer perceptron is built from at least three layers: an input layer, one or more so-called hidden layers and an output layer. Perceptrons in the input layer only have one input which are the network's inputs. Perceptrons in the subsequent layers take the output of all perceptrons in the previous layer as inputs. In the first formulation of the multilayer perceptron, the original definition of the perceptron was used, but nowadays, the term is applied to networks built from the more general neurons as well. Multilayer perceptrons with more than one hidden layer are also referred to as deep neural networks. A deep neural network with four inputs, four hidden layers with different layer sizes and two outputs is shown in Figure 3.3.




Figure 3.3: Example for a deep neural network with an input layer (green), four hidden layers (blue) and an output layer (red). Each node represents a perceptron. The input layer takes four inputs x_i and the output layer produces two outputs y_0 . The colour of the interconnections denotes the value of the weights.

Each layer in the network can be of a different type and use a different activation function. The number of layers and the configuration per layer depends on the problem at hand. In the following sections, the layer types relevant to this thesis and a selection of activation functions are explained in more detail.

3.2 Layer types

In the example of a multilayer perceptron shown before all neurons are connected with each other. In practice, however, a large selection of different layer types is available in modern neural network techniques. The choice of which layers to use for building a model depends heavily on the problem at hand. Due to the large number of layer types this section only covers a selection that are used in this thesis. A full list of available layers in the machine learning framework used in this thesis can be found in the framework documentation [37].

3.2.1 Dense

The simplest layer type is the dense layer. It connects each output of the previous layer to each input of the dense layer. Therefore, it is also referred to as a fully-connected layer. Each connection is assigned a unique weight w_{ij} and a bias b_i is added to each weighted input. Mathematically the dense layer can be written as

$$y_i = \sigma \left(\sum_j (w_{ij} x_j) + b_i \right). \quad (3.3)$$

While at the start all nodes are densely interconnected, during the training connections that have little influence on the networks output are effectively removed by the accompanying weight becoming 0. This makes the dense layer quite flexible, but as a drawback it comes with a large number of trainable parameters with its $n_{\text{in}} \times n_{\text{out}}$ input weights and n_{out} bias values. For dense layers the number of trainable parameters grows exponentially with the number of dense layers making it increasingly harder to train the network.

3.2.2 Dropout

Another very important concept in neural networks is called dropout where during training neurons are randomly removed in each iteration by setting the value of this neuron to 0. This is done to prevent the network from overfitting which occurs when the network effectively remembers the training samples and therefore, not generalizing.

In Keras this technique is implemented with the Dropout layer [38]. Dropout layers have the same number of output nodes as input nodes and just feed the inputs through to the corresponding output node.

Adding dropout layers leads to the network architecture to change during training for every iteration and ensures that all parts of the network are trained evenly. As a result more iterations in the training are necessary to achieve the same performance as without any dropout, but at the same time the risk of overfitting is reduced leading to a model better at generalization. A common practice is to add a dropout layer behind each layer with dropout rates between 10 % and 25 %.

3.2.3 Convolutional layers

Convolutional layers are quite different to the layers discussed before. While the previous layers directly formed connections between input and output layers, convolutional layers scan the input in slices of configurable size. For each slice one or more filters, also called kernels, are applied which can detect features within the slice. This filter is shifted over the inputs in a given step size called strides. The technique is very effective at finding patterns within images even if the pattern arises at different positions in the image. In this thesis convolutional layers are not used to process images but rather to find patterns in data from particle detectors.

Each convolutional layer has four parameters: the number of filters, the size of the filters, strides and the padding type. While the first three parameters have already been discussed, padding is best explained using the visualization in Figure 3.4. It shows the same convolution layer once with valid and once with zero padding. When using a stride of one and zero padding, the dimensions of the output is the same as the input except for one axis due to the filters. When using valid padding on the other hand, the output's dimension is reduced.

One advantage of convolutional layers is that the number of trainable parameters is independent of the size of the input data. It depends only on the filter size $s_{x,y}$ and the number of filters applied n_{filters} . For a two-dimensional convolution there are $n_{\text{filters}} \times s_x \times s_y$ parameters.




Figure 3.4: Visualization of a two-dimensional convolution with a filter size of 3×3 and 12 filters with valid padding (left) and with zero padding (right).

3.2.4 Pooling layers

Pooling layers help to reduce the amount of data to process by trying to select the most important features in the input. This is done by dividing the input data in groups and transforming each group by either taking the average or maximum value within the group. The pooling layer has no trainable weights and only two parameters: the pooling size and the metric used for pooling.

An example for average and maximum pooling is shown in Figure 3.5. It can be seen that the output of the layer is smaller than the input effectively removing less important information.

3.2.5 Flatten layers

In some cases it is necessary to manipulate the structure of the data flowing through the network. One common problem is the need to reduce the number of dimensions to only a single dimension in order to put together datasets. Inserting a flatten layer is achieving just that. It has no trainable weights and only a single parameter to define the order of the flattened dataset.

3.2.6 Concatenate layers

When a network has either multiple inputs or the network branches at some point it becomes necessary to combine the outputs of both branches together. The concatenate layer is used for this task. As the flatten layer it just manipulates the structure of the data but not the data itself. Hence, it also has no trainable weights and the only parameter is along which axis the data should be concatenated.




Figure 3.5: Example for average pooling (left) and maximum pooling (right). The pooling groups are shown in different colours with the pooled layer on top. As can be seen the size of the output layer is only a quarter of the input layer.

3.3 Activation functions

Depending on the problem at hand the choice of an activation function differs. For example when doing a binary classification, meaning that the network should assign the input to one of two categories by yielding either 0 or 1, the activation of the output layer should clamp the value to be within 0 and 1. This requirement only applies to the output layers, the activations for any hidden layers is fully independent. Another criteria when choosing the activation function may depend on the computing resources available with some functions requiring more resources than others. For real time applications it is best to choose a linear function of more complex functions.

A selection of limited and unlimited activation functions is given in Figure 3.6. However, this covers only the most commonly used activation functions in neural networks and is far from complete. Depending on the problem one might even define custom activation functions.




Figure 3.6: Selected activation functions used in neural networks with a limited output range (left) and an unlimited range (right).

For classification the commonly used functions are sigmoid for binary classification with a single output or the softmax for multi-class classification having multiple outputs. In hidden layers often the rectified linear unit (ReLU) [39] or variations of it are used.

In this thesis all networks only use the ReLU in the hidden layers and the sigmoid function in the output layer. The ReLU is defined as

$$f(x) = \begin{cases} x & \text{for } x > 0 \\ 0 & \text{else.} \end{cases} \quad (3.4)$$

While the sigmoid function is given by

$$f(x) = \frac{1}{1 + e^{-x}}. \quad (3.5)$$

3.4 Training of neural networks

Training of a neural network is the process of optimizing the weights and biases of the neural network to yield an output as close as possible to the expected output. To train a neural network four main ingredients are needed: the definition of the model's architecture, a training dataset including the expected outputs¹, a function to evaluate the performance and an optimizer.

3.4.1 Model architecture

The definition of the model architecture describes the flow of information from the data inputs to the outputs of the network. Depending on the available data the model has either a single or multiple input layers. These input layers are connected to the output using one or more of the previously discussed layers. The exact composition lays the foundation for the performance of the network and hence, for complicated networks great care has to be taken in the design phase.

¹At least for supervised learning which is the only kind of learning discussed here.

3.4.2 Training data

The second crucial ingredient is the input dataset used for training. To ensure that the neural network actually learns to solve the intended problem, the dataset must be carefully prepared. This process is often the most time-consuming step in developing neural networks. In the following, the most important requirements for the datasets are outlined.

The dataset must contain diverse data ideally also covering corner cases. At the same time, it is important that the dataset is unbiased. The importance of having no bias can be demonstrated by a study aiming to train a neural network to recognize cancerous lesions in pictures. While the initial results looked promising, the research group noticed that the network based its decision heavily on the presence or absence of a ruler in the photographs. The researchers afterwards noticed, that doctors are more likely to put a ruler next to lesions which look concerning. [40, 41]

The dataset also should be uniformly distributed over the feature space and be free of duplicates. Having duplicates or not uniformly distributed features leads to those samples to be deemed more important than others. If the amount of non uniformities is known this may also be remedied by including weights for each sample in the dataset to correct for the relative importance.

The last requirement is that the dataset must be sufficiently large to perform proper training. As the required number of samples heavily depends on the networks architecture and trainable parameters there is no universally applicable method to predict how many samples are required at minimum. Depending on the nature of the data it might be possible to artificially increase the statistics by transforming the samples and thus creating new samples. For example if there are symmetries in the input data one could translate, mirror or rotate the samples to create additional training samples.

With all of those requirements in mind one can construct a training dataset. However, not all data should be used for training the network. More precisely, training of neural networks requires multiple independent datasets in the process. In general one splits the available dataset into three separate datasets with each containing samples uniformly distributed in the feature space. The three datasets are usually called:

- Training dataset,
- Validation dataset,
- Test dataset.

The largest fraction of the available data is used for the training dataset. It is used for optimizing the networks weights during training. The training is usually performed on this dataset multiple times. Each iteration over the full set is called an epoch.

The validation dataset is a smaller sample and is used to evaluate the performance of the network during the training process. This is done after each epoch. As the datasets are independent, this allows to get an unbiased measure for the network's performance. In cases where the network becomes too specific or starts to remember the samples in the training dataset the validation performance will become smaller than the performance measured with the training dataset.

As the validation dataset indirectly introduces a bias on the training process by influencing the adjustments made by the developer. For example if a specific network architecture is found to perform good on the validation dataset one might prefer it over another performing slightly worse. This effect is even stronger when dedicated tools are used to optimize the network

architecture and training parameters in what is called a Hyperparameter scan as the validation performance is often used as a measure in that optimization.

3.4.3 Performance metric

In order to assess the performance of a trained network some kind of metric is required. The performance measure used in the training process is called the loss function. The choice of the loss function depends on the specific problem. The networks trained in the context of this thesis are all classification problems with only two classes. This means that each sample has a label with either 1 or 0 depending on the output class. For example this could be 0 for background and 1 for signal. For this kind of problem the binary cross-entropy loss function is used as a loss function being defined as

$$L(y_i, y_i^{\text{true}}) = -\frac{1}{N} \sum_{i=1}^N y_i^{\text{true}} \cdot \log(y_i) + (1 - y_i^{\text{true}}) \cdot \log(1 - y_i), \quad (3.6)$$

where N is the total number of inputs, y_i the predicted label and y_i^{true} the true label for the inputs.

Effectively, the loss is the negative mean of the logarithm of the distance between the true and the predicted label. The better the network is the closer the summands become to one and the smaller the loss becomes.

As an easier to interpret metric for the performance of networks for classification the accuracy can be used. It gives the percentage of correctly predicted samples with a threshold of 0.5. The accuracy is defined as

$$\text{accuracy} = \text{mean}(y_i == (\hat{y}_i > 0.5)). \quad (3.7)$$

3.4.4 Optimizer

The optimizer is the component that defines how the training process is performed. It takes the feedback from the loss function and determines the required adjustment of the weights and biases to minimize the loss. One optimizer commonly used is the gradient descent. It will be explained in detail in the next section on the training procedure.

3.4.5 Training procedure

With the four main ingredients in place, it is time to put them together and start the training procedure. A schematic drawing of the process is given in Figure 3.7. It shows the flow of information through the network. The input is propagated through the network in batches of samples with tunable size in order to yield a prediction on the label y_i for each event. The loss function takes the predicted label and the true label and computes the loss score which in turn is given to the optimizer. After each batch, the optimizer updates the weights such that the loss function is minimized and the next batch is processed. The optimum batch size depends on the problem at hand, but common values are in the range of 64 and 256.

This process is repeated for the full dataset multiple times with the dataset being shuffled after every epoch to prevent bias of the ordering on the update of the weights.

A commonly used optimizer is the gradient descent [42]. Its goal is to find the right weights and biases at which the loss function is minimum. The gradient descent is divided into two steps:




Figure 3.7: Schematic drawing of the neural network training process for a network with two layers. The flow of information through the nodes is denoted with arrows. Based on [42].

The forward pass and the backward pass. During the forward pass the predictions and loss is calculated for a batch of training samples. The backward pass then calculates the gradient of the loss with regard to the networks parameters. By moving the weights in the opposite direction from the gradient by a small amount the loss should be minimized. This can be written as

$$w_{ij}^{\text{new}} = w_{ij} - \alpha \cdot \nabla L(y_i, y_i^{\text{true}}) = w_{ij} - \alpha \frac{dL(y_i, y_i^{\text{true}})}{dw_{ij}}, \quad (3.8)$$

where w_{ij} are the weights used in the forward pass, w_{ij}^{new} are the updated weights and $L(y_i, y_i^{\text{true}})$ the loss function (e.g. the binary cross-entropy loss). The indices i and j denote the neurons in the corresponding set of layers. The parameter α is called the learning rate. It controls the magnitude of the update of the weights. The influence of the learning rate is shown for a one-dimensional parameter space in Figure 3.8.

Nowadays, multiple improved extensions of the gradient descent method exist. One of the most commonly used methods is the Adaptive Moment Estimation (Adam) which extends the gradient descent with momentum terms to prevent oscillations in the parameter space [43].

3.5 Variable Importance

Understanding how a neural network comes to its decision is no easy task. This is because the trained weights have in general no obvious meaning. To get some insights on what the network bases its decision on a method called partial dependence functions can be used [44].

The partial dependence functions are computed by evaluating the networks with modified data-sets. For each variable x with values $x_i, i \in [1, N]$ the network is evaluated N times. For each evaluation i the dataset is copied and the value x_i is used for all samples. The standard deviation of the networks output can then be used as a measure for the variable importance. For a variable which has low influence on the networks decision the standard deviation is expected to be small while it becomes larger the more important a variable is.




Figure 3.8: Gradient descent with a learning rate of $\alpha = 0.1$ (left) and $\alpha = 0.4$ (right) for an example loss function in a one-dimensional parameter space.

4. CERN & the LHC

The CERN is a European research organization located close to Geneva (Switzerland) on the border to France and operates a multitude of laboratories, accelerators and large scale experiments. It was founded in 1954 by twelve countries in Europe [45]. The main purpose of CERN is providing the accelerators and accompanying infrastructure needed for high-energy physics experiments. At the time of writing in total 23 countries are members of CERN with over 12000 users [46].

4.1 The accelerator complex

Even though by most people CERN is first associated with having the largest accelerator in the world, namely the Large Hadron Collider (LHC), there is not only a single accelerator but a whole complex of different accelerators located at CERN. An overview of this complex is shown in Figure 4.1 where all active accelerators are included together with a selection of experiments.

Looking at this schematic closely it is obvious that some accelerators are linked to other ones, larger one building a chain of accelerators with increasing size. This is owed to the fact that each accelerator is only able to store and accelerate particles within a certain energy range. Hence, before protons reach the LHC they start their journey in a linear accelerator called LINAC 2 (replaced by LINAC 4 in 2020). Following this stage the protons are further accelerated in three circular accelerators (Proton Synchrotron Booster, Proton Synchrotron and Super Proton Synchrotron (SPS)) after which they have reached an energy of 450 GeV and are injected into the LHC. To provide particle beams to other experiments not located at the LHC multiple extraction points exist in the accelerators to direct particles to other beam lines. As both experiments featured in this thesis are located at the LHC, the following section gives a more detailed introduction while ignoring the other accelerator complexes.

4.2 Large Hadron Collider

With a circumference of 27 km the LHC is the largest particle accelerator built so far. Two proton beams are accelerated in opposite direction and brought to collision at crossing-points located in the four large experiments: ATLAS, CMS, LHCb and Alice. By design each beam can reach previously unprecedented energies of 7 TeV, however, up to this point the maximum energy used was 6.5 TeV as the superconducting magnets required to bend the beams have not been commissioned to the required currents for running at the 7 TeV. As previously stated, the protons are injected at an energy of 450 GeV. The protons are gradually accelerated to the target energy and kept at this energy from there on. Hence, the LHC is not only an accelerator but acts as a storage ring at the same time. To keep the particles on the curved track in total 1232




Figure 4.1: Schematic overview of the CERN accelerator complex and a selection of experiments using the accelerators. Modified from [47] to include location of the the FASER experiment.

superconducting dipole magnets are used with 392 quadrupole magnets allowing for focussing the beam. Over time the number of protons in the beam is decreasing due to the proton collisions and losses in the beam pipe and once it decreased below a certain threshold the beam is dumped and another fill cycle is started by injecting protons from the SPS. The threshold is chosen to allow for as many collisions as possible and on average an injection cycle is started once to twice per day.

The beam is not a continuous current of protons but is composed of discrete bunches. Each bunch contains approximately 1.1×10^{11} protons and is spaced with a distance of 25 ns. The bunches are further grouped to so-called trains with empty bunch positions in between to allow for abort gaps which can be used to ramp up the kicker magnets steering the beam out of their circular orbit into a special beam dump target in case of an emergency or in a nominal dump of the beam before an injection cycle. As a side effect it allows for the experiments to record background events during the runs at brief intervals where it is ensured that no collisions are taking place.

The different data taking campaigns of the LHC are divided into Runs. Run 1 is the first operation period from 2009-2013 with the first year being needed for machine commissioning and the physics program starting on 30 March 2010. During this time energies of 3.5 TeV to 4 TeV per

beam. The Runs are separated by Long Shutdowns (LS) needed to perform maintenance and upgrades on both the accelerators and the experiments. Run 2 lasted from 2015 to 2018 during which the LHC operated at 6.5 TeV per beam.

4.3 Heavy ions and the LHC as a photon accelerator

The previous explanations focussed only on accelerating protons in the LHC. However, it is also capable of running with two ion beams as well as a mixed configuration with one proton beam and one ion beam. So far two different ions were accelerated in the LHC: Xenon- and Lead-ions. While the Xenon-ion run was only conducted for a few hours, Lead-ions are used on a regular basis. For future runs other lighter ions are under consideration, but so far no definite plans exist. For Run II two runs with lead ions lasting one month each were conducted in 2015 and 2018. Lead-lead collisions are performed at a centre-of-mass energy of $\sqrt{s_{NN}} = 5.02$ TeV.

Charged particles travelling at relativistic energies produce electromagnetic fields surrounding the particles. Due to the coherent action of the protons within the nucleus of the lead ions the electromagnetic field is very strong [48] being in the order of 10^{14} T when accelerating lead ions. One can imagine this field being a cloud of photons surrounding the ions and consequently as a photon beam. This model is called Weizsäcker-Williams method or equivalent photon approximation (EPA) and is often used to simplify the calculation of photon induced interactions [49, 50]. The photons are considered to be quasi-real with a small virtuality of $Q^2 < 1/R^2$, with R being the radius of the charge distribution, giving $Q^2 < 10^{-3} \text{ GeV}^2$.

Clearly observing photon induced signatures is only feasible if it is not overlapped by particles created by a collision of the beam particles. Therefore, only events are considered where the beam particles are not colliding but passing each other closely, which are called Ultra-Peripheral Collisions (UPC) events. In UPC events the impact parameter is larger than the sum of the particles radii and at least one of the two colliding particles stays intact after the collision.

As the photon flux per beam is scaling with the square of the ion charge Z^2 the resulting cross-sections are significantly enhanced for heavier ions with a factor of $Z^4 \approx 5 \times 10^7$ for lead-lead beams ($Z = 82$) compared to proton-proton beams ($Z = 1$). Hence, the LHC is suited much better as a photon collider when running with lead beams than with proton beams.

4.4 Measurement of the luminosity

An important quantity for studying particle interactions in a high-energy collider experiment is the number of total collisions taking place. This measure is the number of collision per unit time and unit area and is called instantaneous luminosity. For Gaussian shaped beam profiles colliding head-on with beam profiles of σ_x and σ_y in x and y direction it is given by [51]:

$$\mathcal{L} = \frac{\dot{N}}{\sigma_{\text{tot}}} = \frac{N_1 N_2 f N_b}{4\pi\sigma_x\sigma_y}, \quad (4.1)$$

with N_1 , N_2 being the number of particles in the bunches, N_b the number of bunches and f the revolution frequency. Integrating the instantaneous luminosity yields the so-called integrated luminosity $L = \int \mathcal{L} dt$.

Another important quantity is the cross-section which describes the probability for a process to occur. It is measured as an area and is usually given in barn with $1 \text{ b} = 10^{-28} \text{ m}^2$. It is

further differentiated between the total and the fiducial cross-section. The total cross-section considers that the interaction occurs at all while the fiducial cross-section is taking into account the limited phase space of the measurement. As the detector cannot measure the full phase space only the fiducial cross-section can be measured in the experiment and must be extrapolated to the full phase space to produce the total cross-section.

The fiducial cross-section for a process under investigation can be calculated with:

$$\sigma_{\text{fid}} = \frac{N_{\text{process}}}{L \cdot A \cdot \epsilon}, \quad (4.2)$$

where ϵ is the detector efficiency, A the geometrical acceptance and N_{process} the number of observed events.

A precise knowledge of the integrated luminosity is important to confirm predictions of the SM and even more to probe for beyond SM occurring at very low rates. This requires measuring and monitoring the beam parameters over time. The effective beam profiles are measured every few months using the van der Meer method [52]. In this method the two beams are offset with respect to each other and the overlap is varied over time while monitoring the rate in the luminosity detectors. Each of the large experiments monitors the luminosity independently and hence, the use different detectors for this purpose. The detectors used by the ATLAS experiment are briefly listed in the ATLAS detector description in section 5.4.

Part I

Light-by-light scattering

5. The ATLAS experiment

The ATLAS detector is a general purpose detector at the LHC at CERN. It consists of a range of different sub-detectors covering a solid angle of nearly 4π . A computer generated view of the ATLAS detector is shown in Figure 5.1 with parts of the detector systems cut away to allow for a view inside the detector. The detector can be divided into three large parts: the barrel and the two end-caps. In the barrel all subsystems are arranged in a cylindrical shape around the beam pipe. The end-caps are oriented perpendicular to the beam pipe and enclose the barrel on both sides with a gap in between. The overlap between barrel and end-caps provides coverage over a large solid angle, while the gap leaves enough space for cables and other appliances like cooling pipes.

The figure also shows the coordinate system used in the ATLAS detector. The origin of the coordinate system is the Interaction Point (IP). The z-axis is following the beam axis with the x-axis pointing towards the centre of the LHC and the y-axis pointing up. Vectors are described in cylindrical coordinates using the distance to the z-axis r , the polar angle θ and the azimuthal angle φ . For convenience, the so-called pseudorapidity η is often used instead of the polar angle θ . The pseudorapidity has the advantage of being Lorentz invariant under boosts along the longitudinal axis and is defined by $\eta = -\ln(\tan \theta/2)$.




Figure 5.1: Computer generated view of the ATLAS detector [53]. The detector is shown with a cut-out to make the various sub-detectors visible more clearly. The coordinate system defined for the detector is shown with a vector depicting the momentum of a particle. Its origin is in the centre of the detector, with the z-axis following the beam axis and the x-axis pointing towards the centre of the LHC.

Like most general purpose detectors, the ATLAS detector follows the so-called onion shell-principle which means that the subsystems are forming layers around the IP. This allows the detection of all particles created in a collision but neutrinos. The amount of material is increasing with each layer, except for the muon system. The schematic view in Figure 5.2 shows a perpendicular cross-section view of one sector of the detector where the various detector systems are clearly visible. The interaction of various kinds of particles is also depicted in the drawing showing that each has a clear signature allowing to identify them during the reconstruction process.

The first layer is the so-called Inner Detector (ID) which is surrounded by a solenoid magnet. It is used to measure the path of charged particles at multiple points in space to allow reconstruction of tracks for each particle. The next layers are the two calorimeters, which absorb all particles except for muons and neutrinos. The first calorimeter layer is the electromagnetic calorimeter which stops electrons and photons while hadrons, muons and neutrinos travel through it without much interaction. Hadrons are then absorbed in the second calorimeter layer called hadronic calorimeter. Hence, only muons and neutrinos reach the muon spectrometer which forms the last layer of the detector. It is located in a toroidal magnetic field which causes the trajectories of charged particles to bend. By measuring the track of the muon through the magnetic field its charge and momentum can be reconstructed. The only particles basically invisible to the detector are neutrinos which escape the detector without interacting with it at all.

All subsystems of the detector and their basic working principle will be discussed in the following sections in more detail.

5.1 Inner detector and magnetic field

The innermost detector is a tracking detector recording the trajectories of charged particles. It consists of three subsystems using different detector technologies as can be seen in Figure 5.3. It is 6.2 m long and has a diameter of 2.1 m.

Directly surrounding the beam pipe are four layers of high precision silicon pixel detectors. The inner three layers feature a pixel size of $50 \times 400 \mu\text{m}^2$ while the most outer layer has a pixel size of $50 \times 250 \mu\text{m}^2$ [54, 55]. The high resolution of the pixel detector is required to allow the extrapolation of particle tracks to the interaction region to find the vertex position of the interactions. In total there are 1744 pixel sensors in the pixel detector with a total of over 92 million readout channels.

The pixel detector is followed by the Semiconductor Strip Tracker (SCT) detector. It is built from 15912 silicon microstrip sensors. In the barrel region each module has two layers, each consisting of two 6 cm-long sensors with a strip pitch of $80 \mu\text{m}$. The two layers are slightly rotated against each other to provide a measurement in two dimensions. Due to the small stereo angle between the two layers the resolution in one coordinate is much better than in the other giving an intrinsic resolution of $17 \mu\text{m}$ in the $R - \varphi$ plane and $580 \mu\text{m}$ in z for the SCT detector [54].

The outermost part of the inner detector is the Transition Radiation Tracker (TRT). It is built from up to 73 layers of straw tubes. The tubes are made of two $35 \mu\text{m}$ multi-layer films bonded back-to-back which are coated with a thin layer of aluminium protected by graphite-polyimide layer on the inside to produce a conducting layer. In the centre of each tube a $31 \mu\text{m}$ diameter tungsten wire plated with 0.5 to $0.7 \mu\text{m}$ gold is stretched as the anode. A voltage of -1530 V is




Figure 5.2: Schematic drawing of a sector of the ATLAS detector visualizing the interaction of particles with the different sub-detectors. The circle in the bottom is depicting the beam pipe and the coloured lines are particles measured by the detector. In regions where the particle lines are dashed, little to no interaction with the detector material occurs.

applied to the cathode producing a radial electric field. The tubes are filled with a mixture of around 70 % xenon or argon with CO_2 and O_2 . Charged particles traversing the straw tubes ionize the gas and thereby create electron-ion-pairs. Due to the electric field the electrons drift towards the wire while the ions slowly drift towards the wall of the tube creating measurable currents in the anode and the cathode respectively. By measuring the time difference between the two currents the minimum distance of the particle to the drift wire can be calculated with an accuracy of $\approx 110 \mu\text{m}$ [56]. With each charged particle creating signals in around 30 tubes the track can be reconstructed by combining all drift radius measurements.

Another purpose of the straw tubes is the discrimination between different charged particles. To achieve this the material of the tube walls is constructed to have different refractive indices which causes charged particles crossing the detector to create transition radiation. The amount of transition radiation depends on the particle's Lorentz factor and, thus, depends on its mass [57]. Therefore, measuring the amount of transition radiation allows for distinguishing between electrons, muons and pions.

The three Inner Detector (ID) subsystems are surrounded by a superconducting solenoid magnet system. It is aligned with the beam axis and provides an axial magnetic field with a field strength of 2 T for the inner detector. It has a diameter of 2.46 m and is 5.8 m while weighing a total of 5.4 t. To ensure stable superconductivity the magnet coils need to be cooled to 4.5 K by continuously pumping liquid helium through the cryogenic system.

The paths of charged particles in the inner detector are bent due to the magnetic field allowing




Figure 5.3: Computer generated drawings of the inner detector. The left shows a cut-away view of the full inner detector while the right image shows a detailed section of the barrel section. Images modified from [54].

for a measurement of the charged particles' momentum by reconstructing the curvature of the tracks. The inner detector and magnet system are designed to allow for a resolution of $\sigma_{p_T}/p_T \approx 0.05\% p_T \oplus 1\%$ [54].

5.2 Calorimeters

The purpose of a calorimeter is to measure the energy of a particle. This is achieved by completely stopping the particle and absorbing all of its energy within the calorimeter. The fact that the particles are stopped in the calorimeter explains why it has to be placed after the tracking detector. Stopping a high energetic particle requires a large amount of material. To reduce the material cost and the physical size of calorimeters they are often built as sampling calorimeters. Sampling calorimeters are built from alternating layers of absorbers made from a dense material and active detectors used to measure the energies of particles.

In case of ATLAS there are two calorimeters: the electromagnetic calorimeter is optimized to stop electrons and photons and the hadronic calorimeter to measure hadrons. The size of both calorimeters is optimized to ensure that the showers are always fully stopped at the expected particle energies. A schematic drawing of the calorimeters in the setup is shown in Figure 5.4.

Electromagnetic calorimeter

The electromagnetic calorimeter is using lead as the absorber material which is mechanically stiffened with 0.2 mm thick stainless-steel sheets and liquid argon as the active detector material. It has an accordion structure to allow for a full coverage in φ without any gaps. The readout planes for the liquid argon are made out of conductive copper layers on a polyimide film and are placed in the gaps between the lead absorber layers. To provide spatial information, patterns are etched into the readout layers both radially and in η . As can be seen in Figure 5.4 there are three layers with a very fine segmentation closest to the beam line and larger cells for the outer two layers. The fine strip layer is $4.3X_0$ long and consists of strips with a size of $\approx 4.7 \times 147.3 \text{ mm}^2$. This layer is providing important information required for the reliable

identification of photons and electrons and discriminating against noise. The second layer is slightly coarser using cells instead of strips and covers $16X_0$. The last layer is also featuring cells which are twice as large as the ones used in the second layer.

In the barrel the electromagnetic calorimeter consists of two half barrels split at $\eta = 0$ with each spanning a range up to $|\eta| < 1.375$. The end-cap electromagnetic calorimeters cover a region of $1.375 < |\eta| < 3.2$. In the transition regions on both sides from the electromagnetic barrel calorimeter to the electromagnetic end-cap calorimeter a large amount of material totalling 5 to $10X_0$ is located in front of the first calorimeter layer [58]. As this large amount of material will produce showers before the calorimeter the strip layer would not be useful and is therefore omitted. This region is referred to as the calorimeter crack and defined by $1.37 \leq |\eta| \leq 1.52$. Due to the missing strip layer the region is excluded from the fiducial range for measurements where the particle identification in the electromagnetic calorimeter is critical.




Figure 5.4: Cut-away view of the ATLAS the electromagnetic and hadronic calorimeters (left) and a section of the electromagnetic calorimeter (right). In the cut-away view, the inner detector is shown in grey in the centre. Images modified from [54].

Hadronic calorimeter

The hadronic calorimeter is surrounding the electromagnetic calorimeter and is divided into the barrel calorimeters, the liquid argon hadronic end-cap and the Forward Calorimeter (FCal) in the very forward region. The tile barrel is divided again into three sections: one central section with a length of 5.8 m and two extended barrel sections with a length of 2.6 m each in the forward and backward region. Together they cover a range of up to $|\eta| < 1.7$.

In the barrel region the hadronic calorimeter uses alternating layers with tiles of steel and polystyrene scintillators. To convert the scintillation light from ultraviolet into visible light the polystyrene is doped with wavelength-shifting fluors. Using optical fibres the light produced is guided to Photomultiplier Tubes (PMTs) located on the outside of the calorimeter. The granularity in the tiles is coarser than in the electromagnetic calorimeter as the showers are larger by nature and hence, do not require a high resolution.

The end-cap and the Forward Calorimeter (FCal) use liquid argon as the active material and

copper and tungsten as the absorber materials which have a larger density than steel. The difference in the choice of the absorber material is explained by the space constraints in the forward region. The FCal covers the range $3.1 < |\eta| < 4.9$ and as the electromagnetic calorimeter is not extending that close to the beam pipe it must measure both electromagnetic and hadronic showers.

5.3 Muon spectrometer

The Muon Spectrometer (MS) is the outermost detector system in the ATLAS detector. The central part of the muon spectrometer is the huge barrel toroid magnet system which is built from eight superconducting magnet coils. Overall it is 25.3 m long and has an inner and outer diameter of 9.4 m and 20.1 m respectively and produces a magnetic field with up to 2.5 T.

The spectrometer covers a pseudorapidity range of $|\eta| < 2.7$ and allows triggering on muons for $|\eta| < 2.4$. It is designed to measure the transverse momentum of muons with a resolution of $\sigma_{p_T}/p_T < 10\%$. This translates to the requirement of a resolution of $\leq 50 \mu\text{m}$ for the track reconstruction. To be able to meet the requirements on spatial and timing resolution while also covering the large area four different gaseous detector technologies are employed.

For the high precision measurements Monitored Drift Tubes (MDT) are used which consist of three to four layers of drift tubes each. The drift tubes have an outer diameter of $\approx 30 \text{ mm}$ and a length of 1 to 6 m. They are filled with a gas mixture of 93:7 Ar:CO₂ and pressurized to 3 bar. In total 1088 chambers with 18 different geometries are utilized in three layers with a radius of 4 – 6, 8 and 10 m respectively. The large number of different geometries is owed to the large number of support structures and appliances, e.g. cooling systems, that share the same space as the muon spectrometer. To reduce the number of holes in the active area around 70 custom tailored chambers complete the Monitored Drift Tubes (MDT) system. In its entirety the chambers cover an area of 5500 m².

In the forward region with $|\eta| > 2$ the particle flux in the innermost Muon Spectrometer (MS) layer exceeds the maximum rate for the MDT chambers of 150 Hz cm⁻². Therefore, this region uses Cathode-strip chambers (CSCs) instead of MDT chambers. The Resistive Plate Chambers (RPCs) are multiwire proportional chambers [59] with the wires oriented parallel to the z-axis and the two plates with segmented cathodes. One cathode has the strips oriented parallel and the other orthogonal to the wires. This allows for a reconstruction of the position in the transverse plane. In the bending direction the resolution achieved is 60 μm while the resolution in the non-bending direction is only 5 mm due to a coarser segmentation.

For triggering on muons by transverse momentum and event topology RPCs and Thin Gap Chambers (TGCs) are used. They have a worse spatial resolution compared to the MDT and CSC detectors but a much better timing resolution with 5 ns for the RPC and 25 ns for the TGC. In the barrel three layers of RPCs are used. They are placed on both sides of the MDT chambers in the second layer and on the inside or outside of the MDT chambers in the third layer. The RPCs are built from two parallel plates filled with a detection gas with an electric potential applied across them. Strips on the outside surface of the plates capacitively pick up the signals produced by avalanches on the conductive plates. The TGCs are again multiwire proportional chambers with the speciality that the distance between the wires is larger than the distance between the wire and the plates. This gives excellent time resolution with the drawback of a reduced spatial resolution.

5.4 Luminosity detectors

The ATLAS experiment also possesses detectors for measuring the luminosity. Those detectors are located as close as possible to the beamline in the forward and backward region of the detector to detect proton fragments coming from a collision. The two primary detectors for measuring the luminosity are the Beam Conditions Monitor (BCM) [60] and the Luminosity measurement using Cherenkov Integrating Detector 2 (LUCID2) [61] which allow a deadtime-free measurement of the luminosity for every bunch. The BCM is built from four $8\text{ mm}^2 \times 8\text{ mm}^2$ diamond sensors located around the beampipe at a distance of 1.84 m on both sides of the interaction points. The LUCID2 detector is measuring Cherenkov light produced by particles originating from the inelastic collisions of the bunches. The Cherenkov light is produced in either quartz fibres or directly in the quartz windows of the PMTs used for the measurement.

5.5 Reconstructing objects

Measurements from all relevant sub-detectors are combined to reconstruct particles in an event. By combining the measurements of a particle through the detector up to where it is absorbed a hint on the kind of particle can be extracted as discussed before.

5.5.1 Charged particle tracks

Charged particles traversing the detector deposit energy in the tracking detector. As a single particle deposits charge in multiple adjacent pixels of the detector, a clustering algorithm groups these charge deposits to so-called hits and calculates three-dimensional space-points for the clusters. An iterative algorithm is used to find so-called seeds, which are a combination of three space-points that could potentially belong to a valid track. This process is performed independently in the different sub-detectors of the ID to reduce the computational workload. Assuming a perfect helical trajectory in a uniform magnetic field the momentum and impact parameters are estimated for each of the seeds allowing for filtering the seeds and subsequently the use of a combinatorial Kalman filter [62] to build track candidates utilizing all systems of the ID.

At this stage a single hit can be part of multiple track candidates and more processing is needed to extract the real tracks. To resolve these ambiguities a track score is calculated for each track candidate based on e.g. the number of compatible space-points and the logarithm of the track's transverse momentum. The score is reduced if it passes an active detector layer at a position where no hits are present. Due to the finite size of the detector pixels/strips clusters potentially can belong to more than one particle and hence, shared clusters are allowed. However, tracks containing more than two shared clusters are considered invalid and are removed from the list of candidates. Using the track scores and some more coarse requirements the final selection of track candidates is formed. The additional requirements are chosen to minimize the occurrence of bad track candidates and hence, reducing both the storage space and the complexity for the individual analyses, while still being coarse enough not to limit the possibilities of the analysers.

As muons pass the calorimeters their properties are not only measured in the ID but additionally the MS is used to further improve momentum resolution and the overall muon reconstruction efficiency. The track candidate search in the MS is first performed completely independently of the ID. It starts with a search for hit patterns inside the muon chambers to find straight segments which could be part of a muon trajectory. These straight segments are matched from MDT chambers to close-by trigger chambers. Similar to the reconstruction in the ID matching

segments are identified by their kinematic properties. A χ^2 -fit over the matched segments is then used to reconstruct the trajectory with high accuracy. Four different combinations between ID tracks and MS tracks are considered, however, in most cases only the combined muon and the segment-tagged muon variants are relevant. The combined muon case means that a track from the ID is compatible with a track from the MS. In this case the track is fitted once again including the hits from both the ID and the MS. In the segment-tagged muon case an ID-track can be extrapolated to match a single segment in the MS. This is mostly the case for muons with a low transverse momentum that do not reach the outer layers of the MS or when they cross the detector in a region with a reduced coverage in the MS. The addition of the MS allows the ATLAS detector to reach a muon reconstruction efficiency close to 99 % over most of the pseudorapidity range of $|\eta| < 2.5$ for muons with a transverse momentum $p_T > 5 \text{ GeV}$. The relative transverse momentum resolution is measured to be 1.7 % at small pseudorapidities and 2.3 % for pseudorapidities $|\eta| \approx 2.3$ [63].

5.5.2 Electromagnetic particles

Reconstructing particles from the electromagnetic (EM) calorimeter requires finding cells in the EM calorimeter with energy deposits and combining them correctly into clusters. To achieve this, multiple steps of cluster finding are performed [58]. Tracks from the inner detector are extrapolated to the cluster positions to differentiate between electrons and photons. By evaluating the properties of associated tracks, e.g. where in which layer of the ID the first hit occurs, photons that convert to electron-positron pairs before reaching the calorimeter can be identified [64]. To reduce clusters with large amount of their energy coming from poorly functioning or noisy calorimeter cells additional requirements are imposed. By evaluating the timing associated with the calorimeter cells energy deposits from particles not matching to the collisions are removed.

Finally, the energy associated to the reconstructed photons/electrons is corrected using a dedicated calibration [65]. This corrects for energy loss due to interaction with material in front of the calorimeter as well as for energy losses within the calorimeter itself due to lateral and longitudinal shower leakage. Along with the kinematic properties of the electromagnetic particle a set of so-called shower shapes are calculated and stored. These are used during the analysis step to identify the particles and further remove potential noise. A lower limit on the electromagnetic particles transverse energies of $E_T = 2.5 \text{ GeV}$ is imposed during the reconstruction procedure as particles with lower energies potentially lose too much energy in the material in front of the calorimeter. This has little to no impact in most physics analyses as the energy regime is much higher there and electromagnetic particles with such low energies are not relevant in those cases.

5.5.3 Jet reconstruction

As discussed in chapter 2, quarks can not exist freely and therefore, quarks produced in the collision quickly hadronize forming so-called particle-jets. These jets are produced in a cone-shape structure and leave clusters of energy deposits in the calorimeters. Similar to the procedure used for the electromagnetic particles, calorimeter cells are combined to clusters belonging to a particular jet [66]. The reconstruction procedure in ATLAS includes a set of different jet reconstruction algorithms with the standard being the anti- k_t jet clustering algorithm [67]. A dedicated calibration procedure is applied to the reconstructed jets to correct for material in front of the calorimeters, dead cells in the calorimeter and shower leakage [68]. By combining various properties measured in the calorimeter and exploiting the ID to measure the vertex

position special algorithms exist to identify to some extent the kind of hadron that initiated the jet [69].

5.6 Trigger system

Reading out the full detector continuously is impossible due to the dead time of various detector readout electronics and the large amount of data that would need to be transferred without interruptions. However, even if reading out the detector continuously would be possible one would record lots of events that do not contain any interesting physics and exceed both the available storage space and the computing power available for performing the reconstruction procedure.

Therefore, a decision if an event potentially contains interesting processes must be made already during data taking in a process called triggering. For Run II the ATLAS Trigger System consists of a two level decision where both have to fire for the event to be stored [70]. The first level trigger, called Level-1, is completely done in hardware logic. It takes information from the calorimeters and the muon-system. As the Level-1 trigger must deliver the trigger decision within a time window of only $2.5\ \mu\text{s}$ the inputs provided are only of a coarse granularity. After the Level-1 trigger, the event rate is reduced from the LHC bunch crossing rate of 30 MHz to 100 kHz. Together with its trigger decision it forwards regions-of-interest to the next trigger stage to limit the search space. It also triggers the detector to be read out storing all event information in buffers for further processing.

The second trigger level is the so-called High level trigger (HLT) and is a software based system. The trigger decision on Level-2 must be ready after 200 ms allowing for more complicated analyses of the event topology. As the Level-1 trigger already started the detector readout, the Level-2 trigger has access to more detailed information. The regions-of-interest from the Level-1 trigger allow the High level trigger (HLT) to only look at the interesting event data without having to load the information for all sub-systems. Once events are accepted by the HLT the buffered data is forwarded for processing and storing on disk. The final trigger rate is on average around 1 kHz.

5.7 Event simulation

An important foundation of science is the ability to model and predict the outcome of an experiment in order to use measurements to validate the theory. The vast complexity of the detectors used in particle physics experiments makes this even more crucial. In order to predict the result for a given process all steps from the interaction of the particles from the beam, the interaction with the detector material, the corresponding detector response to the readout electronics must be fully simulated. To allow for accurate simulations the conditions of the detector used for the simulation need to be closely matched to the real conditions. Therefore, the simulation has to be run for the different data taking periods independently while also taking into account input from the ATLAS conditions database to allow correct modelling of e.g. dead channels in the tracker modules or different high voltage bias levels.

A simplified flow diagram of the event simulation chain used in the ATLAS experiment is shown in Figure 5.5 with each step being explained in more detail in the following sections.



Figure 5.5: Simplified flow diagram of the ATLAS event simulation chain with the algorithms drawn in boxes and the objects created written over the arrows.

5.7.1 Event generation

The first step in this chain is the generation of the events on what is called particle level. This step is performed by various generators, which take probabilities derived from the standard model processes and produce the outgoing particles. The event generator describes all final state particles before any interaction with the detector. As this process is using random numbers to reproduce the probabilities, the generators are often referred to as Monte-Carlo event generators. In the context of this thesis one can name examples like **SuperChic** 3 [71] or **STARlight** 2 [72]. Due to limited knowledge of the underlying processes or computational resources, the generators are often employing simplifications limiting the accuracy of such generated events. Often this is specified by the maximum perturbation order of the underlying process considered in the calculations. Given these limitations it is important to choose a suitable generator for the process under investigation and compare the results produced by different generators to reduce and understand the bias originating from the event generation. During the following steps this particle level information, which is also referred to as truth information/truth level, is kept together with the additional data generated to allow for a comparison of the detector response to the event.

