The light at the end of the tunnel gets weaker:

Observation and measurement of photon-induced W^+W^- production at the ATLAS Experiment

Savannah E. Clawson

A thesis submitted to
The University of Manchester
for the degree of Doctor of Philosophy
in the Faculty of Science and Engineering

School of Natural Sciences
Department of Physics and Astronomy

2022

Contents

Word count: 56,160

List of Figures

List of Tables

Abstract

This thesis presents the observation of photon-induced W^+W^- production using proton-proton collision data at a centre-of-mass energy of 13 TeV and with an integrated luminosity of 139 fb^{-1}, recorded by the ATLAS detector. Particular emphasis is given on data-driven corrections to pileup distributions in simulation and calculation of the signal efficiency to pass the exclusivity selection, which requires that no tracks be reconstructed within a ± 1 mm region of the signal vertex. Studies into track multiplicity distributions due to multi-parton interactions are also presented, detailing corrections to leading-track transverse momentum. The fiducial $\gamma\gamma \to WW$ cross-section is measured to be

$$\sigma^{\text{fid.}}_{\text{meas.}} = 3.13^{+0.32}_{-0.31}\ (\text{stat.})^{+0.29}_{-0.26}\ (\text{syst.})\ \text{fb}.$$

when considering the $e^{\pm}\nu_e\mu^{\mp}\nu_\mu$ decay channel and the background hypothesis is rejected with a significance of 8.4 standard deviations. The $\gamma\gamma \to WW$ process is then measured differentially in various kinematic distributions and an unfolding procedure is presented to remove detector effects from the data.

Finally, prospects for future $\gamma\gamma \to WW$ analyses at the High-Luminosity LHC are presented. The analyses are split into different categories, depending on whether selections are made on forward protons measured with the ATLAS Forward Proton detector. It is concluded that reconstructing tracks below nominal kinematic thresholds is an important area of research and development for such analyses to be fruitful in the future of LHC data-taking. A novel analysis method of matching reconstructed forward protons to vertex timing information is presented and shown to possibly improve signal to background ratios relative to an analysis in which no requirement on forward proton tags is made. Detector performance work on the inter-plane alignment of the ATLAS Forward Proton detector and forward proton reconstruction is also presented.

Declaration

No portion of the work referred to in this thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

Savannah E. Clawson

Copyright

Acknowledgments

They say it takes a village to raise a child and it has certainly taken an international collaboration and more to finish this PhD. There are so many people that I would like to thank for seeing me through to this point and I apologise to anyone that I have not mentioned by name.

First of all, to my family. My mum for being my biggest cheerleader in all I do, despite not understanding any of it. My dad for first getting me into science as a kid by sitting me down to watch documentaries with him in the evenings — I know that you will still tell me you know more about particle physics than I do, even when I have a PhD on the subject. My brother whose birthday shall now forever be the day that I passed my viva — I like to think that I inspired you to follow at least partway in my footsteps. To Adam senior (a.k.a. Boxy), thanks for always treating me like one of your own. To my two dogs, Mylo and Hazel, who both sadly passed away whilst I was writing this thesis. You provided me with over a decade of endless entertainment and comfort and it brings me some peace that the concept of particle physics was never within your grasp.

To my original Manchester gang. Alice and George for all the fun times and all of the support during the not-so-fun times. Chris for all of the bouldering trips. Annie for always being available at the end of the phone when I was having a work- or life-induced crisis. To you, I owe most of my degree and PhD.

To the friends that I have made here in the Manchester HEP group, too many to name, without you I truly don't think I would have made it this far. From a different themed drinking occasion every night during my first year, Niels Bohr karaoke, lunchtime and coffee break chats. I struggle to see how I will ever find another department where I feel so much at home. A special shout-out goes to Patrick who didn't particularly help in the writing of this thesis but provided many much-needed distractions from it.

To the friends made at CERN, again too many to name, thanks for all the cycling, skiing, climbing, hiking and lake-swimming trips. I will never forget my time spent in Geneva. Particular thanks go to Mat, Aaron and Melissa, who welcomed me so graciously into their LTA lives and continue to provide me with infinite wisdom.

To the colleagues within ATLAS and beyond. The members of the $\gamma\gamma \to WW$ team, past and present, who have been with me since the start of my PhD. The members of the forward physics community, particularly Rafał Staszewski for being the most patient and technical supervisor, as well as Lydia Beresford and Jesse Liu for trusting me with part of their analysis. I'm further grateful to Lydia for trusting me enough to offer me a job and Jesse for always telling me my plots were beautiful, even when they definitely weren't. A huge thank you to Ste who has provided me with my own personal statistics lessons throughout my PhD and dealt with many of my breakdowns during the final few weeks of writing this thesis. I hope to one day be an honourable guest on #StatsWithSte.

To my supervisors. Brian for originally making me realise physics was cool back when I was a teenager — I never imagined that I would be working with you one day. Andy for providing me with endless support and encouragement (and many beers) over the last four years. Any success that I have had in my PhD, I owe to Andy. For this and so much more, I will be forever grateful.

To Adam junior for supporting me through so much of my PhD, after we accidentally moved in together during the pandemic. You have always been there for me, even when I did not deserve it.

Finally to myself, for finishing this thing.

Chapter 1

Introduction

The earliest documented studies on the nature of light can be attributed to the Greek and Hellenistic traditions, due to a desire to understand vision. These initial theories were somewhat distant from the modern field of optics, with Plato advocating that light consisted of rays emitted by the eyes, believing that vision was initiated by the eyes reaching out to "feel" objects at a distance. It wasn't for over a millennium that thought turned more to light as an independent entity, where arguments on its nature usually fell into two categories: those that considered light to be a wave, and those that considered it to be a particle.

In the mid 1800s, building on the work of Faraday and Ampère, James Clerk Maxwell demonstrated the unification of electricity and magnetism into one fundamental force, electromagnetism. Importantly, Maxwell showed that the laws of electromagnetism obey wave equations, thus firmly engraving the wave picture of light. However, in 1900, Max Planck introduced the idea of quantised energy in attempts to solve the ultraviolet catastrophe. Planck's self-described "act of desperation" was a huge deviation from the previous classical descriptions of energy. In his 1905 paper on the photoelectric effect [1], Albert Einstein concluded that light must come in discrete bundles, or quanta, giving new meaning to Planck's hypothesis and leading to the modern concept of the photon.

The true nature of light remained a contentious issue at the beginning of the 20th century. On one hand were the interference and diffraction experiments that required a wave nature of light for their explanation and, on the other, was the photoelectric effect that could only be understood by invoking a particle-like picture. A formal theory that would simultaneously explain these phenomena would have to wait almost a quarter of a century, until the birth of quantum mechanics in the summer of 1925.

This was a revolutionary period in the history of science and had dramatic implications for the understanding of light and of nature as a whole. In a seminal paper published in 1927 [2], Dirac synthesised the wave and particle natures of light in a single theory and lay the groundwork for the theory of quantum electrodynamics, which was independently developed in the 1940s by Richard Feynman, Julian Schwinger and Tomonaga Shin'ichirō. Forbidden in classical electromagnetism, an early prediction of QED was that two photons should be able to indirectly scatter from each other by exchanging heavy particles inside a virtual loop. The first direct observation of light-by-light scattering was announced by the ATLAS experiment in 2019 [3].

We have come a long way from the earliest studies on the nature of light, from attempts to understand vision, to the description of light as a particle, then as a wave, and finally as a wave-particle duality. A year before his death, Einstein said *"All the fifty years of conscious brooding have brought me no closer to the answer to the question: What are light quanta? Of course today every rascal thinks he knows the answer, but he is deluding himself."*. This thesis, however deluded, aims to add a small drop to the pool of knowledge of photon interactions, using the Large Hadron Collider (LHC) as a TeV-scale photon-photon collider.

The main work in this thesis documents the observation and measurement of photon-induced W^+W^- production with the ATLAS experiment at the LHC. Photon collisions occur in peripheral proton-proton (*pp*) collisions at the LHC, where the highly boosted electromagnetic fields of each proton can be treated as a source of coherent, quasi-real photons. Using the LHC as a photon-photon collider in *pp* runs provides a unique laboratory probe of quantum electrodynamics and electroweak theory, complementing lower-energy probes in heavy-ion runs and high-intensity laser beam experiments. Despite their small cross-sections, photon-fusion processes are typically very experimentally clean, allowing us to perform high precision measurements and searches for new physics. A hallmark prediction of photon-induced processes is intact protons scattered at small angles, which can be detected by dedicated spectrometers such as the ATLAS Forward Proton (AFP) detector.

Chapter 2 briefly summarises the theoretical grounding needed to understand the work presented in this thesis, with a review of the Standard Model, the Equivalent Photon Approximation, and physics at hadron colliders. An overview of state-of-the-art Monte Carlo event generators is given, with additional information on generators specific to photon-fusion processes.

Chapter 3 contains a description of the LHC accelerator complex before describing the ATLAS detector and its subdetector systems. Then the process of object reconstruction within ATLAS is discussed, focusing on the objects relevant to this thesis, namely tracks and leptons.

The ATLAS Forward Proton detector is introduced in Chapter 4. The author has been heavily involved with detector performance work within the AFP group, developing the inter-plane alignment algorithm of the silicon tracking system and calculating the alignment across Run 2 data. The author also developed the working points of proton reconstruction quality in AFP and calculated the efficiencies of proton reconstruction using a per-plane tag-and-probe method.

Chapter 5 details the first observation of the $\gamma\gamma \to WW$ process [4], performed in the fully-leptonic decay channel. Observing this flagship process is important as it only occurs via triple and quartic gauge-boson couplings at leading order, making it an ideal probe for anomalous couplings. This is a rare process with a huge background from quark-antiquark annihilation. The background cross-section is several orders of magnitude larger than the signal and would be impossible to distinguish if it were not for the presence of additional charged particles originating from multi-parton interactions. The author developed a data-driven method to calculate the efficiency of the signal and backgrounds to pass the most important analysis requirement, which aims to remove backgrounds with additional reconstructed tracks, other than those from the W-boson decays. The employed approach is very sensitive to the presence of nearby interactions in the same or neighbouring bunch crossings, referred to as pileup. The author developed many other data-driven corrections to improve pileup and charged-particle multiplicity modelling, including reshaping the beamspot in MC to better match data and studies of charged particle kinematic spectra. The author also calculated the acceptance and correction factors necessary for the cross-section measurement. The author's work was vital in reducing the leading systematic uncertainties in this analysis, enabling the observation of this process.

A differential measurement of the $\gamma\gamma \to WW$ process is being carried out by the author within ATLAS and is the subject of Chapter 6. Fiducial differential cross-sections are unfolded to remove detector effects, with the ultimate aim of using these unfolded cross-sections to constrain new physics in an effective field theory framework. Preliminary studies are presented and the unfolding method is optimised for several kinematic observables.

Chapter 7 details studies into prospects of future $\gamma\gamma \to WW$ measurements at the High-Luminosity LHC (HL-LHC) [5], which the author initiated and led. Rare photon-fusion processes could benefit substantially from the increased dataset size at the HL-LHC. However, the number of interactions per bunch crossing is expected to increase by a factor of 5–6 compared to Run 2, threatening the expected gain in statistical precision due to increased luminosity. Performing these measurements in dense pileup environments like the HL-LHC brings many challenges in background modelling, requiring the development of novel analysis methods. These challenges will only become more relevant in the future as we aim to perform precision measurements and searches for new physics with ever increasing interactions per bunch crossing.

In addition to the work presented in this thesis, the author applied data-driven corrections to pileup in the first cross-section measurement of photon-induced dilepton pairs with an intact proton measured by the AFP detector [6], ultimately reducing the dominant central detector systematic uncertainty by 50 %. This was the first physics result to utilise the AFP detector in high-luminosity runs at the LHC.

Chapter 2

Theory

"You dig deeper and it gets more and more complicated, and you get confused, and it's tricky and it's hard, but... it is beautiful."

– Brian Cox

2.1 The Standard Model

The formalism of this section follows that of Refs. [7–11]. The Standard Model (SM) is formulated as a *quantum field theory* (QFT) which describes the interactions of 12 fermion, 12 antifermion and 5 boson fields. Quanta of these fields are observed as elementary particles. Fermions have spin-1/2 and are further subdivided into quarks and leptons, where quarks carry colour charges, C, but leptons do not. Both quarks and leptons carry weak isospin charge, I_3 and electric charge, Q. Together these give the hypercharge $Y = 2(Q - I_3)$. The list of known fermions is given in Table 2.1, documenting their electric charge, Q, and mass, m. The properties of the anti-fermions can be derived from the same table since anti-particles have the same mass as their particle counterparts but opposite charges. For example, the anti-electron (the positron) has the same properties as the electron apart from it has an electric charge of $+e$ instead of the electron's $-e$[1].

		Generation I	Generation II	Generation III
Quarks	Up-type	$Q = 2/3$ u up $m_u = 2.2^{+0.5}_{-0.4}\,\mathrm{MeV}$	$Q = 2/3$ c charm $m_c = 1.28^{+0.03}_{-0.04}\,\mathrm{GeV}$	$Q = 2/3$ t top $m_t = 173.0 \pm 0.4\,\mathrm{GeV}$
Quarks	Down-type	$Q = -1/3$ d down $m_d = 4.7^{+0.5}_{-0.3}\,\mathrm{MeV}$	$Q = -1/3$ s strange $m_s = 95^{+9}_{-3}\,\mathrm{MeV}$	$Q = -1/3$ b bottom $m_b = 4.18^{+0.04}_{-0.03}\,\mathrm{GeV}$
Leptons	Charged	$Q = -1$ e^- electron $m_e = 0.510999\,\mathrm{MeV}$	$Q = -1$ μ^- muon $m_\mu = 106.658\,\mathrm{MeV}$	$Q = -1$ τ^- tau $m_\tau = 1777.86 \pm 0.12\,\mathrm{MeV}$
Leptons	Neutral	$Q = 0$ ν_e electron neutrino $m_{\nu_e} < 1.1\,\mathrm{eV}$	$Q = 0$ ν_μ muon neutrino $m_{\nu_\mu} < 0.19\,\mathrm{MeV}$	$Q = 0$ ν_τ tau neutrino $m_{\nu_\tau} < 18.2\,\mathrm{MeV}$

Table 2.1: The classification, symbol, name, electric charge, mass and quantum numbers of all known spin-1/2 fermions, with the exception of anti-fermions. All matter is made out of the lightest fermions: protons and neutrons consist, to first order, of up and down quarks. Atoms are made from protons, neutrons and electrons. The masses are free parameters in the Standard Model and are obtained by experiment. Their values are taken from Ref. [12], with the 95 % confidence interval given as the bound on the neutrino mass.

[1] e represents an electron Volt (eV), equivalent to the amount of energy that an electron gains when accelerated in a vacuum by an electric field with a potential difference of 1 V. The charge of all fundamental particles, Q, are given as a number of electron volts.

The interactions of the SM are the result of gauge symmetries described by the

$$\mathrm{SU}(3)_C \otimes \mathrm{SU}(2)_L \otimes \mathrm{U}(1)_Y$$

gauge group, with $\mathrm{SU}(3)_C$ corresponding to the strong force mediated by the massless, spin-1 boson called the gluon[2] and $\mathrm{SU}(2)_L \otimes \mathrm{U}(1)_Y$ corresponds to the electroweak interaction. The $\mathrm{SU}(2)_L \otimes \mathrm{U}(1)_Y$ symmetry is spontaneously broken to give the weak and electromagnetic interactions, mediated by three massive bosons, $W^{\pm}$, Z and the massless photon γ, respectively. An additional complex doublet field is added to induce spontaneous symmetry breaking of the $\mathrm{SU}(2)_L \otimes \mathrm{U}(1)_Y$ symmetry. This field is known as the Brout-Englert-Higgs field, or simply Higgs field, and has a non-zero vacuum expectation value (VEV), resulting in mass terms for the $W^{\pm}$, Z and charged fermions. The excitation of the Higgs field is the Higgs boson, which has been experimentally measured to have a mass of $m_H = 125.25 \pm 0.17\,\mathrm{GeV}$, spin-0 and no electric charge [12]. As alluded to by the subscripts of each gauge group, each symmetry of the SM results in a conserved quantum number according to Noether's theorem [13]. For the $\mathrm{U}(1)_Q$ symmetry, this is the electric charge, Q. For the $\mathrm{SU}(2)_L$ symmetry, this is the weak isospin, I_3, and for the $\mathrm{SU}(3)_C$ symmetry, this is colour, C. The information above is summarised in Table 2.2.

The formulation of the SM builds on developments in classical field dynamics,

Symmetry	Interaction	Mediator	Mediator Mass	Quantum Number
$\mathrm{U}(1)_Q$	Electromagnetic	Photon (γ)	0	Electric charge, Q
$\mathrm{SU}(2)_L$	Weak	Z boson	91.1876(21) GeV	Electric charge, Q Weak Isospin, I_3
		$W^{\pm}$ bosons	80.377(12) GeV	Weak Isospin, I_3
$\mathrm{SU}(3)_C$	Strong	Gluon (g)	0	Colour, C

Table 2.2: The gauge symmetries of the Standard Model and their corresponding mediators (in the mass basis) and conserved quantum numbers. In the gauge basis, the fields of $\mathrm{SU}(2)_L$ are $W_\mu^{1,2,3}$ and the field of $\mathrm{U}(1)_Y$ is B_μ. There are eight gluonic fields, G_μ^{1-8} which carry different colour charges. As we cannot measure colour experimentally, we generally refer to excitations of all of the gluonic fields as the gluon. In the mass basis after spontaneous symmetry breaking described in Section 2.1.4, the photon is a linear combination of the B_μ and W_μ^3 fields and the Z is the orthogonal combination, resulting in the Z boson also conserving electric charge Q. Masses are taken from Ref. [12].

[2] The strong force is actually mediated by eight gluonic fields with different colour charge but, as we cannot measure colour and distinguish these fields, it is common to describe them as resulting in a single mediator.

quantum mechanics, and special relativity. It is constructed to respect Lorentz invariance as well as the various gauge symmetries. The dynamics of the SM can be determined through the *principle of least action*, analogous to classical mechanics. The *action* is defined as the integral of the Lagrangian, L, over time

$$S = \int L\mathrm{d}t = \int \mathcal{L}(\phi, \partial_\mu \phi)\mathrm{d}^4 x, \tag{2.1}$$

where we have defined the Lagrangian density, $\mathcal{L}$, which is more natural to use when discussing a quantum field, ϕ. The benefit of the Lagrangian formalism is that it can be easily generalised to any number of dimensions or any number of degrees of freedom. By enforcing the principle of least action, $\delta S = 0$, we arrive at the *Euler-Lagrange* equation of motion for a field,

$$\frac{\partial \mathcal{L}}{\partial \phi} - \partial_\mu \left(\frac{\partial \mathcal{L}}{\partial(\partial_\mu \phi)} \right) = 0. \tag{2.2}$$

Just as we use the Lagrangian density instead of the Lagrangian, we too wish to use the Hamiltonian density, $\mathcal{H} = \pi\dot{\phi} - \mathcal{L}$, where $\pi = \partial\mathcal{L}/\partial\dot{\phi}$ is the conjugate momentum density. The Hamiltonian can be split as $\mathcal{H} = \mathcal{H}_0 + \mathcal{H}_1$, where $\mathcal{H}_0$ describes the free fields and $\mathcal{H}_1$ their interactions. Provided interactions are sufficiently weak and short lived, they can be approximated as perturbations between the free initial, $|i; t_0 = -\infty\rangle$, and final, $|f; t = +\infty\rangle$, states. The transition amplitude from $t_0 = -\infty$ to $t = +\infty$ is given by

$$S_{fi} = \langle f; t = +\infty | \, \hat{U}(t, t_0) \, |i; t_0 = -\infty\rangle \,, \tag{2.3}$$

where $\hat{U}(t, t_0)$ is the time-evolution operator. The probability of transitioning from initial state $|i\rangle$ to final state $|f\rangle$ is then proportional to $|S_{fi}|^2$. The operation of $\hat{U}(t, t_0)$ on the initial state can be represented by an infinite series with increasing powers of a dimensionless coupling constant, α, mapping all possible routes between $|i\rangle$ and $|f\rangle$. If $\alpha \ll 1$, each successive order of the series expansion has a smaller impact than the preceding one. Therefore, the scattering amplitude can be approximated by truncating the series expansion at a finite order.

Each term in the expansion of S_{fi} can be represented pictorially by Feynman diagrams governed mathematically by Feynman rules, examples of which can be found in Ref. [14]. The matrix element of a scattering process at a fixed order in perturbation theory can then be obtained by summing over all possible Feynman diagrams. However, problems arise when unconstrained momenta appear in loops, leading to divergences in the limit of infinite momenta or infinitesimal distance. These divergences are called *ultraviolet* (UV) divergences and are removed by the processes of regularisation and renormalisation. Regularisation imposes a cut-off scale to prevent integration up to infinite momenta and renormalisation acknowledges that quantities appearing in the Lagrangian do not correspond to physical observables and we can therefore scale these terms and introduce counter terms which cancel out these divergences when computing physical quantities. There is another set of divergences that correspond to the momenta of additional emissions and within loops approaching zero. This set of divergences are called *infrared* (IR) divergences and cancel between the full set of Feynman diagrams. We will come back to IR divergences when discussing simulating physics at hadron colliders in Section 2.5.

Now that we have the relevant tools to determine equations of motion and scattering amplitudes we can discuss the development of the SM, beginning with its precursor *quantum electrodynamics.*

2.1.1 Quantum Electrodynamics

Quantum electrodynamics (QED) is the extension of classical electromagnetism to a QFT and is the precursor to the fully fledged SM as it is today. QED describes the interactions between charged spin-1/2 fermions, mediated by the massless spin-1 photon, γ, and is represented by the Abelian U(1) gauge group. The Lagrangian of the free fermionic fields is the Dirac Lagrangian,

$$\mathcal{L}_{\text{Dirac}} = i\bar{\psi}\gamma^{\mu}\partial_{\mu}\psi - m\bar{\psi}\psi, \tag{2.4}$$

where γ^{μ} are the Dirac matrices and the field ψ represents the bispinor field of the fermion-antifermion pair with mass m. The conjugate of the fermion field is given by $\bar{\psi}$. The U(1) gauge transformation is a local rotation of the field, $\psi(x) \rightarrow \exp(i\alpha(x))\psi(x)$, and results in an additional term appearing in the Lagrangian in Equation (2.4), spoiling the gauge invariance. In order to restore gauge invariance to the Lagrangian, an additional term is required to be added. This is done by introducing a new bosonic spin-1 field, $A_{\mu}(x)$, that transforms as

$$A_{\mu}(x) \rightarrow A_{\mu}(x) + \frac{1}{e}\partial_{\mu}\alpha(x) \tag{2.5}$$

and replacing the partial derivative ∂_μ with the covariant derivative

$$\mathrm{D}_\mu \equiv \partial_\mu - ieA_\mu(x), \tag{2.6}$$

where the additional term represents an interaction between the fermionic field and this new vector[3] field with a coupling strength equal to the electric charge, e. An additional kinetic term for this new field must also be added to our initial Dirac Lagrangian and must also satisfy gauge invariance. This terms equates to $-\frac{1}{4}F_{\mu\nu}F^{\mu\nu}$, where

$$F_{\mu\nu} \equiv \frac{i}{e}[\mathrm{D}_\mu, \mathrm{D}_\nu] = \partial_\mu A_\nu - \partial_\nu A_\mu \tag{2.7}$$

is the electromagnetic field tensor. Note that at no point does a mass term for this new field enter the Lagrangian — doing so would spoil the gauge invariance. It follows that an excitation of this new A_μ field must be massless and this is observed to be the case with the photon. We can now write our full QED Lagrangian as

$$\mathcal{L}_{\mathrm{QED}} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + i\bar{\psi}\gamma^\mu \mathrm{D}_\mu\psi - m\bar{\psi}\psi. \tag{2.8}$$

As a result of renormalisation in QED, the coupling of the electromagnetic field and the electron charge itself are not actually constant and depend on energy scale. This can be understood physically as charge screening caused by the spontaneous production and annihilation of fermion-antifermion pairs. This so-called *fine structure* becomes more apparent at higher energy scales and the coupling must be scaled to account for this. The QED coupling can be written as

$$\alpha_{\mathrm{QED}}(Q^2) = \frac{e^2(Q^2)}{4\pi} = \frac{\alpha_{\mathrm{QED}}(\mu_R^2)}{1 - [\alpha_{\mathrm{QED}}(\mu_R^2)/3\pi]\ln(Q^2/\mu_R^2)}, \tag{2.9}$$

where μ_R^2 is introduced as a *renormalisation scale* [15], to prevent divergence at large values of Q^2. In the case of $Q^2 \to 0$, we obtain $\alpha_{\mathrm{QED}}(Q^2) \approx 1/137$.

2.1.2 Quantum Chromodynamics

The strong interaction is described by the theory of quantum chromodynamics (QCD) which is built as a gauge QFT with SU(3) gauge group which conserves colour charge. We arrive at the QCD Lagrangian through a similar procedure as with QED above, albeit now complicated by the non-Abelian gauge group. A generalised SU(N) group has $N^2 - 1$ degrees of freedom after imposing the

[3]Vector field, in this context, means that the field has spin-1 and transforms as a vector under the Lorentz group.

conditions of unitarity and tracelessness. Therefore, an arbitrary SU(N) transformation can be parametrised by a combination of N^2-1 orthogonal matrices. The SU(3) group is generated by eight 3×3 *Gell-Mann* matrices, λ^a, with generators $t^a=\frac{1}{2}\lambda^a$ ($a=1$–8). The requirement of gauge invariance leads to the addition of one bosonic spin-1 gauge field per generator, G^a_μ, with covariant derivative

$$\mathrm{D}_\mu \equiv \partial_\mu - ig_S G^a_\mu t^a, \tag{2.10}$$

where g_S is the coupling of the strong interaction. The gluonic field strength tensor is given by

$$G^a_{\mu\nu} \equiv \frac{i}{g_S}[\mathrm{D}_\mu, \mathrm{D}_\nu] = \partial_\mu G^a_\nu - \partial_\nu G^a_\mu + g_S f^{abc} G^b_\mu G^c_\nu, \tag{2.11}$$

where f^{abc} are so-called structure constants given by $[t^a, t^b] = if^{abc}t^c$. The non-commutation of the SU(3) group generators means that QCD is a non-Abelian theory which practically means that gluon self-interactions are possible. Such a term was not present in QED as the U(1) generators commute and therefore photon self-couplings are not possible. The QCD Lagrangian is given by

$$\mathcal{L}_{\mathrm{QCD}} = -\frac{1}{4}G^a_{\mu\nu}G^{a,\mu\nu} + \sum_f i\bar{q}_f\gamma^\mu \mathrm{D}_\mu q_f - m_f\bar{q}_f q_f, \tag{2.12}$$

where q_f and $\bar{q_f}$ are the individual quark and anti-quark fields and the summation is performed over all quark flavours, f. Note that Equation (2.12) is the pure QCD Lagrangian which only includes the strong interaction of the quarks and gluons. The weak and the electromagnetic interactions of the quarks must be added separately which we will attempt in Section 2.1.3. Similarly to QED, virtual corrections result in a running coupling constant that varies with scale. However, it takes on a different form from the QED coupling[4] due to the presence of gluon self-interactions, which are absent for the photon. These self interactions lead to a running coupling that becomes stronger at larger distance scales (or equivalently, lower energies). The form of the QCD coupling constant is, at leading order,

$$\alpha_S(Q^2) = \frac{12\pi}{(33-2n_f)\ln\left(Q^2/\Lambda^2_{\mathrm{QCD}}\right)}, \tag{2.13}$$

[4]In fact, it is possible that the QED and QCD couplings could merge at very high momentum transfers, far beyond our current reach. This would represent the unification of all fundamental forces in the Standard Model [16].

where n_f is the number of quark flavours[5] and Λ_{QCD} is the scale below which α_S becomes too large for perturbative calculations to be appropriate in QCD. Below this scale, quarks are bound by the strong force in colour-neutral states. This is the principle of *colour confinement*. As α_S becomes smaller with increasing Q^2, we can treat colour charges as free at high energy scales (large momentum transfer). This is the principle of *asymptotic freedom*. In practice, $\Lambda_{\text{QCD}} \sim 1\,\text{GeV}$. An extensive review of the theoretical and empirical knowledge of the QCD running coupling can be found in Ref. [17].

2.1.3 The weak interaction and the electroweak sector

The weak interaction is modelled by an SU(2) gauge symmetry, whose transformations are generated by $N^2 - 1 = 3$ matrices, t^i (i = 1–3), related to the *Pauli matrices* by $t^i = \frac{1}{2}\sigma^i$ and which obey the commutation relation $[t^i, t^j] = i\varepsilon^{ijk}t^k$, where ε^{ijk} is the rank-3 Levi-Civita antisymmetric tensor. The conserved charge of the weak interaction is the weak isospin, I_3. As we are now accustomed to, we must introduce three new gauge fields, W_μ^i, to preserve gauge invariance. The weak interaction theory is non-Abelian like QCD and therefore the field strength tensor is

$$W_{\mu\nu}^i \equiv \frac{i}{g}[\mathrm{D}_\mu, \mathrm{D}_\nu] = \partial_\mu W_\nu^i - \partial_\nu W_\mu^i + g\varepsilon^{ijk}W_\mu^j W_\nu^k, \tag{2.14}$$

where g is a coupling constant. The weak interaction is known experimentally to violate parity invariance [18], namely symmetry under inversion of all vectors $\mathbf{v} \to -\mathbf{v}$. Therefore it must interact differently with states of different chirality. We now rewrite our fermion spinors as states of definite chirality with left- and right-handed components, $\psi = \psi_\text{L} + \psi_\text{R}$ and $\psi_{\text{L,R}} = \hat{P}_{\text{L,R}}\psi$, where we have introduced the *projection operators*

$$\hat{P}_{\text{L,R}} = \frac{1 \mp \gamma_5}{2}, \quad \gamma_5 = i\gamma^0\gamma^1\gamma^2\gamma^3, \tag{2.15}$$

for states of left- and right-chirality respectively and $\gamma^{0,1,2,3}$ are the usual Dirac matrices. Applying a parity transformation transforms left-chiral states into right-chiral ones and vice versa.

In the case of the weak interaction, the left-chiral fermion fields form weak isospin doublets, corresponding to flavours of fermions in the same generation that differ by one unit of electric charge. This results in doublets of left-chiral charged lep-

[5] It should also be noted that if there were enough flavours ($n_f > 17$), the QCD coupling would behave like the QED coupling and grow with the energy.

tons and their corresponding neutrinos, as well as doublets of left-chiral up-type and down-type quarks from the same generation, written as

$$\psi_L^{\text{lepton}} = \begin{pmatrix} \nu_e, L \\ e_L \end{pmatrix}, \begin{pmatrix} \nu_{\mu,L} \\ \mu_L \end{pmatrix}, \begin{pmatrix} \nu_{\tau,L} \\ \tau_L \end{pmatrix}; \quad \psi_L^{\text{quark}} = \begin{pmatrix} u_L \\ d_L \end{pmatrix}, \begin{pmatrix} c_L \\ s_L \end{pmatrix}, \begin{pmatrix} t_L \\ b_L \end{pmatrix}.$$

The non-interacting right-chiral fermion fields are represented by singlets. These fields are then

$$\psi_R^{\text{lepton}} = \nu_{e,R},\ \nu_{\mu,R},\ \nu_{\tau,R},\ e_R,\ \mu_R,\ \tau_R; \quad \psi_R^{\text{quark}} = u_R,\ d_R,\ c_R,\ s_R,\ t_R,\ b_R.$$

Right-handed neutrinos do not interact with the SM as they do not carry any conserved charge. The third component of the weak isospin is equal to $I_3 = +1/2, -1/2, 0$ for up-type ψ_{L}, down-type ψ_{L} and ψ_{R} respectively. To apply this paragraph to anti-fermions, simply switch left- and right-handed. For example, the weak interaction couples to right-handed anti-fermion doublets but not to left-handed anti-fermion singlets. Since the weak interaction couples only to left-handed fermions, we refer to it as a maximally chiral theory.

Ultimately, we would like to have a unified theory of all fundamental forces just as electricity and magnetism were unified by Maxwell to form the theory of electromagnetism. In the 1960s, Glashow, Salam and Weinberg developed a unified theory of QED and the weak interaction. The unified electroweak interaction is represented by the $\mathrm{SU}(2)_L \otimes \mathrm{U}(1)_Y$ gauge group and is associated to conservation of weak isospin, I_3 and hypercharge, Y. In order to maintain gauge invariance when considering the new unified electroweak gauge group, the covariant derivative must now be defined as

$$\mathrm{D}_\mu = \partial_\mu - ig\sigma_a W_\mu^a - \frac{i}{2} g' Y B_\mu, \tag{2.16}$$

where W_μ^a with $a = 1–3$ are three gauge fields associated with the weak $\mathrm{SU}(2)_L$ symmetry and B_μ is a single field associated with the $\mathrm{U}(1)_Y$ symmetry. g and g' are dimensionless coupling constants. This turns out to describe the interactions observed in nature. As is standard, we add kinetic terms to our Lagrangian for these gauge fields

$$\mathcal{L}_{\text{Gauge}} = -\frac{1}{4} W_{\mu\nu}^i W^{i,\mu\nu} - \frac{1}{4} B_{\mu\nu} B^{\mu\nu}, \tag{2.17}$$

where $W_{\mu\nu}^i$ has the form of Equation (2.14) and $B_{\mu\nu}$ the form of Equation (2.7). As we found with both QED and QCD, introducing mass terms for these new gauge fields would spoil gauge invariance so for now we leave them as massless. Additionally, if we rewrite the Dirac Lagrangian for the fermionic fields (Equation (2.4)) in terms of the left- and right-handed components,

$$\mathcal{L}_{\text{Dirac}} = i\bar{\psi}_{\text{L}}\gamma^{\mu}\partial_{\mu}\psi_{\text{L}} + i\bar{\psi}_{\text{R}}\gamma^{\mu}\partial_{\mu}\psi_{\text{R}} - m(\bar{\psi}_{\text{R}}\psi_{\text{L}} + \bar{\psi}_{\text{L}}\psi_{\text{R}}), \tag{2.18}$$

we see that the term associated with the mass, m, mixes the two components. This means that the Lagrangian is not gauge invariant if the left- and right-handed components transform differently which we have already decided must be the case for the $\text{SU}(2)_L$ group. Therefore, we have discovered that non-zero fermion mass breaks the gauge symmetry of our chiral theory but we also need our theory to be chiral to reproduce the effects that we observe in nature. Additionally, we know from experiment that the bosons of the weak interaction are also massive, causing the short-range nature of the weak force, which also breaks the gauge symmetry. Therefore, we must find a way to break the symmetry of the Standard Model in order to agree with what we see from experiment. The mechanism through which the symmetry is broken, and hence the charged fermions and weak bosons acquire mass, shall be described in the next section.

2.1.4 Spontaneous symmetry breaking

The unfortunate consequence of our formalism so far is that there is no mechanism through which the gauge bosons or fermions can acquire mass without explicitly breaking the gauge symmetries upon which our theory was built. A solution to this problem is to *spontaneously* break the symmetry instead. Spontaneous symmetry breaking (SSB) within a QFT refers to the phenomenon by which the Lagrangian is invariant under the symmetry, but the vacuum of the theory, i.e. the state that minimises the Hamiltonian, is not. As we will show in the next paragraphs, this allows for non-vanishing masses of gauge bosons and fermions in a chiral theory.

SSB is achieved in the SM by introducing a complex scalar field, the Higgs field, in the form of an SU(2) doublet, ϕ. Introducing the Higgs field breaks the gauge symmetry of the SM down to $\text{SU}(3) \otimes \text{U}(1)$. An additional term is introduced to the Lagrangian,

$$\mathcal{L}_{\text{Higgs}} = (D_{\mu}\phi)^{\dagger}(D^{\mu}\phi) - V(\phi), \tag{2.19}$$

with the potential

$$V(\phi) = -\mu^2\phi^\dagger\phi + \frac{\lambda}{4}(\phi^\dagger\phi)^2, \tag{2.20}$$

where μ^2 and λ are positive coefficients, $\mu^2 > 0$ and $\lambda > 0$. This potential has an unstable maximum at $V(\phi) = 0$ and a stable but degenerate set of non-zero minima. By defining $v \equiv 2\mu/\sqrt{\lambda}$, we can then define the vacuum expectation value (VEV) of the field in such a way that the $W^\pm$ and Z bosons acquire mass whereas the photon remains massless, as required. The conventional form of the VEV, as suggested by Weinberg [19], is

$$\langle 0|\,\phi\,|0\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 \\ v \end{pmatrix}, \tag{2.21}$$

where v is a free parameter in the SM, determined experimentally to be $\sim$ 246 GeV [12]. We can then express our field as an expansion about the VEV as

$$\phi = \begin{pmatrix} G^+ \\ \frac{1}{\sqrt{2}}(v + h + G^0) \end{pmatrix}, \tag{2.22}$$

where G^+ and G^0 are massless Goldstone bosons associated with the broken symmetry [20]. By using the unitary gauge, we can remove these Goldstone modes and write Equation (2.22) as

$$\phi = \frac{1}{\sqrt{2}}\begin{pmatrix} 0 \\ v + h \end{pmatrix}, \tag{2.23}$$

where the remaining scalar field, h, is the real physical Higgs field whose excitation is the Higgs boson, and the Goldstone bosons have been absorbed into the longitudinally-polarised degrees of freedom of the massive weak bosons. The discovery of the Higgs boson by ATLAS and CMS in 2012 [21, 22] was the final missing piece of the Standard Model. However there are many reasons to believe that it is an incomplete theory, some of which are presented in Section 2.1.6.

It should be noted that the W^i_μ and B_μ fields presented in Section 2.1.3 do not correspond to the observed mass eigenstates of the electroweak bosons which were presented in Table 2.2. Instead, these are obtained as

$$W^\pm_\mu = \frac{1}{\sqrt{2}}\left(W^1_\mu \mp iW^2_\mu\right) \tag{2.24}$$

and

$$\begin{pmatrix} A_\mu \\ Z_\mu \end{pmatrix} = \begin{pmatrix} \cos\theta_W & \sin\theta_W \\ -\sin\theta_W & \cos\theta_W \end{pmatrix} \begin{pmatrix} B_\mu \\ W^3_\mu \end{pmatrix}, \tag{2.25}$$

where A_μ and Z_μ are the fields of the physical photon and Z boson, respectively. The Weinberg angle is θ_W, with $\cos\theta_W = g/\sqrt{(}g^2 + g'^2)$ and $\sin\theta_W = g'/\sqrt{(}g^2 + g'^2)$. The SM then predicts that the masses of the $W^\pm$ and Z bosons are, at tree level,

$$m_{W^\pm} = \frac{1}{2}vg \quad \text{and} \quad m_Z = \frac{1}{2}v\sqrt{g^2 + g'^2}. \tag{2.26}$$

The relationship between the neutral Z-boson and charged W-boson mass is, to first order,

$$m_Z = \frac{m_W^\pm}{\cos\theta_W}. \tag{2.27}$$

2.1.5 Putting it all together: The Standard Model Lagrangian

The full SM Lagrangian can be divided into distinct sectors,

$$\mathcal{L}_{\text{SM}} = \mathcal{L}_{\text{Gauge}} + \mathcal{L}_{\text{Fermion}} + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{Yukawa}}, \tag{2.28}$$

where each term is documented explicitly with consistent sign convention in Ref. [14]. A pedagogical overview of each term is given here:

- The $\mathcal{L}_{\text{Gauge}}$ term describes all gauge bosons via their field strength tensors which includes the kinetic terms and the interactions between them. Importantly for this thesis, it contains triple and quartic couplings between photons and W-bosons.

- The $\mathcal{L}_{\text{Fermion}}$ term describes the interactions between the gauge bosons (the force carriers) and fermions (the matter particles), and the fermion kinetic terms. This term describes the electroweak and strong force together as all basic interaction vertices share many similarities due to the construction of electroweak theory and QCD as gauge QFTs.

- The $\mathcal{L}_{\text{Higgs}}$ terms describes the interaction between the Higgs field and the gauge bosons, and the Higgs kinetic term. Gluons and photons do not couple to the Higgs field and are therefore massless whereas the massive $W^\pm$ and Z bosons do.

- The $\mathcal{L}_{\text{Yukawa}}$ term describes how fermions couple to the Higgs field and therefore obtain mass. The mass of each fermion is proportional to its coupling with the Higgs field and the VEV: $m = yv/\sqrt{2}$, where y is the Yukawa coupling with the Higgs doublet of the corresponding fermion. The Yukawa couplings are free parameters in the SM that must be obtained experimentally.

It should be stated that the SM is a hugely successful theory that has been rigorously tested and found to have extremely high predictive power. This is demonstrated in Figure 2.1, which shows cross-section measurements made by the ATLAS Collaboration of many physics processes predicted by the SM at different energy scales. On the whole, theory and experiment agree well. Nevertheless, there are some phenomena which cannot be explained within the SM framework and the general consensus is that the SM is a low energy approximation to some higher scale "true" physics model.

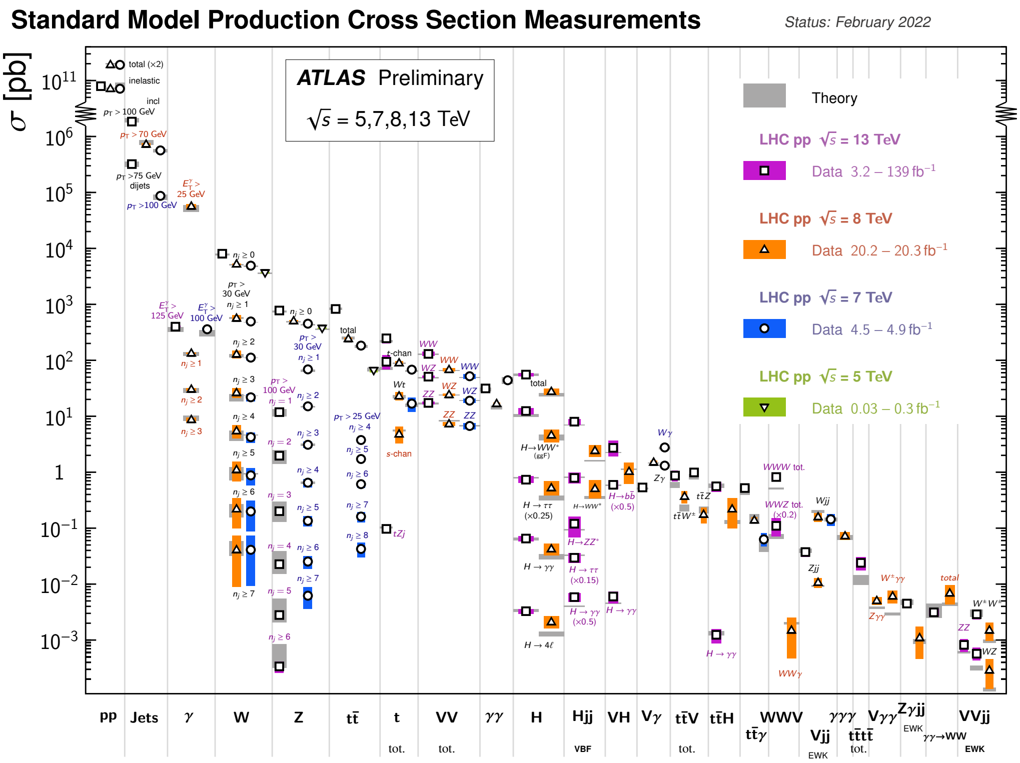

Figure 2.1: Summary of several Standard Model total and fiducial production cross-section measurements made by the ATLAS experiment at centre-of-mass energies, $\sqrt{s}$, of 5, 7, 8 and 13 TeV [23]. The measurements are corrected for branching fractions and compared to the corresponding theoretical expectations. In some cases, the fiducial selection is different between measurements in the same final state for different centre-of-mass energies $\sqrt{s}$, resulting in lower cross-section values at higher $\sqrt{s}$.

2.1.6 Where the Standard Model falls short

The main takeaway of this chapter so far is that the SM is the current best description of all of the fundamental particles in the Universe (that we know of) and the interactions between them. However, we already know that it is far from complete. For one, it makes no attempt to describe one of the four fundamental forces: gravity. Practically this does not deduct from the predictive power of the SM as gravity is very weak relative to the three other forces in the energy ranges that we have explored so far. However, it is ultimately desirable to have a *Grand Unified Theory* which unifies all four fundamental forces. The best description of gravity to date is Einstein's theory of General Relativity which is considered incompatible with the Standard Model. Furthermore, cosmological observations such as galactic rotation curves and measurements of the cosmic microwave background (CMB) [24] also imply that the SM only accounts for $\sim 5\,\%$ of our observable Universe, with the remainder made up of dark matter ($\sim 27\,\%$) and dark energy ($\sim 68\,\%$). It is thought that dark matter could be described by a new fundamental particle which interacts via gravity and could possibly very weakly interact with the SM, although not with photons otherwise it would not be dark! However, there is no candidate particle for dark matter in the SM description other than neutrinos which are not abundant enough to explain the cosmological observations. Additionally, attempts to explain dark energy with the vacuum energy of the SM lead to discrepancies with cosmological measurements of around 120 orders of magnitude. This is referred to as the *cosmological constant problem* [25]. Perhaps the most direct evidence of all that the SM is incomplete is the observation of neutrino oscillation [26, 27] necessitating that neutrinos have mass, albeit this is constrained by measurement to be very small. Neutrinos are strictly massless in the SM description.

There are other experimental inconsistencies with the SM. For example, it is predicted that the Big Bang should have produced an equal amount of matter and antimatter which would have annihilated to leave a Universe full of pure energy. However, we live in a Universe comprised of matter which needs three conditions to explain: baryon number violation, C and CP-symmetry violation, and interactions out of equilibrium, collectively referred to as the *Sakharov conditions* [28]. The SM includes a single source of CP violation in the form of a complex phase appearing in the coupling between quarks and $W^{\pm}$ bosons and has been experimentally verified both indirectly in mixing [29] and directly in decay [30, 31]. However, current measurements suggest that it is insufficient to account for the observed baryon asymmetry of the Universe.

There is another set of arguments as to why the SM is an incomplete theory that fall under the term of *naturalness*. These are not so much based in experimental contradictions to theoretical predictions but rather the feeling that certain details of the SM require or invite a deeper and more satisfactory explanation. This usually materialises in parameters that have been measured to have theoretically unexpected values, which are then said to be *fine-tuned*. An example of fine-tuning in the SM is the *electroweak hierarchy problem* which refers to the experimentally measured value of the Higgs mass and is particularly relevant to the search for new physics at LHC energy scales. If m_H is governed by some new physics at high energies we would expect the Higgs mass to depend on the energy scale of new physics, Λ, according to $m_H^2 \sim \Lambda^2$. Λ is at most the Planck mass ($\sim 10^{19}\,\mathrm{GeV}$) as this is where gravitational effects must be taken into account and the SM needs to be replaced with a description that includes quantum gravity. If new physics only emerges at these scales, then the Higgs mass must be suppressed by a very small number which is argued to be unnatural. It can loosely be argued that the fact that we have not observed any Beyond the Standard Model (BSM) physics at our current reach, $\mathcal{O}(10^3\,\mathrm{GeV})$, already implies a level of fine-tuning for the Higgs mass. It should be noted that arguments surrounding fine-tuning are often highly subjective and often rely on previous experiences of observations resulting from underlying symmetries.

There are several proposed extensions to the SM that seek to solve some of these outstanding issues. This thesis shall not explore any particular BSM theory but shall briefly introduce the idea of parametrising the effects of BSM physics through effective field theories in Section 2.1.7. The main aim of this thesis is to perform precision tests of rare SM processes, where any observed deviation in measured fiducial and differential cross-sections could indicate BSM physics.

2.1.7 A brief introduction to Effective Field Theories

Chapters 6 and 7 discuss current and future differential $\gamma\gamma \to WW$ measurements in the context of differentially measuring tails of distributions for input into Effective Field Theory (EFT) fits. Therefore, a brief introduction to EFTs is given here to provide context for these discussions.

Historically, hints of new physics are sometimes seen indirectly before a direct discovery is possible. In order to indirectly search for new physics in mostly model-independent ways Effective Field Theory frameworks are used. EFTs are applicable in any theory with large scale separation which allows an interaction mediated by a heavy particle to be described as a point-like interaction at low energies. When we operate at energies well below the mediator mass, M, we are blind to physics at distance scales $L \ll M^{-1}$. Therefore, we may observe exchange of a new heavy particle instead as a very small modification to couplings of already known SM vertices. An example of an EFT in this context is Fermi's four-fermion interaction which was postulated to describe β-decay and later superseded by the theory of weak interactions involving the exchange of the massive W-boson.

The main premise of an extension to the SM via EFT is the assumption that any new physics is at some energy scale much higher than currently accessible and higher than the EW scale which is the highest scale in the SM formalism. The SM is then assumed to be a low energy approximation to the true physics. The effective Lagrangian can then be written as

$$\mathcal{L}_{\mathrm{EFT}} = \mathcal{L}_{\mathrm{SM}} + \sum_i \frac{c_i^{(5)}}{\Lambda} \mathcal{O}_i^{d=5} + \sum_i \frac{c_i^{(6)}}{\Lambda^2} \mathcal{O}_i^{d=6} + \sum_i \frac{c_i^{(7)}}{\Lambda^3} \mathcal{O}_i^{d=7} + \sum_i \frac{c_i^{(8)}}{\Lambda^4} \mathcal{O}_i^{d=8} + \dots, \tag{2.29}$$

where $\mathcal{O}_i^d$ represent higher dimension operators with associated dimensionless coupling coefficients c_i. If the operators of Equation (2.29) are built using only SM fields and SSB occurs as in the SM, then the EFT is referred to as SMEFT. The scale of new physics is given by Λ and each higher order term is suppressed by a factor of $1/\Lambda^{d-4}$. The dominant deviations from the SM will therefore be described by the lowest-dimension operators, when $c_i/\Lambda^{d-4} \ll 1$. Odd-dimension terms are generally not considered in the SMEFT as they have been shown to violate lepton (L) and/or baryon (B) number [32–34]. From experimental constraints on neutrino mass, the scale of this new physics can be assumed to be at least $\Lambda \sim 10^{15}$ GeV. It is usually assumed in SMEFT that, in addition to the sector generating the neutrino masses, there exists a new physics sector conserving $B - L$ with a much lower mass scale. In this case, certain dimension-6 operators will have a much larger impact on weak scale observables than dimension-5 ones, and may potentially be observable in current or future experiments. For this reason, many current SMEFT efforts at colliders focus on constraining the coupling coefficients of dimension-6 and (to a lesser extent) dimension-8 operators. There are eight dimension-6 operators in the Warsaw

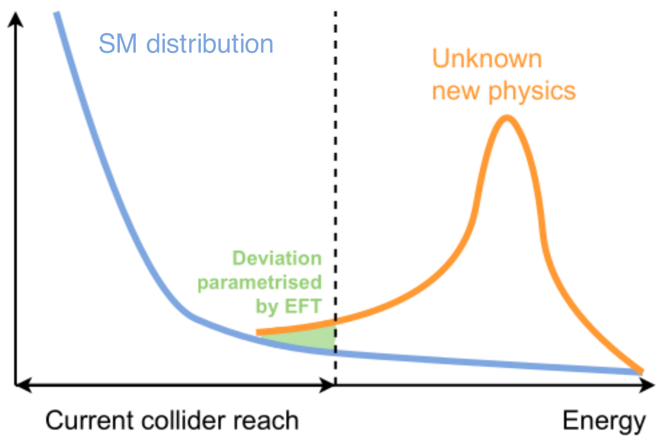

Figure 2.2: A diagram showing possible measurable effects in the tails of current energy reaches, due to new physics present beyond current accessible scales. Differences in measured cross-sections can be explored differentially and parametrised in EFT frameworks.

basis [35] and thirteen dimension-8 operators described by the Eboli model [36] that involve γWW and $\gamma\gamma WW$ couplings.

EFTs are an expression of the bottom-up approach to describe physics as the coupling constants are free parameters that must be constrained by experimental data. The effects of an EFT can be seen as changes in measured cross-sections due to contributions from the pure EFT or due to interference between the EFT and the SM. Both the EFT and the true UV model can deviate from the SM prediction at energies well below the mass of the heavy new particle, M_{UV}. The EFT describes the effect of this new particle well, up to energies near M_{UV}. Only at energies near the new particle production threshold does the EFT deviate from the true UV model. This is shown pictorially in Figure 2.2. By performing high-precision measurements of the tails of distributions predicted to be impacted by EFT operators, EFT effects can be constrained and limits placed on coupling coefficients.

2.2 Particle physics at hadron colliders

The Standard Model can be tested experimentally in high energy particle collisions. This thesis focuses on TeV-scale photon collisions that occur in 13 TeV proton-proton (pp) runs at the Large Hadron Collider (LHC). Therefore, we must also discuss the physics of proton collisions.

Protons are bound states of three *valence quarks*, *uud*. The quarks remain bound via the strong interaction, mediated by the exchange of gluons. It is possible for these gluons to produce quark-antiquark pairs ($q\bar{q}$) on short timescales, called *sea quarks*. During a proton-proton collision at LHC energies, the proton is considered resolved and it can be any of the valence quarks, sea quarks, or gluons that are involved in the hard-scatter (HS) process. We collectively refer to the particles within the proton as partons. The probability of resolving a particular parton within the proton to have a momentum fraction x of the total proton momentum, at a given energy scale, is quantified by a *parton distribution function* (PDF).

The probability of any given process occurring in a collision is quantified by a *cross-section*. The cross-section of a process occurring in a pp collision, $pp \to X$, is dominated by $2 \to N$ parton scattering processes and can be calculated perturbatively by summing over all incoming partons and integrating over allowed momenta, making use of the factorisation theorem [37]. The total cross-section can be written as

$$\sigma_{pp\to X} = \int \mathrm{d}x_i \mathrm{d}x_j \sum_{i,j} f_i(x_i, \mu_F^2) f_j(x_j, \mu_F^2) \sigma_{ij\to X}, \tag{2.30}$$

where the $f(x, \mu_F^2)$ are the relevant PDFs for the initial partons i and j with *factorisation scale*, μ_F. The PDFs are not physical by themselves and it is only the total hadronic cross-section, $\sigma_{pp\to X}$ that is directly measurable. This allows us to move infrared singularities from the partonic cross-section, $\sigma_{ij\to X}$, caused by soft and collinear emissions, into the PDFs if we so wish, allowing exact calculation of the partonic cross-section as a scattering amplitude. The partonic cross-section is calculated from the scattering matrix element (ME) amplitude by summing over all allowed Feynman diagrams for the process. This calculation can now be done perturbatively and can be performed to a varying degree of accuracy, or order, in the relevant coupling constant α. The calculation with the lowest possible order of α allowed by the Feynman rules is referred to as leading order (LO), the next is next-to-leading order (NLO), and the n^{th} order is N^{n}LO. The order of a process is increased by adding additional vertices in the form of initial/final state radiation (legs) and virtual corrections (loops). Renormalisation schemes are applied to remove UV divergences, which is essentially saying that we make no assumption that our theories hold up to infinite energy (which we have already decided is the case for the SM). It should be noted that partonic cross-sections by themselves are unphysical in QCD; they must be paired with PDFs in order to predict physical observables like hadronic cross-sections. We can then think of μ_F as the scale separating perturbative from non-perturbative physics i.e. emissions with momentum below μ_F should then be absorbed into the PDF rather than form part of the partonic cross-section. The choice of μ_F is in principle arbitrary but it is typically taken to be $\mu_F \sim Q$, where Q is the momentum transfer of the hard-scatter. As PDFs include non-perturbative physics, they cannot be calculated from first principles but their evolution with Q^2 is calculable. PDFs are constrained from global fits to data at different energy scales. The main PDF sets considered in this thesis are provided by the CTEQ [38] and NNPDF [39–41] collaborations, accessed through the LHAPDF library [42].

To summarise, in order to calculate the hadronic cross-section to any fixed order in perturbation theory, there are several steps:

1. Calculate the relevant Feynman diagrams for the process of interest up to the desired order to obtain the partonic cross-section. Remove UV divergences via renormalisation by fixing a renormalisation scale.

2. Absorb any soft emissions below a factorisation scale, μ_F, into the PDFs.

3. Integrate the partonic cross-section and PDFs over all possible momentum fractions and sum over all initial partons to obtain the hadronic cross-section.

The hadronic cross-section in Equation (2.30) tells us the overall probability of measuring the process $pp \to X$ but does not tell us anything about the kinematics of X. Therefore, we often wish to measure *differential cross-sections*, $\mathrm{d}\sigma/\mathrm{d}x$ where the process is measured as a function of some distribution, x.

Unfortunately, the total picture of pp collisions is further complicated by the interactions of the spectator partons that did not participate in the hard-scatter process, producing additional activity called the *underlying event*. We have also ignored the fact that at the LHC we collide bunches of protons, which means that we must also account for multiple interactions per bunch crossing, called *pileup*. Both of these effects typically involve interactions with low momentum-transfer and therefore fall under the category of *soft-QCD* which cannot be modelled perturbatively due to large α_S. The underlying event and pileup are discussed in Section 2.5.2.

2.3 The physics of photon fusion

At current experimental precision, QED corrections become important and we must adapt our PDFs to also include photons. In this case, photons can be considered partons inside the proton just like quarks and gluons. In the case of a photon being emitted from a quark line in the proton, we simply adapt the cross-section calculation in Equation (2.30) to use the photon PDF. Such PDFs are provided by several collaborations. Both the NNPDF and CTEQ collaborations parametrise the input photon PDFs and determine parameters along with quark and gluon PDFs from global data. NNPDF start with unbiased starting distributions and constrain the photon PDFs mainly from Drell-Yan LHC data. CTEQ apply a normalisation parameter $p_0(\gamma)$ which is expressed as the momentum fraction of the proton carried by the input photon and constrain the PDF from fits to $ep \rightarrow e\gamma X$ data. The CT14qed [43] PDF set is shown in Figure 2.3 for two different factorisation scales. It can be seen that the valence quarks dominate the PDFs at high values of momentum fraction x whereas gluons dominate at lower x. The photon PDFs are much smaller than those of the quarks and gluons, with the exception of being comparable to the sea quark PDFs at high x. Constraining PDFs

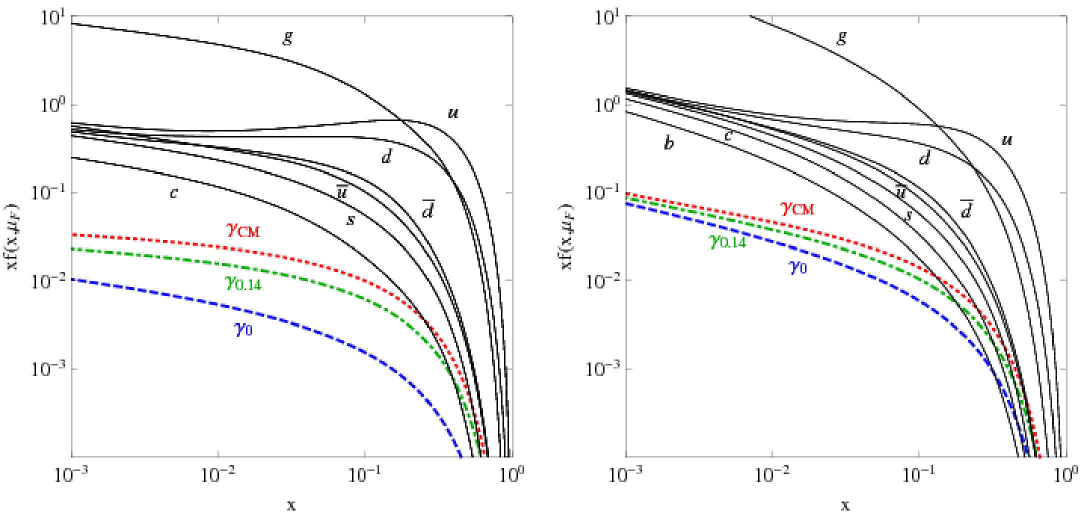

(a) $\mu_F = 3.2\,\text{GeV}$ (b) $\mu_F = 85\,\text{GeV}$

Figure 2.3: The photon PDF of the proton compared to other partons for (a) $\mu_F = 3.2\,\text{GeV}$ and (b) $\mu_F = 85\,\text{GeV}$ [43]. Three representative photon PDFs are plotted: the "Current Mass" photon PDF (γ_{CM}) and the PDFs for photons with inelastic momenta fractions of 0 % (γ_0) and 14 % ($\gamma_{0.14}$), which is an arbitrary normalisation parameter. The effects of the different initial photon PDFs on the quark and gluon PDFs are imperceptible in these plots. The inelastic photon PDF is constrained by deep inelastic scattering data from the ZEUS collaboration.

in this way assumes *incoherent* emission of the photon from individual quark lines within the proton. However, there is another contribution to the photon-induced cross-section, especially at low values of x. This is *coherent* photon emission, where the photon can be treated as emitted from the elastic proton as a whole. This contribution can be treated as an additional PDF, as attempted in Ref. [44]. However, this elastic contribution is usually treated separately, as described in Section 2.3.1.

2.3.1 Equivalent Photon Approximation

In the case of coherent photon emission at $Q^2 \approx 0$, we can treat the photon as being radiated from a point-like proton and the use of photon PDFs are no longer necessary for cross-section calculations. Instead, we can calculate elastic photon-induced cross-sections from a convolution of *equivalent photon* fluxes due to the proton's electromagnetic field. Treating the electromagnetic field of a moving charged particle as a composition of photons was first proposed by Fermi [45] and extended to the relativistic case by Weizsäcker and Williams [46, 47]. A proton at rest has a symmetric radial electric field which can be treated as a complicated superposition of low-energy photons. If the proton is undergoing ultra-relativistic motion, things simplify and we can treat the electric field as a source of coherent high-energy photons, travelling in the same direction as the proton. This is due to the form of the electromagnetic field strength tensor after undergoing a Lorentz boost. If this boost is in the z direction (along the beam axis), then the transverse field (in the x–y plane) is scaled by a factor of γ where $\gamma = (1 - v^2)^{-1/2} \approx 7460$ at the LHC. The field in the direction of motion, E_z, is unaffected by the Lorentz transformation and can therefore be neglected. Ref. [48] shows that this approximation is well justified. This effect is demonstrated pictorially in Figure 2.4.

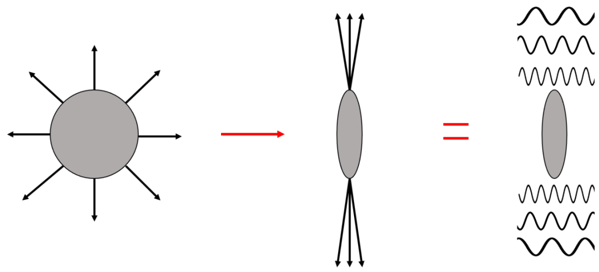

Figure 2.4: The difference in electromagnetic field lines of a charged particle at rest and when travelling at ultra-relativistic speeds. In the case of an ultra-relativistic charge, the electric field can be treated as a source of coherent photons with energy dependent on impact parameter.

The cross-section of an elastic two-photon process in pp collisions can be quantified in the Equivalent Photon Approximation (EPA), where the cross-section of the $\gamma\gamma \to X$ process, $\sigma_{\gamma\gamma\to X}$, is convoluted with the equivalent photon fluxes. The cross-section for a two-photon process [49] can be written as

$$\sigma_{pp\to pp(\gamma\gamma)\to ppX} = \int \mathrm{d}\omega_1 \mathrm{d}\omega_2 \frac{n\,(\omega_1)}{\omega_1}\frac{n(\omega_2)}{\omega_2}\sigma_{\gamma\gamma\to X}(\omega_1, \omega_2), \tag{2.31}$$

where $n(\omega)$ is the number of equivalent photons of energy ω, given by

$$n(\omega) = \frac{2}{\pi} Z^2 \alpha_{\mathrm{EM}} \ln\left(\frac{\gamma}{\omega R}\right) \tag{2.32}$$

in natural units for a relativistic process, where Z is the atomic number of the nucleus (equal to 1 for a proton), $\alpha_{\mathrm{EM}} = 1/137$, R is the charge radius and the photon energies are ω_1 and ω_2. The total cross-section of X is obtained by integrating over all possible photon energies and accounting for the photon flux given in Equation (2.32). It is assumed that the photons have low-virtuality, i.e. $Q^2 \approx 0$ and therefore Equation (2.31) is not applicable for processes in which multiple virtual photons are exchanged. A very thorough discussion of the two-photon production mechanism and EPA can be found in Ref. [50]. Higher momentum-transfer photon emission is considered to be from within a parton inside the proton, causing the proton to break apart after the emission. This is referred to as *dissociative production* and is modelled using photon PDFs, as discussed at the beginning of Section 2.3. Even in elastic scattering, the protons do not always remain intact due to non-zero probability of rescattering after the photon emission. The probability of the proton surviving the interaction without rescattering is discussed in Section 2.3.2.

This approach can be extended to collisions of other charged particles. The factor of Z^2 in Equation (2.32) results in a cross-section enhancement of $\sim 5 \times 10^7$ in lead-lead (PbPb) collisions relative to pp collisions. However, the maximum photon energy is given by $\omega_{\mathrm{max}} \approx \frac{\gamma}{R}$ and therefore higher photon energies can be reached in pp collisions due to higher γ and smaller R. In order to produce two on-shell W-bosons, it is necessary to perform a $\gamma\gamma \to WW$ analysis in pp collisions at the LHC. Lower mass photon-induced analyses can be performed in ultraperipheral heavy ion collisions (UPC), benefiting from the higher photon flux. For example, the observation of light-by-light scattering [3] and measurements of $\gamma\gamma \to \ell\ell$ in PbPb UPC [51–53].

2.3.2 Final states with intact protons

It is possible for protons in LHC collisions to remain intact after an interaction, as long as the exchange does not change any quantum numbers of the proton(s). This is possible in QCD via the emission of a so-called *pomeron*, which can be thought of as a collection of soft gluons forming a colour-singlet in the non-perturbative regime. At higher scales, the pomeron structure becomes measurable and plays an important role in our understanding of QCD and total cross-sections. Processes involving the exchange of a pomeron are referred to as a *diffractive* processes, which make up a large yet theoretically poorly understood part of the total *pp* cross-section at the LHC. Diffractive processes can be further subdivided into elastic, single-diffractive and double-diffractive where both protons remain intact or one or both protons break apart, respectively. It is also possible for a double pomeron exchange (DPE) to produce a central state which is disconnected in colour from the initial protons. If all of the pomerons' mass goes into the production of the central state, referred to as central exclusive production (CEP), the mass of the central products is directly correlated to the energy loss of the initial protons.

The QED equivalent of CEP is a double photon emission, in which it is also possible to have direct and resolved interactions, leading to both elastic and dissociative photon emission, where the proton remains intact or breaks apart. As the photon is fundamental, photon-photon fusion via this mode always produces a central state that is fully kinematically correlated to the initial protons. This allows constraint on any missing mass in the central detector in the case of elastic production with two intact protons in the final state, which could be hugely advantageous for e.g. dark matter searches [54]. Photon-induced processes are not diffractive by definition although they can share many similar characteristics to pomeron-induced processes. One classic hallmark of a CEP or photon-induced process is a region of the detector void of any activity, quantified by a so-called *rapidity gap*. Rapidity gaps occur because any hadronic activity associated to these processes through proton dissociation is expected to be produced in the forward direction. However, in practice rapidity gaps can be obscured by the presence of additional soft interactions between protons in the same or neighbouring bunch crossings, called *pileup*, that are not associated to the interaction vertex of interest but do produce additional particles. The final state particles that define the rapidity gap may also fall outside of the detector acceptance, rendering rapidity gap magnitudes immeasurable. One way to recover inefficiencies from these means is to reconstruct any intact protons in the final state that originate from exclusive processes, which are typically scattered at very small angles (∼few µrad) relative to the beamline. The proton fractional energy loss, ξ, is given by

$$\xi = 1 - E_{\text{proton}}/E_{\text{beam}}, \tag{2.33}$$

where E_{proton} is the energy of the scattered proton and E_{beam} is the nominal beam energy. In the case of a fully reconstructed central system, the proton energy loss can be related to the mass m and rapidity y of the central system by

$$\xi^{\pm} = \left(m/\sqrt{s}\right) e^{\pm y}, \tag{2.34}$$

where $\sqrt{s}$ is the centre-of-mass (CoM) energy and the $\pm$ refers to the proton with positive or negative rapidity.

In elastic photon-induced processes, the photons are typically emitted with very low Q^2 and therefore there is little transverse momentum transfer and the protons are scattered with very small $p_{\text{T}} \approx 0$. Dissociative processes sourcing a higher-virtuality photon and diffractive pomeron-induced processes typically have higher momentum transfers and the protons will be scattered at slightly larger angles.

Proton soft survival

Even in the case of elastic photon-emission, there is also the possibility that the protons will rescatter and produce secondary particles, the probability of which is quantified by proton *soft survival* factors. The term survival historically refers to the survival of the rapidity gap associated with the exchange of a colour-neutral object which any hadronic activity from the proton rescatterings would spoil. A rapidity gap is an angular region of phase-space void of hadronic activity and is a hallmark of *diffractive* processes [55]. The effect of proton rescattering is to lower the observed elastic cross-section when relying on definitions of exclusivity surrounding lack of additional hadronic activity. The simplest way to model proton soft survival is simply to modify the predicted cross-section from EPA (Section 2.3.1) by a multiplicative factor, $S \leqslant 1$. In fact, the majority of Monte Carlo generators do not model soft survival at all and this factor has to be inferred from experiment and applied to predicted EPA cross-sections. This is the method applied in the $\gamma\gamma \to WW$ observation (Chapter 5) to compare the measured cross-sections to theory. Strictly speaking the value of the factor depends on the particular process and the cuts imposed in the experiment. It has been demonstrated that proton soft survival depends on proton p_{T} and therefore impact parameter, i.e. how close the photons are when the photon exchange occurs [56]. As photons in the EPA are emitted with virtuality close to zero and therefore small p_{T}, the impact parameter between the protons is expected to be larger than in a diffractive central exclusive process which have low survival factors [56] and the survival factor in photon-induced processes is expected to be between 80–100 %. However, as the rescatterings are non-perturbative they cannot be calculated analytically and instead phenomenological "eikonal" screening models are used [57, 58].

2.4 $\gamma\gamma \to WW$ production mechanisms

This thesis aims to observe and measure the process $\gamma\gamma \to WW$ at the LHC. The production of two W-bosons through photon-photon fusion offers a unique opportunity to precisely probe the EW sector as the process can only occur via vertices with EW gauge boson couplings at leading order. The $\gamma\gamma \to WW$ interaction can involve triple gauge couplings (TGCs) through a t- or u-channel exchange of a virtual W-boson between the two photons or a quartic gauge coupling (QGC). Independently, each contribution to the cross-section diverges at high energies but the QGC vertex amplitude cancels both the leading quadratic and sub-leading linear divergences of the TGC amplitudes. Hence measuring the $\gamma\gamma \to WW$ process directly tests the non-abelian SU(2)⊗U(1) gauge structure of the Standard Model [59]. Additionally, a (differential) cross-section measurement can set limits on anomalous couplings, parametrised in effective field theories (EFTs) with additional dimension-6 and dimension-8 operators [35, 59, 60]. A brief introduction to EFTs was presented in Section 2.1.7. The t-channel exchange and quartic coupling interaction are shown in Figure 2.5. Not shown in Figure 2.5 are the protons, from which the photons are emitted. When including the initial protons as part of the process of interest, there are two types of production mode we can consider:

Elastic production: At low virtualities ($Q^2 \approx 0$), the photon can be treated as being radiated coherently from the proton as a whole and the proton remains intact after the emission approximately 80 % of the time. As the photon is typically emitted with very small transverse-momentum, p_{T}, the intact protons travel approximately parallel to the beam axis. In these events, there is no additional activity accompanying the hard-scatter, other than the products of the W-boson decays. The remaining 20 % of elastic events undergo additional soft-QCD interactions due to proton rescattering, addi-

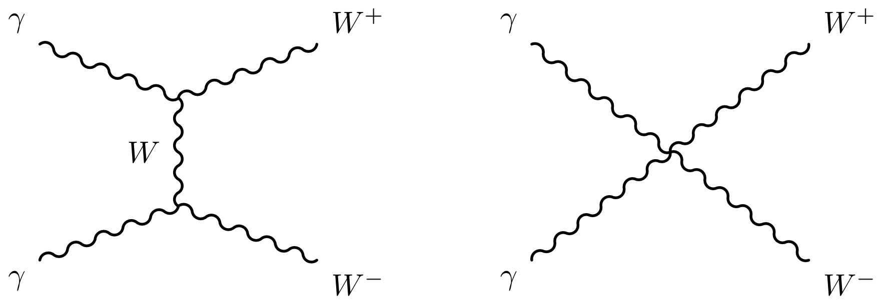

Figure 2.5: Feynman diagrams of photon-induced WW production, $\gamma\gamma \to WW$. The figure on the left shows the t-channel interaction with two triple γWW couplings and the figure on the right shows the quartic $\gamma\gamma WW$ coupling. The two W-bosons must have opposite sign to adhere to the law of charge conservation. The u-channel diagram is not shown here but also contributes to the leading order process.

tional hadronic activity accompanying the hard-scatter. The probability of the proton remaining intact is quantified by soft-survival factors, as discussed in Section 2.3.2.

Inelastic production: At higher virtualities, we can consider the photon as having been radiated from a parton within the proton. In this case, it is the parton that undergoes the hard-scattering, causing the proton itself to break apart, or *dissociate*. This type of inelastic production can be divided into two production modes: *single-dissociative* (SD) and *double-dissociative* (DD) production, where one or both protons break apart, respectively. Inelastic production can also occur at low virtualities if the photon is produced in the decay of the proton from an excited state.

The elastic and dissociative modes are demonstrated in Figure 2.6. The characteristic feature of photon-induced processes is that they are typically produced without any additional hadronic activity. Therefore, so-called *exclusive* $\gamma\gamma \to WW$ production can be selected experimentally by requiring no activity in the vicinity of the signal vertex, other than from the W-decays. This selection is referred to as the *exclusivity requirement* and suppresses *inclusive* backgrounds which mimic the signal, the largest being $qq/qg/gg \to WW$ production which has a cross-section of the order 10^3 times that of the $\gamma\gamma \to WW$ process. Quark and gluon induced processes are usually accompanied by an underlying event and therefore are expected to have additional activity surrounding the production vertex.

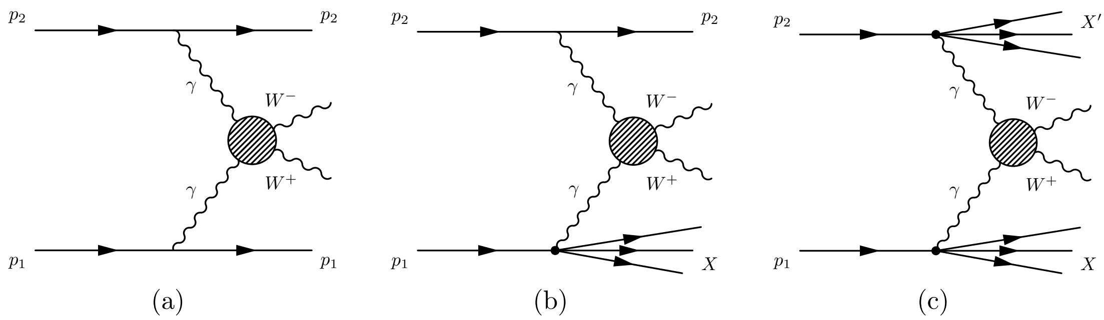

Figure 2.6: Feynman diagrams for the (a) elastic, (b) single- and (c) double-dissociative contributions to the signal. The hatched circle represents all leading-order contributions to the signal production. The W-decays are not shown.

2.5 Simulating physics at hadron colliders

Particle collisions and their measurements are inherently stochastic processes and must be quantified using probabilistic methods. On top of this, a high dimensionality phase-space must be integrated over, necessitating the need for numerical methods. Monte Carlo (MC) generators [61] employ Monte Carlo techniques to simulate collisions at particle detectors. MC generators rely on *factorisation* which is the ability to separate parts of the generation into different stages according to some characteristic energy scale. Starting from the highest scales, i.e. the shortest distances, sub-processes are added which will populate different phases of particle emission and creation at lower scales or larger distances. There are several stages of MC production:

1. **Matrix element generation:** The matrix element of the hard-scatter, considering the production of heavy and/or highly energetic particles, is generated at the desired accuracy (LO, NLO, etc) and from this the hard-scatter cross-section is computed, $\sigma_{ij\to X}$ in Equation (5.15). The complexity of this stage scales hugely with the generation order. The ME is then calculated by summing all possible Feynman diagrams contributing to the hard-scatter process at the desired order. The hard-scatter cross-section is proportional to the square of the ME, $\sigma \propto |\sum_i \mathcal{M}_i^2|$, averaging over all possible spin states of the initial particles.

2. **Parton shower and hadronisation:** Many radiative corrections can be applied after the matrix element generation in the form of a *parton shower* (PS). The parton shower models additional radiation down to a certain cut-off scale. Once this scale is reached, *hadronisation* algorithms group the remaining quarks and gluons to form hadrons (colour-neutral particles).

3. **Decay of unstable hadrons:** The final step of the event generation is performing the decay of any unstable hadrons (and tau leptons). The definition of unstable is somewhat arbitrary and usually analysis-specific but can be considered as particles that decay before reaching the innermost layers of the detector.

4. **Overlaying the effect of multi-parton interactions (MPI):** Hadron colliders are messy environments where not only do we need to worry about the hard-scatter and associated ISR and FSR, but also the impact of additional interactions that originate both from spectator partons in the hard-scatter *pp* interaction and from other *pp* interactions in the same or neighbouring bunch crossings.

5. **Detector response:** The interactions of collision decay products with the detector are also simulated using MC methods. The ATLAS detector is simulated [62] using GEANT4 [63]. The ATLAS detector and object reconstruction shall be discussed in Chapter 3. This stage is not necessary for parton or particle-level calculations.

2.5.1 Parton shower, hadronisation and decay

Showering and hadronisation are non-perturbative processes which must rely on phenomenological models tuned to experimental observations in data. Parton shower MC methods are applied after the ME generation and simulate the effects of both initial and final state radiation. In QED, this represents the emission of photons from the charged quark lines. In QCD, this represents a splitting of quarks and gluons as $q \to qg$ and $g \to gg$, respectively. Parton shower algorithms perform the splitting as a cascade until a cut-off scale is reached. Applying a cut-off scale prevents divergences which are not infrared and collinear safe. Physically, this can be understood as being unable to resolve very soft (low energy) emissions and emissions at very small angles. Once this cut-off scale is reached, colour confinement becomes important and non-perturbative hadronisation takes over. The effects of the parton shower are absorbed into the hard-scatter cross-section and therefore an upper-scale on the parton shower is also set to avoid double counting of emissions already calculated in the matrix element. The emissions can be ordered in transverse momentum or by angle of emission. The factorisation and renormalisation scales are to some degree arbitrary and can be varied to define systematic uncertainties on theoretical predictions from MC.

Hadronisation is usually modelled by the Lund string [64] or cluster model [65–67] and describes the transition from parton-level to particle-level, where the individual partons come together to form colour-singlet hadronic bound states. Hadronisation will produce several unstable hadrons which must be subsequently decayed into lighter particles that are stable on time-scales to be measured by the detector. It is these decay products that are ultimately measured in data and thus hadronic decays is an important part of modelling necessary for comparing phenomenological models to data for tuning purposes.

2.5.2 Multi-parton interactions and pileup

The vast majority of collisions at the LHC are soft, leading to diffractive scattering or multi-particle production with low transverse momenta. Soft-processes can be subdivided into elastic, single-diffractive, double diffractive, and non-diffractive and are all needed to offer an inclusive description of the total pp cross-section at the LHC. These soft processes also need to be simulated but their non-perturbative nature means that we have to resort to phenomenological models with tunable parameters to describe the data. On top of soft-QCD effects already discussed, there are additional effects that are especially relevant in this thesis, namely the *underlying event* (UE) and the effect of multiple interactions per bunch crossing in the ATLAS detector, *pileup* (PU). It should be noted that it is possible to have hard MPI and hard diffraction so both effects are not fully encompassed by the term soft-QCD although as the majority of these events are soft they shall be classed as such here. The probability of more than one hard-scatter occurring in an LHC bunch crossing is extremely small.

The calculation of the hadronic cross-section does not account for the fact that there are other spectator partons within the proton that, although they may not directly contribute to the hard-scatter, are likely to impact the final state that we observe in our detector. At small impact parameters, which are typically necessary for hard processes, there are usually multiple parton interactions (MPI) due to these spectator partons in the proton. Including the effects of MPI is necessary to match distributions observed in data as MPI increases the charged particle multiplicity due to showering and hadronisation from these additional interactions. Modelling the underlying event is not as simple as overlaying minimum bias data (minimum-bias refers to the most inclusive sample of events that can be measured) with the hard-scatter process as a hard-scatter process is more likely to have large MPI activity, resulting in the so-called *pedestal effect* observed in data. In order for an MPI model to accurately describe the shape of the pedestal (e.g. to describe both minimum bias data and underlying event distributions simultaneously) it must include impact parameter dependence. Various MPI models have been developed, each of which needs to specify how often MPIs occur as a function of the relative p_{T} of the protons, which is inversely related to the impact parameter. Colour connections between the MPI and the hard-scatter should also be taken into account, as well as possible *colour reconnection* due to interactions between MPI and partons produced in the hard-scatter itself. The presence of MPI activity can affect hard-scatter object reconstruction and therefore the UE should be simulated on a per-event basis. Mismodelling of charged particle multiplicity in the UE is a known and well-documented problem [68–70].

An important point to make here is that there will not always be an underlying event accompanying the hard-scatter process. Photon-induced processes do not involve colour exchange and also tend to occur at larger impact parameters than QCD hard-scatter processes and therefore it is possible to have an event void of any MPI. The presence (or lack) of MPI from the underlying event can be used to separate photon-induced processes from QCD backgrounds, in which the hard-scatter process itself may be indistinguishable. In practice, this is complicated by other effects such as proton dissociation, proton rescattering (soft survival) and additional activity from nearby pileup interactions. Pileup interactions refer to proton interactions in the same or neighbouring bunch crossing of the hard-scatter process of interest. Pileup interactions are modelled by a combination of diffractive and MPI models. The diffractive contributions are described using *Regge trajectories* [71–73] whereas MPI models are used for the non-diffractive components.

2.5.3 State-of-the-art MC event generators

There are many different MC generators available, with some performing the ME generation, some the parton showering and hadronisation, and some capable of performing both. Different generators may employ different phenomenological models and it is desirable to test multiple generators for important physics processes within an analysis. Matching/merging the order to which the ME is generated to the order used for the parton shower is also incredibly important to avoid double counting emissions and ensure the full phase-space is covered. Matching refers to approaches in which high-order corrections to an inclusive process are integrated with the parton shower. Merging involves introducing a scale where any parton produced above that scale is generated with a corresponding higher-order matrix element and, conversely, any parton produced below is generated by the shower. A multi-purpose generator is one that can perform MC generation for a large range of SM (and possibly BSM) processes. Some of the important multi-purpose generators used in this thesis are discussed in this section, with a focus on generators specific to photon fusion processes discussed in Section 2.5.4.

Herwig

Herwig [74, 75] is a multi-purpose event generator, capable of generating matrix elements, MPI and performing the combination of the ME and PS. It can also be interfaced with other generators. Herwig is capable of generating MEs at NLO accuracy in the QCD coupling and LO in the QED coupling. Angular ordered parton showering is used as default but the Catani-Seymour dipole factorisation [76] can be used as an alternative, where partons are bound into colour connected dipoles which undergo sequential splitting ordered in p_{T}. Matching and merging between the matrix element and parton shower are performed using the MatchBox framework [77]. Simulation of MPI in the underlying event is based on an eikonal model and hadronisation is performed using the cluster model. In this thesis, the elastic component of the $\gamma\gamma \to WW$ and $\gamma\gamma \to ll$ processes are modelled at LO accuracy using Herwig7 (v.7.1.5p3) interfaced with the BudnevQED photon flux [78] through ThePEG software [79]. The photon flux is differential in both x and Q^2. Herwig is also used for parton showering in alternative background samples.

Sherpa

Sherpa [80] is multi-purpose event generator similar to Herwig in that it is capable of performing ME generation as well as showering and hadronisation. The default matrix element generator is AMEGIC++ [81] with Comix [82] preferred for high-multiplicity SM process generation. The matrix elements are generated at NLO accuracy in QCD for up to one additional parton and LO for up to three additional partons. In contrast to Herwig, the default parton showering scheme in Sherpa is the Catani-Seymour dipole scheme. Matching is based on the MC@NLO [83] method which calculates the approximate parton shower contribution that is matched to the ME generation and subtracts this from the exact NLO calculation. Merging is performed with the MEPS@NLO [84] prescription.

Also in contrast to Herwig, Sherpa uses the Lund string model for hadronisation. Virtual corrections to NLO predictions are provided by external libraries like OpenLoops [85]. In this thesis, background diboson processes are simulated using the Sherpa 2.2.2 generator with the NNPDF3.0 NNLO PDF set along with the dedicated set of tuned parton-shower parameters developed by the Sherpa authors.

Powheg-Box

The MC@NLO matching process used in Sherpa can lead to large negative event weights which could in principle lead to unphysical negative cross-section predictions, prompting the development of the Powheg method [86, 87]. Powheg-Box [88] is a general framework for implementing processes at NLO using the Powheg method. The matrix element generated in Powheg-Box can then be interfaced to other generators for showering, production of soft interactions and hadronisation. In this thesis, the inclusive $q\bar{q} \to WW$, WZ and $Z(\to ee/\mu\mu/\tau\tau)$ background processes are modelled at NLO accuracy using the Powheg-Box v2 event generator [89, 90] using the CT10 PDF set [91] and interfaced with Pythia. For the generation of $t\bar{t}$ and Wt processes, the Powheg-Box v2 generator at NLO is used with the NNPDF3.0 NLO PDF set. In alternative background samples, Powheg-Box is interfaced to Herwig7 (v.7.1.6), using the H7UE set of tuned parameters and the MMHT2014LO PDF set [92]. Gluon-induced Higgs production with a subsequent decay into W-bosons is modelled at NNLO accuracy in QCD using the Powheg NNLOPS program [93, 94] with the NNPDF3.0 NNLO PDF set.

Pythia

Pythia is a multi-purpose MC event generator [95, 96] providing hard and soft components to the collisions of hadrons and leptons. It is used in this thesis as a parton shower MC interfaced to other matrix element generators, with the exception of simulating double-dissociative $\gamma\gamma \to ll$ events. Pythia is capable of generating both elastic and inelastic interactions. Parton showering in Pythia is based on the dipole evolution ordered in decreasing transverse momentum [97]. Several options for matching and merging of the parton shower with the matrix element calculation are available in Pythia including the aforementioned MC@NLO and Powheg methods but also CKKW-L [98, 99]. MPI in the underlying event is described by the Sjöstrand-Zijl model [100] and the Lund string model is used for showering and hadronisation.

Pythia is also used to simulate pileup interactions as inelastic (diffractive and non-diffractive) proton-proton collisions using the A3 tune [101], which was developed to reproduce ATLAS minimum bias data at 0.9, 7 and 13 TeV. The pileup simulation is combined with hard-scatter simulation using an overlay method [102].

MadGraph

MadGraph (MG5_aMC@NLO) [103, 104] is a matrix element generator that generates processes at either LO or NLO accuracy with high flexibility in matching and merging possibilities. Matching and merging is performed with the MC@NLO and FxFx [105] methods. Madgraph is one of the few generators capable of simulating elastic, single- and double-dissociative $\gamma\gamma \to WW$ and $\gamma\gamma \to ll$ events. It is then interfaced with Pythia for parton showering and hadronisation using the CT14qed PDF. EPA methods are used in elastic event generation. The in-built photon flux used in MadGraph is integrated over the virtuality Q^2 which results in the initial state photons being generated with zero transverse momentum. This is somewhat mitigated by giving events additional intrinsic p_{T} during the hadronisation step which has been tuned to reproduce the elastic sample from Herwig, which treats the flux deferentially in Q^2. Particle multiplicity distributions predicted by MadGraph are reweighted to match those predicted by the photon-fusion generator LPair.

2.5.4 Photon fusion generators

Traditionally, theoretical predictions for elastic photon-induced production are calculated using the equivalent photon approximation (EPA) described in Section 2.3.1 whereas predictions for dissociative production are obtained using photon PDFs determined in global parton analysis in which the photon is treated as an additional point-like parton [43, 106, 107]. Many of the general-purpose MC generators listed above are capable of modelling elastic photon-fusion processes at the LHC and some can also model dissociative processes, for example MadGraph. However, some specialised MC event generators also exist and are briefly discussed in this section if relevant to this thesis.

LPair

LPair [108, 109] is a matrix element generator devoted to photon-induced production of lepton pairs in lepton-lepton, lepton-hadron and hadron-hadron interactions. It is capable of generating both elastic and dissociative events. The single-dissociative $\gamma\gamma \rightarrow ll$ process is modelled in this thesis using LPair interfaced with Pythia using the NNPDF3.1 NLO QED PDF set [110].

SuperChic

SuperChic [56, 111, 112] is an event generator for exclusive and photon-initiated production in pp and heavy ion collisions. A range of SM final states are implemented with spin correlations where relevant. The program is interfaced to Pythia for showering and hadronisation. An important and unique feature of SuperChic is that it includes a fully differential treatment of the soft survival factor, discussed in Section 2.3.2, for both elastic and dissociative processes. The $\gamma\gamma \rightarrow WW$ process was not fully implemented and validated at the time of the analysis so SuperChic samples are not used in this thesis but are being validated for the differential measurement discussed in Chapter 6.

Chapter 3

The Large Hadron Collider and the ATLAS detector

"Anyone who thinks the Large Hadron Collider will destroy the world is a twat."

– Brian Cox

3.1 The Large Hadron Collider

The Large Hadron Collider (LHC) [113] is a (roughly) circular hadron accelerator situated 100 m underground across the border of France and Switzerland at CERN[6]. It is located in the tunnel of its predecessor, the Large Electron-Positron collider (LEP) [114], which collided electron-positron beams up to an energy of 209 GeV before its decommission at the start of the millennium to make way for the LHC. The LHC is fed by a series of accelerators, shown in Figure 3.1, accelerating protons (and heavy ions) to higher and higher energies before being injected into the LHC ring.

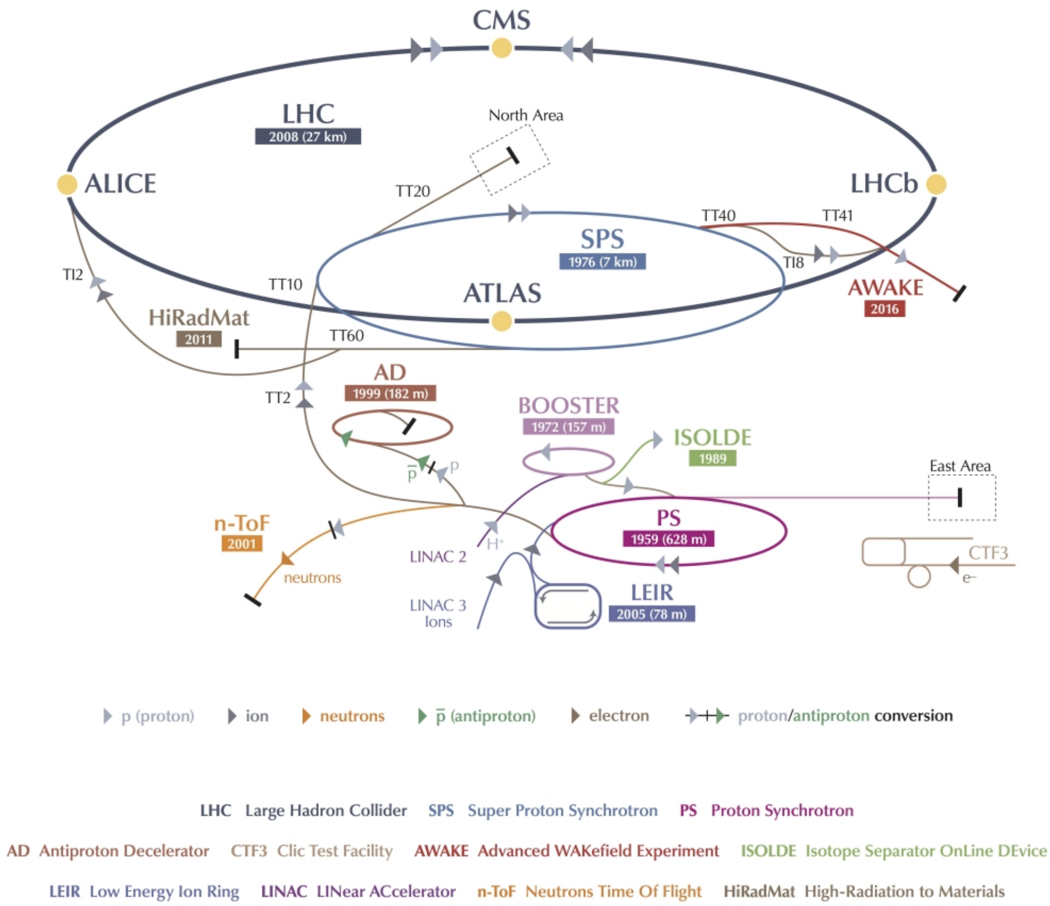

Figure 3.1: The CERN accelerator complex and experiments in Run 2 [115]. The LHC is the largest accelerator ring with a circumference of 27 km, shown here in navy blue. In Run 3, Linac2 has been replaced by Linac4 which accelerates hydrogen anions (H^-).

[6] CERN stands for *Conseil Européen pour la Recherche Nucléaire* and was founded as a European hub for science in 1954.

Each proton begins its journey inside a small bottle of hydrogen. From this bottle, hydrogen atoms are stripped of their electrons, leaving protons which are then inserted into Linear accelerator 2 (Linac2). Linac2 was the starting point for the protons used in experiments at CERN for 40 years until it was switched off for the last time at the end of 2018. Linear accelerators use radiofrequency (RF) cavities with oscillating fields to accelerate charged particles. For Linac2, this was up to a proton energy of 50 MeV before injection into the Proton Synchrotron Booster (PSB), Proton Synchrotron (PS) and Super Proton Synchrotron (SPS). The SPS became CERN's flagship particle physics programme when it switched on in 1976 and was the discovery place of the W and Z bosons in proton-antiproton collisions. It can accelerate protons up to energies of 450 GeV for insertion into the LHC ring. As of 2020, Linac4 became the source of proton beams for the CERN accelerator complex and is able to accelerate protons up to 160 MeV. The remaining linear accelerator in Figure 3.1 is Linac3 which accelerates heavy ions for injection into the Low Energy Ion Ring (LEIR) and PS before being sent to experiment areas and the other accelerators.

The final stage of acceleration is the LHC. The design centre-of-mass (CoM), $\sqrt{s}$, energy of the LHC (= 2× the beam energy) for pp collisions is 14 TeV and the LHC has been gradually ramping up in energy over the distinct periods in its lifetime: 2011–2012 marked Run 1 with $\sqrt{s}$ = 7–8 TeV and 2015–2018 marked Run 2 with $\sqrt{s}$ = 13 TeV. Like the other accelerating rings, the LHC uses RF cavities to accelerate the proton beams. Each cavity operates at 400 MHz, providing an electric field of 5 MV m^{-1} that flips polarity with the RF frequency. This alternating field leads to different proton accelerations. A synchronous proton has revolution frequency that matches the radiofrequency of the RF cavity and is not accelerated. Protons arriving slightly out of phase with this synchronous frequency will either be accelerated or decelerated toward the synchronous protons, creating segmented bunches of protons in the ring. In a maximum LHC fill, there are 2556 bunches of protons in the ring with a temporal separation of 25 ns. The spatial dimensions of the bunches depends on the location in the LHC ring and are determined by the magnetic focusing optics. The proton beams are steered around the LHC ring by a collection of copper-clad niobium-titanium dipole magnets which ramp up in current as the beams accelerate to provide a field of roughly 8 T at nominal beam energy. There are 1232 main dipoles, and each one is 15 m long. The beams are focused in the transverse plane with quadrupole magnets and erroneous trajectories and beam halos are scraped with collimators. Due to the high currents required to produce such high magnetic fields, the magnets must be operated in a superconducting state with negligible resistivity to prevent extreme heating. This is achieved by placing the machine components in a cryostat system containing superfluid helium at a temperature of 1.9 K.

After the acceleration stage, the LHC circulates bunches containing $\sim 10^{11}$ protons up to beam energies of 6.5 TeV[7] before bringing them to collide at four interaction points around the 27 km accelerator ring. Surrounding these interaction points are the four main experiments of the LHC: ATLAS, CMS, LHCb and ALICE. ATLAS [116, 117] and CMS [118] are general purpose detectors, built to have high angular acceptance and be sensitive to as many physics processes as possible. LHCb [119] is dedicated to studying heavy flavour states, such as B-mesons and charm hadrons, and searching for charge-parity (CP) violation. ALICE [120] is dedicated to studying heavy ion collisions during special heavy ion runs of the LHC in order to study the conditions of the early universe. These runs primarily collide lead-lead (Pb-Pb) and proton-lead (p-Pb) beams but xenon pilot runs have also been tested [121].

A key concept at a collider is that of *luminosity*. The rate of collisions, $\mathrm{d}N/\mathrm{d}t$, is given by

$$\frac{\mathrm{d}N}{\mathrm{d}t} = \sigma_{\mathrm{tot.}}\mathcal{L}, \tag{3.1}$$

where $\sigma_{\mathrm{tot.}}$ is the total cross-section and $\mathcal{L}$ is the instantaneous luminosity which measures the number of particles per unit area per unit time [122]. Assuming Gaussian bunches with identical parameters, the instantaneous luminosity is given by

$$\mathcal{L} = \frac{N_b N_p^2 f_{\mathrm{rev}} F}{4\pi\sigma_x\sigma_y}, \tag{3.2}$$

where N_b is the number of proton bunches, N_p is the number of protons within each bunch, f_{rev} is the revolution frequency of the beams, F is a geometrical factor related to the angle of beam crossing ($F \leqslant 1$, with $F = 1$ for head-on collisions) and $\sigma_{x/y}$ are the transverse sizes of the beam at the collision point [123]. The design value of the instantaneous luminosity at the LHC is $\mathcal{L} \sim 10^{34}\,\mathrm{cm}^{-2}\,\mathrm{s}^{-1}$ for pp collisions. The total integrated luminosity is the integral of the instantaneous luminosity over some period of time, T, and is used as a measure of the total data collected,

$$\mathcal{L}_{\mathrm{int.}} = \int_0^T \mathcal{L}(t)\mathrm{d}t. \tag{3.3}$$

[7] First collisions at a record beam energy of 6.7 TeV have recently been achieved in July 2022.

The higher the integrated luminosity over any given time period, the higher the number of events for any given physics process of interest. This means that processes with small cross-sections necessitate high-luminosity running in order to have sufficient statistical precision to perform hypothesis testing in the presence of backgrounds. The drawback of high-luminosity running is that the probability of multiple proton interactions increases. In Run 2 standard running conditions, the average number of interactions per bunch crossing,

$$\langle \mu \rangle = \frac{\mathcal{L}\sigma_{\text{inel}}}{N_b f_{\text{rev}}}, \tag{3.4}$$

was roughly 35 and peaked at around 70. The symbols in Equation (3.4) have the same meaning as in Equation (3.2) and σ_{inel} is the inelastic cross-section in pp collisions, taken to be 80 mb at $\sqrt{s} = 13\,\text{TeV}$. The majority of these pileup collisions are soft and produce additional hadronic activity that populates the detector environment, negatively impacting detector performance and object reconstruction. For this reason, the LHC also has dedicated low-luminosity runs for processes with relatively high cross-sections that ordinarily would suffer highly from pileup contamination. These "low-μ" runs typically have $\langle \mu \rangle \leqslant 1$.

The physics programme of the LHC spans multiple decades, demonstrated in Figure 3.2. This thesis is focused on pp collision data collected with the

Figure 3.2: The schedule of the LHC and future high-luminosity upgrade of the LHC (HL-LHC), accurate as of January 2022 [124].

ATLAS detector during LHC Run 2 between 2015–2018. During this run, the LHC delivered 156 fb^{-1} of integrated luminosity, with ATLAS recording 147 fb^{-1} and 139 fb^{-1} being labelled good for physics analysis based on the detector status. In addition, future prospects of measurements at the High-Luminosity LHC (HL-LHC) will be discussed in Chapter 7. The HL-LHC is a major upgrade planned for the LHC over the second half of this decade in order to increase its instantaneous luminosity by a factor of five and the integrated luminosity by a factor of ten beyond the original LHC design values [125]. Data taking is currently scheduled to commence in 2029 for 12 years. The HL-LHC will be a high-precision machine, allowing studies of electroweak symmetry breaking mechanisms through measurements of the properties of the Higgs boson and searches for new physics through the study of rare SM processes. The ATLAS detector will go through major upgrades during Long Shutdown 3 (LS3) in preparation for this phase of data-taking.

3.2 The ATLAS detector

A general purpose detector at the LHC requires:

- Fast, radiation-hard electronics and sensor elements;
- High granularity;
- Large acceptance in pseudorapidity with almost full azimuthal angle coverage;
- Good charged-particle momentum resolution and reconstruction efficiency;
- Vertex detectors close to the interaction point to observe secondary vertices;
- Very good electromagnetic calorimetry for electron and photon identification and measurements, complemented by full-coverage hadronic calorimetry for accurate jet and missing transverse energy measurements;
- Good muon identification and momentum resolution over a wide range of momenta and the ability to determine unambiguously the charge of high-p_{T} muons;
- Highly efficient triggering on low transverse-momentum objects with sufficient background rejection to achieve an acceptable trigger rate.

A Large Toroidal LHC Apparatus (ATLAS) is a general purpose detector designed to meet the above requirements and be sensitive to as many physics processes as possible over the lifetime of the LHC. It is designed to have an almost 4π solid angle acceptance and is comprised of several sub-detector systems, each with a different purpose. Charged particles are brought to collide at the interaction point (IP, specifically IP1 for ATLAS) in the centre of the detector volume. The innermost detector subsystem is the inner detector (ID) which aims to reconstruct the tracks of charged particles and consists of the pixel detector, semiconductor tracker (SCT) and transition radiation tracker (TRT). The ID is surrounded by a solenoid magnet with central 2 T magnetic field, designed to bend the trajectory of charged particles in the inner detector for the purpose of measuring both their charge and momentum. The next layer is formed of the calorimeter, designed to measure particle position and energy for all incident photons, electrons and hadrons. The calorimeter is divided into an electromagnetic (ECal) and hadronic (HCal) calorimeter to aid particle identification. The outer layer of ATLAS is the muon spectrometer (MS), designed to measure the trajectories of muons which pass through the detector material without being stopped. The MS is immersed in an up to 4 T field from a superconducting toroid magnet, again designed to bend the trajectories of the muons so that both their charge and momentum can be measured. A diagram of the ATLAS detector is shown in Figure 3.3, with the sub-components labelled. ATLAS uses a cylindrical coordinate system, with the origin at IP1 and the z-axis along the beamline. The x-axis points inwards, towards the centre of the LHC ring, and the y-axis points upwards to the ground's surface. A diagram of the ATLAS coordinate system with respect to the LHC ring is shown in Figure 3.4. The two opposite ends of the detector are labelled A-side and C-side. With $z = 0$ at IP1 at the centre of the detector, A-side is in the $+z$ direction and C-side $-z$[8]. As the colliding partons can have different momenta in a hadron collider, it is not possible to know the boost of the initial colliding system and it is conventional to instead use the transverse momentum of the collision, $p_{\mathrm{T}} = \sqrt{p_x^2 + p_y^2}$, which must be conserved to be zero. As the boost is unknown, it is also desirable to define quantities that are invariant to boosts along the beam axis (in z). One such example of a Lorentz invariant quantity is differences in *rapidity*, Δy, where rapidity is defined by

$$y = \frac{1}{2} \ln \left(\frac{E + p_z}{E - p_z} \right), \tag{3.5}$$

[8]For those familiar with the layout of Geneva and CERN, A-side points in the direction of the airport and C-side points towards Saint-Genis-Pouilly.

where p_z is the momentum measured along the beam axis. At high-energy particle colliders such as the LHC, where the colliding particles are ultra-relativistic, it is commonplace to instead use *pseudorapidity*,

$$\eta = -\ln\left(\tan\frac{\theta}{2}\right), \tag{3.6}$$

where θ is the angle with respect to the beamline measured in the $z-y$ plane. Due to the presence of the beampipe, it is not possible to measure particles as $\theta \to 0$ or $\theta \to \pi$. In the ATLAS detector, this corresponds to a maximum pseudorapidity acceptance of around $|\eta| = 5$ but the ability to reconstruct particular physics objects is highly dependent on the acceptance of each sub-detector, discussed in more detail in the following subsections. The angular separation between two objects measured in the detector is given by

$$\Delta R = \sqrt{\Delta\eta^2 + \Delta\phi^2}. \tag{3.7}$$

3.2.1 Inner Detector

The inner detector (ID) [128, 129] is the closest sub-detector to the interaction point. It is 7 m long and has an outer radius of 115 cm (limited by the cryostat

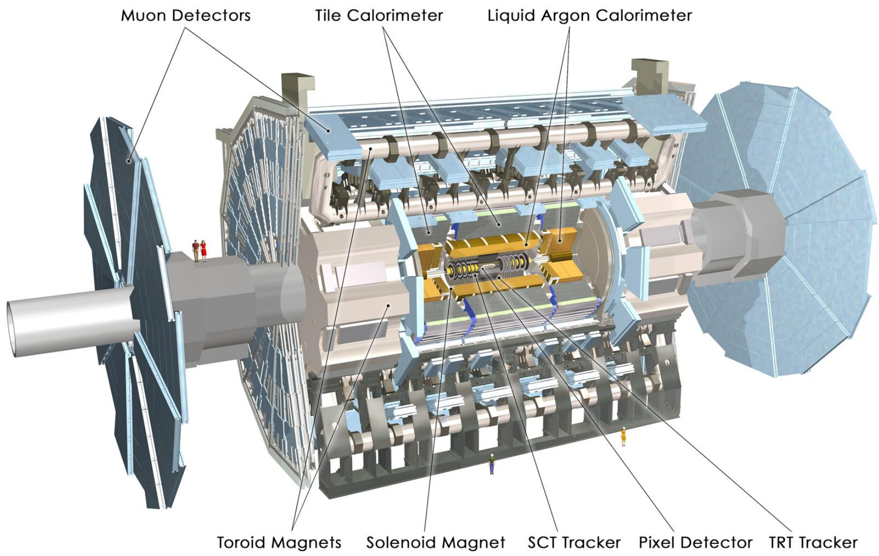

Figure 3.3: A computer generated image of the whole ATLAS detector, with sub-components labelled [126]. A slice is removed to show each detector component, including the interaction point in the centre of the image. The detector is cylindrically symmetric and is 44 m long with 25 m diameter. Four people are shown for scale.

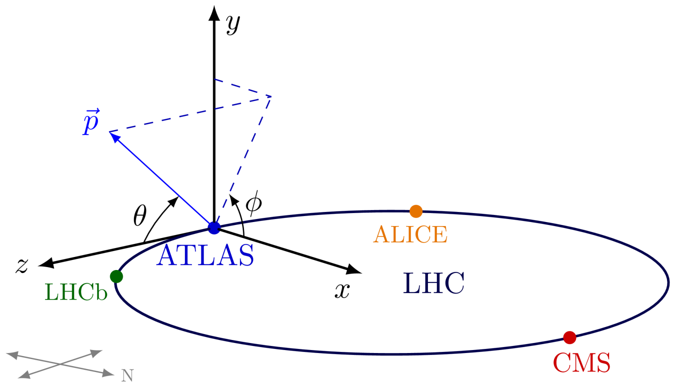

Figure 3.4: The ATLAS coordinate system with respect to the LHC ring [127]. $+x$ points towards the centre of the LHC ring, $+y$ points towards the ground surface. The z-direction is the beam axis at the interaction point. The rotation in the $x-y$ plane is given by ϕ and the rotation in the $z-y$ plane is given by θ.

containing the EM calorimeter), with an inner radius of approximately 33 mm. As the ID is very close to the interaction point, it must be extremely radiation hard and have high granularity to perform precision track and vertex reconstruction in the dense LHC interaction environment. The ID is further subdivided into three components. The pixel detector forms the innermost layer, followed by the semiconductor tracker (SCT). The outer layer is the transition radiation tracker (TRT). Each component has a cylindrical barrel region arranged in concentric cylinders around the beam axis and two disk end-caps perpendicular to the beam axis, covering a total pseudorapidity range of $|\eta| < 2.5$. All ID subsystems are shown in Figure 3.5.

The pixel detector [131] consists of three barrel layers and two end-caps with a total of 1,744 silicon sensors. Each sensor consists of roughly 46,000 pixels with an area of $50 \times (400–600)\,\mu\text{m}^2$, where the longer pixels cover the gaps between adjacent readout electronics. Silicon is a semiconductor which has a depleted region when a potential difference is supplied. A passing ionising particle creates electron-hole pairs which drift to the electrodes to produce an electrical signal. Silicon is used in many tracking applications due to its fast response time and high radiation hardness. The Insertable B-layer (IBL) [132] is an additional inner layer that was added to the pixel detector during Long Shutdown 1 (LS1), improving vertex reconstruction, especially for B-hadron decays. It consists of a combination of planar [133] and 3D silicon pixels [134]. Planar pixels have each electrode implanted on the surface of the active material, whereas 3D silicon pixels have both electrodes created inside the detector bulk. The maximum charge drift and depletion distance is then set by electrode spacing which results

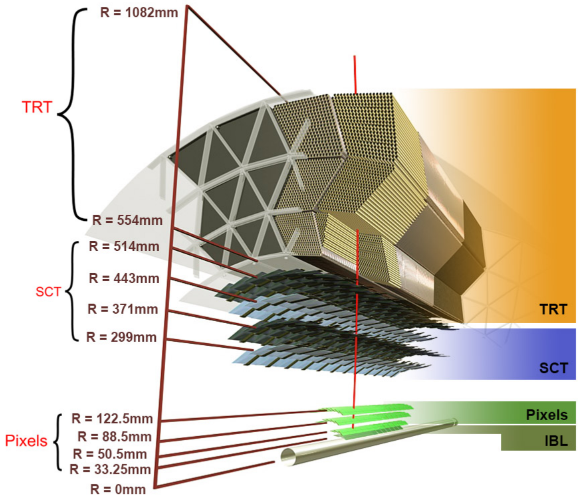

Figure 3.5: A 3D visualisation of the structure of the barrel of the ID. The beam pipe, the IBL, the Pixel layers, the four cylindrical layers of the SCT and the three layers of TRT barrel modules consisting of 72 straw layers are shown [130].

in reduced collection time and required depletion voltage compared to planar pixels. In total, one IBL and three pixel hits are expected from each charged particle passing through the ID.

The pixel detector is followed by the SCT [135] which is a silicon strip tracker. The strip sensors have a larger area of 126 mm × 80 µm which means that they are not as good as pixels at solving hit ambiguities in dense tracking environments and have position resolution along the beam axis roughly ten times worse than the pixels. However, pixels cannot be used for the entire ID due to high material budget and high cost. This lower strip resolution is partly mitigated by having two pairs of single-sided strip sensors glued back-to-back in each module with a 40 mrad angle between them in order to provide a 2D measurement. The SCT consists of 4088 modules in four barrels and two end-caps of nine disks each. Both the pixel and SCT must be cooled to around −10 °C to reduce noise caused by thermal excitations.

The TRT [136–138] consists of gas-filled drift tubes, called *straws*. Each straw has a diameter of 4 mm and has a 30 µm thick gold-plated tungsten wire anode at the centre. In a similar principle to the silicon trackers, when an ionising particle passes through the gas-filled straws, the gas becomes ionised and the charges are collected by the anode wire at the centre, providing an electrical signal. In the barrel region of $|\eta| < 1$, there are roughly 50,000 straws parallel to the beam which have a maximum length of 150 cm. In the two end-cap regions combined, $1 < |\eta| < 2$, there are roughly 320,000 straws in the radial direction. Each channel provides a drift-time measurement, giving a spatial resolution of 170 µm per straw. The high hit multiplicity per ionising particle (around 36 hits per track) of the TRT improves the spatial resolution of the ID as a whole and the radial extent improves momentum resolution. In addition, the TRT has the ability to aid in particle identification. When charged particles pass through a medium that has two different dielectric constants, transition radiation can be produced. Polypropylene-polyethylene foils are placed between each straw to induce this radiation, the intensity of which is proportional to the Lorentz γ factor of the transitioning particle [139], particularly improving the ability to distinguish between electrons and light charged-hadrons like pions. A downside of the TRT compared to the silicon tracking is the slow drift time of charge carriers, meaning the TRT is not well equipped to deal with high particle multiplicity.

3.2.2 Calorimeters

The calorimetry system [140] is designed to measure particle energies and position and is split into two subsystems: those using liquid argon (LAr) technology [141] and those using tile scintillator technology [142]. Liquid argon is a common choice of active material due to its fast and uniform response and radiation hardness. It is also common to hear the calorimetry system discussed separately as the Electromagnetic Calorimeter (ECal) and the Hadronic Calorimeter (HCal). The ECal uses solely LAr technology whereas the HCal uses a combination of LAr and tile scintillators in different η-regions. Additionally, there is a forward calorimeter (FCal) which also uses LAr technology in the forward region ($3.1 < |\eta| < 4.9$). Together, the calorimeters are designed to make accurate measurements of the energy and position of electrons and photons, energy and direction of jets, particle identification, and reconstruction of missing transverse energy, $E_{\mathrm{T}}^{\mathrm{miss}}$, in an event over the maximum possible acceptance in η. They also have trigger capabilities and form part of the Level-1 trigger system described in Section 3.2.5. A detailed view of the calorimetry system is shown in Figure 3.6. The ECal (green) covers the region $|\eta| < 3.2$, the HCal barrel (light blue) covers $|\eta| < 1.7$ and the end-caps (orange) cover $1.5 < |\eta| < 3.2$. All calorimeters are sampling calorimeters which have active and passive layers interleaved. The passive layers are designed to induce either electromagnetic or hadronic showers such that the propagating particle is eventually stopped within the calorimeter volume. As the passive layers absorb some of these shower particles, the active layers only sample some of

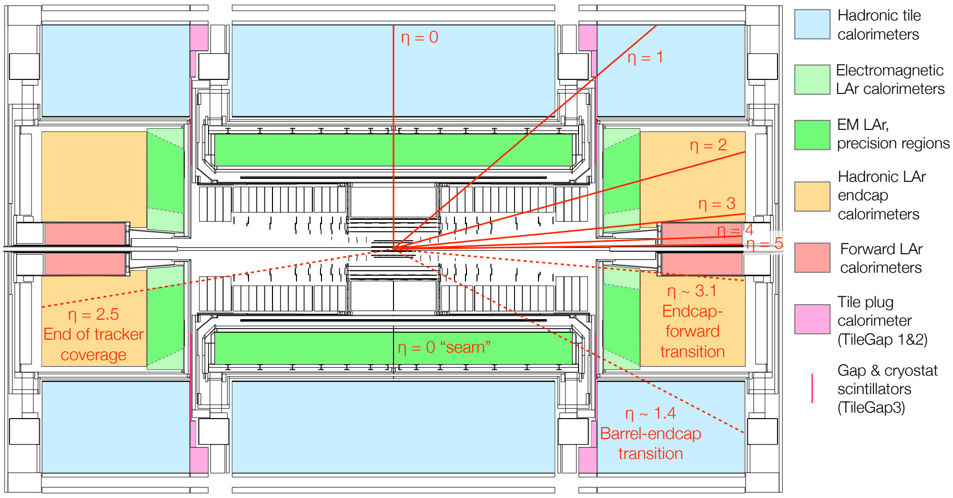

Figure 3.6: Layout of the ATLAS calorimeters with pseudorapidity (η) values marked for reference. The electromagnetic (EM) barrel has a seam in the construction at $\eta = 0$ which can impact jet energy resolution. The transition between the EM end-cap and extended region around $\eta \sim 2.5$ involves a dramatic change in the material layers due to the absence of ID layers above $|\eta| = 2.5$. The regions corresponding to the transition from barrel to end-cap ($\eta \sim 1.4$) and from end-cap to forward calorimeter ($\eta \sim 3.1$) are also given for reference [143].

the total shower energy. There are gap-regions (crack-regions) in the calorimetry system to provide passage for electronics and service pipes.

The ECal is a lead-liquid-argon detector with zig-zag "accordion" geometry, designed to trigger on and provide precision measurements of electrons and photons. The ECal itself cannot distinguish electrons from photons, this is only possible when paired with tracking information from the ID. An accordion geometry is used to maintain high acceptance whilst limiting dead space and shortening electronic cables for fast readout. Electrons with energies above $\sim 1\,\text{GeV}$ interact electromagnetically with the passive lead layers and are decelerated, leading to bremsstrahlung emission. High energy photons passing through the material create electron-positron pairs which produce further photons via bremsstrahlung. This results in a cascade of emission, referred to as an electromagnetic shower. The shower continues until photon energies fall below $2m_e$ and pair production is no longer possible. Once charged particles lose most of their energy via bremsstrahlung, they ionise the active material and created pairs drift towards the electrodes. The amount of collected signal is then calibrated to the particle energy.

In the central region of $|\eta| < 1.8$, The ECal is preceded by a presampler detector which is used to correct for the energy lost in the material before the calorimeters. The total thickness of the ECal is between 24–26 radiation lengths, where the radiation length, X_0, is the average distance that an electron travels through a material before its energy is reduced by a factor of $1/e$ or equivalently, 7/9th of the distance a photon travels through a material before its intensity is reduced by the same factor[9]. The radiation length is a material property, depending on the atomic number, Z, and atomic mass, A of the material. The barrel of the ECal is contained within the cryostat which surrounds the inner detector and solenoid magnet. Two end-cap cryostats house the end-caps of the ECal, HCal and forward calorimeters which are kept at a temperature of $-184\,^{\circ}$C to keep the argon in its liquid form.

Not all particles are stopped in the ECal. Hadrons tend to penetrate further than electrons and photons through the detector. Hadrons can interact both electromagnetically and hadronically with the calorimeter material, leading to both electromagnetic and hadronic showers. The requirements for hadronic calorimetry are mainly related to jet reconstruction and $E_{\mathrm{T}}^{\mathrm{miss}}$. The characteristic distance in the HCal is the *interaction length* or *mean free path*, λ, which is proportional to $A^{1/3}$. The hadronic calorimeter thickness is approximately 11 interaction lengths (including support structures) at $\eta = 0$ for shower containment, both for energy resolution reasons and to reduce the background in the muon spectrometer, referred to as *punch-through*. At low luminosity, electronic noise in the calorimeter must be minimised and at high luminosity, fast response is necessary to reduce the impact of pileup.

The HCal barrel is a cylindrical detector divided into three layers. It consists of 3 mm thick plastic scintillator plates called *tiles*, interleaved with 14 mm thick iron absorber plates to induce showers. The tiles are perpendicular to the beam axis and staggered in depth. The signal from the active scintillator material is read out by wavelength shifting fibres into photomultipliers. The HCal barrel also acts as the main flux return for the solenoid magnet. Due to the high multiplicity of jets in the forward region, the HCal end-caps must have high radiation hardness and a copper-LAr detector with parallel plate technology is used. Each HCal end-cap consists of two wheels of equal diameter but different thickness. The innermost wheels have a thickness of 25 mm whereas the outer wheels have a thickness of 50 mm.

[9] where e is Euler's number ($= 2.71828\ldots$), not the electric charge.

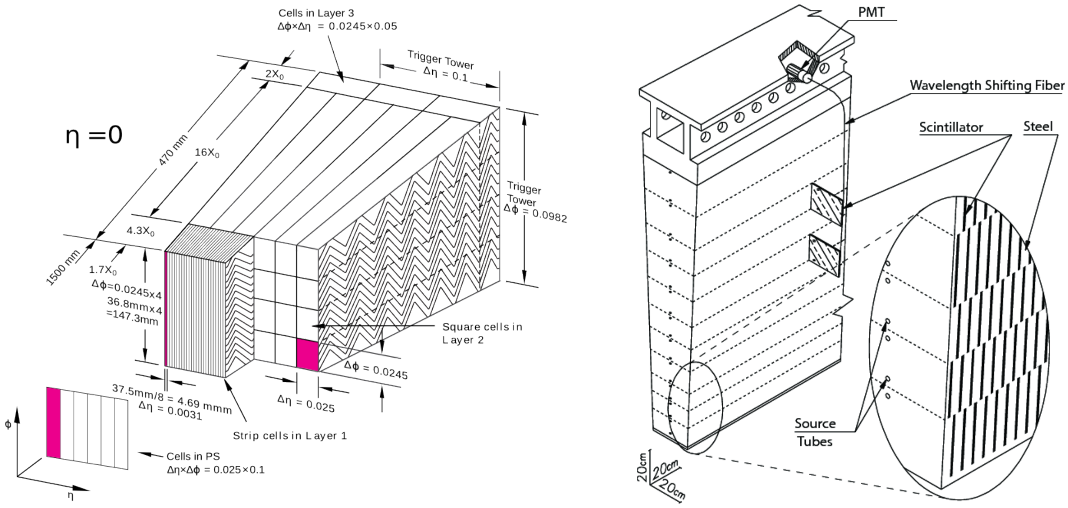

(a) ECal barrel module (b) HCal tile module

Figure 3.7: (a) An accordion LAr module that forms part of the ECal barrel region [141]. The higher granularity layer is the closest to to the beam axis. (b) A module of the HCal tile calorimeter [142].

The forward calorimeter inner face is roughly 4.7 m from the interaction point. Space is limited in the FCal region and therefore it is a high density LAr detector, formed of three layers. Each layer consists of a metal matrix with regularly spaced longitudinal channels filled with concentric rods and tubes. Each FCal end-cap consists of three modules. The first module is designed to measure electromagnetic showers and uses copper as the passive layer. The other two modules use tungsten for the passive layer and measure the energy of hadronic interactions.

The combination of the ECal, HCal end-caps, and forward calorimeter are collectively referred to as the LAr calorimeter, and the HCal barrel is the tile calorimeter. The LAr calorimetry system has a radius of 2.25 m and a total length of 13.3 m. The Tile calorimeter has an outer radius of 4.23 m and is slightly shorter in length than the LAr calorimeter with a total length of 12.2 m. The combined weight of all calorimeters is around 4000 tonnes. Diagrams of a LAr module in the ECal barrel and a tile module of the HCal barrel are shown in Figure 3.7a and Figure 3.7b, respectively.

3.2.3 Muon Spectrometer

As muons are roughly 200 times heavier than electrons, they are less subject to bremsstrahlung radiation and high energy muons penetrate through the calorimetry systems. The Muon Spectrometer (MS) [144] is the outermost layer of the ATLAS detector designed to measure these highly penetrating muons. It has stand-alone triggering and momentum measurement capabilities over a wide range of transverse momentum, pseudorapidity, and azimuthal angle.

The MS layout is based on the magnetic deflection of muons by the toroid magnet. In the barrel region, muon tracks are measured by three layers of chambers arranged in a cylindrical pattern around the beam axis and chambers are installed perpendicular to the beam axis in the end-cap regions. Monitored drift tubes (MDTs) provide precision measurements of track coordinates in the principle bending direction of the toroid field over most of the pseudorapidity coverage. In the forward regions close to the beampipe, Cathode Strip Chambers (CSCs) with high granularity are used. The MS forms part of the level-1 trigger system, covering the $|\eta| \leqslant 2.4$ region and implemented by resistive plate chambers (RPCs) in the barrel and thin gap chambers (TGCs) in the end-caps. The RPCs and TGCs are designed to provide fast-response measurements of muon position and momentum and complement the better resolution but slower response measurements of the MDTs and CSCs in the offline reconstruction. All of the components of the MS are shown in Figure 3.8.

The precision MDTs are filled with a mix of argon and carbon dioxide and operate on a similar principle as the TRT straws with a wire anode at the centre. A single MDT is 30 mm in diameter with a 50 µm diameter gold-plated tungsten-rhenium

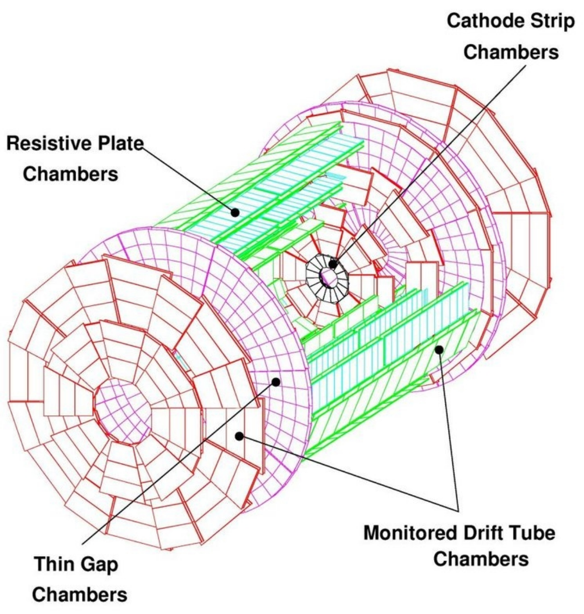

Figure 3.8: A schematic of the Muon Spectrometer components [145].

wire and has a spatial resolution of around 80 µm when operated at a gas pressure of 3 bar. The MDT lengths vary from 70–630 cm depending on their position in the detector. To improve the overall precision muon tracking beyond 80 µm resolution, multiple layers of MDTs are used to achieve a spatial resolution of roughly 30 µm. In total, there are around 1,200 MDT chambers with roughly 355,000 individual MDTs. The CSCs have a similar spatial resolution to the MDTs but better performance at the higher rates expected in the forward regions. They cover the region of $2.0 < |\eta| < 2.7$. The chambers are filled with gas and have wire anodes like the MDTs, also containing an argon and carbon dioxide gaseous mix. However, the read-out is performed by cathode strips orthogonal to the anode wires that are located at the edge of the chamber. Important characteristics of the CSCs are small charge drift times $\leqslant 30$ ns (i.e. fast response time), good time resolution of $\sim$ 7 ns, good multi-track resolution, and low neutron sensitivity (i.e. radiation hard). The CSCs provide measurements in η and ϕ. Diagrams of an MDT chamber and CSC are shown in Figures 3.9a and 3.9b, respectively.

The LHC bunch-crossing period of 25 ns is much shorter than the maximum drift time of the MDTs. Therefore the RPCs and TGCs form the fast muon triggering system. The RPC is a gaseous (mostly tetrafluorethane) detector with poor spatial resolution but excellent timing resolution of $\sim$ 1 ns. A trigger chamber is made from two rectangular detector layers, each one read out by two series of pick-up strips arranged orthogonally, providing two coordinate measurements which are required for the offline pattern recognition. The TGCs have a similar design to the CSCs, with an electric field configuration and the small wire distance providing short drift time and therefore good timing resolution.

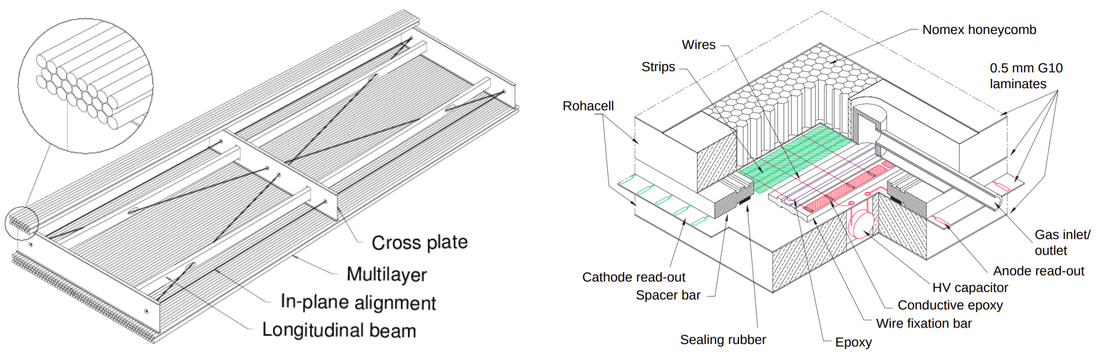

(a) MDT chamber (b) CSC

Figure 3.9: The precision detectors of the MS: (a) Monitored Drift Tube (MDT) chamber and (b) Cathode Strip Chamber (CSC) in the Muon Spectrometer [144].

The muon spectrometer has a momentum resolution of typically 2–3 % across most of the kinematic range. This increases to around 10 % for muons with very high momenta of $p_{\mathrm{T}} = 1\,\mathrm{TeV}$, limited by detector resolution. The resolution at lower momenta is limited by energy loss fluctuations in the calorimeters and multiple scattering effects.

3.2.4 Magnet System

The ATLAS detector has two separate magnet systems that aim to bend the trajectories of charged particles such that their momentum can be measured to high precision and the sign of their charge determined. The innermost magnet is a superconducting solenoid [146], providing an axial field of 2 T central strength for the inner detector. The iron yoke for magnetic flux return (the tile calorimeter elements) is separate from the coil end, resulting in a non-uniform magnetic field along the beam axis. For this reason, the alignment of the magnet structure must be known to within ±1 cm along the beam axis. The solenoid is located in the space between the inner detector vacuum vessel and inner cold vessel of the barrel LAr cryostat. Due to its location, the solenoid material budget must be low to avoid affecting calorimeter performance. The technology chosen uses an indirectly-cooled aluminium-stabilised superconductor and a vacuum vessel wall is shared with the LAr calorimeter to reduce material. Also due to its location within the detector, the solenoid must be robust, radiation hard, and quench safe to prevent unnecessary interventions.

The second magnet system in ATLAS is the air-core superconducting toroid magnet, located outside of the calorimeter systems and amongst the muon spectrometer components. It is designed to produce a large-volume magnetic field covering the rapidity range $0 \leqslant |\eta| \leqslant 2.7$, with an open structure to minimise multiple scatterings. The toroid consists of the barrel toroid [147] and end-caps [148]. The barrel toroid consists of eight flat coils assembled radially and symmetrically around the beam axis. It extends over a length of 25 m with an inner bore of 9.4 m and an outer diameter of 20.1 m. Each coil is contained within its own cryostat. The assembly of the toroid requires a very strong and rigid mechanical support due to both the weight of the magnet and the magnetic forces exerted during operation, with each barrel coil subject to a radial force of 1,100 tonnes. During normal operation, the end-caps experience an axial force of 320 tonnes toward the centre of the detector. The end-caps are constrained by the geometry of the full detector and occupy the axial space between 7.63–12.63 m,

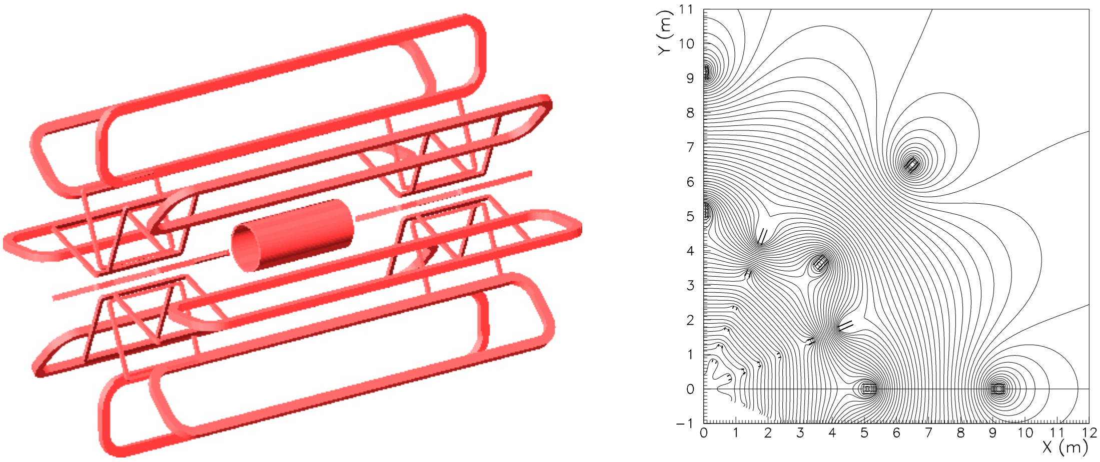

(a) The magnet system

(b) The magnetic field

Figure 3.10: The (a) ATLAS magnet system with central solenoid and barrel plus end-cap toroid coils, and (b) one quadrant of the magnetic field in the plane perpendicular to the beam axis and at the transition region between the barrel and end-cap toroid [149].

with a radial span from 1.2–5.0 m. The operating temperature of the toroid is 4.5 K.

The aim of the toroid system is to bend muon tracks for precision measurements in the MS. For muons with $|\eta| \leqslant 1.0$, magnetic bending is provided by the barrel toroid. For $1.4 \leqslant |\eta| \leqslant 2.7$, the bending is provided by the end-caps. For the intermediate $1.0 \leqslant |\eta| \leqslant 1.4$ transition region, it is a combination of the barrel and end-cap fields. This magnet configuration provides a field that is mostly orthogonal to the muon trajectories, while minimising the degradation of resolution due to multiple scattering. The combination of all magnet systems is shown in Figure 3.10a and the resulting field pattern is shown in Figure 3.10b. This field leads to helical trajectories of particles within the detector and any changes in the field must be understood to accurately reconstruct particle momenta.

3.2.5 Trigger and Data Acquisition

It is neither possible nor desirable to process and store all of the information from collisions within the ATLAS detector. Proton bunches collide in the LHC at a rate of 40 MHz with an average number of interactions per bunch crossing around $\langle \mu \rangle = 35$ at IP1 in Run 2. The majority of these collisions are soft processes not usually considered interesting for physics analyses which tend to search for a hard-scatter or high-mass interaction. ATLAS employs a two-level trigger and data acquisition (TDAQ) system to quickly decide which events to process and store for further analysis. Triggers allow fast event rejection based on combined information from individual subsystems. Such decisions are made based on the characteristics of the event such as if energy or momentum of an object is below some pre-defined threshold. Each subsequent selection increases in granularity at the cost of processing time.

The first level, Level-1 (L1) [150], is a hardware-based trigger which reduces the data-rate from 40 MHz to 100 kHz. The L1 trigger works on a subset of information from the calorimeters (L1Calo) and the MS (L1Muon) and operates within a latency of 2.5 µs. It delivers geometric information on where signals were detected, called a *region of interest* (RoI), as well as the candidate object type and the energy thresholds passed. Additional information is provided by a Minimum Bias Trigger Scintillator (MBTS) and Beam-Pickup Monitors (BPMs) for triggering on filled and unpaired bunches. If an event passes an L1 trigger selection, it is passed to the central trigger processor (CTP) which applies any pre-scaling before passing the selected event onto the next stage of processing. Pre-scales reduce the trigger rates of common signatures, such as low-p_{T} jets, by only processing a fraction of that signature and discarding the rest.

The second level is the software-based High-Level Trigger (HLT) [151] which further reduces the rate from 100 kHz down to 1 kHz. The HLT uses fast versions of object reconstruction algorithms, referred to as *online* reconstruction as it is done in real-time as new data is processed by the detector. However, the latency is much higher than the L1 trigger at around 500 ms. The full *offline* object reconstruction takes even longer and is described in Section 3.3. To meet the HLT requirements, reconstruction algorithms are staged in several steps to form trigger *chains*. Groups of chains are collected into groups with related signatures, called trigger *menus*.

ATLAS aims to use inclusive triggers as much as possible in its physics selection strategy in order to maximise physics coverage and to be as open as possible to new and possibly unforeseen signatures. The configuration of the entire trigger system is designed to be flexible so that algorithms, their thresholds, and pre-scales can be easily changed between data-taking runs.

3.3 ATLAS object reconstruction

The ATLAS detector does not directly register particles but rather electrical signals from different subdetector systems which have to be combined to create physics objects. Various reconstruction techniques are used, depending on the object signature. The standard reconstructed objects in ATLAS are tracks, photons, electrons, muons, taus, jets and missing transverse energy, $E_{\mathrm{T}}^{\mathrm{miss}}$, which is related to particles which do not interact with the detector such as neutrinos or possible BSM signatures. The *offline* reconstruction of these objects is described in this section, where offline means that the reconstruction can be performed at any time and does not form part of the triggering procedure. When reconstructing objects, it is common to assign categories or *working points* (WPs). Three standard selection WPs are designed to cover the needs of the majority of physics analyses. In order of increasing purity and decreasing reconstruction efficiency, generic working points are the loose, medium, and tight WPs, where the objects passing the medium WP requirements constitute a subset of those passing loose WPs, and tight are a subset of those passing medium. Additional WPs might be developed for specific analysis requirements. The selection criteria of the WPs used in this thesis are documented in the relevant subsections. This thesis is only concerned with the reconstruction of tracks, electrons and muons.

3.3.1 Track reconstruction

In a high-pileup event, there can be of the order of 1000 charged particles passing through the ATLAS detector. The information from the Inner Detector (Section 3.2.1) is used to reconstruct the trajectories of charged particles, called *tracks*. A large part of this thesis revolves around measuring charged particle multiplicity within a region in z and therefore the number of tracks reconstructed in an event and their reconstructed position with respect to the beam axis are very important quantities.

Charged particles traversing a homogeneous magnetic field have helical trajectories which can be described by a reference point plus five parameters, shown in Figure 3.11. ATLAS uses a perigee representation where track properties are measured at the point of closest approach to the beam line. The reference point chosen is the average position of the interactions in the bunch crossing which is an event-dependent quantity. In principle, this reference point can be shifted to any position along the beam axis and, in this thesis, it is often taken to be the reconstructed vertex z_0-position for a given track of interest. The

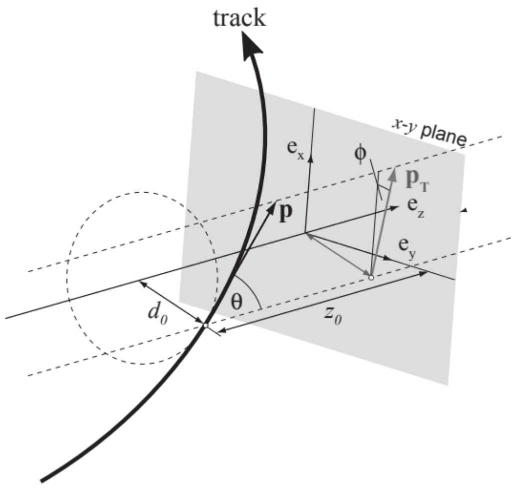

Figure 3.11: Track parameters in the perigee representation: d_0, z_0, ϕ, θ, and q/p [152]. The e_z axis is the beam axis and the point of reference is the centre of the interaction region of the bunch crossing.

first two parameters are the transverse and longitudinal impact parameters, d_0 and z_0, defined as the transverse and longitudinal distances of the single point of closest approach to the reference point. The other three parameters are the azimuthal angle, ϕ, and the polar angle, θ, of the track momentum at the reference point, and the ratio q/p of the charge of the reconstructed track divided by the magnitude of its momentum.

There are two types of tracking algorithms, referred to as *inside-out* and *outside-in*:

- **Inside-out:** As the name suggests, the inside-out algorithm starts by finding track candidates in the innermost layers of the ID and then extends tracks out into the TRT if consistent TRT hits are found. This is the default track reconstruction for high-quality tracks originating from primary interactions in *pp* collisions. It is described in detail below.

- **Outside-in:** To increase acceptance to particles produced at a greater distance from the beamline, such as electrons originating from photon conversions in the detector material, a secondary back-tracking pass is performed using the detector hits not already assigned to tracks from the primary pass. Outside-in track reconstruction is only attempted in RoI already determined by the ECal. The outside-in reconstruction begins with hits in the TRT and works backwards to match them to pixel+SCT hits that have not already been used in the inside-out track reconstruction.

There is an additional track-category called *TRT-standalone* which forms tracks from TRT segments not used in the default methods above. These are typically low quality tracks with high fake rates as they are very short tracks that only rely on 1D measurements. TRT-standalone tracks are not used at any point in this thesis and shall not be discussed further. Other special tracking algorithms exist, e.g. large-radius tracking for long-lived particle searches, but are also not relevant to this thesis and shall not be discussed.

The inside-out process begins by clustering pixel and SCT hits which are interpreted as deposits left by individual traversing charged particles. Pairs of one-dimensional clusters on either side of a SCT sensor module or single pixel clusters are then converted into 3D *space-points* within the inner detector geometry, with position uncertainties determined by the detector resolution and alignment uncertainties. The track seeding and finding then proceeds as follows:

(a) Track seeds are formed of triplets of space-points in either of the pixel or SCT detector within some loose selection criteria, e.g. a maximum ΔR between space-points to create track seeds consistent with a smooth trajectory.

(b) Confirmation space-points are added to the track seeds where possible. A confirmation space-point is a fourth space-point from an adjacent radial layer which results in a new seed that has compatible curvature to the original when replacing the outermost space-point from the original seed. This indicates that the confirmation space-point is likely from the same trajectory.

(c) Then search roads are built through the remaining detector based on the estimated seed trajectory. Search roads are formed of sets of detector modules that can be expected to contain clusters compatible with the seed. The use of search roads reduces computational time by reducing combinatorics.

These steps are demonstrated in Figure 3.12. The seeds are extended along the search roads using a combinatorial Kalman filter [153], which searches for adjacent clusters both outwards and inwards radially while attempting to smooth the trajectory. If a seed candidate fails the track finding, it is checked for calorimeter compatibility. If the seed is within a calorimeter RoI, the track-finding procedure is redone, allowing a loosening of some of the quality requirements. This allows recovery of efficiency for electron reconstruction, in which electron trajectories are bent by bremsstrahlung emission. The end result of this track finding/pattern recognition process is a set of track candidates that then undergoes further refinement.

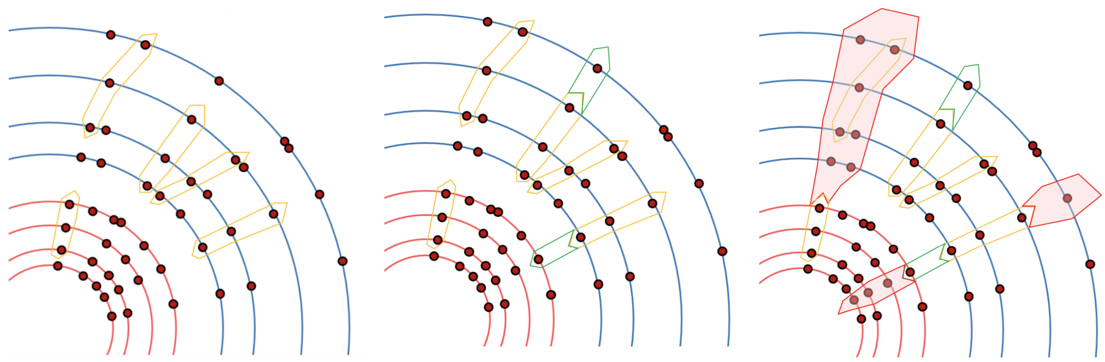

(a) Track seeding (b) Confirmation space-points (c) Search roads

Figure 3.12: A pictorial representation of the track finding process [152]. The red layers represent the pixel detector and the blue layers represent the SCT. The red circles are reconstructed cluster hits in the silicon modules.

The Kalman filter approach is computationally fast but relatively imprecise and does not solve track ambiguities, where several track fits could be consistent with a set of seeds and some space-points might contribute to multiple tracks. Ambiguities in tracking are common in dense interaction environments like the LHC. Therefore, track candidates produced by the Kalman filter are passed through an ambiguity solver which scores track candidates based on a range of quality criteria and rejects lower-quality candidates sharing a large number of associated hits with higher-quality tracks. In certain topologies such as high multiplicity jets where the separation of charged particles is expected to be less than detector resolution, a limited number of shared hits is permitted to maintain high performance. After the ambiguity solver, the refined and purified track candidates are then re-fit using a global χ^2 method [154] to obtain the final, high-precision track parameter estimate.

Finally, an extension into the TRT subdetector is attempted with a procedure similar to detailed above. This is done in order to profit from additional measurements on the track, in particular for momentum resolution and particle identification as described in Section 3.2.1. A road search is performed into the TRT volume starting from the position estimation of the pixel+SCT track candidate and a Kalman filter is used again for the track finding. Then the TRT hits are added to the tracks and the whole tracks are refit with the global χ^2 fitter if the extension is successful. Silicon-only tracks are kept if the fit worsens with the addition of the TRT hits.

Outside-in tracking is performed after the inside-out pass, using detector hits not already assigned to tracks. Outside-in track reconstruction is only attempted in RoI determined by deposits in the electromagnetic calorimeter and begins by finding segments of hits in the TRT compatible with the RoI. If compatible TRT segments are found, short silicon track seeds are constructed in the SCT, only including hits near to the TRT. These seeds are then extended into track candidates using the same procedure as for the inside-out pass. Resulting track candidates are then again extended (back) into the TRT using the collected hit segments.

In MC, we have access to both the reconstructed tracks and the charged particles in the generator record and we can assign so-called *truth-match probabilities* to the reconstructed tracks which are measures of ID components input to the track fit that also had energy deposited by the particle in simulation [155]. The probability that a reconstructed track is directly matched to a given charged particle trajectory, P_{match}, is defined as

$$P_{\mathrm{match}} = \frac{2N^{\mathrm{pix}}_{\mathrm{common}} + N^{\mathrm{strip}}_{\mathrm{common}}}{2N^{\mathrm{pix}}_{\mathrm{track}} + N^{\mathrm{strip}}_{\mathrm{track}}}, \tag{3.8}$$

where $N^{\mathrm{pix/strip}}_{\mathrm{common}}$ is the number of pixel/strip detector clusters associated to the track and the particle being matched, and $N^{\mathrm{pix/strip}}_{\mathrm{track}}$ is the number of pixel/strip detector hits assigned to the track. If a track has a truth-match probability greater than 0.5, we call this a true track and vice versa is considered a fake track. The performance of the track reconstruction algorithms can then be gauged from MC, primarily using two performance criteria; track efficiency and fake rate. The latter is the rate at which fakes are generated. Fakes occur either by false associations of hit to tracks or from fake noise hits. Tracking efficiency is defined as the number of selected reconstructed tracks matched to a truth particle with $P_{\mathrm{match}} > 0.5$ divided by the number of selected truth particles. True tracks in MC can be separated into track categories depending on whether they originate from pileup or a hard-scatter process, based on the information of the corresponding charged particle.

At each step of the reconstruction chain, track parameter cuts are defined and configured for the tracking algorithms. At the seeding stage, cuts are applied to estimates of track impact parameters and p_{T}, rejecting seeds failing to meet the cuts. Further tracking cuts are also applied during candidate finding and ambiguity resolution, rejecting any tracks failing to meet thresholds on impact parameters, p_{T} and the number of holes and hits per track. Due to the presence of the solenoid field, the bending of tracks in the ID can be used for charge identification and to calculate the track-p_{T}. The lower the track-p_{T}, the more it will bend in the magnetic field and generally there will be more tracking ambiguities which makes track fitting more difficult, especially in high pileup environments. Very low-p_{T} tracks will not reach all layers of the tracker and be impacted more by rescattering. For this reason, the spatial and momentum resolution of low-p_{T} tracks is worse than for high-p_{T} tracks. Resolution also worsens with increasing $|\eta|$ due to a longer extrapolation distance and shorter transverse distance from the beam axis.

Tracking recommendations and quality cuts are determined by the tracking combined performance group. The tracks used in this thesis pass the 2018 *tight primary* quality selections unless specified otherwise. These are:

- $p_{\mathrm{T}} > 500\,\mathrm{MeV}$, $|\eta| < 2.5$;
- Number of pixel+SCT hits $\geqslant 9$ for $|\eta| \leqslant 1.65$ and $\geqslant 11$ for $|\eta| > 1.65$;
- 1 or fewer hits shared with another track;
- Number of track *holes*, i.e. missing track hits in an active sensor where hits are expected, in the pixel detector (including IBL) $= 0$;
- Number of holes in the SCT $\leqslant 2$.

These selections result in a reconstruction efficiency of 84 % for central tracks ($|\eta| \leqslant 0.1$) and 62 % for $2.3 \leqslant |\eta| \leqslant 2.5$. The rate of fake tracks with $P_{\mathrm{match}} < 0.5$ is dependent on μ and the kinematic phase-space but is typically below 10 % for tracks with $p_{\mathrm{T}} > 500\,\mathrm{MeV}$ and below 5 % for tracks with $p_{\mathrm{T}} > 1\,\mathrm{GeV}$.

Vertex and beamspot reconstruction

Reconstructed tracks are then used to reconstruct primary and secondary interaction vertices, where a primary vertex (PV) is defined as the point in space where an interaction of the beam occurred. A secondary vertex is the point in space where an unstable particle produced in a beam collision decays and is usually selected as having larger d_0 parameter, i.e. its location is further away from the beam axis than PVs. Standard ATLAS vertex reconstruction in Runs 1 and 2 uses an iterative procedure to fit track positions [156]. In each iteration, less compatible tracks are downweighted and the vertex position is recomputed. All vertices with at least two associated tracks are retained as PV candidates. In Run 3, the iterative vertex finding method will be replaced with an adaptive multi-vertex fitter [157] in response to the increase in expected number of interactions per bunch crossing.

The luminous region of interactions at the IP is called the *beamspot* and is assumed to follow a 3D Gaussian distribution with longitudinal length of approximately 40 mm along the beam axis and width of around 8 µm in the transverse plane. Its parameters are reconstructed based on an unbinned maximum likelihood fit to the distribution of PVs sampled over many events. The beamspot fit uses one PV per event, chosen as the PV with the highest sum of square transverse momentum associated to it, $\sum p_{\mathrm{T}}^2$, and with at least five associated tracks. The beamspot is reconstructed from about 10 minute periods of data-taking and several thousand vertices are typically used in the fit.

3.3.2 Electron reconstruction

As briefly mentioned in Section 3.2.2, the ECal cannot distinguish between electromagnetic showers originating from an electron or a photon and information must be combined with tracking information from the ID to make this distinction. As photons have no electric charge, they do not leave a track in the ID unless they covert to an electron-positron pair. This thesis is primarily concerned with electron reconstruction and does not include final states with photons, other than for some minor background processes.

At the start of Run 2, ATLAS used a fixed cluster size, sliding-window reconstruction method for electrons and photons [158] which was later superseded by a dynamic cluster size (supercluster) reconstruction to recover energy loss from EM showering, described in this subsection. An electron is defined as an object consisting of a supercluster and matched track(s) from the ID. A converted photon (a photon which has produced an electron-positron pair) is a cluster matched to a conversion vertex (or vertices), and an unconverted photon is a supercluster that is not matched to any ID tracks. About 20% of photons at low η convert in the ID, and up to about 65 % convert at $|\eta| \approx 2.3$ due to increased material budget in the forward direction making EM showers more probable [159]. The reconstruction procedure of electrons and photons within $|\eta| < 2.5$ is as follows:

1. Clusters of energy deposits measured in topologically connected ECal cells are reconstructed, called *topo-clusters*;

2. The topo-clusters are matched to ID tracks which are refitted to account for bremsstrahlung;

3. Conversion vertices are built and matched to selected topo-clusters;

4. Topo-clusters are merged to form superclusters. The merging is performed differently for electrons and photons and depends on whether matched tracks or conversion vertices are present;

5. The objects are built and their energies are calibrated.

The topo-cluster algorithm [160–162] forms clusters in the ECal only, using thresholds of energy to noise ratios in calorimeter cells. Topo-cluster seeds are formed if the cell energy to noise ratio is above four. The clustering algorithm then iteratively adds neighbouring cells around the seed with successively smaller energy to noise ratios, and the new additions become seeds for the next iteration. The final topo-clusters are required to have total energy greater than 400 MeV to reduce selection of spurious clusters which come from pile-up interactions, as well as $\pi_0 \to \gamma\gamma$ decays. Electron track-matching is then performed by creating RoIs in the calorimeter from topo-clusters. As mentioned in Section 3.3.1, if a track seed in the ID fails the tracking reconstruction criteria but is within the RoI of an ECal topo-cluster, an adapted tracking algorithm is run, allowing for up to 30 % energy loss at each material intersection. The loosely-matched, refitted tracks are then matched to topo-clusters with tighter requirements in η and ϕ. If multiple tracks are matched, tracks with precision hits in the pixel detector and tracks with the smallest ΔR difference are preferred.

Satellite clusters are then added to recover inefficiencies from bremsstrahlung. This is the process of forming *superclusters*. Topo-clusters with energies above 1 GeV and that are matched to a track with at least four silicon hits are selected as supercluster seeds and then all remaining unused clusters are examined for association to the seed cluster. These satellite clusters must have lower p_{T} than the seed cluster and be within a window of $\Delta\eta \times \Delta\phi = 0.075 \times 0.125$ of the seed cluster barycentre. For electrons, this window is then extended to $\Delta\eta \times \Delta\phi = 0.125 \times 0.3$ to search for additional satellite clusters that share matched silicon tracks with the seed cluster and can be attributed to distant bremsstrahlung emission. The direction of the solenoid magnetic field results in

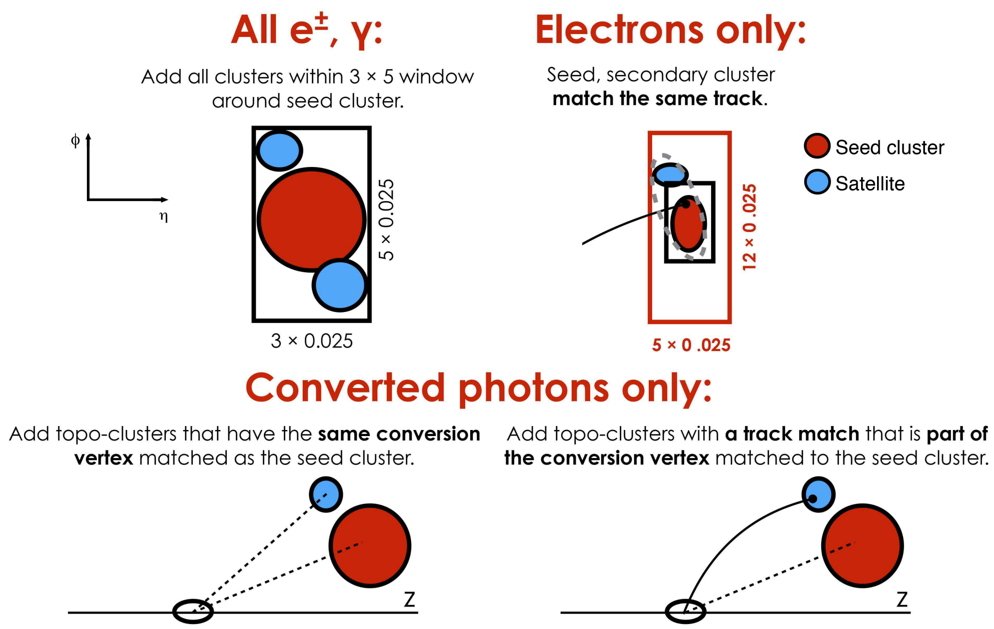

Figure 3.13: Diagram of the superclustering algorithm for electrons and photons. Seed clusters are shown in red, satellite clusters in blue [159].

the separation of the electron and bremsstrahlung photons occurring mainly in the ϕ direction.

The clustering procedure for electrons and photons is shown in Figure 3.13. If a particular object can be easily identified as a photon (a cluster with no good track attached) or as an electron (a cluster with a good track attached and no good photon conversion vertex), then only a photon or an electron object is created for analysis. Otherwise, both an electron and a photon object are created and explicitly marked as ambiguous. A multivariate approach is used to calibrate the reconstructed energy [163], where regression boosted decision trees (BDT) are trained to correct the uncalibrated energies of electron and photon candidates.

Electron identification and isolation criteria

Leptons originating from prompt decays of SM bosons can be discriminated from leptons originating from hadronic decays by measuring the amount of hadronic activity in their vicinity, referred to as *isolation* and defined by the transverse energy measured in a ΔR cone around the lepton, divided by the lepton-p_{T}. This thesis uses electrons which pass medium identification and loose isolation criteria. The electron identification quantities are chosen according to their ability to discriminate prompt isolated electrons from hadronic jets, from converted photons and from non-prompt electrons produced in the decays of heavy-flavour hadrons.

The electron identification is determined from variables describing the EM shower shape, track selections and track matching to superclusters in the ECal and the medium selection has an efficiency of roughly 88 %. The electron isolation is determined from a combination of the tracks of nearby charged particles in the ID and from energy deposits in the calorimeters. The loose criterion requires that the ratio of the energy in the calorimeter and track p_{T} in the ID within a cone of $\Delta R < 0.2$ to the p_{T} of the reconstructed electron must be smaller than 0.20 and 0.15, respectively. The reconstructed electron energy and p_{T} are subtracted from the cone before performing the calculation. The loose isolation requirement has an efficiency of roughly 95 % for electron $E_T = 20\,\mathrm{GeV}$ but this quickly rises to around 99 % for $E_T > 30\,\mathrm{GeV}$. The reconstruction of the electric charge of an electron relies solely on the measurement of the curvature of its associated track in the ID. Bremsstrahlung emission can lead to distortions in the primary electron track and sometimes hits from shower particles being included in the fit of the primary track, possibly resulting in the charge of the electron being misidentified. A BDT is used to suppress electron charge misidentification based on properties of the track and calorimeter information.

3.3.3 Muon reconstruction

Full muon reconstruction and identification relies on information from the ID, MS and calorimeters. Several types of muon object can be reconstructed depending on the subdetector information used in the reconstruction:

- **Combined (CB):** Tracking is performed individually in the ID and MS and then matching is performed between the two. Matched tracks are then refit using inside-out or outside-in fitting. Outside-in reconstruction extrapolates MS tracks to the ID and performs a combined fit on both tracks. The q/p resolution of muon tracks in the MS is higher than that in the ID due to the larger radial extent. CB muons have the highest purity across the largest acceptance range of all muon types. For $|\eta| > 2.5$, MS tracks may be combined with short track segments reconstructed from hits in the pixel and SCT detectors, leading to a subset of CB muons referred to as silicon-associated forward (SiF) muons which allows better use of the ID near the boundaries of its acceptance.

- **Muon Spectrometer extrapolated (ME) / Standalone (SA):** This category is used when no matching ID track can be found for a corresponding track in the MS. In this case, the MS track is loosely extrapolated to the beamline. ME muons are most common in the region of $2.5 < |\eta| < 2.7$ where the MS has acceptance but the ID does not.

- **Segment-tagged (ST):** This reconstruction aims to recover efficiency losses for low-p_T muons which might not reach all layers of the MS or muons which are not fully within the MS acceptance. An ID track with at least one associated segment hit in the MS is used but only ID information is used for the determination of the track parameters.

- **Calo-tagged (CT):** Similar to ST muons but the ID track is matched to calorimeter deposits consistent with an isolated minimally ionising particle. Only ID information is used for the determination of the track parameters. This category is used in the MS gap region of $|\eta| < 0.1$ and the lack of MS hits results in this being the lowest purity category.

A summary of the different muon categories is shown in Figure 3.14. Overlaps between different muon types are resolved before producing the collection of muons used in physics analyses. When two muon types share the same ID track, preference is given to CB muons, then to ST, and finally to CT muons. The overlap with ME muons in the muon system is resolved by analysing the track hit content and selecting the track with better fit quality and larger number of hits [165]. The muon energy loss whilst traversing the calorimeters must be quantified for accurate measurement in the MS. An analytic parametrisation of the average energy loss is derived from a detailed description of the detector geometry and combined with the measured energy deposited in the calorimeter

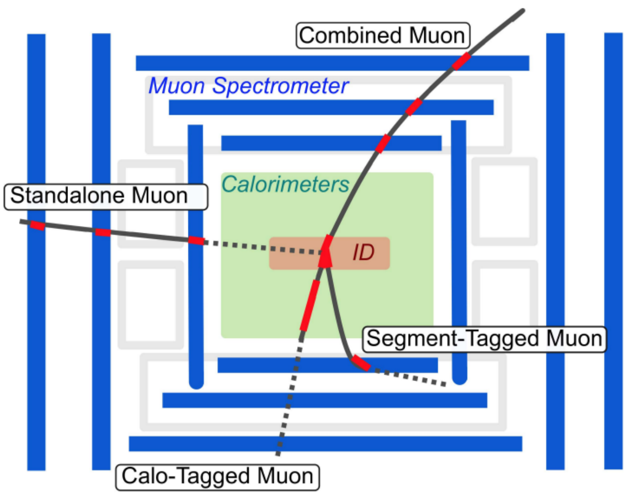

Figure 3.14: The different categories of reconstructed muons in the ATLAS detector [164], with the absence of silicon-associated forward muons. This thesis only uses Combined and Standalone Muons.

to provide the final estimate.

MS tracks use the same perigee representation as in Section 3.3.1 and their reconstruction is performed as follows:

1. Identify short straight-line local track segments reconstructed from hits in an individual MS station. Segments are identified in the individual stations by means of a Hough transform which is a method for performing shape or motion analysis in environments containing noise and missing data [166];
2. Segments in the different stations are combined into preliminary track candidates using a loose pointing constraint to the IP and a first-order approximation to the parabolic trajectory of the muon through the magnetic field;
3. Information from precision measurements in the bending plane is combined with measurements of the second coordinate from the fast-response trigger detectors to create three-dimensional track candidates;
4. A global χ^2 fit of the muon trajectory through the MS is performed, taking into account the effects of possible interactions in the detector material and any misalignment of the MS detector chambers;
5. Hits that are outliers with respect to the fitted track are removed and newly consistent hits are added into the track candidate and the fit is repeated;
6. Ambiguities are resolved by removing lower-quality tracks that share a large fraction of hits with higher-quality tracks;
7. The final set of tracks is then re-fitted with a loose IP constraint and taking into account the energy loss in the calorimeters. The trajectory is then back-extrapolated to the beam line and the p_{T} of the track is expressed at the point of closest approach to the beam axis.

Muon identification and isolation criteria

This thesis uses the medium identification WP for muons, which accepts only CB muons (inside-out and outside-in) for $|\eta| < 2.5$. The CB muons must have at least two precision station measurements, i.e. in the MDTs or CSCs described in Section 3.2.3, with the exception of $|\eta| < 0.1$ where one precision station measurement is allowed. The q/p compatibility is defined as

$$q/p \text{ compatibility} = \frac{|(q/p)_{\mathrm{ID}} - (q/p)_{\mathrm{MS}}|}{\sqrt{\sigma^2(q/p)_{\mathrm{ID}} + \sigma^2(q/p)_{\mathrm{MS}}}}, \tag{3.9}$$

where $(q/p)_{\mathrm{ID/MS}}$ are the ratios of the muon charge to momentum for the ID and MS tracks, expressed at the IP, and σ are the corresponding uncertainties. The q/p compatibility is required to be less than seven to ensure a loose agreement between the ID and MS measurements. Among prompt muons passing the Medium WP in $t\bar{t}$ events, more than 98% are outside-in CB muons. Muon isolation can be measured using ID tracks, calorimeter topological cell clusters, or a combination of the two called *particle flow* [167]. This thesis uses a loose WP for muon isolation with $\Delta R = 0.2$ and selects muons with $p_{\mathrm{T}} > 20\,\mathrm{GeV}$, where overall muon reconstruction efficiencies are higher than 90 % [168].

3.3.4 Other reconstructed objects within ATLAS

In addition to the objects listed above, ATLAS also reconstructs tau leptons, jets and $E_{\mathrm{T}}^{\mathrm{miss}}$.

Tau leptons differ from electrons and muons in that they decay before reaching the detector due to their larger mass. Therefore, taus must be reconstructed via their decay products rather than directly. These decays can either be leptonically to electrons/muons and corresponding neutrinos ($\sim$ 35 % of the time) or hadronically into charged and neutral hadrons and neutrinos ($\sim$ 65 % of the time). Leptonic tau decays are indistinguishable from prompt electrons and muons and it is only hadronic taus that have dedicated reconstruction. The only tau decays relevant to this thesis are leptonic and therefore tau reconstruction shall not be discussed further.

Hadronic jets are formed by QCD showers and are reconstructed from combined ID and calorimeter information. $E_{\mathrm{T}}^{\mathrm{miss}}$ is calculated from all other reconstructed objects by balancing the transverse momentum in the detector which must be conserved to be zero:

$$\mathbf{p}_{\mathrm{T}}^{\mathrm{miss}} = -\sum_{\mathrm{sel.}\ e} \mathbf{p}_{\mathrm{T}}^{\mathbf{e}} - \sum_{\mathrm{sel.}\ \gamma} \mathbf{p}_{\mathrm{T}}^{\gamma} - \sum_{\mathrm{sel.}\ \mu} \mathbf{p}_{\mathrm{T}}^{\mu} \sum_{\mathrm{sel.}\ \tau} \mathbf{p}_{\mathrm{T}}^{\tau} - \sum_{\mathrm{sel.\ jet}} \mathbf{p}_{\mathrm{T}}^{\mathrm{jet}} - \sum_{\mathrm{un.\ track}} \mathbf{p}_{\mathrm{T}}^{\mathrm{track}}, \quad (3.10)$$

where “un. track” stands for “unused tracks”, i.e. track objects that have not been used in any other object reconstruction. Neither jets nor missing transverse energy are considered in this thesis and shall not be discussed further.

Chapter 4

The ATLAS Forward Proton Detector

"The way to a deeper understanding is to move forward with small steps."

– Brian Cox

4.1 Introduction

The ATLAS Forward Proton (AFP) detector [169] is an extension of the ATLAS detector in the very forward region ($|\eta| \approx 10$), with the aim of measuring the momentum and scattering angle of protons originating from diffractive and photon-induced processes at IP1. Reconstructing forward protons in AFP is referred to as forward proton tagging. This thesis is concerned primarily with measuring forward protons associated to photon-induced processes and therefore diffractive physics shall not be discussed further. A summary of diffractive physics at the LHC can be found in Ref. [55].