5.7.2 Detector simulation

Using the final state particles from the event generator as an input the interaction of particles with the detector material of the ATLAS detector as well as the decay of particles within the detector is simulated [73]. This step exploits the **Geant4** simulation toolkit [74] together with a detailed geometrical description of all detector components. Within **Geant4** a particle's trajectory is calculated by moving it through the detector in discrete steps and calculating the energy loss, changes in trajectory due to interactions or a magnetic field and possibly the creation and decay of new particles. The particle's trajectory is followed until it either completely leaves the detector or the energy of the particle falls below the minimum detectable energy. The value of the threshold is chosen low enough to ensure that further contributions from the particle would not influence the detector response any more [74]. The simulation of the full detector is computationally costly and, hence, some simplifications can be enabled optionally. An example for this is the **FastCaloSim**-algorithm [75] which uses parametrizations of electromagnetic and hadronic showers to produce a relatively accurate shower simulation without the need to simulate all individual shower particles. This allows for much higher speeds at the cost of accuracy and therefore, such options must not be enabled without detailed considerations.

Information on particle decays occurring in the detector simulation is combined with truth information provided by the generator. In addition, truth information on tracks and other objects is collected which is mainly used for detector upgrade and performance studies.

5.7.3 Digitization

In the following step, the hits and energy deposits are run through a tool which is generating signals compatible with what is produced by the real detector readout. The term signals means here the low level information like collected charges in the trackers or calorimeters, drift times or pulses as generated by PMTs. As this step is basically emulating the readout drivers in the ATLAS detector, it is called the digitization step. To include pile-up (e.g. multiple pp interactions in a single event) the events can be overlaid with a randomly chosen pre-generated set of so-called low bias events or recorded events. This is only done within the digitization step to reduce the amount of resources that would otherwise be required for the previous simulation step. From here on simulated events are indistinguishable from real events.

5.7.4 Reconstruction

After the previous step, the generated data has the exact same format as the data recorded with the real detector except for the additional truth data. This allows using the same algorithms to reconstruct the event objects for both simulated and recorded data. The reconstruction procedure is run as explained in section 5.5. The information available at this point excluding the truth data is commonly referred to as reconstruction level or detector level.

6. Measurement of light-by-light scattering and ALP limits

Following the classical Maxwell's equations, scattering of two photons is not possible due to the superposition principle, which is a result of the linearity of the equations [76]. In Quantum Electrodynamics (QED), however, this process is possible at high energies and albeit very rare, can be observed. The lowest possible order for the light-by-light scattering process $\gamma\gamma \rightarrow \gamma\gamma$ in QED is the scattering of photons via a virtual loop involving fermions or the $W^\pm$ as shown in Figure 6.1.

While the scattering phenomenon was already observed indirectly in multiple experiments, it was never directly observed with quasi-real photons before the measurement described in this thesis. Most notably, it was measured in the anomalous magnetic moment of a single electron and positron [77] and the anomalous magnetic moment of the muon [78].

However, just observing and measuring properties of the light-by-light scattering process directly for the first time is not the only motivation for this analysis. In addition, some postulated extensions of the SM predict contributions to the light-by-light scattering process and hence, possibly make it sensitive to new physics. In principle, every particle coupling to photons could be produced by and decay into two photons through an s-channel process: $\gamma\gamma \rightarrow X \rightarrow \gamma\gamma$. This would produce a resonance peak in the invariant mass spectrum of the di-photon final state. One possible candidate in the probed mass range are Axion-Like-Particles (ALPs) [79] which were introduced in section 2.3. The Feynman diagram for the s-channel process involving an Axion a ($\gamma\gamma \rightarrow a \rightarrow \gamma\gamma$) is shown in Figure 6.1.




Figure 6.1: Feynman diagram of the lowest possible order for scattering of photons in QED (left) and the production of axion-like particles (right).

A study from 2013 [80] concludes that the light-by-light scattering process could possibly be observed in UPC events during Pb+Pb runs. As explained in section 4.3, in UPC events the beam

particles are not colliding but are passing each other closely. This enables photons surrounding the lead ions to interact and allows for photon induced processes, like light-by-light scattering, at the LHC. In two analyses exploiting data from the 2015 lead-lead campaign first evidence for light-by-light scattering was found by both the ATLAS and the CMS experiments in 2017 and 2018 [81, 82]. Due to the limited size of the datasets from this campaign only 13 and 14 event candidates were found by ATLAS and CMS respectively with expected background contributions of 2.6 and 4 events. Even with this small statistics the excess of data relative to the background-only hypothesis has a significance of 4.4σ (3.7σ) being just below the 5σ mark commonly considered the threshold for an observation in particle physics.

The chapter opens with an overview of the measurement procedure. It is followed by a short description of the simulated and experimental data samples included in the analysis. Afterwards, the event and object selection is discussed. The development and testing of a dedicated photon identification procedure is laid out in detail next. After determining the background contributions, the kinematic distributions are shown. Following that, the procedure of evaluating the systematic uncertainties on the measurement is explained along with the various sources for systematic deviations. The next two sections focus on the fiducial and differential cross-section calculations. The extraction of limits on the ALPs process is discussed next. Finally, the full chapter is summarized and a conclusion on the analysis is given. The results have been published in [2] and [3].

6.1 Measurement procedure

Due to the kinematic of the process it is far from trivial to measure light-by-light scattering with the ATLAS detector. The small di-photon invariant mass $m_{\gamma\gamma}$ results in photons with transverse momenta in the range of 0 to 10 GeV which is much smaller than the requirement of 40 GeV normally used. In Figure 6.2 the normalized particle level distributions from simulation are shown with no selections except for standard cuts applied on generator level. As can be seen there the di-photon invariant mass is falling steeply, making it important to include photons down to the lowest energies possible to maximize the light-by-light candidate yield.




Figure 6.2: Normalized particle level distribution from MC simulation for the di-photon invariant mass $m_{\gamma\gamma}$ (left) and the single photon transverse momentum $p_{T,\gamma}^Y$ (right) without any fiducial selection except for cuts applied on generator level.

As a consequence, the planning for this analysis begins before any data taking is performed as the process is so different from what is normally measured at the ATLAS experiment that a dedicated trigger is needed to ensure any candidate events being recorded at all. This is especially crucial here as the cross-section for $\gamma\gamma \rightarrow \gamma\gamma$ scattering is expected to be only 80 nb [83] and hence, a very efficient trigger strategy and reconstruction method is required. Extensive studies on this were performed in the context of another PhD thesis by a collaborator and a set of dedicated triggers was put in place for the 2018 data taking campaign.

After the data is recorded, the relevant quantities are reconstructed and saved in so called Analysis Object Data (AOD) files. To reduce the amount of data to be processed in the analysis the dataset is filtered further. The filter criteria applied in this step are e.g. requirements on the trigger which caused the event to be recorded or the reconstructed particles in the event. In addition, objects not needed in the analysis are removed to reduce the size of the dataset. In this specific analysis this reduced the number of events processed in the analysis phase from 28.4 million events (≈ 115 TB) to 1.35 million events (≈ 1.5 TB). Up to here the standard ATLAS toolchain is used with a dedicated configuration for the derivation step.

To further reduce the number of events a program is developed to convert the AOD files to so called NTuples. In this step a stricter selection is imposed on the events and the object selection criteria relevant for this analysis are imposed. The NTuples are saved in a ROOT [84] file as branches of a tree. At this point the data to be analysed is reduced to less than 0.65 kB / event. Up to here all the processing is performed on high performance clusters allowing to process many samples in parallel. To run jobs in parallel on a node the GNU parallel [85] tool is used.

The final analysis step is written in Python 3 [86] using PyROOT to read the NTuples. Here the final object selection is performed and the histograms of the measured distributions are calculated for each sample. A separate script is used to correctly normalize the simulated samples and merge them to construct the expected distribution. Additional programs are written for the unfolding and ALPs limit extraction.

6.2 Measured and simulated samples

For this measurement Pb+Pb collisions with a centre-of-mass energy of $\sqrt{s_{NN}} = 5.02$ TeV are used, which were recorded in two data taking periods in the 2015 and 2018 Heavy Ion collision campaigns of the LHC. The combined dataset corresponds to a total integrated luminosity of 2.22 nb^{-1} with an uncertainty of 3.2 %. In the following plots being created for a specific data taking period are labelled either with “Data 15” or “Data 18”. Whenever the label just states “Data”, it is based on the combined dataset.

In order to study the light-by-light scattering process and the possible backgrounds in a controlled way, simulated samples of the processes are required. These are generated using the procedure described previously in section 5.7.

For the light-by-light scattering signal process two different approaches are utilized. The first approach uses a generator called **SuperChic 3** [71] to generate the events. It includes box diagrams with leptons or quarks and $W^\pm$ bosons. The second approach is based on calculations from [83]. These calculations are folded with the Pb+Pb photon flux simulated by a generator called **STARlight 2.0** [72, 87]. For both generators the theoretical uncertainties on the cross-section are dominated by the limited knowledge of the nuclear form-factors and hence, the initial photon fluxes giving an uncertainty on the prediction of 10 % [81, 88]. A smaller

contribution to the uncertainty is arising from not including higher-order corrections into the calculations. For the invariant di-photon mass range studied here this contribution is on the order of 1 to 3 % [89, 90].

The Central Exclusive Production (CEP) of the di-photon final state is produced via the strong interaction through a quark loop in the exchange of two gluons in a colour-singlet state. For this **SuperChic 3** is used once again. As the theory uncertainty on the cross-section of this $gg \rightarrow \gamma\gamma$ process is of $\mathcal{O}(100\%)$ the normalization is not taken from the generator, but is performed using a data-driven method. In addition, an alternative sample for the CEP process is generated using **STARlight 2.0**.

As the $\gamma\gamma \rightarrow e^+e^-$ process can possibly fake the signal signature if both electrons are misidentified as photons it is generated as well. This process is also used for multiple detector performance studies. It is modelled using **STARlight 2.0** which combines the Pb+Pb photon flux with the leading-order diagram for $\gamma\gamma \rightarrow e^+e^-$.

Lastly, for the ALP-related studies, the $\gamma\gamma \rightarrow a \rightarrow \gamma\gamma$ process is simulated using **STARlight 2.0** for different axion masses m_a in the range of $5 \text{ GeV} \leq m_a \leq 100 \text{ GeV}$ with a spacing of 1 GeV for $m_a < 30 \text{ GeV}$ and a spacing of 10 GeV above.

6.3 Event selection

The selection of events is a two-step process. Firstly, the events have to be preselected to be recorded during the data taking process. Secondly, in the analysis procedure the events have to be filtered to yield the events containing signals from light-by-light scattering. In addition, events with bad detector conditions need to be removed based on the so called good-run-list, which contains information when parts of the detector were not functioning correctly.

6.3.1 Trigger

In a collision of lead ions a multitude of processes can occur leading to signals in the detector. Therefore, the events have to be carefully identified and selected already at the trigger level. For the light-by-light scattering this selection is not trivial as the momenta of the outgoing particles are much smaller than the particles usually studied in ATLAS. Hence, a set of dedicated triggers is used in the data taking. These triggers are sensitive to events with moderate activity in the calorimeter system and very little signal in the rest of the detector. The triggers for the two data taking periods are not identical due to optimizations performed in advance of the 2018 run based on experience gained in earlier data takings. The trigger strategy and all trigger related studies were done by a collaborator and hence, are only shortly summarized in this section. All information is taken from [2, 91].

2015 trigger requirements

In 2015 a single trigger called **HLT_gg_upc_L1TE5_VTE200** was used, which imposes the following requirements [92]:

1. at Level-1 the seed requires a total transverse energy E_T between 5 GeV and 200 GeV in the entire detector,

2. veto on Minimum Bias Trigger Scintillators (MBTS) activity, meaning a maximum of one hit in the inner ring of the MBTS,
3. a final HLT requirement that allows a maximum of 10 hits in the Pixel detector, which vetoes events containing tracks with more than 10 hits.

The `HLT_gg_upc_L1TE5_VTE200` trigger was running unprescaled in the full 2015 data-taking period, which means that no events were discarded to limit the rate and all events where the trigger fired were read out and stored to disk.

2018 trigger requirements

For the 2018 Pb+Pb run a set of two new triggers (`HLT_hi_upc_FgapAC3_hi_gg_upc_L1TAU1_TE4_VTE200` and `HLT_hi_upc_FgapAC3_hi_gg_upc_L12TAU1_TE50`) was introduced [91]. These triggers require the event to fulfil the subsequent criteria and were also used without any prescale:

1. one of the following Level-1 seeds activated:
 - a minimum of one EM cluster of $E_T > 1$ GeV and a total E_T between 4 and 200 GeV in the entire calorimeter,
 - at least 2 EM clusters with $E_T > 1$ GeV and a total E_T in the calorimeter below 50 GeV.
2. little signal in the FCal ensured by rejecting events with $\sum E_T^{\text{FCal}} > 3$ GeV on any side,
3. low activity in the ID by vetoing events with more than 15 hits in the Pixel Detector.

2018 supporting trigger requirements

For the 2018 data taking period multiple supporting triggers were in place to record events with two electrons in the final state. These events are used for various efficiency and detector calibration studies.

Trigger performance

To measure the Level-1 trigger performance $\gamma\gamma \rightarrow e^+e^-$ events with exactly two reconstructed tracks and two EM clusters with a minimum transverse energy of $E_T > 1$ GeV are used. Noise is suppressed by requiring the reduced acoplanarity $A = 1 - \Delta\varphi/\pi$ between the two clusters to be below 0.01. As the photons from the $\gamma\gamma \rightarrow \gamma\gamma$ scattering process are balanced in E_T the trigger efficiency is parametrized as a function of the sum of the energy of the two clusters $\sum E_T = E_T^{\text{cluster1}} + E_T^{\text{cluster2}}$. The extracted efficiency is shown in Figure 6.3. For the 2018 data taking period the efficiency is around 65 % at $\sum E_T = 5$ GeV and reaches a plateau of full efficiency at $\sum E_T = 8$ GeV. In contrast, the efficiency for the 2015 data taking period is only 50 % at $\sum E_T = 8$ GeV reaching the plateau at $\sum E_T = 10$ GeV. This difference is caused by the much higher threshold for the transverse energy in the 2015 trigger configuration.

To extract smooth functions which can be used to reweight the simulated samples to reflect the measured trigger efficiency an error function is used to fit the measured points for both data taking periods. The fit is also shown in Figure 6.3 with the uncertainties shown in a shaded band.

The MBTS veto efficiency is estimated using $\gamma\gamma \rightarrow \mu\mu$ events passing a dedicated control trigger and is found to be $98 \pm 2\%$. For the FCal veto efficiency $\gamma\gamma \rightarrow e^+e^-$ events passing a control trigger are used which does not impose the FCal veto and its efficiency is found to be $99.1 \pm 0.6\%$.




Figure 6.3: Level-1 trigger efficiency as a function of the sum of two electron clusters extracted from $\gamma\gamma \rightarrow e^+e^-$ events [3]. Data is shown as points for the 2015 (blue) and 2018 (black) data taking periods. The fit to data is shown as dashed line for 2015 and solid for 2018 with uncertainties given by the two coloured bands.

Lastly, the pixel veto imposed at the HLT is found to be highly efficient with the inefficiency being neglectable.

6.3.2

Object selection

Light-by-light scattering events contain only two photons in the final state. Hence, exactly two photons are required to be measured in the detector. The invariant mass of the di-photon system must be larger than 2.5 GeV. Due to the similar detector signature of the $\gamma\gamma \rightarrow e^+e^-$ process with two electrons in the final state a strict veto on charged particle tracks is imposed. As a measure to discriminate against fake-photon sources the transverse momentum of the di-photon system $p_T^{\gamma\gamma}$ needs to be smaller than 1 GeV for an invariant mass of the di-photon system $m_{\gamma\gamma}$ below 12 GeV and below 2 GeV for $m_{\gamma\gamma} > 12$ GeV.

The photons are required to be emitted in a back-to-back topology. This can be measured by using the so called reduced acoplanarity A , which is a measure for the deviation of the two particles from being coplanar. The reduced acoplanarity A is given by

$$A = 1 - \frac{\Delta\phi_{\gamma\gamma}}{\pi}, \quad (6.1)$$

where $\Delta\phi_{\gamma\gamma}$ is the difference in the azimuthal angle of the two photons.

Details on the procedure of reconstructing and identifying the objects is given in the following section.

6.4 Object and particle reconstruction

6.4.1 Track reconstruction

Tracks are objects describing the path of a charged particle through the detector. They are reconstructed using information from the various tracking detectors using the default ATLAS track reconstruction explained in section 5.5. The standard ATLAS identification methods are however too strict to ensure a sufficient reduction of background induced by $\gamma\gamma \rightarrow e^+e^-$ where both electron tracks are not reconstructed. There are two track object definitions used in the analysis: Standard tracks and so-called pixel tracks. The standard tracks are tracks reconstructed using the full inner detector while the pixel tracks only make use of the two innermost silicon detectors (Pixel + SCT).

The standard tracks need to have a transverse momentum $p_T > 1 \text{ GeV}$, a pseudorapidity $|\eta| < 2.5$, at least one hit in the Pixel detector and at minimum a total of six hits in silicon detectors. For the pixel tracks the requirements are even looser with a minimum transverse momentum of $p_T > 0.5 \text{ GeV}$, $|\eta| < 2.5$ and at least three hits in the Pixel detector. To suppress pixel tracks originating only from noise in the pixel detector, they are only kept if they are close to a photon candidate fulfilling $\Delta R < 0.5$ with ΔR being a measure for the distance in the $\eta - \varphi$ plane defined as

$$\Delta R = \sqrt{(\eta_1 - \eta_2)^2 + (\varphi_1 - \varphi_2)^2}.$$

6.4.2 Photon reconstruction

For the reconstruction of the photons the standard ATLAS algorithm is utilized as described in section 5.5. The reconstruction is setting a lower limit on the minimum measurable photon transverse momentum with $p_T > 2.5 \text{ GeV}$ [64]. This lower limit is given by the amount of material in front of the calorimeter. The requirements for a photon candidate for this analysis are $p_T > 2.5 \text{ GeV}$ and a pseudorapidity $|\eta_\gamma| < 2.37$. In addition, it must not be within the calorimeter crack region defined by $1.37 < |\eta_\gamma| < 1.52$.

The photon also needs to fulfil identification criteria to ensure that the energy deposit is not originating from noise or another particle passing through the detector. As the identification method requires a longer description, this is explained in more detail in section 6.5.

To measure the photon reconstruction efficiency at low photon energies in data, $\gamma\gamma \rightarrow e^+e^-$ events can be used as sometimes one of the electrons will emit a hard-bremsstrahlung photon due to interacting with material in the ID. These photons have a similar energy to the photons produced in light-by-light scattering. Taking the ratio of events, where a hard-bremsstrahlung photon was reconstructed and where it was not, although one is expected, the reconstruction efficiency can be calculated. Hence, events with exactly two back-to-back charged-particle tracks and one electron with a minimum energy of $p_T > 4 \text{ GeV}$ are selected. To ensure that the two tracks are back-to-back the acoplanarity of the two is required to be smaller than 0.3. One of the tracks has to match the electron within a window of $\Delta R < 1$. Due to emitting the hard-bremsstrahlung photon the energy of the second electron is too small to be properly reconstructed. To make the selection cleaner a cluster inside the EM calorimeter is required within $\Delta R < 0.3$ around the second track.

For low transverse momenta the electron track is bent strong enough by the magnetic field to lead to a good separation between the electron cluster and the hard-bremsstrahlung photon.

Using an approximation for the bending radius [93] the separation can be estimated:

$$R [\text{m}] = \frac{p_T [\text{GeV}]}{0.3B [\text{T}]},$$

where R is the bending radius in m, p_T the transverse momentum of the charged particle in GeV and B the field strength of the magnetic field in T. With this it can be estimated that with a transverse momentum of $p_T < 1.5 \text{ GeV}$ a minimum separation of $\Delta R > 0.4$ can be ensured when extrapolating to the front of the EM calorimeter. Hence, the second track is required to fulfil $p_T < 1.5 \text{ GeV}$. The transverse momentum of the photon is approximated by the difference of the transverse energy of the electron and the momentum of the track emitting the hard-bremsstrahlung photon $p_T^\gamma = E_T^e - p_T^{\text{track}}$.

In total 3454 hard-bremsstrahlung events are selected in the combined dataset. To validate the selection is sufficient to get a clean sample, the momentum of the hard-bremsstrahlung photon and the transverse momentum of the system comprised of the two tracks and the photon is compared with simulated $\gamma\gamma \rightarrow e^+e^-$ events. The comparison is shown in Figure 6.4. It shows reasonable agreement between data and simulation. Therefore, the selection of the hard-bremsstrahlung events is considered pure.

To calculate the reconstruction efficiency the number of hard-bremsstrahlung events with a reconstructed photon is divided by the total number of hard-bremsstrahlung events. This is performed in two parametrizations: by the expected photon transverse momentum $p_T^\gamma = E_T^e - p_T^{\text{track}}$ and by pseudorapidity as given by the track η^{track} . The resulting efficiencies are shown in Figure 6.5. For lowest photon energies of 2.5 GeV the reconstruction efficiency is $\approx 60 \%$. It rises quickly with the photon energy reaching a maximum of 90 %.

For the efficiency as a function of η it is visible that the reconstruction efficiency is high around $|\eta| \approx 0$ with 85 % and falling towards the forward region reaching a minimum around the calorimeter crack. In the far forward region the efficiency is reaching a maximum with 90 %.

To correct the reconstruction efficiency from the simulation to reflect the efficiencies measured on data the simulated events are weighted by so-called Scale Factors (SF). The Scale Factors (SF) are given by the ratio between the observed values for data and simulation. The SF are parametrized by the photon transverse momentum $p_{T,\gamma}$ for the nominal analysis. A second set of SF parametrized by the photon pseudorapidity is used to derive uncertainties on the correction. This is discussed in more detail in section 6.8.

6.5 Photon identification

In most ATLAS analyses one of two reference photon Particle Identifications (PIDs) – called *loose* and *tight* – is used to identify photons and differentiate them from noise. They take various, so-called shower shape variables, into account and use rectangular cuts to make a selection of photons. [64] The *loose* selection is cutting only on variables from the second layer ($w_{\eta,2}$) and the hadronic layer (R_η and R_{had}). The *tight* selection is using all shower shape variables discussed in subsection 6.5.1. As most analyses within ATLAS are targeting prompt photons with a high energy the cuts are optimized to have the best efficiency above 40 GeV.

Both selections are tested for this analysis with the single photon identification efficiency shown in Figure 6.6. To study the efficiency of the two reference PIDs the $\gamma\gamma \rightarrow \gamma\gamma$ signal simulation is used. For all reconstructed photons with $p_T > 2.5 \text{ GeV}$ and $|\eta| < 2.37$ that match a truth




Figure 6.4: Transverse momentum of the hard-bremsstrahlung photons (top) and of the electron, track and photon system (bottom) for the 2015 (left) and 2018 (right) data taking periods compared to simulated events.




Figure 6.5: Photon reconstruction efficiency in dependence of the photons transverse momentum (left) and the photons pseudorapidity (right). Efficiencies measured in data (red) are compared to the efficiencies from the $\gamma\gamma \rightarrow e^+e^-$ simulation (blue).

photon with $\Delta R < 0.3$ the different identification results are checked. The efficiency is then given by the ratio of identified photons over the total number of reconstructed (and truth-matched) photons. As expected, the *tight* identification gives significantly lower efficiency than the *loose* identification due to the stricter cuts applied. At lowest photon energies the efficiency with *loose* PID is $\approx 50\%$ and only 15% for *tight* PID. Given that most signal is expected at low photon energies these efficiencies would lead to a large loss of signal.

Another important aspect is the background rejection capability of the identification method. It is defined as $1 - \epsilon_{\text{background}}$ where $\epsilon_{\text{background}}$ is the efficiency for background events being classified as signal. To achieve a clean selection of signal photons the background rejection must be as high as possible. Clearly, the signal efficiency and background rejection are not independent of each other. While one could accomplish a signal efficiency of 100% easily by just accepting all photons, this would also lead to a background rejection of 0. Thus, tuning the identification method is always a trade-off between signal efficiency and background rejection.

For the default identification methods the single photon background rejection is found to be quite low with 39% for *loose* and 81% for *tight* PID.




Figure 6.6: Efficiency for the identification of a single photon as determined from the simulated $\gamma\gamma \rightarrow \gamma\gamma$ sample for the two default ATLAS photon PID methods. The efficiency is shown as a function of the photon transverse momentum p_T (left) and the pseudorapidity η (right).

Given the poor performance of the default selection methods for both signal efficiency and background rejection two different approaches are investigated. The aim is to improve the identification efficiency while keeping a strong background rejection down to the lowest energies. Thus, maximizing the yield for $\gamma\gamma \rightarrow \gamma\gamma$ events while suppressing background contamination.

The first approach is to use low energy samples to derive rectangular cuts similar to the reference identification methods but optimized for lower energies. The second approach uses neural networks trained to discriminate between photon and background. After introducing the full list of shower shapes and the photon and background samples the two newly developed methods are discussed.

6.5.1 Shower shape variables

The photon shower shapes are variables calculated from measurements in the different parts of the calorimeters and characterize the shape of the clusters [64]. The variables used in the scope of this thesis are shown schematically in Figure 6.7 and summarized in the following:

- f_1 – longitudinal shape of the shower given by the fraction of the energy deposit associated with the photons cluster in the first layer over the total energy of the cluster
- E_{ratio} – ratio of the difference between the maximum energy deposit ($E_{\text{max},1}$) and the second-largest energy deposit ($E_{\text{max},2}$) over the sum of the two deposits:

$$|E_{\text{max},1} - E_{\text{max},2}| / (E_{\text{max},1} + E_{\text{max},2})$$
- $w_{\eta,2}$ – measure for the lateral width of the shower given by $w_{\eta,2} = \sqrt{\frac{\sum E_i \eta_i^2}{\sum E_i} - \left(\frac{\sum E_i \eta_i}{\sum E_i}\right)^2}$ with E_i being the energy in cell i and η_i the pseudorapidity. The sum is calculated in a 3×5 ($\eta \times \phi$) region of cells in S_2
- R_η – ratio of the energy in the 3×7 ($\eta \times \phi$) cells and the 7×7 ($\eta \times \phi$) cells around the cluster position
- $R_{\text{had},1}$ – leakage into the hadronic calorimeter calculated by the ratio of the energy deposit in the first sampling layer of the hadronic calorimeter and the total cluster energy in the EM calorimeter
- R_{had} – leakage into the hadronic calorimeter calculated by the ratio of the energy deposit in the hadronic calorimeter and the total cluster energy in the EM calorimeter




Figure 6.7: Visualizations of the variables from the calorimeter used for the neural network photon identification. Inspired by [94].

6.5.2 Photon and background samples

To develop a method to discriminate between signal and background samples for both cases are required. In this specific case the fake photon sample is extracted from a fake enriched control region. This region is defined by imposing the following selections:

- passes signal trigger requirements,
- exactly two good-quality photons with $p_T > 2.5$ GeV and $|\eta| < 2.37$ (without any identification requirements),
- no tracks originating from a primary vertex,
- invariant mass for the di-photon system $m_{\gamma\gamma} > 5$ GeV,
- transverse momentum of the di-photon system $p_T^{\gamma\gamma} > 2$ GeV.

It is mostly dominated by fake signals generated from cosmic muons crossing the EM calorimeter and also includes noise by the calorimeter itself.

The signal photon samples are extracted from the two simulated signal samples using the signal cuts described in subsection 6.3.2 but additionally requiring the two reconstructed photons to be matched to a photon on truth level.

Before any optimizations are performed the datasets are split into a training and a testing set with a ratio of approximately 70 % for training to 30 % for testing. This strict splitting of the datasets allows to validate the performance of the selection method on a “new” sample and helps to detect if the selection criteria are over-specific.

Figure 6.8 shows the normalized shower shape distributions for the full signal sample as well as the background sample. The clearest separation can be seen in the variables f_1 and E_{ratio} . However, the other variables show some discriminating power as well.

The shower shapes change significantly with the pseudorapidity η . The identification is therefore performed in multiple pseudorapidity bins independently. The bins are chosen such that two bins cover the barrel calorimeter and two bins cover the end-cap region.

6.5.3

Custom cut-based identification

For a previous analysis of light-by-light scattering a custom photon PID using rectangular cuts was utilized for the photon identification [82]. This used the same three shower shape variables as the *loose* PID. In a similar fashion, a set of rectangular cuts is determined for this analysis exploiting the larger dataset available for this analysis with respect to the earlier analysis. Early studies showed that using the full set of six shower shape variables does not reduce the efficiency and hence, for this analysis all six variables are used in the rectangular cuts.

To establish the linear cuts a framework for multivariate data analysis is used called Toolkit for Multivariate Data Analysis (TMVA) [95]. It allows building identification methods using various different technologies. However, for this analysis only the rectangular cuts are utilized. The optimization inside the framework is done on a simple Monte Carlo (MC) sampling method. This randomly samples cut values within the range of the discriminating variables and determines the efficiency and background rejection for each set of cut values. There are more efficient fit methods available at the cost of increased complexity. As the dataset available for the fitting is small it was therefore decided to use the MC fitting. All other options were left at the TMVA default settings.

The resulting cut values for a single photon PID efficiency of 95 % as well as the corresponding background rejection for the four $|\eta_\gamma|$ -bins are given in Table 6.1. Additionally, in Figure 6.9 the background rejection as a function of the single photon PID efficiency is plotted.




Figure 6.8: Normalized distributions of the photon shower shapes extracted from the simulated signal sample (blue) and the background sample from data (red).

Table 6.1: Cut values for the cut-based photon PID at a single photon identification efficiency of 95 %. The background rejection values are determined by evaluating the testing dataset.

Variable	$ \eta_\gamma < 0.6$	$0.6 \leq \eta_\gamma < 1.37$	$1.52 \leq \eta_\gamma < 1.81$	$1.81 \leq \eta_\gamma < 2.37$
$w_{\eta,2}$	< 0.0174	< 0.0162	< 0.0165	< 0.0145
f_1	> 0.0588	> 0.0694	> 0.0768	> 0.0185
E_{ratio}	> 0.5448	> 0.1236	> 0.1279	> 0.6822
R_η	> -142.7355	> -57.3662	> 0.3259	< 0.2671
R_{had}	< 0.1573	< 8.2703	< 7.6526	< 4.6064
R_{had1}	< 2.5479	< 0.4005	< 4.0244	< 0.1505
Background rejection	98.0 %	92.6 %	70.0 %	84.2 %

6.5.4 Neural network photon identification

The identification method based on rectangular cuts is limited to a single threshold per shower shaper variable and hence, does not take the values of the other shower shapes into account. To overcome this, the second approach uses a neural network which can base the identification not only on the value itself but includes the relation between the variables.

For training the neural network one of the most used machine learning toolkit called **Keras** is

used in combination with **Tensorflow** [96] as the backend. To optimize the network multiple different architectures are tested and their performance compared. Due to the small number of input variables, the finally selected network is fairly simple consisting out of only two hidden layers with 20 and 10 neurons respectively. Those hidden neurons all use a rectified linear unit as activation.

The output layer consists of a single neuron with a sigmoid activation function to ensure that the output value y is in the range of 0 to 1. When the output value is larger than a given threshold t_{signal} ($y > t_{\text{signal}}$) the input is classified as a photon while for $y < t_{\text{signal}}$ it is considered background. Using this single output opposed to a multi neuron output as often done for classification allows for an easy tuning of the signal threshold and in turn setting a desired identification efficiency. To prevent overfitting a dropout layer is included after each hidden layer with a dropout factor of 0.1 during training. A summary of the networks' architecture is given in Table 6.2.

Table 6.2: Neural network architecture for the photon identification. The dropout rate of 0.1 results in 10 % of the connections being randomly dropped during training.

Layer	Neurons	Activation function	Parameters
Dense	20	ReLU	140
Dropout			0
Dense	10	ReLU	210
Dropout			0
Output	1	Sigmoid	11
Total			361

As done for the cut-based identification for each $|\eta_\gamma|$ -region a dedicated network is trained. This approach is chosen over the alternative of using $|\eta_\gamma|$ as an additional input due to limited statistics of the background sample in the forward region. The limited statistics would make it more complicated to get consistently good results over the full $|\eta_\gamma|$ range with just a single network.

To prepare the datasets for training each variable x is normalized using Gaussian normalization with the following transformation:

$$T(x) = \frac{x - \bar{x}}{\sigma} \quad (6.2)$$

For every input variable the mean μ and the standard deviation σ are determined such that the mean of all transformed training samples equals to 0 and the standard deviation is 1. Therefore, every input variable has specific parameters for the transformation. Additionally, every set of input variables is labelled with 1 for a photon and 0 for background. To have a balanced training set a random subset of each category (signal, background) is chosen so that the number of samples in each category is the same.

The training dataset is again split into two datasets. 90 % are used for the actual training and 10 % to validate the networks' performance after each iteration. As an optimizer Nesterov-accelerated Adaptive Moment Estimation (Nadam) is used and the mean squared error is taken as the loss function. The training is performed over 100 epochs, however the validation accuracy is determined after each epoch and the model with the best validation accuracy is saved in the

end. This acts as a second safeguard to prevent overfitting, which would lead to the validation accuracy to drop with respect to the training accuracy. Increasing the number of epochs for the training process is evaluated but does not improve the accuracy any further.

After the training, the network is evaluated on the testing dataset to ensure that the network is neither over- nor under-fitted. Comparing the evaluation on the validation and the testing dataset shows pretty good agreement. Hence, the network is not overfitted and can give good predictions for new data as well.

As discussed before, the performance of a classification method is rated based on the combination of signal efficiency and background rejection. One method to visualize this is to evaluate the signal efficiency and background rejection for different thresholds in the range between 0 and 1. Plotting the background rejection versus the signal efficiency for different binary classifiers makes it easy to compare them.

In Figure 6.9 this is shown for the neural network and the rectangular cut based approach. As a reference the working point used for the previous light-by-light analysis is also shown. It can be clearly seen that the neural network performs better over all $|\eta_\gamma|$ -bins. While in the central region ($|\eta_\gamma| < 0.6$) the difference in background rejection at 95 % is only 0.6 %, in the forward region the neural network rejects up to 25 % more single photons from background than the rectangular cuts. The results for the ATLAS reference *loose* and *tight* PID are not included in the plot as they are by far outside the range of the other methods. Table 6.3 summarized the single photon background rejection capabilities for the two new methods in all $|\eta_\gamma|$ -bins at a fixed single photon efficiency of 95 %. It also shows the consistency between the training and testing datasets with the difference between the two being below 1.3 % giving no indication for overfitting.

The thresholds for the networks are selected to set the working point to a single photon PID efficiency of ≈ 95 % based on evaluating the photons in the testing sample.

Table 6.3: Background rejections for the different selection methods and datasets at a single photon signal efficiency of 95 %.

Selection	Data	$ \eta_\gamma < 0.6$	$0.6 \leq \eta_\gamma < 1.37$	$1.52 \leq \eta_\gamma < 1.81$	$1.81 \leq \eta_\gamma < 2.37$
Neural network	train	98.7 %	98.3 %	87.6 %	90.4 %
Neural network	test	98.5 %	97.7 %	87.2 %	90.1 %
TMVA Cuts	train	98.1 %	93.6 %	71.1 %	83.7 %
TMVA Cuts	test	98.0 %	92.4 %	70.0 %	84.2 %

To gain some understanding on the impact of the different shower shape variables on the identification decision the partial dependence functions discussed in section 3.5 are evaluated for all models. The procedure described there is performed for all six used shower shapes and the four networks. The resulting ranking is shown in Figure 6.10. There it is visible that the longitudinal shower shape f_1 is the most important variable for discriminating signal from background photons. The hadronic leakage is less important in the central region while it gains more discriminative power in the forward region.




Figure 6.9: Comparison of the background rejection as a function of the photon PID efficiency for a single photon in the four $|\eta_\gamma|$ -regions. Shown are the results for the neural network (blue), rectangular cuts (red) and the rectangular cuts used in the 2016 publication (green dot).




Figure 6.10: Relative importance ranking on the photon particle identification decision for the neural network for the different models.

6.5.5 Measurement of the photon identification efficiency

To validate the photon PID efficiencies extracted from simulated samples, the efficiency is also measured in data. This is done on two different types of events. Firstly, $\gamma\gamma \rightarrow e^+e^-$ events where one of the electrons is emitting an Final-State Radiation (FSR) photon are used. A second method is exploiting the $\gamma\gamma \rightarrow e^+e^-$ events containing a hard-bremsstrahlung photon discussed in subsection 6.4.2.

Final-state radiation photons

One method to select photons in data is to use $\gamma\gamma \rightarrow e^+e^-$ events where one of the electrons radiates a photon in the final state. The selection criteria for the FSR events were determined and studied by a collaborator and hence, this section is only shortly explaining the selection as well as the results from studying the photon identification performance based on the FSR sample. [91]

The following selection criteria are imposed to preselect such events with an FSR photon:

- no extra trigger requirement (except for the ones used for the pre-filtering of the analysis dataset)
- exactly two opposite-charge tracks with a minimum transverse momentum $p_T > 0.5 \text{ GeV}$ passing the *loose* track identification selection
- invariant mass of the two-track system $m^{TT} > 1.5 \text{ GeV}$

In addition to the preselection a photon candidate with a minimum energy of $E_T > 2.5 \text{ GeV}$ and $|\eta_\gamma| < 1.37$ or $1.52 < |\eta_\gamma| < 2.37$ is required. The photon candidate must be separated from both tracks with $\Delta R > 0.35$ to ensure the clusters of the electron and the photon candidate are isolated and no energy leakage between them occurs. From simulation, it is determined that a further cut on the transverse momentum of the two track + photon system $p_T^{T\gamma} < 1 \text{ GeV}$ is required to suppress contamination with other processes.

This selection yields 212 events in the 2015 dataset and 1333 events in the 2018 dataset. In Figure 6.11 and Figure 6.12 the resulting single photon PID efficiencies in dependence of the photons transverse momentum p_T and pseudorapidity η are shown for the 2015 and 2018 dataset respectively.

For the 2015 dataset the efficiency from data is estimated to be larger than in the simulated $\gamma\gamma \rightarrow \gamma\gamma$ sample for smaller transverse photon momenta $p_{T,\gamma}$ and over the full range in η . The difference between the two is below 5 % for $p_{T,\gamma}$ and 7 % for η which is considered not significant. However, the ratio between data and simulation is used as MC-to-data SFs to correct the simulation. For the nominal analysis the SF parametrized by $p_{T,\gamma}$ are used and the SF parametrized by η are taken as a systematic variation to derive uncertainties.

For the 2018 dataset the differences are even smaller with a maximum deviation of 4 % in the transverse photon momentum. The efficiencies in dependence of η show deviations close to 10 % in the very forward region. Due to the low number of FSR photons available in that range the uncertainties are large and hence, this difference is considered acceptable. As for the 2015 dataset the ratios are used as SF for the nominal and systematic variation.

To check that the sample is not contaminated the shower shapes of the selected photon candidates in data are compared to the shower shapes of photons in the simulated signal sample. The comparison plots can be found in the appendix in Figure A.1 and acceptable agreement between data and MC is observed.

Hard-bremsstrahlung photons

The second method is exploiting the $\gamma\gamma \rightarrow e^+e^-$ events with hard-bremsstrahlung already used to determine the photon reconstruction efficiency. As a cross-check, this dataset is also used to determine the photon identification efficiency. In data the hard-bremsstrahlung selection contains a total of 2875 reconstructed photon candidates which are known to correspond to real photons. These can be used to extract the photon identification efficiency by taking the ratio of identified photons over the total number of reconstructed photons. This study is done on data and simulated events with the results shown in Figure 6.13 in dependence of the photon's momentum p_T and pseudorapidity η . No large deviation between data and MC can be seen. As for the FSR sample the comparison of the shower shapes between data and MC shows no significant deviations as can be seen in Figure A.2.




Figure 6.11: Single photon PID efficiency extracted from final-state radiation photons radiated by an electron in dependence of the photon's momentum p_T (left) and pseudorapidity η (right) for the 2015 data taking period. Shown are values for data (black) and simulated events (blue).




Figure 6.12: Single photon PID efficiency extracted from final-state radiation photons radiated by an electron in dependence of the photon's momentum p_T (left) and pseudorapidity η (right) for the 2018 data taking period. Shown are values for data (black) and simulated events (blue).




Figure 6.13: Single photon PID efficiency extracted from hard-bremsstrahlung photons emitted by an electron in dependence of the photon's momentum p_T (left) and pseudorapidity η (right). Shown are values for data (black) and simulated (blue).

6.6 Background estimation

Multiple background processes exist which produce a similar event signature as the light-by-light scattering. Therefore, it is important to identify and quantify the contributions of those processes inside the signal selection. The largest contribution originates in the production of an electron-positron pair ($\gamma\gamma \rightarrow e^+e^-$). These events can fake signal signatures if the electrons are misidentified as photons or in case a hard-bremsstrahlung photon is emitted. Another large contribution comes from CEP $gg \rightarrow \gamma\gamma$ which also produces a di-photon final state. Lastly, fake photons could be a source of background. In Figure 6.14 shows the Feynman diagrams for the relevant background processes.

In the following subsections, each of those contributions will be discussed in detail followed by a short summary of additional processes where a negligible or no background contribution is expected.




Figure 6.14: Feynman diagrams of the light-by-light scattering process (left) as well as the $\gamma\gamma \rightarrow e^+e^-$ (middle) and CEP $gg \rightarrow \gamma\gamma$ (right) processes yielding a relevant amount of background events in the signal region.

6.6.1 Electron-positron background

The largest background process is the decay of the photons into an electron-positron-pair $\gamma\gamma \rightarrow e^+e^-$. This process is much more likely as it has two fewer vertices compared to $\gamma\gamma \rightarrow \gamma\gamma$. At the same time the signal in the detector looks similar to two photons if the charged particle tracks are not properly reconstructed. Hence, the suppression of this background is of large importance to be able to extract good light-by-light scattering events.

The suppression is made possible by vetoing any track activity in the detector. As previously stated two different track definitions are used: standard tracks and pixel tracks. The latter uses much less stringent cuts giving a large efficiency to reconstruct them. Using a veto only on standard tracks the background efficiency decreases to 0.7% compared to no cuts on the number of tracks with the signal efficiency decreasing by 7%. Using the additional requirement of no pixel tracks close to the photon candidates, the background is suppressed by a factor of 10^4 with a loss of signal efficiency of only 1%.

Due to the inaccurate modelling of pixel tracks in the simulation, a data-driven measurement of the background contribution from misidentified $\gamma\gamma \rightarrow e^+e^-$ events is done. This is done using events with two photons with $p_{T,\gamma} > 2.5$ GeV each and an invariant mass of $m_{\gamma\gamma} > 5$ GeV and three control regions defined as following:

1. $CR(N_{\text{track}} = 1)$: one standard track, any number of pixel tracks and $A < 0.01$ for the two photons
2. $CR(N_{\text{pix track}} = 2)$: signal selection except for requiring 2 pixel tracks
3. $CR(N_{\text{pix track}} = 1)$: signal selection except for requiring 1 pixel track

Using the first control region $CR(N_{\text{track}} = 1)$ the conditional probability p_{mistag} is derived to miss a pixel track if the corresponding electron standard track is not reconstructed. The standard track needs to match one of the electrons with $\Delta R < 0.5$. The same requirement is imposed for the pixel tracks, however, in the case of two pixel tracks they must not match the same photon. The mistag probability can then be extracted by taking the ratio of events with one pixel track over the total number of events in the control region.

In Figure 6.15 the number of pixel tracks for events in $CR(N_{\text{track}} = 1)$ are shown for the signal region ($A < 0.01$) and for the CEP control region ($A > 0.01$). It can be seen that the simulation differs quite strongly from the measured results with a tendency to a higher number of reconstructed pixel tracks. This is caused by the mismodelling of the pixel track reconstruction efficiencies and shows that the fake electron background can not be extracted from simulation but must be determined using a data-driven method. The mistag probability p_{mistag} is calculated to be 47.55 %. To assess the systematic uncertainty on p_{mistag} the cut on the acoplanarity A is dropped for $CR(N_{\text{track}} = 1)$ leading to a change of 9 %.




Figure 6.15: Number of reconstructed pixel tracks in events with two photons and exactly one track matching one of the photons in the signal region (left) and the CEP control region (right). Data (black points) is compared with results from the $\gamma\gamma \rightarrow e^+e^-$ simulation (blue).

Knowing the mistag probability p_{mistag} the probabilities for different event compositions can be calculated:

$$p_{N_{\text{pix}}=2} = (1 - p_{\text{mistag}})^2 \quad (6.3)$$

$$p_{N_{\text{pix}}=1} = 2 \cdot p_{\text{mistag}} \cdot (1 - p_{\text{mistag}}) \quad (6.4)$$

$$p_{N_{\text{pix}}=0} = (p_{\text{mistag}})^2. \quad (6.5)$$

Table 6.4: Results for the data-driven e^+e^- background study. The extrapolated expected e^+e^- background contribution is shown in the last row.

	Data nominal	p_{mistag} variation	MC
p_{mistag}	47.55 %	38.25 %	11.69 %
$p_{N_{\text{pix}}=2}$	27.51 %	38.13 %	77.99 %
$p_{N_{\text{pix}}=1}$	49.88 %	47.24 %	20.64 %
$p_{N_{\text{pix}}=0}$	22.61 %	14.63 %	1.37 %
$\text{CR}(N_{\text{track}} = 1)$	36	36	5.36
$\text{CR}(N_{\text{track}} = 2)$	15	15	27.86
$N_{\text{expected,low}}^{\gamma\gamma \rightarrow e^+e^-}$	12.33		0.46
$N_{\text{expected,high}}^{\gamma\gamma \rightarrow e^+e^-}$	16.32		0.46
$N_{\text{expected}}^{\gamma\gamma \rightarrow e^+e^-}$	14.90	8.74	0.46

Combining those probabilities with the measured numbers of pixel tracks in the control regions the expected number of background events inside the signal region can be extrapolated. This is performed taking the average of the two control regions with at least one pixel track as the central value and the following formula:

$$N_{\text{expected}}^{\gamma\gamma \rightarrow e^+e^-} = \left(N_{\text{CR}(N_{\text{pix}}=1)} + N_{\text{CR}(N_{\text{pix}}=2)} \right) \cdot \frac{p_{N_{\text{pix}}=0}}{1 - p_{N_{\text{pix}}=0}}. \quad (6.6)$$

The lower and upper bounds are given by the following relations:

$$N_{\text{expected,low}}^{\gamma\gamma \rightarrow e^+e^-} = N_{\text{CR}(N_{\text{pix}}=1)} \cdot \frac{p_{N_{\text{pix}}=0}}{p_{N_{\text{pix}}=1}} \quad (6.7)$$

$$N_{\text{expected,high}}^{\gamma\gamma \rightarrow e^+e^-} = N_{\text{CR}(N_{\text{pix}}=2)} \cdot \frac{p_{N_{\text{pix}}=0}}{p_{N_{\text{pix}}=2}}. \quad (6.8)$$

In Table 6.4 a full summary of the extracted probabilities, event counts and expected background contribution is given. Based on these numbers the expected number of e^+e^- background events in the signal region is estimated to be $N^{\gamma\gamma \rightarrow e^+e^-} = 15 \pm 7$. The uncertainty covers both the uncertainty on p_{mistag} and the limited statistics in the two pixel track control regions. The lower and upper bounds for the expected background contributions are also covered within the uncertainty.

The same procedure is repeated with an inverted acoplanarity cut (hence, requiring $A > 0.01$) to determine the e^+e^- background in the CEP control region. From this the number of $\gamma\gamma \rightarrow e^+e^-$ events inside this control region is estimated to be $N^{\gamma\gamma \rightarrow e^+e^-} = 30 \pm 14$.

6.6.2 Central exclusive production

CEP $gg \rightarrow \gamma\gamma$, as shown in Figure 6.14, has a di-photon final state just as the light-by-light scattering. However, the kinematics of CEP differs from $\gamma\gamma \rightarrow \gamma\gamma$ as the gluons can radiate

off gluons in the initial state. To estimate the contribution of CEP MC simulation generated with **SuperChic 3** is used. Due to limited knowledge on the cross-section for the CEP process, the normalization is extracted from a dedicated CEP enriched control region in data and then extrapolated to the signal region. This control region is defined by taking the light-by-light scattering signal selection except for inverting the acoplanarity cut effectively requiring $A > 0.01$. Using the following condition the normalization inside the control region is determined:

$$N^{\text{data}}(A > 0.01) = N^{gg \rightarrow \gamma\gamma}(A > 0.01) + N^{\gamma\gamma \rightarrow \gamma\gamma}(A > 0.01) + N^{\gamma\gamma ee}(A > 0.01). \quad (6.9)$$

Here N^{data} is the number of events observed in data, $N^{gg \rightarrow \gamma\gamma}$ the number of expected events originating from CEP, $N^{\gamma\gamma \rightarrow \gamma\gamma}$ the number of expected light-by-light signal events and $N^{\gamma\gamma \rightarrow e^+e^-}$ the number of events from e^+e^- as determined by the data-driven method explained in the previous section.

Figure 6.16 shows the di-photon acoplanarity distribution for events passing the signal selection except for the $A < 0.01$ requirement for measured data and the expected background and signal yields. It can be clearly observed that the acoplanarity is steeply falling for the photons in the signal process while it is only slowly falling for the CEP. For $A > 0.01$ the relative yield in background events exceeds the yield in light-by-light scattering events. This motivates a cut on the acoplanarity with $A < 0.01$ to reduce the amount of CEP background inside the signal region.




Figure 6.16: Di-photon acoplanarity distribution for events passing the signal selection except for the $A < 0.01$ requirement for measured data (black) and the expected background and signal yields (lines). Error bars for measured data show the statistical uncertainties while the shaded band gives the the uncertainties on the signal and background predictions, excluding uncertainty in the luminosity.

Using Equation 6.9 the expected background contribution from CEP inside the signal region is calculated to be 12 ± 3 events. The uncertainty includes the limited statistics in the CEP control region as well as uncertainties on the modelling of the CEP process. The uncertainty on the shape of the CEP signal is evaluated by utilizing a second sample generated with **SuperChic 2**

including extra gluon interactions but no absorptive effects. The second sample shows a 20 % difference in the background yield inside the signal region with respect to the nominal sample and hence, is included as an additional systematic.

6.6.3 Fake photon background

Cosmic muons passing the detector and noise in the calorimeter can create fake photon signatures in the calorimeter. If the cosmic muons do not cross the Pixel detector no charged tracks will be reconstructed for the muons. However, the fake signatures will produce clusters with different shower shapes than real photons for both cases. To estimate the background contribution from muons a control region which requires at least one track reconstructed in the muon spectrometer: $N_{MS} > 0$. A dedicated trigger is used which records events only during empty bunches, hence, no lead ions are passing the collision point and all signals in the detector must come either from noise in the detector components or cosmic particles passing through the detector. In total of 8 events are found which pass the empty bunch trigger and contain two photons with the di-photon invariant mass $m_{\gamma\gamma} > 5 \text{ GeV}$. A normalized distribution of the number of tracks in the muon spectrometer for those events is shown in Figure 6.17. From this it can be seen that approximately 60 % of the fake di-photon events are expected to have at least one track in the muon spectrometer. However, when applying the $p_T^{\gamma\gamma}$ cut no events with $N_{MS} > 0$ are found. From this it is concluded that there is no significant background from fake-photons in the signal region. For completeness, it should be noted that this study is only performed on the 2018 dataset as the empty bunches trigger was not in place for the 2015 dataset, but no differences are expected for with respect to the 2018 dataset.

To cross-check this assumption the time information associated to the photon clusters is studied. The calorimeter time measurement in ATLAS is calibrated such that particles travelling from the IP at the speed of light will give a value of 0 ns [97]. The photons from the light-by-light scattering process are falling into this category and therefore, the measured photon cluster time is expected to be around 0 ns. In Figure 6.18 the measured time distribution for the control region with $A > 0.01$ is shown and compared to the expectation from simulation. A good agreement between data and expectation can be observed for the timing information. Next to it the timing information of the events triggered on empty bunches containing two photons passing the selection criteria with a di-photon invariant mass larger than 5 GeV. For the latter the measured time appears to be distributed flat over the range with no clustering around 0 as would be expected for real prompt photons. This further supports the assumption that the background in the signal region due to fake di-photon events.




Figure 6.17: Number of tracks in the muon spectrometer for di-photon events with $m_{\gamma\gamma} > 5 \text{ GeV}$ in empty bunches (left) and for events passing the signal trigger with the additional requirements of $N_{\text{MS}} > 0$ and $N_{\text{track}} = 0$ (right).




Figure 6.18: Distribution of the time information from the cluster associated to the sub-leading photon for events in the control region with $A > 0.01$ (left) and for events triggered on empty bunches (right). For the events in the control region the measured data is compared to the expected distribution from simulation showing good agreement.

6.7 Kinematic distributions

A selection of kinematic distributions is used to compare the measured results with expectation from simulation. Three distributions related to the energy of the process are chosen: the leading photon transverse momentum $p_{T,\gamma}$, the transverse momentum of the di-photon system $p_T^{\gamma\gamma}$ and the mass of the di-photon system $m_{\gamma\gamma}$. The leading photon is defined as the photon with the larger transverse momentum.

To study the angular distribution of the di-photon system, the di-photon rapidity $|y_{\gamma\gamma}|$ and the $|\cos \theta^*|$ are used. The definition for $|\cos \theta^*|$ is as follows:

$$|\cos \theta^*| = \left| \tanh \left(\frac{\Delta y_{\gamma 1, \gamma 2}}{2} \right) \right|, \quad (6.10)$$

where θ^* is the scattering angle of the two photons in their centre-of-mass system given by the polar angle with respect to the beam axis and $y_{\gamma 1, \gamma 2}$ the difference between the rapidities of the two photons. The last distribution studied is the leading photon $|\eta_\gamma|$, which is the pseudo-rapidity of the photon with the larger transverse momentum.

Figure 6.19 shows these kinematic distributions of the measured $\gamma\gamma \rightarrow \gamma\gamma$ scattering event candidates and compares them to the sum of all expected signal and background contributions within the signal region. The dataset contains a total of 97 events where 45 signal events and 27 background events are expected from simulation.

In total, the expectation from simulation describes the measured data well, however, the overall event yield is found to be slightly larger than expected.




Figure 6.19: Measured distributions (black points) for various kinematic variables are compared with the prediction from MC simulation. The systematic uncertainty on the prediction is shown by the grey area.

6.8 Systematic and statistical uncertainties

Any result is worthless without knowledge on its uncertainties. Hence, gaining knowledge on the uncertainties involved in the analysis is a vital part of the process. This means that for each number, simulated sample and analysis method going into the results the uncertainty associated to it has to be determined. On a first level one distinguishes between systematic and statistical uncertainties. The latter are all uncertainties which are introduced by statistical fluctuations due to limited statistics be it in the measured data or simulated MC samples. They can be reduced by increasing the number of simulated events or measuring over a longer period of time. Systematic uncertainties on the other hand can be introduced by many different sources.

The two main categories are experimental and theoretical systematic uncertainties. Experimental systematic uncertainties stem from limited knowledge on various aspects of the experimental setup. This includes the efficiency of the various detector components, the energy calibration and resolution. Theoretical uncertainties arise from the physics modelling for the measured process as well as the possible background processes.

To evaluate the systematic uncertainties in this analysis the recommendations by the ATLAS collaboration are followed. For all systematic uncertainties evaluated with the ATLAS analysis tools an up and down variation is available which shift the parameters used in the various reconstruction or correction tools by one standard deviation in either direction. The analysis of all simulated samples is run once for every systematic uncertainty yielding the full set of distributions for each.

However, due to the unusual energy range of the light-by-light scattering process with respect to the average ATLAS analysis, some non-standard methods for estimating systematic uncertainties are used.

The influence of each systematic variation is determined by calculating the difference of the varied results to the results obtained without any systematic shift applied, also referred to as nominal, in each bin. The total systematic uncertainty is calculated by adding the different contributions quadratically for each bin i , assuming they are uncorrelated:

$$\sigma_{\text{total sys},i} = \sqrt{\sum_{\text{sys}} \sigma_{\text{sys},i}^2}.$$

In the following sections the different sources of systematic uncertainties and the methods used to extract them are explained in detail. They are bundled into two groups of signal systematic uncertainties and background systematic uncertainties, where the latter are needed for the unfolded differential cross-sections. The trigger efficiency, photon angular resolution and background systematics were extracted by other members of the analysis team. Hence, only a short summary of the results is given there.

6.8.1 Systematic uncertainties on the signal

Trigger efficiency

As described in section 6.3.1 the trigger efficiency is derived from $\gamma\gamma \rightarrow e^+e^-$ events. The trigger efficiency is extracted as a function of the sum of the electron transverse energy $E_T = E_T^{\text{cluster } 1} + E_T^{\text{cluster } 2}$. To receive a continuous description of the trigger efficiency the distribution is fitted

with the following function:

$$\frac{\text{erf}\left(\left(E_{\text{T}}^{\text{cluster } 1} + E_{\text{T}}^{\text{cluster } 2} - p_1\right) / p_2\right) + 1}{2},$$

where p_1 and p_2 are the two fit parameters.

By using slightly modified selection criteria an up and down systematic is estimated. The resulting efficiencies for the 2015 and 2018 conditions are shown in Figure 6.3. The relative uncertainty of the trigger efficiency on the signal yield is found to be 5%.

Photon reconstruction and identification efficiency

As explained in subsection 6.4.2 and 6.5, the differences between the photon reconstruction and photon identification efficiency measured in data and in the simulated samples are corrected for by applied scale factors to each simulated event. In the baseline the scale factors are parametrized as a function of the photons transverse momenta. To estimate the systematic uncertainty on the efficiencies the scale factors are extracted as a function of the photons rapidity and the difference between the nominal and varied results are taken as the uncertainty. For the photon reconstruction efficiency the uncertainty is found to be 4% while the uncertainty on the photon identification efficiency is below 1%. As there is no dedicated up and down variation possible here, the systematic uncertainty is considered to be symmetric.

Photon energy scale and energy resolution

To investigate the uncertainty from the photon energy scale and resolution the ATLAS tool from the combined performance group is used. It applies corrections to the measured properties of photons taking into consideration the detector conditions (e.g. high voltage uniformity and gain variations). To evaluate the influence of the uncertainties for the energy scale and resolution the parameters used for the corrections are varied up and down by one standard deviation. The impact on the overall signal yield is $\pm 1\%$ for the up/down variation of the energy scale and 2% for the energy resolution.

Photon angular resolution

The limited knowledge of the photon angular resolution is again extracted using $\gamma\gamma \rightarrow e^+e^-$ events. Using the fact that the angular resolution of the electron tracks is much better than resolution of the cluster the angular resolution of the cluster can be extracted from measuring $|\varphi^{\text{cluster},1} - \varphi^{\text{track},1}| - |\varphi^{\text{cluster},2} - \varphi^{\text{track},2}|$ which is then given by:

$$\sigma_{\varphi^{\text{cluster}}} \approx \frac{|\varphi^{\text{cluster},1} - \varphi^{\text{track},1}| - |\varphi^{\text{cluster},2} - \varphi^{\text{track},2}|}{\sqrt{2}}$$

Using this method it is found that an additional smearing of the resolution is needed in simulation with $\sigma^{\text{cluster}} \approx 0.006$ to describe the measured resolution from data. When applying this extra smearing a difference of 2% is seen which is taken as an additional systematic uncertainty.

Alternative signal simulation

The choice of the event generator used to produce the event input for the production of simulated samples introduces a model dependency into the analysis. To check the impact of the generator on the signal shape the nominal simulation is compared with an alternative sample

generated with **STARlight** 2.0. The samples are found to be in agreement with a relative difference between the two samples of only 1 %.

6.8.2 Systematic uncertainties on the background

Not only the signal yield is prone to systematic uncertainties, but also the background yield for the backgrounds derived in section 6.6. In this section the sources for uncertainties for the $\gamma\gamma \rightarrow e^+e^-$ and CEP background yield are discussed. They are only relevant for the calculation of the differential cross-sections.

Electron-positron background uncertainties

As discussed earlier the e^+e^- background contribution in the signal region is determined from a control region in data leveraging the probability to miss one electron pixel track if the standard track is not reconstructed. The calculation of the e^+e^- background yield is based on the electron mistag probability, the choice of the control region and the statistical uncertainties in the control region.

To evaluate the uncertainty on the mistag probability a cut is dropped from the definition of the control region used to calculate the mistag probability. This results in an uncertainty of 10 % on the mistag probability. The calculation of the total background yield is therefore repeated with the mistag probability varied up and down by 10 % and the difference to the background yield for the nominal mistag probability is taken as the corresponding systematic uncertainty. The variation leads to a change in the background yield of ± 2.75 events or 10.3 % compared to the nominal background expectation. In the following this uncertainty contribution is referred to as “ee mistag”.

The choice of the control region for the extrapolation into the signal region also influences the overall e^+e^- background normalization. Using varied definitions for the control region as shown in equations 6.7 and 6.8 the background contribution is calculated once again yielding an up and down shift in the overall background yield of ± 2.6 events or 9.8 % compared to the nominal value. This uncertainty is later referred to as “ee CR”.

The last uncertainty in the overall e^+e^- background contribution is due to the limited statistics in the two control regions ($CR(N_{\text{pix}} = 1)$ and $CR(N_{\text{pix}} = 2)$). To estimate this uncertainty the statistical uncertainty from the events in the control regions is propagated to the signal region. Again, this is only impacting the overall normalization of the background contribution and leads to a change of the total background contribution of ± 4.1 events or 15.4 %. It is abbreviated with “ee stat” in the following.

In addition to the previous three uncertainty contributions, an uncertainty on the shape of the e^+e^- background must be considered. As no events pass the signal selection in the simulated samples for the $\gamma\gamma \rightarrow e^+e^-$ process, the shape is extracted from data using the control region $CR(N_{\text{pix}=1})$. To estimate the uncertainty on the shape an alternative control region with $CR(N_{\text{pix}=2})$ is used. For the distributions and binnings used in the unfolding procedure the changes in the shape lead to a deviation of at most 25 % in single bins. This uncertainty is denoted as “ee shape” in the following plots.

CEP background uncertainties

While the shape for the CEP background contribution is taken from simulation, the overall normalization was determined using a data-driven method, calculated from a control region

with $A > 0.01$ as shown in Figure 6.16.

The statistical uncertainty from the control region is propagated to the signal region by calculating the CEP background with the number of events varied within the statistical uncertainty, yielding an uncertainty on the normalization. This results in a change of the overall background yield of $\pm 4.6\%$ and is referred to as “CEP CR Stat.”

To evaluate the influence of the generator on the modelling of the $gg \rightarrow \gamma\gamma$ process, another sample is simulated using a different version of the generator (**SuperChic 2** instead of **SuperChic 3**), which uses a different modelling of the process. With this new sample both uncertainties on the shape as well as on the normalization of the CEP background are derived. The change in the normalization is caused by the change of the shape in the acoplanarity distribution leading to a different amount of events in the $CR(A > 0.01)$. The different normalization changes the overall background yield by $\pm 2.5\%$. In the following the shape uncertainty is called “CEP shape” and the normalization “CEP SC2”.

6.8.3 Statistical uncertainties

In addition to the systematic uncertainties, a second source of uncertainties has to be taken into account which are statistical uncertainties. These are a concern both in simulated samples and in data. The method to extract the statistical uncertainties for the unfolding are described in this section.

Data statistics

To determine the statistical uncertainty on data a so-called toy method is used, where 1000 Poisson distributed variations around the measured values are created. For all variations the background is subtracted and the unfolding procedure is performed. The Root Mean Square (RMS) of the differences between the varied and nominal unfolded result is taken as the uncertainty coming from limited statistics of the data.

MC statistics

The uncertainty from the limited statistics of the MC prior is determined in a similar way. This time when creating the histograms for the detector response each event is given an additional randomly generated event weight following a Poisson distribution around 1. Using those histograms the unfolding procedure is executed on the nominal data and the RMS of the relative difference between the varied results and nominal is taken as the statistical uncertainty from the MC prior.

6.8.4 Luminosity uncertainty

The measurement of the integrated luminosity is performed by the ATLAS Luminosity Group as described in section 4.4. The luminosity measurement for the 2015 dataset is finalized [52] while the analysis for the 2018 dataset is still ongoing. A preliminary value for 2018 is used in this analysis. The change of the central value for the 2018 luminosity is expected to be negligible while the systematic uncertainties might potentially be improved slightly.

The preliminary luminosity uncertainty for the combination of the 2015 and 2018 datasets is found to be 3.2% , which is propagated to all results.

6.9 Integrated fiducial cross-section

To measure the inclusive cross-section for the light-by-light scattering process in a fiducial phase space the following formula is used:

$$\sigma_{\text{fid}} = \frac{N_{\text{data}} - N_{\text{bckg}}}{C \times \int L dt}, \quad (6.11)$$

where N_{data} is the number of observed events, N_{bckg} the number of expected background events, $\int L dt$ the integrated luminosity of the dataset and C an overall correction factor to account for the efficiency of the detector.

The fiducial phase space is defined by two photons with a minimum transverse momentum of $p_T > 2.5 \text{ GeV}$ and an absolute pseudorapidity $|\eta_\gamma| < 2.4$. The invariant mass of the di-photon system $m_{\gamma\gamma}$ has to be larger than 5 GeV and the transverse momentum $p_T^{\gamma\gamma}$ must be smaller than 1 GeV. To ensure that the photons are back-to-back the di-photon acoplanarity A must be smaller than 0.01.

The overall correction factor C is determined by taking the ratio of events passing the fiducial selection on detector level over the ones passing on particle level in the signal MC sample. To estimate the uncertainty on the correction factor it is calculated varying the data/MC scale factors within their uncertainties. The correction is determined to $C = 0.263 \pm 0.021$ with the uncertainties listed in Table 6.5.

Table 6.5: Contributions on the uncertainty of the overall detector correction factor C .

Source of uncertainty	Impact on correction
Trigger efficiency	5 %
Photon reconstruction efficiency	4 %
Photon PID efficiency	2 %
Photon energy scale	1 %
Photon energy resolution	2 %
Photon angular resolution	2 %
Alternative signal MC	1 %
Signal MC statistics	1 %
Total	8 %

The combined 15+18 dataset contains $N_{\text{data}} = 97$ event candidates with an expected background contribution of $N_{\text{bckg}} = 27 \pm 5$. Together with the integrated luminosity $\int L dt = 2.20 \pm 0.07 \text{ nb}^{-1}$ this yields a fiducial cross-section of $\sigma_{\text{fid}} = 120 \pm 17 \text{ (stat.)} \pm 13 \text{ (syst.)} \pm 4 \text{ (lumi.) nb}$. The theoretical predictions give $80 \pm 8 \text{ nb}$ [83] and $78 \pm 8 \text{ nb}$ [71]. The ratio between the measured fiducial cross-section and the theoretical prediction is 1.50 ± 0.31 and 1.54 ± 0.32 . The excess can not be explained with the current model, however, it is not statistically significant.

6.10 Differential fiducial cross-sections

Measuring an observable always yields a result that is distorted by limited detector resolution. This is caused by e.g. imperfections or noise induced signals in the detector. To compensate for detector effects a detailed model of the detector, as implemented in the event simulation (c.f. section 5.7), is used in a procedure called unfolding to correct for the detector response.

In other words, the goal is to obtain the particle level distribution from the detector level distribution. Particle level denotes the final state of the event before any interaction with the detector while detector level refers to the distribution as measured by the detector.

In this analysis two different methods are used to derive the particle distributions from the measured distributions. One method – called bin-by-bin unfolding – is used as a simple baseline and cross-check, while the second more advanced Bayesian unfolding is used to derive the final unfolded distributions.

The explanation of the unfolding methods refers to various distributions by the following notation:

- **d**: the measured data distribution on detector level,
- **b**: the expected background contribution on detector level,
- **s**: the background free signal distribution on detector level,
- **x**: the unfolded (particle level) distribution.

All distributions are measured in n bins and the content of a bin i is referred to by adding the index i , e.g. d_i for bin i of the measured data distribution.

6.10.1 Binning of distributions

Before performing the unfolding, the binning of the distribution has to be fixed taking into account the following aspects. The largest limiting factor in this analysis is the low number of events observed in data, which allows for only a small number of bins. Secondly, the purity of the distributions is considered with aiming for at least $\approx 80\%$ in each bin. It is defined as the fraction of events reconstructed in the same bin i that originate in the same bin on particle level:

$p_i = \frac{N_{(\text{detector} \rightarrow \text{particle})i}}{N_{\text{detector},i}}$. Lastly, each bin must contain a number of events both on data and MC. For MC, each bin should contain at minimum around 1000 events. Due to the last criterion the bin size is not constant but increased at the edges (e.g. at higher transverse photon momentum).

The resulting binning is given in Table 6.6 and the corresponding purities are shown in Figure 6.20. It can be seen that the purity is over 95 % for the absolute di-photon rapidity and $|\cos\theta^*|$. For the di-photon invariant mass and the average photon momentum it is at least 84 % and 73 % respectively.


6.10.2 Bin-by-bin unfolding

The simplest method of unfolding is called bin-by-bin unfolding. It relies on the same principle as used in the calculation of the fiducial cross-section but using multiple bins. In fact, one could think of the fiducial cross-section calculation as a bin-by-bin unfolding in a single bin. Hence, instead of a single correction factor, bin-by-bin unfolding uses one correction factor per bin.

Same as for the fiducial cross-section calculation the first step is to subtract the expected background **b** from the measured distribution **d** in each bin to calculate the background free

Table 6.6: Overview of the binnings used in the unfolding procedure.

variable	bin edges
absolute di-photon rapidity	0, 0.6, 1.2, 1.8, 2.4
invariant mass $m_{\gamma\gamma}$ [GeV]	5, 10, 15, 20, 30
average photon transverse momentum $\langle p_{T,\gamma} \rangle$ [GeV]	2.5, 3.5, 4.5, 6, 8, 30
$ \cos \theta^* $	0, 0.25, 0.5, 0.75, 1

**Figure 6.20:** Purity as a function of di-photon rapidity (upper left), di-photon invariant mass (upper right), average photon transverse momentum (lower left) and $|\cos \theta^*|$ (lower right).

signal distribution $s_i = d_i - b_i$.

In the next step, each bin is multiplied with a correction factor a_i yielding the unfolded distribution x_i . The correction factors are calculated as the ratio between the particle level and detector level distributions in the simulation. With N_i^{particle} being the number of entries in bin i on truth level and N_i^{detector} the number of entries in bin i on detector level the corrections are given by

$$a_i = \frac{N_i^{\text{particle, fid.}}}{N_i^{\text{detector, fid.}}} \quad (6.12)$$

Putting all together allows calculating the unfolded distribution x_i with

$$x_i = a_i \cdot s_i = \frac{N_i^{\text{particle, fid.}}}{N_i^{\text{detector, fid.}}} \cdot (d_i - b_i). \quad (6.13)$$

The distributions unfolded using the bin-by-bin method are shown later in a comparison of the results produced by the two methods.

6.10.3 Methodology of Bayesian unfolding

In the previous section it was shown how the unfolding can be performed using the bin-by-bin method. However, as bin-by-bin unfolding treats every bin separately, it does not account for migrations between bins. To take into account those migrations a different approach must be chosen, which is known as Bayesian unfolding.

While for bin-by-bin unfolding a single correction factor for each bin is sufficient, Bayesian unfolding requires the corrections to be performed in two steps.

The fiducial corrections f_i correct for events that are inside the fiducial volume on detector level but are outside the fiducial volume on particle level. These events are often referred to as fakes. In this analysis a fake event could for example be induced by noise in the calorimeter resulting in a photon being reconstructed. The factors are calculated as the ratio of events in bin i within the fiducial volume on particle and on detector level $N_i^{\text{particle \& detector}}$ and all events inside the fiducial volume on detector level N_i^{detector} :

$$f_i = \frac{N_i^{\text{particle \& detector}}}{N_i^{\text{detector}}}. \quad (6.14)$$

The efficiency corrections c_i correct for inefficiencies of the detector, reconstruction and analysis selection. It is determined from events that are missed on detector level but are present on particle level, which are often called misses. The factors are given by the ratio of events within the fiducial volume on detector level N_i^{detector} and the events within the fiducial volume on particle level N_i^{particle} . The corrections are calculated independently for every bin i which is defined by the particle level value of the variable.

$$c_i = \frac{N_i^{\text{detector, fid.}}}{N_i^{\text{particle, fid.}}}. \quad (6.15)$$

The basic idea of Bayesian unfolding is to find the probability that an event was generated in bin i of the distribution on particle level under the condition that it is reconstructed in bin j of the measured distribution. This probability is written as $P(x_i|s_j)$ and using Bayes' theorem can also be expressed by:

$$P(x_i|s_j) = \frac{P(s_j|x_i)P_0(x_i)}{P(s_j)} \quad (6.16)$$

$$= \frac{P(s_j|x_i)P_0(x_i)}{\sum_i P(s_j|x_i)P_0(x_i)}, \quad (6.17)$$

with $P(s_j|x_i)$ being the conditional probability that an event generated in bin i is reconstructed in bin j and $P_0(x_i)$ the initial probability of the event being generated in bin i . From the previously discussed probability the so-called response matrix $R_{ji} = P(x_i|s_j)$ can be constructed and is estimated from MC simulation. The response matrix describes all correlations between all

bins on detector and particle level is dependent on the a-priori probability $P_0(x_i)$ and hence, a strong model dependency is introduced in the unfolding procedure. This can be mitigated by unfolding in multiple iterations and thereby, reducing the dependency on the prior at the cost of an increase in the uncertainties caused by statistical bin-by-bin fluctuations.

The best estimate of the true number of events in bin i on particle level for a measured distribution $\mathbf{s}^{\text{data}}$ is given by [98]:

$$x_i^{\text{data}} = \frac{1}{c_i} \prod_{j=1}^n \left(f_j \cdot s_i^{\text{data}} \right) \cdot P(x_i | s_j). \quad (6.18)$$

Here f_j and c_i are the fiducial and efficiency corrections defined earlier in Equation 6.14 and Equation 6.15.

For the practical implementation of the unfolding the calculation is split into multiple steps as depicted schematically in Figure 6.21. Starting with the measured data distribution $\mathbf{d}$ the first step is to subtract the expected background contribution as derived in section 6.6. Afterwards fiducial correction factors are applied, which gives the signal distribution $\mathbf{s}$. Written as a mathematical expression this is given by:

$$s_i = (d_i - b_i) \cdot f_i.$$

In the next step the bin migrations are corrected by applying the unfolding on the signal $\mathbf{s}$ resulting in a new distribution $\mathbf{y}$. This distribution is then corrected for limited acceptance and inefficiencies using efficiency corrections c_i finally yielding the efficiency corrected unfolded distribution $\mathbf{x}$.

$$x_i = y_i / c_i.$$

6.10.4 MC prior

The nominal light-by-light scattering signal simulation generated with **SuperChic 3** is used as the MC prior for the unfolding. The simulation is done separately for the 2015 and 2018 data-taking campaign and preprocessed using the according data-to-MC scale factors. They are merged taking the measured luminosities for the periods into account. The truth level selection is done using the signal selection as given in subsection 6.3.2 except for allowing the truth photons to be within the calorimeter crack ($1.37 \leq |\eta_\gamma| < 1.52$). For reconstruction level the signal selection is used without any changes.

6.10.5 Fiducial and efficiency corrections and response matrices

Using the definitions from subsection 6.10.3 the fiducial and efficiency corrections are calculated. The resulting correction factors are shown in Figure 6.22. The fiducial corrections are flat for all variables except for the average photon p_T where a drop is visible at lowest energies. The influence of systematically varied prior is studied by calculating the correction factors for all variations. The total systematic uncertainty is calculated by quadratically summing the differences between the result with varied and the nominal correction factors. The systematic uncertainty on the fiducial corrections is found to be at most 2.2 % for all variables.

The efficiency corrections are relatively flat for the di-photon rapidity and di-photon $|\cos \theta^*|$ with a slight drop for higher $|\eta_\gamma|$ in the di-photon rapidity. The average photon p_T and di-




Figure 6.21: Schematic showing the basic principle of the Bayesian unfolding procedure. The numbers shown in the drawing are the numbers from the unfolding of the average photon momentum.

photon invariant mass however, show a strong turn on behaviour. This turn on behaviour can be explained by the low trigger efficiency and photon reconstruction efficiency at low energies as presented in subsection 6.3.1 and subsection 6.4.2. The systematic uncertainties are again extracted by calculating the factors for all the variations. The largest impact can be seen from the trigger systematics with a maximum deviation of 11 % in the average photon p_T and 8 % in the di-photon invariant mass. For all other systematic variations the maximum deviation is in the order of 6 %.

The detector response matrices are built using events from the signal MC sample which pass both on truth level and on reconstruction level. The resulting matrices are shown in Figure 6.23. It can be seen that the detector response is not having a large effect on the reconstructed quantities for the absolute di-photon rapidity and the di-photon $|\cos\theta^*|$ as most events land on the diagonal of the matrix. For the average photon p_T and the di-photon invariant mass however, there is bin migration visible especially at low p_T .

6.10.6 Closure studies

An important step in the unfolding is to verify the validity of the unfolding procedure. The main method to ensure the validity is to test the unfolding procedure on a dataset where the distributions are known both on detector level and on particle level. This is the case for MC samples and hence, the first quick check is to unfold the same MC sample used to derive the correction factors and the response matrix and compare it with the particle level distribution of the sample. The unfolded distribution should exactly match the particle level distribution after a single iteration. This check is done for the unfolding and passed with no deviations.




Figure 6.22: Efficiency and fiducial corrections as applied in the unfolding procedure as a function of di-photon rapidity (upper left), di-photon invariant mass (upper right), average photon transverse momentum (lower left) and $|\cos \theta^*|$ (lower right). The statistical uncertainties are given as bars while systematic uncertainties are shown as bands.




Figure 6.23: Detector response matrices as a function of di-photon rapidity (upper left), di-photon invariant mass (upper right), average photon transverse momentum (lower left) and $|\cos \theta^*|$ (lower right). The numbers represent the corresponding bin values.

MC non-closure

The next step is to use a different MC sample as fake data. Typically, this is done by splitting the signal sample in two and using one half to create the unfolding and the second half to confirm it. For this analysis a second, statistically independent MC sample is available generated with a different generator (**SuperChic 3.05** instead of **SuperChic 3**). Both samples agree within statistical precision as shown in Figure A.3. The difference between the known particle level distribution and the unfolded distribution is called the non-closure. This method is also used to determine if the number of iterations chosen for the unfolding, which in this case is a single iteration, is sufficient by determining the non-closure for different numbers of iterations. This is shown in Figure 6.24 for the di-photon invariant mass $m_{\gamma\gamma}$. The results for the other distributions can be found in Figure A.4 in the appendix. It can be seen that the non-closure is already small after one iteration and at the same level as the statistical uncertainty. This implies that part of the non-closure is caused by statistical artefacts and leads to an increase of the non-closure with more iterations. At the same time the statistical uncertainties rise with the number of iterations as expected. The uncertainties are combined by adding them quadratically and as shown is the smallest for a single iteration. The non-closure is used as an uncertainty on the unfolding procedure and listed as MC non-closure uncertainty.

Data-driven non-closure

The data-driven closure test is following the same scheme as the MC closure test. The difference is that the MC is modified to create a toy dataset that looks similar to the measured data by reweighting the sample. Applying the same scale factors on both detector and particle level results in a sample resembling the measured data on detector level for which the particle level distribution is known as well. To extract those scale factors the measured distributions are normalized and divided by the normalized detector level distributions from simulation. A smooth function is fitted to these scale factors and used for reweighting in dependence of the particle level quantity. Due to the low number of bins in this analysis a simple first order polynomial is used for the fit.

In this data-driven method the fiducial and efficiency corrections are skipped as the fake data is constructed from events which pass the selection both on detector and particle level. Hence, only the unfolding matrix from the nominal unfolding is applied to perform the unfolding.

Figure 6.25 shows the plots related to the data-driven closure test for the di-photon invariant mass




Figure 6.24: Relative non-closure, statistical and combined uncertainty from the MC non-closure test for the di-photon invariant mass for different numbers of iterations.

mass starting with the scale factors used for reweighting along with the fit used to determine the factors. Also shown are the model uncertainty calculated by the difference between the distribution after applying the unfolding matrix and the reweighted particle level distribution. The third plot shows the uncertainties yielding from the limited statistics of the MC sample which are much larger than the model uncertainties and the last plot shows the combination of both uncertainties. As for the MC based non-closure test, the model uncertainty is small compared to the statistical uncertainties. This explains why they are not drastically decreasing with higher numbers of iterations. The same can be observed for the other distributions as shown in Figure A.4.

The difference between the unfolded distribution and the reweighted particle level distribution is taken as an additional systematic uncertainty under the name *model uncertainty*.

As an additional cross-check, the procedure is repeated with the second (SuperChic 3.05) sample resulting in consistent results (refer to Figure A.5).




Figure 6.25: Scale factors with fit used for reweighting (upper left), non-closure from comparing unfolded and particle level distribution (upper right), statistical uncertainty from the limited MC statistics (lower left) and the combination of non-closure and statistical uncertainty (lower right) are shown for the data-driven closure test for the di-photon invariant mass.

Influence of fiducial correction factors

As the distributions of the di-photon invariant mass $m_{\gamma\gamma}$ and the average photon $\langle p_{T,\gamma} \rangle$ are steeply falling, the influence of the fiducial corrections on the unfolded result is studied in more detail to ensure that this behaviour is not leading to large changes at the lower edge of the fiducial volume.




Figure 6.26: Comparison between the unfolded results with nominal (line) and varied (points) fiducial corrections for the di-photon invariant mass $m_{\gamma\gamma}$ (left) and the average photon momentum $\langle p_{T,\gamma} \rangle$ (right). The lower half of the plots give the ratio of the varied over the nominal results.

To estimate the sensitivity of $m_{\gamma\gamma}$ and $\langle p_{T,\gamma} \rangle$ shapes around the fiducial region boundary, first the alternative signal sample is reweighted to data using scale factors extrapolated to outside the fiducial region by the same method as described for the data-driven closure test, except that this time also events outside fiducial region are included in the reweighting procedure.

In the second step the fiducial corrections for this reweighted sample are calculated. Using those varied fiducial corrections together with the unfolding matrix and efficiency corrections from the nominal unfolding, the data is unfolded and compared to the results from the nominal unfolding. The results from this study are shown in Figure 6.26. It can be seen that the influence of the modified fiducial correction factors is small, being below 0.5 % and 0.3 % for the di-photon invariant mass and the average photon p_T respectively, and only present in the first bin.

6.10.7 Comparison of the two unfolding methods

Figure 6.27 shows a comparison between the differential cross-sections for all measured variables along with the ratio of the values for bin-by-bin over Bayesian. As can be seen the differences between the two methods are small in most bins with the relative deviation being below 5 % for all distributions except for the di-photon rapidity where it is up to 10.5 % in one bin.

The good agreement confirms that the Bayesian unfolding works as expected and no major mistakes are made in the unfolding procedure.




Figure 6.27: Comparison of the differential cross-sections as a function of absolute di-photon rapidity (upper left), di-photon invariant mass (upper right), average photon transverse momentum (lower left) and absolute di-photon $|\cos \theta^*|$ (lower right) as determined by Bayesian (line) and bin-by-bin (points) unfolding. The ratio in the bottom of each plot shows the ratio of the values for bin-by-bin over Bayesian. In this comparison only statistical uncertainties are shown. It should be noted, that the statistical uncertainties are highly correlated.

6.10.8 Results

In this section the differential fiducial cross-sections unfolded from detector to particle level using the Bayesian unfolding procedure as a function of the absolute di-photon rapidity, the di-photon invariant mass, average photon p_T and the absolute di-photon $|\cos\theta^*|$ are presented. The unfolded distributions are shown in Figure 6.28 with all uncertainties, compared to the prediction from **SuperChic 3.0**.

The measured differential cross-sections show fair agreement with the prediction except for differences in the overall normalization. The shapes however, are in good agreement with the prediction. This can be seen clearly in the normalized differential cross-sections shown in Figure A.6. A global χ^2 comparison of data and prediction confirms the good agreement, yielding $\chi^2/\text{degrees of freedom}$ of at most 5.8/4 for the $|\cos\theta^*|$ distribution.

The total uncertainties for all variables are dominated by statistical uncertainties ranging from 25 % to 75 %. Given the low number of measured event candidates this is expected. The systematic uncertainties are dominated mostly by the background uncertainties with up to 45 % in the upper bins of the di-photon rapidity and up to 20 % in the other distributions. This can be explained by limited statistics within the control regions.

The theoretical uncertainty of the prediction are mainly caused by the limited knowledge of the nuclear form-factors and the related initial photon fluxes which is discussed in [88]. From this a 10 % flat uncertainty is estimated as the theoretical uncertainty.




Figure 6.28: Measured differential cross-sections as a function of absolute di-photon rapidity (upper left), di-photon invariant mass (upper right), average photon transverse momentum (lower left) and absolute di-photon $|\cos\theta^*|$ (lower right). The measured cross-section are shown as points with error bars giving the statistical uncertainty. The grey band is denoting the combination of systematic and statistical uncertainties. The measurement is compared with the prediction from **SuperChic 3.0** (red line) and the corresponding band giving the theoretical uncertainty on the prediction.

6.11 Limits on the production of axion-like-particles

Any particle created in a photon induced s-channel process which is decaying into two photons would create a resonance peak in the di-photon invariant mass $m_{\gamma\gamma}$ spectrum. One prominent particle postulated to couple to photons are ALPs which would produce a narrow resonance [79]. To search for any resonance consistent with ALPs a multi-bin fit is performed on the di-photon invariant mass spectrum at various mass points. The expected ALP signal shape for that is determined by simulation using **STARlight 2.0** for various axion masses in the range $5 \text{ GeV} \leq m_a \leq 100 \text{ GeV}$. The background for this process is composed by the same backgrounds as for light-by-light scattering, however, the light-by-light scattering process is now being considered a background as well. No interference effects have to be taken into consideration as long as the resonance is narrow enough which is the case for ALPs as detailed in [99].

Signal samples

The samples are generated for masses in steps of 1 GeV for $m_a < 30 \text{ GeV}$ and in steps of 10 GeV for higher invariant masses. In Figure 6.29 the expected di-photon invariant mass $m_{\gamma\gamma}$ on detector level after applying the full event selection is shown for a selection of samples together with a Gaussian fit for each point to better visualize the expected shape. For each mass point 10000 events have been generated.




Figure 6.29: Di-photon invariant mass $m_{\gamma\gamma}$ distribution on detector level for a selection of the $\gamma\gamma \rightarrow a \rightarrow \gamma\gamma$ samples generated with **SuperChic 3** after applying the full event selection. For each mass point a Gaussian is fitted to the distribution to better show the expected shape.

As a cross-check for the **STARlight 2.0** samples an additional set of samples is generated using **SuperChic 3.06**. Both were found to be in good agreement within the statistical uncertainties for all mass points. This is shown exemplary in Figure 6.30 for the samples generated with an axion mass of $m_a = 8 \text{ GeV}$ for the di-photon mass $m_{\gamma\gamma}$ and the di-photon transverse momentum $p_T^{\gamma\gamma}$.