The CMS and TOTEM collaborations have an analogous Precision Proton Spectrometer (PPS) [170] with an extremely similar physics programme to AFP and the results from the two experiments can be compared. One of the flagship analyses planned for both AFP and PPS in high-luminosity running is measuring exclusive $\gamma\gamma \to WW$ production in association with forward protons. The first observation of the $\gamma\gamma \to WW$ process without forward proton tagging is presented in Chapter 5. The additional information from forward protons can help to further suppress backgrounds whilst probing the high mass tails of distributions that are expected to be sensitive to any new physics introduced by anomalous gauge couplings. PPS recently set limits on non-linear dimension-6 and linear dimension-8 anomalous quartic couplings for diboson masses above 1 TeV using 100 fb^{-1} of 13 TeV data, for hadronic final states [171]. The possibility of performing a fully-leptonic $\gamma\gamma \to WW$ measurement at the High-Luminosity LHC with the addition of AFP forward proton tags will be evaluated in Section 7.6.

The author developed the inter-plane alignment algorithm documented in Section 4.5 and the proton reconstruction working points discussed in Section 4.3. The author also calculated the per-plane cluster reconstruction efficiencies detailed in Section 4.6.

4.2 Detector location and components

At roughly 200 m along the beamline from the ATLAS IP, protons that have undergone scattering but remain intact, such that $E_{\text{proton}} < E_{\text{beam}}$, are sufficiently separated from the nominal beam orbit, due to the magnetic optics, that they can be intercepted by so-called *Roman Pot* (RP) systems that are inserted into the beam pipe aperture. The general locations of the AFP stations are chosen for their acceptance in proton energy loss but the precise locations are determined by constraints on free space along the beamline, away from essential beam elements and instrumentation. The AFP detector consists of four separate RP stations, with two "NEAR" stations located at $z = \pm 205$ m and two "FAR" stations located at $z = \pm 217$ m from ATLAS IP1. Two stations per side separated by a distance in z allows not only the fractional energy loss, ξ, of the proton to be measured with respect to beam but also allows measurement of the proton p_{T} or the four-momentum transfer squared, t, which is related to scattering angle. A schematic of the AFP stations, their location and the ATLAS coordinate system is shown in Figure 4.1. Each of the four stations house four silicon tracking planes, which an aim to reconstruct proton track positions in x relative to the beam axis to an accuracy of less than 20 µm. This aim was far surpassed in practice. Details of the RP system and detector acceptance are given in Section 4.2.1. The SiT modules are described in Section 4.2.2. The two FAR stations also house Cherenkov timing detectors, described in Section 4.2.3, with the aim of measuring the forward proton time-of-flight (ToF) and reconstructing the proton interaction vertex in z.

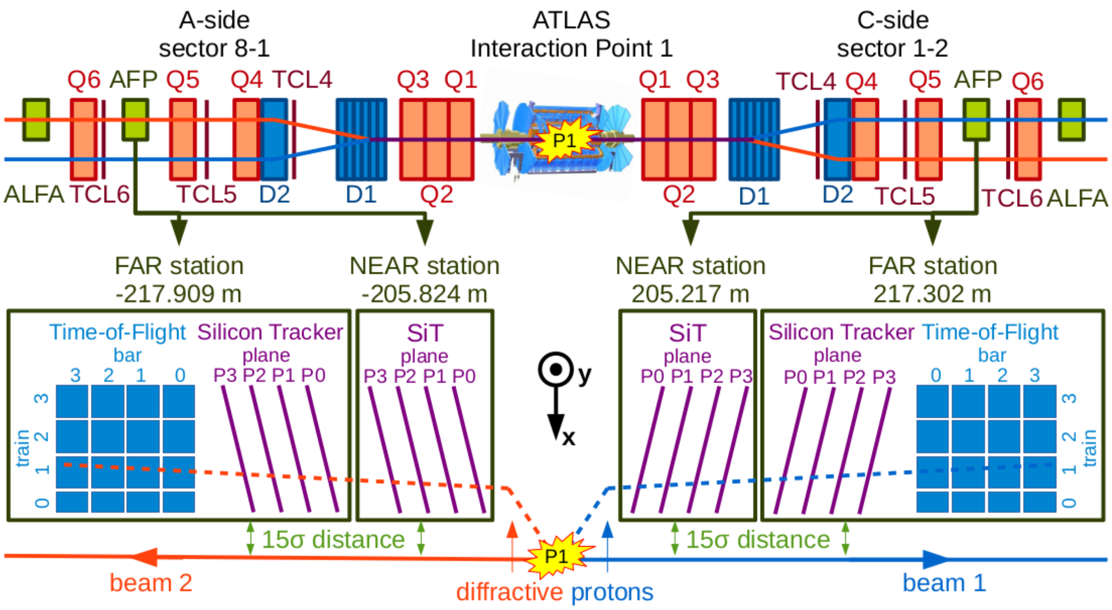

Figure 4.1: A schematic of the ATLAS Forward Proton (AFP) detector stations and their location relative to the ATLAS interaction point (IP1). The NEAR stations are located at roughly $z = \pm 205$ m from IP1 and the FAR stations are roughly $z = \pm 217$ m from IP1. All four stations have 4 planes of 3D silicon pixel trackers (SiT) whilst the outer stations also have additional Cherenkov radiation time-of-flight (ToF) detectors. The aim of all stations is to measure intact protons that have scattered at small angles with respect to the beam and accurately reconstruct their four-vectors [172].

4.2.1 AFP-beam interface and detector acceptance

The AFP stations contain the RP and the mechanics allowing it to enter the beam pipe aperture. The pots are connected to the vacuum chamber of the collider by bellows, which are compressed as the pots are pushed towards the circulating beam. The cylindrical pot's orientation and motion are horizontal, transverse to the beam direction. The PPS stainless steel cylindrical pot design [173], with small modifications, was adopted for the AFP RP design. Due to the thin RP bottom window, the protection of the ultra-high LHC vacuum (10^{-10}–10^{-11} mbar) is a major safety aspect. For this purpose and to minimise the deformation induced on the RP window, the interior of the RP is kept in a secondary vacuum.

The acceptance of each AFP station is highly dependent on beam parameters, the AFP station locations along the beamline and the machine components located between these locations and P1. The size of the beam envelope at the AFP station locations depends on the crossing angle, θ_c, and the β^* parameter, which represents the distance over which the transverse width of the beam halves, as measured at IP1. In standard runs at nominal luminosity in 2017, it was set to $\beta^* = 0.3\,\text{m}$ [174]. The crossing angle in 2017 was $\theta_c = -150\,\mu\text{rad}$ [175]. The size of the beam at AFP is also determined by collimators which scrape the beam to remove beam halo and prevent magnet quenches. There are three collimators on each side between ATLAS P1 and the AFP stations: TCL4, 5 and 6, shown in Figure 4.1. To keep reasonable AFP acceptance, the collimator aperture is set to $\pm 30\sigma$, $\pm 50\sigma$, and $\pm 40\sigma$ for TCL4, 5, and 6 respectively, where σ is the standard deviation of the transverse beam width at each position. The dependence of AFP acceptance on beam optics is discussed in detail in Ref. [176].

A summary of the acceptance in ξ for all AFP stations in given in Table 4.1. The

Station no.	Station name	Sector	Position z [m]	ξ acceptance
0	A-FAR	8–1	+217.9	[0.018, 0.12]
1	A-NEAR	8–1	+205.8	[0.028, 0.115]
2	C-NEAR	1–2	-205.2	[0.026, 0.115]
3	C-FAR	1–2	-217.3	[0.019, 0.12]

Table 4.1: Location and ξ acceptance of each AFP station, in 2017 data-taking in standard runs. The position in z is measured from IP1 in central ATLAS. The acceptance range is where the proton reconstruction efficiency is above roughly 80 %.

total acceptance of all AFP stations in 2017 is quoted as $0.02 < \xi < 0.12$. This is reduced to $0.035 < \xi < 0.08$ when requiring double station proton reconstruction, in the region where station efficiencies are well understood. The AFP acceptance also depends on the global alignment of each RP with respect to the beam, discussed in more detail in Section 4.4.

4.2.2 Silicon tracking

The requirements for the SiT planes include high and non-uniform radiation resistance as the detector is regularly placed within a few millimetres of the beam, subject to an maximum expected 1 MeV equivalent neutron fluence per cm^2 of $3 \times 10^{15}\ n_{\mathrm{eq}}/\mathrm{cm}^2$ at full luminosity [169]. They must also have spatial resolution capable of determining the forward proton position accurately. 3D silicon pixel tracking technology used for the ATLAS IBL tracker [134] was chosen to be suitable, with the additional requirement that any dead edge of the SiT planes was kept to a minimum on the beam side to maximise detector acceptance, as cross-sections for physics processes of interest decrease rapidly with increasing forward proton distance from the beam. The dead edge on the beam side is around 100 µm. Each plane consists of 336×80 pixels with an individual pixel area of $50 \times 250\,\mu\mathrm{m}^2$, resulting in a total active area of $20 \times 16.8\,\mathrm{mm}^2$. The planes have a depth of 230 µm in the beam direction and are separated by a distance of $\Delta z = 9\,\mathrm{mm}$. The reduced charge drift distance for 3D sensors compared to planar silicon sensors reduces the chances of charge trapping due to radiation damage. Beam tests have shown that the 3D sensors accommodate uneven irradiation without degradation [177], which is important in AFP due to the large flux gradient when in close proximity to the beam.

The short-pixel direction is set so that it is the axis along which the dipole magnets steer the protons and is hence the predominant axis along which the protons that have undergone energy loss are deflected. Any divergence of protons along the long-pixel direction arises primarily from non-zero crossing angle at IP1. The resolution in each direction is determined by the pixel size/$\sqrt{12}$, resulting in a per-plane resolution of 14 µm in the short-pixel (x-) direction and 72 µm in the long-pixel (y-) direction. The resolution on the proton position with respect to the beam is further improved by using multiple planes in the reconstruction and by tilting the planes at a 14° angle about the y-axis so as to maximise charge sharing between pixels. This allows interpolation between two neighbouring pixels, improving the resolution in the short-pixel direction to roughly 6 µm. The

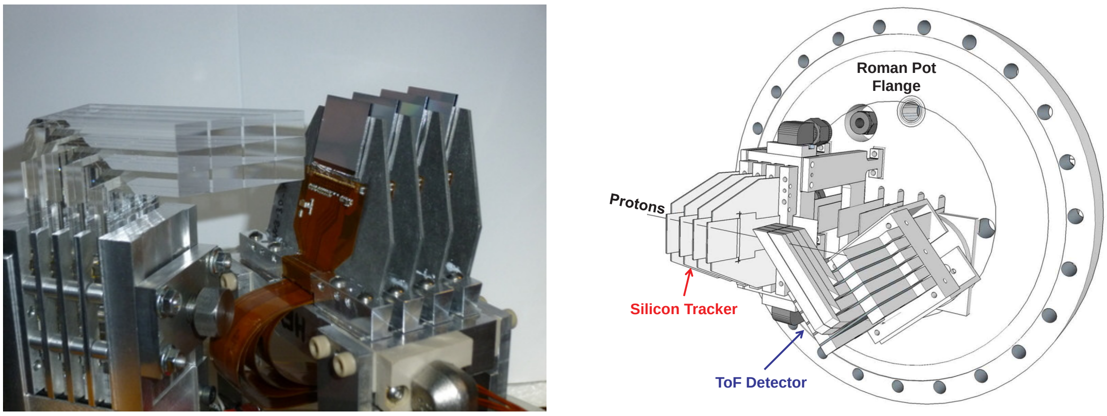

Figure 4.2: On the left is a photograph of the AFP SiT and ToF detectors in the C-FAR station. On the right is a diagram of an AFP-FAR station, showing both detector modules with respect to the forward proton direction and the RP flange [172].

same is not possible in the long-pixel direction, leading to discretised hit locations and a lower resolution of roughly 30 µm. Figure 4.2 shows a photograph (left) and sketch (right) of one AFP-FAR station, with both the SiT and ToF detectors.

4.2.3 Time of flight

The high-resolution Time-of-Flight (ToF) detector is important in high-luminosity data-taking, providing additional rejection against "pileup" protons that originate from single-diffractive processes. Such protons result in combinatorial backgrounds in analyses that require a proton tag in both AFP arms. These backgrounds can be suppressed by measuring the difference in the two protons' time-of-flight to AFP, which allows an interaction vertex to be reconstructed in z under the assumption that both protons originate from the same vertex, and comparing this to the hard-scatter vertex reconstructed in the central detector.

The requirements for the AFP-ToF detector are:

- Timing resolution of $\mathcal{O}(10\,\mathrm{ps})$,
- Acceptance that fully covers the acceptance of the SiT detectors with segmentation for multi-proton timing,
- High efficiency nearing 100 % and high rate capability of $\mathcal{O}(10\,\mathrm{MHz/segment})$,
- L1 trigger capability within ATLAS latency,

- Radiation hard.

The adopted solution to meet these requirements is a Cherenkov-based detector consisting of L-shaped light-guiding quartz bars (LQ-bars) [178], photo-sensitive devices, electronic readout and a reference timing system to allow correlation between the detector stations. The ToF detectors are housed in the AFP-FAR stations, located at $z = \pm 217\,\text{m}$, and are located just behind the SiT planes on the outer side of the station. The L-shape was designed to optimise collection of light whilst adhering to space constraints set by the RP size. Each LQ-bar, shown in Figure 4.3, has two arms glued together at a 90° angle with mirror tape at the bend to direct photons along the light-guide arm. Four such LQ-bars are placed one after the other to form a "train" parallel to the forward proton trajectory and each AFP-ToF detector has four trains, equaling sixteen bars in total. The end of each light-guide arm is attached to a micro-channel plate multi-anode photo-multiplier (MC-PMT) with 16 channels which converts the light signal into an electric pulse for readout. The geometry of the bars is designed so that the optical path in all bars at any given position in the train is equalised and the size of each bar at subsequent train positions decreases with increasing distance from IP1, so as to ensure synchronous arrival of the signal to the MC-PMT for any one proton.

The ToF detector efficiencies in 2017 data-taking were measured in an AFP calibration run with $\mu \sim 1$. The train efficiency was 6–9 % in the A-FAR station

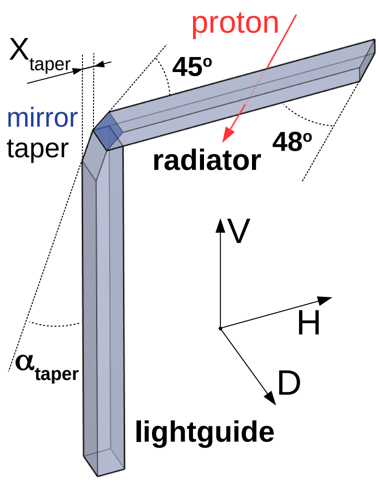

Figure 4.3: Design of the AFP-ToF quartz-Cherenkov LQ-bar. The L-shape with mirror tape was designed to optimise light collection whilst adhering to space constraints of the Roman Pot stations. Each arm of the bar is glued together using UV transparent glue and the end bar is attached to a micro-channel plate multi-anode photo-multiplier [172].

and 3–5 % in the C-FAR station [179], far below the required efficiency. The train efficiencies were found to show a time dependence due to ageing effects while the dependence on pileup was negligible. The low and continuously decreasing efficiencies measured in the single channels, and therefore the trains, were due to low signal amplitudes at the output caused by insufficient high-voltage and exceeded lifetimes of the PMTs (only expected to be 2–4 months based on operating conditions). However, the timing resolution was measured to be excellent at

$$\sigma^{\text{ToF}}_{\text{A-FAR}} = 20.2 \pm 4.0\,\text{ps}, \quad \sigma^{\text{ToF}}_{\text{C-FAR}} = 25.7 \pm 4.7\,\text{ps},$$

corresponding to an expected forward proton vertex resolution in z of

$$\sigma^{\text{ToF}}_{\text{expected}}(\Delta z) = 5.5 \pm 2.7\,\text{mm},$$

when also accounting for smearing due to the unknown y-coordinate of the proton track. The very low efficiencies of the AFP-ToF detectors in 2017 rendered them ineffectual in high-μ analyses. A new design [180] has been developed for LHC Run 3 to improve PMT lifetime and prevent electrical failures. The excellent timing resolution means that the AFP-ToF detectors have promising potential for use in Run 3 physics analyses. They shall not be discussed further until considering future $\gamma\gamma \to WW$ measurements with AFP in Section 7.6.

4.2.4 Triggering

Many physics processes with well-defined signatures in the central detector can rely on the standard ATLAS menu for triggering. However, AFP L1 triggers are needed for processes that produce forward protons when there is either no well defined central signature, central detector backgrounds are large, or where triggers would be heavily prescaled in other menus. The AFP L1 trigger must be very fast due to the signal propagation delay between the AFP station locations and ATLAS P1 and accounting for the fact that the ATLAS L1 trigger latency is around 2.15 µs. The AFP trigger signal is formed by both far SiT and ToF detectors and the decision on which one to use is made by a radio frequency switch before the signal is sent to the ATLAS Central Trigger Processor (CTP). In special low-μ runs, there is a collection of trigger chains based on AFP as part of the minimum bias menu. AFP HLT triggers focus on signatures which would otherwise be heavily prescaled, e.g. for exclusive dijet production [181].

4.3 Forward proton reconstruction

The reconstruction flow in the AFP-SiT detectors progresses from silicon hits to clusters to station tracks to proton objects. Firstly, SiT pixel hits within each plane are clustered using neighbouring functions. Pixel hits are classed as neighbours if they are in the same silicon plane and in adjacent pixel columns in x and within the same pixel row in y, using the coordinates as shown in Figure 4.1. The dimensions of the SiT pixels mean that it is likely for a track to pass through two adjacent pixels in the same column in x (short-pixel direction), but not within two adjacent pixels in the same row in y (long-pixel direction). The local SiT cluster positions are corrected for inter-plane misalignment before performing the track reconstruction, discussed further in Section 4.5. Figure 4.4 shows a comparison of SiT hits to reconstructed clusters in the first plane of the A-FAR station in one run in 2017, corresponding to a few hours of data-taking.

Tracks are reconstructed by finding clusters that are within a user-defined maximum separation in the x-y plane. The default value of this maximum separation is $\sqrt{(\Delta x)^2 + (\Delta y)^2} < 0.5\,\text{mm}$. This distance can be set separately for clusters within the same plane and/or for clusters in different planes. A track candidate is created from all clusters that fall within this allowed distance. If there is one cluster in the track, the track (x, y) position is set to the (x, y) position of the

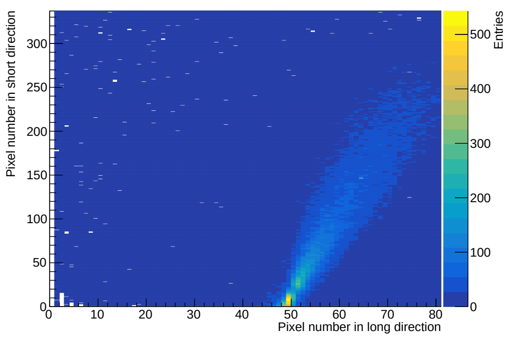

(a) Pixel hits

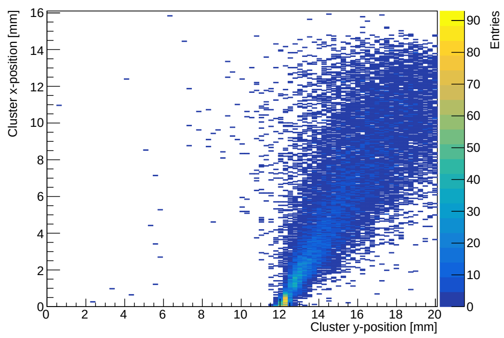

(b) Reconstructed clusters

Figure 4.4: A comparison of (a) pixel hits to (b) reconstructed clusters in plane-0 of the AFP-FAR station in run 336506. (a) is binned according to pixel number within the plane whereas (b) is binned in distance from the SiT plane corner. The nature of the clustering algorithm removes many hits that are due to showering or noise. Coordinates here are ATLAS coordinates, with the possible exception of sign differences. The $x = 0\,\text{mm}$ (ordinate) boundary in (b) represents the active edge of the SiT plane closest to the beam.

cluster (not possible in default reconstruction as clusters are required in at least two planes). For more than one cluster in the track candidate, linear regression is performed to fit the best track to match this set of clusters. The χ^2 test statistic of the track is used as a measure of the goodness of fit and is calculated using the local cluster uncertainties set by the pixel resolution. Any known global misalignment of the AFP station is applied to the calculated track coordinates.

Forward protons within AFP acceptance are reconstructed from SiT tracks that pass certain quality requirements. The quality of the proton objects can be decided by the user from a list of *loose*, *medium* (default), and *tight*:

Loose: Minimum number of clusters in track = 2,

Medium: Loose + Minimum number of planes = 2,

Tight: Medium + No more than one cluster hit per plane.

On each side of AFP, protons can be reconstructed using either tracks in a single station or tracks in both stations. If double-station reconstruction is used, a maximum distance in the transverse plane, $r_{\text{max}} = 2\,\text{mm}$, between the two track candidates is imposed as

$$r = \sqrt{(x_{\text{FAR}} - x_{\text{NEAR}})^2 + (y_{\text{FAR}} - y_{\text{NEAR}})^2} < r_{\text{max}}.$$

If single station reconstruction is allowed and there are reconstructed tracks in both stations but they fail the r_{max} selection, the track in the FAR station is used as the proton candidate.

The goal of the forward proton reconstruction is to convert sets of spacial coordinate measurements $(x, y, z)_{\text{AFP}}$ to outgoing proton four-vectors $(E, p_x, p_y, p_z)_{\text{proton}}$. It is assumed that the protons remain ultra-relativistic with $E/m >> 1$ and are scattered at small angles with $p_{\text{T}}/p_z << 1$. In this case, $E \approx p_z$. The proton ξ is calculated from a parametrised transport function relating the measured position in AFP to the expected ξ value due to the magnetic optics, $x = T(\xi)$ [169, 182, 183]. For example, the transport function used in Ref. [6] is $x(\xi) = -119\xi - 164\xi^2$, determined from a fit to x versus ξ in simulation.

The main sources of systematic uncertainty on ξ reconstruction are the global alignment at low ξ values, discussed in the next section, and uncertainties related to beam optics at higher ξ values. The most important beam optics consideration is the dynamical variation of the crossing angle, which is at the level of 50 µrad over the course of an LHC run. The total systematic uncertainty on ξ is roughly 16 % at the lower edge of the acceptance and 10 % at higher ξ values.

4.4 Global alignment

The accuracy of the reconstructed proton kinematics depends largely on the global alignment of the AFP stations, which must be monitored over time due to the movement of the AFP RPs and the varying beam dynamics. The RPs are restricted in how close they can approach the beam, in order to protect the machine and to not degrade beam quality for the other experiments. The distance of the SiT edge to the beam centre is set to around 12–15σ of the beam transverse width when in the nominal data-taking position, equivalent to around 2 mm. Before stable beams are reached at the start of an LHC fill, AFP is positioned further out of the beam into the "parked" position at roughly 40 mm.

The position of a forward proton track in AFP, $x(s,t)$, depends on AFP station, s, and time, t, and can be broken down into several components:

$$x(s,t) = x_{\text{pre-align}} + x_{\text{tracker}} - x_{\text{beam}}(s) + x_{\text{RP}}(s,t) + \delta x_{\text{corr}}(s), \tag{4.1}$$

where $x_{\text{pre-align}}$ is the measured track position before corrections. The additional parameters are

- Tracker position, x_{tracker}: An estimate of the distance between the edge of the active region of the SiT planes and the outer side of the Roman Pot. It is taken as -0.5 mm, independent of station.
- Beam position, $x_{\text{beam}}(s)$: The offset of the beam centre relative to the beam axis is measured using beam-based alignment (BBA) procedures [184] and cross-checked on a run-by-run basis using beam position monitor (BPM) measurements [185]. In 2017, the central beam position was typically within 1 mm of the beam axis. The beam position throughout the year was found to be stable to better than 100 µm so time-dependence is not considered.
- Roman Pot position, $x_{\text{RP}}(s,t)$: This parameter corresponds to the distance between the beam axis and the edge of the AFP pots when in the data-taking position. Changes in x_{RP} took place at two distinct points during

2017 data taking, corresponding to three insertion settings of $12\sigma + 0.8$ mm, $12\sigma + 0.3$ mm, $11.5\sigma + 0.3$ mm. This equates to distances of 2–3 mm for the FAR stations and 3–4 mm for the NEAR stations. Investigations of x_{RP} with increased time granularity are ongoing.

- In situ corrections, $\delta x_{\mathrm{corr}}(s)$: This is a global correction that accounts for any remaining misalignment seen once all other corrections are applied. The correction is obtained using exclusive dimuon events by comparing the kinematics of the central dimuon system to the kinematics of forward protons measured in AFP. Very tight quality cuts are applied to ensure high correlation between the dimuons measured in ATLAS and the protons measured in AFP. The corrections obtained are roughly $-0.2 < \delta x_{\mathrm{corr}}(s) < -0.4$ mm for all stations. A conservative uncertainty on the global alignment from these corrections is estimated to be 300 µm which partly covers neglected time dependence [6].

4.5 Inter-plane alignment

A smaller effect on reconstructed proton kinematics is the relative alignment of each silicon plane within each AFP station. Each plane is expected to have small offsets and rotations relative to each other of the order of 10–100 µm and a few mrad, respectively. An iterative algorithm was developed by the author to determine the relative alignment of SiT planes within each AFP station and correct for any offsets and rotations, as well as checking the stability of the alignment parameters over time in Run 2 data. The algorithm uses data from standard runs across 2017 with a dedicated AFP "Good Runs List" (GRL) applied to measure residual differences between the pixel cluster and reconstructed track positions. The AFP-GRL ensures the detector was functioning sufficiently to perform proton reconstruction for physics analyses.

The SiT modules use a local coordinate system that differs to the one used by the central ATLAS detector discussed in Section 3.2. The interplane alignment algorithm uses the local SiT coordinate system, where x refers to the long-pixel direction (equivalent to $\pm y$ in central ATLAS coordinates), y refers to the short-pixel direction (equivalent to $\pm x$ in ATLAS coordinates), and $\pm z$ remains along the beam axis. The sign depends on whether the track is reconstructed on the A- or C-side of AFP. A diagram of the SiT planes with local coordinates is shown in Figure 4.5. The three Tait-Bryan angles, α,

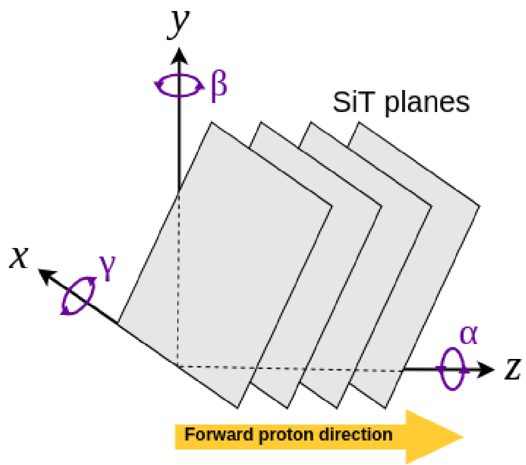

Figure 4.5: A diagram showing the local coordinate system of one AFP-SiT detector. Although not shown here, the origin is typically measured from the corner of the first SiT plane per station. Absolute positions are not important for the inter-plane alignment procedure as it is the difference between measured cluster and reconstructed track positions that is used as input. The local SiT coordinates can be converted to ATLAS coordinates by switching x and y, with sign differences depending on whether the station is on A-side or C-side.

β, and γ, refer to rotations of a given plane about the z-, y-, and x-axis respectively.

The measured position vector, $\mathbf{r_m}$, of a reconstructed cluster within a SiT plane can be expressed as a rotation of the forward proton position vector relative to a perfectly aligned detector, $\mathbf{r}$, plus linear offsets, $\delta\mathbf{r}$. The rotation about all axes is given by

$$\mathbf{R} = \mathbf{R}_z \cdot \mathbf{R}_y \cdot \mathbf{R}_x = \begin{pmatrix} \cos\alpha & -\sin\alpha & 0 \\ \sin\alpha & \cos\alpha & 0 \\ 0 & 0 & 1 \end{pmatrix} \begin{pmatrix} \cos\beta & 0 & \sin\beta \\ 0 & 1 & 0 \\ -\sin\beta & 0 & \cos\beta \end{pmatrix} \begin{pmatrix} 1 & 0 & 0 \\ 0 & \cos\gamma & -\sin\gamma \\ 0 & \sin\gamma & \cos\gamma \end{pmatrix}. \tag{4.2}$$

We can simplify this by assuming very small rotations α, β, $\gamma \ll 1$ and then the measured position of a cluster hit in a SiT plane can be written as

$$\mathbf{r_m} = \mathbf{R} \cdot \mathbf{r} + \delta\mathbf{r} = \begin{pmatrix} 1 & -\alpha & \beta \\ \alpha & 1 & -\gamma \\ -\beta & \gamma & 1 \end{pmatrix} \begin{pmatrix} x \\ y \\ z \end{pmatrix} + \begin{pmatrix} \delta x \\ \delta y \\ \delta z \end{pmatrix}. \tag{4.3}$$

The assumption of small rotations is further validated by applying an iterative procedure to estimate the rotations and offsets, resulting in each successive correction being smaller than the last until a stable solution is found.

As we do not know the true position of the forward proton trajectory, we begin by taking this as the reconstructed track coordinates. The differences between the measured cluster positions and the reconstructed track position in x and y for a given plane are then

$$\begin{aligned}\Delta x &= \mathbf{r_m}[0] - x = -\alpha y + \beta z + \delta x,\\ \Delta y &= \mathbf{r_m}[1] - y = \alpha x - \gamma z + \delta y,\end{aligned} \tag{4.4}$$

where $\mathbf{r_m}[0]$ is the x component and $\mathbf{r_m}[1]$ is the y component of the measured cluster position, respectively. Differentiating these expressions with respect to the coordinate directions gives

$$\frac{\partial \Delta x}{\partial y} = -\alpha, \quad \frac{\partial \Delta x}{\partial z} = \beta, \quad \frac{\partial \Delta y}{\partial x} = \alpha, \quad \frac{\partial \Delta y}{\partial z} = -\gamma, \quad \text{etc.} \tag{4.5}$$

Measurements of cluster-track residuals are calculated for many events over one run of data-taking so that the spread of residuals can be plotted. By plotting the residuals measured in the same plane against the coordinate directions, the rotation of the planes can be found from the gradient of a linear fit to the residuals versus each coordinate direction. There are two possibilities for calculating α. The measurement of Δy versus x is chosen due to the better resolution in the (local) y-direction. The linear offsets are calculated by plotting the residuals against their respective coordinate axis and taking the standard deviation as the offset. For a perfectly aligned detector, this should result in a sharply peaked distribution centred on zero, for the short-pixel direction. Any deviation from zero is taken to be the offset and is corrected.

In reality, one cannot know the true position of the forward proton trajectory: track reconstruction is smeared by the SiT resolution and can be biased by noisy pixels and showering particles. Furthermore, track reconstruction depends on the alignment and hence the reconstructed track position changes as the alignment changes. For this reason, the alignment algorithm is iterative and the process is repeated until the final alignment parameter values are no longer changing across iterations. Detailed steps of the algorithm are given below.

The inter-plane alignment algorithm:

1. Begin by assuming perfect alignment: α, β, γ, δx, δy, $\delta z = 0$.

2. Reconstruct tracks from SiT hits. The maximum allowed distance between clusters in a track is set to the default 0.5 mm, equal to twice the long-pixel

length, and the tracks are forced to have zero slope in z, i.e. removing the linear regression and instead taking the average position of the clusters in x and y as the position of the track.

3. Require that there is only one track per station per event and that this track has no more than one cluster per SiT plane, with a minimum of three clusters per station.

4. Calculate the difference between the measured cluster hit position and the reconstructed track position in x and y, Δx and Δy, for each SiT plane.

5. For linear offsets in x and y, plot a histogram of the residuals and use the mean value as the offset in the respective direction.

6. For rotations about the coordinate axes, fit a linear function to a profile plot of the residuals versus the respective coordinate axis. The gradient of the fit is the correction to the rotation, as shown in Equation (4.5).

7. Update the alignment setting with the new parameters.

8. Iterate the process from (2) until the parameter calculations converge to a steady value, i.e. all corrections are zero. From testing, this was chosen to be 10 iterations.

9. Apply a reduced-χ^2, χ^2_R, selection to reject all tracks with $\chi^2_R > 2$ and repeat step 8.

10. Use the ATLAS Standard Model bootstrap generator [186] to calculate statistical uncertainties on the final parameter values, using 100 replica histograms. In each replica, events are weighted with random numbers generated from a Poisson distribution with mean equal to one.

A simplified view of the alignment procedure in one-dimension is shown in Figure 4.6. In reality, one must correct for three rotations and three offsets per plane. One degree of freedom is removed by assuming that the z-position is known exactly which is appropriate as small offsets in z will not impact track reconstruction. The problem is also simplified by only considering rotations about the z-axis (α), fixing $\beta = \gamma = 0$ and leaving three free parameters per plane. Rotations about the z-axis are the most likely to result in hits being registered in an adjacent pixel and therefore have the largest impact on track reconstruction. All alignment values are calculated relative to the first plane in each station to have a fixed point of reference. As the algorithm is considering

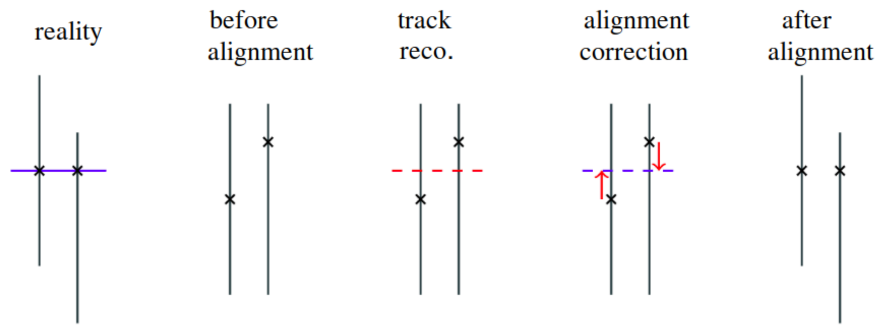

Figure 4.6: A simplified example of the SiT alignment correction in one-dimension. Before alignment, the planes will have an offset (and rotation) that are unknown. By assuming perfect alignment, clusters in the SiT planes result in reconstructed tracks that are not parallel to the beam. Tracks are reconstructed to always be parallel to the beam and are placed at the average position between cluster hits in the x-y plane. The alignment is corrected using the residuals between the cluster positions and reconstructed track position. The track reconstruction changes with the updated alignment and so the process is iterated until the alignment parameters converge to a steady value.

misalignment between SiT planes only, the correction is calculated for each station independently of all other stations.

As described in step (2) of the algorithm, the allowed distance between clusters used in track reconstruction is set to 0.5 mm. This removes tracks at large angles, which are likely to be due to proton showering rather than intact forward protons. Proton showers are also suppressed by requiring exactly one track per station per event and by forcing the reconstructed track to be parallel to the z-axis. This assumption is appropriate as forward proton trajectories will have a very small angle to the beam, as the differential cross-section decreases exponentially with increasing momentum transfer squared. These approximations are put in place to remove the effect of weak modes in the analysis. Weak modes are alignments of the SiT planes which also produce “good” tracks during reconstruction that pass track quality requirements. In these cases, one cannot distinguish which possible set of tracks is the true set, leading to problems with convergence of the algorithm. Examples of weak modes can be seen in Figure 4.7.

The decision to apply a χ^2 selection as a measure of goodness of fit to tracks after ten iterations resulted from consistency tests of the alignment procedure, using orthogonal data subsets created by rejecting either even or odd numbered events within one run and comparing the results. It was found that there were differences in the measured alignment parameters depending on which set of tracks were used. This was due to a few tracks with very large χ^2 values even after several iterations of the algorithm, demonstrated in Figure 4.8. The presence of high-χ^2 tracks possibly indicates that the forward proton trajectory is at a larger angle than the reconstructed track, which is forced to have zero angle with respect to the beam. Such tracks are likely to be due to showers rather than forward protons and can largely bias the alignment results. Another cause of tracks with large χ^2 values will be noisy pixels biasing cluster reconstruction. A reduced χ^2 selection, $\chi^2_R = \chi^2/\text{NDF}$ where NDF is the number of degrees of freedom per track, was applied to tracks after the tenth iteration and then the alignment procedure was repeated. Applying this selection removed $\sim 0.1\,\%$ of tracks but resulted in all profile gradients being consistent with zero after a further ten iterations of testing. The initial ten iterations allows an alignment estimate to be found before applying the χ^2_R selection, otherwise only a few tracks pass the selection in the first iteration. This is because perfect alignment is assumed, resulting in only "bad" tracks that match this (generally) incorrect alignment being used in the alignment analysis.

4.5.1 Parameter corrections

Figure 4.9 shows an example of the results obtained in the alignment procedure, before (blue) and after (red) the alignment algorithm has been run for plane-0 of the A-NEAR station in one run of data-taking. Corrections for all other planes are calculated relative to the alignment of plane-0 in each station. Note that the plotting coordinates have been converted to the central ATLAS coordinate system so now x represents the short pixel direction and y the long pixel direction. Figure 4.9a shows the distribution of residuals in x, Figure 4.9a in y, and Figure 4.9c shows the residuals in x plotted against y, from which the correction to the rotation angle α is calculated from the gradient.

The mean and standard deviation presented in the legend of Figure 4.9a are returned from a Gaussian fit to both histograms. The values before alignment corrections show that there is an average difference between the measured cluster hit position and the reconstructed track position of (15.21 ± 0.04) µm, indicating an offset of the plane in the x-direction. After correcting the alignment over 20 iterations, it can be seen that the average residual between the cluster and reconstructed position is (-0.03 ± 0.04) µm, which is consistent with zero. In comparison to Figure 4.9a, the distributions in the y-direction form a series of sharp peaks due to the discrete nature of hits in the long-pixel direction,

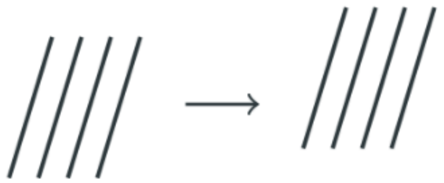

(a) A global offset of all planes within a station

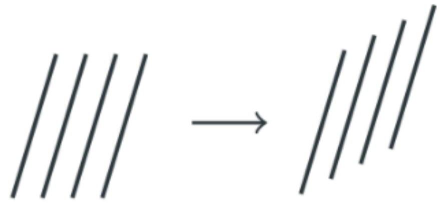

(b) A linearly increasing offset of planes

Figure 4.7: Two examples of weak alignment modes that can lead to problems with parameter convergence. It is impossible for the alignment algorithm to decide which is the best estimate of the alignment out of the examples shown in this diagram as they all lead to tracks which would be considered "good". The example shown in (a) can be eliminated by fixing the position of one plane. The example shown in (b) can be eliminated by requiring the all tracks are reconstructed parallel to the beam.

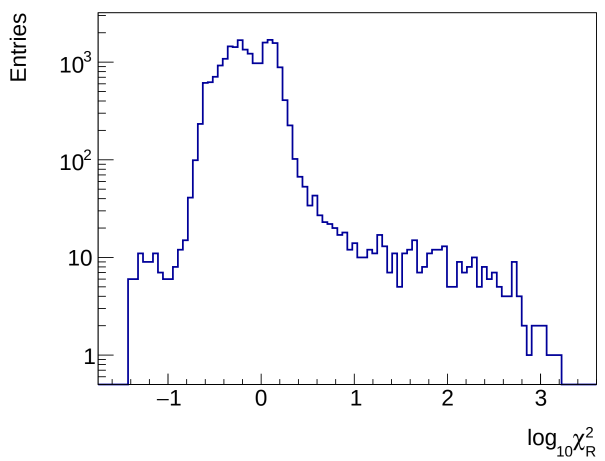

Figure 4.8: The reduced-χ^2 value of track fits in the A-FAR station after 10 iterations of the alignment algorithm, where χ^2 is a test statistic used as a measure of goodness of fit. It was chosen to apply a $\chi^2_R < 2$ selection, $\log_{10}\chi^2_R < 0.3$, to the tracks to remove the tail of the distribution thought to be causing inconsistencies in the alignment results.

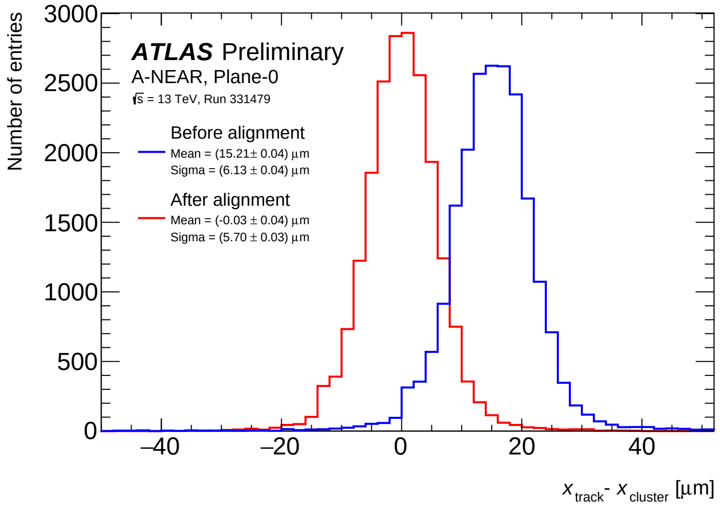

(a)

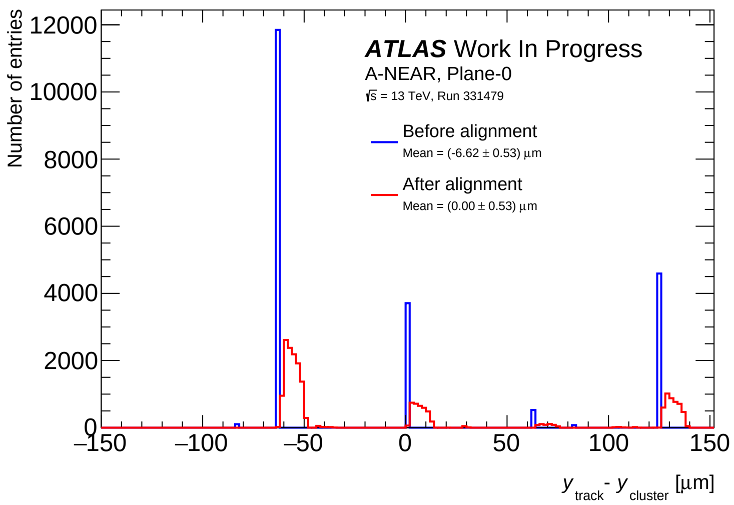

(b)

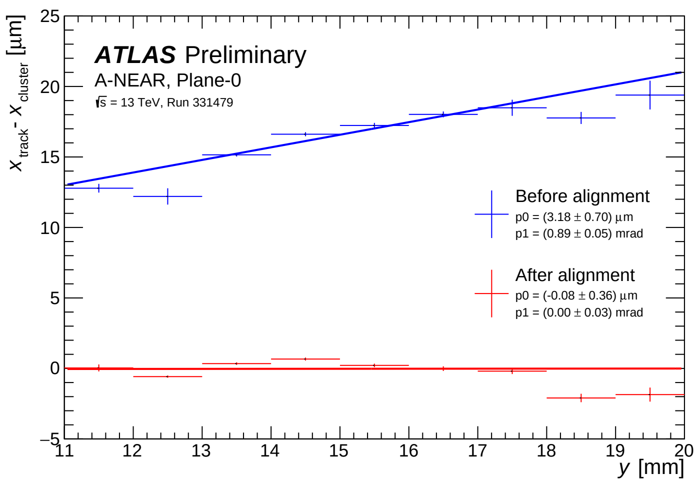

(c)

Figure 4.9: Corrections to alignment parameters calculated from data for (a) the offset in x, (b), the offset in y, and (c) the rotation about z. The coordinate system used is the ATLAS coordinate system. Results are shown for the first plane (plane-0) of the A-NEAR station, before (blue) and after (red) the alignment algorithm has been applied. Uncertainties given are statistical only.

demonstrated in Figure 4.9b. This is not seen in the x-distribution due to both the shorter pixel length in that direction and the 14° tilt about the y-axis, which is designed to maximise charge sharing in x, allowing interpolation of the exact hit location. Despite the discrete peaks in y, the requirement of the algorithm remains that the mean value of the distribution should be zero. Before the alignment, the mean difference in measured cluster hit to reconstructed track position in y is (-6.62 ± 0.53) µm and after the alignment correction it is (0.00 ± 0.53) µm.

The mean residual Δx varies with y before the rotation of the plane is corrected, shown in Figure 4.9c, representing an increasing divergence of the measured hit position and reconstructed track position as the y-position increases. This suggests that the plane is rotated about the z-axis, by an angle of (0.89 ± 0.05) mrad. After the alignment has been corrected, the gradient of the fit has been reduced to (-0.00 ± 0.03) mrad, consistent with zero. Uncertainties here are statistical only. There seems to be a slight oscillatory pattern in Figure 4.9c which is potentially due to rotations about x and y not being considered and is a source of systematic uncertainty. However, the uncertainties on the interplane alignment are assumed to be negligible compared to the uncertainty on the global alignment (300 µm).

Similar calculations are carried out for all four planes of all four stations. Fit ranges are adjusted to ranges determined by the position of forward proton hits. For example, the forward proton hits in the C-side stations have a typical y-range between 0–10 mm, whereas this is roughly between 10–20 mm on A-side.

4.5.2 Alignment parameter evolution across iterations

By calculating the correction to alignment parameters at each iteration, the evolution of each offset and rotation can be monitored to check for convergence. This is shown in Figure 4.10 for α, Δx and Δy in one run of data-taking for the A-NEAR station (again converted to central ATLAS coordinates). The parameters begin at zero for all planes as the algorithm begins by assuming perfect alignment. These values quickly change as the iterations increase but converge quickly to steady values. A small change is seen around iteration 11 as this is when the χ^2_R selection is applied, removing anomalous tracks from the analysis that were biasing results. Occasionally, parameter oscillation around a steady value was observed in the convergence plots and therefore a damping

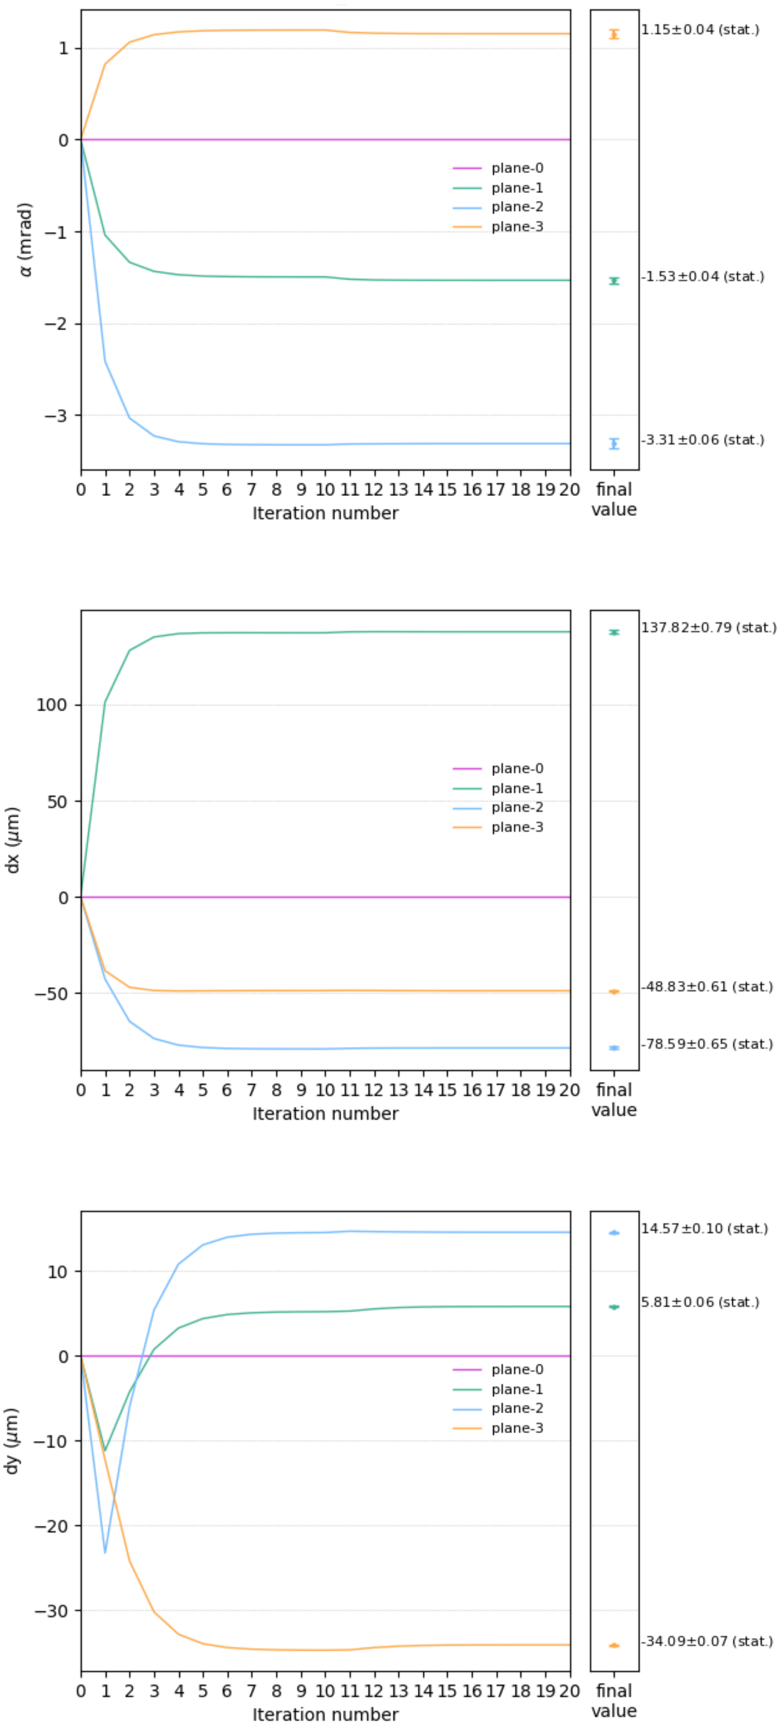

Figure 4.10: The evolution of the (top) α, (middle) Δx and (bottom) Δy alignment parameters for the A-NEAR station in run 337052.

parameter of 0.7 $\times$ correction is applied at each iteration in an attempt to force convergence. The uncertainty on the final value is calculated from a bootstrap analysis in the final iteration.

Parameter convergence plots for all stations using this run can be found in Appendix A. It should be noted that the alignment algorithm had to be adapted for the A-FAR station due to only two working planes in this station for the majority of 2017. It can be seen that the 10+10 iteration method is unnecessary for most parameters and the alignment algorithm could be improved in the future by automating the application of the χ^2-selection when the change in parameter value between iterations is smaller than a chosen value to save on computational time.

4.5.3 Stability of alignment over time

Almost all standard runs in which AFP was used in nominal data-taking were used to evaluate the stability of the alignment over time. There was no intervention to the SiT detectors between the full AFP installation at the start of 2017 and the year-end technical stop (YETS) so it was predicted that the alignment should remain fairly stable across these runs, with possible exceptions due to detector heating and other mechanical stresses.

Figure 4.11 shows how the Δx parameter changed over this period for the A-NEAR station, again in ATLAS coordinates. Each marker represents the final parameter values after 20 iterations of the algorithm for each run with associated statistical uncertainties calculated in a bootstrap analysis. The final values shown on the right-hand side of the plot are the weighted average of the runs, with the uncertainty given by the weighted standard deviation. Similar plots for all alignment parameters and all stations can be found in Appendix A.

Figure 4.11 is an example of reliable alignment, where the parameter remains fairly steady within the SiT resolutions across all of 2017. If parameter stability is not observed, this generally highlights an issue with a particular plane, especially if instability is seen for all parameters of the plane in question. Therefore, the alignment analysis can be used to cross-check the proper functioning of SiT modules across data-taking periods. Although the AFP GRL rejects lumiblocks with station or global issues, it can be worthwhile to also note individual problematic SiT planes so that these planes can be avoided in proton

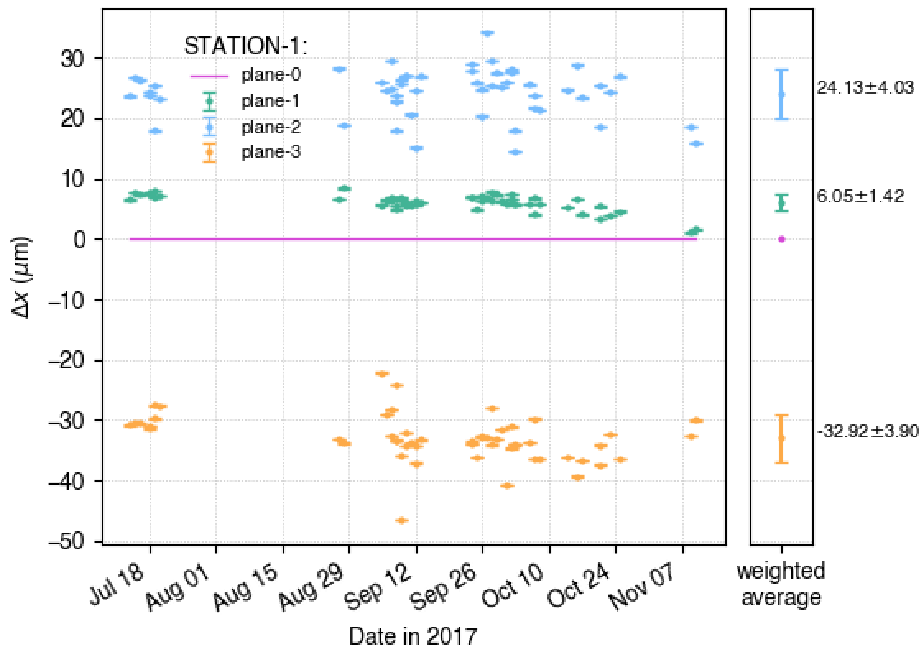

Figure 4.11: The stability of the Δx alignment parameter over standard runs with AFP inserted in 2017, with the AFP GRL applied. Values and uncertainties in the central plot are the final Δx value and bootstrap uncertainty calculated after 20 iterations of the algorithm, and the results on the right-hand plot show the weighted average across these runs, with the uncertainty given by the weighted standard deviation. Uncertainties are statistical only.

reconstruction. The alignment analysis highlighted that plane-1 of the C-FAR station was not behaving as expected in 2017 and it was discovered that this plane had no high-voltage across the runs investigated and it was subsequently removed from the alignment analysis. It is advised not to use this plane for proton reconstruction in 2017 data.

The weighted average alignment values for all parameters are input to a general AFP reconstruction toolbox and automatically applied in any analysis using AFP data from 2017. The alignment parameters for the first plane of each station are set to zero, as well as for the β and γ rotations. Any planes that are highlighted to be problematic also have their alignment parameters set to zero and are advised to not be used in proton reconstruction.

4.6 Proton reconstruction efficiencies

In order to perform a cross-section measurement with AFP, proton reconstruction efficiencies must be fully understood and quantified. In this context, a station efficiency refers to the probability that a proton object is reconstructed in AFP when a true proton is within the AFP acceptance. The nominal way to calculate the station efficiencies is to "tag-and-probe" the proton tracks within each station. The near-station efficiency is estimated by first selecting events with exactly one track in the far (tag) station in the acceptance common to both stations. The efficiency is the fraction of these events that also have one or more tracks in the near (probe) station satisfying $r < r_{\mathrm{max}} = 2\,\mathrm{mm}$. The tag and probe stations are inverted to measure the far-station efficiency.

The author performed a cross-check of the nominal station tag-and-probe calculation by calculating the efficiency of each SiT plane per station to reconstruct a cluster hit if a proton was reconstructed in the station. In order to calculate the station efficiencies independently in this method, the default proton reconstruction was altered slightly to allow reconstruction in a single station only. Only tracks associated to reconstructed proton objects in data were used, as these already have various data quality requirements applied. The efficiency of each SiT plane is denoted by ε_i, with $i = \{0, 1, 2, 3\}$, and was calculated using a tag-and-probe method where three planes are used as the tag and then the fourth plane is probed to calculate its efficiency. For example, if we want to calculate the efficiency of plane-2 within a station, ε_2, and we know that the probability of seeing hits in all four planes is simply the product of all plane efficiencies, $\varepsilon_0\varepsilon_1\varepsilon_2\varepsilon_3$, and that the probability of seeing hits in all other planes with no requirement on plane-2 is $\varepsilon_0\varepsilon_1\varepsilon_3$, then we can calculate ε_2 by taking the ratio of the number of tracks with hits in all planes to the number of tracks which have hits in plane-0, 1, and 3:

$$\varepsilon_2 = \frac{N_{\mathrm{track}}(\text{hits in planes 0, 1, 2, 3})}{N_{\mathrm{track}}(\text{hits in planes 0, 1, 3})}. \tag{4.6}$$

Inherently, this method only selects protons with hits in a minimum of three planes within a given station. Although this differs from the nominal proton reconstruction, it does not have any impact on the calculation of the reconstruction efficiency of each plane. This tag-and-probe method results in efficiencies which are independent of the performance of other planes.

Figure 4.12 shows the calculated plane efficiencies for all stations calculated over several runs across 2017, with the exception of A-FAR which only had two working planes and therefore the per-plane tag-and-probe method could not be used as any protons reconstructed in this station will have hits in both planes, by design. A general degradation of efficiency over time can be seen for all planes, which was expected due to the large radiation doses that the AFP detectors receive during data-taking. Some of this loss can be recovered by applying a SiT tuning to the planes which involves masking noisy pixels as well as changing amplification thresholds to account for detector degradation. Some tuning values are set globally for the whole sensor but some for each individual pixel channel. A tuning was employed before run number 336505 (represented by the dashed line in Figure 4.12) and a resulting gain in all plane efficiencies can be seen. The tuning is applied within the hardware and not in any post-processing so it is difficult to conclude whether the jump in efficiencies seen is due to the tuning alone.

The efficiencies on C-side are generally lower than those on A-side and typically fall faster with time. This was also seen in the station tag-and-probe proton reconstruction efficiencies. More background has been observed in AFP C-side during special low-μ running with the same beam parameters but we would not expect this to translate to any observable difference in standard high-μ runs. More pixels were masked on C-side, suggesting more radiation damage to the SiT planes on this side. This can be explained by the fact that the C-side stations were installed in 2016, one year before A-side, and the same SiT modules were reused. Another observation is that the efficiencies of plane-0 and plane-3 are generally lower than plane-1 and plane-2, for all stations. This can be explained for plane-0, which is the plane closest to IP1 for each station, due to tracks with small angles that are right at the lower edge of the station acceptance in ξ, resulting in missing hits in the first plane. The lower efficiencies in the outermost plane, plane-3, could be due to the increased probability of proton showers as the proton traverses more planes, causing higher occupancy of hits and reducing reconstruction efficiencies. However, this is unlikely to cause a large difference in plane efficiency relative to plane-2.

The proton reconstruction efficiency within a station can then be calculated as a sum of all possible ways in which a proton could be reconstructed in the default reconstruction, i.e. if there are hits in all four planes, three out of four

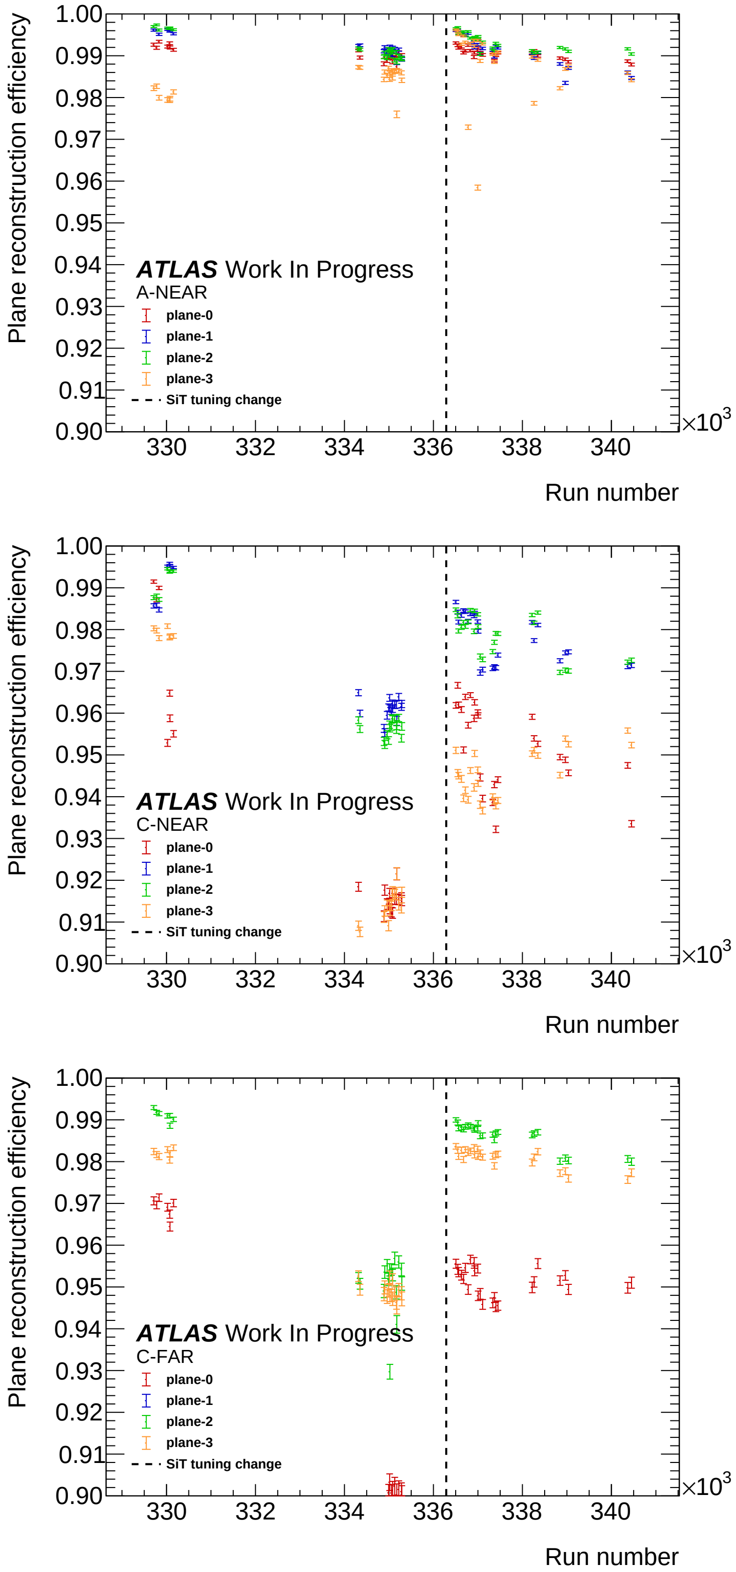

Figure 4.12: The cluster reconstruction efficiencies of each SiT plane in the (top) A-NEAR, (middle) C-NEAR, (bottom) C-FAR station across 2017. A general downward trend can be seen for all planes but some loss in efficiency is recovered by a SiT tuning change before run 336506, represented by the dashed line. The efficiency of plane-1 in C-FAR is much lower at around 30 % due to a lack of high-voltage in this plane.

planes, or two planes. The station efficiency, ε_S, can be written as

$$\begin{aligned}\varepsilon_S = {} & \varepsilon_0\varepsilon_1\varepsilon_2\varepsilon_3 \text{ (4 planes)} \\ & + \varepsilon_0\varepsilon_1\varepsilon_2(1-\varepsilon_3) + \varepsilon_0\varepsilon_1(1-\varepsilon_2)\varepsilon_3 + \ldots \text{ (3 planes)} \\ & + \varepsilon_0\varepsilon_1(1-\varepsilon_2)(1-\varepsilon_3) + \varepsilon_0(1-\varepsilon_1)\varepsilon_2(1-\varepsilon_3) + \ldots \text{ (2 planes).}\end{aligned} \tag{4.7}$$

With each ε_i generally above 90 %, the resulting station efficiencies are very high.

Calculating the station efficiencies from Equation (4.7) results in NEAR station efficiencies of over 99.9 % which is higher than calculated from the direct station tag-and-probe NEAR efficiencies but these are also very high at around 99 %. A larger difference is seen between the two methods for the C-FAR station, which was 99.7 % from Equation (4.7) but roughly 96 % from the direct station tag-and-probe. This difference can be explained by proton showering between the NEAR and FAR stations. If a proton showers between the NEAR and FAR station, this will be measured as an inefficiency in the direct station tag-and-probe whereas the station efficiencies calculated from the combination of SiT plane efficiencies is not affected by proton showering between stations. Although proton showers can still cause tracks to be reconstructed in the FAR stations, proton quality requirements (mainly the requirements on minimum cluster distance in the x-y plane) are designed to remove these tracks. The only published analysis to date to utilise AFP in standard runs is the observation and measurement of forward proton scattering in association with lepton pairs produced via the photon fusion mechanism at ATLAS [6]. This analysis used the default proton reconstruction where a proton track must be reconstructed in both the NEAR and FAR station on any one side of AFP when calculating cross-sections and therefore it makes sense to use the direct tag-and-probe station efficiencies that also account for proton showering between the stations. Single station reconstruction was allowed for the observation only. Combining the NEAR and FAR station efficiencies gave an overall proton reconstruction efficiency of 0.92 ± 0.02 for this analysis. Most of the uncertainty on the method is due the dependence of the reconstruction efficiency on ξ. Although there were differences between A- and C-side found in the plane efficiency method, this is counteracted by the lower A-FAR efficiency due to only two working planes in this station. Another benefit of the direct station tag-and-probe is that it can measure the A-FAR efficiency which the plane efficiency method cannot.

Chapter 5

Observation of the $\gamma\gamma \rightarrow WW$ process

"No experimental scientist really wants observation to agree with theory after all :-)"

– Brian Cox

5.1 Introduction and motivation

This chapter details the first observation of the electroweak process $\gamma\gamma \to WW$ in the fully-leptonic decay mode [4]. Evidence of this process had previously been seen in the analysis of Run 1 data by ATLAS [187] at centre-of-mass (CoM) energy of $\sqrt{s} = 8\,\text{TeV}$ and by CMS [188, 189] in a combined 7 TeV and 8 TeV dataset. The Run 2 dataset is the largest integrated luminosity recorded to date at the LHC, allowing several previously unobserved processes to be probed with increased statistical precision. The higher CoM energy of $\sqrt{s} = 13\,\text{TeV}$ in Run 2 also has the advantage of increasing the cross-sections of processes of interest compared to previous data-taking. As the photon PDF is generally much smaller than quark and gluon PDFs and $\gamma\gamma \to WW$ is a photon-initiated process with lowest order α^2_{QED}, it has a very small cross-section relative to many quark-induced processes at the LHC and benefits hugely from this increased dataset. The dominant background to the $\gamma\gamma \to WW$ process is quark-induced WW production, $q\bar{q} \to WW$, which is indistinguishable from the signal other than for the presence of an underlying event.

Within this analysis, the author developed many data-driven pileup corrections in simulation, discussed in Section 5.4, performed studies into discrepancies between different generators in simulating the underlying event (Section 5.5.2), and calculated the acceptances and correction factors required for the cross-section measurement (Section 5.8.4). The author also applied the data-driven exclusive efficiency calculation that she developed, described in Section 5.4.1, to an exclusive dilepton analysis in association with forward protons reconstructed in the AFP detector [6].

5.2 Data and Simulated Event Samples

The analysis uses proton-proton collision data recorded with the ATLAS detector during the Run 2 data taking period (2015–2018) at $\sqrt{s} = 13\,\text{TeV}$. The data correspond to an integrated luminosity of $\mathcal{L} = 139\,\text{fb}^{-1}$ with an uncertainty of $\pm 1.7\,\%$. The data have been cleaned from events with non-functional detector components [190]. The number of interactions per bunch crossing, μ, also referred to as *pileup*, varied roughly between $10 < \mu < 60$ with an average value of $\langle\mu\rangle = 33.7$ [174]. The average longitudinal width of the beamspot, σ_{BS}, was 35 mm.

Signal and background processes were modelled using MC event generators

Process	Matrix element (alternative)	PDF set	UEPS (alternative model)	UE tune (alternative model)	Prediction order for total cross-section
$\gamma\gamma \to WW$ (EE)	BudnevQED	n/a	Herwig7	n/a	LO
	(MG5_aMC@NLO)		(Pythia8)	n/a	
$\gamma\gamma \to WW$ (SD)	MG5_aMC@NLO	CT14qed	Pythia8	n/a	
$\gamma\gamma \to WW$ (DD)	MG5_aMC@NLO	CT14qed	Pythia8	n/a	
$\gamma\gamma \to ll$ (EE)	BudnevQED	n/a	Herwig7	n/a	LO
	(MG5_aMC@NLO)		(Pythia8)	n/a	
$\gamma\gamma \to ll$ (SD)	Lpair	n/a	Lpair	n/a	LO
	(MG5_aMC@NLO)	(CT14qed)	(Pythia8)	n/a	
$\gamma\gamma \to ll$ (DD)	Pythia8	NNPDF23QED	(Pythia8)	n/a	LO
	(MG5_aMC@NLO)	(CT14qed)	(Pythia8)	n/a	
$qq, qg \to WW$	Powheg-Box v2	NNPDF3.0 NLO	(Pythia8)	AZNLO	NNLO QCD
	(Sherpa v2.2.2)	(NNPDF3.0 NNLO)	(Sherpa v2.2.2)	(Sherpa v2.2.2)	
$gg \to WW$	Sherpa v2.1.1	CT10	Sherpa v2.1.1	Sherpa v2.1.1	NLO
$gg \to H$	Powheg-Box v2	NNPDF3.0 NLO	Pythia8	AZNLO	NNLOPS
WW VBS	Sherpa v2.2.2	NNPDF3.0 NNLO	Sherpa v2.2.2	Sherpa v2.2.2	LO
$WZ/W\gamma^\star/ZZ/Z\gamma^\star$	Sherpa v2.2.2	NNPDF3.0 NNLO	Sherpa v2.2.2	Sherpa v2.2.2	NLO
$W\gamma/Z\gamma$	Sherpa v2.2.2	NNPDF3.0 NNLO	Sherpa v2.2.2	Sherpa v2.2.2	NLO
$t\bar{t}$	Powheg-Box v2	NNPDF3.0 NLO	Pythia8	A14 tune	NNLO+NNLL
Wt	Powheg-Box v2	NNPDF3.0 NLO	Pythia8	A14 tune	NLO
W	Sherpa v2.2.1	NNPDF3.0 NNLO	Sherpa v2.2.1	Sherpa v2.2.1	NNLO
$Z/\gamma^\star$	Powheg-Box v1	NNPDF3.0 NLO	Pythia8	AZNLO	NNLO
	(Sherpa v2.2.1)		Sherpa v2.2.1		

Table 5.1: Overview of the simulation tools used to generate signal and background processes, and to model the underlying event and parton shower (UEPS). EE, SD, and DD refer to elastic, single-dissociative and double-dissociative contributions to the process, respectively. VBS stands for Vector Boson Scattering. Where a PDF set is listed as "n/a", this indicates that a photon flux or similar was used in place of a PDF set. Alternative event generators and configurations used to estimate systematic uncertainties are shown in parentheses.

to study kinematic distributions, evaluate background contamination in the signal region and to interpret the results. Inside ATLAS, separate MC samples are generated for the different data-taking years to match the different data-taking conditions, referred to as MC campaigns. A list of the MC generators used in this thesis were detailed in Sections 2.5.3 and 2.5.4. Table 5.1 summarises the samples used in the $\gamma\gamma \to WW$ observation. All standard ATLAS corrections are applied to improve the general data/MC agreement. These scale factors and reweightings are derived by dedicated combined performance groups to correct for known inefficiencies in the detector and differences in data-taking conditions that were unknown before MC production. The corrections applied are the trigger, reconstruction, identification and isolation scale factors for the leptons and a pileup reweighting (PRW), which is required as the beamspot profile varies with time over data-taking and cannot be known a priori. The standard ATLAS PRW does not provide sufficient data/MC agreement for the required precision, resulting in many additional pileup corrections being derived within this analysis. It will be seen in the proceeding sections that any distribution of charged particles arising from non-perturbative, or *soft*, interactions is very difficult to model accurately and a large part of the analysis procedure is attempting to mitigate these shortcomings, in both pileup interactions and the underlying event for quark-induced backgrounds.

5.3 Event Selection

The $\gamma\gamma \to WW$ production modes were discussed in Section 2.4. This chapter considers the leptonic decays of the W-bosons into an opposite-sign, different-flavour final state as signal, with the exception of $W \to \tau\nu$ decays[10]. The $W \to e\nu$ and $W \to \mu\nu$ branching fractions are both roughly 11 % [12]. An isolated $e^{\pm}\mu^{\mp}$ final state can occur experimentally for processes with two W-bosons, two tau-leptons, or for other final states in which there are either missed leptons, leptons reconstructed with the wrong flavour/sign, or other physics objects wrongly reconstructed as leptons. For example a jet could be wrongly reconstructed as an electron shower in the ECal. The latter two mis-reconstructions are generally referred to as *fakes*. The opposite-sign $\gamma\gamma \to WW$ selection is a physical requirement as the initial diphoton system has zero charge, which must be conserved. The opposite-flavour final state is an experimental choice to reduce backgrounds that produce a same-flavour final state like neutral current Drell-Yan (DY) production, $q\bar{q} \to Z/\gamma^* \to \ell^+\ell^-$, and photon-induced dilepton production, $\gamma\gamma \to \ell^+\ell^-$. Pomeron-induced backgrounds are heavily suppressed due to much smaller proton survival factors in these processes.

The two neutrinos in the final state are not reconstructed by the ATLAS detector but do carry away momentum, resulting in the two leptons not being produced exactly back-to-back in the detector. This allows selections on dilepton p_{T}, $p_{\mathrm{T}}^{\ell\ell}$, to separate the signal from dilepton final states without neutrinos, e.g. DY and $\gamma\gamma \to ll$.

The $\gamma\gamma \to WW$ process can be split into three distinct contributions: elastic (EE) events in which both protons remain intact, single-dissociative (SD) events in which one proton dissociates, and double-dissociative (DD) events in which both protons dissociate. Considering these production modes, the full process can be written as

$$pp(\gamma\gamma) \to W^+W^-p^{(*)}p^{(*)} \to e^{\pm}\nu_e\mu^{\mp}\nu_\mu p^{(*)}p^{(*)},$$

[10]When the tau decays leptonically, it is indistinguishable from a $W \to \mu\nu$ or $W \to e\nu$ decay.

where p^* represents a proton that has dissociated. For brevity, it shall mostly be referred to as $\gamma\gamma \to WW$. Although the dissociative contributions are produced with additional hadronic activity from the proton remnant hadronisation, the decay products are typically outside of the ATLAS inner detector acceptance and outside of the fiducial region of the analysis. For this reason, it is difficult to distinguish the elastic from dissociative production modes experimentally and therefore all three are considered as signal.

One of the most important event selections, the exclusivity requirement, shall be discussed in detail in Section 5.4.1. The only sizeable background that cannot be suppressed by the exclusivity requirement is exclusive dilepton production which occurs via the same production mechanisms as $\gamma\gamma \to WW$ production but with a larger cross-section due to the mass differences between the dilepton and diboson systems. The $\gamma\gamma \to e^+e^-/\mu^+\mu^-$ processes are hugely suppressed by the different-flavour signal selection. The $\gamma\gamma \to \tau^+\tau^-$ process forms a background when the two tau-leptons decay into an $e^{\pm}\nu_e\nu_\tau\mu^{\mp}\nu_\mu\nu_\tau$ final state. However, as the τ is much lighter than the W, the leptons are expected to be produced with large angular separation and this background can be suppressed by requiring the p_{T} of the dilepton system to be above a certain threshold.

5.3.1 Track selection

Tracks are selected according to the recommendations for tight primary tracks, detailed in Section 3.3.1. Additionally, a $|d_0| < 1\,\mathrm{mm}$ selection is applied to reject tracks from secondary interactions. The final event selection only considers tracks reconstructed with $|z_0^{\mathrm{trk.}} - z_{\mathrm{vtx}}^{\ell\ell}| < 1\,\mathrm{mm}$, where $z_{\mathrm{vtx}}^{\ell\ell}$ is the dilepton vertex defined in the next subsection. The tracks associated to the two highest-p_{T} leptons must pass the track selections but are not considered in track counting when defining analysis regions based on track multiplicity. The definitions of d_0 and z_0 were shown in Figure 3.11.

5.3.2 Lepton selection

Electrons are reconstructed as described in Section 3.3. They are required to have $p_{\mathrm{T}} > 20$ GeV and $|\eta| < 2.47$, excluding the calorimeter crack region of $1.37 < |\eta| < 1.52$. They must satisfy the medium quality and loose isolation criteria, both of which are described in more detail in Section 3.3.2 and Ref. [191]. Overlap removal is applied to any electrons that share the same ID track, with only the higher-p_{T} electron being kept. The significance of the transverse impact parameter calculated with respect to the beam line is required to be $|d_0|/\sigma(d_0) < 5.0$.

Muon candidates are reconstructed as detailed in Section 3.3.3, using combined muons which pass medium ID requirements and loose isolation criteria [165]. Signal muons are required to have $p_{\mathrm{T}} > 20$ GeV, $|\eta| < 2.4$ and $|d_0|/\sigma(d_0) < 3.0$. If any calo-muon[11] shares an ID track with an electron, the calo-muon is removed. If any of the selected muons share an ID track with an electron, the electron is removed.

Primary vertices are reconstructed as described in Section 3.3.1. Within an ATLAS analysis, the hard-scatter vertex is usually selected by choosing the primary vertex with the highest sum-p_{T} associated to it, $\sum p_{\mathrm{T}}^2$. This criterion is based on the assumption that the charged particles produced in hard-scatter interactions have, on average, a harder p_{T} spectrum than those produced in pileup collisions. However, this selection is not optimal for the fully-leptonic $\gamma\gamma \to WW$ process as the signal vertex is only expected to have two associated tracks from the two leptons. Applying the standard PV reconstruction biases the position of the dilepton vertex as it is influenced by the presence of nearby pileup tracks and sometimes will not be selected as the HS vertex at all. Instead, the dilepton vertex, $z_{\mathrm{vtx}}^{\ell\ell}$, is defined to be the weighted average of the two lepton track z_0^{ℓ} positions, given by

$$z_{\mathrm{vtx}}^{\ell\ell} = \frac{z_0^{\ell_1} \sin^2\theta_{\ell_1} + z_0^{\ell_2} \sin^2\theta_{\ell_2}}{\sin^2\theta_{\ell_1} + \sin^2\theta_{\ell_2}}, \tag{5.1}$$

where $1/\sin^2\theta_\ell$ is a proxy for the resolution of the track position based on its angle to the beamline, θ. This gives higher weight to more central leptons ($\theta \to \pi/2$) which have better z_0 resolution. In addition to the track selections listed in Section 5.3.1, the lepton tracks must satisfy $|z_0^{\ell} \sin\theta_\ell - z_{\mathrm{vtx}}^{\ell\ell}| < 0.5$ mm. No selection is made on the lepton vertex position in the x-y plane, other than that already imposed by the selection on individual lepton d_0.

5.3.3 Data acquisition and quality requirements

Only data events in which the detector components were functioning as expected were considered in the analysis. This selection was applied via the ATLAS "Good

[11] Calo-muons have the lowest purity of all the muon types and are therefore only used in the overlap removal but not selected in the analysis. The different types of muons were described in Section 3.3.3.

Runs List" (GRL) which flags blocks of data as good or bad for analysis based on a set of detector checks. The lowest granularity of data which can be flagged as good or bad is one luminosity block (LB), corresponding to roughly one minute of data taking [190]. This analysis used the "All Good" GRL selection, resulting in a data set corresponding to an integrated luminosity of $\mathcal{L} = 139\ \mathrm{fb}^{-1}$.

Candidate events were recorded with either single-electron or single-muon triggers. The lepton that activated the trigger was required to have a transverse momentum above a certain threshold, which varied across different data-taking periods. The single-lepton trigger thresholds are chosen such that the trigger efficiency has reached a plateau at and above the chosen threshold. The efficiencies on the plateau are approximately 70 % for muons with $|\eta| < 1.05$, 90 % for muons in the range $1.05 < |\eta| < 2.4$, and > 90 % for electrons in the range $|\eta| < 2.47$. A trigger matching scheme was employed to ensure that at least one of the fully reconstructed leptons in the event was the object that triggered the event readout.

5.3.4 Event candidate preselection

In addition to the object selections listed above, a simple cut-based selection was applied to select all $\gamma\gamma \to WW$ event candidates for further processing. This selection was as follows:

- Exactly two leptons passing the lepton selection criteria described in Section 5.3.2. No requirement on lepton flavour is made at the preselection stage.

- The *leading* lepton must have $p_{\mathrm{T}} > 27$ GeV and the *subleading* lepton must have $p_{\mathrm{T}} > 20$ GeV. These selections are chosen to maximise trigger efficiencies and reduce fakes.

- The dilepton system must have invariant mass $m_{\ell\ell} > 20$ GeV. This is to avoid the region of hadronic resonances below this threshold, which are challenging to model accurately.

The lepton, vertex, track, and event preselection criteria are summarised in Table 5.2. Additional requirements on the flavour of the leptons, their momentum ($p_{\mathrm{T}}^{\ell\ell}$) and the track multiplicity (n_{trk}) are used to define separate analysis regions, as further detailed in Section 5.3.5.

Selection requirement	Selection value
p_{T}^{ℓ}	$> 27\,\mathrm{GeV}$ (leading), $> 20\,\mathrm{GeV}$ (subleading)
η_ℓ	$\lvert\eta_e\rvert < 2.47$ (excl. $1.37 < \lvert\eta_e\rvert < 1.52$), $\lvert\eta_\mu\rvert < 2.4$
Lepton identification	Medium (Sections 3.3.2 and 3.3.3)
Lepton isolation	Loose (Sections 3.3.2 and 3.3.3)
Dilepton charge	$\mathrm{c}_{\ell 1} \times \mathrm{c}_{\ell 2} < 0$
Number of leptons fulfilling lepton selections	Exactly 2
Vertex selection	Inverse-variance weighted average lepton vertex, $z_{\mathrm{vtx}}^{\ell\ell}$, Equation (5.1)
Lepton-vertex association	$\lvert z_\ell - z_{\mathrm{vtx}}^{\ell\ell}\rvert < 0.5\,\mathrm{mm}$
Track selection	Tight primary (Section 3.3.1), $\lvert z_0^{\mathrm{trk.}} - z_{\mathrm{vtx}}^{\ell\ell}\rvert < 1\,\mathrm{mm}$
Dilepton mass	$m_{\ell\ell} > 20\,\mathrm{GeV}$

Table 5.2: Summary of lepton, track, vertex, and event preselection criteria for $\gamma\gamma \to WW$ candidate events. In the table ℓ stands for e or μ. The definitions of lepton identification and isolation are given in Refs. [191] and [165].

5.3.5 Signal and control region definitions

Many event regions are defined in the $\gamma\gamma \to WW$ analysis, based on characteristics of the dilepton candidate and the underlying event. Three *control regions* (CRs) are defined to check the modelling of processes in a regions of phase-space in which backgrounds are enhanced but which are expected to be kinematically similar to the signal region (SR). A fourth CR is defined to extract a scale factor to account for proton dissociation. All of the control regions satisfy the preselection criteria, as described in Section 5.3.4. The differences in the definitions of the control regions come with different requirements on the flavour of the selected leptons, the multiplicity of tracks within $\pm 1\,\mathrm{mm}$ window around the dilepton vertex and the p_{T} of the dilepton system. The following outlines the details of these requirements and the purpose of each CR, which are also summarised in Table 5.3:

- **SR:** The signal region is defined by requiring different flavour leptons, $e\mu + \mu e$, in the final state in a region with no additional track within $\pm 1\,\mathrm{mm}$ of the lepton vertex and with a large dilepton transverse momentum of $p_{\mathrm{T}}^{\ell\ell} > 30\,\mathrm{GeV}$. The $p_{\mathrm{T}}^{\ell\ell}$ requirement suppresses contributions from exclusive di-tau production ($\gamma\gamma \to \tau\tau$), where each τ-lepton decays to a different flavour leptonic final state. This requirement also helps to further suppress DY and other $\gamma\gamma \to ll$ backgrounds where the flavour of one lepton is wrongly identified. The requirement of no tracks within ± 1 mm of $z^{\ell\ell}_{\mathrm{vtx}}$ is called the *exclusivity selection* or *track-veto selection*, often referred to as the $n_{\mathrm{trk}} = 0$ requirement. This is one of the most important selections within the analysis, as it is the primary way to separate the exclusive $\gamma\gamma \to WW$ signal from the dominant $q\bar{q} \to WW$ background.

- **CR1:** This control region is designed to check the modelling of $\gamma\gamma \to \tau\tau$ process and has identical requirements to the signal region, with the exception of the $p^{\ell\ell}_{\mathrm{T}}$ requirement. CR1 requires $p^{\ell\ell}_{\mathrm{T}} < 30$ GeV.
- **CR2:** The inclusive $q\bar{q} \to WW$ production modelling is tested in a region of low track multiplicity, $1 \leqslant n_{\mathrm{trk}} \leqslant 4$ within ± 1 mm of the dilepton vertex. All other requirements are identical to the signal region.
- **CR3:** This region is designed to check modelling of DY($\tau\tau$) and inclusive $q\bar{q} \to WW$ production. It is very similar to CR2 ($1 \leqslant n_{\mathrm{trk}} \leqslant 4$) but with a low-$p^{\ell\ell}_{\mathrm{T}}$ requirement of $p^{\ell\ell}_{\mathrm{T}} < 30$ GeV.
- **CR4:** The exclusivity scale factor is measured in a region of high dilepton mass, $m_{\ell\ell} > 160$ GeV, and with no tracks within ± 1 mm of the dilepton vertex, $n_{\mathrm{trk}} = 0$. No $p^{\ell\ell}_{\mathrm{T}}$ cut is applied and only the same-flavour leptonic final states $ee + \mu\mu$ are considered to enhance contributions from photon-induced dilepton production, $\gamma\gamma \to ll$.

There is no selection applied to jets and no selection on missing transverse momentum/energy in this analysis.

Control regions	**CR1**	**CR2**	**CR3**	**CR4**	**SR**
Preselection	$p^{\ell_0}_{\mathrm{T}} > 27$ GeV, $p^{\ell_1}_{\mathrm{T}} > 20$ GeV, $m_{\ell\ell} > 20$ GeV, $\lvert z_{\mathrm{trk}} - z^{\ell\ell}_{\mathrm{vtx}} \rvert < 1$ mm				
Lepton flavour	$e\mu + \mu e$	$e\mu + \mu e$	$e\mu + \mu e$	$ee + \mu\mu$	$e\mu + \mu e$
$\mathrm{n}_{\mathrm{tracks}}$	$n_{\mathrm{trk}} = 0$	$1 \leqslant n_{\mathrm{trk}} \leqslant 4$		$n_{\mathrm{trk}} = 0$	
$p^{\ell\ell}_{\mathrm{T}}$	< 30 GeV	> 30 GeV	< 30 GeV	-	> 30 GeV
$m_{\ell\ell}$	-	-	-	> 160 GeV	-

Table 5.3: A summary of the control regions used in the analysis along with the respective definition criteria. The preselection is applied to all candidate events and all other rows represent additional requirements on top of the preselection.

5.4 Corrections to pileup mismodelling in simulation

The exclusivity selection is extremely sensitive to the modelling of pileup. It was found that the standard ATLAS PRW did not sufficiently improve the data/MC agreement for track distributions and therefore, additional corrections were derived within the analysis to ensure that the exclusivity requirement in simulation is not biased. There are two main sources of mismodelling in the distribution of the number of pileup tracks, $n_{\mathrm{trk}}^{\mathrm{PU}}$: The mean position and width of the luminous region of interactions in the bunch crossing, referred to as the *beamspot*, and the modelling of pileup track number and density around an individual pileup interaction vertex. Corrections for both effects are applied to MC and will be described in the following sections. The so-called *exclusive efficiency*, described in Section 5.4.1 is commonly used as a metric to gauge mismodelling of pileup and is applied as a correction to the cross-section measurement. The studies detailed in this section are all the author's own work.

5.4.1 Data-driven method to establish exclusive efficiency

An important quantifier of the impact of pileup tracks is the exclusive efficiency. The exclusivity selection in the SR requires that no tracks can be reconstructed with $|z_0 - z_{\mathrm{vtx}}^{\ell\ell}| < 1\,\mathrm{mm}$ (the track-veto window), other than the tracks matched to the signal leptons. This is referred to as the $n_{\mathrm{trk}} = 0$ requirement and the efficiency to pass this selection is referred to as the *exclusive efficiency*. Nearby pileup vertices and pileup tracks can fall within the track-veto window and cause the signal to fail the exclusivity requirement. This is demonstrated pictorially in Figure 5.1. It is important to quantify the signal inefficiency due to this effect for the cross-section measurement. The probability of an event to pass the $n_{\mathrm{trk}} = 0$ selection is henceforth denoted as $\varepsilon_{\mathrm{excl.}}$ and it is defined as

$$\varepsilon_{\mathrm{excl.}}(z_{\mathrm{vtx}}) = \frac{N_{\mathrm{pass}}^{n_{\mathrm{trk}}=0}(z_{\mathrm{vtx}})}{N_{\mathrm{total}}(z_{\mathrm{vtx}})}, \tag{5.2}$$

where $N_{\mathrm{pass}}^{n_{\mathrm{trk}}=0}(z_{\mathrm{vtx}})$ is the number of events at a given z_{vtx}-position that pass the exclusivity selection and $N_{\mathrm{tot.}}(z_{\mathrm{vtx}})$ is the total number of events with that z_{vtx} value. The exclusive efficiency depends on the track density, which itself is dependent on the beamspot parameters and the number of interactions per bunch crossing, μ.

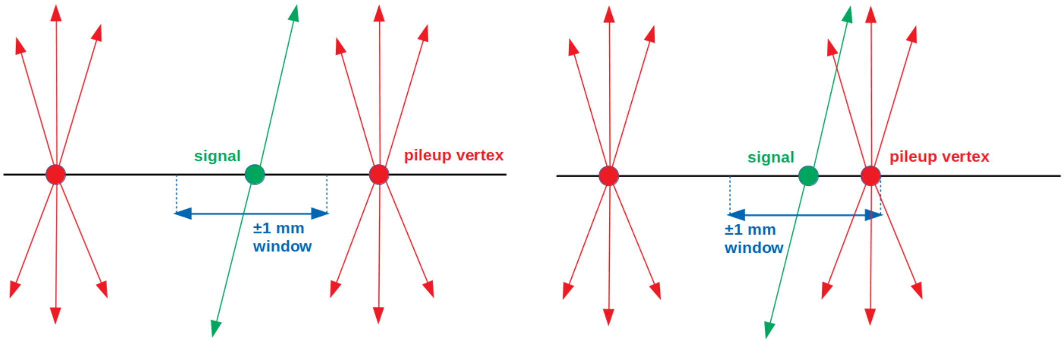

(a) Event passes $n_{\mathrm{trk}} = 0$ selection (b) Event fails $n_{\mathrm{trk}} = 0$ selection

Figure 5.1: The $n_{\mathrm{trk}} = 0$ selection requires that zero tracks are reconstructed with $|z_{\mathrm{trk}} - z^{\ell\ell}_{\mathrm{vtx}}| < 1\,\mathrm{mm}$, other than the tracks associated to the two signal leptons. If a pileup interaction occurs near to the signal hard-scatter, pileup tracks can cause the signal to fail the exclusivity requirement, leading to signal inefficiencies. This is demonstrated in (b).

A data-driven method was developed to measure the exclusive efficiency in the presence of pileup and applied to data and simulation separately. The candidate preselection was applied, as described in Section 5.3.4, to ensure a well-defined signal vertex. No requirement on the dilepton flavour combination was made to reduce statistical uncertainties. The exclusive efficiency was calculated as a function of z_{vtx} by scanning across a z_0-range of $\pm 150\,\mathrm{mm} \approx \pm 5\sigma_{\mathrm{BS}}$, where σ_{BS} is the beamspot longitudinal length, using bins of width $2\,\mathrm{mm}$ that were constructed so that no bin overlapped but all possible values were covered. Each bin in the z_{vtx} distribution was filled on an event-by-event basis if no tracks were found to fall within it, i.e. if $n^{\mathrm{bin}}_{\mathrm{trk}} = 0$, if the bin was farther than $10\,\mathrm{mm}$ from the dilepton vertex. This was repeated for all selected events. Then a ratio was taken to the total number of times each bin was sampled to give the exclusive efficiency as a function of z. The average exclusive efficiency across the beamspot was then calculated by convoluting the exclusive efficiency as a function of z with the normalised lepton vertex distribution and integrating, given mathematically by

$$\varepsilon^{\mathrm{avg.}}_{\mathrm{excl.}} = \sum_{i=1}^{N_{\mathrm{bins}}} \left(\frac{N^{n_{\mathrm{trk}}=0}_{\mathrm{pass}}}{N_{\mathrm{total}}} \right)_i \cdot p_i = \sum_{i=1}^{N_{\mathrm{bins}}} \varepsilon^{i}_{\mathrm{excl.}} \cdot p_i, \tag{5.3}$$

where $\varepsilon^{i}_{\mathrm{excl.}}$ is the exclusive efficiency of bin i (representing one $\pm 1\,\mathrm{mm}$ window in z), calculated using Equation (5.2), and p_i is the probability density of finding the lepton vertex in bin i within the beamspot. A diagram of the method is shown in Figure 5.2.

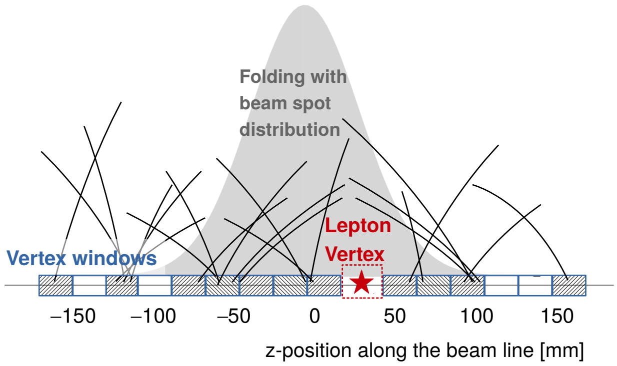

Figure 5.2: A diagram of the exclusive efficiency calculation in pileup. The $\pm 5\sigma_{\mathrm{BS}}$ beamspot longitudinal width is scanned in z-windows of width ± 1 mm, avoiding the region within ± 10 mm of the dilepton vertex, represented by the star, to avoid bias from the signal. The hatched bins shows bins in which additional tracks have been reconstructed and hence would fail the exclusivity requirement [4].

5.4.2 Exclusive efficiency results in data

The exclusive efficiency as a function of z is calculated as given by Equation (5.2) in data, shown in Figure 5.3a. The efficiency is close to unity towards the edges of the distribution as the track density is much lower towards the edges of the beamspot and therefore it is unlikely to find any track with $|z_{\mathrm{trk}} - z_{\mathrm{bin}}| < 1$ mm. The efficiency drops as the sampling window moves closer to the centre of the lepton vertex distribution, which is shown in Figure 5.3b and normalised to unity. The probability of finding the lepton vertex to be at a given position in z, p_i in Equation (5.3), is taken from Figure 5.3b.

The exclusive efficiency can be calculated as a function of any variable that impacts pileup track density, such as the average interactions per bunch crossing and the particular period of data-taking. Figure 5.4 shows the calculated exclusive efficiency across all Run 2 data, as a function of both z-position in the beamspot and number of interactions per bunch crossing. As in Figure 5.3a, the efficiency is lowest around $z = 0$ mm as this is the densest interaction region of the Gaussian beamspot. We now see that the efficiency decreases as μ increases which is expected as a higher number of interactions per bunch crossing increases the average track density within the beamspot. At low values of $\mu < 10$, the exclusive efficiency is higher than 70 % across all values of z. For the highest values of $\mu \approx 70$, the exclusive efficiency drops below 20 % in the centre of the beamspot. The efficiency isn't completely symmetric in z due to the central position of the beamspot in data varying slightly across data-taking.

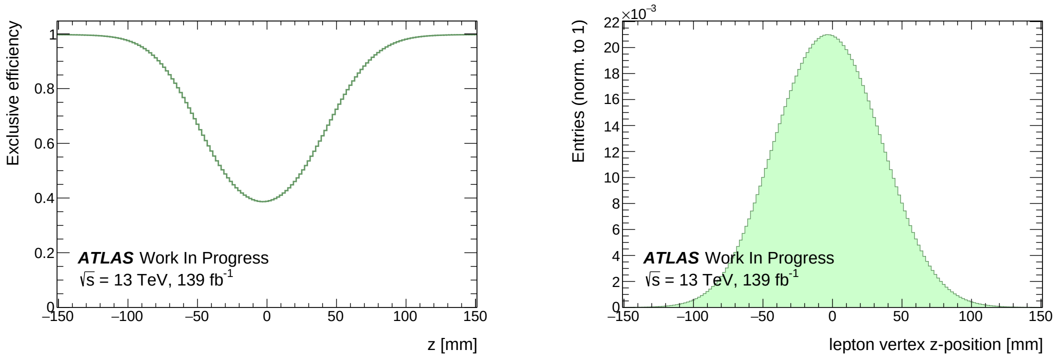

(a) Exclusive efficiency as a function of z (b) Normalised $z_{\mathrm{vtx}}^{\ell\ell}$ distribution

Figure 5.3: (a) The exclusive efficiency calculated in Run 2 data, as a function of position in z. (b) The lepton vertex distribution in z, normalised to unity. The average exclusive efficiency across the beamspot is calculated by multiplying (a) and (b) and summing all bins in z, given by Equation (5.3).

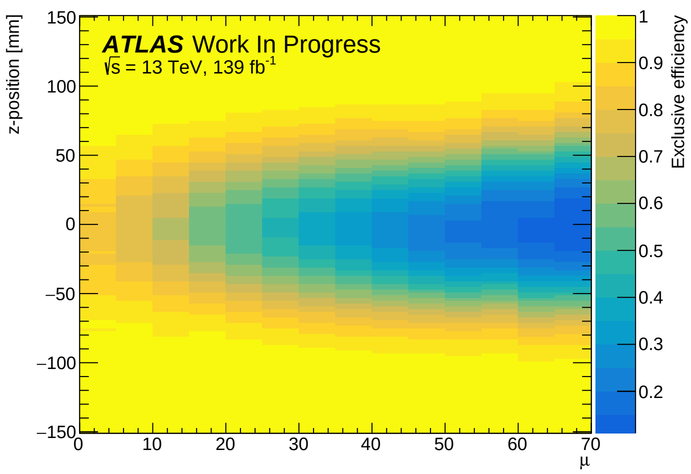

Figure 5.4: The exclusive efficiency in Run 2 data as function of z-position in the beamspot and the number of interactions per bunch crossing, μ.

Data-taking year	Integrated luminosity [fb^{-1}]
2015 + 16	$(3.2 + 33.0) \pm 0.8$
2017	44.3 ± 1.0
2018	58.5 ± 1.2
Combined	139.0 ± 2.4

Table 5.4: The integrated luminosity for each year of Run 2 data taking that passes the "All Good" quality requirements. The values and uncertainties are taken from Ref. [174], with uncertainties ranging from 2.0–2.4 % for the individual years and a total uncertainty on the combined dataset of 1.7 %. 2015 and 2016 are combined as 2015 makes up only 3 % of the total dataset.

By convoluting the measured efficiency with the probability to be at a given z position and integrating over the beamspot width, we can find the average exclusive efficiency as a function of any variable. Figure 5.5a shows the average exclusive efficiency as a function of run number, where the runs have been grouped according to the year of data-taking. One run usually corresponds to a single LHC fill which is $\sim$ hours of data taking and equal to $\sim$ 1000 luminosity blocks. The corresponding μ distribution for each year is shown in Figure 5.5b. The integrated luminosity for each data-taking year is shown in Table 5.4. Although high values of exclusive efficiency are measured in the earlier data-taking years, the larger integrated luminosity in the later years causes the average exclusive

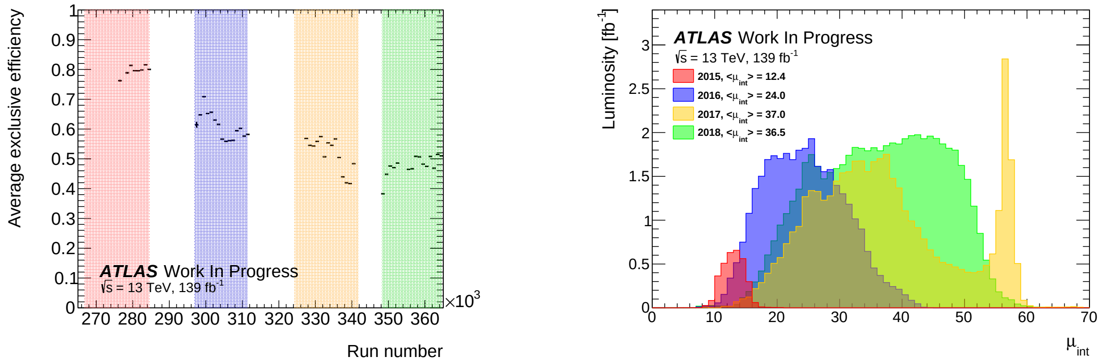

(a) Average $\varepsilon_{\mathrm{excl.}}$ vs. run number (b) μ distribution per data-taking year

Figure 5.5: The average exclusive efficiency as a function of run number in Run 2 is shown in (a). The different coloured regions represent the different data-taking years, i.e. red is 2015, blue is 2016, yellow is 2017, and green is 2018. Uncertainties are statistical only. The corresponding luminosity-weighted μ distribution for each data-taking year is shown in (b), with the mean value for each year given in the legend.

efficiency across all of Run 2 to be pulled towards lower values. The exclusive efficiency was lower in later years due to 2017 and 2018 typically having higher instantaneous luminosity and therefore a higher average number of interactions per bunch crossing. The average exclusive efficiency calculated across all Run 2 data is

$$\varepsilon_{\text{excl.}}^{\text{data}} = 0.52607 \pm 0.00004,$$

where the uncertainty is statistical only and is very small due to each event being sampled multiple times in the data-driven method. This result means that the dilepton vertex position will have a pileup track within $z_{\text{vtx}}^{\ell\ell} \pm 1\,\text{mm}$ 47 % of the time on average.

The same method was applied in simulation, applying standard ATLAS corrections including the PRW. Figure 5.6 compares the full dataset to the simulation, as a function of pileup. The average pileup in data is $\langle\mu\rangle = 33.7$, at which point data/MC disagreement is around 10 % and rises with increasing μ. In order to probe the origins of this discrepancy, a closure test of the exclusive efficiency method was performed in simulation, comparing the data-driven method to the average efficiency extracted directly in the ± 1 mm window around the dilepton

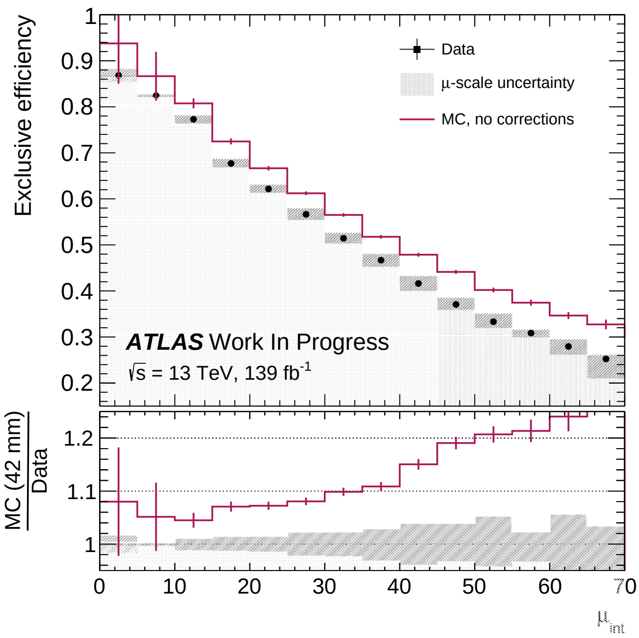

Figure 5.6: The exclusive efficiency calculated as a function of pileup, μ, in both data and simulation. The simulation has the ATLAS pileup reweighting applied but no further corrections. The grey hatched bands on the data represent the uncertainty on the measured and simulated μ. The MC was simulated with beamspot width of 42 mm and the uncertainty on the MC is statistical only.

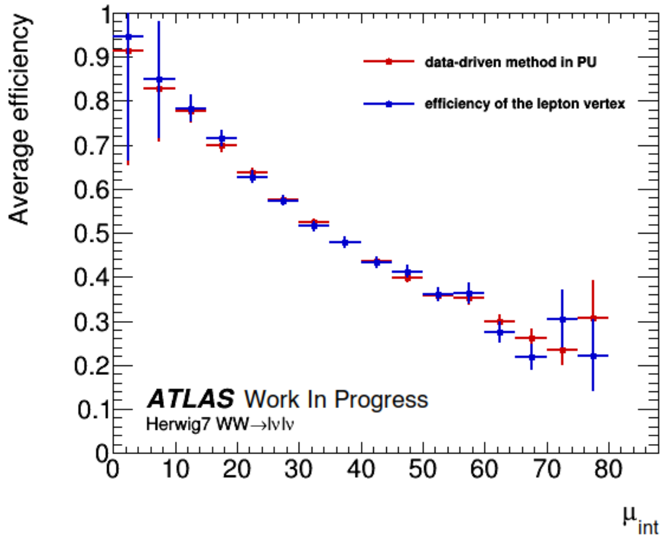

Figure 5.7: Closure test of the data-driven method to calculate the exclusive efficiency, performed in the HERWIG7 signal MC. The data-driven method is applied in the MC to calculate the exclusive efficiency (red) and agrees within statistical uncertainties perfectly with the direct determination of the efficiency measured around the dilepton vertex (blue). It is possible to measure the efficiency around the dilepton vertex directly in exclusive MC as we know that there will be no contamination from additional tracks originating from an underlying event. Uncertainties are statistical only.

vertex, as a function of μ. It is possible to directly measure the number of pileup tracks around the dilepton vertex in simulation, using track truth-match probabilities to assign track categories. Elastic signal MC was used so that no tracks from any underlying event should be present, which could bias the results. Therefore the data-driven method can be applied in MC and compared to the calculation of the exclusive efficiency directly around the dilepton vertex. This is shown in Figure 5.7. The method closes completely within uncertainties and therefore the discrepancy that we see in Figure 5.6 must be due to mismodelling of pileup in the simulation. As this analysis relies so heavily on accurate modelling of track multiplicity and track density in the $n_{\mathrm{trk}} = 0$ selection, it is imperative to correct the simulation to look more like data. Two further data-driven corrections to the pileup distribution in simulation were developed and are described in the following subsections.

5.4.3 Beamspot rescaling in simulation

The luminous region of interactions, the beamspot, has a three-dimensional Gaussian profile with dimensions determined by the operating parameters of the LHC. The longitudinal width in z, σ_{BS}, of the beamspot is of interest in the context of this analysis, since it ultimately determines the average density of interactions per bunch crossing. Both σ_{BS} and the central position of the beamspot, $\langle z_{\mathrm{BS}}\rangle$, vary with time over Run 2, demonstrated in Figure 5.8 which shows the (a) luminosity weighted beamspot longitudinal width and (b) average position measured in data for each year of Run 2 data-taking. The average values of σ_{BS} and $\langle z_{\mathrm{BS}}\rangle$ for each data-taking year are given in the legend. The average across all of Run 2 is $\sigma_{\mathrm{BS}} = 35\,\mathrm{mm}$ and $\langle z_{\mathrm{BS}}\rangle = -3\,\mathrm{mm}$. The transverse width of the beamspot is approximately 8 µm in both x and y and is not relevant in this analysis as its size is much smaller than the selection on the track transverse impact parameter of $|d_0| < 1\,\mathrm{mm}$.

MC is generally produced before the end of data-taking, when typical values of beamspot parameters are unknown. Therefore, it is often the case that MC beamspot parameters do not match those in data. The size of the luminous region in Run 2 was expected to be around 38 mm and all MC samples, including the pileup overlay, were produced with a beamspot longitudinal width of 42 mm to account for possible shifts in the central position. The mean beamspot position in MC was set to $\langle z_{\mathrm{BS}}\rangle = 0\,\mathrm{mm}$ with no variation. A data-driven beamspot rescaling was developed to reshape the beamspot in simulation to look more like data. The rescaling was performed by applying a transform to all primary vertex and pileup track z_0 positions, as so

$$z_0^{\mathrm{scaled}} = \frac{\sigma_{\mathrm{BS}}^{\mathrm{new}}}{\sigma_{\mathrm{BS}}^{\mathrm{old}}} z_0^{\mathrm{unscaled}}, \tag{5.4}$$

assuming an initial Gaussian distribution of longitudinal width $\sigma_{\mathrm{BS}}^{\mathrm{old}}$ which we wish to rescale to have width $\sigma_{\mathrm{BS}}^{\mathrm{new}}$. When rescaling MC to data, the rescaling was performed by randomly sampling the beamspot widths in data, with each sampling point corresponding to 30 runs. The beamspot in MC was then rescaled to data on per-event basis using pseudo-random run numbers to generate a σ_{BS} value corresponding to the correct data-taking period. In principle, the position of the beamspot centre should also be rescaled but the absolute position of the distribution does not matter, only the relative position of tracks and vertices within the beamspot width, so this is left as $\langle z_{\mathrm{BS}}\rangle = 0\,\mathrm{mm}$.

Care must be taken as the aim of the rescaling is to change the density of pileup interactions, not the density of pileup tracks around any given interaction vertex. As tracks generally have z_0 positions that do not match their associated

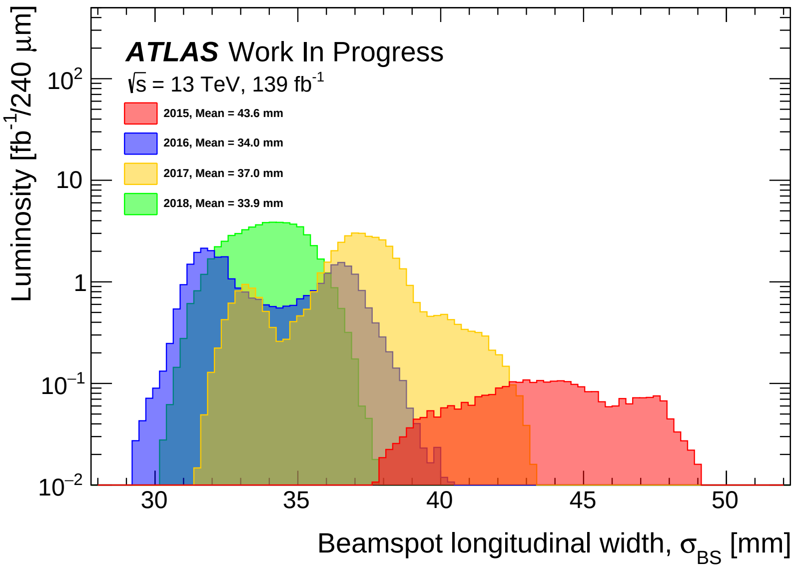

(a)

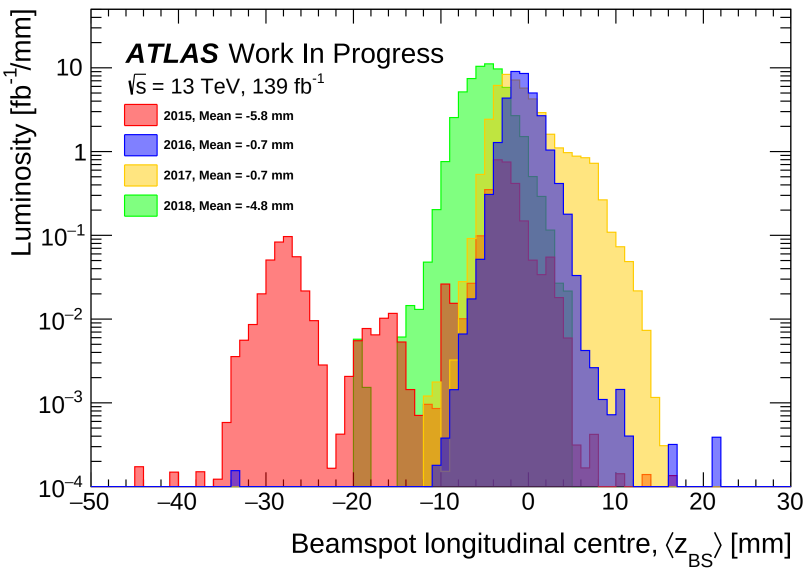

(b)

Figure 5.8: The luminosity weighted distribution of the (a) longitudinal width, σ_{BS}, and (b) mean position, $\langle z_{\mathrm{BS}} \rangle$, of the luminous region in ATLAS during $\sqrt{s}$ = 13 TeV pp collisions in Run 2 of the LHC. The beamspot distribution is determined from a maximum likelihood fit to the spatial distribution of primary vertices collected over a two minute period. The average values for each year of data-taking are given in the legend.

vertex z_0 position, the rescaling method in Equation (5.4) also changes the track density around each vertex when applied to all track and vertex positions. Ideally, Equation (5.4) would only be applied to the interaction vertices and then each track would be shifted by the same amount as its respective vertex. However, it is not generally possible to do this for pileup tracks as the pileup vertices are sometimes incorrectly reconstructed, or not reconstructed at all, due to pileup tracks typically having low p_{T} and therefore larger z_0 resolution. It is also not typical to store pileup track information in the generator record due to storage limitations so lost/incorrect pileup vertices cannot be recovered. Therefore, all pileup vertices and tracks are rescaled using Equation (5.4). Generator-level information is available for the hard-scatter tracks and therefore only the hard-scatter vertex is rescaled using Equation (5.4), whereas the hard-scatter tracks are shifted by the same amount as the HS vertex. This preserves the track density of any underlying event, which should be unaffected by differences in beamspot size.

A closure test of the rescaling method was performed with $Z \to \mu^+\mu^-$ MC samples, using the calculation of the exclusive efficiency as described in Section 5.4.1. A MC sample produced with $\sigma_{\mathrm{BS}} = 35\,\mathrm{mm}$ was compared to a nominal sample with $\sigma_{\mathrm{BS}} = 42\,\mathrm{mm}$ which was rescaled to have $\sigma_{\mathrm{BS}} = 35\,\mathrm{mm}$ using Equation (5.4). The average exclusive efficiency was calculated as a function of the number of interactions per bunch crossing, shown in Figure 5.9. All preselection requirements are applied, as listed in Section 5.3.4, solely for the purpose of having a well-defined dilepton vertex z_0-position. There is some remaining non-closure in Figure 5.9, rising from roughly 2 % at low-μ to 7 % at high-μ. This non-closure is likely to arise from scaling both pileup vertices and pileup tracks in the same way, as discussed above, impacting the track density around each interaction vertex, rather than just the track density of the beamspot as a whole.

Although the simulation produced with $\sigma_{\mathrm{BS}} = 35\,\mathrm{mm}$ should already be a better match to the pileup track distributions in data, the rescaling method was also applied to these samples to account for variations across data-taking. The number of pileup tracks per $\pm 1\,\mathrm{mm}$ window was plotted, $n_{\mathrm{trk}}^{\mathrm{PU}}$, in both the rescaled $\sigma_{\mathrm{BS}} = 42\,\mathrm{mm}$ and rescaled $\sigma_{\mathrm{BS}} = 35\,\mathrm{mm}$ samples and compared to the full Run 2 dataset, shown in Figure 5.10. Again, all preselection requirements are applied with no selection on lepton flavour combination. Neither MC agrees well with data after the rescaling, with data/MC discrepancies for the rescaled

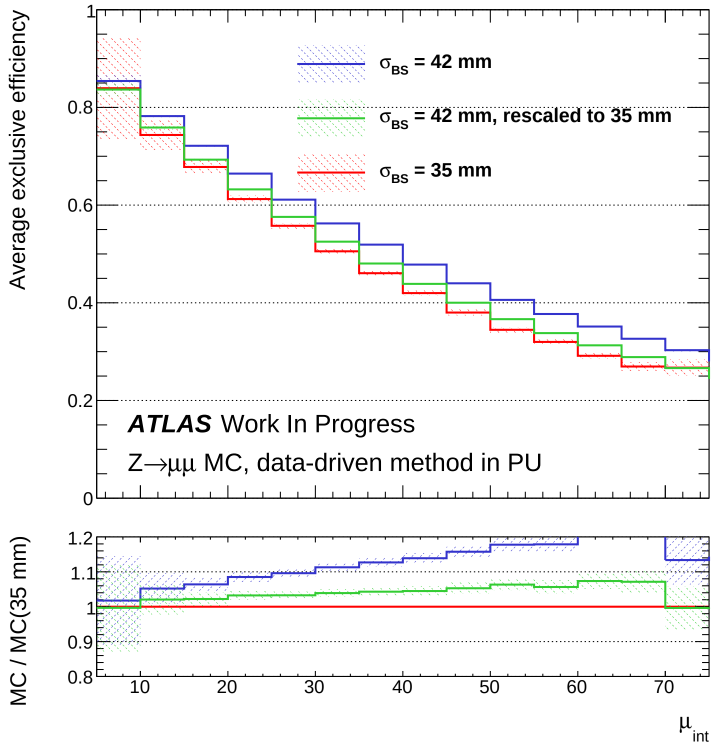

Figure 5.9: Closure test of the rescaling method showing the average exclusive efficiency as a function of the number of interactions per bunch crossing, μ, calculated in pileup using the data-driven method. A $Z \to \mu^+\mu^-$ MC sample with $\sigma_{\rm BS} = 35\,$mm (red) is compared to the same signal process with $\sigma_{\rm BS} = 42\,$mm (blue). The $\sigma_{\rm BS} = 42\,$mm sample was then rescaled to also have $\sigma_{\rm BS} = 35\,$mm (green). The lower plot shows the ratio of the nominal and rescaled sample to the sample produced with $\sigma_{\rm BS} = 35\,$mm. Only regions far from the dilepton vertex are considered to probe pileup tracks only.

42 mm sample of over 10 % at $n_{\rm trk} = 1$ and discrepancies of up to 5 % for the rescaled 35 mm sample. As the non-closure of the beamspot rescaling method was less than 5 % at the average $\langle\mu\rangle = 33.7$ and the discrepancy of the pileup track multiplicity is larger than this at several values of $n_{\rm trk}$, it was concluded that the mismodelling of pileup track multiplicity, and hence mismodelling of the exclusive efficiency, was not due to mismodelling of the beamspot width alone.

To investigate discrepancies between data and MC further, the $p_{\rm T}$ and η distributions of pileup tracks within each $\pm 1\,$mm window were plotted and are shown in Figure 5.11. There is a clear mismodelling in the $p_{\rm T}$ distribution of pileup tracks, shown in Figure 5.11a. This mismodelling is around the level of 10 % at the reconstruction threshold of $p_{\rm T} \approx 500\,$MeV and the level of discrepancy is unaffected by the application of the beamspot rescaling. The mismodelling

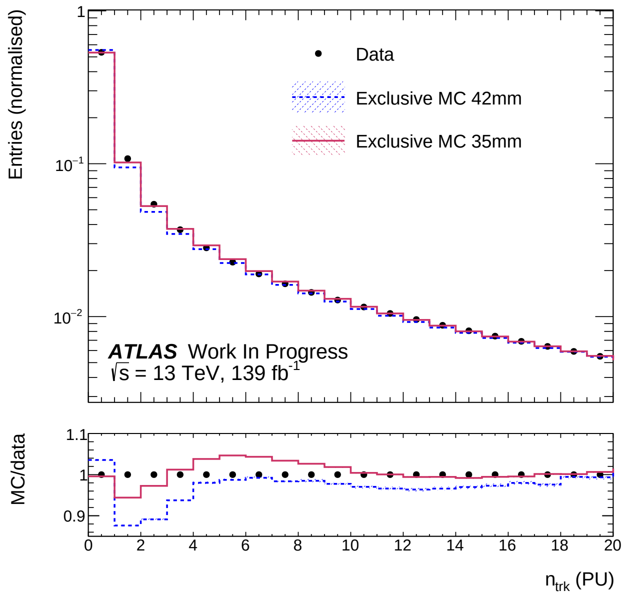

Figure 5.10: Test of the beamspot rescaling method showing the number of pileup tracks reconstructed per ± 1 mm window in z, for windows more than 10 mm from the dilepton vertex. Several simulations samples were used to increase statistics. Simulation with $\sigma_{\mathrm{BS}} = 35$ mm is shown in pink and simulation with $\sigma_{\mathrm{BS}} = 42$ mm is shown in blue with a dashed line. Both samples have been rescaled by generating pseudo-random desired beamspot widths on a per-event basis. Both samples are compared to data.

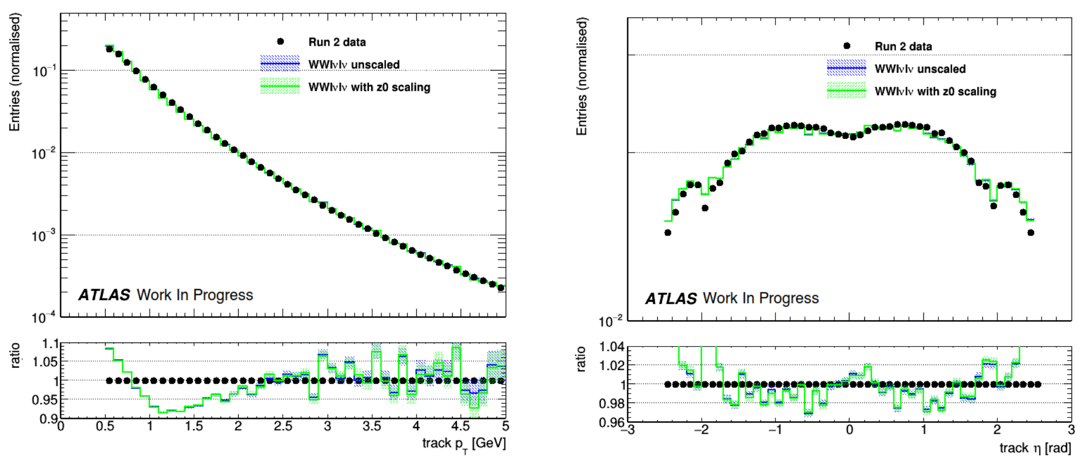

(a) pileup track p_{T} (b) pileup track η

Figure 5.11: Kinematic cross-checks of the beamspot rescaling method: (a) p_{T} distribution and (b) η distribution of pileup tracks within ± 1 mm of the dilepton vertex, for scaled and unscaled elastic $\gamma\gamma \to WW$ MC. Both distributions are compared to the corresponding distribution in Run 2 data. The beamspot rescaling does not impact the pileup track kinematics, which is both expected and desired. However, mismodelling of up to 10 % is seen in the p_{T} distribution.

of pileup track p_{T} at the reconstruction threshold propagates into mismodelling of the pileup track multiplicity around a given interaction vertex. Hence an additional pileup correction is warranted on top of the beamspot rescaling, discussed in the next section.

5.4.4 Correction for pileup track multiplicity mismodelling in simulation

The mismodelling of pileup track p_{T} at the reconstruction threshold results in a mismodelling in the number of reconstructed pileup tracks around each interaction vertex. An additional data-driven correction was developed to correct the $n_{\mathrm{trk}}^{\mathrm{PU}}$ distribution around the dilepton vertex.

The distribution of $n_{\mathrm{trk}}^{\mathrm{PU}}$ was plotted for all ± 1 mm windows in z that satisfy $|z_{\mathrm{window}} - z_{\mathrm{vtx}}^{\ell\ell}| > 10$ mm for both data and MC. The ratio of the data to simulation of the number of pileup tracks vs z-position is shown in Figure 5.12, There is asymmetry in the ratio of data to MC in Figure 5.12a due to the central beamspot position varying in data whereas it is constant and set to zero in MC. To correct for this effect, the sampled z-position in data was shifted by the offset of the mean beamspot position and the ratio was recalculated, shown in Figure 5.12b. This plot was then used to reweight the number of pileup tracks

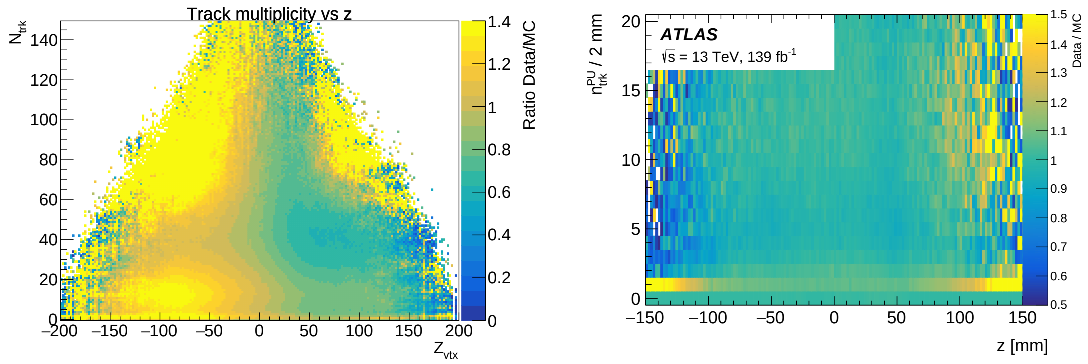

(a) (b)

Figure 5.12: The ratio of the distribution of pileup tracks vs position in the beamspot for data and MC. The ratio in (a) is asymmetric in z due to the beamspot central position varying in data whereas it is always set to $z = 0$ mm in MC. (b) shows the same plot but zoomed in and with this shift corrected. The number of pileup tracks around the hard-scatter vertex in MC is reweighted from (b) by sampling the corresponding z-position of the hard-scatter vertex and the measured number of pileup tracks within the ± 1 mm window around $z_{\mathrm{vtx}}^{\ell\ell}$.

as a function of z, after beamspot rescaling corrections had been applied. The results of the full pileup corrections are presented in Section 5.4.5.

Mismodelling of low-p_{T} pileup tracks is an already known problem. In Run 3, ATLAS will use a new simulation setup for pileup simulation at low-p_{T}, with the aim of improving modelling around tracking thresholds. In Run 2, pileup interactions were generated and showered using PYTHIA8 with the A3 tune. In Run 3, the majority of pileup interactions will be simulated with EPOS [192], with rarer high-p_{T} pileup events generated by PYTHIA8. The author performed similar studies as presented in this section, comparing pileup track kinematic distributions in EPOS and PYTHIA8, which can be found in Appendix B.

5.4.5 Comparison of data and simulation after full pileup corrections

The average exclusive efficiency as a function of the number of interactions per bunch crossing in simulation and in data is compared in Figure 5.13. The results are shown for the simulation before any pileup corrections (other than ATLAS PRW) and at each stage of pileup correction. The bottom ratio panel shows the results in simulation for MC produced with an original beamspot size of $\sigma_{\mathrm{BS}} = 42\,\mathrm{mm}$ and the bottom panel shows the same but for MC produced with an original beamspot size of $\sigma_{\mathrm{BS}} = 35\,\mathrm{mm}$. The solid orange curve in the bottom ratio panel represents the original uncorrected simulation that was also shown in Figure 5.6. As the average longitudinal width of the beamspot was between 34–38 mm for the bulk of the data, the difference between the uncorrected and the corrected predictions for the $\sigma_{\mathrm{BS}} = 35\,\mathrm{mm}$ MC samples is much smaller than for the $\sigma_{\mathrm{BS}} = 42\,\mathrm{mm}$ samples.

The simulation agrees best with data when all corrections are applied in all scenarios, apart from at high-μ for the $\sigma_{\mathrm{BS}} = 35\,\mathrm{mm}$ sample where there is non-closure of around 10 % but with large statistical uncertainty. Simulation and data are consistent within uncertainties for both samples after full corrections are applied, in the main μ-range of interest around the average Run 2 value of $\langle\mu\rangle = 33.7$.

The average exclusive efficiencies for the different data-taking periods in Run 2 are shown in Table 5.5 for data and simulation, where the data-driven method is compared to directly measuring the efficiency around the dilepton

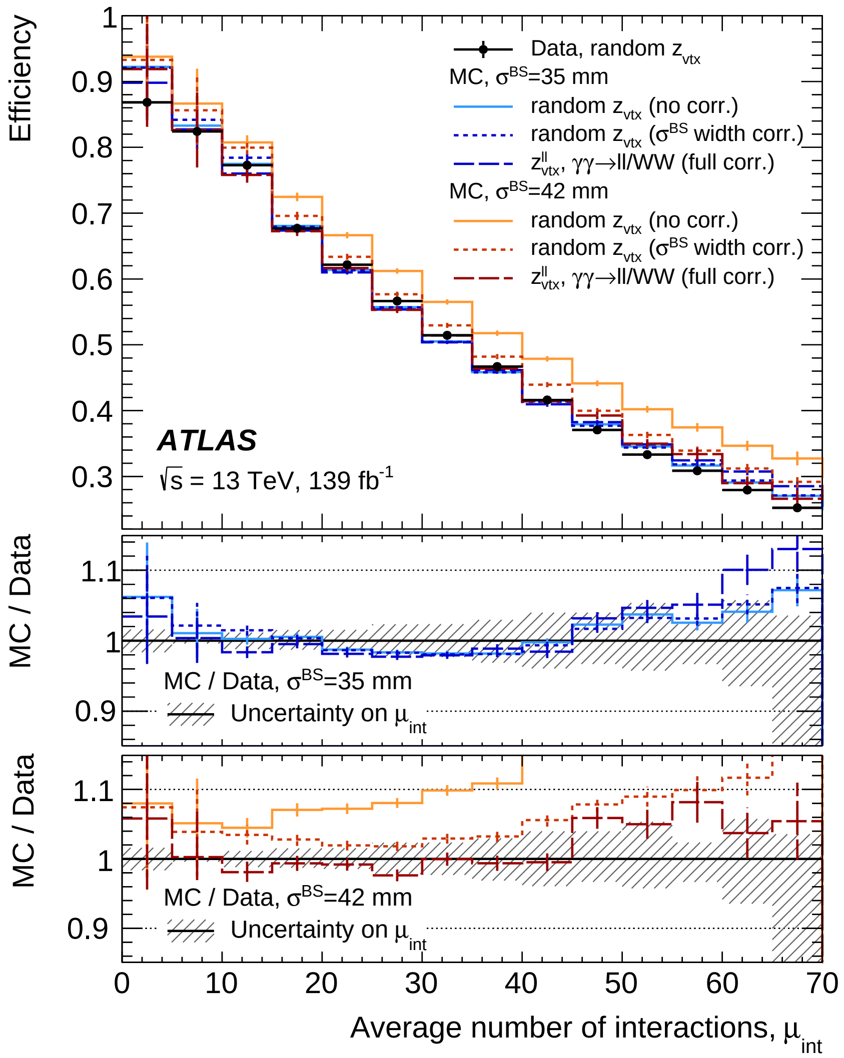

Figure 5.13: The exclusive efficiency of selecting $n_{\text{trk}} = 0$ as a function of the average number of interactions per bunch crossing, μ_{int}. Results are shown for simulation before pileup corrections are applied and at each correction step. The top ratio panel compares simulation produced with 35 mm beamspot to data and the bottom panel compared simulation produced with 42 mm beamspot to data. The error bars show the statistical uncertainty only. The hatched bands show the uncertainty on the average number of interactions. Plot produced by the author for Ref. [4].

vertex in simulation. The two methods in simulation agree within one standard deviation. Residual differences between MC and data are due to a difference in average μ which can be corrected for by applying a scaling of $(1.03 \pm 0.04) \times \mu_{\mathrm{MC}}$. This scaling was applied to produce an accurate comparison of MC to data in Figure 5.13 but is not applied in Table 5.5 as the differences are corrected for automatically in the cross-section calculation described in Section 5.8.4.

Remaining non-closure of the data-driven correction for pileup track density and multiplicity was assigned as a systematic uncertainty, resulting in 1 % and 3 % uncertainty in the efficiency of the $n_{\mathrm{trk}} = 0$ selection for signal and background, respectively. The uncertainty of having a low number of pileup tracks associated with the vertex ($1 \leqslant n_{\mathrm{trk}} \leqslant 4$) is 2 % for photon-induced processes.

5.4.6 Optimising the exclusivity selection

A track-veto window size of ±1 mm was originally chosen as it was the optimal size found in previous ATLAS analyses of exclusive processes [187, 194]. Studies were carried out to check whether the size of ±1 mm was optimal for this analysis and to see if additional cuts could improve the exclusivity selection and increase the exclusive efficiency. In these studies, the ratio $R_{\mathrm{excl}} = \varepsilon_S/\sqrt{\varepsilon_B}$ (with ε_S and ε_B being the efficiency of the signal and the dominant $q\bar{q} \to WW$ background to pass the exclusivity selection in simulation, respectively) is used as a simple

Period	$\langle\mu\rangle$	Data $\varepsilon_{\mathrm{excl.}}$	MC $\varepsilon_{\mathrm{excl.}}$ (data-driven)	MC $\varepsilon_{\mathrm{excl.}}$ ($z^{\ell\ell}_{\mathrm{vtx}}$)
2015–16	13.4 (2015) 25.1 (2016)	0.618	0.637 ± 0.007	0.635 ± 0.010
2017	37.8	0.500	0.512 ± 0.005	0.503 ± 0.008
2018	36.1	0.483	0.503 ± 0.004	0.495 ± 0.007
Full Run 2	**33.7**	**0.526**	$\mathbf{0.541 \pm 0.003}$	$\mathbf{0.536 \pm 0.005}$

Table 5.5: The average exclusive efficiency, calculated in data across Run 2 and compared to simulation. In MC, the efficiency can be calculated both via the data-driven method and also directly measuring the efficiency of the exclusive signal vertex. Uncertainties are statistical only and are not presented for data as they are negligible compared to the uncertainties in MC, being $< \mathcal{O}(10^{-4})$. All pileup corrections have been applied to the simulation, except for a μ-scaling factor of 1.03 ± 0.04 which is used to correct the inelastic cross-section of 80 mb used in MC compared to the results in Ref. [193]. This scale factor is accounted for in the cross-section corrections in Section 5.8.4 by taking the ratio of the calculated exclusive efficiencies in data and MC.

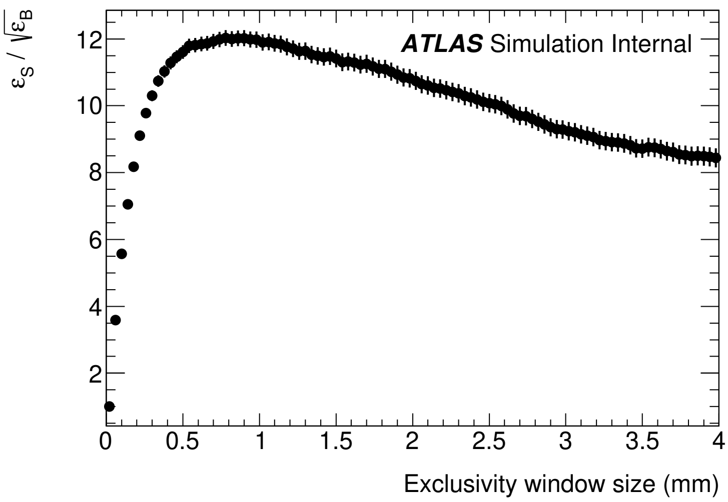

Figure 5.14: Ratio $R_{\text{excl}} = \varepsilon_{\text{S}}/\sqrt{\varepsilon_{\text{B}}}$ as a function of the size of the exclusivity window. Note that the window size here represents half of the total window width in z as the size is measured symmetrically from the dilepton vertex. The errors between bins are strongly correlated because all events without tracks in a given window also have no tracks in any of the smaller windows.

method to gauge the relative improvement of different exclusivity selections. Figure 5.14 shows R_{excl} versus the size of the exclusivity window. As in previous analyses, a shallow maximum is found around a window size of 0.5–1 mm. One possible improvement to the exclusivity selection would be to have a variable window size depending on the number of interactions per bunch crossing. This was tested by splitting the dataset into bins of pileup: $0 \leqslant \mu < 25$, $25 \leqslant \mu < 50$, and $50 \leqslant \mu < 75$. The maximum R_{excl} values and corresponding window sizes are given in Table 5.6, as well as the R_{excl} values for a window size of ± 1 mm. Even though the optimal window size was smaller for higher pileup values, it was decided to opt for the simplest approach of a constant window size as there was very little difference between the optimal sizes and a window size of ± 1 mm due to the shallow maximum. It should also be noted that this simple metric does

μ range	Window size at max. R_{excl} [mm]	Maximum R_{excl}	R_{excl} for 1 mm window
$0 \leqslant \mu < 25$	1.02 ± 0.02	13.9 ± 0.3	13.9 ± 0.4
$25 \leqslant \mu < 50$	0.62 ± 0.02	12.0 ± 0.2	11.7 ± 0.3
$50 \leqslant \mu < 75$	0.66 ± 0.02	10.5 ± 0.5	9.9 ± 0.7

Table 5.6: Maximum R_{excl} and corresponding size of the exclusivity window for different μ ranges, compared to R_{excl} values when using a window size of 1 mm. Here, window size refers to half of the total window width in z.

not take into account any systematic uncertainties, which would alter the optimal window size in a way such as to increase the rejection of backgrounds.

Several ideas were explored to try to improve the performance of a cut-based exclusivity selection, including:

- Accepting events with a single track in the exclusivity window if that track is closer to a track outside of the window than to the lepton vertex.
- Considering lower quality tracks than the default "tight" selection when counting tracks in the exclusivity window.

None of these ideas resulted in a significant performance gain and therefore the exclusivity selection remained no additional reconstructed tracks within ± 1 mm of the dilepton vertex. Optimisation of the track-veto window size shall be discussed in much greater detail in the context of future $\gamma\gamma \to WW$ measurements in Chapter 7.

5.5 Corrections to underlying event mismodelling in simulation

The corrections applied in Section 5.4 correct for mismodelling of pileup interactions but this does not fully match the n_{trk} distribution in MC to look like data around the dilepton vertex due to the presence of tracks originating from the underlying event in inclusive processes ($qq/qg/gg \to X$). Charged particle multiplicity, n_{ch}, mismodelling in the underlying event is particularly prevalent at low charged particle multiplicity [68–70], which is of particular relevance for this analysis. It was imperative to derive a data-driven correction for quark-induced processes, which is described in this section. The author was responsible for performing cross-checks of the charged particle multiplicity, n_{ch}, correction, namely checking differences between the reweighted spectra in both $q\bar{q} \to WW$ and DY simulation, for both the Sherpa and Powheg+Pythia8 (PP8) samples. Discrepancies were seen between the different generators and it was shown that the leading-track p_{T} in the UE was not well modelled. The author developed a lead-track p_{T} reweighting to correct for this.

The majority of charged particles in the underlying event are created in multi-parton interactions at relatively low energy scales [195]. In contrast, the hard-scatter $q\bar{q} \to WW$ process that leads to the dominant background occurs at a high energy scale, as the quarks involved in this interaction must have at least enough energy to create two on-shell or nearly on-shell W-bosons. As the $q\bar{q} \to WW$ process results in a final state that does not carry colour charge (in the case of leptonic W-decays), it is effectively decoupled from any MPI. However, as momentum must be conserved in all collisions, the charged particle multiplicity is correlated to the hard-scatter p_{T}. Therefore, the underlying event kinematics can be assumed to be the same for all such colour-neutral processes that have similar momentum transfers [68]. For this reason, the n_{ch} correction is derived as a function of the p_{T} of the hard-scattering system.

It is not easy to create a pure $q\bar{q} \to WW$ selection region and, due to the neutrinos in the final state of the dominant $q\bar{q} \to WW$ background, it is not possible to reconstruct the total p_{T} of the WW system from the two leptons. Therefore, a selection in data designed to select the also colour-neutral DY dilepton process on the Z-peak ($70\,\mathrm{GeV} < m_{\ell\ell} < 105\,\mathrm{GeV}$) was chosen to derive the correction. Preselections given in Section 5.3.4 were applied to the data and tracks were counted within a $\pm 1\,\mathrm{mm}$ window of the dilepton vertex, such that $|z_{\mathrm{trk}} - z_{\mathrm{vtx}}^{\ell\ell}| < 1\,\mathrm{mm}$.

The purpose of the n_{ch} reweighting is to correct the $n_{\mathrm{ch}}^{\mathrm{UE}}$ spectrum in MC to look like data, i.e. we want to weight the simulated events by something like

$$\text{event weight} = \frac{n_{\mathrm{ch}}^{\mathrm{UE}}(\text{data})}{n_{\mathrm{ch}}^{\mathrm{UE}}(\text{MC})}. \tag{5.5}$$

However, we don't have access to $n_{\mathrm{ch}}^{\mathrm{UE}}(\text{data})$, instead we can only measure the total number of reconstructed tracks,

$$n_{\mathrm{trk}}(\text{data}) = n_{\mathrm{trk}}^{\mathrm{PU}}(\text{data}) + n_{\mathrm{trk}}^{\mathrm{UE}}(\text{data}).$$

Therefore, we need to isolate $n_{\mathrm{trk}}^{\mathrm{UE}}(\text{data})$ and remove detector effects to estimate $n_{\mathrm{ch}}^{\mathrm{UE}}(\text{data})$ at particle-level. Even in such a pure background-enriched region, several data-grooming steps were required to perform this reweighting procedure:

1. Firstly, a data-driven correction was derived in bins of $p_{\mathrm{T}}^{\ell\ell}$ to correct for mismodelling of the p_{T} of the Z-boson, $p_{\mathrm{T}}(Z)$, in MC.

2. Contamination from the $\gamma\gamma \to ll$ process had to be subtracted from the data before weights could be derived. Other backgrounds had a smaller than 1-per-mile contribution in this region and were considered negligible.

3. Pileup tracks were subtracted from the reconstructed n_{trk}(data) distribution to leave only tracks from the underlying event, $n_{\mathrm{trk}}^{\mathrm{UE}}$(data).

4. Detector effects were then removed from the reconstructed $n_{\mathrm{trk}}^{\mathrm{UE}}$(data) distribution to obtain a resulting charged particle distribution, $n_{\mathrm{ch}}^{\mathrm{UE}}$(data), in bins of $p_{\mathrm{T}}^{\ell\ell}$, which is directly correlated to $p_{\mathrm{T}}(Z)$. The contribution of events with large boson p_{T} is small and the n_{ch} reweighting was considered in 5 GeV steps in boson p_{T} up to 30 GeV.

5. Weights were derived as in Equation (5.5) but using the ratio of the unfolded track distribution in data, $n_{\mathrm{ch}}^{\mathrm{UE}}$(data), to the unfolded charged particle distribution in simulation. Although we could obtain $n_{\mathrm{ch}}^{\mathrm{UE}}$(MC) directly from the generator record, it was preferred to use the unfolded $n_{\mathrm{trk}}^{\mathrm{UE}}$ (MC) distribution to at least partially cancel the non-closure of the pileup subtraction in step 3. Any remaining non-closure was taken as a systematic uncertainty.

6. The derived n_{ch} weights were then applied to DY, $q\bar{q} \to WW$, $q\bar{q} \to WZ$ and $q\bar{q} \to ZZ$ MC samples on a per-event basis.

Figure 5.15a compares the underlying event charged particle multiplicity in DY

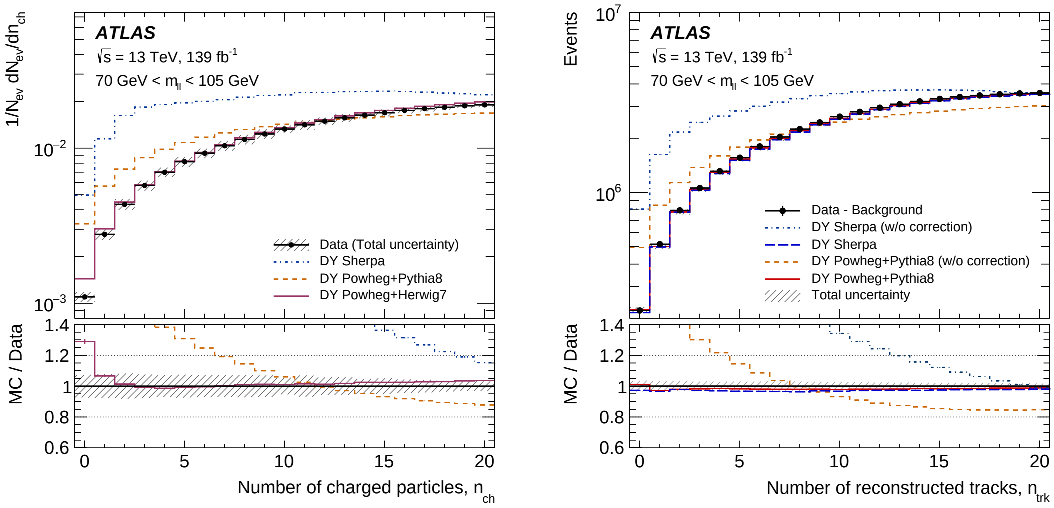

(a) n_{ch} before corrections (b) n_{trk} before and after corrections

Figure 5.15: (a) The normalised charged particle multiplicity distribution, for charged particles satisfying $|z_{\mathrm{ch}} - z_{\mathrm{vtx}}^{\ell\ell}| < 1\,\mathrm{mm}$. Simulation predictions for SHERPA, POWHEG+PYTHIA8 and POWHEG+HERWIG7 are compared with the unfolded data. The effect on the number of reconstructed tracks with $|z_{\mathrm{trk}} - z_{\mathrm{vtx}}^{\ell\ell}| < 1\,\mathrm{mm}$ is illustrated in (b). SHERPA and POWHEG+PYTHIA8 are shown before and after the correction and compared with data. The total uncertainties for SHERPA and POWHEG+PYTHIA8 are very similar [4].

simulation to data, for the different generators before any corrections are applied. and Figure 5.15b shows the huge improvement in the DY reconstructed track multiplicity distribution after the n_{ch} weights are applied, again comparing to data.

5.5.1 Applying the n_{ch} weights to $q\bar{q} \to WW$ samples

The n_{ch} weights were then applied to the $q\bar{q} \to WW$ samples depending on how many charged particles originated from the underlying event on a per-event and per-generator basis, in bins of $p_{\mathrm{T}}(WW)$. The resulting n_{ch} distributions were then compared between DY and $q\bar{q} \to WW$ samples, as well as between different generators. Differences of over 10 % were seen between DY and $q\bar{q} \to WW$ at low n_{ch} for the nominal PYTHIA8 samples, shown in Figure 5.16. This contradicts the assumption that these distributions should be process independent, upon which the correction was derived. Shape differences in the n_{ch} spectra were also seen between PP8 and SHERPA. As the analysis selects on low n_{ch} and low n_{trk} and as there is no physics case to choose one generator over the other, these differences had to be investigated further. A transfer factor, T, between DY and $q\bar{q} \to WW$

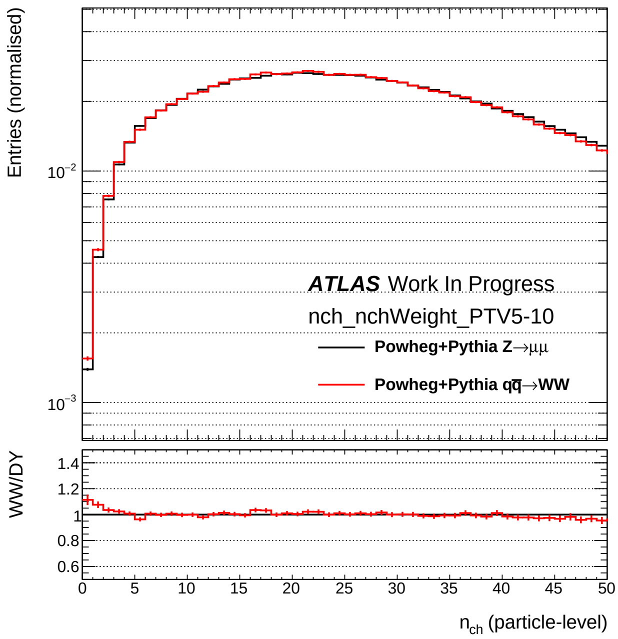

Figure 5.16: The particle-level n_{ch} distributions in DY (black) and $q\bar{q} \to WW$ (red), after the n_{ch} weights have been applied. The distribution is shown for the $5 < p_{\mathrm{T}}(\text{boson}) < 10\,\text{GeV}$ bin, which is where the $p_{\mathrm{T}}(Z)$ distribution peaks. Large differences are seen at low n_{ch} values, which is the important region within the analysis and contradicts the assumption that the n_{ch} spectrum should be the same for the DY and $q\bar{q} \to WW$ process.

simulation was defined as

$$T = \frac{A_{WW}}{A_{\mathrm{DY}}} \text{ with } A = \frac{1}{N}\frac{\mathrm{d}n_{\mathrm{ch/trk}}}{\mathrm{d}N}, \tag{5.6}$$

where A is the normalised n_{ch} or n_{trk} distribution, depending on whether checks were being carried out at particle- or reconstruction-level, and N is the number of events in each bin. To cover all bases, the transfer factor was plotted for the n_{ch} and n_{trk} distributions, both with and without the n_{ch} correction applied. The transfer factor for PP8 and Sherpa were then compared per 5 GeV range of p_{T}(boson). The $5 < p_{\mathrm{T}}(\text{boson}) < 10$ GeV plots are shown in Figure 5.17 as this is where the $p_{\mathrm{T}}(Z)$ distribution peaks. The first check was at particle-level, where

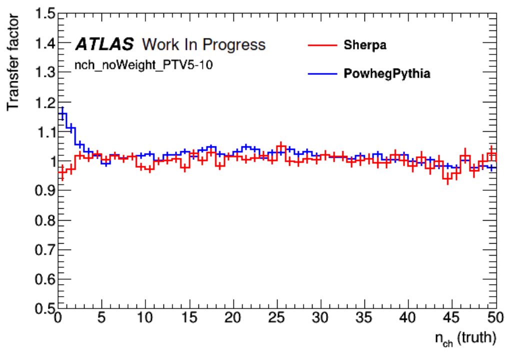

(a) n_{ch}, no weights

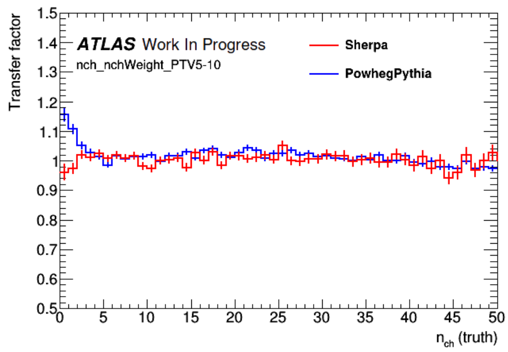

(b) n_{ch}, weighted

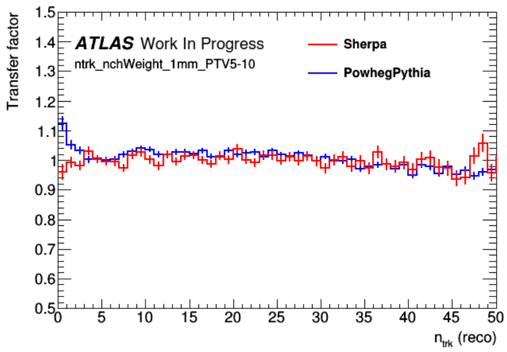

(c) n_{trk}, weighted, $|z_{\mathrm{trk}} - z^{\ell\ell}_{\mathrm{vtx}}| < 1$ mm

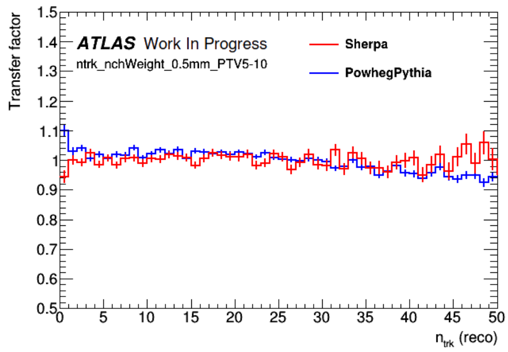

(d) n_{trk}, weighted, $|z_{\mathrm{trk}} - z^{\ell\ell}_{\mathrm{vtx}}| < 0.5$ mm

Figure 5.17: The transfer factor of DY/WW in the n_{ch} and n_{trk} distributions for Powheg+Pythia8 (blue) and Sherpa (red), with and without lead-track p_{T} reweighting applied. The double ratio of PP8/Sherpa in the $n_{\mathrm{ch}} = 0$ bin is 0.84 ± 0.04 for the n_{ch} distribution both before (a) and after (b) the n_{ch} weights are applied. The PP8/Sherpa ratio in the $n_{\mathrm{trk}} = 0$ bin is 0.89 ± 0.04 for the n_{trk} distribution with n_{ch} weights applied (c). A second n_{trk} distribution is also shown varying the size of the track-veto window from ± 1 mm to ± 0.5 mm (d). The PP8/Sherpa ratio in the $n_{\mathrm{trk}} = 0$ bin is also 0.89 ± 0.04. Only the $5 < p_{\mathrm{T}}$(boson) < 10 GeV bin is shown here as that is where the $p_{\mathrm{T}}(Z)$ distribution peaks.

it was seen that the application of the n_{ch} weights had very little effect on the transfer factor (Figures 5.17a and 5.17b). Then the particle-level transfer factor across the n_{ch} distribution was compared to the one across the reconstructed n_{trk} distribution (Figure 5.17c). Again, not much difference was observed but both PP8 and Sherpa were slightly more consistent with $T = 1$ at $n_{\text{trk}} = 0$, meaning that the application of the n_{ch} weights was slightly improving the agreement between the reconstructed track multiplicity distribution between the DY and $q\bar{q} \to WW$ processes for both generators. The effect of varying the window size in which tracks were counted from ± 1 mm to ± 0.5 mm was also investigated (Figure 5.17d), again with very little effect.

Many additional checks were performed to try to find the cause of the differences between DY and WW simulation and between the different generators in an attempt to reduce the associated systematic uncertainty on the n_{ch} correction, including:

- Redefining the particle-level boson p_{T} to better match the reconstructed $p_{\text{T}}^{\ell\ell}$ distributions in data, essentially "dressing" the truth boson by defining its p_{T} using the sum of prompt leptons and prompt neutrinos not coming from hadrons within $0 < |\eta| < 5$.

- Improving the modelling precision of the $p_{\text{T}}(WW)$ distribution by reweighting to higher order theoretical predictions. Both the nominal (Powheg+Pythia8) and alternative (Sherpa) $q\bar{q} \to WW$ MC were produced to NLO precision but more recent NNLO QCD calculations had been performed with Matrix+RadISH [196] in a phase-space very similar to the one used in this analysis, with up to next-to-next-to-next-to-leading-logarithm (N^{3}LL) precision for the p_{T} of the WW system. Weights were derived at particle-level by comparing the Powheg+Pythia8 and Sherpa $p_{\text{T}}(WW)$ distributions to the one from Matrix+RadISH and applying these weights to $q\bar{q} \to WW$ MC samples.

- Alternative generator tunes were tested to observe any differences. The largest differences was observed when varying the multi-parton interaction model used in the simulation but all differences were $< 5\,\%$ for DY and WW individually and did not bring the two generators into agreement.

- Checks on track kinematic distributions, namely the leading-track p_{T} within ± 1 mm of the dilepton vertex (not including prompt leptons), were made to see if any significant differences could be found that may impact the n_{ch} reweighting. These studies are presented in more detail in Section 5.5.2.

Only small improvements, not large enough to cover the differences observed, were found from the majority of checks listed above. As no physics case could be found to choose one generator over another, it was decided to take the average of all distributions (including tune variations) as the nominal prediction of the $q\bar{q} \to WW$ yield, with the envelope of all variations giving the systematic uncertainty on the theory prediction. The systematic uncertainty was previously taken as the absolute difference between PP8 and Sherpa.

Large differences in the lead-track p_{T} distributions between the different MC generators and data were seen, discussed in the following section. However, as these checks were being performed alongside the publication and the expected significance of the result (Section 5.8) was above the threshold for observation, the studies were not finalised and are left open to future improvements.

5.5.2 Checks on the leading-track p_{T} in the underlying event

During investigations into the differences between generator predictions for the n_{ch} spectrum, it was found that there were also differences in the p_{T} distribution of the leading-track p_{T}, i.e. the track with the highest p_{T} that was not associated to the prompt leptons. The ratio of the leading-track p_{T} spectrum between the DY and $q\bar{q} \to WW$ process was taken in both generators. DY/$q\bar{q} \to WW$ (DY/WW) was flat in PP8 but a slope was seen in Sherpa. The difference in the leading-track p_{T} spectrum was unaffected by the application of the n_{ch} reweighting.

Before the decision was made to redefine the $q\bar{q} \to WW$ prediction as the envelope of multiple theoretical predictions, it was decided to reweight the Sherpa leading-track p_{T} spectrum in $q\bar{q} \to WW$ simulation to look like the PP8 spectrum, using the double ratio of

$$\textsc{Sherpa}\left(\frac{\mathrm{DY}}{WW}\right) \div \mathrm{PP8}\left(\frac{\mathrm{DY}}{WW}\right). \tag{5.7}$$

Using the double ratio to reweight the Sherpa $q\bar{q} \to WW$ simulation in this way also accounts for any differences in DY simulation between the two generators. The ratios of DY/WW in Sherpa and PP8, as well as the double ratio of Equation (5.7) are shown in Figure 5.18. Here, lead-p_{T} refers to the particle-level

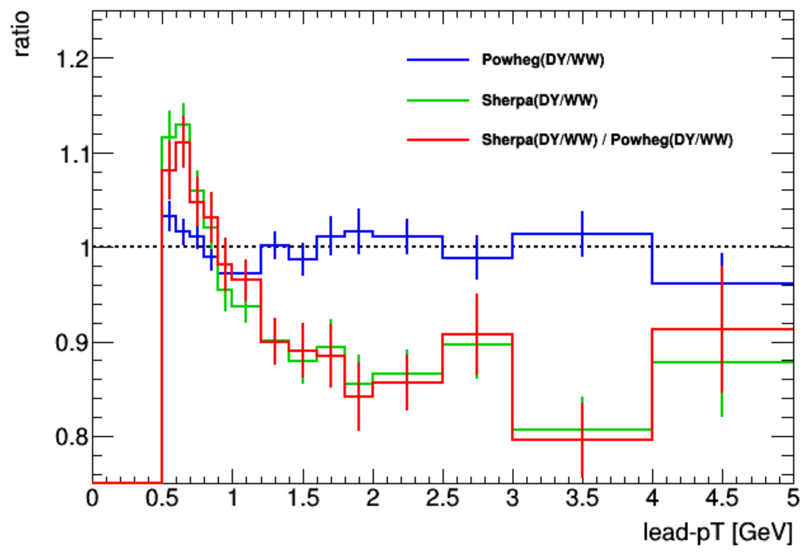

Figure 5.18: The double ratio (red) of SHERPA (DY/WW) (green) to POWHEG+PYTHIA8 (DY/WW) (blue) in the leading-track p_{T} spectrum.

charged-particle p_{T}. All preselection requirements plus a $n_{\mathrm{ch}} < 5$ selection are applied.

After applying the weights derived from the double ratio on a per-event basis, both the charged particle multiplicity and reconstructed track multiplicity distributions, n_{ch} and n_{trk}, were investigated to see if there was any change. These distributions are shown in Figure 5.19. The leading-track p_{T} reweighting brought the two generators into slightly better agreement but differences of up to 10 % are still seen in certain $n_{\mathrm{ch}}/n_{\mathrm{trk}}$ bins. The lead-p_{T} reweighting should not impact events with zero charged-particles/reconstructed-tracks as there is no leading track to reweight in this case. The differences in weighted to unweighted SHERPA distributions are consistent in the $n_{\mathrm{ch}}/n_{\mathrm{trk}} = 0$ bin, within statistical uncertainties.

As it was decided in the analysis to move away from taking PP8 as the nominal $q\bar{q} \to WW$ MC, a data-driven correction to the leading-track p_{T} in both PP8 and SHERPA was derived. As it is difficult to select an enriched $q\bar{q} \to WW$ region in data, event weights were derived using DY MC. A smoothing function was used initially to avoid propagating any statistical fluctuations in data to the reweighting but it was found that this shifted the position of the peak in differences between the distributions. Several alternatives were tested before settling on a combination of a 9th order polynomial (valid for track-$p_{\mathrm{T}} < 2$ GeV) and linear fit (for 2 GeV $\leq$track-$p_{\mathrm{T}} \leq$ 5 GeV) to the ratio of data/MC. Events

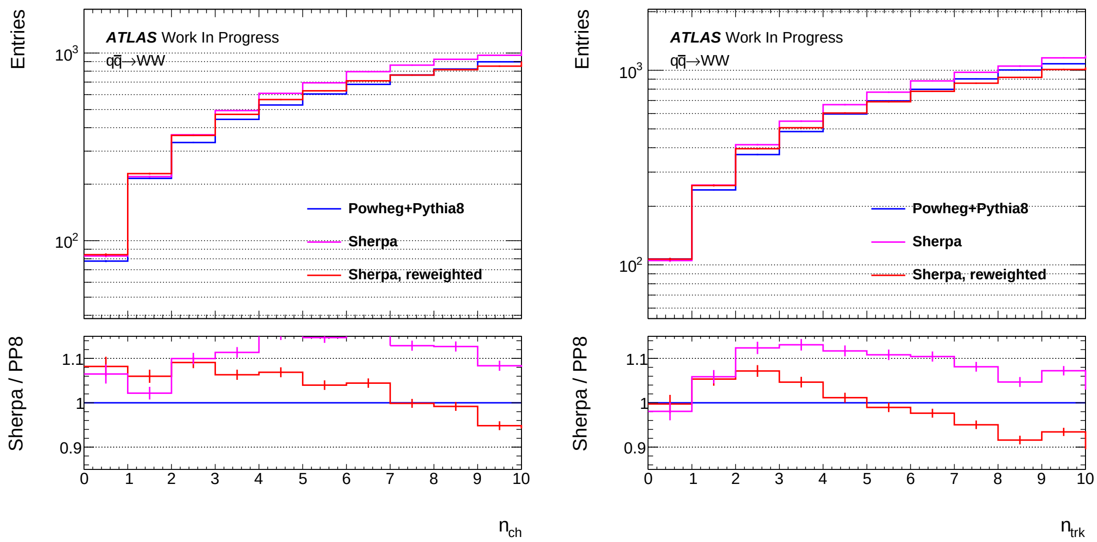

(a) n_{ch} (b) n_{trk}

Figure 5.19: The impact of the double ratio weights when applied to Sherpa $q\bar{q} \to WW$ MC at both particle-level (left) and reco-level (right).

with lead-$p_{\mathrm{T}} > 5\,\mathrm{GeV}$ were weighted with the weight at $p_{\mathrm{T}} = 5\,\mathrm{GeV}$. A weight of one was applied to any events with no additional charged particles with $p_{\mathrm{T}} > 500\,\mathrm{MeV}$. This function was then used to reweight the lead-track p_{T} distribution in MC and observe any differences in the n_{trk} distribution that this caused. The effect of applying this reweighting to the $q\bar{q} \to WW$ samples is shown in Figure 5.20. There is some change in each MC independently when applying the lead-p_{T} reweighting but the ratio between the two remains fairly unchanged.

Ultimately, the reweighting studies documented in this section were not used in the observation analysis. They may be revisited in an ongoing differential measurement if similar discrepancies are observed between different processes, generators and between data and simulation. In general, it is difficult to deconvolute individual aspects of mismodelling in the leading-track p_{T}, $p_{\mathrm{T}}(\text{boson})$, and n_{ch} distributions. This is enhanced by thresholds like the minimum-reconstructed track-p_{T} of $500\,\mathrm{MeV}$ which make any shifts between data and MC difficult to analyse for very soft tracks and also by the presence of pileup tracks when trying to develop data-driven corrections.

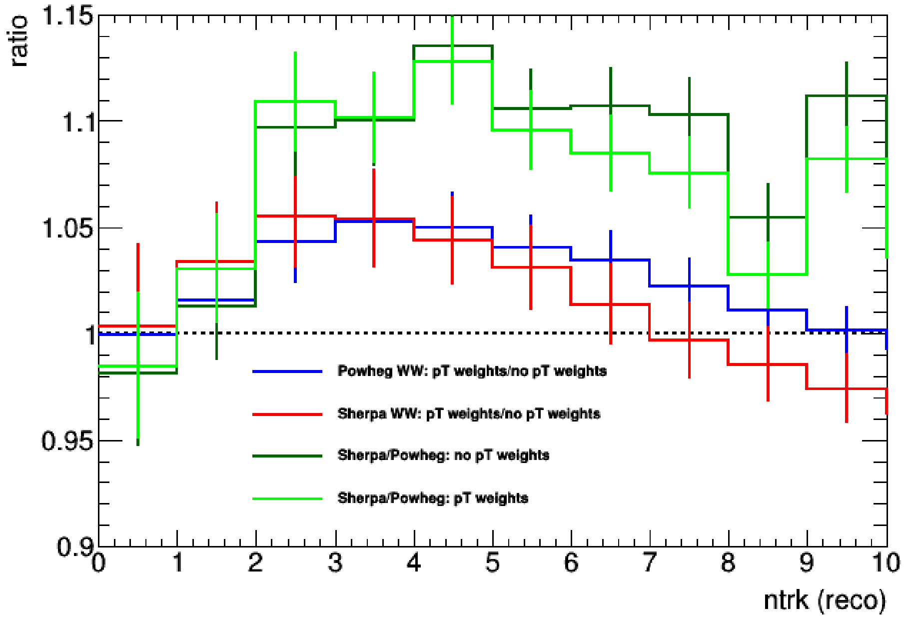

Figure 5.20: The effect of reweighting leading-track p_{T} in the underlying event to data on the n_{trk} distribution in $q\bar{q} \to WW$ simulation. Event weights are derived for POWHEG +PYTHIA8 and SHERPA individually. All preselection cuts are applied as well as requiring $p_{\mathrm{T}}^{\ell\ell} > 30\,\mathrm{GeV}$. No selection on n_{trk} is made.