Figure 6.30: Comparison of the nominal $\gamma\gamma \rightarrow a \rightarrow \gamma\gamma$ sample at $m_a = 8$ GeV generated using **STARlight 2.0** with the alternative sample generated using **SuperChic 3.06**. Shown is the di-photon invariant mass $m_{\gamma\gamma}$ (left) and the di-photon transverse momentum $p_T^{\gamma\gamma}$ (right). Both show good agreement within statistical uncertainties.

Systematic uncertainties on the signal

The systematic uncertainties for the ALP signal are determined identically to the procedure explained in section 6.8. The total systematic uncertainty in the central peak region is around 10%. Given the narrow resonance, in the tails of the resonance the total relative uncertainty rises to up to 90% due to the variation of the energy scale and resolution.

Setup of the multi-bin fit

The extraction of the upper limits is performed using a dedicated framework for multi-bin histogram-based analysis programmed in python called **pyhf** [100]. By using a multi-bin analysis the calculation of the upper limit is taking the shape of the signal into account. The parameter of interest is the signal strength μ which is expressing what fraction of the signal sample can at most fit into the measurement. Nuisance parameters are included accounting for the various uncertainties. The two samples in the fit are the ALP signal and the SM background. As the overall normalization of the expected SM light-by-light scattering does not match the data perfectly, different setups are tested to evaluate the influence of the background normalization on the limits. In the nominal setup the background contribution from light-by-light scattering is set floating by adding an overall normalization parameter for this background. The background contributions from the CEP and e^+e^- processes are fixed as their normalization is determined from the measured data. As a cross-check the fit is repeated without the additional normalization parameter. To study the influence of the non-smooth shapes of the CEP and e^+e^- backgrounds (c.f. Figure 6.19) the fit is repeated with smoothing applied to these two backgrounds.

To reflect the uncertainty on the luminosity a flat 3.2% uncertainty is added. The systematic uncertainties on the signal are evaluated as previously discussed in section 6.8, but as the sources are influencing the signal and background in the same way they are not included in the nominal fit. To cover for the limited knowledge of the initial photon fluxes [88] an additional systematic of 10% is added uncorrelated with the other systematics.

The statistical uncertainties are taken to be flat again with a value of 10% for both signal and

background.

The fit calculates the upper limit on the signal strength μ with a 95 % confidence interval. This is then converted into an upper limit on the $\gamma\gamma \rightarrow a \rightarrow \gamma\gamma$ cross-section $\sigma_{\gamma\gamma \rightarrow a \rightarrow \gamma\gamma}$ with:

$$\sigma_{\gamma\gamma \rightarrow a \rightarrow \gamma\gamma} = \mu \cdot \sigma_{\text{gen}}, \quad (6.19)$$

where σ_{gen} is the generator cross-section for the Axion signal samples. In addition, the coupling of the axion to two photons can be calculated using

$$\frac{1}{\Lambda_a} = \sqrt{\mu} \cdot \frac{1}{\Lambda_a^{\text{gen}}}, \quad (6.20)$$

where Λ_a^{gen} is the coupling used in the generator.

The expected limit is found by testing the background only hypothesis.

The observed and expected upper limits are determined for each mass bin independently yielding an upper limit for the full mass range. As no events are observed above $m_{\gamma\gamma} = 30$ GeV, the limit in that range is calculated in steps of 10 GeV. The resulting observed and expected 95 % confidence level limits on the cross-section $\sigma_{\gamma\gamma \rightarrow a \rightarrow \gamma\gamma}$ and axion to photon coupling are shown in Figure 6.31.




Figure 6.31: 95 % confidence level upper limit on the $\gamma\gamma \rightarrow a \rightarrow \gamma\gamma$ cross-section $\sigma_{\gamma\gamma \rightarrow a \rightarrow \gamma\gamma}$ (left) and on the axion to photon coupling (right). The expected limits (dashed lines) and the $\pm 1\sigma$ and $\pm 2\sigma$ uncertainty bands are shown in addition.

Cross-checks

Comparing the resulting limits to the initial analysis published in [3] the limits presented here are compatible but slightly more stringent. This is expected as including the shape information gives more information for the upper limit calculated. To further verify the procedure and understand the impact of the fit setup, multiple cross-checks are presented in the following section.

Bayesian limits

As a first test, a Bayesian method for calculating the confidence intervals on the signal strength in a signal plus background to data comparison is used. The method is explained in detail in [101].

The observed value n is composed of $n = n_s + n_b$, where n_s and n_b are the numbers for signal and background events. Both are Poisson variables with means v_s and v_b resulting in n to be a Poisson variable as well.

One can construct the likelihood as a function of v_s for a uniform prior density as

$$L(n_{\text{obs}}|v_s) = \frac{(v_s + v_b)^{n_{\text{obs}}}}{n_{\text{obs}}!} e^{-(v_s + v_b)}. \quad (6.21)$$

Using Bayes theorem one can find the upper limit for the number of signal events v_s^{up} :

$$\beta = \frac{\int_0^{v_s^{\text{up}}} L(n_{\text{obs}}|v_s) dv_s}{\int_0^{\infty} L(n_{\text{obs}}|v_s) dv_s} - 1, \quad (6.22)$$

which, after integration by parts n_{obs} times, becomes:

$$\beta = \frac{e^{-(v_s^{\text{up}} + v_b)} \sum_{n=0}^{n_{\text{obs}}} \frac{(v_s^{\text{up}} + v_b)^n}{n!}}{e^{-v_b} \sum_{n=0}^{n_{\text{obs}}} \frac{v_b^n}{n!}}. \quad (6.23)$$

This equation can be solved numerically with for v_s^{up} , yielding for $\beta = 0.05$ the 95 % upper limit on the number of signal events compatible with the observed events. To solve the equation a minimizer from the python library `scipy` is utilized. As n_{obs} , n_s and n_b are plain numbers no shape information is included in this calculation. Therefore, to use the method for the limit extraction the bins containing the resonance peak are summed, and the sum is used for n_{obs} and n_b . The number of bins is chosen to contain at least $\approx 90\%$ of the resonance peak inside the selected bins. This results in 3 GeV wide bins for up to 30 GeV and 10 GeV bins above that. In Figure 6.32 the limits calculated with the Bayes method are compared to the limits from the previous section. It can be seen that those limits are slightly less stringent than the limits from the fit in the previous section. This is expected due to not including any shape information in the Bayes method. As another disadvantage of this method, it does not include any systematic uncertainties. While this makes it inferior to the nominal method, it confirms the nominal results shown before.

Influence of the background normalization and shape

In order to understand the influence from the background normalization and shape on the upper limit calculation two variations of the upper limit fit are considered.

To investigate the influence of the floating light-by-light background compensating for the mismatch of the overall normalization between the theoretical expected and observed event numbers in the light-by-light scattering measurements the upper limit extraction with `pyhf` is repeated with all backgrounds fixed.

Due to the low available statistics for the $\gamma\gamma \rightarrow e^+e^-$ and CEP background shape extraction

from data it is far from smooth, as can be seen in Figure 6.19. Hence, the limit calculation is repeated with a smoothed background shape. To smooth the background shape a sliding window averaging with a width of 3 bins is applied on the histogram and the resulting histogram scaled to keep the number of events constant. During the limit calculation with the smoothed background an additional systematic is added with the differences between the smoothed and nominal background histograms.

In Figure 6.33 the resulting limits for both background variations are shown. It can be seen that the smoothed background has no influence on the limits while the fixed background leads to slightly weaker limits with the largest deviation being 30 %.




Figure 6.32: Comparison of the 95 % confidence level upper limit on the $\gamma\gamma \rightarrow a \rightarrow \gamma\gamma$ cross-section $\sigma_{\gamma\gamma \rightarrow a \rightarrow \gamma\gamma}$ extracted using the Bayes method (solid red) with the results from the nominal method (solid black). The expected limits (dashed line) and uncertainty bands are from the nominal results. A good agreement between both is seen.




Figure 6.33: Comparison of the 95 % confidence level upper limit on the $\gamma\gamma \rightarrow a \rightarrow \gamma\gamma$ cross-section $\sigma_{\gamma\gamma \rightarrow a \rightarrow \gamma\gamma}$ extracted with fixed backgrounds (left) and with a smoothed background (right) to the results from the nominal method (solid red).

Exclusion

Figure 6.34 puts the derived constraints on the ALP mass and coupling to photons into context by comparing them to constraints from other experiments [81, 102–105]. This plot is made under the assumption of a branching fraction for ALP into photons of 100 %. In the exclusion plot it is obvious that for a mass range of $6 \text{ GeV} \leq m_a \leq 100 \text{ GeV}$ the limits presented in this thesis are the strongest limits as of today.




Figure 6.34: Exclusion limits at 95% CL on the ALP to photon coupling $1/\Lambda_a$ versus the ALP mass m_a obtained by different experiments. The existing limits are derived from [81, 102–105]. The limit labelled “LHC (pp)” are based on proton-proton data from ATLAS and CMS. For all measurements a branching fraction into photons of 100 % is assumed.




Figure 6.35: Event display for a light-by-light scattering candidate event showing a rendering of the ATLAS detector. The two photons are clearly visible with very low activity in the inner detector.

6.12 Conclusion

In this thesis the first direct observation and measurement of light-by-light scattering is presented with a total of 97 event candidates where 45 signal and 27 background events are expected. The background contribution is calculated using data-driven techniques for all relevant background processes. The measurement strategy is improved significantly with respect to the strategy used in the previous measurement which resulted in an evidence for light-by-light scattering in 2017. This is achieved by dedicated selection methods for tracks to better suppress the $\gamma\gamma \rightarrow e^+e^-$ background and improving the photon performance. The custom photon identification method vastly improves the performance at lowest energies. This new identification method is based on neural networks trained with events from simulation and from a background enriched region in data. Compared to the standard loose photon identification from ATLAS the single photon identification efficiency is improved by 60 % at the lowest photon energies.

The fiducial cross-section as well as differential cross-sections as a function of selected kinematic properties are measured as well. The measured fiducial cross-section is found to be $\sigma_{\text{fid}} = 120 \pm 17 \text{ (stat.)} \pm 13 \text{ (syst.)} \pm 4 \text{ (lumi.) nb}$. This is slightly higher than the theoretical predictions with $80 \pm 8 \text{ nb}$ and $78 \pm 8 \text{ nb}$, however, the excess is not large enough to be statistically relevant.

Moreover, the search for ALP exploits the measured spectrum of the di-photon invariant mass. In contrast to previously published ALP limits [3, 81] the new limits are obtained including the shape information of the ALP signal. The limits are retrieved by a multi-bin fit of the expected ALP signal from simulation to the measured spectrum with the SM expectation as the background. As of today, the resulting upper limits are the most stringent available in the mass range $6 \text{ GeV} < m_a < 100 \text{ GeV}$.

The measurement takes advantage of the LHC as a photon collider in a genuine manner, opening a large field of possible measurements in a previously mostly unexplored scenario. The results encourage further similar measurements to constrain new physics in the light-by-light scattering processes.

Part II

FASER

7. Introduction

The experiments installed at the LHC in the previous runs have focussed mostly on searching for new particles which are heavy (at TeV-scales) and have large couplings. They are sensitive to new particles with high transverse momenta which would be produced isotropically around the IP. The opposite end of the spectrum are light, weakly interacting particles with masses in the range of MeV-GeV. Those lighter particles would not be produced isotropically but rather with transverse momenta of $p_T \approx 100 \text{ MeV} - \text{GeV}$. However, as the coupling of new particles in that range are expected to be small, a high rate of such particles is required to allow for their observation. Given the large inelastic proton-proton cross section of $\sigma_{\text{inel.}} = 78.1 \pm 2.9 \text{ mb}$ [106] / $\sigma_{\text{inel.}} = 71.3 \pm 3.5 \text{ mb}$ [107] at $\sqrt{s_{\text{NN}}} = 13 \text{ TeV}$ one would expect a huge amount of particles being produced in the far forward region in the low energy range. For future runs of the LHC at $\sqrt{s_{\text{NN}}} = 14 \text{ TeV}$ the number of inelastic proton-proton scattering events is estimated to be

$$N_{\text{inel.}} \approx 1.1 \times 10^{16} \quad (2.2 \times 10^{17}),$$

at an integrated luminosity of 150 fb^{-1} for LHC Run 3 (3 ab^{-1} for High-Luminosity LHC (HL-LHC)) assuming a similar inelastic cross section [108]. From this it is calculated that at the HL-LHC $N_{\pi^0} \approx 4.6 \times 10^{18}$ pions and $N_B \approx 1.4 \times 10^{15}$ B-mesons will be produced per hemisphere.

Given that the coupling of new particles to SM particles could be small, the particles produced by decays of those mesons are long-lived and can travel a large distance before decaying back into SM particles. They may be highly collimated e.g. new particles from pion decays are produced within angles of $\theta \sim \Lambda_{\text{QCD}}/E$, where E is the energy of the particle. Given they are light the energies are expected to be high with $E \sim \text{TeV}$. At energies this high the spread of such particles is only on the order of $10 \text{ cm} - 1 \text{ m}$ after having travelled $\approx 500 \text{ m}$ downstream. Hence, such particles may be observed using a small detector placed in the far-forward region.

The FASER detector is designed to do exactly this. It is located 480 m away from the ATLAS IP centred on the line-of-sight of the beam collision axis and has a decay volume of only 0.047 m^3 . The small size of the detector allowed for a very short timescale from publishing the Letter of Intent [109] and a Technical Proposal [110] in November 2018 to the installation of the finished detector in March 2021. The material cost for the full detector including a later addition of a neutrino detector (FASERv) is below 2 MCHF.

The FASER experiment is located in an unused side tunnel called TI12 which was formerly used for connecting the SPS to the LEP tunnel (which now is the LHC tunnel) and is not used any more. It is located perfectly to position the detector on the line-of-sight with only limited amounts of civil engineering required. Using in situ measurements and simulations done by the CERN Sources, Targets and Interaction group show that the background from high-energetic particles is quite low. In addition, beam radiation measurements show that the radiation levels in the tunnel are low.

Figure 7.1 shows a schematic view of the far-forward region from the ATLAS IP to the FASER location. The top shows the beam pipe with a 270 m long straight section which is the Insertion pipe followed by the arced section of the LHC. As can be seen on the left of the detail view, most charged particles in the forward region are bent away by the beam optics towards shielding material. Other particles hit the absorber which is located directly behind where the proton beams split. To cross from the IP to FASER in total 10 m of concrete and 90 m of rock must be travelled through. As a result the only SM particles produced at the IP which can reach FASER are muons and neutrinos.




Figure 7.1: Schematic drawing showing the location of the ATLAS IP and the location of the FASER detector. The top overview shows the 500 m length from the ATLAS IP to the FASER detector. The detail view on the left shows the region around the ATLAS detector up to the shielding blocks. The right detail view shows the area where FASER is located. Note that the vertical scale on the left detail view is much smaller. Figure based on [108].

As this thesis is targeting the search for axions at the LHC the following sections will focus on the potential production processes of ALP reaching the FASER detector and the decay into two photons. For a detailed evaluation of the physics reach of the FASER experiment for other long-lived particles refer to [108] and [111].

7.1 ALP production

ALPs may be produced from the IP in different processes. In this section the two processes for which the FASER experiment is most sensitive are discussed: the production from photons through the Primakoff effect and through the decay of heavy mesons.

The Primakoff effect denotes the interaction of a high-energy photon with the coulomb field of an atomic nucleus leading to a resonant production of neutral pseudoscalar mesons. This process is also postulated to produce ALPs. Under the assumption that the ALPs couple only to photons, it is the dominant production mechanism in beam dump style experiments [30, 112]. Such experiments work by dumping the beam in a dense material to quickly absorb hadronic particles only leaving weakly interacting particles. It is initiated by photons produced, e.g. from π^0 decays, interacting with nuclei from the various materials separating the FASER experiment

from the ATLAS IP and thereby producing an ALP. The most notable type of material is made up by the beam absorbers called TAN which are located ≈ 130 m from the IP after the beam pipes begin to curve (compare Figure 7.1).

A second postulated production process for ALPs potentially reaching the FASER detector is the production in heavy meson decays, namely the $B^\pm$ meson and the kaon. Both processes are based on an interaction called the flavour-changing neutral current (FCNC) which changes the flavour of a fermion without changing its electric charge. For example a down type quark becomes an up type quark through emission of a charged boson. This charged boson radiates an ALP before interacting again with the up type quark changing its flavour making it a down type quark once again.

While processes involving such interactions are allowed in the SM, they are heavily suppressed by the GIM mechanism [113] resulting in small rates for SM FCNC processes [114]. This gives ALP production in FCNC decays a striking signature and provides excellent prospects for discovery [115].

Figure 7.2 shows the Feynman diagrams for the production of an ALP through both processes discussed. For the production through heavy meson decay, the decays of a charged $B^\pm$ -mesons and kaon K^+ are given as an example. For the kaon it should be noted though, that also neutral kaons K_L can decay via the FCNC process: $K_L \rightarrow \pi^0 a$.




Figure 7.2: Feynman diagrams showing the production of an ALP by the Primakoff process (left), the charged B^+ meson decay (middle) and the K^+ meson decay (right).

7.2 ALP decay

Due to the low coupling constant, ALPs may travel long distances through material before decaying into two photons. Assuming an ALP momentum in the order of $p_a \sim \text{TeV}$, a mass of $m_a \sim 50 \text{ MeV}$ and a coupling constant $g_{a\gamma\gamma} \sim 1 \cdot 10^{-4} \text{ GeV}^{-1}$ the distance $\bar{d}$ is of the order of a few hundred meters [30]. This makes it feasible for an ALP produced from the ATLAS IP to decay within the FASER detector located ≈ 480 m from the IP resulting in a measurable signal in the detector.

The two photons from the ALP decay would be highly collimated due to the high energy and low mass of the ALP. The separation is expected to be in a range between 0 to $2000 \mu\text{m}$. This presents a challenge in cleanly distinguishing between events with one photon and di-photons. As background processes resulting in a single photon in the detector are expected to be low, the observation of high-energy deposits in the calorimeter with no signal in the rest of the detector would already give a strong hint towards new physics. Nevertheless, being able to classify between single and di-photon events reliably would greatly increase the sensitivity of the experiment as shown in Figure 7.3.




Figure 7.3: Reach for FASER in the ALP mass – photon coupling plane depending on the minimal separation δ of the two photons distinguished as two photons in the preshower. This projection assumes $N = 3$ signal events at an integrated luminosity of 3 ab^{-1} . From [30]

7.3 Outline

In the scope of this thesis, multiple components of the FASER detector were designed and fabricated. This includes:

- the design and construction of the upper support frame (c.f section 8.2),
- the initial module testing of the calorimeters (see section 10.1),
- the design of the PMT voltage divider and end-cap (see sections 10.2 and 10.4)
- the development and commissioning of a calorimeter calibration system (see chapter 11)
- the construction, simulation and performance evaluation of a possible future upgrade of the FASER detector (see chapter 12) leading to a publication [4],
- the implementation of an existing procedure for testing the quality of individual tracker modules and automation of the process (see section 9.2),
- the commissioning of several tracker planes following an existing procedure (see section 9.4).

8. | Detector

In this chapter, a detailed description of the design of the FASER detector is given. It starts with a short overview of the different sub-detectors before each individual detector is described in more detail. It is based mostly on the FASER Technical Proposal [110], however, updates from internal FASER meetings and personal design work are included. For these updates, no public documentation exists at the time of writing this thesis, but a formal detector design report will be published by the end of the year.

8.1 Overview of the detector components

Figure 8.1 shows a rendering of the FASER detector. Starting from the left, which is the side facing Interaction Point 1 (IP1) and hence, where particles enter the detector, the first component is the FASER ν neutrino detector. It consists of a scintillator station for vetoing charged particles and a chamber filled with alternating layers of tungsten plates and emulsion film. It is followed by a first tracker station, which is used for correlating signals in the main FASER detector with tracks in the emulsion film.

Next is another set of scintillators separated by a wall of lead bricks. These scintillators are to provide a high efficient veto for discriminating against all event candidates originating from charged particles, e.g. muons, entering the main detector. The lead bricks cause photons to shower and allow to shield against photons entering the detector.

It is followed by the first of three permanent magnets with an integrated magnetic field of 0.55 T m. The magnets are necessary to allow for reasonable separation of two opposite-charged tracks coming from high energetic particles. This first magnet is referred to as decay volume as it is where particles are to decay for optimum reconstruction efficiency. It has a length of 1.5 m.

Directly behind the decay volume another layer with scintillators is located. These are used to measure exact timing information for the event.

It is followed by the two other magnets. Due to the limited space in the tunnel these magnets are only 1 m long each. They are surrounded and separated by the three silicon tracker layers for reconstructing the path of charged particles through the detector. The far end of the detector is occupied by a short and simple preshower detector as well as an electromagnetic calorimeter.

The magnets are fixed onto plates which can be moved using screws allowing for precisely aligning the magnets on the line-of-sight of the beam collision axis. To ensure perfect positioning of the tracker station they are fixed to a rigid steel bar which is then mounted to the top of the magnets in a single unit. All other components do not have as tight tolerances on their location and are fixed to the so-called upper support frame. It is built from extruded aluminium profile allowing for flexible positioning of all parts.

All components are explained in some detail in the following sections. More detailed descriptions of the tracker and the calorimeter are given in chapter 9 and 10.




Figure 8.1: Computer rendering of the FASER detector showing all sub-detectors with the front of the detector on the left. For better visibility the cables are hidden in the render.

8.2 Upper support frame

The upper support frame is holding various sub-detectors as well as serving as a support for many electronics crates and distribution boxes. The frame is built using extruded aluminium profiles with a cross-section of $40 \times 40 \text{ mm}^2$ connected with angle brackets. The number of custom parts is kept as low as possible to reduce both the total cost as well as the manufacturing time. It is broken down into five separate compact stations to simplify transporting and installation of the frame. The individual sections are connected with horizontal beams to fix the position of the stations and to hold accessories in place. It is designed to be as flexible as possible to allow on site adjustments where necessary.

The design of the different support frame stations will be addressed in the section of corresponding sub-detector.

8.3 Scintillator stations

In total there are three scintillator stations on the detector: a veto station, a timing station and a preshower station. Each station will be shortly discussed in the following sections with some remarks on design decisions.

8.3.1 Veto station

The veto station consists of four scintillators arranged in two pairs as shown in Figure 8.2 separated by lead bricks. To allow a close spacing of the two pairs they are rotated by $\pm 3.7^\circ$ with respect to the vertical axis. The impact of rotating the scintillators on physics is negligible. The individual scintillator assemblies are made up from a 20 mm thick and $300 \times 300 \text{ mm}^2$ plastic scintillator covering not only the magnet aperture but nearly the full magnet. The plastic scintillator material used is made by Eljen Technology and is called EJ-200. It is chosen for its good efficiency paired with fast timing and a long optical attenuation length allowing for a good time resolution at high efficiencies [116]. Due to the thickness of 20 mm a very high detection efficiency of well above 99 % for a single layer can be assured. This is important to reduce the background from radiative processes associated with muons coming from IP1. The scintillator is coupled over a tapered plastic lightguide directly to a PMT for the signal acquisition. The PMTs used are Hamamatsu H6410 [117]. The efficiency of each scintillator pair is required to be at least 99.99 % to reduce the background to negligible levels. Due to the redundancy of using two layers the veto efficiency can be measured in situ as the inefficiencies should not be correlated. Each scintillator assembly has an optical fibre glued into the light tight wrapping to allow for monitoring the PMT response. It can also be used to validate the timing information of the various scintillators.

By putting lead bricks in between the two scintillators, high energetic photons from e.g. muon bremsstrahlung can either be fully absorbed or forced to shower and, hence, be easily detected in the second set of scintillators. This results in a highly efficient veto against incoming photons.

The scintillators are mounted using custom brackets fixing the PMT housing to the upper frame. To prevent too much strain on the scintillator, it is also supported at the bottom with an aluminium wedge. This wedge can be moved side to side to adjust the vertical position of the scintillator while keeping the angle constant. The lead blocks are just placed loosely on the bottom plate of the scintillator station.




Figure 8.2: Rendering of the veto station assembly from different views without any of the support structure. The four scintillators with light guides and PMTs are shown. They are grouped in two pairs with two layers of interlocking lead bricks stacked in between.

8.3.2 Timing station

The main purpose of the timing station is to notice the appearance of charged particles originating from a decay within the decay volume. To be able to correlate the event with collisions from IP1 and thus suppressing non-collision backgrounds, the time resolution of the timing station has to be in the order of less than 1 ns. To achieve that, the design varies from the design used for the veto station. The scintillator plates used are only 10 mm thick with a width of 400 mm and a height of 200 mm. The increased width is chosen to register muons entering the detector at an angle missing the veto and the tracker stations and could induce showers when travelling through the magnet.

As visible in Figure 8.3 the height of only 200 mm means that two scintillators need to be stacked on top of each other to cover the height of the magnet. However, there are multiple reasons for using the two individual plates instead of a monolithic approach. For one, the light guides required to concentrate the light from a single plate would be much larger which would be infeasible given the space constraints. Secondly, using two plates decreases the vertical time walk. The stacking is done with a small overlap of the two scintillator plates of around 4 mm to prevent any gaps in the active area. Each of the plates is read out on both sides to increase the sensitivity and to be able to compensate for the longitudinal time walk increasing the timing resolution of the station.

To reduce the width of the whole station as much as possible the light guides coupling the scintillator to the PMTs have a 90° bend in them to allow the PMTs to be oriented parallel to the detector. This more complicated setup is inevitable due to the space limitations in the detector location. As for the other scintillator stations, a Hamamatsu H6410 is used for the readout and optical fibre inputs are attached to the scintillators. As the scintillators are read out from both sides, each scintillator features two individual fibres.

Due to the thin scintillator, the glue joint between scintillator and light guide is quite fragile and the mounting method must not over-constrain the scintillator assembly. To minimize strain on the glue joint, the assembly is held only at the PMT housings and the scintillators are supported on the bottom edge. To hold the PMT housing in place, without putting force on it, common pipe clamps with rubber padding are used. These can be adjusted precisely to clamp the assembly in its natural position. The lower scintillator is resting on an aluminium bar while the upper scintillator is connected to the lower one using a set of 3D-printed brackets.




Figure 8.3: Rendering of the timing station assembly from different views without any of the support structure. It includes the two scintillators with bent light guides on both sides and the PMTs. The top scintillator rests on top of the bottom scintillator with two black brackets.

8.3.3 Preshower station

The purpose of this station is to identify events only consisting of photons with a high efficiency. This signature would otherwise only leave a high-energy deposit in the calorimeter which could not be differentiated from e.g. deep inelastic scattering of a high-energy neutrino, as the calorimeter itself has no longitudinal segmentation. To force the electromagnetic particles to produce showers, tungsten plates with a thickness of 3 mm are used as a radiator. This is equivalent to approximately 1 Radiation Length (X_0) [7] each. Three blocks of graphite with a thickness of 50 mm each are placed around the scintillator layers to absorb backscat from low energetic particles produced in the showers from travelling backwards through the detector. Graphite is chosen as it is able to absorb the low energetic backscat while not degrading the energy measurement.

Figure 8.4 shows a rendering of the preshower station. The scintillator assemblies are of the same design as the veto station. This allows for less individual components to be produced and the spare scintillators can be used for either of the two stations. In addition to being used to identify physics signals, this station also features as a second trigger and timing station. As both previous stations, the scintillators of the preshower station are equipped with an optical fibre.

To hold the scintillators in the support frame, a similar clamping method is used as in the veto station. The bottom edge of the scintillator rests on an aluminium plate on which the graphite and tungsten plates are resting as well. To hold them in place they are just wrapped together with Kapton tape.




Figure 8.4: Rendering of the preshower station assembly from different views without any of the support structure. It shows the two scintillator plates with light guides and PMTs, the graphite blocks (light grey) and the tungsten plates (dark grey).

8.4 Electromagnetic calorimeter

The electromagnetic calorimeter is the last component in the detector. It provides a strong identification criterion between electrons/photons over muons and hadrons. However, even for an electron-positron-pair having travelled through the whole magnetic field, the expected separation in the calorimeter is smaller than a few mm. Hence, it is not feasible to measure the energy of the individual electrons/photons and as a consequence, only the full event energy

will be measured. The tight space constraints in the cavern do not allow for the installation of a leakage detector resulting in degrading performance at the highest energies. The amount of leakage is to be determined in a future testbeam.

To reduce the amount of development and detector costs, the calorimeter is built from spare detector components from the Large Hadron Collider beauty (LHCb) experiment. For that, LHCb provided eight spare modules from their outer electromagnetic calorimeter for FASER to use. These modules are so called Shashlik-style calorimeters [118]. This style of calorimeters is constructed by stacking alternating layers of absorbing and scintillating material with Wavelength Shifting (WLS) fibres going perpendicular through all the layers. In the case of the LHCb modules, 66 layers of 2 mm thick lead and 4 mm thick plastic scintillator plates make up the calorimeter module, resulting in a total of $25 X_0$. The individual plates are $121.2 \text{ mm}^2 \times 121.2 \text{ mm}^2$ in size with the full module (including PMT) being 754 mm long. Figure 8.5 features a technical drawing of such a module where the stacked layers and the WLS fibres are clearly visible. The 65 WLS fibres collecting the light from the various scintillating plates are bundled together in the rear of the module and are coupled to a PMT through a light mixer. To form the calorimeter, four modules are put together in an 2×2 array and wrapped together to prevent light leaking into the modules. To reduce efficiency loss caused by gaps between the separate modules and by the non scintillating fibres going straight through, the modules are rotated by 2.8° both in the vertical and the horizontal axis. This is enough to prevent particles from travelling notable distances through the calorimeter volume without hitting the absorber.

Similar to the scintillators of the various scintillator stations, the calorimeter modules feature a port for attaching an optical fibre. A WLS fibre going straight through to the centre of the module to the PMT makes sure that nearly all incoming light reaches the PMT. The calibration and fibre connection is further discussed in chapter 11.

The calorimeter station consists of a simple frame construction with an aluminium plate to support the calorimeter modules. To set the rotation of the modules correctly, two silver steel pins are pressed into the aluminium plate to align the modules against. To fix the modules together, aluminium brackets are fitted to the edges and everything is tied together with long zip ties.




Figure 8.5: Technical drawing of a module for the LHCb electromagnetic calorimeter (based on [119]). On the left side, a side view of the module is shown with the alternating layers being visible. The green WLS fibres penetrating all layers are bundled at the rear end and are coupled to a PMT through a light mixer. On the front of the module, a port is available for attaching an optical fibre for calibration purposes. On the right, a single scintillator plate is displayed showing the holes for the optical fibres.

8.5 Tracker

For reconstructing the path of charged particles in the detector, it is equipped with three tracking stations. The stations are located at the end of each magnet. The number of stations is optimized for sufficient redundancy and resolution while keeping the amount of material in the detector as low as possible. High energetic low mass particle decay into highly collimated particles with an included angle of $\theta \propto m/E$. For a Long-Lived Particles (LLP) with mass $m = 100$ MeV and an energy $E = 1$ TeV this leads to an opening angle of $\theta \approx 200$ μ rad. This angle is too small to resolve the two tracks clearly, but as the tracks are opposite-charged, the separation can be increased by the magnetic field. This would in principle also give the possibility to reconstruct the momenta of the charged particles using the curvature of the reconstructed tracks. However, considering the limited magnetic field strength and detector length, the achievable measurement is not expected to yield a precise measurement of the particle's momentum.

To save time and cost, the tracker stations are constructed using spare modules from the ATLAS SCT. These modules have been originally developed for the barrel region of the ATLAS ID. The module is constructed from four $6.36 \text{ cm} \times 6.40 \text{ cm}$ silicon detectors with 768 strips resulting in a pitch of $80 \mu\text{m}$. Two 12.8 cm strips are built by wire-bonding two detectors together. The two strips are then glued back-to-back with a thermal pyrolytic graphite baseboard for cooling and mechanical connections in between forming a full module. The layers are rotated by an included angle of 40 mrad to allow for a two-dimensional position reconstruction. Assuming a setup with four SCT layers in perfect alignment, this yields a resolution of $16 \mu\text{m}$ for the precision coordinate and $580 \mu\text{m}$ for the other direction [120]. A bare module is depicted in Figure 8.6.

For the readout, a flexible Printed Circuit Board (PCB) is wrapped around the module. It is housing six Application-Specific Integrated Circuits (ASICs) per module side and the silicon is connected to the PCB with wire-bonds. On one edge of the flexible PCB it has a connector for connecting to the readout electronics.

Each station is built up from three planes with eight SCT modules each. A photograph of an assembled plane is shown in Figure 8.6. It shows the aluminium plane support frame equipped with eight modules. The frame is providing both a stable fixation of the SCT modules as well as cooling. While the cooling is not strictly necessary for the silicon detectors given the low radiation at the FASER location, the ASICs must be actively cooled. To ensure proper heat dissipation, the plane support frame has internal bores for circulating chilled water.

8.6 Magnet system

As stated before, FASER uses dipole magnets to achieve a better separation between opposite-charged tracks and to allow a coarse momentum measurement. As the detector location has no infrastructure for cryogenic cooling, using super conducting magnets is out of question and the use of permanent magnets is the only feasible option for the experiment. The magnets are built from multiple neodymium-magnets (NdFeB) arranged in a Halbach-array [121] to create the dipole field. The size of the magnet and hence, the magnetic field strength possible is highly limited by the available space in the detector location as well as the transport capability within the LHC tunnel. With an aperture of 200 mm diameter and an outer diameter of 430 mm , a field of 0.55 T is achieved with a fixed homogeneity of $\pm 2\%$.




Figure 8.6: Photography of a bare SCT module [122] (left) as well as an assembled plane (right). The top silicon detectors are visible with the edge of one of the bottom detectors peeking out on the lower edge due to the stereo angle of 40 mrad. The wire-bonding can clearly be spotted in the area where the two top layer detectors meet.

8.7 Readout and reconstruction

This section gives a short introduction into the FASER readout system. A more detailed description can be found in [123].

The detector is read out by multiple acquisition boards controlled by a custom software stack. The full setup is designed with an expected rate of 500 to 1000 Hz in mind. Most signals will be induced by muons produced at the ATLAS IP crossing the FASER detector. These will mostly produce signals in the scintillator stations while energetic deposits in the calorimeter are expected at a rate of around 5 Hz.

To read out the scintillators and calorimeter modules, a digitizer (CAEN VX1730) is used. It features 16 input channels with a 14-bit resolution at a sampling rate of 500 MS/s. The digitizer is able to produce triggers for all input channels at individual threshold settings. The thresholds for scintillators will be set below the expected signal for a single minimum ionizing particle while the calorimeter threshold will be set at an energy deposit of 20 GeV.

The trigger signals produced from the digitizer are fed to the Trigger Logic Board (TLB). The TLB performs the central trigger decision and sends out the global trigger accept signal (L1A) to the digitizer and tracker readout boards to start the event readout. To keep the tracker data in sync with the TLB events, the tracker readout boards are supplied with a clock signal and a bunch counter reset (BCR) signal which is pulsed on every LHC orbit signal. Once a L1A signal is received, the full detector is read out and the full raw information is written to a ROOT file. The data is copied from the data acquisition (DAQ) computer to a storage server in regular intervals to keep the local disk space free. This allows to reconstruct the data without interfering with the detector operation.

The reconstruction for FASER is performed using the FASER offline system, called *calypso* [124]. It is based on the *athena* and *Gaudi* frameworks developed by the ATLAS collaboration. *Calypso* processes the raw detector information and associated metadata and produces ROOT files in the AOD format. This includes fits of the PMT waveforms with baseline subtraction, reconstruction of tracks and calculation of track parameters and vertices.

No established method exists so far to analyse the reconstructed files. However, a program was written in the scope of this thesis to produce NTuples for the test beam analysis.

9. | Tracker hardware and tracker commissioning

In this chapter, a detailed description of the FASER tracker is given. It opens with an overview on the working principle and the digital front-end of the SCT modules. The modules used are spare modules from the ATLAS silicon strip tracker, which have been in storage for multiple years under uncertain environmental conditions. Therefore, each module must be tested individually to confirm that all operating parameters are still in the nominal range. The next part is, hence, focussed on the setup used to test the modules and the results from this Quality Assurance (QA) process.

It is followed by a description of the on surface commissioning procedure for the assembled tracker stations performed before installation in the experimental cavern.

9.1 SCT module

The ATLAS SCT modules combine silicon strip sensors with analogue and digital readout in a compact package. Each module has four single sided sensors in two layers. Each sensor is $6.36 \times 6.40 \text{ cm}^2$ in size and has 768 strips with a pitch of $80 \mu\text{m}$. The two sensors on the same layer are wire-bonded together to form 12.8 cm long strips. The two layers are glued back-to-back, with a thermal pyrolytic graphite baseboard for cooling and mechanical connections in between, forming a full module. The two layers are rotated with respect to each other by an angle of 40 mrad to enable reconstruction of two-dimensional hit positions. A flexible PCB is wrapped around the detector housing six ASICs on each side. Wire bonds connect the individual strips to the flexible PCB and, thereby, to the ASICs.

A schematic drawing of an SCT module is shown in Figure 9.1. There, one can also see the mechanical mounting features of the module. The baseboard holding the four square sensor extends over the module boundaries on both long sides to fix the module to a support structure. Beryllium oxide sheets are glued to the exposed baseboard. Beryllium oxide is an electrical insulator and features excellent thermal conductivity enabling an efficient heat dissipation from the silicon to the support structure. The flexible PCB extends even further over the baseboard and has a 36-pin connector for connecting the module to the readout system.

Two washers are fixed to the mounting surface to precisely align the module on the detector support structure. The datum washer is a precise round washer constraining the position in both axes, while the slotted washer has a slot to allow for movement in one axis. Together, the washers perfectly align the module without over constraining it. This is important as otherwise stress in the baseboard could cause the module to bend or in the worst case mechanically damage the module.




Figure 9.1: Schematic drawing of the bare SCT module. Clearly visible is the angle between the two silicon sensor layers. The ASICs are marked with their numbering in the readout chain. Two washers on the bottom right are used to precisely align the modules on the detector support structure. Adapted from [125].

The ASIC, called ABCD, is designed especially for the ATLAS silicon tracker and features all necessary analogue and digital circuits for reading out the silicon detectors [126]:

- integrating the strip charge,
- pulse shaping,
- discriminating based on pulse amplitude,
- storage of data in a pipeline until a trigger signal is received,
- de-randomisation and compression of the data,
- generation and injection of calibration pulses,
- transmission of the data via a serial daisy chain link.

For each channel, the signal is first amplified in a transimpedance preamplifier and, subsequently, shaped in a pulse shaper. To detect if a strip was hit by a charged particle, the resulting signal is fed into a discriminator. The threshold for this discriminator is set with a single 8-bit Digital-to-Analog converter (DAC) and is common to all 128 channels of an ABCD chip. Using a common threshold for all channels arises the problem that all comparators must have the exact same behaviour to guarantee a uniform performance. As this could not be achieved in the first prototype stages, a 4-bit DAC was added to each channel, named TrimDAC. This DAC is used to offset the threshold for each individual channel. The range of the TrimDAC can be adjusted by another 2-bit setting to allow for larger, albeit less precise, corrections.

A capacitor is connected to each channel to inject calibration pulses at various amplitudes. The channels are grouped into four groups with every forth strips capacitor being connected together. The group for the charge injection is selected with a 2-bit register. This allows for studying cross-talk between the strips as the strips surrounding the one under test are always without charge. The calibration strobe signal can be shifted with respect to the main clock by a tunable delay up to two clock periods. The strobe delay setting is common for a whole chip

and thus affecting 128 channels.

The signal from the silicon detector, taking into account any possible smearing of charge by the magnetic field and diffusion effects, is in the order of 1 fC. In order to limit the maximum noise occupancy to 5×10^{-4} , the Equivalent Noise Charge (ENC) must not exceed 1500 electrons. The timing performance is defined by the requirement of clearly associating hits to a specific bunch crossing, which leads to a maximum allowed time-walk of 16 ns. Given that the expected rate for FASER is magnitudes below the bunch crossing rate, this is by far sufficient for the FASER experiment.

The data transmission is realized by chaining the chips on each side in series. Each side has one dedicated master chip, four slaves and one end chip. Each module has two transmission lines to interface the module from the DAQ. The two lines are connected to the master chips on each side of the module, which take care of communicating with the chips on the corresponding side. Two prevent a module rendering unusable in case of a single chip failure, each chip is not only connected to the next chip in the chain but also two the next-to-next chip. In case one of the end chips is broken any slave chip can be configured to act as the end of the chain. The module can function even in the case that one of the master chips is broken which results in the full module being read out over a single master chip.

9.2 Module quality assurance procedure

The silicon tracker modules used for FASER are spare modules left over from the inner detector of the ATLAS experiment. They were individually packaged in a metal transport housing, wrapped in plastic bags and stored inside open boxes. The individual plastic bags were however not sealed and no control over temperature or humidity was in place. Due to this storage situation and the fact that is unclear how the modules were treated after the initial quality control by ATLAS each module needed to be tested.

While the basic test procedure already existed, it was optimized and automated as much as possible in the scope of this thesis.

The FASER readout was not ready in time for the quality assurance performed in spring of 2019 and therefore, a readout system was borrowed from the University of Cambridge [127]. It provides the functionality to write the configuration of the ASICs, inject test pulses with variable charge and readout hits from the silicon strips.

In addition to the readout, three power supplies are used: one for the ABCD chips on the hybrid, one for an adapter card and readout and one for the High Voltage (HV) bias of the silicon. Two general purpose power supplies, namely Agilent E3646A [128], are used for the low voltage of the digital circuits on the hybrid. These are fixed to the necessary voltages and are just turned on/off for switching the module. For the high voltage a precision source meter (Keithley 2400 [129]) is used to both supply the bias voltage as well as precisely measure the current drawn. To automate the quality assurance procedure the power supply for the bias voltage is not controlled manually, but controlled remotely using a serial interface.

To keep the modules sufficiently cooled a water chiller is used to cool down an aluminium plate on which the modules are placed. The aluminium plate is surrounded by a polystyrene box with a detachable lid. This helps keeping the full module at a uniform temperature. To prevent any condensation the temperature and humidity within the box is monitored constantly to ensure operating above the dew point. In addition, the box is constantly flushed with dry air to keep the humidity levels down and approximately constant. The full setup can be seen in Figure 9.2.




Figure 9.2: Setup at CERN used for the quality assurance of the individual SCT modules. The module is kept in its transport housing and placed on top of a water cooled aluminium plate in an insulated box.

All modules were tested using the same procedure. It starts with taking the module from its bag and noting the module's serial number. Next, the cover protecting the flexible PCB housing the module connector is removed and the adapter card for the readout is attached. The module is placed on the cooling plate in the polystyrene box as shown in Figure 9.2. The chiller is switched on with a set point of 12 °C. When the module's temperature has reached a stable level, a scan of the current in dependence of the bias voltage is performed up to a voltage of 300 V. The scan is performed using a dedicated software developed in the context of this thesis which controls the power supply via a serial connection. It is written in Python [86] and based on the framework **PyMeasure** [130]. The parameters for the scan are set in a graphical interface with the nominal testing procedure measuring in steps of 5 V up to 25 V and in 25 V steps up to 300 V. The ramp up and down speed is limited to 1 V/s to keep the current flow low during ramping of the voltage. Once the desired voltage is reached the program waits for 10 s to allow for the current to stabilize and then measures the current five times saving the average value. After the scan is completed the voltage is ramped back down to 150 V for the next steps in the QA procedure. Figure 9.3 shows the interface for this tool after a scan is completed. The scan results are saved in a csv file for further analysis.