5.6 Signal modelling corrections

At the time of the analysis, the available dissociative $\gamma\gamma \to WW$ simulation was generated by MADGRAPH and showered by PYTHIA8 (MG5+Py8). However, in the case of single-dissociative simulation, it was only possible to have ISR and FSR for both protons or for neither which is not physical in the case where one proton scatters elastically and one breaks apart, leading to mismodelling of the n_{ch} and n_{trk} distributions for the single-dissociative signal. This issue is not present in the more physically motivated LPAIR but this generator cannot simulate the $\gamma\gamma \to WW$ process. As dissociation is modelled in the same way for the $\gamma\gamma \to ll$ and $\gamma\gamma \to WW$ process in MG5+Py8, the n_{trk} spectrum was compared between MG5+Py8 and LPAIR for the $\gamma\gamma \to ll$ process in the place of the $\gamma\gamma \to WW$ simulation. Differences in the normalised n_{trk} spectrum of up to 80 % were seen for $n_{\mathrm{trk}} = 0$ and MG5+Py8 predicted much higher average track multiplicity than LPAIR. If ISR was turned off in the MG5+Py8 simulation, MG5+Py8 underpredicted the track multiplicity relative to LPAIR. Independent corrections to the signal event yields were derived in the $n_{\mathrm{trk}} = 0$ and $1 \leqslant n_{\mathrm{trk}} \leqslant 4$ regions, discussed in the following subsections.

5.6.1 $n_{\mathrm{trk}} = 0$ regions

For $n_{\mathrm{trk}} = 0$, a data-driven method similar to that used in Refs. [187] and [197] was developed. It is assumed that the additional activity from proton dissociation is decoupled from the hard-scatter process, other than from invariant mass conservation considerations. Therefore, a $\gamma\gamma \to ll$ control sample can be used to estimate the contributions from single-dissociative (SD) and double-dissociative (DD) production modes, where one or both of the protons break up after the photon emission, respectively. The derivation was performed in the exclusive $\gamma\gamma \to ll$ control region, CR4, defined in Section 5.3.5. A $m_{\ell\ell} > 160\,\mathrm{GeV}$ selection was made so that the two photons in principle have enough energy to produce two on-shell W-bosons. The same-flavour lepton selection is a requirement of flavour conservation in the $\gamma\gamma \to ll$ process and automatically results in a data sample that is completely orthogonal to the signal region data.

The exclusive scale factor, S_{excl}, is defined as

$$S_{\mathrm{excl}} = \frac{N_{\mathrm{data}} - N_{\mathrm{bkgd}}}{N_{\mathrm{EE}}}, \tag{5.8}$$

where N_{data} is the measured number of events in data within the CR4 selections, N_{bkgd} is the expected number of background events in the region, and N_{EE} is the number of elastic (EE) $\gamma\gamma \to ll$ events predicted in MC. The contamination of DY events in CR4 is comparable to the $\gamma\gamma \to ll$ yields and therefore the DY contribution had to be subtracted from the data sample before S_{excl} could be derived. Rather than take an estimation from DY simulation, which has uncertainties associated to the shape and normalisation, the DY background in CR4 was also estimated from data by varying the n_{trk} requirement. The n_{trk} distribution for exclusive processes like $\gamma\gamma \to ll$ and $\gamma\gamma \to WW$ fall off very rapidly as n_{trk} increases above $n_{\mathrm{trk}} = 0$. However, the opposite is true for DY production, which is expected to be produced with additional tracks from the underlying event. A selection of $n_{\mathrm{trk}} = 5$ was chosen to extract the background normalisation in a region around the Z-peak, whilst $n_{\mathrm{trk}} = 2$ was used to evaluate the systematic uncertainty in this selection.

Once the DY background was subtracted, the value of S_{excl} was extracted from a combined $e^+e^- + \mu^+\mu^-$ channel to reduce the statistical uncertainty. The two channels were checked independently and no differences were seen within uncertainties. The measured value of the exclusive scale factor was

$$S_{\text{excl}} = 3.59 \pm 0.06\,(\text{stat.}) \pm 0.14\,(\text{syst.}),$$

with the region used to extract the value shown in Figure 5.21. The advantage of the data-driven method developed in the $n_{\text{trk}} = 0$ regions is that it naturally accounts for proton soft-survival, as well as proton dissociation. On the other hand, a large assumption in this method is that the ratio of elastic to dissociative events is the same between the $\gamma\gamma \to ll$ and $\gamma\gamma \to WW$ process, within the analysis selections. Whether this assumption is physically well-motivated is discussed in detail in Ref. [198] and a systematic uncertainty is assigned to the transferability of the method, described in Section 5.7.3. Ultimately, the value of S_{excl} was extracted from a profile-likelihood fit in the SR and all CRs, described in Section 5.8, with

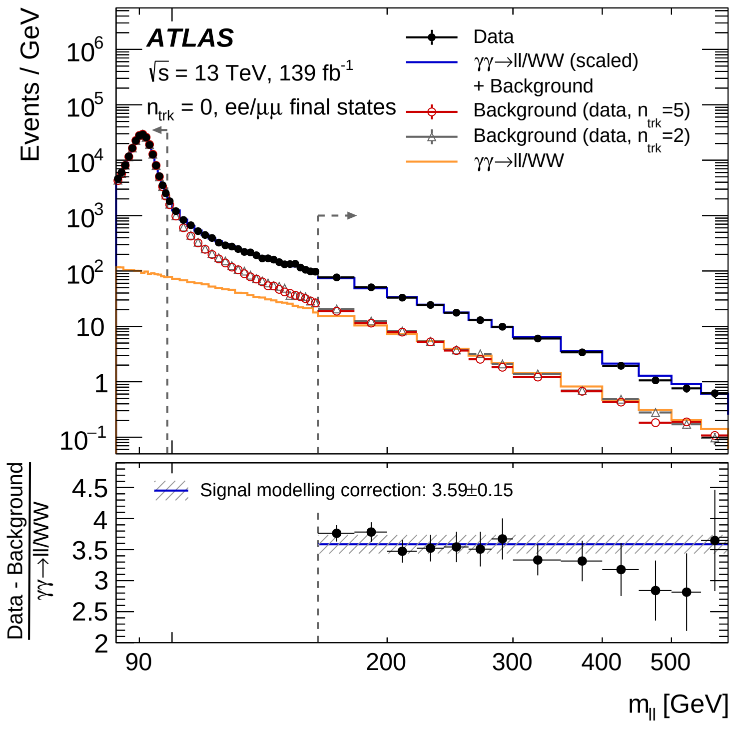

Figure 5.21: The distribution of $m_{\ell\ell}$ in CR4, where the exclusive scale factor is extracted. Shown are the data, where a requirement of $n_{\text{trk}} = 0$ has been applied, and the background templates selected from data using $n_{\text{trk}} = 2$ and $n_{\text{trk}} = 5$. In addition, the $\gamma\gamma \to ll$ and $\gamma\gamma \to WW$ MC predictions are depicted, as well as the sum of the nominal background template ($n_{\text{trk}} = 5$) and the $\gamma\gamma \to ll$ and $\gamma\gamma \to WW$ MC predictions scaled by the signal modelling correction. The normalisation region around the nominal Z-boson mass is indicated with a vertical dashed line, as is the region where the signal modelling correction is extracted ($m_{\ell\ell} > 160\,\text{GeV}$). The uncertainties shown are statistical only [4].

systematic uncertainties as determined in this section. The value calculated in this section was used as a cross-check to the exclusive normalisation extracted from the fit.

5.6.2 $1 \leqslant n_{\mathrm{trk}} \leqslant 4$ regions

The MG5+Py8 simulation was used to estimate the contribution of dissociative events in the $1 \leqslant n_{\mathrm{trk}} \leqslant 4$ control regions. Despite charged particle multiplicity mismodelling still being prevalent in these regions, the impact of the mismodelling is much smaller in these CRs compared to the $n_{\mathrm{trk}} = 0$ regions due to the smaller fraction of exclusive processes relative to inclusive backgrounds at $n_{\mathrm{trk}} > 0$. The issues with ISR and FSR in MG5+Py8 leads to the single-dissociative MG5 samples having more tracks associated to the dilepton vertex than expected when ISR is turned on. The LPAIR n_{trk} distribution in $\gamma\gamma \to ll$ simulation was used to reweight the MG5+Py8 n_{trk} distribution in the single-dissociative $\gamma\gamma \to ll$ samples, in ranges of diphoton mass, $m_{\gamma\gamma}$. This reweighting was then applied to the MG5+Py8 n_{trk} spectra in the single-dissociative $\gamma\gamma \to WW$ samples. The showering problem is not present in the double-dissociative simulation and no correction is applied to these samples.

5.7 Systematic uncertainties

Any prediction or measurement of event yields has associated statistical and systematic uncertainties. Statistical uncertainties arise from finite datasets and systematic uncertainties cover any other unknown in the analysis method, or at least the *known unknowns*. Systematic uncertainties can occur at any point in the analysis chain, from the event generation and the theoretical models that are input to simulation, to object reconstruction and detector calibration. Many systematic uncertainties from early in the analysis chain are already calculated by dedicated combined performance groups and their effects are propagated in the analysis to obtain uncertainties on event yields. This section summarises the sources of systematic uncertainty on event yields in the analysis.

5.7.1 Trigger, lepton and tracking uncertainties

The uncertainty in the modelling of the trigger is evaluated by varying the single-lepton trigger scale factors and efficiencies in simulation and observing the impact on predicted event yields [199, 200]. The resulting variations are below 0.15 % for $\gamma\gamma \to WW$ and $q\bar{q} \to WW$ events, and below 0.25 % for DY events.

Uncertainties in the track reconstruction originate from uncertainties in the impact parameter resolution, residual ID alignment uncertainties and uncertainties in the track reconstruction efficiency and fake-rate. These uncertainties are evaluated, respectively, by smearing the impact parameter of all tracks in simulated events, biasing the d_0, z_0, and p_{T} of tracks in simulated events, by randomly removing good tracks according to the efficiency of their reconstruction, and by removing all fake tracks. The effect of each is symmetrised and considered as a source of uncertainty. The tracks associated to the signal leptons are not included in the tracking systematic uncertainty calculations. For photon-induced processes, these uncertainties are small and below 1% in the $n_{\mathrm{trk}} = 0$ regions. For backgrounds, the tracking uncertainties are a few percent in all kinematic regions.

Lepton systematic uncertainties arise from electron and muon reconstruction efficiencies, energy and momentum resolution, and identification and isolation efficiencies [159, 165]. In the SR, the uncertainties for $\gamma\gamma \to WW$, $q\bar{q} \to WW$ and DY production in the electron and muon reconstruction are approximately 0.5% each and are of similar size in the CRs. The total uncertainty on the background contribution from non-prompt leptons ranges from 50–100 %, depending on the region.

5.7.2 Pileup and underlying event modelling systematics

Uncertainties in the ATLAS pileup reweighting are evaluated by varying the scale factor with which the number of interactions per bunch crossing in data is shifted relative to the value in simulation by $\pm$ 4 %, which corresponds to the uncertainty on the scale factor. As additional pileup modelling corrections are applied to all simulation in the analysis, as described in Section 5.4, it is important to avoid double counting of uncertainties. This is avoided by removing any selection on the number of tracks when calculating the pileup reweighting uncertainty, so that this uncertainty only assesses the effect of pileup on the event, and electron and muon reconstruction. The uncertainties arising from the ATLAS PRW are found to be approximately 0.5% for exclusive processes and 1.5% for $q\bar{q} \to WW$ and DY production. The full effect of the analysis-specific pileup corrections derived in Section 5.4.4 are assigned as a systematic uncertainty, propagating to 3.1 % in the yields in $n_{\mathrm{trk}} = 0$ regions in samples with the beamspot simulated with a longitudinal width of 42mm and 1.0 % in samples with beamspot longitudinal width of 35 mm. The uncertainties in $1 \leqslant n_{\mathrm{trk}} \leqslant 4$ events are 2.0 % for exclusive events (with $\sigma_{\mathrm{BS}} = 35\,\mathrm{mm}$) and 10 % for inclusive events (mostly $\sigma_{\mathrm{BS}} = 42\,\mathrm{mm}$), respectively.

The correction to the underlying event was described in Section 5.5. Uncertainties were derived by studying the non-closure between unfolded and particle-level simulation and varying the mass range of the Z-peak region in which the correction was derived. Systematic uncertainties in the $\gamma\gamma \to ll$ subtraction were calculated from differences in a muon-only and an electron-only fit, and investigating different MC generators for the unfolding and shape templates. The largest sources of uncertainty were found in the non-closure of the pileup subtraction and the transfer from DY to $q\bar{q} \to WW$ events, resulting in around 2.5 % uncertainty in most regions, apart from for DY events in the SR in which the uncertainty due to the unfolding is roughly 5 %. Many of the uncertainties associated to the underlying event correction are correlated and can be further constrained through the profile-likelihood fit, discussed in Section 5.8.

5.7.3 Exclusive scale factor systematics

The method to calculate the exclusive scale factor, S_{excl}, was described in Section 5.6.1. Systematic uncertainties arose in the data-driven background estimation, the $m_{\ell\ell}$ window used to calculate the background normalisation factor and the choice of $m_{\ell\ell}$ range in which S_{excl} was calculated. An alternative background estimation was performed by altering the n_{trk} selection selection, resulting in a systematic uncertainty of 3.61 %. The systematic uncertainty on the background normalisation was calculated by varying the $m_{\ell\ell}$ window used to define the Z-peak, resulting in an uncertainty of 1.11 %.

A transfer uncertainty was also applied to S_{excl} to account for any differences between the ratio of elastic to dissociative events between the $\gamma\gamma \to ll$ and $\gamma\gamma \to WW$ process, within the selections used for the extraction. The uncertainty on the transfer applicability was calculated by varying the lower bound on the $m_{\ell\ell}$ range used for the extraction, nominally $m_{\ell\ell} > 160\,\mathrm{GeV}$. The lower bound was varied between 100 GeV and 400 GeV, in increments of 10 GeV and the uncertainty was taken as the largest difference in the resulting S_{excl} from the nominal value. This gave an uncertainty of 11.20 % and is the largest uncertainty associated with the method. As this uncertainty applies to the scaling of the $\gamma\gamma \to WW$ yields, it consequently impacts the signal strength parameter, as well as cross-section predictions that use the exclusive scale factor to predict the dissociative contribution to the signal. However, it cancels in the fiducial cross-section measurement.

5.7.4 Theoretical systematic uncertainties

The main theoretical uncertainty in this analysis arises due to the modelling of quark-induced backgrounds produced with few additional tracks in the underlying event. These were evaluated by comparing the event yields predicted from the nominal and alternative event samples listed in Tables 5.1. For the Drell-Yan background, a comparison of events simulated with POWHEG+PYTHIA8 and SHERPA were made. The prediction from PP8 is taken as the central value and the difference to SHERPA is assigned as the upward and downward systematic uncertainty. The same procedure is used for backgrounds from $qq \to WZ$ and $qq \to ZZ$ production.

A large non-reducible uncertainty arises from the choice of generator and tune chosen for the prediction of $q\bar{q} \to WW$ yields. The event yields predicted by each generator and tune agree well in the control regions but vary significantly in the signal region, even after the correction to the underlying event is applied. As no preference for any particular model could be deduced from the data, the background yield from the $q\bar{q} \to WW$ process was estimated as the midpoint of the highest and lowest value of the various predictions and the envelope of all predictions was taken as ± 1 standard deviation uncertainty boundary, amounting to 7 % for events selected with $n_{\mathrm{trk}} = 0$ and less than 1 % for events selected with $1 \leqslant n_{\mathrm{trk}} \leqslant 4$. The uncertainties in the total $q\bar{q} \to WW$ cross-section and the shape of the $p_{\mathrm{T}}(WW)$ distribution are taken from the MATRIX+RADISH prediction, amounting to 5–6 %. The total uncertainty on the $q\bar{q} \to WW$ yield is the dominant systematic uncertainty in the analysis due to the size of this background in the signal region.

Detailed studies of subdominant backgrounds were not necessary as uncertainties on these yields had a minimal effect on the overall background uncertainties. An uncertainty of 100 % was assigned to background yields from $W\gamma$ production, with 30 % assigned to the backgrounds from top quark production and WW production through vector boson scattering (VBS).

The uncertainty in the theoretical prediction of contributions of exclusive processes to CR2 and CR3 (the $1 \leqslant n_{\mathrm{trk}} \leqslant 4$ regions) were evaluated by comparing elastic, single- and double-dissociative production between LPAIR and MG5. These differences were found to be 20 % and 30 % for elastic and single-dissociative production. A 100 % uncertainty was assigned for double-dissociative production which translated to a 20 % uncertainty in the predicted yields of exclusive production in CR3 and CR4.

A summary of the dominant systematic uncertainties on the cross-section measurement is presented in Table 5.10.

5.8 Signal extraction and cross-section measurement

After applying all of the corrections described in previous sections, we could be confident that we had modelled the shapes of our distributions well within our assigned uncertainties. The absolute normalisations of signal and background yields due to higher order theoretical predictions were then obtained from a fit to data. In order to normalise the event yields and further validate modelling in the signal and control regions, a profile likelihood fit was performed in the SR and CR1–4. In the fit, systematic uncertainties were included as constrained nuisance parameters. The signal and dominant background normalisations were allowed to float as free parameters and were taken from a simultaneous fit to all five regions as no single region purely contains a single process, i.e. there is cross-contamination from multiple backgrounds and even the signal in the CRs. The fit procedure is detailed in this section before presenting the cross-section measurement and comparing to theoretical predictions. The author calculated the acceptance and correction factors for the cross-section measurement, discussed in Section 5.8.4.

5.8.1 Fit procedure

Event yields are governed by Poisson probabilities,

$$P(k;\lambda) = \frac{\lambda^k e^{-\lambda}}{k!}, \tag{5.9}$$

where $P(k;\lambda)$ is the probability of measuring a discrete random variable to have value k and λ is the expected value of k, i.e. the average value if the measurement were to be repeated many times. The $\gamma\gamma \to WW$ measurement is performed by maximising the product of Poisson probabilities in the SR and CRs for N_{sig} signal events and N_{bkgd} background events, in order to produce the observed number of events measured in data, N_{data}. Hence, we can rewrite Equation (5.9) as the likelihood of observing N_{data} in a single bin, assuming that $N_{\text{sig}} + N_{\text{bkgd}}$ is the true value one would obtain in that bin with enough statistics:

$$L(N_{\text{data}}; (N_{\text{sig}} + N_{\text{bkgd}})) = \frac{(N_{\text{sig}} + N_{\text{bkgd}})^{N_{\text{data}}} \times e^{-(N_{\text{sig}}+N_{\text{bkgd}})}}{(N_{\text{data}})!}. \tag{5.10}$$

We can reframe this thinking slightly to suggest that if we scan different values of $N_{\text{sig}} + N_{\text{bkgd}}$, given that we have already observed some value of N_{data}, we can find the value of $N_{\text{sig}} + N_{\text{bkgd}}$ that maximises this likelihood function. The aim of performing such a profile likelihood fit is to find appropriate normalisations of MC yields that maximise the likelihood across all regions of the analysis. Following this thinking, we write N_{sig} as

$$N_{\text{sig}} = \mu \times \sigma \times \mathcal{L} \times A \times C, \tag{5.11}$$

where σ is the predicted cross-section of the $\gamma\gamma \to WW$ process from MC, $\mathcal{L}$ is the integrated luminosity of the measurement, A and C are acceptance and correction factors that shall be described further in Section 5.8.4, and μ is the signal normalisation factor, or *signal strength*, which is allowed to float as a free parameter. We can write the total expected background yield, N_{bkgd}, as a sum of the individual background process yields in MC, N_B, scaled by a normalisation factor, β_B:

$$N_{\text{bkgd}} = \sum_B \beta_B N_B. \tag{5.12}$$

The normalisation is a free parameter in the fit for the $q\bar{q} \to WW$, DY and $\gamma\gamma \to ll$ background processes. All other β_B values are fixed at 1.

Note that Equation (5.10) makes no mention of systematic uncertainties. It is assumed that each systematic uncertainty, k, can alter the N_{sig} and N_{bkgd} yields by a fractional amount, S_k and B_k, resulting in a modification of the total yields by $(1 + \sum_k x_k S_k)$ for the signal and $(1 + \sum_k x_k B_k)$ for the background, where the summation is performed over all systematic variations. The $\{x_k\}$ are called nuisance parameters and are assumed to be normally distributed such that $x_k \sim N(0, 1)$. Each x_k is allowed to float in the fit but any large deviation from a mean value of 0 and variance of 1 would suggest some bias or mismodelling in the measurement. In order to account for systematics in the fit, we must modify Equations (5.11) and (5.12) so that they now become $N_{\text{sig}} = \mu\sigma\mathcal{L}AC(1 + \sum_k x_k S_k)$ and $N_{\text{bkgd}} = \sum_B \beta_B N_B(1 + \sum_k x_k B_k)$. Then we must also modify Equation (5.10) to use these new definitions of N_{sig} and N_{bkgd} and multiply the likelihood by a Normal probability density function term to represent the probability density of the corresponding nuisance parameter value:

$$L(N_{\text{data}};(N_{\text{sig}}+N_{\text{bkgd}})) = \frac{(N_{\text{sig}}+N_{\text{bkgd}})^{N_{\text{data}}} \times e^{-(N_{\text{sig}}+N_{\text{bkgd}})}}{(N_{\text{data}})!} \times \prod_k \frac{1}{\sqrt{2\pi}} e^{\frac{-x_k^2}{2}}. \tag{5.13}$$

The total likelihood across the SR and four CRs will be a product of the individual likelihoods in each region. It is customary to instead measure the natural log of the likelihood so that these products become summations and multiplicative factors become irrelevant when scanning to find a maximum. By taking the negative value of the log likelihood, we then derive a function that is minimised at the point of best fit:

$$-\ln L(N_{\{\text{SR,CRs}\}},\{x_k\}) = \sum_i -\ln\left(\frac{(N^i_{\text{sig}}+N^i_{\text{bkgd}})^{N^i_{\text{data}}} \times e^{-(N^i_{\text{sig}}+N^i_{\text{bkgd}})}}{(N^i_{\text{data}})!}\right) + \sum_k \frac{x_k^2}{2}, \tag{5.14}$$

where the sum over i represents a sum over each analysis region. The minimum of Equation (5.14) gives the best fit values of the signal and background normalisations, as well as for each nuisance parameter.

As the data was initially blinded in the signal region to avoid the introduction of human bias into the analysis method, the fit procedure was first validated in a CR-only fit and the resulting background normalisations were then input to a full fit using an Asimov dataset, where all observed quantities were set equal to their expected values. An expected measurement significance was extracted from the fit to Asimov data. Finally, a full fit was performed in all analysis regions, to obtain the measured significance with respect to the background-only hypothesis. The significance of rejecting the background hypothesis is calculated from a likelihood ratio between the nominal fit result and the fit with the signal strength fixed at $\mu = 0$ [201], in which uncertainties on signal modelling do not enter. The fit was performed using the HistFactory package [202].

5.8.2 Fit to control region data and Asimov model

The background normalisations were initially obtained by performing the fit in the CRs only, where signal contributions were expected to be small. The SR was excluded from the fit. Here, N_{data} is the data observed in the CRs and N_{sig} and N_{bkgd} are the signal and background yields expected from MC, with all corrections applied, with respective normalisation parameters. The normalisations of the $\gamma\gamma \to WW$, $q\bar{q} \to WW$, DY and $\gamma\gamma \to ll$ backgrounds were allowed to float as free parameters, whereas all other background normalisations and the signal strength were fixed at a value of one. The best-fit results of the floating background normalisations were

$$\begin{aligned} \beta_{q\bar{q}\to WW} &= 1.15^{+0.21}_{-0.23}, \\ \beta_{\gamma\gamma\to ll} &= 3.59 \pm 0.15, \\ \beta_{\text{DY}} &= 1.14 \pm 0.12. \end{aligned}$$

The value for $\beta_{\gamma\gamma\to ll}$ is significantly higher than one as no dissociative samples were used in CR4 so that the exclusive scale factor could be extracted from the fit as $S_{\text{excl}} = \beta_{\gamma\gamma\to ll}$. The value extracted from the fit is consistent with the value derived in Section 5.6.1. Most of the nuisance parameters were fitted very close to their nominal value of zero. A few larger shifts were driven by an over-prediction of the estimated background yields in CR1 but were all still consistent within 1σ of the nominal value.

The fit was then repeated in all analysis regions, constructing an Asimov dataset in the SR with $\mu = 1$, to obtain an expected significance of rejecting the background-only hypothesis and assign uncertainties on the signal strength. The background normalisation factors and nuisance parameters associated to the background estimate from non-prompt leptons were set to the values obtained in the CR-only fit. All nuisance parameters with shifts larger than ± 0.03 were fixed to the value from the CR-only fit. The signal strength resulting from the fit to Asimov data was

$$\mu = 0.99^{+0.13}_{-0.12}\,(\text{stat.})^{+0.17}_{-0.14}\,(\text{syst.})$$

and the background-only hypothesis was expected to be rejected with a significance of 6.7 standard deviations.

The largest source of uncertainty after the fit to Asimov data was the transferability systematic which accounts for the unknown applicability of the exclusive scale factor derived in $\gamma\gamma \to ll$ events to $\gamma\gamma \to WW$ events. This was followed by statistical uncertainties on the background yield which were mainly driven by the estimation of yields due to non-prompt leptons and theoretical uncertainties in the estimation of the $q\bar{q} \to WW$ backgrounds. If systematic uncertainties were excluded from the fit, the background-only hypothesis was expected to be rejected with a significance of 9.5 standard deviations.

5.8.3 Fit to data

After calculating the expected significance of the measurement, the fit was finally performed on unblinded data in all regions, including the SR. The background normalisation factors were found to be

$$\begin{aligned}\beta_{q\bar{q}\to WW} &= 1.21^{+0.19}_{-0.23},\\ \beta_{\gamma\gamma\to ll} &= 3.57 \pm 0.15,\\ \beta_{\mathrm{DY}} &= 1.16^{+0.22}_{-0.12},\end{aligned}$$

all of which are consistent with those found from the CR-only fit. The signal strength was measured to be

$$\mu = 1.33 \pm 0.14\,(\mathrm{stat.})\,^{+0.12}_{-0.17}\,(\mathrm{syst.}).$$

The background-only hypothesis was rejected with a significance of 8.4 standard deviations, making this measurement the first observation of exclusive $\gamma\gamma \to WW$ production. Larger nuisance parameter shifts were found compared to the fit to the Asimov dataset described in Section 5.8.2, related to the larger than expected signal strength parameter which introduces additional tension between data and the sum of signal and background in CR1. During the profiling of the nuisance parameters with Gaussian constraint, none of the fitted values were extreme enough to lead to negative Poisson parameters. This is consistent with the high observed significance of rejecting the background-only hypothesis.

It is worth discussing this relatively high observed signal strength further. The observed signal strength is not consistent with $\mu = 1$, within one standard deviation. A value of $\mu > 1$ means that the signal measurement in data is higher than expected from input models. Although the difference is not significant enough to cause concern (or excitement at possible hints of new physics), there is a discussion to be had on why the $\gamma\gamma \to WW$ yield could be higher than predicted from simulation. The applicability of measuring the exclusive scale factor, S_{excl}, in $\gamma\gamma \to ll$ events and transferring to $\gamma\gamma \to WW$ regions is discussed extensively in Ref. [198], largely in response to the $\gamma\gamma \to WW$ observation measurement presented in this chapter. The authors of Ref. [198] present expected ratios of elastic to dissociative events for both the $\gamma\gamma \to ll$ and $\gamma\gamma \to WW$ process, with the same selections used in the SR of this analysis and including proton soft-survival effects. They calculate that

$$S_{\text{excl}}^{\gamma\gamma\to WW}/S_{\text{excl}}^{\gamma\gamma\to ll} = 4.3/3.5 = 1.2.$$

Importantly, this value does not equal 1 as was assumed in this analysis. Also worth noting is that the calculated value of $S_{\text{excl}}^{\gamma\gamma\to ll}$ in Ref. [198] is consistent with the value measured in this analysis, within uncertainties. Taking the results from Ref. [198] at face value would mean that the expected $\gamma\gamma \to WW$ elastic signal yield should be scaled by an extra 20 % in the SR. This would bring the expected signal strength and measured signal strength in data into much better agreement and should be explored in future measurements of the $\gamma\gamma \to WW$ process.

The post-fit event yields and observed event yields in data are shown in Table 5.7, for the signal region and CR1–3. 307 events are observed in the signal region, with a post-fit $\gamma\gamma \to WW$ yield in this region of 174 ± 20. The post-fit event yields are shown differentially in different variables in Figure 5.22. The dashed grey lines represent boundaries between analysis regions, with the arrows indicating the signal region. Differential distributions in the signal region are explored further in Chapter 6.

	Signal region		Control regions	
n_{trk}	$n_{\mathrm{trk}} = 0$		$1 \leqslant n_{\mathrm{trk}} \leqslant 4$	
$p_{\mathrm{T}}^{\ell\ell}$	$> 30\,\mathrm{GeV}$	$< 30\,\mathrm{GeV}$	$> 30\,\mathrm{GeV}$	$< 30\,\mathrm{GeV}$
$\gamma\gamma \to WW$	174 ± 20	45 ± 6	95 ± 19	24 ± 5
$\gamma\gamma \to ll$	5.5 ± 0.3	39.6 ± 1.9	5.6 ± 1.2	32 ± 7
Drell–Yan	4.5 ± 0.9	280 ± 40	106 ± 19	4700 ± 400
$q\bar{q} \to WW$	101 ± 17	55 ± 10	1700 ± 270	970 ± 150
Non-prompt	14 ± 14	36 ± 35	220 ± 220	500 ± 400
Other backgrounds	7.1 ± 1.7	1.9 ± 0.4	311 ± 76	81 ± 15
Total	305 ± 18	459 ± 19	2460 ± 60	6320 ± 130
Data	307	449	2458	6332

Table 5.7: Summary of the data event yields and the predicted signal and background event yields in the signal region and control regions, as obtained after the fit. The uncertainties shown include statistical and systematic components. The numbers for $qq \to WW$ also contain a small contribution from gluon-induced WW and electroweak $WWjj$ production. The event yields for other backgrounds include contributions from WZ and ZZ diboson production, top-quark production and other gluon-induced processes [4].

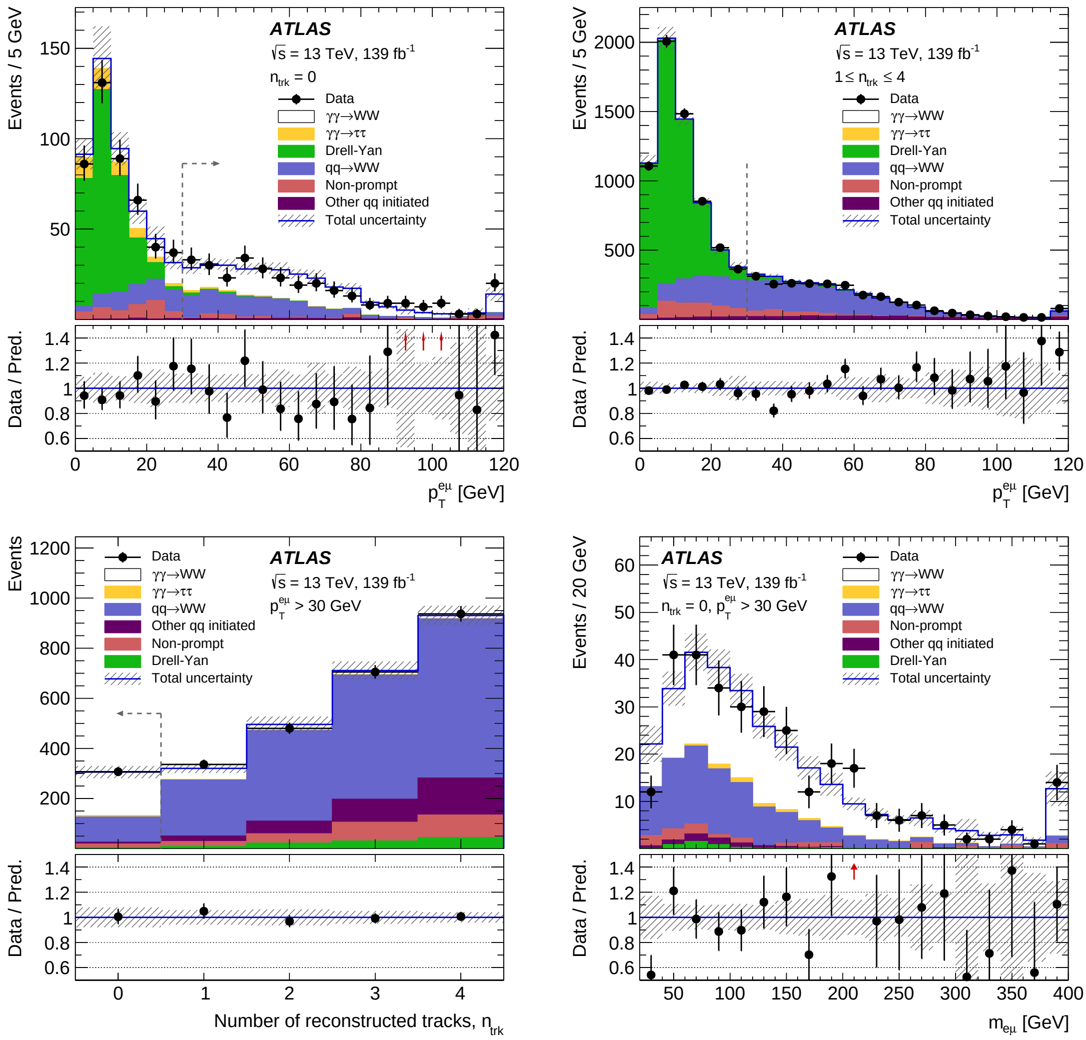

Figure 5.22: The post-fit event yields plotted differentially in (top row) dilepton p_{T} when requiring (left) $n_{\mathrm{trk}} = 0$ and (right) $1 \leqslant n_{\mathrm{trk}} \leqslant 4$, (bottom left) the number of tracks satisfying $|z_{\mathrm{trk}} - z_{\mathrm{vtx}}^{\ell\ell}| < 1\,\mathrm{mm}$ when $p_{\mathrm{T}}^{\ell\ell} > 30\,\mathrm{GeV}$ and (bottom right) dilepton mass when $n_{\mathrm{trk}} = 0$ and $p_{\mathrm{T}}^{\ell\ell} > 30\,\mathrm{GeV}$. The grey dashed lines indicate boundaries between analysis regions, with the arrows indicating the signal region. The $qq \to WW$ and "other qq initiated" components also contains contributions from smaller background processes. The total uncertainties are shown as hatched bands. The lower panel shows the ratio of the data to the prediction, with the total uncertainty displayed as a hatched band [4].

5.8.4 Cross-section measurement

Beyond observing the $\gamma\gamma \to WW$ process, we would also like to measure its cross-section which means that we must understand how the number of events that we observe in data is affected by inefficiencies in our detector and in our analysis methods. Once this is understood, we can correct for these effects to measure a *fiducial* cross-section which can be thought of as our best estimate of the true cross-section within our region of phase-space if it was measured by a perfect version of our real detector, i.e. the ATLAS detector without reconstruction inefficiencies or biases.

The selections defining the fiducial region of the analysis at particle-level are given in Table 5.8. Tau decays of the W-bosons were excluded from the fiducial region by identifying the $W \to \tau\nu_\tau$ vertex in the generator information. The fiducial cross-section, $\sigma_{\text{fid.}}$, can be defined as

$$\sigma_{\text{fid.}} = \frac{N_{\text{data}} - N_{\text{bkgd}}}{\mathcal{L} \cdot C_{WW} \cdot k_{\text{excl.}}}, \tag{5.15}$$

where N_{data} is the number of events observed in the data, N_{bkgd} is the post-fit number of background events predicted by simulation, $\mathcal{L}$ is the integrated luminosity of the dataset, and C_{WW} and $k_{\text{excl.}}$ are correction factors that shall be explained further.

The first correction factor, C_{WW}, is the ratio of the total number of events in signal simulation passing the reconstruction-level selections, $N_{\text{pass}}(\text{reco})$, to the

Selection requirement	Selection value
Primary charged particles	$\lvert\eta\rvert < 2.5$, $p_{\text{T}} > 500\,\text{MeV}$, $n_{\text{ch}}^{\text{UE}} = 0$
Single lepton transverse momentum	$p_{\text{T}} > 27\,\text{GeV}$ (leading), $p_{\text{T}} > 20\,\text{GeV}$ (subleading)
Dilepton selection	$n_{\text{lep}} = 2$, flavour $= e\mu$ or μe, $\text{c}_{\ell 1} \times \text{c}_{\ell 2} < 0$
Dilepton mass	$m_{\ell\ell} > 20\,\text{GeV}$
Dilepton transverse momentum	$p_{\text{T}}^{\ell\ell} > 30\,\text{GeV}$

Table 5.8: Summary of the selections that define the analysis fiducial region at particle-level. The $n_{\text{ch}} = 0$ cut does not consider the leptons themselves.

total number of events passing fiducial selections at particle-level, $N_{\text{pass}}(\text{truth})$:

$$C_{WW} = \frac{N_{\text{pass}}(\text{reco})}{N_{\text{pass}}(\text{truth})}, \tag{5.16}$$

where the selections at reconstruction-level are the same as those listed in Table 5.2. It is not possible to distinguish leptonic $W \to \tau\nu_\tau$ decays in data so no selection on tau decays is made at reconstruction-level. An additional correction factor is included to correct for known differences in the pileup distribution between data and simulation. This is quantified by the ratio of the exclusive efficiency in data to MC, $k_{\text{excl.}}$, defined as

$$k_{\text{excl.}} = \left(\frac{\varepsilon_{\text{excl.}}(\text{data})}{\varepsilon_{\text{excl.}}(\text{MC})}\right)_{\text{data-driven method}}. \tag{5.17}$$

As described in Section 5.4, only small residual differences remain in the exclusive efficiency between data and MC after corrections to pileup modelling are made. Therefore $k_{\text{excl.}}$ is expected to be close to 1.

Values of C_{WW} and $k_{\text{excl.}}$ were calculated from the nominal HERWIG7 elastic $\gamma\gamma \to WW$ simulation with beamspot width $\sigma_z = 35\,\text{mm}$ and all pileup corrections applied, including the beamspot rescaling. The results are presented here for the full signal selection but additional studies can be found in Appendix C which include comparing other dilepton channels, comparing the effects of the beamspot rescaling, and the impact of including dilepton final states originating from tau decays in the fiducial definition. It was shown that the contribution of $W \to \tau\nu_\tau$ decays to the $\gamma\gamma \to WW$ signal yield in the fiducial region is roughly 5 % at both particle- and reconstruction-level.

Different MC campaigns are used for the different data-taking years to account for differences in beam and detector conditions, which impact event reconstruction. The C_{WW} factor was calculated separately for each MC campaign in the signal simulation, corresponding to the periods of 2015–16, 2017 and 2018. The resulting values are given in Table 5.9. The main cause of the differences in C_{WW} between the campaigns is the change in the exclusive efficiency due to varying pileup distributions between the different data-taking years. Before the $n_{\text{trk}} = 0$ selection is applied, the correction factor is roughly 75 % across all years. The combined value for Run 2 was calculated by weighting each campaign by the integrated luminosity for that data-taking period, resulting in a value of

$$C_{WW} = 41.3 \pm 0.3 \text{ (stat.) } \%.$$

The correction factor includes the $n_{\text{trk}} = 0$ requirement, the efficiency of passing which was defined as the exclusive efficiency discussed in Section 5.4.1. There were some residual differences in the average exclusive efficiency calculated in data and MC, the origin of which is thought to be due to the discrepancy between the inelastic cross-section assumed in MC of 80 mb compared to the measured inelastic cross-section in data [193]. The ATLAS pileup task force recommend a scale factor of 1.03 ± 0.04 to be applied to μ measured in data to correct for this difference and for uncertainties on the instantaneous luminosity measured in data. In this analysis, this correction was included as an additional correction to the cross-section, in the form of $k_{\text{excl.}}$, which measures the discrepancy between the exclusive efficiency as measured in data and MC. A summary of the exclusive efficiency calculated in full Run 2 data and in MC was given in Table 5.5. The exclusive efficiency correction factor, $k_{\text{excl.}}$, is calculated from Equation (5.17) to be

$$k_{\text{excl.}} = 97.2 \pm 0.5 \text{ (stat.) } \%.$$

Note that this value is roughly equal to the inverse of the recommended scale factor, as expected.

Putting all of these components together, the measured $\gamma\gamma \to WW$ fiducial cross-section and its uncertainty are

$$\sigma^{\text{fid.}}_{\text{meas.}} = 3.13^{+0.32}_{-0.31} \text{ (stat.)}^{+0.29}_{-0.26} \text{ (syst.) fb}.$$

The measured cross-section has smaller relative uncertainties than the measured signal strength parameter. This is due to correlations between the parameters μ

2015–16	2017	2018
0.492 ± 0.005	0.388 ± 0.005	0.383 ± 0.004

Table 5.9: The correction factor for the cross-section calculation, $C_{WW} = N_{\text{pass}}(\text{reco})/N_{\text{pass}}(\text{truth})$, for signal events passing the fiducial selection. Values were calculated in signal MC with 35 mm beamspot, with all pileup corrections applied. The years represent the different MC campaigns and differing beam conditions that the simulation was designed to replicate.

Source	Impact [%]
Experimental	
Track reconstruction	1.1
Electron energy scale and resolution, and efficiency	0.4
Muon momentum scale and resolution, and efficiency	0.5
Misidentified leptons, systematic	1.5
Misidentified leptons, statistical	5.9
Other background, statistical	3.2
Modelling	
Pileup modelling	1.1
Underlying event modelling	1.4
Signal modelling	2.1
WW modelling	4.0
Other background modelling	1.7
Luminosity	1.7
Total	8.9

Table 5.10: Impact of different components of systematic uncertainty on the measured fiducial cross-section, without accounting for correlations. The impact of one source of systematic uncertainty is computed by first performing the fit with the corresponding nuisance parameter fixed to one standard deviation up or down from the value obtained in the nominal fit, then these up and down variations are symmetrised. The impacts of several sources of systematic uncertainty are added in quadrature for each component [4].

and $\beta_{\gamma\gamma\to ll}$ which are by construction negative, such that uncertainties between the theoretical prediction including dissociative contributions and the signal strength parameter partially cancel. The correlation coefficient obtained in the fit is $\rho(\mu, \beta_{\gamma\gamma\to ll}) = -0.20$. A breakdown of the largest contributions to the systematic uncertainty on the measured cross-section is given in Table 5.10

In principle, we could attempt to measure the inclusive cross-section of the $\gamma\gamma \to WW$ process by also including the efficiency of our analysis selections in sampling the entire phase-space. The acceptance of the signal, A_{WW}, is defined as the ratio of events passing the particle-level cuts in Table 5.8, $N_{\text{pass}}(\text{truth})$, to the total number of generated events, $N_{\text{tot.}}(\text{truth})$, including any filter efficiencies from selections applied at the generation stage:

$$A_{WW} = \frac{N_{\text{pass}}(\text{truth})}{N_{\text{tot.}}(\text{truth})} \times \text{filter eff.} \tag{5.18}$$

For example, the signal MC is produced with a filter to only select leptonic decays of the W-bosons and this selection is recorded as part of the filter efficiency. The

Sample	Cross-section prediction
H7 EE	0.6567 fb
MG5 EE	0.7823 fb
MG5 SD	2.3775 fb
MG5 DD	0.0794 fb

Table 5.11: Overview of the theoretical fiducial cross-sections, without any consideration of proton soft survival factors. Only simulation of elastic $\gamma\gamma \to WW$ events is possible with HERWIG7, whereas MADGRAPH+PYTHIA8 is able to simulate elastic and dissociative events. The MG5+Py8 results in this table use the CT14qed PDF set.

filter efficiency of the HERWIG7 elastic $\gamma\gamma \to WW$ signal sample is 0.30838. The acceptance of the fiducial selection in signal simulation was calculated to be

$$A_{WW} = 7.24 \pm 0.05 \text{ (stat.) } \%.$$

Quoting the signal acceptance allows extrapolation into an inclusive phase-space, if desired, for comparison between measurements with different fiducial volumes. However, care must be taken in performing such extrapolations as we cannot claim to be certain that the measurements we make within our fiducial volume are applicable over the full phase-space. For this reason, the result is quoted as a fiducial cross-section only.

5.8.5 Comparison of results to theoretical predictions

The theoretical prediction of the fiducial cross-section can be calculated at generator level as $A_{WW} \cdot \sigma_{WW}$, where σ_{WW} is the theoretical prediction of the inclusive $\gamma\gamma \to WW$ cross-section. Predictions from the nominal (H7) and alternative (MG5) signal generators are determined with the fiducial selection applied. These predictions are given in Table 5.11. Only the prediction for the elastic contribution to the cross-section was possible with HERWIG7. In order to obtain a theoretical prediction of the H7 total cross-section, including the dissociative production modes, the elastic cross-section is corrected with the exclusive scale factor extracted from the fit:

$$\sigma_{\text{H7,tot.}} = \sigma_{\text{H7,EE}} \times \beta_{\gamma\gamma\to ll}. \tag{5.19}$$

This has the advantage that experimental correlations between the corrected theoretical cross-section and the measured cross-section can be extracted from the fit as well, reducing systematic uncertainties. Another advantage of this approach is that the data-driven approach to measuring the exclusive scale factor automatically accounts for proton soft survival probabilities, which are poorly constrained theoretically. The theoretical fiducial cross-section prediction from HERWIG7, including dissociative contributions and soft survival probabilities is therefore

$$\sigma^{\text{fid.(H7)}}_{\text{el.+diss.}} = \sigma^{\text{fid.(H7)}}_{\text{EE}} \times (3.59 \pm 0.15\,(\text{exp.}) \pm 0.39\,(\text{trans.})) = 2.34 \pm 0.27\,\text{fb},$$

which is consistent with the measured cross-section within two standard deviations.

The theoretical expectation for the fiducial cross-section as evaluated with MG5 and taking into account all production modes is

$$\sigma^{\text{fid.(MG5)}}_{\text{el.+diss.}} = 3.2 \pm 0.6\,(\text{scale}) \pm 0.3\,(\text{PDF}),$$

when using the CT14qed PDF set. This prediction does not include soft survival factor considerations. The probability of an event to still pass the exclusivity requirement of $n_{\text{trk}} = 0$ despite proton rescattering effects has been calculated in phenomenological studies to be in the range of 0.64 [203], 0.65 [204] up to 0.82 [205]. Ref. [205] uses a two-channel eikonal model that also accounts for the helicity structure of the hard scattering process and is taken to be the most physically motivated prediction. Including a soft survival factor of 0.82 into the MG5+Py8 cross-section predictions results in the theoretical fiducial $\gamma\gamma \to WW$ cross-section is predicted to be

$$\sigma^{\text{fid.(MG5)}}_{\text{el.+diss.}} = 2.8 \pm 0.8\,\text{fb},$$

which is consistent with the measured cross-section within one standard deviation. The cross-section prediction using MG5 is slightly higher that the cross-section predicted by H7, which could be due to the different scale choices or the choice of soft survival factor. The scale uncertainty has been determined for the total cross-section and is symmetrised. The impact of PDF uncertainty is evaluated using the CT14qed error PDF sets.

A standalone theory prediction for the fiducial cross-section is computed with MG5+Py8 using the appropriate elastic or inelastic MMHT2015qed PDF sets [206] for each of the contributions by applying the fiducial requirements to all photon-induced contributions, which yields

$$\sigma^{\text{fid.(MG5)}}_{\text{el.+diss.}} = 4.3 \pm 1.0\,(\text{scale}) \pm 0.1\,(\text{PDF}).$$

The contributions to this cross-section prediction from elastic and single-dissociative production are 16 % and 81 %, respectively. Double-dissociative production contributes only 3 %. When applying the soft survival factor of 0.82, the predicted fiducial cross-section is

$$\sigma^{\text{fid.(MG5)}}_{\text{el.+diss.}} = 3.5 \pm 1.0\,\text{fb},$$

which also agrees with the measured cross-section of 3.13 ± 0.4 fb within one standard deviation.

5.9 Conclusions and outlook

This chapter detailed the observation of the photon-induced $\gamma\gamma \to WW$ process in proton–proton collisions at $\sqrt{s} = 13$ TeV recorded with the ATLAS detector at the LHC, corresponding to an integrated luminosity of 139 fb^{-1}. Events with leptonic W-boson decays into $e^{\pm}\nu\mu^{\mp}\nu$ final states were selected by requiring that no tracks except those of the two charged leptons are associated with the production vertex. Several data-driven corrections to charged particle and reconstructed track multiplicity for both pileup interactions and the underlying event were derived in order to perform the measurement. The background-only hypothesis was rejected with a significance of 8.4 standard deviations whereas 6.7 standard deviations was expected. This measurement constitutes the first observation of photon-induced WW production, a process for which only evidence was previously reported. The signal strength and the cross-section for the sum of elastic and dissociative production mechanisms were measured. The cross-section for the $pp(\gamma\gamma) \to p^{(*)}W^+W^-p^{(*)}$ process in the decay channel $W^+W^- \to e^{\pm}\nu\mu^{\mp}\nu$ in a fiducial phase-space close to the experimental acceptance was measured to be

$$\sigma^{\text{fid.}}_{\text{meas.}} = 3.13^{+0.32}_{-0.31}\ (\text{stat.})^{+0.29}_{-0.26}\ (\text{syst.})\ \text{fb}.$$

This result is in agreement with the theoretical predictions within uncertainties.

The observation of the rare $\gamma\gamma \to WW$ process with the large integrated luminosity of the Run 2 dataset opens up avenues for differential measurements of kinematic spectra, with the aim of probing distributions that could be sensitive to new physics parametrised by an EFT. A differential measurement of the $\gamma\gamma \to WW$ process with the ATLAS detector is ongoing and is the subject of the next chapter.

Chapter 6

Differential measurement of the $\gamma\gamma \to \boldsymbol{WW}$ process

"Things inexorably get worse."

– Brian Cox

6.1 Motivation and aim

Now that an observation of the $\gamma\gamma \to WW$ process has been made, it is desirable to measure the process differentially, both to check the modelling of kinematic distributions and to constrain new physics effects. It is expected that the cross-section of the $\gamma\gamma \to WW$ process induced by anomalous couplings increases with respect to the SM cross-section with the invariant mass of the diboson system and with the transverse momentum of each W boson [207]. A differential measurement of the fully-leptonic $\gamma\gamma \to WW$ process is being carried out by the author within ATLAS, with the ultimate aim of measuring fiducial differential cross-sections that have been unfolded to remove detector effects.

This chapter begins by explaining the procedure of unfolding and presents preliminary unfolded distributions. The unfolding is then optimised, focusing on the dilepton invariant mass and transverse momentum distributions.

The data and simulated event samples used in this chapter, as well as the fiducial event selection at both particle- and reconstruction-level, are identical to those presented in Chapter 5. Unfolded data are compared to elastic particle-level distributions predicted by Herwig7, scaled by the exclusive scale factor discussed in Section 5.6.1 and the signal strength measured in Section 5.8.

All work in this chapter is the author's own. As the analysis is a work-in-progress, the analysis chain is not yet finalised and likely to change with respect to the observation analysis. Results are presented as a current status of the measurement but should not be taken as final and many optimisations are ongoing. The chapter concludes with a summary of the studies so far and outstanding analysis tasks that are foreseen.

6.1.1 Distributions to unfold

Several lepton kinematic distributions are unfolded: the dilepton mass, $m_{\ell\ell}$, dilepton p_{T}, $p_{\mathrm{T}}^{\ell\ell}$, the leading-lepton p_{T}, $p_{\mathrm{T}}^{\mathrm{lead}}$, subleading-lepton p_{T}, $p_{\mathrm{T}}^{\mathrm{sublead}}$, the azimuthal angular separation of the leptons, $\Delta\Phi_{\ell\ell}$, and dilepton rapidity, $y_{\ell\ell}$. The $m_{\ell\ell}$, $p_{\mathrm{T}}^{\ell\ell}$, $p_{\mathrm{T}}^{\mathrm{lead}}$, $p_{\mathrm{T}}^{\mathrm{sublead}}$ distributions are expected to be most sensitive to any new physics in the high-end tails of the distributions [208]. The $\Delta\Phi_{\ell\ell}$ and $y_{\ell\ell}$ distributions are measures of event topology, which may allow us to probe kinematic differences between elastic and dissociative contributions to the signal.

The bin boundaries of every kinematic distribution were optimised to have at least 20 events per bin whilst also maintaining sensible bin boundaries within the detector resolution. A requirement of 20 events per bin was so that statistical uncertainties can be assumed to be Gaussian, leading to a maximum statistical uncertainty per bin of 22 %. Under/overflow was added into the edge bins unless the boundary was due to an analysis selection, in which case the under/overflow flow is stored separately to calculate efficiencies related to the unfolding procedure. The binning for each variable was optimised to be:

$$\begin{aligned}
m_{\ell\ell} &= [20,\ 70,\ 100,\ 140,\ 200,\ 400]\ \text{GeV};\\
p_{\mathrm{T}}^{\ell\ell} &= [30,\ 40,\ 50,\ 60,\ 70,\ 85,\ 110,\ 150]\ \text{GeV};\\
p_{\mathrm{T}}^{\text{lead}} &= [27,\ 45,\ 60,\ 75,\ 100,\ 300]\ \text{GeV};\\
p_{\mathrm{T}}^{\text{sublead}} &= [20,\ 25,\ 30,\ 40,\ 60,\ 300]\ \text{GeV};\\
\Delta\Phi_{\ell\ell} &= [0,\ 3\pi/8,\ 5\pi/8,\ 6\pi/8,\ 7\pi/8,\ \pi]\ \text{rad};\\
y_{\ell\ell} &= [-5.0,\ -1.0,\ -0.5,\ 0.0,\ 0.5,\ 1.0,\ 5.0].
\end{aligned}$$

As the analysis is somewhat statistically limited, the bin widths are larger than the detector resolution for all variables and therefore we expect there to be very little migration of events between particle-level and reconstruction-level bins.

The event yields as a function of these observables are shown in Figure 6.1. Each figure shows the data measured with all signal selections applied and compared to the post-fit MC distributions in the signal region. The signal MC prediction is obtained by scaling the elastic Herwig7 prediction by the exclusive scale factor of 3.59 (Section 5.6.1), which accounts for the missing SD and DD contributions, and includes the observed signal strength of 1.33 (Section 5.8). The total uncertainty shown on the simulation prediction is the combination of all systematic uncertainties, including statistical uncertainties in the background predictions. The uncertainty on the data is statistical. Generally, the data and simulation agree well within uncertainties. One exception is the $p_{\mathrm{T}}^{\ell\ell}$ distribution, where there is a relative data/MC discrepancy of 70 % in the $85 < p_{\mathrm{T}}^{\ell\ell} < 110$ GeV bin. The leading lepton p_{T} also appears to be slightly harder in data compared to simulation and there are data/MC discrepancies of roughly 25 % in the highest $\Delta\Phi_{\ell\ell}$ bins.

Now that we are measuring differentially in kinematic spectra, we do not necessarily predict that our signal modelling will be well-motivated, as the exclusive scale factor applied to the elastic prediction does not account for any

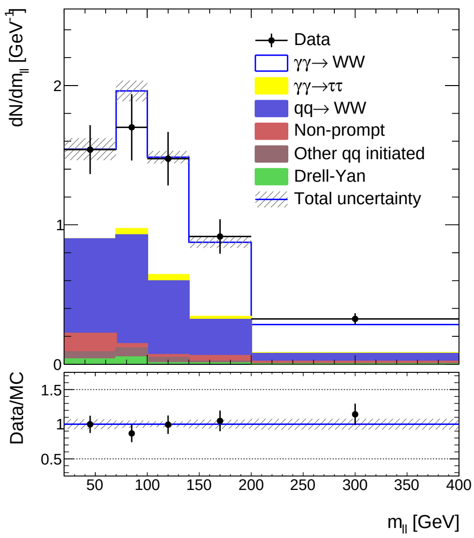

(a) $m_{\ell\ell}$

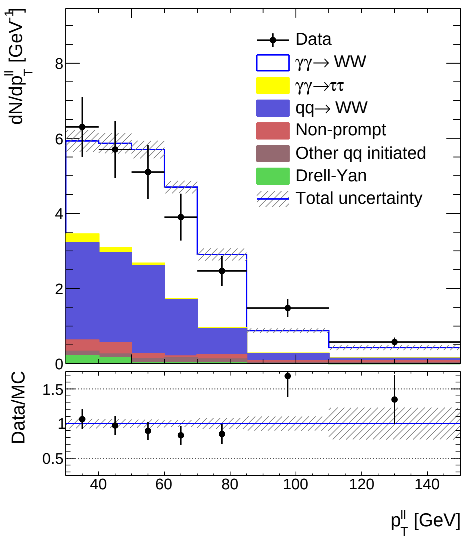

(b) $p_{\mathrm{T}}^{\ell\ell}$

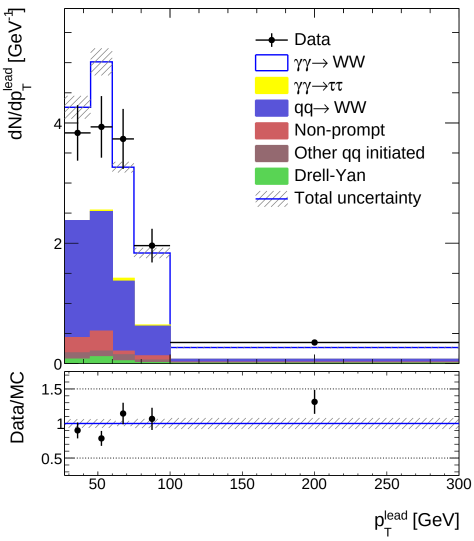

(c) leading lepton p_{T}

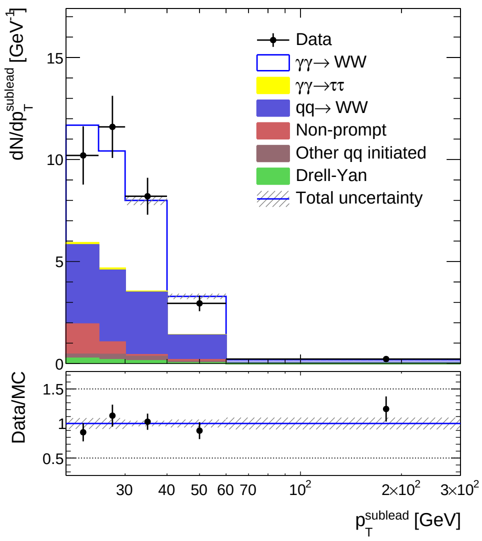

(d) subleading lepton p_{T}

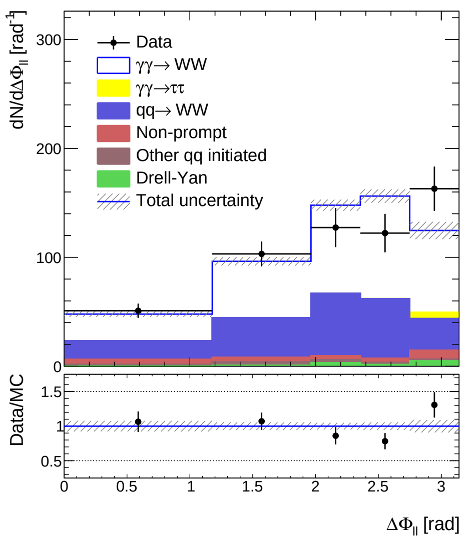

(e) $\Delta\Phi_{\ell\ell}$

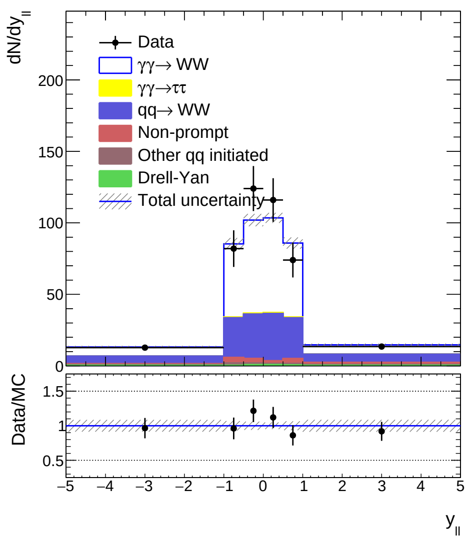

(f) $y_{\ell\ell}$

Figure 6.1: Differential event-yield distributions in the $\gamma\gamma \to WW$ signal region.

kinematic differences between elastic and dissociative $\gamma\gamma \rightarrow WW$ events. Any differences are most likely to manifest in the $p_{\mathrm{T}}^{\mathrm{lead}}$ spectrum (as single-dissociative production is more likely to be boosted with respect to elastic production), as well as topological distributions such as $\Delta\Phi_{\ell\ell}$.

6.2 Unfolding procedure

It is desirable to make any measurement as reproducible and reinterpretable as possible. The process of *unfolding* aims to remove detector effects from a measurement, giving a best estimate of what the true underlying physics is free from reconstruction inefficiencies and detector resolution.

In its essence, an unfolding procedure aims to create a map between two different data spaces: the reconstructed data and the "true" particle-level data. In measuring a process in data, we are applying a *folding* map which transforms the true physics to the observables that we measure. We are usually dealing with histograms of data where we have measured an event yield in a series of bins with finite width. In this case, the folding procedure is a measure of how detector effects migrate observable values between particle-level bins, t_j, and reconstruction-level (reco-level) bins, r_i. We can represent this folding by a matrix, R_{ij}, that encodes the probability, P, of mapping an event in each bin at particle-level to a bin at reconstruction-level.

$$P(r_i) = \sum_j R_{ij} P(t_j). \tag{6.1}$$

The folding matrix is commonly called the *response matrix*, as it is a measure of the detector response.

If we are starting with a measured distribution in data, we wish to create an *unfolding* map to transform the reconstructed distribution to particle-level. At its simplest, the unfolding map is an inversion of R_{ij} and we can represent the unfolding procedure as

$$P(t_j) = \sum_i R_{ji}^{-1} P(r_i) = \sum_i U_{ji} P(r_i), \tag{6.2}$$

where U_{ji} is the unfolding matrix. However, matrix inversion is not always possible and even when it is, statistical fluctuations can lead to large instabilities in unfolded results. In these cases, other unfolding methods must be used.

A commonly used unfolding scheme in high energy physics is the iterative unfolding method proposed by D'Agostini [209, 210]. The method makes use of Bayes' theorem, defining the conditional probability that an event originated in particle-level bin t_j, given we have observed the event to be reconstructed in bin r_i, as

$$P(t_j|r_i) = \frac{P(r_i|t_j) \cdot P(t_j)}{\sum_k P(r_i|t_k) \cdot P(t_k)}. \tag{6.3}$$

Here, $P(t_j|r_i)$ form the elements of the unfolding matrix to apply to our observed data and $P(r_i|t_j)$ form the elements of the response matrix. As we do not have access to the event yields of t_j in data, the elements of the response matrix must be calculated from MC simulation. The $P(t_j|r_i)$ rely on a prior, $P(t_j)$, which we also cannot know with certainty. In principle, the prior we choose can be arbitrary but we (usually) have a good estimate of what our prior distribution looks like and can also estimate this from MC simulation, representing our theoretical knowledge of the underlying physics. The assumption of the specific prior can introduce a bias into the unfolding because the prior we assume may be incorrect.

We can reduce the bias of our assumption on the initial prior by iterating the procedure, using each calculated posterior as the input prior for the next iteration, such that

$$P(t_j|r_i)_n = \frac{P(r_i|t_j) \cdot P(t_j)_{n-1}}{\sum_k P(r_i|t_k) \cdot P(t_k)_{n-1}}, \tag{6.4}$$

where n represents the iteration number and $P(t_j)_{n-1}$ is equal to $P(t_j|r_i)_{n-2}$ for $n \geqslant 2$. The elements $P(r_i|t_j)$ are not updated with the iterations. Each successive iteration reduces the reliance on the initial prior and gives more weight to the observed data. Iterations can be carried out until convergence is observed. However, as with a simple matrix inversion, problems arise when we have statistical fluctuations in measured spectra. The application of the unfolding matrix aims to restore the original underlying spectrum. If migrations are properly modelled, we obtain the correct values if the input data match the expected value of the reconstructed bin contents. However, the unfolding matrix is applied to one instance of a measurement, i.e. one set of data, not its expectation value. As a consequence, the unfolding procedure "assumes" that fluctuations in the data have their origin in real physical structures which are diluted by detector effects (and not of statistical origin). As a result, statistical fluctuations are propagated into the unfolded spectrum. This problem can be reduced by regularisation by, for example, limiting the number of iterations, changing bin widths to mitigate large fluctuations with large statistical uncertainty, or smoothing intermediate distributions in iterative methods. However, in increasing bin sizes, we lose information about our data and in limiting iterations or smoothing intermediate distributions, we introduce bias.

To obtain event yields instead of probability densities, we must also account for events that are within the fiducial volume at either particle- or reconstruction-level but not both. The unfolded event yield in each bin of the particle-level spectrum is given by

$$N(t_j) = \frac{1}{\varepsilon_j} \sum_i f_i \cdot N(r_i)_{\mathrm{obs}} \cdot P(t_j|r_i)_{n_f}, \tag{6.5}$$

where ε_j is the reconstruction efficiency of an event that originates in t_j, representing the loss of events in the fiducial volume at particle-level that do not pass reconstruction selections. The fiducial fraction, f_i, of each reconstructed bin represents the fraction of events in bin r_i that have a corresponding particle-level event within the fiducial volume. The $P(t_j|r_i)_{n_f}$ are calculated using Equation (6.4), where the final iteration number, n_f, is usually chosen so as to minimise the combination of the bias introduced by the assumed prior and the statistical uncertainty on the unfolded result.

The `RooUnfold` [211] implementation of D'Agostini's iterative method is used to correct the $\gamma\gamma \to WW$ differential distributions for detector inefficiency and resolution.

6.3 Unfolding inputs

The inputs to the unfolding procedure are:

- **Response matrix:** The elements of the detector response, as calculated in signal simulation, are

$$R_{ij} = P(r_i|t_j)_{\mathrm{MC}}, \tag{6.6}$$

  representing the probability of an event being reconstructed in bin r_i given that it originated in particle-level bin t_j. The response matrix only includes events that are within the fiducial volume at both particle- and reconstruction-level.

- **Fiducial fraction:** The fiducial fraction per bin, $f^i_{\text{fid.}}$, quantifies the fraction of events that are reconstructed in r_i but fall outside of the fiducial volume at particle-level. The fiducial fraction in bin r_i is given by

$$f^i_{\text{fid.}} = \frac{\sum_j R_{ij}}{P(r_i)}. \tag{6.7}$$

- **Reconstruction efficiency:** The reconstruction efficiency per bin, $\varepsilon^j_{\text{reco}}$, is the probability that the event passes reconstruction-level fiducial selections, given that it originated in particle-level bin t_j. It is given by

$$\varepsilon^j_{\text{reco}} = \frac{\sum_i R_{ij}}{P(t_j)}. \tag{6.8}$$

- **Assumed prior:** The initial prior is taken from the MC theoretical prediction,

$$P(t_j)_0 = P(t_j)_{\text{MC}}. \tag{6.9}$$

The response matrix is created using the nominal HERWIG7 elastic $\gamma\gamma \to WW$ MC, with all signal selections given in the SR column of Table 5.3 applied at reco-level and all selections given in Table 5.8 applied at particle-level. As the events in the response matrix pass both the particle- and reco-level selections, they are weighted with all generator and reconstruction weights. All standard ATLAS weights, including generator weights, lepton reconstruction and isolation efficiencies, and the PRW, are applied. The beamspot rescaling is also applied, as described in Section 5.4.3. The additional pileup and underlying event corrections have not been applied yet but, as they are designed to improve the modelling of charged particle multiplicity, they are not expected to have a large impact on the shape of any distribution presented in this chapter. They will however impact the overall normalisation of event yields.

The response matrices are shown in Figure 6.2 for each observable, with the binning as given in Section 6.1.1. As expected, every response matrix is highly diagonal, with $P(r_i|t_i) > 80\,\%$ across all bins of each distribution, and in most cases much higher at over $90\,\%$. The lowest values of $P(r_i|t_i)$ are seen in distributions that are steeply falling, such as the $p_{\text{T}}^{\text{lead}}$ and $p_{\text{T}}^{\text{sublead}}$ distributions. Migration between bins is also higher around sharp peaks of distributions, due to smaller binning in regions where there are a higher number of reconstructed events. The $P(r_i|t_j)$ are generally lower in the $p_{\text{T}}^{\ell\ell}$ distribution due to the finer

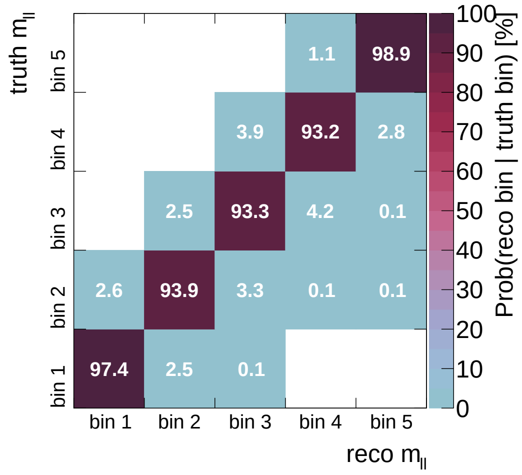

(a) $m_{\ell\ell}$

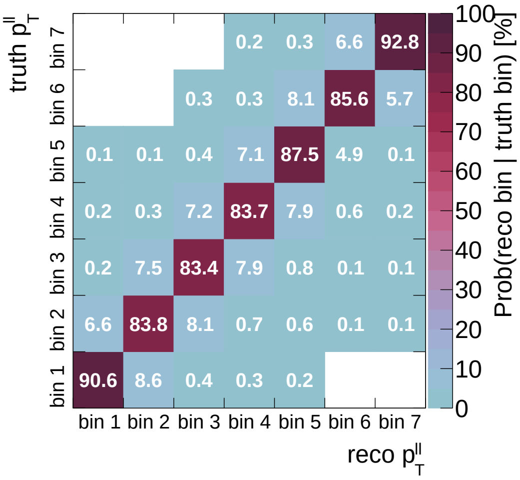

(b) $p_{\mathrm{T}}^{\ell\ell}$

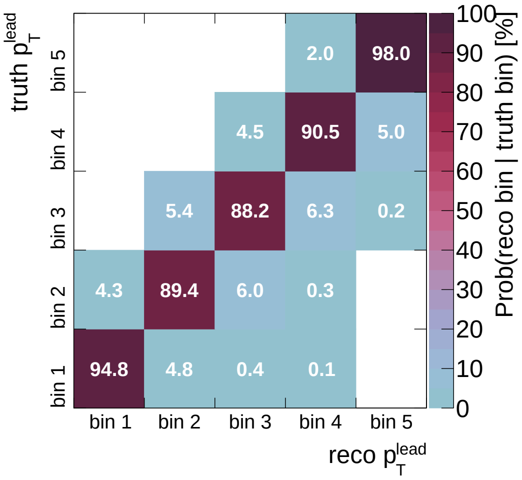

(c) leading lepton p_{T}

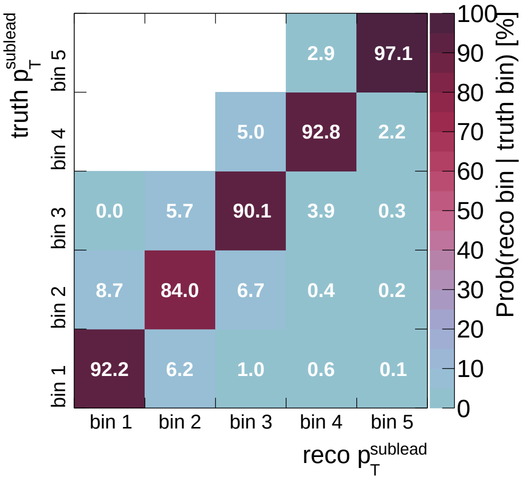

(d) subleading lepton p_{T}

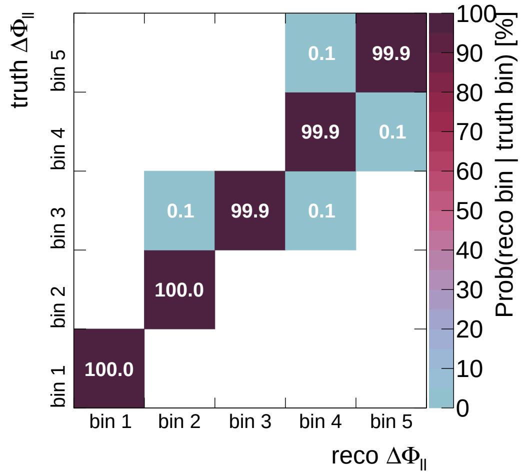

(e) $\Delta\Phi_{\ell\ell}$

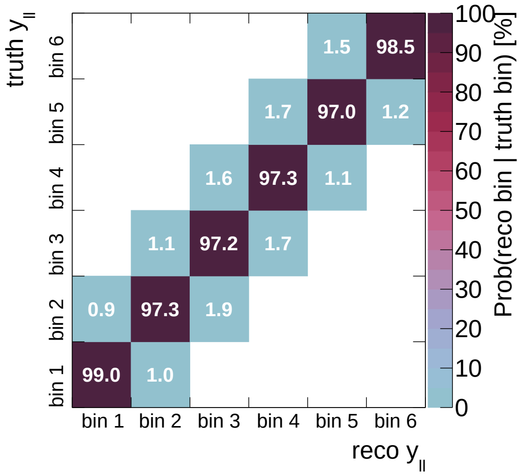

(f) $y_{\ell\ell}$

Figure 6.2: Response matrices, R_{ij}, showing the probability of reconstructing an event in reco-level bin r_i, given that it originated in particle-level (truth-level) bin t_j. All response matrices are highly diagonal due to the large bin widths relative to detector resolutions. All bins are given equal width in these figures to aid readability.

binning in this distribution. As the $P(r_i|t_i)$ are so high, finer binning may be ultimately preferable in all distributions as it allows the shapes of distributions to be better analysed.

Measures of the fiducial fraction and reconstruction efficiency are shown in Figure 6.3. The reconstruction efficiency is generally between 30–40 % across each distribution. This is due to a combination of trigger efficiencies, lepton identification and isolation efficiencies, and the exclusive efficiency of the $n_{\mathrm{trk}} = 0$ requirement. The single-muon triggers have an efficiency of 68 % in the barrel region and 85 % in the end-caps [200]. The single-electron triggers have an efficiency of 75 % for electrons with transverse energy of 31 GeV and 96 % for electrons with 60 GeV [199]. The leading lepton peaks between $60 < p_{\mathrm{T}} < 75$ GeV whereas the subleading lepton peaks much lower that this, between $20 < p_{\mathrm{T}} < 25$ GeV. Electron identification efficiencies range between 90–95 % [191] and muon identification efficiencies range between 96–98 % [165]. The exclusive efficiency as measured in signal MC is 54 %. Putting all of these components together, the reconstruction efficiency is estimated to be roughly $0.7 \times 0.9 \times 0.9 \times 0.96 \times 0.54 = 0.3$, so the observed efficiencies in the unfolding seem reasonable. The efficiency is not flat across $\Delta\Phi_{\ell\ell}$ due to positive correlation between this variable, $m_{\ell\ell}$ and individual lepton p_{T}.

The fiducial fraction is very high for all distributions, meaning that the number of reco-level candidates without a corresponding particle-level candidate within the fiducial volume is small, typically ~ 5 %. The fiducial fraction is often lower at reconstruction boundaries in the kinematic selection, e.g. $p_{\mathrm{T}}^{\mathrm{sublead}} \approx 20$ GeV. This is due to events being smeared into the fiducial region at reco-level that did not pass the particle-level fiducial selections.

Another useful quantity is the purity of each bin, which can be used to optimise bin widths. The purity of a given reco-level bin is the probability that the event originated from the same bin at particle-level, i.e. $P(t_i|r_i)$, calculated in the MC simulation. This differs from the diagonals of the response matrix which are equal to $P(r_i|t_i)$. The purity is also shown in Figure 6.3 and is high for all distributions, with $P(t_i|r_i) > 80$ %. As the response matrices are highly diagonal, $P(t_i|r_i) \approx P(r_i|t_i)$.

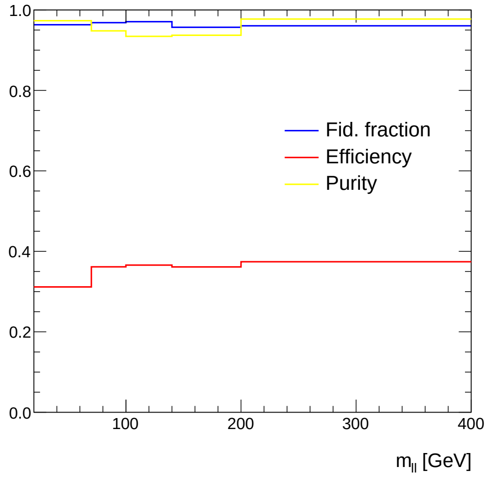

(a) $m_{\ell\ell}$

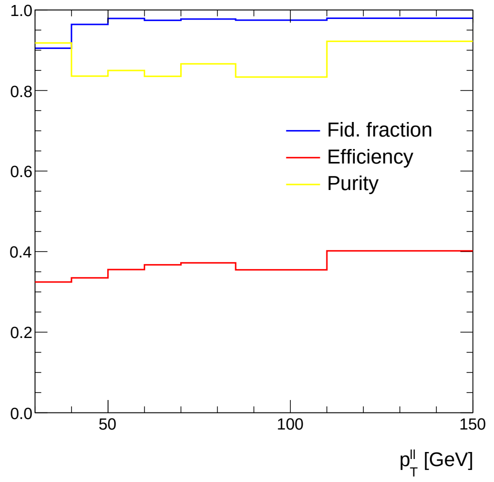

(b) $p_{\mathrm{T}}^{\ell\ell}$

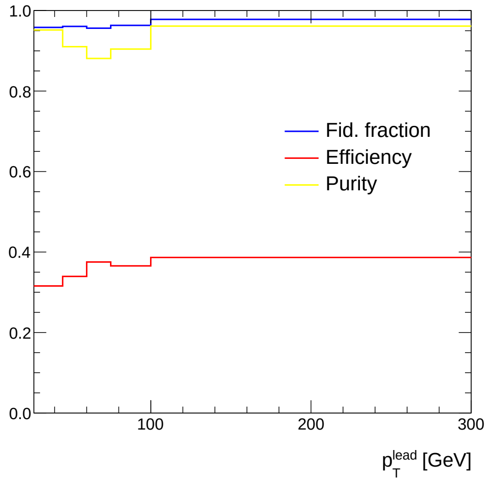

(c) leading lepton p_{T}

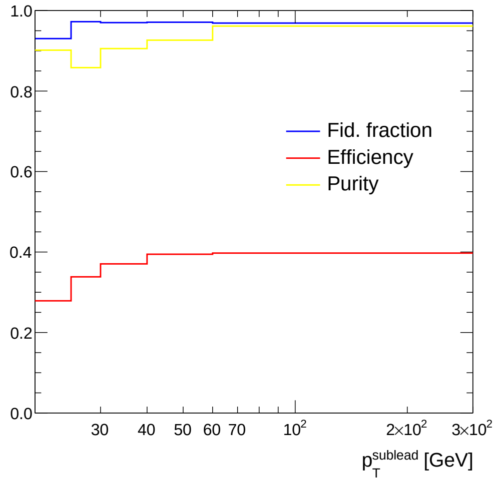

(d) subleading lepton p_{T}

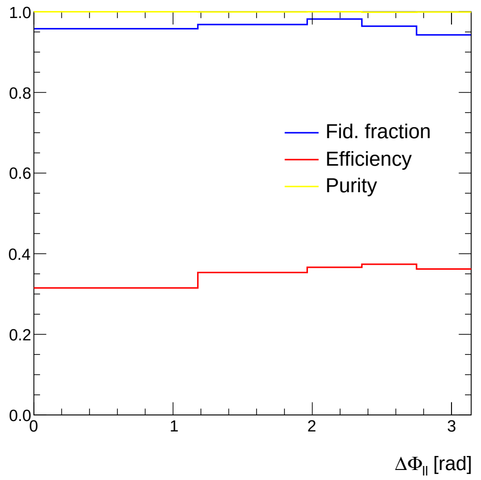

(e) $\Delta\Phi_{\ell\ell}$

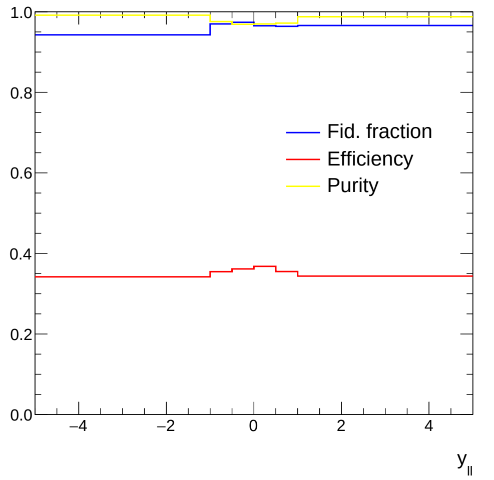

(f) $y_{\ell\ell}$

Figure 6.3: Purity, efficiency and fiducial fraction.

6.3.1 Technical closure test

A technical closure test is implemented to ensure that the unfolding procedure has been set up properly. The reconstruction-level distributions predicted by the signal MC are unfolded to particle-level using the nominal response matrix that was created using the same MC. The unfolding is performed using one iteration of D'Agostini's method. By definition, the unfolded distribution should close perfectly with the particle-level distribution, for all observables. If this is not the case, there is an error in the unfolding procedure. The technical closure test closes perfectly for all distributions, shown for $m_{\ell\ell}$ and $p_{\mathrm{T}}^{\ell\ell}$ in Figure 6.4. Remaining technical closure tests can be found in Appendix D.

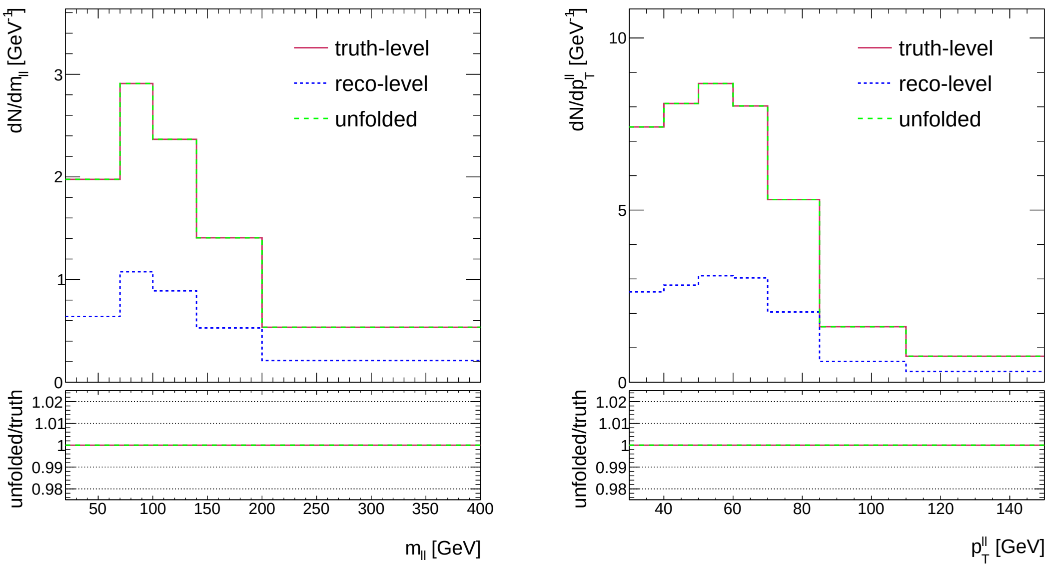

(a) $m_{\ell\ell}$ (b) $p_{\mathrm{T}}^{\ell\ell}$

Figure 6.4: Technical closure test of the unfolding procedure for (a) the $m_{\ell\ell}$ distribution and (b) the $p_{\mathrm{T}}^{\ell\ell}$ distribution. Truth-level refers to particle-level and reco-level refers to reconstruction-level observables.

6.4 Unfolded results

The response matrices in Figure 6.2 were then used in combination with Equation (6.4) to unfold the data distributions shown in Figure 6.1, with the total background predicted by simulation subtracted. Results are presented in Figure 6.5 for each distribution, unfolded with four iterations of D'Agostini's iterative method. The number of iterations is optimised later in Section 6.6.2.

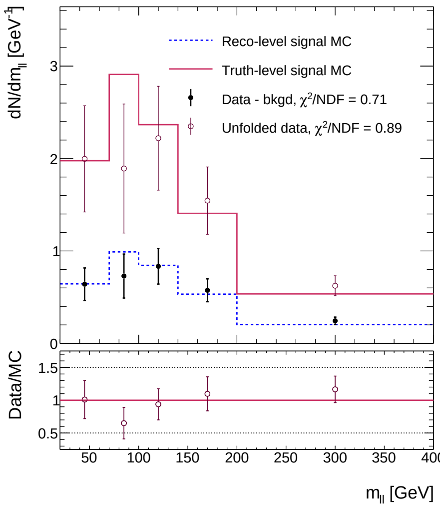

(a) $m_{\ell\ell}$

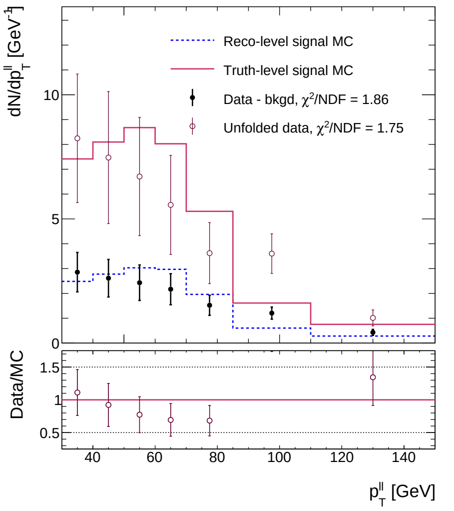

(b) $p_{\mathrm{T}}^{\ell\ell}$

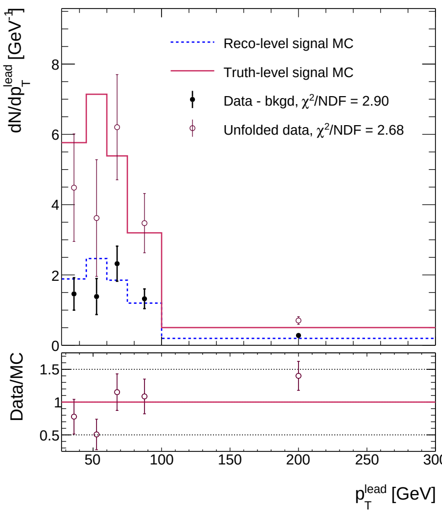

(c) leading lepton p_{T}

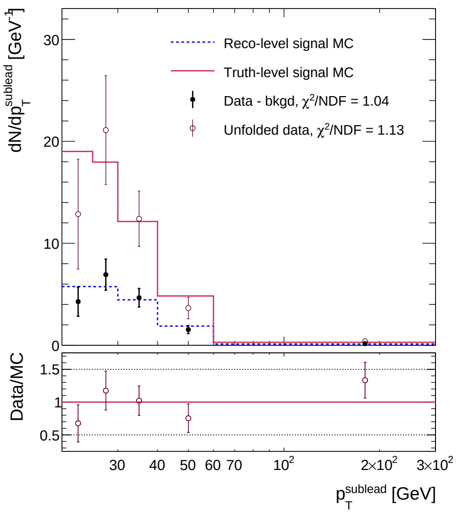

(d) subleading lepton p_{T}

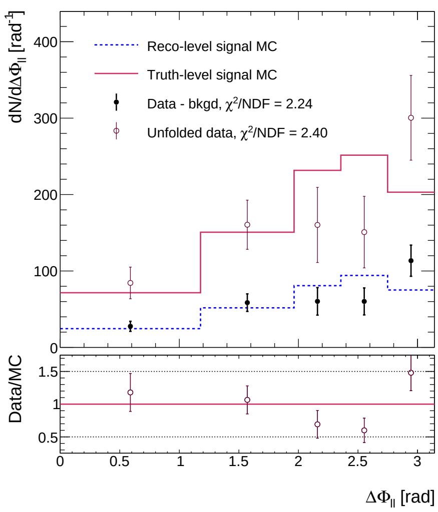

(e) $\Delta\Phi_{\ell\ell}$

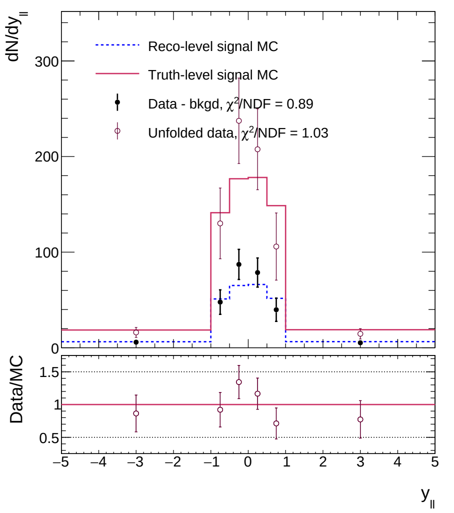

(f) $y_{\ell\ell}$

Figure 6.5: Data unfolded to particle-level using four iterations of D'Agostini's iterative method. Only the statistical uncertainty on the data is included.

Each figure shows the background-subtracted event yield in data (solid black markers) and the predicted signal yield from simulation (dashed blue line). The uncertainty on the data is the statistical uncertainty of the data measurement only, i.e. $\sqrt{N_{\text{data}}}$ before the background subtraction. The statistical uncertainty on the background prediction is considered as a systematic uncertainty. The unfolded data (hollow pink markers) are then compared to the particle-level prediction (solid pink line). The legend of each figure includes the χ^2/NDF test statistic between the data and the corresponding theoretical prediction, which should not change significantly before and after the unfolding procedure if no new information is added by the method. The bottom panels show the ratio of unfolded data to the particle-level prediction.

As the response matrices are very diagonal, the reconstructed distributions do not look very dissimilar to the particle-level distributions, apart from differences in normalisation that correspond to the reconstruction efficiencies in each bin. The statistical uncertainties are around 25–30 % and are increased relative to the statistical uncertainties on the reconstructed yields due to the statistical (anti-)correlations between different bins. The correlation matrices are presented in Appendix D.

Systematic uncertainties are not included. The systematic uncertainties relevant for comparing the unfolded result to theory enter in several ways. Studies into the impact of systematic uncertainties on the unfolded distributions are ongoing and are discussed in more detail in Section 6.5.

6.5 Background systematics

Experimental and theoretical uncertainties on the background yields are passed through the unfolding procedure by varying each contribution to the expected background yield by the $\pm 1\sigma$ uncertainties and repeating the unfolding process. For each source of systematic uncertainty, differences between the unfolded spectrum with systematic variation and the nominal unfolded spectrum is taken as the uncertainty associated to the systematic in question.

The systematics are grouped into different categories by adding contributions in quadrature. "Lepton systematics" refer to uncertainties on lepton identification and isolation, "Track systematics" refer to uncertainties on track reconstruction, "Lumi systematics" refer to the uncertainty on the measured luminosity, "Theory systematics" refer to background modelling uncertainties, "UE correction systematics" refer to uncertainties on the charged particle multiplicity correction in the underlying event and "PU correction systematics" refers to uncertainties on the corrections to pileup track multiplicity mismodelling. More detail on the systematic uncertainties and how they were derived was given in Section 5.7. A breakdown of the sources of systematic uncertainty and their impact on the unfolded $m_{\ell\ell}$ and $p_{\mathrm{T}}^{\ell\ell}$ distributions is shown in Figure 6.6. The lepton and luminosity systematics have a small impact on both distributions, of the order of 1 %. The most important systematic uncertainties are related to the theoretical modelling of backgrounds and to the track reconstruction.

There are some sudden jumps in the systematic uncertainty estimates in Figure 6.6, such as the contribution of lepton systematics in the $140 < m_{\ell\ell} < 200$ GeV bin. This is likely to be due to the limited statistical precision of the data which affects the determination of the systematic uncertainties on the unfolded distributions. The bootstrap method [186], in which several toy histograms are created by randomly varying each bin within Poissonian uncertainties, can be used to evaluate statistical uncertainties on the systematic shifts. This allows only systematic uncertainties to be included in the final result.

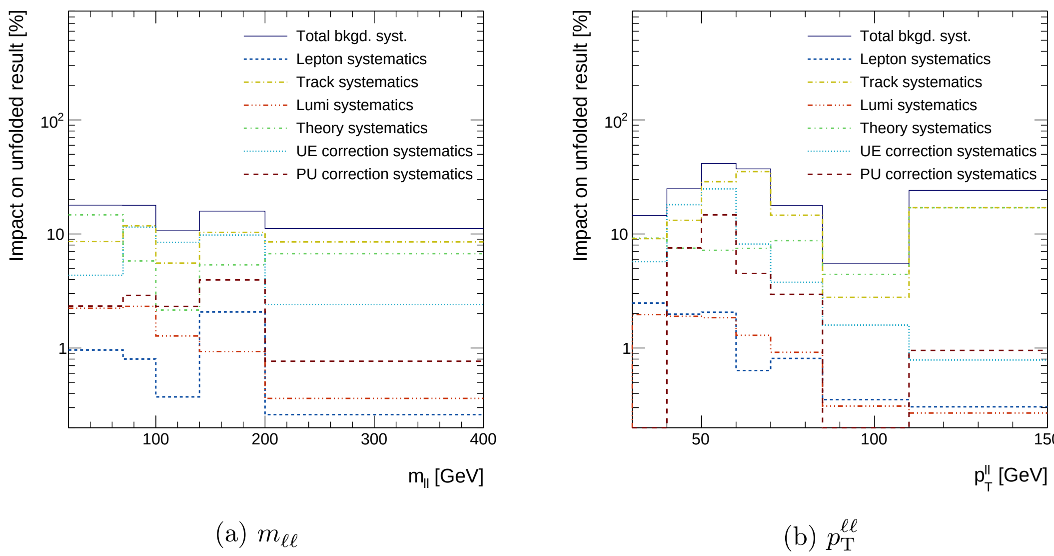

Figure 6.6: The effect of background systematics on the unfolded result for (a) $m_{\ell\ell}$ and (b) $p_{\mathrm{T}}^{\ell\ell}$ when unfolding the background subtracted data with four iterations of unfolding. It is assumed that each source of uncertainty is uncorrelated.

6.6 Addressing sources of bias in the unfolding procedure

Bias can enter the unfolding procedure in several ways. If the background shape templates are mismodelled outside of assigned uncertainties, bias is introduced into the background-subtracted data distribution. Preliminary studies into background shape mismodelling in the $p_{\mathrm{T}}^{\ell\ell}$ distribution are presented in Section 6.6.1. Any unfolding procedure that applies a regularisation scheme introduces bias into the results. The amount of bias can be estimated in a data-driven closure test, presented in Section 6.6.2.

6.6.1 Background $p_{\mathrm{T}}^{\ell\ell}$ mismodelling

The data/MC agreement in the $q\bar{q} \to WW$ control region (CR2 in Table 5.3) for the $m_{\ell\ell}$ and $p_{\mathrm{T}}^{\ell\ell}$ distributions is shown in Figures 6.7a and 6.7b, respectively. The simulation matches the data well for the $m_{\ell\ell}$ distribution but there seems to be a trend in the $p_{\mathrm{T}}^{\ell\ell}$ distribution which suggests that there is some residual mismodelling of the background. Performing a χ^2 test between the data and total MC prediction results in a χ^2/NDF of 0.85 for the $m_{\ell\ell}$ distribution and 2.94 for the $p_{\mathrm{T}}^{\ell\ell}$ distribution. The corresponding p-values are 0.495 and 0.007, respectively.

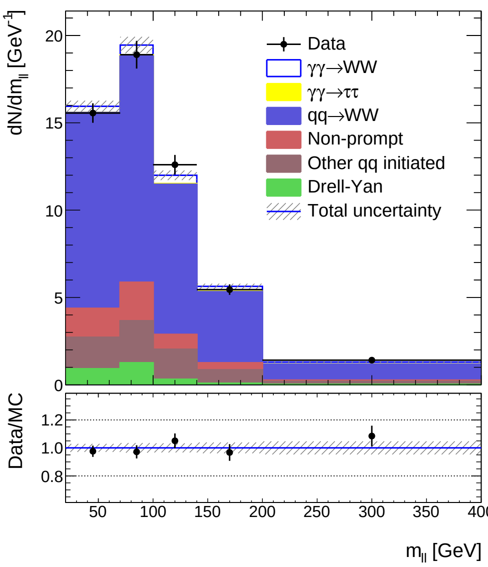

(a) $m_{\ell\ell}$ in CR2

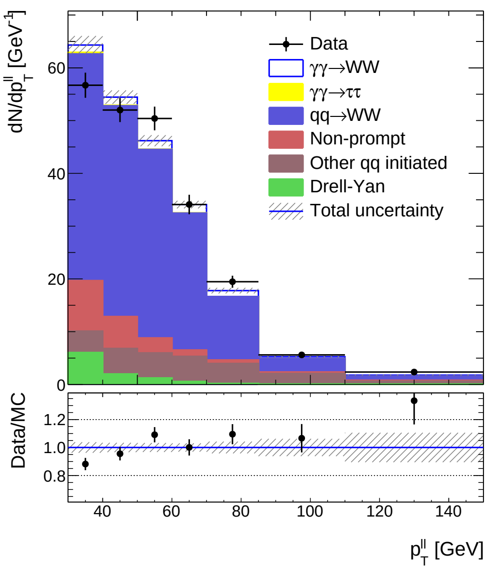

(b) $p_{\mathrm{T}}^{\ell\ell}$ in CR2

Figure 6.7: Comparison of data and MC in the $q\bar{q} \to WW$ control region (CR2) for the (a) $m_{\ell\ell}$ and (b) $p_{\mathrm{T}}^{\ell\ell}$ distributions.

It is assumed that mismodelling in CR2 is due to mismodelling of the $q\bar{q} \to WW$ background. The $q\bar{q} \to WW$ MC can be reweighted using a linear fit to the data/MC ratio to better match the data in this region, after subtracting the signal MC and other backgrounds from the data. The linear fit can then be applied to the $q\bar{q} \to WW$ MC in the signal region as the same mismodelling will be present. The result of the linear fit is $w = (0.6 \pm 0.1)p_{\mathrm{T}}^{\ell\ell}/\mathrm{GeV} + (7 \pm 2) \times 10^{-3}$, where w is the weight to apply to each $q\bar{q} \to WW$ simulated event reconstructed in bin with central value of $p_{\mathrm{T}}^{\ell\ell}$.

The $p_{\mathrm{T}}^{\ell\ell}$ distribution in CR2, after the linear reweighting is applied to the $q\bar{q} \to WW$ simulation, is shown in Figure 6.8a. The $q\bar{q} \to WW$ background was scaled to ensure that the yield was consistent before and after the reweighting. The data and MC agreement is now much better, by construction, with $\chi^2/\mathrm{NDF} = 0.84$ and a p-value of 0.54. The same reweighting is then applied to the $q\bar{q} \to WW$ background in the signal region, shown in Figure 6.8b. The reweighting does act to bring the data and simulation into better agreement in the signal region, especially in the higher $p_{\mathrm{T}}^{\ell\ell}$ bins. The SR $p_{\mathrm{T}}^{\ell\ell}$ distribution in data, with the reweighted background subtracted, was then unfolded to particle level, using four iterations of unfolding. The unfolded distribution before and after the

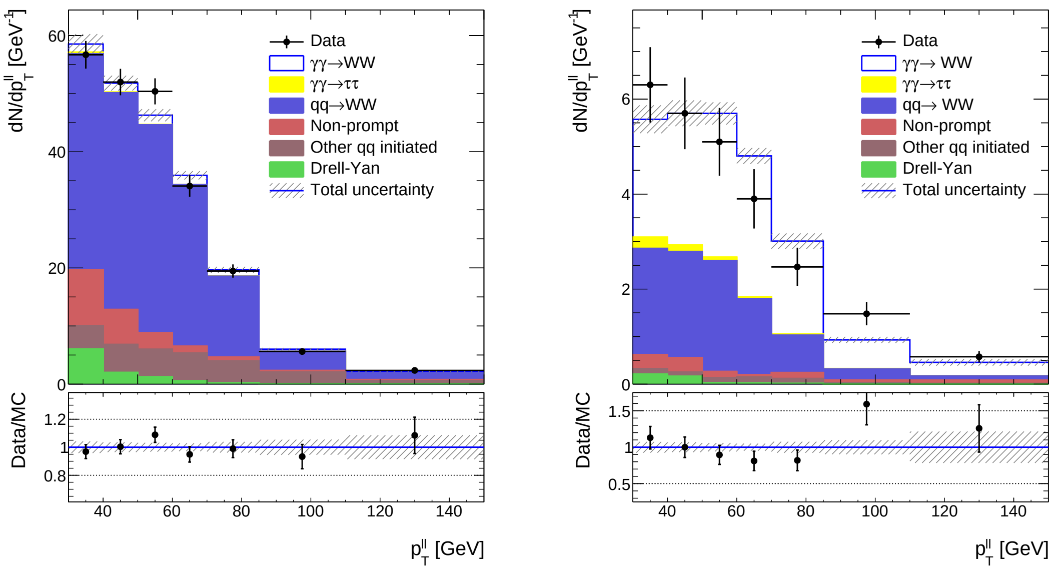

(a) CR2 after reweighting

(b) SR after reweighting

Figure 6.8: The $p_{\mathrm{T}}^{\ell\ell}$ distribution in (a) the $q\bar{q} \to WW$ control region (CR2) and (b) the signal region, after the $q\bar{q} \to WW$ background yield has been reweighted to data using a linear fit. The background normalisation is scaled so that the total background yield is equivalent to the yield before the reweighting is applied.

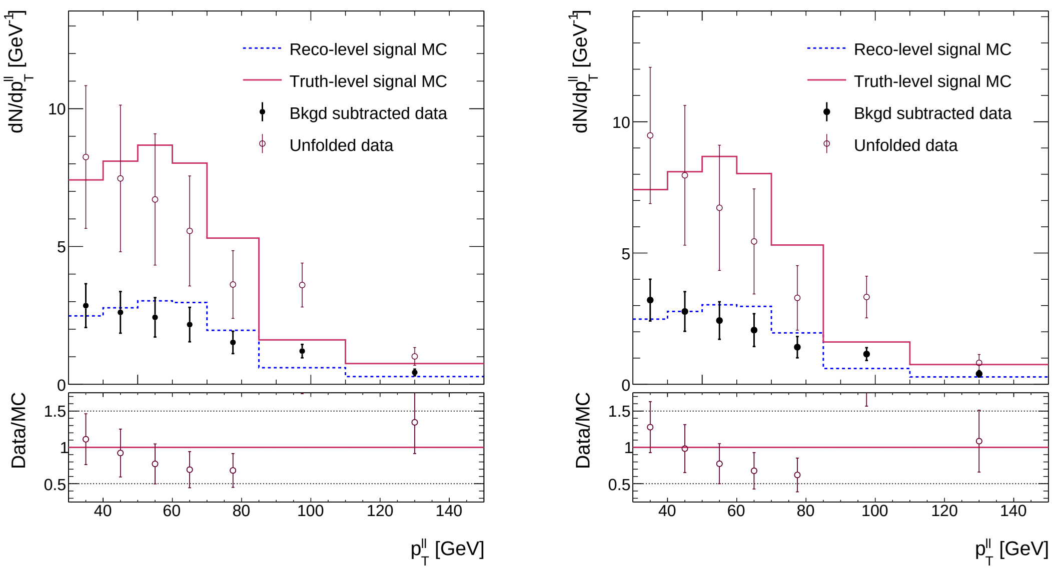

(a) No reweighting

(b) With $q\bar{q} \to WW$ $p_{\mathrm{T}}^{\ell\ell}$ reweighting

Figure 6.9: The unfolded $p_{\mathrm{T}}^{\ell\ell}$ spectrum, (a) before and (b) after reweighting the $q\bar{q} \to WW$ background. The background-subtracted data is reweighted using four iterations.

$q\bar{q} \to WW$ $p_{\mathrm{T}}^{\ell\ell}$ reweighting are compared in Figures 6.9a and 6.9b, respectively. Discrepancies at low $p_{\mathrm{T}}^{\ell\ell}$ are slightly increased after the reweighting but are largely improved in the highest bin, reducing from a relative difference of roughly 35 % to 10 %, albeit within large statistical uncertainties. It is possible that differences between data and MC in the signal region are also due to signal mismodelling, in the way that the dissociative contributions are estimated. Alternative signal simulation is being produced for the differential measurement and differences in signal modelling will be investigated in the future.

The systematic uncertainty on the $q\bar{q} \to WW$ reweighting is determined by varying the background yield within its theory systematic uncertainty and calculating a new reweighting function. The resulting systematic on the unfolded distribution is smaller than 2 % in all bins of $p_{\mathrm{T}}^{\ell\ell}$. However, the difference in the linear fit result when varying the $q\bar{q} \to WW$ theory systematics is within the statistical uncertainties on the fit itself. Therefore, a new estimate of the uncertainty on the $p_{\mathrm{T}}^{\ell\ell}$ reweighting was estimated by instead varying the original fit result within the uncertainty on the fit parameters and passing these variations through the unfolding. The resulting uncertainty on the unfolded result is 3.5 % in the lowest $p_{\mathrm{T}}^{\ell\ell}$ bin and 2.5 % in the highest $p_{\mathrm{T}}^{\ell\ell}$ bin.

If it is decided to reweight the $p_{\mathrm{T}}^{\ell\ell}$ distribution as described in this section for the final measurement, $p_{\mathrm{T}}^{\ell\ell}$ weights must be applied to all distributions on a per-event basis. It may be that the $p_{\mathrm{T}}^{\ell\ell}$ reweighting worsens data/MC agreement in other distributions. A final decision on the reweighting will be made once the unfolding procedure is fully optimised.

6.6.2 Data-driven closure test

The data-driven closure test aims to measure the bias introduced by the unfolding method. Bias is introduced by regularisation as we assume a prior distribution that is equal to the SM prediction. In order to test the amount of bias introduced, we want to reweight the prior distribution to reflect a different but realistic distribution. For this, we choose the measured distribution in data which is by definition a realistic alternative distribution. The weights have to be derived at reconstruction-level, since we don't have access to particle-level data. The weights are then applied on a per-event basis to the simulation, with the weight determined by the particle-level value of the observable. The corresponding weight is then applied to both the particle-level and reconstructed observable, if a reconstructed observable exists. The choice of the level on which to reweight is in principle arbitrary, we simply want to create a new distribution in the simulation that we believe to be realistic. It is chosen to reweight at particle-level because the fiducial fraction ($\sim 95\,\%$) is much higher than the reconstruction efficiency ($\sim 35\,\%$), in all bins and for all observables. If we were to reweight the events based on the reconstructed observable, the majority of events would not be weighted at all.

The bias is then estimated by unfolding the reweighted reco-level spectrum using the nominal response matrix and comparing the unfolded spectrum to the reweighted particle-level spectrum. If the unfolding was fully unbiased, the unfolded spectrum would match the reweighted particle-level spectrum perfectly. Any difference between the two is taken as the bias of the chosen unfolding scheme.

To avoid propagating statistical fluctuations into the data-driven closure procedure, the data/MC ratio is smoothed using Friedman's super smoother [212]. The smoothed functions are shown in Figures 6.10a and 6.10b for $m_{\ell\ell}$ and $p_{\mathrm{T}}^{\ell\ell}$, respectively. The weight for each bin is then applied to particle-level events such that the reconstructed distribution more closely matches the data. The reweighted

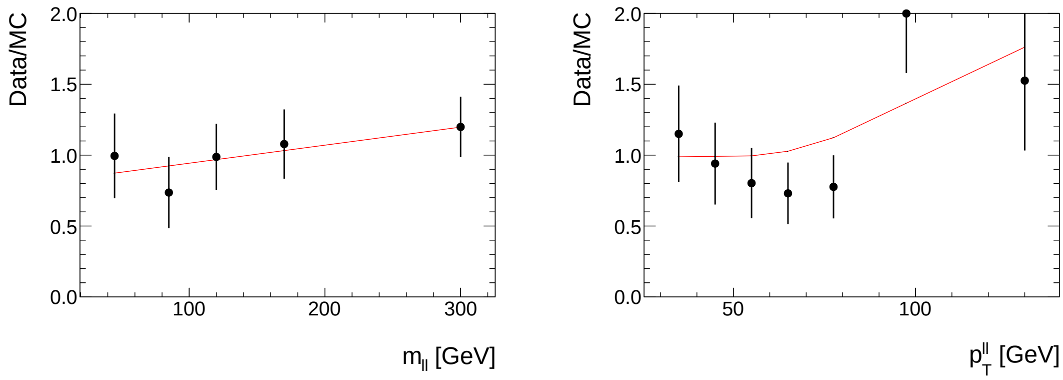

(a) $m_{\ell\ell}$ (b) $p_{\mathrm{T}}^{\ell\ell}$

Figure 6.10: The smoothed ratio (red) of background subtracted data to simulated signal yields in the signal region for (a) $m_{\ell\ell}$ and (b) $p_{\mathrm{T}}^{\ell\ell}$.

reconstructed signal distribution is then unfolded with the nominal response matrix and compared to the reweighted particle-level distribution. This is shown for $m_{\ell\ell}$ and $p_{\mathrm{T}}^{\ell\ell}$ in Figures 6.11a and 6.11b, respectively. Both distributions are unfolded using four iterations. Figure 6.11a has a negligible bias $\ll 1\,\%$ in all bins. The same is true in the $p_{\mathrm{T}}^{\ell\ell}$ distribution, apart from at low values of $p_{\mathrm{T}}^{\ell\ell}$. However, the bias is still insignificant compared to statistical uncertainties on the data and even on the MC itself, with a bias in the lowest $p_{\mathrm{T}}^{\ell\ell}$ bin of $0.6\,\%$. Uncertainties are not shown on the spectra as the statistical uncertainties are much larger than the measured bias, with statistical uncertainties of around $25\,\%$ seen in the unfolded data spectra in Figure 6.5. As the measured bias is significantly smaller than the statistical uncertainty, it is likely that fewer iterations will be preferred.

Optimising the number of iterations

The data-driven closure test can be used to optimise the number of unfolding iterations. By measuring the bias between the unfolded and reweighted spectra and the statistical uncertainty on the unfolded result, the optimal number of iterations can be selected as the number that minimises the combination of the two. The data-driven closure test was performed using 1–10 iterations and the resulting bias and statistical uncertainty in each bin is shown in Figure 6.12. To monitor a realistic statistical uncertainty across the iterations, an Asimov dataset was created using the expected distribution predicted by the reweighted MC, but with the statistical uncertainty on the data assigned as the uncertainty. The bottom panel shows the bias and statistical uncertainty added in quadrature. The statistical uncertainty dominates by 1–2 orders of magnitude and therefore

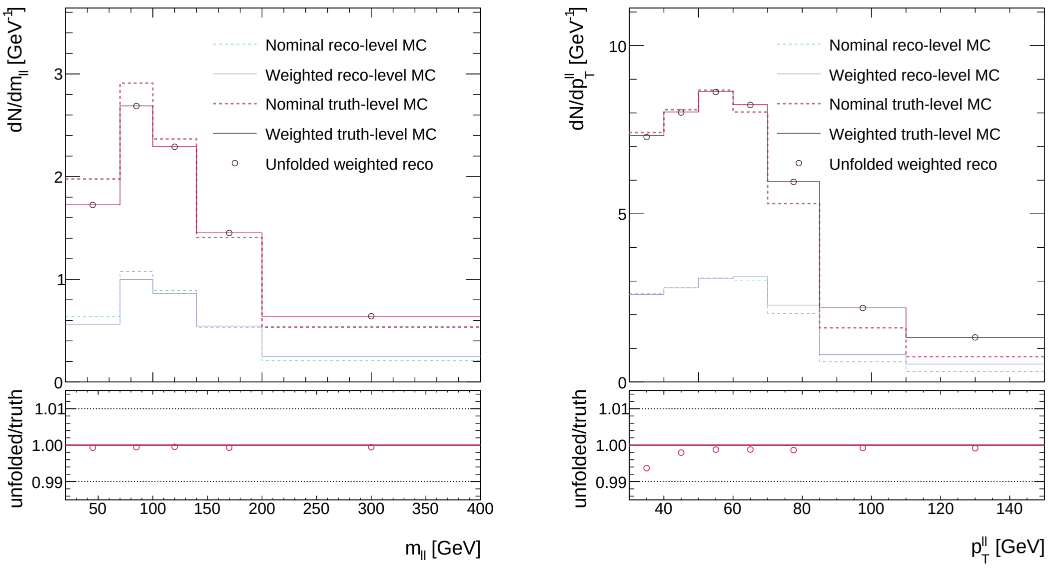

(a) $m_{\ell\ell}$ (b) $p_{\mathrm{T}}^{\ell\ell}$

Figure 6.11: Data-driven closure test for (a) the $m_{\ell\ell}$ distribution and (b) the $p_{\mathrm{T}}^{\ell\ell}$ distribution. Reweighted reconstruction-level simulation is unfolded to particle-level using four iterations and the nominal response matrix.

one iteration of unfolding is optimal for both distributions, in all bins.

The unfolding of the background subtracted data is re-performed with one iteration and compared to the original unfolded spectra when using four iterations, shown in Figures 6.13a and 6.13b for the $m_{\ell\ell}$ and $p_{\mathrm{T}}^{\ell\ell}$ distribution, respectively. The difference in the unfolded result is minimal, but the data points are generally pulled slightly closer to the MC prediction. This is expected as using fewer iterations biases the unfolded spectrum towards the prior, which is the MC prediction.

One iteration being optimal is not surprising, given how diagonal the response matrices are. However, it necessitates the question of whether D'Agostini's iterative unfolding is the optimal method, given that it introduces additional systematic uncertainties in the unfolding procedure. Alternative unfolding methods shall be investigated in the future and compared to the results from the iterative method.

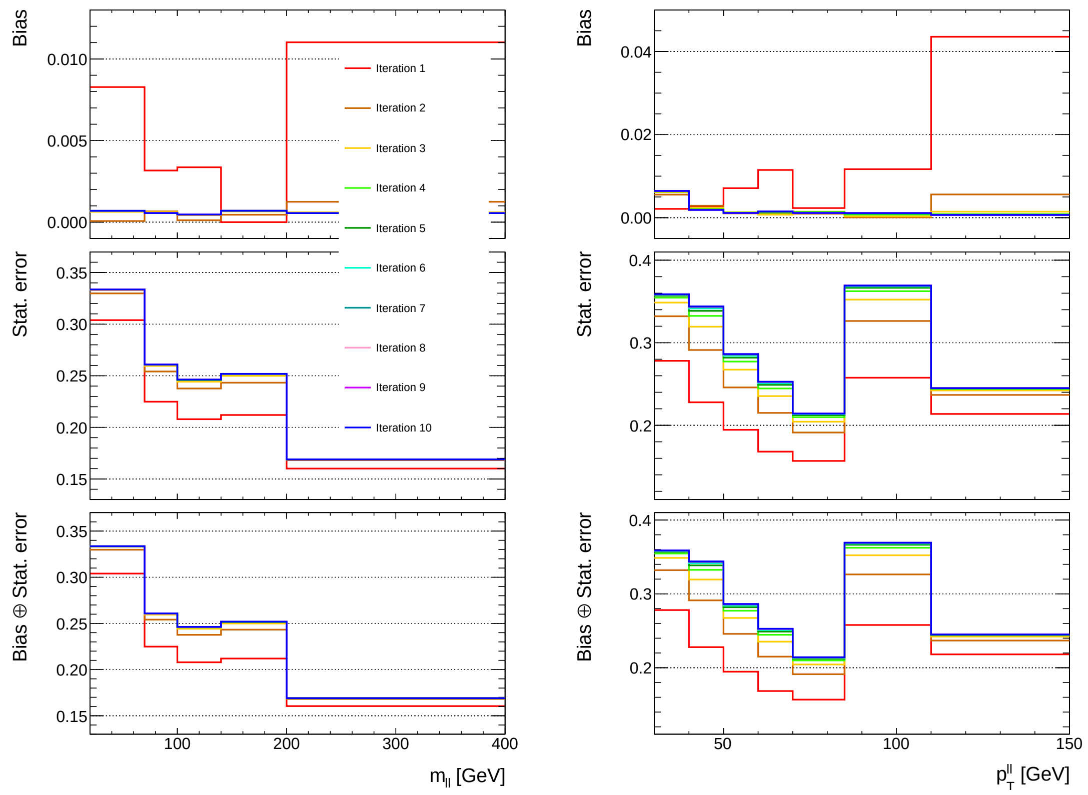

(a) $m_{\ell\ell}$ (b) $p_{\mathrm{T}}^{\ell\ell}$

Figure 6.12: Bias and statistical uncertainty at each iteration of unfolding, using 1–10 iterations. The first iteration is shown in red and the 10th iteration is shown in dark blue. The bottom panel shows the combination of the bias and statistical uncertainty added in quadrature.

Cross-checks and improvements on the data-driven closure

The implementation of the data-driven closure can be cross-checked by redoing the test using a fully-unbiased unfolding method such as matrix inversion, where the response matrix is inverted and used to unfold the data.

An improvement to the data-driven closure test would be to repeat the procedure with several alternative but realistic distributions. This could be done by altering the parameters of the smoothed fit to the data/MC ratio, from which the weights are derived. By fitting the curves as being consistent with the data but as far from the SM prediction as possible, we can define a bias test which reflects the maximal reasonable amount by which the “true” reconstructed spectrum can deviate from the SM and still be consistent with the measured data. Ultimately, it is unlikely that any bias will compare to the relative size of the statistical uncertainty in the analysis.

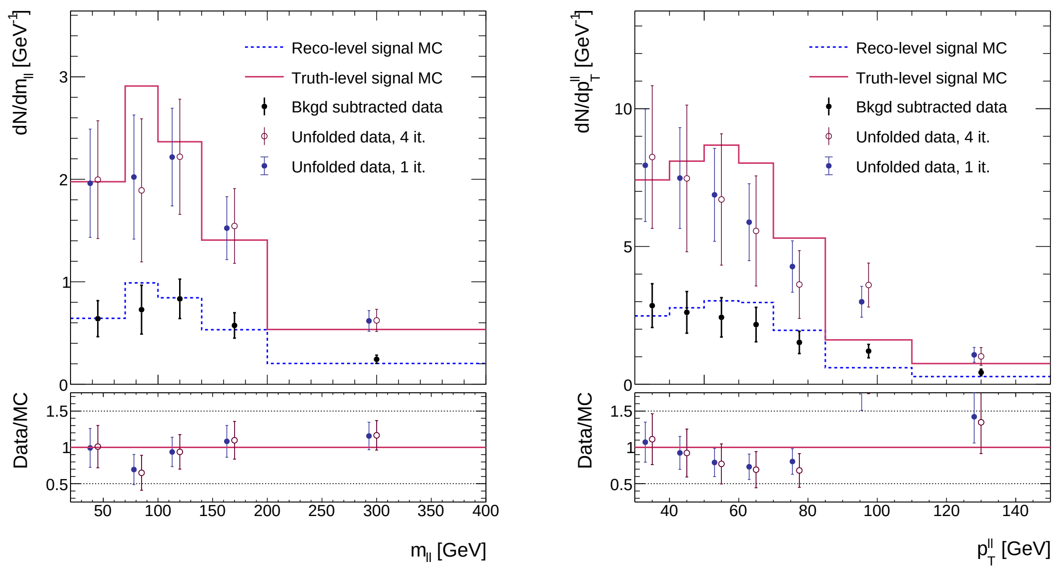

(a) $m_{\ell\ell}$ (b) $p_{\mathrm{T}}^{\ell\ell}$

Figure 6.13: Unfolding the background-subtracted data using four iterations (pink hollow markers) and one iteration (blue solid markers) of the unfolding procedure. Uncertainties are statistical only. The markers for one iteration have been offset along the abscissa to aid readability.

6.7 Summary and outlook for the differential measurement

This chapter presented preliminary unfolding studies for various kinematic spectra in the $\gamma\gamma \to WW$ signal region, as defined in the observation analysis. The post-fit simulation predicted yields generally agree well with the data, within statistical and systematic uncertainties. There is larger disagreement at high dilepton p_{T}, $p_{\mathrm{T}}^{\ell\ell}$, and the leading-lepton p_{T} appears to be slightly harder in data compared to simulation. There is some disagreement also at high values of $\Delta\Phi_{\ell\ell}$, where the leptons have large separation in azimuthal angle. Mismodelling of $p_{\mathrm{T}}^{\ell\ell}$ in the background $q\bar{q} \to WW$ simulation was investigated, and its effects corrected for in the signal region. This brought the unfolded spectra into slightly better agreement with particle-level predictions in simulation but further studies need to be performed to decide if this correction is appropriate. It is possible that tensions between data and MC in the signal region are due to the way in which dissociative contributions are accounted for in the theoretical prediction, where it is assumed that the dissociative signal kinematics are the same as for the elastic contribution.

Unfolded distributions were presented using the iterative unfolding technique proposed by D'Agostini, with four iterations. Attempts to optimise the method when considering statistical uncertainties only were performed for the $m_{\ell\ell}$ and $p_{\mathrm{T}}^{\ell\ell}$ distributions and demonstrated that fewer iterations were optimal, as the measurement has very little migration between particle-level and reconstruction-level bins and has large statistical uncertainties per bin of around 20 % in the data, pre-unfolding. In the case of the $m_{\ell\ell}$ and $p_{\mathrm{T}}^{\ell\ell}$ distributions, studies considering statistical uncertainties only suggest that one iteration would be optimal. Other unfolding techniques shall be tested and compared in the future. The unfolding method will ultimately be determined in tests which include all sources of uncertainty, including systematic uncertainties. The impact of systematic uncertainties on the background yields were propagated through the unfolding procedure, before optimisation. The total impact on the unfolded result was between 10–20 % in the $m_{\ell\ell}$ yields and between 5–40 % in the $p_{\mathrm{T}}^{\ell\ell}$ yields, depending on bin. However, further work needs to be done to understand how the statistical uncertainty on the data affects the predicted impact of background systematics on unfolded distributions.

So far, only systematics affecting the background yields in the signal region have been considered. Theoretical uncertainties on the signal yield enter into the response matrix used for the unfolding procedure. In order to test the impact of these systematics, a new response matrix should be created for every signal modelling uncertainty and the unfolding should be repeated with this new response matrix on Asimov data. Differences between the nominal and systematic unfolded result is then taken as the uncertainty on the unfolded yields. Systematic uncertainties that affect the signal and background, such as lepton reconstruction efficiencies, can be cross-checked by applying them to either the background yields or the response matrix but care should be taken to not double count these uncertainties in the final results. Once all sources of uncertainty are accounted for, pull tests should be performed by generating pseudo-experiments within the full assigned uncertainty to make sure the full unfolding method is not biased due to some other reason specific to the analysis. The response matrix itself often depends on the shape of the underlying distribution, especially when it is steeply falling. This is something that unfolding cannot correct, and therefore pull tests are necessary to ensure that there is not a large bias introduced and that there are not any significant unknown sources of bias.

Alternative theoretical predictions for the signal will also be used to create alternative response matrices. Differences in unfolded results when using these alternative response matrices will be taken as a systematic on the unfolding method, as it relates to the assumption that the predicted distribution in simulation is the true underlying distribution. Moving forward, new simulation is available for dissociative $\gamma\gamma \to WW$ production, including predictions from SuperChic and improved MadGraph5_aMC@NLO simulation. These new samples will also allows us to probe differences between the elastic and dissociative contributions to the signal.

The amount of bias introduced by the unfolding procedure was investigated by performing a data-driven closure test, in which the predictions from simulation were reweighted to look more like data, representing an alternative yet realistic prior distribution. Similar tests can be performed to test how well the unfolding procedure can reproduce new physics signals within the data, referred to as an injection test. This could be tested by unfolding pseudo-data that has been created from a BSM model, such as introducing a non-zero EFT coupling coefficient, and observing whether this signal is recovered when using the nominal response matrix. Another test that could be performed is a stress test, in which the entire distribution is weighted by large and varying weights. These tests are essentially more extreme versions of the data-driven closure test, where we observe the bias in data distributions that we do not necessarily believe will be realistic.

Once the analysis chain has been finalised and the unfolding procedure has been optimised, unfolded distributions will be used to calculate limits on EFT operators that impact $WW\gamma$ and $WW\gamma\gamma$ couplings. It is expected that dimension-6 operators will be better constrained by other VBS and diboson processes and therefore focus will be on setting limits on anomalous contributions to the $\gamma\gamma \to WW$ cross-section arising from dimension-8 EFT operators. From preliminary studies [208], measurement of the $\gamma\gamma \to WW$ process is expected to improve current limits on the $f_{M,2\text{–}5}/\Lambda^4$ and $f_{T,5\text{–}7}/\Lambda^4$ parameters. These preliminary studies only considered quadratic EFT contributions to the cross-section and would likely be further constrained by also considering linear interference terms with the SM. Furthermore, these calculations were performed for a single-bin measurement and further improvement of approximately a factor of three is expected to be found in a differential analysis.

Chapter 7

Future prospects of measuring $\gamma\gamma \to WW$ at the High-Luminosity LHC

"The practice of science happens at the border between the known and the unknown. Standing on the shoulders of giants, we peer into the darkness with eyes opened not in fear but in wonder."

– Brian Cox

7.1 Introduction and motivation

The expected integrated luminosity of between 3000–4000 fb^{-1} at the end of data-taking at the High-Luminosity LHC will allow a large improvement in statistical precision for any given measurement, which is especially important for processes with relatively small cross-sections like $\gamma\gamma \to WW$. However, the increase in instantaneous luminosity comes at a cost. The average number of interactions per bunch crossing is predicted to be between $140 < \mu < 200$ at the HL-LHC [213, 214], compared to an average of 33–35 in Run 2 [215]. This increase in pileup will be detrimental to a $\gamma\gamma \to WW$ analysis due to the way in which the exclusivity requirements are defined, namely the $n_{\mathrm{trk}} = 0$ requirement.

This chapter studies the expected sensitivity to the $\gamma\gamma \to WW$ process at the HL-LHC when using a similar analysis approach as in the Run 2 observation, discussed in Chapter 5. The aim is to understand how the HL-LHC environment and planned detector upgrade will impact such a measurement. Section 7.2 presents the ATLAS detector upgrades most relevant to this measurement and the procedure for reconstructing tracks below the nominal p_{T} reconstruction threshold. Section 7.4 discusses the differences in various reconstruction efficiencies and how they impact the expected signal and background yields between Run 2 and the HL-LHC within the context of a $\gamma\gamma \to WW$ analysis. The additional procedure of reconstructing forward protons in AFP and the associated efficiencies are discussed in Section 7.6.

The projection is first done in the context of an inclusive cross-section measurement and results are presented for many possible tracking configurations and track-veto window sizes, with the aim of determining the optimal analysis strategy in HL-LHC operating conditions. These results are presented in Section 7.5. After the optimal configuration is found, the results are presented as a function of dilepton invariant mass, $m_{\ell\ell}$, in Section 7.7. The results of these studies have been published in Refs. [5, 216]. The impact of requiring a single or double proton-tag in AFP is then investigated and results of these studies are currently being used as part of a physics case to install AFP-like detectors during Phase-II upgrades. The author led and performed all studies in this chapter, with the exception of the derivation of low-p_{T} tracking working points and low-p_{T} track reconstruction efficiencies documented in Section 7.2.1.

7.2 ATLAS Phase-II Upgrades

ATLAS is due to undergo major "Phase-II" upgrades before HL-LHC data taking begins. The main upgrade will be the full replacement of the Inner Detector with a fully-silicon Inner Tracker [217, 218], discussed in Section 7.2.1. New electronics will be installed for both the LAr and Tile Calorimeters [219, 220]. Upgrades to the muon system are foreseen to increase geometrical acceptance and improve triggering [221]. The TDAQ system will be upgraded to cope with increased data rates due to the high number of interactions per bunch crossing [222]. It is planned to install a new High-Granularity Timing Detector (HGTD) [223] in the $2.4 < |\eta| < 4$ region, with the main aim to aid pileup suppression in the forward region. The HGTD detector is described in Section 7.2.2, as it allows a novel $\gamma\gamma \to WW$ analysis approach where a lepton-timing tag from the HGTD is required in conjunction with forward proton ToF measured in AFP. This is described in Section 7.6.3.

7.2.1 Inner Tracker

The current ATLAS ID will be replaced by the all-silicon Inner Tracker (ITk), with the aim to maintain (or improve on) the ID tracking performance but in a much harsher radiation environment. The ITk will comprise of a silicon pixel detector [217] complemented by a silicon strip detector [218], together allowing track reconstruction in the range $|\eta| < 4$. The pixel system is the innermost of the two subsystems, comprising of five barrel layers plus inclined modules and rings in the end-cap regions. The strip system has four barrel layers and six end-cap disks, covering the region $|\eta| < 2.7$. A diagram of the ITk layout with both silicon pixels and strips is shown in Figure 7.1. The pixel system will be composed of more than 10,000 modules with over 10^9 channels. The strip system will have more than 18,000 modules with almost 60 million channels. A minimum number of 9 measurements per track will be assured throughout the angular acceptance range of the ITk, with an average of about 13 pixel plus strip hit measurements for most of the acceptance range.

The design goal of the ITk is to maintain high efficiency for charged particles above $p_{\mathrm{T}} > 1\,\mathrm{GeV}$ and to have a low number of fake tracks. Table 7.1 compares the current ID to the ITk for both the minimum reconstructed track-p_{T} threshold and the minimum number of hits per track.

The ITk allows a much stronger rejection of fake tracks compared to the

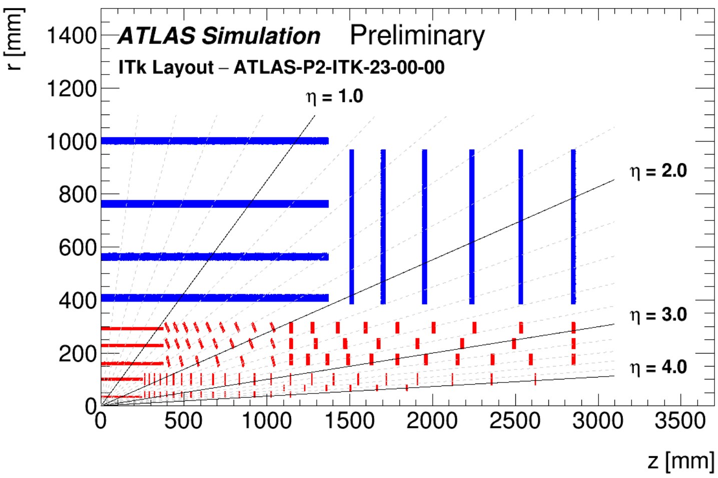

Figure 7.1: A schematic diagram of one quadrant of the ITk, with the most recent updated pixel layout. The active elements of the strip detector are shown in blue, and those of the pixel detector are shown in red. The horizontal axis is along the beam line with zero being the interaction point. The vertical axis is the radius measured from the interaction region [224].

present ID, resulting in a negligible fake-rate for tracks with $p_{\mathrm{T}} > 1\,\mathrm{GeV}$. However, it is desirable to reconstruct below the nominal ITk track-p_{T} thresholds in the $\gamma\gamma \to WW$ analysis, in order to maintain high background suppression through the $n_{\mathrm{trk}} = 0$ requirement. This is because tracks in the underlying event of quark-induced backgrounds are soft and usually have very low p_{T}.

	Run 2 ID	ITk (HL-LHC baseline)		
	$\lvert\eta\rvert < 2.5$	$\lvert\eta\rvert < 2.0$	$2.0 < \lvert\eta\rvert < 2.6$	$2.6 < \lvert\eta\rvert < 4.0$
Min. p_{T} [MeV]	500	900	400	400
Min. number of Si hits	7	9	8	7

Table 7.1: Minimum p_{T} and minimum number of required silicon (pixel + strips) hits for track reconstruction in the Run 2 ID and foreseen for the ITk. The latter is further segmented in three $|\eta|$ regions and is referred to as the "HL-LHC baseline" reconstruction in the text.

Low-p_{T} track reconstruction

The performance of the nominal ITk configuration has been studied extensively in Ref. [225]. Low-p_{T} tracking refers to the process of reconstructing tracks of charged particles with p_{T} below the default threshold. The reconstruction of low momentum particles has been examined in Ref. [226] for the ATLAS ID and is currently being studied in ongoing work for the ITk.

The simplest implementation of low-p_{T} tracking is to apply the default ITk track parameter requirements with reduced track and seed-p_{T} selections. The problem with this approach is that the nominal selections have been designed and optimised for the reconstruction of tracks of charged particles with a $p_{\mathrm{T}} > 900\,\mathrm{MeV}$ and are not optimal for reconstructing lower in p_{T}, leading to a reduction in track efficiency below this nominal threshold. The nominal configuration requires a minimum of nine hits in the low η region ($|\eta| < 1$) for the ITk. Whilst nearly all 1 GeV tracks meet this requirement, a large number of tracks with much lower p_{T} fall below the minimum hit requirement. A modified tracking selection was developed for the purpose of track reconstruction for track-$p_{\mathrm{T}} > 500\,\mathrm{MeV}$ and is compared to the default ITk selections in Table 7.2. The track efficiency and z_0 resolution values have been estimated from a fully-simulated $t\bar{t}$ sample with $\mu = 200$ and the upgraded detector response (including the ITk) fully simulated using GEANT4, where the modified selections in Table 7.2 were applied to provide reasonable efficiency for tracks with $p_{\mathrm{T}} > 500\,\mathrm{MeV}$ and $|\eta| < 2.5$. It is assumed such a reconstruction can be developed for nominal pileup conditions whilst maintaining sufficiently low fake rate.

The longitudinal impact parameter resolution is improved with the ITk relative to the Run 2 ID, going from $\sigma(z_0) \sim 100\,\mu\mathrm{m}$ to $\sigma(z_0) \sim 50\,\mu\mathrm{m}$ in the central region for a representative track-p_{T} of 2 GeV. Figure 7.2 shows how the track-z_0 resolution varies with p_{T} and η. Tracking z_0 resolution degrades with increasing track-η, due to increased material budget and longer extrapolation distance in the forward region. For the ITk, the z_0 resolution is roughly $\sigma(z_0) = 1\,\mathrm{mm}$ at the edge of the angular acceptance for a track with $p_{\mathrm{T}} = 2\,\mathrm{GeV}$. Similarly, tracking

	ITk	Low p_T	
	Default	Default	Modified
Min clusters	[9,8,7]		[6,5,4]
Min Si not shared	[7,6,5]		[6,5,4]
Max Holes	2		4
Max Pixel Holes	2		4
Max Sct Holes	2		4
Max Double Holes	1		4
Min p_T	[900, 400, 400]	[200, 200, 200]	
Min track seed p_T	900	200	

Table 7.2: The ITk default, low-p_{T} default and low-p_{T} modified tracking selections applied for the track reconstruction in a $t\bar{t}$ MC sample with $\mu = 200$. The lists entries correspond to the η-ranges $0 < |\eta| < 2.0$, $2.0 < |\eta| < 2.6$ and $|\eta| < 4.0$.

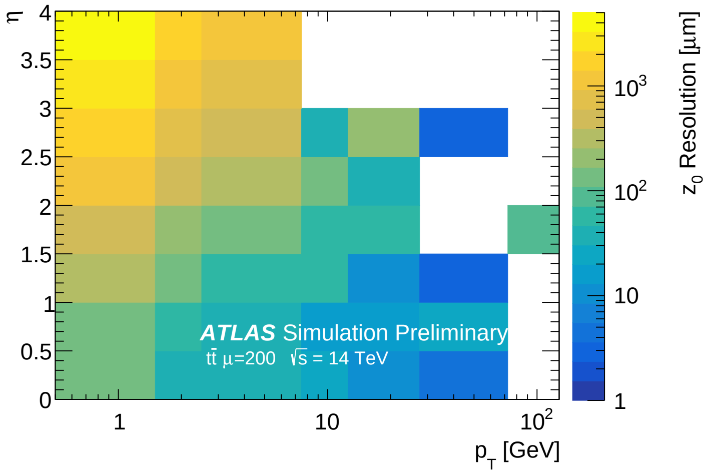

Figure 7.2: The longitudinal impact parameter (z_0) resolution of the ITk as a function of η and p_{T} for $t\bar{t}$ events at pileup $\mu = 200$, for tracks with $p_{\mathrm{T}} > 500\,\mathrm{MeV}$ [5].

p_{T} resolution worsens in the forward region relative to central η values due to the fact that the solenoid field in this region is weaker and the transverse path length of the tracks is shorter. This is why the tracking requirements in Table 7.1 are looser in the forward region to maintain high tracking efficiency for charged particle with $p_{\mathrm{T}} > 1\,\mathrm{GeV}$.

7.2.2 High-Granularity Timing Detector

The effective pileup density at the HL-LHC is predicted to be somewhere between 1–2 mm^{-1} depending on different machine settings and point in the fill [227]. This means that average vertex z_0 positions could be separated by as little as 500 μm which is smaller than the resolution of the ITk for low-p_{T} tracks in the forward region, demonstrated in Figure 7.2. In order to improve track-to-vertex association in this region, ATLAS plans to install a new High-Granularity Timing Detector (HGTD) [223] during the Phase-II upgrades, designed to complement the ITk in the region of $2.4 < |\eta| < 4.0$. The combination of the ITk and HGTD will allow charged particle trajectories to be measured in both time and space, allowing further separation of signal and pileup vertices when the tracks are within the HGTD acceptance. The addition of the HGTD is projected to significantly improve the suppression of pileup jets in VBS topologies [228].

The HGTD will be located in the gap region between the barrel and the end-cap calorimeters, in the volume currently occupied by the Minimum-Bias Trigger Scintillators (MBTS), at a distance in z of approximately $\pm 3.5\,\mathrm{m}$ from the interaction point, shown in Figure 7.3. The entire detector will have a

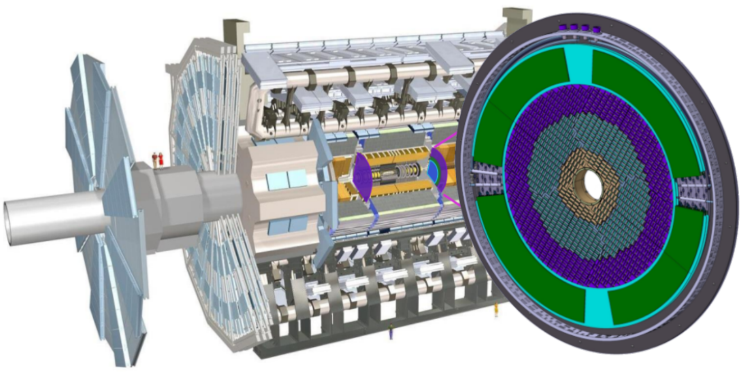

Figure 7.3: A view of the HGTD modules and their position within the ATLAS detector. The HGTD detector spans a radius of 110–1000 mm with the acceptance defined as the surface covered by the HGTD between a radius of 120–640 mm ($2.4 < |\eta| < 4$) at a position of $z = \pm 3.5$ m along the beamline on both sides of ATLAS [223].

radial extent of 110–1000 mm and be 125 mm thick, including a 50 mm neutron moderator placed behind the active area to protect the HGTD and ITk from back-scattered neutrons created by the end-cap calorimeters. Low Gain Avalanche Detector (LGAD) sensor pads [229] of $1.3 \times 1.3\,\text{mm}^2$ area, with an active thickness of 50 µm, were chosen to fulfill space and detector requirements. Each LGAD module contains 30×15 pads, with a total of 8032 modules in the entirety of the HGTD. The chosen pad size ensures occupancies below 10 % at the highest expected levels of pileup, small dead areas between pads, and low sensor capacitance which is important for the timing resolution. The HGTD aims to provide a uniform timing resolution as a function of η of $\sigma_{\text{HGTD}} = 30$ ps per minimum-ionising particle (MIP). This is expected to degrade to ~ 50 ps per MIP at the end of the HGTD lifetime.

The two main experimental challenges of the HGTD are the large amount of detector material between the detector and the interaction point, and the limited η acceptance in the space available between the ITk and the end-cap calorimeter. A high material budget in front of the HGTD increases the probability of secondary scatterings, limiting the ability to accurately associate time to tracks with high confidence. A limited acceptance impacts the ability to determine the global event-vertex time as the hard-scatter interaction needs to produce enough particles within the HGTD acceptance to separate them from forward activity from pileup interactions. An object with multiple associated MIPs has a timing resolution improved by a factor of $1/\sqrt{N_{\text{MIP}}}$, where N_{MIP} is the number of reconstructed particles associated to the object. This is not so useful in an analysis where we expect the signal process to have very few particles associated to it, such as $\gamma\gamma \to WW$, as the timing resolution per MIP is not precise enough to separate the signal process from pileup in the forward region. However, the possibility of matching the timing from a lepton tag in the HGTD to a forward proton in the AFP-ToF detector (which is predicted to have better timing resolution of $\sigma_{\mathrm{ToF}} \leqslant 20\,\mathrm{ps}$ at the HL-LHC) is investigated as a possible extension to the $\gamma\gamma \to WW$ programme in Section 7.6.3.

7.2.3 Installation of an AFP-like detector

The HL-LHC physics programme must be focused on physics topics for which high luminosity is a necessary requirement. While special, low-luminosity runs for diffractive studies might well be possible, the primary goal of the HL-LHC is the study of rare phenomena and a search for new physics beyond the Standard Model. Therefore, the installation of an AFP-like detector at the HL-LHC would focus on measuring photon-induced and central exclusive production processes.

AFP has not yet been commissioned at the HL-LHC but it is likely that the same or very similar technology as the Run 2 AFP detector (Section 4.2) would be reused in a HL-LHC AFP detector. The AFP physics reach at the HL-LHC depends on the mass acceptance, which is governed by the HL-LHC optics and how close the Roman Pots can get to the beam. There are possibilities to have AFP RPs at multiple locations at the HL-LHC based on the location of other machine components. These locations are RP1 at roughly $z = \pm 197\,\mathrm{m}$, RP2 at roughly $z = \pm 218\,\mathrm{m}$, and RP3 at roughly $z = \pm 240\,\mathrm{m}$ from IP1. Each location would house a NEAR and FAR station. Two combinations of these locations are considered: RP1+RP2 with an acceptance in forward proton fractional energy loss of $0.08 < \xi < 0.31$, and RP2+RP3 with $0.05 < \xi < 0.20$. Additionally, the further the distance between the NEAR and FAR stations at each location, the better the resolution in forward proton ξ and p_{T}. The acceptance in ξ is determined by the most likely optics scenario at the HL-LHC, which assumes $\sqrt{s} = 14\,\mathrm{TeV}$, $\beta^* = 0.15\,\mathrm{m}$ and a crossing angle of $\theta_c = 250\,\mu\mathrm{rad}$. The phase of the crossing angle also has an impact on the trajectories of scattered protons. One of IP1 (ATLAS) and IP5 (CMS) will receive a phase of $\phi = 0°$ and the other $\phi = 180°$. It is assumed that ATLAS will receive the phase of $\phi = 0°$.

PPS have released an Expression of Interest to install similar forward detectors surrounding IP5 at the HL-LHC [230].

7.3 Simulated event samples

Dedicated HL-LHC MC samples were produced with $\sqrt{s} = 14\,\text{TeV}$ and $\mu = 200$ for the signal $\gamma\gamma \to WW$ process and the main background process, $q\bar{q} \to WW$. These samples were used to calculate signal and background selection efficiencies compared to Run 2 and these efficiencies were then used to extrapolate the expected Run 2 event yields to the HL-LHC. The HL-LHC samples use the same generators as the Run 2 samples, listed in Section 5.2 and Ref. [4], though in some cases use updated versions of the generators.

The elastic component of the $\gamma\gamma \to WW$ signal process was modelled at leading order using using MG5_aMC@NLO 2.8.1 interfaced to Pythia 8.244, using the default photon flux in MadGraph. The dominant $q\bar{q} \to WW$ background was modelled using Powheg-Box v2 with the CT10 PDF for the matrix element calculation. It was interfaced to Herwig 7.2.1 for the parton shower, using the H7UE tune and the MMHT2014LO PDF set. Pileup was modelled by overlaying the simulated hard-scattering event with inelastic pp events generated with Pythia 8.244p3 using the NNPDF2.3LO set of PDFs and the A3 set of tuned parameters.

The performance of the track reconstruction was studied using $t\bar{t}$ events where the average number of pileup interactions was either set to zero or to $\mu = 200$. The production of these events was modelled using the Powheg-Box v2 generator at NLO with the NNPDF3.0NLO PDF set and the h_{damp} parameter[12] set to $1.5\,m_{\text{top}}$ [231]. The events were interfaced to Pythia 8.240p4 to model the parton shower, hadronisation, and underlying event, with parameters set according to the A14 tune [232] and using the NNPDF2.3LO set of PDFs.

The event samples were processed with a full detector simulation of ATLAS using Geant4, with the ITk layout as described in Ref. [225].

The simulated event samples only include reconstructed tracks using the HL-LHC baseline reconstruction, presented in Table 7.1. A parametrised approach was taken to reconstruct charged particles below the nominal ITk reconstruction thresholds. Charged particles in the analysis samples below the minimum reconstructed p_{T} threshold were weighted by the parametrised

[12] The h_{damp} parameter is a resummation damping factor and one of the parameters that controls the matching of Powheg matrix elements to the parton shower and thus effectively regulates the high-p_{T} radiation against which the $t\bar{t}$ system recoils.

reconstruction efficiencies as a function of p_{T} and η. The production-vertex z-position of each charged particle was smeared by a Gaussian distribution with a width equal to the parametrised track resolution, as given in Figure 7.2, also as a function of p_{T} and η. The efficiency of an event to pass the $n_{\mathrm{trk}} = 0$ selection is then calculated as $\prod_i (1 - \varepsilon^i_{\mathrm{trk}})$, where $\varepsilon^i_{\mathrm{trk}}$ is the probability of reconstructing the i^{th} charged particle in the event. This is an event-by-event exclusive efficiency.

7.4 Event selection and projection strategy

The event selection of the $\gamma\gamma \to WW$ projection [5] follows the same logic as in the observation analysis, given in Section 5.3, with a few important changes and additions described in this section. All extrapolations are performed in the signal region which is defined by the reconstruction-level selections given in Table 5.2.

The most important selection applied in the $\gamma\gamma \to WW$ analysis is the $n_{\mathrm{trk}} = 0$ requirement, which exploits the fact that photon-induced processes do not have an underlying event by requiring exactly zero reconstructed tracks within a symmetric z-window around the dilepton vertex, referred to as the track-veto window. The track-veto window definition was described in detail in Section 5.4. The size of the window was optimised to be ± 1 mm in the Run 2 analysis (see Section 5.4.6) but it is expected that the denser pileup environment at the HL-LHC will result in the optimal window size being smaller than in Run 2 and therefore window sizes ranging between 0.1–1 mm were investigated for the HL-LHC projection.

For the purpose of quantifying track reconstruction performance at the HL-LHC and comparing to Run 2 tracking performance, three different tracking configurations are investigated:

1. **Nominal configuration:** Baseline HL-LHC track reconstruction (see Table 7.1) with track-$|\eta| < 2.5$. This selection was chosen to compare HL-LHC performance to Run 2 as directly as possible.

2. **Forward-tracking configuration:** Baseline HL-LHC track reconstruction with track-$|\eta| < 4.0$. This selection was chosen to see if any gains were observed when using the full tracking capabilities of the ITk.

3. **Low-p_{T} configuration:** Track-$p_{\mathrm{T}} > 500$ MeV, track-$|\eta| < 2.5$. The selection attempts to improve background suppression compared to the nominal

HL-LHC configuration by reducing the reconstructed track-p_{T} threshold. As low-p_{T} tracking in the forward region is not well understood, the pseudorapidity range is limited to the Run 2 acceptance in this configuration.

The $n_{\mathrm{trk}} = 0$ requirement is sensitive to the longitudinal size of the beamspot, as well as track reconstruction efficiencies and thresholds. It is predicted that the HL-LHC beamspot will vary between 40–48 mm [233] and the simulated beamspot size in the HL-LHC MC has $\sigma_{\mathrm{BS}} = 50\,\mathrm{mm}$. All reconstructed tracks are required to have a transverse impact parameter $d_0 < 1.0\,\mathrm{mm}$ to remove those originating from secondary interactions.

The expected signal yield at the HL-LHC, $N_{\mathrm{signal}}^{\mathrm{HL\text{-}LHC}}$, can be calculated as a function of the measured Run 2 yields, $N_{\mathrm{signal}}^{\mathrm{Run\ 2}}$, given in Table 5.7, considering the ratio of various efficiencies related to operating conditions and detector performance between the HL-LHC and Run 2:

$$N_{\mathrm{signal}}^{\mathrm{HL\text{-}LHC}} = N_{\mathrm{signal}}^{\mathrm{Run\ 2}} \times \frac{\mathcal{L}_{\mathrm{HL\text{-}LHC}}}{\mathcal{L}_{\mathrm{Run\ 2}}} \times \frac{\varepsilon_{\mathrm{excl.}}^{\mathrm{HL\text{-}LHC}}}{\varepsilon_{\mathrm{excl.}}^{\mathrm{Run\ 2}}} \times \frac{\varepsilon_{\gamma\gamma WW}^{\mathrm{HL\text{-}LHC}}}{\varepsilon_{\gamma\gamma WW}^{\mathrm{Run\ 2}}} \times \frac{S_{\mathrm{excl}}^{\mathrm{HL\text{-}LHC}}}{S_{\mathrm{excl}}^{\mathrm{Run\ 2}}}, \tag{7.1}$$

where $\mathcal{L}$ is the integrated luminosity, $\varepsilon_{\mathrm{excl.}}$ is the exclusive efficiency as defined in Section 5.4.1, $\varepsilon_{\gamma\gamma WW}$ is the efficiency of passing the event selections (Table 5.2) with the omission of the $n_{\mathrm{trk}} = 0$ requirement, and S_{excl} is the exclusive scale factor discussed in Section 5.6.1. The integrated luminosity is predicted to be 3000–4000 fb^{-1} at the end of HL-LHC data-taking [214]. In this study, $\mathcal{L}_{\mathrm{HL\text{-}LHC}}/\mathcal{L}_{\mathrm{Run\ 2}}$ is taken to equal 20. The exclusive efficiency is highly dependent on pileup environment and the track-veto window size for which the $n_{\mathrm{trk}} = 0$ requirement is applied, and will be discussed in detail in Section 7.4.1. The efficiency of passing the remaining event selections, $\varepsilon_{\gamma\gamma WW}$, is expected to be very close to the Run 2 value; when varying the size of the z-window considered for the exclusivity requirement, an additional correction is considered measuring the probability for both leptons to fall within the z-window of interest. This requirement is more than 98 % efficient for all window sizes investigated, with the exception of the smallest window size of 0.1 mm, where the efficiency drops to 94 %. Therefore, $\varepsilon_{\gamma\gamma WW}$ is a minor consideration, even at the smallest track-veto window size. It is assumed that the ratio of elastic to dissociative signal will not change between Run 2 and the HL-LHC and therefore the ratio of S_{excl} is also expected to be very close to one. This assumption is largely appropriate other than when increasing the tracking acceptance from $|\eta| < 2.5$ to $|\eta| < 4$, discussed briefly in Section 7.4.3.

The background extrapolation contains many of the same considerations as the signal and is given by

$$N_{\text{bkgd.}}^{\text{HL-LHC}} = N_{\text{bkgd.}}^{\text{Run 2}} \times \frac{\mathcal{L}_{\text{HL-LHC}}}{\mathcal{L}_{\text{Run 2}}} \times \frac{\varepsilon_{\text{excl.}}^{\text{HL-LHC}}}{\varepsilon_{\text{excl.}}^{\text{Run 2}}} \times \frac{\varepsilon_{qqWW}^{\text{HL-LHC}}}{\varepsilon_{qqWW}^{\text{Run 2}}}. \tag{7.2}$$

The background efficiency, ε_{qqWW}, is dominated by the probability of the inclusive $q\bar{q} \to WW$ background having no reconstructed tracks originating from the underlying event. The background process can pass the $n_{\text{trk}} = 0$ requirement if additional tracks from the UE are either too soft or too forward to be reconstructed, i.e. if additional tracks are below the minimum-p_{T} threshold or outside of the detector acceptance.

When extrapolating the Run 2 background yields using Equation (7.2), it is implicitly assumed that the total background yield scales in the same way as the $q\bar{q} \to WW$ yield which is only true for inclusive processes like $qq \to Z/\gamma^* \to \ell\ell$ which are expected to have a similar underlying event to the $q\bar{q} \to WW$ process. Exclusive backgrounds like $\gamma\gamma \to \ell\ell$ are not affected by changes in the ε_{qqWW} efficiency due to a lack of underlying event. This assumption is motivated by the fact that $q\bar{q} \to WW$ is the dominant background, making up 75 %[13] of the Run 2 background yield in the SR, with exclusive backgrounds contributing around 10 %. The remaining backgrounds include Drell-Yan and other diboson processes that scale in a similar way to $q\bar{q} \to WW$.

The post-fit expected signal and background yields in the Run 2 analysis are $N_{\text{signal}}^{\text{Run 2}} = 173.56$ and $N_{\text{bkgd.}}^{\text{Run 2}} = 131.7$, and the total systematic uncertainty on the background is taken as $\pm 12\,\%$ [4]. The various HL-LHC efficiencies in Equations (7.1) and (7.2) were calculated directly in the HL-LHC simulation, avoiding extrapolation from low to high-μ. This approach has the advantage that the event selections can be altered to find a new optimal analysis which maximises the expected signal significance. The terms with the largest impact on Equations (7.1) and (7.2) are discussed in more detail below.

7.4.1 Changes in the exclusive efficiency, $\varepsilon_{\text{excl.}}$

The exclusive efficiency is calculated using the data-driven method (Section 5.4.1) in fully-simulated HL-LHC MC and is calculated as a function of the track-veto

[13] The yields for $q\bar{q} \to WW$ also contain a small contribution from gluon-induced WW and VBS WW production.

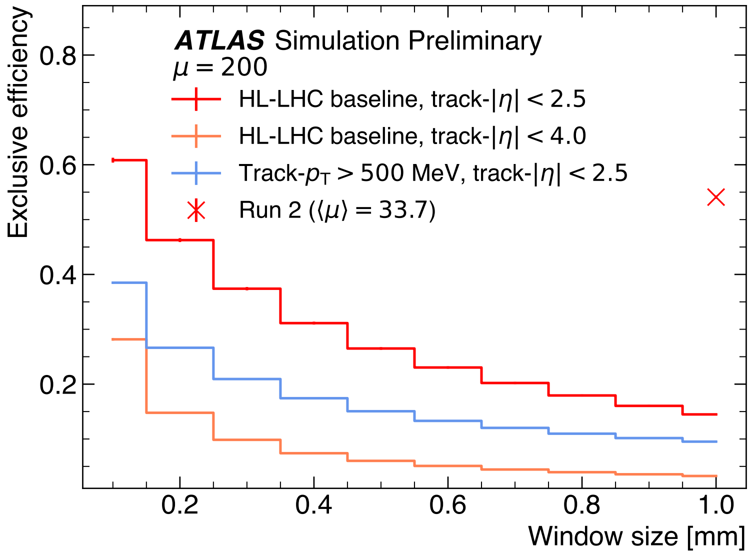

Figure 7.4: The exclusive efficiency, $\varepsilon_{\text{excl.}}$, as a function of window size and for the three different HL-LHC configurations investigated. The efficiency in MC in the Run 2 analysis [4] is marked by a red cross at the optimal Run 2 window size of 1 mm. Uncertainties on each point are statistical only and are generally too small to see. Plot produced by the author for Ref. [5].

window size. The efficiency is shown in Figure 7.4 for the three different HL-LHC tracking configurations. Each configuration is compared to the average exclusive efficiency calculated in Run 2 MC at a window size of ± 1 mm. The smaller the track-veto window, the smaller the probability that a pileup track falls within it. Therefore, the probability of the exclusive signal passing the $n_{\text{trk}} = 0$ requirement increases as the window size decreases. The exclusive efficiency is lower than the Run 2 value at a window size of 1 mm for all configurations investigated due to the much denser pileup environment at the HL-LHC.

For the baseline configuration, the efficiency at a window size of ± 1 mm is roughly 15 %, compared directly to the Run 2 efficiency measured in MC of 54.1 % (see also Table 5.5). This is an almost factor of four drop in signal efficiency and suggests that the analysis strategy must be re-optimised. However, the drop in exclusive efficiency also impacts the background yield so $N_{\text{sig}}/N_{\text{bkgd}}$ will not be impacted but $N_{\text{sig}}/\sqrt{N_{\text{bkgd}}}$ will be reduced. For this reason, other tracking configurations are also investigated. Increasing the tracking acceptance from $|\eta| < 2.5$ to $|\eta| < 4.0$ increases the probability of reconstructing pileup tracks and therefore decreases the exclusive efficiency relative to the baseline $|\eta| < 2.5$ configuration. This is compounded by the fact that tracks in the forward region have worse z_0 resolution and may be smeared in the track-veto window more often. Similarly, reconstructing tracks with lower p_{T} also increases the number of pileup tracks that are reconstructed within the track-veto window, decreasing the exclusive efficiency.

If the exclusive efficiency, as shown in Figure 7.4, was the only consideration, the optimal analysis (out of those investigated) would be one that maximises the exclusive efficiency and therefore maximises $N_{\mathrm{sig}}/\sqrt{N_{\mathrm{bkgd}}}$. This would be the baseline HL-LHC reconstruction with track-$|\eta| < 2.5$ and with an infinitesimal window size.

7.4.2 Changes in the background efficiency, ε_{qqWW}

The background efficiency, ε_{qqWW}, is the probability that both lepton z_0-positions fall within the track-veto window and that the inclusive $q\bar{q} \to WW$ background passes the $n_{\mathrm{trk}} = 0$ requirement. The background efficiency considers tracks originating from the $q\bar{q} \to WW$ vertex only. These tracks are determined by matching reconstructed tracks to true charged particles in the generator record, with an associated truth-match probability > 0.5. This is done to disentangle the efficiency of the background passing the $n_{\mathrm{trk}} = 0$ selection from the exclusive efficiency, which is a consideration of pileup. When making measurements in real data, it is not possible to determine if a track originates from the underlying event and it is instead only $\varepsilon_{qqWW} \times \varepsilon_{\mathrm{excl.}}$ that can be measured around the dilepton vertex.

This background efficiency decreases as the track-veto window size increases, as shown in Figure 7.5. As the background efficiency is defined to be independent of pileup, it is expected that the efficiency should be similar between Run 2 and the HL-LHC if the detector acceptance and performance is similar between the two. This can be seen in Figure 7.5 by comparing the track-$p_{\mathrm{T}} > 500\,\mathrm{MeV}$, track-$|\eta| < 2.5$ configuration to the Run 2 result at a window size of 1 mm. This HL-LHC configuration was chosen to closely match the Run 2 configuration. Differences are due to the different detector reconstruction performance in the Run 2 and HL-LHC simulation.

The background efficiency starts to increase rapidly at the smallest window sizes, due to tracks from the underlying event being smeared outside of the track-veto window when the window size is comparable to the track z_0 resolution. Increasing the η-range of reconstructed tracks from $|\eta| < 2.5$ to $|\eta| < 4$ reduces the background efficiency due to being able to reconstruct additional UE tracks

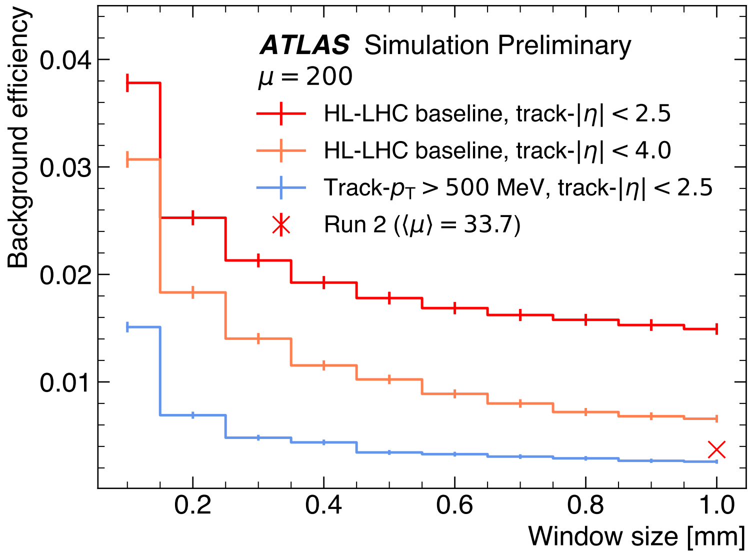

Figure 7.5: The efficiency of the dominant $q\bar{q} \to WW$ background passing the $n_{\mathrm{trk}} = 0$ requirement, once all other event selections have been applied. The efficiency for the three different HL-LHC configurations is shown as a function of the track-veto window size and compared directly to the result in Run 2 which is marked by a red cross. Uncertainties on each point are statistical only. Analysis specific corrections, such as the charged particle multiplicity correction described in Ref. [4], are assumed to cancel between Run 2 and the HL-LHC. Plot produced by the author for Ref. [5].

that are produced in the forward region. However, the largest reduction in the background efficiency is seen when the minimum track-p_{T} threshold can be reduced to 500 MeV from the baseline reconstruction given in Table 7.1. This is because UE tracks are produced from soft interactions, usually with very low-p_{T}. By lowering the track p_{T} threshold, more of these UE tracks can be reconstructed and it is possible to better reject backgrounds like $q\bar{q} \to WW$. If the background efficiency was the only consideration in this projection, the optimal setup would be to reconstruct tracks low in p_{T} (whilst avoiding large increases in fake-rates) and have a window size that covered the entire beamspot. Unfortunately, this is in direct contradiction with the optimal setup when only considering the exclusive efficiency in Figure 7.4. Hence, the analysis must be optimised to find a compromise between these two considerations.

When considering the differential analysis in $m_{\ell\ell}$, a fixed window size is used. It is assumed that all efficiencies affecting the signal, as given in Equation (7.1), are constant at any fixed window size and are independent of $m_{\ell\ell}$. However, ε_{qqWW} decreases with increasing $m_{\ell\ell}$ due to an increase in UE activity at higher partonic centre-of-mass energies. The number of events generated for the HL-LHC samples was not large enough to probe this effect within the statistical uncertainties. Therefore, the ratio of the efficiency at a given window size was taken to the Run 2 efficiency in the inclusive analysis and used to scale the Run 2 efficiencies as a function of $m_{\ell\ell}$. This is demonstrated in Figure 7.6, comparing the nominal ± 1 mm track-veto window in Run 2 (red) to a window size of 0.2 mm at the HL-LHC (blue).

Systematic uncertainties were assumed to remain unchanged between the HL-LHC and Run 2, where the largest source of systematic uncertainty was due to background modelling of the underlying event, giving an uncertainty on the background yield of $\pm 12\,\%$. This is a pessimistic approach as it is likely that modelling uncertainties will be reduced before the start of HL-LHC data-taking and analysis. Therefore reduced background systematic uncertainties of $6\,\%$ and $3\,\%$ were also considered and compared to the nominal projection.

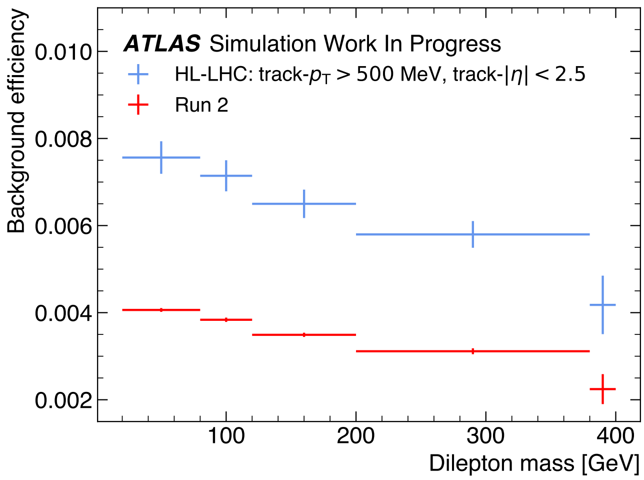

Figure 7.6: The background efficiency, ε_{qqWW}, as a function of dilepton mass, $m_{\ell\ell}$, shown for tracks with $p_{\mathrm{T}} > 500$ MeV and $|\eta| < 2.5$ at a window size of 0.2 mm at the HL-LHC and compared to nominal Run 2 efficiencies (red). The final bin includes overflow. No charged track multiplicity reweighting is applied in the calculation of these efficiencies.

7.4.3 Exclusive scale factor considerations

The exclusive scale factor is a measure of the ratio of elastic to dissociative $\gamma\gamma \rightarrow WW$ signal events, described in Section 5.6. It is assumed that the ratio of elastic to dissociative signal yield remains constant between Run 2 and the HL-LHC, when considering the same tracking acceptance of $|\eta| < 2.5$ and placing no requirement on proton tags with AFP. However, this assumption was reconsidered when investigating the forward tracking configuration at the HL-LHC as tracks from proton dissociation are likely to be produced in the forward region and may be reconstructed with the increased ITk tracking acceptance of $|\eta| < 4$.