When the bias voltage has reached 150 V the low voltage for the readout system and the ASICs on the module are enabled and the temperature is noted once stable. A so-called threshold scan is used to evaluate the noise level on each strip. This scan measures the per strip occupancy, meaning the number of hits per strip, at different thresholds. The threshold is varied between 0 and 60 ADC counts in steps of 1 with 10000 triggers per value. The measured number of hits in dependence on the trigger threshold is fitted with an error function to extract the threshold




Figure 9.3: Screenshot of the program used for measuring the current in dependence of the bias voltage in the module QA.

above which noise is not triggered any more. The error function is defined as:

$$f(t) = A \cdot \operatorname{erf}\left(\frac{t - t_0}{\sqrt{2}\sigma}\right), \quad (9.1)$$

where the t_0 is giving the strip noise. The threshold scan for a single strip of a module is shown in Figure 9.4 together with the per strip noise distribution for the full module. The dual peak structure is caused by the noise level distribution being different between the two silicon detectors of the module.

The last test is a gain measurement. It is basically the same procedure as for the noise scan except for injecting charges into the strips with the charge injection feature of the SCT module. In this scan, the threshold is varied between 0 and 100 ADC counts in steps of 1 with 100 triggers per value and performed once for each charge value. Again the function in Equation 9.1 is used to fit the measured threshold scans and extract the edge. By injecting pulses with 1.5 fC, 2.0 fC and 2.5 fC the gain for each channel can be extrapolated using a linear fit. The fit results for a single strip as well as the gains for a full module are shown in Figure 9.5.

In addition to this procedure, one module was selected and undertaken a set of more advanced tests to confirm that the thermal stability of the modules could still be relied on. In the first step, the standard testing procedure is performed on the module. To ensure that the measurement of the leakage current is stable it was kept powered on over a duration of 72 h and the leakage current and module temperature is measured once per minute. This test did not show any abnormal behaviour. Using a coordinate-measuring machine multiple markers on the modules surface were measured and compared to the nominal values from the ATLAS pro-




Figure 9.4: Results from the noise threshold scan for a single strip fitted with an error function as defined in Equation 9.1 (left) and the distribution of the noise for a full module (right).




Figure 9.5: The measured signal versus three injected charges (1.5, 2 and 2.5 fC) with a linear fit for a single strip (left) and the distribution of the resulting gains for a full module (right).

duction documentation. The module was then put into a climate chamber and the temperature was ramped from 20 °C to 60 °C and back in 100 cycles. After thermocycling the standard testing procedure is performed once again. This revealed that the electrical performance of the module did not change due to thermocycling.

9.3 Results from the module quality assurance

The first quality criterion checked is the leakage current at a voltage of 150 V. The distribution over the modules is shown in Figure 9.6a. Module with a leakage current larger than $1.5\ \mu\text{A}$ are considered bad and are not further considered for the usage for FASER. This removes six modules from the selection. The measured curves of the leakage current versus the bias voltage show no abnormal behaviour for any of the remaining modules.

By evaluating the activity on each strip in the noise and gain scans different failure modes can be tested for. The following four failure modes are considered:

- dead strips: strips that do not respond to injected pulses at all thresholds
- low efficiency strips: strips which register hits only in 10 %-90 % of triggers (evaluated at the most efficient threshold)
- noisy strip: strips with a high level of noise
- cross talk: strips which show a large amount of cross talk from neighbouring strips.

To find strips with high cross talk the threshold is found at which the hit occupancy with charge injection is 95 %. If the hit occupancy at the same threshold for events with no charge injection in that strip is over 10 % the strip is considered prone to cross talk. The most frequent fault is noisy strips with a total of 71 strips over all modules finally selected for the FASER tracker. In total the selected modules showed 108 strips falling into one of the four failure modes during the initial quality assurance. This corresponds to less than 0.1 % of strips with defects in the selected modules for the tracker stations.

The gain is found to be relatively consistent over all modules. In Figure 9.6c the per module average gain is shown and is in the range of 43 to 50 mV/fC. The average over all strips of the installed modules in the quality assurance data is 47.6 mV/fC. While this is slightly lower than the design target of 50 mV/fC [126], it is still acceptable. This difference is not fully understood but most likely it is caused by the readout software; not supporting all tuning features of the ASIC. This is further endorsed by the results observed during commissioning of the assembled tracker stations with the final readout system where an average gain of 54 mV/fC is measured.

The equivalent noise charge is measured to be around 1100 electrons for test pulses of 1 fC. This is far below the nominal value of 1500 electrons at 1 fC [126]. As can be seen in Figure 9.6d only two modules surpass this target and are not usable. An additional noise study using data from the three-point gain measurements where the noise is measured with charge injection enabled leads to a noise of ≈ 1760 electrons on average while the ATLAS quality assurance found an average noise of ≈ 1650 electrons using this method. The difference is only $\approx 6\%$ and could be either explained by differences in the measuring setup or environment or by the fact that ATLAS selected the best modules for their detector and thus the test results are biased towards higher noise levels.

9.4 Plane and station commissioning

Using the results from the QA explained in the previous section the best modules are selected to be assembled into the FASER tracker. The assembly process is performed in two steps: first eight modules are fixed into their frame to form a so-called tracker plane. These planes are then tested individually to ensure that the modules were not damaged during installation and all electrical connections are good. Only if they pass all checks three planes are assembled into a full station. This allows for bad modules to be replaced before assembling a full station and hence, reduces the risk of damaging modules during repeated station assembly and disassembly. Considering the tight timeline for building the FASER detector this also gives the advantage that finished tracker planes can be tested while the other planes of a station are still in assembly.

In the following sections the different setups are explained, followed by the commissioning procedure and finally, the results found from the commissioning are given. In the scope of this thesis, multiple planes and one station have been commissioned using the existing procedure.




Figure 9.6: Leakage current at 150 V (top left), number of bad strips per module (top right), average module gain (lower left) and average module equivalent noise charge (lower right) as measured in the quality assurance used to select the modules for building the FASER tracker stations.

9.4.1

Commissioning setups

To allow commissioning of multiple planes in parallel two commissioning stations in different labs are used. The first lab is only setup to commission a single plane while the other lab has the required equipment to run both tests in parallel on a single plane as well as on a full station. After a full plane is assembled it is first tested in the first lab and once three planes are tested completely they are assembled into a station. This station is then tested in the second lab. The single plane setup in the second lab is mostly used to run tests on an early prototype plane to debug the firmware of the readout boards and the DAQ software on the PC side. It is also used to occasionally test planes to validate the results from the other lab.

Both setups use the same readout and monitoring system and same power supplies as are used during operation in the cavern. The most notable components are:

- low and high voltage power supplies,
- Tracker Interlock Modules (TIM),
- readout card,

- computer for readout and controlling the Finite State Machine (FSM).

The Tracker Interlock Module (TIM) is an important safety system for operating the tracker as it monitors the temperature of each module and the module frame and switches off the power in case the temperature is increasing above the safe operating range. It is using discrete logic to make the interlock decision purely hardware based but also features a microprocessor to enable reading the temperature values over ethernet.

Figure 9.7 shows the setup used for commissioning the assembled tracker stations in the second lab. The station is held in a temporary support frame to allow for manipulating its position and orientation. This is especially useful as it permits to run either the normal commissioning sequence or take cosmic data by rotating the station to be horizontal. Not shown in the picture are the power supplies for the tracker, the chiller and the readout computers. To keep the risk of condensation on the plane as low as possible the station is wrapped with plastic foil and dry air is flushed around it. The setup in the plane commissioning lab is similar except for the planes just laying flat on a table instead of mounted in a frame.




Figure 9.7: Setup used for commissioning of the FASER tracker stations as of July 2020. The station is held in a temporary frame to allow for manipulating its position and orientation while keeping the connections protected.

9.4.2 Testing procedure

The testing procedure is explained only for the station commissioning in the following as the procedure for plane and station commissioning are mostly identical. The main difference is that all tests have to be run for each plane in parallel for the station commissioning.




Figure 9.8: Temperature (left) and analogue rail current (right) for Plane 10 in the powered on and configured state as measured during plane commissioning and station commissioning. Both results agree with exception for module 2 which was replaced during station assembly.

In preparation of running the tests, the station is cabled carefully following the cabling diagram and the connections are noted in a dedicated spreadsheet. The dry air supply and chiller for the water cooling are turned on and the cooling loop is inspected for leaks. To confirm that the cabling of the module power supplies is correct the low voltage is turned on for a short duration module by module. Using the TIM monitoring it is confirmed that the temperature is rising in the correct module and henceforth, the cabling is correct.

This is followed with a scan of the leakage current against bias voltage as done in the module QA up to a voltage of 150 V in steps of 10 V. At each step the voltage is kept constant for 20 s and the current is continuously recorded. When the scan is finished, the bias voltage is switched off again before continuing.

The next steps are focussed on the electrical and thermal behaviour of the ASICs. First the temperatures for all module and frame sensors are noted without any power applied. Next the low voltages are switched on for all modules and after waiting for the temperatures to stabilize the measured voltages and currents as well as the temperatures are noted. Next the modules are configured which increases the power consumption for the analogue rail and consequently the temperature of the module rises. Once again, all values are noted. A comparison between the results from plane and station commissioning for the temperature and analogue rail current is shown exemplary for plane 10, used as layer 0 in station 3, in the powered on and configured state in Figure 9.8. Good agreement is visible there with some deviations in module 2 which was replaced after the plane commissioning. The same is observed for all the other layers.

The next tests are all performed using a dedicated calibration software developed by the FASER DAQ team called *tcCalib*. It is able to run a sequence of tests either on a full plane or a selection of modules. All analogue tests are based on a threshold scan either with or without injecting test charges. Depending on the aim of the scan a different parameter is changed during the threshold scan procedure.

The first test in the sequence is a so-called mask scan. It is searching for strips with high noise level to exclude them from further analysis. The next scan is needed to find the optimal delay for the calibration strobe to ensure that the pulse is properly captured by the digitization. Without tuning this correctly the results from the threshold calibration might be skewed. It is

called the strobe delay scan, in which the value for the strobe delay is varied between 0 and 64 and a threshold scan is run for each setting with 100 pulses per value. The value resulting in the highest measured charge for all strips of a module is taken as the optimum value.

Once the values for the optimum strobe delay are found for all modules a three point gain measurement is executed. As in the QA procedure the three point gain is measured by injecting charges of 1.5, 2 and 2.5 fC and measuring the response. In this range the gain is most linear and the scan can verify that the chip is working properly. Its purpose is only to validate the analogue performance of the chip but not to precisely calibrate it. The measured gains and noise levels from this scan are shown in Figure 9.9 for comparison with the gains measured during the module QA given in Figure 9.6. The noise level is higher compared to the module QA results as the noise in the station commissioning is measured with charges injected where the module QA measured the noise without any charge injection. The difference in gain is due to the final readout system allowing more precise tuning of the SCT modules which is not possible with the system used during the module QA.

The next step is tuning the values for the TrimDAC to ensure a uniform response for the full module. This is done to minimize the threshold variations at a charge injection of 1 fC so that the efficiency is as close to identical as possible. The scan procedure measures the threshold at four different TrimDAC values for each channel which is then fitted with a linear function. Using the linear function the best TrimDAC value is calculated. The hit occupancy for a single chip before and after applying the optimum TrimDAC values is shown in Figure 9.10. There, one can clearly see the importance of tuning the TrimDAC values to get a uniform tracking efficiency.

It is followed by another three point gain test to ensure that the TrimDAC settings are giving the expected result before running a time-consuming response curve scan. The response curve scan is an extended version of the three point scan injecting charges of 0.5, 0.75, 1.0, 1.25, 1.5, 2, 3, 4, 6 and 8 fC and is also referred to as a 10-point gain measurement. As a result from this scan, precise information on the gain and offsets of the preamplifiers are calculated that are used to update the module configuration. A full response curve measurement for a single strip is shown in Figure 9.11.

9.4.3 Results

Performing the extensive testing on each module multiple times allows for gaining a good understanding on the behaviour and stability of the tracker system. A few modules were found that showed problematic behaviour or did not perform consistently. The problems observed were, among others, that chips would sometimes fail and prevent a full group of 128 strips not being read out. Those modules were consequently replaced before the station assembly. In addition to those problematic modules two modules were mechanically damaged during the assembly procedure and were replaced as well.

The overall performance is found to be excellent with a very small amount of defect strips and all modules functioning nominally. In the scans performed after installation of the detector in the cavern a total of 46 dead and 59 noisy strips are found for the three main tracker stations, which corresponds to less than 0.1% of all strips. Cosmic tests both on the surface prior to installation and in the cavern confirm the expected performance.




Figure 9.9: Measured per chip gain (left) and noise (right) for all modules used in the tracker. The noise is calculated from the threshold scan with 1 fC charges injected.




Figure 9.10: Hit occupancy for a single chip of a module before optimizing the TrimDAC (left) and after applying optimum TrimDAC settings for said module (right).




Figure 9.11: Measured response for different injected charges used to determine the response curve for a single channel of the first module of the first station (points) with the fitted response curve (red line).

10. | Calorimeter

As previously discussed, the calorimeter is built of four spare calorimeter modules from the outer electromagnetic calorimeter of the LHCb experiment which kindly provided eight modules in total. This allows to select the best performing modules for installation in the experiment and having additional modules for on surface testing or usage as replacement in case of failure of a module.

This chapter focusses on tests of the calorimeter modules themselves, the PMTs used and the design of a custom PMT voltage divider base. In the scope of this thesis, the initial module tests were performed and the voltage divider and PMT end-cap were designed and fabricated.

10.1 Initial module testing

The integrity and basic functionality of the calorimeter modules is tested at CERN in the same setup originally used for commissioning of the modules for the LHCb experiment. It relies on a radioactive source being moved along the module under test by a linear actuator and thereby scanning over all scintillator plates. This allows to verify that the scintillator plates are working as expected and that the WLS fibres are intact and guiding the scintillation light to the PMT.

The radioactive source used is an ^{137}Cs source emitting photons with an energy of 662 keV. A PMT¹ is used along with the power supply and readout by LHCb. For all eight modules each side is placed down on the fixture and scanned with the radioactive probe. The whole setup is covered in black cloth to prevent light leakage into the module. A picture of the setup without the black cloth and the scan result for one side of a module is shown in Figure 10.1.

Looking at the scan results one can see the response of the calorimeter module versus the position of the radioactive source. The measured signal is slowly rising until the radioactive source reaches the front of the calorimeter module at a distance of ≈ 90 mm. At that point the signal is steeply rising followed by a slight rise overlaid by an oscillation. The rise can be explained by the light attenuation becoming smaller the shorter the distance to the PMT becomes. The oscillation on the other hand is produced by the alternating layers of lead absorber and scintillating material. Within each module the response is found to be nearly identical for all four sides. A comparison between the modules is not possible as the absolute signal yield depends strongly on the positioning of the PMT which was not reproducible among the scans. However, one can state that the overall performance is found to be as expected with no obvious damages in any of the eight modules tested.

¹Hamamatsu R7899-20




Figure 10.1: Setup by LHCb to test the response of a calorimeter module (left) and the scan results for one side of a module (right).

10.2 Photomultiplier and power supply

To detect the scintillation light produced in the calorimeter so called PMTs are used. PMTs are a special kind of vacuum tube and produce a measurable signal from single incident photons. They are made of an evacuated glass tube with electrical connections fed through on one end of the tube. On the other end of the tube, a photocathode is situated which allows for conversion of photons to electrons through the photoelectric effect [131]. The emitted photoelectrons are focused onto a conducting plate, called dynode. An electric potential applied to the dynode causes the electrons to be accelerated and hence, on impact on the dynode secondary electrons are emitted from it. This process repeats in each of the subsequent dynodes increasing the number of electrons in every step. The electrons emitted from the last dynode are collected on the anode producing a measurable electrical signal.

The number of electrons being produced per initial photoelectron is referred to as the gain of a PMT. The gain depends on the average number of secondary emissions δ and the number of dynodes n and can be calculated by:

$$g = \delta^n. \quad (10.1)$$

The average number of secondary emissions itself depends on the material used for the dynodes and the potential difference E between each adjacent dynodes:

$$\delta = a \cdot E^k, \quad (10.2)$$

where a is a constant and k is a material specific constant normally in the range of 0.7–0.8 [131]. Under the assumption that the interstage voltage is equally distributed, one can set $E = \frac{V}{n+1}$, where V is the overall potential difference. Inserting this relation and Equation 10.2 in Equation 10.1 yields:

$$g = a^n \left(\frac{V}{n+1} \right)^{kn} = A \cdot V^{kn}, \quad (10.3)$$

where A is combining all device specific constants and the potential difference between the individual dynodes. Figure 10.2 contains a schematic view of the components of a PMT and the flow of electrons. For the overall efficiency of a PMT, in addition to sufficient gain, one also has to take into account, that not all photons will cause a photoelectron to be emitted




Figure 10.2: Schematic drawing showing the main components and the working principle of a photomultiplier. The electrical connections from the base to the components are left out for clarity (based on a graphic from [131]).

from the photocathode. The probability for the emission of a photon is given by the quantum efficiency. The quantum efficiency depends on the material and quality of the photocathode and the photon's wavelength. A typical range for the quantum efficiency is 10 to 20 %.

To realize the required voltages at the various dynodes the most common solution is to use resistive voltage divider circuits. This allows to operate the PMT with a single HV power supply and to change the gain by varying this voltage. It should be noted though that depending on the use case other means of biasing are more appropriate. For example for high occupancy applications an active voltage generator can be favourable over a fixed voltage divider as it provides better stability under higher current flows.

The PMTs used in FASER to detect the scintillation light are the same type as originally used in the calorimeter and were kindly provided by the LHCb collaboration along with the calorimeter modules. They are produced by Hamamatsu and have the model number R7899-20 [132]. This model is a custom variant for the LHCb experiment and is based on the R7899-01 variant with the key difference of having the pluggable solid connectors replaced with longer flying leads for soldering directly onto an adapter PCB. The R7899-20 has 10 dynode stages and can be operated up to a voltage of 1800 V at which it features a gain of $\approx 2 \cdot 10^7$.

LHCb uses an active voltage generator for their calorimeter [133] which is a good solution when using a large number of PMTs as the cost per channel is lower compared with a voltage divider. As the FASER calorimeter only needs four PMTs in total using this or a similar system is not an option and a classic resistive voltage divider is preferred. Due to the soft leads of the PMT the commercial voltage divider bases by Hamamatsu can not be used and a custom voltage divider is required. The mechanical design of the voltage divider is inspired by the LHCb solution consisting out of a round adapter PCB and a rectangular PCB for the electronics connected by an angled pin-header allowing the assembly to fit within the tube at the end of the calorimeter module. The schematic for the voltage divider is following the recommendations by Hamamatsu for the specific PMT and is shown in Figure 10.3.

The potential difference between the stages is not equidistant, but is tapered to both ends. This partially counteracts charge up effects in the later stages and henceforth, stabilizes the gain and increases the linearity. In addition, capacitors are added to the last four dynodes to compensate voltage drops during short peaks in current flow caused by the electron amplification.

The round adapter PCB is using pin receptacles for the leads of the PMT soldered through holes in the PCB which allows to detach the PMT from the voltage divider for various tests or




Figure 10.4: Transmissivity of a neutral density filter used in the calorimeter in dependence of the light's wavelength in two different positions on the filter (red and blue line). The grey band highlights the specified tolerance by the manufacturer for wavelengths between 400 to 700 nm.

10.4 Mechanical assembly

The full assembly of the light detection is located in a tube attached to the back of the calorimeter. To ensure that light mixer and optical filter are in the correct position a 3D-printed holder is fabricated which holds both the light mixer and the filter in position. A notch in the filter holder allows to attach a string to easily retract it from the tube for removing or inserting the filter in the assembly.

At the end of the tube a 3D-printed end-cap is fixed which holds both an SHV connector for the bias voltage and a BNC connector for the signal output. It is designed to be light tight to prevent any ambient light from altering the energy measurement results or even triggering false measurements. The light tightness of the printed holder is confirmed in a dark room using a torchlight with the PMT bias voltage raised up to the maximum voltage of 1800 V. In this test no influence on the noise or signal distributions could be observed. Three set screws are used to fix the end-cap to the tube which hold it in position even when attaching or removing the cables. To reduce the noise levels the end-cap is shielded by wrapping it with aluminium tape.

A computer rendering of the complete assembly with the tube and end-plug being shown in section view is displayed in Figure 10.5.




Figure 10.5: Rendering of the readout assembly for the calorimeter. The tube and the end-plug are cut open to allow visibility of all components. The components from left to right are: the filter holder (light blue) with both the filter (black) and the light mixer (light green), the photomultiplier connected to the voltage divider via the adapter PCB and the end-plug with the electrical connections.

11. | Calorimeter calibration system

As stated before, the calorimeter system of FASER consists of four spare modules from the LHCb outer electromagnetic calorimeter read out using one photo multiplier per module. To ensure a good and stable energy measurement, the calibration of the PMTs has to be checked from time to time. In experiments which expect a large number of charged particles from known processes this calibration can be done using measurements from those processes without having to rely on additional hardware. For example the energy calibration for photons and electrons in ATLAS is determined by Z boson decays into electron-positron pairs [65].

This is however not the possible for FASER as there is no high rate of known decays which could be used to calibrate the energy measurement. The only known particles which will cross the calorimeter with high rates are muons. However, as they are minimum ionizing their signals are quite small and can not be detected at the low gains the PMT will be operated at. The only way to use the muons to calibrate the energy is to run at higher gains and then extrapolate down to the nominal operating range. This requires a possibility to inject a known amount of light into the module while lowering the voltage in steps and tracking the change in amplitude. To realize this the modules have a calibration port in the front where a fibre can be attached as can be seen in Figure 8.5.

In this chapter the design of a calibration system for this purpose along with the accompanying components is discussed, which was designed and commissioned in the scope of this thesis. First, the overall design requirements are laid out. Next, the circuit for driving the Light emitting diode (LED) and the design of the calibration board is discussed. It is followed by a characterization and results from commissioning the board. Finally, the procedure for the calibration and results from testing the method are given.

11.1 Requirements

The most important feature for the calibration system is the ability to create short light pulses ($\mathcal{O}(10\text{ ns})$) with a variable but known amplitude. The light should be guided to the four calorimeter modules using optical fibres. The amplitude should be stable over time to reduce systematic errors in the calibration due to drifts in the amplitude.

The amount of light produced must be sufficient to create a measurable signal in the PMT with the optical filter installed, and the light fed to the calorimeter module over optical fibres. This must be fulfilled over the full gain range of the PMT demanding a high dynamic range of the calibration pulses. To independently monitor the pulses amplitude in situ originally it was planned to add an additional fibre to feed some light to a PIN photo diode. However, during prototyping the overhead in required equipment and needed development time for this solution was found to be too large compared with the gain in functionality.

Another important requirement is the connectivity of the system. It must be controlled over a standard ethernet connection to allow for connection to the control system from outside the tunnel without any special communication electronics. A simple Application Programming Interface (API) interfaced via standard HTTP is chosen to reduce the complexity of communicating with the system which enables sending commands to the board from basically any programming language without any special libraries. For simple testing it can even be controlled by just opening the website in a standard web browser.

In addition to the channel for the calibration of the calorimeter, a second channel is required to validate the performance of the PMTs in the various scintillator stations. This is less about the gain of those PMTs but more about checking the overall performance and testing the timing resolution as well as relative timing delays between the PMTs.

Regarding radiation hardness of the electronics or the optical fibres no special requirements are imposed as the background radiation at the FASER location is expected to be negligible. However, to ensure that no pulses are accidentally generated during normal data taking operations, the low voltage for the LED driver circuit is to be turned off after a calibration run has been finished.

11.2 LED driver circuit

Before it was decided to build a custom driver circuit existing implementations from other experiments have been inspected (e.g. LHCb, NA62, CALICE). However, all of those implementations were found to be too complex or not fulfilling the requirements. Therefore, a different approach is chosen for the driver circuit.

This section focusses on the general design of the driver circuit, the functional principle of the most important components and the first prototype for the driver circuit.

11.2.1 The Kapustinsky circuit

A popular method to generate ultra-short light pulses that is commonly used in experiments is a circuit known as Kapustinsky circuit which was proposed by J. S. Kapustinsky in [135] back in 1985. It features a very simple, low cost design that only needs a handful of components, but is still very flexible and can be tuned for different scenarios. The original circuit as presented in the original paper is shown in Figure 11.1.

To pulse the LED the circuit relies on a capacitor which is quickly discharged through a thyristor-like element formed by a matching pair of High-Frequency (HF) transistors. To reduce the duration of the pulse further an inductor is put in parallel with the LED which will produce a current opposite to the current from the discharging capacitor and therefore, reduces the decay constant of the light pulse.

In the original design the thyristor-like element is triggered using a negative square pulse riding on a negative DC offset. The negative DC bias of up to -24 V is used to charge the capacitor and together with the main capacitor defines the amplitude of the generated light pulse. The decay constant of the light pulse on the other hand can be tuned by changing the inductance in parallel to the LED.




Figure 11.1: Schematic for a LED driver circuit used to generate fast light pulses with durations of $\mathcal{O}$ (ns) and variable amplitudes as presented by Kapustinsky in [135]. The circuit is triggered by a square pulse riding on a negative DC offset applied to the bias input on the left.

11.2.2 Transistor + thyristor

Transistors are the most basic ‘active’ electronic components and are made from semiconducting materials. Semiconductors are materials with an electrical conductivity between conductors and insulators. When building semiconducting electronic components silicon is generally used as the semiconductor and a small amount of atoms from a different element is added to the material in a controlled way to modify the electrical behaviour. These atoms are called impurities. One differentiates between p-doped and n-doped semiconductors. In p-doped materials an element like Boron is added which has one electron less in the conduction band than silicon providing holes that can accept electrons and effectively allow for transport of positive charges through the material. For n-doping an element like Phosphorus is used which has an additional electron in the conduction band and therefore, increases the density of negative charge carriers. [136]

In the region where two differently doped materials meet, the holes and electrons combine and effectively a small region with no available charge carriers is created called depletion region. Applying a potential across the two layers with the higher potential at the n-doped side causes the holes in the p-doped layer and the electrons in the n-doped layer to be pulled away from the junction effectively increasing the depletion region and preventing nearly any current flow. This changes when the electrical potential is applied across the two layers with the polarity switched. In this configuration the potential at the n-doped side is lower than at the p-doped side causing the holes and electrons to be pushed towards the depletion region and thereby decreasing its size. When the potential is large enough the depletion region becomes sufficiently small to allow for current to flow through the junction. The potential at which the junction becomes conductive is called diffusion potential V_D . This behaviour is used in so-called diodes which allow current to flow only in one direction.

Transistors are made from three layers of semiconducting material with metal connectors on each layer for connecting to an electrical circuit. One differentiates between npn- and pnp-type

transistors. The names arise from the type of doping of the semiconducting layers. A npn-type transistor is built from two n-doped semiconducting layers sandwiched around a thin p-doped layer while the pnp-type transistor has two p-doped layers around a thin n-doped layer. The electronic symbols as well as the layer stack up for both types of transistors are shown in Figure 11.2. The number of impurities in the different layers differs with the emitter being doped stronger than the base and the base being doped stronger than the collector.

In the following the working principle of a npn-type transistor is explained, the same applies for pnp-type transistors with all polarities inverted to reflect the inverted doping of the layers. One can picture transistors as a set of two diodes connected in opposite directions which share one layer in the middle. Therefore, when an electrical potential is applied between the collector and emitter $V_{CE} > 0$ the transition region between the base and emitter is shrinking and current can flow between the two layers. The depletion region between the base and collector however, is increased and, hence, does not allow any current to flow between the collector and emitter. If in addition, a voltage larger than the diffusion potential of the junction is applied between the base and the emitter with $V_{BE} > V_D$ the junction between the base and emitter becomes conductive and electrons are flowing from the emitter into the base and holes from the base to the emitter. Due to the higher doping of the emitter the electron-current surpasses the current from the holes. Given the high number of electrons flowing into the base only few of them recombine with holes in the base while the other electrons can pass through the thin base layer reaching the collector. This means current is flowing between the collector and emitter.

The number of electrons flowing from the emitter into the base depends on the base-emitter voltage V_{BE} . Therefore, the current flow between collector and emitter can be controlled by varying the base-emitter voltage.




Figure 11.2: Electronic symbol and semiconductor layer stack up for npn-type transistors (left) and pnp-type transistors (right).

Using a npn-type and a pnp-type transistor with matching parameters connected in a special configuration forms a so-called thyristor. As shown in Figure 11.3 the collector of the pnp-type transistor is connected to the base of the npn-type transistor and the base is connected to the collector. This effectively creates a component with three p-n-junctions and three electrical contacts. When applying a positive potential V_{AC} between the anode and the cathode, the first junction is forward-biased while the second junction is reverse-biased and hence, no current can flow between anode and cathode in either direction. This holds true as long as V_{AC} is smaller than the breakdown voltage. By applying a positive voltage between gate and cathode the breakdown voltage is lowered substantially and a breakdown of the junction occurs even at low V_{AC} . Once the thyristor is switched to the conducting state it keeps this state as long as the potential difference V_{AC} is applied and a sufficient amount of current is flowing even when the voltage at the gate is removed. Once the potential difference V_{AC} is removed the thyristor switches to the off state and needs to be activated by applying a voltage at the gate once again. This behaviour makes a thyristor ideal to quickly allow to discharge a charge stored in a ca-

pacitor through a component. The current charging the capacitor must be smaller than the current needed to keep the thyristor conducting otherwise the thyristor will not fall back to the off state and prevents from charging the capacitor.




Figure 11.3: Electrical connection of a npn-type and a pnp-type transistor forming a thyristor.

11.2.3 First prototypes

To understand if the amount of light produced with a Kapustinsky style driver is sufficient prototypes have been built and tested with both a single PMT and with a calorimeter module at CERN. The first prototypes were identical to the Kapustinsky circuit except that the trigger line is separated from the bias voltage and is made accessible through a LEMO connector. One problem observed with the first prototypes was caused by instability in the lab power supply utilized for the bias voltage leading to uncontrollable fluctuations in the pulse amplitude. In the final driver circuit prototype this is mitigated using a precision voltage reference (REF01 [137]) outputting $10\text{ V} \pm 0.3\%$. By design the voltage reference is intended to be used as a positive voltage reference, however the current needed for biasing the driver circuit is small enough that it can be operated as a negative voltage reference by connecting the positive output to the boards ground and using the outputs ground connection as the negative pole. To vary the bias voltage a potentiometer is added in a voltage divider configuration allowing to set a bias voltage between 0 and 10 V. The schematic and photographs of this prototype are shown in Figure 11.4. To reduce the turnaround time the boards are isolation milled in the lab using a small CNC machine.

The main goal of this prototype is to ensure that the amount of light producible with this kind of driver is large enough to be detected by the PMT when being fed to the calibration port of a calorimeter module via an optical fibre. The setup and measured results from this test are shown in Figure 11.5. The plot shows the measured PMT response which has an amplitude of 1.5 V at a PMT bias voltage of -1300 V and a pulser bias voltage of $\approx 10\text{ V}$. In this test the optical filter is not installed which would reduce the pulse height to only 150 mV. The test is considered a success anyway as the fibre used in the test has a diameter of only $325\text{ }\mu\text{m}$ and is only coupled using a makeshift fibre holder. In addition, the driver circuit is only biased with 10 V while it can be biased with up to 24 V leading to higher amplitudes.




Figure 11.4: Schematic for the first prototype of a LED driver with adjustable bias voltage and the fabricated board.





Figure 11.5: Picture of the setup for testing the LED driver circuit (left) and the measured PMT response (red line) along with the trigger pulse (blue line) as measured with an oscilloscope (right).




Figure 11.6: Picture of the calorimeter calibration board with the fibre connectors removed and no enclosure.

11.3 Calibration board

The calibration board can be divided into multiple functional blocks as displayed in Figure 11.7. In the following each part of the calibration circuit will be explained in detail. Figure 11.6 shows a picture of the final calibration board.




Figure 11.7: Schematic block diagram of the various blocks in the calibration board. The arrows show the data flow/electrical connections. The power supply connections are not shown for clarity as it connects to all components.

Power supply

The calibration board needs bipolar power with positive polarity for driving the digital electronics and negative polarity for biasing the LED driver circuit. While the first prototype is supplied with 5 V and -24 V the later versions feature a switching power supply on the board producing 5 V to allow running it with a symmetrical ± 24 V power supply. This change is made to reduce the voltage drop over the wires from the FASER power supply box to the back of the detector where the calibration system is installed. A second linear voltage regulator is utilized to produce 3.3 V for the ethernet communication circuit. In total the system uses around 100 mA on the 24 V rail and 2 mA on the -24 V rail.

Main logic unit

The main logic unit is responsible for the control of the pulser and trigger and handling commands received via ethernet. As the controller, an 8-bit AVR microcontroller called ATmega328P manufactured by Microchip is used. This chip is used in a multitude of applications and therefore support and experience is widely available. It runs at a frequency of 16 MHz which is plenty given the task at hand.

The firmware for the main logic is written in the programming language C and compiled using an open source cross-platform compiler from the GNU project called `avg-gcc` [138]. It implements the communication protocols, controls the bias voltage and generates the trigger signals for the trigger outputs and the LED driver circuit.

Bias voltage

To control the negative bias voltage for the LED driver circuit a 12-bit DAC is used, which outputs a set voltage and can be digitally controlled from the main controller. It contains a resistor network with resistors in series between the supplied reference voltage and ground with digital switches to connect the output to one of the resulting voltage divider stages. This allows

setting a voltage between the reference voltage and ground in 4096 steps. The model used is a MCP4922 [139] by Microchip and features two independent outputs making it perfect to drive both driver circuits. In addition to the DAC, a second stage is needed as the output range of the DAC can only produce positive voltages smaller than the 5 V supply rail. The second stage is built from a rail-to-rail Operational Amplifier (OP) which is connected in an inverting amplifying mode with an amplification factor of ≈ -5.35 . This produces a bias voltage in the range from slightly under 0 V down to -21.94 V in steps of ≈ 5.35 mV. It can not reach exactly 0 V as the positive supply voltage of the OP is tied to ground. Originally this should have been tied to 5 V, but needed to be changed due to the OP drawing the 5 V rail below 0 V when being switched on before the switching power supply was stable. This triggered a fault detection mode in the switching power supply and prevented it from starting up. Consequently, when both power rails were plugged in at the same time the 5 V rail would stay at 0 V.

LED driver circuits

While the LED driver is in principle identical to the circuit described already in subsection 11.2.1, the calibration board now houses two independent drivers for the two output LEDs. The main difference is the selection of the transistors due to availability problems of HF PNP transistors which is no longer produced. The best match for a PNP transistor is the BFT93 [140], which proved to be a good replacement for the original choice. It should be noted however, that this model also went out of production and is only available from leftover stock in low quantity with no real alternative being available. This is not a huge problem for the FASER experiment where the number of channels is low but potentially rules out this style of pulse generator for larger experiments.

PIN photodiode feedback circuit

The original plan included adding a PIN photodiode as an additional means to measure the generated pulse amplitudes without the PMTs. It was included in the first prototype of the calibration board. The implementation there was however lacking proper filtering due to a design mistake and was not properly tuned for the photocurrent produced. Also, the amplitude of the signal from the photodiode was quite small requiring huge gains to make the signal easily measurable. Due to the complexity of optimizing the circuit to provide the required gains with low noise and the short timeline it was decided to drop this requirement and instead only ramp the voltage for one PMT while monitoring the pulse height with the other PMTs instead of relying on an additional photodiode.

11.3.1 Tuning the driver behaviour and LED selection

To evaluate the behaviour of the pulser circuit simulating the circuit with LTspice [141] is tested. Unfortunately, no suitable model for the LED is available and simulating the light emission of the LED is not possible. Due to those reasons the simulation was found to be insufficient to predict the performance of the driver circuit.

Therefore, the behaviour is tested by installing various capacitors and inductances in the circuit and the response is measured with a PMT. This is tested for all possible combinations for a range of values¹². The results from these measurements are shown in Figure 11.8. As can be seen the influence of the inductance on the pulse amplitude is minor at higher capacitances and plays a

¹ C = 100 pF, 120 pF, 180 pF, 220 pF, 180 pF, 330 pF, 1 nF, 10 nF

² 22 nF and L = 10 nH, 47 nH, 100 nH and 220 nH

larger role for small capacitances while the opposite is seen for the pulse area. This is because it mainly defines the pulses decay constant and hence, the duration of the pulse. In the final configuration a capacitance of 10 nF and an inductance of 220 nH is used. This configuration provides sufficient light yield with an acceptable overall pulse width.




Figure 11.8: Pulse amplitude (left) and pulse area (right) measured with a PMT with different capacitor and inductance values.

In addition to tuning the circuit for maximum amplitude multiple LEDs are evaluated to find a good signal shape and sufficient light yield. Five different LEDs are tested by directly placing it in the calibration board in front of a PMT. To allow for quick modifications of the board and test the setup, a light-tight box is built from wood covered with a light tight foil called Mavel Guard. This setup and the pulse shapes as measured by a PMT for various LEDs are shown in Figure 11.9.

As can be seen the LED used initially (OVL-5523) is producing the largest light yield with the HLMP-CB1A being slightly smaller with an acceptable pulse shape but with a larger second peak. Hence, the OVL-5523 is used in the final calibration system.

11.4 Fibre optics and mechanics

To transport the light from the calibration system to the calibration port of the calorimeters and the scintillator stations optical fibres are used. For the calorimeter custom fibre bundles are developed due to the large amount of light required there. As the gain in the scintillator stations is much higher, commercial fibre optics cables can be used. In the following sections the production of the fibres and the mechanical connections are discussed.

11.4.1 Calorimeter fibres

To maximize the light yield the fibre bundles for the calorimeter are built from 1 mm diameter cladded fibres. On the LED side four fibres are fixed into an adapter which is held in place using a captive nut with an M22 thread as displayed in Figure 11.10. The tapered fibre holder fits into the outer housing ensuring the fibre holder being centred on the LED and a light tight connection. Using the simple screw connection permits quick changes of the fibre bundles in




Figure 11.9: Photography of the light-tight box used for various measurements with the lid removed revealing the calibration board and the PMT (left) and the measured waveforms for various LEDs (right).

case of failures or for debugging purposes. All mechanical parts are fully 3D-printed from PETG using an FDM-printer.

For building the fibre bundle first four fibres are cut to length and using a dedicated tool the cladding is removed over 3 cm on one side of each fibre. This end is then glued into the fibre holder using a slow setting two part epoxy (Araldite 2011). Different glues were tested with Araldite 2011 showing the best results regarding uniform coupling in all fibres and good fixation of the fibres. To get the fibres as closely packed together as possible the bore in the fibre holder is tapered with only the last few mm being cylindrical. The cylindrical part is printed undersized and opened up with a drill to allow for tight tolerances.

The other end of the fibre is liberated from its cladding and first ground flat using sandpaper and subsequently polished using a jig for polishing fibres for the telecommunication industry. Afterwards, the screw connectors provided by LHCb fitting the calorimeter calibration port are glued onto the fibre cladding using a drop of cyanoacrylate glue.




Figure 11.10: Section view of the fibre connection at the LED side for the calorimeter fibre bundle.

11.4.2 Scintillator fibres

For the scintillator fibres commercially available fan-out cables are used to distribute the light to many output channels. Specifically, a cable is used which has an MTP-connector on one side where 12 fibres with 50 μm diameter are exposed in a row next to each other as seen in Figure 11.11. To couple this connector to the LED another 3D-printed holder is designed. It is a two part design consisting of a main housing with a recess for the MTP-connector and a lid clamping the connector in place and providing a light tight seal.

On the other end of the cable are twelve individual fibres with a single channel LC-connector each. This enables using commonly available fibre optic cables and couplings for the various connections between the splitter bundle and the individual scintillators. Each scintillator station has a dual channel LC-LC adapter mounted next to it to the upper frame where short fibres are connecting to the scintillators. Due to the various couplings and the small fibres the amount of light reaching the scintillators is much smaller than for the calorimeter. However, the PMTs are operated at higher gains compared to the calorimeter which compensates the lower light yield.




Figure 11.11: Picture of a MPO connector (left) and a schematic section view of the opened MPO holder used for coupling it to the LED on the calibration board (right).

11.5 Characterization and commissioning of calibration board

To get a good understanding on how the calibration system is behaving in respect of the overall performance and stability multiple test procedures are used. In this section first the test setup with all hardware used is explained in detail, following the findings from the different studies are discussed.

11.5.1 Measurement setup

For all studies the same overall setup is used. It consists of two PMTs, a HV power supply, two lab power supplies, an oscilloscope, a microcontroller with an environmental sensor and a computer. The two PMTs are of different models with one Hamamatsu R7899-20 with the custom

voltage divider discussed in section 10.2 and one Hamamatsu R2076 with a stock voltage divider. Each photomultiplier is housed in a metal tube with a 3D-printed cover on the front allowing to connect a calibration fibre with the same connector as used in the calorimeter modules. In front of the PMTs optical neutral density filters are placed similar to the one used in the final FASER setup. The signal and bias supply cables are fed out of the rear end of the tube covered with black tape to prevent any environmental light from entering the tube. The bias voltage for the PMTs is provided by a 6 channel negative voltage power supply (CAEN V6533). The signals from the PMTs are digitized using an oscilloscope, which is also calculating the integrated pulse signal and the amplitude. Most measurements are done using a Rigol DS1202Z-E.

Using a calorimeter fibre bundle the PMTs are connected to the calibration system. The power for the calibration board is provided by the two lab power supplies connected in series to achieve a bipolar supply with ± 24 V.

To monitor the environmental conditions an environmental sensor (Bosch BME280) is attached to a microprocessor. It provides measurements for the temperature, pressure and humidity. Using the microprocessor all values are read and stored at an interval of 10 s.

All devices are controlled or read out by a computer using a custom lab control system described in section B.2 and automated using various scripts. For most tests only the accumulated measurements are read from the oscilloscope instead of each individual pulse due to the slow communication protocol available. However, for most tests at least one waveform is fully read out from the oscilloscope and stored to review the shape of the signal.

11.5.2 Linearity of the bias voltage

As the bias voltage for the LED driver circuit is an important part of the calibration system the linearity is measured for one of the boards in each prototype generation. This is the only measurement performed manually due to the oscilloscope not being suited for accurately measuring voltages over such a large range of 0 to 24 V. Hence, wires are soldered to the output of the biasing circuit and using a digital multimeter³ the output voltage at different amplitude settings is measured. During these measurements only the bias voltage is enabled while the pulser itself is disabled. The voltage is measured at different DAC values for the amplitude in steps of 50 DAC counts between 0 and 4050. As can be seen in Figure 11.12, the bias circuit shows perfect linearity in both channels. Using a linear fit the parameters for the bias voltage are extracted which are in perfect agreement with the theoretically expected value of 5.357 mV per DAC step for nominal resistor values in the OP circuit. Small deviations of the expected slope are due to the resistor values being slightly off nominal. The difference observed between the two channels has no impact on the performance of the calibration system as the channels are used independently and no correlation between the two channels is required.

11.5.3 Pulse amplitude scans

The next test is aimed at understanding the pulse amplitudes at different pulser bias voltages. At each DAC value 1500 pulses are triggered and the average integrated charge and pulse height for each of the two PMTs is measured with the oscilloscope. The normalized signal amplitude and integrated charge in dependence of the pulser bias voltage are shown in Figure 11.13 where the conversion between the DAC to a voltage is done using the results from the linearity measurement. The normalization is performed to allow to show the measured values for both PMTs

³A Voltcraft VC175 is used with an accuracy of $\pm(0.8\% + 8 \text{ digits})$.




Figure 11.12: Measured pulser bias voltage on the final calibration board for channel A (left) and B (right) with a linear fit (red line). The uncertainties are determined by the accuracy of the multimeter.

in the same graph and easily compare the response shape⁴. It can be seen that the amount of light produced is following an s-shape where it is slowly increasing at low bias voltages, then rising quickly and falling of again at higher bias voltages. In the setup used for the characterization below a bias voltage of 3.7V no signal is observed, however, the minimum amplitude is limited by the capabilities of the oscilloscope. It can also be nicely seen that the response differs for the two different models of PMTs, however, the overall behaviour is similar. At a bias voltage of around 10 V the upper limit of light emission of the LED is nearly reached and the increase of light yield with higher bias voltages becomes smaller. While one could have limited the voltage and thus, increased the resolution of the bias voltage it was decided to keep the full range of 24 V in case of degradation of the LED light yield over time.