An acceptance study was performed at particle-level in Run 2 single-dissociative $\gamma\gamma \to ll$ simulation modelled by LPAIR, as no reliable dissociative $\gamma\gamma \to WW$ samples existed at the time of the study. The charged particle η distribution is presented in Figure 7.7, demonstrating that more tracks from proton dissociation are likely to be reconstructed with the increased acceptance of the ITk compared to the Run 2 ID acceptance. This will result in the dissociative signal failing the $n_{\text{trk}} = 0$ requirement more often than in Run 2, causing a reduction in the exclusive scale factor. To estimate this reduction, ratios of the number of events with $n_{\text{ch}} > 0$ within the tracking acceptance to the total number of simulated events were calculated for both $|\eta| < 2.5$ and $|\eta| < 4.0$ and when considering charged particles with $p_{\text{T}} > 500\,\text{MeV}$ and with $p_{\text{T}} > 900\,\text{MeV}$. The exclusive scale factor at the HL-LHC when considering the full tracking acceptance of the ITk of $|\eta| < 4$ was calculated to be in the range of $S_{\text{excl}}^{\text{HL-LHC}} = 0.92 \times S_{\text{excl}}^{\text{Run 2}}$ to $0.93 \times S_{\text{excl}}^{\text{Run 2}}$.

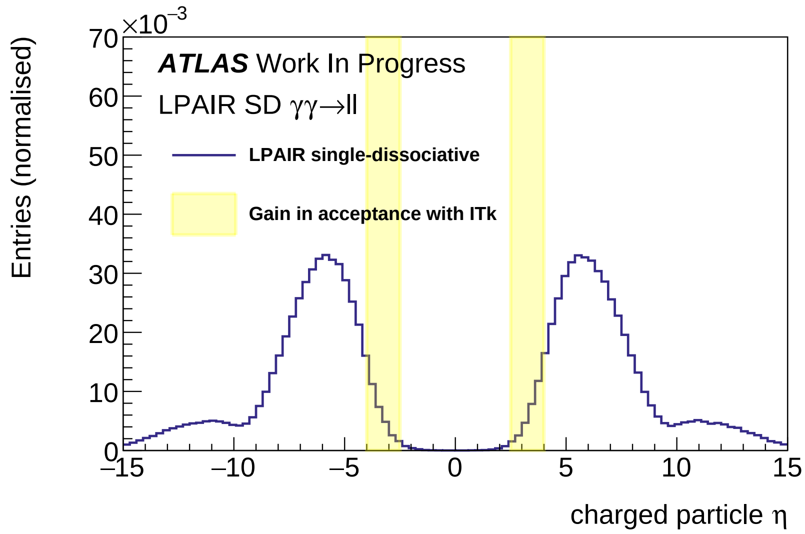

Figure 7.7: The pseudorapidity of charged particles in single-dissociative $\gamma\gamma \to ll$ events simulated with LPAIR. Charged particles originating from pileup are not considered, leaving the central region devoid of activity. LPAIR models the proton dissociation on both sides of the detector, resulting in two large peaks around $|\eta| = 6$. Smaller peaks are seen around $|\eta| = 11$ from the remaining intact proton. Highlighted in yellow is the additional tracking acceptance foreseen with the ITk at the HL-LHC that was not available in Run 2.

There are a few caveats with this method. The first is that only single-dissociative events have been considered whereas S_{excl} accounts for both single- and double-dissociative events. However, double-dissociative events are more likely to have additional tracks in the central region and therefore S_{excl} is dominated by the single-dissociative contribution. Secondly, and more importantly, the changes in the exclusive scale factor could only be calculated at particle-level. Tracking inefficiencies and resolution effects mean that not all of the charged particles will be reconstructed, and those that are may fall outside of the track-veto window or fail the kinematic requirements at selection boundaries. Therefore, it is expected that the true difference in S_{excl} will be smaller than the differences calculated above. It was not possible to perform a more detailed study on the changes in S_{excl} and a value of $S_{\text{excl}}^{\text{HL-LHC}}/S_{\text{excl}}^{\text{Run 2}} = 1$ is assumed for all configurations. The difference in the ratio of measured elastic to dissociative signal when requiring forward protons tags is discussed along with the AFP event selection in Section 7.6.

7.5 Inclusive fiducial measurement without AFP

The section reports results on the expected sensitivity for a fiducial cross-section analysis as a function of track-veto z-window size, with the aim of finding the optimal tracking configuration for a HL-LHC $\gamma\gamma \to WW$ analysis without forward proton tags. Results are presented as both the absolute expected signal and background yield, as well as the relative uncertainties on the signal yield. All results are compared to those in the nominal Run 2 analysis [4], which had a track-veto window size of 1 mm and reconstructed tracks with $p_{\text{T}} > 500\,\text{MeV}$ and $|\eta| < 2.5$.

The metric used to quantify the expected sensitivity at the HL-LHC is the relative uncertainty on the extracted signal yield after selections and background subtraction. The statistical-only uncertainty is then defined as:

$$\sigma_{\text{stat.}}/N_{\text{sig.}} = \sqrt{N_{\text{sig.}} + N_{\text{bkgd.}}}/N_{\text{sig.}}, \tag{7.3}$$

while the total relative uncertainty on the extracted signal yield is defined as

$$\sigma_{\text{tot.}}/N_{\text{sig.}} = \sqrt{N_{\text{sig.}} + N_{\text{bkgd.}} + (xN_{\text{bkgd.}})^2}/N_{\text{sig.}}, \tag{7.4}$$

where x represents the fractional systematic on the background yield, nominally taken to be 0.12 as in the Run 2 analysis.

The expected signal and background yields after selections are calculated separately for the three different HL-LHC configurations listed in Section 7.4. The results are shown in Table 7.3, together with the expected statistical-only and total relative uncertainty on the extracted signal, calculated using Eqs. (7.3) and (7.4) respectively. The Run 2 yields are also shown with uncertainties calculated in the same way. While for Run 2, it was better to have a window width that was many times the tracking resolution in order to consider tracks from the background UE, the denser pileup environment at the HL-LHC results in the optimal window being narrower to reduce signal losses. When considering statistical uncertainties only, the optimal window size is 0.2 mm but becomes wider when introducing systematic uncertainties on the large background yields. The effect of a lower reconstructed track-p_{T} threshold reduces the exclusive efficiency in pileup simply due to more pileup tracks being reconstructed, which in turn reduces both signal and background yields. However, a lower threshold also largely reduces the background efficiency to pass the $n_{\mathrm{trk}} = 0$ requirement due to the ability to reconstruct more tracks from the UE, which generally have very low-p_{T}. Therefore, lowering the track-p_{T} threshold from the HL-LHC baseline to 500 MeV improves the signal to background ratio.

When considering the effect of systematics, Figures 7.8a–7.8d show the relative uncertainty on the signal yield as a function of track-veto window size. The

Configuration	$N_{\mathrm{sig.}}$	$N_{\mathrm{bkgd.}}$	$\sigma_{\mathrm{stat.}}/N_{\mathrm{sig.}}$	$\sigma_{\mathrm{tot.}}/N_{\mathrm{sig.}}$
Run 2	174	132	0.10	0.14
1 mm window				
HL-LHC baseline, track-$\lvert\eta\rvert < 2.5$	929	2 840	0.07	0.37
HL-LHC baseline, track-$\lvert\eta\rvert < 4.0$	209	281	0.11	0.19
Track-$p_{\mathrm{T}} > 500$ MeV, track-$\lvert\eta\rvert < 2.5$	611	323	0.05	0.08
0.2 mm window				
HL-LHC baseline, track-$\lvert\eta\rvert < 2.5$	2 930	15 300	0.05	0.63
HL-LHC baseline, track-$\lvert\eta\rvert < 4.0$	934	3 560	0.07	0.46
Track-$p_{\mathrm{T}} > 500$ MeV, track-$\lvert\eta\rvert < 2.5$	1 684	2 410	0.04	0.18

Table 7.3: The expected number of signal and background events, the expected statistical-only and the total relative uncertainty on the extracted signal are shown for the three HL-LHC tracking configurations. A luminosity of $3000\,\mathrm{fb}^{-1}$ and a background systematic of 12 % are assumed. The HL-LHC baseline track reconstruction is given in Table 7.1.

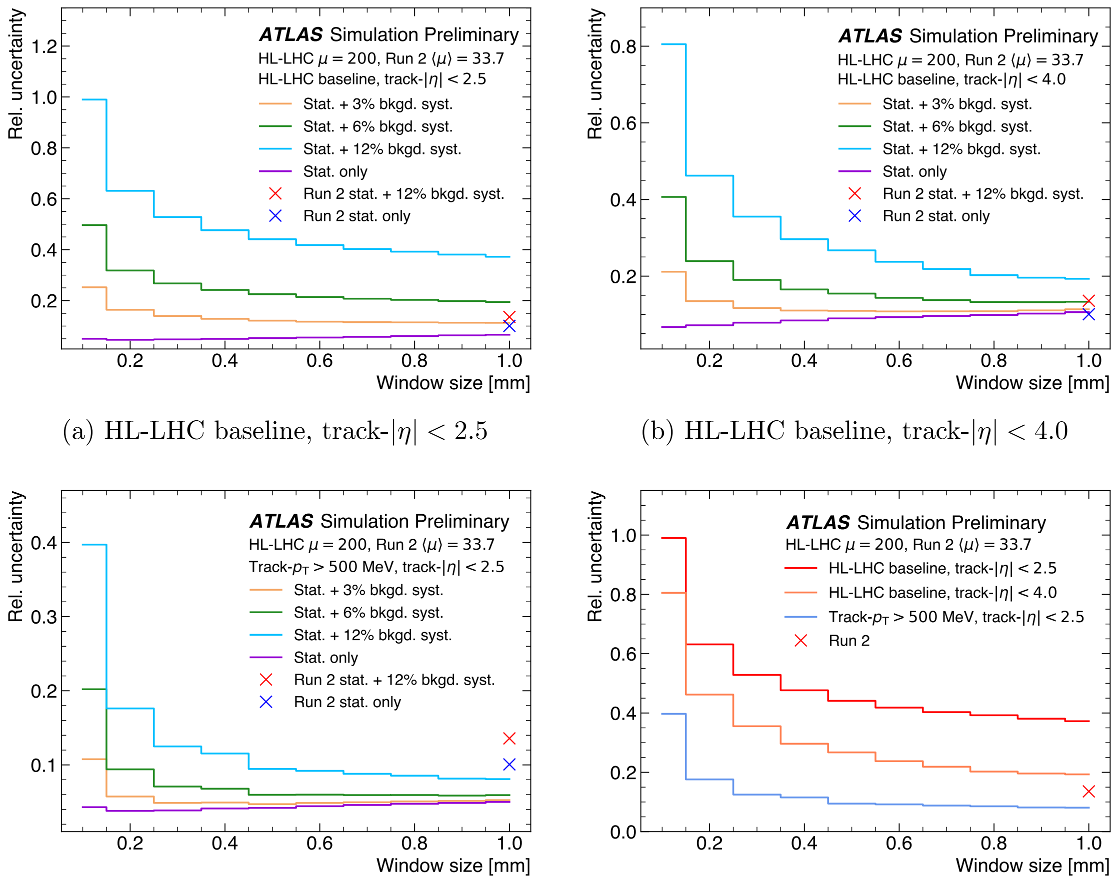

(a) HL-LHC baseline, track-$|\eta| < 2.5$

(b) HL-LHC baseline, track-$|\eta| < 4.0$

(c) HL-LHC track-$p_{\mathrm{T}} > 500\,\mathrm{MeV}$, track-$|\eta| < 2.5$

(d) Comparison of nominal total uncertainty

Figure 7.8: A comparison of the relative uncertainty on the signal for (a) the baseline HL-LHC track reconstruction, when considering the central detector only ($|\eta| < 2.5$), (b) the baseline HL-LHC track reconstruction with an extended tracking acceptance of $|\eta| < 4.0$, (c) low-p_{T} tracking ($p_{\mathrm{T}} > 500\,\mathrm{MeV}$) and $|\eta| < 2.5$, (d) a comparison of all setups when considering a nominal systematic on the background of 12 %. Each plot shows the relative uncertainty for different sizes of background systematics. Plots produced by the author for Ref. [5].

total uncertainty, as given in Equation (7.4), is reduced relative to the Run 2 uncertainty for many window sizes, when considering a background systematic of $\pm$ 12 % and reconstructing tracks with $p_{\mathrm{T}} > 500\,\mathrm{MeV}$, shown in Figure 7.8c. This is not the case if the HL-LHC baseline track reconstruction is used, shown in Figure 7.8a. Figures 7.8a–7.8c also show the impact of reducing the background modelling uncertainty. A lower background systematic uncertainty is seen as a key requirement to make the measurement competitive when the large HL-LHC dataset will be available.

The analysis was repeated with an extended tracking acceptance from the ITk of $|\eta| < 4.0$ when using the baseline HL-LHC track reconstruction and results are shown in Figure 7.8b. The signal yield when using this configuration is reduced to 32 % (55 %) of the $|\eta| < 2.5$ yield for a 0.2 mm window and baseline (low-p_{T}) track reconstruction, as shown in Table 7.3. The background yield in this configuration is reduced to 23 % of the $|\eta| < 2.5$ background yield when reconstructing tracks with baseline p_{T} requirements. This is partly due to the increased tracking acceptance making it more likely to reconstruct tracks from the $q\bar{q} \to WW$ hard-scatter but also due to the poor tracking z_0 resolution in the forward-η region. The degrading resolution with increasing track-η causes an increase in the number of pileup tracks not associated to the hard-scatter that migrate into the track-veto window, providing additional background rejection but also causing large signal losses. However, the background yield is still increased by almost 50 % for this configuration compared to the track-$p_{\mathrm{T}} > 500\,\mathrm{MeV}$, track-$|\eta| < 2.5$ configuration. This is largely due to the increased ITk p_{T} reconstruction thresholds.

For the HL-LHC baseline track reconstruction, an analysis that uses the forward-tracking region would likely bring slight gains in $\gamma\gamma \to WW$ signal significance compared to an analysis only performed in the $|\eta| < 2.5$ region, especially when considering a background systematic of 12 %. However, the reduction in the background yield caused by the increased tracking acceptance is not enough to compensate the large background yields accumulated from the increase in minimum track-p_{T} compared to Run 2. This results in a relative uncertainty on the signal that is still almost roughly 1.4 times the size of the relative uncertainty on the Run 2 result at the nominal ± 1 mm window. The only way to improve on the relative Run 2 signal uncertainty with baseline tracking is to reduce the background systematics by at least a factor of 2. Reducing the background yield by a factor of 2–4 would have a similar effect, although it is extremely difficult to do this without also losing signal, as has been demonstrated. The largest improvement is seen for the low-p_{T} configuration, which reduces the total uncertainty on the signal to less than 60 % of the Run 2 value for the nominal 1 mm window. Figure 7.8d compares the total uncertainty of the three different HL-LHC configurations when considering the nominal background systematic of 12 % and shows that the configuration in which tracks can be reconstructed with $p_{\mathrm{T}} > 500\,\mathrm{MeV}$ is optimal for all window sizes investigated.

It is assumed that $S_{\mathrm{excl}}^{\text{HL-LHC}}/S_{\mathrm{excl}}^{\text{Run 2}} = 1$ for all configurations. This is not necessarily an appropriate assumption for the forward tracking configuration with $|\eta| < 4$. From the particle-level acceptance studies presented in Section 7.4.3, it can be expected to lose up to an additional 8 % of the signal yields quoted above for this configuration. The background yields are unaffected by this correction.

Changes in trigger thresholds and efficiencies were not considered in these studies. The offline p_{T} threshold for single-muon and single-electron triggers was 27 GeV in Run 2. This is expected to be lowered to 20 GeV for single-muon and 22 GeV for single-electron triggers at the HL-LHC [222]. The dilepton electron-muon trigger is expected to have offline p_{T} threshold of 10 GeV for both objects. However, in the Run 2 observation the improved efficiency of the single-lepton triggers outweighed the gain from lower offline p_{T} thresholds. There are significant improvements expected in muon trigger efficiencies due to increased resistive plate chamber coverage in the barrel region of the MS, with a predicted efficiency increase from roughly 70 % in Run 2 to 90 % at the HL-LHC for $|\eta| < 1.05$ [214]. Combined-muon acceptance will be extended from $|\eta| = 2.5$ to 2.7. If any trigger prescales were to be introduced, this would increase the statistical uncertainty on the measurement.

Also not considered in these studies is the expected increase in cross-section going from $\sqrt{s} = 13\,\mathrm{TeV}$ to $\sqrt{s} = 14\,\mathrm{TeV}$. This increase is expected to be roughly 10 % for both the elastic $\gamma\gamma \to WW$ and $q\bar{q} \to WW$ processes and would therefore improve $N_{\mathrm{sig}}/\sqrt{N_{\mathrm{bkgd}}}$.

The conclusion of this section is that it will be possible to perform a measurement of the fiducial $\gamma\gamma \to WW$ cross-section at the HL-LHC with baseline track reconstruction but it will be challenging to improve on the Run 2 precision without improved background suppression or improved background modelling leading to reduced background systematics. It is not unreasonable to assume that modelling will improve before and during HL-LHC data-taking but modelling systematics of the UE will have to be reduced by over a factor of 2 before the relative uncertainty on the signal is halved compared to Run 2, in the best case scenario. It has been demonstrated in this study that dedicated low-p_{T} tracking below the baseline ITk reconstruction will be crucial for the $\gamma\gamma \to WW$ analysis at the HL-LHC and therefore this is an important area of research and development that should be pursued.

One possibility to improve background suppression is to measure the forward protons associated to the $\gamma\gamma \to WW$ process in the AFP detector. The only way the $q\bar{q} \to WW$ background can pass selections requiring forward protons is due to combinatorial backgrounds arising from pileup diffractive processes. The projection above is now extended to include forward proton tagging in the following sections, considering an analysis differentially in dilepton mass, $m_{\ell\ell}$, which is one of the measures of energy scale currently being investigated in the ongoing Run 2 differential measurement (Chapter 6). Requiring selections on forward protons means that the analysis can then be divided into categories: one in which no selection on forward protons is made (a zero-tag analysis), one in which a proton-tag is required on both sides of AFP (a double-tag analysis) and one in which a proton-tag is only required on one side of AFP (a single-tag analysis). The additional AFP event selections are discussed in the next section, before presenting the differential event yield projection.

7.6 AFP event selection

This section considers the impact of using forward protons to further suppress backgrounds, with a detector similar to the ATLAS Forward Proton (AFP) detector discussed in Chapter 4, assuming detector locations and optics at the HL-LHC as discussed in Section 7.2.3. Section 7.6.2 describes the procedure of requiring a double proton-tag in the AFP acceptance and the efficiency of this requirement in the context of a $\gamma\gamma \to WW$ signal. Section 7.6.3 describes the procedure of requiring a single proton-tag in AFP in conjunction with one of the signal leptons being within the HGTD acceptance, a configuration never before discussed. The results of a proton-tagged $\gamma\gamma \to WW$ analysis projection are presented differentially in $m_{\ell\ell}$ along with the central-detector-only differential projection in Section 7.7.

The efficiency of AFP selections is independent of the central detector and therefore the expected yields at the HL-LHC when including forward proton selections are calculated as

$$N^{\text{HL-LHC}}_{\text{signal}} \to N^{\text{HL-LHC}}_{\text{signal}} \times \varepsilon^{\gamma\gamma WW}_{\text{AFP}} \quad \text{and} \quad N^{\text{HL-LHC}}_{\text{bkgd}} \to N^{\text{HL-LHC}}_{\text{bkgd}} \times \varepsilon^{\text{bkgd}}_{\text{AFP}}, \tag{7.5}$$

where $\varepsilon^{\gamma\gamma WW}_{\text{AFP}}$ and $\varepsilon^{\text{bkgd}}_{\text{AFP}}$ are the efficiencies of the signal and background to pass the AFP selection requirements, respectively. The dominant $q\bar{q} \to WW$ background does not have prompt forward protons in the final state. Non-prompt protons may be formed during hadronisation in the underlying event but these usually have ξ and p_{T} such that they do not reach the AFP detectors. Therefore, the $q\bar{q} \to WW$ background will only pass the AFP selections if it coincides with prompt forward protons that pass the AFP selections, originating from diffractive pileup processes in the same or nearby bunch crossings.

Protons in HL-LHC MC are used to calculate the efficiency of a proton-tag analysis and are considered if they have energy loss within the AFP acceptance. Particle-level protons originating from both signal and pileup interactions are included in the MC. The proton fractional energy loss, ξ, is calculated as $\xi = 1 - E_{\text{proton}}/E_{\text{beam}}$. In these studies, a CoM energy of 14 TeV is assumed and therefore $E_{\text{beam}} = 7\,\text{TeV}$. The AFP-ToF detectors (Section 4.2.3) are assumed to be operational with 100 % reconstruction efficiency and proton timing resolutions of $\sigma_{\text{ToF}} = 20\,\text{ps}$ (original design goal) and $\sigma_{\text{ToF}} = 10\,\text{ps}$ (HL-LHC goal) are studied. It is assumed that the transverse momentum of the proton can be measured with a resolution of 0.1 GeV when applying selections on forward proton p_{T}, discussed in the next section, but no resolution effects are considered when measuring the proton ξ.

7.6.1 Requiring small forward proton p_{T}

Requiring small forward proton p_{T} in a proton-tagged analysis helps to further suppress the combinatorial background, as the protons originate from diffractive processes where the momentum transfer is expected to be larger than in photon-induced processes, resulting in larger forward proton p_{T}. Figure 7.9 shows the particle-level proton-p_{T} in the signal HL-LHC sample, including pileup protons. The pileup proton p_{T} distribution is harder than the hard-scatter proton p_{T}. It can be seen that only selecting forward protons below some p_{T} threshold will suppress a large amount of the combinatorial background due to pileup. Several potential p_{T} selections were investigated in the context of a double-tag measurement. Table 7.4 shows the additional background reduction factor and signal efficiency of different

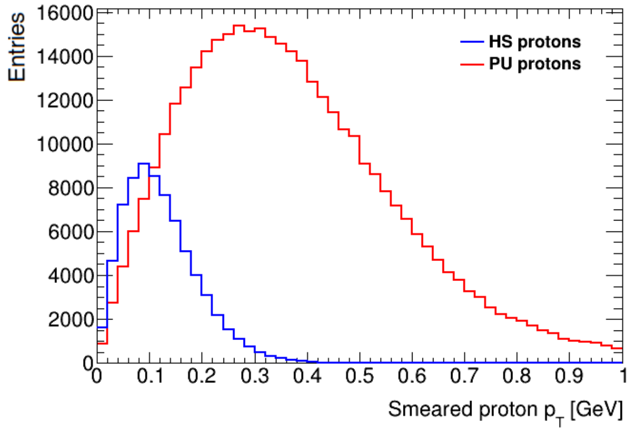

Figure 7.9: Particle-level forward proton-$p_{\rm T}$ for $\gamma\gamma \to WW$ signal (HS, blue) and pileup (PU, red) protons. The proton-$p_{\rm T}$ is smeared assuming a Gaussian resolution of 0.1 GeV. Each histogram bin shows the absolute number of events per 20 MeV in the simulation sample with no selection on ξ.

proton-$p_{\rm T}$ requirements, when assuming Run 2 AFP acceptance ($0.02 < \xi < 0.12$) and an AFP-ToF resolution of 10 ps for a double-tag analysis category. From these assumptions, the optimal requirement for a double-tag analysis was selected to be proton-$p_{\rm T}$ < 0.25 GeV and for a single-tag analysis $p_{\rm T}$ < 0.30 GeV. The optimal requirement was chosen as the one that maximised $N_{\rm sig}/\sqrt{N_{\rm bkgd}}$. Note that this requirement is dependent on acceptance and may need to be re-optimised for more realistic ξ acceptances at the HL-LHC.

$p_{\rm T}$ requirement	background reduction factor	signal efficiency [%]
$p_{\rm T} < 0.35$ GeV	$\times 3$	99
$p_{\rm T} < 0.30$ GeV	$\times 5$	96
$p_{\rm T} < 0.25$ GeV	$\times 7$	90
$p_{\rm T} < 0.20$ GeV	$\times 15$	74

Table 7.4: The approximate impact of different proton $p_{\rm T}$ requirements on pileup background and $\gamma\gamma \to WW$ signal for a double-tag analysis, when assuming an AFP acceptance of $0.02 < \xi < 0.12$ and AFP-ToF resolution of 10 ps.

7.6.2 Double proton tag in elastic $\gamma\gamma \to WW$

If both protons remain intact after photon-emission, they can be tagged on each side of AFP. Under the assumption that the two protons are produced by the same pp interaction, measuring each proton's time-of-arrival in AFP, t_A and t_C, allows reconstruction of the interaction vertex,

$$z_{\rm ToF} = -\frac{c}{2}(t_A - t_C), \tag{7.6}$$

where it is appropriate to assume the protons are travelling at the speed of light, within the timing resolution. Assuming a reconstructed primary vertex from the leptonic W-decays in the central detector, $z^{\ell\ell}_{\mathrm{vtx}}$, one can then define

$$\Delta z = z^{\ell\ell}_{\mathrm{vtx}} - z_{\mathrm{ToF}}. \tag{7.7}$$

On average there will be multiple protons measured in each AFP station per event at $\mu = 200$. In this scenario, z_{ToF} must be used to select the protons that originate from the signal process and not from pileup. The Δz distribution of the signal forms a sharp peak related to the timing resolution, σ_{ToF}, whereas the background forms a much wider distribution related to the beamspot longitudinal width, σ_{BS}. This is demonstrated in Figure 7.10.

By requiring

$$|\Delta z| < 2 \times \frac{c}{\sqrt{2}} \sigma_{\mathrm{ToF}}, \tag{7.8}$$

95 % of the signal within the AFP acceptance can be retained whilst improving background rejection. In the case of multiple protons reconstructed per side of AFP, Δz should be calculated for every possible proton pair that passes

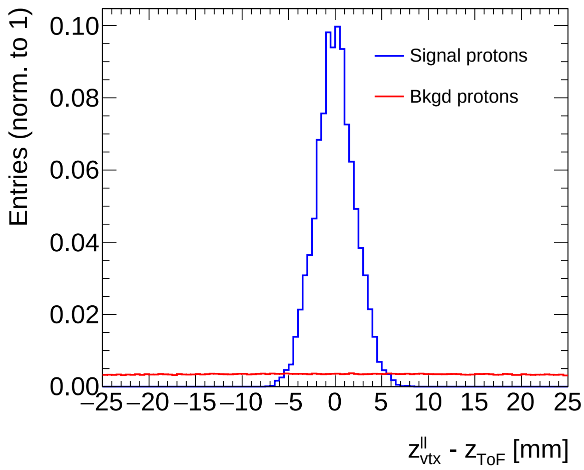

Figure 7.10: The difference between the signal vertex reconstructed from the forward proton time-of-flight to AFP, z_{ToF}, and from the lepton tracks in central ATLAS, $z^{\ell\ell}_{\mathrm{vtx}}$. The vertex reconstructed from the true signal protons originating from the $\gamma\gamma \to WW$ interaction (blue) follows a sharp peak centered on zero, whereas the vertices reconstructed from randomly pairing pileup protons (red) follows a flat distribution in Δz. All protons have $0.02 < \xi < 0.12$ and $p_{\mathrm{T}} < 0.25$ GeV and an AFP-ToF resolution of $\sigma_{\mathrm{ToF}} = 10$ ps is assumed. Both distributions are normalised to unity.

acceptance and p_{T} requirements, and the smallest value of Δz taken as the signal proton pair.

Despite the large fraction of retained events, the double-tag selection will only be possible for events in which both protons remain intact in the final state, causing the immediate 72 % loss of the dissociative contributions to the signal where one or both protons breaks apart after the photon emission. This is an intrinsic loss of efficiency compared to the non-AFP analysis.

Figure 7.11 shows the efficiency of requiring a double-tag for both the $\gamma\gamma \to WW$ signal and background due to pileup protons coinciding with the $q\bar{q} \to WW$ hard-scatter (PU background), as a function of reconstructed dilepton mass. Efficiencies for both sets of detector locations are compared and background efficiencies are compared for an AFP-ToF resolution of both 10 and 20 ps. The background efficiency does not depend on dilepton mass as the background arises from proton pairs that are uncorrelated to the reconstructed dilepton candidate, usually originating from single-diffractive pileup processes. Diffractive backgrounds produce forward protons which typically have higher p_{T} than photon-induced processes and therefore a proton p_{T} selection of $p_{\mathrm{T}} < 0.25$ GeV is also applied to further reject these backgrounds, as discussed in Section 7.6.1. The signal efficiency increases with increasing dilepton mass,

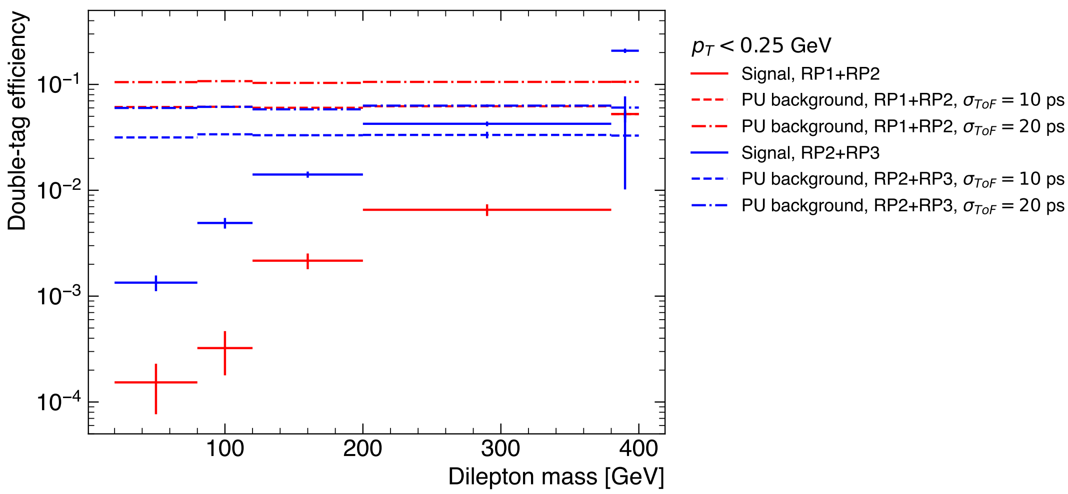

Figure 7.11: The efficiency of requiring a double-tag for both $\gamma\gamma \to WW$ signal (solid) and pileup background (dashed), comparing RP1+RP2 (red) and RP2+RP3 (blue). Background efficiencies are shown for an AFP-ToF resolution of both 10 ps and 20 ps. The forward proton p_{T} is required to be smaller than 0.25 GeV to further reject the pileup background.

as the protons in these events have lost more momentum in the scattering, making them more likely to be deflected into the AFP acceptance. High-mass regions are interesting from a physics perspective as this is where we expect to be most sensitive to any deviations from new physics effects at a high scale [208].

The double-tag signal efficiency is very low at low values of dilepton mass, at the order of 0.01–0.1 %. However, this rises quickly with increasing dilepton mass to 5–20 % for $m_{\ell\ell} > 380\,\text{GeV}$, depending on the detector locations. The signal efficiency is better for the combination of RP2+RP3 as this has lower ξ acceptance than RP1+RP2. RP2+RP3 also has lower efficiency of pileup protons passing the double-tag requirements due to the decreased upper bound in ξ. When considering a ToF resolution of 10 ps, the background double-tag efficiency is 6 % for RP1+RP2 and 2 % for RP2+RP3. It can be seen from Figure 7.11 that the signal double-tag efficiency is only higher than the background efficiency if the detector stations are located at RP2+RP3 and, even in this case, only at the highest $m_{\ell\ell}$ values. The expected signal and background yields for a HL-LHC $\gamma\gamma \to WW$ analysis when requiring a double-tag are summarised in Section 7.7.

7.6.3 Single proton tag in $\gamma\gamma \to WW$

The HGTD subdetector was introduced in Section 7.2.2. It is possible to match lepton timing information from the HGTD to forward proton timing information from AFP, to reconstruct the position of the candidate signal vertex and compare to the reconstructed dilepton vertex from the ITk. Once the lepton arrival time in the HGTD, $t_{\text{vtx.}}$, has been measured, the proton arrival time on both A- and C-side of AFP can be predicted:

$$t^{\text{A/C}}_{\text{pred.}} = \frac{1}{c}|z^{\text{A/C}}_{\text{AFP}} - z^{\ell\ell}_{\text{vtx}}| + t_{\text{vtx.}}, \tag{7.9}$$

where $z^{\text{A/C}}_{\text{AFP}}$ is the AFP detector location in z relative to IP1, with $z^{\text{C}}_{\text{AFP}} = -z^{\text{A}}_{\text{AFP}}$, assuming a negligible lepton ToF between the interaction point and the hit in the HGTD relative to the proton ToF to AFP. The measured proton arrival time in AFP, $t_{\text{meas.}}$, must be consistent with the predicted arrival time, within detector resolutions:

$$|t_{\text{meas.}} - t_{\text{pred.}}| < 2 \times \sqrt{\sigma^2_{\text{AFP-ToF}} + \sigma^2_{\text{HGTD}}}, \tag{7.10}$$

where it is assumed that the vertex z-resolution is negligible compared to ToF resolutions. Again, this requirement retains 95 % of the signal. Similarly to the double-tag, the forward protons associated to the signal form a peak centred on

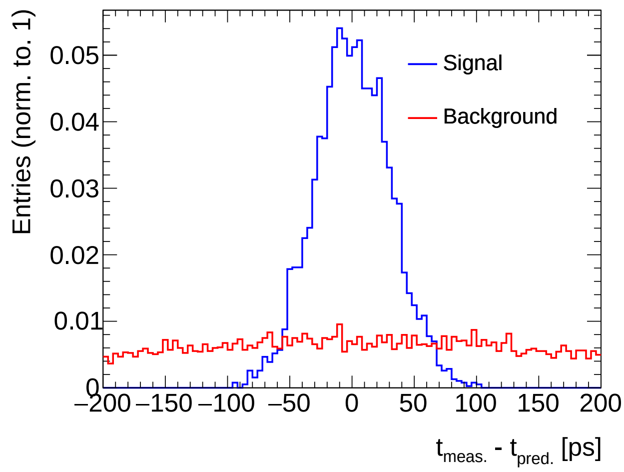

Figure 7.12: The difference in predicted proton time-of-flight to AFP, $t_{\text{pred.}}$, calculated from the lepton time-of-flight in the HGTD, to measured proton time-of-flight in AFP, $t_{\text{meas.}}$. Distributions are shown for the signal protons (blue) and pileup protons (red). Both distributions are normalised to unity.

zero, related to the ToF resolutions of AFP and the HGTD, whereas the pileup protons forming the combinatorial background have a much wider distribution. This is demonstrated in Figure 7.12.

In such an analysis, at least one of the leptons from the W-decays would have to be within the HGTD acceptance of $2.4 < |\eta| < 4.0$. This limits the signal acceptance but has the effect of selecting events with boosted topologies, in which an intact forward proton with the same sign of η as the forward lepton is more likely to be within the AFP acceptance. This effect is demonstrated in Figure 7.13. This means that by measuring a lepton in the HGTD, the side of AFP in which the forward proton is reconstructed can be predicted and any protons reconstructed on the opposite side rejected. Enforcing this selection halves the background from pileup in such a single-tag analysis. The fraction of signal due to elastic events in which both protons are within the AFP acceptance is small for such a topology. The leptons have sufficiently high p_{T} ($p_{\text{T}} > 20\,\text{GeV}$) that the ITk resolution is still acceptable for the lepton in the HGTD acceptance.

The elastic signal acceptance, A_{EE}, of AFP plus the HGTD can be written as a combination of the relative difference in lepton acceptance in the central detector, $A_{\ell\ell}$, between $|\eta| < 2.5$ and $2.4 < |\eta| < 4.0$, and the acceptance of

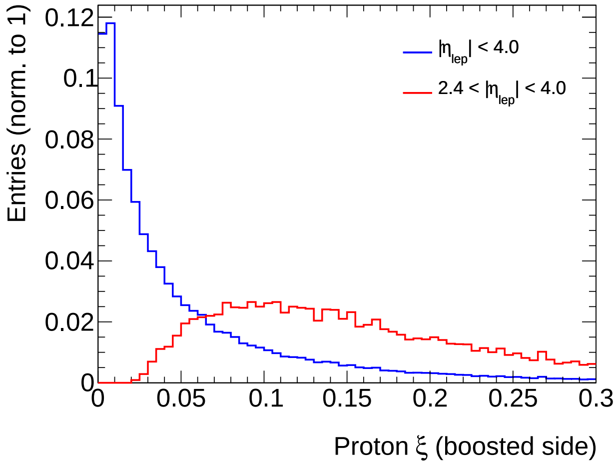

Figure 7.13: The proton fractional energy loss distribution for different single-lepton kinematics. The inclusive distribution for all leptons with $|\eta| < 4$ is shown in blue, where a sharp peak at very low ξ values mostly outside of the lower edge of the AFP acceptance can be seen. The distribution when at least one lepton is within $2.4 < |\eta| < 4$ is shown in red, demonstrating how the proton with the same sign in p_z as the forward lepton is pushed to higher ξ values and more likely to be within AFP acceptance. Both histograms are normalised to unity.

requiring a single forward proton in AFP, A_{AFP}:

$$A_{\text{EE}} = \frac{A_{\ell\ell}(2.4 < |\eta| < 4.0)}{A_{\ell\ell}(|\eta| < 2.5)} \times A_{\text{AFP}}. \tag{7.11}$$

Figure 7.14 shows the elastic $\gamma\gamma \to WW$ acceptance when requiring at least one lepton to be within $2.4 < |\eta| < 4.0$, relative to all signal candidates with both leptons in the range $|\eta| < 4.0$. The latter distributions show the effect of additionally requiring a single-proton to be in the AFP acceptance, on the same side as the lepton in the HGTD (i.e. A- or C-side). Interestingly, the upper edge of the AFP acceptance becomes important for a single-tag analysis, as the kinematics of the boosted system begin to push the forward proton out of the upper edge of the acceptance at higher dilepton masses. For this reason, RP1+RP2 performs slightly better than RP2+RP3 in a single-tag analysis. It is assumed that the lepton acceptance in Figure 7.14 applies to both the signal and dominant $q\bar{q} \to WW$ background.

The overall signal efficiency of the AFP single-tag selection can be calculated as

$$\varepsilon^{\gamma\gamma\to WW}_{\text{single-tag}} = 0.64 \times 0.95 \times A_{\text{EE}}, \tag{7.12}$$

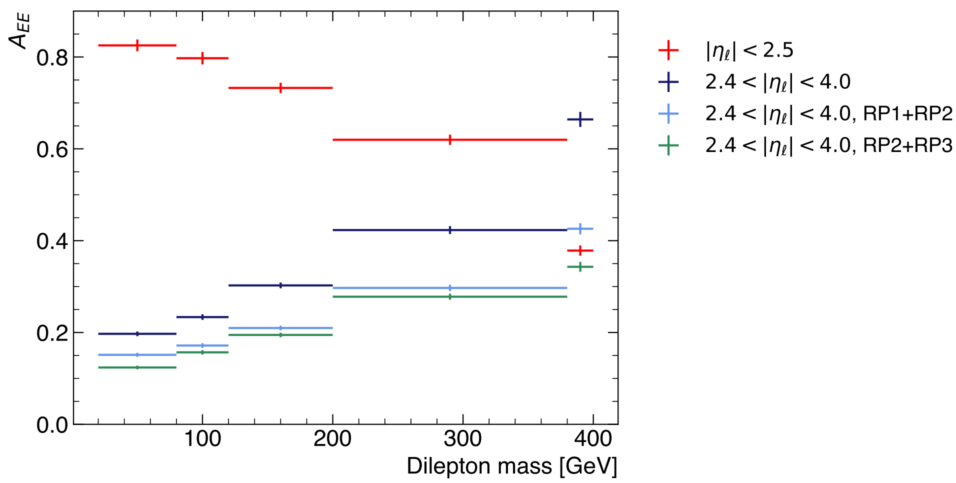

Figure 7.14: The signal acceptance of requiring at least one lepton to be within the HGTD acceptance of $2.4 < |\eta| < 4.0$ (dark blue), relative to all leptons with $|\eta| < 4.0$. The HGTD lepton acceptance is then multiplied by the efficiency of requiring the proton on the same side as the forward lepton to be within the AFP acceptance, A_{AFP}. Note that this acceptance is for elastic events only, where it is assumed that the proton on the same side as the forward lepton has not dissociated. In reality, only half of the single-dissociative events will have this intact proton in the final state, causing an additional reduction factor of $(0.28 + 0.5 \times 0.72)$, ignoring the loss of double-dissociative events.

where A_{EE} is the elastic signal acceptance from Equation (7.11), the factor of 0.95 is the efficiency of the ToF selection in Equation (7.10), and the factor of 0.64 is because the proton on the same side as the lepton could be the one that has dissociated, resulting in only half the single-dissociative (SD) events being reconstructed. Ignoring the loss of double-dissociative events, which is a much smaller component of the signal, this factor is calculated as $N_{\mathrm{EE}} + N_{\mathrm{SD}}/2$. Figure 7.15 shows the single-tag efficiency, $\varepsilon_{\text{single-tag}}$, in Figure 7.14 for both signal and background. The forward proton p_{T} is required to be below 0.3 GeV to further suppress the background. In Figure 7.15, the signal efficiency is higher than the background efficiency for all $m_{\ell\ell}$ and is large enough that a single-tag analysis will be statistically viable at the HL-LHC, rising from 7–9 % at low $m_{\ell\ell}$ and rising to 20–25 % for $m_{\ell\ell} > 380$ GeV. The background efficiency also increases with $m_{\ell\ell}$ due to the probability of the leptons associated to the $q\bar{q} \to WW$ process being in the HGTD acceptance being higher at high $m_{\ell\ell}$. The difference between RP1+RP2 and RP2+RP3 is small until the highest $m_{\ell\ell}$ bin is reached.

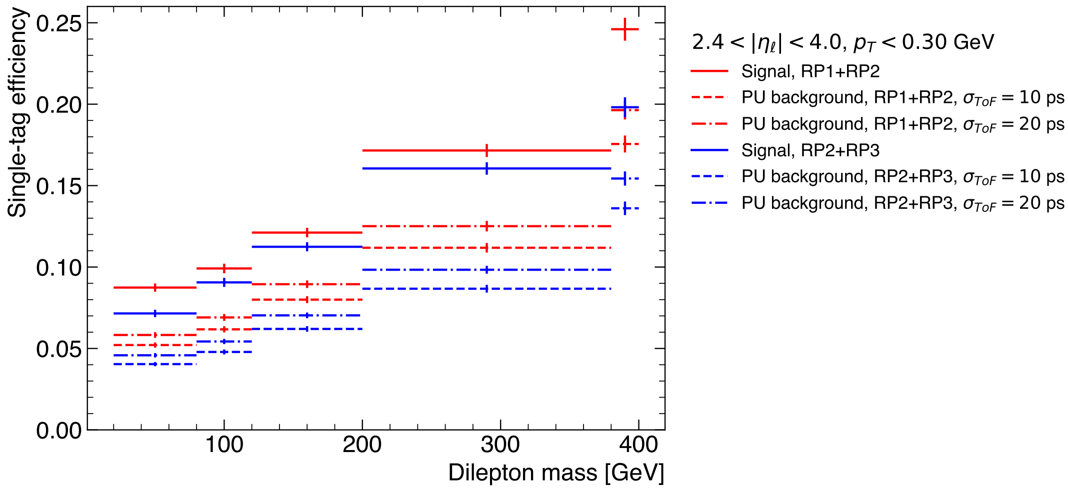

Figure 7.15: The efficiency of requiring a single-proton-tag matched to lepton vertex timing from the HGTD, in a HL-LHC $\gamma\gamma \to WW$ analysis. This figure contains HGTD and AFP kinematic acceptance considerations, as well as ToF matching between the two sub-detector systems. The additional signal efficiency of 64 % due to the loss of half of the single-dissociative events is included.

7.7 Projected event yields for the differential measurement in $m_{\ell\ell}$

The AFP selection efficiencies were then combined with central detector efficiencies to evaluate the projected event yields in a differential $\gamma\gamma \to WW$ analysis at the HL-LHC, with and without forward proton tags. The tracking configuration was chosen to be the optimal one found in the inclusive measurement, with $|\eta| < 2.5$ for tracks, other than the leptons which are allowed to have $|\eta| < 4$ in the single-tag analysis, and with track-$p_{\mathrm{T}} > 500$ MeV. Although Figure 7.8 shows that the relative uncertainty on the signal is fairly flat with increasing window size when background systematics are considered, a window size of 0.2 mm is chosen for the differential analysis in $m_{\ell\ell}$ to minimise the relative statistical uncertainty in each bin, which becomes particularly important at high-$m_{\ell\ell}$ where yields are low.

The differential analysis was first performed for a central-detector only analysis, i.e. with no AFP selection requirements (a “zero-tag” analysis). The predicted HL-LHC yields in the SR are shown, with their relative statistical and total uncertainty, in Figure 7.16a and a comparison of the statistical to total uncertainties with different levels of background systematics is given in Figure 7.16b, as well as a comparison to the Run 2 values. The statistical uncertainty is

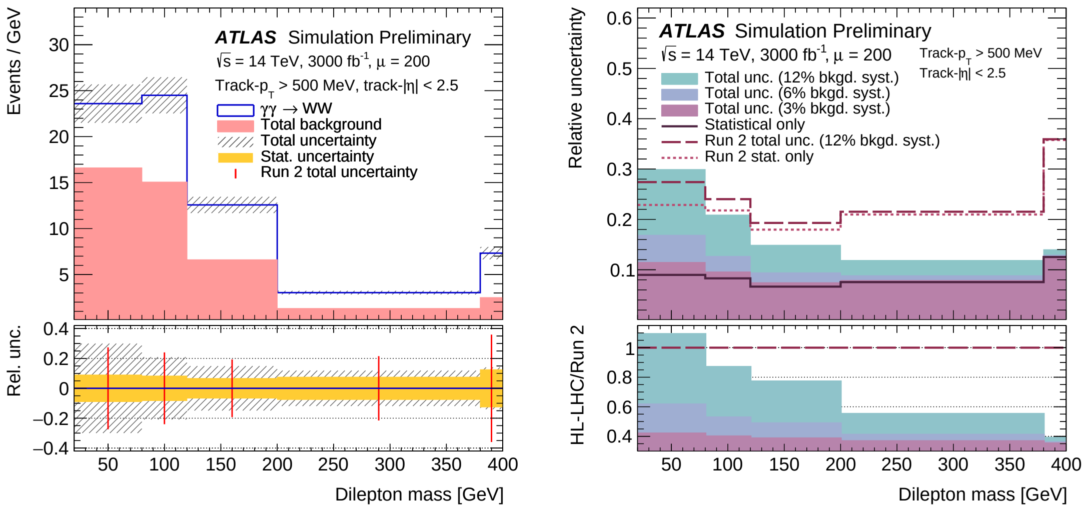

(a) Expected HL-LHC yields (b) Comparison of relative uncertainties

Figure 7.16: The expected signal and background yields at the HL-LHC are shown in (a), stacked and as a function of the dilepton mass. The relative statistical and total uncertainty on the signal are shown in the bottom panel of (a), assuming the same background systematic as in Run 2 of 12 %. A comparison of the total uncertainty when considering a reduction in the nominal background systematic by a factor of 2 and 4 is shown in (b) and compared to the Run 2 relative uncertainties. The bottom panel in (b) shows the ratio of the total relative uncertainty on the signal yield at the HL-LHC as a ratio to the Run 2 uncertainty. The average pileup in Run 2 was $\langle\mu\rangle = 33.7$. Plots produced by the author for Ref. [5].

reduced by a factor of 2.6 in the $200 < m_{\ell\ell} < 380$ GeV bin, when comparing to the Run 2 expected yields separated with the same binning in $m_{\ell\ell}$. However, the relative systematic uncertainty is almost two times larger when considering the nominal 12 % background systematic due to the much larger background yield. The relative total uncertainty on the signal is reduced by almost 50 % in the $200 < m_{\ell\ell} < 380$ GeV bin compared to Run 2, even when considering the nominal 12 % background systematic. An even larger relative improvement is found in the $m_{\ell\ell} > 380$ GeV region due to the very large Run 2 statistical uncertainty in this bin.

Figure 7.16b demonstrates the effect of reducing the background systematics by a factor of 2 and 4. It is seen that reducing the background systematic has a large effect at low $m_{\ell\ell}$ but is not so important at high $m_{\ell\ell}$ where background yields are much smaller. Even with the large increase in integrated luminosity at the HL-LHC, the statistical uncertainty on the signal yield begins to dominate in the highest $m_{\ell\ell}$ bins but is still largely reduced relative to Run 2. A roughly

64 % reduction on Run 2 in expected relative statistical uncertainty is seen in the $200 < m_{\ell\ell} < 380$ GeV bin. Improving background modelling systematics of the underlying event before HL-LHC data analysis begins should remain a priority in order to keep up with improvements in statistical precision but it is not as crucial as for a fiducial measurement when studying tails of distributions where yields are low. The results in the differential analysis assume that dedicated low-p_{T} tracking will be possible and therefore focus on optimising low-p_{T} reconstruction methods at the HL-LHC is required. Not included in the results are changes in cross-section going from $\sqrt{s} = 13$ TeV to $\sqrt{s} = 14$ TeV. These increases are expected to range from roughly 2 % in the lowest $m_{\ell\ell}$ bin to 16 % in the highest $m_{\ell\ell}$ bin, for the elastic $\gamma\gamma \to WW$ process.

Ultimately, any analysis at the HL-LHC that has intact forward protons in the final state can be split into three categories; zero-tag (central-detector only), single-tag, and double-tag. The expected signal and background yields for the three categories are presented in Table 7.5 and compared to the expected yields in the Run 2 analysis. Although both the signal and background yields decrease with increasing $m_{\ell\ell}$, the background yield falls faster due to the decreasing background efficiency and therefore the signal to background ratio increases with increasing $m_{\ell\ell}$.

All proton-tagged results are presented for both combinations of detector location considered, RP1+RP2 or RP2+RP3, and APF-ToF resolutions of 10 and 20 ps are compared. The analysis requirements are tailored to retain 95% of the signal for any given ToF resolution. The forward proton p_{T} is required to be smaller than 0.25 GeV for the double-tag analysis and smaller than 0.30 GeV in the single-tag analysis to further suppress the pileup background, studies of which were given in Section 7.6.1. Kinematic matching between the protons measured by AFP and leptons measured by the central detector cannot be performed in either analysis due to missing information from the undetected neutrinos, as well as any proton that has dissociated or is outside of the AFP acceptance.

The large signal losses due to small AFP acceptance and loss of dissocative signal result in a double-tag analysis that would be far too statistically limited to be viable at the HL-LHC, with no more than a couple of signal events predicted in each dilepton mass range. The signal to background ratio is also much worse than in any other analysis category, despite improved background suppression compared to a central-detector-only analysis. This is due partly to the increased

	Dilepton mass [GeV]				
	20–80	80–120	120–200	200–380	380+
Run 2					
Signal	43	39	49	32	10
Background	54	33	29	13	3
HL-LHC zero-tag					
Signal	419	377	477	310	97
Background	997	603	529	236	50
HL-LHC single-tag					
RP1+RP2					
Signal	44	47	79	86	63
Background 10 ps	63	47	58	43	23
Background 20 ps	70	52	65	48	26
RP2+RP3					
Signal	36	43	73	80	51
Background 10 ps	49	36	45	33	18
Background 20 ps	55	41	51	37	20
HL-LHC double-tag					
RP1+RP2					
Signal	< 0.1	< 0.1	0.3	0.6	1.4
Background 10 ps	61	37	32	15	2.6
Background 20 ps	105	65	55	25	5.3
RP2+RP3					
Signal	0.2	0.5	1.9	3.7	6.6
Background 10 ps	31	20	18	7.9	1.6
Background 20 ps	60	37	31	15	3.0

Table 7.5: Comparison of expected signal and background yields for three categories of $\gamma\gamma \to WW$ analysis at the HL-LHC and the expected yields in the Run 2 analysis [4], in bins of dilepton mass, $m_{\ell\ell}$. All results assume central detector selections of track-$p_{\mathrm{T}} > 500\,\mathrm{MeV}$ and track-$|\eta| < 2.5$ with a HL-LHC track-veto window-size of 0.2 mm and Run 2 window size of 1 mm. An integrated luminosity of $3000\,\mathrm{fb}^{-1}$ and $\mu = 200$ are assumed at the HL-LHC. The single-tag analysis requires at least one lepton to be within the HGTD acceptance of $2.4 < |\eta| < 4.0$ but all other categories assume leptons within an acceptance of $|\eta| < 2.5$. Different combinations of AFP station locations are considered: RP1+RP2 ($0.08 < \xi < 0.31$) and RP2+RP3 ($0.05 < \xi < 0.20$). The background yields are presented for an AFP-ToF resolution of both 20 ps (current best) and 10 ps (goal for HL-LHC). The HGTD resolution is taken as 30 ps.

number of pileup protons at the HL-LHC but mainly the loss of dissociative signal and the increased lower edge of the AFP acceptance compared to Run 2 and 3. In order to benefit from double-proton-tagging in a $\gamma\gamma \to WW$ analysis, we need very low acceptance in ξ, as the energy loss of each proton is expected to be very small. The single-tag category actually improves the signal to background ratio compared to the zero-tag analysis. When considering that the background would be calculated in an independent way (data-driven event mixing) in a proton-tagged analysis compared to zero-tag, a single-tag analysis will be useful at further constraining backgrounds and reducing background systematics. The single-tag analysis would also be performed in a different kinematic phase-space, allowing further constraints of any EFT effects.

Not considered here is the possibility of relaxing central detector signal requirements in response to requiring forward proton tags. For example, the strict track-veto requirement that is imposed in the central-detector-only analysis may be able to be relaxed as reconstructing forward protons in the final state is an independent method of defining exclusivity.

7.8 Conclusion on future $\gamma\gamma \to WW$ prospects

The results presented in this chapter have shown that the higher pileup environment expected at the HL-LHC will give rise to important challenges for accurately identifying exclusive processes like $\gamma\gamma \to WW$. The expected sensitivity to the process $\gamma\gamma \to WW \to e^{\pm}\nu\mu^{\mp}\nu$ with the ATLAS detector at the HL-LHC, when assuming an integrated luminosity of $3000\,\mathrm{fb}^{-1}$ and $\mu = 200$, was estimated by projecting the Run 2 analysis [4] and correcting for the main expected differences in performance after the ATLAS Phase-II upgrade. Full-simulation HL-LHC MC samples were used to derive the expected performance of the upgraded ATLAS Inner Tracker. Furthermore, several potential optimisations of the HL-LHC baseline analysis have been investigated, including narrowing the track-veto z-window size, reducing the minimum-reconstructed track-p_{T} to the nominal Run 2 value of $p_{\mathrm{T}} > 500\,\mathrm{MeV}$ and increasing the tracking acceptance from $|\eta| < 2.5$ to $|\eta| < 4.0$. The expected sensitivity to background modelling uncertainties was also investigated. The prospects for a differential analysis in $m_{\ell\ell}$ were investigated, with the aim of targeting the high-$m_{\ell\ell}$ region for improved constraints on EFT operators. The increased statistical precision at the HL-LHC results in improved prospects in this region, with a roughly 64 % reduction in expected relative statistical uncertainty and almost 50 % reduction in relative total uncertainty in the $200 < m_{\ell\ell} < 380\,\mathrm{GeV}$ bin, when using a minimum reconstructed track-p_{T} of $p_{\mathrm{T}} > 500\,\mathrm{MeV}$ and a track-veto window size of 0.2 mm.

The theoretical modelling uncertainty on the main background becomes a dominant uncertainty with large expected increases in background yields and efforts to reduce it would ensure it will not limit the overall expected precision on the extracted cross-section. The reconstructed track-p_{T} threshold will be a key experimental consideration in the prospects for $\gamma\gamma \to WW$ analyses at the HL-LHC and further research and development is required to optimise low-p_{T} tracking methods for HL-LHC operating conditions.

In addition to the central detector studies, two additional categories of $\gamma\gamma \to WW$ analysis were studied in the expected operating conditions at the HL-LHC: a single- and double-proton-tag analysis with an AFP-like detector. AFP has the potential to increase sensitivity in these analyses by both suppressing backgrounds originating from pileup interactions and allowing central detector cuts to be relaxed to increase signal acceptance. However, requiring a proton tag in AFP also largely reduces signal acceptance due to limited AFP kinematic acceptance, as well as the loss of dissociative events depending on whether a single- or double-tag is required.

It was found that the central-detector-only analysis outperformed an analysis in which a double proton tag was required in AFP. This is largely due to the loss of signal acceptance when considering limited AFP kinematic acceptance, and the loss of dissociative events which make up 72 % of the signal. The results presented in this chapter assume an AFP-ToF resolution of $\sigma_{\mathrm{ToF}} = 20$ ps which is the best achieved resolution to date [179] and compare to the HL-LHC goal of $\sigma_{\mathrm{ToF}} = 10$ ps. In order to make a double-tag analysis viable, the AFP-ToF resolution would have to be reduced to $\sigma_{\mathrm{ToF}} < 10$ ps which would require a huge amount of timing detector research and development. Another possibility would be to reduce the lower edge of the AFP acceptance in ξ, however it is highly unlikely that an improvement will be able to be made on the Run 2 value of $\xi_{\mathrm{AFP}}^{\mathrm{min.}} = 0.02$ due to machine and detector-location limitations at the HL-LHC. Therefore, it is concluded that a double-tag $\gamma\gamma \to WW$ analysis at the HL-LHC will not be viable, under the assumptions and analysis selections made in these studies.

On the other hand, a single-tag $\gamma\gamma \to WW$ analysis, in which one lepton is in the acceptance of the HGTD, can improve the signal to background ratio compared to a central-detector-only analysis. Furthermore, a single-tag analysis could be interesting in two ways: Firstly, backgrounds are calculated in a different way in analyses with forward proton tags (data-driven event mixing), allowing cross-checks and reducing systematic uncertainties on backgrounds which was shown to be vital in order to not limit the measurement precision. Secondly, the different topology of the single-tag selection will possibly allow further constraints on any EFT effects, possibly improving limits on EFT coupling coefficients.

Although not a crucial factor in the $\gamma\gamma \to WW$ single-tag analysis, from these studies it is recommended that the combination of RP2+RP3 AFP detector locations be chosen over RP1+RP2 as the lower ξ acceptance would allow any analysis which requires a double-tag to have lower reach in central mass.

Not considered in this thesis are prospects in Run 3. Operating conditions in Run 3 are expected to be similar to Run 2 but with higher average interactions per bunch crossing of around $\langle\mu\rangle \approx 60$ and an increased CoM energy of $\sqrt{s} = 13.6\,\text{TeV}$. ATLAS and CMS aim to more than double their total integrated luminosity over the course of Run 3, allowing improvements on statistical precision for rare electroweak processes. As shown in the HL-LHC studies presented in this chapter, the increase in pileup will be detrimental to exclusive analyses such as $\gamma\gamma \to WW$ but the increase in integrated luminosity and CoM energy will help to reduce statistical uncertainties associated to measurements. It is also hoped that efficient time-of-flight detectors will be available in AFP, which aims to more than double its collected dataset in the first year of Run 3 data-taking alone.

Chapter 8

Conclusions and outlook

The work presented in this thesis covered the observation and measurement of photon-induced W^+W^- production in the fully-leptonic decay mode. The observation was performed in proton-proton collision data at a centre-of-mass energy of 13 TeV and with an integrated luminosity of $139\,\mathrm{fb}^{-1}$, recorded by the ATLAS detector. Particular emphasis was given on data-driven corrections to charged particle multiplicity distributions in simulation, which were vital for understanding signal and background processes. The efficiency of the signal to pass the exclusivity selection in data was determined to be 52.6 %. The measured $\gamma\gamma \rightarrow WW$ cross-section in the fiducial volume of the analysis was

$$\sigma^{\text{fid.}}_{\text{meas.}} = 3.13^{+0.32}_{-0.31}\ (\text{stat.})^{+0.29}_{-0.26}\ (\text{syst.})\ \text{fb}.$$

and the background hypothesis was rejected with a significance of 8.4 standard deviations. This is the first observation of the $\gamma\gamma \rightarrow WW$ process in any decay channel.

Distributions were then presented differentially in many kinematic observables. The post-fit simulation predicted yields generally agree well with the data, within statistical and systematic uncertainties. Larger disagreement was seen at high dilepton p_{T}, $p_{\mathrm{T}}^{\ell\ell}$, and the leading-lepton p_{T} appeared to be slightly harder in data compared to simulation. There was also some disagreement also at high values of $\Delta\Phi_{\ell\ell}$, where the leptons have large separation in azimuthal angle. Preliminary unfolding studies were then presented, using D'Agostini's iterative unfolding method implemented in the `RooUnfold` package. As the measurement has very little migration between particle-level and reconstruction-level bins and has large statistical uncertainties, it was found that one iteration was optimal to minimise the variance of the unfolded result. Optimisation studies are ongoing and other unfolding techniques shall be tested and compared in the future. Ultimately, unfolded distributions will be used to calculate limits on EFT operators that impact $WW\gamma$ and $WW\gamma\gamma$ couplings.

In addition to the observation and measurement of the $\gamma\gamma \rightarrow WW$ process, performance work on the ATLAS Forward Proton detector was presented, detailing the inter-plane alignment algorithm and the alignment results calculated in 2017 data. The working points of proton reconstruction within AFP were developed and reconstruction efficiencies were calculated to be over 90 % in all stations.

Finally, prospects for future $\gamma\gamma \rightarrow WW$ analyses at the High-Luminosity LHC were presented. It was demonstrated that the higher pileup environment expected at the HL-LHC will give rise to important challenges for accurately identifying exclusive processes like $\gamma\gamma \rightarrow WW$. Several potential analysis optimisations were investigated and it was determined that the reconstructed track-p_{T} threshold will be a key experimental consideration in the prospects for $\gamma\gamma \rightarrow WW$ analyses at the HL-LHC and further research and development is required to optimise low-p_{T} tracking methods for HL-LHC operating conditions. The prospects for a differential analysis in $m_{\ell\ell}$ were investigated, with the aim of targeting the high-$m_{\ell\ell}$ region for improved constraints on EFT operators. The increased statistical precision at the HL-LHC results in improved prospects in this region relative to the Run 2 observation, with a roughly 64 % reduction in expected relative statistical uncertainty and almost 50 % reduction in relative total uncertainty for dilepton masses of $200 < m_{\ell\ell} < 380\,\mathrm{GeV}$, when using a minimum reconstructed track-p_{T} of $p_{\mathrm{T}} > 500\,\mathrm{MeV}$.

In addition to the central detector studies, two additional categories of $\gamma\gamma \rightarrow WW$ analysis were studied in the expected operating conditions at the HL-LHC: a single- and double-proton-tag analysis with an AFP-like detector. It was determined that a novel single-tag $\gamma\gamma \rightarrow WW$ analysis, in which one lepton is in the acceptance of the High-Granularity Timing Detector, can improve the signal to background ratio compared to a central-detector-only analysis. The expected signal and background yield when selecting a dilepton mass above 380 GeV and with no selection on AFP is 97 and 50 events, respectively, whereas it is 51 and 18 in a single-tag analysis (assuming an AFP-ToF resolution of 10 ps and an acceptance of $0.05 < \xi < 0.20$). Furthermore, a single-tag analysis would likely allow additional constraints on background systematic uncertainties through data-driven predictions and on any new physics effects due to the different topology of the single-tag events. Therefore, AFP should be seen as an addition to the ATLAS HL-LHC physics programme.

Appendix A

Additional figures for the AFP inter-plane alignment

This appendix contains additional figures for the AFP inter-plane alignment, first presented in Section 4.5. The coordinate system used in this appendix is central ATLAS coordinates.

Figures A.1–A.3 show the evolution of alignment parameters over time for the A-FAR, C-NEAR and C-FAR stations. The results for the A-NEAR station were presented in Section 4.5.2. The alignment analysis can highlight issues with individual planes if convergence is not seen. For example, STATION-3 (C-FAR), plane-1 had no high voltage across 2017 which resulted in poor track reconstruction. The effect of this can clearly be seen in Figure A.3. STATION-0 (A-FAR) only had two working planes during most of 2017, impacting track reconstruction. Fewer planes result in worse position resolution per track and increase the relative fraction of fake tracks reconstructed. This effect can be seen in Figure A.1.

Figures A.4–A.6 show the AFP inter-plane alignment stability plots for the α, Δx and Δy parameters across standard runs in 2017 in which AFP was inserted and the data passed the AFP GRL. Individual planes have been removed from the analysis if there were known issues. The weighted mean and standard deviation of all runs are taken as the final parameter values. There is a jump seen in parameter values in some figures, such as the top-left of Figure A.6, which correspond to a change in SiT tuning, discussed in Section 4.6.

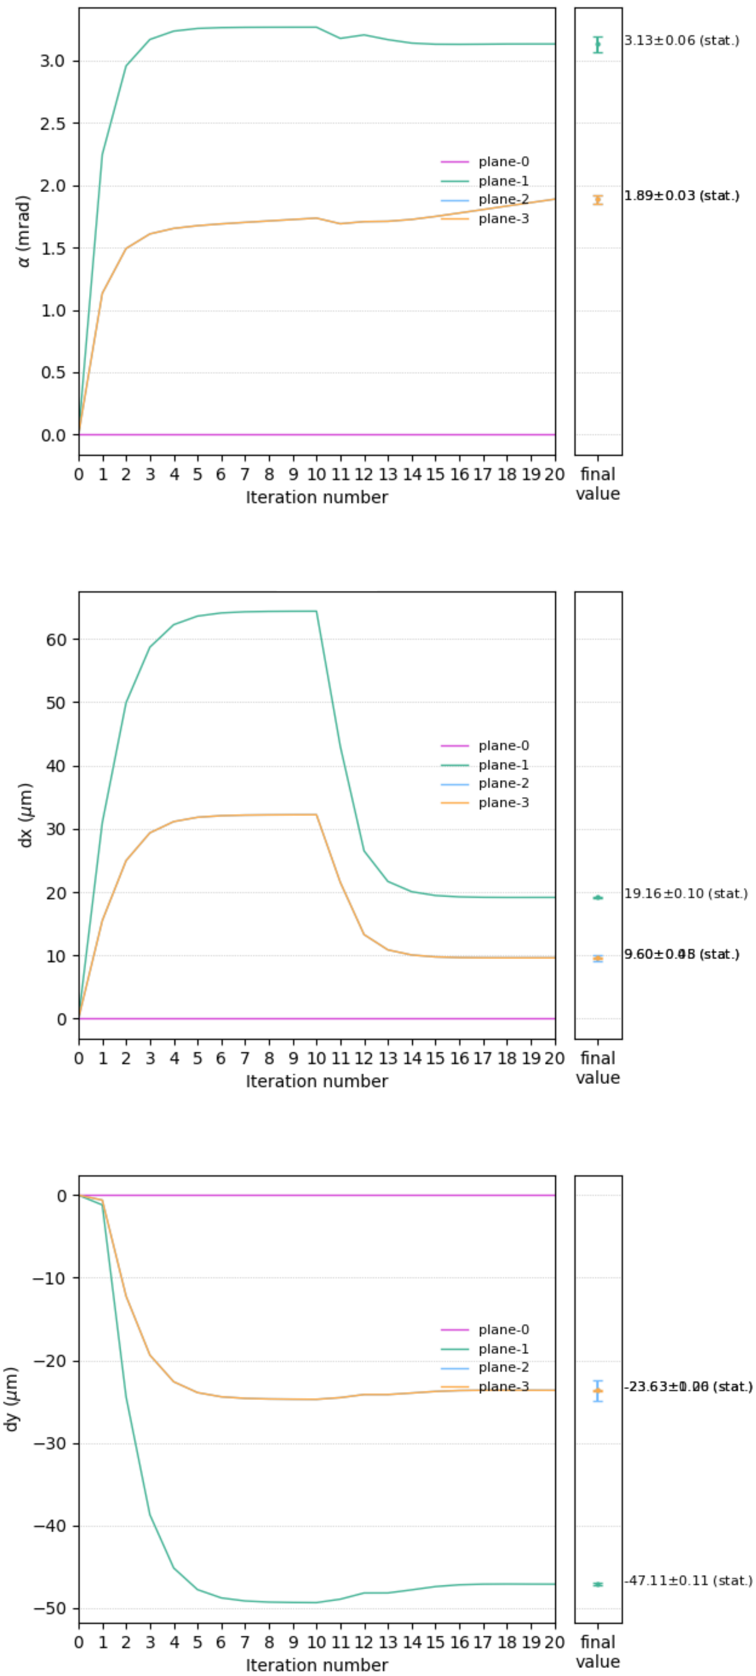

Figure A.1: The evolution of the (top) α, (middle) Δx and (bottom) Δy alignment parameters for the A-FAR station in run 337052.

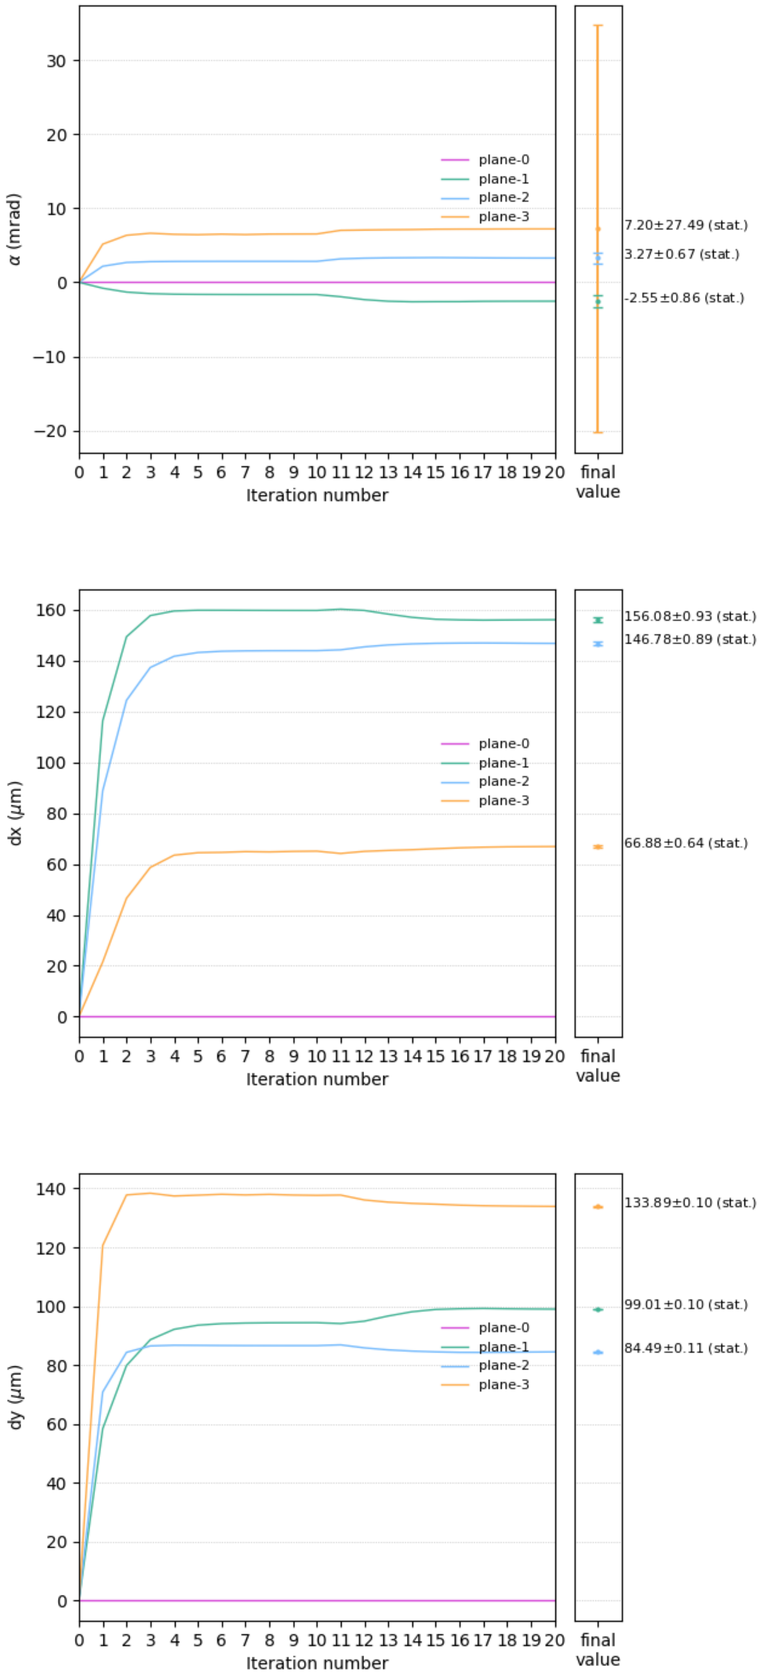

Figure A.2: The evolution of the (top) α, (middle) Δx and (bottom) Δy alignment parameters for the C-NEAR station in run 337052.

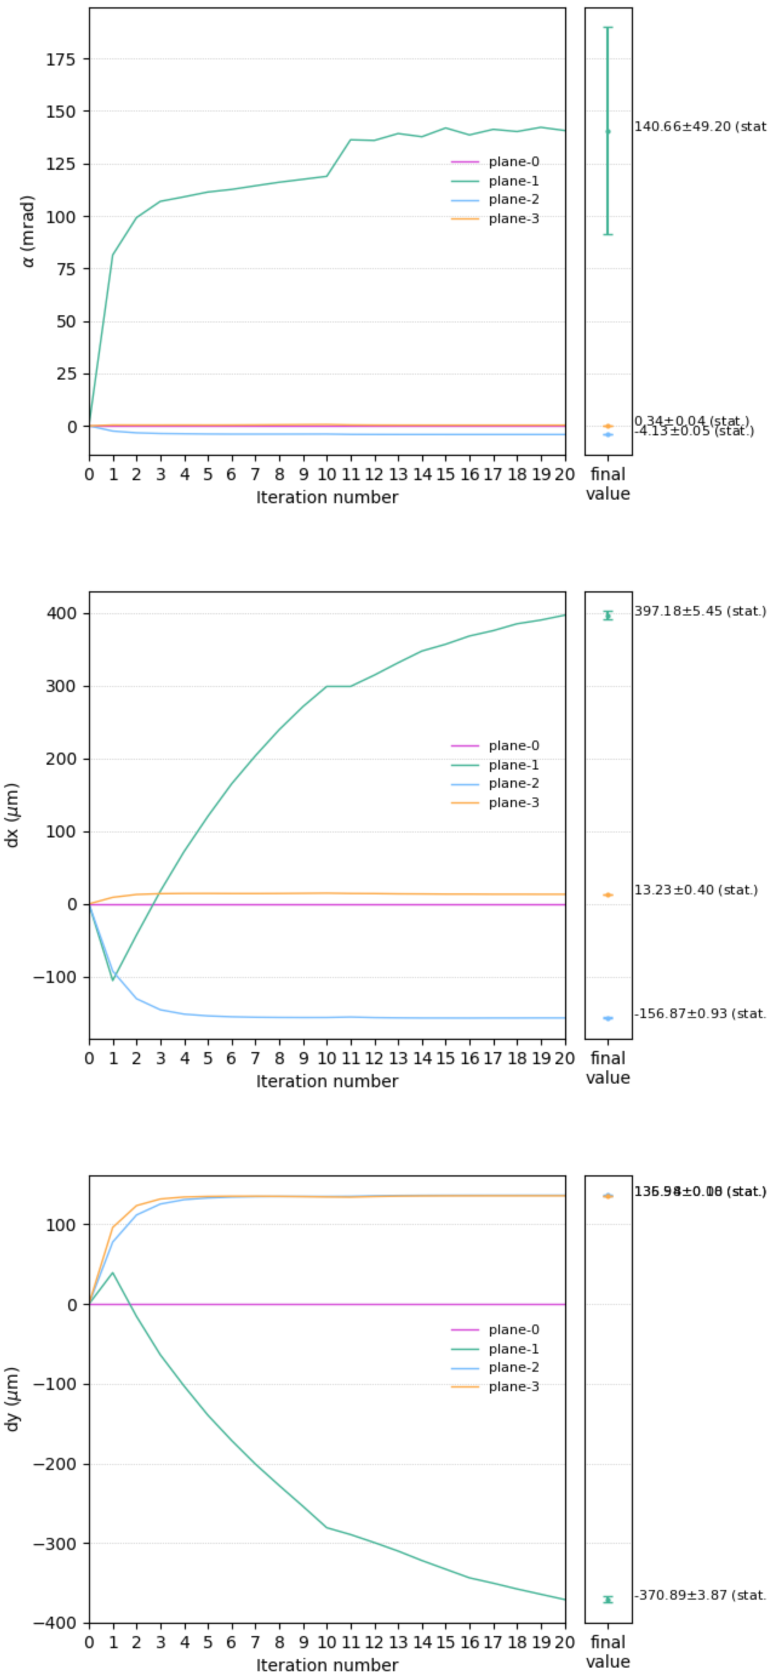

Figure A.3: The evolution of the (top) α, (middle) Δx and (bottom) Δy alignment parameters for the C-FAR station in run 337052.

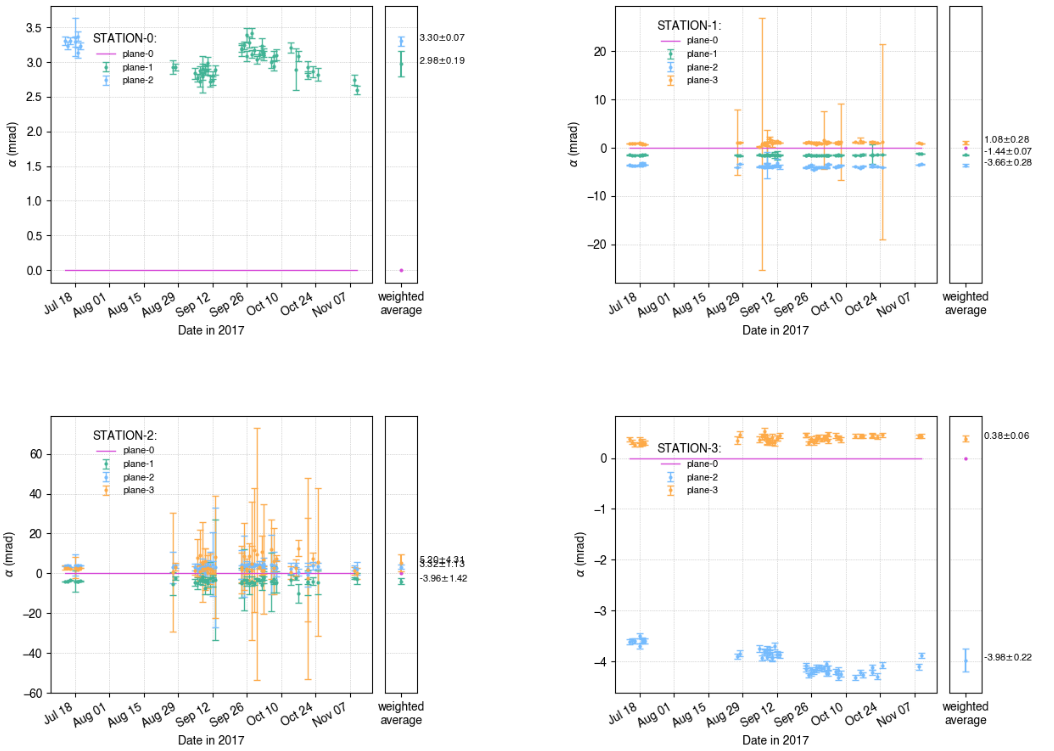

Figure A.4: Stability of the α alignment parameter across 2017 data-taking

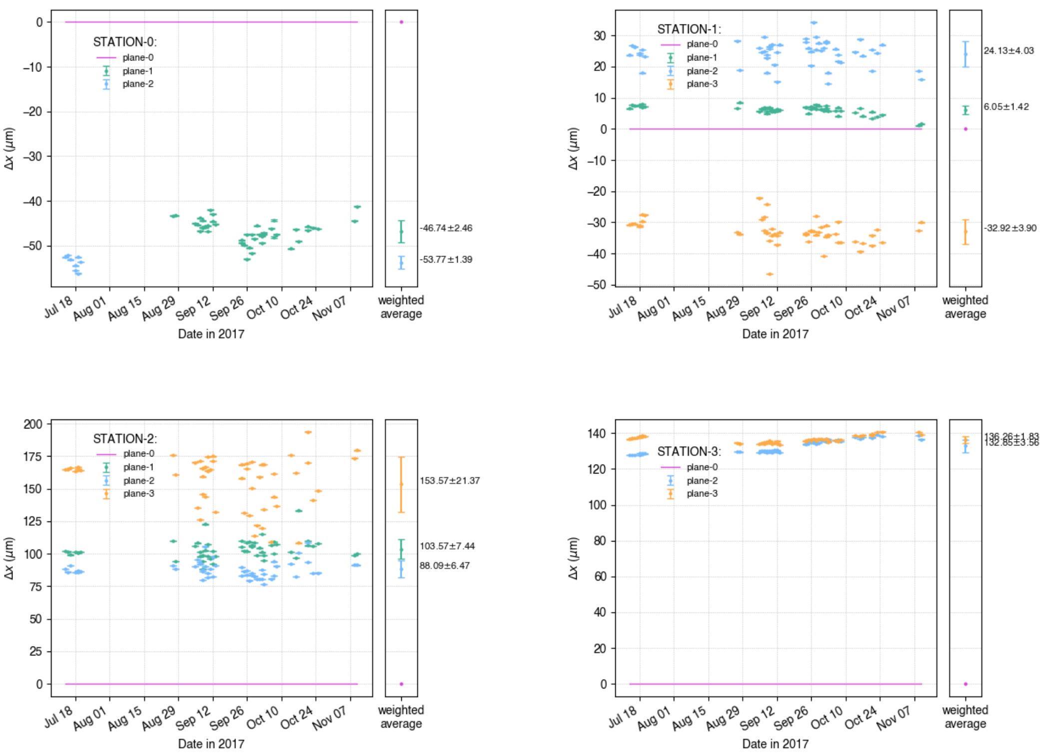

Figure A.5: Stability of the Δx alignment parameter across 2017 data-taking

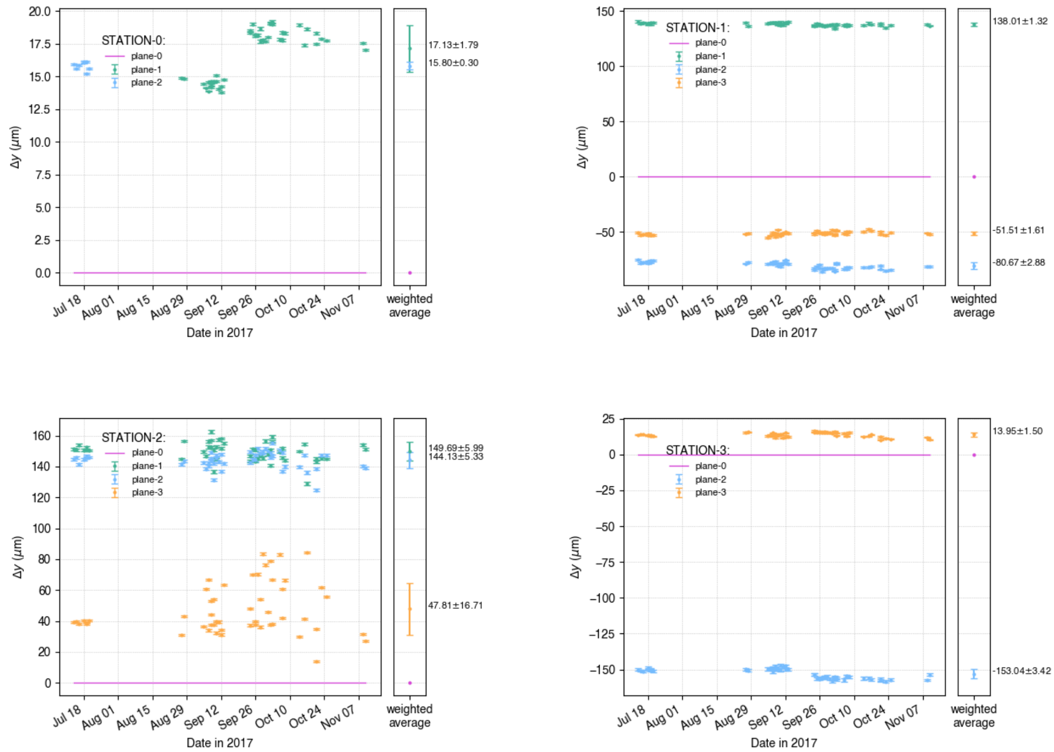

Figure A.6: Stability of the Δy alignment parameter across 2017 data-taking

Appendix B

Pileup simulation comparisons: PYTHIA8 vs EPOS+PYTHIA8

The author performed studies into pileup track kinematic distributions comparing the Run 2 simulation setup to the new simulation setup for Run 3. In Run 2, ATLAS generated and showered pileup interactions using PYTHIA8 (Py8) with the A3 tune [101]. However, this setup was known to model pileup track distributions poorly for low-p_{T} tracks [234, 235]. Therefore a new setup was developed for Run 3, using EPOS [192] for over 99 % of minimum-bias event generation, while keeping PYTHIA8 for rarer high-p_{T} activity (jet-$p_{\mathrm{T}} > 35\,\mathrm{GeV}$). In PYTHIA8 inclusive pp interactions are described by a model that splits the total inelastic cross-section into non-diffractive processes, dominated by t-channel gluon exchange, and diffractive processes involving a colour-singlet exchange. The simulation of the non-diffractive component uses MPI models. The diffractive component is further divided into single- and double-dissociative processes, described using a pomeron-based approach [236]. EPOS provides an implementation of a parton-based Gribov–Regge theory [237], which is an effective QCD-inspired field theory describing hard and soft scattering simultaneously.