In the final calibration procedure it is currently planned to set the amplitude of the calibration pulses such that it is just within the bounds of the digitizer used for the readout at the highest PMT gain and keep it constant while ramping the PMT HV down. Therefore, it is not crucial to precisely know the relation between the pulser bias voltage and the amplitude. The measurements are useful nonetheless to test that the boards are working properly over the full range.

11.5.4 Frequency scans

The capacitor which is discharged through the LED to produce the light pulse is charged over a resistor. Hence, after a pulse is triggered the capacitor takes some time to be fully charged again. When the trigger frequency of the circuit is larger than the time it takes to charge the capacitor the amount of light emitted will decrease as well. To evaluate how strong the frequency dependence of the light output is in reality, a similar scan as described for the pulser bias voltage is performed where this time the pulse frequency is altered.

The measurements are performed in a frequency range from 1 Hz to 5 kHz with steps of 2 Hz up to 200 Hz and in steps of 50 Hz above. After setting the frequency the signal response is measured for 30 s and the average values are saved. For frequencies below 10 Hz the recording time is increased to 150 s to increase the number of pulses.

⁴The unnormalized graphs are given in the appendix on page 185.




Figure 11.13: Measured normalized signal amplitude (left) and integrated charge (right) for both PMTs at different pulser bias voltages. The values are normalized to allow for better comparison of the two PMTs to the value at a DAC setting of 4000.

In Figure 11.14 the measured normalized signal amplitude and integrated charge in dependence of the pulse frequency⁵. The normalization is performed using the values at a pulse frequency of 1 Hz. While the frequency response between the two channels of a board are similar at low frequencies, the behaviour is different at high frequencies. This can be explained by tolerances in the properties of the defining components which mainly are the LED and the inductance. However, as the differences only start at frequencies where the light yield is steeply decreasing this effect is not relevant for the experiment. The behaviour of the signal amplitude and integrated charge is slightly different with the integrated charge starting to fall at lower frequencies with respect to the amplitude. This is an expected trend as the amplitude is saturated above some bias voltage level where the pulse length and, hence, the integrated charge is still increasing with higher bias voltages.

At low frequencies the decrease in light yield is visible, but at a frequency of 200 Hz the fall of is less than 1%. This is negligible, but it should be noted that a full calibration sequence should be performed at a constant frequency to prevent introducing systematic effects from the small change.

11.5.5 Uniformity of the fibre bundles

As the fibre bundles are custom-built for the purpose the uniformity of the light output is not given by default. Due to the calibration procedure being performed relative to the initial signal in each calorimeter module a perfect matching of the fibres is not needed. However, having the light yield for each fibre being close to uniform prevents problems with one PMT being saturated with another PMT barely detecting the pulse at all.

To measure the difference in light yield each fibre is connected to the same PMT one after the other and the amplitude is measured. Using the same PMT reduces variations in gain, however, the bias voltage must be deactivated for the fibre to be switched, which can have a small effect on the measurement. This effect is expected to be small compared to the variations expected in the light yield of the fibres.

⁵The unnormalized graphs are given in the appendix on page 185.




Figure 11.14: Measured normalized signal amplitude (left) and integrated charge (right) for two channels of a calibration board at different pulse frequency settings. The values are normalized to allow for better comparison of the different channels to the value at a frequency of 1 Hz.

In Figure 11.15 the results of this test for a fibre bundle are shown. It can be seen that the light yield is not perfectly uniform along all fibres but sufficiently close. The difference in the integrated charge between the fibres with the lowest and highest light yield is $\approx 27.5\%$.




Figure 11.15: Measured waveforms (left) and integrated charges (right) for each fibre in the same PMT one at a time. The waveforms are shifted in time to fit them in a common plot.

11.5.6 Long-term stability

To evaluate the stability of the system and possible degradation of the LED performance due to the high currents the pulser is run at stable settings over a period of ≈ 357 hours and the signal in the PMTs is monitored. The pulser is run with a bias setting of 1200 DAC counts, a frequency of 100 Hz and the pulse heights and area are averaged in intervals of 10 s. In Figure 11.16 the integrated charge normalized to the average of the first measurements is shown along with the environmental temperature. For the normalization the first 20 measurements are

chosen to ensure a stable baseline. Some correlation can be seen between the environmental temperature, but it is obviously not the only parameter responsible for the drift. However, given that the amplitude not only depends on the amount of light produced but is also influenced by the PMTs' bias voltage and the stability of the oscilloscope's Analogue-to-digital converters (ADCs) no definite source for the drift can be named. As the deviation on the relative integrated charge is below 1.7% and it is only increasing over the measurement time no obvious degrading of the LED's performance is observed.




Figure 11.16: Measurement of the long-term stability of the calibration board with the pulser at a constant setting over ≈ 357 hours. The top plot shows the integrated charge normalized to the average of the first 20 measurements while the bottom plots shows the environmental temperature.

11.6 Test of the calibration system in a test beam

The calibration system performance and procedure was also tested with a high energy particle beam at a test beam in August 2021 at a beamline in the North Area at CERN. This beamline is supplied by a proton beam from the SPS which can produce a secondary beam of different particles by selecting appropriate targets. Using a set of chicanes and collimators the energy of the secondary beam can be adjusted.

The test beam setup is shown in Figure 11.17 and consists of a pair of trigger scintillators, the Interface Tracker (IFT), the preshower station and six calorimeter modules. Four of the six calorimeter modules are the selected modules for the FASER experiment and were extracted from the TI12 cavern for the test beam while the other two modules are spare modules. The DAQ is done using the same setup as in the final FASER setup with some minor changes to the configurations to account for the different channel mapping in the test beam setup. Three

different triggers are set in the test beam: coincidence of trigger scintillators, LED calibration pulser and a random trigger. Of those only the first two are used in the described analysis. The calibration pulser is set to continuously fire at a rate of 10 Hz.

All components are placed on a platform which can be moved in both axis perpendicular to the beam axis allowing to scan the full setup. During the test beam various measurements with this detector were taken to extract an energy calibration of the calorimeter and preshower as well as studying the performance of the tracker station. By now only preliminary results on the performance of the calibration system itself can be presented as no simulated samples of the setup are available yet and studies on the exact tracker alignment needed for the energy calibration are still ongoing. The focus of this section is therefore on showing that the calibration system is performing as expected and that the overall calibration procedure is feasible.

To be able to inject calibration pulses to all six calorimeter modules, the connector for the scintillator stations is removed from the calibration board and replaced by a second connector for calorimeter fibre bundles. While this means that not all modules are fed by the same LED it allows using the same fibre as in the final experiment setup. Thereby, the four main calorimeter modules are connected to the first channel and the remaining two are connected to the second channel.

To test the calibration procedure the energy response at the centre of each of the six calorimeter module is measured at different PMT bias voltages starting from the nominal setting for each PMT. At each position the bias voltage is set to 100, 200, 300, 400, 500, 550, 600, 650, 700, 750, 800 V below the nominal working point when reducing the voltage and to 700, 600, 500, 250 V on the way back up.




Figure 11.17: Rendering (left) and photograph (right) of the detector setup used for the combined calorimeter and tracker test beam.

11.6.1 Calibration procedure

The calibration procedure for the in-situ calibration follows a simple scheme. First the PMTs are operated at a sufficiently large gain to observe signals produced from muons crossing the calorimeter. After recording an adequate number of muon events, the calibration system is turned on with an amplitude to produce a large signal which is still safely within the range of the digitizer. Once the calibration pulses are observed to be stable the bias voltage for one PMT at a time is lowered in steps of around 100 V with the remaining PMTs being used as a reference to verify the stability of the calibration pulses. Using the integrated charge at each voltage step the gain can be extrapolated from the high voltage where the muons are observed

to the nominal operating voltage of the PMTs.

To finally convert the measured signal into an energy a cross calibration between the muon signals and measured responses at different electron energies is applied while taking into account the gain as measured by the LED calibration procedure. The model for cross calibrating is yet to be finalized, but will be extracted from the test beam measurements. Therefore, in the following studies instead of starting from muon signals the response is tested with electrons.

11.6.2 Analysis of the test beam

First, the recorded raw data is processed with the Calypso framework using a simplified track reconstruction method. The produced files are converted into NTuples to simplify and speed up the subsequent analysis procedure. The most important variable is the integral calculated from the fit to the waveform.

For each run the events triggered on the calibration system and the coincidence of the trigger scintillators (beam events) are treated separately. To select clean events a series of selection criteria is imposed on the events.

For the calibration events the following set of criteria are used:

- event triggered by the calibration system,
- remove events where the waveform fit did not succeed,
- veto against events with signal in the trigger scintillators.

For the beam events the selection criteria are:

- event triggered by the coincidence trigger,
- signal in the preshower scintillators,
- rectangular cut to ensure the track is pointing to the module under test.

For each run the mean of the integral is calculated and the $\text{RMS} / \sqrt{N}$ is used as the uncertainty on the integral. In addition, the values are determined using a fit with a Gaussian for the calibration pulses and a Crystal ball function. Both values agree within less than 1%. Hence, the simple mean method is used for the further analysis. To evaluate that the calibration system is working as expected the ratio between the response to the calibration system and the beam events can be evaluated. This is shown in Figure 11.18 where the ratio between the calibration and beam events is shown for the upper centre module at different PMT bias voltage settings. The histogram in the same figure shows the integrated charges measured in the other five modules where the voltage is kept constant during the measurement. It shows that the signal from the calibration system is constant with the largest deviation being below 0.5%. This shows that the pulses generated by the calibration system are very stable for the duration of a calibration run.

To extract the gain curve from the measured data the normalized integral is fitted with the expected gain curve from Equation 10.3 except for replacing $A = b^n$. The integral is normalized to the value at the nominal bias voltage of each PMT. This is chosen as the voltages have been previously adjusted to achieve the same Minimum Ionizing Particle (MIP) response.

The response to the calibration system and beam events in dependence of the PMT bias voltage for a single module as well as the fitted results for all modules are shown in Figure 11.19. It shows that the gain of the two spare modules (modules 0 and 5) is smaller than the gain of the other modules as expected from the initial commissioning. Only minor deviations between the observed and expected gain curve are visible at lowest bias voltages.




Figure 11.18: Ratio between the integrated charge for calibration pulses over the integrated charge for beam events at different PMT bias voltages of the upper centre module (left) and the integrated charge for calibration pulses relative to the integrated charge of the first run in the other five modules (right).




Figure 11.19: Normalized integral in dependence of the PMT bias voltage for the upper centre calorimeter module for LED pulses (blue) and beam events (red). In green a fit to the data for the beam events with a typical gain curve function is shown. The right plot shows the fit results for all modules. It can be clearly seen that the gain of the modules 0 and 5 is smaller than the gain of the other modules.

11.7 Conclusion

In this chapter the development of a scintillator and calorimeter calibration system for the FASER experiment is discussed. It includes the schematic and circuit design as well as firmware development. The driver circuit for flashing the LED is tuned to match the requirements of the FASER experiment and the full system is extensively studied to gain a good understanding of its behaviour. To connect the calibration system to the calorimeter custom fibre bundles are produced.

The overall performance during characterization and also during the first commissioning studies is found to be within the expected specifications. However, due to the LHC not being in operation, full in-situ calibration runs of the system have yet to be performed. In a combined tracker and calorimeter test beam campaign the performance of the calibration system is tested and found to be within the expected working parameters. The full energy calibration and the exact calibration sequence is yet to be finalized once the relevant accompanying studies are concluded.

12. | Preshower upgrade

Clearly detecting the signature of an ALP decaying into two photons with the currently installed FASER detector is difficult because both the preshower detector and the calorimeter are lacking any transversal segmentation. Hence, it is close to impossible to clearly discriminate between events containing two high energetic photons from a single photon produced by a different process. Possible processes producing background with potentially a single photon in the calorimeter are [111, 142]:

- high energetic neutrinos interacting with a neutron in the preshower material via Deep Inelastic Scattering (DIS),
- off-orbit protons from the LHC beam initiating particle showers close to the FASER detector,
- particles produced in beam-gas interactions by the beam entering the FASER detector from the back, where no substantial rock shielding is available.

In principle, knowing the expected number of events with high energy deposits in the calorimeter with no signal in the charged-particle veto, the scintillators and the tracker would allow determining limits on the ALP production cross-section even without clearly differentiating between single and di-photon events. However, the cross-sections for the background processes are not precisely known resulting in huge uncertainties on the expected number of background events.

Therefore, it would be advantageous to replace the simple preshower with a more powerful detector featuring transversal segmentation to allow separating one and two photon events. The discrimination between one or two photon events is to be done using a neural network. A possible design for a potential future upgrade of the preshower detector is studied in this thesis. The studies are based on simulated samples generated using a dedicated simulation setup which is tuned with data from an electron test beam.

This chapter opens with an explanation of the prototype design and the working principle of the utilized detector technology and readout. It is followed by the setup and results from the electron test beam used for tuning the simulation. In section 12.3 the simulation setup, the tuning of the simulation to data and a comparison between data and MC is given in detail. Afterwards, the neural network for discriminating between single and di-photon events is explained. The chapter closes with a conclusion summarizing the results and giving a prospect for further steps in the development of the preshower upgrade. The study of the preshower upgrade is published in [4].

12.1 Prototype design

With the preshower upgrade having to be a drop-in replacement of the currently installed simple preshower station, the size of the upgrade is limited to the space taken by the current preshower detector. This limits the longitudinal extent of the preshower upgrade to 240 mm. Because of this limitation, the preshower prototype is built from three layers of interleaved tungsten plates and active detector. To allow for varying the amount of absorber material, 1 mm thick tungsten plates are used and clamped into a 3D-printed holder in front of the detectors. A schematic drawing of the detector is shown in Figure 12.1. In this design, it would be in principle possible to add more layers of absorbers and active detectors. However, for the final detector most likely graphite blocks would be inserted in the currently free space to reduce backslash from the preshower reaching the tracker station directly next to it upstream.

As the active detector, gas detectors called micromegas are utilized. For the prototype, a set of small test chambers available from unrelated efforts with an active area of $90 \times 90 \text{ mm}^2$ are used.

Originally, it was planned to use $1X_0$ worth of material in front of the first two layers and $2X_0$ in front of the last layer. This was inspired by initial studies from another group. However, during simulation of the setup it was found that the photon conversion rate is improved significantly when doubling the amount of material in front of the first station. This improved geometry is therefore used in the development of the event discrimination algorithm.

The working principle of said detectors as well as the readout and acquisition setup are explained in the following two sections.




Figure 12.1: Rendering of the preshower prototype used for the test beam as seen from the front (left) and the side (right). The 3D-printed holder and clamps for the absorber plates are shown in green in the front view.

Micromegas detectors

Micromegas detectors can be seen as a successor of the famous multi-wire proportional chambers (MWPC) [143, 144]. Hence, to explain the working principle of micromegas, it is beneficial to first remember how MWPCs function.

A MWPC relies on a volume filled with a gas which can easily be ionized by crossing charged particles [145]. The volume is enclosed by conducting plates at the top and bottom and thin wires are stretched in between the plates parallel to each other at a distance of a few mm. By applying a high potential to the wires and grounding the plates, an electrical field is created within the gas volume. This causes electrons from an ionization process to drift towards the wires while the ions are drifting towards the cathode-plates. In the direct vicinity of the wires, the electrical field is sufficiently large to accelerate the electrons to high enough energies that they ionize the gas once again and create more free electrons. With these new electrons being accelerated, the process repeats and an avalanche of electrons is produced. The drifting electrons produce a measurable signal in the wire allowing for detecting the passage of a charged particle through the detector.

Instead of stretching wires in a gas volume, in micromegas detectors the anode-wires are replaced by copper traces etched on the surface of a PCB with a pitch in the order of $250\text{ }\mu\text{m}$. To achieve sufficiently high field strengths for producing avalanche amplification, the gas volume is separated into two field regions by a fine stainless steel mesh with a $128\text{ }\mu\text{m}$ high amplification region and a larger ($\mathcal{O}(\text{cm})$) drift region where the main ionization is to take place. The position of said mesh is ensured by pressing it onto small support pillars added on top of the PCB as shown in Figure 12.2. One common problem is the PCB bowing away from the mesh and, thereby, varying the amplification gap size. To prevent this, a common practice is to add a high precision flat support structure built from an aluminium honeycomb structure laminated in between two plates of blank PCB material. This lamination provides great stiffening with a relatively small material budget.




Figure 12.2: Schematic drawing of a two-dimensional micromegas detector with a charged particle (teal line) passing the detector. The drift lines of the produced electrons are shown in red. The drawing is not to scale and in reality the number of electrons is much higher. For better clarity, the ions drifting back towards the cathode and mesh are left out.

By grounding the mesh and applying a negative voltage of 300 V to the cathode, a homogeneous electric field is produced between the mesh and the cathode preventing electron-ion-pairs from recombining and forcing the electrons to drift towards the mesh and the ions to drift towards the cathode. A positive voltage of $\approx 550\text{ V}$ is applied to the strips on the PCB generating a much stronger electric field between the PCB and the mesh with a field strength of $\approx 42\text{ kV cm}^{-1}$. The

arising field in the amplification region extends slightly through the holes of the fine mesh. This guides the electrons through the holes resulting in an electron transparency of $\approx 90\%$ although the optical transparency of the mesh is only 40% .

In the amplification region the field is sufficiently large to cause avalanches creating signals on the strips. The advantage of the small amplification region is the short path the ions have to travel before recombining at the mesh. This is important because the ions are roughly 1000 times heavier than the electrons and hence, have a much lower drift velocity. By choosing a small amplification gap size, the dead time of the detector is reduced, and the detector can be operated at high rates without charge up effects causing the electric field to degrade. However, the small gap size comes with the disadvantage that electric sparks can easily form between the strips and the mesh causing the high voltage to drop and hence, reduce the electric field strength. This is further complicated under the presence of small dust or dirt particles within the amplification region, imposing high demands on cleanliness during the assembly of such detectors. To mitigate the effect of discharges in newer generations of micromegas detectors the strips on the PCBs surface are covered with an insulating layer, followed by another set of so-called resistive strips made from a material with a high surface resistance in the order of 5 to $20\text{ M}\Omega\text{ cm}^{-1}$ [146]. This limits the current flow through a discharge and causes a local voltage drop on the resistive strip without tripping the power supply. Thus, the electric field is weakened only locally while most of the detector is still operational.

Another benefit of the resistive strips is that due to the high resistance, the charge inserted on the strip by the avalanche is only slowly dissipated to the ends of the strip. The moving charges induce measurable currents on the readout strips in the layer below due to inductive and capacitive coupling. By orientating the strips on the surface of the PCB (y-layer) perpendicular to the resistive strips and adding a second layer of strips parallel to the resistive strips (x-layer) in the inner layer of the PCB, a two-dimensional readout is possible.

For the detectors used in the prototype, the resistive strips and the strips in the x-layer are $200\text{ }\mu\text{m}$ wide while the strips in the y-layer have a width of only $80\text{ }\mu\text{m}$. For all layers, the pitch, meaning the centre-to-centre distance between two strips, is $250\text{ }\mu\text{m}$. The difference in strip width between the two readout layers is to compensate for the large distance of the x-strips to the resistive strips and hence, the coupling being weaker than for the y-layer [147]. In total each layer has 360 strips resulting in an active area of $90 \times 90\text{ mm}^2$.

The detectors are operated with a mixture of Argon: CO_2 with a ratio of 93:7. While argon is used due to its low ionization energy, the CO_2 is added as a so-called quencher. The quencher gas is added to absorb any photons emitted from excited atoms or the anode material to prevent them from producing secondary ionizations. For this purpose, a polyatomic gas is best suited as it can dissipate the absorbed energy through elastic collisions or dissociation.

Readout electronics

To read out the micromegas detectors, a data acquisition (DAQ) system based on the Scalable Readout System (SRS) is used, which was developed by the RD51 collaboration [148]. It consists of an ADC card and a Front-End Concentrator card (FEC). On receiving a trigger signal, the FEC card initiates the digitization process in the ADC card and assembles the event information. Using standard HDMI cables, multiple pre-amplifier cards are connected to the ADC card. This connection is used both for the power supply as well as the analogue and digital signals.

The pre-amplifier cards used in the prototype setup are called APV hybrid and are based on a custom analogue pipeline ASIC named APV25, originally developed for the CMS experiment [149].

It is designed for reading out 128 input channels. For each channel, the ASIC features a dedicated low-noise amplifier and pulse shaper. The APV hybrids come in two variants: a master and a slave version. Each master hybrid has the ability to control a slave hybrid and synchronizes all communications with the SRS.

As the name suggests, the analogue pipeline ASIC is not digitizing the signals on the pre-amplifier boards but instead integrates the charge in windows of 25 ns for every strip and stores it in an analogue buffer on the chip. On receiving a trigger signal from the SRS, up to 30 samples with 25 ns each are sent to the ADC card over an analogue pipeline. On the ADC card the received signals are demultiplexed and digitized using the ADCs. The digitized samples are sorted and assembled into events by the SRS and transmitted to the readout computer using the User Datagram Protocol (UDP) over a standard ethernet connection. The readout computer runs a software called `mmdaq`, also developed by RD51 [150], for receiving the events as well as a slow-control program for configuring the SRS and APV hybrids. The DAQ program takes care of translating the APV number and channel number to the corresponding physical strip location and detector. In addition, a per channel pedestal subtraction is performed based on pedestal information recorded with a random trigger. This ensures a good baseline for each individual channel suppressing electrical noise. The events are written to a tree stored in a standard ROOT-file. For each event, the local event time and event counter, the APV ID, detector ID, strip and the integrated charges are stored.

As shown in Figure 12.3 each detector needs six amplifier cards, meaning 18 APVs are necessary for a full readout. With a single SRS, however, only 16 APVs (8 master and 8 slave) can be used. As the available electronics does not support synchronized running of two SRS, the first detector layer is only read out with two APVs per layer, reducing the active area to $60 \times 60 \text{ mm}^2$. To keep the reduced area centred with respect to the subsequent layers, the first detector is shifted in the transverse plane by a few cm.




Figure 12.3: Photography of a micromegas detector seen from the front. Around the entrance window the printed tungsten plate holder is visible (green). The six APV chips are found on the top and left sides of the detector.

12.2 Electron test beam

In order to properly simulate the response of the prototype to high energetic electromagnetic particles, measured data is required to allow tuning the simulation accordingly. As during the long shutdown of the LHC and the pre-accelerators no particle accelerators delivering electrons with highest energies are available, the studies are performed at much lower energies. A test beam is conducted at the particle accelerator of the University of Mainz called Mainz Microtron (MAMI) in 2020 [151]. Data is recorded with different amounts of tungsten absorber plates at the maximum available energy at the test beam location of 855.1 MeV.

In the next two sections, the test beam setup and the analysis procedure are described in detail. It is followed by a section summarizing the results.

12.2.1 Test beam setup

For the test beam, the prototype assembly discussed in section 12.1 is augmented by a set of trigger scintillators to detect electrons entering the prototype detector. At MAMI the setup is placed at a specialized facility for detector tests where a piece of the beam pipe is removed. The system is clamped firmly to a movable table which allows to position the setup relative to the beam in the two axis perpendicular to the beam axis. The frame of the preshower prototype is put on wooden blocks to raise it high enough for the beam to pass through the centre of the detectors. The ability to move the setup precisely is important to temporarily move the detector out of the beam line during regularly performed accelerator machine optimizations, which require running with high beam currents. Exposing the detector to such high beam currents would introduce the risk of damaging the readout structure even with the high voltage disabled. Two cameras are placed to monitor the setup remotely. One camera is showing the movable table to enable operating it from the control room safely. The second camera is targeted on the exit window of the beam pipe where, at high currents, the beam is producing a visible point used to validate its alignment after machine optimization.

The micromegas detectors and trigger PMTs are powered using a CAEN V6533 power supply for the negative rail and a desktop power supply (CAEN DT1471HET) for the positive rail. The signal from the trigger PMTs is fed into a leading edge discriminator (CAEN V895). Using a programmable logic unit (CAEN V2495) and a custom firmware, the discriminated trigger signals are put into coincidences. All components are controlled from the computer using the LabControl project described in appendix B.2.

The gas is mixed on site using a precise mass-flow based mixing system being fed from an Argon and a CO₂ bottle. The gas is flushed constantly with a rate of ≈ 3 L/h. The constant flushing ensures that the gas composition and humidity are stable over time. The pressure in the chamber is only slightly above the environmental pressure to prevent gas from flowing into the detector from possible leaks in the seals.

Pictures of the full setup and a close view on the prototype preshower are displayed in Figure 12.4. It shows a full overview over the setup with the movable table as well as a detailed view of the prototype detector under test. In the detailed view, the two scintillators used for triggering are clearly visible.

12.2.2 Analysis procedure

The recorded test beam data is processed event-by-event using a python based analysis framework. To filter out signals arising from noise, a series of cuts is applied at strip level. For de-




Figure 12.4: Photographs showing the test beam setup at the MAMI accelerator. The left picture shows a complete overview of the setup with the beam direction being from left to right. On the right, a detailed view of the prototype detector from the front is shown with the trigger scintillators being clearly visible.

termining where a particle hit the detector, it is however not sufficient to look at the signal of single strips. This is because the avalanches produced by particles passing through the detector do not only create a signal in a single strip but rather produce signals in neighbouring strips where the signal is approximately Gaussian distributed around the actual hit position. Hence, to extract the actual hit positions, neighbouring strips with a charge deposit need to be combined to clusters. For alignment studies, the found clusters are then matched in between the layers of a detector to reproduce the two-dimensional information. Due to ambiguities with higher number of clusters, this is however only possible for low hit multiplicities and is only used for alignment studies performed to validate the analysis.

The analysis procedure for each event therefore consists out of the following steps:

1. read event information from ROOT file and collect the hit strips per detector layer,
2. remove all events where the maximum charge is lower than a specified threshold to filter empty events,
3. search for cluster candidates and select good quality clusters in each layer,
4. if hits can be matched between the two layers of a chamber, a two-dimensional hit map is constructed (only for reference runs).

The information available for every event consists of the charge measured on every strip in 27 time slices of 25 ns each. This allows a close inspection of the signal development over time. Figure 12.5 shows part of the raw data from the middle detector for a single event. The measured charge is encoded by colour and normalized such that the maximum charge observed is equal to 1. One can see two clearly separated clusters in both layers with noise on a single strip in the X-layer in strip 130. The signal shape in the Y-layer shows a characteristic V-shape with long tails which is produced by the charge dissipating slowly to the ends of the resistive strips crossing multiple strips in the Y-layer. This effect is not visible in the X-layer as it is oriented parallel to the resistive strips and hence, the signal is spanning over fewer strips than in the Y-layer.

A series of criteria are imposed on the raw strip charges to reject bad or empty events. To ensure the event is not empty, the maximum charge must pass a configurable threshold in one of the layers at least. To reject events where the waveform is only partially recorded due to the

trigger arriving either too early or too late, the time slice containing the maximum charge in each layer must be within 5 and 18. In addition, an upper limit is imposed on the maximum charge in the first time slice to enforce a stable baseline.




Figure 12.5: Raw data for the X-layer (left) and Y-layer (right) from the middle station detector for a single event. The charge is normalized such that the maximum observed charge is equal to 1.

To find clusters, the maximum charge on each strip is extracted to remove the time dependence and reduce the complexity of the cluster search. For low hit multiplicities with good separation, the clusters can be identified by just grouping all neighbouring strips with only one strip not carrying charge in between to take into account dead/low efficiency strips. This simple method, however, does not work if the clusters are close to each other and are partially overlapping.

To separate overlapping clusters, a different approach is taken using a peak finding method implemented in the ROOT analysis framework. With this peak finder the position of the peak maxima is located. To determine the extent of each cluster, the strips next to the peak maximum are iterated in both directions until either a neighbouring peak is reached or two consecutive strips are found with no observed charge. In case the iteration is stopped due to reaching a neighbouring peak, the strip with the smallest charge between the two peaks is taken as the edge of the cluster. In the latter case, the edge is given by the last strip carrying charge.

The cluster candidates are required to contain at least two strips and, as a baseline, at most 10 strips in the X- and 20 strips in the Y-layer. Given that the cluster width is influenced by the gain and the exact detector properties, the parameter for the maximum cluster width is tuned for the different detectors. Additionally, a lower limit is imposed on the maximum charge in the cluster as well as the integrated charge of the cluster. This eliminates clusters consisting of noise.

In Figure 12.6 the maximum charges per strip are shown for a single event in the Y-layer of the middle detector. Each cluster found by the cluster finder is drawn in a different colour. The cluster candidate around strip 121 is not coloured as it does not pass the cluster selection criteria on account of the charge being below the required threshold.

The final step is matching the clusters in the two layers of each detector. However, with two hits in the detector, already four possible hit positions are found at the four positions where the strips cross. This can be resolved to some extent by comparing the charge ratio of the




Figure 12.6: Maximum charges on each strip in the Y-layer of the middle detector in an event from the test beam with eight identified clusters (each drawn in a different colour) demonstrating the separation of overlapping clusters. The two strips at strip position 121 are from noise and do not pass the cluster selection.

clusters in the X- and Y-layer. For detector assemblies designed for track reconstruction, more than three detectors, this could be resolved by initially allowing all possible combinations and subsequently using more advanced track finding algorithms to find the most probable track candidates [152].

As the preshower setup only consists of three detectors and the shower particles are mostly seen in only one or two detectors, none of those methods are suitable for this setup. Therefore, for the preshower studies, only the charge deposits per layer are evaluated without any information on matching between X- and Y-layer. For alignment and performance studies, only events with exactly one hit in each detector are selected to include only good quality tracks.

12.2.3 Resolution and alignment

Although perfect alignment is not necessary for the comparison between data and simulation, alignment and detector resolution studies are performed to test the performance of the detectors and the analysis framework. These studies use measurements with no tungsten absorber plates installed in front of the detectors.

Having three detectors in a telescope configuration with straight particle trajectories allows for an easy approximation of the detector's resolution. The method relies on calculating the expected hit position $\bar{x}_2$ in a detector under test by extrapolating the linear trajectory using two detectors as reference as shown in Figure 12.7. With x_1 and x_3 being the measured hit positions in the reference chambers and d_{12} and d_{13} being the distance between the respective chambers along the beam axis, the expected position for a straight trajectory is given by:

$$\bar{x}_2 = \frac{x_3 - x_1}{d_{13}} \cdot d_{12} + x_1. \quad (12.1)$$

Consequently the residual δ between the measured position x_2 and the expected position $\bar{x}_2$ is given by:

$$\delta = x_2 - \bar{x}_2 = x_2 - x_1 \frac{d_{12}}{d_{13}} - x_3 \left(1 + \frac{d_{12}}{d_{13}} \right). \quad (12.2)$$




Figure 12.7: Schematic drawing showing a misaligned middle detector and the definition of the residual δ between the expected (yellow dots) and observed hit position (green dots). The residual is used both for alignment and resolution studies.

Evaluating this residual for every good quality event produces a distribution which for a straight track is expected to be Gaussian distributed with a width of $\Delta\delta$. To determine the exact resolution of each chamber, a setup with four detectors is required in which case a set of equations is derived by applying Gaussian error propagation to Equation 12.2 for each triplet combination. As in this measurement only three detectors are used, the resolution is approximated directly by the width of the residual distributions for every layer.

In reality, the track of the particles is not perfectly straight, but instead small changes in direction are expected caused by multiple scattering [153]. Due to this, the distribution of the residuals is not following a single Gaussian distribution but it is rather the sum of two Gaussian distributions with a common mean but different amplitudes and widths. To separate the effects from the detector resolution and the multiple scattering, the distribution is fitted with the sum of two Gaussians given by:

$$f(\delta) = A_{\text{core}} \exp\left(-\frac{(\delta - \mu)^2}{2\sigma_{\text{core}}^2}\right) + A_{\text{tail}} \exp\left(-\frac{(\delta - \mu)^2}{2\sigma_{\text{tail}}^2}\right), \quad (12.3)$$

where A_{core} and A_{tail} are the amplitudes and σ_{core} and σ_{tail} are the widths of the core and tail respectively while μ is the expected value of the residuals.

Performing this procedure for every detector yields a good approximation for each layer's resolution. The residuals with the fits are shown in Figure 12.8. All resolutions are found to be within 119 and 168 μm which is within the expected range. In the Y-layer a dual peak structure is visible in the residuals. This is caused by the beam being centred exactly on a dead strip in the first detector layer. This does not pose a problem though for the measurements with tungsten plates installed as the showers are more spread over the active area.

Using the test beam results, the alignment of the detectors can be extracted. The simplest alignment method is measuring shifts of detectors in the transverse plane. It assumes that the particles travel on a straight track through the detector and uses the distribution of the hit positions.




Figure 12.8: Residuals for the X-layers (top) and Y-layers (bottom) of the preshower prototype detector fitted with the sum of two Gaussians. The resolution of the detector is given by the width of the core function σ_{core} .

In principle the rotation of the detectors can be extracted as well. This is not done here as the detectors are ensured to be parallel by the support frame. The rotation around the beam axis is also set sufficiently precise before the test beam takes place.

12.2.4 Results

To test the impact of the amount of absorber material in front of the detectors on the shower development, multiple measurements are performed with different numbers of plates installed. The configurations range from a baseline with no plates installed up to the intended prototype configuration of approximately $1X_0$ each in front of the first two detectors and $2X_0$ in front of the last detector. With the 1 mm thick tungsten plates this corresponds to a configuration of 3 plates each in front of the first detectors and 6 plates in front of the last one.

In Figure 12.9 the beam widening in the three detector layers is shown for various amounts of absorber material. As the first two setups were run at different gains, the beam widening is normalized to the third measurement. The standard deviation of the cluster positions in each layer is used as a measure for the beam widening. As expected, it can be seen that the beam is widening with increasing amount of absorber material. The beam widening is compared with results from the detector simulation discussed in the next section. Reasonable agreement between the test beam data and the simulation is observed with the simulation slightly overestimating the beam widening.




Figure 12.9: The beam widening as determined by the standard deviation of the hit position normalized to the third run taken during the test beam for all three detectors for data and simulated samples. The amount of absorber material is the full absorber material in front of the detector layer.

12.3 Detector simulation

In order to simulate the expected behaviour of the preshower detector, a dedicated simulation chain is set up. It consists of multiple steps to produce raw signals in all detectors for various incidence particle configurations. Those samples of raw signals can then either be analysed by the analysis framework described in subsection 12.2.2 or converted into input files used for training the discrimination algorithm.

In a first step, the interaction of particles with the various materials is simulated. This is done using the **Geant4** simulation toolkit [74]. As a result of this simulation step, all energy deposits within the detector gas volumes are saved along with the position where the interactions take place. In addition, the output includes truth information on the initial particles and an association between hits and the corresponding initial particle. In the next step, those energy deposits are used to create the digitized raw signals. The last step is the conversion into the different formats used for machine learning or analysis with the analysis framework.

In the next two sections the simulation steps are explained in more detail. It is followed by the tuning of the simulation to data and a comparison between both. Finally, an overview over the various different samples generated for the neural network training or evaluation is given.

12.3.1 Material interaction

The basic idea of simulating the passage of particles through matter in **Geant4** is that the particle is moved through the material in discrete steps. In each step, the probability for various possible physics processes is calculated and based on the likeliness, either no interaction or one of the possible interactions is chosen. For that, the user has to define a set of processes to take into consideration. However, in most cases this does not need to be defined from scratch, but instead one chooses a so-called physics list which compiles a set of processes. **Geant4** comes with a range of such physics lists optimized for specific applications. For simulating detectors for high energy particles, and especially calorimeters, a physics list called **FTFP_BERT** is recommended [154]. It includes all relevant interactions and is thoroughly tested

and validated.

Detector geometry

As the possible interactions of the tracked particle highly depend on the material the particle is in, one of the most important aspects in the simulation is the correct description of the detector geometry. This not only includes the position and dimensions of the various detector parts but also requires correctly defining the material of the different elements.

In this simulation all parts of the detector are modelled as squares with $90 \times 90 \text{ mm}^2$. This simplifies designing the geometry to a one dimensional problem as only the thickness and position have to be specified. The geometry for the simulation of the preshower prototype as seen from the side is shown in Figure 12.10. Going from the front (left) to the rear (right) of the setup, the first component is a Mylar foil representing the exit window of the beam pipe. It is followed by the two trigger scintillators with a Mylar foil on each side for the light tight wrapping. Next, are the three detector layers which consist of the absorber plate, one Mylar foil for the entry window and another one for the cathode, the gas volume and the readout PCB. The last two detectors have a stiffener made from blank FR4 and aluminium honeycomb. All parts are placed into a volume filled with air which is approximated by 70 % nitrogen and 30 % at a density of 1.2041 kg/m^3 .

Precisely modelling the honeycomb structure would greatly increase the complexity of the geometry, massively increasing the computing time. As an approximation, the honeycomb is modelled as a solid block with the density set to only 1 % of that of aluminium to account for the honeycomb structure.




Figure 12.10: Geometry of the test beam setup used in the **Geant4** simulation. The Mylar foils are drawn thicker than they are in the simulation to make them more visible.

Primary particle generation

The first step for each event is the generation of one or multiple primary particles. In **Geant4** this is done by supplying a class responsible for generating primary particles. In this class a so called ParticleGun is configured to the wanted behaviour by setting its position, direction, energy and the kind of particle generated. As the properties of the primary particles changes for the different samples needed for the preshower studies, the primary generator is implemented to be easily configurable using the **Geant4** messenger protocol. This enables loading different macro files at the beginning of the simulation program to specify the desired generator settings.

The parameters of the particle generator not only support fixed values but also randomly generating values from a given range and distribution for each event. The direction for the particles is always set perpendicular to the detectors. This is chosen also for the di-photon simulation as the angle between the two photons from the axion decay is small enough to be neglected here.

Sensitive detector

After the particles are generated, **Geant4** takes care of stepping the particles through the previously defined geometry. Thereby, it produces information about each interaction with the air or detector materials. The only active parts in the prototype are the gas volumes of the three detectors. Hence, the interactions with all other parts are not of interest for the subsequent simulation. To store only interactions within the active material, the gas volumes are defined as “Sensitive Detectors” in **Geant4**. This causes **Geant4** to call a function every time a particle deposits energy in one of the gas volumes. The function has access to information about the current volume, the energy deposited, the type of the interaction and the position both in the global frame and in the local frame of the hit volume. This information is compiled into a Hit object stored in a collection. At the end of each event, all hits are collected and written to a tree stored in a ROOT file.

12.3.2

Hit digitization

The most accurate way of simulating the process in the micromegas detector is fully simulating the drift of the electrons and the gas amplification in the electric field using a dedicated gas transport simulation like **Garfield++** [155]. This uses a high resolution representation of the electric field and simulates the movement of the charged particles through the field including the production of secondary electrons. The resulting charges at the anode can then be used to simulate the signal formation and the response of the readout. This process obviously results in the most detailed detector simulation possible.

The downside of said approach is the extremely high requirement on computational resources with processing of a single incidence electron already taking several minutes to simulate. While such a setup makes sense for detailed research of the signal formation in the detector for a relatively low number of events, it is not feasible to simulate the full signal formation process for the preshower prototype studies presented here. This is due to the high hit multiplicity and the large number of events needed for training and evaluating the neural network.

As an alternative approach, the energy deposits from the **Geant4** simulation are converted directly into charge measurements in the readout electronic. This is achieved by sampling the event shapes from a large collection of clusters extracted from data. To make sure that the cluster samples are as clean as possible, only events with exactly one cluster in each of the six layers are used to build the samples. As the final preshower detector should use a slightly different readout system which does not sample the signal in multiple time slices but rather stores only the maximum charge for each strip, the cluster information extracted from data contains only the maximum charge on each strip normalized such that the integral over the charges is 1.

To simulate the behaviour of the different detector layers more accurately, for each of the six layers a unique collection is used. These collections are further filtered to find a cluster matching the properties of the hit as close as possible. This includes the offset of the hit from the centre of the closest strip to correctly incorporate the charge distribution on the

neighbouring strips. In addition, the deposited energy is matched to the integrated charge of the cluster in the collection. After filtering the collection for compatible clusters, one of the clusters is randomly selected. As the charge per strip in the cluster samples is normalized, it can then easily be scaled by multiplying the energy deposit with a gain factor. The resulting charges are added to the corresponding strips stored as a one dimensional array. The gain factor is specific to each detector with an additional factor to account for the difference in signal strength between the X- and Y-layer.

The digitization program is written in Python with focus put on performance. As a first optimization, `numpy` [156] arrays are used for all data structures. As `numpy` already uses native C types, this alone already greatly improves performance compared to bare Python. In addition, the data structures for the cluster samples are optimized to allow performing the filtering in bins reducing the amount of comparisons needed to find compatible clusters. As a last optimization, the main functions are compiled to machine code using Cython [157].

12.3.3 Tuning of the simulation

The beam characteristics in the **Geant4** simulation and parameters like gain factors used in the digitization are tuned by iteratively changing parameters and comparing the results to data. The comparison is performed after digitizing the raw hits from **Geant4** into charges on the strips, converting them into the **mmdaq** event format and, subsequently, running the micromegas analysis over the events.

Using this approach the final simulation results do not perfectly describe the expected data, but yield distributions sufficiently close to use the simulation chain for first studies. In Figure 12.11-12.13 the comparison between data and simulation for different quantities is shown for the original preshower setup with $1X_0$, $1X_0$, $2X_0$ of tungsten in front of the three detectors. The main deviation is found in the charge distribution which is mainly caused by the simulation modelling a slightly different readout system as well as not including a detailed noise simulation. As the front detector in the preshower prototype shows different behaviour than the other two layers, the focus during tuning is set on the last two layers. The different behaviour of the first detector can be explained by being operated with too high bias voltages during the test beam.

To probe the influence of the charge mismodelling, rescaled samples are generated where the strip charges are varied up and down by 10 %.

For samples with high energy particles, the factors representing the gain in the detectors are lowered to prevent the signals from saturating. The shape templates and the gain ratio between the X- and Y-layers are unchanged. This is equivalent to reducing the amplification voltage in a real micromegas detector.

12.3.4 Simulated samples

Multiple different samples are simulated for both tuning the simulation to match the measured data and for studying the behaviour at higher energies for training the neural network.

For the test beam simulation, the position of the generated electrons is set to be Gaussian distributed in the transversal plane around the centre of the detectors with a width of 0.5 mm and a fixed energy of 855.1 MeV corresponding to the values from the test beam. The distance of the particle gun to the first detector is set to 225 mm, which corresponds to the distance between the end of the beam pipe and the first layer in the test beam.




Figure 12.11: Comparison of the number of clusters per layer and event found in data (blue points) and simulation (red bars). The upper row shows the X-layers and the lower row shows the Y-layers for the detectors from front (left) to back (right).




Figure 12.12: Comparison of the number of strips per cluster for data (blue points) and simulation (red bars). The upper row shows the X-layers and the lower row shows the Y-layers for the detectors from front (left) to back (right).




Figure 12.13: Comparison of the total charge per cluster for data (blue line) and simulation (red bars). The upper row shows the X-layers and the lower row shows the Y-layers for the detectors from front (left) to back (right).

For the photon events at higher energies required for the neural network training and evaluation, the same simulation is used as for the test beam. However, now the particle origin is randomly sampled uniformly in the transversal plane within an area of $60 \times 60 \text{ mm}^2$ around the centre of the detectors. When generating di-photon events, the origin of the first photon is randomly sampled. Subsequently, a distance d in the configured range and an angle $\varphi \in [0, 2\pi]$ is randomly chosen from a uniform distribution to calculate the origin of the second photon using

$$\begin{pmatrix} x_2 \\ y_2 \end{pmatrix} = \begin{pmatrix} x_1 \\ y_1 \end{pmatrix} + d \cdot \begin{pmatrix} \sin \varphi \\ \cos \varphi \end{pmatrix}, \quad (12.4)$$

where $x_{1,2}$ and $y_{1,2}$ are the coordinates of the first and second photon respectively. This ensures that both the distances and positions are uniformly distributed within the specified range.

To generate the energies of the photons, multiple strategies are implemented and can be selected for the different scenarios. For the neural network training all quantities are sampled uniformly. For single photon events the energy is chosen from a range of 200 to 7000 GeV while for di-photon events the range is limited to 100 to 3500 GeV. This ensures that the total event energy is consistent for both classes of events. The distance between the photons in di-photon events is sampled in a range of 0 to 2 mm.

Samples for single and di-photon events are always produced and processed independently being merged only in the preparation of the samples for machine learning. In this merging process various requirements can be imposed on the events. This allows, for example, to only consider events where the photon(s) converts in the absorber material in front of the first detector based on truth information. During this merging step, the samples are randomly permuted to remove any potential patterns in adjacent events. Weights are applied for the di-photon events accounting for the non uniformly distributed total event energy.

The same amount of single and di-photon events are chosen to guarantee the same fraction of both event classes. This is a common practice in machine learning as it results in the weights being updated for both classes equally. The selected events are distributed into three files used for training, validation and testing as explained in subsection 3.4.2. The training sample contains 70 %, the validation sample 10 % and the testing sample 20 %.

As previously stated, the original detector geometry used $1X_0$ of tungsten in front of the first two layers and $2X_0$ in front of the last layer. Based on truth information, however, it was found that only 50.5 % of all photons convert within the first absorber layer with this configuration. As a consequence, the amount of material in front of the first layer is doubled to $2X_0$ increasing the conversion probability in the first layer to 74.3 %.

In total 6 million single photon and 12 million di-photon events are simulated for the training dataset. An additional 1 million single and 2 million di-photon events are generated for the evaluation of the physics reach. The higher number of di-photon events is chosen to compensate for the lower event yield when imposing a selection on the photon conversion and the larger parameter space when determining the discrimination efficiencies.

12.4 Neural network

The toolchain for training the neural network is similar to the setup used for the neural network used in the photon identification in the light-by-light scattering analysis discussed in section 6.5. It is again using **Keras** with **Tensorflow** as the back end. The architecture of the neural network, however, is slightly more complex due to the larger number of input variables.

The process of training the neural network, the tested architectures and the resulting performance are discussed in the following sections.

12.4.1 Training setup

As previously discussed, **Keras** is used for performing the neural network training. Due to the large amount of events in each sample and the many input variables in each event, special care has been taken to handle reading the data for training. While for smaller samples a common approach is to load all data into memory at the beginning of the training, this is not possible here. Instead, a library for reading the training samples from the hard drive on-the-fly is utilized. It supports running the training and data loading processes in parallel using multiple threads. In addition, it takes care of caching and pre-fetching a variable amount of samples to optimize for run-time and memory usage. The sequence is also randomized to prevent training each epoch with the same sequence of events leading to a better training performance.

To reduce the amount of processing required during loading of the samples, all normalizations and the equalization of the number of events per event class are performed in a separate step before the training. That way, the data can be loaded from disk and be used for training as is.

Binary cross-entropy as defined in Equation 3.6 is used as the loss function for training of the network. The performance of the network is evaluated with the accuracy defined in Equation 3.7. The training is run for up to 100 epochs with the best model being stored after each epoch.

12.4.2 Selection of training data

To train and evaluate the neural network, a clean dataset is needed. Not all photons interact in the absorber material and hence, some do not produce any hits in the micromegas detectors. Including such events in the neural network training would result in bad performance of the network as it would train on events which only have hits from a single photon but are labelled as a di-photon event or even contain no hits.

To ensure that the training consists mostly of events where the photons converted, at least 25 hits per photon are required across the three detectors on truth level. This requirement is not met by 10 % of photons and hence, removes 10 % of the single photon events and ≈ 20 % of di-photon events.

It should be noted that this selection is only imposed on the training and validation dataset while it is not included in the evaluation of the performance or physics reach.

12.4.3 Neural network architecture

The neural network has two separate inputs for the detector signals and the total event energy. The detector signals are input as a two-dimensional array with dimensions of 360×6 containing the charge deposits on each strip of each layer. As the shape of the showers depends on the energy of the particles and, in the FASER detector the total energy is provided by the calorimeter, the energy is included as an additional input. To protect from overfitting, each layer is followed by a dropout layer with a dropout rate of 0.1.

The first input of the network is accepting the detector signals. To find local patterns in the spatial information, the data is processed by a 1D convolution layer followed by a maximum pooling layer. As explained in chapter 3, this allows recognizing features independently of their location. To reduce the number of dimensions, a flattening layer is added which concatenates the neurons and reduces it to a single dimension.

Next, the second input, containing the total event energy, is concatenated with the results from the convolution. This is followed by a set of dense layers and an output layer with a single node.

The activation function for all layers is set to a rectified linear unit except for the output layer which uses a sigmoid activation function. This ensures that the output of the network is fixed in a range of 0 to 1 and pushed towards 0 and 1 due to the form of the sigmoid function.

The specific number of dense layers and the number of neurons per layer are determined again by using **Tune** [158] to test various different values in a defined search space. The search space here is limited by the long training duration of up to 6 hours per epoch.

The search space is including the number of convolutional layers and their kernel and filter size as well as the number and size of dense layers before and after adding the second input with the total event energy.

The final architecture is summarized in Table 12.1.

Table 12.1: Neural network architecture for the discrimination between single and di-photon events. The dropout rate of 0.1 results in 10 % of the connections being randomly dropped during training. As the parameters depend on the layer's type, empty cells are on purpose.

Layer type	Neurons	Kernel size	Filters	Activation function	Parameters
Input layer (charges)					
1D Convolution		14	20	ReLU	6800
1D Max Pooling					
Dropout					
Flatten					
Input layer (energy)					
Concatenate					
Dense	40			ReLU	576 080
Dropout					
Dense	20			ReLU	820
Dropout					
Dense	10			ReLU	210
Dropout					
Output	1			Sigmoid	11
Total					583 921

12.4.4 Performance evaluation

To study the discriminative power of the neural network, the signal efficiency and the background rejection are evaluated for various threshold values. By plotting the resulting values, the so-called Receiver operating characteristic (ROC) curve is obtained. It is commonly used in the optimization of binary classifier systems to quantify its quality. The larger the area under the curve, the better the classifier works. For a ROC curve with an area of 0.5, the classifier has no discriminative power and the decision is random.

The ROC curve for the preshower prototype is shown in Figure 12.14 for different distance ranges for two photons with an energy of 500 GeV each and for an asymmetric kinematic with energies of 800 GeV and 200 GeV respectively. Visually, the difference is minor but the area under the curve (AUC) for the symmetric event kinematic is slightly larger than for the asymmetric kinematic.

Comparing the ROC curves for the different photon distances, it becomes obvious that the network can identify di-photon events with a large photon separation much better than di-photon events with a small photon separation. This is expected as the closer the photons are, the more their showers are overlapping and consequently, the harder they can be separated.

To come to a decision on the best working point, the expected background needs to be taken into account in order to calculate the signal significance. As a detailed background simulation is not yet available, two working points are evaluated as an example. They are chosen to provide background rejections of 90 % and 99.5 % and are named *loose* and *tight*.




Figure 12.14: Background rejection over the signal efficiency for two photons with an energy of 500 GeV each (left) and for an asymmetric kinematic with energies of 800 GeV and 200 GeV respectively (right). While the difference is small, the AUC shows that the performance for symmetric events is slightly better than for asymmetric energy distributions.

Using these two working points, the efficiency in dependence of the event kinematics are evaluated. The event kinematic is defined in three dimensions: the energies of the two photons and the distance between the two.

Figure 12.15 shows the resulting efficiency maps in dependence of the two photons energies for two photon separation ranges for the *loose* working point. From comparing the two, it becomes obvious that the signal efficiency increases with larger di-photon separation. Evaluating the

efficiencies with the samples where the strip charges are systematically varied up or down by 10 % yields to a relative change of the signal efficiencies of at most 4 % and a relative change of the background rejection of 3.7 %.

Given that 10 % of the photons produce less than 25 hits across all detectors, the inefficiency in the event classification partially stems from one of the photons not producing any significant signals in the preshower.

Figure 12.16 illustrates the improvement in background rejection and the effect on the signal efficiency when increasing the discrimination threshold. It becomes obvious that the better background rejection comes at the price of a strong reduction of the signal efficiency.




Figure 12.15: Signal efficiency for di-photon events evaluated in dependence of the two photons energy in multiple bins for di-photon events with a separation of $d < 300 \mu\text{m}$ (left) and for $500 \mu\text{m} < d < 1000 \mu\text{m}$ (right) at an average background rejection of 90 %.




Figure 12.16: Background rejection in dependence of the single photon energy (left) and signal efficiency in dependence of di-photon distance (right) for both working points.

12.4.5 Layer importance

To gain some understanding into how the neural network decision is made, the method explained in section 3.5 is used to determine the influence of the input variables on the classification results. Due to the large set of input variables, this is performed only on a per layer level and only using 10000 events from the testing dataset. As can be seen in Figure 12.17, the decision is most sensitive to the contents of the last detectors. No consistent trend for the importance of the X and Y layer is visible. In the last detector, however, the X-layer is having a stronger influence on the decision than the Y layer. This could be explained by the cluster size in the X layer being smaller and henceforth, allowing for a better separation of hits at large hit multiplicities present at the last layer.

All together, in a next iteration at least one additional layer should be introduced to measure the shower profile at a higher longitudinal resolution. In addition, the large importance in the last detector suggests that including a detector with higher resolution and better 2D reconstruction at high multiplicities in the last layer of the preshower detector would be beneficial.




Figure 12.17: Relative importance of the detector layer inputs and the energy for the event classification. The inputs of the first detector have relatively low influence as the showers are small there. The most influential variables are from the layers from the second and last detector.

12.4.6 Physics reach

The most important metric to evaluate the abilities of the micromegas based preshower is the parameter space in which the detector is sensitive to ALPs. This can be estimated by taking the convolution of the expected number of events with the efficiency of the detector as extracted from simulation.

The ALP production is extensively studied and simulated in [159]. The generated datasets were kindly provided by the author and are used in this estimation. The dataset contains the expected number of events in dependence of the photon energies and their distance for a minimum separation of 200 μm and considers 1452 different combinations of the ALP mass m_a and coupling g_{aWW} :

$$m_a = 0.1 - 2 \text{ GeV} \quad \text{and} \quad g_{aWW} = 5 \cdot 10^{-7} - 8 \cdot 10^{-4} \text{ GeV}^{-1}. \quad (12.5)$$

The expected number of events for an integrated luminosity of 90 fb^{-1} is available in a three-

dimensional histogram with a binning of

$$E_{\gamma,1/2} = 0, 200, 300, 400, 500, 900, 1250, 1750, 2250, 3500 \text{ GeV}$$

$$d_{\gamma\gamma} = 0, 200, 300, 500, 1000, 2000 \text{ } \mu\text{m}.$$

Multiplying the expected number of events for each event kinematic with the corresponding efficiency yields the expected number of events observable with the preshower detector. This is repeated for each of the ALP mass and coupling points available and the integral is calculated to get the number of expected events for each model.

Based on the number of expected signal and background events, the physics potential can be evaluated. As a detailed background simulation is not available at the time, the following calculations are performed for the assumption of 10 and 200 background events for an integrated luminosity of 90 fb^{-1} . The reach is calculated for a signal-over-background significance of 3σ .

This translates to a requirement of observing at least $3 \cdot \sqrt{\epsilon \cdot N_{\text{background}}}$ signal events with a background efficiency of ϵ .

Three different scenarios are studied. The baseline is the current FASER setup without any preshower capable of discriminating single and di-photon events. This assumes a 100 % signal efficiency, but at the same time no background suppression ($\epsilon = 1$). The other two scenarios include the preshower prototype with the loose and tight working points providing a background rejection of 90 % ($\epsilon = 0.1$) and 99.5 % ($\epsilon = 0.005$), respectively.

The resulting sensitive regions in the ALP mass and coupling plane are shown in Figure 12.18 for the integrated luminosities expected for the partial LHC Run 3 (90 fb^{-1}) and for HL-LHC (3 ab^{-1}) assuming a minimum di-photon separation of $200 \text{ } \mu\text{m}$.

To improve sensitivity in this region, the classification efficiency at low photon distances would need to be improved on. Nevertheless, the region where the micromegas based preshower detector is sensitive is quite large given the simple and comparably cheap design.




Figure 12.18: Maximum sensitive region for 3σ signal significance in the ALP mass and coupling plane for integrated luminosities of 90 fb^{-1} (blue) and 3 ab^{-1} (red) assuming at most 10 / 333 (top) and 200 / 6660 (bottom) background events. The two working points of the preshower prototype reflecting background rejections of 90 % (dashed line) and 99.5 % (dotted line) are compared to the current FASER setup with no single/di-photon discrimination capabilities (solid line). As the expected number of events is only available for a minimum distance of $200 \mu\text{m}$, the reach is only calculated for this fiducial region. Already excluded regions are shown in grey and are taken from [160].

12.5 Conclusion and further prospects

First studies on a potential preshower upgrade and its performance are presented. It is shown that discriminating between single and di-photon events is possible with the micromegas based preshower prototype. This provides a good baseline for a possible preshower upgrade which can be produced at relatively low costs by reusing already existing detector designs. In addition, the studies reveal promising improvements for later iterations of the experiment.

The most important information stems from the last detector layer, as seen earlier in the evaluation of the influence of the different detector layers on the prediction results of the network. This observation can be explained by the shower being widest in this layer and the probability that photons converted before reaching the last layer is maximum. Consequently, it would be worth to experiment with adding one or two additional layers with $1X_0$ of material in front of each. To allow for additional layers while still fitting the preshower in the available space, the micromegas detector could be built with a shared gas volume as demonstrated in [161] and the absorber material be placed within the gas volume. When increasing the material of the preshower, great care has to be taken to understand the influence on energy resolution of the calorimeter.

Additionally, instead of only using detectors with a 2D-strip-readout, one or more detectors could be replaced with a detector featuring an anode segmented in pads. The implementation would be similar to [162] but with a higher granularity to increase the resolution. Due to the significant increase in required readout channels, this would result in significantly higher costs for the detector electronics. This could not be tested in the scope of this thesis as no test beam data is available to tune the simulation chain for such a pad detector design.

A second option for reducing ambiguities and henceforth, delivering more information to the neural network on the 2D-position could be adding readout planes with the strips oriented at a different angle as done in [163]. At such high multiplicities, both options might yield better cluster separation but the problem of ambiguously reconstructed hits would remain.

At the time of writing, the micromegas based solution is only considered as a fallback solution for FASER. The baseline preshower upgrade will be built using dedicated monolithic silicon pixel detectors which will have superior resolution and hit position reconstruction. However, the price of the monolithic silicon detector is magnitudes larger compared with a micromegas based solution.

13. | Summary

In the first part of this thesis, a search for Axion-like-particles using the ATLAS experiment is presented. The second part focusses on the FASER experiment. The design of the upper support frame, the design and testing of the calorimeter modules, the development of a calibration system for the calorimeter and contributions to the FASER tracker commissioning are presented. In addition, a potential detector upgrade is developed and studied to provide a better sensitivity to ALPs.

The search with the ATLAS experiment is based on measurements of photon scattering in Lead-Lead UPC events from the measuring campaigns in 2015 and 2018. The analysis involves a dedicated photon identification based on neural networks to select photons at lowest energies with high efficiency and background rejection (cf. section 6.5). As detailed in chapter 6, this measurement results in the first direct observation of light-by-light scattering. In the full data-set 97 signal event candidates are observed with an expectation of 45 signal and 27 background events.

The fiducial cross-section for light-by-light scattering is found to be $\sigma_{\text{fid}} = 120 \pm 17 \text{ (stat.)} \pm 13 \text{ (syst.)} \pm 4 \text{ (lumi.) nb}$. The observed cross-section is approximately 1.5 times higher compared with the theoretical predictions of $80 \pm 8 \text{ nb}$ [83] and $78 \pm 8 \text{ nb}$ [71] the observed cross-section is approximately 1.5 times higher, which corresponds to an excess of 1.72σ and 1.81σ , respectively. In addition, differential cross-sections are derived as a function of the absolute di-photon rapidity, di-photon invariant mass, average photon transverse momentum and absolute di-photon $|\cos\theta^*|$. These results are published in [2, 3].

The di-photon invariant mass spectrum is surveyed for a resonant production of Axions using shape information from simulated samples. No significant excess is found and hence, exclusion limits in the ALP mass-coupling plane are extracted as shown in Figure 6.34. Those limits are the most stringent limits for the mass range of 6 to 100 GeV published to date.

The observation of light-by-light scattering at the LHC demonstrates that it can be used as a high energy photon collider when operating with heavy ions. The measured excess in the fiducial cross-section suggests an extension of the heavy ion operation at the future LHC run time. As the uncertainties of the current results are dominated by the limited statistics, longer measuring campaigns should allow to significantly reduce the statistical uncertainties.

In the second part the development of a new experiment called FASER is laid out. It aims to search for new light and weakly interacting particles like ALPs. As the first physics data taking for the experiment will only start once the LHC resumes operation in 2022. The design and commissioning of various detector components as well as a possible upgrade are covered.

A detailed description of the quality assurance procedure for selecting the tracker modules to use from the available ATLAS spares and the subsequent tracker commissioning is given (cf. chapter 9). The testing of the calorimeter modules and development of a custom voltage

divider is presented (cf. chapter 10). A custom LED based calibration system is developed from scratch including the schematic and PCB design. In various measurements including a test beam the functionality of the calibration system is confirmed (cf. chapter 11).

Finally, the possibility of using Micromegas detectors in an upgrade of the preshower detector is investigated. For this a simulation toolchain is implemented based on and verified with results from a dedicated test beam campaign. It is shown that the proposed design is able to achieve a reasonable separation of single and di-photon events at high energies. This result is published in [4]. For di-photon events with a separation of $200\text{ }\mu\text{m}$ an efficiency of 32.4 % at a background rejection of 90 % is achieved. For a separation above $1000\text{ }\mu\text{m}$ the efficiency becomes 80 %. The efficiency is mainly limited by the resolution and the conversion probability. It could be further improved by adding one or two additional active layers with $1X_0$ of absorber material in front of each. The results suggest further studies, whether the replacement of one strip detector with a pixel detector improves the performance. Once opportunities for high energy test beams become available again, an improved version of the preshower prototype should be tested again to validate that the signal characteristics at lower gains match that of the simulation.

All installed components of the FASER detector including the calibration system function are ready for first beam in 2022. The foreseen and validated performance of FASER puts it into a good position to either discover new physics within the next few years or to further constrain the parameter space for multiple BSM models. In addition, with FASERv the experiment will deliver measurements of neutrinos at the highest human-made energies ever recorded, broadening the understanding of neutrino interactions.

Part III

Appendix

Acknowledgements

I would like to thank everyone who helped make this thesis possible!

First and foremost, I want to thank [REDACTED] for always being there and cheering me up, when things did not go as they should! I would also like to thank my family for always supporting me throughout my studies.

Special thanks goes to my supervisor, [REDACTED], for the guidance and support over the last 6 years. You always had an idea on how to proceed when I was stuck. I appreciate that I was able to work on so many different topics and initiate some projects by myself.

I would also like to thank all colleagues I had the pleasure of working with during my time in the group. Especially to [REDACTED], [REDACTED], [REDACTED] and [REDACTED]. Not only for discussions related to work but also for coffee- and lunch-breaks and team events. Furthermore, I want to thank [REDACTED] for his input on the support frame design and assistance in producing the support frame.

No less thanks goes to [REDACTED] and [REDACTED] from the light-by-light scattering analysis group for their patience and support during the analysis.

I also want to thank my colleagues from the FASER experiment for giving me the exceptional opportunity to work on such interesting projects and the wonderful experiences. I really enjoyed all work related to the experiment with the highlight being the installation in T112. Thanks to [REDACTED] for trusting me with various projects, to [REDACTED] and [REDACTED] for helping with all PMT related tasks and [REDACTED] for the dinners.

A. Light-by-light scattering




Figure A.1: Normalised photon shower shapes photon candidates radiated by an electron due to final-state radiation in $\gamma\gamma \rightarrow e^+e^-$ events with no photon PID imposed. Data (black points) is compared with simulation (blue) and the two are in good agreement.




Figure A.2: Normalised photon shower shapes from hard-bremsstrahlung photon candidates emitted by an electron in $\gamma\gamma \rightarrow e^+e^-$ events with no photon PID imposed. Data (black points) is compared with simulation (blue) and the two are in good agreement.




Figure A.3: Comparison of the alternative $\gamma\gamma \rightarrow \gamma\gamma$ signal sample used for the closure tests in the unfolding to the nominal $\gamma\gamma \rightarrow \gamma\gamma$ signal sample. Shown are the di-photon transverse momentum $p_T^{\gamma\gamma}$ (left) and the di-photon invariant mass $m^{\gamma\gamma}$ (right). Both are well in agreement as can be clearly seen by the ratio plot.




Figure A.4: Relative non-closure, statistical and combined uncertainty from the MC non-closure test for the di-photon rapidity, average photon p_T and di-photon $|\cos \theta^*|$ for different numbers of iterations.




Figure A.5: Relative non-closure for the data-driven closure test using the alternative signal sample instead of the nominal signal sample for the di-photon rapidity, average photon p_T and di-photon $|\cos \theta^*|$ for different numbers of iterations.




Figure A.6: Measured normalised differential cross sections as a function from absolute di-photon rapidity (upper left), di-photon invariant mass (upper right), average photon transverse momentum (lower left) and absolute di-photon $|\cos \theta^*|$ (lower right). The measured values for the cross-section are shown as points with the error bars showing the statistical uncertainty. The grey band is denoting the combination of systematic and statistical uncertainties. The measurement is compared with the prediction from **SuperChic 3.0** (red line).

B. | FASER

B.1 Calibration system measurements




Figure B.1: Measured signal amplitude (left) and integrated charge (right) for both PMTs at different pulser bias voltages.




Figure B.2: Measured signal amplitude (left) and integrated charge (right) for two channels of a calibration board at different pulse frequency settings.

B.2 Lab control

Many measurements related to the FASER related hardware design work and upgrade studies involved the usage of various lab equipment in the Lichtenberg Group lab in Mainz. Before the development of the lab control system presented here, all instruments had to be controlled manually at the device itself and some new instruments had no way to control them at all. This short chapter outlines the features and implementation details of the new control system.

Architecture and communication




Figure B.3: Diagram of the lab control architecture showing the MQTT Broker as the central node (blue) with two control clients (red) and one web client (orange) connected. The control clients themselves communicate with the devices directly connected to them (green).

The control system is built in a modular fashion allowing for multiple devices to send state information and receive controls from the main control server. This is required as the devices to be controlled are located at different locations in the laboratory and can not easily be connected to a single computer. The user interface should be accessible from multiple locations in the lab and support multiple users at once.

Figure B.3 shows a diagram of the control architecture with all devices currently setup in the lab. The full project is written in Python except for HTML and JavaScript used in the user interface.

Each computer with devices connected to it runs a control software and thereby acts as a control client. Each client is responsible for transmitting the current state for one or multiple connected devices. The information is shared between the control clients and server using a publish-subscribe network protocol called MQTT [164]. It is based on messages with two fields: a topic and a message content. The topic is a string with the possibility to group topics in a directory style fashion by using slashes in the name. The message itself is transmitted as binary and can thus store various data types. The messages are sent to the server, also called broker, which in turn distributes them to all connected clients. To ensure that the clients only receive

message relevant to them, each client has to register what topics it wants to be updated on reducing unnecessary transmissions.


To make the control system easily accessible from every computer in the lab the control system is available as a website. This enables monitoring and controlling all devices from a standard web browser. The website is dynamically created using a web framework called Flask [165]. The website then uses JavaScript to directly communicate with the MQTT broker to receive updated values from the instruments or send control commands. Figure B.4 shows a screenshot of the web interface. Each device configured has its own tab selectable at the top. In the displayed view the settings page of a high voltage power supply is open.



Figure B.4: Screenshot of the web interface for the lab control system. Each device has its own tab selectable on the top. The screenshot shows the tab for a high voltage power supply with two channels switched on.

State Information

As said all state information is transmitted over MQTT. The topic for the messages follows a fixed template to prevent name collisions and keep track of the various devices. The topics are always structured in the following form:



The device name is a unique identifier for each device in the system. The channel number is optional and only used for devices with multiple channels. Devices with multiple channels may also have global state variables, e.g. the measured temperature of the device or a device interrupt. The suffix `set` is appended to the topic for messages intended to change the state of the device. This exemplary message would update the current limit for channel 1 of the power supply named `caen_v6533_1` to the value specified in the message body.

The control client will acknowledge that it processed the message by sending out a message with the same topic except for the `set`-suffix removed. The reply always contains the read back value after trying to set the demanded value. This ensures that in case the new value is invalid the web interface will update to the value actually set.

The message with the updated value is received by all open instances of the control website leading to a nearly instantaneous update for all users. Given the flexible message architecture, the devices can not only be controlled through the web interface, but also from scripts enabling automated measurements. This is used in some of the measurements of the calibration system to set the desired bias voltages for the PMTs.

In addition, each control client updates the measured values for each device in a configurable interval defaulting to 5 s.

Archiving

Often it is quite useful to have access to the conditions of the setup at a given point in time. This helps significantly when investigating problems in the measured data as problems in the configuration or unexpected power consumption are easy to spot.

To realize an automated logging of all conditions a simple bridge is written which subscribes to all topics matching the device status updates and stores them into a database. As the database, an InfluxDB is used which is optimized at storing time series along with metadata. To monitor the live data or check the history a dashboard software called Grafana is used [166]. It can quickly display graphs for a given variable in the database, works in the browser and is easy to configure.

An example for a monitoring page is shown in Figure B.5. It includes the configured and measured values for all relevant systems: gas flow, detector bias voltages and the x-ray tube.




Figure B.5: Screenshot of a Grafana dashboard for monitoring the micromegas x-ray setup. The page shows the configured and measured values for all relevant systems: gas flow, detector bias voltages and the x-ray tube.

Bibliography

- [1] F. Neuhaus. 'Results from Commissioning and/or Prototypes for New Small (HL-)LHC Experiments'. In: LHCP2021 (2021), p. 320. DOI: 10.22323/1.397.0320.
- [2] ATLAS Collaboration. 'Observation of Light-by-Light Scattering in Ultraperipheral Pb+Pb Collisions with the ATLAS Detector'. In: *Physical Review Letters* 123.5 (31st July 2019), p. 052001. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.123.052001. arXiv: 1904.03536.
- [3] ATLAS Collaboration. 'Measurement of Light-by-Light Scattering and Search for Axion-like Particles with 2.2 Nb⁻¹ of Pb+Pb Data with the ATLAS Detector'. In: *Journal of High Energy Physics* 2021.11 (13th Oct. 2021), p. 050. DOI: 10.1007/JHEP03(2021)243. arXiv: 2008.05355.
- [4] F. Neuhaus, E. R. Cholis and M. Schott. *Detection of Collinear High Energetic Di-Photon Signatures with Micromegas Detectors*. 29th Nov. 2021. arXiv: 2111.14939.
- [5] M. Thomson. *Modern Particle Physics*. Cambridge, United Kingdom ; New York: Cambridge University Press, 2013. 554 pp. ISBN: 978-1-107-03426-6.
- [6] D. Galbraith. *UX: Standard Model of the Standard Model*. URL: <http://davidgalbraith.org/portfolio/ux-standard-model-of-the-standard-model/> (visited on 13/07/2021).
- [7] Particle Data Group et al. 'Review of Particle Physics'. In: *Progress of Theoretical and Experimental Physics* 2020.8 (14th Aug. 2020), p. 083C01. ISSN: 2050-3911. DOI: 10.1093/ptep/ptaa104.
- [8] ATLAS Collaboration. 'Observation of a New Particle in the Search for the Standard Model Higgs Boson with the ATLAS Detector at the LHC'. In: *Physics Letters B* 716.1 (Sept. 2012), pp. 1–29. ISSN: 03702693. DOI: 10.1016/j.physletb.2012.08.020. arXiv: 1207.7214.
- [9] CMS Collaboration. 'Observation of a New Boson at a Mass of 125 GeV with the CMS Experiment at the LHC'. In: *Physics Letters B* 716.1 (Sept. 2012), pp. 30–61. ISSN: 03702693. DOI: 10.1016/j.physletb.2012.08.021. arXiv: 1207.7235.
- [10] R. P. Feynman. 'The Theory of Positrons'. In: *Physical Review* 76.6 (15th Sept. 1949), pp. 749–759. ISSN: 0031-899X. DOI: 10.1103/PhysRev.76.749.
- [11] F. Zwicky. 'Die Rotverschiebung von Extragalaktischen Nebeln'. In: *Helvetica Physica Acta* 6 (Jan. 1933), pp. 110–127.
- [12] J. de Swart, G. Bertone and J. van Dongen. 'How Dark Matter Came to Matter'. In: *Nature Astronomy* 1.3 (Mar. 2017), p. 0059. ISSN: 2397-3366. DOI: 10.1038/s41550-017-0059. arXiv: 1703.00013.
- [13] M. Markevitch et al. 'Direct Constraints on the Dark Matter Self-Interaction Cross Section from the Merging Galaxy Cluster 1E 0657–56'. In: *The Astrophysical Journal* 606.2 (10th May 2004), pp. 819–824. ISSN: 0004-637X, 1538-4357. DOI: 10.1086/383178.
- [14] D. Clowe et al. 'A Direct Empirical Proof of the Existence of Dark Matter'. In: *The Astrophysical Journal* 648.2 (10th Sept. 2006), pp. L109–L113. ISSN: 0004-637X, 1538-4357. DOI: 10.1086/508162. arXiv: astro-ph/0608407.

- [15] Planck Collaboration et al. ‘Planck 2018 Results: I. Overview and the Cosmological Legacy of *Planck*’. In: *Astronomy & Astrophysics* 641 (Sept. 2020), A1. ISSN: 0004-6361, 1432-0746. DOI: 10.1051/0004-6361/201833880.
- [16] G. ’t Hooft. ‘Symmetry Breaking through Bell-Jackiw Anomalies’. In: *Physical Review Letters* 37:1 (5th July 1976), pp. 8–11. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.37.8.
- [17] G. ’t Hooft. ‘Computation of the Quantum Effects Due to a Four-Dimensional Pseudo-particle’. In: *Physical Review D* 14:12 (15th Dec. 1976), pp. 3432–3450. ISSN: 0556-2821. DOI: 10.1103/PhysRevD.14.3432.
- [18] J. M. Pendlebury et al. ‘A Revised Experimental Upper Limit on the Electric Dipole Moment of the Neutron’. In: *Physical Review D* 92:9 (4th Nov. 2015), p. 092003. ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.92.092003. arXiv: 1509.04411.
- [19] R. D. Peccei and H. R. Quinn. ‘CP Conservation in the Presence of Pseudoparticles’. In: *Physical Review Letters* 38:25 (20th June 1977), pp. 1440–1443. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.38.1440.
- [20] W. A. Bardeen, R. Peccei and T. Yanagida. ‘Constraints on Variant Axion Models’. In: *Nuclear Physics B* 279:3 (1987), pp. 401–428. ISSN: 0550-3213. DOI: 10.1016/0550-3213(87)90003-4.
- [21] Y. Asano et al. ‘Search for a Rare Decay Mode $K^+ \rightarrow \pi^+ \nu \bar{\nu}$ and Axion’. In: *Physics Letters B* 107:1-2 (Dec. 1981), pp. 159–162. ISSN: 03702693. DOI: 10.1016/0370-2693(81)91172-2.
- [22] J. E. Kim. ‘Weak-Interaction Singlet and Strong CP Invariance’. In: *Physical Review Letters* 43:2 (9th July 1979), pp. 103–107. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.43.103.
- [23] M. Shifman, A. Vainshtein and V. Zakharov. ‘Can Confinement Ensure Natural CP Invariance of Strong Interactions?’ In: *Nuclear Physics B* 166:3 (Apr. 1980), pp. 493–506. ISSN: 05503213. DOI: 10.1016/0550-3213(80)90209-6.
- [24] M. Dine, W. Fischler and M. Srednicki. ‘A Simple Solution to the Strong CP Problem with a Harmless Axion’. In: *Physics Letters B* 104:3 (Aug. 1981), pp. 199–202. ISSN: 03702693. DOI: 10.1016/0370-2693(81)90590-6.
- [25] A. R. Zhitnitsky. ‘On Possible Suppression of the Axion Hadron Interactions. (in Russian)’. In: *Soviet Journal of Nuclear Physics [translation of Yadernaya Fizika]* 31 (1980), p. 260.
- [26] R. D. Peccei. ‘The Strong CP Problem and Axions’. In: *Axions: Theory, Cosmology, and Experimental Searches*. Ed. by M. Kuster, G. Raffelt and B. Beltrán. Berlin, Heidelberg: Springer Berlin Heidelberg, 2008, pp. 3–17. ISBN: 978-3-540-73518-2. DOI: 10.1007/978-3-540-73518-2. arXiv: hep-ph/0607268. URL: <https://doi.org/10.1007/978-3-540-73518-2>.
- [27] S. L. Cheng, C. Q. Geng and W.-T. Ni. ‘Axion-Photon Couplings in Invisible Axion Models’. In: *Physical Review D* 52:5 (1st Sept. 1995), pp. 3132–3135. ISSN: 0556-2821. DOI: 10.1103/PhysRevD.52.3132. arXiv: hep-ph/9506295.
- [28] J. Preskill, M. B. Wise and F. Wilczek. ‘Cosmology of the Invisible Axion’. In: *Physics Letters B* 120:1-3 (Jan. 1983), pp. 127–132. ISSN: 03702693. DOI: 10.1016/0370-2693(83)90637-8.
- [29] L. D. Duffy and K. van Bibber. ‘Axions as Dark Matter Particles’. In: *New Journal of Physics* 11:10 (16th Oct. 2009), p. 105008. ISSN: 1367-2630. DOI: 10.1088/1367-2630/11/10/105008.
- [30] J. L. Feng, I. Galon, F. Kling and S. Trojanowski. ‘ALPs at FASER: The LHC as a Photon Beam Dump’. In: *Physical Review D* 98:5 (18th Sept. 2018), p. 055021. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.98.055021. arXiv: 1806.02348.

- [31] W. S. McCulloch and W. Pitts. 'A Logical Calculus of the Ideas Immanent in Nervous Activity'. In: *The Bulletin of Mathematical Biophysics* 5.4 (Dec. 1943), pp. 115–133. ISSN: 0007-4985, 1522-9602. DOI: 10.1007/BF02478259.
- [32] C. M. Bishop. *Pattern Recognition and Machine Learning*. Information Science and Statistics. New York: Springer, 2006. 738 pp. ISBN: 978-0-387-31073-2.
- [33] J. Schmidhuber. 'Deep Learning in Neural Networks: An Overview'. In: *Neural Networks* 61 (Jan. 2015), pp. 85–117. ISSN: 08936080. DOI: 10.1016/j.neunet.2014.09.003.
- [34] F. Rosenblatt. 'The Perceptron: A Probabilistic Model for Information Storage and Organization in the Brain.' In: *Psychological Review* 65.6 (1958), pp. 386–408. ISSN: 1939-1471, 0033-295X. DOI: 10.1037/h0042519.
- [35] M. Minsky and S. Papert. *Perceptrons: An Introduction to Computational Geometry*. Expanded ed. Cambridge, Mass: MIT Press, 1988. 292 pp. ISBN: 978-0-262-63111-2.
- [36] D. E. Rumelhart, J. L. McClelland and R. J. Williams. 'Learning Internal Representations by Error Propagation'. In: *Parallel Distributed Processing: Explorations in the Microstructure of Cognition*. Cambridge, Mass: MIT Press, 1986. ISBN: 978-0-262-18120-4 978-0-262-13218-3.
- [37] Keras Layers API. URL: <https://keras.io/api/layers/> (visited on 14/09/2021).
- [38] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever and R. Salakhutdinov. 'Dropout: A Simple Way to Prevent Neural Networks from Overfitting'. In: *Journal of Machine Learning Research* 15.56 (2014), pp. 1929–1958.
- [39] V. Nair and G. E. Hinton. 'Rectified Linear Units Improve Restricted Boltzmann Machines'. In: *Proceedings of the 27th International Conference on International Conference on Machine Learning*. ICML'10. Madison, WI, USA: Omnipress, 2010, pp. 807–814. ISBN: 978-1-60558-907-7.
- [40] A. Esteva et al. 'Dermatologist-Level Classification of Skin Cancer with Deep Neural Networks'. In: *Nature* 542.7639 (2nd Feb. 2017), pp. 115–118. ISSN: 0028-0836, 1476-4687. DOI: 10.1038/nature21056.
- [41] N. V. Patel. *Why Doctors Aren't Afraid of Better, More Efficient AI Diagnosing Cancer*. URL: <https://www.thedailybeast.com/why-doctors-arent-afraid-of-better-more-efficient-ai-diagnosing-cancer>.
- [42] F. Chollet. *Deep Learning with Python*. Shelter Island, New York: Manning Publications Co, 2018. 361 pp. ISBN: 978-1-61729-443-3.
- [43] D. P. Kingma and J. Ba. *Adam: A Method for Stochastic Optimization*. 29th Jan. 2017. arXiv: 1412.6980 [cs].
- [44] B. M. Greenwell, B. C. Boehmke and A. J. McCarthy. *A Simple and Effective Model-Based Variable Importance Measure*. 12th May 2018. arXiv: 1805.04755 [cs, stat].
- [45] A. Hermann, L. Belloni, J. Krige and E. O. for Nuclear Research, eds. *History of CERN*. Amsterdam ; New York: North-Holland Physics Pub, 1987. 1 p. ISBN: 978-0-444-87037-7 978-0-444-88207-3.
- [46] CERN. 'CERN Annual Report 2019'. In: (2020). DOI: 10.17181/ANNUALREPORT2019.
- [47] E. Mobs. 'The CERN Accelerator Complex - 2019. Complexe Des Accélérateurs Du CERN - 2019'. In: (July 2019).
- [48] M. Kłusek-Gawenda and A. Szczurek. 'Exclusive Muon-Pair Production in Ultrarelativistic Heavy-Ion Collisions: Realistic Nucleus Charge Form Factor and Differential Distributions'. In: *Physical Review C* 82.1 (19th July 2010), p. 014904. ISSN: 0556-2813, 1089-490X. DOI: 10.1103/PhysRevC.82.014904. arXiv: 1004.5521.

- [49] C. F. v. Weizsäcker. 'Ausstrahlung bei Stößen sehr schneller Elektronen'. In: *Zeitschrift für Physik* 88.9-10 (Sept. 1934), pp. 612–625. ISSN: 1434-6001, 1434-601X. DOI: 10.1007/BF01333110.
- [50] E. J. Williams. 'Nature of the High Energy Particles of Penetrating Radiation and Status of Ionization and Radiation Formulae'. In: *Physical Review* 45.10 (15th May 1934), pp. 729–730. ISSN: 0031-899X. DOI: 10.1103/PhysRev.45.729.
- [51] B. Muratori and W. Herr. 'Concept of Luminosity'. In: (2006). DOI: 10.5170/CERN-2006-002.361.
- [52] ATLAS Collaboration. 'Luminosity Determination in Pp Collisions at $\sqrt{s} = 8$ TeV Using the ATLAS Detector at the LHC'. In: *The European Physical Journal C* 76.12 (Dec. 2016), p. 653. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-016-4466-1. arXiv: 1608.03953.
- [53] J. Pequeno. 'Computer Generated Image of the Whole ATLAS Detector'. Mar. 2008. URL: <https://cds.cern.ch/record/1095924>.
- [54] ATLAS Collaboration. 'The ATLAS Experiment at the CERN Large Hadron Collider'. In: *Journal of Instrumentation* 3.08 (Aug. 2008), S08003–S08003. DOI: 10.1088/1748-0221/3/08/s08003.
- [55] ATLAS Collaboration. *Production and Integration of the ATLAS Insertable B-Layer*. 16th May 2018. DOI: 10.1088/1748-0221/13/05/T05008. arXiv: 1803.00844. URL: <http://arxiv.org/abs/1803.00844> (visited on 18/11/2021).
- [56] ATLAS Collaboration. 'Performance of the ATLAS Transition Radiation Tracker in Run 1 of the LHC: Tracker Properties'. In: *Journal of Instrumentation* 12.05 (May 2017), P05002–P05002. ISSN: 1748-0221. DOI: 10.1088/1748-0221/12/05/P05002.
- [57] E. Kröger. 'Transition Radiation, Cerenkov Radiation and Energy Losses of Relativistic Charged Particles Traversing Thin Foils at Oblique Incidence: Theoretical Calculations and Numerical Computations'. In: *Zeitschrift für Physik* 235.5 (Oct. 1970), pp. 403–421. ISSN: 1434-6001, 1434-601X. DOI: 10.1007/BF01394931.
- [58] ATLAS Collaboration. 'Electron and Photon Performance Measurements with the ATLAS Detector Using the 2015–2017 LHC Proton-Proton Collision Data'. In: *Journal of Instrumentation* 14.12 (Dec. 2019), P12006–P12006. ISSN: 1748-0221. DOI: 10.1088/1748-0221/14/12/P12006.
- [59] F. Sauli. 'Principles of Operation of Multiwire Proportional and Drift Chambers'. In: (1977). DOI: 10.5170/CERN-1977-009.
- [60] ATLAS Collaboration. 'Luminosity Determination in Pp Collisions at $\sqrt{s} = 8$ TeV Using the ATLAS Detector at the LHC'. In: *The European Physical Journal C* 76.12 (Dec. 2016), p. 653. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-016-4466-1.
- [61] G. Avoni et al. 'The New LUCID-2 Detector for Luminosity Measurement and Monitoring in ATLAS'. In: *Journal of Instrumentation* 13.07 (23rd July 2018), P07017–P07017. ISSN: 1748-0221. DOI: 10.1088/1748-0221/13/07/P07017.
- [62] R. Frühwirth. 'Application of Kalman Filtering to Track and Vertex Fitting'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 262.2-3 (Dec. 1987), pp. 444–450. ISSN: 01689002. DOI: 10.1016/0168-9002(87)90887-4.
- [63] ATLAS Collaboration. 'Muon Reconstruction Performance of the ATLAS Detector in Proton-Proton Collision Data at $\sqrt{s} = 13$ TeV'. In: *The European Physical Journal C* 76.5 (May 2016), p. 292. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-016-4120-y.