The tests were performed comparing $Z \rightarrow \ell^+\ell^-$ simulation to 2018 data, corresponding to an integrated luminosity of $58.5\,\mathrm{fb}^{-1}$, which most closely matches the beam parameters set in the simulation. Figures B.1a and B.1b are used as sanity checks, to make sure distributions that shouldn't be affected by the different modelling are consistent between the two setups. Both Py8 and EPOS+Py8 agree for these two distributions. The agreement with data in Figures B.1a and B.1b is not perfect but this is not a concern as the μ distribution will be corrected with pileup reweighting and the beamspot is expected to have slightly different dimensions and position in data and MC.

Figures B.1c and B.1d show the $n_{\mathrm{trk}}^{\mathrm{PU}}$ distribution for tracks that satisfy $|z_{\mathrm{track}} - z_{\mathrm{bin}}| < 1\,\mathrm{mm}$ where each z-bin samples across the beamspot longitudinal width. Figure B.1d is zoomed into the low-n_{trk} region of $n_{\mathrm{trk}}^{\mathrm{PU}} < 10$. Neither setup models the track distribution perfectly but EPOS+Py8 performs significantly better for $n_{\mathrm{trk}}^{\mathrm{PU}} < 10$.

Figure B.1e shows the pileup track-p_{T} distribution. Again, neither setup models the distribution perfectly but EPOS+Py8 generally performs better than standalone PYTHIA8, apart from in the region of roughly $1.5 < p_{\mathrm{T}} < 3\,\mathrm{GeV}$. It should be noted that EPOS data/MC agreement within 5 % is impressive considering the difficulty in modelling diffractive pileup interactions, particularly at low-p_{T}.

Figure B.1f shows the exclusive efficiency (discussed at length in Section 5.4) as a function of μ. Note that the $n_{\mathrm{trk}} = 0$ bin in Figures B.1c and B.1d is this distribution integrated over all μ. Therefore, Figure B.1f is a more accurate representation of the $n_{\mathrm{trk}} = 0$ modelling as it removes the effect of differences in μ between data and MC. Although both setups are generally within a few % of data in Figure B.1f, the average μ in data is slightly higher which coincidentally improves the data/MC agreement.

From the results shown here, it can be concluded that EPOS+Py8 outperforms standalone Py8 when considering the number of reconstructed tracks per region in z and better models the pileup track-p_{T} across the spectrum.

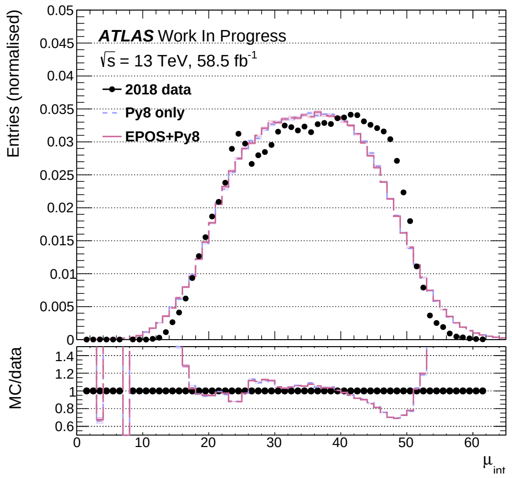

(a) μ

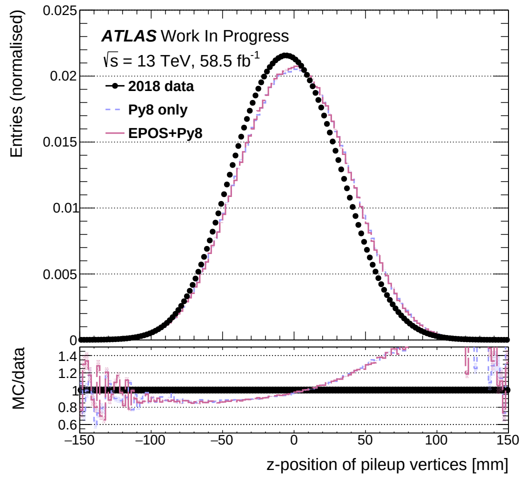

(b) $z_{\mathrm{vtx}}^{\mathrm{PU}}$

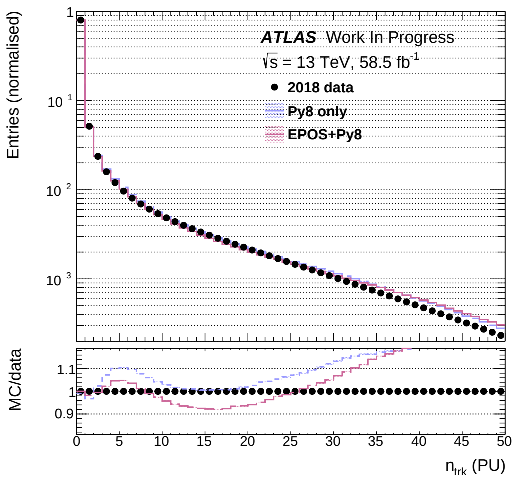

(c) $n_{\mathrm{trk}}^{\mathrm{PU}}$

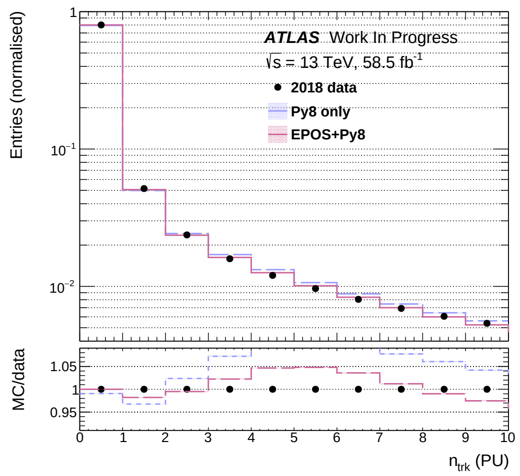

(d) $n_{\mathrm{trk}}^{\mathrm{PU}} < 10$

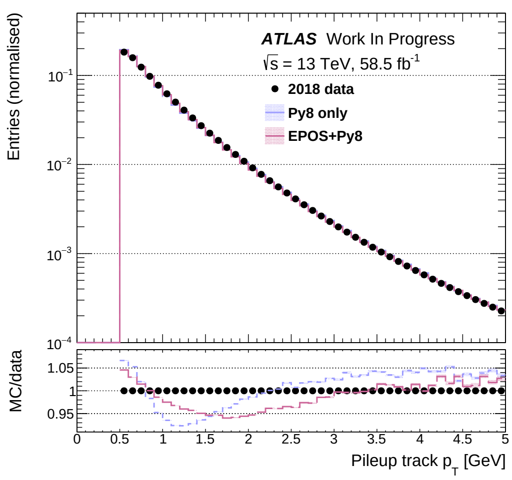

(e) Pileup track-p_{T}

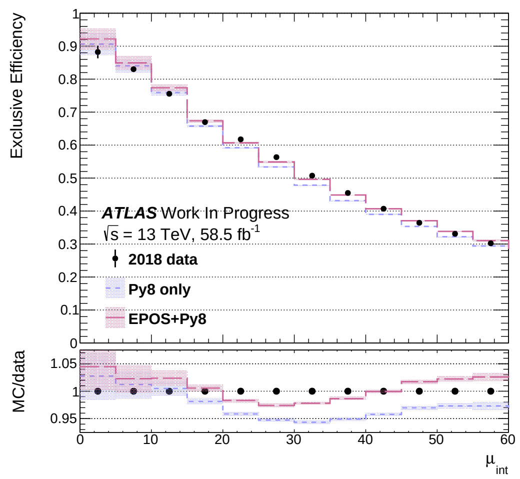

(f) Exclusive efficiency vs μ

Figure B.1: Comparisons of standalone PYTHIA8 to a combination of EPOS and PYTHIA8 for track-$p_{\mathrm{T}} < 1\,\mathrm{GeV}$ and track-$p_{\mathrm{T}} > 1\,\mathrm{GeV}$, respectively. All figures are normalised to unity. PYTHIA8 uses the A3 tune in both setups.

Appendix C

Further studies on the corrections to the $\gamma\gamma \to WW$ cross-section

In addition to the results presented in Section 5.8.4, further studies were carried out into the corrections applied to the cross-section calculation. Firstly, A_{WW} and C_{WW} were calculated in signal MC with two different beamspot sizes but which had both been scaled to better match the beamspot found in data. The acceptance, $A_{WW} = (N_{\text{pass}}(\text{truth})/N_{\text{tot.}}(\text{truth})) \times \text{filter eff}$, should be independent of z_0-scaling as no selection on pileup tracks is made at particle-level. Therefore, any differences will be due to the process in which the alternate beamspot samples are produced. Table C.1 shows the acceptance values calculated in both samples. The acceptance is consistent between the two samples to within one standard deviation.

Table C.2 compares the C_{WW} factor for the two different beamspot sizes, across the different MC campaigns and combined. The 35 mm sample is chosen as nominal as the beamspot size is closer to that seen in data and therefore, less correction is needed when applying z_0-scaling to tracks. The discrepancy between the two samples can be attributed to the systematic uncertainty on the beamspot rescaling method.

BS σ_z [mm]	A_{WW}
42	0.0724 ± 0.0005
35	0.0732 ± 0.0004

Table C.1: The signal acceptance calculated in signal MC. The BS σ_z refer to the beamspot size before z_0-scaling is applied. Both samples are z_0-scaled to better match BS distributions seen in data.

Studies into the effect of including/excluding tau decays in the fiducial region were also carried out. Taus are excluded from the fiducial definition at particle-level in the cross-section measurement. Tables C.3, C.4, C.5, and C.6 show values for C_{WW} under different conditions; including/excluding tau decays and with/without z_0 scaling applied. Results are also given for different dilepton flavour channels: ee, $\mu\mu$ and $\mu e/e\mu$. The latter is the signal selection. The relative contribution of taus is 15% both at particle- and reconstruction-level. This equates to an absolute event yield contribution of around 5 % of tau-decays in the fiducial volume.

	C_{WW} (τ decay excluded in fiducial selection, no z_0 scaling)			
channel	2015–16	2017	2018	**combined**
$\mu\mu$	0.5818 ± 0.0112	0.4598 ± 0.0100	0.4767 ± 0.0090	0.4997 ± 0.0058
ee	0.4738 ± 0.0119	0.3727 ± 0.0101	0.4125 ± 0.0092	0.4159 ± 0.0059
$\boldsymbol{\mu e/e\mu}$	0.5278 ± 0.0084	0.4330 ± 0.0072	0.4394 ± 0.0065	$\mathbf{0.4602 \pm 0.0042}$

Table C.3: C_{WW} calculated in the full fiducial region, with tau decays excluded and no z_0 scaling applied.

	C_{WW} (τ decay included in fiducial selection, no z_0 scaling)			
channel	2015–16	2017	2018	**combined**
$\mu\mu$	0.5080 ± 0.0106	0.3995 ± 0.0092	0.4163 ± 0.0083	0.4357 ± 0.0054
ee	0.4128 ± 0.0109	0.3271 ± 0.0092	0.3576 ± 0.0084	0.3624 ± 0.0054
$\boldsymbol{\mu e/e\mu}$	0.4601 ± 0.0078	0.3752 ± 0.0066	0.3844 ± 0.0060	$\mathbf{0.4010 \pm 0.0039}$

Table C.4: C_{WW} calculated in the full fiducial region, with tau decays included and no z_0 scaling applied.

	C_{WW}			
BS σ_z [mm]	2015–16	2017	2018	**combined**
42	0.502 ± 0.007	0.393 ± 0.006	0.397 ± 0.005	0.423 ± 0.003
35	0.471 ± 0.007	0.367 ± 0.006	0.362 ± 0.005	0.392 ± 0.003

Table C.2: C_{WW} calculated for different MC campaigns in Herwig7 elastic $\gamma\gamma \to WW$ simulation, with initial beamspot sizes of 42 mm and 35 mm, where both have been z_0-scaled to match the beamspot size in data.

	C_{WW} (τ decay excluded in fiducial selection, z_0 scaling to data)			
channel	2015–16	2017	2018	**combined**
$\mu\mu$	0.5536 ± 0.0113	0.4373 ± 0.0100	0.4419 ± 0.0090	0.4706 ± 0.0058
ee	0.4472 ± 0.0118	0.3532 ± 0.0100	0.3751 ± 0.0091	0.3870 ± 0.0058
$\boldsymbol{\mu e/e\mu}$	0.4995 ± 0.0084	0.4158 ± 0.0072	0.4041 ± 0.0064	$\mathbf{0.4324 \pm 0.0042}$

Table C.5: C_{WW} calculated in the full fiducial region, with tau decays excluded and z_0 scaling applied to better match the beamspot distribution in data.

	C_{WW} (τ decay included in fiducial selection, z_0 scaling to data)			
channel	2015–16	2017	2018	**combined**
$\mu\mu$	0.4834 ± 0.0106	0.3800 ± 0.0091	0.3859 ± 0.0082	0.4103 ± 0.0053
ee	0.3897 ± 0.0108	0.3101 ± 0.0091	0.3252 ± 0.0082	0.3373 ± 0.0053
$\boldsymbol{\mu e/e\mu}$	0.4354 ± 0.0078	0.3603 ± 0.0065	0.3535 ± 0.0059	$\mathbf{0.3768 \pm 0.0038}$

Table C.6: C_{WW} calculated in the full fiducial region, with tau decays included and z_0 scaling applied to better match the beamspot distribution in data.

Appendix D

Additional figures for the differential $\gamma\gamma \to WW$ measurement

This appendix includes additional technical closure test figures and correlation matrices related to the differential $\gamma\gamma \to WW$ measurement presented in Chapter 6. All technical closure tests close perfectly, as expected. The statistical correlation between the event yields in bins i and j, N_i and N_j, is defined as

$$\mathrm{corr}(N_i, N_j) = \frac{\mathrm{cov}(N_i, N_j)}{\sigma_{N_i}\sigma_{N_j}}, \tag{D.1}$$

where σ is the standard deviation of N_i and N_j and the covariance, $\mathrm{cov}(N_i, N_j)$, is calculated as described in Ref. [238]. The correlation matrices have small anti-correlations on the non-diagonals as there is very little migration between bins.

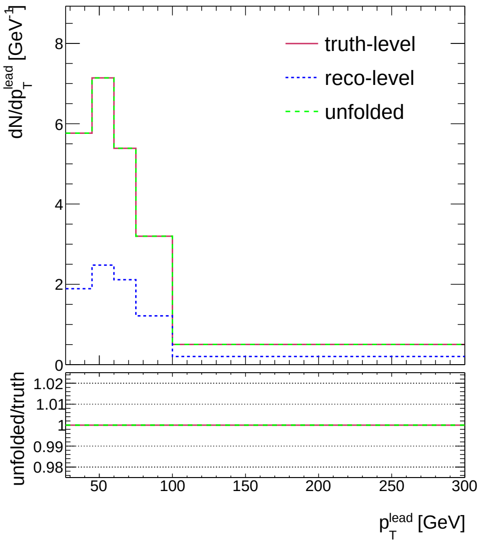

(a) leading lepton p_{T}

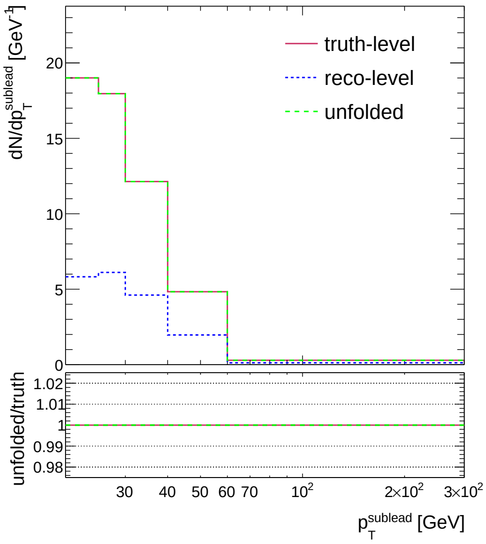

(b) subleading lepton p_{T}

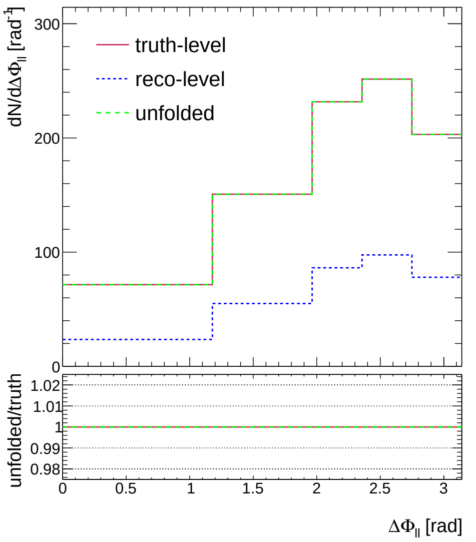

(c) $\Delta\Phi_{\ell\ell}$

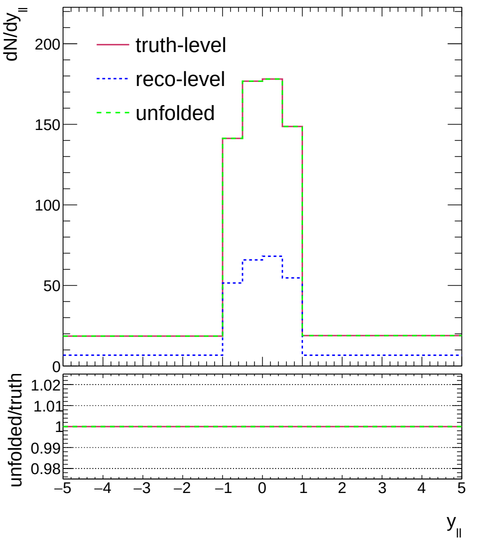

(d) $y_{\ell\ell}$

Figure D.1: Technical closure tests.

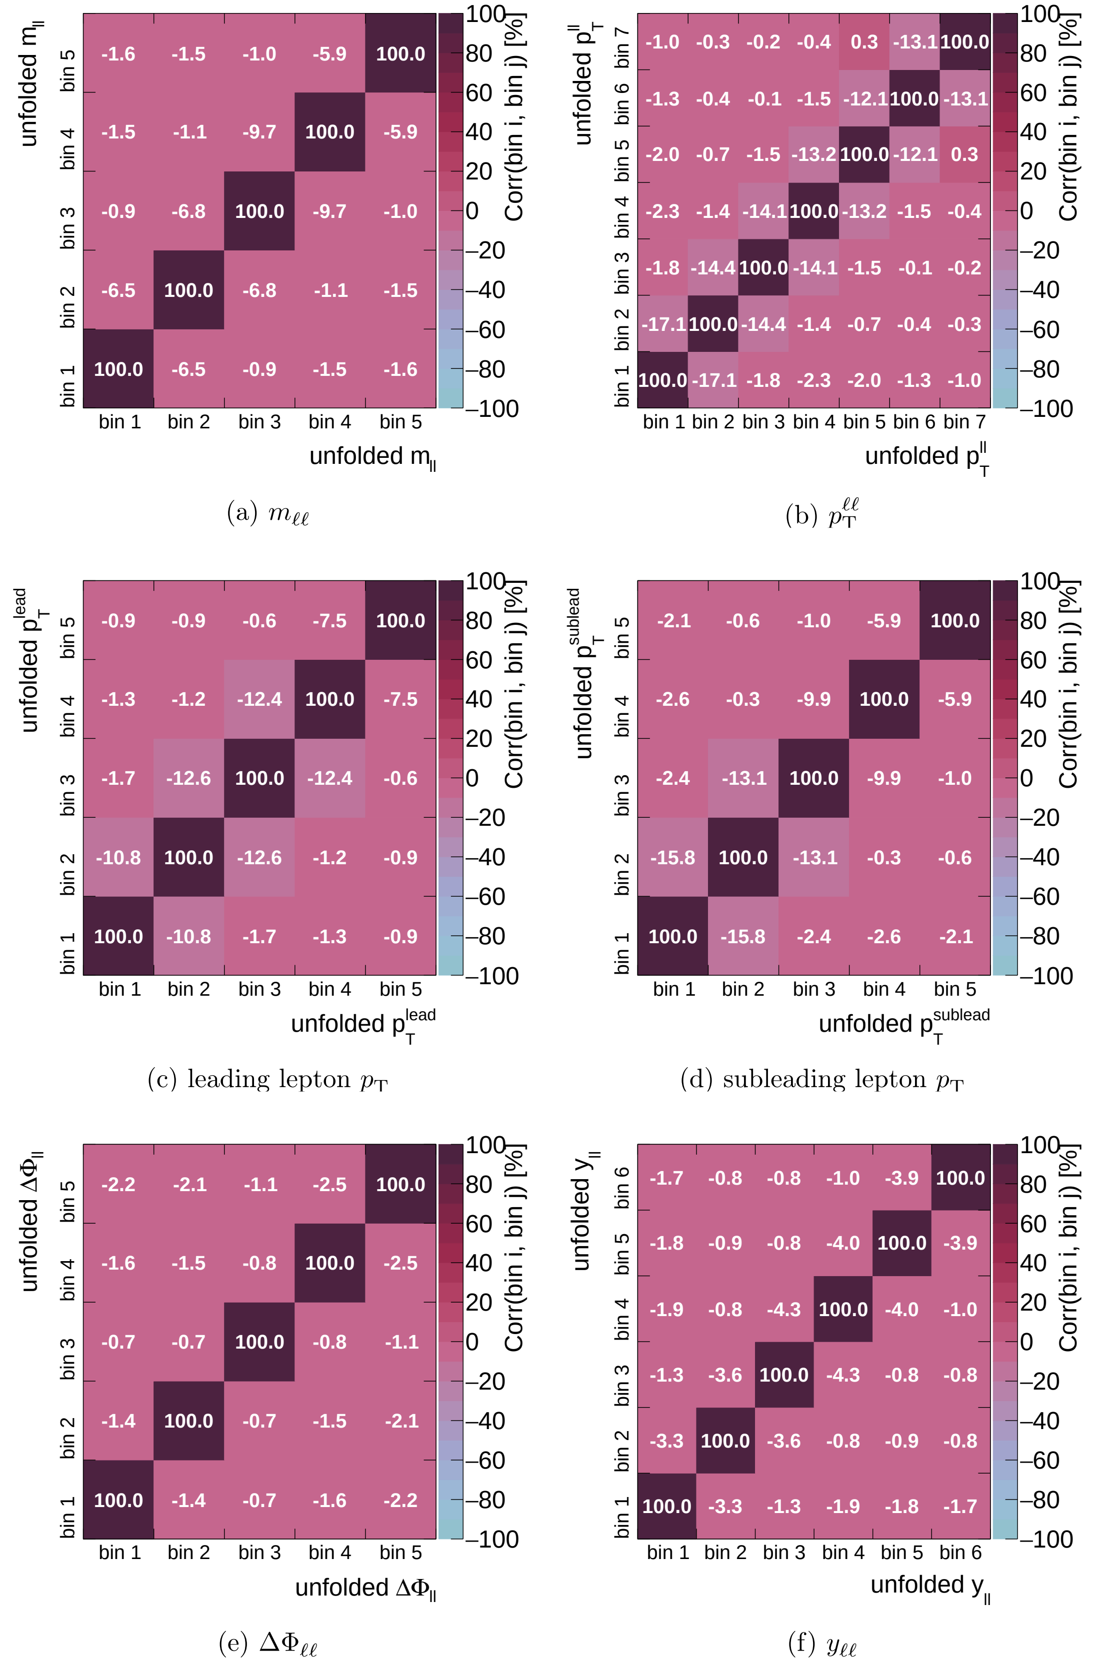

Figure D.2: Correlation matrices showing the statistical correlation between bins in the unfolded distributions. All bins are given equal width in these figures to aid readability.

References

[1] A. Einstein. "Über einen die Erzeugung und Verwandlung des Lichtes betreffenden heuristischen Gesichtspunkt". In: *Annalen der Physik* 322.6 (1905), pp. 132–148. DOI: 10.1002/andp.19053220607.

[2] Paul Adrien Maurice Dirac and Niels Henrik David Bohr. "The quantum theory of the emission and absorption of radiation". In: *Proceedings of the Royal Society of London. Series A, Containing Papers of a Mathematical and Physical Character* 114.767 (1927), pp. 243–265. DOI: 10.1098/rspa.1927.0039.

[3] ATLAS Collaboration. "Observation of Light-by-Light Scattering in Ultraperipheral Pb+Pb Collisions with the ATLAS Detector". In: *Phys. Rev. Lett.* 123 (2019), p. 052001. DOI: 10.1103/PhysRevLett.123.052001. arXiv: 1904.03536 [hep-ex].

[4] ATLAS Collaboration. "Observation of photon-induced W^+W^- production in pp collisions at $\sqrt{s} = 13\,\text{TeV}$ using the ATLAS detector". In: *Phys. Lett. B* 816 (2021), p. 136190. DOI: 10.1016/j.physletb.2021.136190. arXiv: 2010.04019 [hep-ex].

[5] ATLAS Collaboration. *Sensitivity to exclusive WW production in photon scattering at the High Luminosity LHC*. ATL-PHYS-PUB-2021-026. 2021. URL: https://cds.cern.ch/record/2776764.

[6] ATLAS Collaboration. "Observation and Measurement of Forward Proton Scattering in Association with Lepton Pairs Produced via the Photon Fusion Mechanism at ATLAS". In: *Phys. Rev. Lett.* 125 (2020), p. 261801. DOI: 10.1103/PhysRevLett.125.261801. arXiv: 2009.14537 [hep-ex].

[7] Christoph Englert. *Quantum Field Theory*. 2014. URL: https://conference.ippp.dur.ac.uk/event/785/attachments/3688/4135/QFTnotes.pdf (visited on 08/08/2022).

[8] Michael E. Peskin and Daniel V. Schroeder. *An Introduction to quantum field theory*. Reading, USA: Addison-Wesley, 1995. ISBN: 978-0-201-50397-5.

[9] Elena Accomando. *An Introduction to QED & QCD*. 2019. URL: https://conference.ippp.dur.ac.uk/event/785/attachments/3688/4138/QEDQCD-notes.pdf (visited on 08/08/2022).

[10] Fedor Bezrukov. *The Standard Model*. 2019. URL: https://www.hep.manchester.ac.uk/u/bezrukov/STFC_HEP_2019/SMlectures.pdf (visited on 08/08/2022).

[11] Martin White Andy Buckley Chris White. *Practical Collider Physics*. Institute of Physics Publishing, 2021. DOI: 10.1088/978-0-7503-2444-1.

[12] Particle Data Group. "Review of Particle Physics". In: *Progress of Theoretical and Experimental Physics* 083C01.8 (2022). DOI: 10.1093/ptep/ptac097.

[13] Max Bañados and Ignacio Reyes. "A short review on Noether's theorems, gauge symmetries and boundary terms". In: *International Journal of Modern Physics D* 25.10 (2016), p. 1630021. DOI: 10.1142/s0218271816300214.

[14] Jorge C. Romão and Joã P. Silva. "A resource for signs and Feynman diagrams of the Standard Model". In: *International Journal of Modern Physics A* 27.26 (2012), p. 1230025. DOI: 10.1142/s0217751x12300256.

[15] R. P. Feynman. "A Relativistic Cut-Off for Classical Electrodynamics". In: *Phys. Rev.* 74 (8 1948), pp. 939–946. DOI: 10.1103/PhysRev.74.939.

[16] Michael Binger and Stanley J. Brodsky. "Physical renormalization schemes and grand unification". In: *Phys. Rev. D* 69 (9 2004), p. 095007. DOI: 10.1103/PhysRevD.69.095007.

[17] Alexandre Deur, Stanley J. Brodsky, and Guy F. de Téramond. "The QCD running coupling". In: *Progress in Particle and Nuclear Physics* 90 (2016), pp. 1–74. DOI: https://doi.org/10.1016/j.ppnp.2016.04.003.

[18] C. S. Wu et al. "Experimental Test of Parity Conservation in Beta Decay". In: *Phys. Rev.* 105 (4 1957), pp. 1413–1415. DOI: 10.1103/PhysRev.105.1413.

[19] Steven Weinberg. "A Model of Leptons". In: *Phys. Rev. Lett.* 19 (1967), pp. 1264–1266. DOI: 10.1103/PhysRevLett.19.1264.

[20] Jeffrey Goldstone, Abdus Salam, and Steven Weinberg. "Broken Symmetries". In: *Phys. Rev.* 127 (3 1962), pp. 965–970. DOI: 10.1103/PhysRev.127.965.

[21] ATLAS Collaboration. "Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC". In: *Phys. Lett. B* 716 (2012), p. 1. DOI: 10.1016/j.physletb.2012.08.020. arXiv: 1207.7214 [hep-ex].

[22] CMS Collaboration. "Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC". In: *Phys. Lett. B* 716 (2012), p. 30. DOI: 10.1016/j.physletb.2012.08.021. arXiv: 1207.7235 [hep-ex].

[23] ATLAS Collaboration. *Standard Model Summary Plots February 2022*. ATL-PHYS-PUB-2022-009. 2022. URL: http://cds.cern.ch/record/2804061.

[24] Planck Collaboration. "Planck 2013 results. I. Overview of products and scientific results". In: *Astronomy & Astrophysics* 571 (2014), A1. DOI: 10.1051/0004-6361/201321529.

[25] Sean M. Carroll. "The Cosmological Constant". In: *Living Reviews in Relativity* 4.1 (2001). DOI: 10.12942/lrr-2001-1.

[26] Super-Kamiokande Collaboration. "Evidence for Oscillation of Atmospheric Neutrinos". In: *Physical Review Letters* 81.8 (1998), pp. 1562–1567. DOI: 10.1103/physrevlett.81.1562.

[27] SNO Collaboration. "Direct Evidence for Neutrino Flavor Transformation from Neutral-Current Interactions in the Sudbury Neutrino Observatory". In: *Phys. Rev. Lett.* 89 (1 2002), p. 011301. DOI: 10.1103/PhysRevLett.89.011301.

[28] A. D. Sakharov. "Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe". In: *Pisma Zh. Eksp. Teor. Fiz.* 5 (1967), pp. 32–35. DOI: 10.1070/PU1991v034n05ABEH002497.

[29] J. H. Christenson et al. "Evidence for the 2π Decay of the K_2^0 Meson". In: *Phys. Rev. Lett.* 13 (4 1964), pp. 138–140. DOI: 10.1103/PhysRevLett.13.138.

[30] kTeV Collaboration. "Observation of Direct CP Violation in $K_{S,L}$ to $\pi\pi$ Decays". In: *Physical Review Letters* 83.1 (1999), pp. 22–27. DOI: 10.1103/physrevlett.83.22.

[31] NA48 Collaboration. "A new measurement of direct CP violation in two pion decays of the neutral kaon". In: *Physics Letters B* 465.1-4 (1999), pp. 335–348. DOI: 10.1016/s0370-2693(99)01030-8.

[32] Steven Weinberg. "Baryon- and Lepton-Nonconserving Processes". In: *Phys. Rev. Lett.* 43 (21 1979), pp. 1566–1570. DOI: 10.1103/PhysRevLett.43.1566.

[33] Céline Degrande et al. "Effective field theory: A modern approach to anomalous couplings". In: *Annals of Physics* 335 (2013), pp. 21–32. DOI: 10.1016/j.aop.2013.04.016.

[34] Landon Lehman. "Extending the standard model effective field theory with the complete set of dimension-7 operators". In: *Phys. Rev. D* 90 (12 2014), p. 125023. DOI: 10.1103/PhysRevD.90.125023.

[35] B. Grzadkowski et al. "Dimension-six terms in the Standard Model Lagrangian". In: *Journal of High Energy Physics* 2010.10 (2010). DOI: 10.1007/jhep10(2010)085.

[36] O. J. P. Éboli and M. C. Gonzalez–Garcia. "Classifying the bosonic quartic couplings". In: *Phys. Rev. D* 93 (9 2016), p. 093013. DOI: 10.1103/PhysRevD.93.093013.

[37] J C Collins and D E Soper. "The Theorems of Perturbative QCD". In: *Annual Review of Nuclear and Particle Science* 37.1 (1987), pp. 383–409. DOI: 10.1146/annurev.ns.37.120187.002123.

[38] CTEQ Collaboration. *CTEQ: The Coordinated Theoretical-Experimental Project on QCD*. URL: https://www.physics.smu.edu/scalise/cteq/ (visited on 08/16/2022).

[39] NNPDF Collaboration. *NNPDF: Neural Network Parton Distribution Functions*. URL: http://nnpdf.mi.infn.it/ (visited on 08/16/2022).

[40] Richard D. Ball et al. “Parton distributions with LHC data”. In: *Nucl. Phys. B* 867 (2013), p. 244. DOI: `10.1016/j.nuclphysb.2012.10.003`. arXiv: `1207.1303 [hep-ph]`.

[41] Richard D. Ball et al. “Parton distributions for the LHC run II”. In: *JHEP* 04 (2015), p. 040. DOI: `10.1007/JHEP04(2015)040`. arXiv: `1410.8849 [hep-ph]`.

[42] Andy Buckley et al. “LHAPDF6: parton density access in the LHC precision era”. In: *The European Physical Journal C* 75.3 (2015). DOI: `10.1140/epjc/s10052-015-3318-8`.

[43] Carl Schmidt et al. “CT14QED parton distribution functions from isolated photon production in deep inelastic scattering”. In: *Phys. Rev. D* 93.11 (2016), p. 114015. DOI: `10.1103/PhysRevD.93.114015`. arXiv: `1509.02905 [hep-ph]`.

[44] A. D. Martin and M. G. Ryskin. “The photon PDF of the proton”. In: *The European Physical Journal C* 74.9 (2014). DOI: `10.1140/epjc/s10052-014-3040-y`.

[45] Enrico Fermi. “On the Theory of Collisions between Atoms and Electrically Charged Particles”. In: *Nuovo Cimento* (1925). Translated from Italian by M. Gallinaro and S. White (2001). DOI: `2:143-158,1925`.

[46] C. F. v. Weizsäcker. “Ausstrahlung bei Stößen sehr schneller Elektronen”. In: *Zeitschrift fur Physik* 88 (1934), pp. 612–625. DOI: `10.1007/BF01333110`.

[47] E. J. Williams. “Nature of the high-energy particles of penetrating radiation and status of ionization and radiation formulae”. In: *Phys. Rev.* 45 (1934), pp. 729–730. DOI: `10.1103/PhysRev.45.729`.

[48] C. A. Bertulani. “Electromagnetic interaction of ultrarelativistic heavy ions”. In: *Physical Review A* 63.6 (2001). DOI: `10.1103/physreva.63.062706`.

[49] C. A. Bertulani. “Photon exchange in nucleus-nucleus collisions”. In: *International Journal of Modern Physics A* 18.05 (2003), pp. 685–723. DOI: `10.1142/s0217751x03012357`.

[50] V.M. Budnev et al. “The two-photon particle production mechanism. Physical problems. Applications. Equivalent photon approximation”. In: *Physics Reports* 15.4 (1975), pp. 181–282. DOI: `https://doi.org/10.1016/0370-1573(75)90009-5`.

[51] ATLAS Collaboration. *Exclusive dielectron production in ultraperipheral Pb+Pb collisions at* $\sqrt{s_{NN}} = 5.02$ *TeV with ATLAS*. arXiv: `2207.12781 [hep-ex]`.

[52] ATLAS Collaboration. “Exclusive dimuon production in ultraperipheral Pb+Pb collisions at $\sqrt{s_{NN}} = 5.02$ TeV with ATLAS”. In: *Phys. Rev. C* 104 (2021), p. 024906. DOI: `10.1103/PhysRevC.104.024906`. arXiv: `2011.12211 [hep-ex]`.

[53] ATLAS Collaboration. *Observation of the $\gamma\gamma \to \tau\tau$ process in Pb+Pb collisions and constraints on the τ-lepton anomalous magnetic moment with the ATLAS detector*. 2022. arXiv: 2204.13478 [hep-ex].

[54] Lydia Beresford and Jesse Liu. "Search Strategy for Sleptons and Dark Matter Using the LHC as a Photon Collider". In: *Physical Review Letters* 123.14 (2019). DOI: 10.1103/physrevlett.123.141801.

[55] M. Trzebiński. "Diffractive Physics at the LHC". In: *Ukrainian Journal of Physics* 64.8 (2019), p. 772. DOI: 10.15407/ujpe64.8.772.

[56] L. A. Harland-Lang, V. A. Khoze, and M. G. Ryskin. "Exclusive physics at the LHC with SuperChic 2". In: *The European Physical Journal C* 76.1 (2016). DOI: 10.1140/epjc/s10052-015-3832-8.

[57] L. A. Harland-Lang et al. "Standard candle central exclusive processes at the Tevatron and LHC". In: *The European Physical Journal C* 69.1-2 (2010), pp. 179–199. DOI: 10.1140/epjc/s10052-010-1404-5.

[58] L. A. Harland-Lang et al. "The phenomenology of central exclusive production at hadron colliders". In: *The European Physical Journal C* 72.9 (2012). DOI: 10.1140/epjc/s10052-012-2110-2.

[59] E. Chapon, C. Royon, and O. Kepka. "Anomalous quartic $WW\gamma\gamma$, $ZZ\gamma\gamma$, and trilinear $WW\gamma$ couplings in two-photon processes at high luminosity at the LHC". In: *Physical Review D* 81.7 (2010). DOI: 10.1103/physrevd.81.074003.

[60] O. J. P. Éboli et al. "Anomalous quartic gauge boson couplings at hadron colliders". In: *Physical Review D* 63.7 (2001). DOI: 10.1103/physrevd.63.075008.

[61] Andy Buckley et al. "General-purpose event generators for LHC physics". In: *Physics Reports* 504.5 (2011), pp. 145–233. DOI: https://doi.org/10.1016/j.physrep.2011.03.005.

[62] ATLAS Collaboration. "The ATLAS Simulation Infrastructure". In: *Eur. Phys. J. C* 70 (2010), p. 823. DOI: 10.1140/epjc/s10052-010-1429-9. arXiv: 1005.4568 [physics.ins-det].

[63] GEANT4 Collaboration, S. Agostinelli, et al. "GEANT4 – a simulation toolkit". In: *Nucl. Instrum. Meth. A* 506 (2003), p. 250. DOI: 10.1016/S0168-9002(03)01368-8.

[64] B. Andersson et al. "Parton fragmentation and string dynamics". In: *Physics Reports* 97.2 (1983), pp. 31–145. DOI: https://doi.org/10.1016/0370-1573(83)90080-7.

[65] Manuel Bähr, Stefan Gieseke, and Michael H Seymour. "Simulation of multiple partonic interactions in Herwig++". In: *Journal of High Energy Physics* 2008.07 (2008), pp. 076–076. DOI: 10.1088/1126-6708/2008/07/076.

[66] Stefan Gieseke, Frashër Loshaj, and Patrick Kirchgaeßer. "Soft and diffractive scattering with the cluster model in Herwig". In: *The European Physical Journal C* 77.3 (2017). DOI: 10.1140/epjc/s10052-017-4727-7.

[67] B.R. Webber. "A QCD model for jet fragmentation including soft gluon interference". In: *Nuclear Physics B* 238.3 (1984), pp. 492–528. DOI: https://doi.org/10.1016/0550-3213(84)90333-X.

[68] ATLAS Collaboration. "Measurement of distributions sensitive to the underlying event in inclusive Z-boson production in pp collisions at $\sqrt{s} = 7\,\text{TeV}$ with the ATLAS detector". In: *Eur. Phys. J. C* 74 (2014), p. 3195. DOI: 10.1140/epjc/s10052-014-3195-6. arXiv: 1409.3433 [hep-ex].

[69] ATLAS Collaboration. "Measurement of charged-particle distributions sensitive to the underlying event in $\sqrt{s} = 13\,\text{TeV}$ proton–proton collisions with the ATLAS detector at the LHC". In: *JHEP* 03 (2017), p. 157. DOI: 10.1007/JHEP03(2017)157. arXiv: 1701.05390 [hep-ex].

[70] ATLAS Collaboration. "Measurement of distributions sensitive to the underlying event in inclusive Z boson production in pp collisions at $\sqrt{s} = 13\,\text{TeV}$ with the ATLAS detector". In: *Eur. Phys. J. C* 79 (2019), p. 666. DOI: 10.1140/epjc/s10052-019-7162-0. arXiv: 1905.09752 [hep-ex].

[71] T. Regge. "Introduction to complex orbital momenta". In: *Nuovo Cim.* 14 (1959), p. 951. DOI: 10.1007/BF02728177.

[72] Geoffrey F. Chew and Steven C. Frautschi. "Regge Trajectories and the Principle of Maximum Strength for Strong Interactions". In: *Phys. Rev. Lett.* 8 (1 1962), pp. 41–44. DOI: 10.1103/PhysRevLett.8.41.

[73] E. Levin. *Everything about reggeons. Part I:Reggeons in "soft" interaction.* 1997. arXiv: 9710546 [hep-ph].

[74] Manuel Bähr et al. "Herwig++ physics and manual". In: *The European Physical Journal C* 58.4 (2008), pp. 639–707. DOI: 10.1140/epjc/s10052-008-0798-9.

[75] Johannes Bellm et al. "Herwig 7.0/Herwig++ 3.0 release note". In: *Eur. Phys. J. C* 76.4 (2016), p. 196. DOI: 10.1140/epjc/s10052-016-4018-8. arXiv: 1512.01178 [hep-ph].

[76] Steffen Schumann and Frank Krauss. "A parton shower algorithm based on Catani–Seymour dipole factorisation". In: *JHEP* 03 (2008), p. 038. DOI: 10.1088/1126-6708/2008/03/038. arXiv: 0709.1027 [hep-ph].

[77] Christian Reuschle et al. "NLO efforts in Herwig++". In: (2016). arXiv: 1601.04101 [hep-ph].

[78] V.M. Budnev et al. "The two-photon particle production mechanism. Physical problems. Applications. Equivalent photon approximation". In: *Physics Reports* 15.4 (1975), pp. 181–282. DOI: https://doi.org/10.1016/0370-1573(75)90009-5.

[79] Leif Lonnblad. "THEPEG: Toolkit for High Energy Physics event generation". In: *Proceedings, HERA and the LHC Workshop Series on the implications of HERA for LHC physics: 2006-2008*. 2009, pp. 733–736. DOI: 10.3204/DESY-PROC-2009-02/42.

[80] T. Gleisberg et al. "Event generation with SHERPA 1.1". In: *JHEP* 02 (2009), p. 007. DOI: `10.1088/1126-6708/2009/02/007`. arXiv: `0811.4622 [hep-ph]`.

[81] Frank Krauss, Ralf Kuhn, and Gerhard Soff. "AMEGIC++ 1.0, A Matrix Element Generator In C++". In: *Journal of High Energy Physics* 2002.02 (2002), pp. 044–044. DOI: `10.1088/1126-6708/2002/02/044`.

[82] Tanju Gleisberg and Stefan Höche. "Comix, a new matrix element generator". In: *Journal of High Energy Physics* 2008.12 (2008), pp. 039–039. DOI: `10.1088/1126-6708/2008/12/039`.

[83] Stefano Frixione and Bryan R. Webber. "Matching NLO QCD computations and parton shower simulations". In: *JHEP* 06 (2002), p. 029. arXiv: `hep-ph/0204244`.

[84] Stefan Höche et al. "QCD matrix elements + parton showers. The NLO case". In: *JHEP* 04 (2013), p. 027. DOI: `10.1007/JHEP04(2013)027`. arXiv: `1207.5030 [hep-ph]`.

[85] F. Cascioli, P. Maierhöfer, and S. Pozzorini. "Scattering Amplitudes with Open Loops". In: *Phys. Rev. Lett.* 108 (11 2012), p. 111601. DOI: `10.1103/PhysRevLett.108.111601`.

[86] Paolo Nason. "A new method for combining NLO QCD with shower Monte Carlo algorithms". In: *JHEP* 11 (2004), p. 040. DOI: `10.1088/1126-6708/2004/11/040`. arXiv: `hep-ph/0409146`.

[87] Stefano Frixione, Paolo Nason, and Carlo Oleari. "Matching NLO QCD computations with parton shower simulations: the POWHEG method". In: *JHEP* 11 (2007), p. 070. DOI: `10.1088/1126-6708/2007/11/070`. arXiv: `0709.2092 [hep-ph]`.

[88] Simone Alioli et al. "A general framework for implementing NLO calculations in shower Monte Carlo programs: the POWHEG BOX". In: *JHEP* 06 (2010), p. 043. DOI: `10.1007/JHEP06(2010)043`. arXiv: `1002.2581 [hep-ph]`.

[89] Tom Melia et al. "W^+W^-, WZ and ZZ production in the POWHEG BOX". In: *JHEP* 11 (2011), p. 078. DOI: `10.1007/JHEP11(2011)078`. arXiv: `1107.5051 [hep-ph]`.

[90] Paolo Nason and Giulia Zanderighi. "W^+W^-, WZ and ZZ production in the POWHEG-BOX-V2". In: *Eur. Phys. J. C* 74.1 (2014), p. 2702. DOI: `10.1140/epjc/s10052-013-2702-5`. arXiv: `1311.1365 [hep-ph]`.

[91] H.-L. Lai et al. "New parton distributions for collider physics". In: *Phys. Rev. D* 82 (2010), p. 074024. DOI: `10.1103/PhysRevD.82.074024`. arXiv: `1007.2241 [hep-ph]`.

[92] L. A. Harland-Lang et al. "Parton distributions in the LHC era: MMHT 2014 PDFs". In: *Eur. Phys. J. C* 75.5 (2015). DOI: `10.1140/epjc/s10052-015-3397-6`. arXiv: `1412.3989 [hep-ph]`.

[93] Giuseppe Bozzi et al. "Transverse-momentum resummation and the spectrum of the Higgs boson at the LHC". In: *Nucl. Phys. B* 737 (2006), pp. 73–120. DOI: 10.1016/j.nuclphysb.2005.12.022. arXiv: hep-ph/0508068 [hep-ph].

[94] Daniel de Florian et al. "Transverse-momentum resummation: Higgs boson production at the Tevatron and the LHC". In: *JHEP* 11 (2011), p. 064. DOI: 10.1007/JHEP11(2011)064. arXiv: 1109.2109 [hep-ph].

[95] Torbjörn Sjöstrand et al. "An introduction to PYTHIA 8.2". In: *Comput. Phys. Commun.* 191 (2015), p. 159. DOI: 10.1016/j.cpc.2015.01.024. arXiv: 1410.3012 [hep-ph].

[96] Christian Bierlich et al. "A comprehensive guide to the physics and usage of PYTHIA 8.3". In: (2022). DOI: 10.48550/ARXIV.2203.11601. URL: https://arxiv.org/abs/2203.11601.

[97] T. Sjöstrand and P. Z. Skands. "Transverse-momentum-ordered showers and interleaved multiple interactions". In: *The European Physical Journal C* 39.2 (2005), pp. 129–154. DOI: 10.1140/epjc/s2004-02084-y.

[98] Leif Lönnblad. "Correcting the Colour-Dipole Cascade Model with Fixed Order Matrix Elements". In: *JHEP* 05 (2002), p. 046. DOI: 10.1088/1126-6708/2002/05/046. arXiv: hep-ph/0112284.

[99] Leif Lönnblad and Stefan Prestel. "Matching tree-level matrix elements with interleaved showers". In: *JHEP* 03 (2012), p. 019. DOI: 10.1007/JHEP03(2012)019. arXiv: 1109.4829 [hep-ph].

[100] Torbjörn Sjöstrand and Maria van Zijl. "A multiple-interaction model for the event structure in hadron collisions". In: *Phys. Rev. D* 36 (7 1987), pp. 2019–2041. DOI: 10.1103/PhysRevD.36.2019.

[101] ATLAS Collaboration. *The Pythia 8 A3 tune description of ATLAS minimum bias and inelastic measurements incorporating the Donnachie–Landshoff diffractive model.* ATL-PHYS-PUB-2016-017. 2016. URL: https://cds.cern.ch/record/2206965.

[102] Tadej Novak. *New techniques for pile-up simulation in ATLAS.* Tech. rep. 2018. DOI: 10.1051/epjconf/201921402044. URL: https://cds.cern.ch/record/2649298.

[103] J. Alwall et al. "The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations". In: *JHEP* 07 (2014), p. 079. DOI: 10.1007/JHEP07(2014)079. arXiv: 1405.0301 [hep-ph].

[104] R. Frederix et al. "The automation of next-to-leading order electroweak calculations". In: *Journal of High Energy Physics* 2018.7 (2018). DOI: 10.1007/jhep07(2018)185.

[105] Rikkert Frederix and Stefano Frixione. "Merging meets matching in MC@NLO". In: *JHEP* 12 (2012), p. 061. DOI: 10.1007/JHEP12(2012)061. arXiv: 1209.6215 [hep-ph].

[106] A. D. Martin et al. "Parton distributions incorporating QED contributions". In: *Eur. Phys. J.* C39 (2005), pp. 155–161. DOI: 10.1140/epjc/s2004-02088-7. arXiv: hep-ph/0411040 [hep-ph].

[107] NNPDF Collaboration. "Parton distributions with QED corrections". In: *Nucl. Phys.* B877 (2013), pp. 290–320. DOI: 10.1016/j.nuclphysb.2013.10.010. arXiv: 1308.0598 [hep-ph].

[108] J. A. M. Vermaseren. "Two Photon Processes at Very High-Energies". In: *Nucl. Phys.* B229 (1983), pp. 347–371. DOI: 10.1016/0550-3213(83)90336-X.

[109] S. P. Baranov et al. "LPAIR: A generator for lepton pair production". In: *Workshop on Physics at HERA Hamburg, Germany, October 29-30, 1991.* 1991, pp. 1478–1482.

[110] Valerio Bertone et al. "Illuminating the photon content of the proton within a global PDF analysis". In: *SciPost Physics* 5.1 (2018). DOI: 10.21468/scipostphys.5.1.008.

[111] 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 (2019). DOI: 10.1140/epjc/s10052-018-6530-5.

[112] L. A. Harland-Lang et al. "A new approach to modelling elastic and inelastic photon-initiated production at the LHC: SuperChic 4". In: *The European Physical Journal C* 80.10 (2020). DOI: 10.1140/epjc/s10052-020-08455-0.

[113] Oliver Sim Brüning et al. *LHC Design Report.* CERN Yellow Reports: Monographs. Geneva: CERN, 2004. URL: https://cds.cern.ch/record/782076.

[114] R W Assmann, M Lamont, and S Myers. "A Brief History of the LEP Collider". In: *Nucl. Phys. B, Proc. Suppl.* 109 (2002), 17–31. 15 p. DOI: 10.1016/S0920-5632(02)90005-8.

[115] Julie Haffner. "The CERN accelerator complex. Complexe des accélérateurs du CERN". In: (2013). URL: https://cds.cern.ch/record/1621894.

[116] ATLAS Collaboration. *ATLAS Detector and Physics Performance: Technical Design Report, Volume 1.* ATLAS-TDR-14; CERN-LHCC-99-014. 1999. URL: https://cds.cern.ch/record/391176.

[117] ATLAS Collaboration. *ATLAS Detector and Physics Performance: Technical Design Report, Volume 2.* ATLAS-TDR-15; CERN-LHCC-99-015. 1999. URL: https://cds.cern.ch/record/391177.

[118] CMS Collaboration. "The CMS experiment at the CERN LHC". In: *JINST* 3 (2008), S08004. DOI: 10.1088/1748-0221/3/08/S08004.

[119] *LHCb reoptimized detector design and performance: Technical Design Report.* Technical design report. LHCb. Geneva: CERN, 2003. URL: https://cds.cern.ch/record/630827.

[120] *ALICE: Technical proposal for a Large Ion collider Experiment at the CERN LHC*. LHC technical proposal. Geneva: CERN, 1995. URL: https://cds.cern.ch/record/293391.

[121] Michaela Schaumann et al. "First Xenon-Xenon Collisions in the LHC". In: *9th International Particle Accelerator Conference*. 2018. DOI: 10.18429/JACoW-IPAC2018-MOPMF039.

[122] Werner Herr and B Muratori. "Concept of luminosity". In: (2006). URL: https://cds.cern.ch/record/941318.

[123] P. Grafström and W. Kozanecki. "Luminosity determination at proton colliders". In: *Progress in Particle and Nuclear Physics* 81 (2015), pp. 97–148. DOI: https://doi.org/10.1016/j.ppnp.2014.11.002.

[124] CERN. *The HL-LHC project*. URL: https://hilumilhc.web.cern.ch/content/hl-lhc-project (visited on 08/23/2022).

[125] I Béjar Alonso et al. *High-Luminosity Large Hadron Collider (HL-LHC): Technical design report*. Ed. by I Béjar Alonso. CERN Yellow Reports: Monographs. Geneva: CERN, 2020. DOI: 10.23731/CYRM-2020-0010. URL: http://cds.cern.ch/record/2749422.

[126] Joao Pequenao. "Computer generated image of the whole ATLAS detector". 2008. URL: https://cds.cern.ch/record/1095924.

[127] Izaak Neutelings. *CMS coordinate system*. URL: https://tikz.net/axis3d_cms/ (visited on 08/22/2022).

[128] ATLAS Collaboration. *ATLAS Inner Detector: Technical Design Report, Volume 1*. ATLAS-TDR-4; CERN-LHCC-97-016. 1997. URL: https://cds.cern.ch/record/331063.

[129] ATLAS Collaboration. *ATLAS Inner Detector: Technical Design Report, Volume 2*. ATLAS-TDR-5, CERN-LHCC-97-017. 1997. URL: https://cds.cern.ch/record/331064.

[130] ATLAS Collaboration. "Alignment of the ATLAS Inner Detector in Run-2". In: *Eur. Phys. J. C* 80 (2020), p. 1194. DOI: 10.1140/epjc/s10052-020-08700-6. arXiv: 2007.07624 [hep-ex].

[131] ATLAS Collaboration. *ATLAS Pixel Detector: Technical Design Report*. ATLAS-TDR-11; CERN-LHCC-98-013. 1998. URL: https://cds.cern.ch/record/381263.

[132] ATLAS Collaboration. *ATLAS Insertable B-Layer: Technical Design Report*. ATLAS-TDR-19; CERN-LHCC-2010-013. 2010. URL: https://cds.cern.ch/record/1291633.

[133] C. Goessling et al. "Planar n+-in-n silicon pixel sensors for the ATLAS IBL upgrade". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 650 (2011), pp. 198–201. DOI: 10.1016/j.nima.2010.11.186.

[134] Cinzia et al Da Via. "3D silicon sensors: Design, large area production and quality assurance for the ATLAS IBL pixel detector upgrade". In: *Nucl. Instrum. Meth. A* 694 (2012), pp. 321–330. DOI: 10.1016/j.nima.2012.07.058.

[135] ATLAS Collaboration. "Operation and performance of the ATLAS semiconductor tracker". In: *JINST* 9 (2014), P08009. DOI: 10.1088/1748-0221/9/08/P08009. arXiv: 1404.7473 [hep-ex].

[136] The ATLAS TRT collaboration. "The ATLAS TRT Barrel Detector". In: *Journal of Instrumentation* 3.02 (2008), P02014–P02014. DOI: 10.1088/1748-0221/3/02/p02014.

[137] The ATLAS TRT collaboration. "The ATLAS TRT end-cap detectors". In: *Journal of Instrumentation* 3.10 (2008), P10003–P10003. DOI: 10.1088/1748-0221/3/10/p10003.

[138] The ATLAS TRT collaboration. "The ATLAS Transition Radiation Tracker (TRT) proportional drift tube: design and performance". In: *Journal of Instrumentation* 3.02 (2008), P02013–P02013. DOI: 10.1088/1748-0221/3/02/p02013.

[139] X. Artru, G. B. Yodh, and G. Mennessier. "Practical theory of the multilayered transition radiation detector". In: *Phys. Rev. D* 12 (5 1975), pp. 1289–1306. DOI: 10.1103/PhysRevD.12.1289.

[140] ATLAS Collaboration. *ATLAS Calorimeter Performance: Technical Design Report*. ATLAS-TDR-1; CERN-LHCC-96-040. 1996. URL: https://cds.cern.ch/record/331059.

[141] ATLAS Collaboration. *ATLAS Liquid Argon Calorimeter: Technical Design Report*. ATLAS-TDR-2; CERN-LHCC-96-041. 1996. URL: https://cds.cern.ch/record/331061.

[142] ATLAS Collaboration. *ATLAS Tile Calorimeter: Technical Design Report*. ATLAS-TDR-3; CERN-LHCC-96-042. 1996. URL: https://cds.cern.ch/record/331062.

[143] ATLAS Collaboration. "Jet energy scale and resolution measured in proton–proton collisions at $\sqrt{s} = 13\,\text{TeV}$ with the ATLAS detector". In: (2020). arXiv: 2007.02645 [hep-ex].

[144] ATLAS Collaboration. *ATLAS Muon Spectrometer: Technical Design Report*. ATLAS-TDR-10; CERN-LHCC-97-022. CERN, 1997. URL: https://cds.cern.ch/record/331068.

[145] A. Aloisio et al. "The trigger chambers of the ATLAS muon spectrometer: production and tests". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 535.1 (2004). Proceedings of the 10th International Vienna Conference on Instrumentation, pp. 265–271. DOI: https://doi.org/10.1016/j.nima.2004.07.130.

[146] ATLAS Collaboration. *ATLAS Central Solenoid: Magnet Project Technical Design Report, Volume 4*. ATLAS-TDR-9; CERN-LHCC-97-021. CERN, 1997. URL: https://cds.cern.ch/record/331067.

[147] ATLAS Collaboration. *ATLAS Barrel Toroid: Magnet Project Technical Design Report, Volume 2*. ATLAS-TDR-7; CERN-LHCC-97-019. 1997. URL: https://cds.cern.ch/record/331065.

[148] ATLAS Collaboration. *ATLAS End-Cap Toroids: Magnet Project Technical Design Report, Volume 3*. ATLAS-TDR-8; CERN-LHCC-97-020. CERN, 1997. URL: https://cds.cern.ch/record/331066.

[149] Roger Ruber ATLAS Magnet Group. *ATLAS Magnet System*. URL: https://atlas-magnet.web.cern.ch/atlas-magnet/ (visited on 08/23/2022).

[150] ATLAS Collaboration. *ATLAS Level-1 Trigger: Technical Design Report*. ATLAS-TDR-12; CERN-LHCC-98-014. 1998. URL: https://cds.cern.ch/record/381429.

[151] ATLAS Collaboration. *ATLAS High-Level Trigger, Data Acquisition and Controls: Technical Design Report*. ATLAS-TDR-16; CERN-LHCC-2003-022. 2003. URL: https://cds.cern.ch/record/616089.

[152] ATLAS Collaboration. *ATLAS Tracking Software Tutorial*. URL: %5Curl%7Bhttps://atlassoftwaredocs.web.cern.ch/trackingTutorial/%7D (visited on 08/25/2022).

[153] 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 (1987), pp. 444–450. DOI: https://doi.org/10.1016/0168-9002(87)90887-4.

[154] T G Cornelissen et al. "The global χ^2 track fitter in ATLAS". In: *Journal of Physics: Conference Series* 119.3 (2008), p. 032013. DOI: 10.1088/1742-6596/119/3/032013.

[155] ATLAS Collaboration. "Performance of the ATLAS track reconstruction algorithms in dense environments in LHC Run 2". In: *Eur. Phys. J. C* 77 (2017), p. 673. DOI: 10.1140/epjc/s10052-017-5225-7. arXiv: 1704.07983 [hep-ex].

[156] ATLAS Collaboration. "Reconstruction of primary vertices at the ATLAS experiment in Run 1 proton–proton collisions at the LHC". In: *Eur. Phys. J. C* 77 (2017), p. 332. DOI: 10.1140/epjc/s10052-017-4887-5. arXiv: 1611.10235 [hep-ex].

[157] Izaac Sanderswood. "Development of ATLAS Primary Vertex Reconstruction for LHC Run 3". In: *Connecting the Dots and Workshop on Intelligent Trackers*. 2019. arXiv: 1910.08405 [hep-ex].

[158] Walter Lampl et al. *Calorimeter Clustering Algorithms: Description and Performance*. ATL-LARG-PUB-2008-002. 2008. URL: https://cds.cern.ch/record/1099735.

[159] ATLAS Collaboration. "Electron and photon performance measurements with the ATLAS detector using the 2015–2017 LHC proton–proton collision data". In: *JINST* 14 (2019), P12006. DOI: `10.1088/1748-0221/14/12/P12006`. arXiv: `1908.00005 [hep-ex]`.

[160] ATLAS Collaboration. "Topological cell clustering in the ATLAS calorimeters and its performance in LHC Run 1". In: *Eur. Phys. J. C* 77 (2017), p. 490. DOI: `10.1140/epjc/s10052-017-5004-5`. arXiv: `1603.02934 [hep-ex]`.

[161] ATLAS Collaboration. *Electron and photon reconstruction and performance in ATLAS using a dynamical, topological cell clustering-based approach*. ATL-PHYS-PUB-2017-022. 2017. URL: `https://cds.cern.ch/record/2298955`.

[162] W Lampl et al. *Calorimeter Clustering Algorithms: Description and Performance*. Tech. rep. 2008. URL: `https://cds.cern.ch/record/1099735`.

[163] ATLAS Collaboration. *Electron and photon energy calibration with the ATLAS detector using data collected in 2015 at $\sqrt{s}$ = 13 TeV*. ATL-PHYS-PUB-2016-015. 2016. URL: `https://cds.cern.ch/record/2203514`.

[164] Sébastien Rettie. "Muon identification and performance in the ATLAS experiment". In: *PoS* DIS2018 (2018), p. 097. DOI: `10.22323/1.316.0097`.

[165] ATLAS Collaboration. "Muon reconstruction performance of the ATLAS detector in proton–proton collision data at $\sqrt{s} = 13\,\mathrm{TeV}$". In: *Eur. Phys. J. C* 76 (2016), p. 292. DOI: `10.1140/epjc/s10052-016-4120-y`. arXiv: `1603.05598 [hep-ex]`.

[166] J. Illingworth and J. Kittler. "A survey of the hough transform". In: *Computer Vision, Graphics, and Image Processing* 44.1 (1988), pp. 87–116. DOI: `https://doi.org/10.1016/S0734-189X(88)80033-1`.

[167] ATLAS Collaboration. "Jet reconstruction and performance using particle flow with the ATLAS Detector". In: *Eur. Phys. J. C* 77 (2017), p. 466. DOI: `10.1140/epjc/s10052-017-5031-2`. arXiv: `1703.10485 [hep-ex]`.

[168] ATLAS Collaboration. "Muon reconstruction and identification efficiency in ATLAS using the full Run 2 pp collision data set at $\sqrt{s} = 13\,\mathrm{TeV}$". In: (2020). arXiv: `2012.00578 [hep-ex]`.

[169] ATLAS Collaboration. *ATLAS Forward Proton Phase-I Upgrade: Technical Design Report*. ATLAS-TDR-024; CERN-LHCC-2015-009. 2015. URL: `https://cds.cern.ch/record/2017378`.

[170] Collaboration CMS-TOTEM. *CMS-TOTEM Precision Proton Spectrometer*. Tech. rep. 2014. URL: `https://cds.cern.ch/record/1753795`.

[171] TOTEM Collaboration CMS Collaboration. *Search for exclusive $\gamma\gamma \to WW$ and $\gamma\gamma \to ZZ$ production in final states with jets and forward protons*. Tech. rep. 2022. URL: `https://cds.cern.ch/record/2803716`.

[172] The ATLAS Collaboration. *AFP Figures*. URL: `https://twiki.cern.ch/twiki/bin/viewauth/Atlas/AFP_Figures` (visited on 07/25/2022).

[173] Collaboration CMS-TOTEM. *CMS-TOTEM Precision Proton Spectrometer*. Tech. rep. 2014. URL: https://cds.cern.ch/record/1753795.

[174] ATLAS Collaboration. *Luminosity determination in pp collisions at $\sqrt{s} = 13$ TeV using the ATLAS detector at the LHC*. ATLAS-CONF-2019-021. 2019. URL: https://cds.cern.ch/record/2677054.

[175] LHC Optics Working Group. *LHC Optics Web: LHC Run II pp physics optics*. URL: https://abpdata.web.cern.ch/abpdata/lhc_optics_web/www/opt2017/.

[176] Maciej Trzebiński. "Machine optics studies for the LHC measurements". In: *SPIE Proceedings*. 2014. DOI: 10.1117/12.2074647.

[177] J. Lange et al. "3D silicon pixel detectors for the ATLAS Forward Physics experiment". In: *Journal of Instrumentation* 10.03 (2015), pp. C03031–C03031. DOI: 10.1088/1748-0221/10/03/c03031.

[178] L. Nozka et al. "Design of Cherenkov bars for the optical part of the time-of-flight detector in Geant4". In: *Opt. Express* 22.23 (2014), pp. 28984–28996. DOI: 10.1364/OE.22.028984.

[179] ATLAS Collaboration. *Performance of the ATLAS Forward Proton Time-of-Flight Detector in 2017*. ATL-FWD-PUB-2021-002. 2021. URL: https://cds.cern.ch/record/2749821.

[180] Tom Sykora. "ATLAS Forward Proton Time-of-Flight Detector: LHC Run2 performance and experiences". In: *Journal of Instrumentation* 15 (2020), pp. C10004–C10004. DOI: 10.1088/1748-0221/15/10/C10004.

[181] ATLAS Collaboration. *Exclusive Jet Production with Forward Proton Tagging. Feasibility Studies for the AFP Project*. ATL-PHYS-PUB-2015-003. 2015. URL: https://cds.cern.ch/record/1993686.

[182] R. Staszewski and J. Chwastowski. "Transport simulation and diffractive event reconstruction at the LHC". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 609.2-3 (2009), pp. 136–141. DOI: 10.1016/j.nima.2009.08.023.

[183] R. Staszewski et al. "Alignment-related effects in forward proton experiments at the LHC". In: *Nuclear Instruments and Methods in Physics Research Section A: Accelerators, Spectrometers, Detectors and Associated Equipment* 801 (2015), pp. 34–43. DOI: 10.1016/j.nima.2015.08.010.

[184] Gianluca Valentino et al. "Semiautomatic beam-based LHC collimator alignment". In: *Phys. Rev. ST Accel. Beams* 15 (5 2012), p. 051002. DOI: 10.1103/PhysRevSTAB.15.051002.

[185] Gianluca Valentino et al. "Final implementation, commissioning, and performance of embedded collimator beam position monitors in the Large Hadron Collider". In: *Phys. Rev. Accel. Beams* 20 (8 2017), p. 081002. DOI: 10.1103/PhysRevAccelBeams.20.081002.

[186] ATLAS Collaboration. *Evaluating statistical uncertainties and correlations using the bootstrap method*. ATL-PHYS-PUB-2021-011. 2021. URL: https://cds.cern.ch/record/2759945.

[187] ATLAS Collaboration. "Measurement of exclusive $\gamma\gamma \to W^+W^-$ production and search for exclusive Higgs boson production in pp collisions at $\sqrt{s} = 8$ TeV using the ATLAS detector". In: *Phys. Rev. D* 94 (2016), p. 032011. DOI: 10.1103/PhysRevD.94.032011. arXiv: 1607.03745 [hep-ex].

[188] CMS Collaboration. "Study of exclusive two-photon production of W^+W^- in pp collisions at $\sqrt{s} = 7$ TeV and constraints on anomalous quartic gauge couplings". In: *JHEP* 07 (2013), p. 116. DOI: 10.1007/JHEP07(2013)116. arXiv: 1305.5596 [hep-ex].

[189] CMS Collaboration. "Evidence for exclusive $\gamma\gamma \to W^+W^-$ production and constraints on anomalous quartic gauge couplings in pp collisions at $\sqrt{s} = 7$ and 8 TeV". In: *JHEP* 08 (2016), p. 119. DOI: 10.1007/JHEP08(2016)119. arXiv: 1604.04464 [hep-ex].

[190] ATLAS Collaboration. "ATLAS data quality operations and performance for 2015–2018 data-taking". In: *JINST* 15 (2020), P04003. DOI: 10.1088/1748-0221/15/04/P04003. arXiv: 1911.04632 [physics.ins-det].

[191] ATLAS Collaboration. *Electron efficiency measurements with the ATLAS detector using the 2015 LHC proton–proton collision data*. ATLAS-CONF-2016-024. 2016. URL: https://cds.cern.ch/record/2157687.

[192] T. Pierog et al. "EPOS LHC: Test of collective hadronization with data measured at the CERN Large Hadron Collider". In: *Phys. Rev. C* 92 (3 2015), p. 034906. DOI: 10.1103/PhysRevC.92.034906.

[193] ATLAS Collaboration. "Measurement of the Inelastic Proton–Proton Cross Section at $\sqrt{s} = 13$ TeV with the ATLAS Detector at the LHC". In: *Phys. Rev. Lett.* 117 (2016), p. 182002. DOI: 10.1103/PhysRevLett.117.182002. arXiv: 1606.02625 [hep-ex].

[194] ATLAS Collaboration. "Measurement of the exclusive $\gamma\gamma \to \mu^+\mu^-$ process in proton–proton collisions at $\sqrt{s} = 13$ TeV with the ATLAS detector". In: *Phys. Lett. B* 777 (2018), p. 303. DOI: 10.1016/j.physletb.2017.12.043. arXiv: 1708.04053 [hep-ex].

[195] ATLAS Collaboration. "Charged-particle distributions at low transverse momentum in $\sqrt{s} = 13$ TeV pp interactions measured with the ATLAS detector at the LHC". In: *Eur. Phys. J. C* 76 (2016), p. 502. DOI: 10.1140/epjc/s10052-016-4335-y. arXiv: 1606.01133 [hep-ex].

[196] Stefan Kallweit et al. "Accurate single- and double-differential resummation of colour-singlet processes with MATRIX+RadISH: W^+W^- production at the LHC". In: *JHEP* 147 (2020). DOI: 10.1007/JHEP12(2020)147.

[197] ATLAS Collaboration. "Evidence of $W\gamma\gamma$ Production in pp Collisions at $\sqrt{s} = 8$ TeV and Limits on Anomalous Quartic Gauge Couplings with the ATLAS Detector". In: *Phys. Rev. Lett.* 115 (2015), p. 031802. DOI: 10.1103/PhysRevLett.115.031802. arXiv: 1503.03243 [hep-ex].

[198] S. Bailey and L. A. Harland-Lang. "Modeling W^+W^- production with rapidity gaps at the LHC". In: *Physical Review D* 105.9 (2022). DOI: 10.1103/physrevd.105.093010.

[199] ATLAS Collaboration. "Performance of electron and photon triggers in ATLAS during LHC Run 2". In: *Eur. Phys. J. C* 80 (2020), p. 47. DOI: 10.1140/epjc/s10052-019-7500-2. arXiv: 1909.00761 [hep-ex].

[200] ATLAS Collaboration. "Performance of the ATLAS muon triggers in Run 2". In: *JINST* 15 (2020), P09015. DOI: 10.1088/1748-0221/15/09/p09015. arXiv: 2004.13447 [hep-ex].

[201] Glen Cowan et al. "Asymptotic formulae for likelihood-based tests of new physics". In: *Eur. Phys. J. C* 71 (2011), p. 1554. DOI: 10.1140/epjc/s10052-011-1554-0. arXiv: 1007.1727 [physics.data-an]. Erratum: in: *Eur. Phys. J. C* 73 (2013), p. 2501. DOI: 10.1140/epjc/s10052-013-2501-z.

[202] ROOT Collaboration. *HistFactory: A tool for creating statistical models for use with RooFit and RooStats*. Tech. rep. 2012. URL: https://cds.cern.ch/record/1456844.

[203] Laurent Forthomme et al. "Rapidity gap survival factors caused by remnant fragmentation for W^+W^- pair production via $\gamma^*\gamma^* \to W^+W^-$ subprocess with photon transverse momenta". In: *Phys. Lett. B* 789 (2019), pp. 300–307. DOI: 10.1016/j.physletb.2018.12.018. arXiv: 1805.07124 [hep-ph].

[204] Mateusz Dyndal and Laurent Schoeffel. "The role of finite-size effects on the spectrum of equivalent photons in proton–proton collisions at the LHC". In: *Physics Letters B* 741 (2015), pp. 66–70. DOI: 10.1016/j.physletb.2014.12.019.

[205] L. A. Harland-Lang, V. A. Khoze, and M. G. Ryskin. "The photon PDF in events with rapidity gaps". In: *The European Physical Journal C* 76.5 (2016), p. 255. DOI: 10.1140/epjc/s10052-016-4100-2.

[206] L. A. Harland-Lang et al. "Ad Lucem: QED parton distribution functions in the MMHT framework". In: *The European Physical Journal C* 79.10 (2019). DOI: 10.1140/epjc/s10052-019-7296-0.

[207] C. Baldenegro et al. "Central exclusive production of W boson pairs in pp collisions at the LHC in hadronic and semi-leptonic final states". In: *JHEP* 12 (2020), p. 165. DOI: 10.1007/JHEP12(2020)165. arXiv: 2009.08331 [hep-ph].

[208] Alix Fell. "The observation of photon-induced W^+W^- production and studies of electron identification with the ATLAS experiment". PhD thesis. University of Sheffield, Sheffield U., 2022.

[209] Giulio D'Agostini. "A multidimensional unfolding method based on Bayes' theorem". In: *Nucl. Instrum. Meth. A* 362.2 (1995), pp. 487–498. ISSN: 0168-9002. DOI: 10.1016/0168-9002(95)00274-X.

[210] G. D'Agostini. "Improved iterative Bayesian unfolding". In: (2010). arXiv: 1010.0632 [physics.data-an].

[211] Tim Adye. "Unfolding algorithms and tests using RooUnfold". In: *Proceedings, 2011 Workshop on Statistical Issues Related to Discovery Claims in Search Experiments and Unfolding (PHYSTAT 2011)* (CERN, Geneva, Switzerland, Jan. 17–20, 2011), pp. 313–318. DOI: 10.5170/CERN-2011-006.313. arXiv: 1105.1160 [physics.data-an].

[212] Jerome H. Friedman. "A Variable Span Smoother". In: (1984). URL: https://www.slac.stanford.edu/pubs/slacpubs/3250/slac-pub-3477.pdf.

[213] ATLAS Collaboration. *Expected pile-up values at the HL-LHC*. ATL-UPGRADE-PUB-2013-014. 2013. URL: https://cds.cern.ch/record/1604492.

[214] ATLAS Collaboration. *Expected performance of the ATLAS detector at the High-Luminosity LHC*. ATL-PHYS-PUB-2019-005. 2019. URL: https://cds.cern.ch/record/2655304.

[215] ATLAS Collaboration. *Characterization of Interaction-Point Beam Parameters Using the pp Event-Vertex Distribution Reconstructed in the ATLAS Detector at the LHC*. ATLAS-CONF-2010-027. 2010. URL: https://cds.cern.ch/record/1277659.

[216] ATLAS and CMS Collaborations. *Snowmass White Paper Contribution: Physics with the Phase-2 ATLAS and CMS Detectors*. ATL-PHYS-PUB-2022-018. 2022. URL: https://cds.cern.ch/record/2805993.

[217] ATLAS Collaboration. *ATLAS Inner Tracker Pixel Detector: Technical Design Report*. ATLAS-TDR-030; CERN-LHCC-2017-021. 2017. URL: https://cds.cern.ch/record/2285585.

[218] ATLAS Collaboration. *ATLAS Inner Tracker Strip Detector: Technical Design Report*. ATLAS-TDR-025; CERN-LHCC-2017-005. 2017. URL: https://cds.cern.ch/record/2257755.

[219] ATLAS Collaboration. *ATLAS LAr Calorimeter Phase-II Upgrade: Technical Design Report*. ATLAS-TDR-027; CERN-LHCC-2017-018. 2017. URL: https://cds.cern.ch/record/2285582.

[220] ATLAS Collaboration. *ATLAS Tile Calorimeter Phase-II Upgrade: Technical Design Report*. ATLAS-TDR-028; CERN-LHCC-2017-019. 2017. URL: https://cds.cern.ch/record/2285583.

[221] ATLAS Collaboration. *ATLAS Muon Spectrometer Phase-II Upgrade: Technical Design Report*. ATLAS-TDR-026; CERN-LHCC-2017-017. 2017. URL: https://cds.cern.ch/record/2285580.

[222] ATLAS Collaboration. *ATLAS TDAQ Phase-II Upgrade: Technical Design Report*. ATLAS-TDR-029; CERN-LHCC-2017-020. 2017. URL: https://cds.cern.ch/record/2285584.

[223] ATLAS Collaboration. *A High-Granularity Timing Detector for the ATLAS Phase-II Upgrade: Technical Design Report*. ATLAS-TDR-031; CERN-LHCC-2020-007. 2020. URL: https://cds.cern.ch/record/2719855.

[224] ATLAS Collaboration. *ITk Pixel Layout Updates*. 2020. URL: https://atlas.web.cern.ch/Atlas/GROUPS/PHYSICS/PLOTS/ITK-2020-002/ (visited on 09/05/2022).

[225] ATLAS Collaboration. *Expected Tracking Performance of the ATLAS Inner Tracker at the HL-LHC*. ATL-PHYS-PUB-2019-014. 2019. URL: https://cds.cern.ch/record/2669540.

[226] ATLAS Collaboration. *Low-p_{T} tracking for ATLAS in nominal LHC pileup*. ATL-PHYS-PROC-2020-041. 2020. URL: https://cds.cern.ch/record/2718583.

[227] Luis Eduardo Medina Medrano et al. *Effective pile-up density as a measure of the experimental data quality for High-Luminosity LHC operational scenarios*. Tech. rep. 2018. URL: https://cds.cern.ch/record/2301928.

[228] ATLAS Collaboration. *Prospective study of vector boson scattering in WZ fully leptonic final state at HL-LHC*. ATL-PHYS-PUB-2018-023. 2018. URL: https://cds.cern.ch/record/2645271.

[229] M. Carulla et al. "Technology developments and first measurements on inverse Low Gain Avalanche Detector (iLGAD) for high energy physics applications". In: *Journal of Instrumentation* 11.12 (2016), pp. C12039–C12039. DOI: 10.1088/1748-0221/11/12/c12039.

[230] CMS Collaboration. *The CMS Precision Proton Spectrometer at the HL-LHC – Expression of Interest*. Tech. rep. 2021. DOI: 10.48550/ARXIV.2103.02752. URL: https://arxiv.org/abs/2103.02752.

[231] ATLAS Collaboration. *Studies on top-quark Monte Carlo modelling for Top2016*. ATL-PHYS-PUB-2016-020. 2016. URL: https://cds.cern.ch/record/2216168.

[232] ATLAS Collaboration. *ATLAS Pythia 8 tunes to 7 TeV data*. ATL-PHYS-PUB-2014-021. 2014. URL: https://cds.cern.ch/record/1966419.

[233] Elias Metral et al. "Update of the HL-LHC operational scenarios for proton operation". In: (2018). URL: https://cds.cern.ch/record/2301292.

[234] ATLAS Collaboration. "Charged-particle distributions in pp interactions at $\sqrt{s} = 8\,\mathrm{TeV}$ measured with the ATLAS detector". In: *Eur. Phys. J. C* 76 (2016), p. 403. DOI: 10.1140/epjc/s10052-016-4203-9. arXiv: 1603.02439 [hep-ex].

[235] ATLAS Collaboration. "Charged-particle distributions in $\sqrt{s} = 13\,\mathrm{TeV}$ pp interactions measured with the ATLAS detector at the LHC". In: *Phys. Lett. B* 758 (2016), p. 67. DOI: 10.1016/j.physletb.2016.04.050. arXiv: 1602.01633 [hep-ex].

[236] Richard Corke and Torbjörn Sjöstrand. "Interleaved parton showers and tuning prospects". In: *JHEP* 2011.3 (2011). DOI: 10.1007/jhep03(2011)032.

[237] H.J. Drescher et al. "Parton-based Gribov-Regge theory". In: *Physics Reports* 350.2-4 (2001), pp. 93–289. DOI: 10.1016/s0370-1573(00)00122-8.

[238] Tim Adye. *Corrected error calculation for iterative Bayesian unfolding.* 2011. URL: https://hepunx.rl.ac.uk/~adye/software/unfold/bayes_errors.pdf (visited on 10/19/2022).