- [64] ATLAS Collaboration. ‘Measurement of the Photon Identification Efficiencies with the ATLAS Detector Using LHC Run 2 Data Collected in 2015 and 2016’. In: *The European Physical Journal C* 79.3 (Mar. 2019), p. 205. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-019-6650-6. arXiv: 1810.05087.
- [65] ATLAS Collaboration. ‘Electron and Photon Energy Calibration with the ATLAS Detector Using 2015-2016 LHC Proton-Proton Collision Data’. In: *Journal of Instrumentation* 14.03 (18th Mar. 2019), P03017–P03017. ISSN: 1748-0221. DOI: 10.1088/1748-0221/14/03/P03017. arXiv: 1812.03848.
- [66] ATLAS Collaboration. ‘Topological Cell Clustering in the ATLAS Calorimeters and Its Performance in LHC Run 1’. In: *The European Physical Journal C* 77.7 (July 2017), p. 490. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-017-5004-5. arXiv: 1603.02934.
- [67] M. Cacciari, G. P. Salam and G. Soyez. ‘The Anti- k_t Jet Clustering Algorithm’. In: *Journal of High Energy Physics* 2008.04 (16th Apr. 2008), pp. 063–063. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2008/04/063. arXiv: 0802.1189.
- [68] ATLAS Collaboration. ‘Determination of Jet Calibration and Energy Resolution in Proton-Proton Collisions at $\sqrt{s} = 8$ TeV Using the ATLAS Detector’. In: *The European Physical Journal C* 80.12 (Dec. 2020), p. 1104. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-020-08477-8. arXiv: 1910.04482.
- [69] ATLAS Collaboration. ‘Measurements of b -Jet Tagging Efficiency with the ATLAS Detector Using $t\bar{t}$ Events at $\sqrt{s} = 13$ TeV’. In: *Journal of High Energy Physics* 2018.8 (Aug. 2018), p. 89. ISSN: 1029-8479. DOI: 10.1007/JHEP08(2018)089. arXiv: 1805.01845.
- [70] A. R. Martínez and ATLAS Collaboration. ‘The Run-2 ATLAS Trigger System’. In: *Journal of Physics: Conference Series* 762 (Oct. 2016), p. 012003. ISSN: 1742-6588, 1742-6596. DOI: 10.1088/1742-6596/762/1/012003.
- [71] L. A. Harland-Lang, V. A. Khoze and M. G. Ryskin. ‘Exclusive LHC Physics with Heavy Ions: SuperChic 3’. In: *The European Physical Journal C* 79.1 (Jan. 2019), p. 39. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-018-6530-5.
- [72] S. R. Klein, J. Nystrand, J. Seger, Y. Gorbunov and J. Butterworth. ‘STARlight: A Monte Carlo Simulation Program for Ultra-Peripheral Collisions of Relativistic Ions’. In: *Computer Physics Communications* 212 (Mar. 2017), pp. 258–268. ISSN: 00104655. DOI: 10.1016/j.cpc.2016.10.016.
- [73] ATLAS Collaboration. ‘The ATLAS Simulation Infrastructure’. In: *The European Physical Journal C* 70.3 (Dec. 2010), pp. 823–874. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-010-1429-9. arXiv: 1005.4568.
- [74] S. Agostinelli et al. ‘Geant4—a Simulation Toolkit’. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 506.3 (July 2003), pp. 250–303. ISSN: 01689002. DOI: 10.1016/S0168-9002(03)01368-8.
- [75] ATLAS Collaboration. ‘The ATLAS Calorimeter Simulation FastCaloSim’. In: *Journal of Physics: Conference Series* 331.3 (23rd Dec. 2011), p. 032053. ISSN: 1742-6596. DOI: 10.1088/1742-6596/331/3/032053.
- [76] J. C. Maxwell. ‘VIII. A Dynamical Theory of the Electromagnetic Field’. In: *Philosophical Transactions of the Royal Society of London* 155 (31st Dec. 1865), pp. 459–512. ISSN: 0261-0523, 2053-9223. DOI: 10.1098/rstl.1865.0008.

- [77] R. S. Van Dyck, P. B. Schwinberg and H. G. Dehmelt. ‘New High-Precision Comparison of Electron and Positron g Factors’. In: *Physical Review Letters* 59.1 (6th July 1987), pp. 26–29. ISSN: 0031-9007. DOI: 10.1103/PhysRevLett.59.26.
- [78] H. N. Brown et al. ‘Precise Measurement of the Positive Muon Anomalous Magnetic Moment’. In: *Physical Review Letters* 86.11 (12th Mar. 2001), pp. 2227–2231. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.86.2227.
- [79] S. Knapen, T. Lin, H. K. Lou and T. Melia. ‘Searching for Axionlike Particles with Ultraperipheral Heavy-Ion Collisions’. In: *Physical Review Letters* 118.17 (26th Apr. 2017), p. 171801. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.118.171801.
- [80] D. d’Enterria and G. G. da Silveira. ‘Observing Light-by-Light Scattering at the Large Hadron Collider’. In: *Physical Review Letters* 111.8 (22nd Aug. 2013), p. 080405. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.111.080405.
- [81] CMS Collaboration. ‘Evidence for Light-by-Light Scattering and Searches for Axion-like Particles in Ultraperipheral PbPb Collisions at $\sqrt{s_{NN}} = 5.02$ TeV’. In: *Physics Letters B* 797 (Oct. 2019), p. 134826. ISSN: 03702693. DOI: 10.1016/j.physletb.2019.134826. arXiv: 1810.04602.
- [82] ATLAS Collaboration. ‘Evidence for Light-by-Light Scattering in Heavy-Ion Collisions with the ATLAS Detector at the LHC’. In: *Nature Physics* 13.9 (Sept. 2017), pp. 852–858. ISSN: 1745-2473, 1745-2481. DOI: 10.1038/nphys4208.
- [83] M. Kłusek-Gawenda, P. Lebedowicz and A. Szczurek. ‘Light-by-Light Scattering in Ultraperipheral Pb-Pb Collisions at Energies Available at the CERN Large Hadron Collider’. In: *Physical Review C* 93.4 (21st Apr. 2016), p. 044907. ISSN: 2469-9985, 2469-9993. DOI: 10.1103/PhysRevC.93.044907.
- [84] F. Rademakers et al. *ROOT - An Object Oriented Data Analysis Framework (v6.20/06)*. Version v6-20-06. Zenodo, 14th June 2020. DOI: 10.5281/ZENODO.3895852.
- [85] O. Tange. ‘GNU Parallel - the Command-Line Power Tool’. In: *login: The USENIX Magazine* 36.1 (Feb. 2011), pp. 42–47. DOI: 10.5281/zenodo.16303.
- [86] G. Van Rossum and F. L. Drake. *Python 3 Reference Manual*. Scotts Valley, CA: CreateSpace, 2009. ISBN: 1-4414-1269-7.
- [87] S. R. Klein and J. Nystrand. ‘Photoproduction of Quarkonium in Proton-Proton and Nucleus-Nucleus Collisions’. In: *Physical Review Letters* 92.14 (9th Apr. 2004), p. 142003. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.92.142003.
- [88] C. Azevedo, V. P. Gonçalves and B. D. Moreira. ‘Exclusive Dilepton Production in Ultraperipheral PbPb Collisions at the LHC’. In: *The European Physical Journal C* 79.5 (May 2019), p. 432. ISSN: 1434-6044, 1434-6052. DOI: 10.1140/epjc/s10052-019-6952-8.
- [89] Z. Bern, A. D. Freitas, A. Ghinculov, H. Wong and L. Dixon. ‘QCD and QED Corrections to Light-by-Light Scattering’. In: *Journal of High Energy Physics* 2001.11 (15th Nov. 2001), pp. 031–031. ISSN: 1029-8479. DOI: 10.1088/1126-6708/2001/11/031.
- [90] M. Kłusek-Gawenda, W. Schäfer and A. Szczurek. ‘Two-Gluon Exchange Contribution to Elastic $\Upsilon \rightarrow \Upsilon$ Scattering and Production of Two-Photons in Ultraperipheral Ultrarelativistic Heavy Ion and Proton-Proton Collisions’. In: *Physics Letters B* 761 (Oct. 2016), pp. 399–407. ISSN: 03702693. DOI: 10.1016/j.physletb.2016.08.059.
- [91] I. Grabowska-Bold et al. *Light-by-Light Scattering with 2018 Pb+Pb Data - Supporting Note*. ATL-COM-PHYS-2019-014. Geneva: CERN, Jan. 2019. URL: <https://cds.cern.ch/record/2653766>.

- [92] M. Dyndal et al. *Light-by-Light Scattering in Ultra-Peripheral Pb+Pb Collisions at $\sqrt{s_{NN}} = 5.02$ TeV – Supporting Note*. ATL-COM-PHYS-2016-857. Geneva: CERN, July 2016. URL: <https://cds.cern.ch/record/2195943>.
- [93] C. Grupen and B. Shwartz. *Particle Detectors*. 2nd ed. Cambridge: Cambridge University Press, 2008. ISBN: 978-0-511-53496-6. DOI: 10.1017/CB09780511534966. URL: <http://ebooks.cambridge.org/ref/id/CB09780511534966> (visited on 13/04/2021).
- [94] J. Saxon. ‘Discovery of the Higgs Boson, Measurements of Its Production, and a Search for Higgs Boson Pair Production’. July 2014. URL: <https://cds.cern.ch/record/1746004>.
- [95] A. Hoecker et al. *TMVA - Toolkit for Multivariate Data Analysis*. 7th July 2009. arXiv: [physics/0703039](https://arxiv.org/abs/physics/0703039).
- [96] M. Abadi et al. *TensorFlow: Large-scale Machine Learning on Heterogeneous Systems*. 2015. URL: <https://www.tensorflow.org/> (visited on 30/11/2021).
- [97] T. Davidek and ATLAS Collaboration. ‘ATLAS Tile Calorimeter Time Calibration, Monitoring and Performance’. In: *Journal of Physics: Conference Series* 928 (Nov. 2017), p. 012003. ISSN: 1742-6588, 1742-6596. DOI: 10.1088/1742-6596/928/1/012003.
- [98] G. D’Agostini. ‘A Multidimensional Unfolding Method Based on Bayes’ Theorem’. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 362.2-3 (Aug. 1995), pp. 487–498. ISSN: 01689002. DOI: 10.1016/0168-9002(95)00274-X.
- [99] A. A. Natale, C. G. Roldao and J. P. V. Carneiro. ‘Two-Photon Final States in Peripheral Heavy Ion Collisions’. In: *Physical Review C* 65.1 (30th Nov. 2001), p. 014902. ISSN: 0556-2813, 1089-490X. DOI: 10.1103/PhysRevC.65.014902. arXiv: [nuc1-th/0107034](https://arxiv.org/abs/nuc1-th/0107034).
- [100] L. Heinrich, M. Feickert and G. Stark. *Pyhf: V0.6.1*. Version 0.6.1. DOI: 10.5281/zenodo.1169739.
- [101] G. Cowan. *Statistical Data Analysis*. Oxford Science Publications. Oxford : New York: Clarendon Press ; Oxford University Press, 1998. 197 pp. ISBN: 978-0-19-850156-5 978-0-19-850155-8.
- [102] Belle II Collaboration et al. ‘Search for Axion-Like Particles Produced in e^+e^- Collisions at Belle II’. In: *Physical Review Letters* 125.16 (14th Oct. 2020), p. 161806. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.125.161806. arXiv: 2007.13071.
- [103] D. Aloni, C. Fanelli, Y. Soreq and M. Williams. ‘Photoproduction of Axion-like Particles’. In: *Physical Review Letters* 123.7 (12th Aug. 2019), p. 071801. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.123.071801. arXiv: 1903.03586.
- [104] M. Bauer, M. Neubert and A. Thamm. ‘Collider Probes of Axion-Like Particles’. In: *Journal of High Energy Physics* 2017.12 (Dec. 2017), p. 44. ISSN: 1029-8479. DOI: 10.1007/JHEP12(2017)044. arXiv: 1708.00443.
- [105] D. Banerjee et al. ‘Search for Axionlike and Scalar Particles with the NA64 Experiment’. In: *Physical Review Letters* 125.8 (17th Aug. 2020), p. 081801. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.125.081801. arXiv: 2005.02710.
- [106] ATLAS Collaboration. ‘Measurement of the Inelastic Proton-Proton Cross Section at $\sqrt{s} = 13$ TeV with the ATLAS Detector at the LHC’. In: *Physical Review Letters* 117.18 (26th Oct. 2016), p. 182002. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.117.182002. arXiv: 1606.02625.

- [107] H. Van Haevermaet. ‘Measurement of the Inelastic Cross Section at $\sqrt{s} = 13$ TeV’. In: *Proceedings of XXIV International Workshop on Deep-Inelastic Scattering and Related Subjects — PoS(DIS2016)*. XXIV International Workshop on Deep-Inelastic Scattering and Related Subjects. DESY Hamburg, Germany: Sissa Medialab, 9th Nov. 2016, p. 198. DOI: 10.22323/1.265.0198.
- [108] FASER Collaboration. *FASER: ForwArd Search ExpeRiment at the LHC*. 11th Jan. 2019. arXiv: 1901.04468.
- [109] FASER Collaboration. *Letter of Intent for FASER: ForwArd Search ExpeRiment at the LHC*. 26th Nov. 2018. arXiv: 1811.10243.
- [110] FASER Collaboration. *Technical Proposal for FASER: ForwArd Search ExpeRiment at the LHC*. 21st Dec. 2018. arXiv: 1812.09139.
- [111] FASER Collaboration. ‘FASER’s Physics Reach for Long-Lived Particles’. In: *Physical Review D* 99.9 (15th May 2019), p. 095011. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.99.095011. arXiv: 1811.12522.
- [112] B. Döbrich, J. Jaeckel, F. Kahlhoefer, A. Ringwald and K. Schmidt-Hoberg. ‘ALPtraum: ALP Production in Proton Beam Dump Experiments’. In: *Journal of High Energy Physics* 2016.2 (Feb. 2016), p. 18. ISSN: 1029-8479. DOI: 10.1007/JHEP02(2016)018.
- [113] S. L. Glashow, J. Iliopoulos and L. Maiani. ‘Weak Interactions with Lepton-Hadron Symmetry’. In: *Physical Review D* 2.7 (1st Oct. 1970), pp. 1285–1292. ISSN: 0556-2821. DOI: 10.1103/PhysRevD.2.1285.
- [114] P. Ball and R. Zwicky. ‘New Results on $B \rightarrow \pi, K, \eta$ Decay Form Factors from Light-Cone Sum Rules’. In: *Physical Review D* 71.1 (18th Jan. 2005), p. 014015. ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.71.014015.
- [115] E. Izaguirre, T. Lin and B. Shuve. ‘Searching for Axionlike Particles in Flavor-Changing Neutral Current Processes’. In: *Physical Review Letters* 118.11 (15th Mar. 2017), p. 111802. ISSN: 0031-9007, 1079-7114. DOI: 10.1103/PhysRevLett.118.111802.
- [116] Eljen Technology. *General Purpose Plastic Scintillator: EJ-200, EJ-204, EJ-208, EJ-212*. URL: https://eljentechnology.com/images/products/data_sheets/EJ-200_EJ-204_EJ-208_EJ-212.pdf.
- [117] *Photomultiplier Tube Assembly H6410*. Photomultiplier tube assembly H6410. URL: <https://www.hamamatsu.com/us/en/product/type/H6410/index.html> (visited on 08/10/2020).
- [118] A. Benvenuti et al. ‘A Shashlik Calorimeter with Longitudinal Segmentation for a Linear Collider’. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 461.1-3 (Apr. 2001), pp. 373–375. ISSN: 01689002. DOI: 10.1016/S0168-9002(00)01247-X.
- [119] S. Amato et al. *LHCb Calorimeters: Technical Design Report*. Technical Design Report LHCb CERN-LHCC-2000-0036-[sic!] Geneva: CERN, 2000. URL: <https://cds.cern.ch/record/494264>.
- [120] ATLAS Collaboration. *Inner Detector: Technical Design Report*. S.n.: ATLAS Inner Detector Community, 1997. ISBN: 978-92-9083-102-0. URL: <http://cds.cern.ch/record/331063/files/ATLAS-TDR-4-Volume-I.pdf?version=1>.
- [121] K. Halbach. ‘Design of Permanent Multipole Magnets with Oriented Rare Earth Cobalt Material’. In: *Nuclear Instruments and Methods* 169.1 (Feb. 1980), pp. 1–10. ISSN: 0029554X. DOI: 10.1016/0029-554X(80)90094-4.

- [122] SCT Barrel Module. URL: <https://jsdhp1.kek.jp/~unno/SCTbarrelmod.html> (visited on 06/07/2021).
- [123] FASER Collaboration. *The Trigger and Data Acquisition System of the FASER Experiment*. 28th Oct. 2021. arXiv: 2110.15186 [hep-ex, physics:physics]. URL: <http://arxiv.org/abs/2110.15186>.
- [124] FASER Collaboration. *Calypso - FASER Offline Software System*.
- [125] ATLAS Collaboration. 'The Barrel Modules of the ATLAS Semiconductor Tracker'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 568.2 (Dec. 2006), pp. 642–671. ISSN: 01689002. DOI: 10.1016/j.nima.2006.08.036.
- [126] F. Campabadal et al. 'Design and Performance of the ABCD3TA ASIC for Readout of Silicon Strip Detectors in the ATLAS Semiconductor Tracker'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 552.3 (Nov. 2005), pp. 292–328. ISSN: 01689002. DOI: 10.1016/j.nima.2005.07.002.
- [127] F. Keizer, A. Gorbach, M. Parker, C. Steer and S. Wotton. 'A Compact, High Resolution Tracker for Cosmic Ray Muon Scattering Tomography Using Semiconductor Sensors'. In: *Journal of Instrumentation* 13.10 (23rd Oct. 2018), P10028–P10028. ISSN: 1748-0221. DOI: 10.1088/1748-0221/13/10/P10028.
- [128] Keysight. *Keysight E364xA Dual Output DC Power Supplies*. URL: <https://www.keysight.com/de/de/assets/9018-01166/user-manuals/9018-01166.pdf> (visited on 26/05/2021).
- [129] Keithley. *Series 2400 SourceMeter*. URL: https://download.tek.com/manual/2400S-900-01_K-Sep2011_User.pdf (visited on 26/05/2021).
- [130] C. Jermain et al. *Ralph-Group/PyMeasure: PyMeasure 0.8*. Version v0.8. Zenodo, 29th Mar. 2020. DOI: 10.5281/ZENODO.3732545.
- [131] Hamamatsu. *Photomultiplier Tubes - Basics and Applications*. URL: https://www.hamamatsu.com/resources/pdf/etd/PMT_handbook_v3aE.pdf (visited on 08/06/2021).
- [132] Hamamatsu. *Photomultiplier Tubes and Assemblies*. URL: https://www.hamamatsu.com/resources/pdf/etd/High_energy_PMT_TPMZ0003E.pdf (visited on 08/06/2021).
- [133] Y. Gilitsky et al. 'LHCb Calorimeters High Voltage System'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 571.1-2 (Feb. 2007), pp. 294–299. ISSN: 01689002. DOI: 10.1016/j.nima.2006.10.086.
- [134] Edmund Optics. *OD 1.0 VIS, 25mm Dia. Non-Reflective ND Filter*. URL: <https://www.edmundoptics.com/p/od-10-vis-25mm-dia-non-reflective-nd-filter/18001/> (visited on 10/06/2021).
- [135] J. Kapustinsky et al. 'A Fast Timing Light Pulser for Scintillation Detectors'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 241.2 (1985), pp. 612–613. ISSN: 0168-9002. DOI: 10.1016/0168-9002(85)90622-9.
- [136] C. Hu. *Modern Semiconductor Devices for Integrated Circuits*. Upper Saddle River, N.J.: Prentice Hall, 2010. 351 pp. ISBN: 978-0-13-608527-0.
- [137] Analog Devices. *Precision 2.5 V, 5.0 V, and 10.0 V Voltage References*. URL: https://www.analog.com/media/en/technical-documentation/data-sheets/REF01_02_03.pdf (visited on 22/06/2021).

- [138] Free Software Foundation / GNU. *Documentation of AVR-GCC*. URL: <https://gcc.gnu.org/wiki/avr-gcc> (visited on 29/11/2021).
- [139] Microchip. *MCP4922 Datasheet*. Datasheet. URL: <https://ww1.microchip.com/downloads/en/DeviceDoc/22250A.pdf> (visited on 30/11/2021).
- [140] NXP. *BFT93, PNP 5GHz Wideband Transistor - Datasheet*. Datasheet. URL: https://www.nxp.com/docs/en/data-sheet/BFT93_CNV.pdf (visited on 30/11/2021).
- [141] Analog Devices. *LTspice Simulator*. URL: <https://www.analog.com/en/design-center/design-tools-and-calculators/ltspice-simulator.html> (visited on 23/06/2021).
- [142] J. A. Formaggio and G. P. Zeller. 'From eV to EeV: Neutrino Cross Sections across Energy Scales'. In: *Reviews of Modern Physics* 84.3 (24th Sept. 2012), pp. 1307–1341. ISSN: 0034-6861, 1539-0756. DOI: 10.1103/RevModPhys.84.1307.
- [143] Y. Giomataris, P. Rebougeard, J. Robert and G. Charpak. 'MICROMEGAS: A High-Granularity Position-Sensitive Gaseous Detector for High Particle-Flux Environments'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 376.1 (June 1996), pp. 29–35. ISSN: 01689002. DOI: 10.1016/0168-9002(96)00175-1.
- [144] G. Charpak, J. Derré, Y. Giomataris and P. Rebougeard. 'Micromegas, a Multipurpose Gaseous Detector'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 478.1-2 (Feb. 2002), pp. 26–36. ISSN: 01689002. DOI: 10.1016/S0168-9002(01)01713-2.
- [145] F. Sauli. 'Principles of Operation of Multiwire Proportional and Drift Chambers'. In: (1977). DOI: 10.5170/CERN-1977-009.
- [146] T. Alexopoulos et al. 'A Spark-Resistant Bulk-Micromegas Chamber for High-Rate Applications'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 640.1 (June 2011), pp. 110–118. ISSN: 01689002. DOI: 10.1016/j.nima.2011.03.025.
- [147] T.-H. Lin et al. 'Signal Characteristics of a Resistive-Strip Micromegas Detector with an Integrated Two-Dimensional Readout'. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 767 (Dec. 2014), pp. 281–288. ISSN: 01689002. DOI: 10.1016/j.nima.2014.09.002. arXiv: 1406.6871.
- [148] M. Chefdeville and the RD51 collaboration. 'RD51, a World-Wide Collaboration for the Development of Micro Pattern Gaseous Detectors'. In: *Journal of Physics: Conference Series* 309 (10th Aug. 2011), p. 012017. ISSN: 1742-6596. DOI: 10.1088/1742-6596/309/1/012017.
- [149] M. Raymond et al. 'The APV25 0.25 Mm CMOS Readout Chip for the CMS Tracker'. In: *2000 IEEE Nuclear Science Symposium. Conference Record (Cat. No.00CH37149)*. 2000 IEEE Nuclear Science Symposium. Conference Record. Vol. 2. Lyon, France: IEEE, 2000, pp. 9/113–9/118. ISBN: 978-0-7803-6503-2. DOI: 10.1109/NSSMIC.2000.949881. URL: <http://ieeexplore.ieee.org/document/949881/>.
- [150] M. Byszewski. 'MMDAQ Q&A'. URL: https://indico.cern.ch/event/218341/contributions/1519619/attachments/352429/491014/mmdaq_QA.pdf (visited on 29/06/2021).
- [151] H. -. Arends. 'The Mainz Microtron MAMI'. In: *Proceedings of The IX International Conference on Hypernuclear and Strange Particle Physics*. Ed. by J. Pochodzalla and T. Walcher. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 1–5. ISBN: 978-3-540-76365-9 978-3-540-76367-3. URL: http://link.springer.com/10.1007/978-3-540-76367-3_1.

- [152] J. Bortfeldt. *The Floating Strip Micromegas Detector: Versatile Particle Detectors for High-Rate Applications*. 1st ed. 2015. Springer Theses, Recognizing Outstanding Ph.D. Research. Cham: Springer International Publishing : Imprint: Springer, 2015. 1 p. ISBN: 978-3-319-18893-5. DOI: 10.1007/978-3-319-18893-5.
- [153] J. Beringer et al. ‘Review of Particle Physics’. In: *Physical Review D* 86.1 (20th July 2012), p. 010001. ISSN: 1550-7998, 1550-2368. DOI: 10.1103/PhysRevD.86.010001.
- [154] J. Allison et al. ‘Recent Developments in Geant4’. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 835 (Nov. 2016), pp. 186–225. ISSN: 01689002. DOI: 10.1016/j.nima.2016.06.125.
- [155] Garfield++. URL: <https://garfieldpp.web.cern.ch/garfieldpp/about/> (visited on 03/07/2021).
- [156] C. R. Harris et al. ‘Array Programming with NumPy’. In: *Nature* 585.7825 (Sept. 2020), pp. 357–362. DOI: 10.1038/s41586-020-2649-2.
- [157] S. Behnel et al. ‘Cython: The Best of Both Worlds’. In: *Computing in Science Engineering* 13.2 (Mar.–Apr. 2011), pp. 31–39. ISSN: 1521-9615. DOI: 10.1109/MCSE.2010.118.
- [158] R. Liaw et al. *Tune: A Research Platform for Distributed Model Selection and Training*. 2018. arXiv: 1807.05118.
- [159] T. Moretti. ‘Looking for Axion-Like-Particle Di-Photon Decay in the FASER Experiment Using SiGe BiCMOS Pixel Detectors’. 2021.
- [160] F. Kling and S. Trojanowski. ‘Looking Forward to Test the KOTO Anomaly with FASER’. In: *Physical Review D* 102.1 (31st July 2020), p. 015032. ISSN: 2470-0010, 2470-0029. DOI: 10.1103/PhysRevD.102.015032. arXiv: 2006.10630.
- [161] B. Brickwedde, A. Döder, M. Schott and E. Yildirim. ‘Design, Construction and Performance Tests of a Prototype MicroMegas Chamber with Two Readout Planes in a Common Gas Volume’. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 864 (Aug. 2017), pp. 1–6. ISSN: 01689002. DOI: 10.1016/j.nima.2017.04.010. arXiv: 1610.09539.
- [162] T.-H. Lin, A. Döder, M. Schott and C. Valderanis. *Development and Study of a Micromegas Pad-Detector for High Rate Applications*. 3rd June 2015. arXiv: 1506.01164 [physics]. URL: <http://arxiv.org/abs/1506.01164> (visited on 04/07/2021).
- [163] M. Bianco et al. ‘Construction of Two Large-Size Four-Plane Micromegas Detectors’. In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 814 (Apr. 2016), pp. 117–130. ISSN: 01689002. DOI: 10.1016/j.nima.2016.01.049. arXiv: 1511.03884.
- [164] R. A. Light. ‘Mosquitto: Server and Client Implementation of the MQTT Protocol’. In: *The Journal of Open Source Software* 2.13 (26th May 2017), p. 265. ISSN: 2475-9066. DOI: 10.21105/joss.00265.
- [165] Pallets Projects. *Flask Documentation*. URL: <https://flask.palletsprojects.com/en/2.0.x/> (visited on 19/09/2021).
- [166] Grafana Labs. *Grafana: The Open Observability Platform*. URL: <https://grafana.com/> (visited on 20/09/2021).

Acronyms

X_0		Radiation Length
ADC	. . .	Analogue-to-digital converter
ALP		Axion-Like-Particle
ANN	. . .	Artificial Neural Network
AOD	. . .	Analysis Object Data
API		Application Programming Interface
ASIC	. . .	Application-Specific Integrated Circuit
ATLAS	. .	A Toroidal LHC ApparatuS
BCM	. . .	Beam Conditions Monitor
CEP		Central Exclusive Production
CERN	. . .	European Organization for Nuclear Research
CMS	. . .	Compact Muon Solenoid
CSC		Cathode-strip chambers
DAC	. . .	Digital-to-Analog converter
DAQ		data acquisition
DIS		Deep Inelastic Scattering
EM		electromagnetic
ENC	. . .	Equivalent Noise Charge
EPA		equivalent photon approximation
FCal	. . .	Forward Calorimeter
FCNC	. . .	flavour-changing neutral current
FSM	. . .	Finite State Machine
FSR		Final-State Radiation
HF		High-Frequency
HL-LHC	. .	High-Luminosity LHC
HLT		High level trigger
HV		High Voltage

ID		Inner Detector
IFT		Interface Tracker
IP		Interaction Point
IP1		Interaction Point 1
LED		Light emitting diode
LHC		Large Hadron Collider
LHCb		Large Hadron Collider beauty
LLP		Long-Lived Particles
LUCID2		Luminosity measurement using Cherenkov Integrating Detector 2
MBTS		Minimum Bias Trigger Scintillators
MC		Monte Carlo
MDT		Monitored Drift Tubes
MIP		Minimum Ionizing Particle
MS		Muon Spectrometer
Nadam		Nesterov-accelerated Adaptive Moment Estimation
OP		Operational Amplifier
PCB		Printed Circuit Board
PID		Particle Identification
PMT		Photomultiplier Tube
QA		Quality Assurance
QCD		Quantum Chromodynamics
QED		Quantum Electrodynamics
QFT		Quantum Field Theory
RMS		Root Mean Square
ROC		Receiver operating characteristic
RPC		Resistive Plate Chambers
SCT		Semiconductor Strip Tracker
SF		Scale Factors
SM		Standard Model of Particle Physics
SPS		Super Proton Synchrotron
SRS		Scalable Readout System
TGC		Thin Gap Chambers

TIM		Tracker Interlock Module
TLB		Trigger Logic Board
TMVA		Toolkit for Multivariate Data Analysis
TRT		Transition Radiation Tracker
UDP		User Datagram Protocol
UPC		Ultra-Peripheral Collisions
WLS		Wavelength Shifting

List of Figures

2.1.	Diagram showing the elementary particles within the Standard Model.	5
2.2.	Feynman diagram for the electron-positron annihilation and pair-productin. . .	8
2.3.	Diagrams of the vertices described by the Standard Model of particle physics. .	8
3.1.	Graphical representation of a perceptron.	14
3.2.	Implementation of logic operations using simple perceptrons.	15
3.3.	Example for a deep neural network with multiple hidden layers.	15
3.4.	Visualization of a two-dimensional convolution.	17
3.5.	Visualisation for pooling layers in convolutional neural networks.	18
3.6.	Selected activation functions used in neural networks.	19
3.7.	Schematic drawing of the neural network training process.	22
3.8.	Gradient descent with different learning rates.	23
4.1.	Overview of the CERN accelerator complex.	26
5.1.	Computer generated view of the ATLAS detector.	31
5.2.	Schematic drawing of a sector of the ATLAS detector visualizing the interaction of particles with the different sub-detectors.	33
5.3.	Computer generated drawings of the ATLAS inner detector.	34
5.4.	Cut-away view of the ATLAS the electromagnetic and hadronic calorimeters and a section of the electromagnetic calorimeter.	35
5.5.	Simplified flow diagram of the ATLAS event simulation chain with the algorithms drawn in boxes and the objects created written over the arrows.	40
6.1.	Feynman diagram of the lowest possible order for scattering of photons in QED and the production of axion-like particles.	43
6.2.	Normalized particle level distribution from MC simulation for the di-photon invariant mass $m_{\gamma\gamma}$ and the single photon transverse momentum $p_{T,\gamma}$ without any fiducial selection except for cuts applied on generator level.	44
6.3.	Level-1 trigger efficiency as a function of the sum of the two electron clusters. .	48
6.4.	Transverse momentum of the hard-bremsstrahlung photons and of the electron, track and photon system.	51
6.5.	Photon reconstruction efficiency in dependence of the photons transverse momentum and the photons pseudorapidity.	51
6.6.	Efficiency for the identification of a single photon for the default ATLAS photon PID methods.	52
6.7.	Visualizations of the variables from the calorimeter used for the neural network photon identification.	53
6.8.	Normalized distributions of the photon shower shapes extracted from the simulated signal sample and the background sample from data.	55

6.9. Comparison of the background rejection as a function of the photon PID efficiency for a single photon in the four $ \eta_\gamma $ -regions.	58
6.10. Relative importance ranking on the photon particle identification decision for the neural network for the different models.	58
6.11. Single photon PID efficiency extracted from final-state radiation photons radiated by an electron for the 2015 data taking period.	61
6.12. Single photon PID efficiency extracted from final-state radiation photons radiated by an electron for the 2018 data taking period.	61
6.13. Single photon PID efficiency extracted from hard-bremsstrahlung photons emitted by an electron.	61
6.14. Feynman diagrams of the light-by-light scattering process as well as the background processes.	62
6.15. Number of reconstructed pixel tracks in events with two photons and exactly one track matching one of the photons.	63
6.16. Di-photon acoplanarity distribution for events passing the signal selection except for the acoplanarity cut.	65
6.17. Number of tracks in the muon spectrometer for di-photon events with $m_{\gamma\gamma} > 5$ GeV in empty bunches and for events passing the signal trigger with the additional requirements of $N_{\text{MS}} > 0$ and $N_{\text{track}} = 0$	67
6.18. Distribution of the time information from the cluster associated to the sub-leading photon for events in the control region with $A > 0.01$ and for events triggered on empty bunches.	67
6.19. Measured distributions for various kinematic variables are compared with the prediction from MC simulation.	69
6.20. Purity as a function of di-photon rapidity, di-photon invariant mass, average photon transverse momentum and $ \cos\theta^* $	76
6.21. Schematic showing the basic principle of the Bayesian unfolding procedure. . .	79
6.22. Efficiency and fiducial corrections as applied in the unfolding procedure as a function of di-photon rapidity, di-photon invariant mass, average photon transverse momentum and $ \cos\theta^* $	80
6.23. Detector response matrices as a function of di-photon rapidity, di-photon invariant mass, average photon transverse momentum and $ \cos\theta^* $	80
6.24. Relative non-closure, statistical and combined uncertainty from the MC non-closure test for different numbers of iterations.	81
6.25. Results from the data-driven closure test for the di-photon invariant mass. . .	82
6.26. Comparison between the unfolded results with nominal and varied fiducial corrections for the di-photon invariant mass $m_{\gamma\gamma}$ and the average photon $\langle p_{T,\gamma} \rangle$. . .	83
6.27. Comparison of the differential cross-sections as determined by Bayesian and bin-by-bin unfolding.	84
6.28. Measured differential cross-sections as a function of absolute di-photon rapidity, di-photon invariant mass, average photon transverse momentum and absolute di-photon $ \cos\theta^* $	85
6.29. Di-photon invariant mass $m_{\gamma\gamma}$ distribution on detector level for a selection of the $\gamma\gamma \rightarrow a \rightarrow \gamma\gamma$ samples.	86
6.30. Comparison of the nominal $\gamma\gamma \rightarrow a \rightarrow \gamma\gamma$ sample generated using STARlight 2.0 with the alternative sample.	87
6.31. 95 % confidence level upper limit on the $\gamma\gamma \rightarrow a \rightarrow \gamma\gamma$ cross-section $\sigma_{\gamma\gamma \rightarrow a \rightarrow \gamma\gamma}$ and on the axion to photon coupling.	88

6.32. Comparison of the 95 % confidence level upper limit on the $\gamma\gamma \rightarrow a \rightarrow \gamma\gamma$ cross-section $\sigma_{\gamma\gamma \rightarrow a \rightarrow \gamma\gamma}$ extracted using the Bayes method with the results from the nominal method.	90
6.33. Comparison of the 95 % confidence level upper limit on the $\gamma\gamma \rightarrow a \rightarrow \gamma\gamma$ cross-section $\sigma_{\gamma\gamma \rightarrow a \rightarrow \gamma\gamma}$ extracted using fixed backgrounds with the results from the nominal method.	90
6.34. Exclusion limits at 95 % CL on the ALP to photon coupling $1/\Lambda_a$ versus the ALP mass m_a obtained by different experiments.	91
6.35. Event display for a light-by-light scattering candidate event showing a rendering of the ATLAS detector. The two photons are clearly visible with very low activity in the inner detector.	92
7.1. Schematic drawing showing the location of the ATLAS IP and the location of the FASER detector.	98
7.2. Feynman diagrams showing the production of an ALP by the Primakoff process, the charged B^+ meson decay and the K^+ meson decay.	99
7.3. Reach for FASER in the ALP parameter space depending on the minimal separation δ of the two photons distinguished as two photons in the preshower.	100
8.1. Computer rendering of the FASER detector showing all sub-detectors with the front of the detector on the left. For better visibility the cables are hidden in the render.	102
8.2. Rendering of the veto station assembly from different views without any of the support structure.	103
8.3. Rendering of the timing station assembly from different views without any of the support structure.	104
8.4. Rendering of the preshower station assembly from different views without any of the support structure.	105
8.5. Technical drawing of a module for the LHCb electromagnetic calorimeter.	106
8.6. Photography of a bare SCT module as well as an assembled FASER tracker plane.	108
9.1. Schematic drawing of the bare SCT module.	110
9.2. Setup at CERN used for the quality assurance of the individual SCT modules.	112
9.3. Screenshot of the program used for measuring the current in dependence of the bias voltage in the module QA.	113
9.4. Results from the noise threshold scan for a single strip fitted with an error function as defined in Equation 9.1 (left) and the distribution of the noise for a full module (right).	114
9.5. The measured signal versus three injected charges with a linear fit for a single strip and the distribution of the resulting gains for a full module.	114
9.6. Various results from the quality assurance used to select the modules for building the FASER tracker stations.	116
9.7. Setup used for commissioning of the FASER tracker stations as of July 2020.	117
9.8. Temperature and analogue rail current for Plane 10 in the powered on and configured state as measured during plane commissioning and station commissioning.	118
9.9. Measured average module gain and noise for all modules used in the tracker.	120
9.10. Hit occupancy for a single chip of a module before optimizing the TrimDAC and after applying optimum TrimDAC settings for said module.	120

9.11. Measured response for different injected charges used to determine the response curve for a single channel of the first module of the first station with the fitted response curve.	120
10.1. Setup by LHCb to test the response of a calorimeter module and the scan results for one side of a module.	122
10.2. Schematic drawing showing the main components and the working principle of a photomultiplier.	123
10.3. Schematic for the resistive voltage divider as implemented for the calorimeter PMTs.	124
10.4. Transmission curves of a neutral density filter used in the calorimeter in dependence of the light's wavelength in two different positions on the filter.	125
10.5. Computer rendering of the readout assembly for the calorimeter.	125
11.1. Schematic for a LED driver circuit used to generate fast light pulses.	129
11.2. Electronic symbol and semiconductor layer stack up for npn-type transistors and pnp-type transistors.	130
11.3. Electrical connection of a npn-type and a pnp-type transistor forming a thyristor.	131
11.4. Schematic for the first prototype of a LED driver with adjustable bias voltage and the fabricated board.	132
11.5. Picture of the setup for testing the LED driver circuit and the measured PMT response.	132
11.6. Picture of the calorimeter calibration board with the fibre connectors removed and no enclosure.	132
11.7. Schematic block diagram of the various blocks in the calibration board.	133
11.8. Pulse amplitude and pulse area measured with a PMT with different capacitor and inductance values.	135
11.9. Photography of the light-tight box used for various measurements and the measured waveforms for various LEDs.	136
11.10. Section view of the fibre connection at the LED side for the calorimeter fibre bundle.	136
11.11. Picture of a MPO connector and a schematic section view of the opened MPO holder used for coupling it to the LED on the calibration board.	137
11.12. Measurement of the linearity of the pulser bias voltage.	139
11.13. Measured normalized signal amplitude and integrated charge in dependence of the pulser bias voltage	140
11.14. Measured normalized signal amplitude and integrated charge at different pulse frequency settings.	141
11.15. Measured waveforms and integrated charges for each fibre in the same PMT one at a time.	141
11.16. Measurement of the long-term stability of the calibration board with the pulser at a constant setting.	142
11.17. Rendering and photograph of the detector setup used for the combined calorimeter and tracker test beam.	143
11.18. Ratio between the integrated charge for calibration pulses over the integrated charge for beam events at different PMT bias voltages of the upper centre module and the integrated charge for calibration pulses relative to the integrated charge of the first run in the other five modules.	145

11.19. Normalized integral in dependence of the PMT bias voltage for the upper centre calorimeter module for LED pulses and beam events.	145
12.1. Rendering of the preshower prototype used for the test beam.	148
12.2. Schematic drawing of a two-dimensional micromegas detector with a charged particle passing the detector.	149
12.3. Photography of a micromegas detector seen from the front.	151
12.4. Photographs showing the test beam setup at the MAMI accelerator.	153
12.5. Raw data for the X-layer and Y-layer from the middle station detector for a single event.	154
12.6. Exemplary results from cluster finder demonstrating the separation of overlapping clusters.	155
12.7. Schematic drawing showing a misaligned middle detector and the definition of the residual δ between the expected and observed hit position.	156
12.8. Residuals for the X-layers and Y-layers of the preshower prototype detector fitted with the sum of two Gaussians.	157
12.9. The beam widening as determined by the standard deviation of the hit position normalized to the third run taken during the test beam for all three detectors for data and simulated samples.	158
12.10. Geometry of the test beam setup used in the Geant4 simulation. The Mylar foils are drawn thicker than they are in the simulation to make them more visible.	159
12.11. Comparison of the number of clusters per layer and event found in data and simulation.	162
12.12. Comparison of the number of strips per cluster for data and simulation.	162
12.13. Comparison of the total charge per cluster for data and simulation.	163
12.14. Background rejection over the signal efficiency for two different event kinematics.	167
12.15. Signal efficiency for di-photon events evaluated in dependence of the two photons energy.	168
12.16. Background rejection in dependence of the single photon energy (left) and signal efficiency in dependence of di-photon distance (right) for both working points.	168
12.17. Relative importance of the detector layer inputs and the energy for the event classification.	169
12.18. Sensitive region in the ALP mass and coupling plane for the preshower prototype.	171
A.1. Normalised photon shower shapes photon candidates radiated by an electron due to final-state radiation in $\gamma\gamma \rightarrow e^+e^-$ events with no photon PID imposed. Data (black points) is compared with simulation (blue) and the two are in good agreement.	179
A.2. Normalised photon shower shapes from hard-bremsstrahlung photon candidates emitted by an electron in $\gamma\gamma \rightarrow e^+e^-$ events with no photon PID imposed. Data (black points) is compared with simulation (blue) and the two are in good agreement.	180
A.3. Comparison of the alternative $\gamma\gamma \rightarrow \gamma\gamma$ signal sample used for the closure tests in the unfolding to the nominal $\gamma\gamma \rightarrow \gamma\gamma$ signal sample. Shown are the di-photon transverse momentum $p_T^{\gamma\gamma}$ (left) and the di-photon invariant mass $m_{\gamma\gamma}$ (right). Both are well in agreement as can be clearly seen by the ratio plot.	180
A.4. Relative non-closure, statistical and combined uncertainty from the MC non-closure test for the di-photon rapidity, average photon p_T and di-photon $ \cos\theta^* $ for different numbers of iterations.	181

A.5.	Relative non-closure for the data-driven closure test using the alternative signal sample instead of the nominal signal sample for the di-photon rapidity, average photon p_T and di-photon $ \cos\theta^* $ for different numbers of iterations.	182
A.6.	Measured normalised differential cross sections as a function from absolute di-photon rapidity (upper left), di-photon invariant mass (upper right), average photon transverse momentum (lower left) and absolute di-photon $ \cos\theta^* $ (lower right). The measured values for the cross-section are shown as points with the error bars showing the statistical uncertainty. The grey band is denoting the combination of systematic and statistical uncertainties. The measurement is compared with the prediction from SuperChic 3.0 (red line).	183
B.1.	Measured signal amplitude and integrated charge in dependence of the pulser bias voltage.	185
B.2.	Measured signal amplitude and integrated charge in dependence of the pulse frequency.	185
B.3.	Diagram of the lab control architecture.	186
B.4.	Screenshot of the web interface for the lab control system.	187
B.5.	Screenshot of a Grafana dashboard for monitoring the micromegas x-ray setup.	189

List of Tables

6.1. Cut values for the cut-based photon PID at a single photon identification efficiency of 95 %. The background rejection values are determined by evaluating the testing dataset.	55
6.2. Neural network architecture for the photon identification. The dropout rate of 0.1 results in 10 % of the connections being randomly dropped during training.	56
6.3. Background rejections for the different selection methods and datasets at a single photon signal efficiency of 95 %.	57
6.4. Results for the data-driven e^+e^- background study. The extrapolated expected e^+e^- background contribution is shown in the last row.	64
6.5. Contributions on the uncertainty of the overall detector correction factor C	74
6.6. Overview of the binnings used in the unfolding procedure.	76
12.1. Neural network architecture for the discrimination between single and di-photon events. The dropout rate of 0.1 results in 10 % of the connections being randomly dropped during training. As the parameters depend on the layer's type, empty cells are on purpose.	166

